Chapter 22
Automatic Generation of Residential Areas using Geo-Demographics

Paul Richmond and Daniela Romano

Abstract

The neighbourhood aspect of city models is often overlooked in methods of generating detailed city models. This paper identifies two distinct styles of virtual city generation and highlights the weaknesses and strengths of both, before proposing a geo-demographically based solution to automatically generate 3D residential neighbourhood models suitable for use within simulative training. The algorithms main body of work focuses on a classification based system which applies a texture library of captured building instances to extruded and optimised virtual buildings created from 2D GIS data.

22.1 Introduction

Virtual environments, games and serious games often require the use of large scale urban models which capture essential attributes of real world locations. For example to build a drug dealers serious game, that is virtual training environment for drugs enforcement officers and community support workers, the key requirements are that the training environment must demonstrate realistically linked residential neighbourhood areas as well as being automatically generated to alleviate the need for complex CAD style of modelling. In addition to this the models created need to be optimised for use within a game engine to make use of its animating facilities and platform, and hence should have a limited number of polygons and be easily ported to existing game engines.

The key focus of this paper is on examining how residential neighbourhood profiling systems (geo-demographics) can be used to model realistic neigh-

---

University of Sheffield
`paul@dcs.shef.ac.uk, d.romano@dcs.shef.ac.uk`

bourhoods and automatically produce 3D models for simulative training. Its novel contribution focuses on how geo-demographics can be used within a classification system to realistically create virtual neighbourhood areas. First the existing methods of model production are reviewed, and then the implementation of an extrusion and building based classification algorithm which links the geo-demographical and physical properties of building footprint data to an extendable texture library of captured building examples is described.

22.2 Previous Work

Existing work in generating city models can be partitioned into two distinct styles which correlate well with the two concluding approaches highlighted by Shiode [1]. The first centres around creating Geographic Information System (GIS) based descriptive models and the second focuses on more detailed constructive models generated either manually or automatically. The next section examines how these two differing methods have been used and highlights some of the problems and limitations with regards to the proposed system.

22.2.1 Descriptive Models using GIS data

Urban Planning models use Geographic Information System (GIS) data to create realistic models of real world locations using high resolution data sets. Typically with highly descriptive methods of city modelling aerial photography is used as a texture draped over an extruded model. In order to extrude the models accurately previous techniques have focused on the use of photogrammetry methods of building outline detection [2] and extrusion focusing on GIS building boundary maps [3,4]. Whilst these techniques have been successful the improved availability and structure of GIS data in new Extensible Markup Language (XML) formats has lead to improved model generation methods allowing extremely large datasets to be visualised even in real-time [5].

With the advantage of new GIS datasets there are also new problems introduced regarding copyright, none more apparent than with the Virtual London project [6, 7], which uses Ordnance Survey (OS) Mastermap data in combination with Infoterra's LIDAR (*Light-Imaging Detection and Ranging*) data to extrude buildings. Unfortunately the government funded project which aimed to 'help Londoners visualise what is happening to their city' [8] is currently unavailable to the public through its intended interface (the Google earth application) due to OS licensing actively protecting their intellectual property. Despite a national incentive launched from the UK newspaper the guardian who has proposed the 'free our data' campaign it has recently been suggested that talks between OS and Google have collapsed and it is unlikely that the virtual London model will be available to the public any time soon.

As an alternative to using GIS building boundaries Google offers a community based alternative through the use of Google Sketch-up for Google Earth. This open modelling approach allows the community of Google Earth users to quickly create 3D content, which is easily integrated over the core imagery and terrain data. Despite overcoming the requirement of using OS data neither the Google earth platform or the before mentioned building extrusion methods are particularly well suited to a first person perspective. The reason for this is a matter of resolution as typical aerial photography data is only available at several pixels per meter [9] making distantly viewed urban planning models visually acceptable but lacking in resolution suitable for detailed texture application. In order to create more high resolution and detailed models of cities (and buildings in particular) the attention of using GIS data has shifted to constructive modelling techniques which use either manual or automatic methods to specify building details.

22.2.2 Constructive Models using Manual and Automatic City Modelling

With regards to manual modelling the games industry has recently demonstrated its interest in creating realistic virtual cities through the recent games releases of Grand Theft Auto IV and The Getaway, which feature detailed models of the cities of New York and London respectively. The techniques used in this instance are likened by Peter Edward (senior producer) as 'like a western movie. They don't have wooden slates at the back but they are just the fronts' [10]. Whilst the London model makes the use of photorealistic textures captured during an 18 month period the budget and time required to first manually create a wireframe model from the photographic data and then map the photography is far beyond the reach of most modelling projects.

The use of automatic methods to generate virtual cities elevates the budgetary requirement of manually modelling a city but does however introduce a degree of guesswork in representing cities or buildings structure. Despite rule based systems being able to produce entirely procedural (pseudo) infinite cities [11] the models rely heavily on the rules specifying the grid like road structure and hence lack the realism of capturing an actual cities structure, making its application to training somewhat limited. Parish and Muller [12] describe a more intuitive alternative for procedural generation of cities which uses L-systems (inspired by Prusinkiewicz and Lindenmayer [13]) with a basic set of data defining land and water boundaries to subdivide into realistic roadmaps, land, lots and buildings. Although the results of this approach are relatively successful at defining a cities structure in terms of road layout they are limited in application towards building generation, as highlighted in a later paper by Wonka et al [14] which states that a buildings structure;

'does not follow a growth like process in the same way that plants and streets do, but instead follow a sequence of partitioning steps.' (p670)

For this reason the automatic generation of buildings tends to focus on the use of shape and split grammars.

22.2.3 Building Generation

Split grammars originally introduced by Stiny [15] are used by Wonka et al. [14] as a grammar operating on shapes. By using the split grammar approach, basic shapes (referred to as building blocks) are then recursively replaced by further shapes which have attributes that describe the material and help to determine further subdivision steps. In addition to the extensive set of split grammar rules the notion of a control grammar is also introduced. The aim of this control grammar is described as a way to specify design decisions spatially in a way that corresponds to architectural principles. More simply this control grammar is used to set attributes in a spatially related way ensuring that split grammar rules are propagated in a way which represents realistic architecture. Despite the visual success of this method its complexity requires collaboration between a designer (with a clear understanding of the workings of a split grammar) and an architect, which has lead to the development of similar but slightly less complex split grammar implementations such as the one from Larive [16] who concentrates on 2.5D building frontages for extruded buildings.

More recent work by Muller et al [17] extends the CityEngine application to apply the idea of a shape grammar to create building scripts capable of capturing intricate detail. This is achieved by combining the previous split grammar approach with a complex mass modelling system allowing more complex primitives to be split without the previous limitations of axis aligned shapes. In addition to this the CGA shape system allows the division and placement of other aspects of the environment allowing the generation of pathways and placement of greenery including trees and shrubs.

Although there is a considerable amount of work focusing on the generation of buildings it is important to consider the full spectrum of city generation. Larive et al. [18] uses the notion of 'Urban Zones' to present a hierarchical division of seven stages of city generation; namely, Urban Zones, Road Networks, Blocks, Lots, Exteriors, Building Plans and Furnished Buildings. Whilst the majority of research is concerned with stages 1 through to 5 the focus is primarily on the building stage and concentrates on creating realistic buildings. Whilst this is important in generating a residential area little consideration is given to the creation of realistic neighbourhoods which in the case of the previously described project is more essential than the exact representation of building models.

22.3 Implementation

22.3.1 Overview

Whilst previously successful implementations are able to build realistic virtual cities and neighbourhoods either procedurally or through using attributes to control architectural variance there is little attention paid to the importance of architectural discrepancies between certain types of housing estates. In addition to this, methods of texture application are either limited to being highly descriptive, in which a large amount of time must be invested to capture buildings or procedural, which although may not require the same investment of time regarding texture capture certainly requires in depth knowledge of architecture and grammars.

The following section of this paper focuses of the implementation of a system which offers the following contributions;

1. Shows how geo-demographics can be used within a classification system to build realistic virtual neighbourhood areas.
2. Demonstrates how an extendable captured building library and classification system can be used to crudely determine a buildings appearance and hence define visual zoning of houses

By achieving the above it is our intention to offer a proof of concept of how geo-demographics can be used to make assumptions of vast city models. In addition to this our work will demonstrate a method somewhere between the descriptive and constructive methods previously described which allows buildings to be represented realistically in a way which creates observable neighbourhood areas. The subsequent section proceeds as follows; Section 3.2 discusses the general extrusion process of the buildings from the floor plan data including how the extrusion process, roofing, polygon face tessellation and any assumptions made of the building structure. The role of captured building textures or Building Instances (BI's) is then discussed in section 3.3 which describes how these are used to influence a buildings visual appearance. Finally section 3.4 describes the classification system used to determine the suitability of each member of the BI library with regards to each building in the dataset.

22.3.2 Building Structure & Extrusion

As there is no guarantee that LIDAR data is available for the areas being extruded a method combining a LIDAR data and floor height estimation is used. This is implemented by pre-processing any ASCII LIDAR data (which could be replaced by alternative ASCII data such as cadastral or observed building heights) to calculate an average building height for each building structure in the data set which is then stored in a lookup table. Where LIDAR data is unavailable some assumptions are made towards the heights of buildings. As a rough guide the square of the buildings ground floor surface area is used to calculate a preliminary height, either the averaged LIDAR height or estimated height is rounded to an integer number of floors by division by the standard building floor height of 2.75m. As the majority of buildings within a residential area constitute two and three floor family homes and flats, most buildings have a primary floor calculation of between two and three floors before any adjustments are made to the extreme exceptions. If the calculated number of floors is however less than one then it is assumed that the structure is not a building and is ignored (it is not uncommon that OS classify out buildings such as barns, greenhouses, temporary buildings and some garages as buildings). Likewise buildings with an estimated floor height greater than fifteen are assumed to be larger than average residential dwellings and are reduced to a warehouse/ council flat style of building with at most three floors and a flat roof.

Where as exiting techniques have used split grammars to split a wall face into smaller face areas, allowing procedural textures, the technique employed in this approach is more commonly used within game level design. In particular this method is used in the half life 2 game using the source engine (http://www.valvesoftware.com) for specifying wall brushes (roofs are however described by model entities). In order to reduce the number of faces for each building, a building is first split into floors and then each wall is split (where possible) into a grid of uniform squares. The remainder of the wall which does not constitute a full squares horizontal length is then split evenly between the walls first and last wall section symmetrically. After the walls have been split into wall sections the texture coordinates are then calculated presuming that a square texture is to be applied. For the smaller split wall edge sections the texture coordinates are calculated allowing tessellation where the section meets a full square. This method, although not requiring the manual specification of a split grammar or likewise does require that a reasonable texture library is available providing a rich background of textures to create suitably complex building facades.

Like previous methods of building extrusion and roof generation [19] we have followed Felkel and Obdrzalek's Straight Skeleton Implementation method to apply a simple hipped roof which has been extended to allow a gable/cross gable roof style which is applied to the majority of buildings.

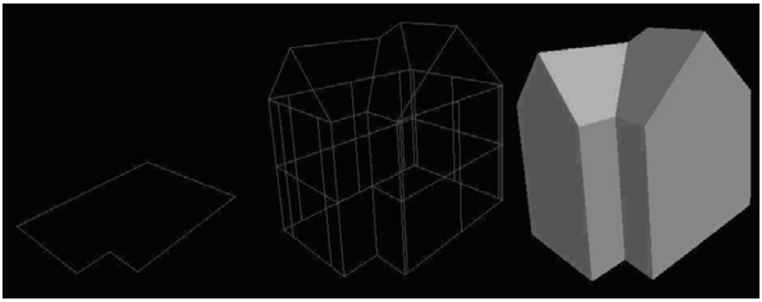

Fig. 22.1 The wireframe example in demonstrates how a building is extruded, split and applied a gable roof

22.3.3 Building Instances

Although there has been significant research into the use of grammatical approaches of texture application to buildings this is not the focus of this paper and as such a simple method has been used to apply textures to buildings. This is achieved by considering a number of captured building textures from the real world which make up what is described in this paper as a Building Instance. Although the choice of which buildings should be captured for the library is somewhat subjective, this method allows the user to capture buildings which stereotype the surrounding area well (see figure 3 for an example of houses from differing ACORN types) and chose buildings where foreground distraction are not present. The subjectivity of BI choices is also considered in section 4.3 which considers the accuracy of the classification for each building in the data set. In addition to capturing a number of textures for any captured building, a number of physical building properties are also recorded for use within the classification system described in section 3.4. For each captured building (or BI) in the texture library, a square texture must be provided for description of a plain brick wall, a windowed area, a doorway and a gable roof section (if the building is flat roofed this is not required). A simple list based description of the building for the ground and above ground floors is then required to provide a base to map each texture to the buildings within the dataset. An example of a captured buildings data is given below where the buildings structure consists of a two floor and four wall structure, where each of the four walls contained room for two descriptive wall elements.

In order to apply the texture to buildings of varying size the description is simply repeated around the building.

Fig. 22.2 The example building above has a texture description for brick, a doorway, a window and a gable roof section

Fig. 22.3 Example of two Sheffield houses in close proximity and of a similar structure but with varying ACORN types

22.3.4 Neighbourhood Profiling and Geo-Demographics

The socio-demographic system which has been used in this work is the ACORN (A Classification of Residential Neighborhood) classification system from CACI Limited that contains a hierarchy of classifications spanning down from category, group and type. The categories range from wealthy achievers to hard pressed with the group level giving further definition such as for the wealthy achievers category; wealthy executives, affluent greys and flourishing families. The ACORN type ranges from type 1) ‘Affluent mature professionals’ in large houses to type 56) ‘Multi-ethnic crowded flats’. For each classification there are a number of lifestyle topics and demographic topics which are used to classify UK postcodes. Whilst the lifestyle topics are concerned with interests such as internet cars and shopping the demographic topics consider in addition to other interests dwelling height, size, house hold structure and tenure all of which are particularly useful with regards to classifying the style of a building or neighbourhood. In addition to loading a GML (Geography Markup Language) data file in the World Generator application a table of postcode acorn values must be provided where the postcode is specified as a spot location. In order to calculate the ACORN classification for any building loaded into the system the buildings proximity to each postcode spot location is considered and the appropriate ACORN type code is returned.

In order to determine the most appropriate look of a neighbourhood and hence each building within it, each building read from the GML data is compared to each of the building instances which contain the following additional information which is pre loaded with the textures;

- Name – for reference
- Floors – the number of floor levels above ground level of the building
- Surface area – the surface area of building in square meters at ground level
- External walls - number of external walls (i.e. complexity of the building)
- Roof type – at the moment this is limited to either flat or gable
- ACORN category – the acorn category within the range of 1 to 5
- ACORN group – the acorn group within the range of 1 to 17
- ACORN type – the acorn type within the range of 1 to 56

Buildings and instances are then compared directly by using a distance measure of similarity for each of the above attributes (excluding name). The distance measure is a weighted combination of the summed distance measures between each of the different attributes. This is formalised below by the following expression.

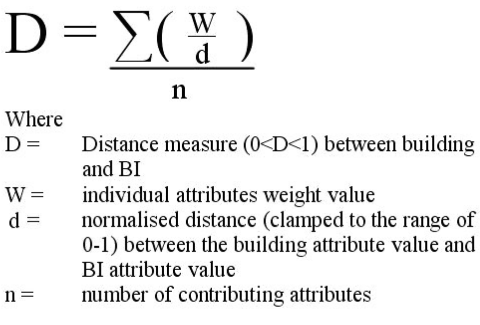

$$D = \frac{\sum\left(\frac{W}{d}\right)}{n}$$

Where

D = Distance measure ($0<D<1$) between building and BI

W = individual attributes weight value

d = normalised distance (clamped to the range of 0-1) between the building attribute value and BI attribute value

n = number of contributing attributes

By considering each attribute in this way (excluding the roof type, ACORN category and ACORN Type attribute) a value between 0 (infinitely dissimilar) and 1 (exactly similar) is attributed to each attribute. Unlike numeric attributes the roof type attribute which can have two values is either the same (value of 1) or different (value of 0). By considering more than just the ACORN classification building instances with architectural variances but sharing the same ACORN type can therefore co-exist and be attributed to buildings within the model providing realistic variances within the neighbourhoods themselves.

22.4 Results & Conclusion

In order to test the building classification method above and its ability to create realistic neighbourhoods a limited number of three building instances were created and applied to a small GML data set of Basingstoke which consisted of four varying ACORN classifications.

22.4.1 Dataset & Performance

The GML data set which was used consists of 23037 unique topographic elements of these 6005 are topographic areas (others mainly constitute topographic lines and point data which are usually a repetition of fully defined polygon areas of unclosed polygons which are not used). Within the 6005 topographic areas there are eighteen unique featured codes each with a different combination of the hierarchical elements; theme, descriptive group and descriptive term (a specification of the OS GML specification is available on their website http://www.ordnancesurvey.co.uk), of these eighteen the key feature code (10021) which is used for building generation has the theme ‘Building’ and descriptive group ‘Buildings’ (there are 2354 in total). All other features within the GML set are treated according to an XML file which can set the following properties of the hierarchical theme, descriptive group and descriptive term elements.

- materialColour - The surfaces rgb material colour in format 'r,g,b' where r,g,b are java ints
- extrusion - The height in km of a vertical extrusion of the layer above ground level (java double format)
- texture - The texture name to be used to texture objects in the layer
- cleanUpPointTolerance - The tolerance used to clean up points which do not specify any additional detail (i.e. cause an angular difference between the previous and next point of the tolerance value).

This allows the buildings to be set within an environment which displays real world boundaries such as road, pavement and housing boundaries as well as defining areas of the natural environment such as treed areas and scrubland.

In order to generate the 3D model the system operates in two stages first all data is read from the GML file using a SAX parser and stored as a simple object containing the objects properties and polygon points. After all objects are read the surfaces are then extruded, textured, cleaned up and in the case of buildings, constructed using the described method. Running on a Pentium Xeon 2.33 Mhz with 2 Gb of RAM the reading stage takes roughly 3.4 seconds and the extrusion stage takes 2.1 seconds with the production of 195375 faces. Currently all objects are in the visual system are stored in memory for the purposes of displaying them, however it is possible to generate the buildings serially and output them directly to .obj format using a command line interface.

22.4.2 Discussion

As demonstrated by figure 4 and figure 5 (which display only the building objects) the effect of using a geo-demographic classification system creates neighbourhoods which reflect the geo-demographic area and assigns buildings a building instance which closely reflects the buildings structure. Despite using only three building instances in this example the complexity of the final model regarding distribution of the building instances is indicative of a realistic neighbourhood. The extendibility of the system also makes the integration of variances in a particular area a matter of simply creating a number of new building instances which reflect new building styles.

In addition to the creation of realistic neighbourhoods, the building/object extrusion and texture application methods create a realistic model of a residential area which despite the lack of integration of decorative objects such as trees, lampposts, etc. provides an ideal platform for a realistic training environment. The use of real GIS data to create the model not only means that the road layouts and housing placements are rational but also offers the possibility of using further GIS data to power the game logic and intelligence models within a simulation.

In order to make a direct comparison of both the classification algorithm and the texturing system a direct comparison is made in figure 7. Although there are some discrepancies, the orientation of the automatically generated terrace roofing being one, the virtual representation demonstrates the correct application of building textures. Despite the texture restriction of only tessellating square textures the low cost texturing method allows the virtual buildings to provide an accurate representation or their real world counterparts.

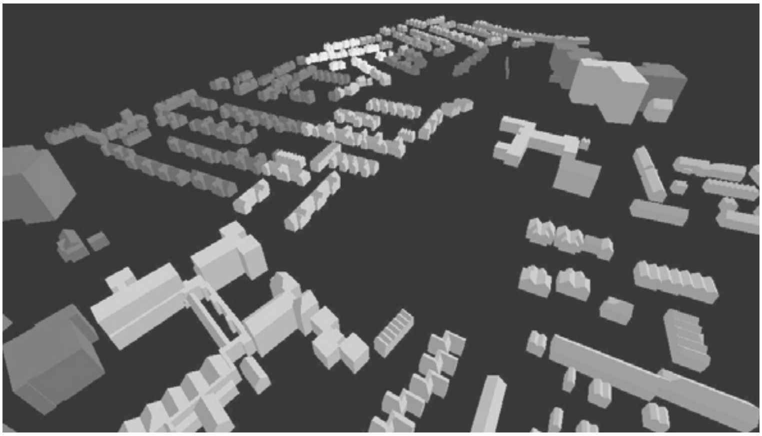

Fig. 22.4 ACORN Types

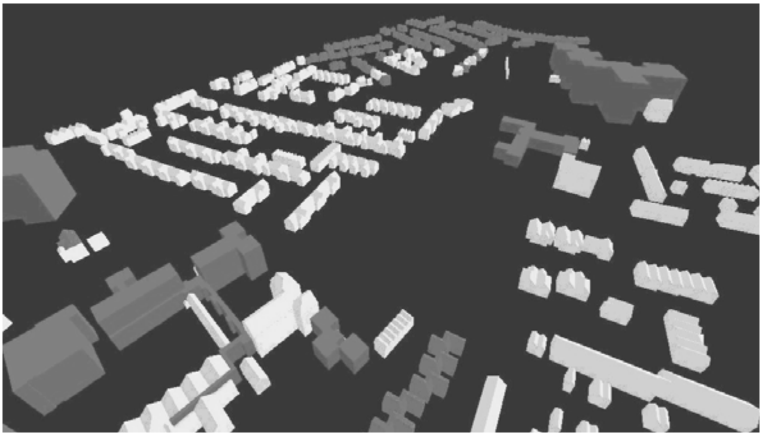

Fig. 22.5 Building Classification

22.4.3 Conclusion and Further Work

The generation of virtual cities seems to have a clear division between descriptive models generated using aerial photography combined with geographic data such as LIDAR and models built using constructive methods such as modelling and shape grammars. Whilst constructive methods offer the only realistically detailed solution of model generation suitable for first person games, the models produced begin to differ from the real world as a function of the time spent either physically modelling, constructing descriptive grammars [14] or manually applying ground captured textures.

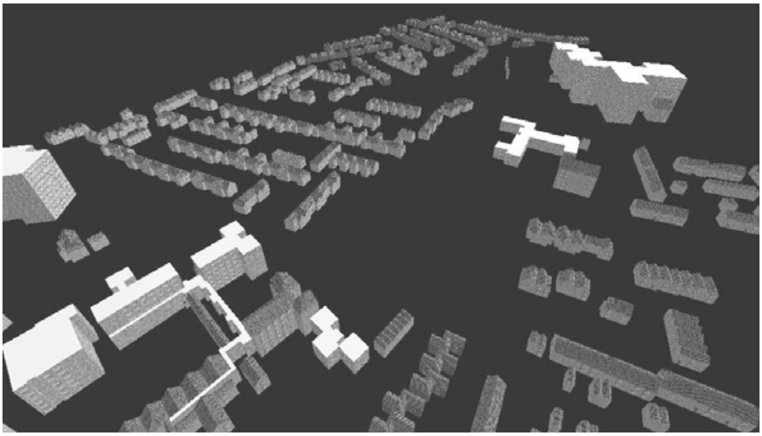

Fig. 22.6 Textured Buildings

The solution offered by this work proposes a method which meets the two objectives set out in section 3.1 by allowing the generation of virtual residential areas and by proving an extendable texturing system. Whilst the building creation methods do not differ greatly from existing techniques the integration of geo-demographics and a classification system allows assumptions to be made of building appearances whilst still maintaining realistic virtual neighbourhood areas. The main weakness of the method described in this paper is the subjective choice of the buildings captured (BI's) within the texture library as a simple extension the system has been implemented to display a buildings colour as a function of its classification accuracy (figure 8). This therefore provides an excellent indication of where extensions to the BI library should be made to improve the model.

With respect to spatial extendibility, it is fair to assume that without the addition of geo-spatial attributes to the classification system it is unlikely that buildings from one area in the UK could be realistically represented by those from another (although the boundaries of neighbourhood areas may still be realistically defined). On a larger scale it is therefore suggested that geo-spatial attributes are included a part of future work. In addition to this it would also be interesting to consider using a more advanced method of texture application for each BI. As a grammar system would be required to achieve this there would be a significantly longer cost of time in preparation of the BI library however it is likely that the building would achieve a more realistic and less granular appearance (a trade off that was made for generation speed and reduced polygon counts in the current implementation).

Fig. 22.7 Image of Sheffield road compared to Virtual Representation (above with sky post processed)

References

1. N. Shoide, 3D urban models: Recent developments in the digital modelling of urban environments in three-dimensions, Centre for Advanced Spatial Analysis, University College London, GeoJournal Volume 52, Number 3, pp 263-269, 2000

Fig. 22.8 Classification strength of buildings in Sheffield example

2 I. Suveg, G. Vosselman, 3D reconstruction of Building Models, Int. Archives of Photogrammetry and Remote Sensing, vol XXXIII, part B2, pp. 538-545, 2000

2. R. Tse, M. Dakowicz, C.Gold, D. Kidner, Automatic Building Extrusion from a TIN model Using LiDAR and Ordnance Survey Landline Data, University of Glamorgan, Treforest, Mid Glamorgan
3. R. Laycock, A. Day, Automatically generating large urban environments based on the footprint data of buildings. In IProceedings of the Eighth ACM Symposium on Solid Modeling and Applications, ACM Press, New York, NY, pp. 346-351, 2003
4. A. Steed, S. Spinello, B. Croxford, Richard Milton, Data Visualization within Urban Models, Theory and Practice of Computer Graphics 2004 , pp. 9-16, 2004
5. A.Steed and E.Frecon, "Building and Supporting a Large-Scale Collaborative Virtual Environment", Proceedings of 6th UKVRSIG, University of Salford, pp 59-69, 1999
6. A.Steed, E.Frecon, D. Pemberton and G. Smith, "The London Travel Demonstrator", Proceedings of the ACM Symposium on Virtual Reality Software and Technology, pp 50-57, ACM Press, ISBN 1-58113-141-0, 1999
7. Online Article, `http://technology.guardian.co.uk/weekly/story/0,,1981821,00.html`
8. Wikipedia Link, `http://en.wikipedia.org/wiki/Google_Earth#Resolution_and_accuracy`
9. Blog article on "The Getaway" PS2 Game, `http://digitalurban.blogspot.com/`
10. S. Greuter, J. Parker, N. Stewart, G. Leach, Real-time procedural generation of `pseudo infinite' cities, In Proceedings of the 1st international

Conference on Computer Graphics and interactive Techniques in Australasia and South East Asia, 2003

11. Y. Parish, P. Müller, Procedural modeling of cities, In Proceedings of the 28th Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH '01, pp 301-308, 2001
12. P. Prusinkiewicz and A. Lindenmayer, The algorithmic Beauty of Plants, 1991
13. P. Wonka, M. Wimmer, F. Sillion, W. Ribarsky, Instant Architecture, ACM Transactions on Graphics, Volume 22, no. 4, pp. 669-677, 2003
14. G. Stiny, Pictorial and formal aspects of shape and shape grammars: on computer generation of aesthetic objects, Birkhauser, 1975
15. M.Larive and V. Gaildrat, Wall grammar for building generation, Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia, pp429-437, 2006
16. P. Muller, P. Wonka, S. Haegler, A. Ulmer, L. Gool, Procedural Modeling of Buildings, in Proceedings of ACM SIGGRAPH 2006 / ACM Transactions on Graphics (TOG), ACM Press, Vol. 25, No. 3, pp 614-623, 2006
17. M. Larive, Y. Dupuy, V. Gaildrat, Automatic Generation of Urban Zones, WSCG'2005, Plzen, Czech Republic, 2005
18. R. Laycock and A. Day, Automatically generating roof models from building footprints, The 11-th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision, 2003