Microscopic simulations of complex metropolitan dynamics

Mark Birkin and Nick Malleson
School of Geography, University of Leeds, Leeds LS2 9JT

December 2011

Paper presented at the *Complex City* workshop, 5th–6th December, Amsterdam

Contents

Abstract

In this paper we will introduce an individual-based model of a British city. The approach draws its inspiration from both microsimulation and agent-based modelling. Microsimulation is used to reconstruct the entire population of a city region at both the household and individual scale. We illustrate this for Leeds, a city with three-quarters of a million inhabitants grouped into more than 300 thousand households. The resulting population is profiled by demographic and social attributes which are richly specified.

In order to incorporate dynamic individual behaviour, we argue that agent-based simulations are more appropriate, and rules will be presented which allow the identification of five essential behaviours which we term domestic living, education, work, recreation and shopping. Through this modelling process we seek not just to understand residential patterns within the city, but the dynamic ebb and flow of the population in everyday metropolitan life.

The novel feature of our research is that we will use up-to-date social network data to calibrate our agent behaviours. Although social network data are likely to be somewhat skewed and unreliable, they are abundant and continually refreshed and also provide temporally-accurate daily behavioural information that are often absent from traditional sources (such as censuses). We will make an attempt to evaluate the robustness and (potential) value of this approach.

1 Introduction

When we consider the long-term evolution of urban areas then, naturally, relative location in space is a crucial factor. In the classic theories, competition for space gives rise to zoning of urban activities (Burgess, 1925) which can be reinforced by the manipulation of transport networks or land use plans, and standard equilibrium models to approximate these processes have long been available (see, for example, Putman (1983) or Miller et al. (2004)) for a slightly more modern twist).

Recent work recognising the character of cities as complex systems has made it clear that at a disaggregate level cities are in a constant state of flux and aggregate models have the effect of smoothing out this underlying dynamism (Batty, 2005). The phenomenon of traffic congestion is an excellent example. Although congestion emerges consistently at known times and places and appears to be in equilibrium, observing the phenomenon at a higher resolution reveals a complex, fluctuating event were thousands of individual units make unique decisions to drive a phenomenon that is far from stable. To avoid this smoothing effect, individual-level models have emerged as more realistic representations of cities. Common methods include cellular automata (CA) (Keith and Gaydos, 1998), agent-based modelling (ABM) (Parker et al., 2003) and microsimulation (Birkin et al., 2006; Ballas and Clarke, 2001).

Some of the difficulty here simply reflects the fact that the critical interrelationship between time and space is undervalued in the classical models of urban structure. Although time may be recognised to the extent that lack of accessibility or long journey times generates a reduction in utility, a more fundamental consideration is that people themselves cannot be in two places at one time; and more generally where they are and what they are doing at one point in time is a restriction on where they can go and what they do next. Of course this observation is not especially novel and forms the central pillar of Torsten Hägerstrand's notion of time geography (Hägerstrand, 1970), which is also of central importance to the related developments of behavioural geography in the work of scholars such as Alan Pred, Reg Golledge and others (Pred, 1967; Golledge et al., 1972).

Consider, for example, the famous 'space-time prism' which is represented in Figure 1. Here an individual has a known location at two points in time, t_i and t_j – suppose that the individual is initially at home and later at work. The prism enfolds the set of all locations which it is possible to reach under these constraints, within a fixed region denoted in the shaded area of the illustration. Whilst these frameworks have been of conceptual and empirical interest to geographers for a long period of time, one of the contentions which we wish to explore in this paper is that for the first time new data sources, and perhaps to some extent the availability of associated modelling methods and frameworks, makes it possible to think about testing and applying these concepts on a much more universal scale – for example, for an entire city.

As individual-level models directly simulate the behaviour of individual city occupants, they are ideally suited to capturing Hägerstrand's theoretical framework and providing a strong theoretical base to a spatial urban model. Agent-based modelling, in particular, potentially offers the most substantial benefit because it can provide the most "natural description" (Bonabeau, 2002) of a system. Specific advantages of the methodology, over aggregate or statistical techniques, include:

- **Capturing complexity**. Linear models face difficulties with modelling complex systems (such as cities) because they generally use simple functional relationships; failing to capture factors such as the the historical path of individuals and its effect on their later behaviour.

- **Spatial realism**. It is easy to incorporate a realistic spatial environment in an agent-based model. This is not the case with many other techniques which revert

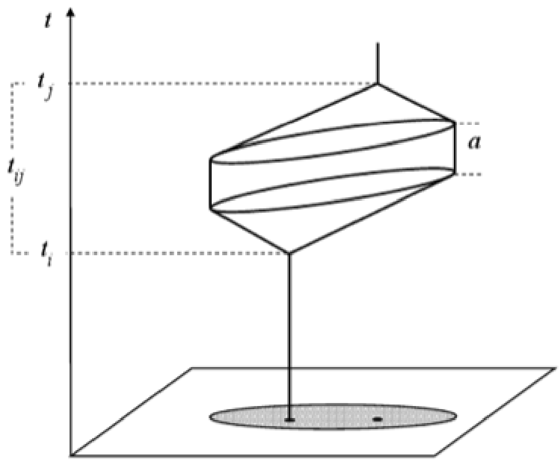

Figure 1: The 'space time prism' (Hägerstrand, 1970). Source: Miller (2004)

to Euclidean distance measures. These do not take items such as road networks, congestion or impassable barriers into account which will strongly influence where, and how, people travel.

- **Non-monotonically changing variables**. It is likely that some variables in a complex social system will have both positive and negative effects on an outcome depending on their context. With crime, for example, a traffic volume variable is neither positively or negatively correlated to burglary risk – it could increase risk because a larger number of people will pass a house and be aware of it, but it could decrease risk because the footfall might make it more difficult to actually break into. An agent-based model can cope with these effects directly simulating peoples' movement, rather than attempting to work to an average, general rule. On the whole, the functional complexity of the problem is greater than statistical modelling will cope with.

Prior research (Birkin et al., 2006) has used UK census data to synthesise a richly specified population of individual people and households for the city of Leeds, UK and a dynamic spatial microsimulation to model demographic change. Advantages of the spatial microsimulation approach are that it allows an efficient representation of highly disaggregate populations, and is well suited to modelling transitions of 'slowly evolving' individual and household characteristics, such as age, employment, home address, income etc (Birkin and Clarke, 2011). However, microsimulation is not appropriate for more general social phenomenon such as retailing, crime, transport, etc. which involve characteristics that evolve more quickly. In these cases, an approach is required that focusses much more heavily on an individuals' cognition, choice, interactions and overall behaviour. For this reason, agent-based modelling is an essential methodological component for this research.

The most contentious issues surrounding the use of agent-based models are how they should be calibrated and validated and in many respects agent-based models of social phenomena are lagging behind other fields in terms of how they approach these problems. Meteorology models, for example, use incoming weather data to improve their predictions as new data become available in real time (Collins, 2007). Similarly, it is increasingly common for transport control systems to rely on real-time feedback between current traffic systems and control networks such as traffic lights or variable direction of flow (e.g. Prothmann et al., 2008). Agent-based models rarely, if ever, use dynamic data streams to improve their predictions as new social data become available.

The reason that social models are lacking in these respects is largely due to data availability; models often use population censuses and social surveys that tend to deal with aggregate groups rather than individuals and occur infrequently. They are focused on the attributes and characteristics of the population, rather than attitudes and behaviours, and they offer a snapshot view rather than a dynamic and continuous perspective. In recent years, however, new data sources have become available that contain a wealth of information about peoples' spatio-temporal behaviour at an individual-level and are being updated continuously. These sources are commonly referred to as "crowd-sourced data" (Savage and Burrows, 2007) or "volunteered geographical information" (Goodchild, 2007) and have the potential to revolutionise our understanding of social phenomena and our approach to model evaluation.

The focus of this research is to develop a framework for calibrating an agent-based model of daily travel behaviour using a novel set of crowd-sourced data from Twitter. In order to incorporate dynamic individual behaviour, we identify five essential behaviours: domestic living, education, work, recreation and shopping. We begin with a highly detailed synthetic population of individuals and outline the steps towards constructing a detailed agent-based model of daily behaviour using available crowd-sourced data. Through this modelling process we seek not just to understand residential patterns within the city, but the dynamic ebb and flow of the population in everyday metropolitan life.

2 Crowd-Sourced Data from Twitter

2.1 Crowd-Sourced Data

The capacity to capture, process and analyse data has been increasing dramatically as technologies have continued to evolve and progress. In the field of computational science and engineering, commentators have gone so far as to suggest the emergence of a new 'fourth paradigm' of data intensive experimental reasoning (Bell et al., 2009). These trends are perhaps even more acutely manifest with respect to social data. Savage and Burrows (2007) have noted a "crisis" in "empirical sociology" arising from the fact that survey-based enquiries into the behaviour of small and selected groups of individuals are now potentially superseded by massively "crowd-sourced" databases which monitor the activity patterns on either a discrete or a continuous basis. Commentators with a more specific interest in the representation of spatial phenomena have noted a rise in "volunteered geographical information" (Goodchild, 2007) including both map and feature data. The growth of OpenStreetMap as a credible alternative to proprietary datasets (e.g. Ordnance Survey) is a notable example in this context.

2.2 Data Overview

For this research, crowd-sourced data from the Twitter service is being used. Twitter is a social networking / microblogging service that allows users to broadcast short messages of up to 140 characters called *tweets*. Although it is possible to broadcast private messages to other users the most common use of the service is for public information sharing; a public tweet can be read by any Twitter user. In order to allow access to a larger volume of data than a single person could analyse, Twitter also provide the 'Streaming API'. This service allows a computer to listen to all tweets filtered in a particular manner, such as by their geographic location. Not all tweets have an associated location, however. For privacy reasons a user must enable the 'Tweeting With Location' feature (it is inactive by default) and it is possible that the device used to create the tweet has no means of establishing its exact location in the first place. Therefore the data used in this research do not represent all tweets in the Leeds area, but rather all tweets in the area that also have exact geographical information.

Tweets that can be geolocated to an area surrounding Leeds (as shown in Figure2) have been collected for the time period 22nd June – 12th October 2011. This consists of 290,215 individual tweets from 9,223 different users. The data are highly skewed; the top 10% most prolific users (922) are responsible for 78% of the total tweets (227,652) and 28% of the users (2,612) only generate one tweet. Although some of the tweet data clearly originate from businesses – examples include advertising cars or night club events – a cursory inspection suggests that most data have been generated by real people. Although formal cleaning routines will be required if the data are to be used as model inputs, in the analyses later we focus on individual user behaviour and thus disregard superfluous data.

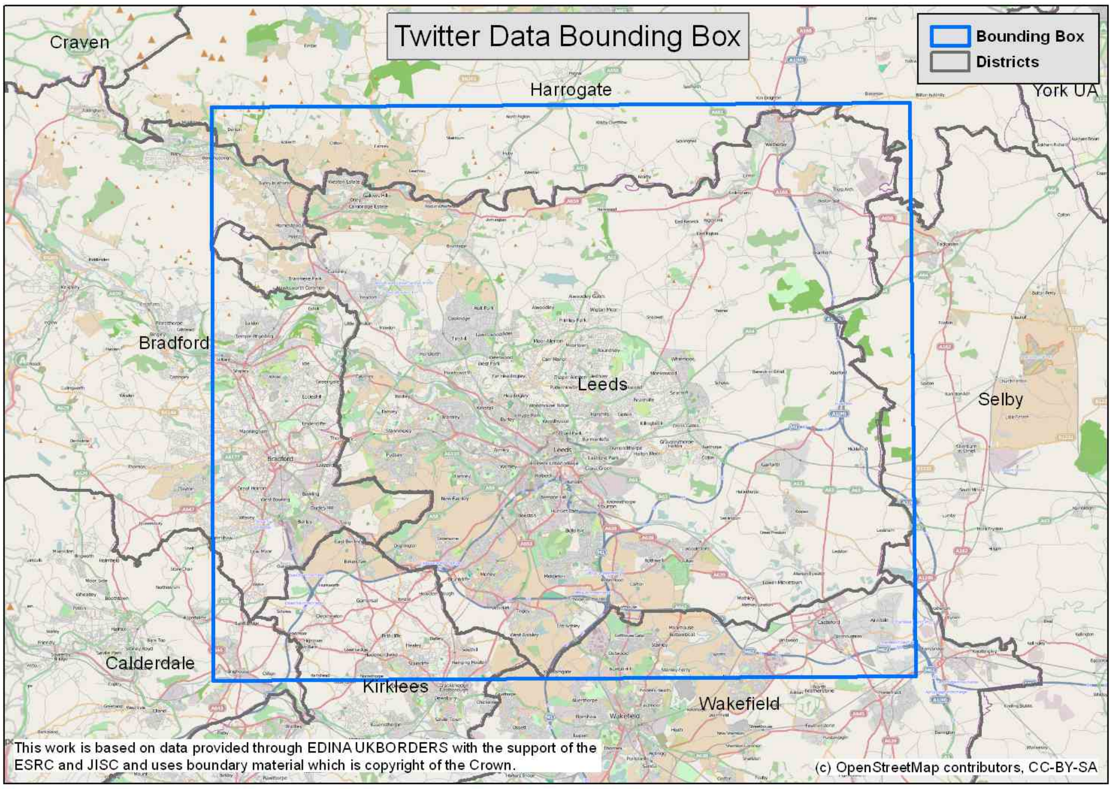

Figure 2: The boundary area for tweet data collected by this research.

Figure 3 presents the density of all tweets, generated calculated using the kernel density estimation algorithm. As would be expected, areas of high population density such as central Leeds, Bradford and surrounding smaller towns (e.g. Otley, Weatherby, Guiseley) exhibit the highest tweet density. By observing the three-dimensional view of the twitter data, it becomes clear that a single location is responsible for a substantial spike in the density. This spike relates to a very large number of tweets from a single user in a single location. Fortunately, as we will discuss shortly, many prolific users send tweets from different locations, making it possible to elucidate information about their daily spatial behaviour.

It should be noted that, at the time of writing, data are still being collected and, as discussed in Section 3.2, future work will move towards making use of streamed data as it becomes available, rather than storing it for later analysis.

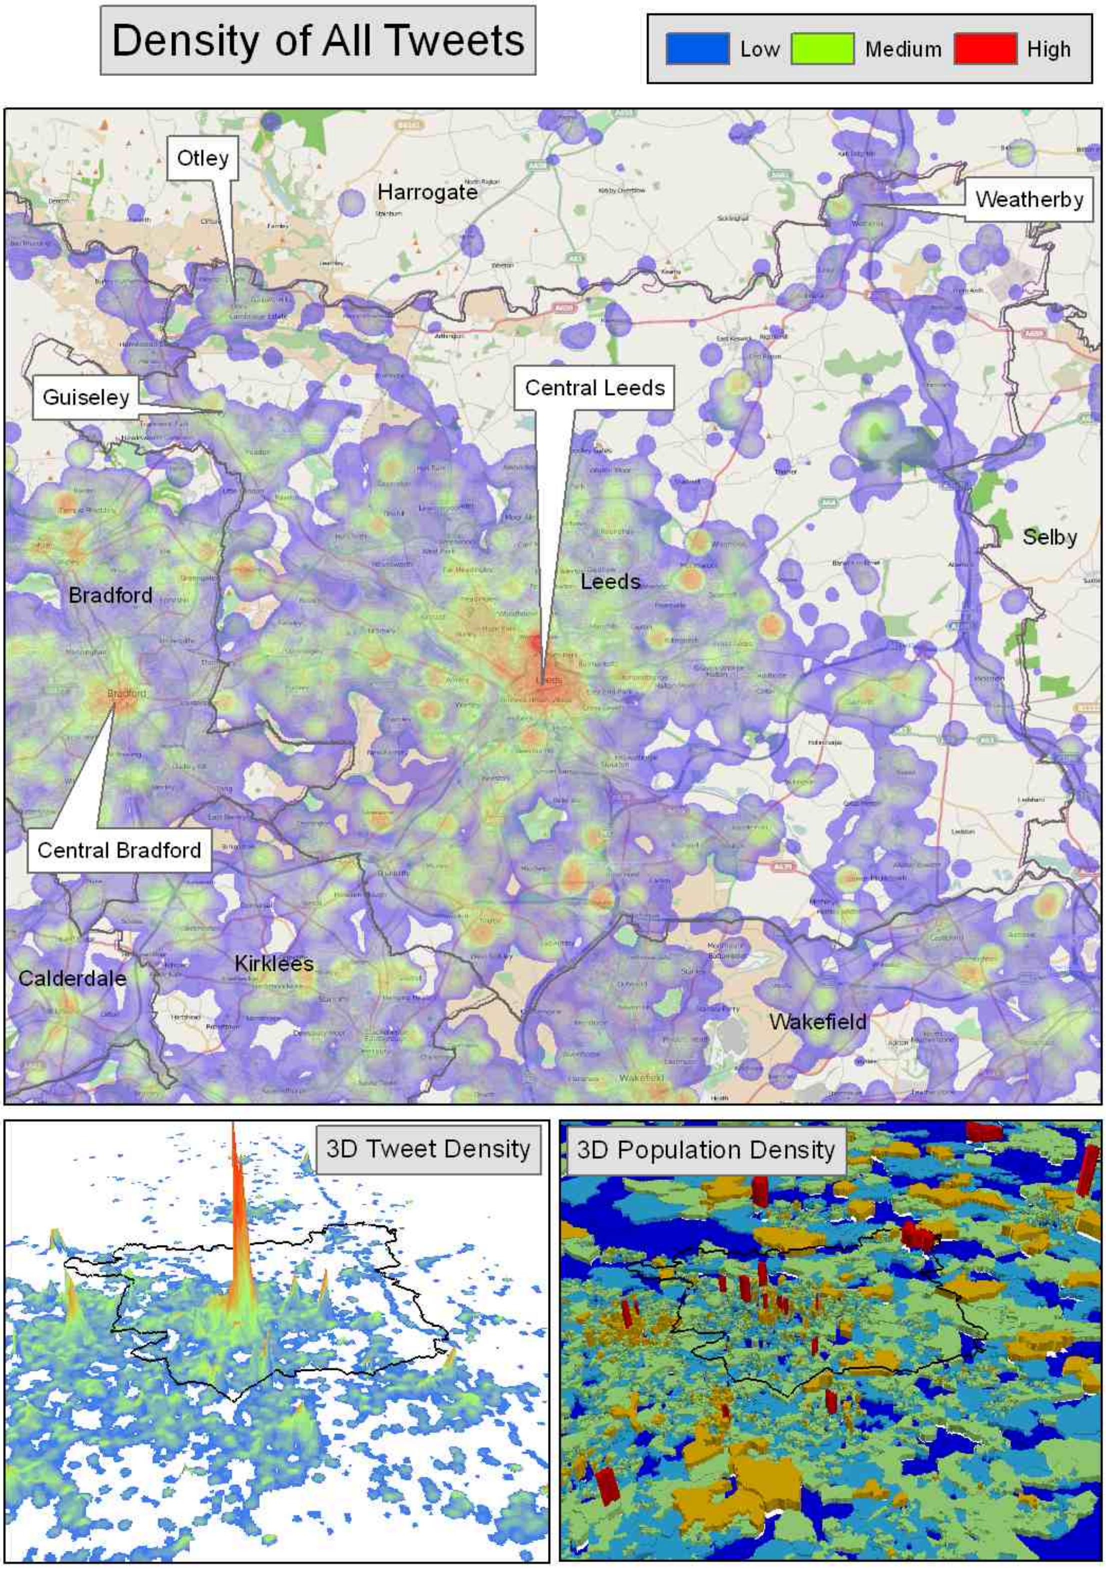

Figure 3: The density of all tweets in the data.

2.3 Tweeting Communities

Figure 3 also compares the density of tweets to the size of the residential population as recorded in the 2001 UK census. Interestingly, there appears to be no relationship between tweet density and population density and this is verified by the scatterplot in Figure 4. This finding is somewhat unexpected because, given the obvious tweet density spikes around population centres, a relationship to residential population would be an understandable assumption. The fact that tweets do not correlate to residential density is encouraging as it suggests that twitter data could be used as a means of improving non-residential population estimates which are extremely difficult to estimate from traditional sources such as censuses. On the other hand, it could be simply that the use of tweets as the independent variable allows a small number of randomly located high volume users to influence this pattern, and that an ability to discriminate between tweets and tweeters – in particular, to define the normal residence for tweeters/ users – could help to refine this pattern. Section 3.2 will discuss this in more detail.

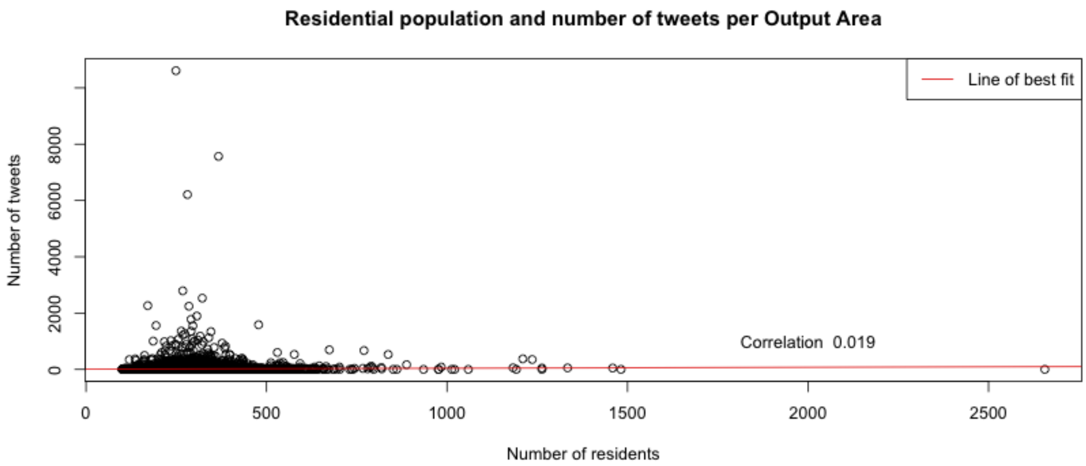

Figure 4: The number of tweets and the number of people (from the UK census) in each output area.

Further information about where tweets originate from, which will be extremely valuable in assessing how representative the data are of the whole Leeds population, can be found by examining the geodemographic profiles of the areas that the tweets are located in. The output area classification (OAC: Vickers and Rees, 2007) is a demographic classification based on the 2001 UK census. The classification organises each output area (the smallest geographic boundary at which census data are released) into a number of distinct categories. To assess which OAC groups have the largest number of tweets, Table 1 provides analysis of tweets for seven high level OAC groups. The ratio of tweet proportions to OAC group proportions is indexed so that OAC groups with index values above 100 have more tweets than would be expected if tweets were distributed evenly across the whole city and those with an index under 100 have fewer tweets than would be expected.

Interestingly, 'city living' areas have three times as many tweets in them than would be expected if tweets were distributed evenly; a considerable amount. This adds support to the finding that most tweets emerge from non-residential areas because this group largely represents city-centre communities. However, as Section 3.2 will discuss, the subset of users who produce the largest number of tweets do not follow this pattern so strongly. Table 1 also exhibits a notable dip in tweeting activity in areas characterised by a high concentration of elderly ('prospering suburbs'), manual employment ('blue collar') and

Table 1: The number of tweets originating from different OAC groups in Leeds.

	Counts		Proportions		
OAC Group	Tweets	Groups	Tweets	Groups	Index
2 - City Living	34482	211	0.268	0.087	309.5
3 - Countryside	4440	62	0.034	0.025	135.6
7 - Multicultural	16883	273	0.131	0.112	117.1
6 - Typical Traits	26283	571	0.204	0.234	87.2
5 - Constrained by Circumstances	19818	443	0.154	0.182	84.7
4 - Prospering Suburbs	17524	530	0.136	0.217	62.6
1 - Blue Collar	9349	349	0.073	0.143	50.7
Grand Total	128779	2439	1	1	

deprived ('constrained by circumstances') populations.

2.4 Temporal Patterns

In order to assess whether or not temporal behaviour is being exhibited by the tweet data, Figure 5 illustrates the tweet rate per day for part of the data series and tweets per hour for all tweets. Surprisingly there is no obvious daily trend apparent – it might be expected that weekends displayed a different number of tweets than weekdays for example. There is a clear hourly pattern though; the number of tweets increases gradually until approximately 8pm when it drops to its lowest levels overnight.

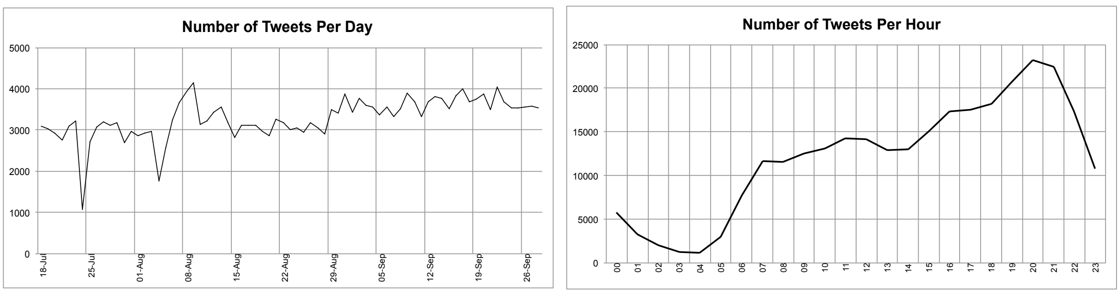

Figure 5: Temporal variations in the total number of tweets published.

It is also possible to explore the spatial distribution of tweet densities at different times of day. Figure 6 illustrates the density of all tweets in the data that occurred on a Tuesday during different time periods throughout the day. Tuesday was chosen because it is likely to be representative of general weekday behaviour. It is striking that although Figure 5 provides clear evidence for daily temporal trends, the spatial distribution of tweets does not vary in such a clearly defined manner. For example, there is no obvious trend of density increasing during the day as people travel to the city centre for work and then spreading out to the suburbs again as people go home. A way forward might be to isolate only the people who tweet the most regularly (i.e. the 10% of users responsible for 78% of the tweets) and explore the density of their tweets. As these prolific users are more likely to exhibit understandable spatio-temporal behaviour, the density of their tweets in isolation might provide a clearer picture of aggregate behaviour. This will be explored as future work.

Fortunately, because the data are available at the level of the individual person, it is possible to analyse individual users in detail without the reliance on aggregate patterns

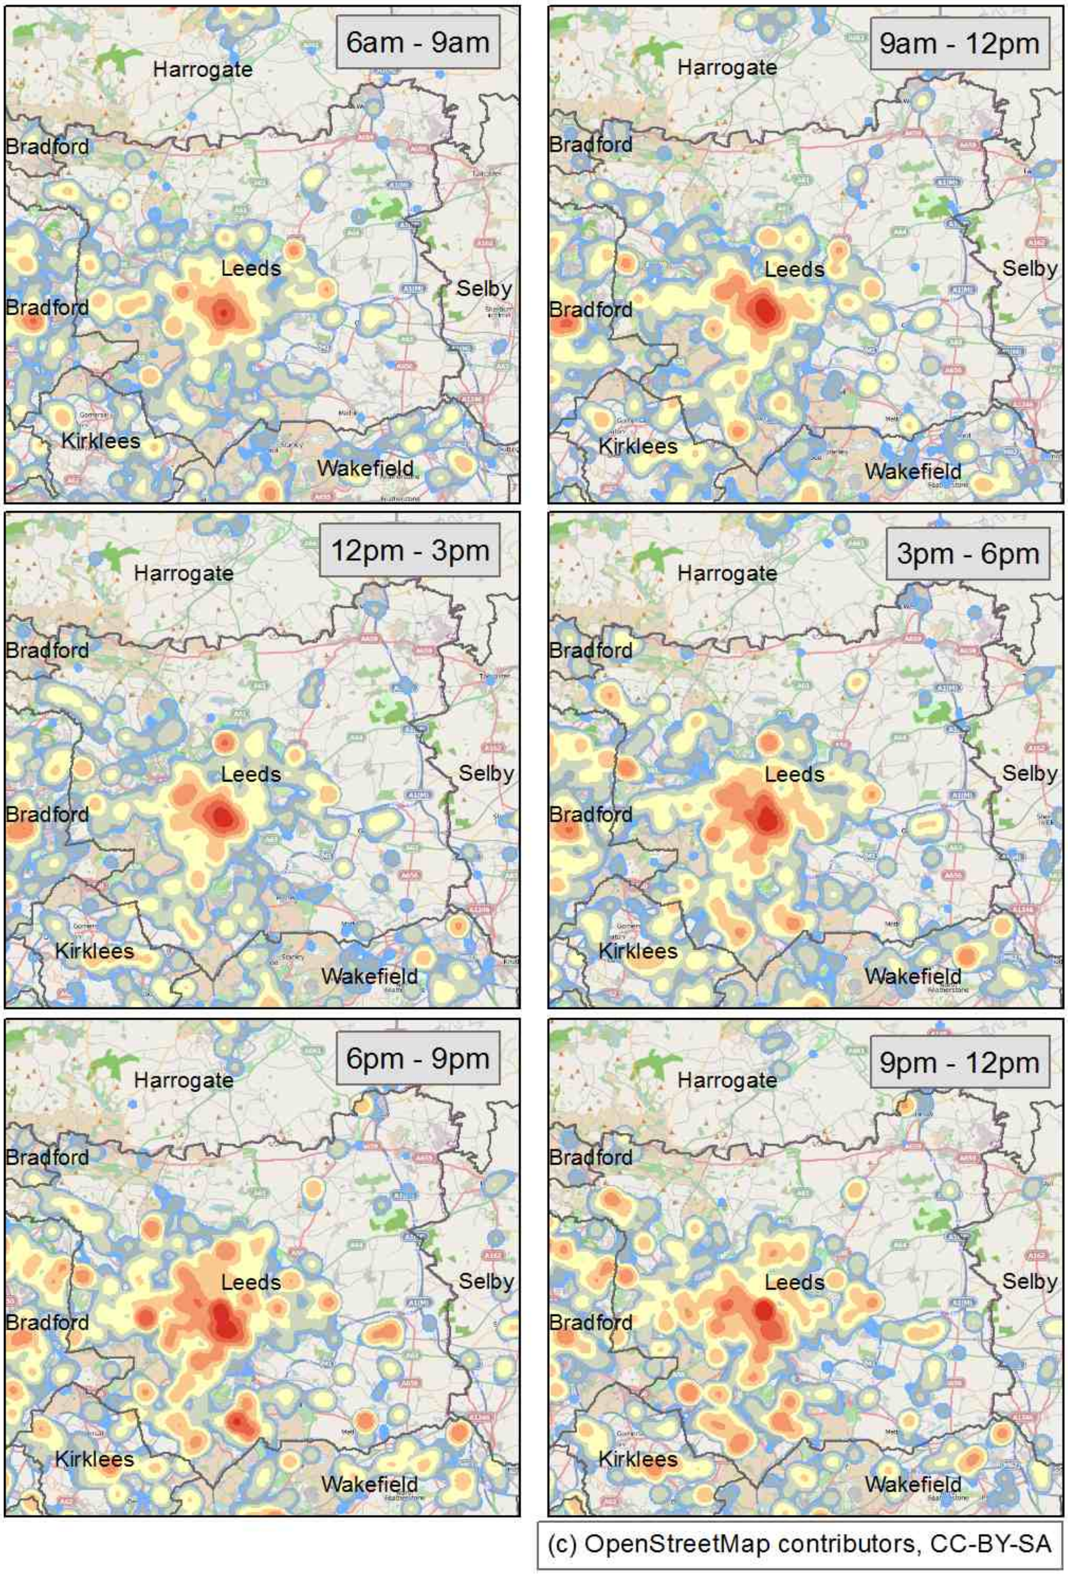

Figure 6: The density of all tweets produced on a Tuesday.

for an understanding of the spatio-temporal dynamics inherent in the data. The following section will explore this avenue in more detail.

3 Identifying Behaviour

3.1 Evidence for Theoretical Principles

In the introduction we examined the idea of Hägerstrand's space-time prisms, which in practice are manifest as individual space-time paths. By observing the spatio-temporal locations of individuals' tweets, it is a relatively straightforward matter to monitor the space-time paths of a large number of actors as they move around the urban environment. For example, if we assume that a user's 'home' is the place where they publish the largest number of tweets (Section 3.2 discusses this assumption in more detail) then Figure 7 is able to illustrate the distance that a single user travels away from home over the course of 1½ days. It appears that they remain at home until 07:00 when they travel approximately 2.5km, returning home at 16:50. This diagram is effectively a representation of what time geography would refer to as a 'space-time path' (i.e. an individual section through the space-time prism) and could be used to better understand the constraints under which a particular user must operate. Over time, for example, it could be possible to estimate average speed of travel for different journeys and hence the type of transport that the user employs. Information of this sort would be extremely useful for the calibration of an agent-based simulation.

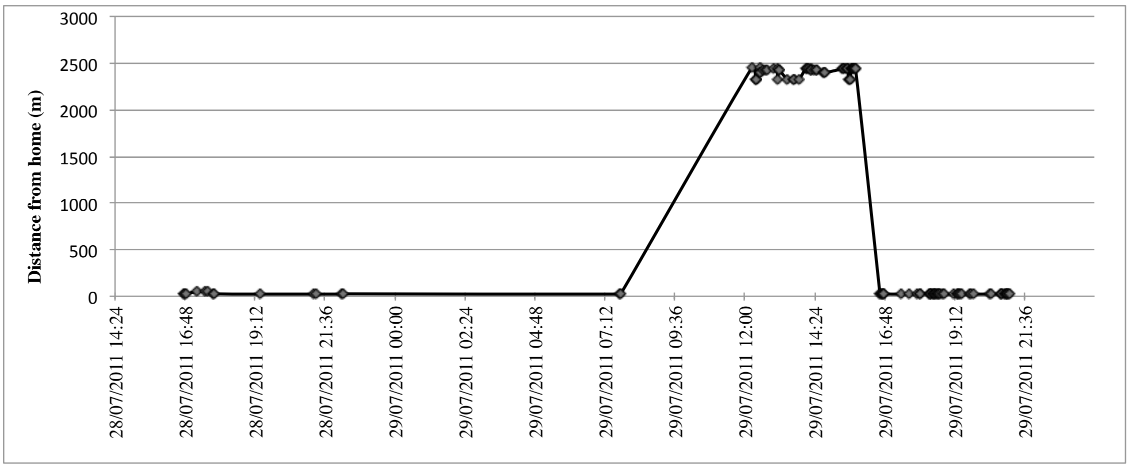

Figure 7: The distance from home over a short period of time for a single user.

Given the richness of the data it is possible to extend this analysis and generate a complete 3D space-time prism constrained to a real-word geography, as illustrated by Figure 8. Visualising the data in this manner shows that it holds great promise for modelling travel behaviour. The two locations are linked by a major road which would most likely be used by the user if they had access to a car. Or, alternatively, by comparing the journey to local bus timetables it would be possible to predict, given the known time interval that the journey occurred in, whether public transport was used instead.

Obtaining data to this level of detail is novel for research that does not recruit human participants; traditionally a researcher would have to equip a group of participants with a location-tracking device and observe their daily behaviour (e.g. Wiehe et al., 2008). All the data used here, however, are publicly available. Hence the use of this type of data in research has the capacity to revolutionise the field of social science (Savage and Burrows, 2007), moving away from expensive, small scale, in depth studies to a methodology that

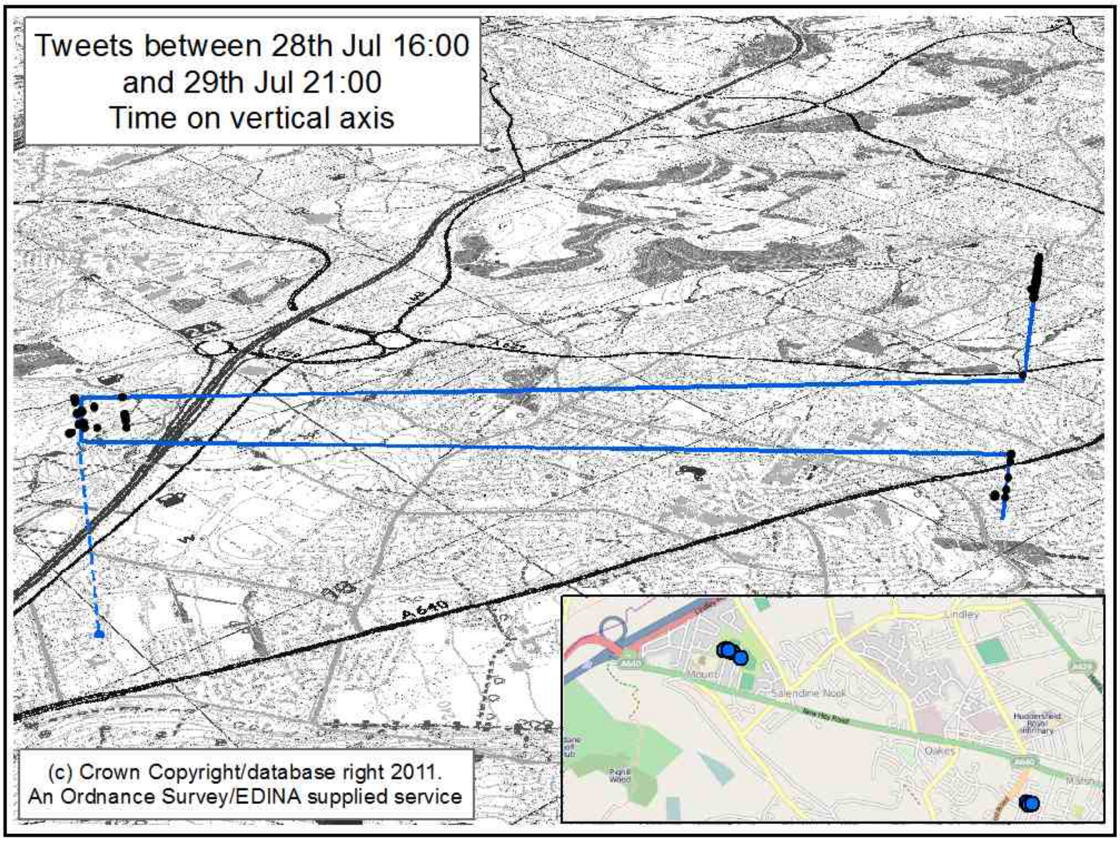

Figure 8: A space-time graph with real geography for a chosen user over a period of time.

embraces the new forms of individual-level data that are becoming readily available. It is important to note that this course of action inevitably has substantial ethical consequences, which will be discussed further in Section 5.

3.2 Classifying Behaviour

The main aim of this on going research is to identify daily behavioural patterns and use these to calibrate an agent-based model. The analysis in the previous section illustrates that the data is of a high enough quality to elucidate behavioural patterns that correspond with theoretical understanding of spatio-temporal behaviour, but even though the movements themselves appear to be accurately defined (both in time and space) it is not so clear what actions should be attached to the movements. For example, in relation to Figure 8 what inferences can be made about what the individual is actually doing at each individual reference point in space and time?

To this end, it is possible to analyse the text of individual tweets in an attempt to identify the action that they were currently undertaking when the tweet was created. Figure 9 plots the locations of all tweets for another user as well as those tweets whose text mentions the words ‘home’, ‘work’ or ‘shop’. These words were used because they correspond with three of the five key behaviours for the agent-based model as outlined in Section 1 (and given the difficulty of the task, further investigation of ‘education’ and ‘leisure’ is postponed for later work).

The vast majority of tweets that mention ‘home’ occur in a very small area, which is also the area of the highest tweet density for the user (indicated on the map). There are also a large number of tweets along the main road between their home and the centre of Leeds so a fair assumption might be that the user works in the city centre. Both of these

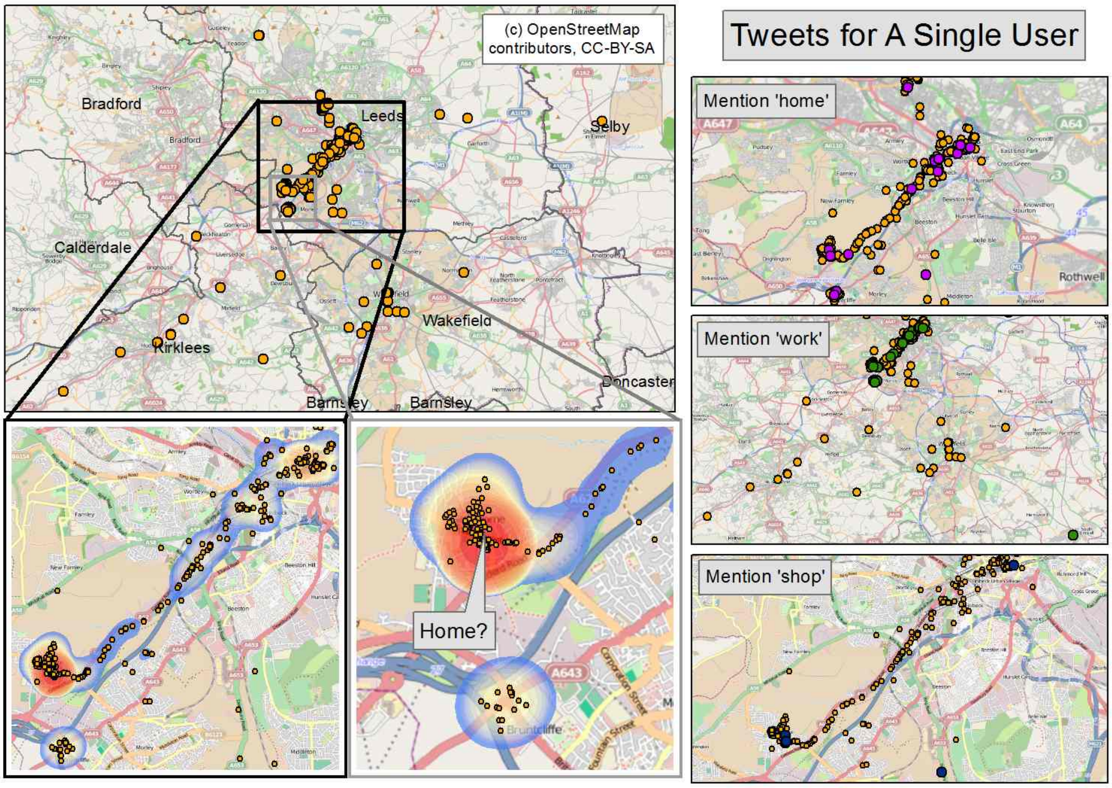

Figure 9: Attempting to classify the behaviour of a single user based on a keyword search of their tweet contents.

assumptions could be investigated further by examining the times that tweets were sent and verifying whether or not the times correspond to those that we would expect a person to be at home or at work.

It should be noted that, as Section 2.3 pointed out, most tweet data originate from non-residential areas. This casts doubt on using the data for behavioural modelling because it might be impossible to establish the main anchor point for a person's travel behaviour: their home. However, it is likely that for the small subset of users who generate enough data for us to understand their travel behaviour (only 6% of users produce more than 100 tweets) most of their tweets will not originate from non-residential areas. This can be verified by Table 2 which provides the ratio index for the number of tweets in different communities for the 6% of users who generate more than 100 tweets. With a comparison to Table 1, it is clear that the index of tweets in 'city living' communities has reduced from 310 (more than three times as many tweets) to 250. This is encouraging as it suggests that it will be possible to estimate the home locations for the most important users (those with sufficient data to analyse). Hence a safe assumption, at this stage, is that a person's home location is that which has the greatest density of their tweets. Future work will, of course, aim to improve this assumption through the use of text mining or other methods (e.g. analysing the surrounding environment).

In terms of identifying a persons' actions, a naïve keyword search such as that illustrated in Figure 3 presents a number of problems. The most obvious drawback is that the 'home' location is the most common place for all tweets, even those that have been identified by 'shop' and 'work' key words. Hence simply searching for a key word is insufficient for establishing the action that the person was undertaking when they published the tweet. By analysing individual tweets it becomes obvious why this is the case. Table 3

Table 2: The number of tweets generated by 'prolific' users (those 100 or more tweets in total) originating from different OAC groups.

	Counts		Proportions		
OAC Group	Tweets	Groups	Tweets	Groups	Index
2 - City Living	18775	211	0.216	0.087	249.8
7 - Multicultural	13703	273	0.158	0.112	140.9
3 - Countryside	2240	62	0.026	0.025	101.4
5 - Constrained by Circumstances	14672	443	0.169	0.182	93.0
6 - Typical Traits	18810	571	0.217	0.234	92.5
4 - Prospering Suburbs	12023	530	0.138	0.217	63.7
1 - Blue Collar	6651	349	0.077	0.143	53.5
Grand Total	86874	2439	1	1	

provides some examples of tweets that include one of the chosen keywords but clearly cannot be classified by that activity. The names of other twitter users have been replaced by '[*user*]' to preserve anonymity.

Table 3: Examples of the text captured with a keyword search.

Keyword	Text
Work	"Cheese on toast for breakfast before I venture into the rain to go to work." "Does anyone fancy going to work for me? Don't want to get up" "Beer does work. Having blurted out to a girl that she is the spit of [*user*] I've bagged her number. And she likes rugby!"
Home	"Pizza ordered ready for ones arrival home" "[user] that's not a nice thing to say about your home. It did look grim at that time of night though" "Ah the good old sight of The White Rose shopping centre. Means I'm nearly home" "Home time at long last"
Shop	"Ah the good old sight of The White Rose shopping centre. Means I'm nearly home" "[user] haha well played. Was she taking him shopping?" "Are the shops in town open today? I really need a hammer, picture hooks, Stanley knife and spirit level. DIY bonanza day"

Although the examples in Table 3 illustrate where the keyword search fails to classify behaviour, it also suggests that most text could actually be classified by a more intelligent routine. Phrases such as "Don't want to get up", and "I'm nearly home" are indicative of clearly defined activities – preparing to leave home or arrive home respectively. Also, through a manual inspection of tweets it was possible to identify when the user was spending time in a local pub and when they were shopping in nearby town centres. Hence an advanced algorithm should be able to successfully classify a large number of the data points.

4 Discussion: Towards an Agent-Based Model of Urban Dynamics

We have argued above that individual agents could prove to be the most effective way of representing entities and their behaviour patterns, especially when we are concerned with actions which are discrete in space and time – such as visiting a bar or restaurant. In the paper we have classified such behaviours into some broad types and aim to extend them to include all of the following: *at home*, *at work*, *shopping*, *leisure* and *education*. Note that of course much more detail could be introduced into many of these categories, so for example the behaviour patterns associated with a regular weekly shop for groceries could be very different to an occasional excursion for shoes or clothes. Activities undertaken at home could take a wide range of forms (gardening, sleeping, reading, cooking, and so on) – and we begin to speculate why these might be of particular interest to the analyst or policy-maker below. It is also worth noting however that the value of agent-based models in simulating the 'slow dynamics' associated with more strategic decisions (e.g. moving home – see Jordan et al. (2011)) have also been championed in certain quarters.

We have begun to develop illustrations of real activity patterns which have a distinct resonance with the theoretical frameworks of behavioural geography and time geography. Through detailed analysis of individual entries in a corpus of geo-located twitter messages it has been shown that constrained intra-urban movement patterns, e.g. from home to work to shop and home again, and so on, can be detected and illustrated. In order to extend this work in the direction of a more general framework for agent-based modelling of urban dynamics, at least two further steps are required:

- First, it is not feasible to attempt the detection of behaviour patterns in a data set with hundreds of thousands of entries through manual inspection. To take this further step requires some kind of automated mining procedure which is able to infer, with reasonable confidence, what a person is doing at a particular point in time; or which common sets of entries together imply that a subject is 'at home' or 'at work'. There are some parallels in the text mining software which has been under development for some time. For example, the National Centre for Text Mining (NaCTeM) has undertaken projects relating to x and y. The use of algorithms to detect sentiment in financial markets is surely a subject of interest, and not just in popular fiction (Harris, 2011). Indeed the Guardian Newspaper has commissioned and published a range of analyses of social network data relating to the London riots of 2011 (Lewis, 2011) with a view to a better understanding of the causes, dynamics and chronology of these events. For all this, one suspects that the software tools available are still far from maturity and that further work is needed to construct a set of methods which can extract maximum value from information which is specifically geo-temporal in its character – again the recurrence of activity at particular locations such as home, work, or a preferred leisure spot; or in the context of urban unrest then perhaps the gathering and dispersal of groups that the notion of a riot suggests.

- Secondly, the kind of data which comes from crowds is not the subject of regular stratified samples, and it is not controlled for quality, completeness and accuracy. We have already seen in our sample dataset that there is an enormous skew in which a very large portion of the data is generated by a small number of prolific individuals. In order to smooth out the significant biases in these data we require at the least some powerful reweighting tools. The idea of geodemographics may have some currency here. Whenever a tweet is geo-located then this location itself could be used as a basis for estimating the character of the individual from whom this message is sourced. Linking the twitter data to some kind of spatial microsimulation model would be a more refined way of attempting something similar. In this case we assume that our target class is not a homogeneous geodemographic type at each location but a distribution of individuals whose profile of socio-economic and demographic characteristics can be estimated. This approach could be made more powerful still if appropriate inferences could be made to qualify the characteristics of the individuals in the twitter data. For example, is it possible to make inferences about age or ethnicity from the names posted by individuals – certainly gender should be identifiable with reasonable confidence. Maybe it is possible to identify family and relationship status from textual analysis, while the use of specific words and phrases could be associated with particular ages or social and income groups.

In order to articulate this model framework more fully, we propose that a population of agents could be initiated through an existing microsimulation model. This population would be provided with a set of attributes which can now be regarded as reasonably standard i.e. age, gender, ethnicity, occupation, household composition, health status – see, for example, Wu et al. (2010) or Malleson and Birkin (2011). Other characteristics can be synthesised if necessary. The dynamics of a spatial microsimulation of this type can be analysed over long time periods using demographic projection methods (Wu et al., 2008) in which the transition between locations and states is monitored in discrete annual time intervals. What we suggest here is that the behaviour of a set of agents could be traced over a much finer timescale, perhaps even dividing each day into individual hourly segments. At each point in time both the location and current activity of the agent would be simulated (where at its most straightforward these activities comprise our standard categories of home, work, shop, leisure etc). From one time period to the next, the transition probabilities are estimated using a function of the form:

$$P(j, b, (t+1)|i, a, t) = f(a, t, t*, y_0) \tag{1}$$

In other words, the probability that an individual at location i, engaged in activity a at time t moves to location j and activity b at time $t+1$ is related to the current activity (a), the time of day (t), the length of time over which the current activity has been conducted ($t*$) and the characteristics of the individual agent (y_0). Through the implementation of the generic textual and spatial analyses which have been described above we would be optimistic about the possibility of establishing a reasonable approximation to these functional relationships.

The idea of data linkage and microsimulation could itself be extended to examine the synthesis not just of traditional sources such as the census and government surveys, but also to other new forms of data. For example, Savage and Burrows (2007) particularly note the rise of commercial data sources, including lifestyle data, as a potentially valuable source. If twitter data could be linked to a microsimulation, then it could also be linked to something like a lifestyle database (Thompson et al., 2012), for example to add extra intelligence about how people actually spend their time (and money) whilst at home, or out and about in the urban system. This kind of information (enjoys reading, gardening etc) could be of value not just to the marketers of retail businesses, but to planners of services such as libraries, civic advice centres, health clinics etc. For example, if the prevalence of smoking is high then maybe there is a greater need for education or intervention (cf Tomintz et al., 2009).

In general, although this paper has mostly focused on behavioural modelling as an intellectual challenge, the potential practical benefits in pursuing this research agenda are substantial. Whilst it may be convenient to assume that urban environments are slow moving and well-structured (and thus, for example, that something like a decennial population census can adequately capture its essential characteristics) in practice the daily ebb and flow of the population is surely critical in understanding its needs, whether for healthcare, education, crime prevention, retailing, or any of the myriad of infrastructure and services which are critical to the lubrication of the urban ecosystem.

5 Conclusion

This paper has explored the contents of new social networking data from Twitter to gauge its applicability to a model of urban dynamics. Although there will be substantial challenges for identifying a person's actions from their individual tweets, the problems are not insurmountable. Much of the data have high-resolution coordinates associated with them which lends them nicely for the calibration and validation of a model. Furthermore, these types of streaming data sources could potentially be used as a means of improving social models *in situ* as they are produced in a similar to that used in other fields such as meteorology (Collins, 2007). Observing social systems over time is a "fundamental problem" (Janssen and Ostrom, 2006) which this type of data might offer a solution to.

Clearly there are considerable ethical issues surrounding the use of these data that must be addressed. The traditional means of obtaining such high quality spatial/behavioural data would be through a formal qualitative interview process in which the subject gave their consent for the research and where any data collected were strictly guarded by data protection / privacy legislation. However, this approach is not feasible with the use of crowd-sourced data; even if it were possible to contact the individuals directly, the sheer quantity of data prohibits this. Also, even though the data is publicly available and users must opt-in to the location services that publish their location it is not clear whether or not they are fully aware of the implications. For example, it is not clear whether the users depicted in Figures 8 and 9 are aware of how much data they have made available to the public. Hence the need to, at this stage, hide user names and draw maps at a resolution that makes it impossible to identify individual houses.

With such an under-developed field the prospects for future work are broad. However, the two most significant areas for immediate exploration are: developing spatial text mining techniques to identify people's actions from their tweet data and working towards a formal means of updating a working agent-based model with streamed social data. Also, by searching for user names in the text of tweets it is possible to construct a social network of Twitter users who communicate via tweets – work along these lines could potentially be extremely interesting.

References

Ballas, D. and G. Clarke (2001, August). The Local Implications of Major Job Transformations in the City: A Spatial Microsimulation Approach. *Geographical Analysis 33*(4), 1–21.

Batty, M. (2005). Agents, cells, and cities: new representational models for simulating multiscale urban dynamics. *Environment and Planning A 37*, 1373–1394.

Bell, G., T. Hey, and A. Szalay (2009). Beyond the data deluge. *Science 323*, 1297–1298.

Birkin, M. and G. Clarke (2011). Spatial Microsimulation: challenges of dealing with demographics at the small area level. In J. Stillwell and M. Clarke (Eds.), *Population Dynamics and Projection Methods: Essays in Honour of Philip Rees*, pp. 193–208. Berlin: Springer.

Birkin, M., A. Turner, and B. Wu (2006). A synthetic demographic model of the UK population: Progress and problems. In *2nd International Conference on e-Social Science*, Manchester.

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences 99*(90003), 7280–7287.

Burgess, E. W. (1925). The Growth of the City. In R. E. Park, E. W. Burgess, and R. D. McKenzie (Eds.), *The City*, pp. 47–62. University of Chicago Press.

Collins, M. (2007). Ensembles and probabilities: a new era in the prediction of climate change. *Philosophical Transactions of the Royal Society A 365*(1857), 1957–1970.

Golledge, R. G., L. A. Brown, and F. Williamson (1972, March). Behavioural approaches in geography: an overview. *Australian Geographer 12*(1), 59–79.

Goodchild, M. (2007). Citizens as sensors: the world of volunteered geography. *GeoJournal 69*, 211–221.

Hägerstrand, T. (1970). What about people in regional science? *Paper in Regional Science 24*, 6–21.

Harris, R. (2011). *The Fear Index.* London: Hutchinson.

Janssen, M. A. and E. Ostrom (2006). Empirically Based, Agent-based models. *Ecology and Society 11*, 1–13.

Jordan, R., M. Birkin, and A. Evans (2011). Agent-based Simulation Modelling of Housing Choice and Urban Regeneration Policy'. In T. Bosse, A. Geller, and C. Jonker (Eds.), *Multi-Agent-Based Simulation XI*, pp. 152–166. Berlin: Springer.

Keith, C. C. and L. J. Gaydos (1998). Loose-coupling a cellular automaton model and GIS: long-term urban growth prediction for San Francisco and Washington/Baltimore. *International Journal of Geographical Information Science 12*, 699–714.

Lewis, P. (2011, September). Reading the Riots study to examine causes and effects of August unrest. *The Guardian.*

Malleson, N. and M. Birkin (2011). Towards victim-oriented crime modelling in a social science e-infrastructure. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 369*(1949), 3353–3371.

Miller, E., J. Douglashunt, J. Abraham, and P. Salvini (2004, January). Microsimulating urban systems. *Computers, Environment and Urban Systems 28*(1-2), 9–44.

Miller, H. J. (2004). Activities in Space and Time. In P. Stopher, K. Button, K. Haynes, and D. Hensher (Eds.), *Handbook of Transport 5: Transport Geography and Spatial Systems*. Pergamon/Elsevier Science.

Parker, D. C., S. M. Manson, M. A. Janssen, M. J. Hoffmann, and P. Deadman (2003). Multi-Agent Systems for the Simulation of Land-Use and Land-Cover Change: A Review. *Annals of the Association of American Geographers 93*(2), 314–337.

Pred, A. (1967). *Behaviour and Location: Foundations for a Geographic Dynamic Location Theory*. Lund Studies in Geography Series B. Lund.

Prothmann, H., F. Rochner, S. Tomforde, J. Branke, C. Müller-Schloer, and H. Schmeck (2008). Organic Control of Traffic Lights. *Lecture Notes in Computer Science 5060*(Chapter 19), 219–233.

Putman, S. (1983). *Integrated urban models Policy analysis of transportation and land use*. London: Pion.

Savage, M. and R. Burrows (2007). The Coming Crisis of Empirical Sociology. *Sociology 41*(5), 885–899.

Thompson, C., G. Clarke, M. Clarke, and J. Stillwell (2012). Modelling the Future Opportunities for Deep Discount Food Retailing in Great Britain. *International Journal of Retail Distribution and Management*.

Tomintz, M. N., G. P. Clarke, and J. E. Rigby (2009). Planning the Location of Stop Smoking Services at the Local Level: A Geographic Analysis. *The Journal of Smoking Cessation 4*(2), 61–73.

Vickers, D. and P. Rees (2007). Creating the UK national statistics 2001 output area classification. *Journal of the Royal Statistical Society A 170*, 379–403.

Wiehe, S. E., A. E. Carroll, G. C. Liu, K. L. Haberkorn, S. C. Hoch, J. S. Wilson, and J. D. Fortenberry (2008). Using GPS-enabled cell phones to track the travel patterns of adolescents. *International Journal of Health Geographics 7*(1), 22.

Wu, B., M. Birkin, and P. Rees (2008). A spatial microsimulation model with student agents. *Computers, Environment and Urban Systems 32*(6), 440–453.

Wu, B. M., M. H. Birkin, and P. H. Rees (2010). A Dynamic MSM With Agent Elements for Spatial Demographic Forecasting. *Social Science Computer Review 29*(1), 145–160.