


This article was downloaded by: [Winchester School of Art]

On: 22 May 2015, At: 09:18

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Annals of the Association of American Geographers

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/raag20>

Developing a Flexible Framework for Spatiotemporal Population Modeling

David Martin^a, Samantha Cockings^a & Samuel Leung^b

^a Geography and Environment, University of Southampton

^b School of Civil Engineering and Surveying, University of Portsmouth

Published online: 13 May 2015.



[Click for updates](#)

To cite this article: David Martin, Samantha Cockings & Samuel Leung (2015): Developing a Flexible Framework for Spatiotemporal Population Modeling, *Annals of the Association of American Geographers*, DOI: [10.1080/00045608.2015.1022089](https://doi.org/10.1080/00045608.2015.1022089)

To link to this article: <http://dx.doi.org/10.1080/00045608.2015.1022089>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Versions of published Taylor & Francis and Routledge Open articles and Taylor & Francis and Routledge Open Select articles posted to institutional or subject repositories or any other third-party website are without warranty from Taylor & Francis of any kind, either expressed or implied, including, but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement. Any opinions and views expressed in this article are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor & Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

It is essential that you check the license status of any given Open and Open Select article to confirm conditions of access and use.

Developing a Flexible Framework for Spatiotemporal Population Modeling

David Martin,* Samantha Cockings,* and Samuel Leung†

*Geography and Environment, University of Southampton

†School of Civil Engineering and Surveying, University of Portsmouth

This article proposes a general framework for modeling population distributions in space and time. This is particularly pertinent to a growing range of applications that require spatiotemporal specificity; for example, to inform planning of emergency response to hazards. Following a review of attempts to construct time-specific representations of population, we identify the importance of assembling an underlying data model at the highest resolution in each of the spatial, temporal, and attribute domains. This model can then be interrogated at any required intersection of these domains. We argue that such an approach is necessary to moderate the effects of what we term the *modifiable spatiotemporal unit problem* in which even detailed spatial data might be inadequate to support time-sensitive analyses. We present an initial implementation of the framework for a case study of Southampton, United Kingdom, using bespoke software (SurfaceBuilder247). We demonstrate the generation of spatial population distributions for multiple reference times using currently available data sources. The article concludes by setting out key research areas including the enhancement and validation of spatiotemporal population methods and models. **Key Words:** GIS, *modifiable spatiotemporal unit problem*, *population*, *spatiotemporal*.

本文为模式化人口于时间及空间中的分布，提出一个一般性的架构，而这与逐渐增加中的需要明确时空的应用范围特别有关；例如为灾害紧急回应的规划提供信息。我们首先回顾建立特定时间的人口再现之企图，随后指认在每一个空间、时间与属性领域，以最高的解析度组合基础数据模型的重要性。此一模型，可接着被整合进这些领域的任何必要交叉。我们主张，此一取径，对缓和我们称之为可调整的时空单元问题之效应而言是必要的，在该问题中，即便是详细的空间数据，也可能不足以支持对时间具有灵敏度的分析。我们使用定製的软件（SurfaceBuilder247），呈现将该架构运用于英国南安普敦案例研究的初步实践。我们运用目前可获得的数据来源，展现多重参照时间中的空间人口分布之生产。本文以阐述包含增进并确认时空人口方法与模型的关键研究领域作结。 **关键词：** 地理信息系统，可调整的时空单元问题，人口，时间—空间。

En este artículo se propone un marco general para modelar las distribuciones de la población en el espacio y el tiempo. En particular, lo anterior es pertinente para un creciente ámbito de aplicaciones que requieren especificidad espacio-temporal; para informar, por ejemplo, la planificación de respuestas de emergencia a situaciones de riesgo. En seguimiento a un estudio sobre intentos de construir representaciones de la población centradas en el tiempo, identificamos la importancia de estructurar un modelo de datos subyacentes a la más alta resolución en cada uno de los dominios espaciales, temporales y de atributos. Posteriormente, este modelo puede ser examinado en cualquiera de las intersecciones requeridas de estos dominios. Abogamos por la necesidad de tal enfoque para moderar los efectos de lo que nosotros denominamos *problema de la unidad espacio-temporal modificable*, circunstancia en la que incluso los datos espaciales detallados podrían resultar inadecuados para soportar análisis sensibles al tiempo. Presentamos una implementación inicial del marco propuesto en un estudio de caso de Southampton, Reino Unido, mediante el uso de software desarrollado especialmente para esa tarea (SurfaceBuilder247). Demostramos la generación de distribuciones espaciales de la población para múltiples tiempos de referencia usando fuentes de datos actualmente disponibles. El artículo concluye precisando áreas de investigación claves que incluyen el mejoramiento y validación de métodos y modelos espacio-temporales de la población. **Palabras clave:** SIG, *problema de la unidad espacio-temporal modificable*, *población*, *espacio-temporal*.

Conventional population mapping is based on the representation of static count data within static zone boundaries. Even when such maps are created and presented through the apparently dynamic media of geographical information systems (GIS) or online mapping, these essential characteristics remain unchanged. Population count data are typically derived from census, survey, register (Rhind 1991), or, increasingly, remotely sensed (Deng and Wu 2013) sources and are associated with specific reference times. An important but rarely acknowledged deficiency of census and administrative sources is that they not only relate to a particular date but are primarily based on “nighttime” residential location assumptions; in other words, they represent a notional time when all members of the population are at their residential address. This is explicit in remote sensing approaches focused on the distribution of nighttime light (Sutton et al. 1997). In reality there is no time when the entire population is at residential addresses, although it will be a close approximation to the true distribution in most neighborhoods in the middle of the night. Recognition of very different nighttime and daytime population distributions can be traced back at least as far as Wirth (1938). There is increasing demand for more realistic distributions, however, and not simply nighttime and daytime, but ideally a full representation of population time, to provide statistics and maps relevant to the specific temporal scope of the application. This is reflected in a growing number of studies concerned with, for example, emergency planning, environmental risk assessment, and accessibility modeling. These share a need to assess population exposure at very specific times or over time ranges that closely match the temporal characteristics of a hazard such as a chemical release (McPherson and Brown 2004; McPherson et al. 2006), natural disaster (Aubrecht et al. 2013), or terrorist incident (Ahola et al. 2007) or the distribution of population at the time when a service is in greatest demand (Turnbull et al. 2008).

Much work has been done in time GIS to develop relevant ontologies, data structures, and query types (Peuquet 2001; Yuan 2008; Pultar et al. 2010), but there has been little overall advance with regard to integration of mainstream population data sources or the specific modeling required to estimate time-specific population distributions. Kwan (2004) provides an individual-level view, bringing population trajectories into GIS using modern data sources and drawing

strongly on the seminal work of Hagerstrand (1970), but this has not yet led to general conventions for spatiotemporal representations of entire populations. This article addresses the research challenge of developing and implementing a framework for time-specific population modeling, building particularly on the structure presented by Ahola et al. (2007), itself based on Yuan (1996) and Peuquet (1994). The objective is to set out a conceptual framework and practical approach that can be demonstrated within the contemporary data environment but also offering rich opportunities for further development. Attention is drawn to a challenge that we term the *modifiable spatiotemporal unit problem*, an extension of the familiar modifiable areal unit problem. We illustrate our conceptual framework with a city-scale example from the United Kingdom, but the approach is internationally applicable and of increasing relevance in the face of rapid growth in spatiotemporal data sources.

This article is structured as follows. The next section reviews the literature concerning time-specific population modeling, tracing growing recognition of the need for better spatiotemporal information. In the third section we propose a new integrated framework for spatiotemporal population modeling, setting out concepts and principles. The fourth section presents our initial implementation of this framework, using a purposely developed software tool and an empirical example for Southampton, United Kingdom. The fifth section discusses the limitations and potential of this new approach. The article concludes by identifying key issues and areas for further research.

A Review of Spatiotemporal Population Modeling

Population mapping is a basic input to a wide range of research and policy applications concerned either with the spatial distribution of people or understanding population-related processes. In many cases, the ability to answer spatial queries is more important than cartographic visualization. Numerous application examples require more time-specific population distributions than those presented by conventional population maps. These include emergency planning (McPherson and Brown 2004) and the organization of accessible services and facilities (Turnbull et al. 2008). These types of analysis require population distributions relating to the same time periods as the analytical scenario; for example, the population affected

by a hazardous event or available to use a service. In the absence of such time-specific models, there continues to be widespread and inappropriate use of generic residential population distributions. Attempts to address this challenge have included both more time-specific data collection and more time-specific methods.

With regard to time-specific data collection, some national censuses and administrative systems collect information about individuals' places of work and education and might even incorporate alternative geographical zoning systems for these data (Coombes 2010; Martin, Cockings, and Harfoot 2013). This can provide insights into two important daily spatial redistributions of population but is severely limited by the division of time into only two periods (work or school time and all other times) and the exclusion of all individuals not engaged in those activities. By contrast, Sutton et al. (1997) employed the de facto pattern of light emissions as a proxy for aggregate nighttime population distribution, but this does not offer the finest spatiotemporal resolution, nor is light directly related to the location of human populations. Sutton, Elvidge, and Obremski (2003) compared three methods for estimating ambient population density, which represents an average of the night-time and daytime distributions. Here, we are concerned not only with nighttime, daytime, or ambient distributions but with development of a rich temporal model that allows identification of fine-grained population distributions at any specified time.

We are not the only ones interested in modeling dynamic populations. Emergent approaches could be divided into those that essentially extend mapping technologies within a time-enabled framework, as here, and those that attempt to track population movements through "big data" (Manyika et al. 2011), such as mobile telephony or social media interactions (Ahas et al. 2010; Birkin and Malleson 2013; Stefanidis, Crooks, and Radzikowski 2013). Although such observational data are extremely powerful, complex analysis is required to estimate population characteristics and activity. Intensities of mobile telephone calls, social media posts, and other trackable activities themselves vary by time of day and participation rates of population subgroups, making them challenging as a basis for estimation of a complete dynamic population distribution. A fuller understanding of the potential use of such data in monitoring population dynamics will require novel methods, beyond the scope of this article.

Attempts to produce time-specific methods have mostly been based on the piecemeal reallocation of population counts between locations in a conventional population map or GIS. Schmitt (1956) was concerned with the location and size of bomb shelters in the Cold War era: "One of the most important and difficult problems now facing city planners is the development of accurate, usable techniques for estimating the current daytime population of census tracts in urban areas" (83). The most remarkable aspect of this quotation is that the problem still persists after nearly sixty years. Schmitt's approach was to identify data series that provided potential proxies for daytime population, including the volume of telephone calls originating from each census tract, in combination with a variety of methods for converting these into population estimates. Most complex of these was termed the *component method*, involving division of population into subgroups whose behavior could be modeled and estimated through time and space. Few data were available and the relative merits of possible approaches were largely untestable. In another early study, Foley (1952) focused on surveys of vehicle movements for sixty-three medium and large U.S. cities, from which indexes of weekday population inflow into the central business district were derived. Although the scale of diurnal population change due to vehicle travel could be estimated, the results related only to this single movement and could be produced only for cities with suitable surveys.

More recently, researchers attempting to develop time-specific population distributions have adopted solutions in which a census base is used for nighttime and some combination of census data and relocation of specific population subgroups forms the basis for a daytime model. Sleeter and Wood (2006) used U.S. census data for small areas, transferring working populations out of home areas during the daytime and redistributing these onto workplace locations derived from a business directory. School populations are similarly reallocated. Their underlying spatial model is dasymetric (Eicher and Brewer 2001; Mennis 2009), based on the intersection of land parcels and census areas, but the study area is small. McPherson and Brown (2004) presented static daytime and nighttime models by allocating population to residential and employment locations. This work was developed by McPherson et al. (2006) into a national gridded model for the United States, but again the temporal division is only daytime and nighttime. Importantly, the latter study also attempts to model people in the transportation

system, with particular emphasis on population transfer to hospital following a hypothetical airborne release of a hazardous substance. Their modeling includes explicit estimation of indoor and outdoor populations.

The Landscan USA project (Bhaduri et al. 2007) introduced some very powerful ideas and concepts but has not to date published a generalized modeling framework. The approach depends on specifics of the available data sources, with a strong reliance on high-resolution remotely sensed data. Work for the UK Health and Safety Executive (Smith and Fairburn 2008) consists of a building-level GIS database with population characteristics interpolated from small census zones and only descriptive temporal attributes for most spatial objects rather than a comprehensive spatiotemporal model. Zhang, Sunila, and Virrantaus (2010) offered a more explicit, object-oriented model of building types such as office buildings, old people's homes, hospitals, hotels, and shops with the objective of weighting population within different building types at different times, again to inform emergency planning.

Other recent studies have produced detailed time-specific population models for large facilities such as airports, cruise ship terminals (Jochem et al. 2013), and universities (Charles-Edwards and Bell 2013). The former study employs flight schedules, seating configurations for planes, passenger load factors, and modeled cumulative passenger arrival times to estimate the population landside within an airport terminal at fifteen-minute intervals. Passenger statistics are also used to estimate the daily population within ports. The facilities chosen illustrate distinct temporal patterns: regular, high-frequency activity (airports), versus lower frequency, seasonal trends (cruise terminals). Similarly, in Charles-Edwards and Bell (2013), routinely collected public transport passenger statistics and vehicle counts are combined with survey data to produce estimates of service populations, taking account of time of day, day of week, and university term dates. These models incorporate a far more sophisticated view of time than the conventional daytime–nighttime classification, but geographical representation is limited to a single, well-defined, site. To date, these intensive local studies have not been fully integrated with regional or national models, although Jochem et al. (2013) forms part of the research contributing to the Landscan USA project (Bhaduri et al. 2007), discussed earlier.

Ahola et al. (2007) shared the widespread interest in modeling time-specific populations for emergency planning, specifically for fire and rescue services in

Helsinki, Finland. Their work is notable for use of a sophisticated space–time model of population, based on Yuan's (1996) three-domain model. Spatial, attribute, and temporal information are treated as three separate, linked domains to describe the spatial behavior of a population. Spatial locations relate to streets and buildings, the attributes of which are defined in terms of usage by different population subgroups associated with different times. They modeled fourteen different time periods such as “week morning” and “Sat. evening” (Ahola et al. 2007, 946) based on temporal variations in the available data. Their population was divided into ten subgroups such as students and children. Most of the values used to allocate populations to activities and times are based on expert knowledge rather than empirical data.

Despite the wide variety of geographical referencing found on input data sets, several of these studies estimate time-specific populations for cells in a regular geographical grid. The advantage of the grid over irregular areal units is stability over time, whereas units devised for purposes such as census enumeration tend to be periodically revised and also to relate strongly to population distributions at specific times, such as the implicit connection between census geographies and nighttime residential population distributions. Further, gridded population models facilitate integration with the results of modeled environmental data (Fielding 2007).

In summary, we have identified a range of studies concerned with producing more time-sensitive spatial population distributions. These adopt a range of spatial resolutions but are typically comparable with the smallest geographical units used for contemporary static population mapping. With regard to time, the principal division is simply into daytime and nighttime, although Ahola et al. (2007) achieved a considerably more sophisticated fourteen-way classification. Nevertheless, the use of a discrete representation of time inherently limits the flexibility of the modeled data. Spatial locations display wide variation in temporal behaviors, such as slightly different lengths of working day, holiday patterns, and seasonal variations, many of which are cyclical in nature. Any specific spatiotemporal query needs to be able to interrogate the intersection of all these temporal patterns and not be constrained by a small number of predetermined classifications of time, even if the spatial resolution is high. At present, a building-level model of a city may be assigned only generalized daytime and nighttime activity patterns.

This is analogous to a concept already very familiar to geographers. The modifiable areal unit problem (Openshaw 1984) refers to the sensitivity of analysis to imposed aggregation units. The temporal equivalent has been termed the modifiable temporal unit problem (Çöltekin et al. 2011; de Jong and de Bruin 2012), in which an imposed aggregation of time into discrete categories (whether two-way or fourteen-way) has the potential to distort and restrict subsequent analysis. We suggest that any serious attempt to build spatiotemporal population distributions needs to engage with the combination of these two phenomena, something that might best be termed the modifiable spatiotemporal unit problem, a phrase coined by Jacquez (2011) but which has not been adopted in the geographical literature. Analysis that is spatially detailed but temporally coarse might be just as likely to impair analysis as that which is temporally detailed but spatially coarse. The modifiable areal unit problem cannot be “solved,” but resilience is achieved when data are referenced to the smallest possible areal units, allowing purpose-specific aggregations (Openshaw 1984). By extension, what is needed is an equivalent model for spatiotemporal data, permitting the finest possible spatial and temporal divisions to be retained and aggregated in ways best suited to specific analyses. Current attempts to extend time-blind population GIS applications into only daytime and nighttime models fall far short of this ideal.

Yuan (2008) observed that GIS technology is particularly lacking in the ability to handle spatiotemporal data. It is not the objective of this research to develop new approaches to time GIS in general, but it is likely that the relative weakness of mainstream GIS in this respect has inhibited development of more sophisticated time-specific population models. The GIS data structure and query architecture proposed by Pultar et al. (2010) are very much in sympathy with the approach proposed here.

A Proposed Framework for Spatiotemporal Population Modeling

Building on the approaches already reviewed, which have offered largely bespoke methods for the generation of time-specific population maps, we propose a general framework for spatiotemporal population estimation. Our intention in specifying such a framework is to set out the essential concepts and requirements, which could subsequently be implemented using a

variety of algorithms and data sources. We offer one such implementation in the following section. Like Ahola et al. (2007) and Pultar et al. (2010), we draw a distinction between the data model and its derived outputs. Importantly, we argue that it is necessary to employ a fine-grained representation not only of space but also of time, to address the modifiable spatiotemporal unit problem and achieve the flexibility to undertake sophisticated spatiotemporal analyses. We propose an essentially dasymetric (Wright 1936; Eicher and Brewer 2001; Mennis 2009) and volume preserving (Tobler 1979) approach, in that total population is redistributed across space subject to a series of weights and constraints based on ancillary data. A novel feature of our approach lies in the adjustment of these spatial constraints based on temporal profiles associated with human activity at each location.

Our general framework treats the population present (PP) as the sum of three categories: resident population (PR), nonresident population (PNR), and population in transit (PT). Population could be divided into subgroups (e.g., by age or economic activity) and there is continuous movement of these subgroups over time among the three categories and between locations. In generalized form, $PP_{I,c,t}$ represents the population of subgroup c present at location I at time t :

$$PP_{I,c,t} = \sum_i^I PR_{i,c,t} + \sum_i^I \sum_j^K PNR_{i,c,t} w_{ij,t} - \sum_i^I \sum_j^K PNR_{j,c,t} w_{ji,t} + \sum_j^K \sum_k^K PT_{jk,c,t} v_{jk,I,t}, \quad (1)$$

where I , the unit of analysis, could itself contain multiple locations i . Both I and i could be referenced as points, areas, or grid cells, depending on the scale of analysis. In the implementation presented in this article, I are grid cells, each containing i points. $PR_{i,c,t}$ is the population of subgroup c resident at location i at time t . The sum of population present in I is increased by the sum of the nonresident population at each i (e.g., incoming employees at their places of work), drawn from all other locations j in the entire study area, where $j = 1, \dots, K$, based on the specific interaction weighting $w_{ij,t}$ of location i with respect to location j at time t . Conversely, the population present is decreased by the sum of the nonresident population

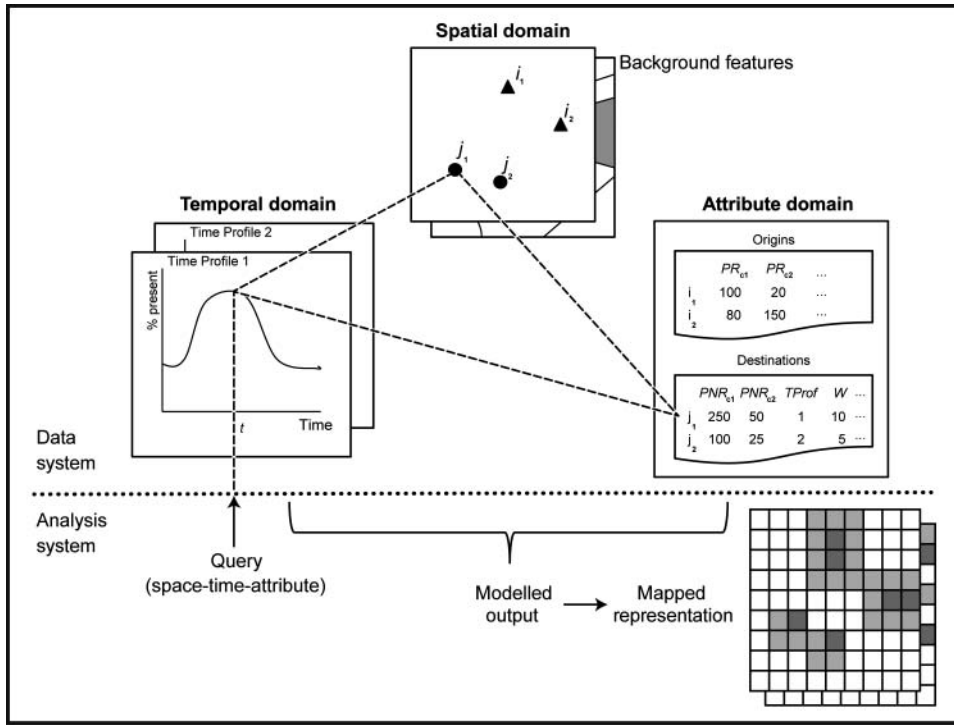


Figure 1. Proposed framework for spatiotemporal population modeling.

(e.g., residents who have gone elsewhere to work) at all other locations j drawn from each i . The final term relates to the population in transit through I at time t between every pair of locations j and k (but not i) based on the specific weighting $v_{jk,I,t}$ of each flow through location I at time t . As every population member can be counted only once, the population in transit does not include those whose journeys begin or end within I ; these are accounted for within the PR and PNR terms. This generalized model shares characteristics with many of the methods reviewed earlier, although there have been various foci of attention; for example, t taking values of only daytime or nighttime, identification of differing subgroups c , or omission of the population in transit PT .

Our proposed framework includes both a data system and an analytical system to support this generalized model. The data system provides a structure for representation of all the spatiotemporal objects of interest. The analytical system provides the means to interrogate these data to answer specific queries. These systems are represented diagrammatically in Figure 1, with the data system above the dotted line and the analysis system below. The numerical values in Figure 1 are purely illustrative. We argue that to advance spatiotemporal population modeling, it is necessary to have a consistent approach to the design

of all the elements in these systems. In this section of the article, we focus on concepts and methods. Data sets, software tools, and implementation are covered in the following section. Similar to Yuan (1996), our data system in Figure 1 includes three domains: spatial containers (top), temporal characteristics (left), and corresponding attributes (right). In practice, these are closely related and interpretation of each is only fully possible with reference to the others; hence the triangular structure.

The first domain of this data system relates to spatial containers of human activity (Hagerstrand 1970; Ahola et al. 2007), which are our basic spatial objects. These are denoted by i, j , and k in Equation 1 and illustrated by i_1, i_2, j_1 , and j_2 in Figure 1. These containers could be represented at a variety of scales, from small aggregations such as census zones or street blocks to individual buildings, and georeferenced by point coordinates or boundaries. Where point coordinates are used, additional information could indicate spatial extents (e.g., the extent of a large site or building or a census area).

Spatial containers could be of two types, which we term origins and destinations. *Origins* (i_1 and i_2 in Figure 1) are a special set of containers, the sum of whose resident populations (PR) represents the entire population to be modeled. Examples would be output

zones from a national census, providing estimates of total population at places of residence. The sum of the population in these containers, plus or minus any known population flows into or out of the study region, is key to the requirement of volume preservation: The total population present (*PP*) may be reallocated between containers during the modeling but can be neither gained nor lost.

We term the second type of population container *destinations* (j_1 and j_2 in Figure 1). These represent additional locations with nonresident populations (*PNR*). Major categories include places of work, education, health care, leisure, and retail activity. Their population capacities reflect the maximum numbers of workers, students, patients, or customers present at any one time. Individual locations could be of either origin i , destination j , or both types, depending on the presence of *PR*, *PNR* or both. The data describing these groups will often come from different sources.

At any point in time, there will be a large number of people who are not present at any fixed location (*PR* or *PNR*) but are in transit between containers (*PT*). The size of this population could be estimated by examining the numbers of people arriving and departing from destination containers in adjacent time periods. The background space between containers is an aspect of the spatial domain not fully addressed by other researchers. A layer of background features (indicated as a second layer of the spatial domain in Figure 1) informs the likely spatial distribution of *PT*. Relevant features include links in the transportation system, through which population moves between origins and destinations, and land use classes that cannot contain population, such as open water or extremely remote and inaccessible regions.

Central to the framework proposed here is the treatment of temporal characteristics, the second domain of the data system. Ahola et al. (2007) noted that temporal understanding of population can only be obtained by collecting details of daily and weekly activities and that this is not a complete record of population movement but, rather, a set of observations and assumptions. Each spatial container may be allocated one or more time profiles, represented by the graphs on the left side of Figure 1, describing the population present as a proportion of total capacity, over time. A key difference from previous approaches is that rather than adopting predefined reference periods (e.g., daytime or week morning), any number of specific and continuous time profiles could be used. These might span a wide range of timescales, from

seasons and term dates through to clock times describing the working day. Thus, a term-time school day could be given a profile involving 95 percent of enrolled students being present from 8:30 a.m. to 4:00 p.m. with phased arrivals and departures at the beginning and end of the day. If sufficiently detailed information is available, different profiles could be used to describe the same location during holiday periods, weekends, or different terms. Individual schools could have similar, but unique, profiles reflecting their different timetables, term dates, or attendance rates. Conversely, a single time profile could be applied to multiple containers. Most activity patterns are cyclical but could be disrupted by special events such as public holidays, for which specific time profiles could be constructed. The objective of retaining this detailed information is to maximize flexibility for reaggregation as required by any (unknown) future analysis.

The third domain of the data system comprises the attributes of the spatial containers. These include population subgroups c that share important activity patterns, indicated by columns PR_{c1} , PR_{c2} and PNR_{c1} , PNR_{c2} in the right attribute tables in Figure 1. For example, children in a given age group might be assumed to participate in a particular level of school education. In some cases, a subgroup could be identified as immobile, in the sense that all their activity is restricted to their origin location. Prisoners are an obvious example, but others could exhibit extremely limited mobility, such as the very elderly. Each destination container has one or more time profiles and associated catchments representing the area from which origin populations could be drawn to participate in activities. These catchments may take a variety of forms and provide the interaction weights $w_{i,j,t}$ in Equation 1, simplified as column W in Figure 1. Additional attributes associated with each container such as its spatial extent or mobility of population subgroups could also be included.

A hierarchical structure applies in all three of the spatial, temporal, and attribute domains, reflecting the principle of indivisibility (Hagerstrand 1975). It is critical for overall volume preservation that the entire population be accounted for in the spatial domain at any time. Ideally, a high resolution should be modeled in all three domains, but the overall framework is able to encompass multiple levels of hierarchical subdivision. Thus, the approach could be implemented whether residential populations are available for census output zones, street blocks, or individual buildings.

Similarly, the students enrolled at a university could be assigned to one single location and time profile or subdivided into different faculties, buildings, and time profiles, but the subdivided components must always sum correctly to the next level in the hierarchy. One advantage of this property is its extensibility: As more detailed spatial, temporal, or attribute data become available, they can be incorporated by subdivision of existing elements.

The data system presented here incorporates all aspects of population distribution in space and time, held at the highest possible resolution, while recognizing that this might still include some aggregation. This resolution is essential to support the diversity of potential spatiotemporal analyses. The system is able to accommodate a great variety of real-world complexity; for example, a school could be occupied by local children engaged in education during the day, by adults from a wider area engaged in sports and social activities during the evening, and unoccupied at night.

The second system in our framework is that which relates to spatiotemporal analysis and appears below the dotted line in Figure 1. Analysis begins with the specification of a spatiotemporal query (at its simplest, a study area comprising units of analysis I and query time, t). Equation 1 summarizes the estimation of population distribution across a unique intersection of spatial, temporal, and attribute domains. This intersection identifies the population subgroups present at each container at the specified time, as illustrated for j_I by the dashed lines in Figure 1. Answering a query is not simply a retrieval from a database but involves evaluation of every relevant element and aggregation to the spatial units of analysis I .

In the implementation illustrated here, population is reallocated from origin to destination containers by simple weighted allocation to nearest destinations or to meet known proportions of population traveling from successive distance bands (recorded in travel-to-work data). This approach accommodates overlapping catchment areas and preserves total population volume but does not take account of complex flows to nonnearest facilities. Time profiles and catchment areas could either be derived from formal sources such as school timetables and official catchment definitions or obtained through administrative or survey data, such as customer numbers and distances traveled to retail outlets. The potential for using more sophisticated spatial interaction models is considered later in the discussion section.

Background weights and outputs must all be calculated for the chosen units of analysis I . Background weights V are calculated from the background feature layer to represent the distribution of population in transit for each unit of analysis I at time t . In the mapped representation of the modeled output in Figure 1 (lower right) and in our implementation here, the analysis units I are grid cells and thus a background weight is calculated for each cell. Depending on the nature of the query, the output could then be explored in a variety of forms, such as mapped representations, data tabulations, or statistical analyses.

Even a relatively simple scenario reveals the complexity of spatiotemporal analysis. An emergency planner might be concerned to estimate the number of workers present in a mixed residential and industrial district on a weekday at lunchtime to assess the impacts of an airborne chemical release, the scenario explored by McPherson et al. (2006). This would be very hard to estimate using conventional population data sources and representations. It requires each of the spatial (business and residential locations i, j, k), temporal (weekday lunchtime t), and attribute (factory worker and resident population PNR_c, PR_c) domains to be interrogated and population to be estimated for this unique space–time–attribute combination. More complex analyses might require accumulation of data over time ranges; for example, to assess population exposure to environmental pollutants over the period of a hazardous event.

Implementation

In the previous section, we proposed a novel flexible framework for spatiotemporal population modeling. In this section we present our initial implementation, using a software tool called SurfaceBuilder247, which we have developed in a .NET environment (Martin 2011). We use for illustration a simple example for the city of Southampton on the south coast of the United Kingdom. The 25 km × 25 km study area is shown in Figure 2, although the input data sources are all available at the national level and it is only the example that is restricted to this study area, for reasons of presentational clarity. The main urban area of Southampton, with a 2006 residential population of 228,700 (Office for National Statistics 2010), is included in its entirety, along with surrounding settlements. The focus here is not on specific details of the area but on practical application of the framework, with the aim

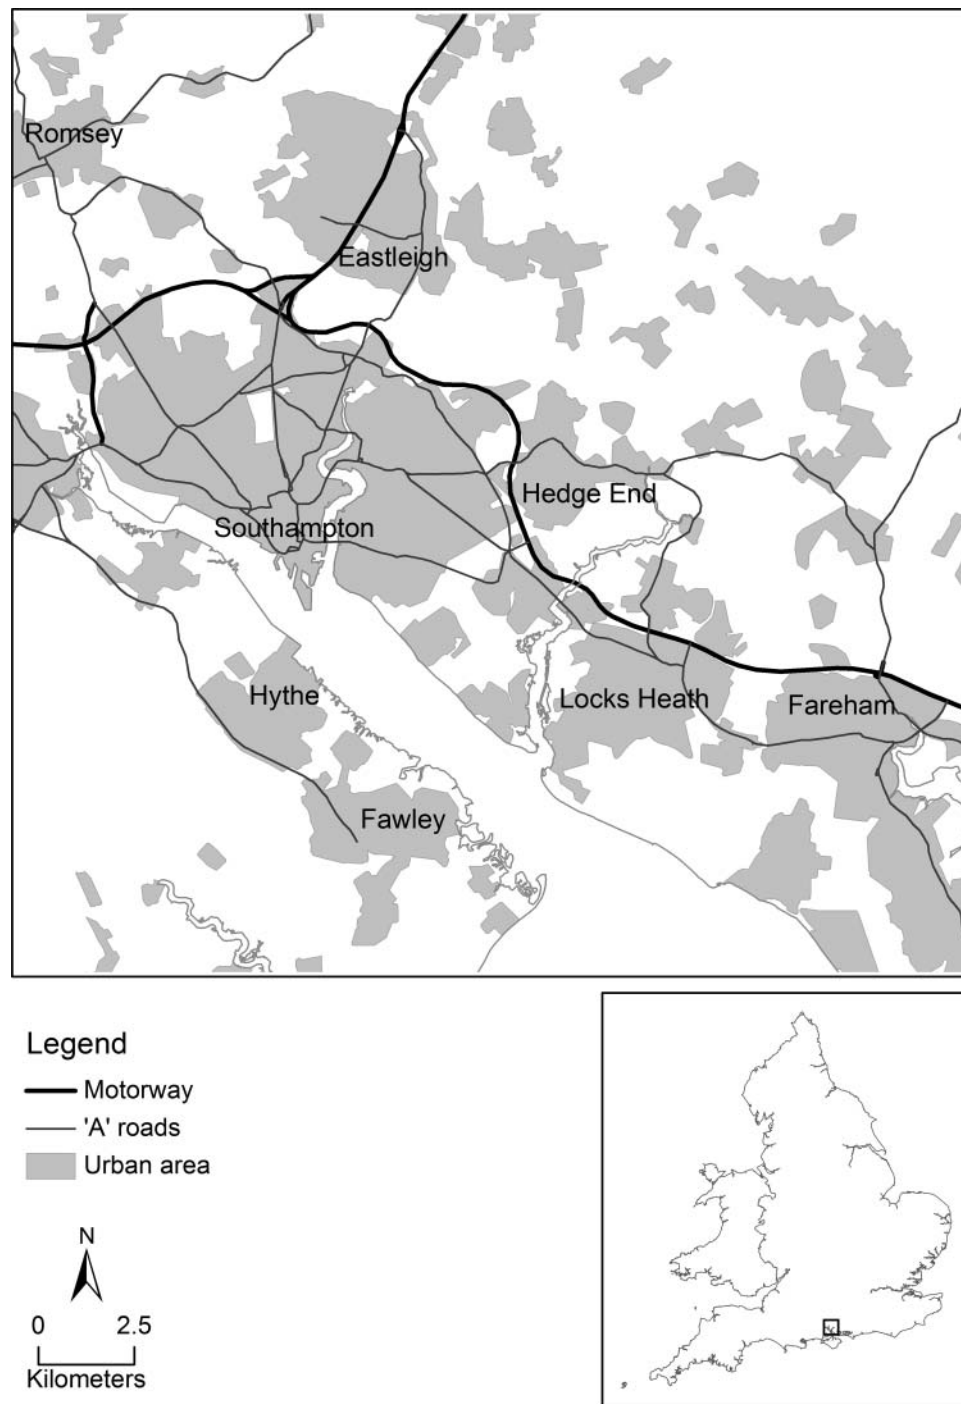


Figure 2. Southampton study area, with inset showing location within England and Wales (contains Ordnance Survey data © Crown copyright and database right [2005, 2013]).

of demonstrating the feasibility of our proposed methods. We do not provide detailed evaluation of the many data-set-specific decisions that would be necessary precursors to a substantive application.

We use a gridded population modeling strategy for the outputs, gaining the advantages of stability over

time and ease of integration with environmental models that have been noted earlier. Computationally, we build directly on previous algorithms that operated in only the spatial and attribute domains, redistributing population from centroid locations onto a grid but without a temporal component (Martin 1989, 1996).

Table 1. Data sets used in Southampton case study

Population subgroups/ containers	Reference date	Geographical reference	Data source URLs (accessed 13 January 2014)	Notes	Time profile
Origins					
Residential locations	2006	2001 Census OA centroids	2001 Census; 2006 MYE http://www.ons.gov.uk/ ; 2006 NSPD http://edina.ac.uk/census/ (registration required)	2006 MYE populations weighted onto 2001 OA centroids using 2006 NSPD delivery point counts	N/A
Immobile populations	2001	2001 Census	2001 Census communal establishment residents http://www.ons.gov.uk/	Prison, care home populations treated as immobile	N/A
Destinations					
Workplaces	2006	2001 OA centroids	2001 Census https://www. nomisweb.co.uk/ 2006 ABI data set http://www.ons. gov.uk/ (subscription required) 2006 NSPD http://edina.ac.uk/census/ QLFS http://www.ukdataser vice.ac.uk	2006 ABI workplace population estimates, Lower Super Output Area level, weighted onto 2001 OA centroids using 2006 NSPD business postcodes; OA- level census travel to work by distance bands	Working patterns by standard industry classification groups derived from QLFS
Educational establishments					
Schools	2005	Grid reference of establishment	2005 EduBase data http:// www.education.gov.uk/edu base/ Neighbourhood Statistics locations http:// www.neighbourhood.stat istics.gov.uk/	Enrolled student counts for school and college, September 2004	Terms and hours estimated clerically from education authority Web sites
Universities	2006	Postcode grid reference of main site	2006 Higher Education Statistics Agency data http://www.hesa.ac.uk/	Enrolled student counts for 2006–2007	Terms and hours based on University of Southampton calendar and travel survey
Hospitals	2006	Postcode grid reference of hospital site	2006 Hospital Episode Statistics http://www.hscic. gov.uk/hes	Inpatient, outpatient, and A&E patient numbers for provider/trust weighted to hospital site by number of beds	Inpatients allocated by length of stay; outpatients by working day; A&E over 24 hours
Background Roads	2006	Road network	Ordnance Survey Meridian https://www.ordnancesur vey.co.uk/opendatadown load/products.html DfT AADF by area, road, and vehicle type for 17 time periods http://www.dft.gov. uk/traffic-counts/	AADFs interpolated onto roads by vehicle, road, and area type from National Transport Model	17 DfT traffic flow time periods covering 7-day week
Land surface	1991	Mean high water	1991 Census MHW polygon http://edina.ac.uk/census/	MHW not available for 2001 Census; 1991 coastline used as mask	N/A

Note: OA = output area, MYE = midyear estimates; NSPD = National Statistics Postcode Directory; ABI = Annual Business Inquiry; QLFS = Quarterly Labour Force Survey; A&E = accident and emergency; DfT = Department for Transportation; AADF = annual average daily flows; MHW = mean high water.

Our background weights are therefore calculated as raster GIS layers, with grid cells I being the spatial units of analysis.

The data sources used in the case study are listed in Table 1. For illustration, we here use a model based only on residential, workplace, educational, and health care containers, with a reference year of 2006. This year represents the midpoint between censuses and is one of the first years for which sufficient data regarding these activities became available. The availability of national data has developed rapidly since 2006, but at the time of writing, publication of the relevant 2011 census outputs was not complete. Our principal data sources are the 2001 census, 2006 mid-year estimates (MYE), Annual Business Inquiry (ABI), Edubase and Higher Education Statistics Agency (HESA) educational institution data, Hospital Episode Statistics (HES), National Statistics Postcode Directory (NSPD), with additional data from Department for Transport (DfT) and Ordnance Survey for background mapping and Quarterly Labour Force Survey (QLFS) for industry-specific time profiles. Web sites and notes for each data source are provided in Table 1: Some are openly downloadable, whereas others require registration or subscription. Each requires data linkage and estimation prior to use in our spatiotemporal modeling. Original counts from the input sources are not recoverable by analysis of the output layers.

The data system presented in the previous section is implemented as a library of .csv format data files following a standard structure, described in detail in Martin (2011). The principal input file type is a list of spatial containers, with one record per container. The same basic data are required for all containers, albeit with some attributes being specific to origins or destinations. The core structure is illustrated schematically in Figure 1 and includes spatial location and population counts for one or more population subgroups. For origins, these counts include the source populations for the model, whereas for destinations they are population capacities. Additional attributes unique to origins include information about the mobility of each subgroup. Additional attributes unique to destinations include reference to relevant time profiles, the spatial extent of the destination itself, and its catchment area. Each file includes meta-data on sources and format. Default values are provided for each variable (e.g., age-specific mobility rates), which can be replaced by container-specific values, if available.

In this example, population origins are 2001 census output areas (OAs). These are the smallest zones for which tabular census results are published, having a mean of 300 usual residents. 2006 MYE data for larger units known as Lower Layer Super Output Areas (LSOA), with mean of 1,500 usual residents, have been allocated to OAs weighted by address counts drawn from the 2006 NSPD to produce OA population estimates for 2006. Each OA is represented by a population-weighted centroid and seven age and activity groups (preschool, three divisions of school age, working age divided into students and nonstudents, retirement age). The example uses a term-time population definition when locating students with different home and term-time addresses.

We exemplify the destination files with reference to education containers. Locations for all schools and universities have been obtained from national data sets (Table 1). Most are available as unit postcodes (the smallest units in the UK postal system), which are readily georeferenced to points using standard lookup tables. Each container has population capacities (numbers of students enrolled) in one or more age groups, reflecting the type of institution. Default values of site sizes and catchment radii have been estimated based on typical values for each type of institution. In this case, simple circular catchment areas have been defined in terms of proportions of the student population expected to travel over specified distance bands (discussed, e.g., by Martin and Atkinson 2001), although a range of more detailed catchment area descriptions could be incorporated where local data are available. Comparable procedures have been followed for employment and health locations using the data sources shown in Table 1. Exact reference dates vary, but 2006 midyear values have been used wherever possible.

Weights (V) for the background layer are here prepared as simple GIS raster layers matching the intended analysis resolution (I). For the Southampton study area, these are based on locations of roads from the OS OpenData Meridian layer and average annual daily flow (AADF) traffic data from the DfT National Transport Model (NTM; Department for Transport 2009). Coastal and open water areas are zero-weighted and all other areas assigned estimated counts of people traveling, using the DfT data on traffic flows by vehicle occupancy and vehicle, road, and area type, thus estimating the spatial distribution of population in transit. This example takes no account of other modes of transport.

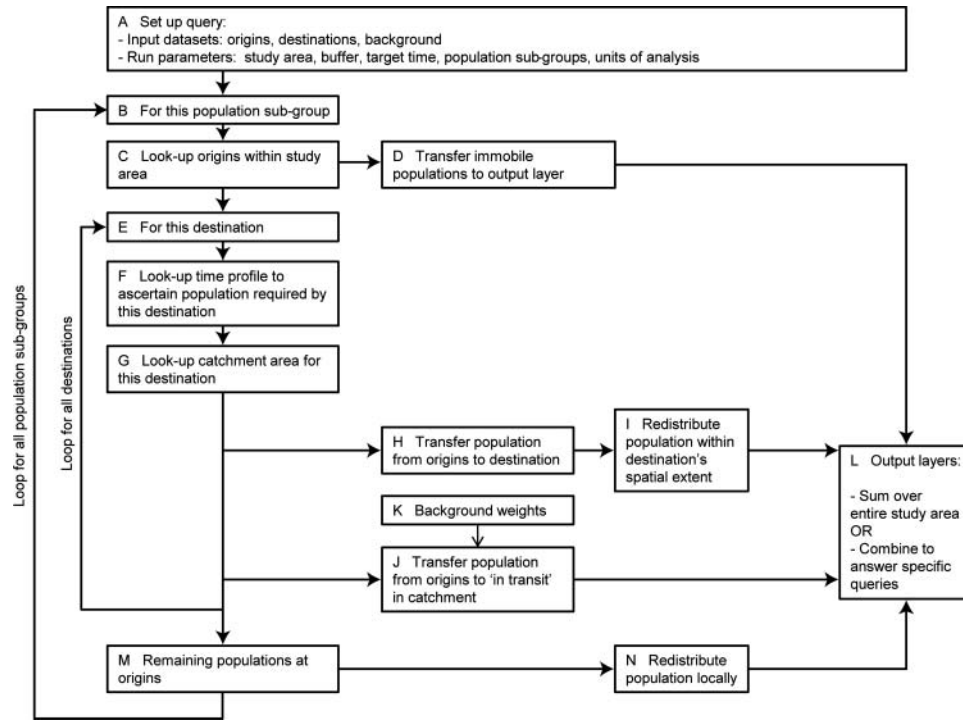


Figure 3. Processing sequence implemented in SurfaceBuilder247 program.

Time profiles are provided in separate files, listing the proportions of population present and in transit associated with different destination types at each time of day. Time intervals are not predetermined and here we use fifteen-minute intervals. For each of primary, secondary, and university working days, profiles are derived from published local education authority and survey data. For employment, these are based on analysis of the QLFS, which includes questions about hours worked and is coded by Standard Industrial Classification (SIC). We estimate time profiles for each SIC code reflecting differing patterns of work and assign these to places of employment. Population in transit is derived from the population arriving and departing each destination and the time taken to travel within its catchment area. Our time profiles are illustrative rather than definitive and could be replaced with detailed local information wherever it is available.

The SurfaceBuilder247 software executes a spatio-temporal query on these files, following the processing sequence illustrated in Figure 3. In the following explanation, letters refer to the elements in Figure 3. A run begins with the specification of all of the parameters necessary to describe the query, input data library, modeling, and outputs (A). Each identified population subgroup is processed separately in a

complete run (B). Processing proceeds for the user-specified study area plus a geographical buffer of a width chosen to capture all local population movements into or out of the study area. A first pass through all origin data sets (C) identifies any immobile populations (D) and transfers these directly to an output layer (E). The remaining population in the origin containers is available for potential redistribution by the model. The principal processing loop visits each destination container in turn (F) and interrogates its associated time profile (G). The time profile indicates the proportion of its capacity population that is expected to be present at the query time, both at the destination and in transit within its catchment area (H). The destination population is handled in two stages. First, the population present at the destination is transferred from origins within the specified catchment area (I). This population is allocated within the spatial extent of the destination (J). Second, the population in transit is allocated across the catchment area (K) in proportion to the weights in the background layer (L). As the populations associated with each destination are redistributed, they are accumulated in an output grid layer (E). Once all destinations have been processed, any population remaining at the origin locations (M) is locally redistributed around the input locations using the original

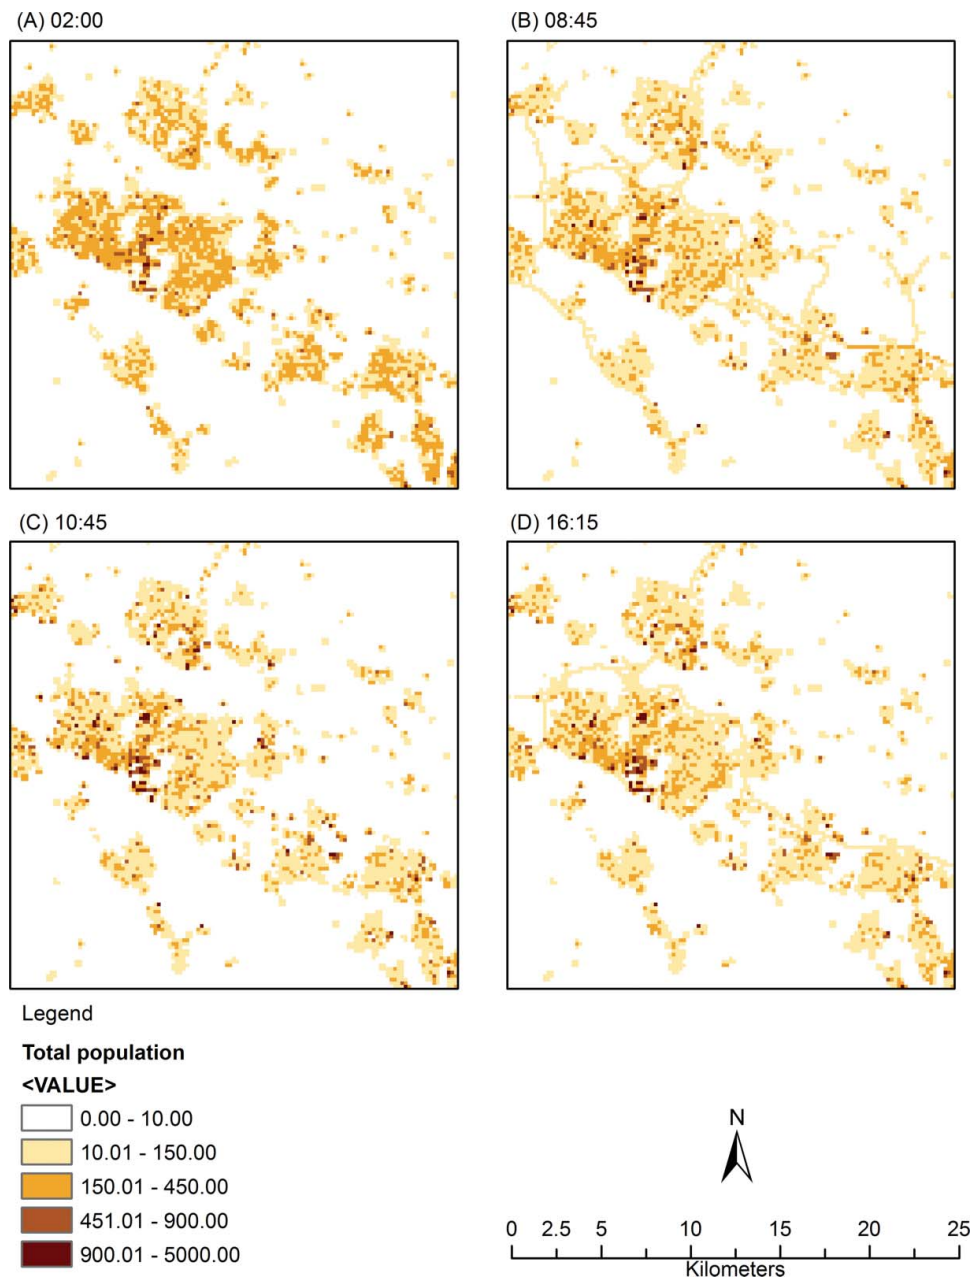


Figure 4. Modeled outputs for Southampton case study for a weekday during school and university term time at (A) 2:00 a.m., (B) 8:45 a.m., (C) 10:45 a.m., and (D) 4:15 p.m.

SurfaceBuilder algorithm (N) and it, too, is transferred to a further output layer (E). At the end of the process, the total population represented by the input origins has been transferred into a series of output layers (E) that sum to the original total. Outputs for the separate subgroups and layers can be combined in a variety of ways to answer specific analytical queries such as the count of all persons at work or count of children at home at a given time.

The preceding processing sequence is illustrated briefly in relation to the Southampton case study.

Processing is undertaken for each of the seven subgroups, for the study area plus a 25-km buffer area, reflecting typical local commuting distances, for different times of day on 8 March 2006, a Wednesday during school and university term time. Raster processing is based on a 200 m \times 200-m grid, reflecting the spatial resolution of the source data sets. Origin populations are drawn from census OA centroid locations. Any immobile populations recorded in census OAs are first transferred to an output layer. Each of the destinations (workplaces, education sites, hospitals) is

then considered and their time profiles are compared to the model target time. Some destinations will have zero population at some times (e.g., schools in the early morning and at weekends), whereas others will have varying but nonzero populations at all times (e.g., workplaces in some industries, hospitals). Each destination is associated with a catchment area and nonresident populations are subtracted from OAs within the catchment areas and added to the relevant *PNR* and *PT* populations, thus decreasing the remaining *PR* count at each origin. In this example, workplaces have catchment areas defined as proportions of workers in successive distance traveled to work bands from 2001 census data. Population at each destination is redistributed across its spatial extent; for example, primary schools being set to 100 m (within a single cell of the output grid). The population in transit is distributed across the catchment area in proportion to the weights in the background layer. The coastal location of the city means that the background layer contains many zero-weighted cells representing the sea. When all destinations have been processed, the remaining origin populations are redistributed from OA centroids into the grid, reflecting local residential population densities. The output layers can be summed to obtain the distribution of the total population at target times, as illustrated in Figure 4, which shows models for a weekday during school and university term time at each of 2:00 a.m., 8:45 a.m., 10:45 a.m., and 4:15 p.m.

The 2:00 a.m. map (Figure 4A) is essentially a residential population distribution, with the densities in the map reflecting the geography of residential neighborhoods. The darkest shading indicates central neighborhoods with the highest density housing. The 8:45 a.m. map (Figure 4B), reflecting the morning travel peak, shows lower population densities across residential areas but with increased concentrations on the road network and central and local business districts. The 10:45 a.m. map (Figure 4C) represents a working day with peak numbers of employees at work and students in education. The highest concentration densities are in business districts but also in prominent sites such as schools, colleges, universities, and hospitals. The final map, 4:15 p.m. (Figure 4D), indicates the temporal asymmetry of a typical weekday, part way through the less peaked evening travel period. By this time most schools and colleges have closed but workplaces remain open and traffic levels are increased, but not yet to the level of 8:45 a.m. This simple sequence serves to illustrate just four

possible outputs from the model, which could be directly interrogated to produce any day of week, time of day, or more complex sequences, within the scope of the source data sets.

Discussion

Our objective has been to introduce a framework for spatiotemporal population modeling and present an initial implementation. We have demonstrated a system capable of estimating population distributions from available data for any desired target time. We here consider a range of issues relating to methods employed, the Southampton case study, data availability, and validation.

Our approach is based on the redistribution of small area aggregate data and we make no attempt to explicitly replicate the behavior of individuals, as would be the case in spatial microsimulation (Birkin and Clarke 2012) or agent-based modeling (Crooks and Wise 2013), or to directly estimate movements between locations, as in spatial interaction modeling (Nakaya et al. 2007). The proposed framework provides the necessary elements for spatiotemporal population modeling, many of which could be achieved using alternative tools and models, thereby improving on our initial implementation. For example, the simplified distance-based catchment areas employed here could be replaced with more sophisticated spatial interaction models to allocate workers to workplaces or children to schools.

All of the spatial containers in our example have been georeferenced and modeled as points with a transformation onto a regular grid for background layers and output. This is not a fundamental requirement of the approach, but for most of the current UK data sets, the input points are the best available spatial locations having associated population characteristics. Smaller objects such as postcodes and buildings lack definitive polygon or population data. This georeferencing strategy will have an impact on some distance-based calculations, such as whether or not a point is included in a specific catchment area, but not on the population counts allocated to each container, which are driven by its capacity and time profile. Our example has used counts and locations from original sources rather than undertaking additional spatial transformations of unknown accuracy. Higher resolution data could readily be incorporated as they become available.

It is possible to directly interrogate the model at the container level, but the point data do not readily

lend themselves to cartographic representation, especially for comparison over time. We have therefore opted to build on previous work for redistribution of population from centroid locations onto a regular grid, reaping specific representational advantages identified in our earlier review. More broadly, we have not attempted to implement an entire spatiotemporal GIS of the type proposed by Pultar et al. (2010), yet we believe that these two lines of development could be fruitfully combined.

The Southampton example does not cover all population activity and is intended to demonstrate the feasibility of the approach using selected existing data sources. Data sources and activity categories are features of the application example rather than limitations imposed by the modeling framework. Notably, the case study excludes leisure and retail customers who are major contributors to urban population movement (although the balance between categories will itself vary with time and day). It also does not include flows to or from neighboring regions. It would be possible to add further activities or to increase resolution in any of the spatial, temporal, or attribute domains, as has been discussed earlier. Each enhancement would further improve the allocation of population into more clearly defined locations at a given time, thereby improving the accuracy of the distributions seen in Figure 4. It would be possible to repeat these models for more recent years, but a limitation remains that reliance on published secondary sources will always introduce time lags of up to one to two years between the most recently available data and the present.

A major consideration for any implementation of our framework is the preparation of the data library. Data availability remains an important issue, but there is currently enormous growth taking place in relevant data sources internationally; for example, through Open Data initiatives such as data.gov.uk and data.gov (Alani et al. 2007). This is making possible the assembly of a coherent population model at a spatial resolution comparable to the smallest areas for which census data are published. These data are increasingly freely available, often produced from sources such as administrative and transactional records that describe attendance or activity in key locations such as schools and hospitals. Available data are mostly aggregated in the spatial and temporal domains (or both), although administrative systems are increasingly able to provide temporal activity patterns, such as patient arrival times and durations in hospital accident and emergency departments (Health and Social Care Information Centre 2013). It

is still necessary to apply distributional assumptions derived from aggregated sources to individual sites. To date, there has been much less effort applied to the collation and publication of temporal information in common formats compared to the spatial aspects.

At present, the validation of our model outputs is extremely challenging. Fundamentally, this is because we are attempting to estimate time-specific population distributions not directly captured by any other measurement systems. We produce detailed estimates of population subgroups, locations, and times, for most of which no true values are available. One approach to validation would be to directly count the population present at certain times and places. There have been some interesting recent attempts to do this. Greger (2015) used observation of three commercial buildings over a seven-hour period to validate a building-level spatiotemporal model of central Tokyo. Charles-Edwards and Bell (2013) used a range of technologies to implement cordon counts of the population entering and leaving a large university campus over a twelve-hour period. In both cases there were difficulties associated with the separation of residential and nonresidential populations. At the building level, it was necessary to select only commercial buildings that could reasonably be assumed to be unoccupied at the start of the observation period. For the university, it was necessary to use administrative information to estimate the number of residents at the start of the period. Further challenges were presented by the need to identify and simultaneously observe all possible entry and exit routes, greatly restricting the choice of candidate buildings. The university study was made possible by its being bounded on three sides by a river, with only a small number of access points. Greger's (2015) approach provided a detailed validation exercise but focused on only one type of population activity over a very limited range of time and space. Many of the differences observed between the model and validation are attributable to these limitations. Charles-Edwards and Bell (2013) developed a hybrid approach but not an independent validation of modeled data. Both studies acknowledge the substantial cost and resource requirements to cover very restricted time periods. In the context of this study it would not be possible to scale up these observational approaches to assess the entire population present even to a single neighborhood, given our concern with variation over multiple timescales in complex multiuse areas with residential, nonresidential, and in-transit populations and multiple access points.

A second approach to validation would be to use alternative data sources, but we have here proposed a framework, the objective of which is to integrate all of the available sources to produce the best possible estimate of the entire population distribution. Administrative data sources tend to be activity-specific; hence, hospital activity data on patient numbers are a definitive source, which we would want to use in the modeling. Patient numbers in the model cannot be validated by removing the hospital data and seeking to reestimate the missing values from sources relating to workplaces or schools. Nevertheless, it should be noted that the approach presented here is akin to an exact interpolation method (Lam 1983) in which known input counts will be directly replicated in the output model. Thus, whenever an additional data source records the population engaged in an activity at a particular point in time and space, those values will be replicated in the output, improving the overall accuracy of the model.

A third approach to the validation challenge would be to employ data that continuously track populations and monitor activity, such as mobile telephony, shopping center footfall counts, traffic sensors, or georeferenced social media posts. At present, research access to these big data types is limited, but they have the potential to provide proxies for population redistribution over space and time, independent of traditional administrative data sources. These data present additional ethical constraints and calibration challenges (e.g., the ratios between population and phone calls or vehicle movements are not fixed and sociodemographic characteristics are not directly measured). It is clear that no one source can be used for the validation of all the others. Rather, it will be necessary to triangulate multiple sources to produce indicators of uncertainty. The science of interpreting and calibrating big data is at a relatively early stage and would most likely still require integration with baseline data of known quality such as residential counts and school registration numbers. Despite these difficulties, we consider the cross-validation and integration of these very different sources to be a key research challenge and potentially the only sufficient means of validating the outputs of comprehensive spatiotemporal population models.

Conclusion

In this article we have proposed a novel conceptual framework for spatiotemporal population modeling,

demonstrated a specific implementation, and illustrated its application with a case study. Our approach involves a spatiotemporal data system, with a continuous representation of time, supporting a separate analysis system. It meets a widely expressed need for population models that are temporally as well as spatially detailed and that have the potential to support a wide range of new analytical uses. This work addresses the central problem that static nighttime residential population distributions continue to be used for the majority of analyses when reality is enormously complex and continuously varying over multiple timescales. Incorporating greater temporal specificity has already been established as having the potential to deliver massively more accurate assessments of population exposure to hazard, demand for services, and emergency preparedness (Bhaduri et al. 2007; Aubrecht et al. 2013).

The approach described here provides an important step forward, but we have identified four key areas in which further research is needed. Implementation of the framework could be enhanced by the incorporation of alternative methods, such as the use of more sophisticated spatial interaction models to define catchments and flows. More generally, our approach lends itself to integration with other contemporary developments in spatiotemporal GIS. We have demonstrated an application using readily available data but have identified potential data enhancements in each of the spatial, temporal, and attribute domains. One particularly attractive development would be the integration of near-real-time data derived from a variety of continuously sensed data sources to augment the predominantly administrative systems employed here. There is also a need for the development of new validation approaches able to handle the inherent uncertainties associated with all of the potential comparator data sources.

Finally, consideration of spatiotemporal population modeling brings to the fore the importance of what we have called the modifiable spatiotemporal unit problem. This is relevant to every application of population mapping although the literature has, somewhat illogically, focused almost exclusively on the modifiable areal unit problem. When using spatial population data, analytical results will be heavily dependent not only on the spatial units but also the reference time to which the data relate. It is clear that a nighttime residential population map will be highly misleading if used to assess the population exposed to a daytime emergency, regardless of the choice of spatial units. Our approach is already able to reconstruct the very different population distribution at, for example,

8:45 a.m. compared to 10:45 a.m. As researchers begin to use new and richer data sources, particularly those captured by continuous tracking rather than formal enumeration, consideration of the modifiable spatio-temporal unit problem will necessarily become much more important. Working rigorously in this exciting but challenging new environment will require appropriate methods for handling spatiotemporal data such as the modeling framework presented here.

Acknowledgments

The authors gratefully acknowledge the comments of four anonymous referees. Data sets contain public-sector information licensed under the Open Government Licence v2.0 and Ordnance Survey OpenData Licence, Contains Ordnance Survey and National Statistics data © Crown copyright and database right 2013. Contains Royal Mail data © Royal Mail copyright and database right 2006. Annual Business Inquiry Service, National Online Manpower Information Service, licence NTC/ABI07-P3020. Quarterly Labour Force Survey, UK Data Service, usage number 40023. Meridian 2 National [SHAPE geospatial data], Scale 1:50000, Tiles: GB, Updated: July 2013, Ordnance Survey (GB), Using: EDINA Digimap Ordnance Survey Service, <http://digimap.edina.ac.uk>, Downloaded: Friday, December 13 22:19:27 GMT 2013.

Funding

The authors gratefully acknowledge Economic and Social Research Council Award RES-062-23-0081.

References

- Ahas, R., A. Aasa, S. Silm, and M. Tiru. 2010. Daily rhythms of suburban commuters' movements in the Tallinn metropolitan area: Case study with mobile positioning data. *Transportation Research C* 18:45–54.
- Ahola, T., K. Verrantaus, J. M. Krisp, and G. J. Hunter. 2007. A spatio-temporal population model to support risk assessment and damage analysis for decision-making. *International Journal of Geographical Information Science* 21 (8): 935–53.
- Alani, H., D. Dupplaw, J. Sheridan, K. O'Hara, J. Darlington, N. Shadbolt, and C. Tullo. 2007. Unlocking the potential of public sector information with semantic Web technology. Paper presented at the 6th International Semantic Web Conference, Busan, Korea. <http://oro.open.ac.uk/20028/1/> (last accessed 14 January 2014).
- Aubrecht, C., D. Özceylan, K. Steinnocher, and S. Freire. 2013. Multi-level geospatial modeling of human exposure patterns and vulnerability indicators. *Natural Hazards* 68:147–63.
- Bhaduri, B., E. Bright, P. Coleman, and M. Urban. 2007. LandScan USA: A high resolution geospatial and temporal modeling approach for population distribution and dynamics. *GeoJournal* 69:103–17.
- Birkin, M., and G. Clarke. 2012. The enhancement of spatial microsimulation models using geodemographics. *Annals of Regional Science* 49:515–32.
- Birkin, M., and N. Malleson. 2013. Investigating the behaviour of Twitter users to construct an individual-level model of metropolitan dynamics. Working Paper 05/13, National Centre for Research Methods, Southampton, UK. http://eprints.ncrm.ac.uk/3166/1/complex_city_twitter_Birkin_Malleson.pdf (last accessed 20 December 2013).
- Charles-Edwards, E., and M. Bell. 2013. Estimating the service population of a large metropolitan university campus. *Applied Spatial Analysis and Policy* 6 (3): 209–28.
- Çöltekin, A., S. De Sabbata, C. Willi, I. Vontobel, S. Pfister, M. Kuhn, and M. Lacayo. 2011. Modifiable temporal unit problem. Paper presented at Persistent problems in geographic visualization, International Cartographic Conference 2011 Workshop, Paris. <http://www.geo.uzh.ch/~arzu/publications/coltekin-et-al-icc-geo-vis2011.pdf> (last accessed 14 January 2014).
- Coombes, M. 2010. Defining labour market areas by analysing commuting data: Innovative methods in the 2007 review of travel-to-work areas. In *Technologies for migration and commuting analysis: Spatial interaction data applications*, ed. J. Stillwell, O. Duke-Williams, and A. Dennett, 227–41. Hershey, PA: IGI Global.
- Crooks, A., and S. Wise. 2013. GIS and agent-based models for humanitarian assistance. *Computers, Environment and Urban Systems* 41:100–11.
- de Jong, R., and S. de Bruin. 2012. Linear trends in seasonal vegetation time series and the modifiable temporal unit problem. *Biogeosciences* 9:71–77.
- Deng, C., and C. Wu. 2013. Improving small-area population estimation: An integrated geographic and demographic approach. *Annals of the Association of American Geographers* 103 (5): 1123–41.
- Department for Transport. 2009. National transport model: High level overview. <http://webarchive.nationalarchives.gov.uk/20110202223628/http://www.dft.gov.uk/pgr/economics/ntm/highleveloverview.pdf> (last accessed 14 January 2014).
- Eicher, C. L., and C. A. Brewer. 2001. Dasymetric mapping and areal interpolation: Implementation and evaluation. *Cartography and Geographic Information Science* 28 (2): 125–38.
- Fielding, J. 2007. Environmental injustice or just the lie of the land: An investigation of the socio-economic class of those at risk from flooding in England and Wales. *Sociological Research Online* 12 (4): 4.
- Foley, D. L. 1952. The daily movement of population into central business districts. *American Sociological Review* 17 (5): 538–43.
- Greger, K. 2015. Spatio-temporal building population estimation for highly urbanized areas using GIS. *Transactions in GIS* 19 (1): 129–50.
- Hagerstrand, T. 1970. What about people in regional science? *Papers in Regional Science* 24 (1): 6–21.

- . 1975. Space, time and human conditions. In *Dynamic allocation of urban space*, ed. A. Karlquist, L. Lundquist, and F. Snickars, 3–12. Westmead, UK: Saxon House.
- Health and Social Care Information Centre. 2013. Provisional monthly topic of interest: Trends in accident & emergency and workforce data. <http://www.hscic.gov.uk/catalogue/PUB11051/prov-mont-hes-admi-outp-ae-apr-feb-12-13-toi-ae-rep.pdf> (last accessed 10 February 2014).
- Jacquez, G. 2011. It's about space and time: From the modifiable areal unit problem (MAUP) to the modifiable temporal unit problem (MTUP) to the modifiable spatio-temporal unit problem (MSTUP). <http://www.biomedware.com/blog/2011/it%E2%80%99s-about-space-and-time-from-the-modifiable-areal-unit-problem-maup-to-the-modifiable-temporal-unit-problem-mtup-to-the-modifiable-spatio-temporal-unit-problem-mstup/> (last accessed 20 December 2013).
- Jochem, W., K. Sims, E. Bright, M. Urban, A. Rose, P. Coleman, and B. Bhaduri. 2013. Estimating traveller populations at airport and cruise terminals for population distribution and dynamics. *Natural Hazards* 68 (3): 1325–42.
- Kwan, M.-P. 2004. GIS methods in time-geographic research: Geo-computation and geovisualization of human activity patterns. *Geografiska Annaler B* 86:267–80.
- Lam, N. S.-N. 1983. Spatial interpolation methods: A review. *The American Cartographer* 10 (2): 129–49.
- Manyika, J., M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. Hung-Byers. 2011. *Big data: The next frontier for innovation, competition, and productivity*. McKinsey Global Institute. http://www.mckinsey.com/~media/McKinsey/dotcom/Insights%20and%20pubs/MGI/Research/Technology%20and%20Innovation/Big%20Data/MGI_big_data_full_report.ashx (last accessed 12 June 2014).
- Martin, D. 1989. Mapping population data from zone centroid locations. *Transactions of the Institute of British Geographers* 14:90–97.
- . 1996. An assessment of surface and zonal models of population. *International Journal of Geographical Information Systems* 10:973–89.
- . 2011. *SurfaceBuilder247: User guide*. Southampton, UK: University of Southampton. <http://www.esrc.ac.uk/my-esrc/grants/RES-062-23-1811/outputs/read/ece508b5-6438-4e96-99eb-8fd68f1d3b99> (last accessed 30 January 2014).
- Martin, D., and P. M. Atkinson. 2001. Investigating the spatial linkage of primary school performance and catchment characteristics. *Computers, Environment and Urban Systems* 5:67–83.
- Martin, D., S. Cockings, and A. Harfoot. 2013. Development of a geographical framework for census workplace data. *Journal of the Royal Statistical Society: Series A: Statistics in Society* 176 (2): 585–602.
- McPherson, T. N., and M. J. Brown. 2004. Estimating daytime and nighttime population distributions in U.S. cities for emergency response activities. Paper presented at the Symposium on Planning, Nowcasting, and Forecasting in the Urban Zone, American Meteorological Society, Seattle, WA. <http://ams.confex.com/ams/pdfpapers/74017.pdf> (last accessed 20 December 2013).
- McPherson, T. N., J. F. Rush, H. Khalsa, A. Ivey, and M. J. Brown. 2006. A day-night population exchange model for better exposure and consequence management assessments. Paper presented at the Sixth Symposium on the Urban Environment, American Meteorological Society, Atlanta, GA. <http://ams.confex.com/ams/pdfpapers/105209.pdf> (last accessed 20 December 2013).
- Mennis, J. 2009. Dasymetric mapping for estimating population in small areas. *Geography Compass* 3 (2): 727–45.
- Nakaya, T., A. S. Fotheringham, K. Hanaoka, G. Clarke, D. Ballas, and K. Yano. 2007. Combining microsimulation and spatial interaction models for retail location analysis. *Journal of Geographical Systems* 9:345–69.
- Office for National Statistics. 2010. Population estimates for UK, England and Wales, Scotland and Northern Ireland, mid 2006. <http://www.ons.gov.uk/ons/publications/re-reference-tables.html?edition=tcn%3A77-213674> (last accessed 13 January 2014).
- Openshaw, S. 1984. *The modifiable areal unit problem*. Norwich, UK: Geo Books.
- Peuquet, D. 1994. It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers* 84 (3): 441–61.
- . 2001. Making space for time: Issues in space-time data representation. *Geoinformatica* 5 (1): 11–32.
- Pultar, E., T. J. Cova, M. Yuan, and M. F. Goodchild. 2010. EDGIS: A dynamic GIS based on space time points. *International Journal of Geographical Information Science* 24 (3): 329–46.
- Rhind, D. W. 1991. Counting the people. In *Geographical information systems: Principles and applications*, ed. D. J. Maguire, M. F. Goodchild, and D. W. Rhind, 123–37. New York: Wiley.
- Schmitt, R. C. 1956. Estimating daytime populations. *Journal of the American Planning Association* 22 (2): 83–88.
- Sleeter, R. and N. Wood. 2006. Estimating daytime and nighttime population density for coastal communities in Oregon. Paper presented at the 44th Urban and Regional Information Systems Association annual conference, British Columbia. http://www.researchgate.net/profile/Nathan_Wood2/publication/228421037_Estimating_daytime_and_nighttime_population_density_for_coastal_communities_in_Oregon/links/00b4952d96dd725f78000000.pdf (last accessed 21 March 2015).
- Smith, G., and J. Fairburn. 2008. *Updating and improving the National Population Database to National Population Database 2*. Research Report 678, Health and Safety Executive, London.
- Stefanidis, A., A. Crooks, and J. Radzikowski. 2013. Harvesting ambient geospatial information from social media feeds. *GeoJournal* 78:319–38.
- Sutton, P., C. Elvidge, and T. Obremski. 2003. Building and evaluating models to estimate ambient population density. *Photogrammetric Engineering and Remote Sensing* 69:545–53.
- Sutton, P., D. Roberts, C. Elvidge, and H. Melj. 1997. A comparison of nighttime satellite imagery and population density for the continental United States.

- Photogrammetric Engineering and Remote Sensing* 63: 1303–13.
- Tobler, W. R. 1979. Smooth pycnophylactic interpolation for geographical regions. *Journal of the American Statistical Association* 74 (367): 519–30.
- Turnbull, J., D. Martin, V. Lattimer, C. Pope, and D. Culliford. 2008. Does distance matter? Geographical variation in GP out-of-hours service use: An observational study. *British Journal of General Practice* 58 (552): 471–77.
- Wirth, L. 1938. Urbanism is a way of life. *The American Journal of Sociology* 44 (1): 1–24.
- Wright, J. K. 1936. A method of mapping densities of population with Cape Cod as an example. *The Geographical Review* 26:103–10.
- Yuan, M. 1996. Modelling semantic, temporal and spatial information in geographic information systems. In *Geographic information research: Bridging the Atlantic*, ed. M. Craglia and H. Couclelis, 334–47. London and New York: Taylor & Francis.
- . 2008. Adding time into geographic information system databases. In *The handbook of geographic information science*, ed. J. P. Wilson and A. S. Fotheringham, 169–84. Oxford, UK: Blackwell.
- Zhang, Z., R. Sunila, and K. Verrantaus. 2010. A spatio-temporal population model for alarming, situation picture and warning system. *International Society for Photogrammetry and Remote Sensing Archives* 38 (2): 69–74.

Correspondence: Geography and Environment, University of Southampton, Southampton, SO17 1BJ, UK, e-mail: D.J.Martin@soton.ac.uk (Martin); S.Cockings@soton.ac.uk (Cockings); School of Civil Engineering and Surveying, University of Portsmouth, Portsmouth, PO1 3AH, UK, e-mail: samuel.leung@port.ac.uk (Leung).