Methods for analysing emerging data sources to understand variability in traveller behaviour on the road network

Fiona Crawford

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds

Institute for Transport Studies

September 2017

The candidate confirms that the work submitted is her own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The thesis has the following structure. Firstly, an introductory section positions the research within the relevant literature and details the objectives of the research. The next four chapters consist of original research where the candidate was the lead researcher, written in the format of journal papers. In all four cases, the research is based on ideas developed by the candidate during the Ph.D. with the support and guidance of her two supervisors who are the co-authors. The candidate undertook all of the analysis in the papers and wrote the papers. The co-authors provided extensive feedback on the content and presentation of the papers. The thesis ends with a discussion and conclusion section which brings together the findings of the four papers and discusses the implications of the research.

Chapter 2 has been published as follows: Crawford, F., Watling, D. P. & Connors, R. D. (2017) A statistical method for estimating predictable differences between daily traffic flow profiles. *Transportation Research Part B: Methodological,* 95, p196-213.

Chapter 3 was submitted for review to the journal Transportation and the candidate has received an invitation to revise and resubmit with major revisions.

Chapter 4 was submitted for review to Transportation Research Part A: Policy and Practice, and the candidate has received an invitation to revise the manuscript.

Acknowledgements

Firstly, I would like to thank my fantastic supervision team: David Watling and Richard Connors. I am extremely grateful for all of the guidance, support and advice that they have given me during this journey. They have taught me to aim high and opened up numerous opportunities for me to meet new people and learn new things. Their enthusiasm is inspiring.

I would also like to thank other members of staff at ITS, particularly those in the Spatial Modelling and Dynamics research group, who have provided constructive feedback on work I have presented. A special thank you goes to Judith Wang who has always made sure I am well fed! I am also indebted to all of the researchers who have provided feedback after my conference and workshop presentations, including INSTR 2015 and UTSG 2017. Takamasa Iryo and Jan-Dirk Schmöcker were particularly generous in letting me visit their universities and meet with them and their students.

This research would not have been possible without the funding provided by the UK Engineering and Physical Sciences Research Council and Highways England through an Industrial CASE PhD Scholarship. All of the data used was provided by Transport for Greater Manchester and I am extremely grateful that they allowed me to access their data, answered all of my questions and engaged in discussions regarding my research.

I could not have made it through the past four years without the support, encouragement and cake provided the other PhD students at ITS, particularly Joanna, Julian, Alvaro, Jingyan, Nur, Haruko and Anna. The supportive environment makes ITS a fantastic place to study.

Finally, I would like to thank my family for their constant support and Stephen for his unwavering belief in me.

Abstract

This thesis argues that while simplifications are a necessary part of the modelling process, there is a lack of empirical research to identify which types of variability should be included in our models, and how they should be represented. This research aims to develop methodologies to undertake empirical analyses of variability on the road network, focusing specifically on traveller behaviour. This is particularly timely given the emergence of rich new data sources.

Firstly, a method is proposed for examining predictable differences in daily link flow profiles by considering both magnitude and timing. Unlike previous methods, this approach can test for statistically significant differences whilst also comparing the shapes of the profiles, by applying Functional Linear Models to transportation data for the first time.

Secondly, a flexible, data-driven method is proposed for classifying road users based on their trip frequency and spatial and temporal intrapersonal variability. Previous research has proposed methodologies for identifying *public transport* user classes based on repeated trip behaviour, but equivalent methods for data from the road network did not exist. As there was not an established data source to use, this research examines the feasibility of using Bluetooth data. Spatial variability is examined using Sequence Alignment which has not previously been applied to Bluetooth data from road networks, nor for spatial intrapersonal variability. The time of day variability analysis adapts a technique from smart card research so that *all* observations are classified into travel patterns and, therefore, systematic and random variability can be measured.

These network- and traveller-focused analyses are then brought together using a framework which uses them concurrently and interactively to gain additional insights into traveller behaviour. For each of the methods proposed, an application to at least one year of real world data is presented.

Table of Contents

List of Tables

List of Figures

1 Introduction

1.1 Context

Aggregate-level transport models seek to provide a representation of a transport network which can replicate interactions between traveller choices and congestion on the network so that traffic conditions on unmonitored links or under different forecasting scenarios can be estimated (Sheffi, 1985). Any mathematical model is a *simplified* representation of reality (Bender, 1978); for example, ignoring the different free flow speeds of different road users is described by Beckmann et al. (1956, p47) as a “convenient fiction”. Network models abound with convenient fictions, although many are not explicitly discussed. Many of these relate to variability. As in the example from Beckmann et al. (1956), early transport network models often used the convenient fiction that all travellers and all days were identical, or at least that there was sufficient “within-period stationarity” (Cascetta, 2009, p10) that any variability in transport characteristics in the modelled period would be negligible. For example, when referring to Wardrop’s User Equilibrium (Wardrop, 1952), Daganzo and Sheffi (1977, p254) state that “implicit in the U-E criterion is that motorists are identical and infallible individuals”. Dial (1971, p85) states that to overcome the shortcomings of modelling approaches, transport planners are “highly skilled at taking the computerized model’s output with an appropriate grain of salt and adjusting its figures with his pencil”. Such an approach is not robust or replicable, however, and is unlikely to adequately reflect the impact of uncertainty in a non-linear system.

Models have, therefore, been developed which progressed from assumptions of homogeneity by substituting random variables with standard statistical distributions for single values in the model. For example, Stochastic User Equilibrium was developed for probabilistic traffic assignment, where travellers' perceived travel times are random variables with a given distribution and traveller choices are based on random utility theory (Daganzo and Sheffi, 1977). Alternative ways have been proposed to introduce heterogeneity into different aspects of assignment models, including using multiple road user classes to represent different traveller characteristics, and varying demand by the time of day.

Assumptions regarding stationarity or homogeneity are made on multiple levels and in multiple dimensions of network models, and are not always obvious. Intrapersonal variability, the variability in an individual's behaviour from day to day, is usually disregarded in equilibrium models. Such static models are based on simplified representations of utility maximisation in trip choices and therefore it is irrelevant whether it is the same individuals who are travelling, or different travellers making the same choices. This is consistent with the observation of Ortúzar and Willumsen (2011, p2), that models are "viewpoint specific", and in this case the repeated behaviour of individual travellers is not the focus of the model. Research which explored traffic conditions when networks were not at equilibrium, however, used the adaptive behaviour of individual travellers to justify the day-to-day dynamics in network performance. While some research implicitly suggested that the same people were making the same trip each day with potential route switching, for example Smith (1979), Horowitz (1984) make the explicit assumption that it is the same people travelling in the same time period on each day. Much of the data collection relating to day-to-day dynamics implicitly strengthened this assumption as it involved laboratory-style research where participants were asked to make hypothetical route or departure time choices on a series of 'days', for example Mahmassani et al. (1986), Iida et al. (1992) and Bogers et al. (2007). More complex information sharing models have been proposed more recently, involving sharing between individuals (Xiao and Lo, 2016, Shang et al., 2016) or through automated information systems (Ben-Elia and Shiftan, 2010, Li et al., 2017). Such mechanisms do not require the assumption that the same people are travelling each day in order to be credible. Models have not yet been developed, however, which include alternative assumptions, for example that a subset of travellers only travel on a fixed subset of days of the week.

Assumptions are also made regarding the variability between travellers. Examples of such assumptions include homogeneity in free flow speeds (as in Beckmann et al. (1956)), their preference for different routes, their knowledge of possible routes and information availability. Researchers have long argued that models should not assume that all travellers are identical and make 'optimal' choices (Dial, 1971, Dafermos, 1972). For example, Stochastic User Equilibrium (SUE) was developed to include variation in route choices by assuming that travellers have *perceived* travel times which are represented by a known statistical distribution around the actual travel time (Daganzo and Sheffi, 1977). Although SUE is a well-established concept, research continues into the assumptions to use for the random component within the route choice model (Prashker and Bekhor, 2004, Shahhoseini et al., 2015, Paz et al., 2016) and how the principles can be suitably applied to models which include within day dynamics (Han, 2003, Wei et al., 2014, Paz et al., 2016). Statistical distributions representing heterogeneous values of time have also been included in static (Huang and Li, 2007) and dynamic models (Zhang et al., 2013). Heterogeneity has also been introduced by applying different parameters to discrete "market segments" (Cascetta, 2009, p17) known as user classes. These classes could relate to a wide variety of characteristics, although the most obvious example, perhaps, relates to vehicle types (Dafermos, 1972). Multiple user classes have also been proposed to deal with different attitudes to risk (Shao et al., 2006, Szeto et al., 2011, Sun et al., 2015, Miralinaghi et al., 2016, Liu et al., 2017), level of network experience (Han et al., 2016), availability of advanced traveller information systems (Huang and Li, 2007, Lou et al., 2016) and income (Shahhoseini et al., 2015). Recently, vehicle classes have also been used to model the effect of mixed fleets including electric vehicles (Agrawal et al., 2016, Xu et al., 2017) and autonomous vehicles (Bagloee et al., 2017). While innumerable different user classes could be defined, the question remains: “what are the fundamentally different user classes?” (Peeta and Ziliaskopoulos, 2001, p252).

Convenient fictions also exist at the aggregate level. Even where multiple user classes are used to describe different subsets of the travelling population, the segmentation proportions are assumed to be fixed over time. The two key inputs to network models, namely demand and capacity, were traditionally assumed to be fixed over the modelled reference period. Different fields of research have emerged, however, which seek to represent variability in these two inputs. Dynamic Traffic Assignment extends static models by specifying demand as a function of the time of day and thus the modelled reference periods do not have constant, evenly distributed demand. Peeta and Ziliaskopoulos (2001) provide an overview of this field. In contrast to this ‘within-day’ variability, other researchers have considered the demand input to a static equilibrium model as a random variable as opposed to a single value. These random variables have been represented by standard statistical distributions such as the Binomial, Poisson, Beta-binomial or Negative binomial distribution (Nakayama and Watling, 2014). Capacity was also traditionally assumed to be fixed (for the modelled period) but particularly due to increased interest in travel time reliability, this assumption has been challenged. Equilibrium models exist which include link capacities as random variables representing unpredictable disturbances (Lo and Tung, 2003) or including unpredictable as well as predictable disturbances (Lam et al., 2008), for example due to adverse weather conditions. Stochasticity can also be generated endogenously by the model, for example through using route choice probabilities (Watling, 2002).

This thesis argues that while convenient fictions, for example relating to stability or a known probability distribution, are necessary and an inevitable part of the modelling process, there is a lack of empirical research which seeks to steer the research agenda and identify which types of variability should be included in our models, and how they should be represented (for example by including systematic and random variations, or by identifying the most suitable probability distribution). This is particularly timely given the emergence of rich data sources in recent years which can provide different kinds of information which we could not have expected to obtain in the past from traditional data sources, such as detailed information about travel behaviour over long periods of time. Data sources which have been used within transportation research more recently include mobile phone data (Calabrese et al., 2011, Iqbal et al., 2014, Gundlegård et al., 2016), Bluetooth data (Barceló et al., 2013, Kieu et al., 2015a, Yu et al., 2015) and GPS traces, for example from taxis (Liu et al., 2009, Liu et al., 2012, Yang et al., 2017).

The motivation for this research, and for much of the previous research looking at variability in network models (including Clark and Watling (2005), Lo et al. (2006), Shao et al. (2006) and Szeto et al. (2011)), is the study of travel time *reliability*. Research into travel time reliability covers a broad range of topics (Taylor, 2013) including the measurement and valuation of reliability, risk analysis and types of variability. Understanding the relevant sources of variability is essential, however, so that they can be appropriately accounted for in the calculation of measures of reliability, appropriately generated by models which seek to do so endogenously (i.e. by using mechanisms which replicate underlying relationships) and appropriately represented (for example, using the most suitable probability distribution). More generally, Jones and Clarke (1988, p65) also argue that “an understanding of variability is central to the modelling of travel behaviour and the assessment of policy impacts". As there is much debate regarding the definition of travel time reliability (Taylor, 2013), it may be prudent to start exploring what we mean by 'understanding' by examining the reliability measures used in practice. The Planning Time Index is used to measure travel time reliability on the Strategic Road Network in England (Department for Transport, 2017). The use of this index is supported by the U.S. Department of Transportation (Federal Highway Administration (FHWA), 2014) and consists of the 95th percentile of travel times divided by the free flow travel time. In England, the free flow travel time is defined as the time it would take a vehicle to traverse a stretch of road if travelling at the speed limit on the given road. As the name 'Planning Time Index' suggests, this measure is supposed to relate to the additional amount of time a traveller should allow for travelling along a given link so that the link does not make them late 95% of the time. It does not, however, take into account that travellers' expectations are likely to be influenced by experiences of recurrent congestion (Pu, 2011) or, in recent times, the availability of real time travel information. As travel time reliability is closely aligned with traveller expectations, it is a dynamic concept which depends upon which travellers are on the network at any given time. Therefore, to supplement the information obtained from a measure of travel time reliability such as the Planning Time Index, information about systematic and random variations between days in terms of who is using the network, what choices they are making and the effective network capacity is required. In Section 1.2, previous research using empirical analyses to examine these additional aspects of variability will be discussed.

1.2 Empirically based research into variability

Road network performance can vary from day to day due to a range of demand or supply-side factors (Carrion and Levinson, 2012). This research will focus on demand-side factors, specifically those aspects related to traveller behaviour; these are assumed to be factors which transport planners and policymakers may be able to influence, unlike exogenous factors such as the weather. The impact of supply-side factors on capacity (assuming a stochastic capacity as in Brilon et al. (2005)) cannot be measured directly and, therefore, future work will be required which could build upon outputs from the current research.

In this thesis, the term ‘travel behaviour’ is used to describe the following collection of choices:

- To travel or not to travel,
- Mode choice,
- Origin and destination choices,
- Departure time choice and
- Route choice.

The variability considered in this research is the variability between days as this is usually what is measured when assessing travel time reliability. As travel behaviour comprises a complex combination of traveller choices, this variability includes both spatial and temporal components. The spatial aspects include the choice of origin and destination, but also the route choice. When considering day-to-day variability, the temporal component is the time of day at which a trip is made. Variability in travel behaviour can be analysed from different perspectives; this research will concentrate on network-focused and traveller-focused analyses. Empirical research undertaken from these two different perspectives will now be considered in turn.

1.2.1 Network-focused analysis

Network-focused analyses utilise data aggregated across travellers to examine network usage and performance. This could include measuring variability in travel times on a link, or volumes of vehicles crossing a bridge. Bates et al. (2001) propose that variability in demand can be separated into a predictable component, for example due to seasonal or day of the week effects, and a random component. In this research that notion is extended to all types of variability, whether it is variability in an individual's departure time or in a link's travel times. This distinction is particularly important for network-focused analyses as the predictable differences are more likely to be driven by variability in travel behaviour. Identifying such systematic differences can provide the basis for scenario testing of future policies, for example in Kim et al. (2013), and can also inform the development of policies which include predictable differences in charges or service frequency by the day of the week or season, for example.

Table 1-1 includes network-focused empirical research which examines predictable or systematic variability, for example by the day of the week or due to different weather conditions. The research generally falls into one of three categories. The first category uses descriptive statistics to suggest predictable differences do or do not exist, for example Kaltenbrunner et al. (2010) and Tao et al. (2014). The second category uses statistical tests to *test* whether predictable differences exist, for example Watling et al. (2012) and Calvert et al. (2016). The third category uses data driven methods to identify systematic differences between days and then tries to explain the different groups identified, for example Weijermars and van Berkum (2005) and Guardiola et al. (2014). The second category has the advantage of being able to determine which predictable differences are sufficiently important to justify scenario testing or specific policy development, and is also guaranteed to be based on factors known to the analyst, for example the day of the week or the amount of precipitation.

The treatment of the time of day differs among the network-focused analyses shown in Table 1-1. In some cases, the time of day is not, or is assumed not to be, relevant, for example when considering the number of trips per day (Arana et al., 2014) or route choice (Watling et al., 2012). Other research considers both time of the day and systematic differences between days, for example according to the day of the week or season, but they do so independently so that systematic differences in timings between days are not considered, for example Stathopoulos and Karlaftis (2001) and Yeon et al. (2009). In other research, the analysis is undertaken separately for different periods during the day, for example using hourly data in Datla and Sharma (2008) and morning, mid-day and evening data in Gao and Niemeier (2007). Only the research which has involved data driven clustering has taken the shape of daily flow profiles into account when examining the differences between days (Weijermars and van Berkum, 2005, Guardiola et al., 2014).

1.2.2 Traveller-focused analysis

In contrast, traveller-focused analyses use data from individual travellers to measure how travel behaviour varies for the same person from day to day (intrapersonal variability). By ignoring intrapersonal variability and using the convenient fiction of a constant set of travellers who make the same trip every day, the assumptions shown in Figure 1-1 are made implicitly. These assumptions are particularly important in any modelling which includes behaviour change, for example adaptation to changing circumstances or policies to encourage behaviour change. The explicit and implicit assumptions should, therefore, not be taken for granted, but should be tested in a wide range of circumstances.

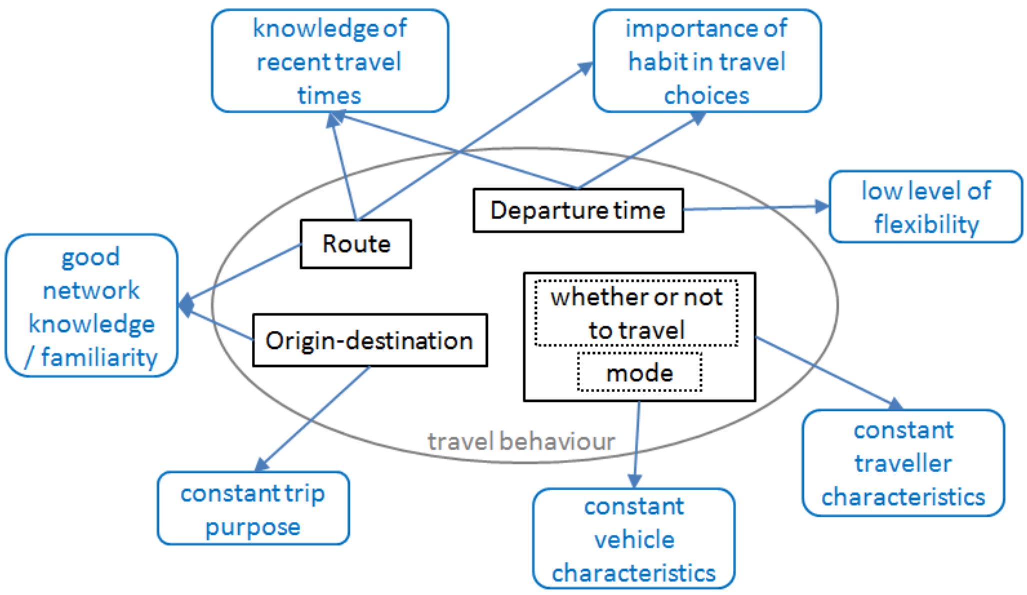

Figure 1-1: Implicit assumptions when assuming high levels of regularity

To test these assumptions, suitable data is required. The current research considers travel behaviour *on the road network* and therefore the data collected should have an explicit connection to the transport network. Previous research relating to smart card data (Ma et al., 2013, Kieu et al., 2015b) and in-vehicle GPS trackers (Elango et al., 2007, Venter and Joubert, 2013) are therefore more relevant than research using mobile phone data (Järv et al., 2014) or traditional travel diaries (Bayarma et al., 2007, Heinen and Chatterjee, 2015). As described for public transportation using smart card data by Pelletier et al. (2011), insights from traveller-focused analyses can inform strategic, tactical and operational planning by providing information about user types and needs. Empirical analyses have been used to define public transport user classes based on day-to-day trip making characteristics (Kieu et al., 2015b, Goulet Langlois et al., 2016), but similar analyses have not been undertaken for road users only, perhaps due to a difficulty in collecting suitable data.

Table 1-2 provides a summary of recent traveller-focused variability research. While research traditionally used travel diary data and so focused on all modes, more recently research has also included analyses of public transport users, using smart card data, and car trips, using GPS surveys. As can be seen in Table 1-2, new technologies such as smart cards or mobile phones result in much larger sample sizes and longer study periods than was previously possible. Smart card data has the advantage that it usually has high levels of coverage, for example over 90% of users in Beijing (Ma et al., 2013) and 80% of ticket sales in Brisbane (Tao et al., 2014). Traveller-focused research into day-to-day variability on the road network, however, has generally focused on personal travel as opposed to travel for commercial purposes, for example Li et al. (2004), Elango et al. (2007) and Spissu et al. (2011). Personal travel makes up the majority of vehicles on the road network, but as shown in a recent report on London, other types of vehicles, particularly Light Goods Vehicles are becoming increasingly important (INRIX, 2016). There has also been a 12% increase in the number of Private Hire Vehicles licenced in England between 2013 and 2015, although this has largely been driven by a 26% increase in London, perhaps exacerbated by the congestion charging zone. There is, therefore, a gap in the research relating to day-to-day variability in travel behaviour relating to all vehicle types on the road network.

Table 1-2 also demonstrates the variety of aspects of intrapersonal variability and the variety of methods used in previous research. Spatial measures of variability are strongly influenced by the type of data available, for example route choice can be considered using GPS data (Spissu et al., 2011), activity spaces using mobile phone data (Järv et al., 2014) and origin-destination (OD) comparisons using smart card data (Ma et al., 2013). Therefore, while some researchers have focused on OD variability (Buliung et al., 2008, Dill and Broach, 2014) and others on route variability (Li et al., 2004, Spissu et al., 2011), there is a gap in research which examines both types of variability together. Shen et al. (2013) examined both spatial and route flexibility, but they only considered commuting trips and they defined OD and route choices simply as either the same or different.

Comparisons of whether a trip is made at the same time of day also vary. For example, the measure of intrapersonal regularity could be based on the percentage of trips within 20% of the median departure time (Muthyalagari et al., 2001), the percentage of trips deviating from the median departure time by more than 5, 10 or 30 minutes (Li et al., 2004) or whether trips begin in the same 10 minute interval of the day (Minnen et al., 2015). Other research has maintained the continuous nature of time (Chikaraishi et al., 2009, Kieu et al., 2015b). It is preferable to treat the time of day as a continuous variable, as this enables an analysis of variability in greater detail, but it only makes sense to do so where the time of day is recorded at a high level of precision, for example when using emerging data sources. Time of day variability analysed in continuous time has not been undertaken on data from road users only and therefore adaptations to existing approaches would be required, depending on the type of data used.

There is, therefore, very little research examining the repeated travel behaviour of road users and a gap in empirical research examining *all types* of road user (not just cars used for non-commercial trips). As much of the previous research has been undertaken on travel diary or smart card data which do not usually contain route information, methods do not exist for quantifying spatial variability in both OD and route choices together. Methods for analysing time of day variability in data from emerging sources with high levels of temporal resolution exist, but would need to be adapted for use on road users. Despite the recognised advantages of analysing data for the purpose of defining network user classes based on multi-day travel behaviour, such analyses have yet to be undertaken on data from road users only.

1.2.3 Combining perspectives

Although empirical analyses usually take either the network-focused or the traveller-focused perspective, this thesis proposes that both approaches are equally valuable and that additional benefits can be obtained by undertaking the two processes simultaneously. These additional insights are obtained for two reasons. Firstly, network-focused analyses can be used to demonstrate whether or not any variability observed at the individual traveller level has an impact on network performance. Secondly, undertaking traveller-focused analyses may help us to understand why aggregate level systematic variability is occurring, for example by identifying that people leave work systematically earlier on certain days of the week. The gaps in the network- and traveller-focused research undertaken to date are, therefore, not completely distinct areas of research as a methodology could be developed to link the two processes in order to obtain a greater understanding of traveller behaviour.

Table 1-1: Literature on network-focused empirical analyses of systematic variability

Paper	Focus	Data type	Systematic variability type	Method for comparison	Relevant findings
Rakha and Van Aerde (1995)	To examine variability within and between days of the week, and also the impact of incidents	Multiple loop detectors	Day of the week	ANOVA Regression Analysis	Flows were similar on weekdays but different on weekend days Monday flows were different from core weekdays (Tuesday-Thursday), but speeds and occupancy were not statistically different at 95% confidence level Fridays were statistically different from core weekdays with respect to all three measures
Stathopoulos and Karlaftis (2001)	To test for differences in flow by year, month, day of the week, time of day and direction of travel	Multiple loop detectors	Year, month and day of the week	Kruskal-Wallis test Wilcoxon rank-sum U test	Flows exhibited little variation between days of the week and months (with the exception of the summer months)
Weijermars and van Berkum (2005)	To cluster days based on the shape and height of flow profiles	One loop detector	Day of the week and holiday periods	Hierarchical clustering	Working days (non-public holiday weekdays) were easier to cluster than non-working days Four clusters of working days were identified: Mondays, core weekdays, Fridays and days within holiday periods
Li et al. (2006b)	To examine sources of travel time variability on	Electronic toll tag and	Day of week and weather	Multiple regression	Weather and incidents had more of an impact on travel times in the afternoon

	a toll road	video data		analysis	peak than the morning peak
Liu and Sharma (2006)	To examine the impact of different national holidays on weekly, daily and hourly traffic flows	Multiple traffic counters	Before, during and after holiday periods	Wilcoxon matched pair test, Friedman method and Chi Square and Binomial tests	Holidays in winter months had a weaker effect Effects varied based on road usage type and direction of travel
Gao and Niemeier (2007)	To compare the day of the week effect at different times of day on the volumes and ratios of Light and Heavy Goods Vehicles	Multiple weigh-in-motion stations	Day of the week	Nonparametric factorial analysis of longitudinal data ANOVA-Type Statistic	Differences were observed by the day of the week in all three time of day periods
Zhang et al. (2007)	To compare link and path flows by day of the week	Multiple loop detectors	Day of the week	ANOVA	Path flows differed on Fridays to other weekdays, but link flows did not
Datla and Sharma (2008)	To examine the effect of snow and cold temperatures on traffic volumes	Multiple traffic counters	Weather (by day of the week)	Regression model and t-tests	Effects varied by the time of day, day of the week and road type Larger impacts were observed on recreational routes and during the off-peak Fridays were affected differently by very cold weather than Monday-Thursdays
Billot et al. (2009)	To examine the impact of rain on the performance of the road network at the micro, meso and macroscopic	Multiple loop detectors	Weather	Descriptive statistics	During adverse weather conditions, speeds dropped, headways increased, there was more platooning and capacity decreased

	levels				Jam density was not affected by rain
Yeon et al. (2009)	To test for differences in road capacity by the time of day, day of the week and road segment type	Multiple traffic microwave sensors	Day of the week	Levene's test and ANOVA	In most cases, mean capacity flows did not differ according to the day of the week Capacity flows differed between morning and evening peaks and off-peak periods
Kaltenbrunner et al. (2010)	To detect temporal and spatial patterns in community bicycle scheme usage	Cycle hire station occupancy data	Day of the week	Descriptive statistics	Demand profiles differed between weekdays and weekend days Weekday profiles also differed between station locations
Miranda-Moreno and Nosal (2011)	To examine the relationship between weather and cycle usage and also temporal trends in cycling	Multiple loop detectors for bicycles	Weather and day of the week	Descriptive statistics Regression models	Bicycle volumes were affected by rain at the time or earlier in the day Temperature had an effect, but it was not a linear effect Usage was higher on weekdays and during the summer months
Watling et al. (2012)	To assess the impact of network capacity reductions and to examine the effectiveness of model predictions	Partial licence plate data	With and without a planned capacity reduction	Two-sample, unequal variance t-test for mean flows, proportions and travel times	"considerable ambient daily variation in flows could easily mask any systematic effect" (Watling et al., 2012, p187) Flow proportions were useful in

					demonstrating changes in route choices
Yazici et al. (2012)	To analyse three travel time reliability measures based on the time of day and day of the week	GPS taxi data	Day of the week and time of day	Classification and Regression Tree methodology	The reliability-based categorisations of time of day did not conform to the typical peak/off-peak periods The reliability-based categorisations did not always combine data from all weekdays
Arana et al. (2014)	To examine the impact of weather on the number of shopping, personal business and leisure trips made by bus	Smart card	Weather	Student's t-test Multiple linear regression	Weather did have a significant impact on the total number of trips per day Higher temperatures resulted in higher usage on both Saturdays and Sundays, but wet or windy weather had more of an impact on Saturdays
Guardiola et al. (2014)	To cluster daily flow profiles based on their shape and height	One loop detector	Data generated	Functional Principal Component Analysis	Three principal components were identified which approximately correspond to days of the week, years and months
Singhal et al. (2014)	To explore the relationship between weather conditions and subway usage based on hourly and daily flows	Subway Automated Fare Collection system data	Weather (by day of the week)	Regression models	The impact of the weather depended on the time of day and the location More variation in total ridership was observed at the weekend, where weather had a greater impact Some differences were observed in the weather variables which were significant in the hourly versus the daily

					usage model
Tao et al. (2014)	To compare usage of the Bus Rapid Transit system compared to other buses in the same city based on temporal and spatial variations	Smart card	Weekday, weekend and national holidays	Descriptive analyses and geo-visualisation	Spatial analyses of school holiday days showed that they were more similar to working days than weekends or national holidays Usage was lower on weekends and national holidays
Schmöller et al. (2015)	To examine the temporal and spatial patterns in usage of a car-sharing scheme	Car share scheme booking data	Day of the week and weather	Descriptive statistics Principal Component Analysis	Bookings were highest on Fridays and Saturdays There was a spatial asymmetry between supply and demand on Mondays
Calvert et al. (2016)	To test for day type specific variations in road capacity	Multiple loop detectors	Work days, weekend days and national holidays	Levene's test and the t-test for equality of means	Motorway breakdown capacity was lower on weekend days than work days Breakdown capacity was not lower on holiday days than on work days

Table 1-2: Literature on traveller-focused empirical analyses of variability

Paper	Mode	Data type	Length of time	Number of participants	Variability type	Method	Relevant findings
Muthyalagari et al. (2001)	Car	GPS in cars of participants plus demographic and trip specific survey data	Approx. 6 days	100 cars	Trip frequency, travel time and distance, first departure and final departure and arrival time per day	Descriptive statistics	All characteristics showed relatively low levels of repetition between days Higher levels of stability were observed in the final arrival time at home each day
Schlich and Axhausen (2003)	All	Travel diary (Mobidrive)	6 weeks	361 people from 162 households	Combined intrapersonal variability	Comparison of different measures of similarity	“travel is neither totally repetitious nor totally variable” (Schlich and Axhausen, 2003, p13)
Li et al. (2004)	Car	GPS in cars	7 days	56 cars	Departure time (for morning commute only) and route choice	Descriptive statistics Chi Square tests	60% of the commuters made at least one stop on their way to work on at least one survey day
Kitamura et al. (2006)	All	Travel diary (Mobidrive)	6 weeks	116 people	Prism vertex location for commuters	Stochastic frontier models and least-squares models ANOVA	Individuals’ time-space prism origin vertices vary more systematically, whereas their departure times vary more randomly

							Suggests that there is high levels of flexibility in morning commutes
Bayarma et al. (2007)	All	Travel diary (a sample from Mobidrive)	6 weeks (weekdays only)	317 people from 139 households	Trip frequency, duration and purpose, time expenditure, mode, spatial and situational factors (including group size and expenditure)	Principal Component Analysis to reduce the dimensionality of the data K-means clustering to identify representative daily patterns across all data Markov Chain models to examine transitions between daily patterns	Heterogeneity is observed in the daily patterns utilised by each traveller during the survey period The transition probabilities between daily travel patterns also differ greatly between individuals
Elango et al. (2007)	Car	GPS in cars of participants plus demographic data	3 years	153 vehicles in 98 households	Trip frequency	Mann–Whitney U-test and visual comparisons	Household income and the number of people and vehicles in the household had a significant impact on the amount of intra-household variability in the number of trips per day
Morency et al. (2007)	Public transport (bus)	Smart card	277 days	7,118 smart cards	Spatial, time of day and trip frequency	Descriptive statistics	On average, approximately 0.7 of a new bus stop is used by each traveller each

							week Zero-boarding days are not evenly distributed by the day of the week. Adult-interzone card holders rarely use the bus on the weekend, whereas senior card holders are more likely to be observed at the weekend
Buliung et al. (2008)	All	Panel survey including travel diary	7 days	262 households including 416 adults	Spatial	Descriptive statistics and visualisation (using a Minimum Convex Polygon metric)	Trips are less variable in terms of destination choice than they are by time of the day High levels of spatial re-use were observed, particularly for public transport users Activities cover a wider geographical area on weekdays than weekend days
Chikaraishi et al. (2009)	All	Travel diary (Mobidrive)	6 weeks	361 people from 162 households	Departure time	Multi-level modelling	For almost all trip purposes, intrapersonal variations contributed

							more to the variability in departure times than inter-individual, inter-household, spatial (OD) variations or aggregate level temporal variations
Spissu et al. (2011)	Car	GPS on smartphones of participants, together with online verification and augmentation of data	2 weeks	12 people	Route choice	Descriptive statistics	More intrapersonal variability was observed in discretionary trips More interpersonal variability was observed for trips to/from work or education 93% of trips made by the same person between the same OD pair used the same route
Stopher and Zhang (2011)	All	Panel survey including GPS data from personal devices	between 7 and 15 days in three waves (2005, 2006 and 2007)	202, 308 and 197 households for waves 1, 2 and 3 respectively	Travel time, distance, activity time and trip timing by trip purpose	Descriptive statistics	Complex series of trips are observed in the data. Even fairly simple sequences, such as home -> work -> home, contributed to intrapersonal variability

Ma et al. (2013)	Public transport (bus, subway)	Smart card	5 days	37,001 smart cards	Trip frequency, timing, route and spatial (OD) variability	K-means clustering of travellers Rough set theory is proposed for larger data sets	The authors suggest that travellers in the High and Very High regularity clusters may be regular users, and these correspond to 41% of travellers
Shen et al. (2013)	All	Personal GPS tracking devices with online verification of data	7 days	96 people	Spatial (OD), time (work start and end times and commute duration), mode and route	Descriptive statistics Example 3D geovisualisations for a sample of travellers	The commuting trips are flexible and complex More variation in the time of the trip is observed than in the other three aspects
Venter and Joubert (2013)	HGV, LGV and car	In-vehicle GPS trackers	3 days	42,000 trucks and light goods vehicles and 720 cars	Vehicle kilometres travelled	Descriptive statistics	The lowest income group had the highest variability in vehicle kilometres travelled per day On average, car users had higher levels of variability than the commercial vehicles
Järv et al. (2014)	All	Mobile phone (call detail records)	12 months	1310 phones	Spatial	Multiple Linkage Analysis to identify activity spaces then General Linear Models to examine monthly	Some seasonality was observed in the size of activity spaces ‘Outlier’ months (atypical behaviour for

						variations	an individual) are observed throughout the year although more occur during the summer
Heinen and Chatterjee (2015)	All	Travel diary and demographics	7 days	14,607 people	Mode choice	Regression analyses	Of the people who made at least one trip, 69% were multi-modal in the week of data collected Females exhibited greater modal variability
Kieu et al. (2015b)	Public transport (bus, train and ferry)	Smart card	4 months (working days only)	Approx. 1 million	Spatial and time of day	Density-based clustering algorithms to identify regular origin and alight stops and to identify habitual trip timings Market segmentation based on binary measures of spatial and temporal variability	Four passenger segments were identified: spatially regular only, temporally regular only, spatially and temporally regular and irregular travellers 64% of the travellers were irregular passengers and 99% of them were infrequent users Many student ticket holders were in the

							spatial regularity only group, perhaps due to more flexible schedules
Minnen et al. (2015)	All	Activity diary	7 days	3,253 people	Trip recurrence (related to frequency) and timing of trips	ANOVA	Only 7.4% of travel had regular tempo and timing Work trips varied more in terms of recurrence (people may always travel at the same time, but may not travel every day), whereas discretional trips varied more in timing
Raux et al. (2016)	All	Travel diary and demographics	7 days	707 people	Trip frequency, time per activity and activity sequence	Descriptive statistics	Socio-demographics had very little influence on intrapersonal variability Systematic variability was only a small component of the intrapersonal variability observed
Xianyu et al. (2017)	All	GPS smartphone app	7 day	46 people	Activity-travel sequences	Panel effects regression models Nested F-test to compare model	Thursdays differed from other weekdays Socio-demographics

						specifications	strongly influenced individual activity-travel sequences Men had lower intrapersonal variability than women

1.3 Research objectives

The aim of this research is to develop robust methodologies which could be used to undertake empirical analyses relating to variability in travel behaviour on the road network using emerging data sources. As described in the previous section, there is a gap in the literature for a network-focused method to test for predictable differences in traveller-aggregated data which accounts for within day dynamics. Suitable methods for analysing spatial intrapersonal variability, taking in to account both OD and route variability, and temporal intrapersonal variability on the road network specifically do not exist in the literature either. A method for classifying all road users based on their repeated trip making behaviour, for example by combining measures of spatial and temporal intrapersonal variability, is also lacking. Previous transportation research has also not explored whether a multiple method type approach (Davis et al., 2010) could be used to obtain additional insights by undertaking both network-focused and traveller-focused analyses concurrently.

This research aims to fill these gaps by utilising emerging data sources, which, as described above, are likely to provide data for longer periods of time and may provide new kinds of evidence of variability. The overarching research question defining this thesis is as follows.

> How can emerging data sources be used to gain insights into variability in travel behaviour on the road network?

The underlying research objectives are:

1) To develop a methodology for identifying statistically significant predictable differences in aggregated travel behaviour observed on the network which takes into account differences in magnitude and timing.

2) To identify methods for measuring spatial and temporal intrapersonal variability in travel behaviour on the road network using data from emerging data sources.
3) To develop a methodology for comparing and/or classifying road users based on the intrapersonal variability in their travel behaviour.
4) To develop a framework for using network- and traveller-focused analyses together to gain additional insights into variability in travel behaviour.
5) To apply the methods to real world data in order to demonstrate the insights which can be achieved.

1.4 Thesis overview

The remainder of the thesis is organised as follows.

In Chapter 2, a network-focused method which tests for predictably different traffic flows is presented. The method utilises data from loop detectors which are commonly used to collect data on the road network. Recent technological advances mean that large historical databases, often at low levels of aggregation, are now available. In this methodology, day-to-day variability is measured according to both the total daily volume and the distribution of the flows throughout the day. As in Guardiola et al. (2014), a Functional Data Analysis approach is used so that flows can be analysed according to the time of day.

Chapter 3 discusses the feasibility of using Bluetooth data to examine intrapersonal variability on the road network. A significant amount of literature exists relating to the measurement of travel times using fixed Bluetooth detectors (Haseman et al., 2010, Martchouk et al., 2011, Bhaskar and Chung, 2013, Díaz et al., 2016, Mathew et al., 2016). Bluetooth data has also been used to examine route choice (Hainen et al., 2011, Carpenter et al., 2012) and in OD matrix estimation (Barceló et al., 2010, Barceló et al., 2013), but previous analyses have not matched the unique identifiers associated with Bluetooth devices (known as MAC addresses) between days. Data from fixed Bluetooth detectors are an increasingly popular way to measure travel times in urban areas due to the low costs involved and their passive nature. There are, therefore, advantages in examining this data source as it could provide data on intrapersonal variability at very little additional cost for cities which already collect this type of data. This chapter concludes that Bluetooth data can provide valuable insights into intrapersonal variability on the road network, provided that suitable techniques are developed to extract meaningful information and the appropriate caveats are presented.

In Chapter 4, a methodology is proposed for analysing the repeated road travel behaviour of individuals using point to point sensor data such as Bluetooth data. The frequency of travel and both spatial and temporal aspects of trip making are considered. Travellers are then clustered based upon these three characteristics to identify road user classes. Such market segmentation has proved useful in the analysis of smart card data for public transport users (Goulet Langlois et al., 2016, Kieu et al., 2015b). By taking a traveller-focused perspective, the methodology provides insights into traveller flexibility, levels of exploration and knowledge of the network.

Chapter 5 proposes an approach for undertaking network-focused and traveller-focused analyses alongside one another. Using this approach, findings are not only triangulated, but the ‘following a thread’ technique (Moran-Ellis et al., 2016) can be used to identify questions or themes from each type of analysis to explore in more detail with the other type. The case study from Chapter 2 is used, along with an application of the methods presented in Chapter 4 to the same road link.

Chapter 6 discusses the progress which has been made towards achieving the objectives listed in Section 1.3. The original contributions to knowledge of this thesis are also described alongside the limitations of the approach taken and the data used.

1.5 References

Agrawal, S., Zheng, H., Peeta, S. & Kumar, A. (2016) Routing aspects of electric vehicle drivers and their effects on network performance. *Transportation Research Part D: Transport and Environment,* 46**,** 246-266.

Arana, P., Cabezudo, S. & Peñalba, M. (2014) Influence of weather conditions on transit ridership: A statistical study using data from Smartcards. *Transportation Research Part A: Policy and Practice,* 59**,** 1-12.

Bagloee, S. A., Sarvi, M., Patriksson, M. & Rajabifard, A. (2017) A Mixed User-Equilibrium and System-Optimal Traffic Flow for Connected Vehicles Stated as a Complementarity Problem. *Computer-Aided Civil and Infrastructure Engineering,* 32**,** 562-580.

Barceló, J., Montero, L., Bullejos, M., Serch, O. & Carmona, C. (2013) A Kalman Filter Approach for Exploiting Bluetooth Traffic Data When Estimating Time-Dependent OD Matrices. *Journal of Intelligent Transportation Systems,* 17**,** 123-141.

Barceló, J., Montero, L., Marqués, L. & Carmona, C. (2010) Travel Time Forecasting and Dynamic Origin-Destination Estimation for Freeways Based on Bluetooth Traffic Monitoring. *Transportation Research Record: Journal of the Transportation Research Board,* 2175**,** 19-27.

Bates, J., Polak, J., Jones, P. & Cook, A. (2001) The valuation of reliability for personal travel. *Transportation Research Part E: Logistics and Transportation Review,* 37**,** 191-229.

Bayarma, A., Kitamura, R. & Susilo, Y. (2007) Recurrence of Daily Travel Patterns: Stochastic Process Approach to Multiday Travel Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2021**,** 55-63.

Beckmann, M., McGuire, C. B. & Winsten, C. B. (1956) Studies in the economics of transportation. New Haven: Yale University Press.

Ben-Elia, E. & Shiftan, Y. (2010) Which road do I take? A learning-based model of route-choice behavior with real-time information. *Transportation Research Part A: Policy and Practice,* 44**,** 249-264.

Bender, E. A. (1978) *An introduction to mathematical modeling,* New York, John Wiley and Sons.

Bhaskar, A. & Chung, E. (2013) Fundamental understanding on the use of Bluetooth scanner as a complementary transport data. *Transportation Research Part C: Emerging Technologies,* 37**,** 42-72.

Billot, R., El Faouzi, N.-E. & De Vuyst, F. (2009) Multilevel Assessment of the Impact of Rain on Drivers' Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2107**,** 134-142.

Bogers, E. A. I., Bierlaire, M. & Hoogendoorn, S. P. (2007) Modeling learning in route choice. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 1-8.

Brilon, W., Geistefeldt, J. & Regler, M. Reliability of freeway traffic flow: A stochastic concept of capacity. Proceedings of the 16th International Symposium on Transportation and Traffic Theory, 2005 College Park, Maryland. 125-144.

Buliung, R. N., Roorda, M. J. & Remmel, T. K. (2008) Exploring spatial variety in patterns of activity-travel behaviour: initial results from the Toronto Travel-Activity Panel Survey (TTAPS). *Transportation,* 35**,** 697-722.

Calabrese, F., Colonna, M., Lovisolo, P., Parata, D. & Ratti, C. (2011) Real-Time Urban Monitoring Using Cell Phones: A Case Study in Rome. *IEEE Transactions on Intelligent Transportation Systems,* 12**,** 141-151.

Calvert, S. C., Taale, H. & Hoogendoorn, S. P. (2016) Quantification of motorway capacity variation: influence of day type specific variation and capacity drop. *Journal of Advanced Transportation,* 50**,** 570-588.

Carpenter, C., Fowler, M. & Adler, T. (2012) Generating Route-Specific Origin-Destination Tables Using Bluetooth Technology. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 96-102.

Carrion, C. & Levinson, D. (2012) Value of travel time reliability: A review of current evidence. *Transportation Research Part A: Policy and Practice,* 46**,** 720-741.

Cascetta, E. (2009) *Transportation systems analysis: models and applications,* New York, Dordrecht, Heidelberg, London, Springer.

Chikaraishi, M., Fujiwara, A., Zhang, J. & Axhausen, K. (2009) Exploring Variation Properties of Departure Time Choice Behavior by Using Multilevel Analysis Approach. *Transportation Research Record: Journal of the Transportation Research Board,* 2134**,** 10-20.

Clark, S. & Watling, D. (2005) Modelling network travel time reliability under stochastic demand. *Transportation Research Part B: Methodological,* 39**,** 119-140.

Dafermos, S. C. (1972) The Traffic Assignment Problem for Multiclass-User Transportation Networks. *Transportation Science,* 6**,** 73-87.

Daganzo, C. F. & Sheffi, Y. (1977) On Stochastic Models of Traffic Assignment. *Transportation Science,* 11**,** 253-274.

Datla, S. & Sharma, S. (2008) Impact of cold and snow on temporal and spatial variations of highway traffic volumes. *Journal of Transport Geography,* 16**,** 358-372.

Davis, D. F., Golicic, S. L. & Boerstler, C. N. (2010) Benefits and challenges of conducting multiple methods research in marketing. *Journal of the Academy of Marketing Science,* 39**,** 467-479.

Department for Transport (2017) Travel time measures for the Strategic Road Network, England: Year ending March 2017. Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/614583/travel-time-measures-srn-england-year-ending-march-2017.pdf [Accessed: 13 July 2017].

Dial, R. B. (1971) A probabilistic multipath traffic assignment model which obviates path enumeration. *Transportation Research,* 5**,** 83-111.

Díaz, J. J. V., González, A. B. R. & Wilby, M. R. (2016) Bluetooth Traffic Monitoring Systems for Travel Time Estimation on Freeways. *IEEE Transactions on Intelligent Transportation Systems,* 17**,** 123-132.

Dill, J. & Broach, J. (2014) Travel to Common Destinations. *Transportation Research Record: Journal of the Transportation Research Board,* 2413**,** 84-91.

Elango, V., Guensler, R. & Ogle, J. (2007) Day-to-Day Travel Variability in the Commute Atlanta, Georgia, Study. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 39-49.

Federal Highway Administration (FHWA) (2014) The Urban Congestion Report (UCR): Documentation and Definitions. Available: https://ops.fhwa.dot.gov/perf_measurement/ucr/documentation.htm [Accessed: 13 July 2017].

Gao, H. O. & Niemeier, D. A. (2007) The impact of rush hour traffic and mix on the ozone weekend effect in southern California. *Transportation Research Part D: Transport and Environment,* 12**,** 83-98.

Goulet Langlois, G., Koutsopoulos, H. N. & Zhao, J. (2016) Inferring patterns in the multi-week activity sequences of public transport users. *Transportation Research Part C: Emerging Technologies,* 64**,** 1-16.

Guardiola, I. G., Leon, T. & Mallor, F. (2014) A functional approach to monitor and recognize patterns of daily traffic profiles. *Transportation Research Part B: Methodological,* 65**,** 119-136.

Gundlegård, D., Rydergren, C., Breyer, N. & Rajna, B. (2016) Travel demand estimation and network assignment based on cellular network data. *Computer Communications,* 95**,** 29-42.

Hainen, A. M., Wasson, J. S., Hubbard, S. M. L., Remias, S. M., Farnsworth, G. D. & Bullock, D. M. (2011) Estimating Route Choice and Travel Time Reliability with Field Observations of Bluetooth Probe Vehicles. *Transportation Research Record: Journal of the Transportation Research Board,* 2256**,** 43-50.

Han, L., Sun, H., Wang, D. Z. W. & Zhu, C. (2016) A stochastic process traffic assignment model considering stochastic traffic demand. *Transportmetrica B: Transport Dynamics***,** 1-21.

Han, S. (2003) Dynamic traffic modelling and dynamic stochastic user equilibrium assignment for general road networks. *Transportation Research Part B: Methodological,* 37**,** 225-249.

Haseman, R., Wasson, J. & Bullock, D. (2010) Real-Time Measurement of Travel Time Delay in Work Zones and Evaluation Metrics Using Bluetooth Probe Tracking. *Transportation Research Record: Journal of the Transportation Research Board,* 2169**,** 40-53.

Heinen, E. & Chatterjee, K. (2015) The same mode again? An exploration of mode choice variability in Great Britain using the National Travel Survey. *Transportation Research Part A: Policy and Practice,* 78**,** 266-282.

Horowitz, J. L. (1984) The stability of stochastic equilibrium in a two-link transportation network. *Transportation Research Part B: Methodological,* 18**,** 13-28.

Huang, H.-J. & Li, Z.-C. (2007) A multiclass, multicriteria logit-based traffic equilibrium assignment model under ATIS. *European Journal of Operational Research,* 176**,** 1464-1477.

Iida, Y., Akiyama, T. & Uchida, T. (1992) Experimental Analysis of Dynamic Route Choice Behavior. *Transportation Research Part B: Methodological,* 26**,** 17-32.

INRIX (2016) London Congestion Trends: March 2016.

Iqbal, M. S., Choudhury, C. F., Wang, P. & González, M. C. (2014) Development of origin–destination matrices using mobile phone call data. *Transportation Research Part C: Emerging Technologies,* 40**,** 63-74.

Järv, O., Ahas, R. & Witlox, F. (2014) Understanding monthly variability in human activity spaces: A twelve-month study using mobile phone call detail records. *Transportation Research Part C: Emerging Technologies,* 38**,** 122-135.

Jones, P. & Clarke, M. (1988) The significance and measurement of variability in travel behaviour. *Transportation,* 15**,** 65-87.

Kaltenbrunner, A., Meza, R., Grivolla, J., Codina, J. & Banchs, R. (2010) Urban cycles and mobility patterns: Exploring and predicting trends in a bicycle-based public transport system. *Pervasive and Mobile Computing,* 6**,** 455-466.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015a) Empirical modelling of the relationship between bus and car speeds on signalised urban networks. *Transportation Planning and Technology,* 38**,** 465-482.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015b) Passenger Segmentation Using Smart Card Data. *IEEE Transactions on Intelligent Transportation Systems,* 16**,** 1537-1548.

Kim, J., Mahmassani, H., Vovsha, P., Stogios, Y. & Dong, J. (2013) Scenario-Based Approach to Analysis of Travel Time Reliability with Traffic Simulation Models. *Transportation Research Record: Journal of the Transportation Research Board,* 2391**,** 56-68.

Kitamura, R., Yamamoto, T., Susilo, Y. O. & Axhausen, K. W. (2006) How routine is a routine? An analysis of the day-to-day variability in prism vertex location. *Transportation Research Part A: Policy and Practice,* 40**,** 259-279.

Lam, W. H. K., Shao, H. & Sumalee, A. (2008) Modeling impacts of adverse weather conditions on a road network with uncertainties in demand and supply. *Transportation Research Part B: Methodological,* 42**,** 890-910.

Li, H., Guensler, R., Ogle, J. & Wang, J. (2004) Using global positioning system data to understand day-to-day dynamics of morning commute behavior. *Transportation Planning and Analysis 2004.* Washington: Transportation Research Board National Research Council.

Li, M., Rouphail, N. M., Mahmoudi, M., Liu, J. & Zhou, X. (2017) Multi-scenario optimization approach for assessing the impacts of advanced traffic information under realistic stochastic capacity distributions. *Transportation Research Part C: Emerging Technologies,* 77**,** 113-133.

Li, R., Rose, G. & Sarvi, M. (2006) Using Automatic Vehicle Identification Data to Gain Insight into Travel Time Variability and Its Causes. *Transportation Research Record: Journal of the Transportation Research Board,* 1945**,** 24-32.

Liu, K., Yamamoto, T. & Morikawa, T. (2009) Feasibility of Using Taxi Dispatch System as Probes for Collecting Traffic Information. *Journal of Intelligent Transportation Systems,* 13**,** 16-27.

Liu, Y., Kang, C., Gao, S., Xiao, Y. & Tian, Y. (2012) Understanding intra-urban trip patterns from taxi trajectory data. *Journal of Geographical Systems,* 14**,** 463-483.

Liu, Y., Li, Y. & Hu, L. Departure Time and Route Choices with Bottleneck Congestion: User Equilibrium under Risk and Ambiguity. 22nd International Symposium on Transportation and Traffic Theory, 2017 Chicago, USA. 571-590.

Liu, Z. & Sharma, S. (2006) Statistical Investigations of Statutory Holiday Effects on Traffic Volumes. *Transportation Research Record: Journal of the Transportation Research Board,* 1945**,** 40-48.

Lo, H. K., Luo, X. W. & Siu, B. W. Y. (2006) Degradable transport network: Travel time budget of travelers with heterogeneous risk aversion. *Transportation Research Part B: Methodological,* 40**,** 792-806.

Lo, H. K. & Tung, Y.-K. (2003) Network with degradable links: capacity analysis and design. *Transportation Research Part B: Methodological,* 37**,** 345-363.

Lou, X.-M., Cheng, L. & Chu, Z.-M. (2016) Modelling travellers' en-route path switching in a day-to-day dynamical system. *Transportmetrica B: Transport Dynamics,* 5**,** 17-41.

Ma, X., Wu, Y.-J., Wang, Y., Chen, F. & Liu, J. (2013) Mining smart card data for transit riders' travel patterns. *Transportation Research Part C: Emerging Technologies,* 36**,** 1-12.

Mahmassani, H. S., Chang, G.-L. & Herman, R. (1986) Individual Decisions and Collective Effects in a Simulated Traffic System. *Transportation Science,* 20**,** 258-271.

Martchouk, M., Mannering, F. & Bullock, D. (2011) Analysis of Freeway Travel Time Variability Using Bluetooth Detection. *Journal of Transportation Engineering,* 137**,** 697-704.

Mathew, J. K., Devi, V. L., Bullock, D. M. & Sharma, A. (2016) Investigation of the Use of Bluetooth Sensors for Travel Time Studies under Indian Conditions. *Transportation Research Procedia,* 17**,** 213-222.

Minnen, J., Glorieux, I. & van Tienoven, T. P. (2015) Transportation habits: Evidence from time diary data. *Transportation Research Part A: Policy and Practice,* 76**,** 25-37.

Miralinaghi, M., Lou, Y., Hsu, Y.-T., Shabanpour, R. & Shafahi, Y. (2016) Multiclass fuzzy user equilibrium with endogenous membership functions and risk-taking behaviors. *Journal of Advanced Transportation,* 50**,** 1716-1734.

Miranda-Moreno, L. & Nosal, T. (2011) Weather or Not to Cycle. *Transportation Research Record: Journal of the Transportation Research Board,* 2247**,** 42-52.

Moran-Ellis, J., Alexander, V. D., Cronin, A., Dickinson, M., Fielding, J., Sleney, J. & Thomas, H. (2016) Triangulation and integration: processes, claims and implications. *Qualitative Research,* 6**,** 45-59.

Morency, C., Trépanier, M. & Agard, B. (2007) Measuring transit use variability with smart-card data. *Transport Policy,* 14**,** 193-203.

Muthyalagari, G. R., Parashar, A. & Pendyala, R. M. Measuring Day-to-Day Variability in Travel Characteristics Using GPS Data. 80th Annual Meeting of the Transportation Research Board, 2001 Washington D.C.: TRB.

Nakayama, S. & Watling, D. (2014) Consistent formulation of network equilibrium with stochastic flows. *Transportation Research Part B: Methodological,* 66**,** 50-69.

Ortúzar, J. d. D. & Willumsen, L. G. (2011) *Modelling Transport,* Chichester, UK, John Wiley & Sons Ltd.

Paz, A., Emaasit, D. & de la Fuente, H. Stochastic Dynamic User Equilibrium Using a Mixed Logit Modeling Framework. IEEE 19th International Conference on Intelligent Transportation Systems (ITSC), 2016 Rio de Janeiro, Brazil. 1206-1211.

Peeta, S. & Ziliaskopoulos, A. K. (2001) Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. *Networks and Spatial Economics,* 1**,** 233-265.

Pelletier, M.-P., Trépanier, M. & Morency, C. (2011) Smart card data use in public transit: A literature review. *Transportation Research Part C: Emerging Technologies,* 19**,** 557-568.

Prashker, J. N. & Bekhor, S. (2004) Route Choice Models Used in the Stochastic User Equilibrium Problem: A Review. *Transport Reviews,* 24**,** 437-463.

Pu, W. (2011) Analytic Relationships Between Travel Time Reliability Measures. *Transportation Research Record: Journal of the Transportation Research Board,* 2254**,** 122-130.

Rakha, H. & Van Aerde, M. (1995) Statistical analysis of day-to-day variations in real-time traffic flow data. *Transportation Research Record: Journal of the Transportation Research Board,* 1510**,** 26-34.

Raux, C., Ma, T.-Y. & Cornelis, E. (2016) Variability in daily activity-travel patterns: the case of a one-week travel diary. *European Transport Research Review,* 8.

Schlich, R. & Axhausen, K. W. (2003) Habitual travel behaviour: Evidence from a six-week travel diary. *Transportation,* 30**,** 13-36.

Schmöller, S., Weikl, S., Müller, J. & Bogenberger, K. (2015) Empirical analysis of free-floating carsharing usage: The Munich and Berlin case. *Transportation Research Part C: Emerging Technologies,* 56**,** 34-51.

Shahhoseini, Z., Haghani, M. & Sarvi, M. (2015) Estimation and application of a multi-class multi-criteria mixed paired combinatorial logit model for transport networks analysis. *Transportmetrica B: Transport Dynamics,* 3**,** 59-78.

Shang, W., Han, K., Ochieng, W. & Angeloudis, P. (2016) Agent-based day-to-day traffic network model with information percolation. *Transportmetrica A: Transport Science,* 13**,** 38-66.

Shao, H., Lam, W. H. K. & Tam, M. L. (2006) A reliability-based stochastic traffic assignment model for network with multiple user classes under uncertainty in demand. *Networks & Spatial Economics,* 6**,** 173-204.

Sheffi, Y. (1985) *Urban Transportation Networks,* Englewood Cliffs, Prentice-Hall Inc.

Shen, Y., Kwan, M.-P. & Chai, Y. (2013) Investigating commuting flexibility with GPS data and 3D geovisualization: a case study of Beijing, China. *Journal of Transport Geography,* 32**,** 1-11.

Singhal, A., Kamga, C. & Yazici, A. (2014) Impact of weather on urban transit ridership. *Transportation Research Part A: Policy and Practice,* 69**,** 379-391.

Smith, M. J. (1979) The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological,* 13**,** 295-304.

Spissu, E., Meloni, I. & Sanjust, B. (2011) Behavioral Analysis of Choice of Daily Route with Data from Global Positioning System. *Transportation Research Record: Journal of the Transportation Research Board,* 2230**,** 96-103.

Stathopoulos, A. & Karlaftis, M. (2001) Temporal and Spatial Variations of Real-Time Traffic Data in Urban Areas. *Transportation Research Record: Journal of the Transportation Research Board,* 1768**,** 135-140.

Stopher, P. R. & Zhang, Y. (2011) Repetitiveness of Daily Travel. *Transportation Research Record: Journal of the Transportation Research Board,* 2230**,** 75-84.

Sun, C., Cheng, L., Zhu, S. & Chu, Z. (2015) Multiclass Stochastic User Equilibrium Model with Elastic Demand. *Transportation Research Record: Journal of the Transportation Research Board,* 2497**,** 1-11.

Szeto, W. Y., Jiang, Y. & Sumalee, A. (2011) A Cell-Based Model for Multi-class Doubly Stochastic Dynamic Traffic Assignment. *Computer-Aided Civil and Infrastructure Engineering,* 26**,** 595-611.

Tao, S., Corcoran, J., Mateo-Babiano, I. & Rohde, D. (2014) Exploring Bus Rapid Transit passenger travel behaviour using big data. *Applied Geography,* 53**,** 90-104.

Taylor, M. A. P. (2013) Travel through time: the story of research on travel time reliability. *Transportmetrica B-Transport Dynamics,* 1**,** 174-194.

Venter, C. & Joubert, J. (2013) Use of Multisource Global Positioning System Data to Characterize Multiday Driving Patterns and Fuel Usage in a Large Urban Region. *Transportation Research Record: Journal of the Transportation Research Board,* 2338**,** 1-10.

Wardrop, J. G. (1952) Some theoretical aspects of road traffic research. *Road Engineering Division Meeting.*

Watling, D. (2002) A Second Order Stochastic Network Equilibrium Model, I: Theoretical Foundation. *Transportation Science,* 36**,** 149-166.

Watling, D., Milne, D. & Clark, S. (2012) Network impacts of a road capacity reduction: Empirical analysis and model predictions. *Transportation Research Part A: Policy and Practice,* 46**,** 167-189.

Wei, C., Asakura, Y. & Iryo, T. (2014) Formulating the within-day dynamic stochastic traffic assignment problem from a Bayesian perspective. *Transportation Research Part B: Methodological,* 59**,** 45-57.

Weijermars, W. & van Berkum, E. Analyzing highway flow patterns using cluster analysis. 8th International IEEE Conference on Intelligent Transportation Systems, 2005 Vienna, Austria. 831-836.

Xianyu, J., Rasouli, S. & Timmermans, H. (2017) Analysis of variability in multi-day GPS imputed activity-travel diaries using multi-dimensional sequence alignment and panel effects regression models. *Transportation,* 44**,** 533-553.

Xiao, Y. & Lo, H. K. (2016) Day-to-day departure time modeling under social network influence. *Transportation Research Part B: Methodological,* 92**,** 54-72.

Xu, M., Meng, Q. & Liu, K. (2017) Network user equilibrium problems for the mixed battery electric vehicles and gasoline vehicles subject to battery swapping stations and road grade constraints. *Transportation Research Part B: Methodological,* 99**,** 138-166.

Yang, L., Kwan, M.-P., Pan, X., Wan, B. & Zhou, S. (2017) Scalable space-time trajectory cube for path-finding: A study using big taxi trajectory data. *Transportation Research Part B: Methodological,* 101**,** 1-27.

Yazici, M., Kamga, C. & Mouskos, K. (2012) Analysis of Travel Time Reliability in New York City Based on Day-of-Week and Time-of-Day Periods. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 83-95.

Yeon, J., Hernandez, S. & Elefteriadou, L. (2009) Differences in Freeway Capacity by Day of the Week, Time of Day, and Segment Type. *Journal of Transportation Engineering,* 135**,** 416-426.

Yu, W., Park, S., Kim, D. S. & Ko, S.-S. (2015) Arterial Road Incident Detection Based on Time-Moving Average Method in Bluetooth-Based Wireless Vehicle Reidentification System. *Journal of Transportation Engineering,* 141**,** 04014084 1-14.

Zhang, K., Mahmassani, H. S. & Lu, C.-C. (2013) Dynamic pricing, heterogeneous users and perception error: Probit-based bi-criterion dynamic stochastic user equilibrium assignment. *Transportation Research Part C: Emerging Technologies,* 27**,** 189-204.

Zhang, W., Mediria, A. & Rakha, H. Statistical analysis of spatiotemporal link and path flow variability. 2007 IEEE Intelligent Transportation Systems Conference, 2007 Seattle, Washington, USA. IEEE, 1080-1085.

2 A statistical method for estimating predictable differences between daily traffic flow profiles

Abstract

It is well known that traffic flows in road networks may vary not only within the day but also between days. Existing models including day-to-day variability usually represent all variability as unpredictable fluctuations. In reality, however, some of the differences in flows on a road may be predictable for transport planners with access to historical data. For example, flow profiles may be systematically different on Mondays compared to Fridays due to predictable differences in underlying activity patterns. By identifying days of the week or times of year where flows are predictably different, models can be developed or model inputs can be amended (in the case of day-to-day dynamical models) to test the robustness of proposed policies or to inform the development of policies which vary according to these predictably different day types. Such policies could include time-of-day varying congestion charges that themselves vary by day of the week or season, or targeting public transport provision so that timetables are more responsive to the day of the week and seasonal needs of travellers. A statistical approach is presented for identifying systematic variations in daily traffic flow profiles based on known explanatory factors such as the day of the week and the season. In order to examine day–to-day variability whilst also considering within-day dynamics, the distribution of flows throughout a day are analysed using Functional Linear Models. F-type tests for functional data are then used to compare alternative model specifications for the predictable variability. The output of the method is an average flow profile for each predictably different day type, which could include day of the week or time of year. An application to real-life traffic flow data for a two-year

period is provided. The shape of the daily profile was found to be significantly different for each day of the week, including differences in the timing and width of peak flows and also the relationship between peak and inter-peak flows. Seasonal differences in flow profiles were also identified for each day of the week.

2.1 Introduction

There is now extensive literature analysing and modelling the extent to which traffic flows systematically vary within a day, due to time-of-day variations in demand, time-of-day variations in capacity (e.g. due to traffic signals), and the temporal and spatial interactions of congestion (Ukkusuri et al., 2012, Du et al., 2015, Han et al., 2015, Long et al., 2016, Ngoduy et al., 2016, Wang and Du, 2016). A corresponding body of work has additionally sought to address the considerable variation observed in traffic flows between days, known as day-to-day variability (Watling and Cantarella, 2013a, Watling and Cantarella, 2013b, Guo et al., 2015, Hazelton and Parry, 2016, Kumar and Peeta, 2015, Xiao et al., 2016). This twin focus, on within-day and day-to-day variation, is the topic of the present paper. Existing models including day-to-day variability usually represent variability by a single probability distribution for each randomly varying component[1], for example in research on demand (Watling, 2002, Clark and Watling, 2005, Sumalee et al., 2006, Shao et al., 2006, Nakayama and Watling, 2014), capacity (Lo and Tung, 2003, Siu and Lo, 2008, Sumalee et al., 2011b) or travel times (Noland and Small, 1995, Clark and Watling, 2005, Pu, 2011, Guo et al., 2012). In contrast to existing models, this research proposes separating the predictable and unpredictable

[1] Whilst Ettema et al. (2005) did develop a day-to-day dynamical model where previously experienced travel times were stored in separate categories within memory, they did not discuss how the categories could be formed or used in practice.

components of variability so that variability could be represented by a set of probability distributions alongside a set of rules specifying which distribution relates to which day type. In this paper ‘day type’ relates to an exhaustive classification based on combining characteristics which would be known far in advance such as the day of the week or season. This classification could be anything from a simple weekday/weekend day split to a complex combination of days of the week and months. This research develops a stochastic model of within-day profiles that includes day types as explanatory factors to identify predictably different day types in a dataset. The functional day type coefficients are estimated, but the development of full probability distributions for models with both within-day and day-to-day dynamics is left for future work. The outputs of the method presented below are also useful in their own right as they can be used by practitioners to better understand travel patterns and perhaps inform day type specific policies in order to better utilise resources. The flow profiles for each day type can also be used to test the robustness of policies and, more importantly, as the day types are known far in advance plans can be made to mitigate potential problems.

The seminal work of Hanson and Huff (1988) provided evidence that cyclical patterns exist in individual travel behaviour. These patterns or cycles are often based around the days of the week or seasons, as evidenced by multi-day surveys (Kitamura and Van Der Hoorn, 1987, Schlich and Axhausen, 2003, Habib and Miller, 2008). More recently, data from emerging data sources have been used to examine activity patterns over a longer period of time, for example Järv et al. (2014) who used mobile phone data to examine monthly variations in activity spaces. Whilst in many cases such patterns will disappear once data has been aggregated, in some cases exogenous factors can cause systematic patterns in travel behaviour which translates into predictably different travel conditions. This could include the widely accepted weekday patterns of peak and off-peak traffic, but also patterns which are more likely to be overlooked, for example lower network demand in winter months or variations in daily flow profiles due to differences in shop opening hours between weekdays. Researchers often try to avoid the component of variability which is predictable by considering "some nominally 'typical' conditions" (Clark and Watling, 2005, p119), such as the 'peak' period of the day on 'non-holiday weekdays' only. Although some researchers have explored the impact of various types of predictable variability when undertaking analysis of flow or travel time data (Rakha and Van Aerde, 1995, Stathopoulos and Karlaftis, 2001, Zhang et al., 2007, Yazici et al., 2012), few have built on this to develop predictive models.

Two exceptions are Kamga and Yazici (2014) and Guardiola et al. (2014). Kamga and Yazici (2014) used GPS data from taxis to classify average travel times per unit distance across the city for each hour of the day and day of the week using regression trees. Guardiola et al. (2014) used Functional Data Analysis on traffic flow profiles from a detector set on a freeway for the purpose of classification and outlier detection. They used Functional Principal Component Analysis to identify the three principal components. The first appears to separate working from non-working days, the second may relate to the year and the third may be a seasonal factor.

The present paper builds on the work of Guardiola et al. (2014) as it also considers variability in daily traffic flow profiles, but it differs by considering day types which would be known in advance, such as the day of the week, rather than using data-driven category selection. Traffic flow data is perhaps the most widely collected data on road networks and therefore there are vast amounts of data to analyse, even when considering just one day type. Flow data informs us about what is occurring on the road network at any given point in time and therefore is particularly of interest to practitioners and those calibrating models.

This research proposes a method for identifying scenarios where traffic flow profiles differ predictably, as a result of characteristics which would be known in advance, such as the day of the week. Day type explanatory models will be estimated for both the magnitude and shape of the daily flow profiles. Data from one loop detector will be considered. The flow at a single location could vary due to many factors, including demand, route choice, departure time and the traffic conditions on other parts of the network. In this research the aim is to identify the day types or seasons where flows at this location are systematically different to those on other days, regardless of the cause, to inform scenario testing. For a single location, this is relevant for modelling localised policies such as capacity reductions (for example due to road works or changes in parking regulations) or congestion charging boundaries. The methodology could also be applied independently to many detectors in an area in order to identify days of the week or times of year where particular problems arise so that policy solutions to target the causes of these problems can be devised. Examples of such targeted policies might include: time-of-day varying congestion charges that themselves vary by day of the week or season; incentives to influence employers or shopping centres to adjust their opening times by day of the week or season; targeting public transport provision so that timetables are more responsive to the day of the week and seasonal needs of travellers.

This paper makes an original contribution by presenting:

1) A method for estimating the effect on daily traffic flow profiles of predictable variability due to known explanatory factors, for example the day of the week.
2) A method for comparing alternative model specifications for the predictable variability through statistical significance.
3) An application of these methods to real-life traffic flow data for a two-year period.

A transferable methodology is presented for identifying predictable variability in both total daily flows and standardised daily flow profiles. The analysis of daily flow profiles involves Functional Linear Regression Models which have not been used for this purpose previously. Previous use of Functional Data Analysis in transportation (for example Guardiola et al. (2014) and Chiou et al. (2014)) has been restricted to Functional Principal Component Analysis, where the components that are predictably different can be identified, but a future day cannot be assigned (a priori) to a component group. In this paper we seek to measure the impact of *known* explanatory variables and hence directly examine the effect of the day of the week or season.

The structure of this paper is as follows. Section 2.2 includes the rationale for using the analytical technique adopted in this paper, namely Functional Data Analysis. Section 2.3 provides a detailed methodology, including a description of a technique to estimate smooth functions from point data and a method to analyse functional data, which in this case will utilise Functional Linear Models. Section 2.4 includes an application to real-life traffic flow data from one site on a commuter route within a large urban area in northwest England. Section 2.5 details the potential for future work relating to predictable variability and the use of Functional Data Analysis in transportation research.

2.2 Variability in daily flow profiles

2.2.1 Requirements

The purpose of this research is to identify day types with significantly different flows so that relevant transport policies can be identified and tested. The method used should also produce sufficiently detailed traffic flow profiles to use as inputs to within-day dynamical models. When comparing daily flow profiles, Weijermars and van Berkum (2005, p832) stress that "it is important to take both differences in shape and in height into account". This is relevant for the exploration of day types with systematically different flow profiles as the separation of differences in the overall volume of traffic (the magnitude) from the shape of the profile could provide indirect information about the underlying cause of the variation. For example, differences in the shape of the profile could be due to constraints such as shop opening hours or hours of daylight. Differences in the magnitude of flows, however, are more likely to be dictated by the demand for activities, for example the total flow on commuter routes may vary based on the time of year. As well as understanding the causes of the systematic differences, exploring whether differences are due to the magnitude or the shape of the profiles can also assist in the formation of policies. If the shape of the flow profile differs between day types, it may be relevant to consider whether time of day varying congestion charging, car parking prices or public transport provision should be tailored to the different day types. Different types of policy may be required for day types with higher magnitudes of flow, for example working with local employers, public transport providers and encouraging changes in route, as was seen for a very short period of time in London during the 2012 Olympic Games (Transport for London, 2013). The overall magnitude and shapes of the profiles are, however, inextricably linked and the two aspects should be considered together to fully understand the patterns observed. A more concentrated peak period, for example, may be a cause for concern if the total daily flow is high, but not if it coincides with a low total daily flow.

Analysis of the magnitude and shape of daily traffic profiles can be undertaken using discrete time periods, as demonstrated in Weijermars and van Berkum (2005). However, Habib et al. (2009, p641) state that "attempting to force time into a discrete framework is inherently limiting, often requiring unrealistic simplifying assumptions … or highly complex models". This is also true for this analysis, where time periods could be assumed to be independent, or a complex model including correlations between adjacent and non-adjacent time periods could be applied. Even a complex model using discrete time periods would have limitations, however, as the arbitrary borders would still indicate, for example, that 8:14 am is more similar to 8:01 than 8:16 when using 15 minute intervals (Habib et al., 2009).

Treating the time of day as a continuous variable would allow "a maximum exploitation of the recorded data" (Guardiola et al., 2014, p133). This could allow a more detailed examination of the timing and widths of peak periods and would also provide a more suitable input for a within-day dynamical model. Information would not be lost in an aggregation process but the resulting profile may contain too much variation so that overall trends are hard to identify. As with the choice of interval width in discrete time analyses, a suitable technique would need to be identified which can retain an appropriate amount of detail. For the current application it is therefore most suitable to represent the time of the day by a continuous variable, but days should be treated as discrete observations.

For use in practice, the method would need to be fairly quick and easy to apply and generate outputs which can be easily interpreted. Also, for scenario testing there must be a way of testing whether a day type variable has enough of an impact on the flow profiles to warrant an additional model run.

In summary, then, the key requirements are:

a) to provide as much indirect evidence as possible on the cause of the differences by considering the magnitude and shape of the flow profiles separately where possible,
b) to consider the time of the day as a continuous variable,
c) to take into account correlations between times of the day,
d) to have a robust way of identifying key features in the profiles, and
e) to have a way of testing the statistical significance of day type variables.

2.2.2 Choice of technique

There is not an obvious technique to use to examine the shape of daily flow profiles. The majority of the research aimed at predicting traffic flow profiles focuses on short term forecasting for real-time applications (see Vlahogianni et al. (2004) and Vlahogianni et al. (2014) for a review of methods applied). Vlahogianni et al. (2014) define short term forecasting as a process using both past and current data to estimate the traffic conditions for a time period either seconds or hours in the future. As the current paper is focused on identifying a typical daily flow profile on a specific day type using past data only, alternative methods are required.

More relevant methods relate to the identification of patterns in past daily profiles only. One option would be to represent the flow profiles as time series, i.e. a sequential set of data points for each profile, for example the average trend or Principal Component Analysis methods described in Li et al. (2015). Other time series based approaches include Jiang and Adeli (2005) who present a ‘time-delay recurrent wavelet neural network model’ which they propose could be used for predicting traffic flow profiles on future days. In this approach, all past data is considered as one long time series and therefore day types would need to be identifiable by a fixed lag. As day types such as public holidays could not be considered using this approach and the information provided by neural network approaches may not aid explanations, this approach is not suitable for the current research. Tang et al. (2014) use a complex network approach to identify patterns in daily traffic flow profiles. In this method, the day types are constructed through an examination of the data and periodicities observed, rather than being specified externally. Whilst all of these methods have their advantages, they are not suitable for the current analysis as they do not satisfy requirement b).

Time of day has been considered as a continuous variable in Functional Data Analysis which has been used by Guardiola et al. (2014) and Chiou et al. (2014) to identify patterns and subsequently classify daily flow profiles into groups. Of the methods previously used in transportation research, Functional Data Analysis has the most potential for the current research as it does not depend on fixed lags between day types, is relatively easy to apply, and provides outputs which are easy to interpret. Also, Li et al. (2015) describe the four main areas in traffic time series analysis as detecting abnormalities, data compression, imputation of missing data and prediction. Functional Data Analysis has already been in transportation for outlier detection (Guardiola et al., 2014) and for missing data imputation (Chiou et al., 2014), and therefore using it for prediction allows us to build on, and perhaps in the future integrate with, existing research.

Functional Data Analysis (FDA) encompasses a broad range of techniques for analysing data where each observation is a curve as opposed to a single point. FDA techniques have been used in many disciplines, including ergonomics (Faraway, 1997), oceanology (Nerini and Ghattas, 2007) and risk response (Lee et al., 2009). Ramsay and Silverman (1997, p8) highlight the use of FDA “to study the important sources of pattern and variation among the data”. In many cases it is used for the former, to classify curves into groups, for example in Ferraty and Vieu (2003), Nerini and Ghattas (2007) and Guardiola et al. (2014). There is an extensive literature surrounding the issues raised by requirement d) and a commonly used method, using a roughness penalty, will be described in Section 2.3.1.

By considering each daily flow profile as an observation, correlations between times of the day are taken into account and, therefore, requirements b) and c) are satisfied. FDA has the added advantage that the times at which the measurements were taken do not need to coincide. This would be relevant if an alternative data source (for example manual count data) was available for a particular day which was crucial to include in the analysis, but was aggregated at a higher level than the rest of the data. Such data could be included in the analysis without having to further aggregate the remainder of the data, although care would need to be taken in interpreting the results. The relationship between time periods can be explored in greater detail by examining the first derivative (for the rate of change) and higher order derivatives of the functions at different points in time. Another advantage of using FDA is the ability to separate timing and intensity, thus making it suitable for the exploration of the magnitude, timing and length of the peak periods in daily flow profiles.

FDA has only occasionally been used in transportation research despite the abundance of time series data. Gao and Niemeier (2008) used FDA to examine ozone and NO_x concentrations throughout the day with the aim of informing transport policies, but the Functional Principal Component Analysis (FPCA) did not include traffic data. Chen and Müller (2014) included an application to transport data in their paper on conditional distributions for functional data. The functions in their work correspond to speed profiles of individual vehicles over a fixed section of road and a method based on FPCA was used to produce prediction regions for the average speed in future time periods. Research using FDA to look at daily traffic flow profiles is limited, to the author’s knowledge, to Guardiola et al. (2014) and Chiou et al. (2014). Guardiola et al. (2014) used FPCA, applied to daily flow profiles for the real-time detection of flows deviating from the expected profile. Chiou et al. (2014) used FPCA for the detection of outliers and the imputation of missing sections of daily profiles. FPCA is a useful technique for identifying patterns in the data, but Ramsay et al. (2009, p100) highlight that “it tends to happen that only the leading eigenfunction has an obvious meaningful interpretation”. While this statement is debateable, particularly since Guardiola et al. (2014) had meaningful interpretations for their three principal components, it is true that the principal components are not *guaranteed* to relate to day type characteristics that are known in advance. It is therefore not suitable for the current research where the aim is to predict the daily profile for any future day where only the date is known for certain.

Therefore, whilst FDA is a promising approach for the examination of daily flow profiles, the most commonly used technique, Functional Principal Component Analysis, is not suitable for the current research. This research will therefore use a different tool within FDA, namely Functional Linear Models, which is an extension to standard linear regression modelling, satisfies requirement e) and will be described in detail in Section 2.3.2. While FDA can account for differences in the magnitudes and shapes of profiles, it would provide greater insight into the causes of differences, and satisfy requirement a), if these two aspects were analysed independently. ANOVA is an obvious option for examining total daily flow as it compares specified groups within the data and determines whether the means differ. In order to apply ANOVA techniques, however, certain assumptions must hold and this will be discussed in more detail in Section 2.3).

2.3 Methodology

Figure 2-1 provides an overview of all of the stages which make up the proposed method. Total daily flows can be analysed using standard ANOVA methods, provided certain assumptions hold. The total daily flow needs to be normally distributed within each group and the population variance within each group should be equal. The observations should also be independent. Data should be tested for normality and for homogeneity of variances before ANOVA is undertaken. Where these assumptions do not hold, alternatives exist such as data transformation or non-parametric tests (for example the Kruskal-Wallis test used in Section 2.4). The independence assumption is required not just for ANOVA but also for nonparametric alternatives. To examine total daily flows, adjustments may need to be applied in order to remove seasonal and longer term trends, assuming a multiplicative relationship. This should provide a more stable basis for the day of the week testing. After applying a suitable test, a plot of the residuals can be examined to determine whether the independence assumption appears to be reasonable.

The process for producing a model based on the shapes of flow profiles is not a standard technique used in transportation research and therefore is described in detail in this section. The flow data is assumed to be available aggregated into time bands, although it should be noted that the narrower the bands, the more detailed the corresponding profiles will be. To remove the magnitude effects from the profiles (as these will be analysed separately), the aggregated flows $u_{ij,}$ for time period j on day i, need to be standardised using ${y_{ij}}^{obs} = u_{ij} / \sum_j u_{ij}$. The process of converting these points into daily flow profiles, $y_i(t)$ for day i, will be described in Section 2.3.1. The analysis of these profiles using Functional Linear Models will be described in Section 2.3.2.

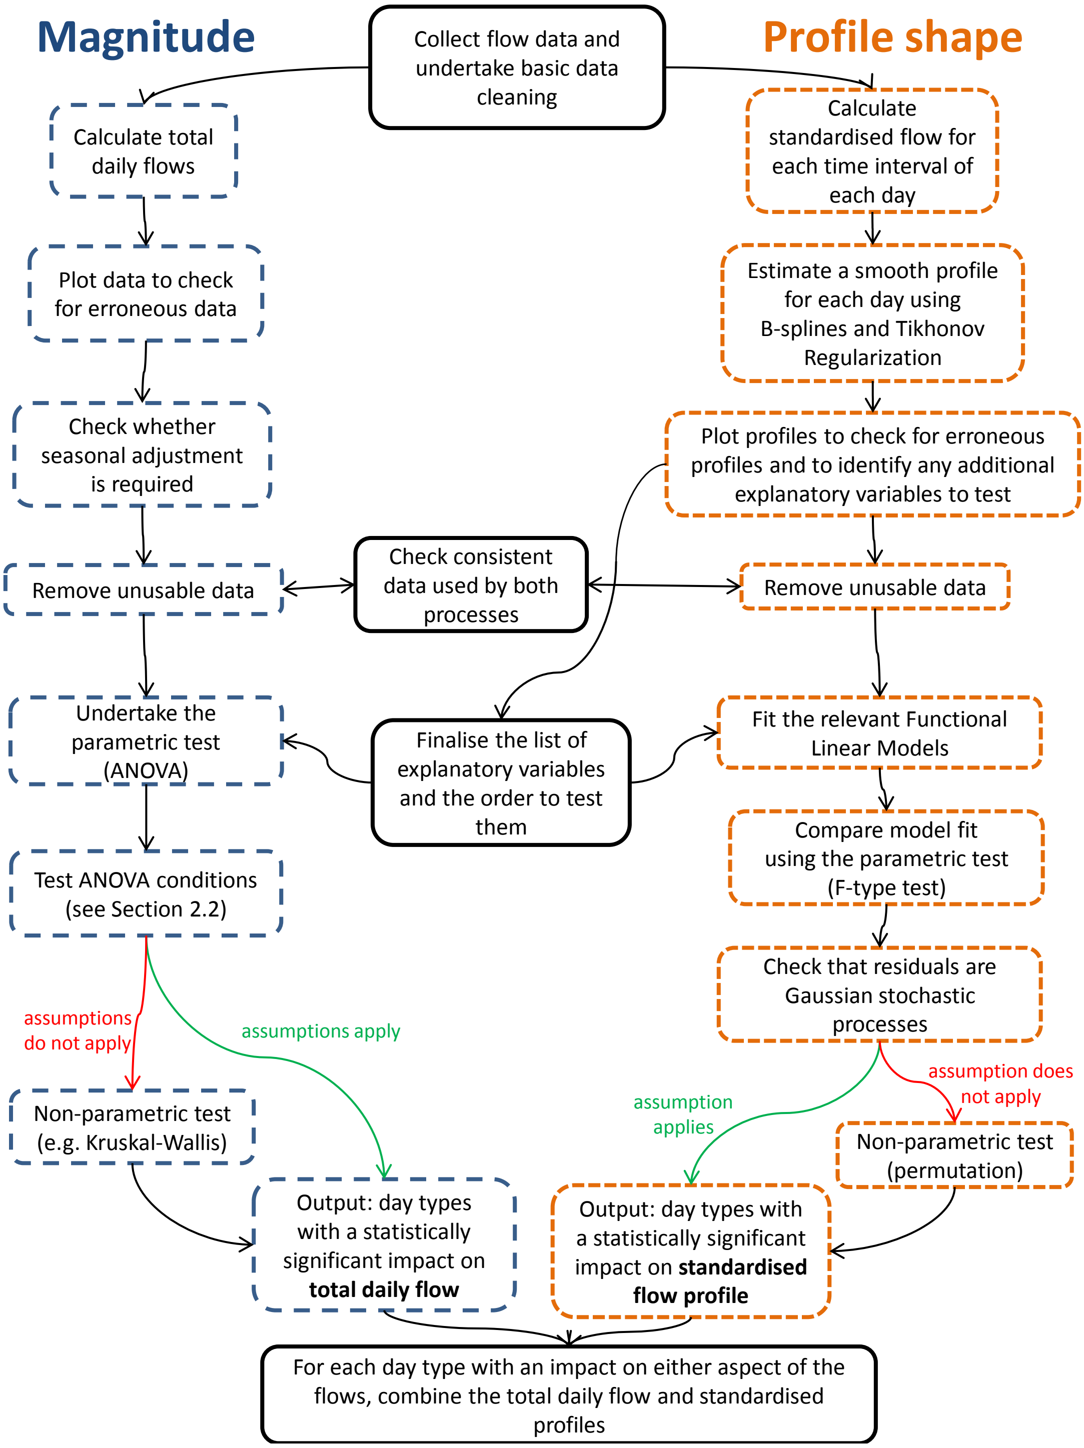

Figure 2-1: Overview of methodology to clean, fit appropriate models and apply statistical significance tests to traffic flow data to produce representative flow profiles for each significantly different day type

2.3.1 Estimating daily flow profiles using B-splines

Although there are many ways in which time series data can be expressed, for the approach used in this paper each daily flow profile needs to be represented by a smooth curve. The success of the functional regression will, to a large extent, depend on the quality of this curve estimation process. The simplest option would be to assume that the profiles form a 'family' of distributions whose parameters vary from day to day (as in Watling et al. (2004, p45 onwards) in relation to travel times). In this research it is important to retain key features in the data such as the start and end times of peak periods and the gradient of the profile through the peak and therefore greater flexibility is required. A commonly used approach is to represent each function as a linear combination of functions, these component functions known as basis functions. Using a linear combination of basis functions is a flexible approach with computational advantages (Ramsay and Silverman, 2005).

The basis function should be chosen so as to best represent the key features of the data. For example, a Fourier series basis would not be an obvious choice in this case as it is a periodic function. Whilst other bases could reasonably be used, for example a wavelet basis, in this paper B-splines will be used. B-splines are often used as de Boor (2001, p95) has shown that all polynomials of order p can be represented as linear combinations of B-splines of that order. This is particularly important given the distinctive M-shaped nature of daily flow profiles. B-splines are piecewise polynomials which are specifically designed so that they have continuous $(p-1)^{th}$ order derivatives (where p is the degree of the polynomial used), even where the pieces join. This means that cubic B-spline representations of the daily flow profiles would have continuous second order derivatives. Each B-spline is piecewise polynomial with compact support, i.e. it is non-zero on one small section of the estimation interval only (Ramsay and Silverman, 1997, p49). This is also an attractive property of B-splines, as demonstrated in Figure 2-2, as the overlapping produces functional estimates which are better able to represent local features in the data than methods which consider the entire interval at once, for example Fourier or polynomial bases (Ramsay and Silverman, 1997, p48). Using B-splines, and cubic B-splines in particular, also makes the curve estimation process consistent with the method used on daily traffic flow profiles in Guardiola et al. (2014).

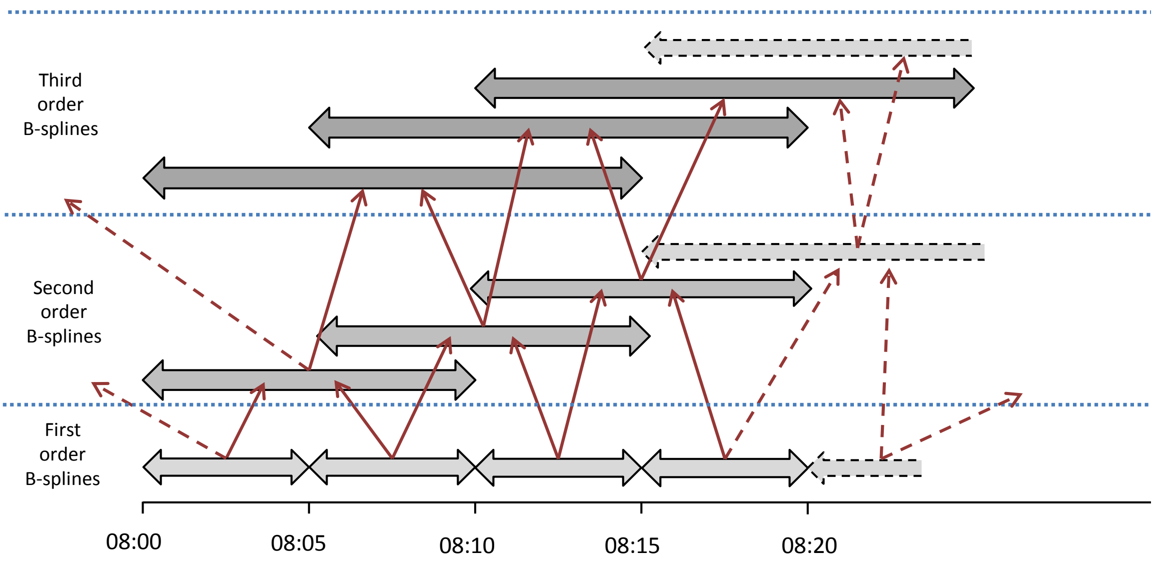

Figure 2-2: B-spline coverage example

The standard process for constructing B-spline basis functions (de Boor, 2001, p89) is used in this research. Knots separate the interval over which $y_i(t)$ is to be estimated into subintervals called knot spans, which in Figure 2-2 are five minute intervals. In this paper, only uniform knot spacing with knots coinciding with data points has been considered. This is to ensure consistency across all functions estimated and because a roughness penalty will be included in the estimation process (see below). The process starts with the construction of a first order B-spline for each knot span as shown in equation (1):

$$B_{k,1}(t) = \begin{cases} 1 & if \quad t_k \le t \le t_{k+1} \\ 0 & otherwise \end{cases} \tag{1}$$

where $B_{k,1}(t)$ is the B-spline with knot *k* as the defining point (on the left hand side) and of order 1. t_k is the time of day relating to the start of the k^{th} subinterval. The first order B-splines, $B_{k,1}(t)$, are simply step functions taking the value 1 in [t_k, t_{k+1}] and 0 elsewhere.

The Cox-de Boor recursion relationship (de Boor, 2001, p90) shown in equation (2) can then be used to define B-splines of higher orders.

$$B_{k,q}(t) = \omega_{k,q}B_{k,q-1}(t) + \left(1 - \omega_{k+1,q}\right)B_{k+1,q-1}(t) \tag{2}$$

where:

$$\omega_{k,q} = \frac{(t - t_k)}{\left(t_{k+q-1} - t_k\right)}$$

Once the B-splines have been generated, the coefficients in the linear relationship need to be estimated based on the standardised flows, ${y_{ij}}^{obs}$, to produce suitable estimates of the daily flow profiles. Even though the daily flow profiles are estimated using cubic B-splines which guarantee continuous second order derivatives, they may still be "'rough' or 'wiggly'" (Green and Silverman, 1994, p4) if estimated using a least squares estimation process. In order to find the optimal balance between capturing local features in the daily profiles and retaining excessive noise, a roughness penalty can be added within the least squares estimation process used to estimate the daily flow profiles from the count data (Ramsay and Silverman, 1997, Chapter 4). This is Tikhonov Regularization applied to functional data. The roughness penalty, denoted by λ, can take many forms in FDA, but the integrated squared second derivative is commonly used (Silverman, 1985, Green and Silverman, 1994, Ramsay and Silverman, 1997), resulting in the following formula for the penalized residual sum of squares to be minimised:

$$PENSSE = \sum_j \left\{ {y_{ij}}^{obs} - y_i(t_j) \right\}^2 + \lambda \int \{ {y_i}''(x) \}^2 \, dx \tag{3}$$

where ${y_{ij}}^{obs}$ is the standardised observed flow at time j on day i, and $y_i(t_j)$ is the estimate of ${y_{ij}}^{obs}$ using the estimated flow profile. Fitting based on PENSSE restricts what functions can be fitted. Those functions with large overall second derivatives will be penalized (by a factor of lambda). Lambda is the roughness penalty which represents the weight given to the fit to the data relative to the weight given to the 'smoothness' of the estimate. Figure 2-3 demonstrates different flow profiles which could be estimated using the same five minute flow data from one day, using different values for the roughness penalty. The choice of lambda should not be made arbitrarily, but should be estimated using the data. This can be done using cross-validation. As in Ramsay and Silverman (2005, p97), the cross-validation score for a given value of λ is the sum of squared errors based on a large sample of estimates of the functional regression coefficients using the leave-one-out method. By calculating the cross-validation scores for a range of suitable values for λ, the λ with the lowest score can be selected to use in the analysis. The suitability of the λ selected using the cross-validation method will be discussed further in Section 2.4.2.2.

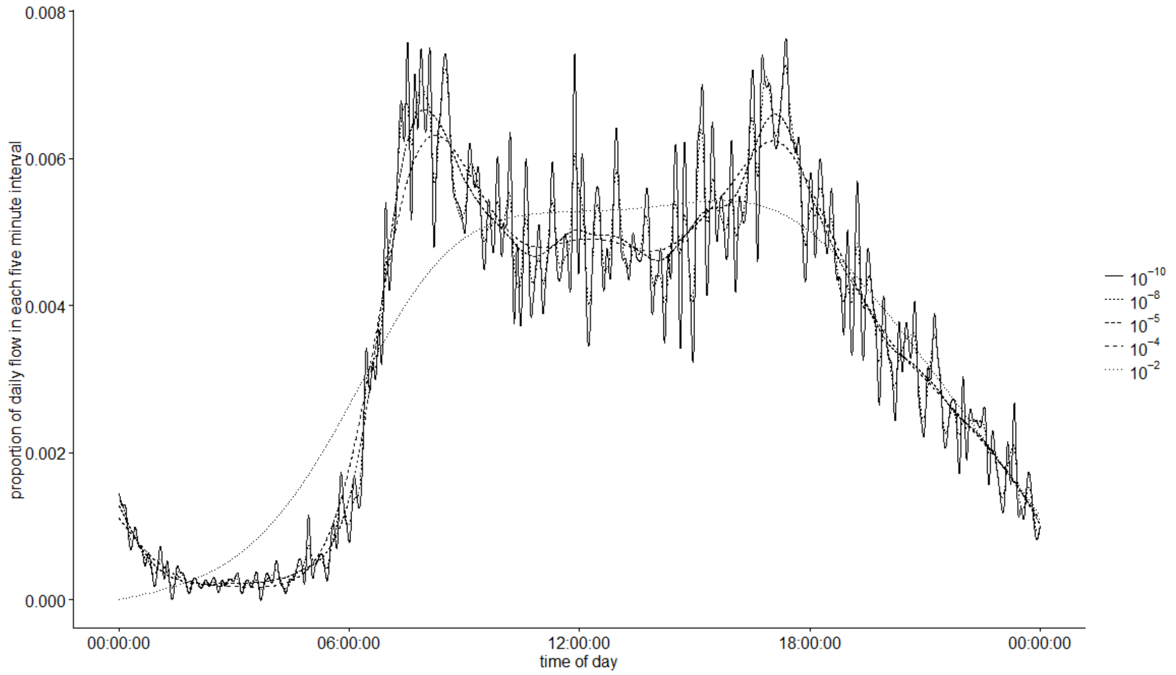

Figure 2-3: Estimated daily flow profiles for one day for a site in Greater Manchester with different roughness penalties applied

2.3.2 Analysing functional data

Whilst Figure 2-3 demonstrates the standardised flow profiles which could be estimated from flow data from one day (using different values of lambda), this process needs to be repeated for each day within the dataset. Figure 2-4 shows the ${y_{ij}}^{obs}$ (i.e. the fifteen minute aggregated flows) on the left hand side and the associated $y_i(t)$ (the estimated flow profiles) on the right hand side, for one month of data from the case study site in Greater Manchester. The value of lambda used for every day was 10^{-6}, as determined by a cross-validation procedure. Although the morning weekday flows follow a similar pattern, the heights of the peaks, between-peak and evening flows vary greatly. Figure 2-4 emphasises that despite the use of a roughness penalty, the functions on the right hand side are ‘messy’ individual observations which will require additional processes to analyse. Even from this one month of data, it is clear that just using an ‘average’ profile for policy testing is likely to be insufficient if the policy impact is sensitive to small changes in the magnitude or shape of daily flow profiles.

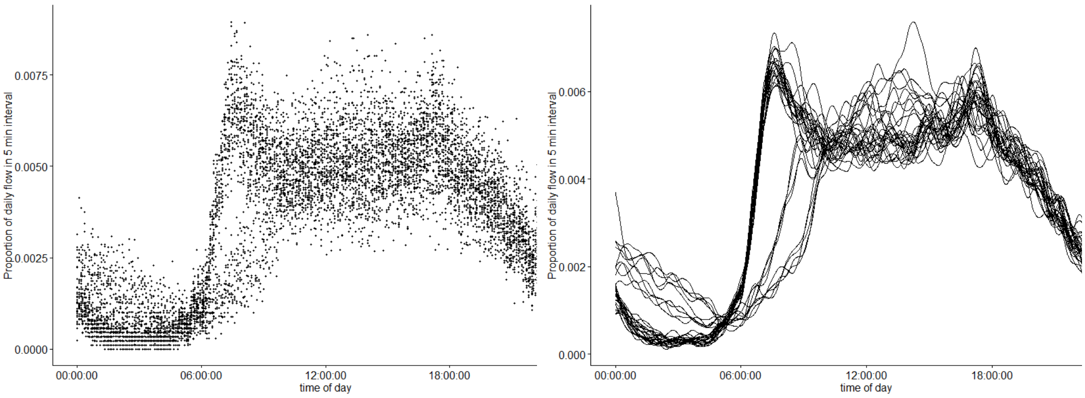

Figure 2-4: One month of standardised traffic flows represented by point data (left side) and by estimated functional observations (right side)

2.3.2.1 Functional linear models with functional responses

The relationships between day type identifiers, such as the day of the week, and daily flow profiles can be explored using an extension to linear modelling where the responses are functions, proposed by Ramsay and Silverman (1997). A Functional Linear Model with a functional response has the following structure:

$$\boldsymbol{y}(t) = \boldsymbol{Z\beta}(t) + \boldsymbol{\varepsilon}(\mathrm{t}) \qquad t\epsilon\mathbb{R} \tag{4}$$

Or in alternative notation:

$$\begin{pmatrix} y_1(t) \\ y_2(t) \\ y_3(t) \\ \vdots \\ y_{n-1}(t) \\ y_n(t) \end{pmatrix} = \begin{pmatrix} z_{1,1} & & z_{1,m-1} & z_{1,m} \\ z_{2,1} & \cdots & z_{2,m-1} & z_{2,m} \\ z_{3,1} & & z_{3,m-1} & z_{3,m} \\ \vdots & \ddots & \vdots & \vdots \\ z_{n-1,1} & \cdots & z_{n-1,m-1} & z_{n-1,m} \\ z_{n,1} & & z_{n,m-1} & z_{n,m} \end{pmatrix} \begin{pmatrix} \beta_1(t) \\ \beta_2(t) \\ \beta_3(t) \\ \vdots \\ \beta_{m-1}(t) \\ \beta_m(t) \end{pmatrix} + \begin{pmatrix} \varepsilon_1(t) \\ \varepsilon_2(t) \\ \varepsilon_3(t) \\ \vdots \\ \varepsilon_{n-1}(t) \\ \varepsilon_n(t) \end{pmatrix} \tag{5}$$

Here, $\boldsymbol{y}(t)$ is a vector of *functional* responses with respect to continuous time, t, which in this case would be the daily flow profiles. ***Z*** is a design matrix consisting of entries $z_{k,l}$, which are 1 if day k is of type l and 0 otherwise. $\boldsymbol{\beta}(t)$ is, therefore, a vector of functional coefficients. $\boldsymbol{\varepsilon}(t)$ is a vector of functional residuals which represent the unexplained variability after the day type variables have been taken into account. These residuals are assumed to be statistically independent (Ramsay et al., 2009, p60).

As the focus here is on predictable variability, only information which could be known far in advance will be used. The indicators in the design matrix will therefore relate to the day of the week and season only. The day of the week was selected as there is a growing body of evidence of the impact on the road network, even between different weekdays (Rakha and Van Aerde, 1995, Weijermars and van Berkum, 2004, Zhang et al., 2007, Ozbay et al., 2014). Similarly, there is evidence of seasonal trends in traffic flows (May, 1990, Stathopoulos and Karlaftis, 2001). The ***Z*** matrix is, therefore, easy to construct, but the $\boldsymbol{\beta}(t)$ need to be estimated by seeking to minimise the residuals (as $\boldsymbol{y}(t)$ is known).

2.3.2.2 Fitting the model

In the current research the aim is not to estimate $\boldsymbol{\beta}(t)$ for all possible day types, but to identify the *most important* day type variables to include in the model and then estimate the relevant coefficients. As testing all possible combinations of the dummy variables is not viable for larger problems, the standard forward stepwise regression process will be used. In practice this could begin with the structure currently used (for example just including a weekend/weekday split) and then test whether adding any other indicators would increase the explanatory power of the model.

Techniques for simultaneous variable selection and estimation, such as LASSO (Tibshirani, 1996), have been applied in functional analysis contexts (Matsui and Konishi, 2011, Lian, 2013, Mingotti et al., 2013). In the current application, however, a non-automated process has been used so that expert knowledge could be used to specify the hierarchy of variables to consider. In the example presented in Section 2.4 the separation of weekdays and weekend days will be considered first, then individual days of the week, and then seasons. The order of the hierarchy matters because each stage of testing is undertaken separately for each significantly different subset identified in the previous stage.

In order to fit the models, a functional extension to the least squares approach will be used to estimate the day type coefficients, $\boldsymbol{\beta}(t)$. As there is an area between the observed and predicted daily flow profiles, the term to be minimised can be expressed (Ramsay and Silverman, 1997, p141) as:

$$Sum\ of\ squared\ residuals = \sum_{i=1}^{n}\sum_{j=0}^{m}\int\left[y_i(t) - Z_{i,j}\hat{\beta}_j(t)\right]^2 . dt \qquad (6)$$

To be consistent with the least squares approach, F-type tests comparing the fit of nested model specifications will be used to determine the most appropriate model to use. This is equivalent to testing the null hypothesis that the ‘reduced’ model, including fewer predictor variables, is preferable to the ‘full’ model which includes one or more additional predictor variables. Ramsay and Silverman (2005) suggest undertaking F-tests at each x-value, to produce point-wise test statistics for functional data. This would not, however, provide information about the statistical significance of the *functional* model, even if all of the point-wise tests are significant (Górecki and Smaga, 2015).

As an alternative to producing point-wise statistics, tests have been proposed to measure the overall significance of a functional model (see Górecki and Smaga (2015)). The F-type test proposed by Shen and Faraway (2004) for linear models with functional responses will be used as it is relatively easy to apply and Shen and

Faraway (2004, p1256) assert that it "examines important rather than trivial differences between models". The test involves calculating a single value for each model comparison as follows:

$$F_{nest} = \left(\frac{RSS_{red} - RSS_{full}}{RSS_{full}}\right) \times \left(\frac{df_{full}}{df_{red} - df_{full}}\right) \qquad (7)$$

where RSS is the residual sum of squares, df is the degrees of freedom and the subscripts show whether the value refers to the reduced or the full model. The residual sum of squares is calculated by laying a fine grid over the profiles as an approximation for the area between the two curves.

The distribution of F_{nest}, the functional F distribution, can be estimated by the ordinary F distribution with degrees of freedom $\varphi\left(df_{red} - df_{full}\right)$ and $\varphi \times df_{full}$, where φ is the degrees of freedom adjustment factor (Shen and Faraway, 2004, p1246). In practice, the method can be easily applied by laying a fine grid over the functions, as demonstrated in Yang et al. (2007). The degrees of freedom adjustment factor can then be estimated (Shen and Faraway, 2004, p1246) by:

$$\varphi = \frac{[Trace(\boldsymbol{E})]^2}{Trace(\boldsymbol{E}^2)} \qquad (8)$$

After laying a fine grid over the functions, it is straightforward to compile ***E,*** the empirical covariance matrix based on the full model, using the covariances between flows at each time of the day corresponding to the fine grid. This F-type test assumes that the residuals are Gaussian stochastic processes. This assumption will be considered further in Section 2.4.2.2.

2.3.3 Combining the total daily flows and standardised flow profiles

The day types identified using the total daily flows and using the standardised flow profiles should then be considered alongside one another in order to observe systematic differences and consider the implications. The standardised flow profile for each day type should then be scaled up by the relevant total daily flow or flows for use as dynamical model inputs. This final stage is represented by the box at the bottom of Figure 2-1 which combines the outputs from the magnitude and profile shape processes.

2.4 Empirical study

The methodology described in the previous section was applied to data from a loop detector on a key arterial route into Manchester. The road is a single lane urban road connecting Stockport (a large town approximately 6 miles south east of Manchester) to the city of Manchester. Two years of data (from 1/05/2013 to 30/04/2015) was used in the analysis. Data relating to public holidays was removed prior to the analysis as the profiles differed from non-public holiday days and yet they were not a homogenous group with sufficient sample size to include in the model. The B-spline estimation and the Functional Data Analysis were undertaken using the ‘fda’ package in R (Ramsay et al., 2014).

The results will be presented, firstly for the analysis of the total daily flow data and then for the analysis of the daily profiles. Section 2.4.2 then includes a discussion of the smoothing parameters used, treatment of outliers and an analysis of the residuals from the same analysis.

2.4.1 Results

A non-parametric test was required to assess the impact of the day of the week on the total daily flow as there is insufficient evidence to reject the null hypothesis in Levene's test for the homogeneity of variances (p=0.40), but the data is not Gaussian distributed (from visual inspection and Shapiro-Wilk test (p<0.001)). The day of the week had a statistically significant impact (Kruskal-Wallis test statistic 495, p<0.05) on total daily flow for this location, with two exceptions: flows on Thursdays and Fridays were not significantly different and neither were flows on Tuesdays and Wednesdays. As monthly adjustment factors have been applied, a formal test of seasonal differences in total daily flows has not been undertaken.

The shape of the daily flow profiles were considered next. To remove the magnitude effect (which was considered above), the flow profiles were transformed into the percentage of flows through the day relative to the total flow for that day. A smooth daily profile was then estimated for every day of data.

The results of the step-wise regression are presented in Figure 2-5. Each row represents a model formulation with the regression coefficients shown. Green arrows indicate statistically significant F-type test results ($\alpha = 0.05$), i.e. where the full model is preferable to the reduced model.

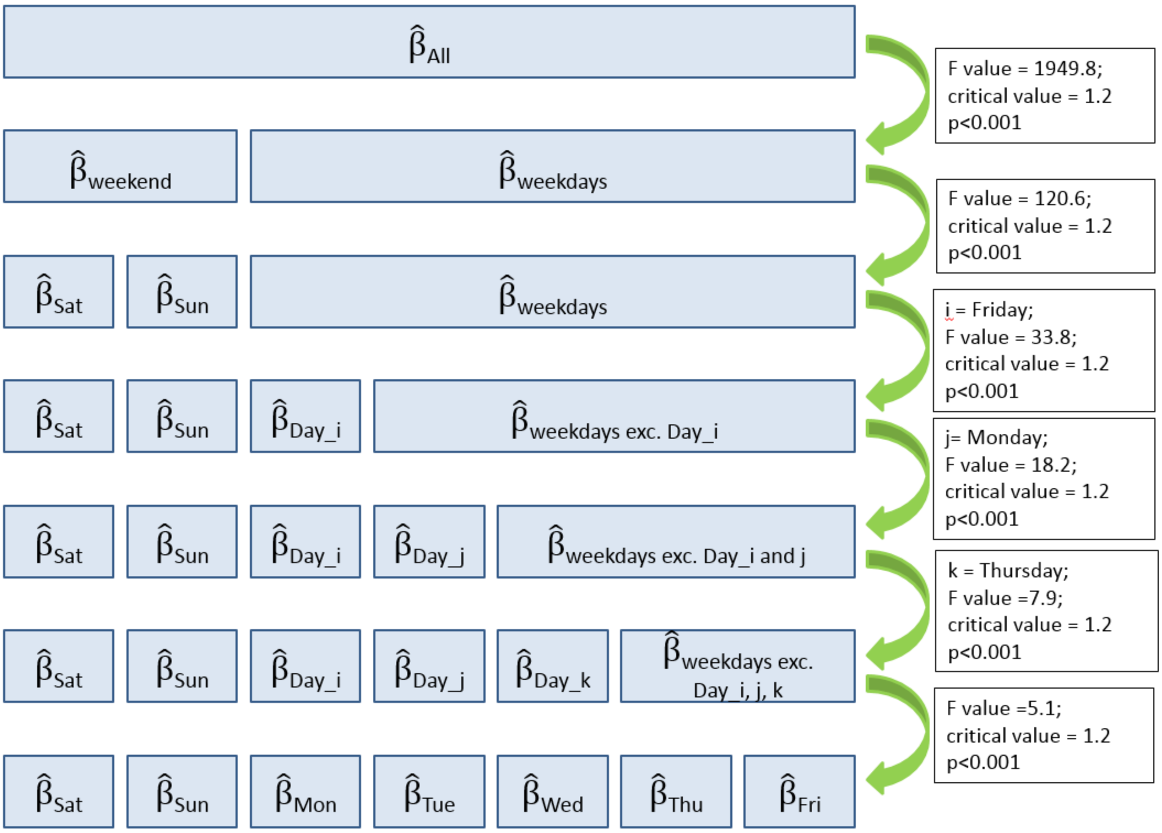

Figure 2-5: Step-wise regression results for the site in Manchester

For the site analysed, all days of the week were identified as having significantly different standardised flow profiles. For this particular site this process is not helpful in reducing the potential number of scenarios to test, but the plots of the coefficients (shown in Figure 2-6 and Figure 2-7) do demonstrate the ways in which the profiles vary. The more concentrated morning peak on a Monday and the variation in the rate of flow decrease in the evening may be of particular interest to explore, depending on the potential policies being explored.

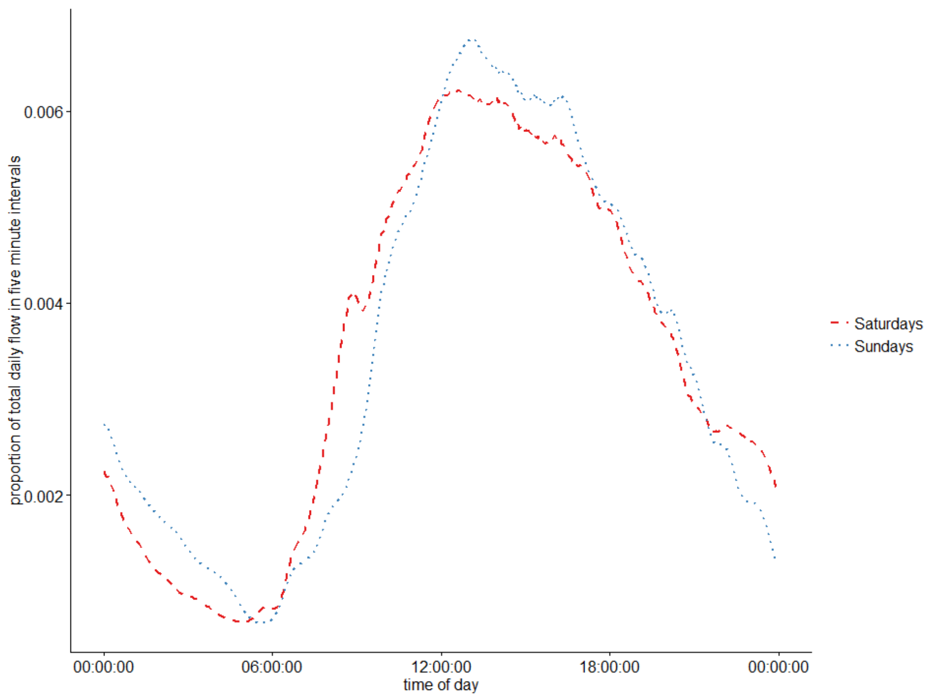

Figure 2-6: Average flow profiles for Saturdays and Sundays

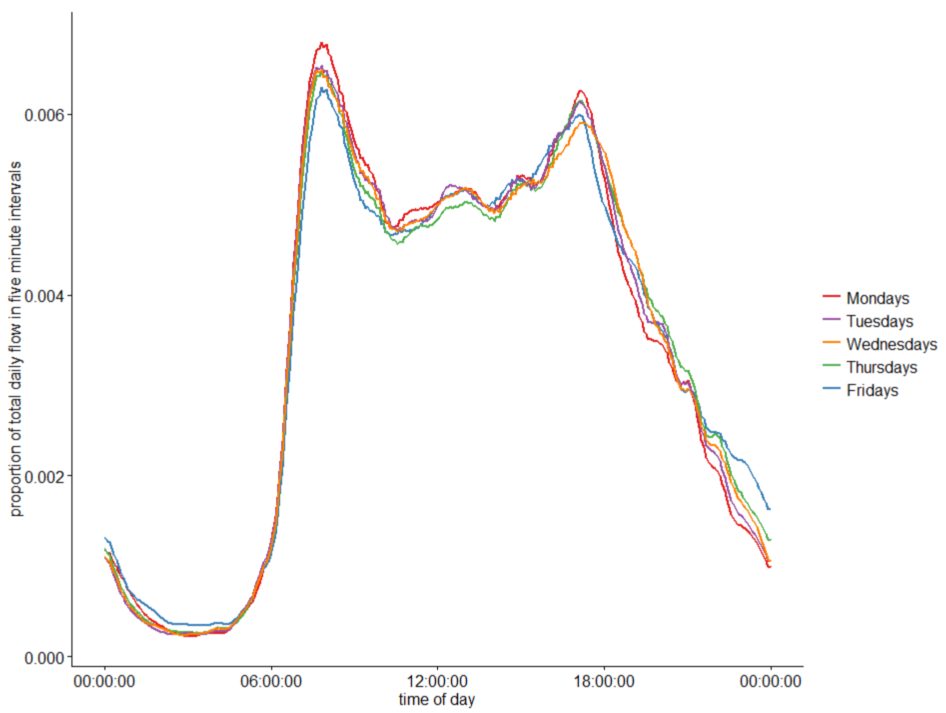

Figure 2-7: Average flow profiles for weekdays

There are many similarities between the average profiles for weekdays shown in Figure 2-7. Although the F-type tests have identified statistically significant differences between the profiles for each day of the week, this does not necessarily mean that all parts of the profiles are significantly different. To explore the differences in more detail, confidence intervals can be plotted around these functional coefficients. The confidence intervals can usually be calculated using the ‘fda’ package in R, but the case study dataset was too large to use this method, even using High Performance Computing, and therefore a suitable approximation was required.

Pointwise confidence intervals were estimated for the average standardised flow at five minute intervals throughout the day for each day of the week. The average flows for each day of the week at each time interval are a rough approximation for the functional coefficients from the FLM. Bootstrapping was used and therefore no distributional assumptions were made. The pointwise 95% confidence intervals for the average standardised flows on Thursdays and Fridays are shown in Figure 2-8. Although the intervals overlap for most of the day, there are times of the day where there is no overlap. The intervals are separate as flows increase before the morning peak, as they decrease after the evening peak (see inset) and at times between the peaks. The fact that the confidence intervals for Thursdays and Fridays do not always overlap is also a reflection of the relatively large amount of data available for this analysis as this has an effect on the interval widths. Other weekday comparisons showed different times of the day at which profiles differ.

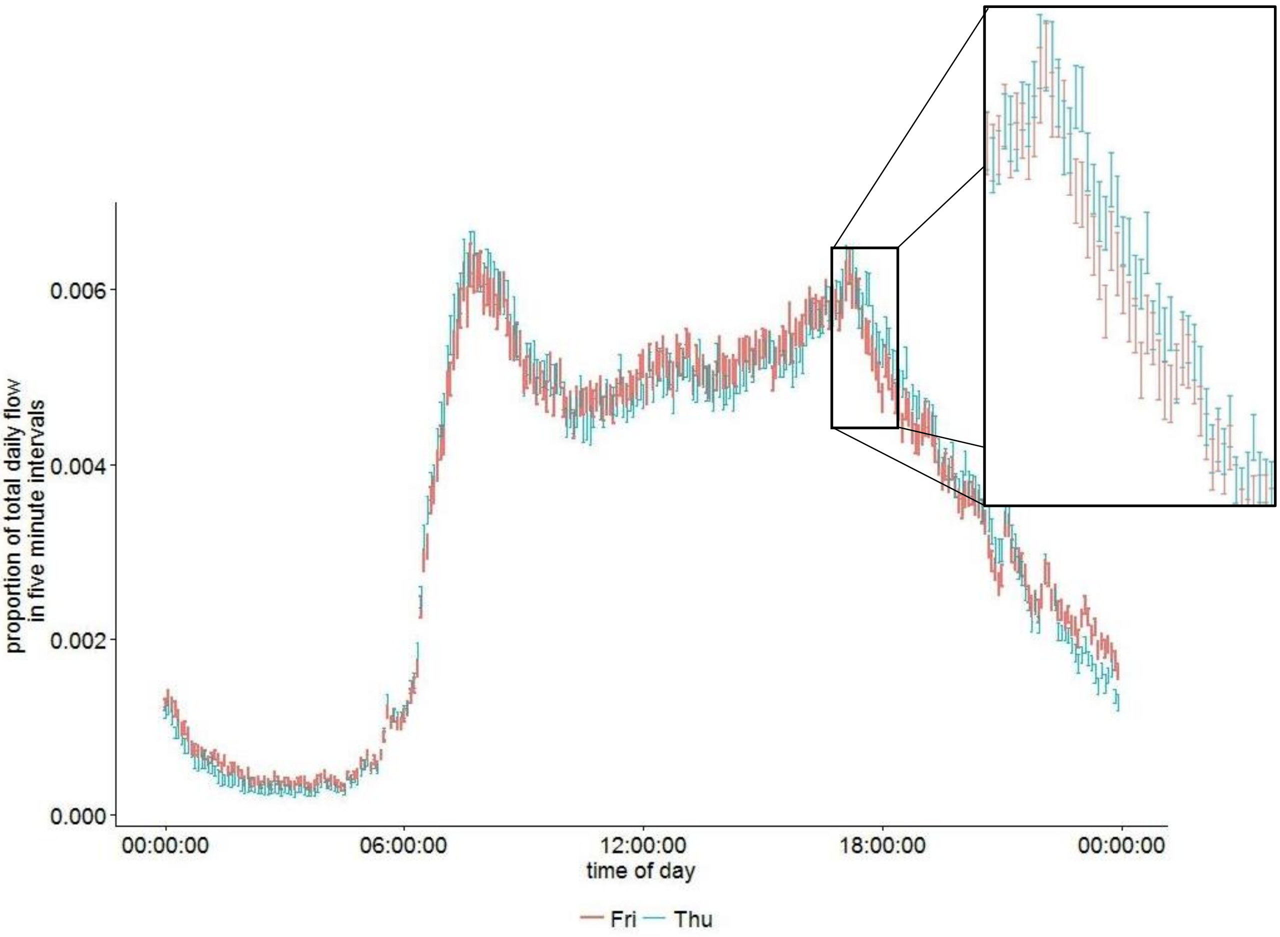

Figure 2-8: Estimated 95% confidence intervals for the standardised flows at each time of day on Thursdays and Fridays

The day of the week may be expected to have an impact on daily flow profiles, but the method can also be used to explore less obvious impacts. Figure 2-9 shows the estimated seasonal coefficients when step-wise regression was undertaken on data from non-holiday Mondays only. In this case, winter and summer have a significant impact on flow profiles, but the profiles in spring and autumn were not statistically significantly different from one another ($\alpha = 0.05$). The winter profile includes a smaller proportion of flows in the morning peak and a larger proportion in the middle of the day, perhaps indicating a higher proportion of non-commuting trips during the holiday season. The seasonal differences are far more pronounced during the morning peak and the inter-peak period than during the evening peak.

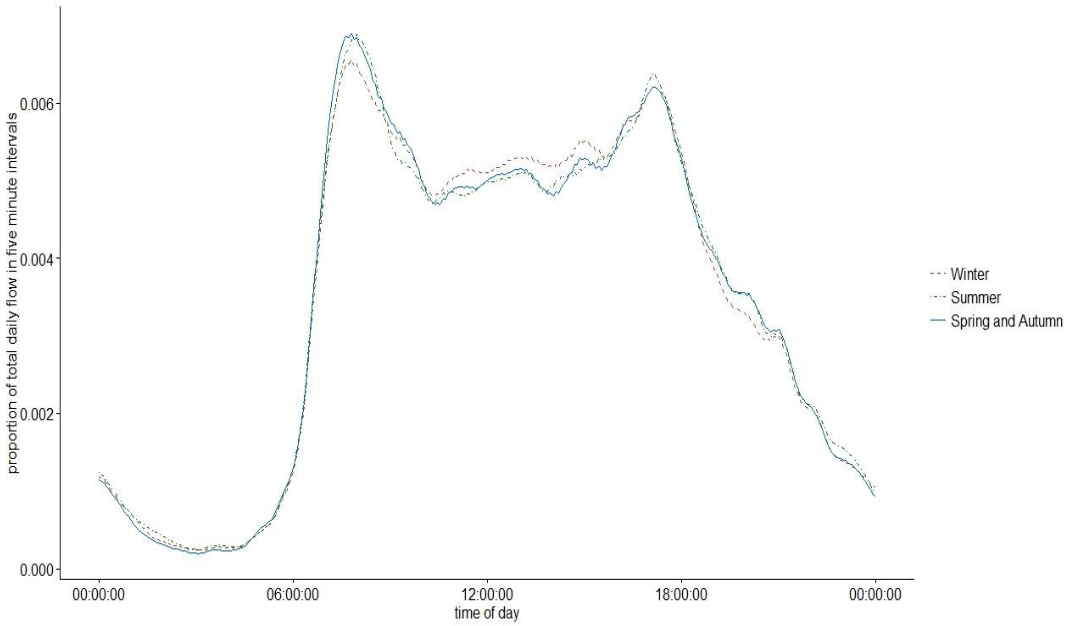

Figure 2-9: Standardised flow profiles by season for non-holiday Mondays

The total daily flows and the flow profiles estimated under the different day types should then be combined to produce a suitable input for within-day dynamical models. Figure 2-10 includes some examples of day types which may be considered for further investigation based on the full results from the analysis above. These have been constructed by combining the estimated total daily flow, using data from the relevant months and days of the week, and regression coefficient (for example $\hat{\beta}_{Mon\_summer}$) for the relevant day type. In this example, the profiles suggest that if the link is close to capacity or a time of day dependent policy is being considered, modelling effects just based on average weekday conditions is unlikely to be sufficient.

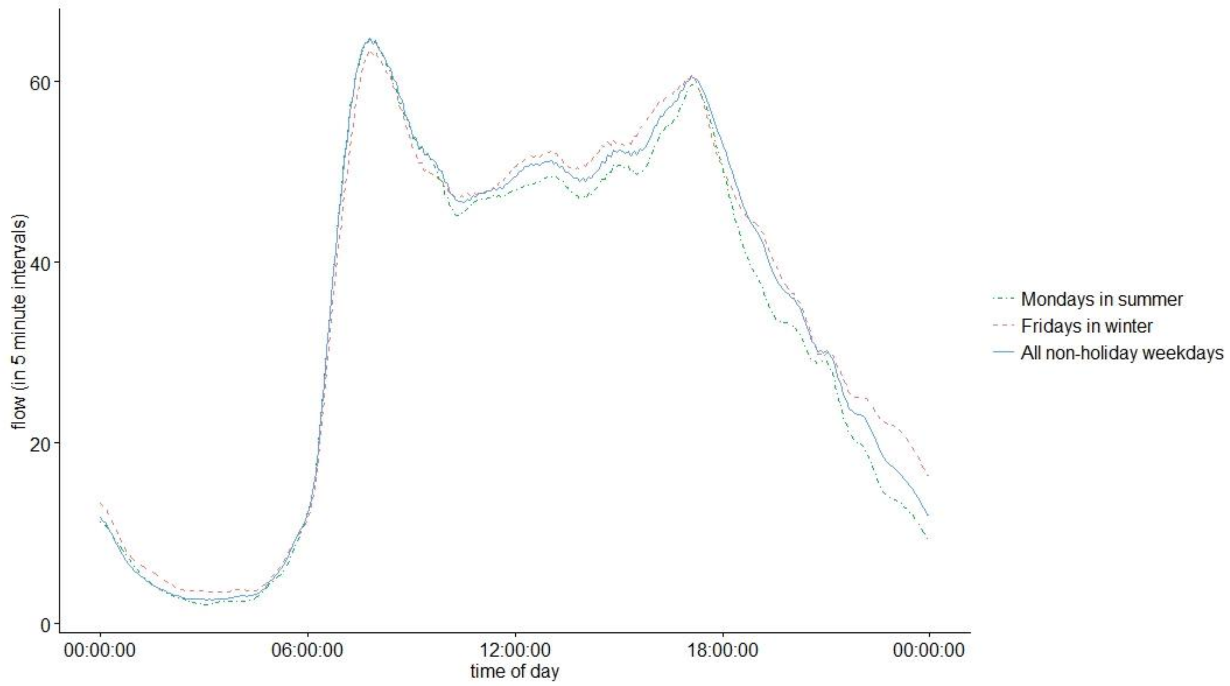

Figure 2-10: Example flow profiles

2.4.2 Discussion

2.4.2.1 Considerations in applying the model

As discussed in Section 2.3.1, the value of the roughness penalty, λ, was estimated using cross-validation. The cross-validation process aims to identify the 'best' value of λ by minimising the differences between the individual data points and the estimated daily profiles using a linear model applied to different subsets of the data. It is not clear, however, whether this objective is suitable for all purposes. In order to examine the robustness of the roughness penalty used, the day of the week analysis was undertaken separately using lambda values of 10^{-8}, 10^{-7}, 10^{-6}, 10^{-5} and 10^{-4}. Although all seven days of the week were identified as being statistically significant under each of these values of lambda, the value of the roughness penalty used does affect the ability to interpret possible reasons for the differences between coefficients using visual inspection. Figure 2-11 demonstrates the different attributes of a 'typical' Saturday identified under different values of lambda. As suggested by Silverman (1985, p5), it may be preferable to use the cross-validation process as a starting point and then examine alternative values whilst considering the purpose of the research. In this case study, for example, the larger values of lambda may be suitable if the profiles are to be used as inputs to a within-day and day-to-day dynamical model for area-wide forecasting where only an approximation of the profiles for each day type are required. Smaller values of lambda could be used for applications requiring very detailed information about parts of the profile, for example for modelling the impact of time of day varying congestion charges on different day types.

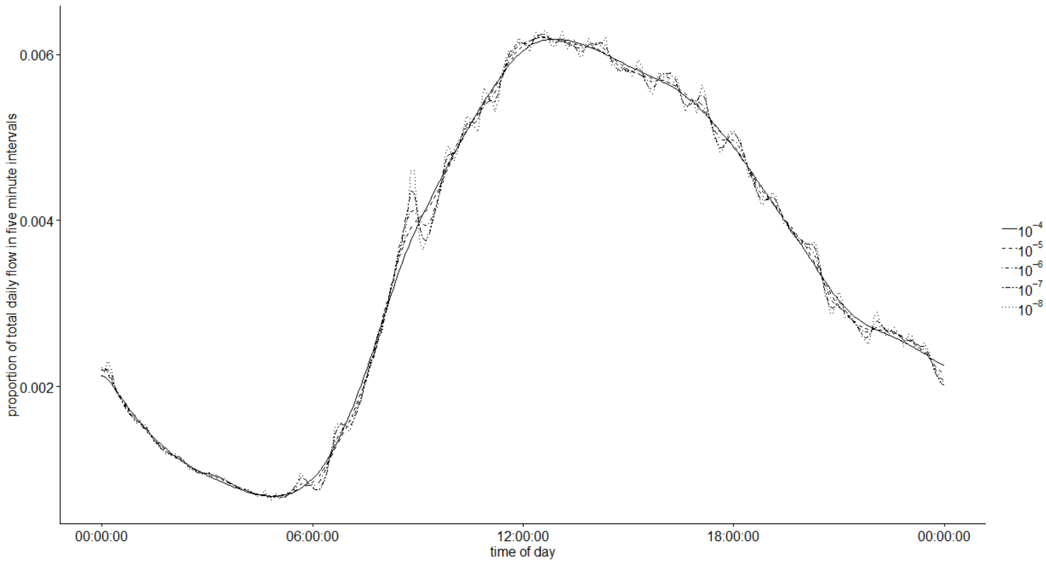

Figure 2-11: Saturday coefficients for lambda between 10^{-8} and 10^{-4}

Another significant decision made in the analysis was to only exclude daily profiles where the data was known to be incorrect due to data collection issues. Other individual profiles which visually would appear to be outliers were included in the analysis. Note that the identification of individual outlier profiles is different from the exclusion of Bank Holidays which could be done a priori using dates alone and therefore is a well defined day type which is not suitable for inclusion in the analysis. Keeping individual ‘outliers’ in the data ensured that the full range of flow profiles actually observed were included so as not to bias any analysis by only including those which were perceived to be near some ‘expected’ profile. If the decision was made to exclude outliers, they could be identified using various methods including methods based on Principal Component Analysis (Chiou et al., 2014), influential observations (Shen and Xu, 2007) or the measurement of the ‘depth’ of a set of functions as used in Guardiola et al. (2014).

2.4.2.2 Analysis of residuals

For the analysis of total daily flows to be valid, the residuals need to be independent. A plot of the residuals after an initial application of the Kruskal-Wallis test revealed trends in the data for December and in early January (up to 5th), with residuals gradually increasing before and after Christmas. The data for these periods in both years was then removed and the residuals in Figure 2-12 were obtained. There are no obvious patterns which suggest the residuals are interdependent. Although this assumption has been satisfied, using the seasonal and trend adjustment process means that it is not reasonable to formally test seasonal differences in the data. An examination of the monthly adjustment factors applied may, however, suggest particular months or seasons to model separately.

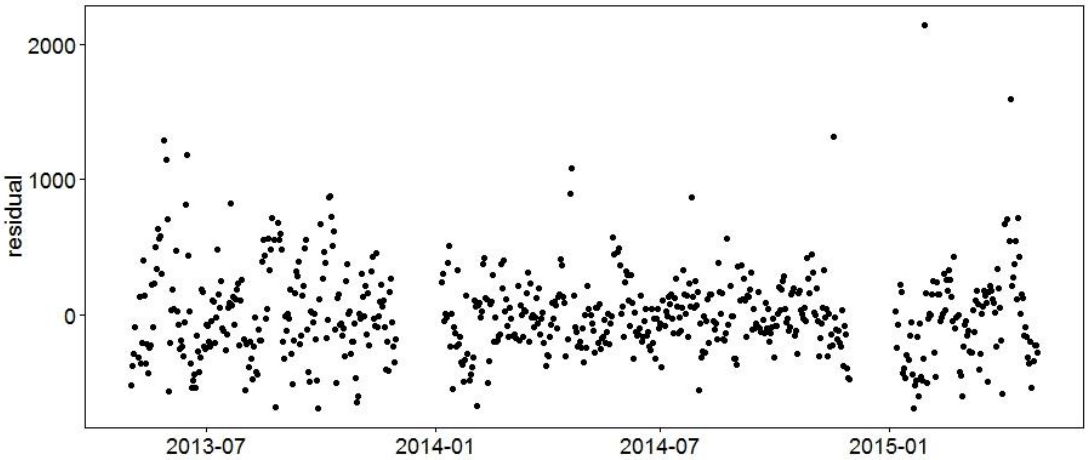

Figure 2-12: Total daily flow residuals after seasonal adjustment, excluding December and early January data

The functional residuals arising from the Functional Linear Model (FLM) also need to be considered. Each residual from the day of the week model can be calculated using:

$$\hat{\varepsilon}_i(\mathrm{t}) = y_i(t) - \boldsymbol{Z}_i{}^T\widehat{\boldsymbol{\beta}}(t) \tag{9}$$

where $\boldsymbol{Z}_i$ is the vector of day type indicators relating to day *i*. These are residuals from the Functional Linear Model only and therefore only consider the smoothed functions and *not* the underlying point data. In this research the curve estimation process used is well established and therefore the residuals generated by moving from the point to the curve data do not require additional investigation.

As proposed by Faraway (1997), Functional Principal Component Analysis (FPCA) of the residuals from the day of the week FLM analysis was undertaken. In contrast to the main analysis, FPCA is appropriate here as the aim is to examine the variability not explained by the model to assess the suitability of the model. The first four components are shown in Figure 2-13. For each residual, the values at different times of the day are not independent, which supports the use of FDA where this is taken into account. The percentage of variability explained by each component (also shown in Figure 2-13) suggests that there may be other explanatory variables affecting the daily flow profile which haven't been accounted for. These could include school term times, sporting events or weather conditions. The impact of weather conditions on transport choices and travel conditions has received a lot of attention (see Böcker et al. (2012) for a summary) and there has been progress in modelling such conditions, for example Lam et al. (2008) and Sumalee et al. (2011a) which could be utilised for scenario testing.

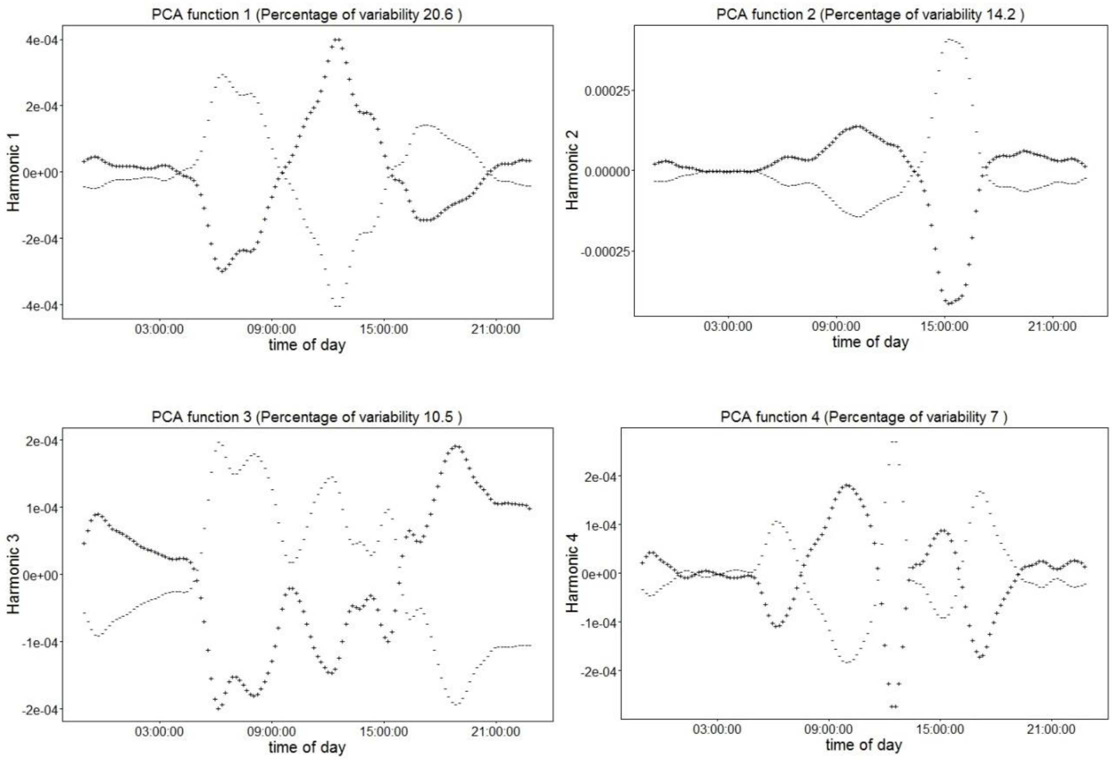

Figure 2-13: The first four principal components of the FLM residuals

The F-type test used in Section 2.4.1 relies on the assumption that the functional residuals are independent Gaussian stochastic processes (Shen and Faraway, 2004). Shen and Xu (2007) proposed visualising the Q-Q plot of studentized residuals against a Chi Squared distribution with φ degrees of freedom. The plot for the data analysed above is shown in Figure 2-14.

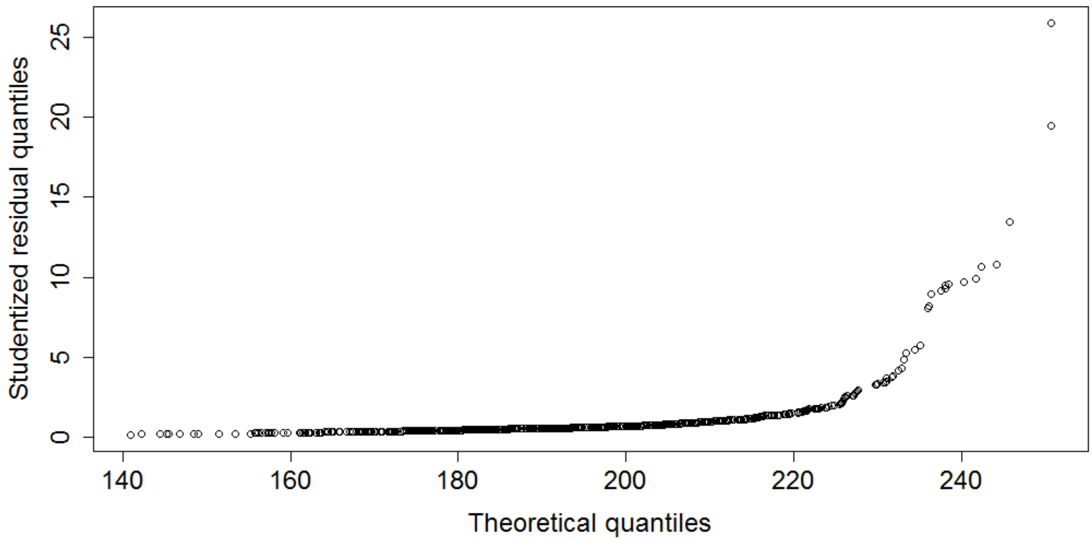

Figure 2-14: Q-Q plot of studentized residuals and the associated Chi Squared distribution

Clearly the plot is not a straight line as expected if the assumption holds. To test whether this was due to outliers in the data, the regression process was repeated after excluding influential outliers which were identified using Cook's Distance (Shen and Xu, 2007). The resulting Q-Q plot was not a straight line either. Further testing identified that the distribution of the studentized residuals had a higher kurtosis than the associated Chi Squared distribution. When the Gaussian assumption does not hold, permutation tests can be used to estimate the distribution of the test statistic rather than using the F distribution (Zhang, 2013). Good (1994) is a comprehensive text on permutation tests for point data, and other authors have applied these techniques to functional data (Muñoz Maldonado et al., 2002, Zhang, 2013, Corain et al., 2014). In a permutation test, the 'labels' connecting explanatory variables to observations are rearranged, and the test statistic is computed for this new, permuted dataset. This process is repeated until a large enough sample of all possible permutations has been collated. This sample provides information about the distribution of the test statistic under the null hypothesis that the 'labels' do not provide any information about the observed value. The test statistic computed using the correct labels is then compared to this distribution. Permutation tests provide a flexible way of testing hypotheses where distributional assumptions do not apply, but they are very computationally intensive in comparison to the F-type test.

A sample permutation test was undertaken for the final stage of the weekday analysis, namely the hypothesis that Tuesdays and Wednesdays have significantly different daily flow profiles. The Tuesday and Wednesday indicators were shuffled 5,000 times and the test statistic was computed for each permutation. This process had a running time of approximately 100 hours, but improvements could be made to speed up the process, for example by running sets of permutations in parallel. The 95th percentile of the distribution, i.e. the critical value, obtained was 1.82 which compares to the estimated critical value based on the F distribution above of 1.17. The value of the test statistic was 5.06 and therefore the outcome of the test is the same under either critical value. This may not always be the case, however, so permutation testing should always be considered if the Gaussian assumption is not satisfied.

2.5 Conclusions and future work

In this paper we present a method for identifying day types, relating to the day of the week or time of year, with systematically different daily flow profiles. The method utilises Functional Data Analysis which is not often used in transportation research. This approach has advantages as it can retain the complexity of within-day flow dynamics whilst having the conceptual simplicity of having one observation (in the form of a profile) to represent each day. This paper describes how data can be transformed into functional data and how linear models can be developed using the data. A statistical method is also presented for identifying the preferred model formulation and thus the day type factors which have a statistically significant effect on flow profiles. The example using real-life traffic flow data identified that all seven days of the week had distinctive differences in the shape of the daily flow profile at this site. These differences included the timing and intensity of peak periods but also differences during the night, which is increasingly of interest.

The methodology could be applied by practitioners to gain a better understanding of traffic flows and would not require the purchase of any specialist software as the analysis in this paper was undertaken using the free software R. By applying the proposed methodology, days of the week or times of year where particular problems arise could be identified and then policy solutions to target the causes of these problems could be devised. Examples of such targeted policies might include: time-of-day varying congestion charges that themselves vary by day of the week or season; incentives to influence employers or shopping centres to adjust their opening times by day of the week or season; targeting public transport provision so that timetables are more responsive to the day of the week and seasonal needs of travellers.

The method in this paper could be built upon to analyse multiple sites. One approach could be to analyse sites independently and then develop a classification of sites based on the daily flow profile coefficients, or extract ‘global’ effects. Alternatively a more complex model could be developed to account for link correlations so that the relative attractiveness of routes under different scenarios could be considered. As well as considering flows, future work could also consider systematic differences in capacity, perhaps due to lighting (van Goeverden et al., 1998, Tenekeci et al., 2010) or/and weather effects (El Faouzi et al., 2010, Calvert and Snelder, 2013). Day-to-day dynamical models (for example from Watling and Cantarella (2013b)) could be extended to incorporate predictable differences, both in terms of input variables such as demand, but also by extending the transition functions between days to account for day type specific learning. The models including stochastic demand in Watling and Cantarella (2013b, Section 4) could also be extended so that the stochastic demand relates to functional data, i.e. randomly selected daily profiles as opposed to single values.

2.6 References

Böcker, L., Dijst, M. & Prillwitz, J. (2012) Impact of Everyday Weather on Individual Daily Travel Behaviours in Perspective: A Literature Review. *Transport Reviews,* 33**,** 71-91.

Calvert, S. C. & Snelder, M. Influence of rain on motorway road capacity – a data-driven analysis. 16th International IEEE Annual Conference on Intelligent Transportation Systems (ITSC 2013), 2013 The Hague, The Netherlands. 1481-1486.

Chen, K. & Müller, H.-G. (2014) Modeling Conditional Distributions for Functional Responses, With Application to Traffic Monitoring via GPS-Enabled Mobile Phones. *Technometrics,* 56**,** 347-358.

Chiou, J.-M., Zhang, Y.-C., Chen, W.-H. & Chang, C.-W. (2014) A functional data approach to missing value imputation and outlier detection for traffic flow data. *Transportmetrica B: Transport Dynamics,* 2**,** 106-129.

Clark, S. & Watling, D. (2005) Modelling network travel time reliability under stochastic demand. *Transportation Research Part B: Methodological,* 39**,** 119-140.

Corain, L., Melas, V. B., Pepelyshev, A. & Salmaso, L. (2014) New insights on permutation approach for hypothesis testing on functional data. *Advances in Data Analysis and Classification,* 8**,** 339-356.

de Boor, C. (2001) *A practical guide to splines,* New York, Springer.

Du, J., Wong, S. C., Shu, C.-W. & Zhang, M. (2015) Reformulating the Hoogendoorn–Bovy predictive dynamic user-optimal model in continuum space with anisotropic condition. *Transportation Research Part B: Methodological,* 79**,** 189-217.

El Faouzi, N.-E., de Mouzon, O., Billot, R. & Sau, J. (2010) Assessing the changes in operating traffic stream conditions due to weather conditions. *Advances in Transportation Studies,* 21A**,** 33-48.

Faraway, J. J. (1997) Regression Analysis for a Functional Response. *Technometrics,* 39**,** 254-261.

Ferraty, F. & Vieu, P. (2003) Curves discrimination: a nonparametric functional approach. *Computational Statistics & Data Analysis,* 44**,** 161-173.

Gao, H. O. & Niemeier, D. A. (2008) Using functional data analysis of diurnal ozone and NOx cycles to inform transportation emissions control. *Transportation Research Part D: Transport and Environment,* 13**,** 221-238.

Good, P. (1994) *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses,* New York, Springer.

Górecki, T. & Smaga, Ł. (2015) A comparison of tests for the one-way ANOVA problem for functional data. *Computational Statistics,* 30**,** 987-1010.

Green, P. J. & Silverman, B. W. (1994) *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach,* London, Chapman and Hill.

Guardiola, I. G., Leon, T. & Mallor, F. (2014) A functional approach to monitor and recognize patterns of daily traffic profiles. *Transportation Research Part B: Methodological,* 65**,** 119-136.

Guo, F., Li, Q. & Rakha, H. (2012) Multistate Travel Time Reliability Models with Skewed Component Distributions. *Transportation Research Record: Journal of the Transportation Research Board,* 2315**,** 47-53.

Guo, R.-Y., Yang, H., Huang, H.-J. & Tan, Z. (2015) Link-based day-to-day network traffic dynamics and equilibria. *Transportation Research Part B: Methodological,* 71**,** 248-260.

Habib, K. M. N., Day, N. & Miller, E. J. (2009) An investigation of commuting trip timing and mode choice in the Greater Toronto Area: Application of a joint discrete-continuous model. *Transportation Research Part A: Policy and Practice,* 43**,** 639-653.

Habib, K. M. N. & Miller, E. J. (2008) Modelling daily activity program generation considering within-day and day-to-day dynamics in activity-travel behaviour. *Transportation,* 35**,** 467-484.

Han, K., Friesz, T. L., Szeto, W. Y. & Liu, H. (2015) Elastic demand dynamic network user equilibrium: Formulation, existence and computation. *Transportation Research Part B: Methodological,* 81**,** 183-209.

Hanson, S. & Huff, J. O. (1988) Systematic variability in repetitious travel. *Transportation,* 15**,** 111-135.

Hazelton, M. L. & Parry, K. (2016) Statistical methods for comparison of day-to-day traffic models. *Transportation Research Part B: Methodological,* 92, Part A**,** 22-34.

Järv, O., Ahas, R. & Witlox, F. (2014) Understanding monthly variability in human activity spaces: A twelve-month study using mobile phone call detail records. *Transportation Research Part C: Emerging Technologies,* 38**,** 122-135.

Jiang, X. & Adeli, H. (2005) Dynamic Wavelet Neural Network Model for Traffic Flow Forecasting. *Journal of Transportation Engineering,* 131**,** 771-779.

Kamga, C. & Yazici, M. A. (2014) Temporal and weather related variation patterns of urban travel time: Considerations and caveats for value of travel time, value of variability, and mode choice studies. *Transportation Research Part C: Emerging Technologies,* 45**,** 4-16.

Kitamura, R. & Van Der Hoorn, T. (1987) Regularity and irreversibility of weekly travel behavior. *Transportation,* 14**,** 227-251.

Kumar, A. & Peeta, S. (2015) A day-to-day dynamical model for the evolution of path flows under disequilibrium of traffic networks with fixed demand. *Transportation Research Part B: Methodological,* 80**,** 235-256.

Lam, W. H. K., Shao, H. & Sumalee, A. (2008) Modeling impacts of adverse weather conditions on a road network with uncertainties in demand and supply. *Transportation Research Part B: Methodological,* 42**,** 890-910.

Lee, K. L., Meyer, R. J. & Bradlow, E. T. (2009) Analyzing risk response dynamics on the web: the case of Hurricane Katrina. *Risk Anal,* 29**,** 1779-92.

Li, L., Su, X., Zhang, Y., Lin, Y. & Li, Z. (2015) Trend Modeling for Traffic Time Series Analysis: An Integrated Study. *IEEE Transactions on Intelligent Transportation Systems,* 16**,** 3430-3439.

Lian, H. (2013) Shrinkage estimation and selection for multiple functional regression. *Statistica Sinica,* 23**,** 51-74.

Lo, H. K. & Tung, Y.-K. (2003) Network with degradable links: capacity analysis and design. *Transportation Research Part B: Methodological,* 37**,** 345-363.

Long, J., Szeto, W. Y., Gao, Z., Huang, H.-J. & Shi, Q. (2016) The nonlinear equation system approach to solving dynamic user optimal simultaneous route and departure time choice problems. *Transportation Research Part B: Methodological,* 83**,** 179-206.

Matsui, H. & Konishi, S. (2011) Variable selection for functional regression models via the L_1 regularization. *Computational Statistics & Data Analysis,* 55**,** 3304-3310.

May, A. D. (1990) *Traffic flow fundamentals,* Englewood Cliffs, Prentice Hall.

Mingotti, N., Lillo, R. E. & Romo, J. (2013) Lasso Variable Selection in Functional Regression. *Working papers* [Online]. Available: http://orff.uc3m.es/bitstream/handle/10016/16959/ws131413.pdf?sequence=1.

Muñoz Maldonado, Y., Staniswalis, J. G., Irwin, L. N. & Byers, D. (2002) A similarity analysis of curves. *The Canadian Journal of Statistics,* 30**,** 373-381.

Nakayama, S. & Watling, D. (2014) Consistent formulation of network equilibrium with stochastic flows. *Transportation Research Part B: Methodological,* 66**,** 50-69.

Nerini, D. & Ghattas, B. (2007) Classifying densities using functional regression trees: Applications in oceanology. *Computational Statistics & Data Analysis,* 51**,** 4984-4993.

Ngoduy, D., Hoang, N. H., Vu, H. L. & Watling, D. (2016) Optimal queue placement in dynamic system optimum solutions for single origin-destination traffic networks. *Transportation Research Part B: Methodological,* 92**,** 148-169.

Noland, R. B. & Small, K. A. (1995) Travel-time uncertainty, departure time choice, and the cost of the morning commute. Irvine: Institute of Transportation Studies, University of California.

Ozbay, K., Yang, H., Morgul, E. F., Mudigonda, S. & Bartin, B. Big Data and the Calibration and Validation of Traffic Simulation Models. Looking Back and Looking Ahead: Celebrating 50 Years of Traffic Flow Theory, 2014 Washington D.C.: Transportation Research Board, 92-122.

Pu, W. (2011) Analytic Relationships Between Travel Time Reliability Measures. *Transportation Research Record: Journal of the Transportation Research Board,* 2254**,** 122-130.

Rakha, H. & Van Aerde, M. (1995) Statistical analysis of day-to-day variations in real-time traffic flow data. *Transportation Research Record: Journal of the Transportation Research Board,* 1510**,** 26-34.

Ramsay, J. O., Hooker, G. & Graves, S. (2009) *Functional Data Analysis with R and MATLAB,* New York, Springer.

Ramsay, J. O. & Silverman, B. W. (1997) *Functional Data Analysis,* New York, Springer.

Ramsay, J. O. & Silverman, B. W. (2005) *Functional Data Analysis,* New York, Springer.

Ramsay, J. O., Wickham, H., Graves, S. & Hooker, G. (2014) *Package 'fda'* [Online]. The Comprehensive R Archive Network (CRAN). Available: http://cran.r-project.org/web/packages/fda/fda.pdf [Accessed: 17th December 2014].

Schlich, R. & Axhausen, K. W. (2003) Habitual travel behaviour: Evidence from a six-week travel diary. *Transportation,* 30**,** 13-36.

Shao, H., Lam, W. H. K. & Tam, M. L. (2006) A reliability-based stochastic traffic assignment model for network with multiple user classes under uncertainty in demand. *Networks & Spatial Economics,* 6**,** 173-204.

Shen, Q. & Faraway, J. (2004) An F Test for Linear Models with Functional Responses. *Statistica Sinica,* 14**,** 1239-1257.

Shen, Q. & Xu, H. (2007) Diagnostics for Linear Models With Functional Responses. *Technometrics,* 49**,** 26-33.

Silverman, B. W. (1985) Some Aspects of the Spline Smoothing Approach to Non-Parametric Regression Curve Fitting. *Journal of the Royal Statistical Society: Series B (Methodological),* 47**,** 1-52.

Siu, B. W. Y. & Lo, H. K. (2008) Doubly uncertain transportation network: Degradable capacity and stochastic demand. *European Journal of Operational Research,* 191**,** 166-181.

Stathopoulos, A. & Karlaftis, M. (2001) Temporal and Spatial Variations of Real-Time Traffic Data in Urban Areas. *Transportation Research Record: Journal of the Transportation Research Board,* 1768**,** 135-140.

Sumalee, A., Uchida, K. & Lam, W. H. K. (2011a) Stochastic multi-modal transport network under demand uncertainties and adverse weather condition. *Transportation Research Part C: Emerging Technologies,* 19**,** 338-350.

Sumalee, A., Watling, D. & Nakayama, S. (2006) Reliable Network Design Problem: Case with Uncertain Demand and Total Travel Time Reliability. *Transportation Research Record: Journal of the Transportation Research Board,* 1964**,** 81-90.

Sumalee, A., Zhong, R. X., Pan, T. L. & Szeto, W. Y. (2011b) Stochastic cell transmission model (SCTM): A stochastic dynamic traffic model for traffic state surveillance and assignment. *Transportation Research Part B: Methodological,* 45**,** 507-533.

Tang, J., Wang, Y., Wang, H., Zhang, S. & Liu, F. (2014) Dynamic analysis of traffic time series at different temporal scales: A complex networks approach. *Physica A: Statistical Mechanics and its Applications,* 405**,** 303-315.

Tenekeci, G., Montgomery, F. & Wainaina, S. (2010) Roundabout capacity in adverse weather and light conditions. *Proceedings of the ICE - Transport,* 163**,** 29-39.

Tibshirani, R. (1996) Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological),* 58**,** 267-288.

Transport for London (2013) Olympic Legacy Monitoring: Personal Travel Behaviour during the Games Available: content.tfl.gov.uk/olympic-legacy-personal-travel-report.pdf [Accessed: 29th July 2016].

Ukkusuri, S. V., Han, L. & Doan, K. (2012) Dynamic user equilibrium with a path based cell transmission model for general traffic networks. *Transportation Research Part B: Methodological,* 46**,** 1657-1684.

van Goeverden, C. D., Botma, H. & Bovy, P. H. L. (1998) Determining Impact of Road Lighting on Motorway Capacity. *Transportation Research Record: Journal of the Transportation Research Board,* 1646**,** 1-8.

Vlahogianni, E. I., Golias, J. C. & Karlaftis, M. G. (2004) Short-term traffic forecasting: Overview of objectives and methods. *Transport Reviews,* 24**,** 533-557.

Vlahogianni, E. I., Karlaftis, M. G. & Golias, J. C. (2014) Short-term traffic forecasting: Where we are and where we're going. *Transportation Research Part C: Emerging Technologies,* 43**,** 3-19.

Wang, D. Z. W. & Du, B. (2016) Continuum modelling of spatial and dynamic equilibrium in a travel corridor with heterogeneous commuters—A partial differential complementarity system approach. *Transportation Research Part B: Methodological,* 85**,** 1-18.

Watling, D. (2002) A Second Order Stochastic Network Equilibrium Model, I: Theoretical Foundation. *Transportation Science,* 36**,** 149-166.

Watling, D., Sumalee, A., Connors, R. & Balijepalli, C. (2004) Advancing methods for evaluating network reliability: A Department for Transport "New Horizons" Project Final Report. Available: http://webarchive.nationalarchives.gov.uk/20070306100043/http:/www.dft.gov.uk/pgr/economics/rdg/jtv/amenr/advancingmethodsforevaluatin3080 [Accessed: 14th May 2014].

Watling, D. P. & Cantarella, G. E. (2013a) Model Representation & Decision-Making in an Ever-Changing World: The Role of Stochastic Process Models of Transportation Systems. *Networks and Spatial Economics,* 15**,** 843-882.

Watling, D. P. & Cantarella, G. E. (2013b) Modelling sources of variation in transportation systems: theoretical foundations of day-to-day dynamic models. *Transportmetrica B: Transport Dynamics,* 1**,** 3-32.

Weijermars, W. & van Berkum, E. Analyzing highway flow patterns using cluster analysis. 8th International IEEE Conference on Intelligent Transportation Systems, 2005 Vienna, Austria. 831-836.

Weijermars, W. A. M. & van Berkum, E. C. (2004) Daily flow profiles of urban traffic. *In:* BREBBIA, C. A. & WADHWA, L. C. (eds.) *Urban Transport X. Urban Transport and the Environment in the 21st Century.* Dresden, Germany: WIT Press.

Xiao, F., Yang, H. & Ye, H. (2016) Physics of day-to-day network flow dynamics. *Transportation Research Part B: Methodological,* 86**,** 86-103.

Yang, X., Shen, Q., Xu, H. & Shoptaw, S. (2007) Functional regression analysis using an F test for longitudinal data with large numbers of repeated measures. *Stat Med,* 26**,** 1552-66.

Yazici, M., Kamga, C. & Mouskos, K. (2012) Analysis of Travel Time Reliability in New York City Based on Day-of-Week and Time-of-Day Periods. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 83-95.

Zhang, J.-T. (2013) Analysis of Variance for Functional Data. Boca Raton, Florida: CRC Press.

Zhang, W., Mediria, A. & Rakha, H. Statistical analysis of spatiotemporal link and path flow variability. 2007 IEEE Intelligent Transportation Systems Conference, 2007 Seattle, Washington, USA. IEEE, 1080-1085.

3 Assessing the feasibility of using Bluetooth data to examine the repeated travel behaviour of road users

Abstract

Information about travellers' repeated trip behaviour, including measures of regularity and variability, provides valuable insights into traveller flexibility, habit and network experience. Such data is usually collected using multi-day travel diaries but this paper explores whether data from fixed Bluetooth detectors could also be used. Consideration is given to sources of bias, for example non-random sampling by traveller demographic, which are relevant when using detector data for this purpose as opposed to measuring travel times. The potential of Bluetooth data with respect to spatial and temporal aspects of intrapersonal variability is explored using a case study in northern England. Although the case study includes just eight Bluetooth detectors, in one year over 2 million trips were recorded. Despite most Bluetooth devices recording very few trips, 1,240 devices were recorded at least once per weekday on average, providing a large sample for examining intrapersonal variability. The study highlights potential pitfalls in analysing intrapersonal variability in the timing of trips that can be remedied by identifying comparable locations or trip types and using a suitable metric to calculate the variability. A suitable approach is presented which exploits the fact that Bluetooth data is collected at fixed points on the network and which does not require departure time or trip purpose information. Similarly for spatial analyses, traditional approaches for analysing origin-destination pairs are not applicable, but by utilising all available spatial information about trips, trajectories can be estimated which can provide even more information about spatial variability. In the case study area the detection rates varied greatly between sensors and therefore collecting

supplementary data on the detection rates of individual sensors once they have been installed is advised.

3.1 Introduction

In reviewing developments in transportation research, both Schlich and Axhausen (2003) and Heinen and Chatterjee (2015) have highlighted the disproportionate amount of attention paid to the variability between people, in contrast to the variability in an individual's behaviour from day-to-day. The latter, known as intrapersonal variability, is important because the regularity (or irregularity) of an individual's trips can provide an insight into their transport needs (Schlich and Axhausen, 2003). These insights, which could relate to flexibility in the timing of trips or to the number of unique users in the observed situation, could inform transport policies. Measuring the regularity of travel behaviours also matters because many transport models implicitly assume that the same individuals travel to work at roughly the same time every weekday. This is particularly relevant for transport project appraisal where assumptions about trip regularity could have a substantial impact on predictions relating to behavioural response. Intrapersonal variability can also inform the parameter values for day-to-day dynamical models which include learning mechanisms, for example the "switching choice probability" described in Cantarella and Cascetta (1995) which relates to travellers reconsidering, but not necessarily changing, their previous route choice. In other research, such parameters are considered to be measures of habit (Arentze and Timmermans, 2005, p16). Understanding intrapersonal variability can also inform the development of user classes, for example based on attitude to risk (Shao et al., 2006) or information availability (Han et al., 2016), and the corresponding parameters for modelling choices. Another application could be to use a model of intrapersonal variability as the basis for generating daily variations in demand, rather than using a standard statistical distribution such as in the models discussed in Nakayama and Watling (2014).

In transportation research, intrapersonal variability has traditionally been explored using multi-day travel diary data from an activity pattern perspective (Huff and Hanson, 1986, Jones and Clarke, 1988, Bayarma et al., 2007). More recently such surveys have utilised newer data collection tools such as GPS trackers (Elango et al., 2007, Stopher and Zhang, 2011) or mobile phone applications (Safi et al., 2015). These surveys are relatively expensive to undertake and place a burden on participants, so sample sizes and the period of time surveyed are usually small. Of the surveys listed above, for example, some collected data for 15 days or less (Jones and Clarke, 1988, Stopher and Zhang, 2011, Safi et al., 2015) and others covered longer periods of time (between 35 days and 1 year), but had relatively few participants (149 people and 139 and 153 households respectively for Huff and Hanson (1986), Bayarma et al. (2007) and Elango et al. (2007)). The availability of data from newly emerging sources, however, provides new opportunities to analyse individual travel patterns for greater numbers of people over much longer periods of time; behaviour over long periods is likely to be much more variable due to seasonal factors and longer term trends. For example, mobile phone data was used by Järv et al. (2014) to analyse changes in individuals' activity spaces over a twelve month period. While mobile phone data can only provide limited geographical information[2], other data sources are more closely connected to the transport network. Smart card data, for example, provides opportunities to gain a better understanding of individuals' use of public transportation. The literature on repeated trips using smart card data includes descriptions of trip frequency

---

[2] The only geographic information provided is the associated network antenna, which in Järv et al. (2014, p126) covered 0.8 km^2 in Tallinn itself and 15.3 km^2 in the surrounding area.

(Utsunomiya et al., 2006) and more complex analyses such as the clustering of travellers based on their trip characteristics (Morency et al., 2007, Ma et al., 2013, Kieu et al., 2015b). Whilst the usefulness of analysing smart card data for strategic, tactical and operational purposes has been recognised (Pelletier et al., 2011), similar data sources have not been utilised to better understand *road* users.

One emerging data source which is particularly relevant for the analysis of road users is Bluetooth data. Fixed Bluetooth sensors can be placed alongside roads and then set to continuously scan for any discoverable Bluetooth devices within their detection zone (see Bhaskar and Chung (2013) for more details). Devices which use Bluetooth include mobile phones, laptops, hands-free devices and in-car audio systems. For any Bluetooth devices detected by a sensor, the unique identifier, known as the MAC address, is recorded along with the time of detection. Bluetooth data is becoming increasingly popular for measuring travel times on the road network (Haseman et al., 2010, Hainen et al., 2011, Moghaddam and Hellinga, 2013), particularly in urban areas, and has also been used in OD estimation (Barceló et al., 2010, Carpenter et al., 2012). Despite this growth in usage relating to spatial matching of observations, the day-to-day matching of MAC addresses to examine intrapersonal variability (or more correctly intra-device variability) has not been addressed, to the knowledge of the authors.

Bluetooth is therefore an established source of data which has more precise location information than mobile phone data and can be used to collect data on a large number of travellers over a long period of time relatively cheaply. There are disadvantages, however, and the limitations of Bluetooth data for travel time estimation are well documented (see Araghi et al. (2014)). The requirements of data for analysing intrapersonal variability are different, however, and the aim of this paper is to explore the feasibility of using Bluetooth data to collect detailed information about intrapersonal variability in travel on the road network. This paper addresses both temporal and spatial variability and Bluetooth data from a case study area in Wigan, northern England, is used to demonstrate the types of analyses which could be attempted. These examples then inform discussions on the suitability of Bluetooth data for each purpose.

The structure of the paper will be as follows. Section 3.2 includes a detailed description of the case study area and a more in-depth discussion of Bluetooth data. Section 3.3 discusses potential sources of bias in the data, including missing individual trips which is particularly problematic for estimating trip frequencies. Intrapersonal variability in the temporal and spatial aspects of trips are considered in Sections 3.4 and 3.5 respectively. Section 3.6 concludes the paper.

3.2 Bluetooth data in the case study area

Fixed Bluetooth sensors, also known as detectors, can record the unique identifier (MAC address) and corresponding timestamp for discoverable Bluetooth devices passing close by. The Bluetooth devices could be associated with the vehicle, for example in-car sound systems or hands-free kits, or with a person in the vehicle, for example a mobile phone, tablet or laptop. It is, however, possible to change the settings within these individual devices so that they are not discoverable by such sensors (Haghani et al., 2010).

In most applications Bluetooth data is used for travel time estimation. This is done by selecting two locations, matching the unique MAC addresses and then calculating the travel times as the differences between timestamps. Once observations are matched from more than two sensors and across days or weeks, however, there is no longer a straightforward rectangular structure to the data. The cleaning of the data needs to consider which observations to chain together and which to separate. The complex data structure arises as a result of retaining all observations within a chain so as not to lose data.

3.2.1 Case study area

This paper uses real data provided by Transport for Greater Manchester (TfGM) to demonstrate how Bluetooth data could be analysed. TfGM co-ordinate transport across Greater Manchester which is a collection of ten boroughs in northwest England.

Since 2011, TfGM have been installing fixed Bluetooth detectors alongside major arterials and orbitals in and around key urban centres such as Manchester, Wigan and Rochdale. They were installed for the purpose of monitoring travel times on strategically important routes. Antennae with 9dBi gain are used, which Bhaskar and Chung (2013) found provided a range of approximately 100m. The developer of the detectors state that it can cover up to 6 lanes of traffic travelling at 70mph. TfGM adjust the strength of detectors on installation so that each is appropriate for the size of the junction. The equipment uses an algorithm to truncate and encrypt MAC addresses prior to storing the data. The process is, therefore, consistent with data protection regulations in relation to this type of data. As of January 2017, the total number of Bluetooth detectors installed within Greater Manchester was over 750.

The case study area is around Wigan town centre, which sits within the Metropolitan Borough of Wigan and has a population of just over 322,000 (Office for National Statistics, 2016). In 2015, 39% of trips into the centre of Wigan were made by car (Transport for Greater Manchester, n.d.) during the morning peak period (7:30-9:30). The majority of vehicles on the roads into Wigan during this period were cars (over 80%), although there were also light goods vehicles and buses (Transport for Greater Manchester, n.d.). Although there are relatively few Bluetooth detectors in and around Wigan, the detectors in place provide good coverage of major roads into the town. The main focus of this paper will be the eight detectors shown in Figure 3-1.

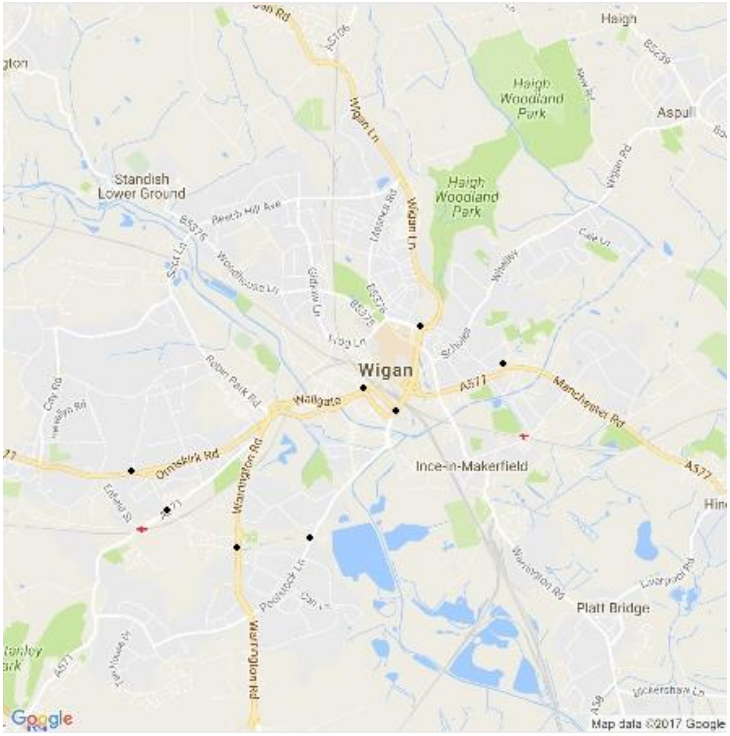

Figure 3-1: Map of case study area including Bluetooth detector locations

Data from all eight sites was analysed for one year, from 1/1/15 to 31/12/15. The examples of missed observations in Section 3.3.3 and route choice in Section 3.5.2 require specific detector characteristics which are not present in this case study area, however, and in those cases all available data from suitable sites close to Wigan has been used. An overview of the data processing steps to transform the raw data into trips is described in Figure 3-2.

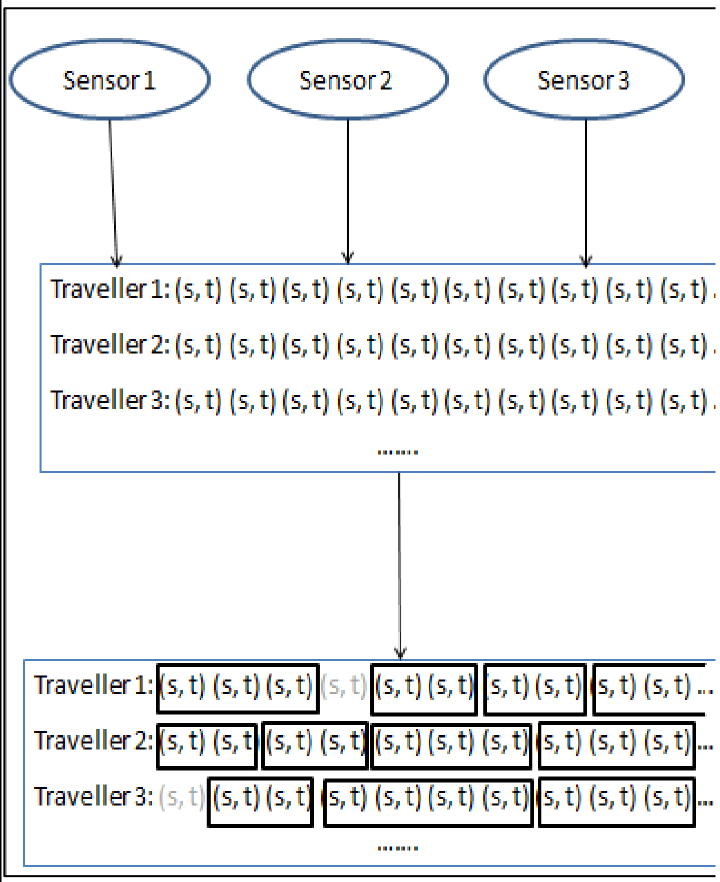

Ensure that a maximum of one observation is retained each time a device passes a detector.

Collate observations (each consisting of a sensor location, denoted by s, and a date-time stamp, denoted by t) by the MAC addresses. For each MAC address order the observations by the date-time stamp.

Compare adjacent observations to see whether they can be linked as part of the same trip. Consider whether the travel time between the two sensors is likely to be due to a direct journey between the two using a motorised vehicle (considering the distance between detectors, speed limits and the travel times of surrounding devices).

Figure 3-2: Processing Bluetooth data into trips

3.3 Potential sources of bias when estimating trip frequency or intrapersonal variability using Bluetooth data

One of the most common questions asked about repeated travel behaviour is the frequency with which trips are made. This informs us about the extent to which each traveller utilises the road network in their daily lives and how familiar each traveller is with the area. The frequency of trips may also be used to make inferences about traveller types, for example whether an individual works full time in the area.

After cleaning and matching the data from all eight sites, 2.3 million trips remained in the dataset for analysis. These trips were made by 196,557 devices, which equates to almost 12 trips per device in the year, on average. The trips were not evenly distributed between the devices, however. Over 81,000 devices recorded just 1 trip. As the data is highly skewed, the frequency is plotted using a log scale in Figure 3-3, which shows trips per device.

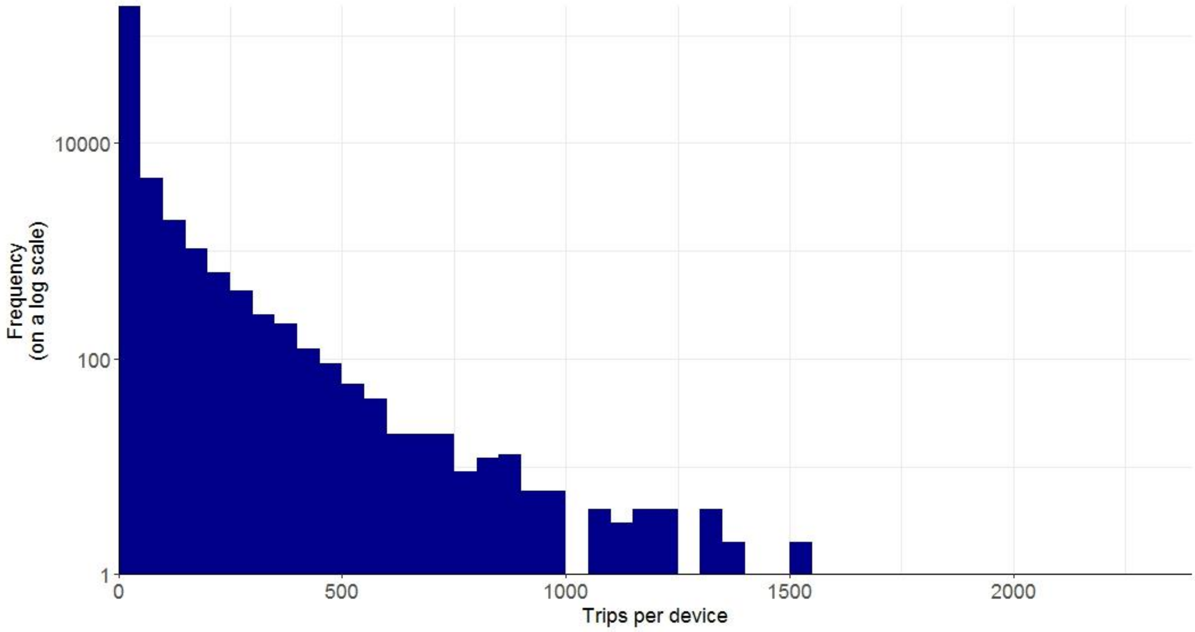

Figure 3-3: Number of trips detected in a one year period for each device

In total, 1,240 devices recorded a trip at least 260 times during the year, which equates to once per weekday on average. Although this is less than 1% of the *devices* recorded, the associated trips represent 22% of the *trips* observed. The number of devices drops to 210 if we consider the devices which were observed at least 520 times, which equates to twice each weekday on average.

Although this shows that in practical terms Bluetooth data can be used to collect data on repeated travel behaviour, it would be incorrect to report these basic findings as representative of road users in and around Wigan. The usefulness of the data first needs to be addressed by considering which trips are being recorded and what is missing.

3.3.1 What Bluetooth data is measuring

When considering trip frequency or intrapersonal variability using Bluetooth data, observations relate to *devices* and not to *people*. While in some cases devices are likely to remain close to one person, for example mobile phones, other types of device may be shared by households or businesses, for example in-car systems. Whilst it may seem preferable to have data at the individual level, in some cases vehicle level data can be more informative, for example by recording what Zhang et al. (2002) call ‘allocated’ household activities, whoever undertakes them. Delafontaine et al. (2012) have demonstrated that additional data collection can be undertaken to obtain details of the types of device being recorded by Bluetooth detectors, for example whether they are smartphones or laptops. If required, this approach could be used to identify data only relating to personal devices, such as smartphones. Scaling this additional data collection over many sites and long periods of time may be challenging as it is not usually part of the standard data collected by Bluetooth detectors developed predominantly for measuring travel times.

Bluetooth data will only ever provide information about a sample of trips. As with travel time estimation, the Bluetooth penetration rate will provide a measure for how meaningful it would be to make generalisations about the population from the data. When estimating travel times between two sensors, the penetration rate can be estimated relatively easily using loop detector data or Automatic Number Plate Recognition (ANPR) data. For example, TfGM compared ANPR and Bluetooth data for one link over a twelve hour period and calculated hourly penetration rates between 16% and 34%. When analysing broader measures of travel behaviour, however, the sampling rate is harder to calculate as it will depend upon the placement of detectors in relation to frequently used routes.

When estimating travel times, if the Bluetooth penetration rate is sufficiently large, any bias in the travellers with Bluetooth devices is assumed to have a minimal impact on the estimates. If the data is being used to examine intrapersonal variability, however, a biased sample of travellers may have a significant impact on results as aggregation does not occur prior to analysis.

3.3.2 Potential bias in people travelling with a discoverable Bluetooth device

Obviously, this method of data collection will only provide information about people travelling with a discoverable Bluetooth device. According to Ofcom (2016), 70% of UK adults use a smartphone, most of which will be Bluetooth-enabled. Also, 65% of UK adults access the internet through a mobile phone and 45% access the internet using a tablet (Ofcom, 2016, p25). The detectable devices are, however, a much smaller subset of devices which are switched on and where Bluetooth is in the discoverable mode. A relatively small survey of 218 business students by Jones and Chin (2015, p565) found that 70% of the students in 2014 claimed to ‘always’ disable Bluetooth on their mobile phones, so that it is not discoverable, when it’s not being used, down from 85% in 2011. This is, of course, self-reported data and the sample has an age and education bias. Phua et al. (2015) found that 34% of shoppers were carrying a phone with Bluetooth in the discovery mode during a survey of shoppers in a city in Australia. There is not, to the authors’ knowledge, a data source providing information about the wide range of devices which may now be Bluetooth-enabled, for example fitness trackers, headphones and in-car sound systems, including how often they are switched off and whether Bluetooth discoverability settings are ever changed.

There is a risk of bias when collecting data based on newer technologies, particularly in the age of travellers recorded. Although 92-93% of 16-34 year olds use a smartphone, this percentage decreases for subsequent age categories, down to just 8% for those 75 or older (Ofcom, 2016, p42). Interestingly, although internet usage via a smartphone follows a similar pattern, internet usage via tablets, E-Book readers and wearable technology does not peak in the lowest age category (16-24 years old). A bias in the age of travellers recorded would be problematic as the number of car trips per year increases between the ages of 17 and 49 and then falls with age (Department for Transport, 2016b). Also, Minnen et al. (2015) found differences in day-to-day variability in travel behaviour by age, where people aged 25-45 had lower levels of variability perhaps due to a higher number of constraints on their time.

Perhaps driven by legislation banning hand-held mobile phone use while driving in many countries, major worldwide car manufacturers offer Bluetooth facilities as standard or as an optional extra in most new cars. Mobile phones can be connected to these cars using Bluetooth to play music, for satellite navigation and for hands-free communication including calls and texts. This could introduce a bias into the data, as newer cars, particularly those with upgrades, are more likely to be driven by people with higher incomes. This is in addition to the bias in smartphone ownership towards people in socio-economic groups AB and C1 (Ofcom, 2016, p42). This would have an impact on results relating to intrapersonal variability as there is evidence that socio-economic classification has an impact on the number of trips made by car and the distance travelled (Department for Transport, 2016c). Also, Elango et al. (2007) found that higher income households have greater variability in travel behaviour. In contrast, however, Minnen et al. (2015) found more variability in travel patterns for unemployed people, compared to employed people, perhaps driven by the differences in trip types made, although only five days of data were analysed.

Another potential source of bias is the introduction of randomised MAC addresses in iPhones from iOS 8, which in principal may prevent devices from being tracked. There is limited evidence exploring the characteristics of iPhone users compared to users of other mobile devices, although Gerpott et al. (2013) have identified only limited differences in age and gender in a survey in Germany. They also speculate that iPhones may no longer be perceived as a highly innovative product which primarily attracts early adopters and therefore any bias in usage may be decreasing over time.

3.3.3 Individual missing trips

As well as missing people entirely, individual trips may also be missed. This is not particularly problematic when calculating travel times, provided overall penetration rates are reasonable, but when examining intrapersonal variability missing trips will result in underestimates of trip frequency. Missing trips could occur when the Bluetooth device is not taken on the trip, is not switched on, or Bluetooth is not in discovery mode. Alternatively, a discoverable Bluetooth device may not be detected when passing a sensor as it may not have spent enough time within the sensor's detection zone, the sensor may have reached its scanning cycle capacity for recording MAC addresses at that particular time, or a large vehicle may have formed a barrier to the sensor.

As detections at at least two sensors are required to record a trip, shorter trips are at greater risk of being missed altogether, either due to the location of Bluetooth detectors or due to not being detected when passing sensors. The increased likelihood of missing short trips is not unique to Bluetooth data, however. Bricka and Bhat (2006) examined under-reporting of trips in telephone travel surveys and reported that trips of under five minutes and discretionary trips were associated with under-reporting. As discussed in Kuwahara and Sullivan (1987), data collected from roadside surveys is likely to suffer from the “double counting” problem, where long trips are more likely to pass at least one survey site and therefore will be over-represented. While appropriate adjustments can be made to aggregated data to counteract these effects, attributing the adjustments to individual travellers would be challenging and perhaps could only be reasonably done for estimates of trip frequency.

To test the consistency of Bluetooth detection when a device passes a sensor, data from three sites on a road to the east of Wigan was analysed from February to July 2015. Only six consecutive months of data was available for all three sites in 2015. MAC addresses were matched between sites A and C (Figure 3-4) and all matches which could feasibly relate to a direct trip by a motor vehicle between the two sites were retained. Site B is on the same road as sites A and C and is on the most direct route between the two sites. The proportion of trips between A and C which were also detected at site B was then calculated and the results are shown in Table 3-1.

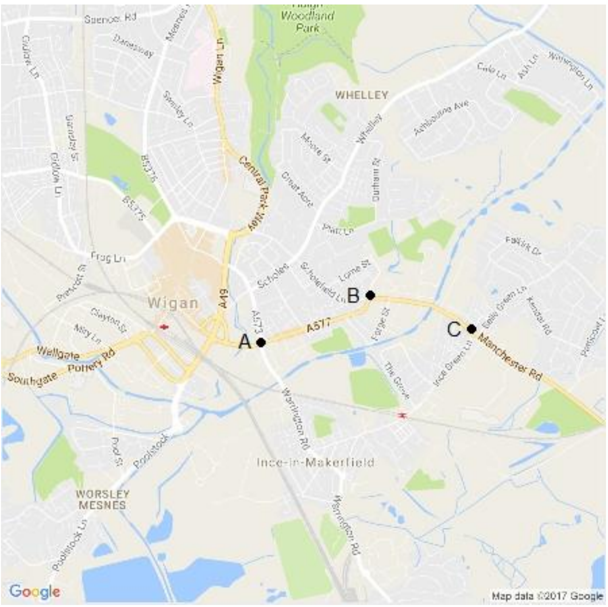

Figure 3-4: Three site case study for Bluetooth device detection

Table 3-1: Summary of the three site case study

Direction	Number of matches (trips)	Proportion observed at B
Eastbound (A to C)	108,915	66%
Westbound (C to A)	142,619	38%

Of the devices detected travelling from A to C, 66% were also detected at B. This is slightly lower than the 80% detection rate found by Araghi et al. (2014), although they did report that the detector position could have an effect on the rate. Although this is a relatively large percentage, it will also apply when devices pass both A and C. If we could assume that the likelihood of a Bluetooth device being switched off or the discoverability settings being changed while travelling along this road is minimal and that the conditional probability of being detected at one site given that the device passed it is independent of the conditional probability of being detected at the other site given that the device passed it, then only 44% of the discoverable Bluetooth devices passing both A and C would have been detected at both sites. If, however, some devices have a tendency more than others to be undetected then the conditional probabilities will not be independent, and we would expect to detect a higher percentage passing both points.

Table 3-1 also shows that 38% of devices travelling from C to A (i.e. from east to west) were detected at site B. The most likely explanation for this is the placement of the detector. It is on the north side of the road and therefore due to left hand traffic and three lanes of traffic at this location, the detector is much closer to the eastbound traffic. A lane bias effect in Bluetooth data may have been identified by Colberg et al. (2014), although the bias may have been due to the speed of vehicles rather than their road position.

To test whether all Bluetooth-enabled devices have an equal probability of being detected when passing a sensor, the detection probability at B was also calculated for each regular traveller between A and C. Devices which were observed travelling between A and C (in either direction) 25 times or more within the six month period were analysed. In total, 826 devices were detected travelling regularly eastbound and 963 devices travelling westbound. This included 602 Bluetooth devices which travelled regularly in both directions. The device-specific probabilities of being detected at site B ranged from 0 to 100% and no distinct patterns were visible in the data. Devices could have different probabilities of being detected due to differences in the type of device, differences in the owners' driving style (including speed and road position) or systematic differences in the placement of the devices within cars.

3.3.4 Scope of data collection

News stories have highlighted the widespread use of Bluetooth within rental cars (for example USA Today (2015)). Whilst this is not a problem when using Bluetooth data for travel time estimation, it could be a problem when considering intrapersonal variability. Millard et al. (2016) found that only 16% of car rentals are for more than a week and therefore Bluetooth devices in these cars are likely to demonstrate more variability over longer periods of time as multiple customers would be associated with the same unique MAC address. In Wigan, there are car rental establishments in the town centre, to the east and the southwest and there are also places which hire out vans. In this area, therefore, it is not possible to isolate Bluetooth sensors which will detect a higher proportion of rental vehicles. Rentals are, however, likely to make up a very small percentage of the Bluetooth sample in most cases.

In this research, only trips made by motorised vehicles on the road network are included. This could include car or motorcycle drivers, or alternatively it could be passengers in cars, buses or taxis. Whilst this may be of interest to road managers, for more general transportation purposes data covering all modes may be required. Bluetooth data is not limited to the analysis of road transportation data, for example it has been used for pedestrian analysis (Delafontaine et al., 2012, Malinovskiy et al., 2012, Versichele et al., 2012) and cyclist travel times (Mei et al., 2012). The difficulty arises in using Bluetooth to collect data on multiple modes. It is not usually possible to differentiate between a trip made by car with a stop en-route, and a trip by a slower mode such as cycling. A more complex trip cleaning procedure would be required, for example a clustering stage as proposed by Araghi et al. (2012), to generate trip data together with the most likely mode for each trip.

Another shortcoming of Bluetooth data compared with travel diary data is that personal travel cannot be separated from business travel. The National Travel Survey in England, for example, collects data only on “travel for private purposes or for work or education, provided the main reason for the trip is for the traveller himself or herself to reach the destination” (Department for Transport, 2016a) . With Bluetooth data, however, this distinction cannot be made. Although this is unlikely to be problematic for estimating travel times, it will affect estimates of trip frequency.

Elango et al. (2007, p39) identified that vehicles in their survey which were used at least occasionally for business purposes showed very different travel patterns to vehicles used for personal trips only. Data which relates to the drivers of buses, taxis and delivery vans, or even to the vehicles themselves, can skew results by suggesting high trip frequencies, in the first case with great regularity and in the latter two cases far less. The maximum number of trips per device over the year within the Wigan area was 2,355. This is over six trips per day for the entire year on average. This is high for personal travel in England, where the average number of trips per year by all modes combined was found in 2016 to be 914 (Department for Transport, 2016b). Although devices with very high trip frequencies could be removed from the data depending on the scope of the study, the difficulty is in deciding the threshold separating regular travellers from vehicles with at least some business use. Setting this threshold too low risks introducing bias by removing very frequent travellers. An alternative approach could be to retain all devices but to be very clear when interpreting the results that the sample includes all types of road users, including buses and taxis. This could provide valuable insights for road network managers which are not available from other surveys of personal travel only.

3.3.5 Additional work required

Bluetooth can, therefore, provide data on a relatively large sample of trips on the network over a long period of time for a relatively low cost. There are unresolved issues regarding potential bias in the data collected, particularly due to age and socio-economic status, which are likely to have a greater impact on measures of intrapersonal variability than travel time estimates. These are not insurmountable problems, however, although more information is required on the characteristics of people carrying various types of Bluetooth devices and the way they use them. If this information was available, the Bluetooth data could be boosted by alternative data collection for under-represented groups. While it may be relatively easy in principle to undertake a GPS travel diary for cars, perhaps using a similar methodology to Elango et al. (2007), it would be challenging to collect data which is consistent with the Bluetooth data. Difficulties could arise in deciding on the geographical boundaries for participants' homes, on the length of the data collection period and in recruiting participants from harder to reach groups.

Research into how people use Bluetooth devices could inform adjustment factors to correct under-reporting of trips due to not carrying devices, switching them off or not having them in discoverable mode. Additional data can also be collected alongside MAC addresses in order to provide additional information about the travellers detected, for example the device type or the brand of device. While this is not practical to do on an ongoing basis, even collecting the data for short periods such as a day or a week would be informative. The proportion of discoverable Bluetooth devices on the road network could also be increased through a publicity campaign encouraging travellers to do so. Sharing of data could be encouraged by providing incentives, for example tailored real time travel information.

The number of devices and trips detected could be increased in other ways too. Firstly, the roadside position of sensors should be carefully considered so as to optimise detections. Brennan et al. (2010) provide information on the impact of the height of placement but also state that with multiple lanes of traffic, sensors on both sides of the roads may be required. The quantity and quality of the trip frequency data will increase if additional Bluetooth sensors are installed at more regular intervals along links and/or routes of interest. This differs from the requirements for travel time estimation where mid-link detectors will not provide additional information and the benefit gained from increasing sample sizes at existing locations can be marginal.

Alternatively, the antennae strength could be increased, to increase the size of the sensor's detection zone and thus the probability of detection. This will decrease the spatial and temporal precision of the observations and therefore is not advisable where the detector will also be used for travel time estimation. For intrapersonal variability, however, the benefit of getting additional data outweighs any small decrease in precision.

Bluetooth data collects information about personal and commercial travel combined. Although this results in comparability issues with other data sources, for example national travel surveys, it could be useful in providing road managers with a more holistic view of road users.

3.4 Time of day trips are made

The time of day at which trips are undertaken is an important aspect of intrapersonal variability. Analysing the variability in the timing of each person's trips can provide a better understanding of how and when travellers use the transport infrastructure. Further work would be required to identify whether travellers with greater variability have greater flexibility in choosing their departure time or whether this occurs because they have greater variability in the needs and constraints underlying their behaviour.

The time of day trips are made is not always included in research on overall measures of intrapersonal variability. Sometimes it is included very broadly, for example in Bayarma et al. (2007) where the only relevant measure is the proportion of trips made between 10pm and 6am. Where intrapersonal variability in the timing of trips is considered, it usually involves separating the day into equally sized bins and then comparing the activity being undertaken in each, as in Minnen et al. (2015) and Goulet Langlois et al. (2016).

The times which can be measured using fixed Bluetooth detectors differ from typical travel diary data in two important respects. First of all, we cannot measure or even estimate an individual’s departure time. What the data can tell us, however, is the regularity at which a device passes a particular location. The reason for any variability could be differences in departure time and/or differences in traffic conditions encountered prior to passing the detector. This may initially seem problematic, but if we consider the viewpoint of local road network managers, then it may not be. Their priorities lie with road users at the point that they enter their jurisdiction and on critical links within the network. The road managers can, therefore, design the Bluetooth detector placement in order to collect information on the important parts of trips only. The use of fixed detectors also means that analysts can examine the variability in the times of day at which a road user passes a fixed point, for example a pinch point such as a bridge or tunnel or a location related to an intervention such as a charging cordon.

The second important aspect in which the recorded times differ, is that unlike travel diaries they do not depend upon the traveller accurately recalling and recording their times. Bluetooth detections are made within a zone surrounding a sensor rather than at a specific point, however, and the accuracy of the data has been considered in detail in relation to estimating travel times (Araghi et al., 2014). For repeated trips, a consistent method for estimating the recorded time at a sensor location is required to ensure that observations at this detector are comparable across days. This is because the sensors are continuously scanning for discoverable Bluetooth devices within a zone of up to 100m for 9dBi gain antennae and therefore each device may be detected several times while passing the sensor, depending on the sensor location, road layout, strength settings and traffic signal cycle lengths, if appropriate. One commonly used approach is to use the first time a device is recorded as it passes through the detection zone. The times will be recorded in continuous time, as opposed to times reported by travellers who automatically tend to register rounded times only (Minnen et al., 2015).

The time of day trips are made can only be meaningfully compared where the trips themselves are comparable. For example, Muthyalagari et al. (2001) considered the times people first leave home, depart work and arrive home each day and Kitamura et al. (2006) considered the first trip per day on all weekdays for workers only. Chikaraishi et al. (2009) went further in their examination of intrapersonal variability in departure times by estimating separate multi-level models by trip purpose and including variables relating to spatial (origin-destination) and temporal (day of the week or season) trip characteristics. With Bluetooth data, trip purpose information is not available and neither is true origin-destination data (see Section 3.5).

One aspect which could be considered is the variability in the time at which each device is *first* detected each day, as used by Muthyalagari et al. (2001) on GPS travel diary data. For the case study area, let us consider the 1,240 devices identified in Section 3.3.3 as being regular travellers. The first detection each day does not appear to be particularly meaningful, as Figure 3-5 shows that some devices have fairly large standard deviations. This could be due to the geographic limitations of the data as these are only the first trips *within the study area* each day and therefore may not be as comparable as the travellers' actual first trip each day.

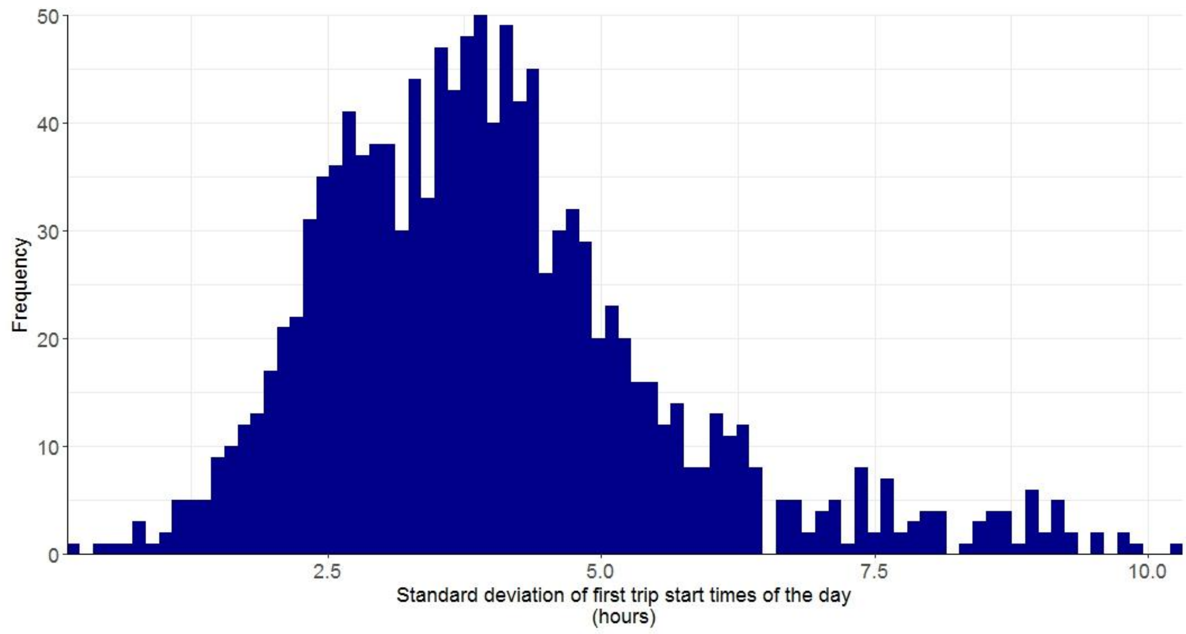

Figure 3-5: Histogram of intrapersonal variability in the first detection time each day

To obtain more comparable trips, the variability in each traveller's trip start times could be calculated within a fixed portion of the day only. For example intrapersonal variability in the start times of trips within the morning peak (between 7am and 9am) is shown in Figure 3-6. This could be assumed to relate to trips for a similar purpose, for example work or education, but it is not clear how the boundaries between different times of the day should be drawn.

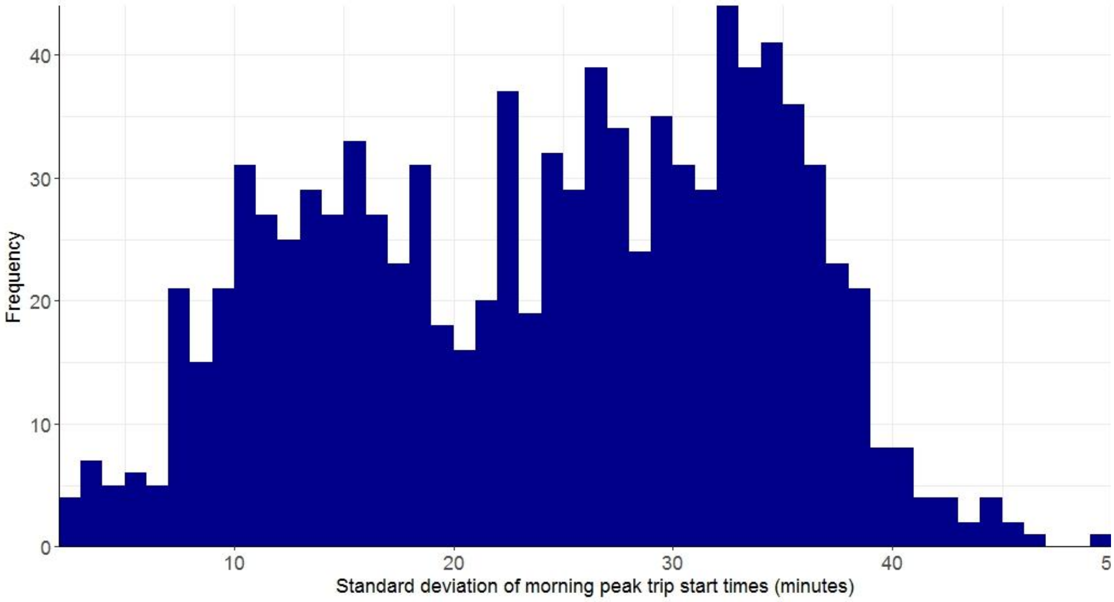

Figure 3-6: Histogram of intrapersonal variability in the time of all trips beginning in the morning peak (7-9am) for regular travellers with 10 or more morning peak trips

It could be hypothesised that people who travel more frequently have a more regular routine and therefore travel at a similar time each day, resulting in less intrapersonal variability. Figure 3-7, however, does not support this hypothesis in the current data. There is not a strong relationship between each traveller's time of day variability and the number of days they are observed, during the morning peak period. The associated correlation coefficient is -0.4.

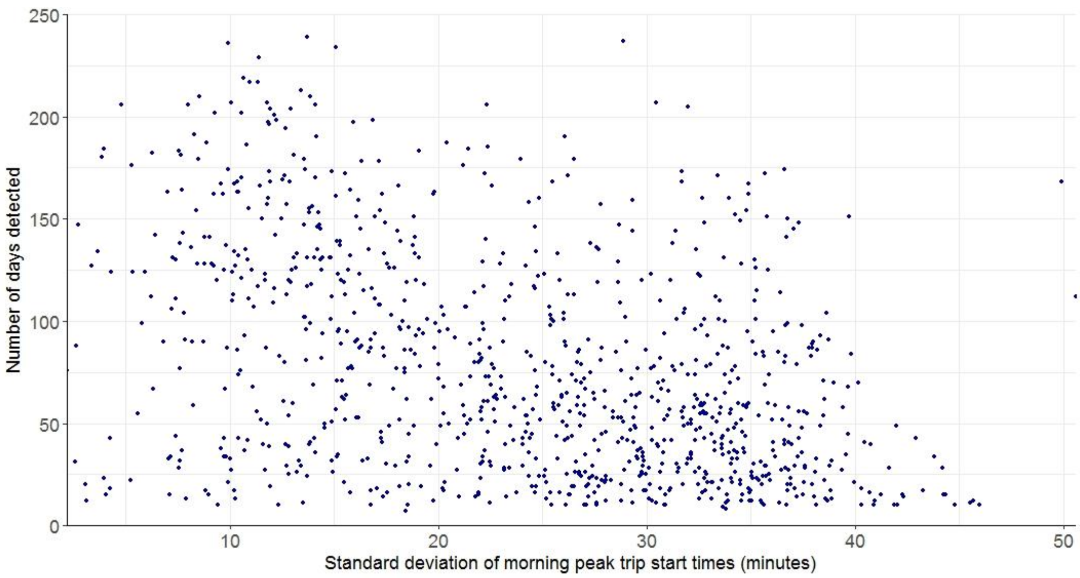

Figure 3-7: Scatterplot of intrapersonal variability in morning peak (7-9am) trip start times and the number of days observed for each traveller

Rather than separating trips by the time of day, an alternative approach could be to compare trips with the same spatial characteristics, i.e. trips with the same start and end detectors. There are two issues with this approach. Firstly, this is likely to result in many time of day variability measures per traveller. This increases the computational burden of the analysis but, perhaps more importantly, it also adds an additional level of complexity to the interpretation of the results. Separating the trips from the 1,240 regular travellers in our case study by OD pair results in 24,312 sets of data to analyse. Secondly, there needs to be a sufficiently large sample for intrapersonal variability in the times to be measured. Specifying a minimum sample size of just 5, reduces the 24,312 subsets to 12,121 in the case study example. A minimum sample size of 10 results in 8,114 subsets remaining. This is a ‘loss’ of data, but it could be assumed that these trips are rarely made by the traveller and therefore, in terms of the intrapersonal variability, they are irrelevant. Such trips would, however, still be important for the analysis described in Section 3.5 which focuses on the spatial variability of trips. An alternative option could be to focus on one particular location which is of interest, for example the time of day that travellers arrive at a bridge toll station.

Once subsets of trips for comparison have been specified, a measure of intrapersonal variability in the times needs to be calculated. As the times in Bluetooth data are continuously distributed, a very simple example of taking the standard deviation of the data was used in Figure 3-5 and Figure 3-6. It is easy to think of common travel patterns where this would not be a sensible approach. For example, some people work half days on some days of the week which would result in the histogram of trip start times from work to home having two different peaks. Alternatively, some people may make the same trip multiple times per day, for example to take children to and from school. The analysis could be undertaken separately on different times of the day, for example the morning peak used in Figure 3-6. There is no guarantee, however, that the separate peaks in trip timings would occur in different time of day categories and the thresholds for splitting up days would be relatively arbitrary.

An alternative option could be to consider data from all times of day together, but to assume that the observed times for an individual traveller correspond to a mixture distribution where the subpopulations relate to systematically different activity patterns and the subpopulation variances tell us about variability in the timing of trips relative to each activity pattern. Model based clustering techniques can be used to fit the most appropriate mixture distribution to the data (Fraley and Raftery, 2002). This approach has been used in other aspects of transportation research, for example flight delays (Tu et al., 2008) and errors in loop detector sensitivity (Corey et al., 2011). The number of clusters and their variances could be compared before and after an intervention for each traveller or across all travellers for one period of time.

Figure 3-8 includes examples of trip timing histograms together with the associated Gaussian mixture distribution for two travellers from the case study area. In each case the trips are for a single OD pair and at least 50 trips are included. For the first traveller, the components of the mixture distribution are overlapping which can be useful for representing fairly complex patterns. This traveller has quite a lot of variability in the timing of trips during the late afternoon, but the large peak around 6pm suggests a constraint limiting how late the trip can occur, in most cases. The plot for the second traveller is dominated by a peak in trips in the morning, with only a small component representing trips taken later in the day. While this traveller does not appear to have systematic differences in the timing of this trip, the variance of the primary component provides information about the random variability in the timing of the trip.

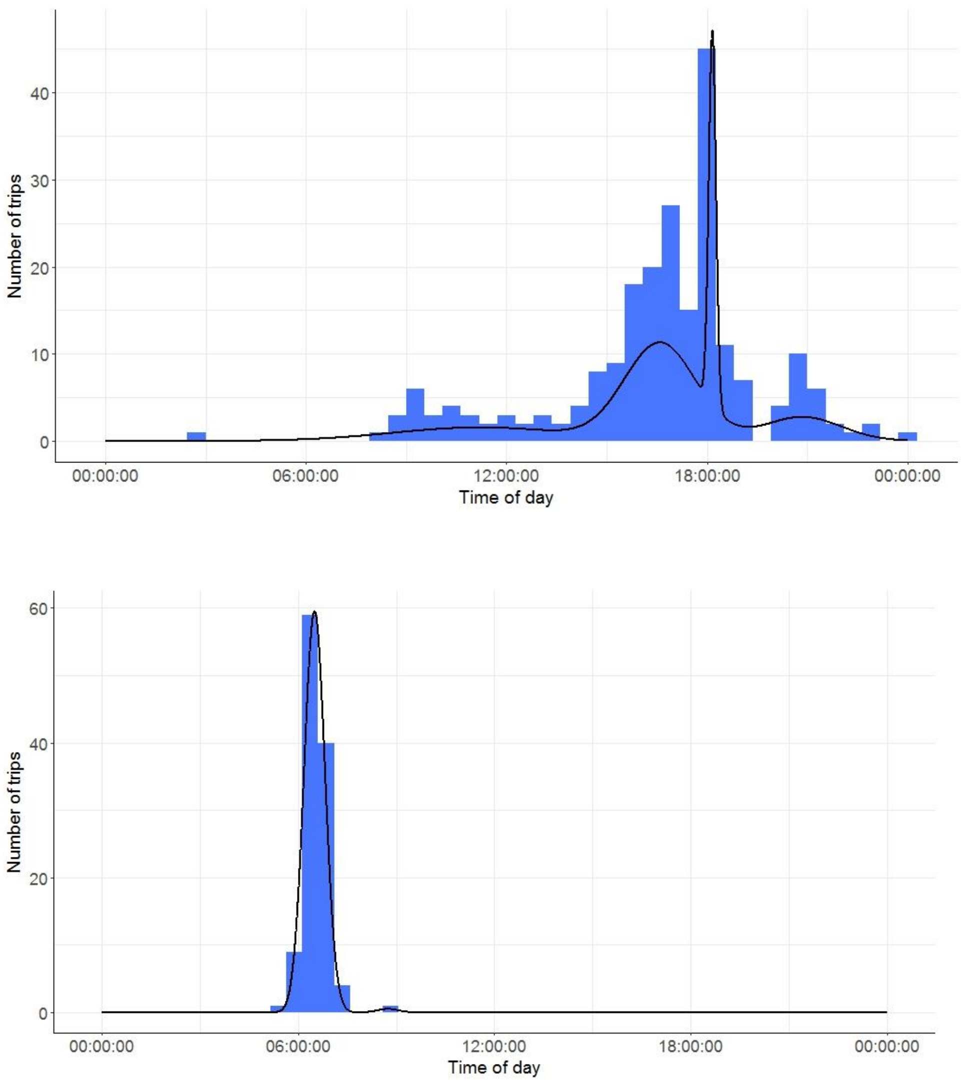

Figure 3-8: Time of day histograms and estimated mixture distributions for trips on one OD pair for two travellers

In summary, Bluetooth data can provide traveller arrival times at specified locations on the road network reported on a continuous scale. These locations are fixed over time and do not have to be the travellers' trip start or end points, thus providing road managers with control over the parts of the network they wish to collect data about. As no data is collected on trip purpose, explanations behind variations in travel times would need to be explored through additional data collection, for example user interviews or focus groups.

3.5 Spatial variability

Traditionally, the spatial dimension of trips is defined by origin-destination and route choice information. Origin-destination information tells us about traveller needs as it represents the places they need to be at different times of the day and week. In most cases, these needs are determined exogenously from the transport system and therefore to have an impact upon these choices external policies and initiatives are often required, including business opening hours and land use policies. Route choices are determined by traveller decisions based on the information available to them. This information could come from many sources including their own experiences and preferences, satellite navigation systems or radio alerts. Due to these distinctions, in this section the two aspects of spatial variability will be considered separately.

3.5.1 Origin and destination of trips

Fixed Bluetooth detectors can inform the construction of OD matrices, as shown by Barceló et al. (2010) and Carpenter et al. (2012), although in these two examples the data relates to one corridor only. As the current paper focuses on intrapersonal variability, it is the distribution of each individual's trips across the OD matrix which is of interest.

Similar to departure times discussed in Section 3.4, it is not normally possible to identify trip origins or destinations using only data from fixed Bluetooth detectors as they only monitor certain links within a network. It is, however, possible to analyse the first and last Bluetooth detector location recorded for each trip. Given that the detector locations are chosen by the relevant road managers in order to monitor important links within their jurisdiction, the first and last detection locations for each trip give estimates of where the traveller is going from and to, within this specific part of the network. Therefore, as in some of the research on smart card data, for example Goulet Langlois et al. (2016) and Kieu et al. (2015b), the first and last detector locations will be referred to as an OD pair in this paper. In smart card data research this is interpreted as the origin and destination stops used on the public transport network and in the current paper it should be interpreted as an estimate of the origin and destination for the part of the trip undertaken by a motorised vehicle within the monitored area. Ensuring the data has been satisfactorily processed is essential as the first and last Bluetooth detections for trips may be sensitive to the trip chaining undertaken during the data processing (see Section 3.2.1).

In this section, the case study sites are considered as 7 separate locations by treating the two sites in the centre of Wigan as one location. This is because general locations as opposed to exact routes are of interest in this section. An OD matrix of the trips detected by the Bluetooth sensors can be constructed for each of the 1,240 devices which were detected in the case study area at least 260 times in the year of data. These matrices can then be used as inputs to calculate measures of spatial intrapersonal variability.

One measure of spatial intrapersonal variability could be the number of OD pairs observed per device over the year. Figure 3-9 demonstrates the wide range in the number of OD pairs observed among the regular travellers over a one year period. The mean number of OD pairs observed per device, for these regular travellers, is 24. The analysis could be taken further by examining the number of ODs observed on weekdays versus weekend days, different days of the week or months of the year in order to explore whether some of the variability in ODs observed is due to systematic differences in activity patterns.

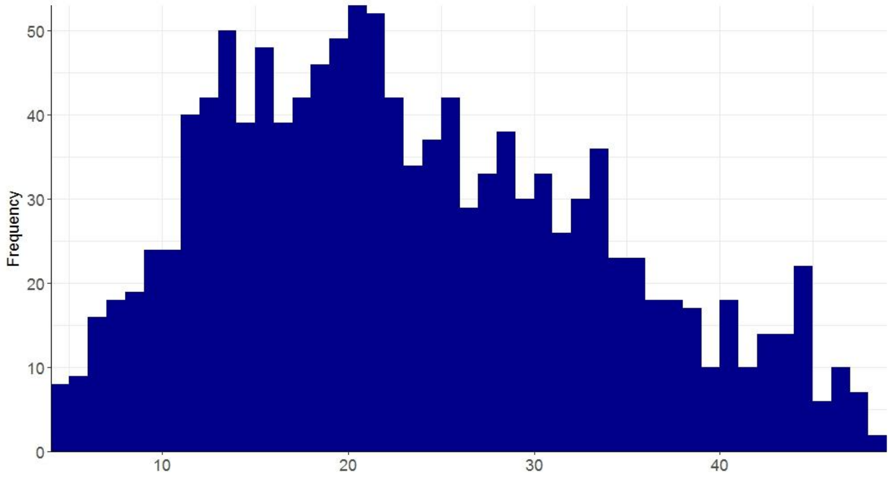

Figure 3-9: Histogram of the number of OD pairs observed per device

The number of OD pairs observed for each traveller is likely to be related to their total number of trips. As shown in Figure 3-10, this is not a linear relationship as the total number of trips increases at a greater rate when 40 or more OD pairs are observed. This could be because these travellers with a high degree of spatial variability are travelling for work, for example making deliveries, which also causes them to make a large number of trips.

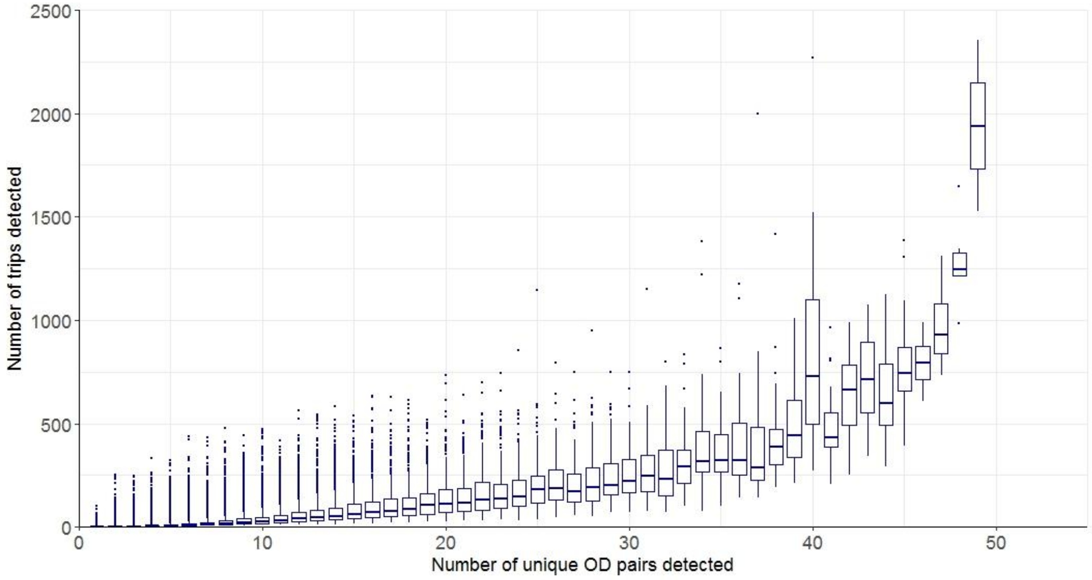

Figure 3-10: Boxplot of the total number of trips observed for travellers grouped by the number of different OD pairs they record

Rather than considering the total number of trips made by each traveller, it may be more appropriate to consider how frequently each traveller makes each OD trip. OD pairs for each traveller could be categorised based on how frequently they are observed. For example, Figure 3-11 shows the proportion of each traveller's trips which are made between OD pairs used frequently and the proportion used rarely, for the regular travellers identified in the case study area. While some travellers make nearly all of their trips between their frequently used ODs, others do not have any OD pairs which they use on a weekly basis. These measures of traveller OD usage could be used to identify clusters of travellers with high or low spatial intrapersonal variability

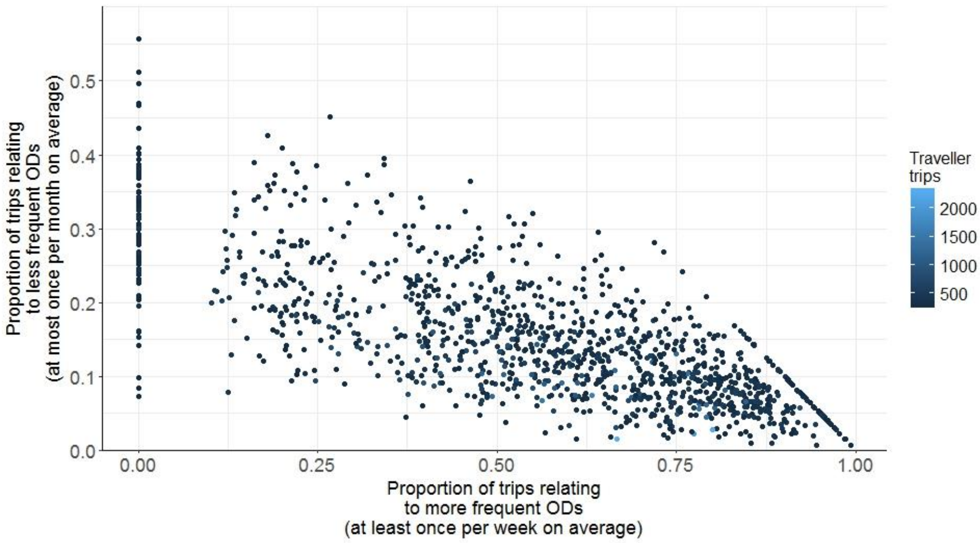

Figure 3-11: Scatterplot showing, for each traveller, the proportion of their trips relating to ODs used at least once per week and at most once per month

For larger and more complex networks, an alternative strategy may be required due to the large number of possible OD pairs, for example combining sensor locations into zones.

An alternative to treating a traveller's trips as independent observations could be to consider each day for a traveller as one observation. By doing so, the complexity of each individual's travel patterns can be measured and compared. Figure 3-12 provides a crude approximation of the complexity of the travel patterns of the regular travellers in the case study area. For each traveller, for every day they are observed, the number of trips in that day have been counted. These counts were then aggregated across all travellers. In the case study area 76% of the traveller days with an observed trip included either 1 or 2 trips. This could be due to the small size of the case study area or perhaps due to the more traditional travel patterns which might be expected in this town. Of the days where a traveller made two trips, 38% consisted of an outbound and return trip with coinciding start and end points. This comparison is not necessarily meaningful, however, as trips may not match due to differing routes to origins or destinations outside the monitored area or missing observations. Meaningful daily patterns for individual travellers are also less likely to be found in Bluetooth data due to the limited network coverage. For example, unlike the highly structured weekly patterns observed in smart card data by Goulet Langlois et al. (2016), in the case study area in the current paper only 37% of travellers had their first detection site matching their last detection on that day on at least half of the days they were observed.

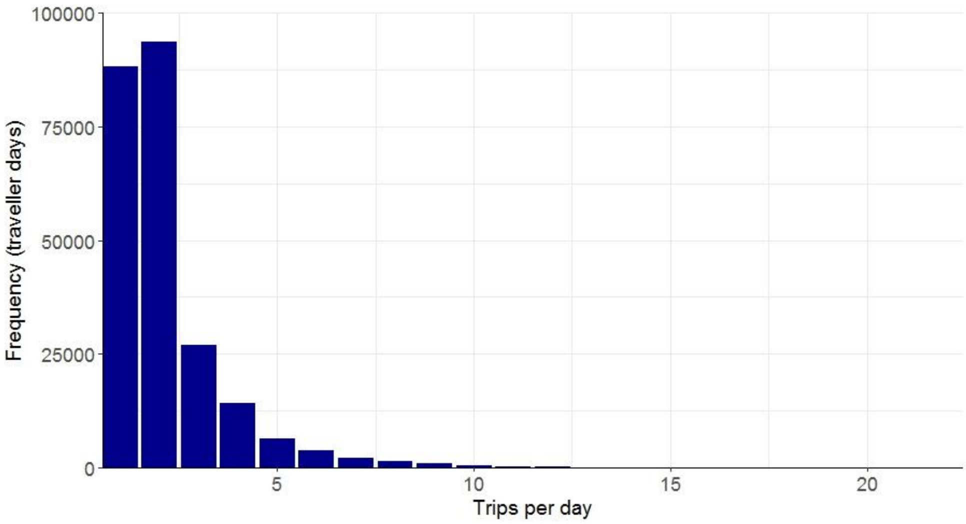

Figure 3-12: Histogram of trips for each traveller per day

3.5.2 Route Choice

Carpenter et al. (2012) have demonstrated that with strategic Bluetooth detector placement, the data can be used to record route choices. Whilst in Carpenter et al. (2012) the aim was to ‘scale up’ this data into OD matrices, in this paper the focus is on measuring intrapersonal variability in those route choices.

The case study used in the current paper provides good coverage of trips into and out of Wigan, but does not present many opportunities for exploring route choice. To explore the usefulness of Bluetooth for route choice analysis, therefore, an alternative case study was identified (Figure 3-13). This may not necessarily be a commonly made trip, but it was selected due to the presence of Bluetooth detectors on alternative routes. To drive between detector H and detector D, there are a number of possible routes, ranging from 4.2 to 5.5 miles long. The shortest routes all go through a short section of Chapel Lane in the town centre halfway between the two locations, but from here there are two possible routes to both H and D. Only one of the routes on the eastern branches has a Bluetooth detector, but both of the western branches have Bluetooth detectors.

Figure 3-13: Route choice case study site

By analysing four months[3] of trips matched between D and H, the proportion of trips which can be assigned to a particular route can be calculated. In total, 3,012

[3] This is the longest period of time in 2015 for which data from all sites was available

trips were recorded travelling westward and 3,029 travelling eastward. Of the 3,012 travelling in a westerly direction, 1,655 trips (55%) were detected at G. Due to the incomplete nature of Bluetooth data (as discussed in Section 3.3.3), we cannot assume that all of the remaining trips were made on the alternative route. However, for comparative purposes, for example before and after an intervention, the stability of the route choice proportions can be tested, as in Watling et al. (2012, p179).

The westerly branches have better coverage as there are Bluetooth detectors on the two most likely routes. Table 3-2 shows the observations at mid-trip detectors relating to trips between D and H.

Table 3-2: Trips detected at interim Bluetooth sites

	East to West (H to D)	West to East (D to H)
Detected at E only	428 *(14%)*	409 *(14%)*
Detected at F only	1,415 *(47%)*	886 *(29%)*
Additional trips detected at both E and F	15 *(0.5%)*	10 *(0.3%)*
Total trips detected	3,012	3,029

Clearly not all devices were detected at either E or F. This could be because another major route between D and H has not been monitored, or it could be because of missing data. Similarly, it is only possible to speculate as to whether the difference in proportion of trips detected at one of these sites between the eastbound and westbound direction (43% versus 62%) is due to a directional difference in route choice or in the probability of being detected at the interim site, for example due to detector placement.

The current analysis includes route choices at all times of the day (and night). The alternative routes do not have a constant relationship with one another, in terms of travel times, throughout the day and hence it is reasonable to assume the route choice proportions will change throughout the day. Future work is required to incorporate such factors into analysis of Bluetooth data for route choice.

Intrapersonal variability in route choice for the travellers observed in the four month period can also be examined. Considering eastbound trips only, 1,866 travellers were observed, 352 of whom were observed more than once. Intrapersonal variability in route choice on the branches to the west of the town centre is considered in Figure 3-14. Clearly some travellers prefer to travel via site F and a small proportion prefer to travel via site E when driving between sites H and D. Almost 29% of the travellers are never observed at either of these interim sites. The ‘sometimes at one site’ category accounts for 24% of the travellers and could include travellers who sometimes take a route not covered by Bluetooth detectors or it could be due to missing observations. The travellers who vary their route account for 13% of the travellers and this category has the highest average number of trips per traveller, at 7.5. This compares to 2.3 trips for travellers always going via E and 3.8 trips for travellers always going via F.

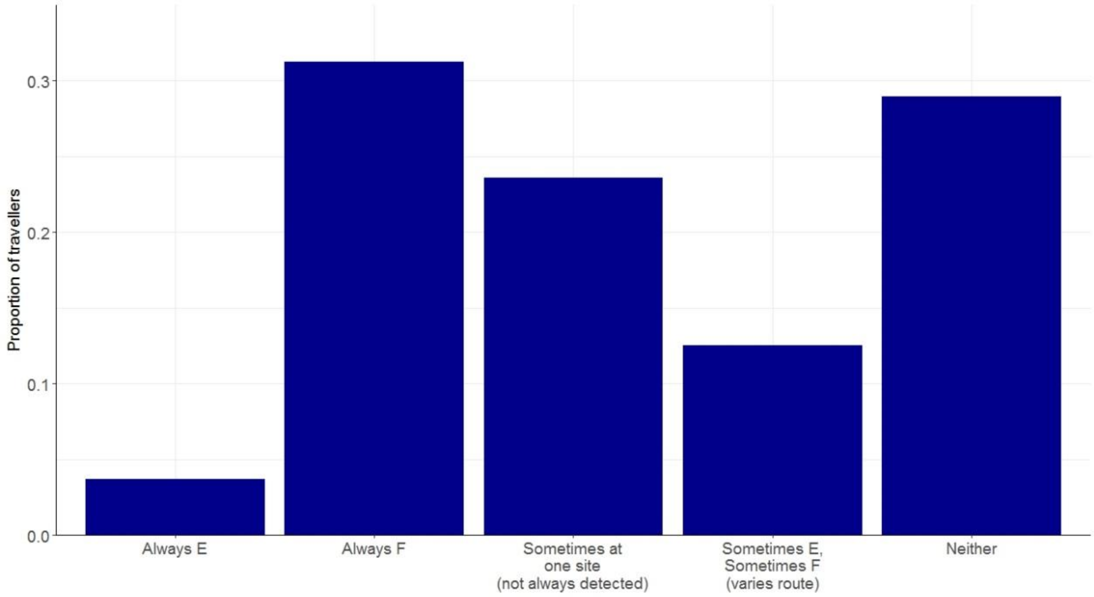

Figure 3-14: Traveller route choice for trips from H to D (travellers recording just one trip have been excluded)

Depending on the placement of Bluetooth detectors, therefore, the data can be used to examine intrapersonal variability in route choice. As demonstrated by Hainen et al. (2011), Bluetooth can be used to monitor traveller response to network disruptions such as bridge closures. In such cases, Bluetooth detectors could be installed temporarily on routes of interest to provide insights into behavioural response.

Whether Bluetooth data is being used to measure intrapersonal variability in route choice for all travellers or to test for systematic differences in route choice before and after an intervention, the same considerations need to be reflected upon. Firstly, are the detectors placed in suitable locations to collect the required data? This includes coverage of routes and also the more precise positioning of the detector, for example the side of the road it is on. Local road user interviews and focus groups could be used to identify whether additional options should be monitored. Secondly, are there sufficient detectors to collect data of the required quality? While mid-link detectors are rarely used for travel time estimation, their inclusion could boost detection rates on key routes of interest for intrapersonal variability.

3.5.3 Overall spatial analysis

The previous two subsections focus on the two aspects of spatial variability traditionally considered in transportation research – ODs and route choices. Both of these approaches are challenging, perhaps because they are not designed specifically around the advantages and issues relating to Bluetooth data.

As discussed in Section 3.3.3, discoverable Bluetooth devices do not have a 100% probability of being detected when they pass a sensor. This raises concerns regarding the use of an OD based technique which assumes that the first and last sensor *detection* of each trip corresponds to the first and last sensor *passed* each trip. It may, therefore, be preferable to undertake a combined spatial analysis of the data which incorporates the start, end and also interim detections.

One option for the combined approach could be to use Sequence Alignment, which was originally developed for analysing biological structures but has been used more recently in the social sciences (Abbott, 1995). This has been used in geographical research to explore tourist movements (Shoval and Isaacson, 2007) and has also been used on Bluetooth detections of trade fair visitors (Delafontaine et al., 2012). Sequence Alignment maintains the order of observations, for example of a Bluetooth device moving through a network, and the analyst can choose how to compare each pair of sequences to get a measure of similarity. Once these similarity measures have been calculated for all pairs in the data, they can be used to cluster ‘similar’ sequences. Various measures of similarity could be used, including local alignment techniques such as Longest Common Subsequence which was used by Kim and Mahmassani (2015) on taxi GPS traces, or global alignment techniques which compare similarities and differences along the full sequences as in Crawford et al. (Under review-b). By identifying clusters of similar sequences, spatial intrapersonal variability can be calculated based on the distribution of trips between clusters for each traveller.

3.6 Conclusions

In this paper the suitability of data from fixed Bluetooth detectors as a means of collecting data on intrapersonal variability has been discussed. As with any source of data, there are limitations and issues to reflect upon, including missing data and bias in the sample. Unlike personal travel diaries, Bluetooth data will not provide information about trip purpose, complete trip ODs nor traveller demographics. The combination of a potentially biased sample with respect to age and income and a lack of traveller information to measure the characteristics of the travellers detected is problematic. Further research is needed to understand the usage of Bluetooth devices other than mobile phones, particularly those related to motor vehicles and also on behaviours relating to disabling Bluetooth on devices. In areas where Bluetooth data is being collected, it would be advisable to undertake sample scans so that the proportion of MAC addresses from each type of device, for example mobile phones, can be estimated.

Despite the issues discussed above, data from fixed Bluetooth detectors has a vast amount of potential for collecting data on repeated trip behaviour on the road network. Large amounts of data can be collected through installing relatively cheap detectors and, as with the data used in the case study in this paper, if the detectors were installed for the purpose of measuring travel times then the only costs in measuring intrapersonal variability relate to the analyses. The cost of collecting alternative types of data is emphasised by the repeated use of a few multi-day travel diary datasets in the intrapersonal variability literature. For example, data from the Mobidrive survey, undertaken on 139 households in Germany in 1999 has been used frequently in the academic literature for more than a decade, including Schlich and Axhausen (2003), Bhat et al. (2004), Kitamura et al. (2006), Bayarma et al. (2007), Habib et al. (2008), Chikaraishi et al. (2009), Tarigan and Kitamura (2009), Cherchi and Cirillo (2014), Susilo and Axhausen (2014). Bluetooth data is not a replacement for such surveys, however, but could be used as a complementary source of data to examine spatial and temporal variability in more detail or changes between survey periods.

Bluetooth data does not collect all of the information in a travel diary. When analysing the data, however, it is important to focus on what *can* be measured using the data and what the objective of the analysis is. For example, Section 3.4 demonstrated how data collected at the same geographic location (chosen by the road manager) on different days can be used to measure intrapersonal variability in the timing of trips. Section 3.5 demonstrated how data on entry and exit points to the part of the road network which is of most interest, together with route choice information on key routes, can be used to explore spatial variability. This research has considered key aspects of intrapersonal variability separately, but to fully understand the variability they could be considered collectively. Such work could use similarity indices such as those discussed by Schlich and Axhausen (2003), or could use clustering to segment the road user population based on intrapersonal variability such as in Crawford et al. (Under review-b).

Further research is required to understand the optimal number and placement of Bluetooth detectors to measure intrapersonal variability in travel on urban road networks. Additional research into the usage and behaviours related to Bluetooth devices should also allow correction factors to be developed which could be used to scale up findings from Bluetooth data to the population of road users.

3.7 References

Abbott, A. (1995) Sequence Analysis: New Methods for Old Ideas. *Annual Review of Sociology,* 21**,** 93-113.

Araghi, B. N., Christensen, L. T., Krishnan, R. & Lahrmann, H. (2012) Application of Bluetooth Technology for Mode-Specific Travel Time Estimation on Arterial Roads: Potentials and Challenges. *Annual Transport Conference at Aalborg University.*

Araghi, B. N., Olesen, J. H., Krishnan, R., Christensen, L. T. & Lahrmann, H. (2014) Reliability of Bluetooth Technology for Travel Time Estimation. *Journal of Intelligent Transportation Systems,* 19**,** 240-255.

Arentze, T. & Timmermans, H. (2005) Modelling Learning and Adaptation in Transportation Contexts. *Transportmetrica,* 1**,** 13-22.

Barceló, J., Montero, L., Marqués, L. & Carmona, C. (2010) Travel Time Forecasting and Dynamic Origin-Destination Estimation for Freeways Based on Bluetooth Traffic Monitoring. *Transportation Research Record: Journal of the Transportation Research Board,* 2175**,** 19-27.

Bayarma, A., Kitamura, R. & Susilo, Y. (2007) Recurrence of Daily Travel Patterns: Stochastic Process Approach to Multiday Travel Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2021**,** 55-63.

Bhaskar, A. & Chung, E. (2013) Fundamental understanding on the use of Bluetooth scanner as a complementary transport data. *Transportation Research Part C: Emerging Technologies,* 37**,** 42-72.

Bhat, C. R., Frusti, T., Zhao, H., Schönfelder, S. & Axhausen, K. W. (2004) Intershopping duration: an analysis using multiweek data. *Transportation Research Part B: Methodological,* 38**,** 39-60.

Brennan, T. M., Ernst, J. M., Day, C. M., Bullock, D. M., Krogmeier, J. V. & Martchouk, M. (2010) Influence of Vertical Sensor Placement on Data Collection Efficiency from Bluetooth MAC Address Collection Devices. *Journal of Transportation Engineering,* 136**,** 1104-1109.

Bricka, S. & Bhat, C. R. (2006) Comparative Analysis of Global Positioning System-Based and Travel Survey-Based Data. *Transportation Research Record: Journal of the Transportation Research Board,* 1972**,** 9-20.

Cantarella, G. E. & Cascetta, E. (1995) Dynamic processes and equilibrium in transportation networks: Towards a unifying theory. *Transportation Science,* 29**,** 305-329.

Carpenter, C., Fowler, M. & Adler, T. (2012) Generating Route-Specific Origin-Destination Tables Using Bluetooth Technology. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 96-102.

Cherchi, E. & Cirillo, C. (2014) Understanding variability, habit and the effect of long period activity plan in modal choices: a day to day, week to week analysis on panel data. *Transportation,* 41**,** 1245-1262.

Chikaraishi, M., Fujiwara, A., Zhang, J. & Axhausen, K. (2009) Exploring Variation Properties of Departure Time Choice Behavior by Using Multilevel Analysis Approach. *Transportation Research Record: Journal of the Transportation Research Board,* 2134**,** 10-20.

Colberg, K., Suh, W., Anderson, J., Zinner, S., Guin, A., Hunter, M. & Guensler, R. (2014) Lane Bias Issues in Work Zone Travel Time Measurement and Reporting. *Transportation Research Record: Journal of the Transportation Research Board,* 2458**,** 78-87.

Corey, J., Lao, Y., Wu, Y.-J. & Wang, Y. (2011) Detection and Correction of Inductive Loop Detector Sensitivity Errors by Using Gaussian Mixture Models. *Transportation Research Record: Journal of the Transportation Research Board,* 2256**,** 120-129.

Crawford, F., Watling, D. P. & Connors, R. D. (Under review) Identifying road user classes based on repeated trip behaviour using Bluetooth data. *Transportation Research Part A: Policy and Practice*.

Delafontaine, M., Versichele, M., Neutens, T. & Van de Weghe, N. (2012) Analysing spatiotemporal sequences in Bluetooth tracking data. *Applied Geography,* 34**,** 659-668.

Department for Transport (2016a) National Travel Survey 2015: Notes and Definitions. Available: https://www.gov.uk/government/statistics/national-travel-survey-2015 [Accessed: 7th April 2017].

Department for Transport (2016b) National Travel Survey 2015: Table NTS0601 Average number of trips (trip rates) by age, gender and main mode: England, 2015.

Department for Transport (2016c) National Travel Survey 2015: Table NTS0708 Travel by National Statistics Socio-economic Classification (NS-SEC) and main mode/stage mode: England, 2015.

Elango, V., Guensler, R. & Ogle, J. (2007) Day-to-Day Travel Variability in the Commute Atlanta, Georgia, Study. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 39-49.

Fraley, C. & Raftery, A. E. (2002) Model-Based Clustering, Discriminant Analysis, and Density Estimation. *Journal of the American Statistical Association,* 97**,** 611-631.

Gerpott, T. J., Thomas, S. & Weichert, M. (2013) Characteristics and mobile Internet use intensity of consumers with different types of advanced handsets: An exploratory empirical study of iPhone, Android and other web-enabled mobile users in Germany. *Telecommunications Policy,* 37**,** 357-371.

Goulet Langlois, G., Koutsopoulos, H. N. & Zhao, J. (2016) Inferring patterns in the multi-week activity sequences of public transport users. *Transportation Research Part C: Emerging Technologies,* 64**,** 1-16.

Habib, K. M. N., Miller, E. & Axhausen, K. (2008) Weekly Rhythm in Joint Time Expenditure for All At-Home and Out-of-Home Activities: Application of Kuhn-Tucker Demand System Model Using Multiweek Travel Diary Data. *Transportation Research Record: Journal of the Transportation Research Board,* 2054**,** 64-73.

Haghani, A., Hamedi, M., Sadabadi, K., Young, S. & Tarnoff, P. (2010) Data Collection of Freeway Travel Time Ground Truth with Bluetooth Sensors. *Transportation Research Record: Journal of the Transportation Research Board,* 2160**,** 60-68.

Hainen, A. M., Wasson, J. S., Hubbard, S. M. L., Remias, S. M., Farnsworth, G. D. & Bullock, D. M. (2011) Estimating Route Choice and Travel Time Reliability with Field Observations of Bluetooth Probe Vehicles. *Transportation Research Record: Journal of the Transportation Research Board,* 2256**,** 43-50.

Han, L., Sun, H., Wang, D. Z. W. & Zhu, C. (2016) A stochastic process traffic assignment model considering stochastic traffic demand. *Transportmetrica B: Transport Dynamics*, 1-21.

Haseman, R., Wasson, J. & Bullock, D. (2010) Real-Time Measurement of Travel Time Delay in Work Zones and Evaluation Metrics Using Bluetooth Probe Tracking. *Transportation Research Record: Journal of the Transportation Research Board,* 2169, 40-53.

Heinen, E. & Chatterjee, K. (2015) The same mode again? An exploration of mode choice variability in Great Britain using the National Travel Survey. *Transportation Research Part A: Policy and Practice,* 78, 266-282.

Huff, J. O. & Hanson, S. (1986) Repetition and Variability in Urban Travel. *Geographical Analysis,* 18, 97-114.

Järv, O., Ahas, R. & Witlox, F. (2014) Understanding monthly variability in human activity spaces: A twelve-month study using mobile phone call detail records. *Transportation Research Part C: Emerging Technologies,* 38, 122-135.

Jones, B. H. & Chin, A. G. (2015) On the efficacy of smartphone security: A critical analysis of modifications in business students' practices over time. *International Journal of Information Management,* 35, 561-571.

Jones, P. & Clarke, M. (1988) The significance and measurement of variability in travel behaviour. *Transportation,* 15, 65-87.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015) Passenger Segmentation Using Smart Card Data. *IEEE Transactions on Intelligent Transportation Systems,* 16, 1537-1548.

Kim, J. & Mahmassani, H. S. (2015) Spatial and Temporal Characterization of Travel Patterns in a Traffic Network Using Vehicle Trajectories. *Transportation Research Procedia,* 9, 164-184.

Kitamura, R., Yamamoto, T., Susilo, Y. O. & Axhausen, K. W. (2006) How routine is a routine? An analysis of the day-to-day variability in prism vertex location. *Transportation Research Part A: Policy and Practice,* 40, 259-279.

Kuwahara, M. & Sullivan, E. C. (1987) Estimating Origin-Destination Matrices from Roadside Survey Data. *Transportation Research Part B,* 21B, 233-248.

Ma, X., Wu, Y.-J., Wang, Y., Chen, F. & Liu, J. (2013) Mining smart card data for transit riders' travel patterns. *Transportation Research Part C: Emerging Technologies,* 36**,** 1-12.

Malinovskiy, Y., Saunier, N. & Wang, Y. (2012) Analysis of Pedestrian Travel with Static Bluetooth Sensors. *Transportation Research Record: Journal of the Transportation Research Board,* 2299**,** 137-149.

Mei, Z., Wang, D. & Chen, J. (2012) Investigation with Bluetooth Sensors of Bicycle Travel Time Estimation on a Short Corridor. *International Journal of Distributed Sensor Networks,* 2012**,** 1-7.

Millard, K., Cairns, S. & Beaumont, C. (2016) Car rental customer survey 2015. Transport Research Laboratory under a contract with BVRLA. Available: http://www.bvrla.co.uk/research/article/rental-customer-survey-2015 [Accessed: 15 September 2016].

Minnen, J., Glorieux, I. & van Tienoven, T. P. (2015) Transportation habits: Evidence from time diary data. *Transportation Research Part A: Policy and Practice,* 76**,** 25-37.

Moghaddam, S. & Hellinga, B. (2013) Quantifying Measurement Error in Arterial Travel Times Measured by Bluetooth Detectors. *Transportation Research Record: Journal of the Transportation Research Board,* 2395**,** 111-122.

Morency, C., Trépanier, M. & Agard, B. (2007) Measuring transit use variability with smart-card data. *Transport Policy,* 14**,** 193-203.

Muthyalagari, G. R., Parashar, A. & Pendyala, R. M. Measuring Day-to-Day Variability in Travel Characteristics Using GPS Data. 80th Annual Meeting of the Transportation Research Board, 2001 Washington D.C.: TRB.

Nakayama, S. & Watling, D. (2014) Consistent formulation of network equilibrium with stochastic flows. *Transportation Research Part B: Methodological,* 66**,** 50-69.

Ofcom (2016) Adults' media use and attitudes. Available: http://stakeholders.ofcom.org.uk/market-data-research/other/research-publications/adults/media-lit-2016/ [Accessed: 9 September 2016].

Office for National Statistics. (2016) *Population Estimates Analysis Tool: Mid 2015* [Online]. Available: https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigr ation/populationestimates/datasets/populationestimatesanalysistool [Accessed: 19 September 2016].

Pelletier, M.-P., Trépanier, M. & Morency, C. (2011) Smart card data use in public transit: A literature review. *Transportation Research Part C: Emerging Technologies,* 19**,** 557-568.

Phua, P., Page, B. & Bogomolova, S. (2015) Validating Bluetooth logging as metric for shopper behaviour studies. *Journal of Retailing and Consumer Services,* 22**,** 158-163.

Safi, H., Assemi, B., Mesbah, M., Ferreira, L. & Hickman, M. (2015) Design and Implementation of a Smartphone-Based Travel Survey. *Transportation Research Record: Journal of the Transportation Research Board,* 2526**,** 99-107.

Schlich, R. & Axhausen, K. W. (2003) Habitual travel behaviour: Evidence from a six-week travel diary. *Transportation,* 30**,** 13-36.

Shao, H., Lam, W. H. K. & Tam, M. L. (2006) A reliability-based stochastic traffic assignment model for network with multiple user classes under uncertainty in demand. *Networks & Spatial Economics,* 6**,** 173-204.

Shoval, N. & Isaacson, M. (2007) Sequence Alignment as a Method for Human Activity Analysis in Space and Time. *Annals of the Association of American Geographers,* 97**,** 282-297.

Stopher, P. R. & Zhang, Y. (2011) Repetitiveness of Daily Travel. *Transportation Research Record: Journal of the Transportation Research Board,* 2230**,** 75-84.

Susilo, Y. O. & Axhausen, K. W. (2014) Repetitions in individual daily activity–travel–location patterns: a study using the Herfindahl–Hirschman Index. *Transportation,* 41**,** 995-1011.

Tarigan, A. & Kitamura, R. (2009) Week-to-Week Leisure Trip Frequency and Its Variability. *Transportation Research Record: Journal of the Transportation Research Board,* 2135**,** 43-51.

Transport for Greater Manchester. (n.d.) *Transport Statistics Greater Manchester Key Centres Section* [Online]. Available: https://data.gov.uk/dataset/transport-statistics-greater-manchester-key-centres-section [Accessed: 19 September 2016].

Tu, Y., Ball, M. O. & Jank, W. S. (2008) Estimating Flight Departure Delay Distributions—A Statistical Approach With Long-Term Trend and Short-Term Pattern. *Journal of the American Statistical Association,* 103**,** 112-125.

USA Today. (2015) *One huge mistake people make when renting cars* [Online]. Available: http://www.usatoday.com/story/tech/columnist/komando/2015/07/03/komando-car-rental-mistakes/29614165/ [Accessed: 22 September 2016].

Utsunomiya, M., Attanucci, J. & Wilson, N. (2006) Potential Uses of Transit Smart Card Registration and Transaction Data to Improve Transit Planning. *Transportation Research Record: Journal of the Transportation Research Board,* 1971**,** 119-126.

Versichele, M., Neutens, T., Delafontaine, M. & Van de Weghe, N. (2012) The use of Bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent Festivities. *Applied Geography,* 32**,** 208-220.

Watling, D., Milne, D. & Clark, S. (2012) Network impacts of a road capacity reduction: Empirical analysis and model predictions. *Transportation Research Part A: Policy and Practice,* 46**,** 167-189.

Zhang, J., Timmermans, H. & Borgers, A. (2002) Utility-Maximizing Model of Household Time Use for Independent, Shared, and Allocated Activities Incorporating Group Decision Mechanisms. *Transportation Research Record: Journal of the Transportation Research Board,* 1807**,** 1-8.

4 Identifying road user classes based on repeated trip behaviour using Bluetooth data

Abstract

Analysing the repeated trip behaviour of travellers, including trip frequency and intrapersonal variability, can provide insights into traveller needs, flexibility and knowledge of the network, as well as inputs for models including learning and/or behaviour change. Data from emerging data sources provide new opportunities to examine repeated trip making on the road network. Point-to-point sensor data, for example from Bluetooth detectors, is collected using fixed detectors installed next to roads which can record unique identifiers of passing vehicles or travellers which can then be matched across space and time. Such data is used in this research to segment road users based on their repeated trip making behaviour, as has been done in public transportation research using smart card data to understand different categories of users. Rather than deciding on traveller segmentation based on a priori assumptions, the method provides a data driven approach to cluster together travellers who have similar trip regularity and variability between days. Measures which account for the strengths and weaknesses of point-to-point sensor data are presented for a) spatial variability, using Sequence Alignment, and b) time of day variability, using Model Based Clustering. The proposed method is also applied to one year of data from 23 fixed Bluetooth detectors in a town in northwest England. The data consists of almost 7.5 million trips made by over 300,000 travellers. Applying the proposed methods allows three traveller user classes to be identified: infrequent, frequent, and very frequent. Interestingly, the spatial and time of day variability characteristics of each user class are distinct and are *not* linearly correlated with trip frequency. The frequent travellers are observed 1-5 times per

week on average and make up 57% of the trips recorded during the year. Focusing on these frequent travellers, it is shown that these can be further separated into those with high spatial and time of day variability and those with low spatial and time of day variability. Understanding the distribution of travellers and trips across these user classes, as well as the repeated trip characteristics of each, can inform the development of policies targeting the needs of specific travellers.

4.1 Introduction

While considering daily snapshots of transport networks is sufficient for many purposes, the benefits of considering the patterns and variability in each individual's behaviour over days, months and even years is receiving increasing research attention. It can inform us about traveller habits (Minnen et al., 2015), predictable differences in travel patterns (Cherchi et al., 2017) and traveller flexibility (Kitamura et al., 2006), all of which are important for developing new policies and modelling traveller *response* to those policies, for example using day-to-day dynamical models which include micro-level learning mechanisms (Chen and Mahmassani, 2004, Liu et al., 2006). Understand the current behaviour of travellers, not just on a single day but over days, weeks and months, also provides information about traveller needs and knowledge of the network.

A common assumption is that urban traffic, particularly the morning peak, consists of commuters who drive between home and work at the same time each weekday. This assumption is often made implicitly and largely for convenience but is rarely challenged despite increases in part time, flexible and home working in recent years. In Great Britain, a 2013 survey (Department for Business, Innovation and Skills, 2014) found that 80% of workplaces with at least 5 employees had part time staff, and other forms of flexible working such as reduced hours, flexitime and compressed hours had all increased since the first comparable survey in 2000. Such an assumption about the regularity of travellers is likely to influence the types of policies formulated to reduce morning peak congestion, some of which may perform differently based on the repeated trip making behaviour of travellers. For example, if the proportion of frequent travellers is overestimated, then the benefits to travellers of switching to public transportation due to savings from weekly or monthly tickets would also be overestimated. Similarly, making an assumption that all travellers have very little departure time flexibility would underestimate the impact of interventions such as time of day specific congestion charging.

One of the reasons why behaviour over multiple days is often overlooked may be the difficulty in collecting data. Detailed information about repeated trip making behaviour has typically been collected using multi-day travel diaries (Muthyalagari et al., 2001, Schlich and Axhausen, 2003, Elango et al., 2007). Such surveys provide data of great depth, but at a cost – both financially and in terms of respondent burden. For example, the National Travel Survey in England involves face to face interviews and 7 day travel diaries for individuals in 7,000 households and costs approximately £2.1 million per year to collect and process (Data.gov.uk, 2012). Respondent burden can be decreased by using GPS devices to track participants (Muthyalagari et al., 2001), but costs remain high, resulting in surveys which are often for short periods of time and/or have small sample sizes. For example, the 7 day travel diaries undertaken annually in England have a relatively large sample size, but sample sizes are usually much smaller for longer surveys, for example the six week Mobi*drive* survey collected in 1999 in Karlsruhe and Halle in Germany had 317 participants in 139 households (Axhausen et al., 2002).

More recently, emerging data sources have been explored to determine their usefulness with respect to measuring repeated trip making behaviour. Mobile phone data has been used to examine activity spaces, as in Järv et al. (2014), but the spatial precision is relatively low. In public transportation research, the availability of smart card data has resulted in researchers identifying different types of user based on their travel behaviour over time (Chu and Chapleau, 2010, Kieu et al., 2015b, Goulet Langlois et al., 2016). Goulet Langlois et al. (2016) analysed four weeks of smart card data from London and identified four types of regular commuter. The daily and weekly activity sequences constructed using the smart card data had distinct patterns for each of these four groups: 'typical' commuters, commuters who sometimes did not take public transportation home at night, commuters who used public transport as their main mode at the weekend and commuters who travelled less during school holiday periods.

The current paper examines data which could be considered the road network counterpart to smart card data, namely point-to-point sensor data, which includes Bluetooth and Automatic Number Plate Recognition (ANPR) data. Point-to-point 'sensors' or 'detectors' collect unique identifiers, either for vehicles or travellers, at fixed locations. It is this "re-identification and tracking" ability which defines this type of data (Antoniou et al., 2011, p140) and as the unique identifiers can be matched over space and time, the data is ideal for examining repeated trip making. Where point-to-point sensors are permanently installed, the amount of data collected can quickly become very large. For example, in Section 4.3 an application to one year of data from just 23 detectors is presented, and that data contains almost 7.5 million trips. These trips are obtained from processing 29.7 million observations, each of which corresponds to a Bluetooth device passing a detector.

Analysing such data with respect to repeated trip behaviour as a whole is not straightforward, however. Previous research on repeated trip making has usually focused on a single aspect, for example trip frequency (Elango et al., 2007, Tarigan and Kitamura, 2009), spatial variability (Buliung et al., 2008, Järv et al., 2014), time of day variability (Kitamura et al., 2006, Chikaraishi et al., 2009) or mode choice (Cherchi and Cirillo, 2014, Heinen and Chatterjee, 2015). Other research has combined different aspects to create a single measure of intrapersonal variability (see Schlich and Axhausen (2003) for an overview). Calculating a single Similarity Index for travellers can be limiting, however, as it cannot account for travellers which differ in terms of different aspects of variability, for example travellers whose trips are spatially predictable but unpredictable in terms of the time of day at which they occur. The current paper uses cluster analysis to segment travellers based on measures relating to multiple aspects of intrapersonal variability, as has been done for public transport users (Goulet Langlois et al., 2016) and with travel diary data (Bayarma et al., 2007). The methods proposed to measure the different aspects are distinctive from previous work, however, due to the nature of the data available from point-to-point sensors. Firstly, point-to-point sensor data does not generally provide origin or destination information due to limited network coverage and the possibility that many trips start and/or end outside the monitored area. It does not provide information about trip purpose either. This means that existing approaches for measuring spatial variability are not suitable. Existing approaches include measuring the distance travelled from home (Bayarma et al., 2007) and comparing daily activity sequences (Goulet Langlois et al., 2016). Secondly, point-to-point sensor data can provide some route choice information, depending on sensor locations, and it would be preferable to have a methodology which takes this additional information into account. Thirdly, for time of day variability, adjustments need to be made since the observations are not departure times.

There is, therefore, a research gap as user classes based on repeated trip behaviour have not, to the authors' knowledge, been considered for *road* users. Addressing this lack of empirical evidence is not trivial since the methods used to measure intrapersonal variability on other modes are not directly transferable.

There is therefore a methodological gap in addition to the empirical one; in the present paper a methodology is proposed which takes into account the strengths and weaknesses of such point-to-point sensor data. The methodology could be applied to any type of point-to-point sensor data, but Bluetooth is probably the most relevant currently due to its increasing popularity for measuring travel times on the road network. It is a data driven approach which clusters together travellers who have similar trip regularity and variability between days without relying on a priori assumptions. The proposed methodology includes using Sequence Alignment to examine spatial variability and Model Based Clustering to measure time of day variability. Sequence Alignment has been used to explore the order in which pedestrians move between attractions (Delafontaine et al., 2012, Shoval and Isaacson, 2007) and on one occasion to classify vehicle trajectories using GPS data (Kim and Mahmassani, 2015). It has not, to the authors' knowledge been used in relation to intrapersonal variability.

The rest of the paper is structured as follows. Section 4.2 describes methods to calculate measures of trip frequency, spatial variability and time of day variability using Bluetooth data. A method for obtaining user classifications based on the measures is also described. Section 4.3 includes an application to one year of data from 23 Bluetooth sensors in and around Wigan in northwest England. Descriptive statistics are presented to demonstrate the distribution of travellers and trips between clusters. Section 0 discusses the limitations of the methodology and the sensitivity of the findings in the case study to choices of parameters and the clustering algorithm. Section 4.5 concludes the paper by describing possible policy implications of the characteristics identified in the empirical study.

4.2 Methodology

The methods presented in this paper utilise point-to-point sensor data, for example Bluetooth, Wi-Fi or ANPR data, where unique identifiers are recorded that can be matched over space (i.e. from point-to-point) and time. Such data can be collected passively, over long periods of time and increasingly cheaply due to technological advances. By definition, the data is available for fixed locations or ‘points’ and therefore observations are directly comparable geographically. This differs from GPS trace data, for example, where observations are not made at consistent locations. The locations are, however, limited by the coverage of the sensors and therefore do not provide origin-destination (OD) information about trips. Also, depending on the type of detector, there is a likelihood of a vehicle/individual not being recorded as it passes a detector. For example an experiment by Araghi et al. (2014) found that discoverable Bluetooth devices passing a sensor were detected 80% of the time. Missing data creates ambiguity as to whether the traveller drove along a link not monitored by sensors or whether they passed a sensor but were not recorded.

Typically point-to-point sensor data only contains unique device identifiers, for example number plates or Bluetooth device identifiers, and the corresponding date-time stamps for each detector. The aspects of repeated trip making which can be measured, therefore, are trip frequency and spatial and temporal patterns of trips.

Point-to-point sensor data requires a significant amount of processing before being used to identify road user classes, as shown in Figure 4-1. First, the data from all sensors needs to be collated by the unique traveller identifiers. The observations for each traveller then need to be ordered according to the date-time stamps, retaining the sensor number (generically referred to by the variable s) and a date-time stamp (t). The time lag between consecutive observations is then considered.

If the distance between the sensors, speed limits and data from surrounding vehicles suggest that the device has travelled directly between the two locations, then the observations should be chained together as part of the same trip. This process should filter out observations relating to travellers of other modes, including pedestrians and cyclists, as these observations will not be chained together at all and therefore they will be dropped from the analysis. The trips for each traveller, i, are then analysed to obtain a traveller specific frequency measure ($freq_i$), spatial measure ($spat_i$) and time of day measure (tod_i). A segmentation of the travellers is then obtained using cluster analysis.

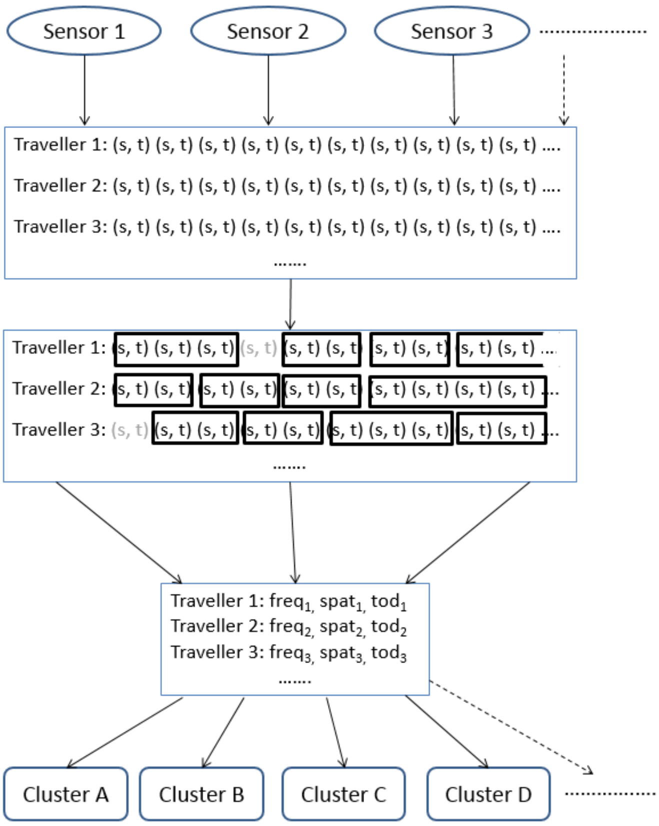

Figure 4-1: Overview of the process to identify road user classes using point-to-point sensor data

The proposed techniques for calculating the repeated trip behaviour measures will now be discussed.

4.2.1 Trip frequency

All types of point-to-point data collection will result in missing observations. A bias in the *travellers* who can be detected could potentially mean that resulting analyses cannot be considered representative of the population of travellers using the roads of interest. This consideration is outside the scope of the current paper. However, missing data may also refer to individual trips which are not detected at all, or an individual sensor not recording all possible data. In the current research, the assumption is made that individual trips are missing at random. This assumption means that for each traveller we have an unbiased, random sample of their trips and so the measure of trip frequency is comparable between travellers, although it will be an underestimate of the actual trip frequency. Alternative assumptions could be made if estimates are available of the bias in the trips recorded, for example by type of Bluetooth device or traffic density. The total number of trips observed per traveller is used as the measure of frequency in this paper.

4.2.2 Spatial variability

For spatial variability it is particularly important to focus on the nature of the data. For example, Järv et al. (2014) used mobile phone data and therefore they focused on individuals' activity spaces over time, as opposed to trip data. Bayarma et al. (2007) used data from a six week travel diary and for spatial variability focused on trip duration and the distance of trip destinations from the individuals' homes. Point-to-point sensor data differs from trip data from other sources as it only contains information about the part of the trip within the monitored part of the network, but it can contain many observations, depending on the sensor locations. Therefore, although OD information is not available, entry and exit points to the part of the network which is monitored can be captured. Route choices, in terms of the ordered sequence of sensors passed between the entry and exit points, can also be captured. To fully utilise the depth of this spatial information, the methodology in the current paper builds on the work of Delafontaine et al. (2012), who examined visitor movements through a large trade fair using Bluetooth data. Pairwise Sequence Alignment is used to calculate similarity measures between trips which can then be used to cluster similar trips. The distribution of trips between these spatial clusters for each traveller is then used to assess the degree of spatial variability.

Sequence Alignment was originally developed to compare protein sequences, but has also been used more recently by social scientists and geographers (Abbott, 1995, Shoval and Isaacson, 2007). It is suitable for point-to-point sensor data as it uses all of the available spatial data for a trip and does not just focus on start and end points. It also provides a systematic way of analysing the data while accounting for missing observations within sequences. Sequence Alignment techniques can be separated into global techniques, which try to match entire sequences, and local techniques which seek to find parts of the sequences which match. Kim and Mahmassani (2015) used one of the latter techniques to identify the Longest Common Subsequences in trace trip data, for example from taxis. Whilst this was a suitable technique in their research as they were aiming to identify ‘representative’ subsequences for clustering travel patterns, it is not suitable for the current research as it can completely ignore data from the start and ends of sequences. Kim and Mahmassani (2015) also had trajectory data which does not have the same problems with missing observations within sequences as point-to-point data. As in Delafontaine et al. (2012), who also considered point-to-point (Bluetooth) data, global sequence alignment will be applied as it considers the similarities and differences across entire sequences.

To calculate a pairwise measure of similarity, the optimal global alignment between two sequences is identified by adding in gaps, known as indels, to both sequences and assessing the similarity of aligned terms. The optimal alignment minimises the pairwise cost which is calculated by aligning sequence x and sequence y and then comparing each pair of aligned letters:

$$Pairwise\ cost = \sum_i dist\_xy_i \qquad (10)$$

where $dist\_xy_i$ is some sort of distance between the letters in position i in sequence x and y (x_i and y_i) with a constant distance used between an indel and any letter.

For example, consider Sequence 1 (ABCEGHIK) and Sequence 2 (BDEFGJK). There are thousands of possible alignments, although many can be instantly dismissed as suboptimal. Three possible alignments are shown in Figure 4-2.

A	B	C	E	-	G	H	I	K
-	B	D	E	F	G	J	-	K

A	B	C	E	G	H	I	K
-	B	D	E	F	G	J	K

-	A	B	C	E	-	G	H	I	K
B	D	-	-	E	F	G	-	J	K

Figure 4-2: Three possible alignments between Sequence 1 and Sequence 2

Rather than calculating all possible alignments, dynamic programming and more specifically the Needleman-Wunsch algorithm (Isaev, 2006, p9) can be used to find an optimal alignment more efficiently.

The specific methodology proposed in this paper for clustering trip sequences is shown in Figure 4-3. When considering point-to-point sensor data, each letter within a trip sequence corresponds to an observation at the sensor assigned with that letter. Each sequence corresponds to a trip, i.e. observations which have been matched between sensor locations using a unique identifier, for example a number plate, and then processed to ensure the traveller drove *directly* between the sensor locations. Substitution costs are denoted by $dist\_xy_i$ and in this paper are calculated as the distance by road between each pair of sensors. Alternatively the geodesic distance between sensors could be used, although this would be less useful where there are parallel routes with little opportunity for switching.

An indel in a trip sequences could represent a missing observation, either due to a slight difference in route or a genuine missing observation. They are also required when comparing sequences of different lengths. The cost associated with indels should be relatively low so as not to excessively punish missing data, which is common in some types of point-to-point sensor data such as Bluetooth. The indel cost should not, however, be less than half of the distance between the two sensors which are furthest apart, otherwise the optimal alignment process would never align observations from those two sites but would align each observation with an indel instead. As the substitution costs are calculated using an evidence-based metric, it is preferable to use this over an indel cost which is chosen more subjectively. Indels could represent devices passing a sensor but not being recorded and in this research it is assumed that the probability of this occurring does not depend on whether the device was recorded at the previous sensor. The same indel cost is, therefore, applied to gaps irrespective of whether they are preceded by gaps or letters (denoting observations).

Each pairwise cost represents the spatial similarity between two trips, where partially overlapping routes and geographical closeness are rewarded. To account for sequences of different lengths, Abbot's normalisation is applied by dividing each pairwise cost by the length of the longer sequence of the pair (Abbott and Tsay, 2000, p13).

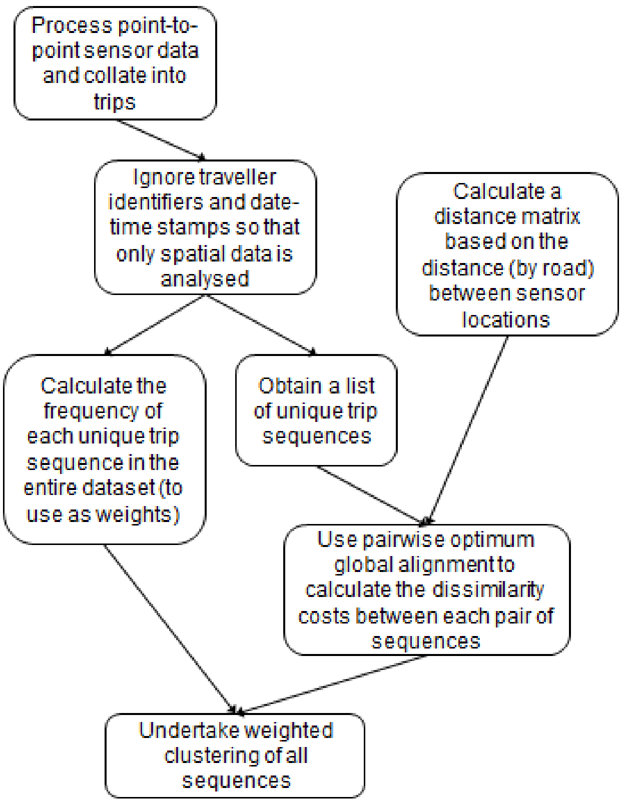

Figure 4-3: Spatial clustering process

The TraMineR package in R (Gabadinho et al., 2011) is used in this research to identify an optimal pairwise alignment for each pair of sequences in the data. Due to the large amount of data involved, the optimum alignments should be computed between all *unique* sequences to prevent duplicating effort. The pairwise costs are then used as the distance metric for clustering the sequences, with weights used based on their frequency in the data as described in Studer (2013). As the number of clusters to use is quite subjective, using hierarchical clustering provides a suitable format of data to identify the most appropriate cut-off to use. After identifying the spatial clusters of trips, each traveller is assessed to see how many of the clusters their trips fall into, and what proportion of their trips fall into their most common spatial cluster. Despite performing Sequence Alignment and clustering on all trip sequences together, therefore, this process will identify spatial variability measures for each traveller.

4.2.3 Temporal variability

Measures relating to intrapersonal variability in the time of day trips are made need to be comparable across all individuals and should give a meaningful insight into the underlying behaviours. Ideally, temporal variability should be measured based on comparable trips, but with point-to-point data this is somewhat ambiguous due to the limited coverage of detectors, missing observations and a lack of trip purpose information. For each traveller, trips which are first detected at matching sensor location and are also last detected at matching sensor locations could be compared but these are not guaranteed to relate to the same trip. This is because of the limited spatial coverage of detectors, i.e. these are not the OD pair of the actual trip, and there may be missing data. Also, for each traveller there may be many start and end detector pairs so there would be a confusing array of measures for each traveller, most of which would have very small sample sizes.

The approach proposed in this research is to consider the time of day that an individual passes a particular detector. For each traveller, the detector they pass most often will be examined. This is somewhat similar to the approach taken by Muthyalagari et al. (2001) on GPS travel diary data. They compare departure times based on location, but also based on trip purpose and so obtain measures of variability for the first departure from home, final departure from work and final arrival at home each day. Using the most common detector location only may be more closely linked to spatial variability as the time they pass a particular point may vary depending on their ultimate destination. It does, however, allow travellers who always travel at the same time of day but go to different locations to be identified in the data.

Having decided upon the observations to compare, a suitable measure of temporal intrapersonal variability needs to be selected. Comparison of observations within time bands can be useful, for example 10 minute intervals were used by Minnen et al. (2015). The results are usually dependent on the (usually arbitrary) choice of time band widths, however, and the relative precision of the time stamps from point-to-point sensors would be wasted. In the current research, therefore, the time of day is treated as a continuous variable and clustering is undertaken for each traveller separately, as was done for public transport users in Kieu et al. (2015b). Kieu et al. (2015b) used a density based clustering algorithm as their aim was to identify the percentage of each traveller's trips which fall within a habitual time cluster, as opposed to other trips which were classified as noise by the algorithm. This algorithm was not suitable for the current paper for two reasons. Firstly, the distribution of the trips classified as noise by a density based algorithm is of interest in this research and therefore a more holistic approach was preferred. Secondly, although the density based clustering algorithm used does not require the specification of the number of clusters to use, it does require a minimum points and density reach parameters. These parameters tell the algorithm how close together points should be if they are to be in the same cluster. For example, Kieu et al. (2015b) used 5 minutes as their density reach parameter for the time of day of trips analysis. An alternative approach which uses all available data and uses a data driven approach to decide on the spread of clusters is Model Based Clustering (see Fraley and Raftery (2002)). This can be used to identify whether there are multiple times of day at which an individual passes a given location.

An outline of how this approach can be used in the current application is shown in Figure 4-4. For regular travellers, the times at which s/he passes their most common location are assumed to have distinct clusters which combine to form a Gaussian mixture distribution. For each traveller with a large enough sample size, Maximum Likelihood Estimation is undertaken on the times at which s/he passes a particular point to identify the parameters of the Gaussian mixture distribution which best fits the times. The measures of each traveller's temporal variability are the number of clusters relating to their most common detector and the variance of the clusters.

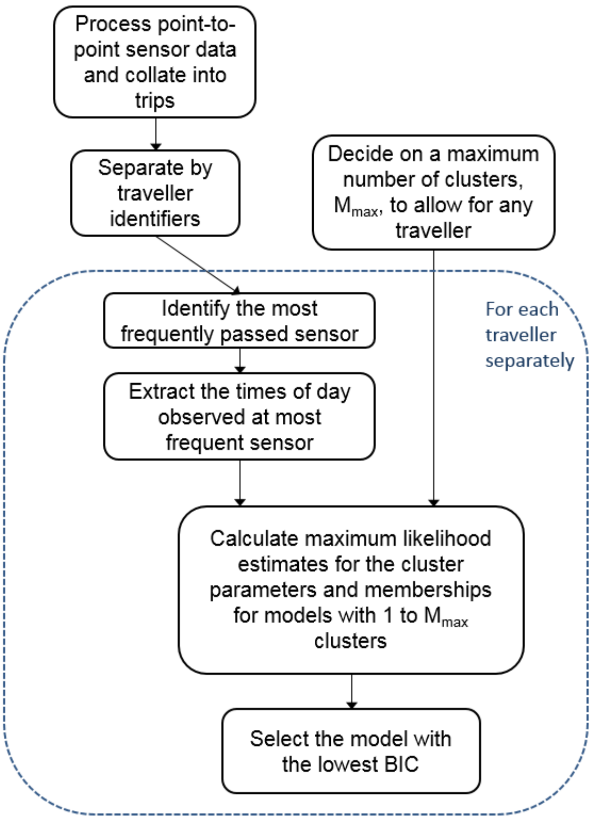

Figure 4-4: Time of day clustering process

4.2.4 Clustering together

Once the measures for the three aspects of intrapersonal variability have been calculated, namely trip frequency, spatial and temporal variability, all values are standardised prior to the clustering process. A number of different clustering methods may be suitable, but in this paper k-means clustering is used as it is relatively computationally fast and k-means was also used by Bayarma et al. (2007) to identify subgroups of travellers using travel diary data.

4.3 Application: a case study in northwest England

In this section the methodology described in Section 4.2 is applied to one year of real-life Bluetooth data from the road network in Wigan, a town in Greater Manchester in northwest England. Transport for Greater Manchester has installed around 770 fixed Bluetooth detectors next to roads in Greater Manchester. Such detectors are an increasingly popular way to measure travel times on the road network (Aliari and Haghani, 2012, Bhaskar and Chung, 2013, Araghi et al., 2014) and they have also been used to estimate OD matrices (Barceló et al., 2010, Chitturi et al., 2014) and measure pedestrian movements (Bullock et al., 2010, Versichele et al., 2012). Types of Bluetooth-enabled devices include smartphones, laptops, hands-free kits and in-car audio systems. By matching unique identifiers for devices between locations, trip data can be generated and filtered to remove travel times not associated with motor vehicles, as described in Section 4.2. As only discoverable Bluetooth devices can be detected, the trip data will only be a sample of trips undertaken in the area. The Bluetooth penetration rate has been measured by comparing ANPR and Bluetooth data for one link within Greater Manchester over a twelve hour period and the hourly penetration rates were calculated to be between 16% and 34%.

Data from 23 fixed Bluetooth detectors in and around Wigan (Figure 4-5) was analysed for all of 2015. The data includes 7,480,204 trips and these trips were associated with 327,264 unique MAC addresses, which for the purposes of this research are assumed to approximately correspond to individual travellers. Almost 28% of these MAC addresses only recorded one trip in this area in the year. Just

2% of the MAC addresses recorded greater than or equal to 260 trips in the year which is equivalent to at least one trip per day, on average, for someone working five out of seven days per week.

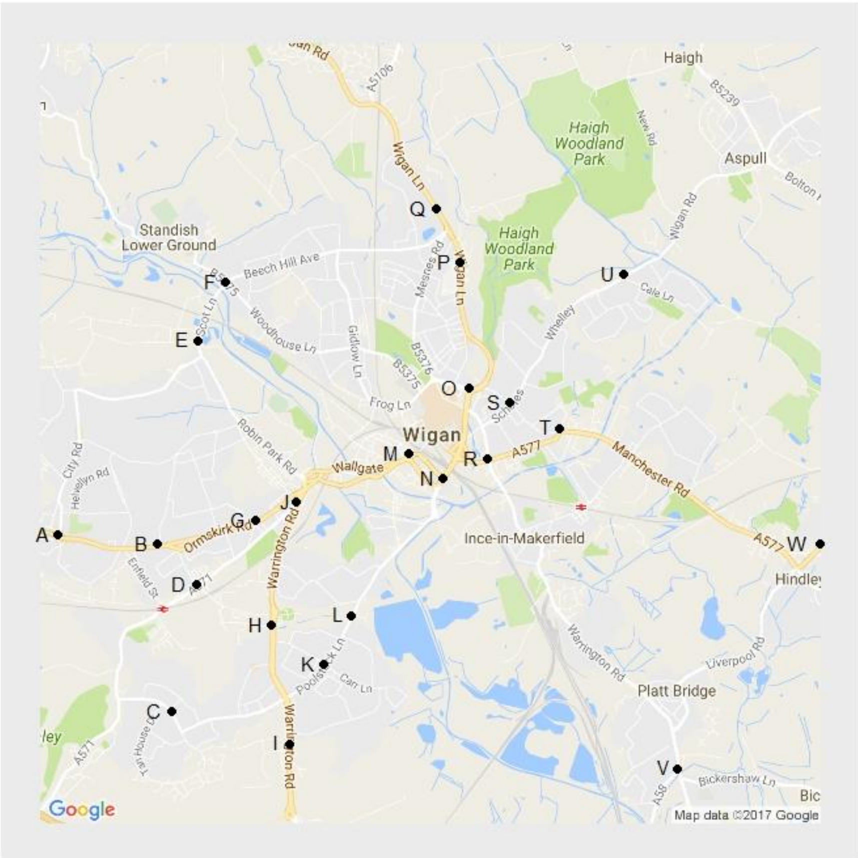

Figure 4-5: Bluetooth detector locations in and around Wigan

Computational limitations make sequence alignment on all unique sequences observed in the year infeasible, but it is possible for a month of data. To select the most appropriate month to use, the unique trip sequences, seq_k, (where $k = 1, \dots n$, the total number of unique sequences) observed in the year of data need to be extracted. For each of these sequences, the months in which it was observed and the total number of occurrences in the year (w_k) should be recorded. Equation (11) can then be used to calculate the coverage of each month. The coverage corresponds to the proportion of trips in the year which have a trip sequence which matches a trip observed in month *l*. The month of data with the highest coverage should be selected. Only sequences observed in the chosen month will undergo the pairwise alignment process and therefore choosing the month with the highest coverage will maximise the number of trips observed in the year which can be assigned to a cluster as others will not be represented in the distance matrix.

$$Coverage\ month\ l = \frac{\sum_k w_k \lambda_{kl}}{\sum_k w_k}$$

$$where\ \lambda_{kl} = \begin{cases} 1 & if\ seq_k\ occurs\ in\ month\ l \\ 0 & otherwise \end{cases} \tag{11}$$

In the case study area, March has the highest coverage (0.98). Pairwise Sequence Alignment was undertaken on March data only. Spatial clusters were then estimated using the sequence alignment scores and finally trip sequences in the full year of data were assigned to spatial clusters.

After calculating the pairwise costs relating to the sequences, standard hierarchical clustering using Ward's method was undertaken. In order to select the most appropriate number of clusters to use, partition quality measures were calculated, following Studer (2013), including Average Silhouette Width and Calinski-Harabasz index. This provided a starting point for testing, but ultimately 150 spatial clusters were used as this provided a useful level of aggregation for the overall intrapersonal variability clustering. The choice of 150 spatial clusters will be discussed further in Section 4.4.

Model Based Clustering was undertaken separately on the times of day at which each traveller passed their most common sensor location, provided there were at least 20 such observations. Temporal clustering could be performed for 33,375 travellers, which is approximately 10% of the travellers observed in the data. As the final intrapersonal variability clustering cannot deal with missing values, the remaining travellers were assumed to have one temporal cluster and the average cluster variance was calculated as the variance of any observations, or assigned a zero if there was only one observation.

Five variables were retained for each traveller in the final cluster analysis:

i. the number of trips observed in the year,
ii. the number of temporal clusters estimated based on the time of day the traveller passed their most common location,
iii. the average variance of their temporal clusters,
iv. the number of spatial clusters their trips were observed in and
v. the percentage of their trips which fell into their most common spatial cluster.

Using the Elbow Method, the number of clusters of travellers to use was set at 12.

4.3.1 Descriptive statistics of the user classes

Overall, the clusters can be separated into the six ‘infrequent traveller’ clusters (A to F), four ‘frequent traveller’ clusters (G to J) and two ‘very frequent travellers’ (K and L).

Table 4-1 shows that the vast majority of travellers assigned to clusters A to F recorded very few trips during the year and the average across these six clusters is just 5 trips in the year. Although the term ‘infrequent travellers’ has been used, these could be people who were visiting Wigan or local people who do not usually travel by road. The low frequency of observations could also be due to the type of data collection, for example a frequent traveller may only occasionally use their Bluetooth enabled hands-free device and thus appear very infrequently in the data. As very little data is available for these travellers, it is not reasonable to try to make distinctions based on the spatial and time of day variability in trips. In the remainder of this section, therefore, clusters A to F will be combined into one class of road users. Figure 4-6 demonstrates the uneven distribution of travellers allocated to the clusters, as clusters A to F contain 89% of travellers but they only account for 19% of the trips observed.

Table 4-1: Cluster membership and trip characteristics

User class	Cluster	Average trips per traveller in 2015	Travellers per cluster	Total trips
Infrequent travellers	A	1	103,340	153,223
	B	3	991	2,767
	C	4	86,473	344,128
	D	6	4,640	28,127
	E	9	23,987	209,480
	F	10	72,042	720,326
Frequent travellers	G	69	16,634	1,144,115
	H	100	8,163	820,221
	I	264	3,089	815,504
	J	274	5,809	1,590,437
Very frequent travellers	K	685	1,901	1,302,874
	L	1,790	195	349,002

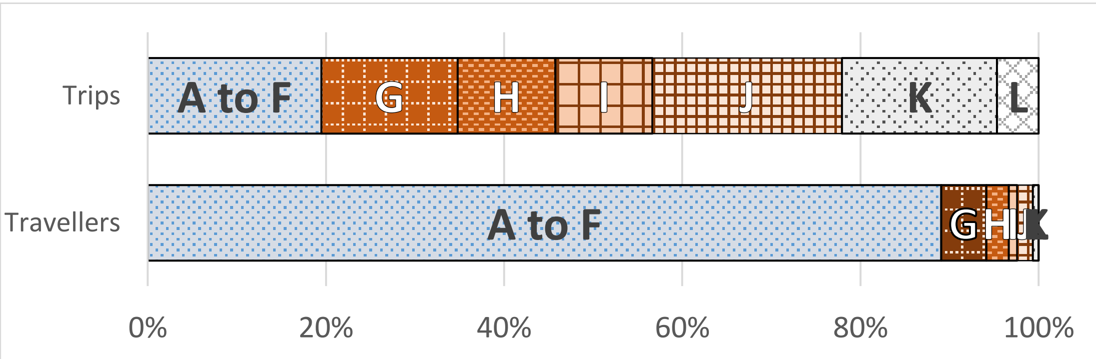

Figure 4-6: Segmentation of trips and travellers into clusters

Within the frequent traveller user class, the four clusters, G to J, have quite different characteristics as shown in Table 4-2. They can be separated into two groups based on their trip frequency. Travellers in clusters I and J are observed almost three times as often as travellers in clusters G and H. Within each pair, one cluster represents more regular trip makers (H and I) and the other represents less regular travellers (G and J). The more regular travellers (H and I) make fewer different kinds of trips (represented by different spatial clusters) and make their most common trip (spatially) a higher percentage of the time, when compared with their pairwise equivalents (G and J respectively). Despite the difference in trip frequency, travellers in both G and J make their most common trip approximately 24% of the time, based on spatial clustering. For travellers in H and I this is around 35%, thus they are classed as more regular travellers. Interestingly, the clusters with higher *spatial* regularity also have higher *time of day* regularity. Table 4-2 shows that the clusters described as being less regular (G and J) have fewer distinct time of day clusters, with greater variances on average. This suggests higher levels of flexibility and lower levels of predictability. Figure 4-7 highlights the connection between the spatial and time of day variability as travellers in clusters G and J have a lower percentage of trips in their most common spatial cluster and also have fewer time of day clusters when compared with their pairwise equivalents (H and I respectively).

Table 4-2: Characteristics of frequent traveller clusters

		Spatial variability			Time of day variability		
Clust	Average trip freq.	Overall	Spatial clusters used	% of trips in most common spatial cluster	Overall	Time of day clusters	Average variance
G	1-2 per week	Less regular	21	24%	Less regular	1.4	0.015
H	1-2 per week	More regular	17	36%	More regular	3.5	0.004
I	5 per week	More regular	26	34%	More regular	6.4	0.003
J	5 per week	Less regular	41	24%	Less regular	2.6	0.009

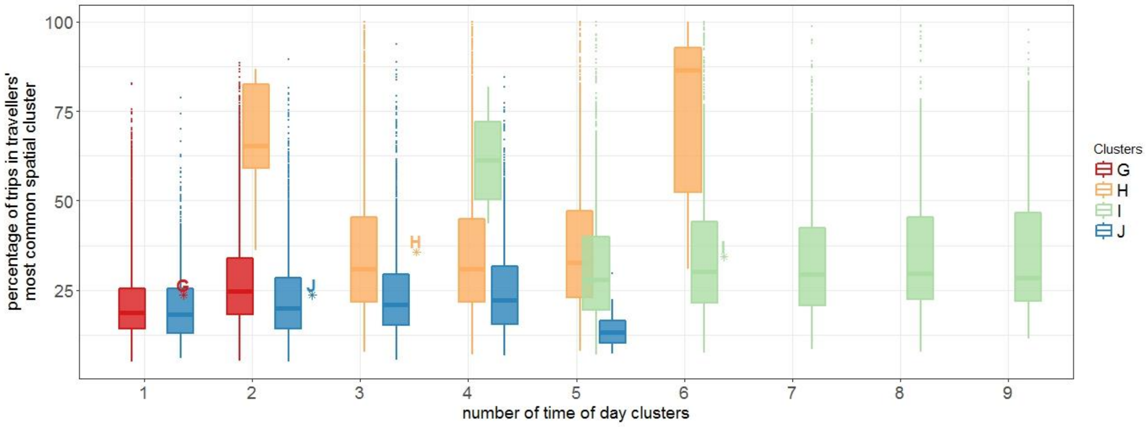

Figure 4-7: Boxplots for each frequent traveller cluster by the number of time of day clusters[4]

The final two clusters (K and L) relate to very frequent travellers. Only 0.64% of travellers, a total of 2,096, are allocated to these clusters, but they make 22% of the trips observed. Clusters K and L contain travellers with 2 and 5 trips per day on average respectively and cluster L has far fewer travellers allocated to it than any other cluster. The very frequent travellers have *time of day variability* characteristics which are similar to a combination of clusters I and J, i.e. the clusters with the next highest trip frequency, separated into more and less regular components. The very frequent travellers are observed in more *spatial* clusters than travellers in I and J, but the amount of spatial variability does not increase at the same rate as the trip frequency. Clusters I and J have 7.6 trips per spatial cluster on average, but clusters K and L have 11.7 and 18.7 respectively. This suggests that as well as making more trips than travellers in the frequent user class, the very frequent user class also have different characteristics in terms of the use they make of different parts of the network.

[4] The stars and accompanying letters represent the average values for both variables for each of the four frequent traveller clusters.

4.3.2 Predictable differences in user class proportions

Identifying predictable differences in the proportion of trips in each road user class by the day of the week or season could help to identify systematic differences in travel behaviour (Crawford et al., 2017). The proportions are relatively stable across days of the week and seasons in this case study area, although some patterns are evident. For example, infrequent travellers are slightly more common on weekend days (see Figure 4-8) which is consistent with Wigan being a trip attractor for weekend activities such as visiting a park, museum or theatre. The proportion of trips made by very frequent travellers is also higher on weekend days than on weekdays. This is particularly surprising on Sundays when there are likely to be fewer buses operating and fewer deliveries being made. For all three user classes, the difference in proportion between weekdays and weekend days is statistically significant ($p<0.00001$).

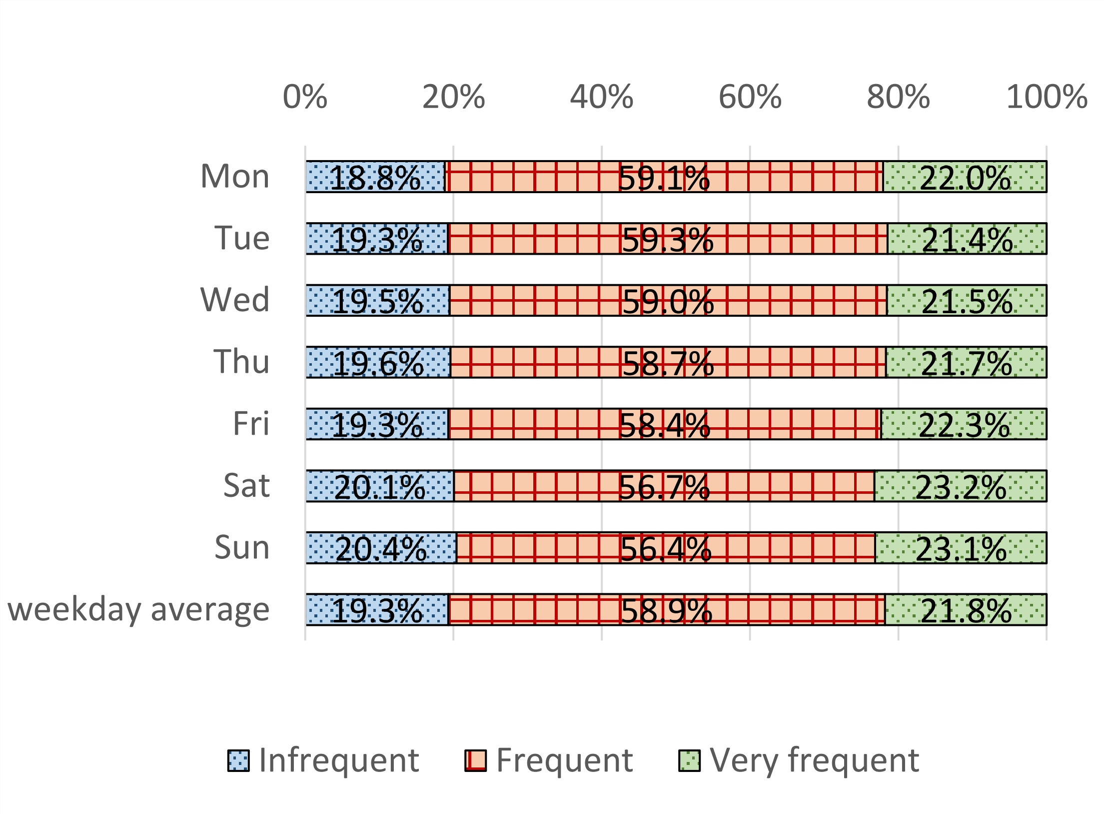

Figure 4-8: Proportion of trips by each road user class in Wigan by the day of the week

Frequent travellers contribute a consistent proportion of trips throughout the year. The proportion is highest at 59.0% in March but remains relatively stable until there is a decrease to 57.1% in December, probably due to the holiday period. As shown in Figure 4-9, the proportions of trips made by infrequent and very frequent travellers vary more throughout the year. For both infrequent and very frequent travellers, the difference in the proportions of trips made by the user class in January and June are statistically significant (p<0.00001). For frequent travellers the same comparison is also statistically significant at the 0.05 significance level (p=0.048), although less convincingly so. A higher proportion of trips are made by infrequent travellers during the winter months. This could be because the leisure trip attractors in the town are more likely to relate to indoor activities, whereas during the summer there may be more competition from outdoor activities which are not within the monitored area of Wigan. It could also be due to an increased reliance on cars during months where the weather is colder and wetter. The proportion of trips made by very frequent travellers peaks during the summer months. This, together with the day of the week analysis, suggests that the very frequent traveller user class is more closely linked to leisure trips than business trips. The current case study is limited as only one year of data was analysed and therefore it is not clear whether the increase in the proportion of trips made by very frequent travellers between January and June is part of a longer term trend, for example an increasing number of taxis operating in the area, or whether it is truly a seasonal pattern.

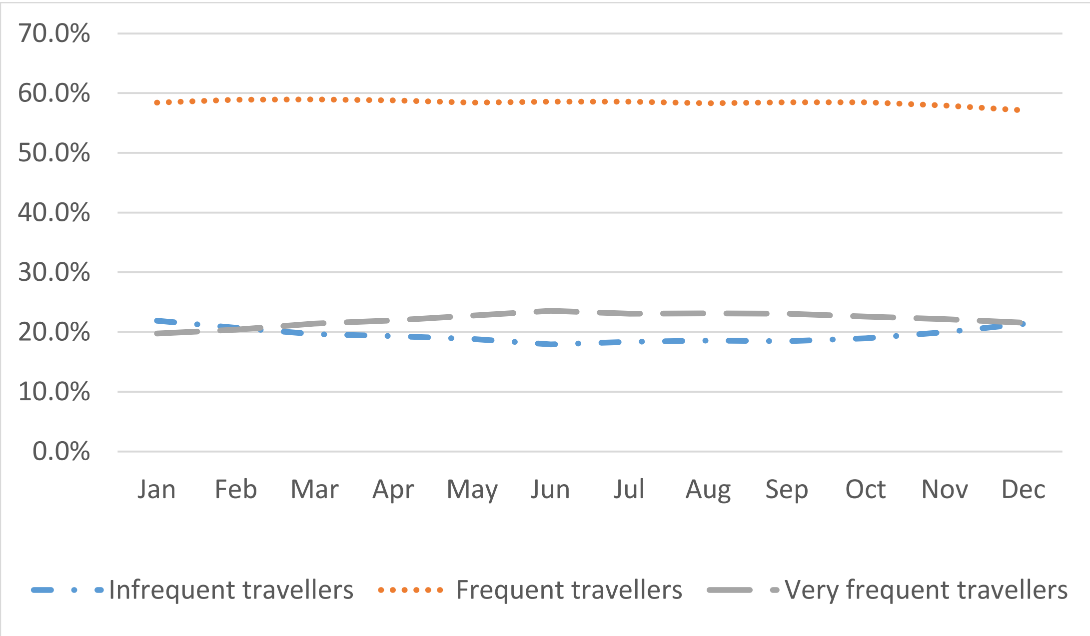

Figure 4-9: Proportion of trips made by each road user class throughout the year in Wigan

4.4 Discussion

Although the trip frequency measure was designed to be a comparative value, it is inevitable that attempts will be made to interpret the user class characteristics using this measure. As fixed Bluetooth detectors do not have a 100% detection rate for discoverable Bluetooth devices and trips could be missed due to not carrying the device or disabling the Bluetooth functionality, the trip frequency data should be considered to be a lower bound for the number of trips actually made. There may also be a bias due to new technologies, for example iPhones from iOS 8 have features which can randomise MAC addresses (the “unique” identifiers) to prevent the tracking of devices. As the data analysed in Section 4.3 is from 2015, the penetration rate of such devices is assumed to be small. In the future, analyses relating to repeated trip making using Bluetooth data may require additional data collection to understand the types of devices being detected and the possible implication for trip frequency measures.

The trip frequency values will be sensitive to the parameters used in cleaning the Bluetooth data, particularly those used as part of the process to connect observations into trips. In the application to one year of data in Section 4.3, for example, 0.4% of the trips observed were circular routes which included at least three observations. Whilst it is virtually impossible to distinguish a very brief drop-off on a route from a stop at a pedestrian crossing, for example, further work should focus on identifying the optimal parameter values for connecting or splitting trip data.

The Sequence Alignment based method was used to identify 150 spatial clusters from the trips observed. Each spatial cluster contains 173 different sequences on average. The heterogeneity of spatial cluster membership is highlighted by the variety in the number of unique sequences assigned to each spatial cluster and the variability in sequence lengths and starting sensor location within clusters.
Figure 4-10 shows the 15 most common sequences, out of 197, in one spatial cluster. These sequences are each observed between 127 and 893 times in the year of data. The sequences go from the west of Wigan to the east via the town centre. The sequences relate to one spatial cluster only and demonstrate the

effectiveness of the method in combining trip sequences where intermediate sites are not observed and those representing slightly different routes, for example trips going past site S or site R.

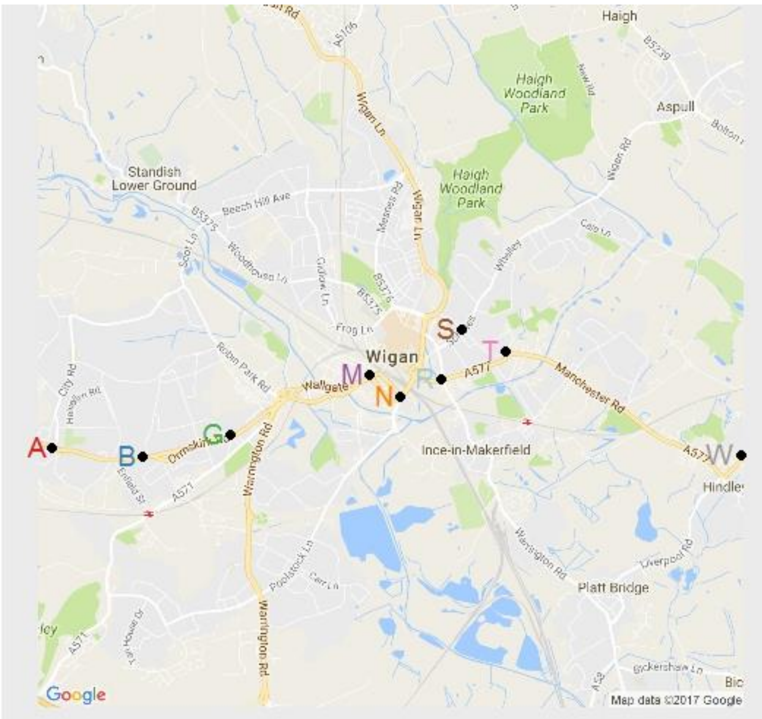

A-B-M-N-R-T-W
A-B-G-M-N-R-T-W
A-B-N-R-T-W
B-G-M-N-R-T-W
A-B-M-N-R-W
A-B-G-N-R-T-W
A-B-R-T-W
A-B-M-R-T-W
A-B-R-W
A-B-N-R-W
A-B-M-R-W
A-B-M-N-S-W
A-B-M-N-S-T-W
A-B-G-M-R-T-W
A-B-G-M-N-R-W

Figure 4-10: A map and table showing the 15 most common sequences assigned to one of the 150 spatial clusters

The choice of 150 spatial clusters was made by considering the partition quality measures implemented in Studer (2013), including the Average Silhouette Width and Calinski-Harabasz index, but also using a more qualitative examination of the trip sequences clustered together. A plot of a range of partition quality measures showed apparent step changes at around 150, 800 and 2,000 clusters. An examination was undertaken of 3 clusters randomly selected at the 150 cluster level to explore whether the clusters at the 800 cluster level were more intuitive. In one case, 95% of the sequences were assigned to a single cluster at the 800 cluster level. In another case, 94% of the trip sequences were split between two clusters at the 800 level, but the separation was not particularly meaningful. The most common sequences in the two clusters were from site W to site R and from site W to site N, which are similar trips but with different sequence lengths. In the third case, the three largest clusters at the 800 cluster level did have meaningful differences: one included trips travelling north to P, one included the reverse trips (travelling south from P) and the other included trips with at least two sites in common with the north-south trips, but which ultimately travelled east-west or further north. This examination was very subjective but it highlights the difficulty in selecting the 'right' level of aggregation overall. In practice, however, a choice has to be made which gives the most meaningful results overall. The final user class clustering was repeated using the spatial variability measures calculated using 150, 800 and 2,000 spatial clusters and the characteristics of the twelve clusters were relatively similar. For example, if we compare using 150 and 2,000 spatial clusters: the percentage of trips by frequent travellers remained fairly constant, the percentage of trips by *in*frequent travellers increased from 19.5% to 20.6% and the percentage of trips by very frequent travellers decreased from 22.1% to 21.0%. This suggests that the overall methodology is fairly robust with respect to the number of spatial clusters selected.

Approximately 10% of the travellers observed in the data had sufficient data to be able to produce a measure of intrapersonal variability for the time of day they pass their most common location. This percentage is determined by the minimum sample size specified for the Model Based Clustering. This parameter has already been set quite low in this case study, at 20, and therefore reducing it further was not feasible. For the travellers with sufficient trips passing their most common sensor location, the proportion of their total trips that this measure represents varies; 29% of these travellers passed their most common location on less than

50% of their trips. Although visual inspections were undertaken for a sample of travellers, further work is required to determine whether the Gaussian assumption is reasonable, or whether another distribution, perhaps a skewed distribution such as the lognormal distribution, would be more appropriate.

The choice of k-means as the clustering algorithm may have an impact on the final clusters identified. Due to the large number of travellers in Section 4.3, standard hierarchical clustering is not possible due to computational limitations. An alternative algorithm which could have been applied is the density-based algorithm DBSCAN (Ester et al., 1996). DBSCAN can identify clusters of arbitrary shape using very few initial parameters and was used by Kieu et al. (2015b) to identify regular ODs and habitual trip timings for travellers using smart card data. DBSCAN was applied to the year of data analysed in Section 4.3, but the results were less satisfactory than those obtained using k-means. DBSCAN identified one very large cluster, which approximately equated to the travellers in the infrequent traveller k-means clusters combined. Irrespective of the parameters applied, this technique resulted in many very small clusters which would not be useful when defining user classes. Also, although it is considered an advantage that DBSCAN can identify noise in the data, it is somewhat problematic in the current application as 4% of travellers have not been assigned to a cluster. DBSCAN does not perform very well when clusters have different densities as a Minimum Points parameter which is suitable for all clusters cannot be specified. It was therefore preferable to use k-means clustering for the case study presented, but alternative algorithms should still be explored in future applications.

The overlap between some of the k-means clusters and the slightly different clusters identified by DBSCAN suggest that it may be more appropriate to use fuzzy, rather than hard, clustering for travellers. Non-fuzzy clustering has been used in the current paper as it provides results which are more intuitive for policy analyses, in relation to travellers belonging to a single user class, but fuzzy clusters could be used in future applications.

4.5 Conclusions

This study has demonstrated the extent to which road user classes based on repeated trip making behaviour can be identified using point-to-point sensor data. The purpose of identifying such classes would be to inform policy development and to use as inputs for economic or behavioural models. The methodology was designed for a specific type of data, namely point-to-point sensor data, and non-traditional techniques (from a transportation research perspective) have been used to extract as much relevant data as possible. The Bluetooth data analysed for the case study area was collected for the purpose of travel time estimation and therefore the marginal costs of using it for this research were minimal.

The results obtained from the proposed method could be used by policy makers and practitioners in several ways. For example, now that an infrequent traveller user class has been defined, road managers may wish to explore whether this user class makes up a greater proportion of travellers on specific days where there are seasonal sales, sporting events or major incidents on other roads in the region. Such insights could inform planning for future special events.

Although the majority of trips are made by frequent travellers (58%), the vast majority are not recorded making two trips per weekday, on average. For the frequent travellers with lower intrapersonal variability, their most common trip spatially makes up only 35% of their total trips. Therefore, while the lower spatial and temporal intrapersonal variability may suggest that ride sharing would be a suitable option to promote, the ability to make a significant proportion of trips to other locations should also be addressed. For the frequent travellers with greater intrapersonal variability alternative options with more flexibility might be more attractive, for example cycling or car clubs.

For the travellers in the very frequent user class, further research is required to examine what sort of trips are being recorded, as the higher proportion of trips at the weekend and during the summer suggest that it is not just related to taxis, buses and delivery drivers. If they are predominantly business trips, then policies promoting mode change for personal travel will not result in a decrease in the 22% of trips made by these very frequent travellers. To have any impact on the trips made by this user class, alternative policies would need to be considered, for example changes to bus routes or encouraging deliveries during off-peak periods.

4.6 References

Abbott, A. (1995) Sequence Analysis: New Methods for Old Ideas. *Annual Review of Sociology,* 21**,** 93-113.

Abbott, A. & Tsay, A. (2000) Sequence Analysis and Optimal Matching Methods in Sociology: Review and Prospect. *Sociological Methods & Research,* 29**,** 3-33.

Aliari, Y. & Haghani, A. (2012) Bluetooth Sensor Data and Ground Truth Testing of Reported Travel Times. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 167-172.

Antoniou, C., Balakrishna, R. & Koutsopoulos, H. N. (2011) A Synthesis of emerging data collection technologies and their impact on traffic management applications. *European Transport Research Review,* 3**,** 139-148.

Araghi, B. N., Olesen, J. H., Krishnan, R., Christensen, L. T. & Lahrmann, H. (2014) Reliability of Bluetooth Technology for Travel Time Estimation. *Journal of Intelligent Transportation Systems,* 19**,** 240-255.

Axhausen, K. W., Zimmermann, A., Schönfelder, S., Rindsfüser, G. & Haupt, T. (2002) Observing the rhythms of daily life: A six-week travel diary. *Transportation,* 29**,** 95-124.

Barceló, J., Montero, L., Marqués, L. & Carmona, C. (2010) Travel Time Forecasting and Dynamic Origin-Destination Estimation for Freeways Based on Bluetooth Traffic Monitoring. *Transportation Research Record: Journal of the Transportation Research Board,* 2175**,** 19-27.

Bayarma, A., Kitamura, R. & Susilo, Y. (2007) Recurrence of Daily Travel Patterns: Stochastic Process Approach to Multiday Travel Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2021**,** 55-63.

Bhaskar, A. & Chung, E. (2013) Fundamental understanding on the use of Bluetooth scanner as a complementary transport data. *Transportation Research Part C: Emerging Technologies,* 37**,** 42-72.

Buliung, R. N., Roorda, M. J. & Remmel, T. K. (2008) Exploring spatial variety in patterns of activity-travel behaviour: initial results from the Toronto Travel-Activity Panel Survey (TTAPS). *Transportation,* 35**,** 697-722.

Bullock, D., Haseman, R., Wasson, J. & Spitler, R. (2010) Automated Measurement of Wait Times at Airport Security. *Transportation Research Record: Journal of the Transportation Research Board,* 2177**,** 60-68.

Chen, R. & Mahmassani, H. S. (2004) Travel time perception and learning mechanisms in traffic networks. *Transportation Research Record: Journal of the Transportation Research Board,* 1894**,** 209-221.

Cherchi, E. & Cirillo, C. (2014) Understanding variability, habit and the effect of long period activity plan in modal choices: a day to day, week to week analysis on panel data. *Transportation,* 41**,** 1245-1262.

Cherchi, E., Cirillo, C. & Ortúzar, J. d. D. (2017) Modelling correlation patterns in mode choice models estimated on multiday travel data. *Transportation Research Part A: Policy and Practice,* 96**,** 146-153.

Chikaraishi, M., Fujiwara, A., Zhang, J. & Axhausen, K. (2009) Exploring Variation Properties of Departure Time Choice Behavior by Using Multilevel Analysis Approach. *Transportation Research Record: Journal of the Transportation Research Board,* 2134**,** 10-20.

Chitturi, M., Shaw, J., Campbell, J. & Noyce, D. (2014) Validation of Origin-Destination Data from Bluetooth Reidentification and Aerial Observation. *Transportation Research Record: Journal of the Transportation Research Board,* 2430**,** 116-123.

Chu, K. & Chapleau, R. (2010) Augmenting Transit Trip Characterization and Travel Behavior Comprehension. *Transportation Research Record: Journal of the Transportation Research Board,* 2183**,** 29-40.

Crawford, F., Watling, D. P. & Connors, R. D. (2017) A statistical method for estimating predictable differences between daily traffic flow profiles. *Transportation Research Part B: Methodological,* 95**,** 196-213.

Data.gov.uk. (2012) *Contracts Finder Archive: National Travel Survey 2013-2018* [Online]. Available: https://data.gov.uk/data/contracts-finder-archive/contract/565163/ [Accessed: 28 February 2017].

Delafontaine, M., Versichele, M., Neutens, T. & Van de Weghe, N. (2012) Analysing spatiotemporal sequences in Bluetooth tracking data. *Applied Geography,* 34**,** 659-668.

Department for Business, Innovation and Skills, (2014) The Fourth Work-Life Balance Employer Survey (2013). Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/398557/bis-14-1027-fourth-work-life-balance-employer-survey-2013.pdf [Accessed: 28th March 2017].

Elango, V., Guensler, R. & Ogle, J. (2007) Day-to-Day Travel Variability in the Commute Atlanta, Georgia, Study. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 39-49.

Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *In:* SIMOUDIS, E., HAN, J. & FAYYAD, U., eds., Second International Conference on Knowledge Discovery and Data Mining (KDD-96), 1996 Portland, Oregon. The AAAI Press, 226-231.

Fraley, C. & Raftery, A. E. (2002) Model-Based Clustering, Discriminant Analysis, and Density Estimation. *Journal of the American Statistical Association,* 97**,** 611-631.

Gabadinho, A., Ritschard, G., Müller, N. S. & Studer, M. (2011) Analyzing and Visualizing State Sequences in R with TraMineR. *Journal of Statistical Software,* 40**,** 1-37.

Goulet Langlois, G., Koutsopoulos, H. N. & Zhao, J. (2016) Inferring patterns in the multi-week activity sequences of public transport users. *Transportation Research Part C: Emerging Technologies,* 64**,** 1-16.

Heinen, E. & Chatterjee, K. (2015) The same mode again? An exploration of mode choice variability in Great Britain using the National Travel Survey. *Transportation Research Part A: Policy and Practice,* 78**,** 266-282.

Isaev, A. (2006) *Introduction to Mathematical Methods in Bioinformatics,* Berlin, Heidelberg, Springer.

Järv, O., Ahas, R. & Witlox, F. (2014) Understanding monthly variability in human activity spaces: A twelve-month study using mobile phone call detail records. *Transportation Research Part C: Emerging Technologies,* 38**,** 122-135.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015) Passenger Segmentation Using Smart Card Data. *IEEE Transactions on Intelligent Transportation Systems,* 16**,** 1537-1548.

Kim, J. & Mahmassani, H. S. (2015) Spatial and Temporal Characterization of Travel Patterns in a Traffic Network Using Vehicle Trajectories. *Transportation Research Procedia,* 9**,** 164-184.

Kitamura, R., Yamamoto, T., Susilo, Y. O. & Axhausen, K. W. (2006) How routine is a routine? An analysis of the day-to-day variability in prism vertex location. *Transportation Research Part A: Policy and Practice,* 40**,** 259-279.

Liu, R., Van Vliet, D. & Watling, D. (2006) Microsimulation models incorporating both demand and supply dynamics. *Transportation Research Part A: Policy and Practice,* 40**,** 125-150.

Minnen, J., Glorieux, I. & van Tienoven, T. P. (2015) Transportation habits: Evidence from time diary data. *Transportation Research Part A: Policy and Practice,* 76**,** 25-37.

Muthyalagari, G. R., Parashar, A. & Pendyala, R. M. Measuring Day-to-Day Variability in Travel Characteristics Using GPS Data. 80th Annual Meeting of the Transportation Research Board, 2001 Washington D.C.: TRB.

Schlich, R. & Axhausen, K. W. (2003) Habitual travel behaviour: Evidence from a six-week travel diary. *Transportation,* 30**,** 13-36.

Shoval, N. & Isaacson, M. (2007) Sequence Alignment as a Method for Human Activity Analysis in Space and Time. *Annals of the Association of American Geographers,* 97**,** 282-297.

Studer, M. (2013) WeightedCluster Library Manual: A practical guide to creating typologies of trajectories in the social sciences with R. *LIVES Working Papers* [Online], 24.

Tarigan, A. & Kitamura, R. (2009) Week-to-Week Leisure Trip Frequency and Its Variability. *Transportation Research Record: Journal of the Transportation Research Board,* 2135**,** 43-51.

Versichele, M., Neutens, T., Delafontaine, M. & Van de Weghe, N. (2012) The use of Bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent Festivities. *Applied Geography,* 32**,** 208-220.

5 Examining day-to-day variability by connecting network- and traveller-focused analyses of travel behaviour

Abstract

The effects of day-to-day variability in travel behaviour are visible on all transport networks, for example in the form of travel time unreliability on road networks. Understanding variability in travel behaviour matters when allocating resources on networks, developing suitable policies and including variability in models. Quantitative analyses relating to variability on transport networks usually adopt a single perspective, either focusing on individual traveller behaviour or focusing on data aggregated for a part of the network such as a road link.

Complex relationships exist between individual traveller behaviour and aggregate level observations on the network. Travellers are assumed to make travel choices based on network level attributes, which in turn result from the aggregated decisions of travellers. Day-to-day patterns on the network can have multiple possible behavioural causes, and variability in traveller behaviour may not be observed at the network level if masked by counteracting behaviour. Undertaking analysis at only one level, therefore, may not provide a complete picture of traveller behaviour. Emerging data sources are providing new opportunities to examine variability from multiple perspectives.

This paper proposes an inductive three stage approach *inspired by* mixed methods techniques to undertake analyses considering multiple perspectives. In Stage 1, analyses are performed independently on data from each perspective for the same geographic area. Stage 2 compares the data used and triangulates the findings from these separate analyses. Stage 3 involves the ‘following a thread’ technique

from mixed methods research, where questions or themes from the analysis of each data type are explored further using the other. Therefore, hypotheses to explain network-level variability are developed then tested using traveller-focused analyses, and vice versa. A real world application is presented consisting of a single road link, where loop detector data provides the network perspective and Bluetooth sensor data provides the traveller perspective.

5.1 Introduction

The effects of day-to-day variability in travel behaviour are visible on all transport networks, for example in the form of travel time unreliability on road networks. It is important to understand such variability in travel behaviour for three reasons. Firstly, it is important for resource allocation. This is particularly relevant when the variations are systematic, for example according to the day of the week or season. A common example could be having different bus timetables for different seasons of the year or days of the week. Secondly, understanding the variability in different aspects of travel behaviour can inform the development of policies, for example Travel Demand Management (TDM) strategies such as those described by Meyer (1999), which aim to influence behaviour in order to improve network performance. TDM strategies may seek to reduce, re-time, re-mode and/or re-route trips (Transport for London, 2013, p18). It is only by having a good understanding of individual's weekly or monthly travel patterns and the variability within those patterns that adaptations to their behaviour can be identified which could result in meaningful long term changes. One example could be to encourage employees to change to a compressed working schedule so that the same number of hours are worked over fewer days, thus resulting in fewer commuting trips. Thirdly, variability can be introduced into models of future scenarios, for example by including variable demand and/or route choice (Nakayama and Watling, 2014). Such models should be based on analyses of real world decisions. There are multiple benefits, therefore, in gaining a better understanding of day-to-day variability in travel behaviour.

One way to gain a greater insight into such variability is by analysing relevant quantitative data. Empirical transportation research can take place from many perspectives, for example the focus may be on individual travellers, parts of the transport network, such as a single road link or underground line, or geographic areas, such as a city or a whole country. The current research concentrates on traveller-focused analyses, to examine individuals' behaviour, and network-focused analyses, where the aggregated effects of the individuals' behaviour can be examined.

Day-to-day variability in behaviour at the traveller level requires data which traditionally would have been collected using a multi-day travel diary (Huff and Hanson, 1986, Jones and Clarke, 1988), although more recently data from emerging data sources has also been utilised, for example from mobile phones (Järv et al., 2014), GPS trackers (Elango et al., 2007, Spissu et al., 2011, Shen et al., 2013) and smart cards (Morency et al., 2007, Ma et al., 2013, Kieu et al., 2015b). Some of the research has focused on a single aspect of travel behaviour, for example spatial variability (Buliung et al., 2008, Spissu et al., 2011, Järv et al., 2014), trip frequency (Elango et al., 2007) or departure time variability (Chikaraishi et al., 2009). A separate strand of research has used empirical data to identify classes of traveller based on their repeated trip characteristics (Ma et al., 2013, Kieu et al., 2015b, Crawford et al., Under review-b).

In contrast to traveller-focused analyses, network-focused analyses often have findings which can inform policy making, for example Singhal et al. (2014) and Heinen and Chatterjee (2015). Network-focused analyses often involve some sort of comparison or the investigation of an ‘effect’, for example, researchers may focus on the day of the week (Zhang et al., 2007, Kaltenbrunner et al., 2010, Yazici et al., 2012) or weather conditions (Datla and Sharma, 2008, Arana et al., 2014, Singhal et al., 2014). They can relate to data from a range of levels, from a single loop detector on the road network (Weijermars and van Berkum, 2005, Crawford et al., 2017) to smart card data from a city-wide public transport system (Singhal et al., 2014, Tao et al., 2014).

Empirical analyses relating to variability on transport networks usually focus on a single type of data and a single level of aggregation, for example the analysis of individual traveller behaviour using travel diaries (Schlich and Axhausen, 2003, Bayarma et al., 2007, Buliung et al., 2008) or the analysis of link flows using loop detector data (Weijermars and van Berkum, 2005, Datla and Sharma, 2008, Crawford et al., 2017). Some research has looked at fusing data from different sources (Kusakabe and Asakura, 2014, Bhaskar et al., 2014), but the analysis uses just one perspective, either a traveller or a network focus.

Transport networks involve complex relationships between the behaviour of individual travellers and the aggregate level trends observed on the network. Travellers are assumed to make travel choices based on network level attributes, such as the flow dependent travel times on the road network (Wardrop, 1952). Network level data such as demand, flows or travel times are generated by the aggregated decisions of all travellers. Variability at the traveller level may not result in variability at the network level, however, if masked by counteracting behaviour. Also, a trend observed at the network level could be caused by any one of a range of possible explanations at the traveller level. Undertaking analysis at only one level, therefore, may not provide a complete picture of traveller behaviour. The unique contribution of the current research is to propose an inductive approach which involves preliminary analyses from both the network and traveller perspective, before making comparisons, then generating hypotheses to test using the alternative perspective. This type of research not only provides insights into traveller behaviour, but may also provide details of bias within the data sources used.

The structure of this paper is as follows. Section 5.2 includes a description of the three stages of the general approach proposed in this paper. In Section 5.3, the case study site which will be used is described. The three stages of the process are applied to the case study area in Sections 5.4, 5.5 and 5.6 respectively. Section 5.7 discusses other applications of this approach including how it could be applied to larger units geographically. Section 5.8 concludes the paper.

5.2 Methodology

Every person makes a series of decisions on each day which includes spatial and temporal choices. As discussed in Section 5.1, day-to-day variability in travel behaviour can be examined from different perspectives. For example, Figure 5-1 is a diagrammatic representation of how data could be aggregated under the network-focused and the traveller-focused perspectives for a very simple case where five days of data are collected on a link used by just three people. Each rectangle in the diagram contains a complex collection of travel choices for one person on one day, which will include the spatial and temporal characteristics of any trips made.

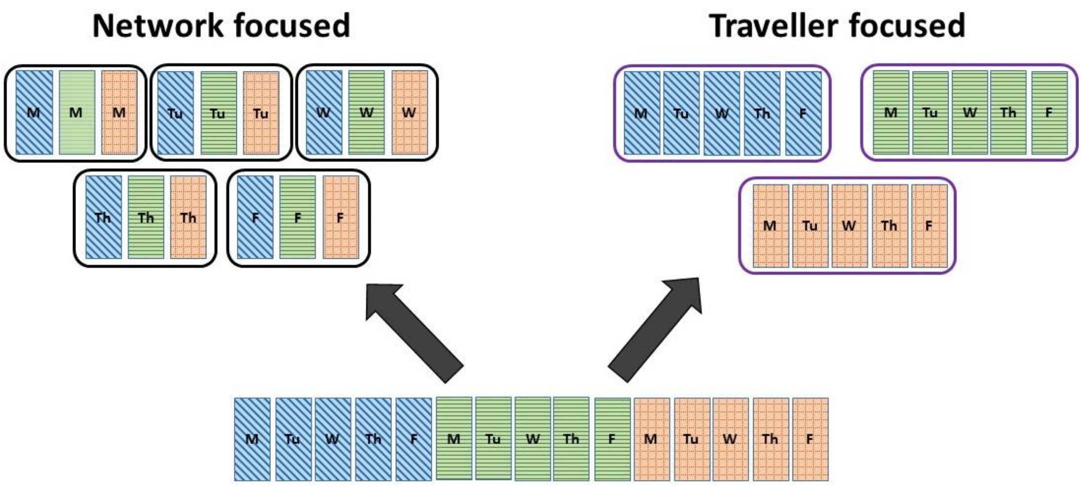

Figure 5-1: Simplistic representation of data aggregation for traveller behaviour analyses

In network-focused analyses, data is aggregated over all travellers for each day. If variability in the time of day at which trips are made is of interest, then the aggregated data should retain some degree of temporal disaggregation. This type of data could include daily flow profiles for a loop detector or hourly counts of vehicles crossing a bridge/tunnel.

Traveller-focused analyses examining day-to-day variability require data which contains unique traveller identifiers so that data can be matched between days, otherwise the aggregation shown in Figure 5-1 would not be possible. The data may therefore be less well defined spatially, but will still retain data about the timing of events during the day. Such data could include Bluetooth, mobile phone, travel diary or smart card data.

For some transport networks it may be possible to have identifiable traveller data for most, if not all, trips, for example for smart card schemes with very high levels of user registrations as seen in Ma et al. (2013). In most cases, however, individual traveller data containing unique identifiers are only available for a sample of users. It is likely, therefore, that different data sources will be used for the network- and traveller-focused analyses, although this is not a requirement. There can be advantages to using different data sources for the two sets of analyses, as they may include a wider range of variables when considered together. As the aim is to make connections between the traveller- and network-focused analyses, they should both contain data for the same (or very close) geographic areas.

As discussed in Section 5.1, further insights may be obtained from connecting the analyses undertaken from the two separate perspectives. The approach proposed in the current paper is inspired by mixed methods techniques and involves the three stages shown in Figure 5-2. In mixed methods approaches, qualitative and quantitative methods are used in the data collection and analysis stages (Johnson et al., 2007). The current approach only uses quantitative analyses and therefore may be considered to be a multiple method approach (Davis et al., 2010) rather than a mixed methods approach, as it does not have the philosophical basis that it is best to use both qualitative and quantitative analyses to understand behaviour (Creswell and Plano Clark, 2011, p5). Johnson et al. (2007, p123) define mixed methods as being "for the broad purposes of breadth and depth of understanding and corroboration", which coincides with the purpose of the current research. Creswell and Plano Clark (2011, p9) also describe how mixed methods can be used to explain or generalise initial results, and more generally *enhance* a study. Given these parallel objectives, it is reasonable to seek inspiration for the approach developed in the current paper from mixed methods research.

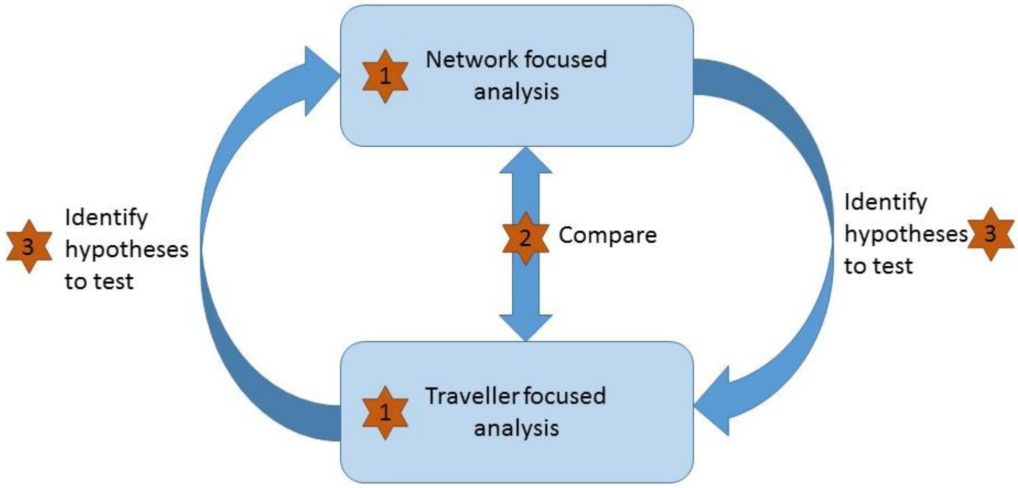

Figure 5-2: Overview of proposed approach for analysing travel behaviour using network- and traveller-focused analyses

The first stage involves performing analysis on data from each perspective independently. The initial analysis is performed concurrently, and interaction will not occur until Stage 2. The two sets of analyses should relate to the same time period and should have the same (or very similar) spatial coverage.

Stage 2 aims to bring together the separate analyses from Stage 1 in two ways. Firstly, the basis for comparison must be investigated by considering the data used. Where the traveller-focused analysis takes place on a sample of travellers, as will usually be the case, the proportion of travellers sampled should be examined for different day types and at different times of day to check for bias in the sample. Secondly, the *findings* of the two sets of analyses should be compared, in a similar way to the ‘triangulation protocol’ (O'Cathain et al., 2010, p1) used in mixed methods approaches. The findings from the two separate analyses should be compared to look for both agreement and disagreement but also to consider whether the findings provide complementary evidence on the same topic (O'Cathain et al., 2010, p2).

Stage 3 is consistent with the ‘Following a thread’ technique (O'Cathain et al., 2010, p2) sometimes used in mixed methods approaches. This involves identifying a question or theme from the analysis of each data type and exploring it in more detail using the other data type (Moran-Ellis et al., 2016). For the current application, this means developing hypotheses or questions to explain the network level variability which can be tested using traveller-focused analyses, and vice versa.

An application to a road link in northern England will now be presented to show how the process works and what kinds of insights can be gained from this approach. The approach is not limited to such small geographic areas or to travel behaviour of road users only, however, and further applications will be discussed in Section 5.7.

5.3 Case study description

The case study area consists of a road link going north from the Manchester Outer Ring Road (M60) and connects Stockport with Manchester (see Figure 5-3). Only the northbound direction has been considered. The purpose of the analysis is to examine day-to-day variability in travel behaviour involving this link. As only one location is involved, the definition of travel behaviour will be limited to whether the traveller uses this road link and if so, at what time. Within-day dynamics must, therefore, also be examined. A fully comprehensive analysis of the data available will not be provided as the aim is to demonstrate the proposed approach in practice. Day of the week effects will be examined in detail, but other variability such as seasonal variations or the impact of sporting events will not be considered.

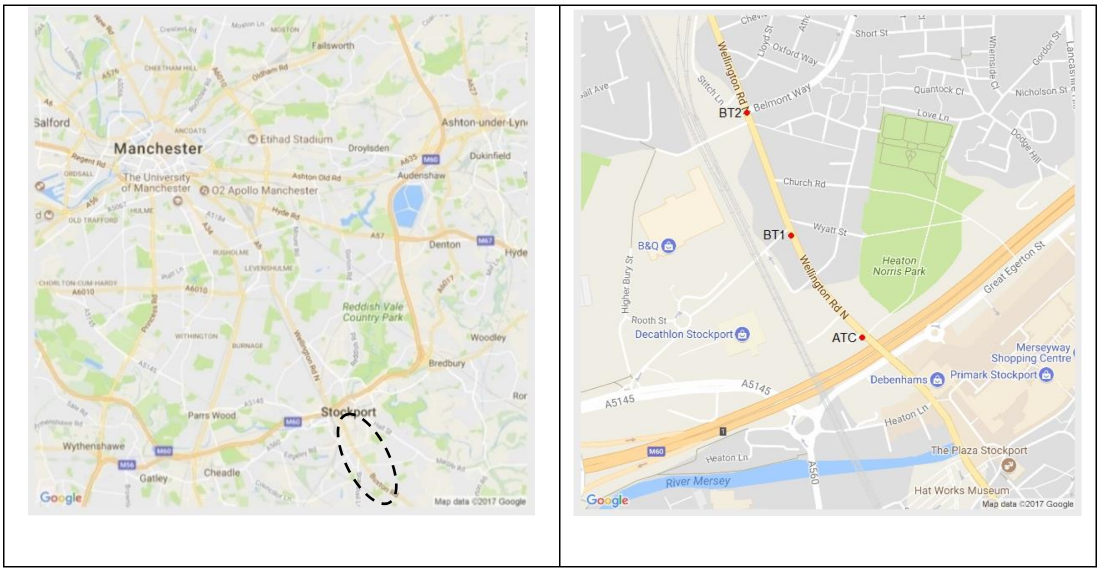

Figure 5-3: Stockport case study area

The network-focused analysis will use data from the induction loop detector labelled ATC in Figure 5-3. The traveller-focused analysis utilises data from fixed Bluetooth detectors. Bluetooth detectors collect unique identifiers of discoverable Bluetooth devices which pass within the communication zone of the detector (Bhaskar and Chung, 2013). The types of Bluetooth devices currently available include mobile phones, fitness trackers, hands-free devices and in-car audio systems. The ability to match the unique identifiers over time and space has resulted in Bluetooth data being used for travel time estimation (Quayle et al., 2010, Moghaddam and Hellinga, 2014, Mathew et al., 2016), OD estimation (Barceló et al., 2010, Carpenter et al., 2012) and for classifying road users based on repeated travel behaviour (Crawford et al., Under review-b). In the current research, data from the two Bluetooth detectors shown in Figure 5-3 are used. After matching observations between the two detectors and cleaning the matched data, only trips travelling from BT1 to BT2, i.e. northbound, were retained. There is assumed to be limited traffic joining or leaving the link between ATC and BT2. Data for the same period of time was used for both types of data, namely 1/5/2013 to 30/4/2015.

5.4 Stage 1: Separate analyses

Firstly, network-focused and traveller-focused analyses are undertaken separately using methods which are appropriate for the data available.

5.4.1 Network-focused analysis of loop detector data

In Crawford et al. (2017), data from the loop detector in this case study was examined for the same period of time (1/5/2013 to 30/4/2015), to test for statistically significant differences in daily flow profiles according to the day of the week or season. Crawford et al. (2017) examined day of the week effects in the loop detector data by separately analysing the magnitude of flows and the within-day distribution. Testing was undertaken for statistically significant differences in magnitudes, using ANOVA, and in the standardised daily flow profiles, using functional ANOVA and permutation tests.

The total counts and average flow profiles by the day of the week are shown in Figure 5-4. Using the Kruskal-Wallis test, the total daily flows were found to differ significantly according to the day of the week, with two exceptions: Thursdays and Fridays were not significantly different, and neither were Tuesdays and Wednesdays. The standardised daily flow profiles were assessed using Functional Linear Models and functional ANOVA and the average profile for each day of the week was found to be significantly different.

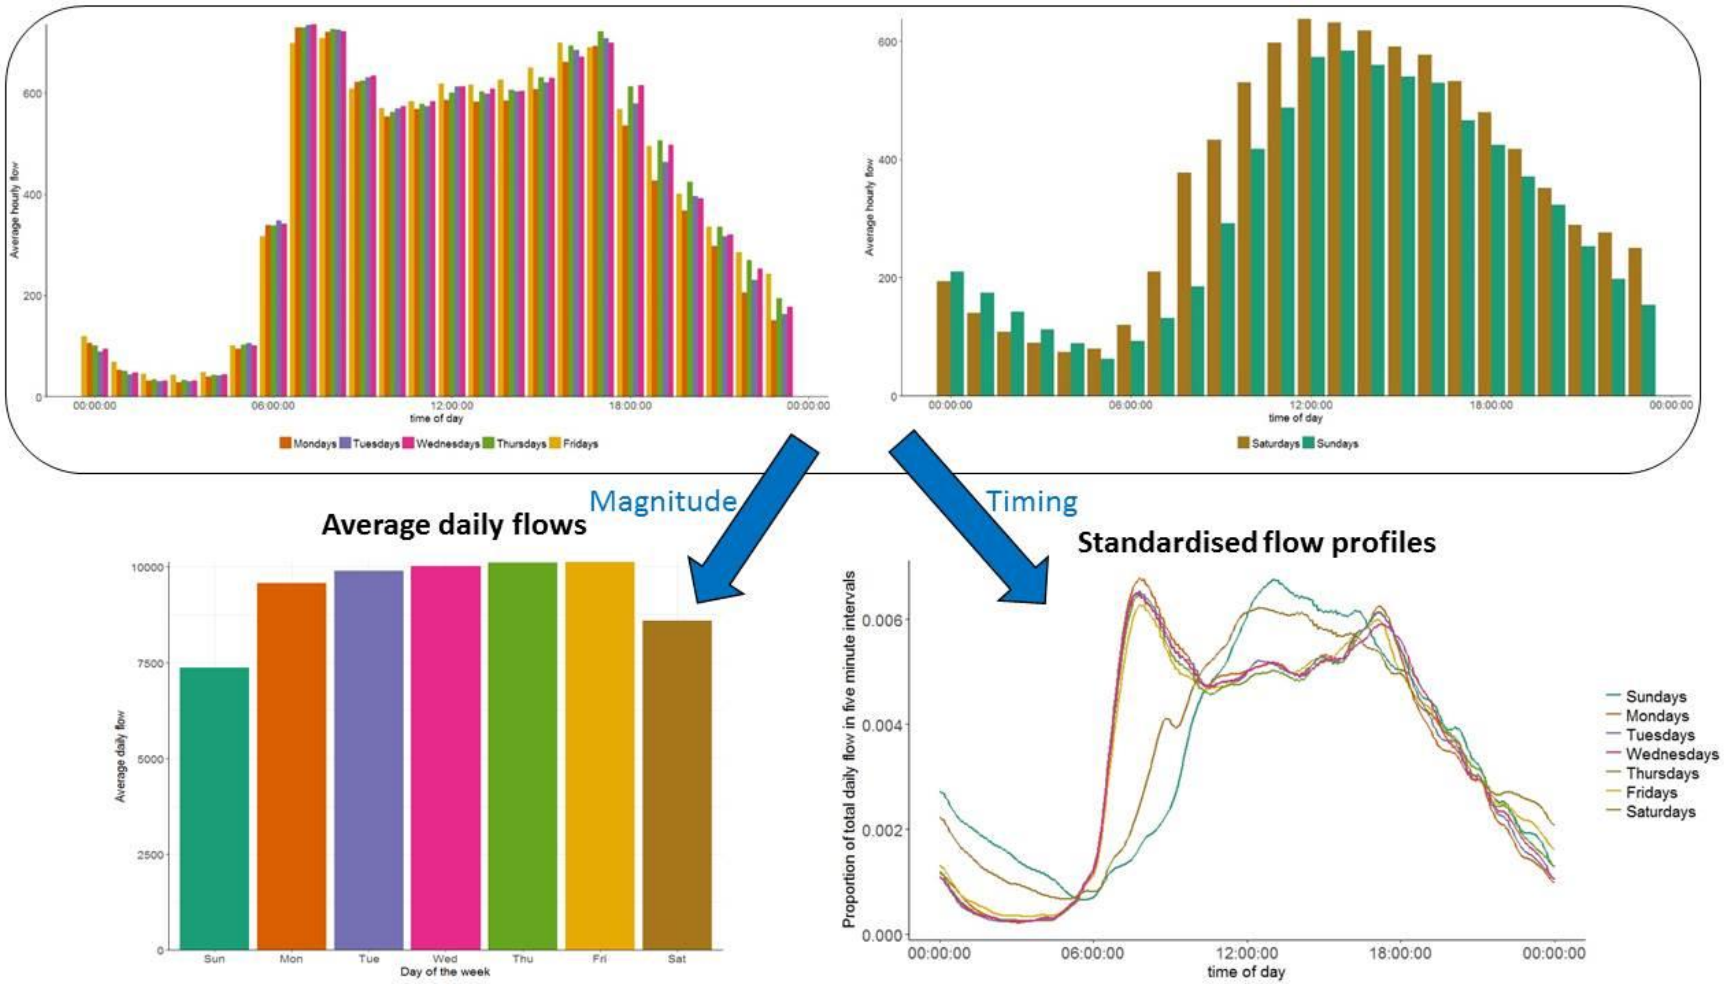

Figure 5-4: Day of the week flow patterns observed at one loop detector site by Crawford et al. (2017)

5.4.2 Traveller-focused analysis of Bluetooth data

The initial traveller-focused analysis was undertaken using the methodology proposed by Crawford et al. (Under review-b). Their method involved calculating trip frequency, spatial and temporal variability measures for each traveller and then using these to cluster travellers into user classes with similar trip characteristics. In the current application, data from only one pair of Bluetooth detectors is used, unlike in Crawford et al. (Under review-b) where 23 Bluetooth detectors were included, and therefore as there is no spatial variability, only the trip frequency and temporal variability measures were included.

The trip frequency measure is simply the number of observations in the two year period. The time of day variability measure is based on the assumption that the times of day a person passes the sensors will form clusters. These clusters could be due to different trip purposes or different daily routines, for example according to the day of the week. As in Crawford et al. (Under review-b), Model Based Clustering (Fraley and Raftery, 2002) is performed separately for each traveller. For each traveller the measures of time of day variability are the number of time of day clusters their trips form and their average variance. The trip frequency and two time of day variability measures are then standardised and used to identify user classes of travellers using k-means clustering.

The 1.1 million trips observed travelling from BT1 to BT2 were made by 197,474 different MAC addresses, which will be referred to as travellers in this research. Four user classes were identified: travellers observed approximately once or twice, a few times per year, a few times per month and a few times per week. Within the first three user classes there are subclasses defined by different degrees of time of day variability, as shown in Table 5-1.

Table 5-1: Road user classes and their characteristics

User class	User subclass	Average trip frequency	Average number of time of day clusters	Average variance of time of day clusters	Number of travellers	Total number of trips
A Annually	A1	1.8	1	0.001	152,690	277,917
	A2	2.3	1	0.330	462	1,060
B Three times per year	B1	5.3	1	0.162	2,338	12,299
	B2	5.4	1	0.073	6,666	36,231
	B3	8.4	1	0.024	27,524	230,247
C Fortnightly	C1	51.5	2	0.008	5,993	308,792
	C2	62.1	5	0.003	1,316	81,713
D Three times per week	D1	298.3	4	0.006	485	144,683

Figure 5-5 shows that while the least frequent user class (A) makes up 78% of the travellers observed, it only accounts for around a quarter of the trips observed. Another quarter of the trips are made by those observed a few times per year (user class B), and the remaining 49% of trips are made by travellers observed approximately once per fortnight or more. One possible explanation is that these observations represent all trips made by those travellers which pass these detectors. Another possible explanation is that due to disabling of Bluetooth devices and incomplete detection of enabled devices, these frequencies only relate to a small proportion of each traveller's trips passing the sensors. The truth will lie somewhere between these two extremes. There are over seven and a half thousand travellers in user classes C and D, however, which provides a fairly large dataset within which to examine intrapersonal variability in travel behaviour. For example, of the travellers in user class C, who are observed approximately once a fortnight, the majority (82%) are in the subclass with more variability in the times of day they pass this site.

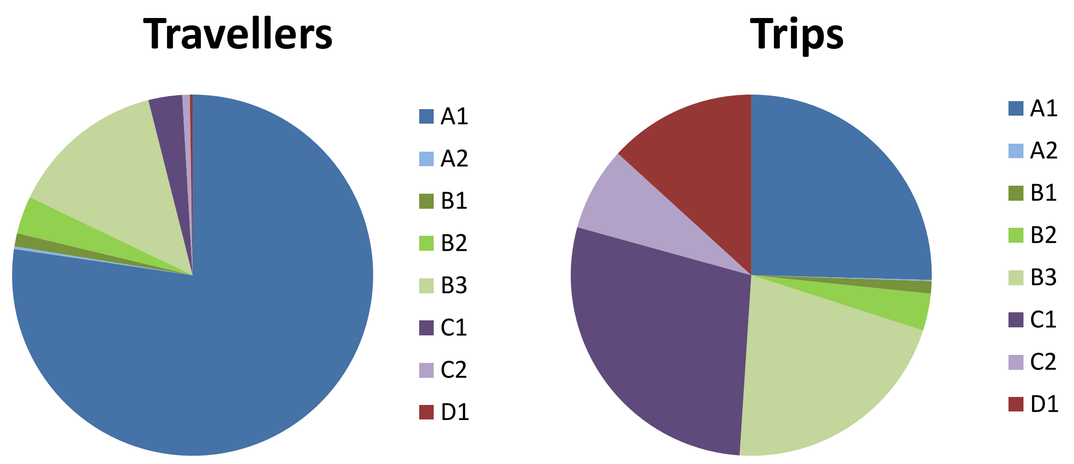

Figure 5-5: Splits of travellers and trips by user subclasses

Figure 5-6 demonstrates that the distribution of trips between user classes is relatively constant by the day of the week. The difference in the proportion of trips made on weekdays compared with weekend days is statistically significant for each of the subclasses ($|z|$ range from 9 to 55). Weekend days include a slightly lower proportion of trips by members of the two most frequent user classes, but the proportion of trips made by travellers who pass the site every three or four months on average is higher. As will be seen in Section 5.5, fewer trips are recorded by the Bluetooth detectors on Saturdays and Sundays than on weekdays, but the absolute number of trips for user subclasses A2, B1 and B2 are actually higher at the weekend. These are the three subclasses generating the fewest trips, however, so the number of travellers involved is relatively low.

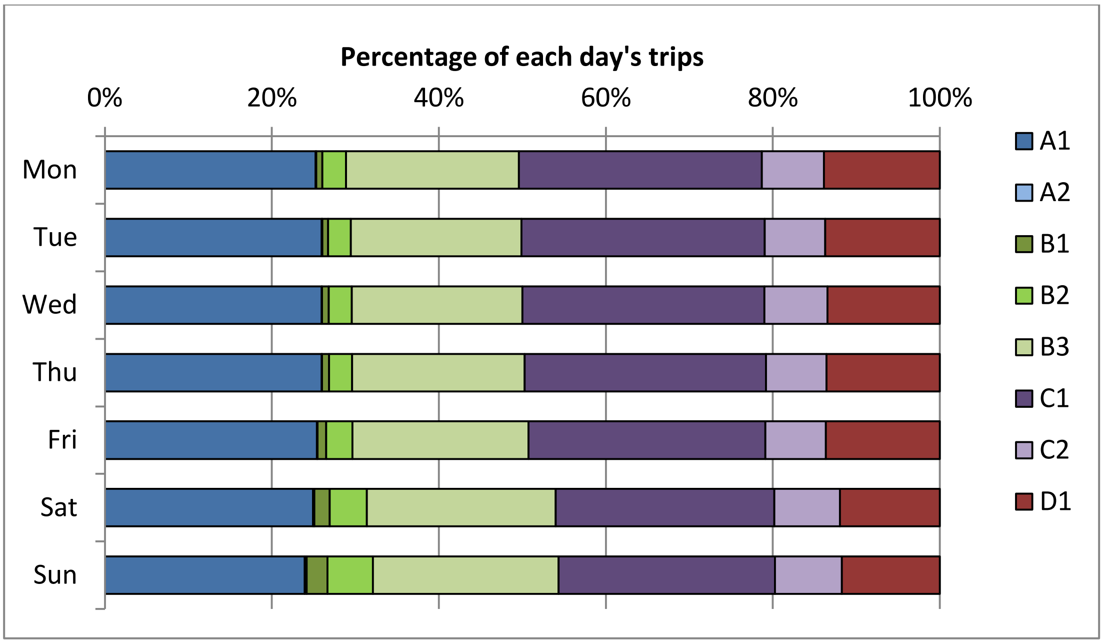

Figure 5-6: User classes by the day of the week

Examining the time of day trips are made by each user class may provide additional insights into what the classes represent. Figure 5-7 includes the total count for each hour of the day according to user classes. The more frequent user classes (D1, C1 and to some extent C2) exhibit a daily profile consistent with the double peak weekday profile observed on many roads used for commuting. This pattern is not present for travellers observed less frequently. For example, the two largest user subclasses (in terms of trips) after C1 are A1 and B3, and both of these have a unimodal daily profile with the most trips observed in the middle of the day. As observed in Figure 5-6, these subclasses are not observed more often during the weekend, but it is likely that more of the trips made by these travellers will be non-commuting trips, compared to travellers in other classes.

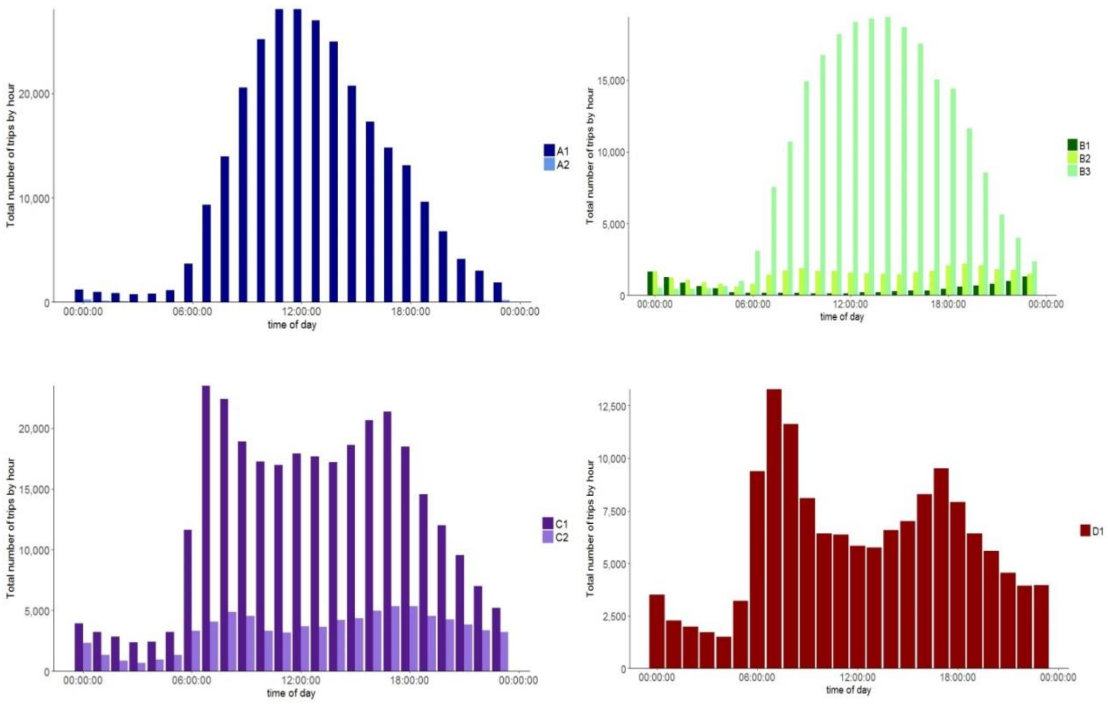

Figure 5-7: Total number of observations of travellers from each user class by the hour of the day

The *proportion* of trips made by each user class in each hour of the day was also examined (Figure 5-8). The most interesting observation is that the proportion of trips made by user class B is relatively stable throughout the day, with the exception of the morning peak period (approximately 5-9am). The distribution amongst the three user *sub*classes, B1, B2 and B3, however, differs according to the time of day. The travellers with lower trip frequencies and higher time of day variability, namely B1 and B2, are more common during the evening, night and early hours (from 6pm to 7am).

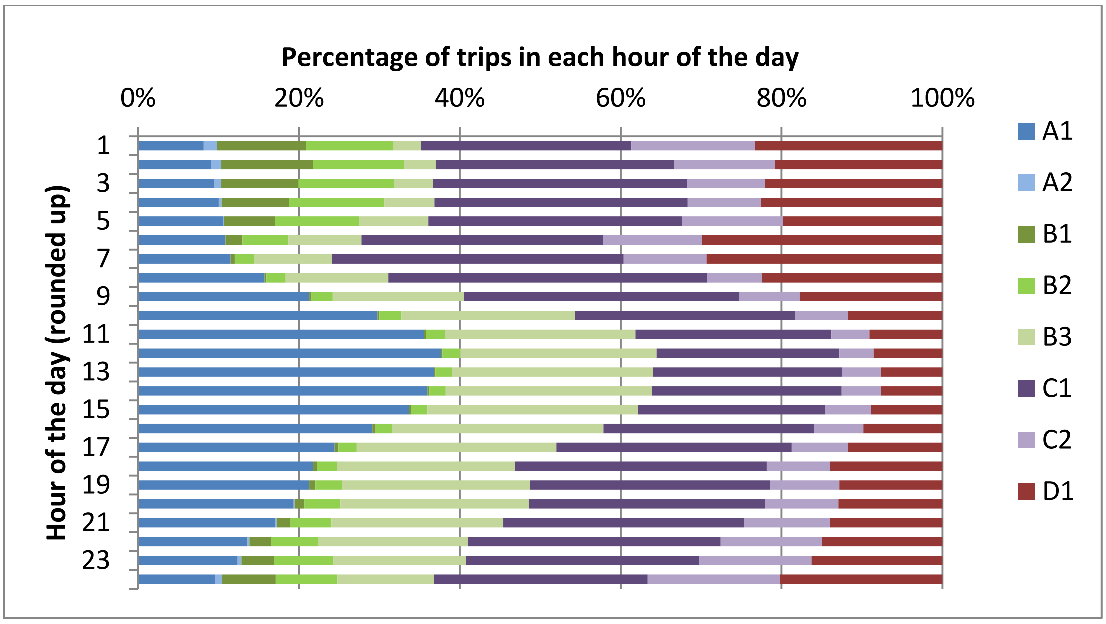

Figure 5-8: Percentage of trips by each user class by the hour of the day

In summary, when clustering the Bluetooth observations according to frequency and time of day variability, four user classes were identified. Almost half of the observations relate to travellers observed at least once a fortnight over the two year period, on average. The user classes can be described based on trip frequency, but the three least frequent traveller user classes are split into subclasses based on the time of day variability in trips. The user class proportions are fairly stable over the days of the week, although on weekend days a smaller proportion of trips are made by the most and the least frequent traveller subclasses. The user classes differ in terms of the times of day that trips are observed. The more frequent traveller classes are more likely to be observed during the typical weekday morning and evening peaks. The less frequent traveller classes are more likely to be observed during the inter-peak period. The proportion of trips made by each *sub*class varies according to the time of day for each user class.

5.5 Stage 2: Comparisons

5.5.1 Data comparison

The extent to which Bluetooth data can be used to explain differences in loop detector data, and vice versa, will depend upon the combined Bluetooth penetration and capture rates and the positioning of the detectors. The total number of trips detected in each day in the two year period (excluding public holidays) was compared. The Bluetooth detectors capture approximately 16% of the trips passing this location. Given that this equates to almost 1.1 million trips travelling northbound in the two year period, this is a relatively large amount of data to use to explore traveller behaviour. Bias in the trips detected should also be considered, however. Prior to comparing the findings of the analyses in Sections 5.4.1 and 5.4.2, therefore, the Bluetooth sample sizes will be compared to the loop detector counts by the two main groupings used in this research, namely the day of the week and the time of day.

The total counts by the day of the week from the loop detector and the Bluetooth detectors over the same period of time are shown in Figure 5-9. The Bluetooth sample is between 16.2% and 16.5% of the loop detector counts across the weekdays, but it is slightly lower at the weekend (Saturday 15.3%, Sunday 15.0%). This could be because there are likely to be more commercial vehicles on the road on weekdays, and due to the longer periods of time drivers spend in such vehicles, they may be more likely to contain a Bluetooth-enabled device such as a hands-free kit.

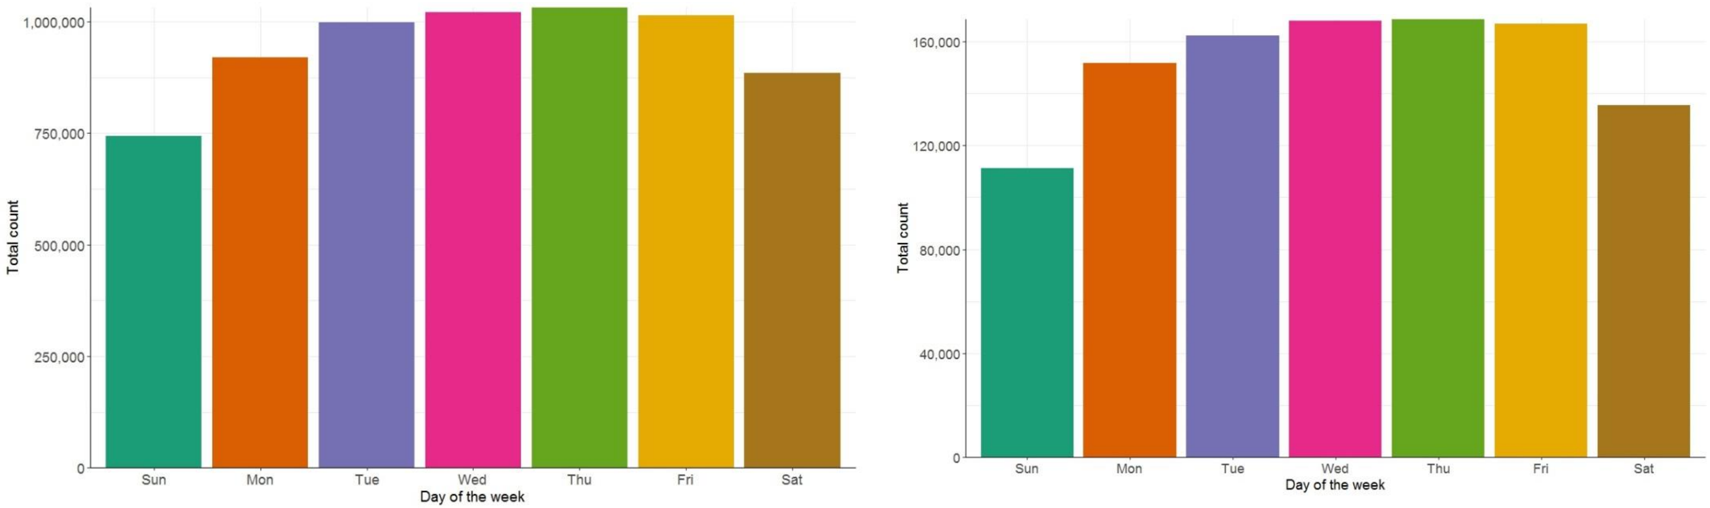

Figure 5-9: Total two year counts by day of the week for a) the loop detector and b) the Bluetooth detector

Figure 5-10 shows how the observations from the loop detector and the Bluetooth sensors are distributed according to the time of day. A two-sided two sample Kolmogorov-Smirnov test was applied, to test whether the Bluetooth and loop detector observations are likely to have the same underlying probability distribution. Bootstrapping was used to calculate the relevant p value as the data contained many ties as the observations had been rounded to the nearest second. The null hypothesis was rejected at the 0.05 significance level (D = 0.030979, p-value < 0.00001). Despite capturing a lower percentage of trips on weekend days, the hourly Bluetooth detections profile is unimodal and peaks between midday and 1pm. The percentage of the loop detector trips which are captured by the Bluetooth detectors is higher than average during two periods of the day. Firstly, it is slightly higher between 1-5am, peaking at 20% between 2-3am. This could be caused by a higher proportion of commercial vehicles where Bluetooth devices may be more prevalent, for example taxis and long distance trucks. The second period of the day with higher than average sampling rates for the Bluetooth detectors is the inter-peak period, from 9am-4pm. After this, sampling rates remain low for the rest of the day, with the minimum value of 13.5% in 10-11pm.

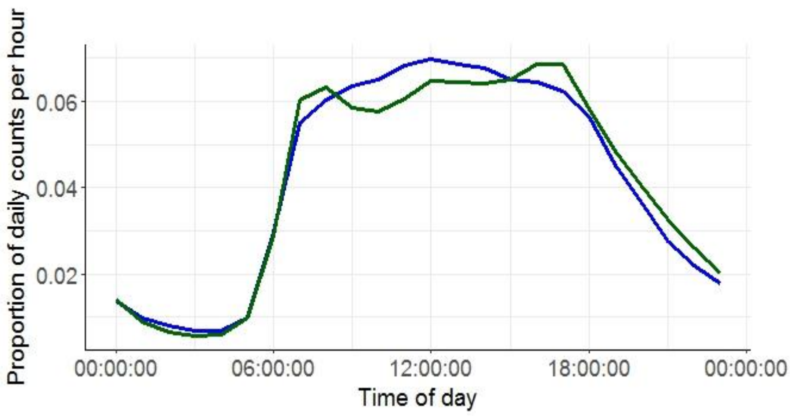

Figure 5-10: Standardised daily profile counts for the loop detector (green) and the Bluetooth detector (blue)

The lower Bluetooth sampling rates during the peak period could be caused by limitations in the capacity of Bluetooth sensors to store sufficient MAC addresses (the unique identifiers for Bluetooth devices) during each scanning cycle to be able to detect the same *proportion* of devices during very busy periods. An examination of individual days of data supports this theory. Figure 5-11, for example, includes histograms of the Bluetooth and loop detector observations for a randomly selected weekday and demonstrates that although there is a relationship between the magnitudes of the two sets of observations, the Bluetooth histogram is much flatter, which is consistent with there being caps on the number of possible observations during busy periods.

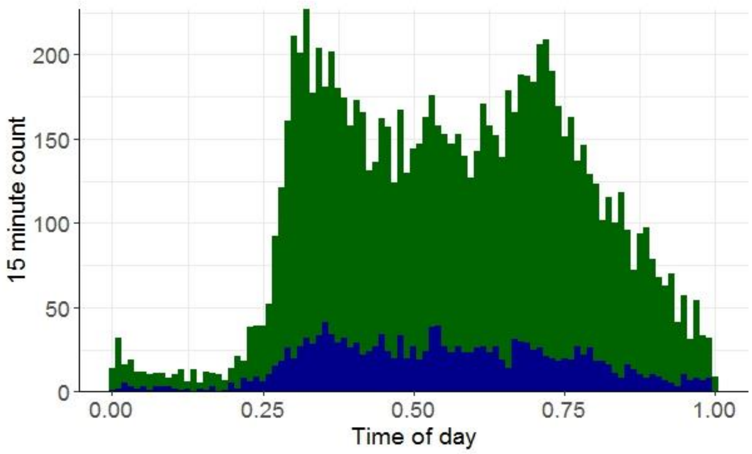

Figure 5-11: Histogram of counts from a randomly selected weekday (23/9/2014) for the loop detector (green) and the Bluetooth detector (blue)

Alternatively, the higher Bluetooth sampling rate during the inter-peak period could be due to a higher sampling rate for subclasses A1 and B3, as this is consistent with the daily distribution of such trips as shown in Figure 5-7. These are the second and third largest subclasses in terms of the number of trips. These are travellers detected quarterly or less frequently, on average, per year and therefore may relate to people who are less familiar with the area and therefore are using satellite navigation systems which may be Bluetooth-enabled.

The sampling rates suggest that the day of the week and the time of day have an impact on the collection of Bluetooth data. This is, of course, assuming that there are very few travellers joining or leaving this link between the loop detector and the most northerly Bluetooth detector. The varying sampling rate could indicate a bias in either the type of travellers and/or the type of trips which are detected using the Bluetooth sensors. Methods for collecting additional data to explore such issues are described in Crawford et al. (Under review-a).

5.5.2 Comparison of findings

Bearing in mind these possible limitations, the findings from Sections 5.4.1 and 5.4.2 will now be compared and contrasted. Despite the statistically significant differences in daily flow profiles between days of the week, the general shape of the weekday profiles from the loop detector data were very similar (Figure 5-4). The Bluetooth data does not suggest that this relative stability is the result of the same people making the same trips each day, however. Half of the trips observed are made by people observed less than once per fortnight, on average. The estimates of travellers' trip frequency from the Bluetooth data do not take into account travellers who vary their route so as not to pass this location, and will be an undercount due to devices being switched off, not carried, Bluetooth disabled or simply not being detected. Without further investigation into each of these factors, it is unclear how many of the times a traveller passes the sensor will actually be recorded by the Bluetooth detector. Also, the estimate of trip frequency is only based on travellers with a detectable Bluetooth device and this may not be a random sample from the wider travelling population.

The network-focused analysis highlighted that each day of the week had a daily flow profile with a different shape and total counts differed between Saturdays, Sundays and three weekday groups. The traveller-focused analysis identified user classes based on frequency and time of day variability of observations, but the proportion of each user class was relatively similar across the days of the week, particularly weekdays. The differences between days of the week are, therefore, relatively subtle as they must be caused by differences in behaviour of the same travellers, or different travellers within the same user class.

While the user classes identified in the traveller-focused analysis provide some insight into the behaviour of travellers passing this location, they are insufficient to explain the variability in trip timings and volumes observed at the link level. Stage 3 of the process is therefore required to gain further insights.

5.6 Stage 3: Following the thread

The third stage uses the findings from the network-focused analysis to identify questions to explore using traveller-focused analyses and vice versa.

5.6.1 From network-focused to traveller-focused analysis

In Section 5.4.1, systematic differences were observed in total daily flows and the shape of the daily flow profiles between days of the week. It is not clear, however, what underlying traveller behaviour is causing these differences. Two key questions which arise are:

- Do different people travel on systematically different days of the week?
- Do individual people travel at systematically different times of the day on different days of the week?

5.6.1.1 Days of week travelled

One explanation for the systematically different volumes and daily profile shapes by the day of the week is that different people are travelling on different days of the week, and those people may have different travel behaviour. The traveller-focused analysis in Section 5.4.2 indicated that user class proportions were consistent between weekdays, but the network-focused analysis showed systematic differences. A more detailed examination is required, therefore, and so individual traveller behaviour will be examined. The null hypothesis to be tested is that each traveller's weekday trips over the two year period are evenly distributed from Monday to Friday.

In this section, weekdays only will be considered as these are usually assumed to have a fairly constant set of travellers. A similar analysis can be undertaken to compare all days, or weekend days only. To ensure sufficient data is available for testing, only the travellers in user classes C and D, namely travellers observed at least once a fortnight on average, will be analysed. For weekdays only, therefore, each traveller's observed counts by day of the week (over the two year period) were compared to the expected counts under the null hypothesis of evenly distributed weekday trips as shown in Table 5-2.

Table 5-2: Format of observed and expected day of the week counts for each traveller

	Monday	Tuesday	Wednesday	Thursday	Friday	Total weekday
Observed counts	x_M	x_{Tu}	x_W	x_{Th}	x_F	$x_A = x_M + x_{Tu} + x_W + x_{Th} + x_F$
Expected counts	$\frac{x_A}{5}$	$\frac{x_A}{5}$	$\frac{x_A}{5}$	$\frac{x_A}{5}$	$\frac{x_A}{5}$	x_A

The comparison was undertaking using a chi squared test. This required the assumption that the observations are independent. This was considered reasonable given that the probability of being detected by a Bluetooth sensor is independent between days, and the probability of making a trip which passes the detector was assumed to depend upon the day of the week, but not on whether the detector was passed on previous or subsequent days. Any reduction in sampling rates when volumes of traffic are high is assumed to be a fairly mild effect, particularly as the test is examining observations for each traveller separately. The chi squared tests could only be applied to travellers with at least 25 weekday trips, so that the chi squared criteria of expected values being at least 5 is satisfied. Of the 7,794 travellers who are in the two most frequent user classes, 5,125 travellers satisfied this criterion.

For 20% of the travellers assessed, the null hypothesis, that weekday trips were evenly distributed over Monday to Friday, was rejected at the 95% level. Therefore, while the majority of travellers are likely to travel equally on all weekdays over the two year period, there are travellers who do not. These travellers, who will be referred to as uneven weekday travellers, make 19% of the trips made by the travellers in user classes C and D. The null hypothesis was rejected for a slightly higher percentage of travellers in user class D (26%) than in the less frequent user class, C. For user subclass C2, the subclass with lower time of day variability, uneven weekday travellers make up a higher percentage of the travellers (24%) than in user subclass C1 (19%).

There are more differences between the uneven weekday and the even weekday travellers than just the distribution of their trips over the days of the week. As shown in Figure 5-12, the uneven weekday travellers have lower variability in the time of day their trips are made. Therefore, although they do not necessarily travel on all weekdays, these travellers pass the sensors at more regular times of day when they do travel.

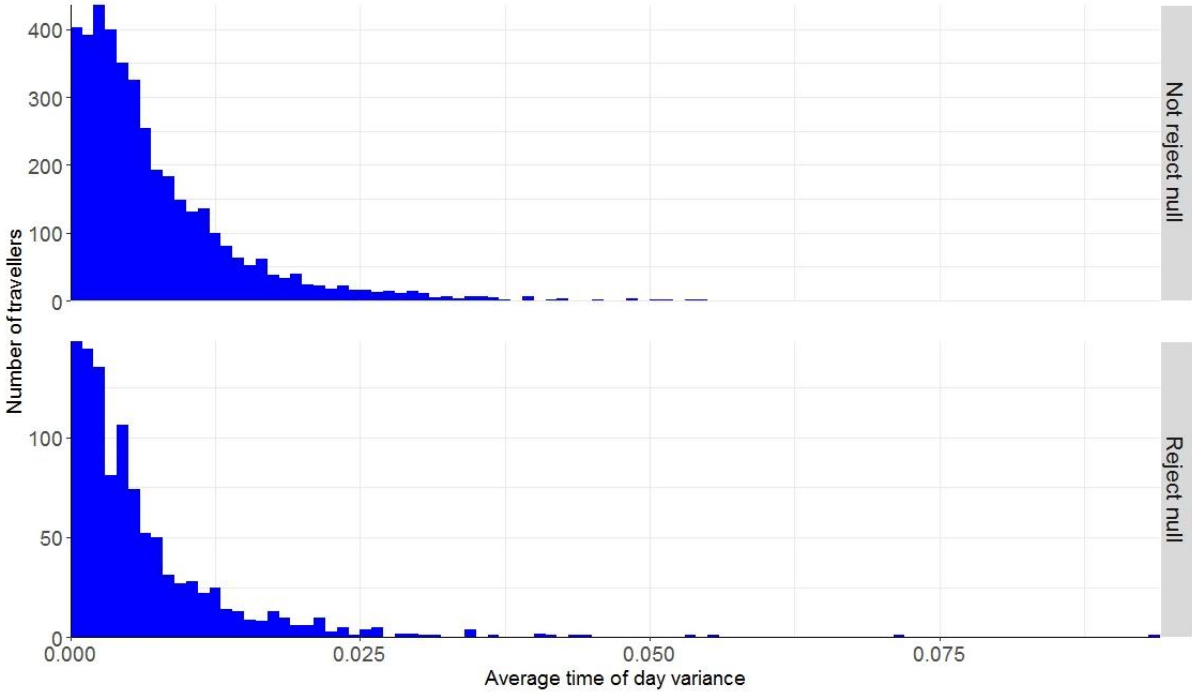

Figure 5-12: Histograms of traveller time of day variance for travellers where there was insufficient evidence to reject the null hypothesis of evenly distributed weekday trips, and for travellers where the hypothesis was rejected

Among the 1,050 travellers where the null hypothesis was rejected there are 30 different combinations of days of the week with higher and lower than expected observed counts. If the observed count exceeded the expected count on only one day, this was most likely to have been a Friday. For two days, they were most likely to be Thursday and Friday, and for three days it was Wednesday, Thursday and Friday.

Therefore, whilst the majority of travellers who are observed approximately once a fortnight or more, on average, distribute their weekday trips evenly over Monday-Fridays, approximately 20% of travellers do not. This figure rises to 26% for the most frequent traveller class. The uneven weekday travellers had different time of day regularity characteristics compared to the even weekday travellers. There were also systematic differences in the days of the week which had a higher percentage of a traveller's trips, with a bias towards days later in the week.

5.6.1.2 Times of day travelled

An alternative explanation for systematic differences in daily flow profiles could be that individuals travel at systematically different times of day on different days of the week due to different activity patterns. This can be explored using the outputs from the Model Based Clustering of times of day used in Section 5.4.2.

To determine whether these time of day clusters relate to different days of the week, chi squared tests were applied. As shown in Table 5-3, the observed values for each traveller are the counts of trips in each time of day cluster, separated by the day of the week. As in all of this section, only weekdays are considered. The observed values are then compared with the expected values using a chi squared test. As shown in Table 5-4, the expected value for each cell is calculated based on the proportion of the traveller's trips in that time of day cluster and that day of the week. The test, therefore, *accounts* for differences across days of the week and *tests* for differences across the time of day clusters.

Table 5-3: Observed counts by time of day cluster and day of the week for one traveller

	Monday	Tuesday	Wednesday	Thursday	Friday	Total weekday
Time of day cluster 1	$x_{M,1}$	$x_{Tu,1}$	$x_{W,1}$	$x_{Th,1}$	$x_{F,1}$	$x_{A,1}$
Time of day cluster 2	$x_{M,2}$	$x_{Tu,2}$	$x_{W,2}$	$x_{Th,2}$	$x_{F,2}$	$x_{A,2}$
….	….	….	….	….	….	….
Total for traveller	x_M	x_{Tu}	x_W	x_{Th}	x_F	x_A

Table 5-4: Expected counts by time of day cluster and day of the week for one traveller

	Monday	Tuesday	Wednesday	Thursday	Friday	Total weekday
Time of day cluster 1	$\frac{x_M * x_{A,1}}{x_A}$	$\frac{x_{Tu} * x_{A,1}}{x_A}$	$\frac{x_W * x_{A,1}}{x_A}$	$\frac{x_{Th} * x_{A,1}}{x_A}$	$\frac{x_F * x_{A,1}}{x_A}$	$x_{A,1}$
Time of day cluster 2	$\frac{x_M * x_{A,2}}{x_A}$	$\frac{x_{Tu} * x_{A,2}}{x_A}$	$\frac{x_W * x_{A,2}}{x_A}$	$\frac{x_{Th} * x_{A,2}}{x_A}$	$\frac{x_F * x_{A,2}}{x_A}$	$x_{A,2}$
….	….	….	….	….	….	….
Total for traveller	x_M	x_{Tu}	x_W	x_{Th}	x_F	x_A

In order to apply the chi squared test, no cells should have an expected value of zero and no more than 20% of cells should have expected values of less than five. To satisfy these criteria, only time of day clusters with at least 25 observations were included. For each traveller, any days of the week with fewer observations than five times their number of time of day clusters were excluded from the analysis.

Of the 7,794 travellers who are in the two most frequent user classes, only 1,077 travellers had more than one time of day cluster of sufficient size to compare. It should be noted, therefore, that this test of systematic variability according to the day of the week could only be undertaken on 0.5% of travellers accounting for 14% of the trips observed in the two year period. The hypothesis that observations were evenly distributed between time of day clusters for each day of the week was rejected for 33% of these travellers. Therefore, for one third of travellers who frequently pass this location at more than one time of day on weekdays, the time of day that they pass appears to vary systematically by the day of the week. There are many reasons why this may be the case, for example some people may work half days on Fridays, or arrive home a little later on Thursday evenings due to extended shopping hours.

In larger datasets, it may be possible to do further examinations. For example, travellers where this hypothesis was rejected could be tested based on more specific hypotheses arising from the network-focused analyses. The analysis in Section 5.4.1, for example, may generate the following questions:

- Do people travel home later on Wednesday evenings?
- Do people travel home earlier on Friday afternoons?

In the current case, however, there was insufficient data to perform a statistical analysis of these hypotheses. Of the 352 travellers where the hypothesis of evenly distributed day of the week counts by time of day cluster was rejected, only 38 had more than one time of day cluster which fell in the evening peak period (defined as 4-7pm).

5.6.2 Traveller-focused to network-focused analysis

In Section 5.4.2, measures of frequency and time of day variability relating to the case study location were calculated and a data driven approach was then used to identify classes of users. In order to understand a little more about the nature of the user classes, and to examine the relationship between traveller and network level variability, the following two questions are proposed:

- Do the traveller user classes correspond to vehicle classes?
- Do times of the day where individuals have more time of day variability correspond to times of day with more variability in flows?

5.6.2.1 Do any of the user classes observed correspond to vehicle classes?

The individual analysis of the Bluetooth data in Section 5.4.2 highlighted the presence of different classes of road user, where trip frequency is a key differentiating factor for the higher level classes. It could be hypothesised that the most frequently observed vehicles relate to commercial vehicles which make many trips, for example buses or delivery vans. Alternatively, the comparisons in Section 5.5 suggested that the higher Bluetooth sampling rates in the early hours of the morning, and perhaps also in the inter-peak, could be due to commercial vehicles which may be more likely to contain Bluetooth-enabled devices.

Vehicle class data was not available from the Bluetooth data. The loop detector data used for the network-focused analysis *does* include vehicle class information, however. The total counts by vehicle class and by the hour of the day could be used to examine the theories proposed in the earlier analyses.

Figure 5-13 includes the distribution of counts for each vehicle class by the hour of the day, for the loop detector in the case study area. It should be noted that the upper plot represents cars and is plotted on a different scale as the counts are approximately 20 times greater than the counts observed for the other vehicle classes. In fact, cars make up 92% of the vehicles counted by the loop detector over the two year period. At this site, therefore, vehicles other than cars make up such a small proportion of trips that they cannot make up a significant proportion of the trips in any user class.

Other than articulated lorries and cars with trailers, both of which represent very few counts, only the rigid lorries have a daily profile which does not peak during the typical morning and evening peak periods. It is possible that the higher proportion of these types of vehicle during the inter-peak period is responsible for some of the increased Bluetooth sampling rate during that time.

Although taxis may also contribute to the commercial vehicles which are theorised to be more likely to contain a Bluetooth-enabled device, these are not differentiated from cars in the loop detector data and no other suitable data was available to examine taxi flows at the time of this analysis.

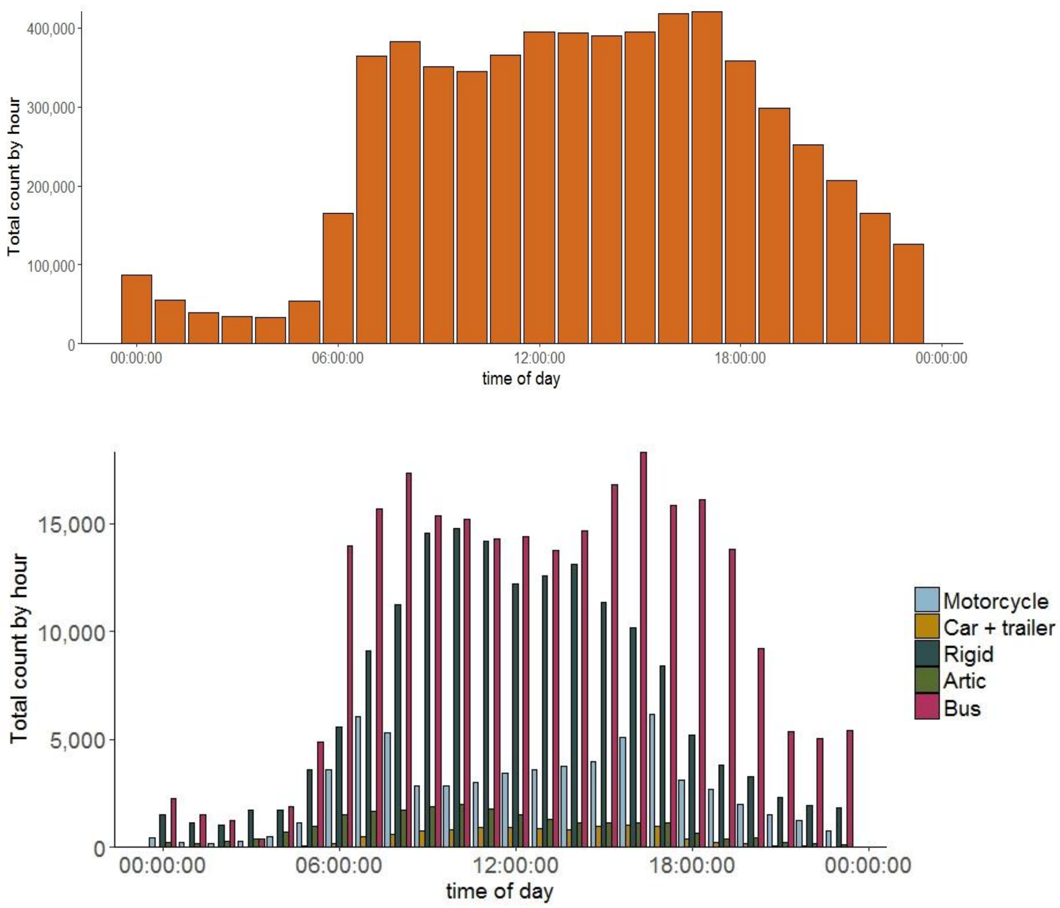

Figure 5-13: Counts of each vehicle class according to the time of day

5.6.2.2 Does individual time of day variability equate to variability in flows?

The method for calculating time of day variability in Section 5.4.2 provides traveller-focused measures which could be compared to network-focused measures. One comparison which could be made it to examine whether the times of day with higher traveller variability in the timing of trips passing the sensors equate to the times of day with higher variability in flows. The variability in times for individual travellers comes from two sources: the different daily travel patterns, represented by different clusters, and the variability within those clusters. There are several different reasons why a traveller may have multiple time of day clusters. The traveller could pass the sensors more than once per day or the clusters could represent trips on different days of the week, for example. Without further investigation, therefore, it is not possible to say whether the between cluster variance for each traveller relates to variability in trip timing. This subsection will, therefore, only consider within cluster variations and so the results should acknowledge that there is likely to be additional systematic variability in individual travel behaviour which has not been taken into account in this analysis.

Figure 5-14 shows how the average time of day cluster variance changes during the day. The morning peak period has particularly low traveller trip timing variability, perhaps reflecting the more rigid schedules which affect morning activity patterns. On average, more variability is observed in individual traveller trip timing in the mid-afternoon than at any other time of day. This may reflect trips to and from schools, which vary due to after-school activities and long holiday periods.

Figure 5-15 shows the variability in flows on the road link, measured using the loop detector. The variability is highest during the morning peak period, with a much lower peak during the evening peak. A similar pattern is observed if 15 or 30 minute intervals are plotted instead of hourly intervals. These do not correspond to the times of day where individual travellers have the most variability in the times of day that they pass the sensor. This suggests that the variability in flows is more likely to result from the decision to make a trip passing the sensors or not, rather than the timing of the trip. Individual trip timing could also contribute to the variability in flows through the *systematic* variability in travel times which were not considered in the analysis in this subsection.

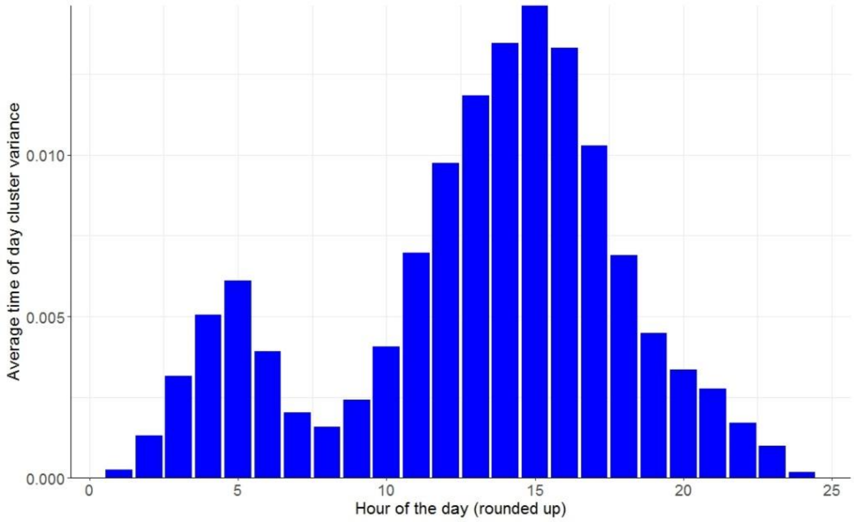

Figure 5-14: Average time of day cluster variance across all travellers by the hourly interval containing the cluster mean

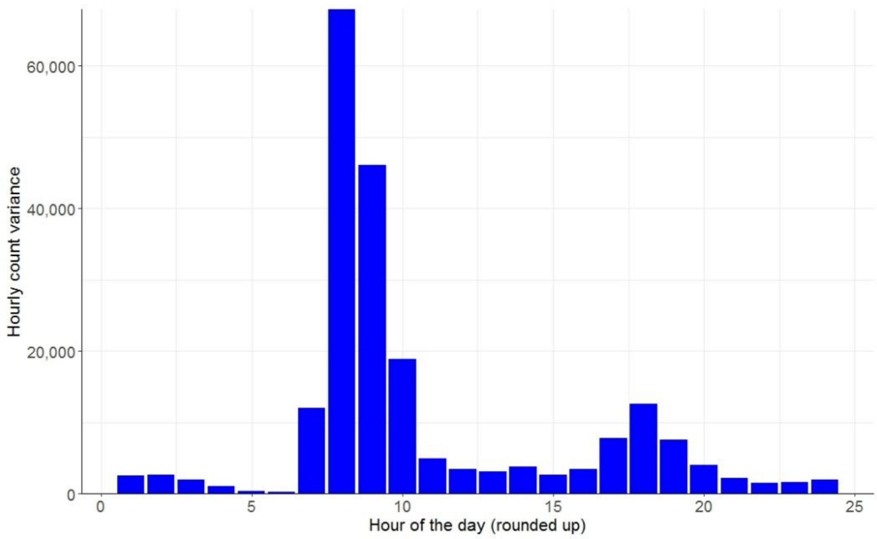

Figure 5-15: Variance of hourly loop detector data

5.7 Discussion

The proposed approach has very broad applications as it could be used for researching any mode of transportation. Different types of network-focused data would be required for different kinds of analyses, for example research into rail travel may involve station usage data. The approach is also not restricted to only using data which explicitly describes movements, but could also include factors which may affect travel patterns. For example, if a large shopping centre or employer can be isolated on the network, data regarding shop opening hours or shift patterns could be examined in conjunction with individual traveller data to analyse travel behaviour. The flexibility of the proposed approach is also a disadvantage, however. A high level framework was developed so that many different sorts of data and different types of analyses could be included, but as it does not include specific techniques, appropriate methods may not exist to analyse the type of data selected. In the case study application in Sections 5.4, 5.5 and 5.6, for example, the analyses of day-to-day variability in flows and traveller behaviour which take into account within-day dynamics would not have been possible without the methods developed in Crawford et al. (2017) and Crawford et al. (Under review-b). The opportunities provided by the proposed approach would be greatly limited if only *established* methodologies were used to analyse *traditional* sources of data.

The flexibility of the approach also means that it will not (and is not intended to) provide a specific output which could be described prior to commencing the research. Whilst the methods used in the analyses in Stage 1 should be justifiable based on the data being analysed and the purpose of the study, the choice of hypotheses to pursue in Stage 3 (‘following the thread’) are more subjective. This framework should, therefore, be considered as an exploratory tool which is likely to result in different findings, even when different analysts use the same data. This does not necessarily mean that the findings will be contradictory, as they may just relate to different kinds of hypotheses.

The analysis in Section 5.6.1.2 demonstrated that the usefulness of the approach could be limited by the amount of data available. Despite the relatively large traveller-focused dataset, there was insufficient data to examine very specific hypotheses, for example considering people who might have more than one evening travel pattern (in terms of the time of travel). This is unsurprising, given that the Bluetooth data is passively collected for the purpose of travel time estimation and therefore is not guaranteed to include a large sample of travellers with every possible repeated trip characteristic. Even if a hypothesis cannot be answered using the data available, however, the process of *identifying* the hypothesis may be helpful. If the hypothesis is deemed to be sufficiently important, then more focused data collection could be undertaken, for example by adding additional questions into user focus group sessions.

As well as not having *sufficient* data, there is also a risk that in Stage 3 the *right kind* of data may not be available. Although the framework shown in Figure 5-2 and applied to the case study suggests that the hypotheses raised in Stage 3 need to be tested using the types of data analysed in Stages 1 and 2, this need not be the case. Other network- and traveller-focused datasets could be used, provided that Stage 2 is repeated in order to examine their comparability with the original types of data analysed. For example, in the case study application, the analyst may have hypothesised that the local maxima at mid-afternoon on the weekday profiles relate to trips escorting children from school. The Bluetooth data does not include trip purpose data nor vehicle occupancy data. If travel diary data was available for a sample of travellers who use this link, then it may be possible to investigate the hypothesis using that data. Whether using the same or different data to Stages 1 and 2, it is possible that the findings from Stage 3 may themselves generate new hypotheses to test.

The application included in this paper only considers one location, although the approach could easily be extended to include multiple locations. This extension could be done in a number of ways. Firstly, the same type of approach could be applied independently to sites close together and on similar routes in order to confirm whether consistent findings are identified. Secondly, Stage 1 of the single site approach could be applied independently to many sites across a city. The outputs of those analyses could then be compared and sites could be clustered based on the network and traveller characteristics observed. For example, clustering could be undertaken based on the day of the week coefficients estimated from the loop detector data. Hypotheses could then be generated based on those clusters of locations, for example:

- Are the same travellers passing sites within the same cluster?
- Are the systematic differences between clusters of sites caused by people travelling to systematically different parts of the city?

It may also be possible to pool data within clusters to test more detailed hypotheses, although suitable adjustments or aggregations of the time of day would need to be considered. Thirdly, the initial analyses in Stages 1 and 2 could be undertaken at an area, as opposed to a single site, level. These areas could be neighbourhoods or cities, or they could be popular transport corridors or bus routes. An example relating to the public transport system for a particular town could include a network-focused analysis examining the total number of passengers by the time of day, day of the week or season, and a traveller-focused analysis examining repeated use of the system including spatial and time of day variability. The hypotheses generated in Stage 3 could relate to the whole system, or may seek to identify more localised patterns.

5.8 Conclusions

In this paper a broad framework has been proposed for undertaking quantitative research considering both the network and the traveller perspective. The inductive nature of the approach was highlighted by an application to one site in Greater Manchester using two years of loop detector and Bluetooth sensor data. As well as providing more information about the data sources used, the following hypotheses were also examined:

- Are each traveller's weekday trips evenly distributed from Monday to Friday?
- Do people travel at systematically different times of day on different days of the week?
- Are repeated trip user classes for travellers related to vehicle types?
- Is random variability in the time of day that individuals travel related to variability in flows on the network?

By using an approach which considers multiple perspectives, therefore, we can discover more about complex daily, weekly or monthly behaviours and their impacts on the road network through interrogating data in greater detail. Such approaches are becoming more feasible due to the availability of new, and often large, data sources. The availability of such datasets also *necessitates* such analytical approaches, since no single source of data provides all of the variables required to undertake a comprehensive analysis of travel behaviour at a large scale. In order to make the best use of the data available, therefore, we need to consider how we can utilise multiple data sources more effectively and not just repeat the sort of analyses undertaken when far less data was available.

5.9 References

Arana, P., Cabezudo, S. & Peñalba, M. (2014) Influence of weather conditions on transit ridership: A statistical study using data from Smartcards. *Transportation Research Part A: Policy and Practice,* 59**,** 1-12.

Barceló, J., Montero, L., Marqués, L. & Carmona, C. (2010) Travel Time Forecasting and Dynamic Origin-Destination Estimation for Freeways Based on Bluetooth Traffic Monitoring. *Transportation Research Record: Journal of the Transportation Research Board,* 2175**,** 19-27.

Bayarma, A., Kitamura, R. & Susilo, Y. (2007) Recurrence of Daily Travel Patterns: Stochastic Process Approach to Multiday Travel Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2021**,** 55-63.

Bhaskar, A. & Chung, E. (2013) Fundamental understanding on the use of Bluetooth scanner as a complementary transport data. *Transportation Research Part C: Emerging Technologies,* 37**,** 42-72.

Bhaskar, A., Tsubota, T., Kieu, L. M. & Chung, E. (2014) Urban traffic state estimation: Fusing point and zone based data. *Transportation Research Part C: Emerging Technologies,* 48**,** 120-142.

Buliung, R. N., Roorda, M. J. & Remmel, T. K. (2008) Exploring spatial variety in patterns of activity-travel behaviour: initial results from the Toronto Travel-Activity Panel Survey (TTAPS). *Transportation,* 35**,** 697-722.

Carpenter, C., Fowler, M. & Adler, T. (2012) Generating Route-Specific Origin-Destination Tables Using Bluetooth Technology. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 96-102.

Chikaraishi, M., Fujiwara, A., Zhang, J. & Axhausen, K. (2009) Exploring Variation Properties of Departure Time Choice Behavior by Using Multilevel Analysis Approach. *Transportation Research Record: Journal of the Transportation Research Board,* 2134**,** 10-20.

Crawford, F., Watling, D. P. & Connors, R. D. (2017) A statistical method for estimating predictable differences between daily traffic flow profiles. *Transportation Research Part B: Methodological,* 95**,** 196-213.

Crawford, F., Watling, D. P. & Connors, R. D. (Under review-a) Assessing the feasibility of using Bluetooth data to examine the repeated travel behaviour of road users. *Transportation*.

Crawford, F., Watling, D. P. & Connors, R. D. (Under review-b) Identifying road user classes based on repeated trip behaviour using Bluetooth data. *Transportation Research Part A: Policy and Practice*.

Creswell, J. W. & Plano Clark, V. L. (2011) *Designing and conducting mixed methods research,* Los Angeles, London, Sage.

Datla, S. & Sharma, S. (2008) Impact of cold and snow on temporal and spatial variations of highway traffic volumes. *Journal of Transport Geography,* 16**,** 358-372.

Davis, D. F., Golicic, S. L. & Boerstler, C. N. (2010) Benefits and challenges of conducting multiple methods research in marketing. *Journal of the Academy of Marketing Science,* 39**,** 467-479.

Elango, V., Guensler, R. & Ogle, J. (2007) Day-to-Day Travel Variability in the Commute Atlanta, Georgia, Study. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 39-49.

Fraley, C. & Raftery, A. E. (2002) Model-Based Clustering, Discriminant Analysis, and Density Estimation. *Journal of the American Statistical Association,* 97**,** 611-631.

Heinen, E. & Chatterjee, K. (2015) The same mode again? An exploration of mode choice variability in Great Britain using the National Travel Survey. *Transportation Research Part A: Policy and Practice,* 78**,** 266-282.

Huff, J. O. & Hanson, S. (1986) Repetition and Variability in Urban Travel. *Geographical Analysis,* 18**,** 97-114.

Järv, O., Ahas, R. & Witlox, F. (2014) Understanding monthly variability in human activity spaces: A twelve-month study using mobile phone call detail records. *Transportation Research Part C: Emerging Technologies,* 38**,** 122-135.

Johnson, R. B., Onwuegbuzie, A. J. & Turner, L. A. (2007) Toward a Definition of Mixed Methods Research. *Journal of Mixed Methods Research,* 1**,** 112-133.

Jones, P. & Clarke, M. (1988) The significance and measurement of variability in travel behaviour. *Transportation,* 15**,** 65-87.

Kaltenbrunner, A., Meza, R., Grivolla, J., Codina, J. & Banchs, R. (2010) Urban cycles and mobility patterns: Exploring and predicting trends in a bicycle-based public transport system. *Pervasive and Mobile Computing,* 6**,** 455-466.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015) Passenger Segmentation Using Smart Card Data. *IEEE Transactions on Intelligent Transportation Systems,* 16**,** 1537-1548.

Kusakabe, T. & Asakura, Y. (2014) Behavioural data mining of transit smart card data: A data fusion approach. *Transportation Research Part C: Emerging Technologies,* 46**,** 179-191.

Ma, X., Wu, Y.-J., Wang, Y., Chen, F. & Liu, J. (2013) Mining smart card data for transit riders' travel patterns. *Transportation Research Part C: Emerging Technologies,* 36**,** 1-12.

Mathew, J. K., Devi, V. L., Bullock, D. M. & Sharma, A. (2016) Investigation of the Use of Bluetooth Sensors for Travel Time Studies under Indian Conditions. *Transportation Research Procedia,* 17**,** 213-222.

Meyer, M. D. (1999) Demand management as an element of transportation policy: using carrots and sticks to infuence travel behavior. *Transportation Research Part A,* 33**,** 575-599.

Moghaddam, S. & Hellinga, B. (2014) Real-Time Prediction of Arterial Roadway Travel Times Using Data Collected by Bluetooth Detectors. *Transportation Research Record: Journal of the Transportation Research Board,* 2442**,** 117-128.

Moran-Ellis, J., Alexander, V. D., Cronin, A., Dickinson, M., Fielding, J., Sleney, J. & Thomas, H. (2016) Triangulation and integration: processes, claims and implications. *Qualitative Research,* 6**,** 45-59.

Morency, C., Trépanier, M. & Agard, B. (2007) Measuring transit use variability with smart-card data. *Transport Policy,* 14**,** 193-203.

Nakayama, S. & Watling, D. (2014) Consistent formulation of network equilibrium with stochastic flows. *Transportation Research Part B: Methodological,* 66**,** 50-69.

O'Cathain, A., Murphy, E. & Nicholl, J. (2010) Three techniques for integrating data in mixed methods studies. *BMJ,* 341.

Quayle, S., Koonce, P., DePencier, D. & Bullock, D. (2010) Arterial Performance Measures with Media Access Control Readers. *Transportation Research Record: Journal of the Transportation Research Board,* 2192**,** 185-193.

Schlich, R. & Axhausen, K. W. (2003) Habitual travel behaviour: Evidence from a six-week travel diary. *Transportation,* 30**,** 13-36.

Shen, Y., Kwan, M.-P. & Chai, Y. (2013) Investigating commuting flexibility with GPS data and 3D geovisualization: a case study of Beijing, China. *Journal of Transport Geography,* 32**,** 1-11.

Singhal, A., Kamga, C. & Yazici, A. (2014) Impact of weather on urban transit ridership. *Transportation Research Part A: Policy and Practice,* 69**,** 379-391.

Spissu, E., Meloni, I. & Sanjust, B. (2011) Behavioral Analysis of Choice of Daily Route with Data from Global Positioning System. *Transportation Research Record: Journal of the Transportation Research Board,* 2230**,** 96-103.

Tao, S., Corcoran, J., Mateo-Babiano, I. & Rohde, D. (2014) Exploring Bus Rapid Transit passenger travel behaviour using big data. *Applied Geography,* 53**,** 90-104.

Transport for London (2013) Olympic Legacy Monitoring: Personal Travel Behaviour during the Games Available: content.tfl.gov.uk/olympic-legacy-personal-travel-report.pdf [Accessed: 29th July 2016].

Wardrop, J. G. (1952) Some theoretical aspects of road traffic research. *Road Engineering Division Meeting.*

Weijermars, W. & van Berkum, E. Analyzing highway flow patterns using cluster analysis. 8th International IEEE Conference on Intelligent Transportation Systems, 2005 Vienna, Austria. 831-836.

Yazici, M., Kamga, C. & Mouskos, K. (2012) Analysis of Travel Time Reliability in New York City Based on Day-of-Week and Time-of-Day Periods. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 83-95.

Zhang, W., Mediria, A. & Rakha, H. Statistical analysis of spatiotemporal link and path flow variability. 2007 IEEE Intelligent Transportation Systems Conference, 2007 Seattle, Washington, USA. IEEE, 1080-1085.

6 Discussion and conclusions

Convenient fictions relating to variability are often used in transportation research. Simplifying assumptions of this kind include assuming homogeneity amongst travellers, vehicles and trips over a fixed period of time, or assuming that variability can be represented by a single standard statistical distribution or by forming multiple user classes. Such assumptions are rarely supported by empirical data and in many cases methods have not been proposed to do so. The motivation for the research undertaken in this thesis, therefore, was to fill this gap by developing methods which could be used to quantify and test for different types of variability in travel behaviour on the road network, which could include systematic or predictable variability from a network perspective, and also the differences in intrapersonal variability between travellers. Whilst these aspects would have been very difficult, if not impossible, to measure in the past, the emergence of new data sources are creating new opportunities to do so, and are simultaneously creating a demand for new methods to be developed in order to do so. The overarching research question for this thesis was, therefore,:

> How can emerging data sources be used to gain insights into variability in travel behaviour on the road network?

In Section 6.1 below, the five research objectives formulated in Chapter 1 in relation to this research question are discussed. This includes a summary of the work undertaken and a critical discussion of the advantages and limitations of the approach taken. In Section 6.2, the contribution to knowledge and practice of the research undertaken is described. In Section 6.3, future work building upon this research is discussed. Section 6.4 concludes the chapter.

6.1 Revisiting the research objectives

Each of the five research objectives identified in Section 1.3 will now be discussed in turn to determine the extent to which they have been achieved by the research undertaken and what limitations exist in the approach taken.

6.1.1 Objective 1

To develop a methodology for identifying statistically significant predictable differences in aggregated travel behaviour observed on the network which takes into account differences in magnitude and timing

In Chapter 2, a methodology was proposed which tested for statistically significant differences in the magnitudes and daily distributions of flows based on predictable day types, such as the day of the week. The daily magnitudes of flows were examined using ANOVA and the standardised daily flow profiles were examined using the functional equivalent, namely fANOVA. Systematic differences between day types are not just observed, they are *tested*. The method identified predictable differences in the magnitude and timing of flows at a case study location based on the day of the week and also the season. By examining the average magnitudes of flows and average daily flow profiles for each statistically significant day type, broad insights into aggregated traveller behaviour were observed. For example, Fridays had the highest daily flows, on average, and the average standardised daily flow profile had less prominent peaks than any other weekday in the case study application.

The methodology utilises induction loop detector data, which is not an emerging data source as it was established in the early 1960s and is now the most commonly used traffic sensor (U.S. Department of Transportation, 2006). Due to decreasing data storage costs and increasing computational power in recent years, data at a higher temporal resolution is being stored for longer periods of time. There is, therefore, a large amount of data available which necessitates the use of techniques more similar to those used for emerging data sources. The success of this methodology relies upon having high temporal resolution data for a relatively long period of time (two years were used in Chapter 2).

Flow data was selected to use for examining aggregated travel behaviour. While a loop detector collects data at a point on the network, variability in flows can be caused by traveller behaviour or capacity variability on the surrounding network. For example, decreased flows could be due to decreased demand for routes passing this point, due to congestion on the link (or further downstream) slowing down vehicles, or due to congestion at an upstream location which prevents vehicles from reaching the observed link. Flows are suitable for analysing predictable differences as only aspects which are *systematically* different between the day types are of interest. The methodology was designed to be used by road management bodies who would have sufficient knowledge of local conditions to be able to separate systematic differences in aggregated travel behaviour from factors beyond travellers' control. Such factors could include predictable differences in signal settings according to the day of the week, or differences in available routes, for example due to bus only lanes on some days of the week.

The average standardised flow profiles for each weekday in the case study application were shown in Figure 2-7. While the fANOVA and the corresponding permutation test found that there were statistically significant differences between all five weekdays, the plot illustrated that the differences were relatively small. It could be argued, therefore, that when using a large amount of data, such tests become too sensitive to differences, although further applications are required to verify whether this is the case. To mitigate the risk of the tests identifying differences which are irrelevant to the practitioner, for example variability during the night time, the analysis could be focused on subsets of the day only, such as the period shortly before, during and after the morning peak only.

The methodology is sufficiently general that it could be used for any flow or usage data collected at a stationary point on a network in continuous time or in very short time intervals. Although Chapter 2 focused on day types which would be known in advance, the methodology could be applied to any dichotomous whole day effects where flow and day type data are available, for example school holiday or heavy flooding indicators. The current methodology only relates to the analysis of data from a single site, but possible extensions to multiple sites were discussed in Section 2.5.

The methodology assumes stable behaviour over the period of analysis but does not account for any long-term trends as this was outside of the scope of the current research. Future research could explore including an independent variable specifying the year of the observation or a more detailed temporal variable in the model.

6.1.2 Objective 2

To identify methods for measuring spatial and temporal intrapersonal variability in travel behaviour on the road network using data from emerging data sources

As previous research has not looked at intrapersonal variability of all road users, there was not an established source of data to develop methods for, to achieve this objective. As the methods should provide insights which apply directly to parts of the road network, rather than larger geographic areas, a data source which is closely related to the road network was required. An emerging data source which collects data on the road network for all vehicle types, at a relatively low cost is Bluetooth data. As Bluetooth data has previously been used to match observations within days on the road network, Chapter 3 explored the suitability of using this type of data for examining travellers' behaviour over multiple days. Chapter 3 concluded that it is *feasible* to use Bluetooth data to examine intrapersonal variability, but as with all types of data, there are possible sources of bias and limitations with the data. For the current application there is a risk that Bluetooth data may have a bias towards travellers making greater numbers of trips, particularly in commercial vehicles, and the sampling rate of trips may differ by Bluetooth device or traveller type.

Chapter 3 highlighted many possible benefits of using Bluetooth data to examine intrapersonal variability, but also emphasised that customised methods are required to make the best use of the data. The customised methods would not only be suitable for Bluetooth data, however, as they could be used on data from other sensors in the "point-to-point" category (Antoniou et al., 2011, p140). Other data sources which involve the collection of unique identifiers from vehicles at fixed points on the road network, which can be matched across sites and over time, include WiFi data, Automatic Number Plate Recognition (ANPR) data and data from electronic tags used for paying tolls. Developing methods to analyse spatial and temporal intrapersonal variability using point-to-point data is, therefore, applicable more widely than just to Bluetooth data.

A method for measuring spatial intrapersonal variability using point-to-point sensor data was presented in Chapter 4. The spatial aspects of trips are compared using Sequence Alignment and then this information is used to produce a distance matrix to cluster trips with similar spatial characteristics. The method examines differences in both origin-destination pairs and routes simultaneously. In Chapter

4, the spatial variability measures for each traveller consisted of the number of spatial clusters their trips were assigned to, and the proportion of the travellers' trips which were assigned to their most frequently used spatial cluster. The core Sequence Alignment process could be used in other ways too, however. For example, measures of how different the spatial clusters used are (in terms of physical distance) could also be incorporated as traveller characteristics.

The Sequence Alignment approach takes into account the order in which Bluetooth sensors are passed and the distance between sensors. It therefore works better for longer routes, for example in the spatial cluster containing trips from the west to the east of Wigan shown in Figure 4-10. In such cases, there are more observations to match (or try to) and sensors are more widely dispersed around the first and last observations. A high concentration of Bluetooth detectors in an urban area may allow greater route disaggregation for longer trips and may result in a higher proportion of short trips being detected, but future work is required to examine the impact of detector placement and concentration on this spatial clustering method. Further work is also required to identify the optimal placement of sensors for measuring intrapersonal variability, as discussed in Section 6.3.4.

A method for measuring temporal intrapersonal variability was also presented in Chapter 4. For each traveller, the method identifies different travel patterns, represented by model-based clusters, and estimates the random variability in the time of travel *within* each of these travel patterns. This method is therefore consistent with one of themes in this research which is the separation of systematic and random elements of variability.

As discussed in Chapters 3 and 4, examining temporal intrapersonal variability using point-to-point sensor data is challenging as observations do not correspond to either departure or arrival times. Departure times are often analysed in traveller behaviour research as they relate directly to the traveller's choice and they are typically what is recorded in travel diaries. A more network-focused approach would be to concentrate on the variability in the time of day a traveller arrives at a specific location which is of interest to the analyst, for example the arrival time at a tunnel or other pinch-point, or the time a particular cordon is passed. This is, arguably, more relevant for a network manager, as they are unconcerned whether a commuter sometimes leaves home earlier to drop a child at school in an outlying area, for example, it is the time at which that traveller enters the congested area around the city centre that is of interest. The method for calculating temporal intrapersonal variability proposed could, therefore, be applied specifically to different pinch-points on the network. In Chapter 4, a broader measure of temporal variability was required and therefore time of day variability measures were calculated based on the Bluetooth sensor which each traveller was detected at most often. In practice, the most common sensor for each traveller could be selected from a subset of sensors identified by the road manager as being of particular interest (for example pinch-points or congested links). If required, the relationship between departure times and the most common sensor location could be examined by collecting GPS tracking data for a small sample of travellers, as in Muthyalagari et al. (2001) or Elango et al. (2007), and matching this to a road network including Bluetooth detectors.

6.1.3 Objective 3

To develop a methodology for comparing and/or classifying road users based on the intrapersonal variability in their travel behaviour

In Chapter 4, a methodology was presented which can identify road user classes based on their repeated trip behaviour. The methodology utilised the methods discussed in Section 6.1.2 and is flexible in terms of the measures of travel behaviour which could be included. As discussed in Chapter 4, different clustering algorithms could be used to identify the user classes, but k-means was deemed to be the most useful for the current application.

The proposed approach is data driven and clustering is used to account for possible non-linear relationships between the different measures of variability. Using a data driven approach prevents a priori assumptions from shaping the resulting user classes, in contrast to some traditional segmentation approaches.

The proposed methodology seeks to describe each traveller's trip frequency, spatial and temporal variability in as few variables as possible so that the user classes obtained can be interpreted fairly easily. The approach used is very flexible, however, and therefore additional variables could easily be included in the clustering process. The additional variables could relate to seasonal trends and changes over longer periods of time in trip frequency, spatial or time of day variability.

6.1.4 Objective 4

To develop a framework for using network- and traveller-focused analyses together to gain additional insights into variability in travel behaviour

In Chapter 5, a high level methodology was proposed as a framework for undertaking network- and traveller-focused analyses concurrently and interactively in order to gain additional insights. The process consists of three stages. Firstly, network- and traveller-focused analyses are undertaken independently using suitable data and methods for each. Different data is likely to be required for each type of analysis, but they should relate to the same geographic area. Secondly, the data used and the results obtained from each type of analysis are compared. Thirdly, queries arising from the network-focused analysis are explored using traveller-focused analyses and vice versa.

The framework describes how analyses from different perspectives can be brought together, but also provides a basis for undertaking exploratory research into variability using multiple data sources. The proposed approach is very general, and therefore could be applied in a wide range of situations involving different modes of transport and/or types of data. This generality is also a limitation, however, because for some types of data, methods may not have been developed to analyse variability and therefore additional work will be required. The ability to test different hypotheses will also be constrained by the type and amount of data available, but the identification of an interesting hypothesis could act as a catalyst to collect additional relevant data.

Chapter 5 demonstrates how such multiple method research can be undertaken, not just by proposing a methodology, but also by applying it to a small case study consisting of one link in Stockport, Greater Manchester. This application utilises the methods proposed in Chapters 2 and 4. Although the application only relates to one link, it provides an insight into the types of hypotheses which could be tested using this approach, and the heterogeneity in traveller behaviour which is evident when testing hypotheses using real world data.

6.1.5 Objective 5

To apply the methods to real world data in order to demonstrate the insights which can be achieved

The methods described in Sections 6.1.1 to 6.1.4, above, were applied to real world data in Chapters 2, 4 and 5. Chapters 2 and 5 used data from Stockport and Chapters 3 and 4 used data from Wigan. Both of these are towns in Greater Manchester and the data was obtained from the same source (Transport for Greater Manchester). The case study areas are expected to be quite different, as the link in Stockport is on a popular arterial route into Manchester city centre whereas Wigan is a satellite town of Manchester, but closer to the town of Bolton.

Applying the method using ANOVA and Functional Data Analysis on loop detector data in Chapter 2 identified differences in the magnitude of flows on Saturdays, Sundays and three categories of weekdays. Differences in the distribution of flows throughout the day (i.e. in the standardised daily profiles) were identified for all seven days of the week. By separating out the magnitude and timing of flows, the application to real world data was able to show that while the morning peak on Mondays looks similar to other days of the week, it is actually caused by lower magnitudes of flows on Mondays which are more highly concentrated in the peak periods than on other days of the week. Also, on the studied link, the flows on Fridays are less concentrated in the peak periods than on other days, but this is less visible when looking at the daily flow profiles as Fridays have higher flows, on average, and therefore the peaks are only slightly lower than on other weekdays. In Chapter 2, estimated 95% confidence intervals were used to identify the times of day at which two day of the week profiles differ. Seasonal differences were also examined for each day of the week profile.

The observation in Section 6.1.1, that the fANOVA may be overly sensitive to differences in the standardised profiles, is somewhat premature, as it is not clear whether all days of the week would be identified as statistically significantly different at a large proportion of sites. Undertaking real world applications of the method at more sites, particularly ones with different characteristics, is therefore required in order to gain a better understanding of how this method works in practice.

The application of the Sequence Alignment process to the real-word data in Chapter 4 provided insights into how the technique works in practice. The freely available R package used (Gabadinho et al., 2011) could not be used on all unique trip sequences in the year as there were too many of them. It is possible that distance matrices could be computed for such large sets of sequences using other more specialised software, but such software may not be available to all practitioners.

The case study in Chapter 4 analysed Bluetooth data and identified three user classes based on repeated travel behaviour. Although the user classes were defined based on measures of trip frequency, spatial and time of day variability, these three user classes could be defined based on the trip frequency alone. The frequent traveller user class contributed 59% of the trips over the one year period, and was separated into four subclasses which *did* depend on the measures of spatial and temporal variability.

In Chapter 5, an application to real world data was presented for the framework for undertaking network- and traveller-focused analyses. The findings from Stage 1 of the case study application provided similar insights to those obtained in Chapters 2 and 4. The comparison of Bluetooth sampling rates by the day of the week and time of day in Stage 2 raised questions consistent with those raised in Chapter 3 regarding possible biases in Bluetooth data. By using the real world application, Stage 2 also highlighted the difficulty in making direct comparisons between analyses undertaken at different levels and this supports the need for a third stage in the process.

In Stage 3, the following hypotheses were also examined:

- Are each traveller's weekday trips evenly distributed from Monday to Friday?
- Do people travel at systematically different times of day on different days of the week?
- Are repeated trip user classes for travellers related to vehicle types?
- Is random variability in the time of day that individuals travel related to variability in flows on the network?

Chi squared tests were successfully used to test the first two hypotheses (generated based on the network-focused analysis). Despite the relatively large number of Bluetooth observations (1.1 million) in the case study, there was insufficient data to explore more specific hypotheses in a robust manner, for example whether people travel home from work earlier on Friday afternoons than on other days of the week. The method is sufficiently flexible, however, to allow the analyst to use any suitable data which is available to them. For example, the comparison of user classes (from the traveller-focused analysis) and vehicle classes (from aggregated network data) showed that it is not necessary to use the same data in Stage 3 of the process that was used in Stage 1, as the vehicle class data had not be examined previously. The vehicle class data came from the same *source* as the flow data, namely the loop detector, but this need not be the case either. For example, in the case study application, the network-focused analysis could have resulted in hypotheses relating to trip purpose which could not be tested using the traveller-focused data used in this application, namely the Bluetooth data, and so a different type of data, for example travel diary data, would be required.

The methodologies developed in this thesis have, therefore, all been applied to real world data. By doing so, the benefits which could be achieved by applying each methodology was demonstrated and practical limitations were identified. The use of case studies with relatively large amounts of data provides additional confidence that the methodologies could be used by practitioners or other researchers to gain a better understanding of variability in travel behaviour on the road network using either network- or traveller-focused analyses or both. All of the case studies in this thesis have been undertaken using the free, open-source software R, thus providing greater opportunities for the methodologies to be applied by practitioners and other researchers as no specialist software is required.

6.2 Contribution to knowledge

In Section 6.1, the research undertaken was summarised and the significant progress made towards satisfying the research objectives was discussed. In this section the contribution which this research makes to the fields of empirical research into variability in travel behaviour and modelling of heterogeneity will be discussed.

Applied Functional Linear Models to traffic flow data for the first time

The proposed methodology applied Functional Linear Models to traffic flow data for the first time, to the author's knowledge, although other research has used Functional Principal Component Analysis to undertake *un*supervised learning on transport data, for example Chiou et al. (2014) and Guardiola et al. (2014). By using known variables within Functional Linear Models, the methodology can be applied much more broadly than the example in Chapter 2 relating to the day of the week and season, provided suitable data is available. For example, analyses of the effect of weather on traffic volumes, such as in Datla and Sharma (2008), could be undertaken so that more subtle differences in the timing of flows and phenomena such as peak spreading can be observed. The method is also applicable to different types of data, for example it could be used to analyse cycle counter data, or demand for bike share stations or public transport stops/stations.

Developed a method which *tests* for predictable variability in both the magnitude and timing of daily flow profiles

The methodology proposed in Chapter 2 can be used to test for statistically significant differences between days of the week, but unlike methods previously used (for example Rakha and Van Aerde (1995), Stathopoulos and Karlaftis (2001) and Liu and Sharma (2006)), this approach takes into account the timing of flows during the day. The methodology can, therefore, be used to identify systematically different flow profiles, which provides opportunities for policies and practice to be developed to take account of the differences in order to make transport systems more efficient. It also has implications for modelling as it may be more meaningful to model each statistically significant day of the week separately rather than assume one ‘typical’ weekday model can be applied to all days (even if that model includes variable demand). Also, the method could be used within the Scenario Manager proposed by Kim et al. (2013), to identify different demand-side scenarios and the associated traffic simulation inputs to be able to estimate a travel time distribution through aggregating the different scenario outputs. The scenario-based approach for modelling travel time variability proposed by Kim et al. (2013) also includes supply-side scenarios, for example traffic crashes and work zones, and these would need to be identified separately.

Demonstrated that Bluetooth data can be used for a new purpose, namely analysing repeated trip behaviour

Chapter 3 identified new uses for Bluetooth data, which is collected for travel time estimation in many cities around the world, for example Brisbane, Australia (Tsubota et al., 2011). Chapter 3 also highlighted additional research which would be required in order to make better use of Bluetooth data for repeated trip analyses, for example research into carrying, switching on and switching to discoverable mode for a wide range of Bluetooth devices. This would build on small studies undertaken by Jones and Chin (2015) and Phua et al. (2015). By undertaking the additional research recommended in Chapter 3, Bluetooth may have the potential to validate or even augment travel diary data in the future. Chapter 3 also demonstrated how Bluetooth data has the potential to inform route choice parameters, in a similar way to the very small study performed by Spissu et al. (2011) using GPS data.

Applied an established technique from Bioinformatics to Bluetooth data from a road network for the first time

In Chapter 4, Sequence Alignment was proposed for comparing the spatial nature of trips as part of a methodology for clustering trip data collected from point-to-point sensors. Sequence Alignment has been used previously in conjunction with Bluetooth data; Delafontaine et al. (2012) used Sequence Alignment to analyse visitor paths through rooms in a major trade fair. It has also been used in transportation research by Kim and Mahmassani (2015) who used it to identify common travel patterns using vehicle trajectory data. The approach proposed in Chapter 4, however, uses a different alignment method than the one used by Kim and Mahmassani (2015), to take into account the different kind of data being analysed, namely observations from fixed sensors as opposed to in-vehicle trace data. It is, therefore, a new application of Sequence Alignment to Bluetooth data on the road network, and it is the first time, to the author's knowledge, that it has been used in relation to the measurement of intrapersonal variability. This method combines OD and route variability into a single measure, whereas previous techniques proposed for analysing spatial intrapersonal variability using other forms of data have focused on either the origins and destinations (Buliung et al., 2008, Dill and Broach, 2014), or the route taken for a fixed OD pair (Li et al., 2004, Spissu et al., 2011). The methodology presented in Chapter 4 has broader applications as it could be used for other types of point-to-point data available now or in the future, for example Automatic Number Plate Recognition (ANPR) data.

Applied model-based clustering to separate systematic and random intrapersonal variability in timestamps from Bluetooth detectors for the first time

The measures of time of day variability presented in Chapter 4 build on the work of Kieu et al. (2015b). Kieu et al. (2015b) used a density-based clustering algorithm to identify habitual (and non-habitual) departure times for public transport users. In the current research, rather than using a density-based clustering algorithm as used by Kieu et al. (2015b) to identify clusters and outliers, model-based clustering (Fraley and Raftery, 2002) was used so that *all* observations for each traveller would be assigned to clusters. For each traveller, the number of clusters and their means provide information about the systematic variability in travel timing and the cluster variances provide information about the random component of their time of day variability. As demonstrated in the case study application in Chapter 5, these clusters can be examined to assess whether any of the systematic differences correspond to predictable factors such as the day of the week. The method would also be applicable to other types of data, including smart card data for public transport users.

Proposed the first methodology for measuring intrapersonal variability and identifying road user classes based on repeated trip behaviour of *all* motorised vehicle users

Previous research has applied clustering techniques to identify user types for public transport users based on their repeated travel behaviour (Kieu et al., 2015b, Goulet Langlois et al., 2016), but Chapter 4 is the first application, to the author's knowledge, to all *road* users (including personal and commercial trips). The methodology utilises the methods proposed for estimating spatial and temporal intrapersonal variability as well as a measure related to trip frequency. The proposed methodology could be used by practitioners to gain a better understanding of road users in their area, or to identify strata of users to collect additional data from (for example using focus groups or interviews). Additional research into bias and sampling rates for Bluetooth data (as discussed in Chapter 3) is required so that the measures of trip frequency from Bluetooth data can be adjusted to give estimates of total road trips for each traveller.

In Chapter 1, some of the typical ways in which variability could be included in network models were described. One of these was the use of multiple user classes to represent heterogeneity in traveller or vehicle characteristics. Building on the work of Han et al. (2016), multiple user classes could be defined based on travel frequency, but instead of using arbitrary categories, the clustering methodology proposed in Chapter 4 could be used to identify the user classes. Han et al. (2016) focus on the difference in knowledge of traffic conditions as a distinguishing factor between regular and irregular travellers, but as described in Figure 1.1 in Chapter 1, different levels of trip regularity between travellers could also have other implications. Other characteristics which could vary between the user classes include the degree of flexibility in trips (for example the width of the departure time interval), the number of routes in the choice set and the weighting given to habitual travel behaviour. In all cases, the parameters associated with these characteristics could be estimated by examining repeated travel behaviour using Bluetooth data.

Used techniques from mixed methods research to develop a framework for undertaking quantitative multi-perspective research on the road network

The framework presented in Chapter 5 for undertaking network- and traveller-focused analyses concurrently and then interactively, demonstrates the synergies in applying the two methods proposed in Chapters 2 and 4 to the same location. The framework is very broad and it could be applied to other modes, for example research looking at public transport users, cyclists or pedestrians, and could involve a wide range of traveller- or network-focused analysis techniques. Given the constantly growing list of possible data sources for transportation research, this general principle of using a mixed method inspired approach to undertake both network- and traveller-focused research together provides an alternative to either trying to find a single ideal data source or directly fusing together different types of data, as in Bachmann et al. (2013) for estimating speeds.

Applied the methods described above to sufficiently large case studies to provide insights into travel behaviour at the sites in Greater Manchester

In Chapters 2, 3, 4 and 5, real world data was used to demonstrate the methods and to inform discussions. Chapter 4 included a case study covering an entire town in northwest England for a one year period, analysing almost 7.5 million trips. Such a large application may be of interest in its own right for practitioners and other researchers focusing on Greater Manchester, or Wigan more specifically.

6.3 Future research

There a number of obvious next steps to take, building on this research. Firstly, as all of the real world applications in this thesis related to towns in Greater Manchester, it would be informative to undertake additional applications of all of the methodologies in more densely populated urban areas, larger rural areas and areas with significantly different vehicle mixes in order to determine how well they work under different circumstances. The user class methodology presented in Chapter 4 could also be applied to different towns so that the resulting user classes can be compared to those obtained in Wigan. It would be particularly interesting to see whether the higher level user classes are always defined by trip frequency alone or whether that was a particular characteristic of the case study location. The flexibility of the methodology should also be explored, for example for towns dominated by tourism additional clustering variables may be required such as the percentage of trips within the tourist seasons.

Four more substantial extensions to the current research will now be discussed.

6.3.1 Applying Functional Linear Models to multiple loop detector sites

The methodology proposed for testing for predictable variability in traffic flow profiles in Chapter 2 could be extended to data from multiple sites in a number of ways.

One possible option is to use one Functional Linear Model for data from all of the sites combined. The model would include overall global effects, for example a Monday effect if days of the week are being assessed, but would also include the remaining site specific effects, such as the Monday effect specific to Site A. This may work for small towns such as Wigan, where there is likely to be a single 'global' effect. For larger or more varied areas, however, more of the variability is likely to be assigned to site specific effects. The greater the proportion of site specific effects relative to global effects, the greater the need to examine the site specific effects. This process would be similar to examining the functional coefficients obtained from applying a separate model to each site and therefore somewhat negates the point of undertaking a combined analysis. This approach also has the disadvantage that the only insights gained into unmonitored parts of the network would relate to the global effects.

It may be preferable to identify groups of sites which have similar day type effects. Capparuccini et al. (2008) discuss three ways that have been proposed to identify road groupings in relation to variability in hourly traffic volumes. Clustering is one of the techniques described and although it could be useful in this context, it has the disadvantage that it may not be possible to *explain* the clusters and, therefore, unmonitored links could not be assigned to clusters. It would be preferable to test whether known variables could be used to form the groups. Capparuccini et al. (2008) discuss grouping based on geographical or functional characteristics, including seasonal patterns and physical properties. Land use variables, including hotel populations or retail employment ratios, may also be helpful in classifying different count sites (Li et al., 2006a).

Further research is required in order to determine in which circumstances there may be sufficient site specific effects to warrant the grouping of sites, and what variables should be used to do so.

6.3.2 Variability influencing behaviour

The current research has focused on traveller behaviour and how this results in day-to-day variability from a network or traveller perspective. This may not be a one-directional relationship, however, as systematic variability at the network level (caused by systematic variability in travel behaviour) also affects travellers' decisions. For example, if a traveller is aware that traffic volumes in a certain town are higher on Fridays then s/he might choose to make their weekly shopping trip on a different day of the week. For this to occur, firstly there needs to be systematic differences at the network level, secondly, at least some travellers need to be aware of these differences, and thirdly, at least some of those travellers need to take this information into account when making their travel decisions. The current research provides a method for testing for systematic differences at the network level, but further research is required to investigate whether the second and third conditions hold. Such research cannot be undertaken using passively collected data.

Other research which explicitly collects data on the decision making processes of individuals related to repeated travel behaviour could also be used in conjunction with the current research. For example, microsimulation software could be used to model different types of behaviour identified during the additional data collection, and then simulated Bluetooth and loop detector data from the model runs could be examined to see whether patterns observed in real world data are visible.

6.3.3 Day-to-day dynamical models

The current research is also relevant for the development of day-to-day dynamical models of the type included in Cascetta and Cantarella (1991), which include traveller adaptation. Although a “day” in day-to-day dynamical models could refer to a day, week or an even longer time period between which travel choices may be reconsidered (Watling and Cantarella, 2013a), the literature tends to focus on days specifically, for example in Iryo (2016) and Djavadian and Chow (2017). Current models of days do not take into account cyclical weekly patterns, however. It is natural to assume that an abnormal travel time on a Friday will have far less impact on the following Monday's travel choices than an abnormal Tuesday would have on the following Wednesday. The current research suggests that systematic differences between days of the week may go further than weekday/weekend day differences. Depending on traveller awareness of these differences, as discussed in Section 6.3.2, it is possible that such models should include day of the week specific travel time expectations.

As alluded to in Chapters 4 and 5, one of the explanations behind systematically different flow profiles according to the day of the week (where they occur) could be that different people are travelling on different days of the week. It may be necessary, therefore, to build a day-to-day dynamical model which includes multiple user classes, each of which includes people travelling in different subsets of the week (or with different probabilities of making a trip on each day of the week). Further research is required in order to understand the learning mechanisms involved in such a system and what implications this has for the stability of the outputs.

6.3.4 Choice of placement of Bluetooth sensors

Methods have been proposed for calculating the optimal number and positioning of Bluetooth detectors for travel time estimation (Asudegi and Haghani, 2013, Park and Haghani, 2015). If Bluetooth data is to be used for alternative purposes, as proposed in this thesis, then a different layout may be required. Some of the criteria will be the same as for travel time estimation (see Asudegi and Haghani (2013, p36)), for example ensuring a high coverage of total volumes and of different OD pairs. Rather than assessing links based on the variability in travel times, however, the aim will be to collect data at locations which provide a representative sample of repeated travel behaviour. This will involve identifying sites on regular commuting routes as well as near to locations which primarily attract less regular, discretionary trips. In order to identify suitable locations, existing Bluetooth data will need to be examined alongside OD matrices, the network topology, land use data and residential and employment counts. By doing so, repeated trip characteristics for all links on the network (monitored and unmonitored) can be estimated so that the optimal location for Bluetooth detectors can be determined.

6.4 Conclusions

Methods have, therefore, been proposed in this thesis which could make better use of emerging data sources to provide insights into traveller behaviour and perhaps challenge some of the 'convenient fictions' which are used liberally and often without justification in transport modelling. Additional research is required in order to make the methods more widely applicable. For example, little is known about when or if people disable Bluetooth on the full range of products now available. The current research has begun to explore how emerging data sources can be used to examine variability in travel behaviour, but this is a vast and constantly changing area which is increasingly important given growing populations and new vehicle technologies. It has also highlighted frameworks within which multiple data sources can be analysed to gain insights into variability from multiple perspectives. Despite the progress made in this thesis, there remains much more work to be done to understand how we can gain better insights into variability in travel behaviour using emerging data sources and how this can inform models in the future.

6.5 References

Abbott, A. (1995) Sequence Analysis: New Methods for Old Ideas. *Annual Review of Sociology,* 21**,** 93-113.

Abbott, A. & Tsay, A. (2000) Sequence Analysis and Optimal Matching Methods in Sociology: Review and Prospect. *Sociological Methods & Research,* 29**,** 3-33.

Agrawal, S., Zheng, H., Peeta, S. & Kumar, A. (2016) Routing aspects of electric vehicle drivers and their effects on network performance. *Transportation Research Part D: Transport and Environment,* 46**,** 246-266.

Aliari, Y. & Haghani, A. (2012) Bluetooth Sensor Data and Ground Truth Testing of Reported Travel Times. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 167-172.

Antoniou, C., Balakrishna, R. & Koutsopoulos, H. N. (2011) A Synthesis of emerging data collection technologies and their impact on traffic management applications. *European Transport Research Review,* 3**,** 139-148.

Araghi, B. N., Christensen, L. T., Krishnan, R. & Lahrmann, H. (2012) Application of Bluetooth Technology for Mode-Specific Travel Time Estimation on Arterial Roads: Potentials and Challenges. *Annual Transport Conference at Aalborg University.*

Araghi, B. N., Olesen, J. H., Krishnan, R., Christensen, L. T. & Lahrmann, H. (2014) Reliability of Bluetooth Technology for Travel Time Estimation. *Journal of Intelligent Transportation Systems,* 19**,** 240-255.

Arana, P., Cabezudo, S. & Peñalba, M. (2014) Influence of weather conditions on transit ridership: A statistical study using data from Smartcards. *Transportation Research Part A: Policy and Practice,* 59**,** 1-12.

Arentze, T. & Timmermans, H. (2005) Modelling Learning and Adaptation in Transportation Contexts. *Transportmetrica,* 1**,** 13-22.

Asudegi, M. & Haghani, A. (2013) Optimal Number and Location of Node-Based Sensors for Collection of Travel Time Data in Networks. *Transportation Research Record: Journal of the Transportation Research Board,* 2338**,** 35-43.

Axhausen, K. W., Zimmermann, A., Schönfelder, S., Rindsfüser, G. & Haupt, T. (2002) Observing the rhythms of daily life: A six-week travel diary. *Transportation,* 29**,** 95-124.

Bachmann, C., Abdulhai, B., Roorda, M. J. & Moshiri, B. (2013) A comparative assessment of multi-sensor data fusion techniques for freeway traffic speed estimation using microsimulation modeling. *Transportation Research Part C: Emerging Technologies,* 26**,** 33-48.

Bagloee, S. A., Sarvi, M., Patriksson, M. & Rajabifard, A. (2017) A Mixed User-Equilibrium and System-Optimal Traffic Flow for Connected Vehicles Stated as a Complementarity Problem. *Computer-Aided Civil and Infrastructure Engineering,* 32**,** 562-580.

Barceló, J., Montero, L., Bullejos, M., Serch, O. & Carmona, C. (2013) A Kalman Filter Approach for Exploiting Bluetooth Traffic Data When Estimating Time-Dependent OD Matrices. *Journal of Intelligent Transportation Systems,* 17**,** 123-141.

Barceló, J., Montero, L., Marqués, L. & Carmona, C. (2010) Travel Time Forecasting and Dynamic Origin-Destination Estimation for Freeways Based on Bluetooth Traffic Monitoring. *Transportation Research Record: Journal of the Transportation Research Board,* 2175**,** 19-27.

Bates, J., Polak, J., Jones, P. & Cook, A. (2001) The valuation of reliability for personal travel. *Transportation Research Part E: Logistics and Transportation Review,* 37**,** 191-229.

Bayarma, A., Kitamura, R. & Susilo, Y. (2007) Recurrence of Daily Travel Patterns: Stochastic Process Approach to Multiday Travel Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2021**,** 55-63.

Beckmann, M., McGuire, C. B. & Winsten, C. B. (1956) Studies in the economics of transportation. New Haven: Yale University Press.

Ben-Elia, E. & Shiftan, Y. (2010) Which road do I take? A learning-based model of route-choice behavior with real-time information. *Transportation Research Part A: Policy and Practice,* 44**,** 249-264.

Bender, E. A. (1978) *An introduction to mathematical modeling,* New York, John Wiley and Sons.

Bhaskar, A. & Chung, E. (2013) Fundamental understanding on the use of Bluetooth scanner as a complementary transport data. *Transportation Research Part C: Emerging Technologies,* 37**,** 42-72.

Bhaskar, A., Tsubota, T., Kieu, L. M. & Chung, E. (2014) Urban traffic state estimation: Fusing point and zone based data. *Transportation Research Part C: Emerging Technologies,* 48**,** 120-142.

Bhat, C. R., Frusti, T., Zhao, H., Schönfelder, S. & Axhausen, K. W. (2004) Intershopping duration: an analysis using multiweek data. *Transportation Research Part B: Methodological,* 38**,** 39-60.

Billot, R., El Faouzi, N.-E. & De Vuyst, F. (2009) Multilevel Assessment of the Impact of Rain on Drivers' Behavior. *Transportation Research Record: Journal of the Transportation Research Board,* 2107**,** 134-142.

Böcker, L., Dijst, M. & Prillwitz, J. (2012) Impact of Everyday Weather on Individual Daily Travel Behaviours in Perspective: A Literature Review. *Transport Reviews,* 33**,** 71-91.

Bogers, E. A. I., Bierlaire, M. & Hoogendoorn, S. P. (2007) Modeling learning in route choice. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 1-8.

Brennan, T. M., Ernst, J. M., Day, C. M., Bullock, D. M., Krogmeier, J. V. & Martchouk, M. (2010) Influence of Vertical Sensor Placement on Data Collection Efficiency from Bluetooth MAC Address Collection Devices. *Journal of Transportation Engineering,* 136**,** 1104-1109.

Bricka, S. & Bhat, C. R. (2006) Comparative Analysis of Global Positioning System-Based and Travel Survey-Based Data. *Transportation Research Record: Journal of the Transportation Research Board,* 1972**,** 9-20.

Brilon, W., Geistefeldt, J. & Regler, M. Reliability of freeway traffic flow: A stochastic concept of capacity. Proceedings of the 16th International Symposium on Transportation and Traffic Theory, 2005 College Park, Maryland. 125-144.

Buliung, R. N., Roorda, M. J. & Remmel, T. K. (2008) Exploring spatial variety in patterns of activity-travel behaviour: initial results from the Toronto Travel-Activity Panel Survey (TTAPS). *Transportation,* 35**,** 697-722.

Bullock, D., Haseman, R., Wasson, J. & Spitler, R. (2010) Automated Measurement of Wait Times at Airport Security. *Transportation Research Record: Journal of the Transportation Research Board,* 2177**,** 60-68.

Calabrese, F., Colonna, M., Lovisolo, P., Parata, D. & Ratti, C. (2011) Real-Time Urban Monitoring Using Cell Phones: A Case Study in Rome. *IEEE Transactions on Intelligent Transportation Systems,* 12**,** 141-151.

Calvert, S. C. & Snelder, M. Influence of rain on motorway road capacity – a data-driven analysis. 16th International IEEE Annual Conference on Intelligent Transportation Systems (ITSC 2013), 2013 The Hague, The Netherlands. 1481-1486.

Calvert, S. C., Taale, H. & Hoogendoorn, S. P. (2016) Quantification of motorway capacity variation: influence of day type specific variation and capacity drop. *Journal of Advanced Transportation,* 50**,** 570-588.

Cantarella, G. E. & Cascetta, E. (1995) Dynamic processes and equilibrium in transportation networks: Towards a unifying theory. *Transportation Science,* 29**,** 305-329.

Capparuccini, D., Faghri, A., Polus, A. & Suarez, R. (2008) Fluctuation and Seasonality of Hourly Traffic and Accuracy of Design Hourly Volume Estimates. *Transportation Research Record: Journal of the Transportation Research Board,* 2049**,** 63-70.

Carpenter, C., Fowler, M. & Adler, T. (2012) Generating Route-Specific Origin-Destination Tables Using Bluetooth Technology. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 96-102.

Carrion, C. & Levinson, D. (2012) Value of travel time reliability: A review of current evidence. *Transportation Research Part A: Policy and Practice,* 46**,** 720-741.

Cascetta, E. (2009) *Transportation systems analysis: models and applications,* New York, Dordrecht, Heidelberg, London, Springer.

Cascetta, E. & Cantarella, G. E. (1991) A Day-to-Day and Within-Day Dynamic Stochastic Assignment Model. *Transportation Research Part A: Policy and Practice,* 25**,** 277-291.

Chen, K. & Müller, H.-G. (2014) Modeling Conditional Distributions for Functional Responses, With Application to Traffic Monitoring via GPS-Enabled Mobile Phones. *Technometrics,* 56**,** 347-358.

Chen, R. & Mahmassani, H. S. (2004) Travel time perception and learning mechanisms in traffic networks. *Transportation Research Record: Journal of the Transportation Research Board,* 1894**,** 209-221.

Cherchi, E. & Cirillo, C. (2014) Understanding variability, habit and the effect of long period activity plan in modal choices: a day to day, week to week analysis on panel data. *Transportation,* 41**,** 1245-1262.

Cherchi, E., Cirillo, C. & Ortúzar, J. d. D. (2017) Modelling correlation patterns in mode choice models estimated on multiday travel data. *Transportation Research Part A: Policy and Practice,* 96**,** 146-153.

Chikaraishi, M., Fujiwara, A., Zhang, J. & Axhausen, K. (2009) Exploring Variation Properties of Departure Time Choice Behavior by Using Multilevel Analysis Approach. *Transportation Research Record: Journal of the Transportation Research Board,* 2134**,** 10-20.

Chiou, J.-M., Zhang, Y.-C., Chen, W.-H. & Chang, C.-W. (2014) A functional data approach to missing value imputation and outlier detection for traffic flow data. *Transportmetrica B: Transport Dynamics,* 2**,** 106-129.

Chitturi, M., Shaw, J., Campbell, J. & Noyce, D. (2014) Validation of Origin-Destination Data from Bluetooth Reidentification and Aerial Observation. *Transportation Research Record: Journal of the Transportation Research Board,* 2430**,** 116-123.

Chu, K. & Chapleau, R. (2010) Augmenting Transit Trip Characterization and Travel Behavior Comprehension. *Transportation Research Record: Journal of the Transportation Research Board,* 2183**,** 29-40.

Clark, S. & Watling, D. (2005) Modelling network travel time reliability under stochastic demand. *Transportation Research Part B: Methodological,* 39**,** 119-140.

Colberg, K., Suh, W., Anderson, J., Zinner, S., Guin, A., Hunter, M. & Guensler, R. (2014) Lane Bias Issues in Work Zone Travel Time Measurement and Reporting. *Transportation Research Record: Journal of the Transportation Research Board,* 2458**,** 78-87.

Corain, L., Melas, V. B., Pepelyshev, A. & Salmaso, L. (2014) New insights on permutation approach for hypothesis testing on functional data. *Advances in Data Analysis and Classification,* 8**,** 339-356.

Corey, J., Lao, Y., Wu, Y.-J. & Wang, Y. (2011) Detection and Correction of Inductive Loop Detector Sensitivity Errors by Using Gaussian Mixture Models. *Transportation Research Record: Journal of the Transportation Research Board,* 2256**,** 120-129.

Crawford, F., Watling, D. P. & Connors, R. D. (2017) A statistical method for estimating predictable differences between daily traffic flow profiles. *Transportation Research Part B: Methodological,* 95**,** 196-213.

Crawford, F., Watling, D. P. & Connors, R. D. (Under review-a) Assessing the feasibility of using Bluetooth data to examine the repeated travel behaviour of road users. *Transportation*.

Crawford, F., Watling, D. P. & Connors, R. D. (Under review-b) Identifying road user classes based on repeated trip behaviour using Bluetooth data. *Transportation Research Part A: Policy and Practice*.

Creswell, J. W. & Plano Clark, V. L. (2011) *Designing and conducting mixed methods research,* Los Angeles, London, Sage.

Dafermos, S. C. (1972) The Traffic Assignment Problem for Multiclass-User Transportation Networks. *Transportation Science,* 6**,** 73-87.

Daganzo, C. F. & Sheffi, Y. (1977) On Stochastic Models of Traffic Assignment. *Transportation Science,* 11**,** 253-274.

Data.gov.uk. (2012) *Contracts Finder Archive: National Travel Survey 2013-2018* [Online]. Available: https://data.gov.uk/data/contracts-finder-archive/contract/565163/ [Accessed: 28 February 2017].

Datla, S. & Sharma, S. (2008) Impact of cold and snow on temporal and spatial variations of highway traffic volumes. *Journal of Transport Geography,* 16**,** 358-372.

Davis, D. F., Golicic, S. L. & Boerstler, C. N. (2010) Benefits and challenges of conducting multiple methods research in marketing. *Journal of the Academy of Marketing Science,* 39**,** 467-479.

de Boor, C. (2001) *A practical guide to splines,* New York, Springer.

Delafontaine, M., Versichele, M., Neutens, T. & Van de Weghe, N. (2012) Analysing spatiotemporal sequences in Bluetooth tracking data. *Applied Geography,* 34**,** 659-668.

Department for Business, Innovation and Skills, (2014) The Fourth Work-Life Balance Employer Survey (2013). Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/398557/bis-14-1027-fourth-work-life-balance-employer-survey-2013.pdf [Accessed: 28th March 2017].

Department for Transport (2016a) National Travel Survey 2015: Notes and Definitions. Available: https://www.gov.uk/government/statistics/national-travel-survey-2015 [Accessed: 7th April 2017].

Department for Transport (2016b) National Travel Survey 2015: Table NTS0601 Average number of trips (trip rates) by age, gender and main mode: England, 2015.

Department for Transport (2016c) National Travel Survey 2015: Table NTS0708 Travel by National Statistics Socio-economic Classification (NS-SEC) and main mode/stage mode: England, 2015.

Department for Transport (2017) Travel time measures for the Strategic Road Network, England: Year ending March 2017. Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/614583/travel-time-measures-srn-england-year-ending-march-2017.pdf [Accessed: 13 July 2017].

Dial, R. B. (1971) A probabilistic multipath traffic assignment model which obviates path enumeration. *Transportation Research,* 5**,** 83-111.

Díaz, J. J. V., González, A. B. R. & Wilby, M. R. (2016) Bluetooth Traffic Monitoring Systems for Travel Time Estimation on Freeways. *IEEE Transactions on Intelligent Transportation Systems,* 17**,** 123-132.

Dill, J. & Broach, J. (2014) Travel to Common Destinations. *Transportation Research Record: Journal of the Transportation Research Board,* 2413**,** 84-91.

Djavadian, S. & Chow, J. Y. J. (2017) An agent-based day-to-day adjustment process for modeling 'Mobility as a Service' with a two-sided flexible transport market. *Transportation Research Part B: Methodological,* 104**,** 36-57.

Du, J., Wong, S. C., Shu, C.-W. & Zhang, M. (2015) Reformulating the Hoogendoorn–Bovy predictive dynamic user-optimal model in continuum space with anisotropic condition. *Transportation Research Part B: Methodological,* 79**,** 189-217.

El Faouzi, N.-E., de Mouzon, O., Billot, R. & Sau, J. (2010) Assessing the changes in operating traffic stream conditions due to weather conditions. *Advances in Transportation Studies,* 21A**,** 33-48.

Elango, V., Guensler, R. & Ogle, J. (2007) Day-to-Day Travel Variability in the Commute Atlanta, Georgia, Study. *Transportation Research Record: Journal of the Transportation Research Board,* 2014**,** 39-49.

Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. *In:* SIMOUDIS, E., HAN, J. & FAYYAD, U., eds., Second International Conference on Knowledge Discovery and Data Mining (KDD-96), 1996 Portland, Oregon. The AAAI Press, 226-231.

Faraway, J. J. (1997) Regression Analysis for a Functional Response. *Technometrics,* 39**,** 254-261.

Federal Highway Administration (FHWA) (2014) The Urban Congestion Report (UCR): Documentation and Definitions. Available: https://ops.fhwa.dot.gov/perf_measurement/ucr/documentation.htm [Accessed: 13 July 2017].

Ferraty, F. & Vieu, P. (2003) Curves discrimination: a nonparametric functional approach. *Computational Statistics & Data Analysis,* 44**,** 161-173.

Fraley, C. & Raftery, A. E. (2002) Model-Based Clustering, Discriminant Analysis, and Density Estimation. *Journal of the American Statistical Association,* 97**,** 611-631.

Gabadinho, A., Ritschard, G., Müller, N. S. & Studer, M. (2011) Analyzing and Visualizing State Sequences in R with TraMineR. *Journal of Statistical Software,* 40**,** 1-37.

Gao, H. O. & Niemeier, D. A. (2007) The impact of rush hour traffic and mix on the ozone weekend effect in southern California. *Transportation Research Part D: Transport and Environment,* 12**,** 83-98.

Gao, H. O. & Niemeier, D. A. (2008) Using functional data analysis of diurnal ozone and NOx cycles to inform transportation emissions control. *Transportation Research Part D: Transport and Environment,* 13**,** 221-238.

Gerpott, T. J., Thomas, S. & Weichert, M. (2013) Characteristics and mobile Internet use intensity of consumers with different types of advanced handsets: An exploratory empirical study of iPhone, Android and other web-enabled mobile users in Germany. *Telecommunications Policy,* 37**,** 357-371.

Good, P. (1994) *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses,* New York, Springer.

Górecki, T. & Smaga, Ł. (2015) A comparison of tests for the one-way ANOVA problem for functional data. *Computational Statistics,* 30**,** 987-1010.

Goulet Langlois, G., Koutsopoulos, H. N. & Zhao, J. (2016) Inferring patterns in the multi-week activity sequences of public transport users. *Transportation Research Part C: Emerging Technologies,* 64**,** 1-16.

Green, P. J. & Silverman, B. W. (1994) *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach,* London, Chapman and Hill.

Guardiola, I. G., Leon, T. & Mallor, F. (2014) A functional approach to monitor and recognize patterns of daily traffic profiles. *Transportation Research Part B: Methodological,* 65**,** 119-136.

Gundlegård, D., Rydergren, C., Breyer, N. & Rajna, B. (2016) Travel demand estimation and network assignment based on cellular network data. *Computer Communications,* 95**,** 29-42.

Guo, F., Li, Q. & Rakha, H. (2012) Multistate Travel Time Reliability Models with Skewed Component Distributions. *Transportation Research Record: Journal of the Transportation Research Board,* 2315**,** 47-53.

Guo, R.-Y., Yang, H., Huang, H.-J. & Tan, Z. (2015) Link-based day-to-day network traffic dynamics and equilibria. *Transportation Research Part B: Methodological,* 71**,** 248-260.

Habib, K. M. N., Day, N. & Miller, E. J. (2009) An investigation of commuting trip timing and mode choice in the Greater Toronto Area: Application of a joint discrete-continuous model. *Transportation Research Part A: Policy and Practice,* 43**,** 639-653.

Habib, K. M. N., Miller, E. & Axhausen, K. (2008) Weekly Rhythm in Joint Time Expenditure for All At-Home and Out-of-Home Activities: Application of Kuhn-Tucker Demand System Model Using Multiweek Travel Diary Data. *Transportation Research Record: Journal of the Transportation Research Board,* 2054**,** 64-73.

Habib, K. M. N. & Miller, E. J. (2008) Modelling daily activity program generation considering within-day and day-to-day dynamics in activity-travel behaviour. *Transportation,* 35**,** 467-484.

Haghani, A., Hamedi, M., Sadabadi, K., Young, S. & Tarnoff, P. (2010) Data Collection of Freeway Travel Time Ground Truth with Bluetooth Sensors. *Transportation Research Record: Journal of the Transportation Research Board,* 2160**,** 60-68.

Hainen, A. M., Wasson, J. S., Hubbard, S. M. L., Remias, S. M., Farnsworth, G. D. & Bullock, D. M. (2011) Estimating Route Choice and Travel Time Reliability with Field Observations of Bluetooth Probe Vehicles. *Transportation Research Record: Journal of the Transportation Research Board,* 2256**,** 43-50.

Han, K., Friesz, T. L., Szeto, W. Y. & Liu, H. (2015) Elastic demand dynamic network user equilibrium: Formulation, existence and computation. *Transportation Research Part B: Methodological,* 81**,** 183-209.

Han, L., Sun, H., Wang, D. Z. W. & Zhu, C. (2016) A stochastic process traffic assignment model considering stochastic traffic demand. *Transportmetrica B: Transport Dynamics***,** 1-21.

Han, S. (2003) Dynamic traffic modelling and dynamic stochastic user equilibrium assignment for general road networks. *Transportation Research Part B: Methodological,* 37**,** 225-249.

Hanson, S. & Huff, J. O. (1988) Systematic variability in repetitious travel. *Transportation,* 15**,** 111-135.

Haseman, R., Wasson, J. & Bullock, D. (2010) Real-Time Measurement of Travel Time Delay in Work Zones and Evaluation Metrics Using Bluetooth Probe Tracking. *Transportation Research Record: Journal of the Transportation Research Board,* 2169**,** 40-53.

Hazelton, M. L. & Parry, K. (2016) Statistical methods for comparison of day-to-day traffic models. *Transportation Research Part B: Methodological,* 92, Part A**,** 22-34.

Heinen, E. & Chatterjee, K. (2015) The same mode again? An exploration of mode choice variability in Great Britain using the National Travel Survey. *Transportation Research Part A: Policy and Practice,* 78**,** 266-282.

Horowitz, J. L. (1984) The stability of stochastic equilibrium in a two-link transportation network. *Transportation Research Part B: Methodological,* 18**,** 13-28.

Huang, H.-J. & Li, Z.-C. (2007) A multiclass, multicriteria logit-based traffic equilibrium assignment model under ATIS. *European Journal of Operational Research,* 176**,** 1464-1477.

Huff, J. O. & Hanson, S. (1986) Repetition and Variability in Urban Travel. *Geographical Analysis,* 18**,** 97-114.

Iida, Y., Akiyama, T. & Uchida, T. (1992) Experimental Analysis of Dynamic Route Choice Behavior. *Transportation Research Part B: Methodological,* 26**,** 17-32.

INRIX (2016) London Congestion Trends: March 2016.

Iqbal, M. S., Choudhury, C. F., Wang, P. & González, M. C. (2014) Development of origin–destination matrices using mobile phone call data. *Transportation Research Part C: Emerging Technologies,* 40**,** 63-74.

Iryo, T. (2016) Day-to-day dynamical model incorporating an explicit description of individuals' information collection behaviour. *Transportation Research Part B: Methodological,* 92**,** 88-103.

Isaev, A. (2006) *Introduction to Mathematical Methods in Bioinformatics,* Berlin, Heidelberg, Springer.

Järv, O., Ahas, R. & Witlox, F. (2014) Understanding monthly variability in human activity spaces: A twelve-month study using mobile phone call detail records. *Transportation Research Part C: Emerging Technologies,* 38**,** 122-135.

Jiang, X. & Adeli, H. (2005) Dynamic Wavelet Neural Network Model for Traffic Flow Forecasting. *Journal of Transportation Engineering,* 131**,** 771-779.

Johnson, R. B., Onwuegbuzie, A. J. & Turner, L. A. (2007) Toward a Definition of Mixed Methods Research. *Journal of Mixed Methods Research,* 1**,** 112-133.

Jones, B. H. & Chin, A. G. (2015) On the efficacy of smartphone security: A critical analysis of modifications in business students' practices over time. *International Journal of Information Management,* 35**,** 561-571.

Jones, P. & Clarke, M. (1988) The significance and measurement of variability in travel behaviour. *Transportation,* 15**,** 65-87.

Kaltenbrunner, A., Meza, R., Grivolla, J., Codina, J. & Banchs, R. (2010) Urban cycles and mobility patterns: Exploring and predicting trends in a bicycle-based public transport system. *Pervasive and Mobile Computing,* 6**,** 455-466.

Kamga, C. & Yazici, M. A. (2014) Temporal and weather related variation patterns of urban travel time: Considerations and caveats for value of travel time, value of variability, and mode choice studies. *Transportation Research Part C: Emerging Technologies,* 45**,** 4-16.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015a) Empirical modelling of the relationship between bus and car speeds on signalised urban networks. *Transportation Planning and Technology,* 38**,** 465-482.

Kieu, L. M., Bhaskar, A. & Chung, E. (2015b) Passenger Segmentation Using Smart Card Data. *IEEE Transactions on Intelligent Transportation Systems,* 16**,** 1537-1548.

Kim, J., Mahmassani, H., Vovsha, P., Stogios, Y. & Dong, J. (2013) Scenario-Based Approach to Analysis of Travel Time Reliability with Traffic Simulation Models. *Transportation Research Record: Journal of the Transportation Research Board,* 2391**,** 56-68.

Kim, J. & Mahmassani, H. S. (2015) Spatial and Temporal Characterization of Travel Patterns in a Traffic Network Using Vehicle Trajectories. *Transportation Research Procedia,* 9**,** 164-184.

Kitamura, R. & Van Der Hoorn, T. (1987) Regularity and irreversibility of weekly travel behavior. *Transportation,* 14**,** 227-251.

Kitamura, R., Yamamoto, T., Susilo, Y. O. & Axhausen, K. W. (2006) How routine is a routine? An analysis of the day-to-day variability in prism vertex location. *Transportation Research Part A: Policy and Practice,* 40**,** 259-279.

Kumar, A. & Peeta, S. (2015) A day-to-day dynamical model for the evolution of path flows under disequilibrium of traffic networks with fixed demand. *Transportation Research Part B: Methodological,* 80**,** 235-256.

Kusakabe, T. & Asakura, Y. (2014) Behavioural data mining of transit smart card data: A data fusion approach. *Transportation Research Part C: Emerging Technologies,* 46**,** 179-191.

Kuwahara, M. & Sullivan, E. C. (1987) Estimating Origin-Destination Matrices from Roadside Survey Data. *Transportation Research Part B,* 21B**,** 233-248.

Lam, W. H. K., Shao, H. & Sumalee, A. (2008) Modeling impacts of adverse weather conditions on a road network with uncertainties in demand and supply. *Transportation Research Part B: Methodological,* 42**,** 890-910.

Lee, K. L., Meyer, R. J. & Bradlow, E. T. (2009) Analyzing risk response dynamics on the web: the case of Hurricane Katrina. *Risk Anal,* 29**,** 1779-92.

Li, H., Guensler, R., Ogle, J. & Wang, J. (2004) Using global positioning system data to understand day-to-day dynamics of morning commute behavior. *Transportation Planning and Analysis 2004.* Washington: Transportation Research Board National Research Council.

Li, L., Su, X., Zhang, Y., Lin, Y. & Li, Z. (2015) Trend Modeling for Traffic Time Series Analysis: An Integrated Study. *IEEE Transactions on Intelligent Transportation Systems,* 16**,** 3430-3439.

Li, M.-T., Zhao, F. & Chow, L.-F. (2006a) Assignment of Seasonal Factor Categories to Urban Coverage Count Stations Using a Fuzzy Decision Tree. *Journal of Transportation Engineering,* 132**,** 654-662.

Li, M., Rouphail, N. M., Mahmoudi, M., Liu, J. & Zhou, X. (2017) Multi-scenario optimization approach for assessing the impacts of advanced traffic information under realistic stochastic capacity distributions. *Transportation Research Part C: Emerging Technologies,* 77**,** 113-133.

Li, R., Rose, G. & Sarvi, M. (2006b) Using Automatic Vehicle Identification Data to Gain Insight into Travel Time Variability and Its Causes. *Transportation Research Record: Journal of the Transportation Research Board,* 1945**,** 24-32.

Lian, H. (2013) Shrinkage estimation and selection for multiple functional regression. *Statistica Sinica,* 23**,** 51-74.

Liu, K., Yamamoto, T. & Morikawa, T. (2009) Feasibility of Using Taxi Dispatch System as Probes for Collecting Traffic Information. *Journal of Intelligent Transportation Systems,* 13**,** 16-27.

Liu, R., Van Vliet, D. & Watling, D. (2006) Microsimulation models incorporating both demand and supply dynamics. *Transportation Research Part A: Policy and Practice,* 40**,** 125-150.

Liu, Y., Kang, C., Gao, S., Xiao, Y. & Tian, Y. (2012) Understanding intra-urban trip patterns from taxi trajectory data. *Journal of Geographical Systems,* 14**,** 463-483.

Liu, Y., Li, Y. & Hu, L. Departure Time and Route Choices with Bottleneck Congestion: User Equilibrium under Risk and Ambiguity. 22nd International Symposium on Transportation and Traffic Theory, 2017 Chicago, USA. 571-590.

Liu, Z. & Sharma, S. (2006) Statistical Investigations of Statutory Holiday Effects on Traffic Volumes. *Transportation Research Record: Journal of the Transportation Research Board,* 1945**,** 40-48.

Lo, H. K., Luo, X. W. & Siu, B. W. Y. (2006) Degradable transport network: Travel time budget of travelers with heterogeneous risk aversion. *Transportation Research Part B: Methodological,* 40**,** 792-806.

Lo, H. K. & Tung, Y.-K. (2003) Network with degradable links: capacity analysis and design. *Transportation Research Part B: Methodological,* 37**,** 345-363.

Long, J., Szeto, W. Y., Gao, Z., Huang, H.-J. & Shi, Q. (2016) The nonlinear equation system approach to solving dynamic user optimal simultaneous route and departure time choice problems. *Transportation Research Part B: Methodological,* 83**,** 179-206.

Lou, X.-M., Cheng, L. & Chu, Z.-M. (2016) Modelling travellers' en-route path switching in a day-to-day dynamical system. *Transportmetrica B: Transport Dynamics,* 5**,** 17-41.

Ma, X., Wu, Y.-J., Wang, Y., Chen, F. & Liu, J. (2013) Mining smart card data for transit riders' travel patterns. *Transportation Research Part C: Emerging Technologies,* 36**,** 1-12.

Mahmassani, H. S., Chang, G.-L. & Herman, R. (1986) Individual Decisions and Collective Effects in a Simulated Traffic System. *Transportation Science,* 20**,** 258-271.

Malinovskiy, Y., Saunier, N. & Wang, Y. (2012) Analysis of Pedestrian Travel with Static Bluetooth Sensors. *Transportation Research Record: Journal of the Transportation Research Board,* 2299**,** 137-149.

Martchouk, M., Mannering, F. & Bullock, D. (2011) Analysis of Freeway Travel Time Variability Using Bluetooth Detection. *Journal of Transportation Engineering,* 137**,** 697-704.

Mathew, J. K., Devi, V. L., Bullock, D. M. & Sharma, A. (2016) Investigation of the Use of Bluetooth Sensors for Travel Time Studies under Indian Conditions. *Transportation Research Procedia,* 17**,** 213-222.

Matsui, H. & Konishi, S. (2011) Variable selection for functional regression models via the L_1 regularization. *Computational Statistics & Data Analysis,* 55**,** 3304-3310.

May, A. D. (1990) *Traffic flow fundamentals,* Englewood Cliffs, Prentice Hall.

Mei, Z., Wang, D. & Chen, J. (2012) Investigation with Bluetooth Sensors of Bicycle Travel Time Estimation on a Short Corridor. *International Journal of Distributed Sensor Networks,* 2012**,** 1-7.

Meyer, M. D. (1999) Demand management as an element of transportation policy: using carrots and sticks to infuence travel behavior. *Transportation Research Part A,* 33**,** 575-599.

Millard, K., Cairns, S. & Beaumont, C. (2016) Car rental customer survey 2015. Transport Research Laboratory under a contract with BVRLA. Available: http://www.bvrla.co.uk/research/article/rental-customer-survey-2015 [Accessed: 15 September 2016].

Mingotti, N., Lillo, R. E. & Romo, J. (2013) Lasso Variable Selection in Functional Regression. *Working papers* [Online]. Available: http://orff.uc3m.es/bitstream/handle/10016/16959/ws131413.pdf?sequence=1.

Minnen, J., Glorieux, I. & van Tienoven, T. P. (2015) Transportation habits: Evidence from time diary data. *Transportation Research Part A: Policy and Practice,* 76**,** 25-37.

Miralinaghi, M., Lou, Y., Hsu, Y.-T., Shabanpour, R. & Shafahi, Y. (2016) Multiclass fuzzy user equilibrium with endogenous membership functions and risk-taking behaviors. *Journal of Advanced Transportation,* 50**,** 1716-1734.

Miranda-Moreno, L. & Nosal, T. (2011) Weather or Not to Cycle. *Transportation Research Record: Journal of the Transportation Research Board,* 2247**,** 42-52.

Moghaddam, S. & Hellinga, B. (2013) Quantifying Measurement Error in Arterial Travel Times Measured by Bluetooth Detectors. *Transportation Research Record: Journal of the Transportation Research Board,* 2395**,** 111-122.

Moghaddam, S. & Hellinga, B. (2014) Real-Time Prediction of Arterial Roadway Travel Times Using Data Collected by Bluetooth Detectors. *Transportation Research Record: Journal of the Transportation Research Board,* 2442**,** 117-128.

Moran-Ellis, J., Alexander, V. D., Cronin, A., Dickinson, M., Fielding, J., Sleney, J. & Thomas, H. (2016) Triangulation and integration: processes, claims and implications. *Qualitative Research,* 6**,** 45-59.

Morency, C., Trépanier, M. & Agard, B. (2007) Measuring transit use variability with smart-card data. *Transport Policy,* 14**,** 193-203.

Muñoz Maldonado, Y., Staniswalis, J. G., Irwin, L. N. & Byers, D. (2002) A similarity analysis of curves. *The Canadian Journal of Statistics,* 30**,** 373-381.

Muthyalagari, G. R., Parashar, A. & Pendyala, R. M. Measuring Day-to-Day Variability in Travel Characteristics Using GPS Data. 80th Annual Meeting of the Transportation Research Board, 2001 Washington D.C.: TRB.

Nakayama, S. & Watling, D. (2014) Consistent formulation of network equilibrium with stochastic flows. *Transportation Research Part B: Methodological,* 66**,** 50-69.

Nerini, D. & Ghattas, B. (2007) Classifying densities using functional regression trees: Applications in oceanology. *Computational Statistics & Data Analysis,* 51**,** 4984-4993.

Ngoduy, D., Hoang, N. H., Vu, H. L. & Watling, D. (2016) Optimal queue placement in dynamic system optimum solutions for single origin-destination traffic networks. *Transportation Research Part B: Methodological,* 92**,** 148-169.

Noland, R. B. & Small, K. A. (1995) Travel-time uncertainty, departure time choice, and the cost of the morning commute. Irvine: Institute of Transportation Studies, University of California.

O'Cathain, A., Murphy, E. & Nicholl, J. (2010) Three techniques for integrating data in mixed methods studies. *BMJ,* 341.

Ofcom (2016) Adults' media use and attitudes. Available: http://stakeholders.ofcom.org.uk/market-data-research/other/research-publications/adults/media-lit-2016/ [Accessed: 9 September 2016].

Office for National Statistics. (2016) *Population Estimates Analysis Tool: Mid 2015* [Online]. Available: https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/populationestimatesanalysistool [Accessed: 19 September 2016].

Ortúzar, J. d. D. & Willumsen, L. G. (2011) *Modelling Transport,* Chichester, UK, John Wiley & Sons Ltd.

Ozbay, K., Yang, H., Morgul, E. F., Mudigonda, S. & Bartin, B. Big Data and the Calibration and Validation of Traffic Simulation Models. Looking Back and Looking Ahead: Celebrating 50 Years of Traffic Flow Theory, 2014 Washington D.C.: Transportation Research Board, 92-122.

Park, H. & Haghani, A. (2015) Optimal number and location of Bluetooth sensors considering stochastic travel time prediction. *Transportation Research Part C: Emerging Technologies,* 55**,** 203-216.

Paz, A., Emaasit, D. & de la Fuente, H. Stochastic Dynamic User Equilibrium Using a Mixed Logit Modeling Framework. IEEE 19th International Conference on Intelligent Transportation Systems (ITSC), 2016 Rio de Janeiro, Brazil. 1206-1211.

Peeta, S. & Ziliaskopoulos, A. K. (2001) Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. *Networks and Spatial Economics,* 1**,** 233-265.

Pelletier, M.-P., Trépanier, M. & Morency, C. (2011) Smart card data use in public transit: A literature review. *Transportation Research Part C: Emerging Technologies,* 19**,** 557-568.

Phua, P., Page, B. & Bogomolova, S. (2015) Validating Bluetooth logging as metric for shopper behaviour studies. *Journal of Retailing and Consumer Services,* 22**,** 158-163.

Prashker, J. N. & Bekhor, S. (2004) Route Choice Models Used in the Stochastic User Equilibrium Problem: A Review. *Transport Reviews,* 24**,** 437-463.

Pu, W. (2011) Analytic Relationships Between Travel Time Reliability Measures. *Transportation Research Record: Journal of the Transportation Research Board,* 2254**,** 122-130.

Quayle, S., Koonce, P., DePencier, D. & Bullock, D. (2010) Arterial Performance Measures with Media Access Control Readers. *Transportation Research Record: Journal of the Transportation Research Board,* 2192**,** 185-193.

Rakha, H. & Van Aerde, M. (1995) Statistical analysis of day-to-day variations in real-time traffic flow data. *Transportation Research Record: Journal of the Transportation Research Board,* 1510**,** 26-34.

Ramsay, J. O., Hooker, G. & Graves, S. (2009) *Functional Data Analysis with R and MATLAB,* New York, Springer.

Ramsay, J. O. & Silverman, B. W. (1997) *Functional Data Analysis,* New York, Springer.

Ramsay, J. O. & Silverman, B. W. (2005) *Functional Data Analysis,* New York, Springer.

Ramsay, J. O., Wickham, H., Graves, S. & Hooker, G. (2014) *Package 'fda'* [Online]. The Comprehensive R Archive Network (CRAN). Available: http://cran.r-project.org/web/packages/fda/fda.pdf [Accessed: 17th December 2014].

Raux, C., Ma, T.-Y. & Cornelis, E. (2016) Variability in daily activity-travel patterns: the case of a one-week travel diary. *European Transport Research Review,* 8.

Safi, H., Assemi, B., Mesbah, M., Ferreira, L. & Hickman, M. (2015) Design and Implementation of a Smartphone-Based Travel Survey. *Transportation Research Record: Journal of the Transportation Research Board,* 2526**,** 99-107.

Schlich, R. & Axhausen, K. W. (2003) Habitual travel behaviour: Evidence from a six-week travel diary. *Transportation,* 30**,** 13-36.

Schmöller, S., Weikl, S., Müller, J. & Bogenberger, K. (2015) Empirical analysis of free-floating carsharing usage: The Munich and Berlin case. *Transportation Research Part C: Emerging Technologies,* 56**,** 34-51.

Shahhoseini, Z., Haghani, M. & Sarvi, M. (2015) Estimation and application of a multi-class multi-criteria mixed paired combinatorial logit model for transport networks analysis. *Transportmetrica B: Transport Dynamics,* 3**,** 59-78.

Shang, W., Han, K., Ochieng, W. & Angeloudis, P. (2016) Agent-based day-to-day traffic network model with information percolation. *Transportmetrica A: Transport Science,* 13**,** 38-66.

Shao, H., Lam, W. H. K. & Tam, M. L. (2006) A reliability-based stochastic traffic assignment model for network with multiple user classes under uncertainty in demand. *Networks & Spatial Economics,* 6**,** 173-204.

Sheffi, Y. (1985) *Urban Transportation Networks,* Englewood Cliffs, Prentice-Hall Inc.

Shen, Q. & Faraway, J. (2004) An F Test for Linear Models with Functional Responses. *Statistica Sinica,* 14**,** 1239-1257.

Shen, Q. & Xu, H. (2007) Diagnostics for Linear Models With Functional Responses. *Technometrics,* 49**,** 26-33.

Shen, Y., Kwan, M.-P. & Chai, Y. (2013) Investigating commuting flexibility with GPS data and 3D geovisualization: a case study of Beijing, China. *Journal of Transport Geography,* 32**,** 1-11.

Shoval, N. & Isaacson, M. (2007) Sequence Alignment as a Method for Human Activity Analysis in Space and Time. *Annals of the Association of American Geographers,* 97**,** 282-297.

Silverman, B. W. (1985) Some Aspects of the Spline Smoothing Approach to Non-Parametric Regression Curve Fitting. *Journal of the Royal Statistical Society: Series B (Methodological),* 47**,** 1-52.

Singhal, A., Kamga, C. & Yazici, A. (2014) Impact of weather on urban transit ridership. *Transportation Research Part A: Policy and Practice,* 69**,** 379-391.

Siu, B. W. Y. & Lo, H. K. (2008) Doubly uncertain transportation network: Degradable capacity and stochastic demand. *European Journal of Operational Research,* 191**,** 166-181.

Smith, M. J. (1979) The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological,* 13**,** 295-304.

Spissu, E., Meloni, I. & Sanjust, B. (2011) Behavioral Analysis of Choice of Daily Route with Data from Global Positioning System. *Transportation Research Record: Journal of the Transportation Research Board,* 2230**,** 96-103.

Stathopoulos, A. & Karlaftis, M. (2001) Temporal and Spatial Variations of Real-Time Traffic Data in Urban Areas. *Transportation Research Record: Journal of the Transportation Research Board,* 1768**,** 135-140.

Stopher, P. R. & Zhang, Y. (2011) Repetitiveness of Daily Travel. *Transportation Research Record: Journal of the Transportation Research Board,* 2230**,** 75-84.

Studer, M. (2013) WeightedCluster Library Manual: A practical guide to creating typologies of trajectories in the social sciences with R. *LIVES Working Papers* [Online], 24.

Sumalee, A., Uchida, K. & Lam, W. H. K. (2011a) Stochastic multi-modal transport network under demand uncertainties and adverse weather condition. *Transportation Research Part C: Emerging Technologies,* 19**,** 338-350.

Sumalee, A., Watling, D. & Nakayama, S. (2006) Reliable Network Design Problem: Case with Uncertain Demand and Total Travel Time Reliability. *Transportation Research Record: Journal of the Transportation Research Board,* 1964**,** 81-90.

Sumalee, A., Zhong, R. X., Pan, T. L. & Szeto, W. Y. (2011b) Stochastic cell transmission model (SCTM): A stochastic dynamic traffic model for traffic state surveillance and assignment. *Transportation Research Part B: Methodological,* 45**,** 507-533.

Sun, C., Cheng, L., Zhu, S. & Chu, Z. (2015) Multiclass Stochastic User Equilibrium Model with Elastic Demand. *Transportation Research Record: Journal of the Transportation Research Board,* 2497**,** 1-11.

Susilo, Y. O. & Axhausen, K. W. (2014) Repetitions in individual daily activity–travel–location patterns: a study using the Herfindahl–Hirschman Index. *Transportation,* 41**,** 995-1011.

Szeto, W. Y., Jiang, Y. & Sumalee, A. (2011) A Cell-Based Model for Multi-class Doubly Stochastic Dynamic Traffic Assignment. *Computer-Aided Civil and Infrastructure Engineering,* 26**,** 595-611.

Tang, J., Wang, Y., Wang, H., Zhang, S. & Liu, F. (2014) Dynamic analysis of traffic time series at different temporal scales: A complex networks approach. *Physica A: Statistical Mechanics and its Applications,* 405**,** 303-315.

Tao, S., Corcoran, J., Mateo-Babiano, I. & Rohde, D. (2014) Exploring Bus Rapid Transit passenger travel behaviour using big data. *Applied Geography,* 53**,** 90-104.

Tarigan, A. & Kitamura, R. (2009) Week-to-Week Leisure Trip Frequency and Its Variability. *Transportation Research Record: Journal of the Transportation Research Board,* 2135**,** 43-51.

Taylor, M. A. P. (2013) Travel through time: the story of research on travel time reliability. *Transportmetrica B-Transport Dynamics,* 1**,** 174-194.

Tenekeci, G., Montgomery, F. & Wainaina, S. (2010) Roundabout capacity in adverse weather and light conditions. *Proceedings of the ICE - Transport,* 163**,** 29-39.

Tibshirani, R. (1996) Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological),* 58**,** 267-288.

Transport for Greater Manchester. (n.d.) *Transport Statistics Greater Manchester Key Centres Section* [Online]. Available: https://data.gov.uk/dataset/transport-statistics-greater-manchester-key-centres-section [Accessed: 19 September 2016].

Transport for London (2013) Olympic Legacy Monitoring: Personal Travel Behaviour during the Games Available: content.tfl.gov.uk/olympic-legacy-personal-travel-report.pdf [Accessed: 29th July 2016].

Tsubota, T., Bhaskar, A., Chung, E. & Billot, R. Arterial traffic congestion analysis using Bluetooth duration data. *In:* TISATO, P., OXLAD, L. & TAYLOR, M., eds., Australasian Transport Research Forum, 2011 Adelaide. p1-14.

Tu, Y., Ball, M. O. & Jank, W. S. (2008) Estimating Flight Departure Delay Distributions—A Statistical Approach With Long-Term Trend and Short-Term Pattern. *Journal of the American Statistical Association,* 103**,** 112-125.

U.S. Department of Transportation (2006) Traffic Detector Handbook. Third edition ed. Available: https://www.fhwa.dot.gov/publications/research/operations/its/06108/06108.pdf [Accessed: 21st August 2017].

Ukkusuri, S. V., Han, L. & Doan, K. (2012) Dynamic user equilibrium with a path based cell transmission model for general traffic networks. *Transportation Research Part B: Methodological,* 46**,** 1657-1684.

USA Today. (2015) *One huge mistake people make when renting cars* [Online]. Available: http://www.usatoday.com/story/tech/columnist/komando/2015/07/03/komando-car-rental-mistakes/29614165/ [Accessed: 22 September 2016].

Utsunomiya, M., Attanucci, J. & Wilson, N. (2006) Potential Uses of Transit Smart Card Registration and Transaction Data to Improve Transit Planning. *Transportation Research Record: Journal of the Transportation Research Board,* 1971**,** 119-126.

van Goeverden, C. D., Botma, H. & Bovy, P. H. L. (1998) Determining Impact of Road Lighting on Motorway Capacity. *Transportation Research Record: Journal of the Transportation Research Board,* 1646**,** 1-8.

Venter, C. & Joubert, J. (2013) Use of Multisource Global Positioning System Data to Characterize Multiday Driving Patterns and Fuel Usage in a Large Urban Region. *Transportation Research Record: Journal of the Transportation Research Board,* 2338**,** 1-10.

Versichele, M., Neutens, T., Delafontaine, M. & Van de Weghe, N. (2012) The use of Bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent Festivities. *Applied Geography,* 32**,** 208-220.

Vlahogianni, E. I., Golias, J. C. & Karlaftis, M. G. (2004) Short‐term traffic forecasting: Overview of objectives and methods. *Transport Reviews,* 24**,** 533-557.

Vlahogianni, E. I., Karlaftis, M. G. & Golias, J. C. (2014) Short-term traffic forecasting: Where we are and where we're going. *Transportation Research Part C: Emerging Technologies,* 43**,** 3-19.

Wang, D. Z. W. & Du, B. (2016) Continuum modelling of spatial and dynamic equilibrium in a travel corridor with heterogeneous commuters—A partial differential complementarity system approach. *Transportation Research Part B: Methodological,* 85**,** 1-18.

Wardrop, J. G. (1952) Some theoretical aspects of road traffic research. *Road Engineering Division Meeting.*

Watling, D. (2002) A Second Order Stochastic Network Equilibrium Model, I: Theoretical Foundation. *Transportation Science,* 36**,** 149-166.

Watling, D., Milne, D. & Clark, S. (2012) Network impacts of a road capacity reduction: Empirical analysis and model predictions. *Transportation Research Part A: Policy and Practice,* 46**,** 167-189.

Watling, D., Sumalee, A., Connors, R. & Balijepalli, C. (2004) Advancing methods for evaluating network reliability: A Department for Transport "New Horizons" Project Final Report. Available: http://webarchive.nationalarchives.gov.uk/20070306100043/http:/www.dft.g ov.uk/pgr/economics/rdg/jtv/amenr/advancingmethodsforevaluatin3080 [Accessed: 14th May 2014].

Watling, D. P. & Cantarella, G. E. (2013a) Model Representation & Decision-Making in an Ever-Changing World: The Role of Stochastic Process Models of Transportation Systems. *Networks and Spatial Economics,* 15**,** 843-882.

Watling, D. P. & Cantarella, G. E. (2013b) Modelling sources of variation in transportation systems: theoretical foundations of day-to-day dynamic models. *Transportmetrica B: Transport Dynamics,* 1**,** 3-32.

Wei, C., Asakura, Y. & Iryo, T. (2014) Formulating the within-day dynamic stochastic traffic assignment problem from a Bayesian perspective. *Transportation Research Part B: Methodological,* 59**,** 45-57.

Weijermars, W. & van Berkum, E. Analyzing highway flow patterns using cluster analysis. 8th International IEEE Conference on Intelligent Transportation Systems, 2005 Vienna, Austria. 831-836.

Weijermars, W. A. M. & van Berkum, E. C. (2004) Daily flow profiles of urban traffic. *In:* BREBBIA, C. A. & WADHWA, L. C. (eds.) *Urban Transport X. Urban Transport and the Environment in the 21st Century.* Dresden, Germany: WIT Press.

Xianyu, J., Rasouli, S. & Timmermans, H. (2017) Analysis of variability in multi-day GPS imputed activity-travel diaries using multi-dimensional sequence alignment and panel effects regression models. *Transportation,* 44**,** 533-553.

Xiao, F., Yang, H. & Ye, H. (2016) Physics of day-to-day network flow dynamics. *Transportation Research Part B: Methodological,* 86**,** 86-103.

Xiao, Y. & Lo, H. K. (2016) Day-to-day departure time modeling under social network influence. *Transportation Research Part B: Methodological,* 92**,** 54-72.

Xu, M., Meng, Q. & Liu, K. (2017) Network user equilibrium problems for the mixed battery electric vehicles and gasoline vehicles subject to battery swapping stations and road grade constraints. *Transportation Research Part B: Methodological,* 99**,** 138-166.

Yang, L., Kwan, M.-P., Pan, X., Wan, B. & Zhou, S. (2017) Scalable space-time trajectory cube for path-finding: A study using big taxi trajectory data. *Transportation Research Part B: Methodological,* 101**,** 1-27.

Yang, X., Shen, Q., Xu, H. & Shoptaw, S. (2007) Functional regression analysis using an F test for longitudinal data with large numbers of repeated measures. *Stat Med,* 26**,** 1552-66.

Yazici, M., Kamga, C. & Mouskos, K. (2012) Analysis of Travel Time Reliability in New York City Based on Day-of-Week and Time-of-Day Periods. *Transportation Research Record: Journal of the Transportation Research Board,* 2308**,** 83-95.

Yeon, J., Hernandez, S. & Elefteriadou, L. (2009) Differences in Freeway Capacity by Day of the Week, Time of Day, and Segment Type. *Journal of Transportation Engineering,* 135**,** 416-426.

Yu, W., Park, S., Kim, D. S. & Ko, S.-S. (2015) Arterial Road Incident Detection Based on Time-Moving Average Method in Bluetooth-Based Wireless Vehicle Reidentification System. *Journal of Transportation Engineering,* 141**,** 04014084 1-14.

Zhang, J.-T. (2013) Analysis of Variance for Functional Data. Boca Raton, Florida: CRC Press.

Zhang, J., Timmermans, H. & Borgers, A. (2002) Utility-Maximizing Model of Household Time Use for Independent, Shared, and Allocated Activities Incorporating Group Decision Mechanisms. *Transportation Research Record: Journal of the Transportation Research Board,* 1807**,** 1-8.

Zhang, K., Mahmassani, H. S. & Lu, C.-C. (2013) Dynamic pricing, heterogeneous users and perception error: Probit-based bi-criterion dynamic stochastic user equilibrium assignment. *Transportation Research Part C: Emerging Technologies,* 27**,** 189-204.

Zhang, W., Mediria, A. & Rakha, H. Statistical analysis of spatiotemporal link and path flow variability. 2007 IEEE Intelligent Transportation Systems Conference, 2007 Seattle, Washington, USA. IEEE, 1080-1085.