4
Laboratories for Visualizing Urban Form

When I wrote the program, I never thought that it would evolve anything more than a variety of tree-like shapes . . . nothing in my 20 years' experience of programming computers, and nothing in my wildest dreams, prepared me for what actually emerged on the screen. I can't remember exactly when in the sequence it first began to dawn on me that an evolved resemblance to something like an insect was possible. (Dawkins, 1986, p. 59.)

4.1 Experimentation as Visualization

The understanding we have already gained about the systematic irregularity of fractal shapes creates a very strong case for judging the success of models by their visual appearance. For example, it is easy to conjecture that the physical properties of land use in terms of plot size, shape and density display an irregularity which is considered to be fractal. From earlier chapters, we know that cities are self-similar in a variety of ways, central place theory being the clearest demonstration of this principle (Arlinghaus, 1985). Thus the idea that actual city structures might be fractal is appealing, but of more import is the possibility that fractal geometry may well contain the basis for linking activity models to their physical context. However, before we launch into the use of fractal geometry in rendering traditional computer models of cities more realistically, we need to formally consider how we might develop this understanding further through designing a consistent and structured set of experiments for the hypothetical model we introduced in the last chapter. Our 'London' sequence provides us with such a model with a strictly limited number of parameters whose variation will generate different urban forms. Here we will extend this model and in a laboratory-like setting, we will manipulate the values of its parameters so that we might explore the complete set of forms which can be generated. The parameter space which bounds this set we will treat as a mathematical space populated by different forms which can be derived from one another, and we will call this the 'space of all cities'. In this space, we will experiment with city patterns whose forms we will assess and evaluate visually, thus establishing a process of experimentation through visualization and vice versa.

Our experiments will consist of sampling different urban forms from the

space of all cities, and simply involve a conventional selection based on different combinations of the limited set of parameter values. The hypothetical models based on our 'London' model from Chapter 3 and which we will explore in the next section, will contain a very limited number of parameters. Such models must reflect the principle of parsimony so that there can be a clear assessment of the effects of different parameter values. Furthermore, throughout this book, we will deal with highly simplified models which in no way approach those operational urban models which are used in the real world by planners and engineers involved in forecasting and designing the future city. Moreover, because our models, although highly structured, are random in that land use is allocated through chance events, we are dealing here with urban forms which portray a general typology of cities rather than anything which is more specific. Indeed in the very title of this book *Fractal Cities*, our emphasis is not upon thinking of some cities as being fractal in contrast to others, which are not, but that all cities display structures and patterns which in certain senses might be fractal, and it is our emphasis on the degree to which their form is fractal which can provide important insights into their functioning.

In the study of form through computer experiments, it is the way certain shapes evolve relative to some baseline which is our essential quest. In one sense, this is the principle which has been used for nearly a century in the study of the evolution of biological forms first exploited by D'Arcy Wentworth Thompson (1917, 1961). His view is cogently illustrated when he says: "In a very large part of morphology, our essential task lies in the comparison of related forms rather than in the precise definition of each; and the *deformation* [his italics] of a complicated figure may be a phenomenon easy of comprehension, though the figure itself has to be left unanalysed and undefined". Thompson's point is of general import to our work here. We can begin as we did at the end of the last chapter with a theoretical model of a city based on the concentric rings of land use around the city's center, the land use being defined according to von Thunen's bid-rent principles which are implied in the land use profiles shown in Figure 3.13. By systematically changing a parameter value, the shape of the city can be deformed or reformed to another, and by systematically charting this deformation, we are engaging in the time-honored tradition of experimentation in which different responses in terms of form are being generated by changing one parameter value at a time. The space in which this occurs is what we have termed the 'space of all cities'.

We have continually alluded to such experimentation in previous chapters. For example, if we have a model whose form is defined in two-dimensional space, in the plane, by various transformations of its x-y coordinates, then as these transformations vary in value, so does the shape which is generated. We noted this in Chapters 2 and 3 where we briefly described how a Koch curve could be related to a Julia set in 2-space using Barnsley's (1988a) IFS method. By interpolating between the transformation values for both objects, we indicated that one object could be slowly transformed into the other. The intermediate forms generated represent the visual trace or trajectory of this process. Whether or not this transition is meaningful will of course depend upon our choice of objects or systems. Here we will not map the complete range of possibilities but simply select some forms which

appear to be good bounds to the space in which our family of cities exists. In fact, an extension to our laboratory which we have not yet developed, could be based on a process in which one form actually evolved to another, the decisions concerning the evolution consisting of single small changes to the parameter values being made by the user on the basis of some visual assessment of the appropriateness of the form. The quote from Dawkins (1986) which prefaced this chapter is based on his response to such a process which initially produced tree-like forms, but through judicious selection of changes in parameter values one at a time, ultimately led to insect-like forms. Dawkins' surprise was over the fact that the kinds of parameters which characterized trees, such as branch orders, widths, bifurcation ratios, branch angles and so on, could quickly develop to shapes which were manifestly insect-like. His amazement is no different from that of the transition between a Koch curve and a Julia set as well as the sorts of deformation between biological systems which was first popularized by D'Arcy Thompson (1917, 1961). We will not develop these possibilities further, but there is considerable potential in our field for such evolution through experimentation, and this represents an important area of work for the future.

In Chapter 1, we introduced, albeit informally, many different examples of urban form, far wider than we will explore here. Later in this book, we will eventually develop more fundamental fractal models, and this will give us some scope for generating a massive variety of urban forms, but these we will leave until Chapter 7. However, at this stage, it is worth reiterating the range of forms which are possible with our 'London' model so that the experimental work of the next section can be put in context. There are a series of dichotomies which characterize such forms. First there is the distinction between monocentric and multicentric cities. The monocentric tend to be industrial cities in that their development in terms of commerce and industry has been centered in and around the CBD, and this is contrast to those multicentric cities where there are several dispersed centers which compete with one other. Multicentric cities characterize the presently emerging post-industrial age where the power of the CBD is no greater, if not less, than many peripheral centers. But there are also multicentric cities which have developed as the fusion of several separate industrial cities; these are called 'conurbations' by Geddes (1915, 1949), and 'megalopolis' by Gottmann (1961) and Doxiadis (1968).

Overlying these distinctions is the notion of concentrated to dispersed which is loosely akin with high to low density cities. Centralized and decentralized also follow this classification, and there are more specific terms such as the 'exploded' city and the 'imploded' which represent growing and declining monocentric forms. The classification of form in this way presents an endless array of different characterizations which are all semantically a little different, but in general, cover a range from concentrated to dispersed. Moreover, there are distinctions which overlie these in terms of shape, from linear to concentric, almost mirroring our distinction between one- and two-dimensional forms which we portrayed in Chapter 2, as for example in Figure 2.10. In our experiments which follow, it is essential that we appreciate the bounds which make possible only a restricted set of forms. In essence, the space of all cities which we define only includes monocentric cities which develop around a single pole or CBD. We will see

how various cities which are concentrated or dispersed can emerge through changes in the way the chance elements operate on land use allocation; this may suggest some multicentric form, but this is still within the bounds of the monocentric assumption and is caused by chance locations of centers which are not related to the CBD. The way we construct our cities is on the assumption that all development is arrayed concentrically around the CBD, and this rules out the possibility of linear cities emerging. This will only be possible with our more fundamental models which we develop from Chapter 7 on.

Finally, there is the possibility that cities might be classified according to the form of their transport networks, usually in terms of the distinction between radially symmetric nets and 'Manhattan' grids. We will not pursue these types of form either, largely because we are continually conscious that the number of parameters which we can deal with and from which we can derive meaningful conclusions must be severely limited. This is not a book about theories of cities. As Crick (1990) says: "The job of the theorist ... is to suggest experiments", and the experiments we will choose are those which are naturally suggested by simple theories of the monocentric city which still compose much of urban economics. We will be content to explore the role of chance in land use allocation and in the shape which such locational patterns display, rather than being concerned with elaborating new or existing theory. We have resisted extending our approach to multicentric urban systems and to models based on spatial interaction, largely because we view our task here as simply a beginning. Moreover, we are conscious that most urban economic theory has also been developed using the monocentric assumption and thus there is more than enough research to develop in first linking fractal patterns to these theories. It is always tempting to add more and more constraints to the models to reflect how cities actually work. Here, however, our concern is not with developing completely realistic models of cities, but with demonstrating that an approach through fractal geometry leads to important insights into their form and functioning.

To set the context, we will also note the idea that the model-building process is based on the loose cycle of inductive explanation, and deductive prediction. For example, some models such as those based on discrete choice are strongly inductive in their specification and estimation, while spatial interaction models are usually specified *a priori*, and are hence deductive. Proponents of either style of modeling rarely pursue the inductive–deductive cycle in any complete sense, but the argument here suggests that fractal simulation can provide a framework for such a process. Large-scale simulation itself establishes such a framework, but there are few attempts which model the entire cycle. The work of Chapin and Weiss (1968) is an exception in that they attempted to explain urban growth using a linear statistical model and then reproduced that growth as a large-scale random simulation. The ideas of this and the last chapter are very much in this spirit but in attempting to model the entire scientific cycle, a number of corners will be cut and only picked up as items for further research.

Another issue which first emerges in this chapter relates to the variety of computer systems and software used to develop these fractally-rendered models. Our work has only been made possible through advances in

computer systems and software, and our demonstrations involve a remarkable mixture of computer and modeling systems and styles. The discrete choice models we use, for example, are estimated using a standard logit package mounted on a mainframe computer, with intermediate processing on a minicomputer which acts as the front end to yet another mainframe on which the spatial mapping is conducted, while the simulations are conducted using a graphics-based micro whose memory is mainly given over to the screen display. In fact, these styles are seen quite clearly in the figures in this chapter and in the various color plates in which the spatial predictions produced for the discrete choice models are presented using standard plotter outputs, in contrast to the simulations which are illustrated in photographs of the raster graphics screen reproduced in the various color plates. Indeed throughout this book, the examples presented have been computed on micros such as PCs and high school computers, Vaxes, Sun workstations, IBM mainframes and so on. Clearly the availability of various machines has influenced what we have used, but it is important not to lose sight of the fact that fractals can be computed using very simple computer programs as Saupe (1991) has noted. Thus the ideas portrayed here should be accessible to a wide variety of readers with different programming skills and with access to very different types of computer.

After extending our hypothetical simulation model, we will take one step back and briefly introduce the model-building process in terms of explanation and simulation, induction and deduction, emphasizing the need to contain both within any complete cycle. We will show how the concept of systems hierarchy, which is so central to fractals, might be exploited through the modeling process, and we will then illustrate how this process can be completed for the fractal simulation of urban structure using the example of residential-housing location in London. The inductive approach we will adopt is based on discrete choice theory, the estimation of a standard multinomial logit model (Hensher and Johnson, 1981) to housing choice, and measurement of its performance using McFadden's (1979) predicted success statistics. We will show how the model is fitted to data which relate choice of house type and location in London to key variables of urban structure based on age and distance. Several models are fitted, some are reestimated and computer maps are used to aid the interpretation process. We are then in a position to begin the fractal simulation of urban structure on the basis of the fitted discrete choice models. The simulations are essentially visual, the data themselves being displayed on the screen and being replaced by predictions of housing types as the simulation proceeds. Two types of simulation are attempted – random and deterministic, and it is shown how it is necessary to develop deterministic procedures to enable the discrete choice models to generate realistic patterns. There are many conclusions to this chapter, and work on both fractal images (Pentland, 1984) and spatial discrete choice models (Lerman, 1985) represent major directions for further research.

4.2 Exploring Urban Forms in the Space of All Cities

The advantage of any computer model of a city is that its parameters can be varied with a view to exploring the effect of such changes on the resulting spatial form. Through visualizing the form using computer graphics, the strength of various relationships which compose urban structure can be assessed, both to improve such models in the laboratory as well as to refine relationships in terms of generating the best fit with reality. In our 'London' model of the last chapter, we could control the distribution and amount of the three land uses generated, by varying their parameters a^u, b^u and R^u. We can also vary the degree of midpoint displacement which is related to the fractal dimension. In fact in this example, if we were to collapse the random allocation of land use to a deterministic one, and use non-random midpoint displacement, we would generate a city in which commercial–industrial land use would entirely dominate the inner suburbs and CBD, and residential the outer suburbs and the periphery, providing a clear example of the von Thunen rings. This is implied in the land use profiles shown in Figure 3.13.

What we will do first is extend our model to encompass other effects. We will increase our land uses by disaggregating the residential sector into three types – high, medium and low density housing, and split the commercial–industrial use into their two separate components. We argued earlier that at different hierarchical levels, various activities might dominate. For example, if we were simulating the growth of an industrial city, then decisions about location depend on the characteristics of a neighborhood. Industrial neighborhoods tend to attract like industries, while residential activity tends to avoid location there, and so on. We will extend our model by allocating land use at each hierarchical level, and using the land use allocation at that level to determine land use allocation at the next. The way this mechanism might work is for a neighborhood to be classed in terms of its dominant land use at the first spatial level of the simulation and then for this dominance to be reflected at the next level of detail down. This involves modification of the probabilities of allocation at that level, these being conditional on the probabilities of land use at the next.

The way we simulate this effect is by modifying equation (3.29) in the following manner:

$$p^u(r_s) = a^u + b^u\,(r_s - R^u), \quad u = 1, 2, \ldots, 6, \tag{4.1}$$

where a^u, b^u and R^u are parameters as before but the distance r_s now reflects the fact that equation (4.1) is applied at each level s of hierarchical simulation. Land use is not actually allocated at level s by the model but at the next level down, $p^u(r_s)$ is used to condition a new probability $P^u(r_{s+1})$ using the following equation:

$$P^u(r_{s+1}) \propto [1 + P^u(r_s)]^{\gamma(u)}\, p^u(r_{s+1}). \tag{4.2}$$

$P^u(r_{s+1})$ is the probability used to effect the allocation. In fact the allocation only takes place at the most detailed level of resolution and this probability

is in fact that passed down from the previous level which acts on the basic probability given by equation (4.1). $\gamma(u)$ is a parameter greater than 0 which measures the importance of the effects at previous levels. If $\gamma(u)$ is equal to 0, then there is no effect from previous levels and $P^u(r_{s+1}) = p^u(r_{s+1})$. The strength of the influence increases as $\gamma(u)$ increases in value and in the simulations to be shown next we have set $\gamma(u)$ at five values, namely $\gamma(u)$ = 0, 1, 5, 25 and 50. Of course equation (4.2) has to be appropriately normalized. We have not made $\gamma(u)$, a^u, b^u or R^u specific to each level because the number of possible combinations of values becomes too large to handle, and in the interest of developing parsimonious but general models, we have begun our experiments with as small a number of parameters as possible.

We will now explore some 20 possible urban forms and the tree of possibilities defining the 'space of all cities' within which different forms exist is shown in Figure 4.1. First we can allocate our land use randomly as in the last section or in deterministic fashion based on the dominant land use predictor at each level and in terms of each basic location. This gives us two choices. As our fractal simulations are based on midpoint displacement of the pure Brownian type with a Hurst exponent $H = 1/2$, then we can either use that level of random displacement of the midpoint or no randomness whatsoever; this also gives us two options. Finally we can use the five values of $\gamma(u)$ and this gives us five possibilities when we keep each value of $\gamma(u)$ the same for each land use u across all hierarchical levels. In total therefore, we have 20 options to simulate. This modified fractal model was simulated on a Sun workstation, and some of its outputs are illustrated in Plate 4.1 (see color section) where the simulations are shown at levels

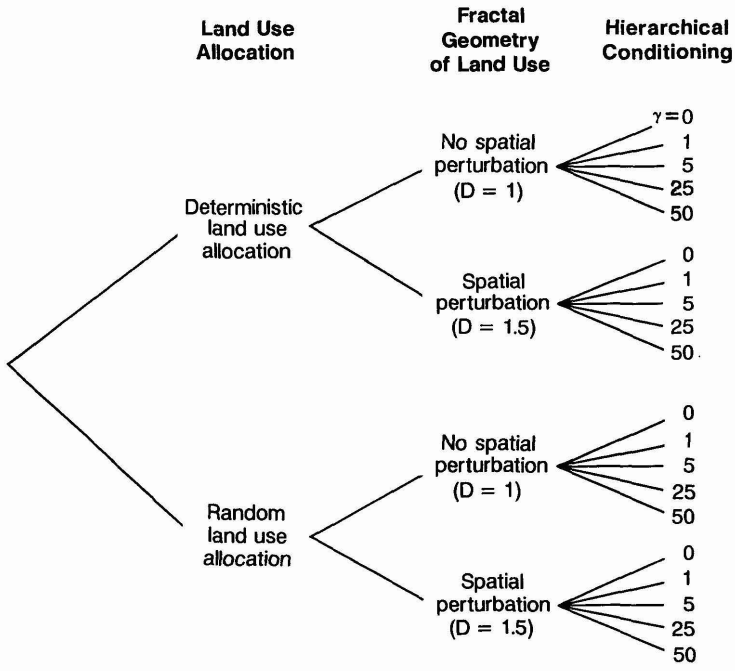

Figure 4.1. The combinatorial map of the space of all cities.

$s = 2$, 4 and 6, thus illustrating once again the extent to which the patterns become more realistic as the scale gets more detailed.

In Plate 4.1, we present six from the 20 simulations to give some flavor of the differences. The deterministic model based on no fractal perturbation, deterministic allocation of land use, and no hierarchical conditioning where $\gamma(u) = 0$, $\forall u$ produces classic von Thunen rings, and as such, represents the theoretical baseline for all our simulations. The CBD and inner areas are dominated by commercial use, while there are three other rings of high, medium and low density housing. Open space and industrial uses do not dominate anywhere, and because any area of the city contains its most dominant land use, these two uses never have a chance of being located. The other feature worth noting is that because there is no fractal perturbation, then the final units for location are identical and perfectly formed triangles. This simulation is illustrated in Plate 4.1(a), and it is just possible to make out the perfect symmetry of the triangles formed by midpoint displacement with no randomness. If we introduce fractal rendering for the same set of parameters, then the von Thunen rings simply appear somewhat cracked due to the fact that the sites are no longer identical in terms of shape and location, although the outputs are still highly reminiscent of von Thunen's theoretical model; this is shown in Plate 4.1(b).

In Plate 4.1(c), we show what happens when heavy hierarchical conditioning is introduced to the model which has fractal rendering but deterministic land use allocation. The heavy conditioning is enabled with $\gamma(u) = 25$ for all six land uses, and the effect is to produce a very strange pattern in which residential land uses dominate everywhere. This is clearly unusual, quite extreme and unlikely to be observed anywhere. Next we show in Plate 4.1(d), the simulation based on no fractal perturbation, hence perfect triangles as sites, randomness in land use allocation, and extreme hierarchical conditioning. Again this produces a slightly more realistic pattern, but one which is sufficiently different from reality to be somewhat unlikely. Note the way the perfect triangles appear at any level here due to the fact that $\gamma(u) = 50$ for all land uses. Reducing the hierarchical factor $\gamma(u)$ to 5 produces more realism as in Plate 4.1(e), while finally in Plate 4.1(f), we show the most realistic simulation we have achieved, based on fractal perturbation, randomness in land use allocation and very slight hierarchical conditioning [$\gamma(u) = 1$, $\forall u$]. This case is interesting in that the simulation at level $s = 4$ is more realistic than $s = 6$ which is, once again, reminiscent of pointillist painting.

The examples we have shown provide a good cross-section of the possible patterns which compose the experiments in our laboratory: we have four examples of hierarchical conditioning and the four possibilities in terms of land use allocation and spatial perturbation. What is quite clear here is that the hierarchical conditioning is far too extreme once it rises much above $\gamma(u) = 1$. Moreover, it is the randomness in land use allocation which seems to provide more of the realism in contrast to fractal perturbation although the fractal perturbation does indicate that the problem of creasing becomes considerably more apparent as hierarchical conditioning is increased. In fact, it is clear that although our model is more realistic than the one we developed in the last section for 'London', it is still fairly

unrealistic in its incorporation of hierarchy and of fractal perturbation where the choice is either $H = 1/2$ or non-random midpoint displacement.

Our last foray into this type of fractal modeling partly resolves the problems of the models just outlined in two important ways. We will now introduce control over the fractal dimension in quite explicit terms and we will reduce the range of hierarchical conditioning to screen out the extreme effects such as those shown in Plate 4.1(c) and (d). Thus, the most extreme conditioning in the new model is where $\gamma(u)$ is equivalent to 5 (but on a new scale where this value is given as 25). Another innovation which we find useful is based on replacing the triangular lattice with a square grid. In Figure 4.2, we show this grid and how successive random midpoint

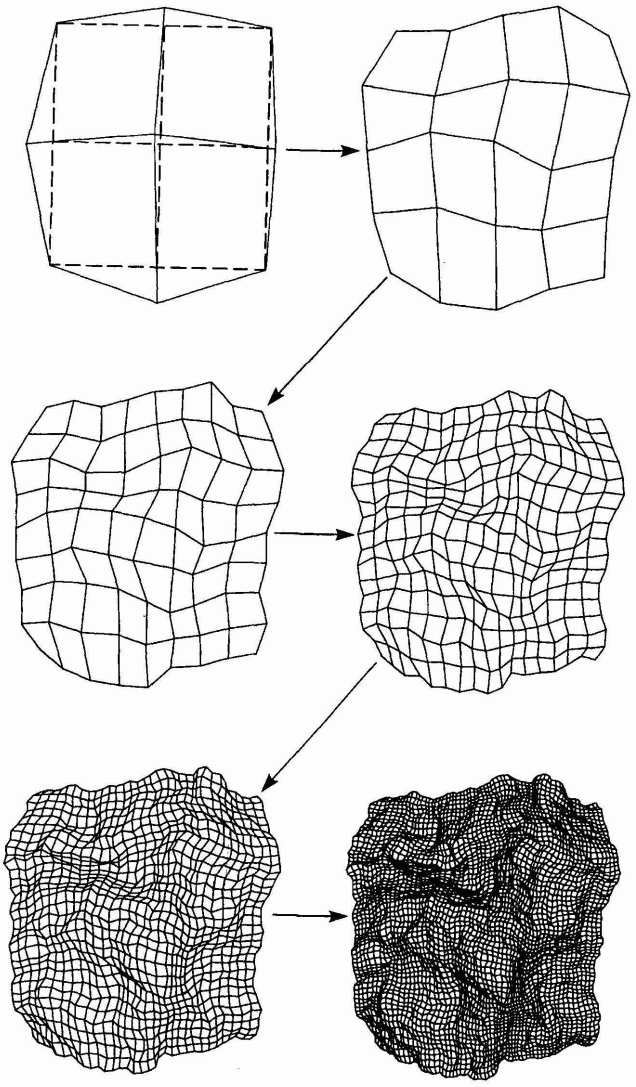

Figure 4.2. Simulating patterns using midpoint displacement of a square grid.

displacement enables a surface to be produced which appears to have less creasing than the triangular. Moreover, as the grid is based on the unit square, there is no measure of distortion present in the first place and thus this would appear a less arbitrary and less biased form of initiator. We have also increased the land uses in the residential sector to be defined over a continuum of population densities, and this makes the simulations more realistic in that densities are determined as a combination of random allocation and distance from the CBD. In effect, this does not add to the number of land uses *per se*, as the simulation is structured to determine the density of the residential use at the stage when this use is allocated.

We show typical simulations from this model in Plate 4.2(a) and (b). In (a), a fractal dimension of 1.26 has been used, hierarchical conditioning is heavy with a value of 5, and randomness in land use allocation has been used. In fact, this simulation is more realistic than any of those shown in Plate 4.1 and it would appear that our new model provides an ideal basis for computer experimentation. We cannot show all the possibilities here, but in Plate 4.2(b), we show the von Thunen rings case where the fractal dimension is unity (no spatial perturbation), there is no hierarchical conditioning and land use allocation is deterministic. In these examples we have also separated out the residential from other land uses, thus showing how the different patterns stand by themselves. In this way, it would appear that both the residential and commercial land uses, which make up most of the city, have distinctive location patterns which are close to those we might observe in existing cities.

There are many more variants we can generate using this model. Clearly we can let the fractal dimension range from 1 to 2, we can explore a range of hierarchical conditioning from $\gamma = 0$ to $\gamma = 5$, and we can make this parameter specific to each level s and/or each land use u. We can even let these parameters take on values outside the range of 0 to 5. But the models developed here simulate only one type of city, the monocentric, and thus it is important to simulate forms other than those which are unipolar and concentric. This would involve introducing mechanisms which measure the accessibility of any point in the city to any other and it would take us towards the mainstream of urban modeling which is based on spatial interaction (Batty, 1976). In fact throughout this book, we will steer well clear of these types of models because these are for a very different purpose. Although we continually allude to fitting fractal models to real situations, in the last analysis, our exposition of fractal cities is motivated by our search for the applicability of these ideas, and the insights they might give to the broad domain of urban studies. Indeed from Chapter 7 onwards, we will begin to demonstrate how fractal models give us a very different perspective on studies of urban density, which suggest that much previous research should be reworked. The value, then, of these forays into hypothetical urban form is in the applicability of the fractal idea and the somewhat more superficial idea of visualizing computer models using fractal rendering. We will now turn to more realistic examples, developing a conventional urban model of housing in London and then showing how its predictions can be best evaluated in terms of the spatial forms that they imply.

4.3 Hierarchical Urban Structure

So far the hierarchical structures we have introduced do not relate to any observable characteristics of city systems except in the most superficial way. Clearly for levels $s < 4$ in Plate 4.1, the images generated show the strong influence of the triangular patches making up the hierarchy and are thus not realistic. When levels with $s \geq 4$ are reached, the images no longer display the influence of the triangular method in that the concatenation of triangles at these levels produces the sorts of irregularity characteristic of land use patterns. Thus in one sense, the triangular subdivision process is scale-dependent. However, in fractal simulation, there is still the need to relate the method of construction to substantive characteristics of the system as in other forms of modeling. Indeed, many examples of fractals can only be modelled coherently by defining their intrinsic properties of self-similarity: trees, for example, are self-similar through their mode of reproduction and growth. In geomorphology, the process of weathering and erosion acts in a self-similar fashion. This is clearly true for cities as well and thus hierarchical structure must reflect this.

We can sketch an idealized process of fractal simulation to which we might aspire. We begin by identifying hierarchy in the system of interest based on our perception of self-similarity in description, and we are then able to measure whether or not the phenomenon is fractal and whether or not the fractal dimension is invariant to changes in scale. Each stage of measurement and description leads to further development of the underlying process through which the structure can be generated, and this in turn leads to models which are consistent with fractal structure. Once appropriate models, applicable to different levels of the spatial hierarchy, have been developed, other fractal structures utilizing such hierarchy and incorporating the application of the underlying process through recursion, can be simulated.

This approach is in fact the classic process of observing a phenomenon, deciding whether or not it meets any theoretical preconceptions we have, developing a 'best' model structure, and then using this to enable new and different predictions to be made. Essentially this is the process of induction followed by deduction, or in a different sense, analysis followed by synthesis. We can think of induction as a process of building theory from the bottom up, from specifics to universals, while deduction is a top-down process in which universals are used to predict specifics. The best expression of this complete process is in the fields of design and problem-solving where problems must be understood (through induction and analysis) prior to their solution (through deduction and synthesis). In fact in design, methods for analysis and synthesis exist which are based on searching for hierarchical structure: problems are decomposed in the quest to induce their structure and thence composed in the quest to synthesize a solution from the elements (Alexander, 1964; Johnson, 1984). There are parallels with the process used here to enable appropriate description and explanation prior to fractal simulation.

A simple example which relates to spatial theory is the rank–size distribution of cities. City size distributions display regular properties which are

consistent with subdivision of a national or regional space into market areas whose decreasing size reflects the frequency of spatial dependence and the rarity value of spatial goods. Idealized size distributions can be developed by taking a primate city and its national market area, generating two next order cities, then four, and eight, and so on, in the manner we illustrated in Chapter 1 in equations (1.1) to (1.4). This is the type of method used in Central Place Theory, and another interpretation using fractal geometry has been developed by Wong and Fotheringham (1990). In terms of our complete cycle of model-building, we first need to identify the hierarchy of market area, transport routes, population centers etc., thus explaining spatial structure at different levels. This is accomplished inductively in bottom-up fashion, possibly using clustering type methods. The simulation then begins from the top-most level in the hierarchy by subdivision and fractal rendering, generating centers and activities at different scales in such a way that lower levels depend on upper. Although there is a sense in which the simultaneity of dependence is treated by correct bottom-up followed by top-down analysis, in terms of fractal simulation which is arbitrarily structured in hierarchial terms, the dependence is only one way. In fact, this is a problem with many hierarchical descriptions for it is clear that any activity at any position in the hierarchy owes its stability to those activities both above and below. This is the concept of 'niche' and it is something which must be explored in considerable depth in further research on fractal simulation.

In spatial modeling, there are some very well-developed techniques to effect this process of hierarchical explanation and simulation. The logical output of a process of continual subdivision is the elemental space which contains the individual, and thus individual behavior lies at the base of the spatial hierarchy. Such models have been widely developed during the last decade to address problems of discrete choice in the economic domain using standard methods of econometric estimation (Lerman, 1985). These are the models which will be used here, and a particularly attractive feature of them is the fact that they can easily and logically incorporate hierarchical structure: these are the so-called sequential or nested logit models (Hensher and Johnson, 1981).

As yet, very few applications exist of truly spatial discrete choice models, and even fewer have been developed in a spatially-nested form. Nevertheless, these models appear promising as the basis of the recursive generation of activity through the spatial hierarchy. The other class of models which will be considered at a later stage of this research, and which are related to discrete choice models, are spatial interaction-entropy models. It is well-known that such models have highly articulate properties of spatial decomposition (Roy, 1983) and this also makes them attractive to hierarchical simulation. There are a variety of methods for enabling hierarchy to be defined and built into spatial models, such as the standard multivariate cluster-type techniques as well as methods based on more subjective comparisons such as Saaty's (1980) analytic hierarchy process; these could also prove useful to further research.

In the sequel, we will not attempt to address the full process of hierarchical description through the identification and use of hierarchical models but we will follow the broad sequence of inductive, then deductive

stages in the modeling process. We will begin by selecting models for individual choice of housing type and location in Greater London which is the urban region we intend to simulate. This first involves a traditional process of formulating, estimating and selecting appropriate discrete choice models. Having accomplished this, we will move onto the simulation in which these discrete choice models are used to predict housing choice at the lowest level of fractal detail generated. In this way, an image of the residential urban structure of Greater London is built up. Hierarchy is still a largely arbitrary affair in these applications, although we will address it in future research. But there are other problems relating to modeling and simulation which emerge and must be dealt with, specifically related to spatial variation. In any case, the logical next step in this work is to develop a 'realistic' version of our hypothetical simulation presented earlier. To this end, we will now sketch the inductive side of this effort, beginning with the theory of discrete choice and its application to residential housing location in Greater London.

4.4 Discrete Choice Models of Urban Structure

To set the context, we must review some fairly standard results but in doing so, we will adapt discrete choice models to our application and thus only select those aspects which are of relevance here. We will first state the multinomial logit model (MNL) in which we can identify the choice by individual i, $i = 1, 2, \ldots, N$, of alternative k, from the set of alternatives $k = 1, 2, \ldots, K$, where there are clearly N individuals in the system making choices from K alternatives. This set of K is referred to as the choice set and in our applications involves types of housing. The MNL model predicts a probability P_{ik} which is the probability of individual i choosing house type k where there are $K = 4$ house types to choose from, and where i implicitly represents the location of the individual in the city. Thus the model is designed to explain choice in terms of location.

First we must associate a utility of choosing alternative k with the individual i. This utility U_{ik} is usually specified as a linear sum of exogenous (input) variables which may be specific to the choice in question or non-specific (generic). In our context, the parameters of these variables are made specific, being referred to as alternative specific constants, but the variables apply to each house type. Then

$$U_{ik} = \sum_m \beta_{km} x_{im} + \epsilon_{im}, \; m = 1, 2, \ldots, M,$$

where the first term on the right-hand side of the equation contains strict utility components made up of parameters β_{km} and independent variables x_{im}, and the error term ϵ_{im} reflects differences in tastes, unobservable influences and such like. The MNL model is derived by assuming that the error components $\{\epsilon_{im}\}$ are identically and independently distributed, and by maximizing utility using the traditional economic logic (Hensher and Johnson, 1981; Ben Akiva and Lerman, 1985). This random utility derivation of the MNL model is subject to the normalization

$$\sum_k P_{ik} = 1,$$

and the model is derived as

$$P_k = \frac{\exp\{U_{ik}\}}{\sum_u \exp\{U_{iu}\}} = \frac{\exp\left\{\sum_m \beta_{km} x_{im}\right\}}{\sum_u \exp\left\{\sum_m \beta_{um} x_{im}\right\}}. \tag{4.3}$$

These sorts of models have been widely applied in transport research, but have also been adapted to a variety of spatial contexts (see Wrigley, 1985). We will not dwell on this, but suffice it to say that equation (4.3) is a particularly flexible and adaptable model structure.

For purposes of estimation and prediction we need to express equation (4.3) somewhat differently. First we must choose one alternative, say k, as the base or numeraire, and express equation (4.3) as

$$P_{ik} = \frac{1}{1 + \sum_{u \neq k} \exp\left\{\sum_m (\beta_{um} - \beta_{km}) x_{im}\right\}}. \tag{4.4}$$

We form the ratio of any two probabilities for different choice alternatives u and k using equation (4.3), and this gives

$$\frac{P_{iu}}{P_{ik}} = \frac{\exp\left\{\sum_m \beta_{um} x_{im}\right\}}{\exp\left\{\sum_m \beta_{km} x_{im}\right\}} = \exp\left\{\sum_m (\beta_{um} - \beta_{km}) x_{im}\right\}. \tag{4.5}$$

We can now express P_{iu} in terms of the numeraire P_{ik} using equations (4.4) and (4.5) which simplify to

$$\begin{aligned} P_{iu} &= P_{ik} \exp\left\{\sum_m (\beta_{um} - \beta_{km}) x_{im}\right\} \\ &= \frac{\exp\left\{\sum_m (\beta_{um} - \beta_{km}) x_{im}\right\}}{1 + \sum_{u \neq k} \exp\left\{\sum_m (\beta_{um} - \beta_{km}) x_{im}\right\}}. \end{aligned} \tag{4.6}$$

When $k = u$, equation (4.6) collapses to equation (4.4). In the sequel, equation (4.5) is used in estimation while model predictions are made using equation (4.6).

4.5 Estimation Methods for the Multinomial Logit Model

The logarithm of equation (4.5) is referred to as the log-odds of alternative u versus k, and this is the actual equation which is used in estimation. Then

$$\log \frac{P_{iu}}{P_{ik}} = \sum_m (\beta_{um} - \beta_{km})\, x_{im} = \sum_m \lambda_{um}\, x_{im}. \tag{4.7}$$

There is a clear interpretation of the parameters in equation (4.7). If λ_{um} is positive, this implies that the choice of alternative u is more important with respect to the variable x_{im} in question than the choice of alternative k. The reverse is true if λ_{um} is negative, while there is no difference in importance between choices if $\lambda_{um} = 0$.

The model parameters in equation (4.7) are usually estimated using weighted least squares or maximum-likelihood, and here we prefer to use the latter because of the availability of Hensher's BLOGIT computer package (Hensher and Johnson, 1981). To assess goodness of fit we also require the data set of actual choices made which is given as F_{ik} where $F_{ik} = 1$ if i actually chose k, while $F_{ik} = 0$ if this choice was not made. We calibrate the model by maximizing the log-likelihood which is given as

$$\Lambda(\beta) = \sum_i \sum_k F_{ik} \log P_{ik}, \tag{4.8}$$

and we can also assess the fit as a variation of this likelihood function. A null hypothesis can be set up in which $\beta_{um} = 0$, $\forall\, u,m$ implying no variation across individuals, that is, $P_{ik} = P_k$, $\forall i$. This can be used to compute the null-likelihood from equation (4.8) which is given as

$$\Lambda(0) = \sum_i \sum_k F_{ik} \log P_k = \sum_k N_k \log P_k, \tag{4.9}$$

where N_k is the actual number of choices of k made by all individuals i. A measure of fit, in some ways similar to the correlation coefficient, is defined as ξ^2. This statistic is defined as

$$\xi^2 = 1 - \frac{\Lambda(\beta)}{\Lambda(0)}, \tag{4.10}$$

which varies between 0 and 1. The statistic can also be modified to reflect degrees of freedom, while typically good value of ξ^2 range between 0.2 and 0.4. In fact Hensher and Johnson (1981) argue that any model with $\xi^2 > 0.2$ is likely to be acceptable. Other measures of fit and diagnostics for log–linear model equations are discussed in Wrigley and Longley (1984), Wrigley (1985), and Ben Akiva and Lerman (1985).

There is a major difficulty in generating less global goodness of fit measures for discrete choice models. Because the observed data represents discrete choices $\{F_{ik}, F_{ik} = 0 \text{ or } 1\}$ while the predictions are given as probabilities $\{P_{ik}, 0 \leq P_{ik} \leq 1\}$, comparisons at the individual level are meaningless. Thus some aggregation is always necessary. One scheme suggested by McFadden (1979) involves computing expected choices, that is, the numbers of individuals who originally chose alternative k and are expected to choose alternative u. In fact, in later simulations we will examine individual predictions, but for the applications to London which follow, comparisons between observations and predictions will be confined to success statistics based on expected choices.

To introduce these statistics, first note the structure of the observed choice set $\{F_{ik}\}$. Then by definition,

$$\sum_k F_{ik} = 1, \sum_i F_{ik} = N_k,$$

$$\sum_i \sum_k F_{ik} = \sum_k N_k = N. \tag{4.11}$$

The first equation in (4.11) implies that any individual can only make one choice, the second is the constraint on the number of choices made for each alternative, while the third simply says that all the number of choices made is the same as the number of individuals N. The analogous structure for the probability set $\{P_{ik}\}$ is

$$\sum_k P_{ik} = 1, \sum_i P_{ik} = \check{N}_k,$$

$$\sum_i \sum_k P_{ik} = \sum_k \check{N}_k = N. \tag{4.12}$$

Similar interpretations for equations (4.12) exist as for those in (4.11), but note that summation of $\{P_{ik}\}$ with respect to individuals yields predicted numbers of choices $\check{N}_k$ in contrast to actual numbers N_k.

For each individual choice F_{ik} (where $F_{ik} = 1$) there is a probability that the same individual will make a different choice P_{iu}. The number of such choices across all individuals is the number of individuals who originally chose k and are expected to choose u, and this is defined as

$$N_{ku} = \sum_i F_{ik} P_{iu}. \tag{4.13}$$

$\{N_{ku}\}$ is the so-called predicted success matrix. Using equations (4.11) to (4.13), the matrix has the following properties:

$$\sum_u N_{ku} = \sum_i F_{ik} \sum_u P_{iu} = \check{N}_k \tag{4.14}$$

and

$$\sum_k N_{ku} = \sum_i \left(\sum_k F_{ik} \right) P_{iu} = \check{N}_u. \tag{4.15}$$

From these definitions it is clear that

$$\sum_k \sum_u N_{ku} = \sum_k N_k = \sum_u \check{N}_u = N.$$

We can devise a variety of statistics relating to proportions and differences between observed and predicted successes using these aggregations. First we can compute the proportion of correct predictions, noting that N_{kk} gives the number of such correct predictions. Then

$$\eta_k = N_{kk}/N_k, \tag{4.16}$$

which varies between 0 and 1. Total predictive success occurs when $N_{kk} = N_k$, $\forall k$ and $N_{ku} = 0$, $k \neq u$. For the entire system the equivalent statistic to equation (4.16) is defined as

$$\eta = \sum_k N_{kk}/N. \tag{4.17}$$

The second index relates to differences between predicted and observed numbers of choices, expressed as proportions or shares. An absolute measure of this index is given by $N_k - \check{N}_k$ while its relative form is defined as

$$\phi_k = (N_k - \check{N}_k)/N, \tag{4.18}$$

which can be positive or negative.

The final index we have computed is called by McFadden (1979) the prediction-success index φ_k. One problem is that if the predicted choices for u were much larger than k, that is if $\check{N}_k \gg N_k$, then the value N_{ku} would be affected accordingly. To account for this, φ_k is defined as

$$\varphi_k = \frac{N_{kk}}{\check{N}_k} - \frac{\check{N}_k}{N},$$

and an overall index φ, appropriately weighted, is defined as

$$\varphi = \sum_k \frac{\check{N}_k}{N} \varphi_k = \sum_k \left\{ \frac{N_{kk}}{N} - \left(\frac{\check{N}_k}{N} \right)^2 \right\}. \tag{4.19}$$

The maximum value of φ occurs when $\Sigma_k N_{kk} = N$, and then

$$\varphi_{\max} = 1 - \sum_k \left(\frac{\check{N}_k}{N} \right)^2. \tag{4.20}$$

A normalized measure is given by $\varphi/\varphi_{\max}$; other applications are given in Wrigley (1985). These indices based on equations (4.13) to (4.20) will be further adapted in the empirical work which follows to aggregations of subsets of individuals located in specific zones; these will be presented below.

4.6 Determinants of Spatial Structure: the Data Base

Conventional descriptions of urban structure tend to be based on disaggregations of urban activities into land use by type and location. One realization of conventional structure was used in the hypothetical 'London' demonstration model presented in the last chapter and its extensions presented earlier in this. In those models, commercial–industrial (work), residential (living) and open space (leisure) activities were treated in a locational framework which emphasized in diverse ways the radial and concentric nature of the contemporary city. It is not possible to take this model further to the applications stage here, largely because we do not have easy access to a comprehensive land use–activity data base. Moreover, we are interested in developing more formally structured discrete choice models which can be embedded within the fractal simulation, thus enabling us to assess the impact of individual spatial choice behavior in the large.

Another consideration which has guided us is not just the absence but the availability of data. We have access to a large-scale housing survey –

the English House Condition Survey (EHCS: DoE, 1978, 1979) – which was conducted in 1976. This was based on a fairly low sample of households in England, something in the order of 1 in 3000, but this represents an easily available, highly disaggregate data source and thus we have chosen to make use of it. Logit models of housing tenure choice have previously been calibrated using this data set (Longley, 1984).

We have chosen housing type as the key variable defining urban structure which is a major category in the EHCS data. Houses are classified into five types: purpose-built flats (apartments), converted flats, terraced (row) houses, detached/semi-detached (single-family) houses, and a miscellaneous group. House type is a particularly clear way of representing urban structure for different areas of the city are often perceived generally in terms of house type: historically, cities have grown reflecting different house types, and house type seems to relate to how far people wish to live from the CBD. Cities are often articulated as spatial patterns with flats near the center, terraced houses occupying the inner suburbs, detached/semi-detached the outer suburbs, each ring reflecting a stage in city growth. Thus density and distance variables are indirectly reflected in house type, and in the case of London, this is particularly relevant in that the city is strongly monocentric, has a well-developed flats market and has been economically buoyant for several centuries. In our applications, we have in fact excluded the miscellaneous category because of the fact that it acts as a residual category and contains less than 2% of the observations available in the data base.

Choice of house type lies at the base of several contemporary theories of urban structure which integrate two important constructs. First, bid-rent theory postulates an implicit trade-off in housing decisions between housing space and type versus proximity to, or distance from, central urban functions; and urban growth and dynamics (as manifest by filtering, suburbanization, urban renewal etc., and as expressed in the age of the stock) exhibit an identifiable correspondence with distinctive dwelling types such as subdivided central city houses, suburban semi-detached homes, purpose-built flats in revitalized inner city neighborhoods and so on. The implication is that dwelling and neighborhood type are clearly related to distance from the CBD and the date at which the land parcel was integrated (or reintegrated) into the contemporary urban development process.

Thus age and distance represent key determinants of urban structure. In designing the models, it was thought important to keep these variables as simple as possible and at the same time, easily measurable. We also considered neighborhood quality at an early stage, but eventually dropped this to keep the model simple; in any case, neighborhood quality was subjectively specified in the EHCS data and thus difficult to predict generally. Age of house in which the household respondent resided was available in the survey, but distance from the CBD was not, and this constitutes a problem. Each individual was not coded by exact location in the data set, but located by Borough, of which there are 33 in Greater London. What we have done in measuring distance is to simply locate a centroid in each Borough and use airline distance from this to a point in the City.

Another consideration involved the fact that when we embed the discrete choice models into the large-scale (fractal) simulation, we require data on

age of housing and distance from the CBD at every conceivable point of residential development in Greater London. These data are amongst the easiest to obtain from independent sources. We used an age distribution for housing measured over seven levels which was available from the (then) Greater London Council (GLC) Intelligence Unit Library. Distance is measurable directly from the map, while neighborhood quality, although available from the GLC, did not appear to match that used in the EHCS, and was thus excluded at an early stage of model estimation.

The general form of the models we have estimated, in log-odds form, is

$$\log \frac{P_{iu}}{P_{i1}} = \lambda_{u0} + \lambda_{u1} R_n + \lambda_{u2} Q_i, \; i \in Z_n, \; u = 2, 3, 4. \tag{4.21}$$

The log-odds equation is normalized with respect to the probability of choosing a purpose-built flat, P_{i1} and the other choices involved converted flats ($u = 2$), terraced houses ($u = 3$) and detached/semi-detached homes ($u = 4$). Q_i is the age of the dwelling in which individual i resided and R_n is the distance from the CBD to the centroid of the Borough in which i resides.

In essence, we assume that R_i is unobserved and that equation (4.21) is an appropriate approximation to the underlying discrete choice model analogous to equation (4.21) in which R_i replaces R_n. Equation (4.21) will only be acceptable if R_n is the mean distance, and the sum of the differences around R_n in the Borough cancel. Formally, if $R_i = R_n + \epsilon_i$ where ϵ_i is the 'error' difference between the mean and the actual distance to individual i, the average R_n can be defined in terms of R_i as

$$\sum_{i \in Z_n} R_i / N_n = R_n + \sum_{i \in Z_n} \epsilon_i / N_n. \tag{4.22}$$

Z_n is the spatial definition of the Borough n and N_n is the number of individuals in Z_n. From equation (4.22), the mean will only be equal to R_n if $\Sigma_{i \in Z} \epsilon_i = 0$, that is the errors around the mean are self-canceling in total. We cannot explore the detailed implications of this aggregation further, but it is important to further research. Discrete choice theory is strangely deficient in clear discussion of the spatial aggregation problem, with the exception of important work by Anas (1981, 1982, 1983).

Before we broach questions of model selection and estimation, we will sketch how the model we are working with could be developed in nested fashion, to account not only for the aggregate form of the distance data, but also for more substantive questions related to the sequence of spatial decision-making. Because distance from CBD is only available at Borough level, it might make sense to conceive the house type–residential location process as one in which a choice of neighborhood type is made first on the Borough (Z_n) level in terms of neighborhood quality and distance from the CBD and then the choice of house type made at the individual location with respect to age. Such a model could be written as

$$P_{iqk} = P_{iq} P_{ik|q},$$

where P_{iqk} is the probability of an individual i choosing neighborhood type q and house type k, P_{iq} the probability that the individual chooses neighborhood type q at Borough level and $P_{ik|q}$ the probability the same individual then chooses house type k, having chosen neighborhood type q. Such a

sequence could be structured so that the fractal simulation enabled neighborhood type to be chosen at an appropriate level of fractal resolution, house-type at a lower level. Although neighborhood type is predicted here, this could be suppressed if it were regarded as only an intermediate variable of little visual significance. There are many issues to resolve here, but some work along these lines in an industrial location context by Hayashi and Isobe (1985) looks promising, as does the theoretical work of Roy (1983). Nested models of this type need to be pursued in extensions to these applications.

4.7 Model Selection and Estimation

We developed a number of preliminary specifications of the model before we decided upon equation (4.21). We first estimated some models based on housing tenure but then dropped these in favor of house type when our ideas relating to urban structure became clearer. We began with five categories of house type including miscellaneous but dropped this when it appeared non-significant in explanation. We then estimated the house type model with all combinations of up to three exogenous variables: age and distance which we eventually selected, but also neighborhood quality. With three variables, there are seven models in all which can be specified and the global fit of each of these seven is given in Table 4.1.

By far the best of the models, indicated in bold type in Table 4.1, are the two which include the age and distance variables. These models in fact are the only ones which reach the threshold of acceptability in which $\xi^2 > 0.2$ suggested by Hensher and Johnson (1981). The best model also includes neighborhood quality but the percentage increase in fit between the model without this variable and that with is less than 5% and thus neighborhood quality has been omitted. Other reasons relate to the fact that neighborhood quality is difficult to produce in a consistent and comprehensive data base for London, and to the fact that in our fractal simulations we have severe memory limitations, which means we need to hold both input and output

Table 4.1. Global fits of models incorporating age, distance and neighborhood quality

Independent variables	ξ^2
Age	0.118
Distance	0.089
Neighborhood quality	0.069
Age and distance	**0.207**
Age and neighborhood quality	0.123
Distance and neighborhood quality	0.095
Age, distance and neighborhood quality	**0.218**

Bold type indicates acceptable models within the Hensher–Johnson Limit $\xi^2 > 0.2$.

data in screen memory simultaneously. This limits the number of variables we can deal with and thus neighborhood quality was felt to be dispensable.

We will now examine the discrete choice model estimated for the age–distance variables in equation (4.21). The three fitted equations are given as follows, with the notation: purpose-built flats: $u = 1$, converted flats: $u = 2$, terraced houses: $u = 3$ and detached/semi-detached houses: $u = 4$.

$$\log \frac{P_{i2}}{P_{i1}} = \begin{matrix} -4.862 \\ \{-7.984\}^* \\ (0.609) \end{matrix} + \begin{matrix} 0.034\, R_n \\ \{0.754\} \\ (0.045) \end{matrix} + \begin{matrix} 0.067 Q_i, \\ \{10.599\}^* \\ (0.006) \end{matrix} \tag{4.23a}$$

$$\log \frac{P_{i3}}{P_{i1}} = \begin{matrix} -3.605 \\ \{-9.818\}^* \\ (0.367) \end{matrix} + \begin{matrix} 0.177 R_n \\ \{6.735\}^* \\ (0.026) \end{matrix} + \begin{matrix} 0.052 Q_i, \\ \{11.439\}^* \\ (0.005) \end{matrix} \tag{4.23b}$$

$$\log \frac{P_{i4}}{P_{i1}} = \begin{matrix} -5.737 \\ \{-12.143\} \\ (0.472) \end{matrix} + \begin{matrix} 0.354 R_n \\ \{11.379\}^* \\ (0.031) \end{matrix} + \begin{matrix} 0.046 Q_i, \\ \{9.102\}^* \\ (0.005) \end{matrix} \tag{4.23c}$$

where $\xi^2 = 0.207$ and $N = 809$. t statistics are shown in curly brackets, significance being denoted by an asterisk; standard errors are shown in parentheses.

Note that the log-odds is essentially the log-likelihood that individual i will select the numerator alternative rather than the denominator alternative. In view of the aggregated nature of the distance data, the ξ^2 of 0.207 indicates a reasonable degree of overall fit, whilst the variable parameters and their corresponding t statistics lend support to our *a priori* expectations. Equations (4.23b) and (4.23c) imply that both terraced and detached/semi-detached are likely to be further from the CBD and to be older than purpose-built flats; and equation (4.23a) suggests that converted flats are likely to be older than their purpose-built counterparts.

These interpretations can only be borne out by a full-scale simulation and the pattern of coefficients suggests that flats of both kinds are nearest to the CBD, while terraced, then detached/semi-detached houses are further away, assuming that terraced are older than detached/semi-detached. As we intend the simulation to be entirely spatial, and spatial structure is not apparent from the model fits presented so far, we need to see how well the models perform spatially at an aggregate level first. The obvious level on which to perform such spatial analysis is similarly aggregate. We will present our analysis visually in the next section where the models' predictive success indices are mapped for the 33 Boroughs.

We have already shown that it is necessary to aggregate individual predicted probabilities so that we can enable some comparison with the observed data. To this end, we introduced McFadden's (1979) predicted success matrix in equations (4.11) to (4.15), and then presented various indices of success in which correct proportions, and differences between

observed and predicted choices were computed in equations (4.16) to (4.20). However, it is possible to compute equation (4.13), the numbers of persons originally choosing k and predicted to choose u, for subsets of individuals, in particular individuals residing in certain zones, in this case Boroughs Z_n. In all the indices which follow equation (4.13), N_{ku} is replaced with

$$N_{kun} = \sum_{i \in Z_n} F_{ik} P_{iu}, \tag{4.24}$$

where N_{kun} is the number of individuals originally choosing k and predicted to choose u in Borough Z_n.

The proportion of correct predictions defined in equations (4.16) and (4.17) for the whole of Greater London can act as a basis for comparison with their zonal equivalents. These statistics were computed using the model in equation (4.23) as

$\eta_1 = 0.533$, $\eta_2 = 0.198$, $\eta_3 = 0.433$, and $\eta_4 = 0.397$.

These indices seem rather low; only in the case of purpose-built flats is there a better than 50% success rate, and converted flats are poorly predicted. The overall percentage of correct predictions from equation (4.17) is computed as $\eta = 0.432$ which is an appropriate average of $\{\eta_k\}$. The spatial (zonal Z_n) equivalents of η_k called η_{kn} are mapped across the 33 Boroughs in Figure 4.3 (note that in all these types of map, the City Borough does not contain any observations and thus is not shaded). These percent-correct predictions show a much wider range of variation. In general, purpose-built flats are better predicted closer to the CBD, while the reverse holds for detached/semi-detached houses. The distribution of converted flats generally shows a low percent prediction with a slight increase towards the CBD while terraced houses show a less distinctive spatial pattern with a slight increase in performance towards the periphery. In fact, Figure 4.3 contains the clearest demonstration we have that individual choice behavior varies spatially. The obvious conclusion is that there are two sets of models, one for inner, the other for outer London, but before we consider these further, we will examine other indices of predictive success.

Indices of the percentage difference between observed and predicted choices given by equation (4.18), $\{\phi_k\}$ have been computed in spatial equivalent form and are mapped in Figure 4.4. The patterns are much less clear than those in Figure 4.3. For purpose-built flats, the largest differences are in the inner suburbs, and the smallest in the center and the west. For converted flats, the pattern is much more random with a slight bias towards higher differences in the inner suburbs. For terraced houses, the inner suburbs show higher levels of under- and over-prediction in the cases of detached/semi-detached houses. These maps are more difficult to interpret than their counterparts in Figure 4.3. What they do show, however, is that there are both sectoral and concentric–geometric spatial biases in the pattern of predictions which can only be accounted for by the addition of new and different explanatory variables (and the possible deletion of one of the existing ones), or the development of models which accept these spatial differences. We will pursue the latter course.

To conclude, it is useful to examine the pattern of overall correct predictions from equation (4.17), computed and mapped spatially, and this is

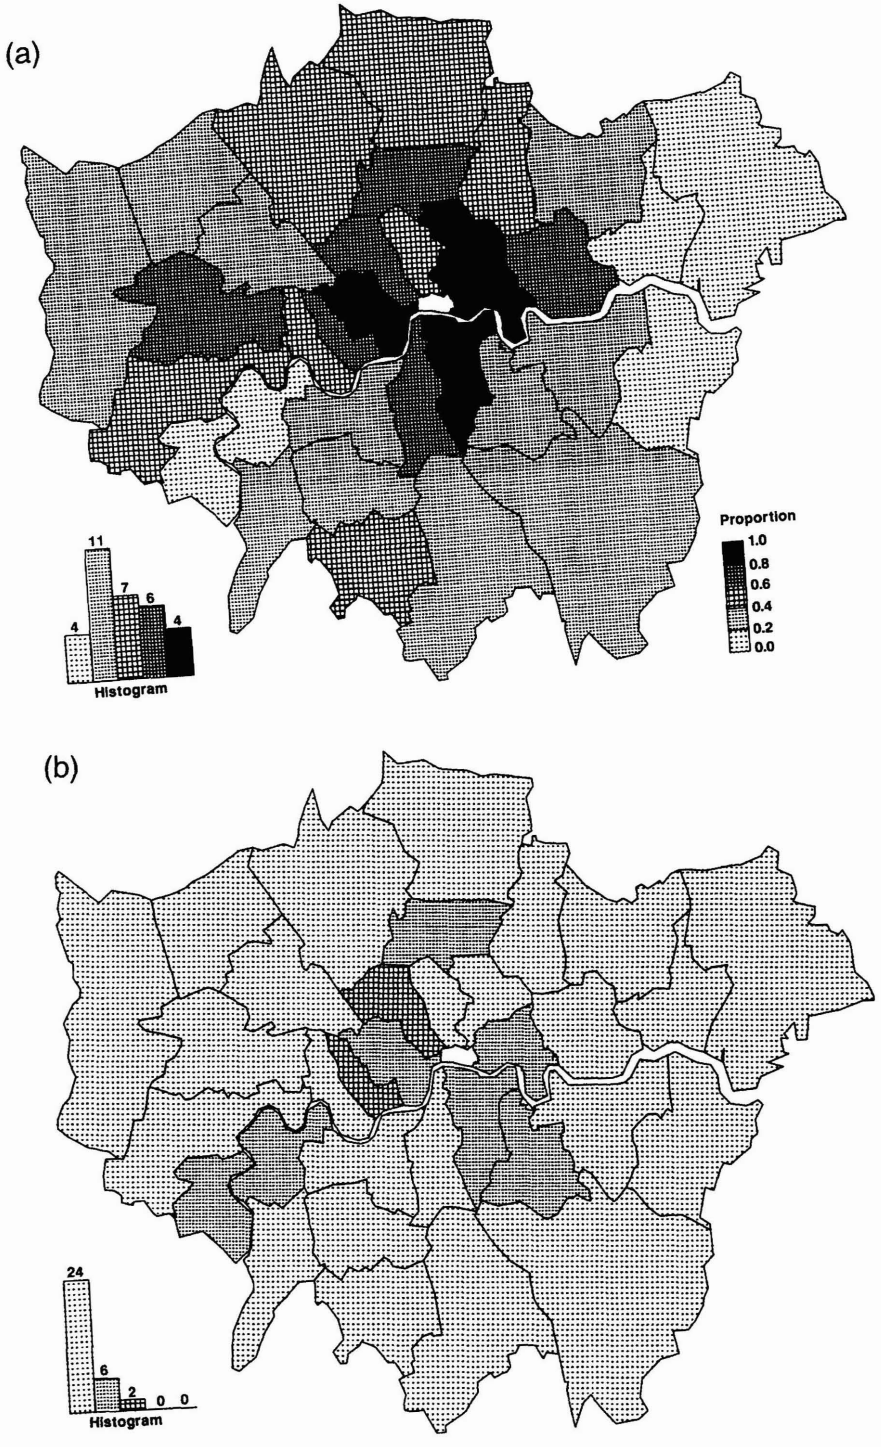

Figure 4.3. Proportions of correct choices of house types (*cont.*): (a) purpose-built flats; (b) converted flats.

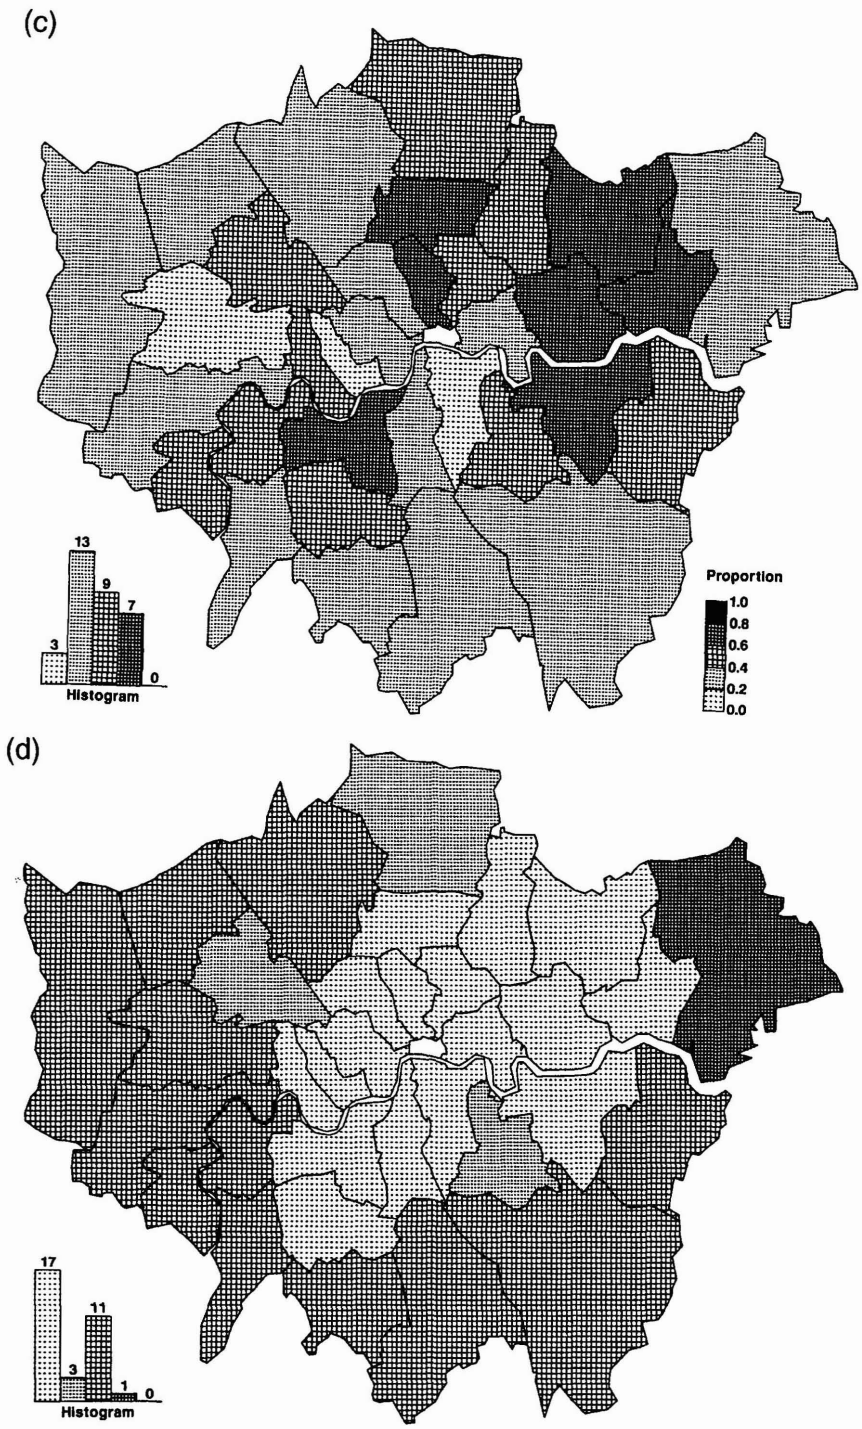

Figure 4.3. Proportions of correct choices of house types: (c) terraced houses; (d) detached/semi-detached houses.

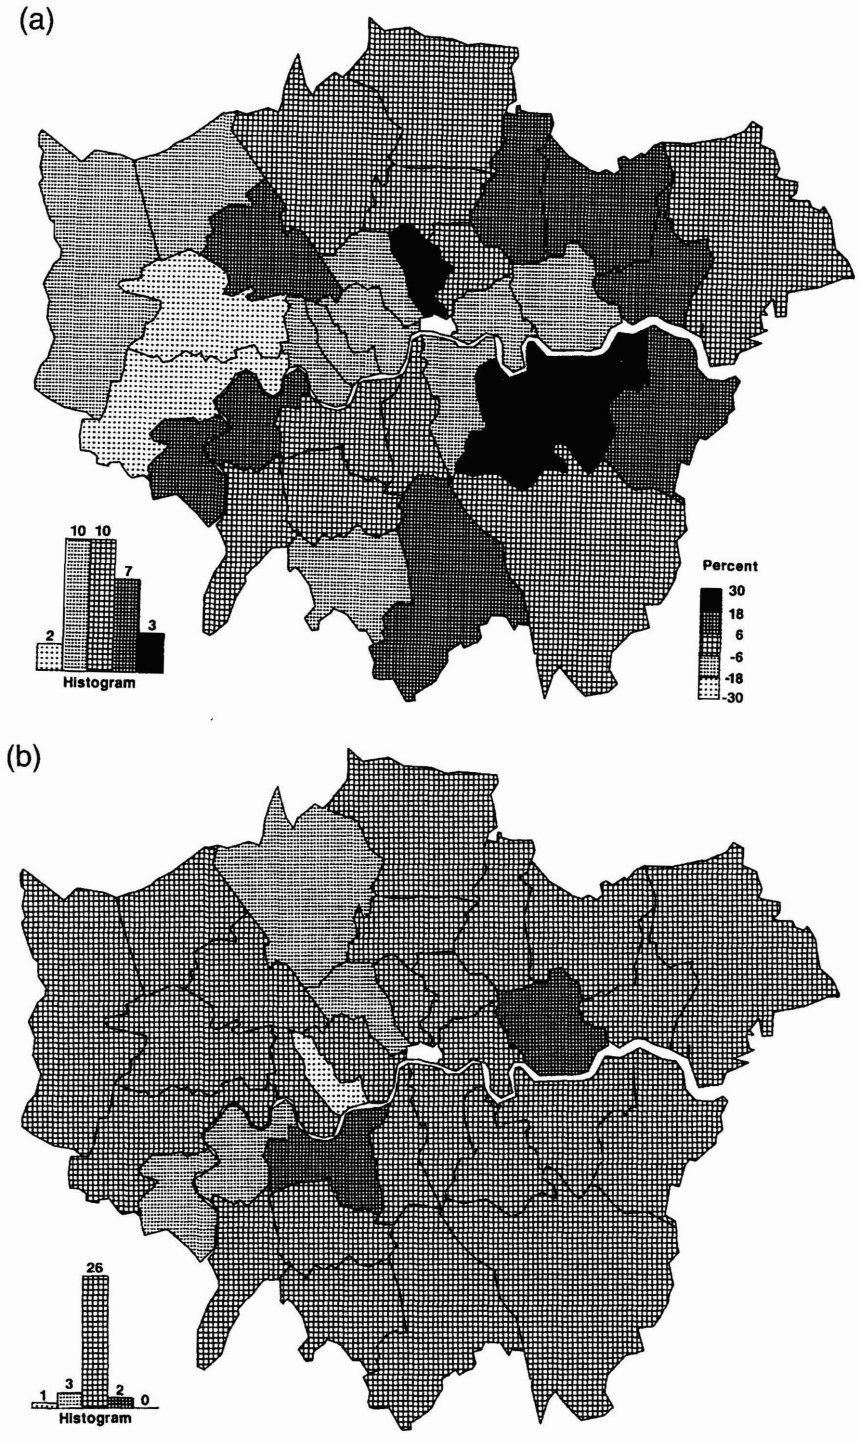

Figure 4.4. Differences in observed and predicted housing choices (*cont.*): (a) purpose-built flats; (b) converted flats.

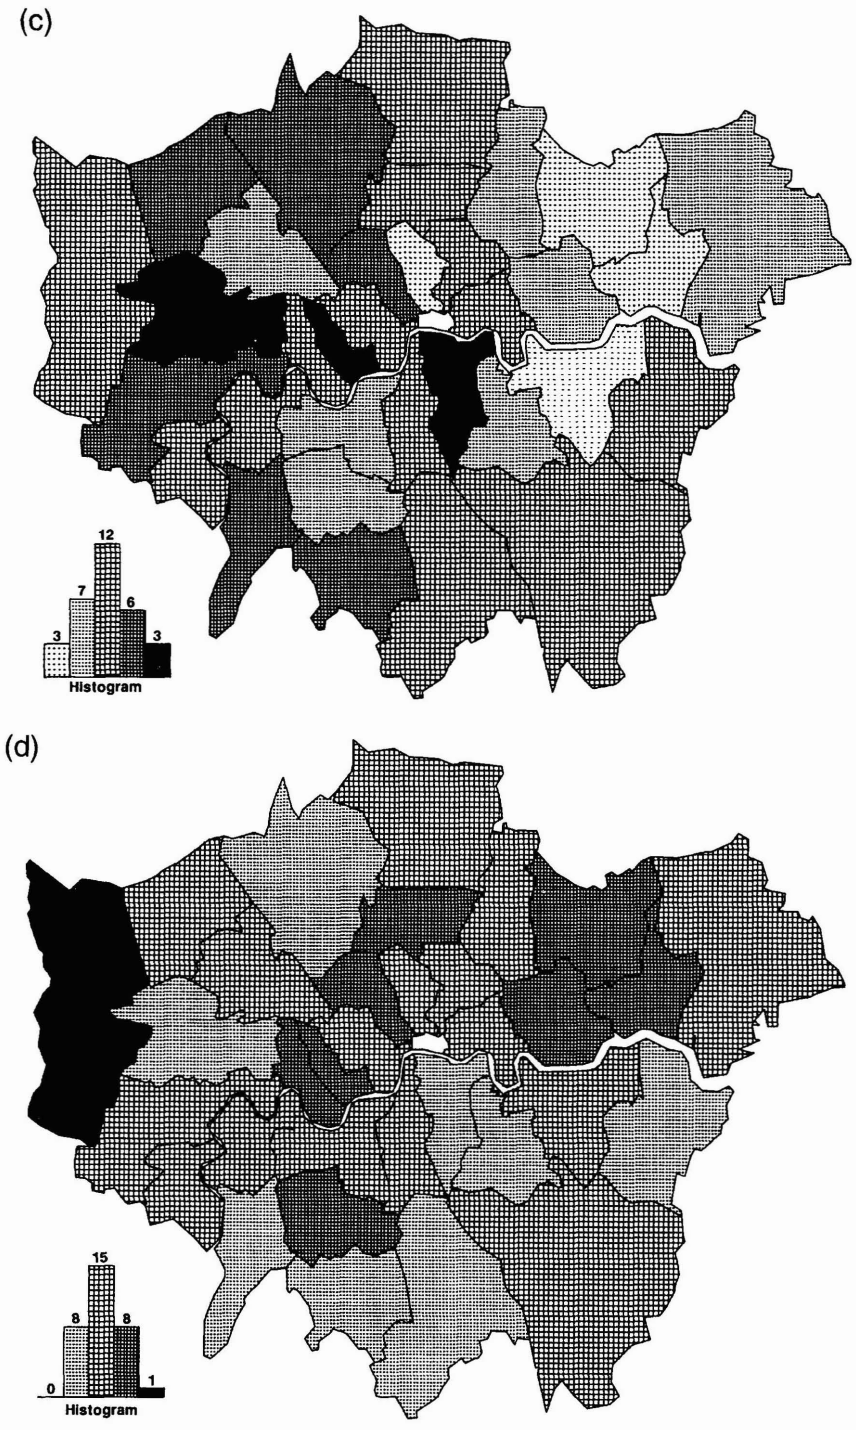

Figure 4.4. Differences in observed and predicted housing choices: (c) terraced houses; (d) detached/semi-detached houses.

presented in Figure 4.5. The best predictions are recorded in and near the center, and in the outermost suburbs. This suggests the need for two separate models of individual choice behavior, one for inner, the other for outer zones. The need for this distinction is even clearer when the normalized success index computed from the spatial equivalents of equations (4.19) and (4.20), and defined as $\varphi_n/\varphi_{n,\max}$, is examined. This is mapped in Figure 4.6, and shows that the best predictions occur nearest the CBD, the worst in the far western and eastern suburbs. On this basis we decided to reestimate our models based on equation (4.21) for inner and outer London, where inner London is based on the 13 Boroughs which compose the Inner London Education Authority (ILEA).

The sample size of 809 observations was divided into 337 based on the inner Boroughs, the remaining 472 comprising the outer Boroughs. First equation (4.21) for the inner Boroughs was estimated as

$$\log \frac{P_{i2}}{P_{i1}} = \begin{matrix} -5.446 \\ \{-5.849\} \\ (0.931) \end{matrix} + \begin{matrix} 0.194R_n \\ \{1.478\} \\ (0.131) \end{matrix} + \begin{matrix} 0.061Q_i, \\ \{8.305\}^* \\ (0.007) \end{matrix} \tag{4.25a}$$

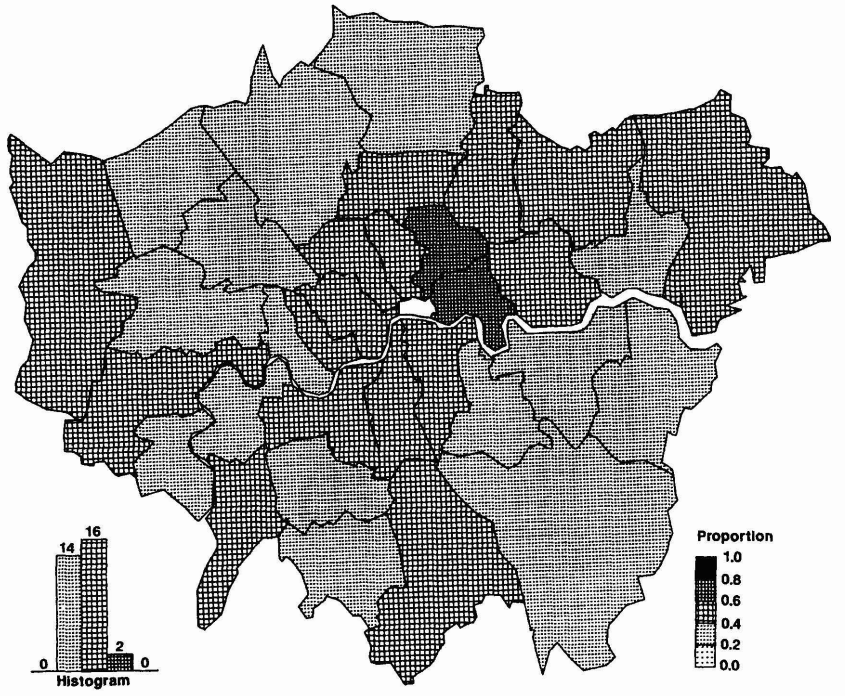

Figure 4.5. Proportions of correct choices for all house types.

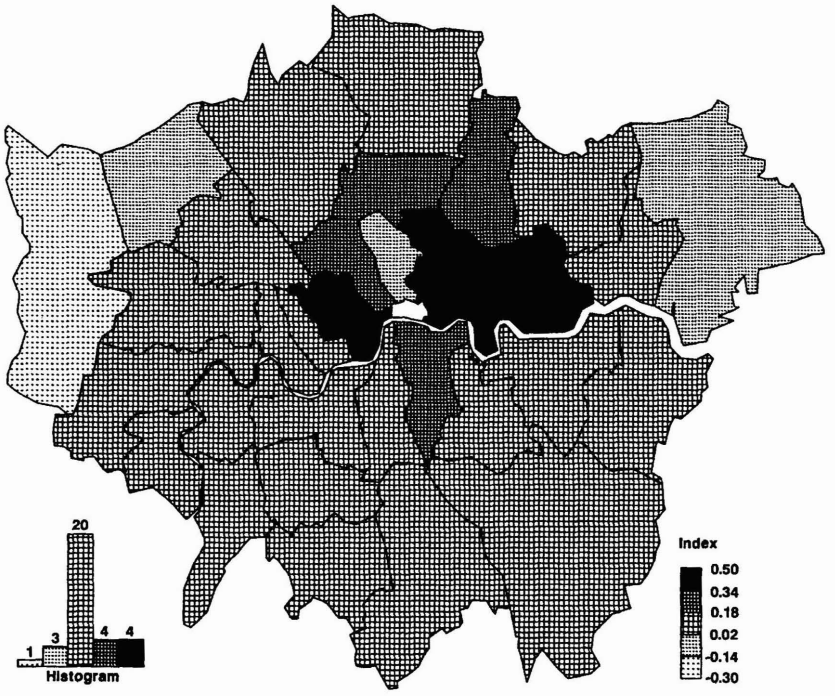

Figure 4.6. Overall normalized success indices for all house types.

$$\log \frac{P_{i3}}{P_{i1}} = \underset{\{-7.046\}^*}{\underset{(0.725)}{-5.106}} + \underset{\{4.106\}^*}{\underset{(0.105)}{0.430R_n}} + \underset{\{8.722\}^*}{\underset{(0.005)}{0.050Q_{i'}}} \tag{4.25b}$$

$$\log \frac{P_{i4}}{P_{i1}} = \underset{\{-5.810\}}{\underset{(1.279)}{-7.430}} + \underset{\{3.203\}^*}{\underset{(0.171)}{0.546R_n}} + \underset{\{5.767\}^*}{\underset{(0.009)}{0.050Q_{i'}}} \tag{4.25c}$$

where $\xi^2 = 0.228$ and $N_{\text{inner}} = 337$; t statistics and standard errors denoted as above; and the appropriate equation(s) for the outer Boroughs estimated as

$$\log \frac{P_{i2}}{P_{i1}} = \underset{\{-3.307\}^*}{\underset{(1.548)}{-5.119}} + \underset{\{0.280\}}{\underset{(0.106)}{0.030R_n}} + \underset{\{6.514\}^*}{\underset{(0.012)}{0.075Q_{i'}}} \tag{4.26a}$$

$$\log \frac{P_{i3}}{P_{i1}} = \underset{\{-1.754\}}{\underset{(0.741)}{-1.300}} - \underset{\{0.170\}}{\underset{(-0.051)}{0.009R_n}} + \underset{\{7.450\}^*}{\underset{(0.008)}{0.059Q_{i'}}} \tag{4.26b}$$

$$\log \frac{P_{i4}}{P_{i1}} = \begin{matrix} -3.537 \\ \{-4.740\} \\ (0.746) \end{matrix} + \begin{matrix} 0.186R_n \\ \{3.816\} \\ (0.049) \end{matrix} + \begin{matrix} 0.051Q_{i'} \\ \{6.538\}^* \\ (0.008) \end{matrix} \tag{4.26c}$$

where $\xi^2 = 0.106$; $N_{\text{outer}} = 472$. What becomes apparent in terms of the t and ξ^2 statistics is that the inner London model (equations (4.25)) performs very much as the original model (equations (4.23)), whilst the outer London results are rather different. Equation (4.26) reflects a diminished role of the distance variable (two of its associated t statistics are insignificant and one parameter exhibits an unexpected sign) which contributes towards a much lower ξ^2 goodness-of-fit measure.

We might rationalize this in terms of our previous land use theory as follows: whilst difficulties of physical accessibility constrained the physical growth of London up until the First World War, the subsequent innovation of mass transit and the automobile rapidly opened up large tracts of land for development. Because most of this development occurred over very large areas, the form of physical development is much less likely to exhibit a very close and identifiable correspondence with distance from the CBD. Reestimation of our model for outer London without the distance variable yields

$$\log \frac{P_{i2}}{P_{i1}} = \begin{matrix} -4.695 \\ \{-7.388\}^* \\ (0.635) \end{matrix} + \begin{matrix} 0.075Q_{i'} \\ \{6.490\}^* \\ (0.011) \end{matrix} \tag{4.27a}$$

$$\log \frac{P_{i3}}{P_{i1}} = \begin{matrix} -1.399 \\ \{-4.990\}^* \\ \{0.280\} \end{matrix} + \begin{matrix} 0.059Q_{i'} \\ \{7.470\}^* \\ (0.008) \end{matrix} \tag{4.27b}$$

$$\log \frac{P_{i4}}{P_{i1}} = \begin{matrix} -0.881 \\ \{-3.347\}^* \\ (0.263) \end{matrix} + \begin{matrix} 0.048Q_{i'} \\ \{6.241\}^* \\ (0.008) \end{matrix} \tag{4.27c}$$

$\xi^2 = 0.078$ and $N_{\text{outer}} = 372$; t statistics and standard errors denoted as above. The ξ^2 statistic is less than that for equations (4.26) and thus the model has not been used in the simulations which follow. At this point we can conclude our section on estimation. Many avenues remain unexplored but several models have been tested and we will take forward those in equations (4.23), (4.25) and (4.26).

4.8 Fractal Simulation of House Type and Location in London

One of the more obscure reasons for developing such a simplified model based on age and distance can now be made clear. Age is a spatially extensive variable, while distance is a property of space itself. Thus it is possible to display a single map shaded according to age from which distance can also be read, in particular distance to some fixed point from any other. If we had more than a single extensive variable, age and neighborhood quality say, these could not be represented on the same map in easily codable form. Clearly it is useful for ease of interpretation to have a single map of input data, for this can be directly associated with a map of the outputs from the models. In fact, the need to store data in map form is essential, for the fractal simulation was run on a very small microcomputer in which only 8K of memory was available for program and data, 20K of memory being given over to the graphics screen. Although it might be possible to store data on disk, and thus include a larger number of independent spatial variables, the continual reading and writing required would make the operation of the model prohibitively slow. In fact, because the data are spatially extensive, it is essential to store them in screen mode, for the resolution we are working with involves 160×256 pixel points which makes any form other than screen storage extremely problematic. The data on age are stored as a screen map, and airline distance is easy to compute from screen coordinates which in turn are a function of the screen addressing.

The age data were made available by the GLC Intelligence Unit in seven age groups which were coded in grid fashion, and colored in the screen memory according to the age group. The screen map is shown in Plate 4.3 where the colors refer to the age of housing. The following average ages in years define the seven ranges in question, $\rightarrow 8 \rightarrow 26 \rightarrow 48 \rightarrow 78 \rightarrow 110 \rightarrow 150 \rightarrow 175 \rightarrow$, and these are colored white, light blue, magenta, dark blue, yellow, green and red respectively. These represent weighted averages which reflect the distribution of housing in any grid square. Distance from the CBD to Borough centroids is measured in kilometers, the GLC boundary being about 24 km maximum from the City and the ILEA boundary used for the inner London model being about 13 km distant. Note also that the shape of urban development in London is coded into the data through grid squares colored on a black background which does not contain housing. These represent 'vacant' land in the sense used earlier, although in these applications, the model in no way predicts this.

The way the simulation works involves first loading the age map into the screen memory from file. Then the fractal simulation begins in the order used previously in the demonstration model, and when the appropriate level of fractal detail is reached, the program retrieves the color of the centroid of the triangle space reached, from the screen, converts this into an age value, computes distance and uses these variables in the model structure based on equation (4.21) to compute the probability of house type. Thus the simulation works by replacing the regular gridded age map by the irregular fractal land use pattern in a literal sense. This rather innovative

technique for input is immediately converted to output and this occurs directly 'before your eyes'. In a sense, it is a version of the WYSIWYG principle ('What You See Is What You Get') which is central to many operations with graphics computers. A note on technical detail is required. The simulation operates on a display with resolution 160×256 pixels in 16 colors. Eight colors are reserved for the age map (seven ages and one vacant land use) and five are used in choosing house type (four types and one vacant land use). The process of replacement is not as clear as it might be because only eight absolute colors are available, hence the replacement of the input map with the output map uses similar colors and is only distinguished in terms of its irregularity.

The process of fractal simulation is essentially the same as that used previously in Chapter 3 in the 'London' sequence. Moreover, we have also used the triangular midpoint displacement technique for fractal rendering which was shown there in Figure 3.14 for the hypothetical demonstrations. The only difference relates to the way the input data are stored and sampled and the way the probability models are developed. Four land use types based on housing, rather than three based on activities, now form the simulated urban structure. The area over which the simulation is operated is fixed and in a sense residential location is already predetermined through the data, and thus it is only house type by location which varies.

We have already noted that two model structures are to be used: that based on the whole of Greater London using equations (4.23) and that based on the distinction between inner and outer London based on equations (4.25) and (4.26). In these simulations, we work at recursive level $s = 5$ which essentially fixes fractal detail at just above the pixel level of the screen. Each simulation takes about three hours and involves examining $10 \times 4^s = 10{,}240$ randomly positioned contiguous triangles which form the network of fractal detail at the lowest level of resolution. In fact, the models are based on 809 data points, and in the area in question there are in excess of three million households, thus the simulation itself is still very much in the nature of a sample-style exercise in which an 'average' individual residing at the lowest level of fractal detail makes a house-type choice which is then assumed to be typical of all individuals at that level and in the space which contains that location.

The other issue involves the conversion of probabilities $\{P_{ik}\}$ into discrete choices. In the demonstration model, a random simulation was adopted in which choice of land use was accomplished according to the probability range fixed by the land use models but ultimately determined using a random number device. The resultant outputs were very satisfactory because the probability profiles were quite distinct, thus enabling fairly clear decisions to be made and characteristic spatial patterns to emerge. Here, however, the probability profiles of the house type models are much less different from one another, and thus to develop clearer spatial patterns, we have also used a deterministic simulation. This simulation is based on choosing a house type according to the rule

$$\text{Type} \leftarrow \max_k \{P_{ik}\} \tag{4.28}$$

which simply makes the choice according to that alternative which has the maximum probability for individual i.

We can now show the simulations. We will first discuss the random simulations which are based on equation (4.23), then equations (4.25) and (4.26), but we will not show these visually as they do not generate much imagery of import. The main impression is one of massive variability of house type in spatial terms. There is almost a complete mix of types everywhere for both types of equation, thus implying that the relative evenness and similarity of the probability profiles gives much greater weight to the lower probabilities in each choice situation than would be the case in a real context. Little spatial pattern can thus be discerned and this suggests that random simulations based on discrete choice models are likely to produce too little spatial discrimination if predicted in this way.

The deterministic simulations which involve equation (4.28) are shown in Plate 4.4(a) and (b) for the full, and inner–outer models respectively, where the four colors – red, yellow, green and blue – reflect converted flats, purpose-built flats, terraced houses, and detached/semi-detached houses respectively. Very clear spatial patterns emerge this time which show the characteristic structure of residential land use in London, but there is little difference between the two sets of model. The clearer of the two patterns is Plate 4.4(a) based on the full model, but there is a ring of purpose-built flats between the terraced and detached/semi-detached areas which is unexpected. In Plate 4.4(b), purpose-built flats are closer in towards the CBD. Note that in the simulations the total number of house choices is not scaled in any way to reflect the scale of housing in London; thus this represents an additional prediction from the model. The patterns in general though are very plausible, reflecting flats, terraced and detached/semi-detached houses at increasing distance from the CBD, with the distribution of purpose-built and converted flats clearly characterizing the flat-market in London. One limitation of the deterministic model is that it does not pick up the degree of local variation one might expect, but a more detailed data base might resolve this.

Finally, we have begun to experiment with these simulations. Running the models at $s > 5$ requires a larger processor because the memory required explodes due to the recursion, and so far we have run the model up to $s = 7$, although the increase in time required is exponential. Level of recursion does affect the patterns we get, but generally these help us to improve the ultimate look of the geometry, not the models themselves. Simple policy-predictive runs of the simulations are possible, for the input data are easy to update. One could assume a process of aging and renewal, varying according to simple rules and policies, which would then enable a pseudo-dynamic simulation to be developed. A series of images of the typical house types in London over the next 50 years could be generated in this way. But these are for the future, and in any case, there are many lines of inquiry that have to be followed up before then.

4.9 Extending the Laboratory for Experimentation and Visualization

The ability to display the overall pattern produced by models with an implicit spatial dimension is a clear advantage of the large-scale simulations adopted here. But these need not be generated within a fractal framework. Simulation could proceed by examining each pixel in turn and building up urban structure in this way on a regular spatial grid. Nevertheless, fractals do generate realistic images, and one of the goals of this chapter has been to make abstract models more visually intelligible and acceptable, and for this, the fractal framework seems promising. As such, the technique is one of generating spatial realism, and this clearly depends upon the display devices used. The main problem emerging from this chapter, however, relates to the development of a more consistent modeling strategy which can be effectively incorporated into the hierarchical method used to structure the simulation. We have already indicated what is involved: in essence, the hierarchy guiding the fractal simulation should be based on characteristics of the city, and this clearly relates to the type of explanation and modeling required. Discrete choice models show promise here, but so do sequential and nested approaches involving entropy models (Batty, 1976).

This reasoning leads us to the conclusion that a more fundamental strategy may be actually to explore land-use models which are themselves fractal. Some examples already exist in physical geography: for example, the sorts of terrain model explored by Goodchild (1982) and illustrated in Chapter 3, and image processing techniques such as those developed by Pentland (1984) are suggestive of the types of stochastic model that might underlie the structure of land use. There are difficulties in that some of the patterns are discontinuous, but it is worth exploring how such ideas could be used to link what we already know about land use, central place, and rank-size together in a fractal framework. With respect to discrete choice models, there may even be the possibility of a fractal interpretation of the underlying mechanisms which give rise to various forms of logit and probit models, and there is clearly a possibility that questions of nesting and aggregation might be reconciled with ideas about recursion and hierarchy. In fact, in this chapter, the whole question of the spatial basis of discrete choice models has emerged as problematic, and this suggests that further research on spatial aggregation and discrete choice is worthy as an end in itself, notwithstanding any fractal interpretations which might emerge.

Many other speculations are possible about where such developments will best be focussed. An interesting project would be to examine the extent to which regular, non-random fractal patterns built from cell-space models (Tobler, 1979b; Couclelis, 1985) could be used as first approximations to city patterns, and there is much work now developing in this domain around concepts of cellular automata and artificial life. We also need to consider how such simulations might be made dynamic, especially as there is an obvious dynamic process underlying a model in which age acts as an independent variable. In one sense, our models might already be seen as explaining urban structure in terms of time and space, age and distance,

and our earlier comments on possible policy simulations endorse this. In particular, the question of redevelopment is central to residential location, and any dynamic extension to the framework should enable such processes to be captured. We will explore these ideas more fully from Chapter 7 onwards, but we also require a firmer empirical basis to our assertion that urban structure is indeed fractal, in order to inform both description and theory. We will begin our assault on this measurement task in the next two chapters.