Research Article

Exploring the Historical Determinants of Urban Growth Patterns through Cellular Automata

Kiril Stanilov
Centre for Advanced Spatial Analysis
University College London

Michael Batty
Centre for Advanced Spatial Analysis
University College London

Abstract

We have adapted METRONAMICA, an established cellular automata (CA) modelling system, to simulate the historical growth of a section of a large world city. Our model is tuned to reflect the morphology of land use patterns more accurately than traditional CA models, which abstract those patterns to more aggregate spatial scales. We explore the spatial determinants of land use patterns with detailed empirical data, documenting the historical growth of West London at an unusually high level of spatial and temporal resolution. The results of the study provide support for our considered speculations: (1) that the spatial relationships between land uses and the physical environment are remarkably consistent through time, showing little variation relative to changes in historical context; and (2) that these relationships constitute a basic code for urban growth which determines the spatial signature of land development in a given metropolitan area.

1 Introduction

The resurgence of urban modelling since the beginning of the 1990s has been partly spearheaded by the development and application of new approaches to understanding the dynamics of urban environments linking global patterns of change to micro-scale processes. Advances in the field of urban simulation related to cellular automata (Batty 1997, Yan, 2009), multi-agent systems (Parker et al. 2003, Crooks et al. 2008), micro-simulation (Birkin et al. 2010) and connectionist models such as artificial neural networks (Li and Yeh 2002, Guan et al. 2005) have re-established urban modelling as a

Address for correspondence: Kiril Stanilov, Centre for Advanced Spatial Analysis (CASA), University College London, 1-19 Torrington Place, London WC1E 6BT, UK. E-mail kiril.stanilov@gmail.com

powerful aid for understanding and managing the complexity of cities at the beginning of the twenty-first century. Like their predecessors, this new crop of urban growth models can be placed in two broad categories – abstract theoretical explorations and operational data-driven applications (Torrens 2001). This division remains distinct to this day, with each group of models exhibiting its own set of methodological strengths and weaknesses. Unlike their predecessors, the models are largely physical in focus, building not on urban economic theory or housing market analysis but on notions about land development and transitions with respect to land use and land cover.

When urban models were first proposed in the late 1950s, this distinction between abstract and operational models was soon established. Urban models then were largely configured as simulating static structures as though the city were in equilibrium with little attention to the detailed dynamics that characterize the new generation of land development and agent-based models that is our focus here. Models then were largely predicated around urban and regional economic theory, based on abstracted land markets and population density profiles linked to movement patterns based on notions of flow and force in classical physics. Their empirical equivalents drew on this theory and operationalized their form through ideas from transportation modelling, input-output analysis, and demographic forecasting, all set within the wider context of the economics of utility maximization (Batty 2008). Then as now there was a tension between the pedagogic value of these models and their policy relevance which resulted in a rather strong reaction against these initial developments that left the field bereft of new ideas for more than a generation (Lee 1973).

As before, the main characteristic and advantage of the new abstract urban growth models that have evolved is that they are more strongly grounded in physical development represented in specific ideas about key determinants of urban form and pattern. These models usually have a long time horizon, relating the explored phenomena to long-term trends in the evolution of the urban environment. Good examples of such models are those assessing the utility of bid rent theory within or beyond the constraints of the monocentric city (Caruso et al. 2009), simulations analysing aspects of residential segregation such as those based on Schelling's (1969) model, or inquiries investigating the fractal dimensions of cities (Batty and Longley 1996) which relate these ideas to density, sprawl, and polycentricity. The majority of these models are powerful pedagogical vehicles for examination of various theoretical assumptions. Their Achilles heel, however, is their reliance on hypothetical or highly aggregate data, founded on plausible but often untestable assumptions which makes their validation a particularly difficult task. A related shortcoming is their limitation in capturing the complexity of the urban pattern. Due to the coarseness of their approach to urban form, they reduce the spatial heterogeneity of urban form and layout to a 'conceptual diagram' of the general distribution of urban activities on a metropolitan-wide scale in which the representation of the complex pattern of the built environment tends to be smoothed out (Batty 2005a).

Unlike these more abstract models, operational urban growth models based on more solid empirical data have been built around intricate investigations of actual development patterns. The main goal of such projects is to check the imprint of contemporary development processes as they happen 'on the ground', so-to-speak, and to establish several growth scenarios for the foreseeable future. These models tend to be most applicable to relatively short time horizons and their main thrust is to serve as the physical development component in the arsenal of tools defined within planning support systems (Ward et al. 2000, Li and Yeh 2000, Yang and Lo 2003, Jantz et al. 2003, Brail

2008). Their lack of theoretical grounding however, ultimately limits their utility as a basis for understanding the generative forces behind urban growth.[1]

In cellular automata modelling, the historical trajectory of the field has been marked by a shift from abstract explorations to operational applications, with few attempts to link such models to urban economic growth theory. Since the late 1990s, the proliferation of CA-based operational models has charted the path for the development of the field in a rather technical direction. Most of the research in CA modelling has been consumed by issues related to improvement of the apparatus of model construction (Yang and Lo 2003), particularly refinements to the techniques of transition, not the mechanics of land development (Yeh and Li 2002b, Liu and Phinn 2003), the automation of calibration, the improvement of model results (Chen and Mynett 2003, Menard and Marceau 2005, Kocabas and Dragicevic 2006), and the integration of CA with GIS (Clarke and Gaydos 1998, Almeida et al. 2005). Within this agenda, little attention has been given to questions of tracking slow and fast dynamics and hence the models reflect transitions in land use which are based on fitting land use change to past trends Good examples are recent efforts to couple CA with ANN where the focus has been on improving the fit of generated to observed patterns through the tuning of cell transitions and neighbourhood size parameters. These experiments have produced promising results in terms of improving the accuracy of model outcomes, but remain fairly limited in their ability to advance our understanding of underlying generative factors and processes shaping urban growth (Yeh and Li 2002a).[2] In contrast, key issues such as the role of transportation and a detailed underpinning of the development process in terms of demand and supply for land, have not been broached in such models, at least in any depth. This has also tended to keep these models at arms length from policy makers, with the older, more operational land use transport models still dominating practice.

The model presented in this article has been conceived as an attempt to bridge what appears to be a widening gap between these theoretical and empirical models of urban growth. It explores a specific theoretical assumption claiming that the underlying forces shaping metropolitan growth patterns are rooted in a set of enduring spatial relationships existing between urban land uses and the built environment in terms of form and layout which show little variation across historical periods. Unlike the majority of the theoretical explorations of urban growth, however, the development of the model is grounded in detailed empirical data documenting the historical growth of a metropolitan area at the possible highest level of spatial resolution, while continuing to employ the detail-oriented analytical gear of operational modelling applications.

The following sections elaborate on these two critical aspects of the model – its theoretical assumptions and data development methodology. The discussion continues with a description of the model's structure, its calibration procedure and results. The conclusions highlight the key findings of the study, recognizing its methodological limitations and sketching a direction for further development of the ideas laid out in this article.

2 Theoretical Assumptions: the Quest for a Basic Urban Generative Code

We begin by exploring the notion that at the core of what matters most in structuring the patterns of urban growth are a set of enduring spatial relationships. These relationships are defined by the forces of attraction and repulsion existing between the major land use

classes, as well as by a number of key spatial characteristics of the built environment. A key argument of this study is that these relationships transcend socio-economic circumstances in the sense that they precede and operate to a large extent autonomously from shifts in economic, political and technological regimes. In a fitting biological analogy, the spatial relationships analysed here can be described as the generative code of urban development – a set of fundamental rules that govern the shape and growth of an urban area over the course of its existence. We refer to this as a generative rather than genetic code for the analogy has not been well-worked out as yet (Silva 2004, Wilson 2010). We prefer to use a more neutral and less controversial term which grounds our approach in the computational generative grammars already developed in urban and built form studies such as shape grammars (Stiny 1980). Socio-economic factors exert an influence on the patterns of growth but as an overlay agency superimposed on the primary set of fixed spatial relationships. To continue with the biological analogy, social agents alter the patterns of urban growth similarly to the way in which the environment impacts the development of an organism whose structure and shape is defined *a priori* by rules that are ultimately reducible to its genetic code.

We postulate that the fundamental spatial relationships shaping metropolitan growth patterns operate on two separate but interrelated levels. On the first level (the local scale), they are defined by longstanding forces of attraction and repulsion exhibited between the various categories of urban land uses. Thus, for instance, all land use classes are attracted to themselves (resulting in the formation of homogeneous land use clusters); some are attracted to each other (e.g. high-density residential and commercial, residential and recreational, etc.); while others are indisposed to co-location (e.g. residential and industrial).[3] On the second level (the regional scale) land use patterns are conditioned by the physical properties of the overall urban spatial frame, which consists of the main elements of the transportation infrastructure (major roads and transit nodes) and the network of activity centres (CBD and suburban activity clusters). The patterning of infrastructure elements and activity centres determines the regional accessibility of every location within the metropolis, exerting decisive influence on the spatial patterns of land use distribution. It is these elements of urban structure that we consider are largely exogenous to the code of development and once determined, the code adapts to the form that is laid down from above.

Spatial interactions between land uses and the impact of accessibility on urban development patterns are well established concepts in urban morphology (Stanilov 2002, 2003). Here we explore and test assumptions at the level of informed speculations which suggest: (1) that the spatial relationships between land uses and the physical environment are remarkably consistent through time, showing little variation relative to historical context; and (2) that these relationships constitute the generative code for urban growth which determines the long-term patterns of land development in a given metropolitan area. These ideas appear to contradict both conventional wisdom and the understanding underlying most urban research that the form of cities and the patterns of urban development are actively shaped by a set of dynamic socio-economic forces. It is widely agreed that changes in economic conditions, demographic trends, technological innovations, political and cultural paradigms continuously mould the built environment. Thus, each historical period leaves its unique spatial signature, layering the urban landscape with patches and pieces which we recognize as components of the 'compact city' of the preindustrial era, the 'garden suburbs' of the late 19th and early 20th centuries, the 'urban sprawl' of the post-war decades, the 'urban renaissance' of the post-industrial age.

Figure 1 Comparative timeframe diagram of traditional land use models (a) and the West London model (b)

The ideas raised here challenge these assumptions to the extent that they attribute the physical patterning of urban uses primarily to a set of spatial relationships that persist through time regardless of changes in socio-economic contexts.

In order to test this hypothesis, we strove to develop a model that draws on the strengths of both exploratory and operational models while avoiding their weaknesses. In other words, we have built an exploratory model which tests the validity of our assumptions based on solid empirical data covering an extended timeframe. It seemed that the only way for us to meet these challenges was to reverse and redefine the thrust of traditional urban growth models. What if rather than trying to predict the future which, as Popper (1957) eloquently argued, is unknowable, we step far back in time and try to predict the present?[4] What if we pretend that the year is 1901, Queen Victoria has just died, the Victorian world is ending, and the peace in Europe continues to be fragile with World War I in sight? Can we determine what cities would look like in the beginning of the 21st century based on data from the late Victorian period? Can we do better than what H. G. Wells (1902) did at the time, relying solely on his astute intuition? We would have the natural advantage of knowing what has actually happened, and a powerful tool in the form of a model that would allow us to test the validity of our assumptions about the consistency and importance of key spatial relationships across time using hard historical evidence rather than speculation.

With this general conceptual framework in mind, we developed a cellular automata model of West London's historical growth, simulating its patterns of land use changes over the course of the last 130 years. The development of the model was based on time series maps which we generated with high spatial and temporal resolution derived from detailed historical Ordnance Survey (OS) records. The idea was to calibrate the model based on data only from the first three map series (1875, 1895, 1915) and then let the model run all the way to 2005 without changing its initial parameter values (as we show in Figure 1). The fit between the post-1915 patterns generated by the model and those recorded from OS maps served as a test for the validity of the model assumptions. The results confirmed our hypothesis, showing a surprisingly robust and accurate 'prediction' of West London's growth patterns for 1935, 1960, 1985 and 2005, thus supporting the

argument that a limited set of spatial variables can be read as an urban generative code, determining to a large extent the long-term physical evolution of an urban area.

3 Data and Methodology

The selection of West London as a case study for this project was based on several considerations. In the history of contemporary urbanization, London has a unique place as the first modern metropolis of the industrial age. As such it served as an inspiration for the development of many large cities around the globe and more specifically for the burgeoning metropolises of the western world during the second half of the 19th and the early part of the 20th centuries. The massive decentralization of population and economic activities which took place in Greater London during this time served as a precedent, setting the tone for the suburbanization of many large cities in the Anglo-Saxon world and beyond, with lasting impacts on contemporary development patterns around the globe. It is widely acknowledged that England was the birth place of the industrial revolution and the industrial cities which grew up there still represent the archetypal form of the city in terms of the structure that has dominated city building until quite recently. Indeed many of the terms that describe the modern city such as 'sprawl' were first used to describe London from the mid 18th century onwards.[5]

This distinctive position of London as a global city of unique historical significance has been reflected accordingly in voluminous studies documenting its growth. The wealth of historical records on the evolution of the metropolis makes possible the reconstruction of its physical expansion with an unprecedented level of detail – building by building and plot by plot. This opportunity to trace the patterns of urban growth with the highest level of precision was deemed critical for the development of our model aimed at analysing the processes of urban form generation at the micro-scale.

A third consideration in selecting London as a case study, besides its key role in urban history and its well-documented past, is related to the characteristics of London's planning regime. The decentralized and fragmented approach to managing urban development has been a consistent feature of London's history (Hall 1989), thus making it easier to isolate the impact of planning on metropolitan development patterns to a limited number of key interventions related, in the case of West London, to the establishment of the Green Belt and a few infrastructure projects of regional significance. Many other countries around the world have developed both planning systems and their instruments such as 'Green Belts' and 'New Towns' based on the example of the UK more generally and London more specifically.

The selection of West London as an area of concentration within the metropolitan territory was dictated by our inability to cover the entire area of Greater London within the time constraints of the project. This area was selected to explore a major axis of London's historic growth, covering one of the most dynamically evolving segments of London's metropolitan fabric. The large spatial extent of the study area (200 km^2) includes a rich mixture of local jurisdictions, which ensures that the recorded development patterns comprise a representative sample of London's metropolitan fabric (notwithstanding intra-metropolitan differences between East, West, South and North London). The east boundary of the study area is defined roughly by the western edge of Hyde Park, coinciding with the edge of the compact city in the initial study year of 1875. The study area expands 20 km westward to the edges of London's Green Belt, which defines the current boundary of Greater London and covers a slice of the metropolis some

Figure 2 Study area of West London

10 km north-south (as we show in Figure 2). It is worth noting that in many respects, this slice has the most variety of any of the possible slices through the metropolis, in that its diversity of land uses is far greater than slices which run out from the centre to the east, north or south or any other compass points that might be chosen. Many detailed analysis of London's urban structure bear this out (see for example Hall 1989, Rasmussen 1937).

A unique feature of the database employed by this project is the extensive time coverage, which spans the last 130 years of London's urban growth. The data set includes time series maps showing slices in the evolution of West London's metropolitan fabric in 20–year increments, starting from 1875 onwards. The time coverage of the project stands in marked contrast to operational urban models, which traditionally go back in time only a decade or two. The unusually long time horizon of this study was framed to reach the pre-urban era in the development of the study area. This allowed us to trace the emergence and evolution of London's suburban fabric from its incipient stages of urbanization to the present.

For the documentation of land use change, we used highly detailed historic Ordnance Survey (OS) maps at a scale of 1:2,500. This allowed us to identify a wide range of land use categories and building types with a high level of spatial and interpretational accuracy. This database presents a significant improvement in data resolution compared to traditional land use and land cover models which rely on historical LANDSAT and SPOT remotely sensed (RS) imagery limited to a pixel size of 10-to-30 m. The data derived from such sources presents significant challenges in the identification of various urban land use classes. In most RS-based land use change models a distinction is made between residential, non-residential, and recreational uses, but even these broad categorizations are vulnerable to inaccuracies of interpretation due to the low resolution of the satellite images. In contrast, the fine-scale OS maps allowed us to identify close to 60 land use classes and building types with an accurate representation of actual parcel boundaries.

In this study a manual interpretation of the maps served as a basis for the creation of ArcGIS polygons of homogeneous land uses. We delineated these polygons covering the entire study area for all of the map series (1875, 1895, 1915, 1935, 1969, 1985 and 2005). The digitizing of these polygon coverages was performed onscreen from geo-referenced TIFF images of the historical OS maps. The TIFF images were downloaded from the online map service, EDINA's Digimap, the JISC national data centre dealing with map coverages at the University of Edinburgh. The co-registration of map series was done by overlaying the land use polygons derived from the earlier time series on the subsequent historical map and adding new or updating existing polygon boundaries in areas where land use change was detected. The process of land use classification involved the interpretation of building footprints from the OS maps; verification of building type (for buildings still in existence) in Microsoft Virtual Earth (now Bing Maps 2D and 3D) and Google Street View; and cross-referencing the results with several land use databases for Greater London (Ordnance Survey; Valuation Office; Virtual London). As the major-ity of the non-residential buildings are clearly labelled in the OS maps, these OS designations were directly used for the assignment of properties in aggregated land use classes. Following this procedure, land use polygon coverages were created for all of the seven map series. The vector-based maps were then converted to grid coverages with a cell size of 25 by 25 m. Throughout this process, we were aggregating vector map data to raster and therefore had complete control over the re-sampling.[6] The small size of the cells employed in our model, compared with most traditional CA land use models where the size of the cells varies from 50 to 1,000 m, allows greater spatial accuracy, reducing the problems associated with cell heterogeneity. Moreover the chosen cell size of 25 m × 25 m is consistent with representing building layout in its most generic yet realistic abstraction.

In addition to the development of the land use database, we recorded the evolu-tion of the infrastructure network in the study area. For each one of the map series, this process included the digitization of each roadway, railway and waterway; and a recording of the location and opening date of each railway and underground station. In addition, we identified the centre of the major suburban clusters as they appeared on the maps using the standard neighbourhood functions in ArcGIS Spatial Analyst. These data were important components in the development of our urban growth model.

4 Model Development: Applying METRONAMICA

The selection of cellular automata as a modelling approach for this project was based on CA's proven ability to deal with spatial phenomena (Torrens 2000) and their capacity to handle high-resolution applications with many cells easily (see, for example, White and Engelen 2000). We chose to use METRONAMICA, a modelling system developed by RIKS (http://www.riks.nl/ available from http://www.metronamica.nl/) on the basis of CA modelling concepts advanced in the 1990s (White et al. 1997) which originated from ideas about fractal cities (Batty and Longley 1994). The system, which has been exten-sively used in various urban contexts, has several advantages compared to the widely popular SLEUTH urban growth model (Clarke et al. 1997) and other more individual-ised packages such as DUEM (Batty 2005b). First and foremost, METRONAMICA has the capacity to model a wide range of urban land uses (the current limit set to 26 classes),

while applications of SLEUTH appear not to have made detailed simulations of urban land use using only a binary distinction between urban and non-urban uses. The second attractive feature of METRONAMICA is its ability, through a well-designed interface, to set parameter values interactively and explore the model behaviour visually and in real time. This ability of the modelling system encourages experimentation, offering immediate feedbacks through built-in features particularly suitable for the calibration of exploratory models in which the impact of determinant forces is tested through a method of trial-and-error. In contrast, DUEM is much more tailored to simulating urban sprawl in the US and was developed specifically for the Detroit / Ann Arbor / southeastern Michigan region.

The first step in the development of our model was to determine the optimal number of land use classes to be included in the model environment. The process involved experiments with different levels of land use aggregation, based on analysis of the spatial behaviour of all land use classes recorded in the map series. The behaviour of the various land uses was evaluated on the basis of their location relative to major roads, railway stations, CBD and suburban activity clusters (town centres); their spatial and functional affiliation with other land uses; and their resilience to conversion to a different type of use. Thus, for instance, parks, cemeteries and active recreation were included in the land use class *recreation*, while due to their relatively short longevity, allotment gardens and plant nurseries were assigned to the *vacant land* category. Functional and locational affiliation, on the other hand, served as a basis for the grouping of institutional, educational and religious uses in the *residential* category. Ultimately, the number of land use classes included in the model was reduced from 35 to nine, leading to an optimized modelling environment consistent with the general requirements of parsimony, and various well-established conventions for aggregating classes when articulating land use activities.

Following the modelling concepts embedded in METRONAMICA, the nine land use classes were divided in three groups. The first group is composed of three classes – residential, commercial and industrial uses – which are actively modelled. The dynamics of these land uses, called *active functions*, respond to exogenous demand for land. In our case the amount of development, or the number of cells in each of the three classes, is set by the area of these land uses as recorded in the map series. In other words, the model takes the number of cells in each one of the study periods and for each one of the three active land uses as a given exogenous constraint, and allocates this growth in the study area for each time step (one year increments).

The second group of land uses, called *passive functions*, is composed of the land use classes that are not controlled by exogenous demand. To this group we assigned the vacant uses and a class called soft development, which is composed of estates, farms and other types of land particularly prone to urban conversion. These passive functions appear or disappear as a result of land being taken or abandoned by the growth or decline of the active functions listed above.

Finally, a third group of land uses is composed of the *static* classes, which appear instantaneously in the landscape and change little over time. Here we placed airports, transportation, water, recreation and large institutional uses (military bases, large hospitals, prisons, etc.) reflecting the fact that these developments are not driven by processes of organic growth but are known to be or at least appear as a result of centralized decisions at certain moments in time. These land uses are therefore introduced in the model at the time when they first appear on the OS maps. In this sense, these uses are not

actively modelled, but they influence the location of the other land uses through their attraction or repulsion effects.

A next critical step in the development of the model was the integration of accessibility parameters, a modelling function that METRONAMICA handles by introducing various infrastructure elements as polyline shapefiles overlaid on top of the land use maps. The resultant accessibility network can be updated in the course of the model runs, which in our case was done in 20-year increments corresponding to the data from the digitized map series. For the road network, we included only the major roads, which we classified as primary arterial roads and secondary collector roads. Local roads were excluded from the model as we considered their properties to be inextricably linked with land development. Our accessibility parameters included also the main transit nodes (railway and underground stations) and the location of the CBD and major suburban activity clusters.

To the extent that our model was calibrated on data from 1875, 1895, and 1915, an era largely preceding the first Town Planning Act of 1909, development regulations were not included as a determinant of land use patterns. This decision was in line with our main goal to test to what extent the patterns of growth in West London could be explained strictly by spatial characteristics inherent in the built environment. However, we introduced restrictions on land development imposed by the establishment of the Green Belt, which was first advanced as a planning concept in the Greater London Plan (Forshaw and Abercrombie 1943, 1944) and first implemented comprehensively in the 1950s by the Ministry of Housing and Local Government.

After plugging in the land use, accessibility and zoning data, the next step in the construction of the model was the establishment of the CA's transition rules. Here we relied on METRONAMICA's built-in allocation algorithms according to which the transition of cells from one land use state to another is based on a value called *transition potential*. This value is calculated for each cell and for each of the active land uses based on data from three sources: the composition of land uses in the cell neighbourhood; the distance of the cell to various elements of the accessibility network; and the overlay of zoning restrictions.[7] This deterministic value is then given a stochastic perturbation: based on the value of its transition potential, a cell changes its state to the land use class for which it has the highest value. If the current land use class has the highest score, the cell remains in the same state. Cell transitions start from the highest ranked cell and proceed downwards until the exogenously determined demand for cells in a particular land use class is satisfied. Transition potential values are dynamically updated in each step of the model run, which in our case have a duration of one year. These times steps are interpolated evenly between the time slices for which the data is available.

The calibration of the model involved the refinement of parameter values related to the interaction among various land uses across distance and the influence of various elements of the accessibility network on the active land use functions. The initial assignment and refinement of parameter weights was done through an interactive process by visually assessing the spatial effects of the parameter weights on the patterns generated by the model. This method is often preferred in urban growth and land use change modelling since general statistical methods of calibration are considered by and large unreliable in modelling phenomena exhibiting high degrees of spatial correlation (Clarke et al. 1997, Ward et al. 2000, Yeh and Li 2002b, Barredo et al. 2003). In fact METRONAMICA produces a wide variety of statistical tests that relate to map pattern

analysis and we tested their relevance as a basis for our interpretations.[8] Our results showed that for our model of long-term land use dynamics, visual evidence is essential, thus confirming Epstein's (2008) argument that statistical tests are but one of many sources for validation and Mandelbot's statement that the most important quality of a model is to 'look right' (Mandelbrot 1983).

In order to test our hypothesis about the significance and consistency of the analysed spatial interactions between land uses and the built environment, we used the map of 1875 as a starting point and calibrated the model parameters with reference to the land use patterns recorded in the 1895 and 1915 map series. Having implemented the calibration based on this 40-year period (from 1875 to 1915), we let the model run from our initial year of 1875 to the present day. Our plan to run the model beyond the calibration period was driven by the project's objective to develop a model which allows us to compare the model's predictions with reality. This latter step is considered the first step in model validation as such (Straatman et al. 2004).

Finally, it is important to underscore that the model which we developed is a constrained cellular automata in the sense that the simulated land use dynamics are influenced by exogenous input. Through the course of the simulation run, spanning 130 years, we updated the model every 20 years with information derived from the map series related to: (1) the demand for development for the next 20-year period for each of the three active functions (residential, commercial and industrial); (2) changes in the accessibility network (new major roads, railway and railway stations, emerging suburban centres); and (3) the introduction of new static land use features (airports, recreation, large institutions). These types of constraints are traditionally employed in most CA-based land use models as a mechanism for adapting the abstract mathematical CA apparatus to the realities of the urban development process, leading to substantial improvement in model outcomes.

5 Model Predictions and Results

The analysis of West London's growth patterns, which we documented in the historical map series, revealed a highly dynamic and complex land use configuration. A comparison of the maps featuring the areas of land absorbed by new development during the six study periods indicated clear qualitative shifts in the patterns of urban growth. We discerned three distinct periods characterised by patterns of nucleation (1875 to 1915), diffusion (1915 to 1960) and infill (1960–2005) as we show in Figure 3. The challenge for our model was to reproduce these patterns based solely on land use data for the first period and a fixed set of parameters reflecting what we deemed to be the key characteristics of spatial interaction between land uses and accessibility.

The results of the spatial analysis of the land use dynamics, which we performed in ArcGIS, confirmed the existence of systematic and consistent relationships between the distribution of land uses and their proximity to major roads, railway stations, suburban activity clusters, and the CBD. The identified consistency of these spatial relationships contradicted some commonly accepted views about the growth of London. For instance, the popular opinion that the interwar period was defined by an unprecedented escalation of decentralization and sprawl (Hall et al. 1973) seemed, on a first glance, to be confirmed by the patterns of residential growth documented in our time series. Compared with the previous periods, the development that took place between 1915 and 1935

**Figure 3**   Patterns of new land development

indeed appeared much more dispersed relative to the location of the CBD and existing railway stations as in Figure 4.

However, when we plotted new residential development (1915 to 1935) on land which was available in 1915 (excluding already developed cells), the distribution of new residential development relative to distance to the CBD and railway stations remained surprisingly similar to the previous periods as in Figures 5a and 5b. The enduring influence of accessibility to key elements of the metropolitan spatial structure on the patterns of growth was even more pronounced in the plots showing the distribution of residential uses relative to arterial and connector roads in Figures 5c and 5d.

Our findings about the importance and consistency of key spatial determinants of urban growth were supported by the outcomes of our simulations. The map of predicted land uses for 2005 generated by our model demonstrated a surprising degree of correspondence with the actual land use patterns for that year as recorded in our map series. In the evaluation of the model results, we looked for the 'correctness' of the overall patterns, exhibited by characteristics such as the general distribution of land uses across the study area, the degree of dispersal relative to the city centre, the location and size of clusters, and the general level of spatial affiliation between pairs of land use classes (residential and commercial, residential and industrial, recreational and residential, commercial and airports, etc.). Our scepticism of the use of statistical tests, which in this case appear as good as many other such applications, is based more on our reactions to what these tests do not measure than what they do.[9] We were cognisant that location-specific estimates based on landscape metrics may not be as useful as having the model reproduce realistic patterns – which is the point made frequently by Mandelbrot (1983) that the outcomes of models must look right – a criterion on which the model surpassed our expectations.

Figure 4   Residential development 1915–1935 relative to the location of railway stations (darker background colour indicates greater distance from railway stations)
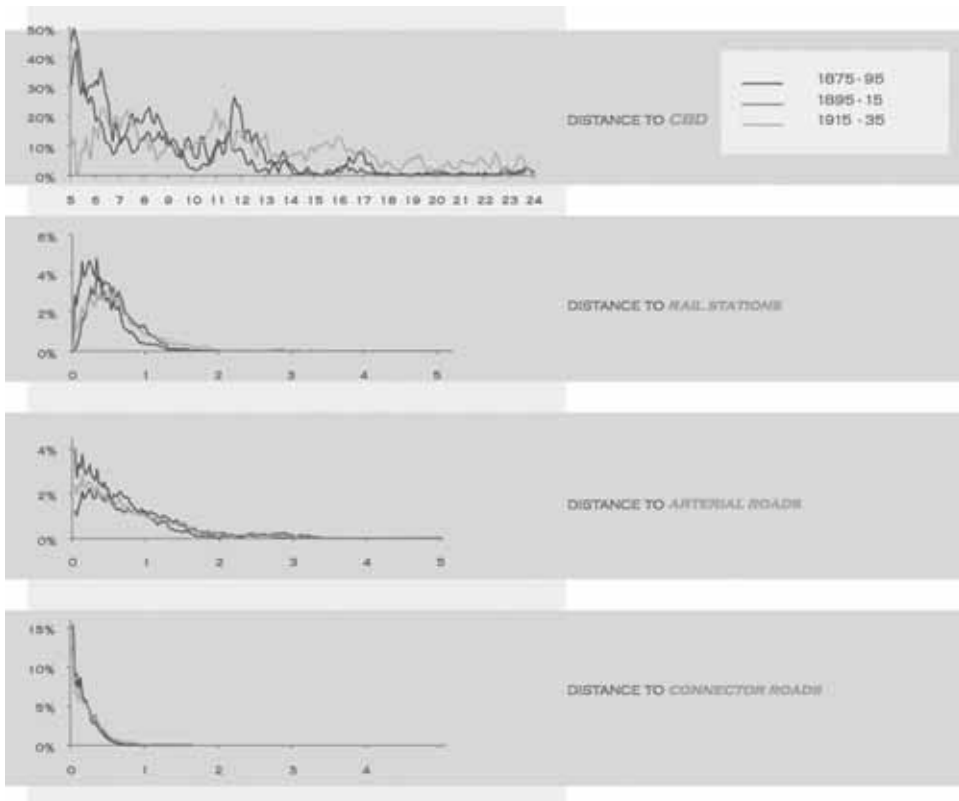

Figure 5   Distribution of residential uses relative to: (a) CBD; (b) railway stations; (c) arterial roads; and (d) connector roads
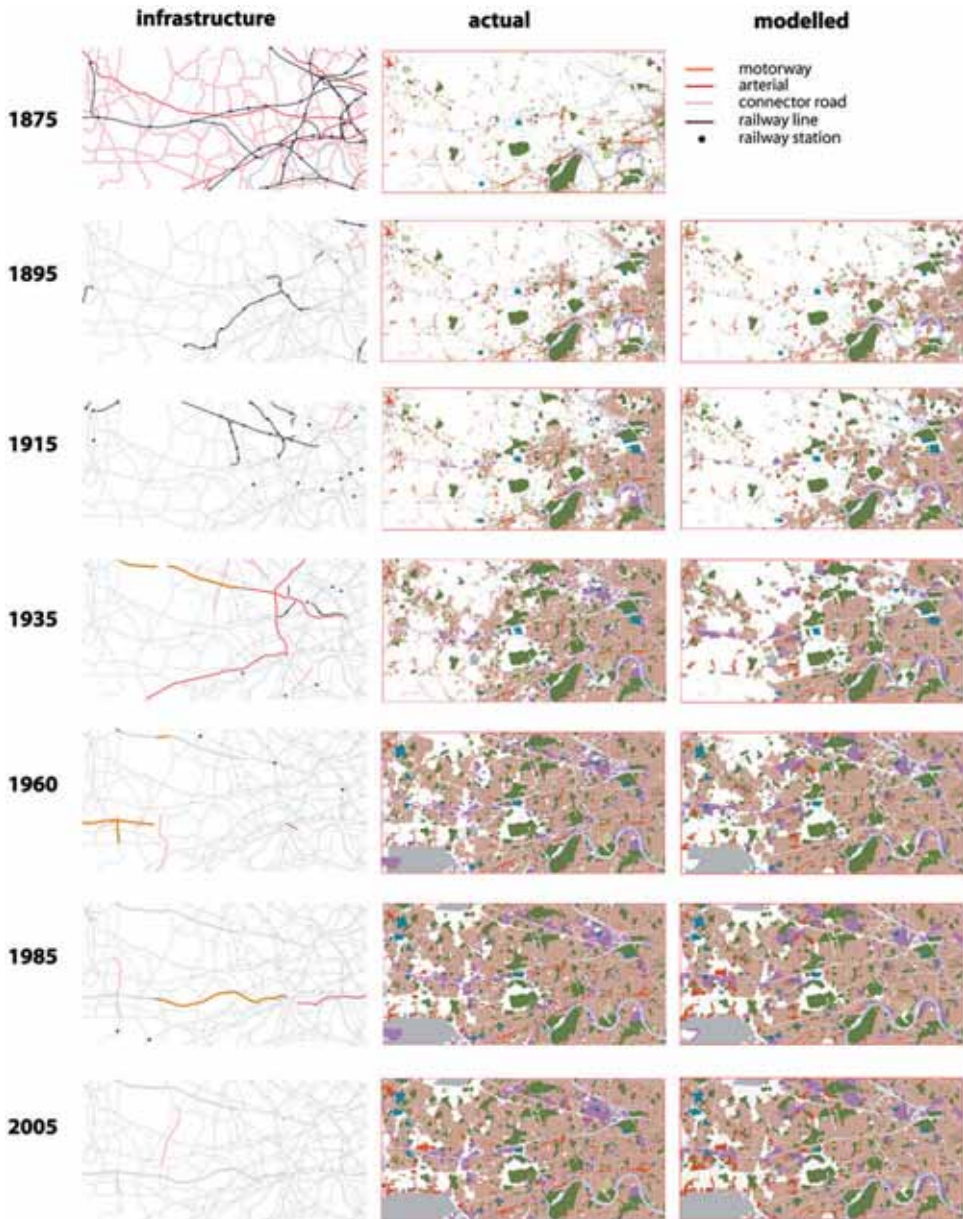
Figure 6   Comparison of 'predicted' and actual land uses by study periods

What is more, the simulation of West London's growth not only generated realistic results for the year 2005 based on data from the 19th century, but it also captured important properties of the urban growth dynamics characterizing the evolution of the urban pattern. The model predicted with high degrees of spatial and temporal accuracy the allocation of land uses in each one of the study periods as we show in Figure 6, capturing the transitions in urban growth from nucleated, to diffused, to infill. In

Figure 7 Emergence of industrial corridors (top) and suburban commercial clusters (bottom) as predicted by the model

addition, our simulation successfully reproduced the emergence of industrial corridors in the first half of the 20th century in Figure 7a, and the emergence of commercial clusters in the suburban outskirts towards the century's end in Figure 7b.

6 Conclusions

Half a century after the development of the first land use models, urban science has yet to cover significant ground in understanding the patterns and processes of urban growth. In this endeavour, CA-based models have proved to be an important tool for exploration, capturing dynamic properties in the urban pattern related to new concepts in complexity science such as self-organization, path dependence, emergence and phase transitions (Couclelis 1985, Cheng and Masser 2004, Batty 2005b). However, CA models are severely limited by the absence of explicit processes that map onto what we know about how land develops and markets clear in urban areas. These models do not capture the details of the built environment that represent the inertia and trace of the historical urban development. In this study, we have not attempted to ground our model in ideas from urban economic theory, but we have attempted to represent the detailed urban form of the city, albeit at the 25m level of resolution, through linking basic ideas of transition to high-resolution empirical data, showing in fact a degree of inertia and resilience to change that has not been characteristic of previous applications.

The spatial analysis which we performed on the data documenting the historical growth of West London provided strong support for our ideas that the patterns of urban growth are underlined by enduring spatial relationships which define the interactions between land uses and accessibility parameters related to the physical framework of the built environment. Calibrating our model on data from 1875 to 1915, we managed to 'predict' with a striking level of realism the dynamics of the growth patterns for the next

90 years. We interpret these results as proof of the enduring nature of the spatial relationships we investigated, which we believe underlie the patterns of urban growth regardless of changes in particular socio-economic circumstances.

Empowered by a better understanding of this generative code of urban growth, we believe that this style of urban modelling could be more actively engaged in long-term forecasting, an area which has been underexplored due to the lack of theoretical and empirical support for the idea that spatial relationships could operate autonomously from social agents. Obviously, it would be inappropriate to draw generalizations based on the results of a single case. The results of this study need to be independently validated for metropolitan areas in other countries and continents. It would be particularly interesting to test if our hypothesis would be valid in urban contexts characterized by stronger or weaker centralized planning regimes. There is mounting evidence, however, that it is possible to re-use qualitative calibrations of the basic spatial relationships in different contexts, suggesting that 'the underlying process that determine urban and regional form are to a large degree universal' (White and Engelen 2000).

Our quest to develop a generative grammar for urban development akin to a genetic code is part of a wider exploration of biological analogies in the built environment. However, while it is unlikely that the application of concepts borrowed from genetics can heal our cities of all their urban ills, we strongly believe that a main reason why so many urban policies and plans continuously fail to achieve their goals is linked to our limited understanding of the spatial dynamics of urban land development and the forces shaping urban growth. Jane Jacobs (1961) made this point almost 50 years ago when she said: 'It is futile to plan a city's appearance, or speculate on how to endow it with a pleasing appearance of order, without knowing what sort of innate, functioning order it has.' We believe that the discovery of an urban generative code could improve the effectiveness of planning and help us meet some of the greatest challenges facing our cities at the beginning of the 21st century.

Acknowledgements

Notes

1   Our focus on lack of theory in CA models should not be meant to suggest we are unaware of the theoretical origins on generative computational systems in the early work of Turing before World War 2 and the sketch of CA developed by Ulam on the Manhattan Project in 1944. Our focus on theory relates to the fact that most CA models, with one or two exceptions (see Caruso et al. 2009) do not relate to the body of urban and regional economic theory which lies at the core of regional science, transportation modelling and urban location theory.
2   We do not underrate the good work done on fitting these kinds of models using map matching techniques (see for example Kuzera and Pontius 2008). Indeed the model we have used has a detailed map comparison toolbox embedded within it but there is still a logical inconsistency

between dynamic models that articulate their goodness of fit between observed and predicted patterns at a cross section in time rather than validating their fit through matching data on actual and simulated land development processes.

3 These forces act at a more local level than the economic forces that drive land uses to agglomerate or disperse – centripetal and centrifugal forces.

4 Some scholars refer to this method as hindcasting, others as backtesting, and both reflect notions of backcasting,

5 There are many quotes attributed to authors such as Defoe, Cobbet, Morris and other social commentators on the state of cities during these years (see http://www.casa.ucl.ac.uk/scatter/

6 The result of the re-sampling of vector to grid coverages was evaluated by overlaying both data layers and an on-screen visual inspection of resultant patterns at various spatial scales. In addition, we compared the total amount of land by land use class in both data formats. None of the tests indicated areas of concern as changes due to format conversion in both the geometry of patterns and volume of development we assessed as negligible for the scale and purpose of analysis.

7 In addition to these three factors, METRONAMICA allows the incorporation of another component called suitability. This factor accounts for the natural characteristics of the site on which the cell is located, such as topography, slope, soil type, flooding, etc. Due to the relatively minor variations of these characteristics within our study area, we decided to exclude this factor. It should be pointed out, however, that suitability could have significant influence on land patterns in cases where variations in the environmental characteristics are more strongly pronounced.

8 These statistical tools are part of Metronamica's Map Comparisons Kit, which provides automated calculations of a range of measures of fit such as Kappa and Percentage-of-agreement as well as some more advanced algorithms such as Fuzzy Set Map Comparison and Hierarchical Fuzzy Pattern Matching.

9 For instance, many of the model runs that produced the highest scores on percentage-of-agreement between modeled and recorded (actual) land use configurations generated patterns that showed greater spatial misfits in terms of general patterns and overall spatial composition.

References

Almeida C M, Monteiro A M V, Câmara G, Soares-Filho B S, Cerqueira G C, Pennachin C L, and Batty M 2005 GIS and remote sensing as tools for the simulation of urban land-use change. *International Journal of Remote Sensing* 26: 759–74

Barredo J I, Kasanko M, McCormick N, and Lavalle C 2003 Modelling dynamic spatial processes: simulation of urban future scenarios through cellular automata. *Landscape and Urban Planning* 64: 145–60

Batty M 2008 Fifty years of urban modelling: Macro statics to micro dynamics. In Albeverio S, Andrey D, Giordano P, and Vancheri A (eds) *The Dynamics of Complex Urban Systems: An Interdisciplinary Approach*. Heidelberg, DE, Physica-Verlag: 1–20

Batty M 2005a Agents, cells, and cities: New representational models for simulating multiscale urban dynamics. *Environment and Planning A* 37: 1373–94

Batty M and Longley P 1996 *Fractal Cities: A Geometry of Form and Function*, London, Academic Press

Batty M 2005b *Cities and Complexity: Understanding Cities with Cellular Automata, Agent-Based Models, and Fractals*. Cambridge, MA, MIT Press

Batty M 1997 Cellular automata and urban form: A primer. *Journal of the American Planning Association* 63: 266–74

Batty M and Longley P 1994 *Fractal Cities: A Geometry of Form and Function*. London, Academic Press

Birkin B, Procter R, Allan R, Bechhofer S, Buchan I, Goble C, Hudson-Smith A, Lambert P, De Roure D, and Sinnott R 2010 Elements of a computational infrastructure for social simulation. *Philosophical Transactions of the Royal Society A* 368: 3797–812

Brail R K (ed) 2008 *Planning Support Systems for Cities and Regions*, Cambridge, MA, Lincoln Institute of Land Policy

Caruso G, Peeters D, Cavailhès J, and Rounsevell M 2009 Space–time patterns of urban sprawl: A 1D cellular automata and microeconomic approach. *Environment and Planning B* 36: 968–88

Centre for Advanced Spatial Analysis 2011 Virtual London. WWW document, http://www.casa.ucl.ac.uk/projects/projectDetail.asp?ID=48

Chen Q and Mynett A E 2003 Effects of cell size and configuration in cellular automata based prey–predator modelling. *Simulation Modelling Practice and Theory* 11: 609–25

Cheng J and Masser I 2004 Understanding spatial and temporal processes of urban growth: Cellular automata modelling. *Environment and Planning B* 31: 167–94

Clarke K C and Gaydos L J 1998 Loose-coupling a cellular automaton model and GIS: Long-term urban growth prediction for San Francisco and Washington/Baltimore. *International Journal of Geographic Information Science* 12(7): 699–714

Clarke K C, Hoppen S, and Gaydos L J 1997 A self-modifying cellular automaton model of historical urbanization in the San Francisco Bay area. *Environment and Planning B* 24: 247–61

Couclelis H 1985 Cellular worlds: A framework for modeling micro-macro dynamics. *Environment and Planning A* 17: 585–96

Crooks A, Castle C, and Batty M 2008 Key challenges in agent-based modelling for geo-spatial simulation. *Computers, Environment and Urban Systems* 32: 417–30

Epstein J M 2008 Why Model? *Journal of Artificial Societies and Social Simulation* 11(4): 12 (available at http://jasss.soc.surrey.ac.uk/11/4/12.html)

Forshaw J H and Abercrombie P 1943, 1944 *County of London Plan*. London, Macmillan

Guan Q, Wang L, and Clarke K 2005 An artificial-neural-network-based, constrained CA model for simulating urban growth. *Cartography and Geographic Information Science* 32 369–80

Hall P G 1989 *London 2001*. London, Unwin Hyman

Hall P G, Gracey H, Drewett R, and Thomas R 1973 *The Containment of Urban England: Urban and Metropolitan Growth Processes or Megalopolis Denied* (Volume 1): *The Planning System: Objectives, Operations, Impacts* (Volume 2). London, George Allen & Unwin

Jantz C A, Goetz S J, and Shelley M K 2003 Using the SLEUTH urban growth model to simulate the impacts of future policy scenarios on urban land use in the Baltimore-Washington metropolitan area. *Environment and Planning B:* 30: 251–71

Jacobs J 1961 *The Death and Life of Great American Cities*. New York, Random House

Kocabas V and Dragicevic S 2006 Assessing cellular automata model behaviour using a sensitivity analysis approach. *Computers, Environment and Urban Systems* 30: 921–53

Kuzera K and Pontius Jr. R G 2008 Importance of matrix construction for multiple-resolution categorical map comparison. *GIS and Remote Sensing* 45: 249–274

Lee D B 1973 Requiem for large-scale models. *Journal of the American Institute of Planners* 39: 163–78

Li X and Yeh A G O 2000 Modelling sustainable urban development by the integration of constrained cellular automata and GIS. *International Journal of Geographic Information Science* 14: 131–52

Li X and Yeh A G O 2002 Neural-network-based cellular automata for simulating multiple land use changes using GIS. *International Journal of Geographic Information Science* 16: 323–43

Liu Y and Phinn S R 2003 Modelling urban development with cellular automata incorporating fuzzy-set approaches. *Computers, Environment and Urban Systems* 27: 637–58

Mandelbrot B B 1983 *The Fractal Geometry of Nature*. New York, W H Freeman

Menard A and Marceau D J 2005 Exploration of spatial scale sensitivity in geographic cellular automata. *Environment and Planning B* 32: 693–714

Ordnance Survey 2011, Address Layer 2. WWW document, http://www.ordnancesurvey.co.uk/oswebsite/products/osmastermap/layers/addresslayer2/

Parker D C, Manson S M, Jansen M A, Hoffmann M J, and Deadman P 2003 Multi-agent systems for the simulation of land-use and land-cover change: A review. *Annals of the Association of American Geographers* 93: 314–37

Popper K R 1957 *The Poverty of Historicism*. London, Routledge

Rasmussen S E 1937 *London, the Unique City*. Cambridge, MA, The MIT Press

Schelling T 1969 Models of segregation. *The American Economic Review* 59: 488–93

Silva E A 2004 The DNA of our regions: artificial intelligence in regional planning. *Futures* 36: 1077–94

Stanilov K 2003 Accessibility and land use: The case of suburban Seattle, 1960–1990. *Regional Studies* 37: 783–94

Stanilov K 2002 Post-war trends, land use changes and patterns of suburban development: The case of Greater Seattle. *Environment and Planning B* 29: 173–95

Stiny G 1980 Introduction to shape and shape grammars. *Environment and Planning B* 7: 343–51

Straatmen B, White R, and Engelen G 2004 Towards an automatic calibration procedure for constrained cellular automata. *Computers, Environment and Urban Systems* 28: 149–70

Torrens P M 2001 Can Geocomputation Save Urban Simulation? Throw some Agents into the Mixture, Simmer and Wait . . . London, University College, CASA Working Paper 32, (available at http://www.casa.ucl.ac.uk/working_papers/paper32.pdf)

Torrens P M 2000 How Cellular Models of Urban Systems Work: 1. Theory. London, University College CASA Working Paper 28, (available at http://www.casa.ucl.ac.uk/working_papers/paper28.pdf)

Valuation Office 2011, Non-domestic Rates Data. WWW document, http://www.voa.gov.uk/business_rates/

Verburg P H, de Nijs T C M, van Eck J R, Visser H, and de Jong K 2004 A method to analyse neighbourhood characteristics of land use patterns. *Computers, Environment and Urban Systems* 28: 667–90

Ward D P, Murray(as in A T), and Phinn S R 2000 A stochastically constrained cellular model of urban growth. *Computers, Environment and Urban Systems* 24: 539–58

Wells H G 1902 *Anticipations: Of the Reaction of Mechanical and Scientific Progress upon Human Life and Thought*. London, Chapman & Hall

White R and Engelen G 2000 High-resolution integrated modelling of the spatial dynamics of urban and regional systems. *Computers, Environment and Urban Systems* 24: 383–400

White R, Engelen G, and Uljee I 1997 The use of constrained cellular automata for high-resolution modelling of urban land-use dynamics. *Environment and Planning B* 24: 323–43

Wilson A G 2010 Urban and regional dynamics from the global to the local: hierarchies, 'DNA', and 'genetic' planning. *Environment and Planning B* 37: 823–37

Yan L 2009 *Modelling Urban Development with Geographical Information Systems and Cellular Automata*. Boca Raton, FL, CRC Press

Yang X and Lo C P 2003 Modeling urban growth and landscape changes in the Atlanta metropolitan area. *International Journal of Geographic Information Science* 17: 463–88

Yeh A G O and Li X 2002a Urban simulation using neural networks and cellular automata for land use planning. In *Proceedings of the Symposium on Geospatial Theory, Processing and Applications*, Ottawa, Canada

Yeh A G O and Li X 2002b A cellular automata model to simulate development density for urban planning. *Environment and Planning A* 29: 431–50