

Raster analysis of urban form

Carlo Ratti

School of Architecture and Planning, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA; e-mail: ratti@mit.edu

Paul Richens

Department of Architecture, University of Cambridge, Cambridge, CB2 2EB, England

Received 11 September 1999; in revised form 3 May 2003

Abstract. A new paradigm for investigating the environment consequences of urban texture is proposed. Using raster-based models and software algorithms derived from image processing the authors develop efficient methods of measuring geometric parameters and predicting radiation exchange. The possibilities of generating synthetic urban textures and of integrating cellular automata are explored. Results suggest the possibility of a raster-based urban model to inform planning and design.

1 Introduction

Over the past few decades many computer tools have been developed to assist in the environmental design of individual buildings. Heat, light, sound, and especially energy consumption can be analysed in many different packages.

This is not generally true for urban design, especially at the medium scale. Although it is widely assumed that urban texture—the pattern of streets, building heights, open spaces, and so on—will determine environmental quality both in the buildings and outside, tools for investigating the connections are sparse. The need for medium-scale understanding is confirmed by Givoni (1989, part 1, page 5):

“The outdoor temperature, wind speed and solar radiation to which an individual building is exposed is not the regional ‘synoptic’ climate, but the local microclimate as modified by the ‘structure’ of the city, mainly of the neighbourhood where the building is located.”

Apart from Hillier’s (1996) space-syntax techniques for investigating urban configuration, there is a lack of tools to connect urban texture to the resulting urban quality. This is partly because of difficulties in describing the urban environment. The representation of the city needs “a geometry of order on many scales, a geometry of organised complexities”, state Batty and Longley (1994), subsequently arguing for the replacement of traditional Euclidean geometry with fractal geometry. The extension of architectural analysis software to the urban scale fails because of the problems of building precise geometric models and the difficulties of processing, as the time is likely to depend on the square of the number of geometric elements.

Our interest is in using very simple raster models of urban form.⁽¹⁾ By borrowing techniques from image processing we have been able to develop a range of promising algorithms for which the processing time is independent of geometric complexity and relates linearly to the area under investigation. In particular, shadow casting, the calculation of sky–ground view factors, the replication of urban texture, and the simulation of pollution migration are presented here. This list should not be considered as comprehensive but simply as a partial review aiming to make a proof of concept of the potential of image-processing techniques in the analysis of urban texture.

⁽¹⁾A ‘raster model of urban form’ may be defined as a two-dimensional, rectangular, pixel grid (such as a digital image) storing some urban attributes (such as building height).

2 Relation to previous work

The analysis of raster maps with image-processing techniques is not new. It has been used in the geosciences for many years, with very successful results in the context of raster geographical information systems (GIS; for instance, see Burrough and McDonnell, 1998). Image-processing techniques are also regularly used in remote sensing, where they allow the extraction of features from collected digital imagery (see Jensen, 1986). Surprisingly enough, however, they have almost never been tested in the urban context—probably because of the difficulties that have existed so far in obtaining three-dimensional datasets in urban areas. A few examples can be found in our earlier work (Ratti, 2002, Richens, 1997). In this paper we build on that line of research and aim at presenting an introductory review of image-processing techniques applied to urban analysis for the planning community.

3 Image-processing basics for urban analysis

The simplest possible representation of urban texture is a figure–ground map such as that shown in figure 1(a). If scanned, the resulting binary image tells us, for each pixel, whether the ground is built on or is open ground. If height information is included we obtain what geographers call a digital elevation model (DEM), which is an image where each pixel has a grey level proportional to the level of the urban surface. A DEM of an urban area in central London is shown in figure 1(b). This DEM was produced with a low-technology approach, based on scanning a digital ground map and colouring each pixel with a level of grey proportional to height (which was estimated from the number of storeys). Today, however, high-resolution remotely sensed DEMs are becoming increasingly available at low cost as a result of the rapid development of techniques such as laser altimetry or synthetic aperture radar (SAR). They can be commercially purchased for all major urban areas.

Urban form represented as a DEM can be analysed with image-processing techniques using a simple package such as NIH Image.⁽²⁾ Many of the ideas in NIH Image should be familiar from popular programs such as Adobe Photoshop. The difference is one of emphasis: Photoshop is concerned mainly with visual appearance, whereas NIH Image concentrates on measurement and analysis.




Figure 1. (a) Urban ground map showing built (black) and unbuilt (white) areas and (b) urban digital elevation model (DEM) in which increased grey scale corresponds to increased height, for a selected case study in central London.

⁽²⁾NIH Image is a public domain image processing software developed by the US National Institutes of Health.

The basic data structure is the image, a two-dimensional (2D), an array of pixels, each storing a single byte (8 bits), typically understood as an integer in the range 0–255. Closely associated with the image is its look-up table (LUT), which translates each of these integers into a colour that is displayed on the screen. The LUT can be freely changed, changing the appearance of the image, without damaging the underlying data stored in each pixel. Image and LUTs can be written to disk as a unit, typically as a TIFF file. Images usually originate with a scanner; video frames can be captured directly if the computer has an audio visual (AV) board. Images can be originated or edited on-screen by using the usual paint-box tools for line drawing, selection, fill, brushing, and erasing.

A large part of the functionality in NIH Image is aimed at image enhancement, using two major techniques. The first is LUT modification, which can alter the brightness and contrast of, or even apply false colour to, the pixel values. For instance, thresholding or density slicing can be performed by LUT modification. Thresholding turns the grey-scale image into a black- and-white image, with interactive control of the threshold value. Density slicing turns a narrow range of pixel values into another colour, again with interactive control. By dragging the slice through the range one can obtain a precise way of exploring the values in the image. Both thresholding and density slicing are performed by LUT modification, but it is possible to force the changes into the image itself, which then becomes a binary image with values of 0 (false, white) and 255 (true, black) only.

The second method of image enhancement is by applying filters. The fundamental idea behind image filters is the kernel, a small window or region of pixels (say 3×3 , or 5×5). The filter operates by centring the kernel on each pixel in the image in turn and performing some computation on the values seen under the kernel. The central pixel is replaced by the result of the computation. So, for example, the ‘max’ filter replaces each pixel by the maximum value in its 3×3 neighbourhood, and the ‘median’ filter replaces each pixel by the median (central) value of its neighbourhood (which is a popular way of removing noise from an image).

Convolution filters work by forming a weighed average of the pixels under the kernel; the kernel cells specify the weights. If the weights are all positive, the result is some sort of smoothing or blurring. More sophisticated kernels, with some negative weights, can sharpen an image, differentiate it, or detect edges. The latter operation can be performed with the so-called Sobel kernels, which construct the derivatives in the y and x directions, f_y and f_x , respectively.

$$\begin{array}{cccccc} 1 & 2 & 1 & -1 & 0 & 1 \\ 0 & 0 & 0 & -2 & 0 & 2 \\ -1 & -2 & -1 & -1 & 0 & 1 \end{array}$$

NIH Image provides some all important features for doing arithmetic on complete images (for example, by adding or multiplying corresponding pixels or taking the maximum or minimum values). The result is a completely new image. In general, these operations are computed at higher precision, scaled and offset by a stated amount, and then clipped to the 0–255 range allowed in each 8-bit pixel.

4 Built-form statistics

Given a simply constructed DEM, NIH Image was able to measure most of the usual built-form statistics. Histograms of building height, arbitrary cross-sections, and axonometric representations [figure 2 (see over)] are standard commands. A very simple macro to measure the whole image by summing the pixel values (with suitable calibration corrections) gives the built volume. We obtain the built area by thresholding the DEM




Figure 2. Axonometric view of the central London digital elevation model.

at $z > 0$, and counting. This works because the ground profile in the sample areas is essentially flat and at zero throughout. If this were not the case we could use a background-removal function to find ground level, which could be subtracted from the image before measuring areas and volumes. By using thresholding to extract an image at any particular level we can repeat area measurements at any horizontal cross-section.

Measurement of external wall area (or the thresholded equivalent, perimeter length) proved more challenging. There is a command to extract edge pixels from the image, which seems a good start. However, summing them and multiplying by the scale does not necessarily give the right answer. Walls aligned to the axes measure correctly, but walls at 45° measure short, by a factor of $2^{1/2}$.

A more accurate answer can be obtained by using the Sobel edge-detector described above. The image is convolved separately with two kernels that construct the derivative in the y direction, f_y , and the derivative in the x direction, f_x . These derivatives are equal to eight times the projected area of the wall in each direction. The total area is therefore obtained by forming the square root of the sum of the squares of the two derivatives [see figure 3(a)]. If the convolutions are carried out separately it is also possible to use the results to classify edges by their orientation, computed as $\text{atan}(f_y/f_x)$ [see figure 3(b)].

Deep-plan areas (those at a certain distance from the perimeter, that are significant consumers of energy as they must have artificial light and ventilation) can be displayed using density slicing to extract the level at which to work and then applying a standard image-erosion function that strips away edge pixels to leave only the deep core.




Figure 3. (a) Edge detection and (b) east–west facade detection for the central London digital elevation model.

5 Urban environmental analysis

5.1 Shadow-casting algorithm

Daylight and sunlight are fundamental constituents of the microclimate. Surprisingly effective algorithms can be devised, using image-processing macros, to compute them from DEMs.

Our central algorithm (Richens, 1997) computes shadows for an arbitrary angle of lighting. The approach taken is to compute shadow volumes as a DEM, that is, the upper surface of the volume of air that is in shadow. We start by defining the three components of the vector pointing towards the sun. Then, we compute the components of an opposite vector, scaled so that the larger of the x and y components is just 1 pixel, and the z component is adjusted to the image calibration. If we translate the DEM by the x and y components, and simultaneously reduce its height by subtracting the z component, we get part of the shadow volume. If we continue translating and lowering by multiples of this vector, and take the maximum of this volume with that previously calculated, we build up the whole shadow volume. The process can be stopped when all levels are zero or the translation has shifted the volume right off the image. This is easily encoded as an image-processing macro. The results are shown in figure 4.




Figure 4. (a) Shadow volume and (b) shadow results for the central London digital elevation model.

To reduce the shadow volume to an actual map of shadows on the roofs and ground of the city, the original DEM is subtracted from the shadow volume. Pixels with negative or zero value are in light; positive values are in shade.

This calculation computes shadows for an arbitrary lighting angle; it is straightforward to add a procedure to calculate shadows from the sun for any given latitude, time of year, and time of day, using the usual astronomical formulae. The next stage is to calculate solar irradiation, the actual amount of energy received at each pixel, from the sun, taking into account the solar altitude, angles of incidence, and shadowing. Computation of the clear-beam normal surface irradiation can follow a standard method—we used that in Page (1986).

These algorithms are simple to write and are impressively fast. The running speed depends on the number of pixels in the image, but not on the geometric complexity, so it is feasible to process acres of city at a time—something unthinkable with traditional geometric models.

5.2 View factors

The view factor is a geometric parameter that is related to many urban environmental processes. Also called shape factor, form factor, or configuration factor, it was originally introduced in heat-transfer theory to model radiant exchange between diffuse surfaces. Later, it was successfully applied to urban climatology (Oke, 1988).

In particular, the view factor from the city to the sky, which represents the openness of the urban texture to the sky, takes part in three major environmental processes: the loss of longwave radiation from city to sky, responsible for the cooling down of the urban surface during the night; the absorption of solar shortwave radiation, causing daytime heating; and the illumination received on a horizontal plane from an overcast sky, which determines daylight potential.

Several methods have been suggested in urban climatology for deriving the view factor from the city to the sky. Johnson and Watson (1984) derived an analytical expression to determine the sky view factor for different urban canyon arrangements. Steyn (1980) suggested a graphical method based on fisheye-lens photographs.

We calculate the view factor by repeated application of the shadow algorithm (figure 5). We simply compute the shadows for a large number of light sources, distributed over the sky, and for each pixel count the number of times they are in light. So, if we use 255 samples, any pixel for which the count is 255 can 'see' all the sky and has a sky view factor 1, whereas a count of 0 means that it cannot see the sky at all.




Figure 5. Sky view factors for the central London digital elevation model.

To get meaningful results, it is necessary to distribute the sample points over the sky in the correct manner. If a uniform distribution is used, then what is measured is not the view factor but the solid angle of sky visible from each point. However, if we distribute our points evenly over a unit circle in the horizontal plane, and then project up to a unit hemisphere, we obtain a cosine-weighted distribution, which is the correct distribution for computing the view factor from the city to the sky.

5.3 Informing urban design: an example

An example of how the above techniques could be used to inform planning and design is given here. A more complete review can be found in Steemers and Ratti (1998; 1999).

The Martin Centre, University of Cambridge, and Cambridge Architectural Research (England), acted as research consultants for an urban development in Greenwich, London, by the Richard Rogers Partnership. The image-processing analysis of DEMs was used to assess the environmental potential of different design schemes.

Most of the analysis focused on the comparison of daylight and sunlight conditions on six urban options figure 6. For each option, images with shadowing contours were produced, showing on each pixel the number of hours of shadow (figure 7).

A simplified parameter, the mean shadow density, was determined for each image. This parameter can vary by approximately 19% in winter and by approximately 21% in summer, following the geometry of the development. Furthermore, the summer maximum does not coincide necessarily with the winter maximum. It was therefore possible to minimise shadowing in wintertime, when the sun is beneficial, while maximising shadowing in summertime. The architects used these results in the preliminary phase of design (or at least we like to think so!) in order to make decisions on which scheme to develop.

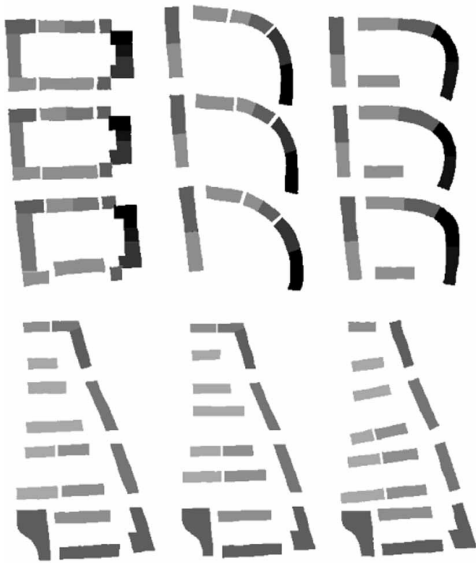


Figure 6. Sketches of development options for the Greenwich peninsula (Richard Rogers Architects Limited) represented as a digital elevation model.




Figure 7. Hours of shadow during summer for each development option presented in figure 6: the image has been obtained by casting shadows at hourly intervals on 21 June, summing all the results and finally applying an edge-detection filter.

6 Urban morphological analysis with DEMs

6.1 Urban ground view factor

Peponis et al (1998) proposed an extension to the theory of space syntax (Hillier 1996; Hillier and Hanson 1984), analysing the visible relations between the surfaces that bound and subdivide space. The study of plan configuration according to the covisibility of surfaces and the introduction of parameters such as visual connectivity would, they suggested allow a more subtle investigation of spatial configuration than Hillier's original space syntax, which is based mainly on the topology of circulation spaces.

We have found DEM techniques highly adapted to the analysis of these measures. We defined a parameter called the 'urban ground view factor' and computed its value on each pixel by modifying the macro for the sky view factor (see section 5.2). The urban ground view factor is a measure of the extension of the view of the ground from each urban pixel and relates directly to the strong visibility polygon (O'Rourke, 1987, Peponis et al, 1998). Results on the selected London case study are shown in figure 8.




Figure 8. Urban ground view factors on a selected portion of the central London digital elevation model.

In addition to the urban ground view factor, other configurational characteristics can be calculated from the DEM. For instance, it is easy to write Matlab functions to compute the following parameters suggested and applied by Batty (1999) to the analysis of the town centre in Wolverhampton:

- (a) the mean distance that can be seen from each point, called 'visibility';
- (b) the furthest distance that can be seen from each point;
- (c) the angle at which this occurs.

It is beyond the scope of this paper to discuss the meaning of these visibility parameters (for a general discussion, see Llobera, 2003). We would simply like to point out that the analysis of DEMs with image-processing techniques is an extremely effective way to compute them.⁽³⁾

6.2 Synthetic cities

"But right there between prediction and game playing, between science and science fiction, lies the realm of sharpened intuition, informed speculation, and educated guess."

Couclelis (1997, page 165)

⁽³⁾When deriving these parameters, we encountered some problems with NIH Image. Its simple image-processing capability has limitations, as the value associated with each pixel is a single byte, and the macro language, a kind of debased Pascal, is inflexible. For more advanced work, therefore, we have used Matlab, which gives us a double-precision variable at each pixel and a complete matrix-oriented programming language. We find this to be much more flexible, though sometimes slower and much more memory intensive, than NIH Image.

Here, Coucleilis is referring to the significance of computer simulation of urban growth using cellular automata (CA). Many other workers have tried to find algorithms for growing city-like structures. Fractal forgery has been used at the very large, almost regional, scale (Batty and Longley, 1994). Multiagent algorithms have been used to investigate the patterns of urban growth and to generate 'fictitious cities' (Cecchini and Viola, 1992).

The interest of these studies does not lie simply in the simulation of city growth (a practical example could be when an empty urban site needs to be filled with a texture that matches its surroundings, as in many historic centres) but also in the help they may provide to characterise urban texture morphologically. In a certain sense, as in the biological sciences, the ability to reproduce fully a complex system clarifies its major determinants. Our own approach, aimed at synthesising textural details down to a scale of about 1 m, corresponding to the resolution of our DEMs, is derived from recent advances in texture replication in the field of computer graphics. This work is at a preliminary stage but is not without promise.

In image synthesis, textures (such as brickwork) are often represented by a sample tile, perhaps a few hundred pixels square. If a large brick wall is to be imaged, then the tile is repeated. Unfortunately, the eye is very receptive to the regularity of the resulting pattern and the area does not look like a brick wall any more. The problem is to synthesise a larger area of brickwork, which 'looks the same' as the original sample but does not involve exact repetitions.

This problem eluded solution until the algorithm of Bonet (1997), which improves on the work of Heeger and Bergen (1995). Both approaches use a Laplacian pyramid to analyse the sample texture at different scales. This is constructed, roughly speaking, by repeatedly blurring the image and halving its resolution. The Laplacian pyramid (figure 9) consists of the differences between successive blurred images and is thus in some sense a kind of spectral analysis of the image, dividing it up into different visual frequency ranges.



Figure 9. Laplacian pyramid of the central London digital elevation model.

Bonet's algorithm uses this analysis to generate a new pyramid, corresponding to the output texture, by a kind of constrained random sampling. The input pyramid provides joint-feature occurrences across multiple scales and is used as a probability density estimator to constrain sampling. So, the output texture contains recognisable fragments of the input texture, placed in similar but not identical contexts.

We implemented a version of this algorithm in Matlab to operate on urban textures represented by DEMs. Given a small piece of city, it synthesises a larger piece. The result, shown in figures 10(a) and 10(b) (see over), is surprisingly good, given

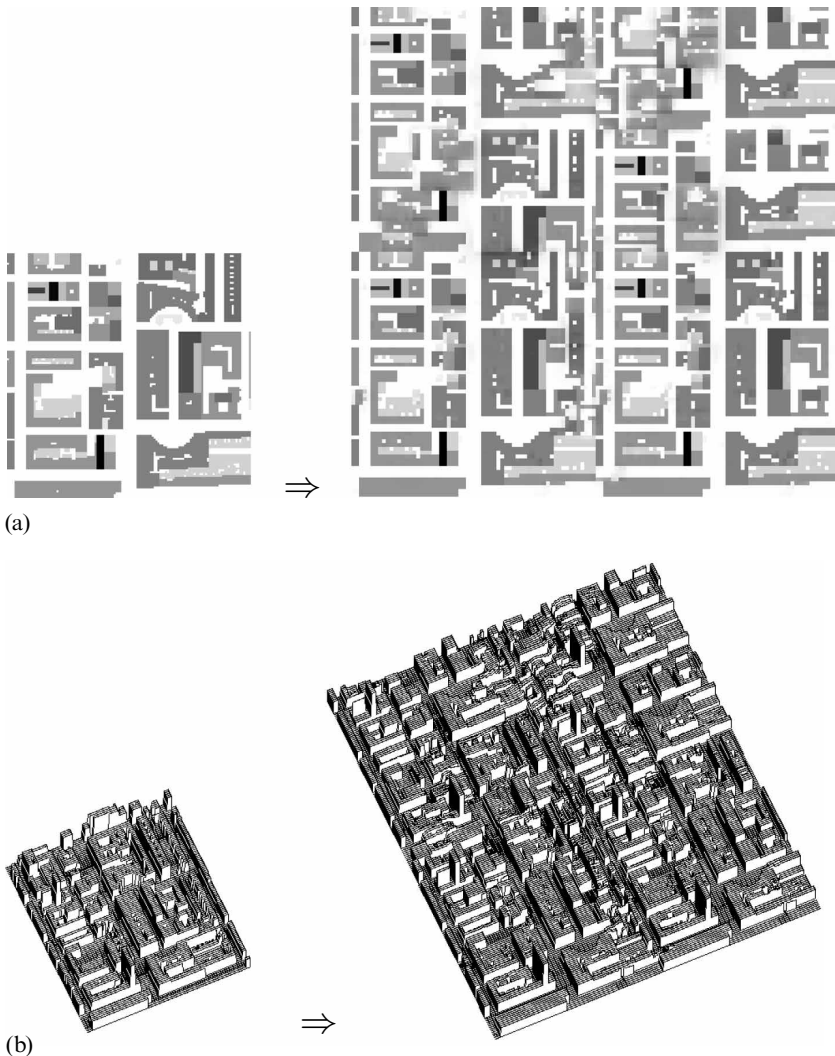


Figure 10. (a) Synthetic texture generation: from the central London digital elevation model (left) a new urban portion with similar textural characteristics is produced (right). (b) Axonometric view of part (a).

that the algorithm has no ‘knowledge’ of urban processes. It preserves the general structure and orientation of the street pattern and the mix of building heights. Features up to the size of a single block can probably be found in the original image, large assemblies probably not. It is remarkably, though not totally, free of unurban conditions such as buildings that block major streets.

7 DEM and cellular automata

“Cellular automata are mathematical idealisations of physical systems in which space and time are discrete, and physical quantities take on a finite set of discrete values. A cellular automaton consists of a regular uniform lattice (or ‘array’), usually infinite in extent, with a discrete variable at each site (‘cell’). The state of a cellular automaton is completely specified by the values of the variables at each site. A cellular automaton evolves in discrete time steps, with the value of the

variable at one site being affected by the values of variables at sites in its ‘neighbourhood’ on the previous time step. The neighbourhood of a site is typically taken to be the site itself and all immediately adjacent sites. The variables at each site are updated simultaneously (‘synchronously’), based on the values of the variables in their neighbourhood at the preceding time step, and according to a definite set of ‘local rules.’”

Wolfram (1983)

Although CA have been used in urban studies mainly to simulate spatial dynamics (Batty et al, 1997), we will consider here their applications in modelling small-scale processes within the urban texture.

According to several authors (for instance, see Chopard and Droz, 1998) the physics of pollution dispersal, wind flow, and noise diffusion can be studied by CA. CA have also been used very successfully in modelling complex situations for which traditional computer techniques are not applicable, such as for flows in porous media, erosion, and transport problems.

In addition to the modelling of physical processes, CA have been used to simulate the local action of individual actors or agents. Batty (1999) applied CA to model how pedestrians move in town centres and buildings. More generally, agent-based models may constitute the base of built-space analysis complementary to space syntax.

CA techniques are highly compatible with DEMs. The DEM, a raster map, is the natural format to input urban geometry into a CA. Operations on cells can simply be viewed as operations on pixels. Image processing and CA can therefore be integrated into a single body of analysis techniques to work with DEMs. As an example, a CA algorithm has been written to simulate the diffusion of pollutants within the urban texture. Results are shown in figure 11: a steady source is positioned at the bottom of the image, and pollutant particles migrate upwards.⁽⁴⁾




Figure 11. Modelling the diffusion of pollutants on a selected portion of the central London case study (left) using cellular automata (right).

8 Conclusions and future work

At present, there are no good analytic techniques to model the built environment at the level above that of the building. The use of DEMs—a format that has a successful track record of application in the geosciences, but that has almost never been applied to urban areas—proved to be very successful for different kinds of analysis. The innovative use of image processing and CA techniques allowed the extraction of a wide range of environmental and morphological parameters, such as shadow casting,

⁽⁴⁾This program was also written in Matlab instead of NIH Image. The Matlab environment is extremely conducive both to standard image processing operations and to CA.

the calculation of sky and ground view factors, the replication of urban texture and the simulation of pollution migration with CA.

This list should not be considered as comprehensive, but simply as a proof of concept of the potential of the DEM for urban analysis. A wider range of analyses could be developed based on the same principles. In fact, at the time of writing, other algorithms are being developed, dealing with aspects such as simulated energy consumption, wind-field characterisation, urban roughness calculation, and pedestrian movement modelling with CA.

Also, a great deal more could be done to extend the modelling capability, using multiple images, rather than simply a DEM—in a way similar to layered systems in the geosciences. If the ground surface is complicated, separate ground and building DEMs can be used. The use of additional channels has been tested in order to classify surface types (vegetation, water, roads, buildings) and to provide more information about buildings, of which the number of storeys would be most valuable, followed by information about glazing, albedo, and materials generally.

In the end, the analysis techniques presented here could become an integrated image-based urban model to inform planning and design—with growing potential in the coming years, when urban DEMs will become increasingly accessible at low cost.

Acknowledgements. The research was partly funded by the European Commission through Project PRECIS (Assessing the Potential for Renewable Energy in Cities). Acknowledgements go to Dr Nick Baker for his suggestions and many informative discussions; to Dr Koen Steemers for supporting the application of image processing techniques to urban design; to the PRECIS team at the Martin Centre, University of Cambridge, for its productive research environment; and, last but not least, to Janet Owers for her invaluable editorial guidance.

References

- Batty M, 1999, "Visual fields" <http://www.casa.ucl.ac.uk/agent-homepage/visual-fields.htm>
- Batty M, Longley P, 1994 *Fractal Cities: A Geometry of Form and Function* (Academic Press, London)
- Batty M, Couclelis H, Eichen M, 1997, "Urban systems as cellular automata" *Environment and Planning B: Planning and Design* **24** 159–164
- Bonet J S, 1997, "Multiresolution sampling procedure for analysis and synthesis of texture images" *Proceedings Computer Graphics ACM SIGGRAPH* 361–368
- Burrough P A, McDonnell R A, 1998 *Principles of Geographical Information Systems* (Oxford University Press, Oxford)
- Cecchini A, Viola F, 1992, "Ficties (fictitious cities): a simulation for the creation of cities", paper presented at the International Seminar on Cellular Automata for Regional Analysis, Venice; copy available from A Cecchini, Laboratorio sulla Simulazione Stratema, Istituto Universitario di Architettura di Venezia, Venice
- Chopard B, Droz M, 1998 *Cellular Automata Modelling of Physical Systems* (Cambridge University Press, Cambridge)
- Couclelis H, 1997, "From cellular automata to urban models: new principles for model development and implementation" *Environment and Planning B: Planning and Design* **24** 165–174
- Givoni B, 1989, "Urban design in different climates", WMO-TD 346, World Meteorological Organisation, Geneva
- Heeger D J, Bergen J R, 1995, "Pyramid based texture analysis/synthesis" *Proceedings Computer Graphics ACM SIGGRAPH* 229–238
- Hillier B, 1996 *Space is the Machine* (Cambridge University Press, Cambridge)
- Hillier B, Hanson J, 1984 *The Social Logic of Space* (Cambridge University Press, Cambridge)
- Jensen J, 1986 *Introductory Digital Image Processing: A Remote Sensing Perspective* (Prentice-Hall, Englewood Cliffs, NJ)
- Johnson G T, Watson I D, 1984, "The determination of view-factors in urban canyons" *Journal of Climate and Applied Meteorology* **23** 329–335
- Llobera M, 2003, "Extending GIS based visual analysis: the concept of visualscape" *International Journal of Geographic Information Science* **17** 25–48
- Oke T, 1988, "Street design and urban canopy layer climate" *Energy and Buildings* **11** 103–113
- O'Rourke J, 1987 *Art Gallery Theorems and Algorithms* (Oxford University Press, Oxford)

-
- Page J K (Ed.), 1986 *Prediction of Solar Radiation on Inclined Surfaces* (D Reidel, Dordrecht)
- Peponis J, Wineman J, Rashid M, Bafna S, Kim S H, 1998, "Describing plan configuration according to the covisibility of surfaces" *Environment and Planning B: Planning and Design* **25** 693 – 708
- Ratti C, 2002 *Urban Analysis for Environmental Prediction* PhD thesis, Department of Architecture, University of Cambridge, Cambridge
- Richens P, 1997, "Image processing for urban scale environmental modelling", in *Proceedings of the International Conference Building Simulation '97, 5th International IBPSA Conference, Prague* Eds J D Spitler, J L M Hensen, <http://www.IBPSA.ORG>
- Stemers K, Ratti C, 1998, "Nuovi strumenti per la progettazione bioclimatica urbana" *Housing* **9** 98 – 100
- Stemers K, Ratti C, 1999, "Informing urban bioclimatic design", paper presented at the International Conference Architecture and Engineering—Plymouth; copy available from the author
- Steyn D G, 1980, "The calculation of view factors from fish-eye lens photographs" *Atmosphere – Ocean* **18** 254 – 258
- Wolfram S, 1983, "The statistical mechanics of cellular automata" *Review of Modern Physics* **55** 601 – 643

