6

The Morphology of Urban Land Use

Nature uses only the longest threads to weave her patterns, so each small piece of her fabric reveals the organization of the entire tapestry. (Feynman, 1965, p. 34.)

6.1 Inside the Fabric of the City

As soon as we turn our attention to the geometrical composition of a city's land use, the urban boundaries with which we have been working, reveal themselves to be both crude and simplistic descriptors of urban form. Inside these envelopes lies a rich mix of heterogeneous activities and uses which are often easier to distinguish from one another than 'urban' is from 'rural' but which belie a level of complexity that threatens to destroy the most sustained attempt to classify their geometry. New problems of boundary definition arise where different land uses, clearly embodying different processes of development, have common edges, and thus the problem becomes one of knowing how to distinguish different processes from a geometry which shows itself in only one form. The problem of defining fractal objects which are spatially adjacent or contiguous to one another becomes central, and thus introduces the tantalizing specter of fractal objects which are clearly different geometrically at one level but when aggregated to the next, compose higher-order objects which have their own integrity and unity. It is in this sense then that the tapestry which Richard Feynman (1965) refers to above is woven from threads which reveal themselves at the lowest level. This chapter will be concerned with identifying how these threads which we defined as entire boundaries to urban development in the last chapter, compose the fabric of the city at the more detailed level of its land use.

So far in our analyses, we have focussed upon the difficulties inherent in measuring geographical boundaries to satisfactory levels of precision, and we have also addressed the difficulties in obtaining objective and consistent definitions of categories of urban land use. We have produced some limited evidence to suggest that the ambiguities inherent in defining 'irreversibly urban' phenomena and the subjective nature of boundary encoding are not in themselves sufficient to impede us in observing temporal trends in the changing fractal dimension of urban boundaries. In the spirit of fractal measurement primarily in the natural sciences, in the last chapter we

developed and extensively tested four different algorithms for measuring irregularity, drawing conclusions as to the strengths and weaknesses of these alternative procedures. In this chapter, we will use these methods to explore the fractal nature of the more detailed urban fabric. This amounts to a broad conceptual treatment and a thorough technical exposition, yet there is an important sense in which our analysis remains a simplistic treatment of geographical tessellations and land use categories. That is, whilst measurement in science can, in many circumstances, be considered to concern physically and geographically isolated structures, this assumption clearly becomes strained where the subjects of our measurements constitute juxtaposed contiguous areal units, which are embedded within an overall geographical matrix.

Viewed from this more holistic perspective, we might expect measurements of line character to reflect predominantly the processes that have molded the form of each pair of adjacent land parcels, or have embodied both sets of processes in more or less equal amount. We are not aware of fractal measurement that explicitly acknowledges the role of boundaries as mediators between adjacent categories, for in the mainstream, such phenomena are considered to be the edges of geographical isolates. From this new standpoint, the coastlines in Chapters 2 and 3 should also reflect characteristics of both adjacent media, that is, the lithology and structural geomorphology of the land mass, and the erosional and/or depositional characteristics of the water body (Kaproff, 1986; Turcotte, 1992). In a system of contiguous land use areas that compose an urban settlement, the boundary to each use will similarly always consist of parts of the boundary of other uses. In this chapter we will begin to address this issue and in so doing, will raise, but not resolve, some severe conceptual problems for the first time.

We will begin by providing a brief summary of the fractal relations we seek to define, in particular the area–perimeter relations which are central to this chapter as well as the perimeter–scale relations which we discussed extensively in the last. We will also present two formulations of scale dependence which we will apply to each of these relations. The application of these methods to land use boundaries in the English town of Swindon is then introduced by first describing the characteristics of the urban area in question. Fractal dimensions based on area–perimeter relations across scales are estimated, and these same dimensions are next derived by examining scale changes within the digital representation of the perimeters themselves. Finally, individual dimensions of each of the land use parcels can be classified on the basis of their fractal dimension. The analysis contains some inevitable ambiguities, but it is clear that careful measurement is required in all such applications, and thus we see this exposition as charting the ground rules for fractal measurement in this domain. Moreover, the sensitivity of the analysis to measurement differences casts considerable doubt on many of the results from applications of fractal geometry presented to date in a variety of other fields.

6.2 Area–Perimeter Relations and Scale Dependence

As we indicated at the end of Chapter 5, in moving to a more detailed level of spatial resolution, we now require methods which will not only compute the fractal dimension of a single object, but of many objects, in this case land uses composing distinct sets. To do this, we must supplement the perimeter–scale relations given in Chapter 2 as equations (2.24) and (2.25), and in Chapter 5 as (5.5) and (5.7) with the relationship between perimeter and its area. This area–perimeter relation was derived in Chapter 2 as equation (2.29) and we will begin by discussing its relevance to the application posed here. In this section, we will also repeat, for the reader's convenience, the perimeter–scale relations which we used in the last chapter.

In Euclidean geometry, a measure of size in a given dimension will scale directly with a measure in another, for example in an adjacent dimension, and this scaling will be some product of the dimensions themselves. Consider area A and volume V based on two and three dimensions respectively. Area has a size calculated as the square of the line measure L, that is L^2, while volume has a size L^3. If it is required to derive area from volume, it is clear that this can be done as $A \propto V^{2/3}$. In the same way, if it is required to derive the line L (which we will henceforth term more familiarly the perimeter) from area A, the relation is

$$L \propto A^{1/2}. \tag{6.1}$$

All relationships such as those implied by equation (6.1) show that size in one dimension can be scaled directly by knowing the dimension of the object in a higher or lower dimension. For example, if $A = \pi r^2$, the area of a circle with radius r, $L \propto r$ and so on for variety of regular forms. These types of relation appear widely in the natural sciences where they form an essential part of the study of relative growth or allometry (Gould, 1966). If the relationship between a line and an area is as postulated in equation (6.1), this is the condition of isometry. If the power of A were greater than $1/2$, this would be positive allometry, if less, this would be negative allometry.

Let us now define the area at a given scale k as A^k. If area is regarded as a measuring device for the perimeter, when the scale is increased to $k + 1$, it is clear that

$$\frac{L_{k+1}}{L_k} > \left(\frac{A_{k+1}}{A_k}\right)^{1/2}, \tag{6.2}$$

because more and more scaled detail concerning the boundary will be picked up. In fact, the equivalent to the coastline conundrum is that in the limit as $k \to \infty$, the ratio of areas in equation (6.2) will converge but the ratio of the perimeters will continue to increase. From equation (6.1), it is clear that to derive L from A, area must be rescaled by a parameter which is greater than 1 but less than 2. That is

$$L \propto (A^{1/2})^D = A^{D/2}, \tag{6.3}$$

where $1 < D < 2$. If $D = 2$, then perimeter would scale up as area which would imply that the area be defined as a space-filling curve, a physically impossible realization for the kinds of geographical systems that we will consider here. If $D = 1$, perimeter would not scale more than the basic unit of measurement, the line, which would imply that no scale effects were present as area increased. The coefficient D is, of course, the fractal dimension. In this context, it again serves as an empirical measure of how much the curve in question departs from a straight line, thus indicating how 'crinkled' or tortuous the boundary across the space is. The relation in equation (6.3) is known as the *area–perimeter relation* and the nature of its scaling clearly implies a way of estimating the value of D (Lovejoy, 1982).

In measuring the boundary of single objects, we have restricted our attention to a single geometrical relation, namely that between a scale defined by a unit r and a measured perimeter L. The general form of this relation was given earlier in equations (2.25) and (5.7) and indexing it now by its scale r, it is

$$L(r) = N(r)r = \alpha r^{(1-D)}, \tag{6.4}$$

where $N(r)$ is the number of chords at scale r which approximate the perimeter $L(r)$. We thus have two relationships for the perimeter $L(r)$, one in terms of area as in equation (6.3), and one in terms of scale as in equation (6.4). Combining these gives

$$L(r) \propto A^{D/2} \propto r^{(1-D)}. \tag{6.5}$$

It is tempting to try to equate these by considering how A relates to the scale r. However, it is not possible to do this in general for it is only meaningful in special cases where the geometry is known or assumed.

Besides the area–perimeter and perimeter–scale relations, there is a third which could be used to estimate the fractal dimension D. This is the number–area rule known as Korcak's law (Mandelbrot, 1983). It relates the number of or fraction of areas $\mathrm{Fr}(A)$ with an area greater than A, to the area itself as

$$\mathrm{Fr}(A) \propto A^{-D/2}. \tag{6.6}$$

Here the characteristic length is again taken as the square root of area A and used in a generalization of the number–scale relation in equation (2.24). We will not use equation (6.6) in any of our analyses for it requires a much larger number of objects, in this case land use parcels, than the level of resolution of the application we have chosen permits. Nevertheless, there may be circumstances amongst the kinds of geographical applications which we will describe where it might be useful (Kent and Wong, 1982).

Both the area–perimeter and perimeter–scale relations in equations (6.3) to (6.5) are intrinsically linear in their parameters D and can thus be estimated by regression techniques after suitable logarithmic transformation. However, the data for these estimations are quite different. For the area–perimeter relation, it would in theory be possible to measure the area and perimeter of an irregular object at different scales and perform the regression on these measurements: but the relation is more suited to estimation using a series of areas and perimeters associated with a set of

objects, all of different sizes. If the relationship holds over many scales, more scaling detail will be picked up in larger objects than in smaller ones. In the case of the perimeter–scale equation, we will use the techniques of Chapter 5 which involve measuring the same object at different scales. In this chapter, we will explore the consequences of these two methods and note some of the conceptual difficulties resulting from their comparison. In the rest of this section, however, we will concern ourselves solely with their estimation as well as techniques for measuring the effects of scale.

First we will write the area–perimeter relation in equation (6.3) as

$$L = \gamma A^{f(D)}, \tag{6.7}$$

where γ is a constant of proportionality and $f(D)$ is some power function involving the fractal dimension D, in this case $f(D) = D/2$. Taking logs of equation (6.7) gives

$$\log L = \log \gamma + f(D) \log A, \tag{6.8}$$

where in the case of equation (6.3), $\log \gamma$ is the intercept and $f(D) = D/2$, the slope of the regression line of the log of perimeter on the log of area. Clearly the slope $f(D)$ can take different functional forms from which D can always be derived, given an estimate of the slope. The perimeter–scale relation in equation (6.4) can also be generalized as

$$L(r) = \alpha r^{g(D)}, \tag{6.9}$$

where α is the constant of proportionality and $g(D)$ a function which in equation (6.4) is $(1 - D)$. Taking logs of equation (6.8) gives

$$\log L(r) = \log \alpha + g(D) \log r, \tag{6.10}$$

where $\log \alpha$ is the intercept and $g(D) = (1 - D)$, the slope of the associated regression line. Note that equations (6.9) and (6.10) are generalized versions of the perimeter–scale relations given in Chapter 5.

The conventional fractal model based on the use of equation (6.3) in (6.7) or equation (6.4) in (6.9), has a linear form implying that D is scale-invariant. However, as we have seen in Chapter 5, in some contexts it can be hypothesized that dimension itself might vary with scale or area and in this case, the linear form would be more complex. A second model based on the notion that fractal dimension does vary systematically with scale was used in Chapter 5 and these variants will also be tested here. In the case of the area–perimeter relation, the fractal dimension D can be hypothesized to be $f(D) = (\zeta + \eta A^{1/2})/2$ which when used in equation (6.8) gives

$$\log L = \log \gamma + \frac{\zeta}{2} \log A + \frac{\eta}{2} A^{1/2} \log A. \tag{6.11}$$

In equation (6.11), the coefficient $\zeta/2$ has an analogous role to $D/2$ in equation (6.3) as applied to (6.7). The third term on the right-hand side of equation (6.11) is a dispersion factor which measures the non-linearity of the area–perimeter relation. It is clear that as $\eta \rightarrow 0$, $\zeta \rightarrow D$ and equation (6.11) collapses back to the logarithmic transformation of equation (6.3) or strictly (6.7) with $f(D) = D/2$.

For the case of the perimeter–scale relation in equation (6.10), we have already seen in Chapter 5 that when $g(D)$ is a function of scale $\lambda + \phi r$, then

$$\log L(r) = \log \alpha + \gamma \log r + \phi r \log r, \tag{6.12}$$

where the fractal dimension D is given by

$$D = 1 - \lambda - \phi r, \tag{6.13}$$

In equation (6.13), as the scale $r \rightarrow 0$, $D \rightarrow 1 - \lambda$. The term $\phi r \log r$ in (6.12) also acts as a dispersion factor which increases the fractal dimension as the scale increases and is a kind of weighted entropy, modulating the effect of the fractal dimension. In the sequel, these 'transient dimension' models based on the systematic variation of dimension with area or scale will be referred to as the 'modified models' in contrast to the 'conventional models' of equations (6.3) and (6.4) where dimension is scale-invariant.

Estimating D for the two models using the area–perimeter relations in equation (6.8) with $f(D) = D/2$ and in equation (6.11) is straightforward. For each land parcel, the area and perimeter can be easily measured and form the dependent and independent variables respectively. The number of parcels in the study obviously affects the fit of the regression, and it may be necessary to identify and exclude outliers. However, the variation in scale within the observations forming the data set is only influenced by the prior selection of land parcels, not by any peculiarity of the area–perimeter measurement. In contrast, the perimeter–scale relations in equation (6.10) with $g(D) = 1 - D$ and in equation (6.12) depend upon the choice of scale and the measurement of the perimeter associated with that scale for each individual object. In some of the applications we present below, we will adapt this procedure to form aggregate perimeters from more than one land parcel. Furthermore, we will make these perimeter–scale measurements using each of the four measurement methods outlined previously in Sections 5.6 and 5.7 of Chapter 5.

6.3 Areas and Perimeters: the Fractal Geometry of Urban Land Use

Swindon, the town chosen for our expository analysis, is located in south central England about 70 miles west of London. The town is quite compact and not affected in its form by the presence of any rapidly growing nearby towns. It has a reasonably buoyant economy which in the 1960s was due to its designation as an expanded town, taking overspill population from Greater London. More recently, its favored location in an expanding area of southern England has led to the location of new service and high technology industries in and around the town itself. Figure 6.1 and Plate 6.1 shows the pattern of land use composing the town in 1981 from which it is clear that as the town has grown, it has absorbed villages in its immediate periphery. This fairly aggregated land use map was compiled using diverse data sources: remotely sensed data and local authority map records used by Rickaby (1987) as part of his studies into the energy requirements of small towns. The five land uses – residential, commercial–industrial, educational, transport and open space – shown in Figure 6.1 constitute the basic

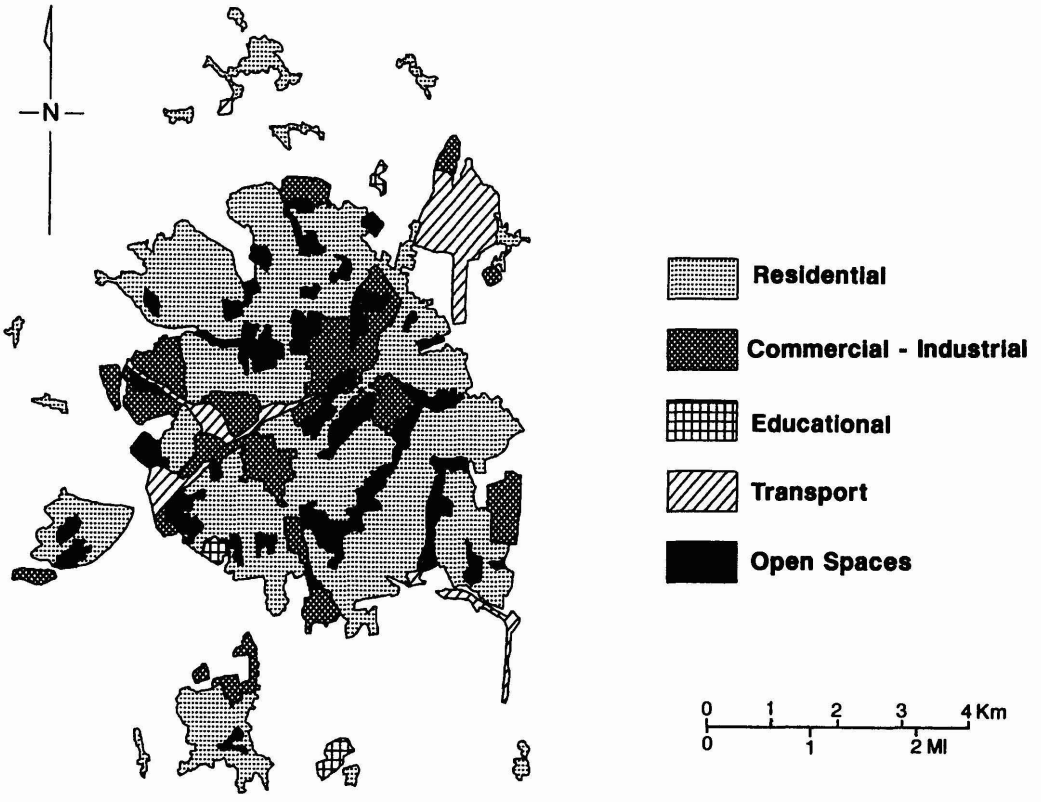

Figure 6.1. Urban land use in Swindon, 1981.

data for this study. The map was digitized using locally available software (Bracken, Holdstock and Martin, 1987) and the land parcels were extracted in polygon form using conventional digital cartographic techniques.

Figure 6.2 shows the polygons which represent the land parcels, drawn to scale and classified according to the five land uses but not arranged in any particular order. Observing how these parcels fit together to form the overall map, the conundrum raised in the introduction to this chapter relating to contiguous boundaries between different uses is immediately apparent. For example, the largest land parcel of all is part of the set of residential land uses shown in Figure 6.2. In one sense, this parcel can be considered as a skeleton for the entire town, but it is clear that about half its boundary is common with other land uses; this raises the conceptual difficulty of making comparisons of the irregularity and form of this boundary with that of adjacent land uses. For the moment, we will assume that the different parcels can be treated separately, and we will pursue the estimation in this manner before commenting further on the problem below.

Some characteristics of the digital representation of the five land uses are presented in Table 6.1. There is considerable variation in the set of land uses, and it is clear that no generalizations can be made about educational land use which comprises only three parcels; and there are limits to how far one can make inferences about the transport land use which comprises

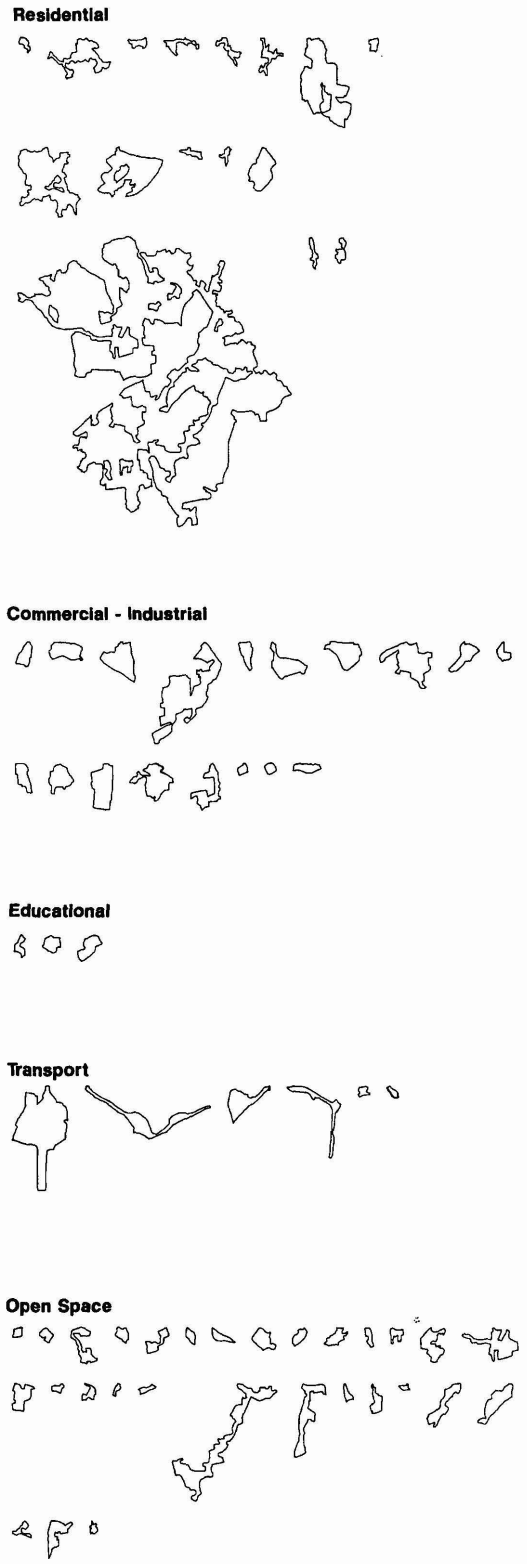

Figure 6.2. Land parcels separated into distinct land uses.

Table 6.1. Characteristics of digitized urban land use in Swindon

Land use	Number of parcels	Number of digitized points N	Number of common points	Percent of common points
Residential	16	2989	1534	51.4
Commercial–industrial	18	1030	626	60.8
Educational	3	109	17	15.6
Transport	6	510	261	51.2
Open space	29	1421	1286	90.5
All land uses	72	6059	3724	61.5

Land use	Mean no. of points per parcel	Average chord length $\bar{d}$	Length of perimeter $L(r) = N(r)r$	Feret's diameter F
Residential	186.8	0.785	2344.9	119.5
Commercial–industrial	57.2	0.837	861.8	42.4
Educational	36.3	0.727	78.5	11.5
Transport	85.0	0.821	417.7	52.8
Open space	49.0	0.776	1106.2	52.9
All land uses	84.2	0.795	4816.3	119.5

only six parcels. However, examining the average chord length of these data which ranges from 0.727 to 0.837 base level units, indicates that the base level digitization is fairly independent of land use type. The number of digitized points given in Table 6.1 for each land use and for the total involves the double counting of common boundaries referred to above in that the points which are common to any pair of land uses are included in each land use.

Of the 6059 points which comprise the total of points in each of the distinct land uses, there are only 2335 points which are not common to adjacent land use boundaries. The remaining 3724 points which are common to various pairs of land uses are in fact counted twice (for each land use in each pair) and thus there are 1862 points which are common in the data set. In total, there are 4197 distinct points in the set, 43% of these being common to adjacent land uses. In terms of the individual uses, 51% of the points defining the residential parcels are common to other uses, while over 90% of the points referring to open space are part of the boundaries of other land uses. These percentages, shown in Table 6.1, give some indication of the position of the land uses within the town. For example, most of the open space is enclosed within the town itself, not on its edge, while educational land use is mainly on the town's edge. In Table 6.1, the perimeter length refers to the sum of all the perimeters relating to a given land use

while the Feret diameter represents the maximum spanning distance found amongst the parcels of any given land use, as defined previously in equation (5.12). It is clear from this and from Figure 6.2 that there is some considerable variation among land use parcels with respect to size.

We are now in a position to estimate the first set of fractal dimensions based on the area–perimeter relation, In Figure 6.3, the log–log plots of perimeter against area are presented as scatter diagrams for each of the five sets of land uses in turn, and then for all five land uses comprising the 72 land parcels in the town. These plots demonstrate strong relationships

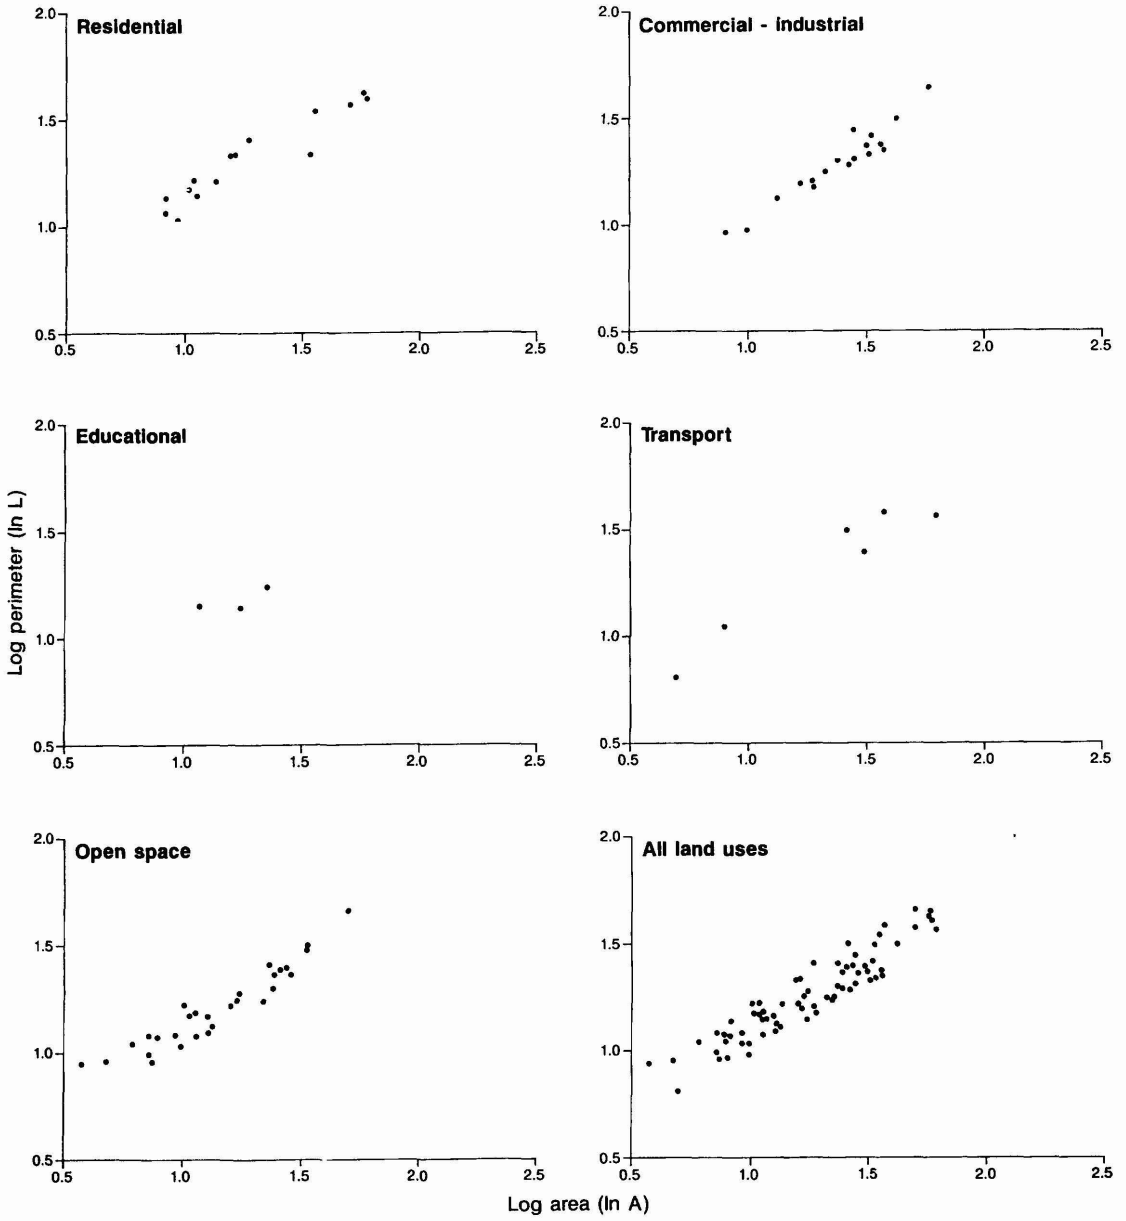

Figure 6.3. Scatter plots of the area–perimeter relations.

between perimeter and area, and it is difficult to detect any significant non-linearity in their form. To these data we have fitted the conventional perimeter–area model based on using equation (6.3) in (6.7) which we can write in the form of equation (6.8) as

$$\log L = \log \gamma_1 + \frac{D}{2} \log A. \tag{6.14}$$

The modified perimeter–area model given earlier in its logarithmic form in equation (6.11) we will repeat for convenience as

$$\log L = \log \gamma_2 + \frac{\zeta}{2} \log A + \frac{\eta}{2} A^{1/2} \log A. \tag{6.15}$$

but note that we now distinguish the intercept terms in equations (6.14) and (6.15) as γ_1 and γ_2 respectively.

The results of these regressions are presented in Table 6.2. With the exception of educational land use where there are only three observations, the adjusted r^2 statistics for both sets of models are acceptable. There are no obvious outliers, for example, whose removal might improve these statistics. The modified model gives a slight improvement over the conventional one, but this is not significant. The fractal dimensions in the conventional model are as postulated, that is, $1 < D < 2$, with the exception of the educational land use which we must exclude from serious analysis. Interpretations of the parameter ζ in the modified model are problematic because of the size of η. As $\eta \to 0$, it is hypothesized that $\zeta \to D$ but none of these results suggest any refined interpretation comparable to the 'transient dimension' perimeter–scale specification explored in Chapter 5. The conventional model is the only one acceptable here and excluding education, the analysis suggests that the commercial–industrial ($D \approx 1.478$) and transport land uses ($D \approx 1.447$) have more tortuous boundaries than those of residential ($D \approx 1.331$) and open space ($D \approx 1.243$). The dimension associated with all the land uses ($D \approx 1.296$) is clearly an average. All these results are consistent with other estimates using aerial data produced by applications of the area–perimeter method (Lovejoy, 1982; Woronow, 1981) but the correlations are not as good. Nevertheless this provides a backcloth and comparison to the perimeter–scale analyses which now follow.

Table 6.2. Parameters associated with the area–perimeter relation

Land use	Conventional model equation (6.14)		Modified model equation (6.15)		
	D	r^2	ζ	η	r^2
Residential	1.331	0.924	0.499	0.229	0.924
Commercial–industrial	1.478	0.923	0.307	0.361	0.926
Educational	0.569	0.111	−15.426	5.461	Not computed
Transport	1.447	0.913	3.996	−0.845	0.950
Open space	1.243	0.892	−0.710	0.693	0.925
All land uses	1.296	0.880	0.339	0.301	0.892

6.4 Perimeters and Scale: Constructing Long Threads from Land Use Parcel Boundaries

Before we introduce the analysis based on perimeter–scale relations, it is worth discussing the degree of irregularity associated with different land use patterns as we perceive it in *a priori* terms. In examining the five land uses, we might argue that open space is more likely to be defined according to the boundaries imposed by natural terrain in contrast to more artificially determined land uses such as the commercial–industrial and transport uses. Residential land use is likely to have a degree of irregularity in its form somewhere between these extremes as might educational use. With respect to the area–perimeter relations, this *a priori* ranking of open space/residential/educational/commercial–industrial/transport from higher to lower degrees of irregularity is not borne out at all by the fractal dimensions. Indeed, Table 6.2 implies somewhat the reverse but the r^2 coefficients are lower than anticipated, and it is possible that area–perimeter relations do not capture scale effects to the same precision as do methods based on perimeter–scale equations.

However, to be consistent with the area–perimeter analysis, it is necessary to devise a way of determining single fractal dimensions for each set of land parcels according to land use types. In a later section, we will look at the variation in fractal dimension across land parcels and types, but here we will begin by defining a global (or total) perimeter for each land use set. Were we to simply calculate a single total perimeter for each land use based on all its parcels, and regress these against scale, this would be similar to our previous analysis as scale would be a proxy for area. What we have done in fact is to calculate a total perimeter for each land use by stringing together the individual land parcel perimeters in the arbitrary order in which the parcels and their coordinate points have been digitized. We have also derived a total of total perimeters in the same way which contains all the points relevant to each land parcel.

In Figure 6.4, we show these total perimeters for each of the five land uses. These are not drawn to the common scales of the parcels contained in Figures 6.1 or 6.2, but are scaled up or down to be roughly comparable in area when displayed on a graphics device. It should be quite straightforward to identify the land parcels from their classification in Figure 6.2. The total perimeters are in fact derived by centering the first digitized point of each land parcel on a common point and producing a string of coordinates in the order in which each land use was digitized. The educational and transport land uses with the fewest land parcels show this most clearly in Figure 6.4. We have not included the total of total perimeters because it is not possible to produce a clear and clean plot due to the continual overlapping of boundaries: we will, however, use this total of totals in the subsequent analysis.

From these base level perimeters, aggregations across the given range of scales yield new perimeters which provide the data for estimating the parameters of the perimeter–scale relation. Two issues are important. First, the order and orientation of the land parcels forming the total perimeter

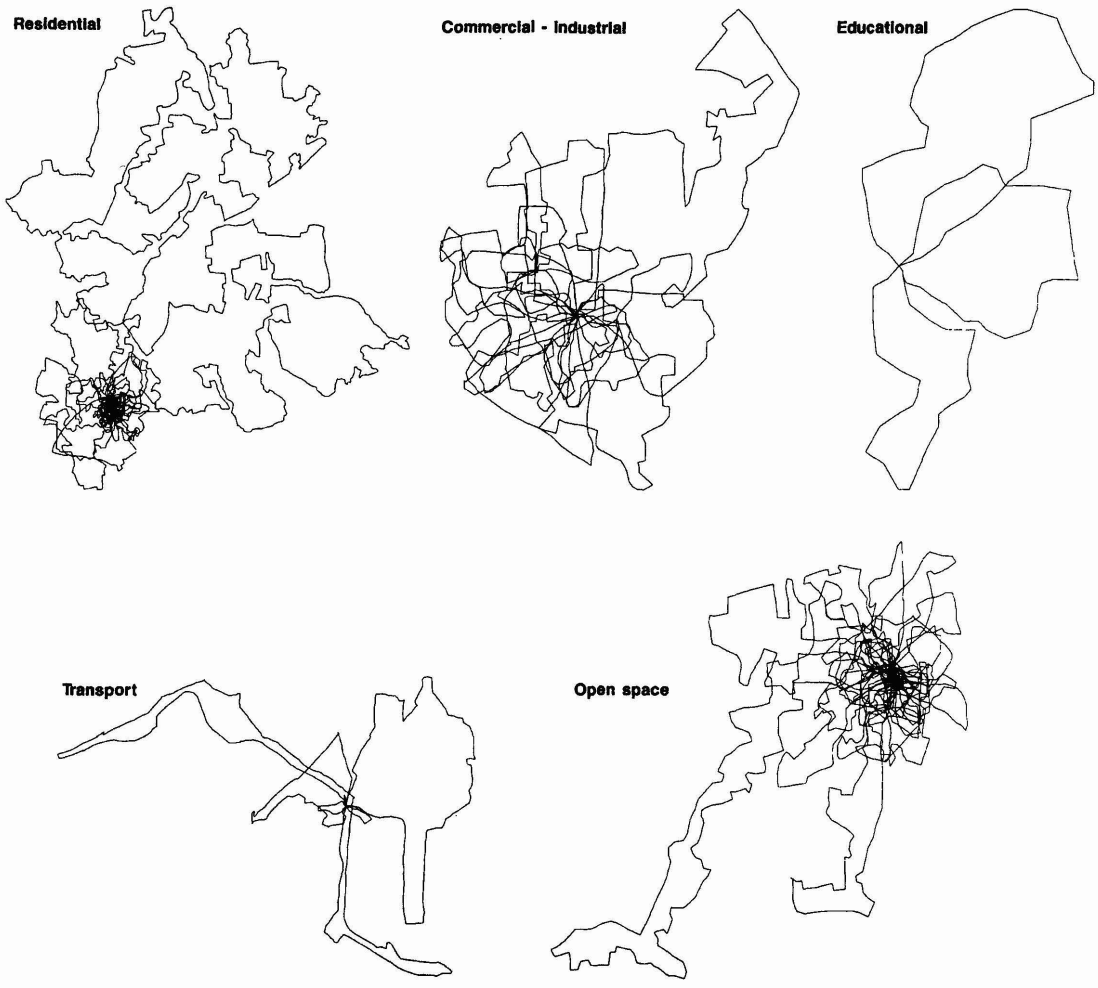

Figure 6.4. Aggregate perimeters for the five land uses.

could be crucial, and second the aggregations should not be as great as to pick up the aggregate shape of these composite perimeters which is clearly quite arbitrary. Order and orientation have been varied and this makes little difference to the subsequent results, but the aggregate shape problem does affect the estimated dimension. In fact, this leads to a reestimation of the perimeter–scale relations using a reduced set of aggregations to be reported in the next section.

As in Chapter 5, the number and scale of the aggregations for each of these perimeters (which provides the set of observations for the log–log regressions) is fixed so that each observation is of equal weight in the estimation. The limits of aggregation for the four measurement methods based on the structured walk, hybrid walk and cell-count are first calculated as follows. The average chord length $\bar{d}$ in these data sets is computed from equation (5.11) and the maximum spanning distance or Feret diameter F from equation (5.12). Note that we assume there are N points in the digitized base level curve, which is thus made up of $N - 1$ straightline segments

or chords. The sequence of aggregations where $\bar{d}$ represents the first chord size and F represents the last, and where m is the number of aggregations, is given by

$$F = \omega^{(m-1)}\,\bar{d}. \tag{6.16}$$

In fact, in this instance the starting point is set as the minimum, not average, chord size, and this can be represented as a fraction ψ of $\bar{d}$. Therefore, equation (6.16) becomes

$$\psi F = \omega^{(m-1)}\,\psi\bar{d}. \tag{6.17}$$

The weight ω scales one chord size to the next in the sequence of aggregations and this is computed from equations (6.16) or (6.17) as

$$\omega = \exp\left(\frac{\log F - \log \bar{d}}{m-1}\right). \tag{6.18}$$

This method of aggregating perimeters can only be applied to the structured walk, hybrid walk and cell-count methods of perimeter approximation, for the equipaced polygon method does not involve distances between points, only the order of points in the base level data set. A similar method of weighting is used, however, involving numbers of base level chords, not length based on distances. As the number of base level chords used to form a new chord increases, the actual length of the new chord increases and this is akin to aggregation to larger distance scales. Then if the number of original points needed to approximate the coarsest acceptable perimeter is N_{max} and the minimum number N_{min}, the sequence of chord sizes in the sequence of aggregations is given as

$$N_{max} = \omega^{(m-1)}\,N_{min}, \tag{6.19}$$

from which ω is determined in the same way as previously; that is, as

$$\omega = \exp\left(\frac{\log N_{max} - \log N_{min}}{m-1}\right). \tag{6.20}$$

In fact, N_{min} is always 2, and N_{max} is set as $N/6$, thus implying that the number of chords defining the most aggregate perimeter is 6; this would make the top level of aggregation consistent with the maximum spanning distance F.

In fact, the algorithm used to aggregate the original chords on each iteration into new perimeters employs ω in equation (6.20) only as a guide. Clearly the number of chords must be integral, not real, thus equation (6.19) involves truncation or addition to create integer numbers. The number of aggregated chords on each iteration $k+1$ is given as $N_{k+1} = \text{int}\,(\omega N_k)$. However, if N_{k+1} is equal to N_k, then N_{k+1} is increased by one chord length, that is $N_{k+1} = N_k + 1$. In the application of these algorithms, we have set m as 100 in each case. In fact, for the equipaced polygon method, although this also applies, the actual number of aggregations made is always less than 100 because of the discrete conditional nature of the aggregation.

The observations produced by applying each of these four methods to the five total perimeters and the total of totals are shown as Richardson

plots in Figures 6.5 to 6.8. Before the associated regressions are discussed there are several points to note. First there are quite clear upper scale effects caused by aggregation to too high a level. These are seen as departures from the trend of each graph and as obvious twists and turns in the tails of some of the plots. Second, these plots show strong evidence of nonlinearity suggesting, as in Section 5.5, that the modified model, where fractal dimen-

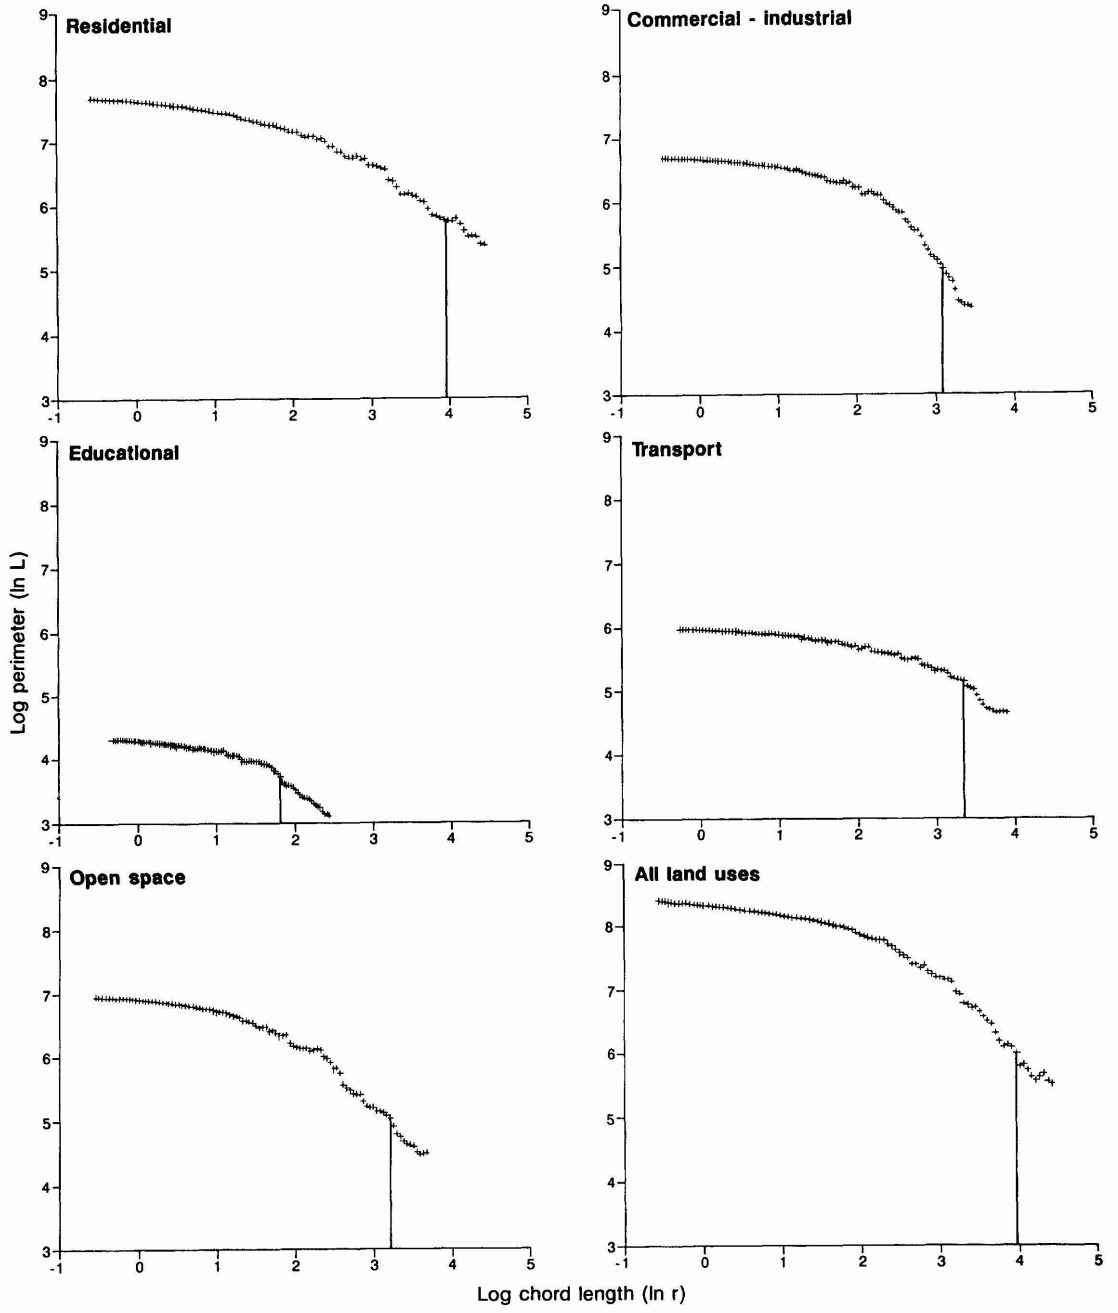

Figure 6.5. Richardson plots of perimeter–scale from the structured walk.

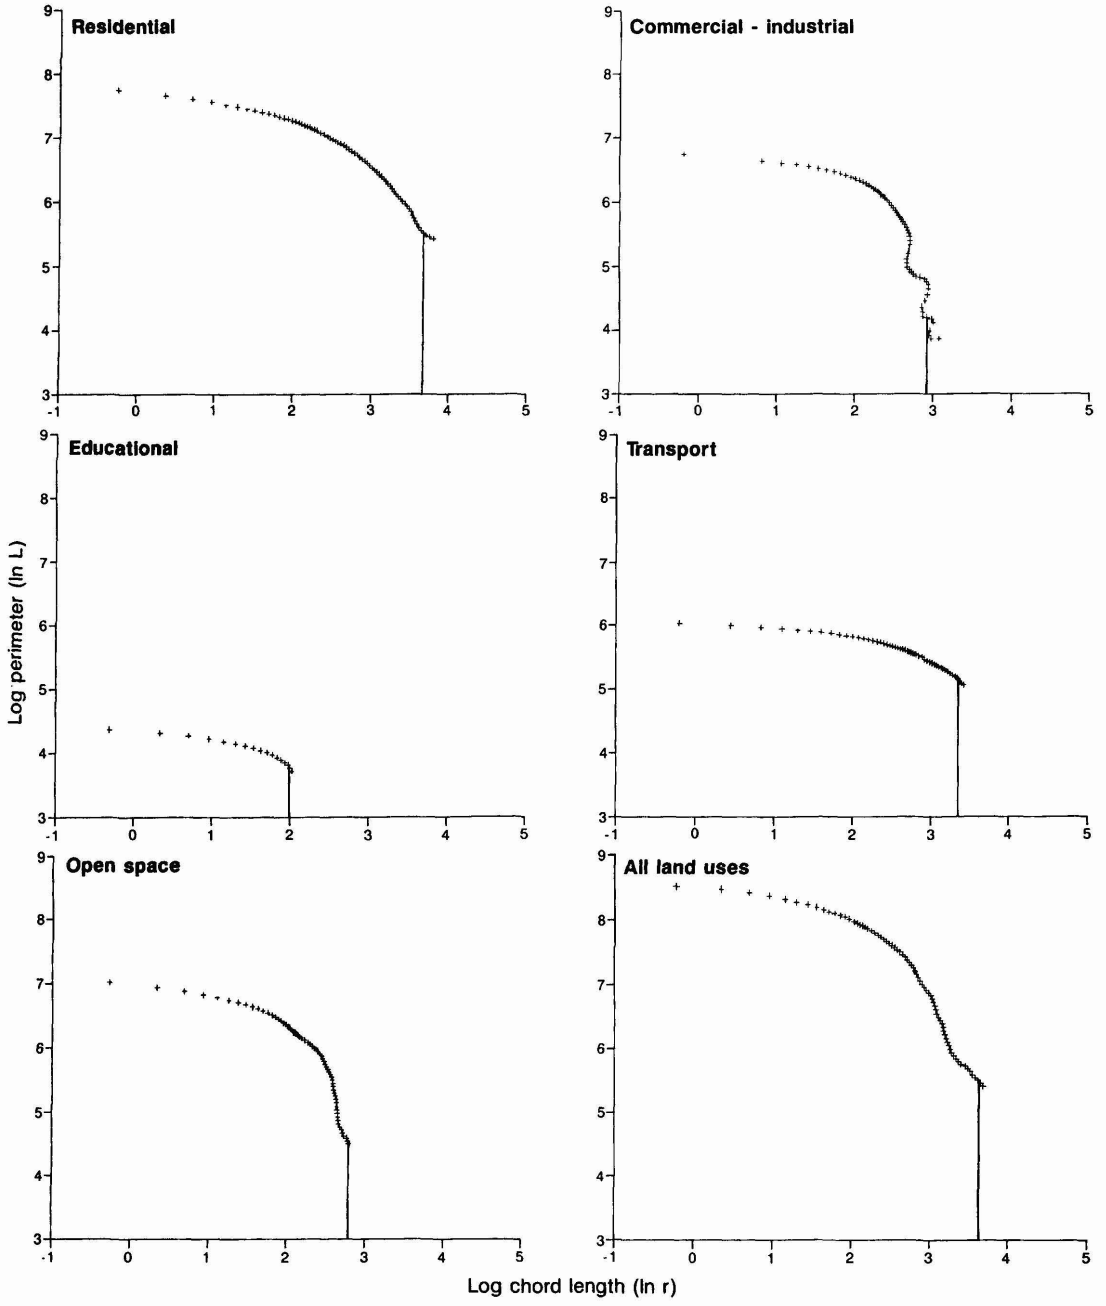

Figure 6.6. Richardson plots of perimeter–scale from the equipaced polygons.

sion varies with scale, is more applicable than the conventional model. Third, and in the context of our evaluation of these different algorithms in Sections 5.6 and 5.7, the equipaced polygon method gives cause for concern in that the algorithm attempting equal weighting does not perform well in establishing equal spacing of observations or in meeting the fixed number ($m = 100$) of aggregations. There are clear twists in the tails of the associated

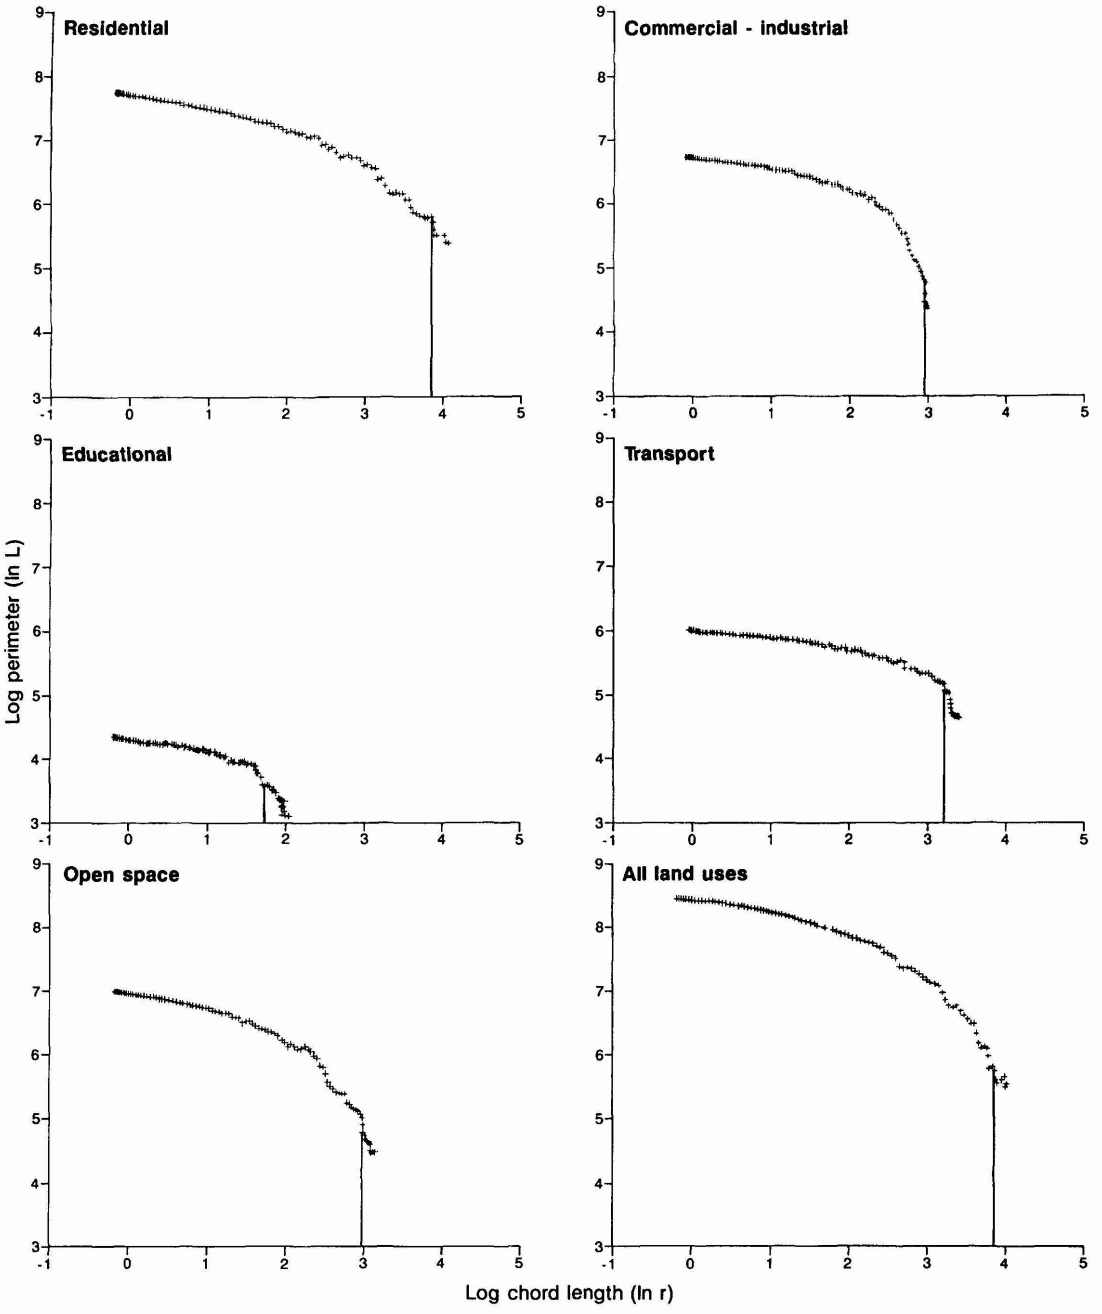

Figure 6.7. Richardson plots of perimeter–scale from the hybrid walk.

plots at the higher levels of aggregation. Finally, the aggregation in the case of educational land use to over 100 levels, is problematic in that there are only 109 digitized points in the total perimeter set as shown in Table 6.1.

We will present the fractal dimensions derived from the conventional and modified models for all the plots shown in Figures 6.5 to 6.8, notwithstanding the fact that the equipaced polygon method and educational land use are, in the sense just described, likely to yield unreliable results. In

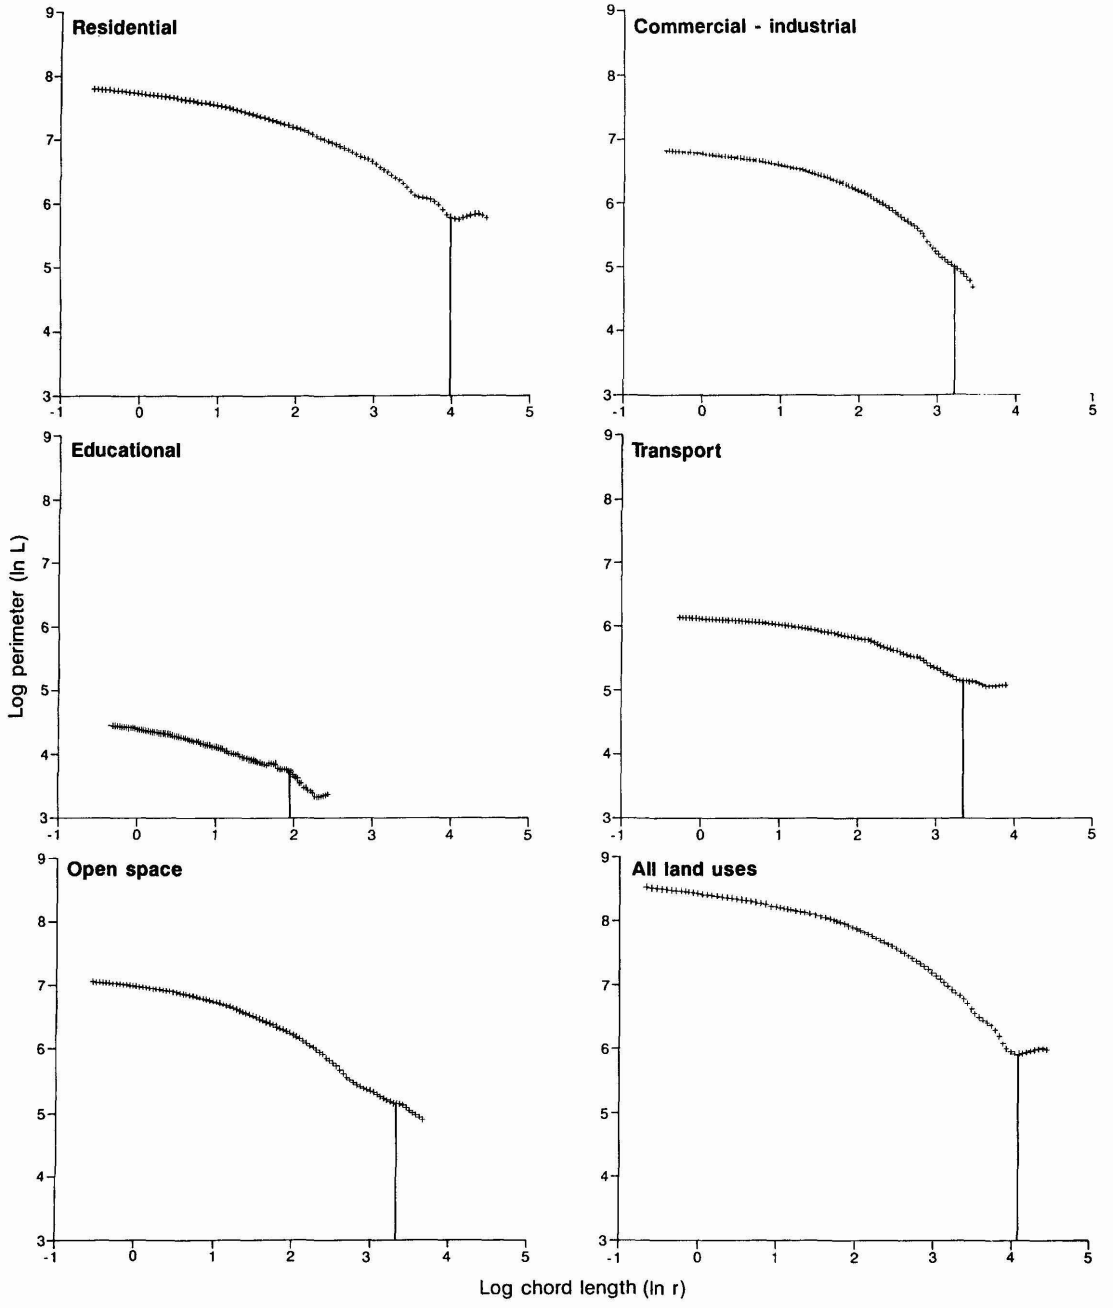

Figure 6.8. Richardson plots of perimeter–scale from the cell-count.

Table 6.3, we show the fractal dimension D, computed from the slope of the regression line $g(D) = 1 - D$ as in equation (6.4) applied to equations (6.9) and (6.10) which involves the conventional model; and we also show the performance of the model in terms of the r^2 statistics. In Table 6.4, we show the same for the modified model as given in equation (6.12). In this table, we first give the fractal dimension D derived from the coefficient λ

Table 6.3. Fractal dimensions based on the conventional model[1,2]

Method	Residential	Commercial–industrial	Educational	Transport	Open space	All land use
Structured walk	1.450 0.903	1.508 0.773	1.396 0.842	1.289 0.836	1.593 0.880	1.570 0.867
Equipaced polygon	1.694 0.853	2.066 0.619	1.274 0.827	1.300 0.811	1.957 0.687	2.021 0.795
Hybrid walk	1.496 0.904	1.566 0.757	1.442 0.799	1.314 0.803	1.666 0.860	1.619 0.862
Cell-count	1.447 0.939	1.499 0.872	1.402 0.946	1.294 0.920	1.543 0.937	1.571 0.909

[1] The structured walk, hybrid walk and cell-count methods are based on $m = 100$ aggregations for each land use. The equipaced polygon method has $m = 72, 68, 17, 53, 65$ and 75 for the five land uses and all land use applications, respectively. These m values also pertain to the modified model results in Table 6.4.
[2] Each cell shows the fractal dimension $D = 1 - g(D)$ with the r^2 statistic beneath.

Table 6.4. Fractal dimensions based on the modified model[1]

Method	Residential	Commercial–industrial	Educational	Transport	Open space	All land use
Structured walk	1.272 –0.004 0.981	1.076 –0.212 0.995	1.058 –0.041 0.991	1.096 –0.006 0.992	1.292 –0.012 0.980	1.291 –0.006 0.981
Equipaced polygon	1.187 –0.011 0.997	0.559 –0.061 0.941	1.006 –0.036 0.988	1.045 –0.007 0.999	0.716 –0.064 0.956	1.176 –0.020 0.974
Hybrid walk	1.249 –0.006 0.994	0.956 –0.036 0.977	0.925 –0.082 0.978	1.041 –0.011 0.974	1.157 –0.027 0.996	1.225 –0.010 0.996
Cell-count	1.333 –0.002 0.972	1.204 –0.014 0.993	1.223 –0.021 0.991	1.190 –0.003 0.968	1.375 –0.007 0.976	1.374 –0.004 0.967

[1] Each cell shows the fractal dimension $D = 1 - \lambda$, the dispersion coefficient ϕ, and the r^2 statistic beneath.

as $1 - \lambda$, and then we give the dispersion coefficient ϕ, noting of course that as $\phi \to 0$, $D \to 1 - \lambda$.

It is immediately clear from Tables 6.3 and 6.4 that the modified model in which dimension is a function of scale gives by far the best performance over all methods and land uses. Yet the equipaced polygon and hybrid walk methods produce strange results for the modified model in that fractal dimensions are less than 1 in four cases. In the case of the conventional model, these methods also appear to give D values higher than anticipated. With respect to the ranking of D values from Table 6.3, there is, however, a fairly consistent order over all methods in which open space, all land uses, and commercial–industrial have higher fractal dimensions than residential which in turn is higher than educational and transport.

A more disordered set of ranks is associated with the modified model although there are some similarities with the conventional model results, and in any case, the dispersion coefficients pick up the nonlinearity in the relations, hence influencing the value of D. In this respect, the dispersion coefficients are quite low for most land uses. To summarize then, if the equipaced polygon and hybrid methods which seem to pick up inappropriate larger scale effects, are ignored, the structured walk and cell-count methods produce a ranking of fractal dimensions across all land uses (with the exception of educational) which accord to our *a priori* expectations. At this stage, it is even possible to say that variations in dimension and coefficients between land uses are clearly wider than between methods, and this implies that the choice of method is less significant than the division into standard types of land use. However, the really important point at issue here is the presence of unwanted and arbitrary scale effects in the data. It is quite clear from Figures 6.5 to 6.8 that we must remove the highest aggregations from all these plots. In doing so, we also immediately remove some of the non-linearity from the data, thus hopefully improving the conventional model estimates as well as resolving some of the anomalous dimensions evident in Tables 6.3 and 6.4.

6.5 Refining the Perimeter–Scale Relations for the Aggregated Land Use Boundaries

The range of aggregations with respect to the structured walk, cell-count and hybrid walk methods given in equations (6.17) and (6.18) begins with the first and last chord lengths set as low as 70% of the average distance $\bar{d}$ and Feret diameter F for residential land use, to as high as 99% of $\bar{d}$ and F for the educational land use. As we have seen in Chapter 5, Shelberg, Moellering and Lam (1982) recommend that the starting points should be no lower than $\bar{d}/2$ while Kaye (1978) recommends the end point be no higher than $F/2$. The lower limits based on $\psi\bar{d}$ we have used do not pose any problem, but the upper limits based on the Feret diameter ψF yield approximations to the total perimeters with as few as two chords and only as many as five in number. As a general rule it is most unlikely that an approximation to the boundary of any irregular object can be made in less than six chords and in the case where we have up to 30 land parcels forming an aggregated perimeter, it could be argued that we should never go below 180 chords. Below this level we unwittingly include scale effects which pick up the arbitrariness of the 'constructed' perimeters; these are also sensitive to order and orientation of the land parcel strings. In these terms, it would appear that we should take an upper limit no greater than 20% of Feret's diameter, that is $\psi = 0.2$.

Examining the Richardson plots in Figures 6.5 to 6.8, it is quite straightforward to determine cut-off limits at their upper tails which would remove those observations clearly sensitive to these unwarranted scale effects. We have defined cut-off limits in these figures, showing the number of obser-

Table 6.5. Reestimation of the fractal dimensions for the conventional model[1,2]

Method	Residential	Commercial–industrial	Educational	Transport	Open space	All land use
Structured walk	1.403 0.892	1.389 0.786	1.229 0.910	1.210 0.897	1.499 0.869	1.486 0.864
Equipaced polygon	1.663 0.850	1.747 0.624	1.244 0.860	1.273 0.823	1.916 0.689	1.993 0.787
Hybrid walk	1.458 0.911	1.477 0.793	1.291 0.853	1.239 0.895	1.573 0.877	1.559 0.869
Cell-count	1.422 0.925	1.452 0.870	1.329 0.980	1.263 0.892	1.516 0.924	1.541 0.891

[1] Format of this table is as Table 6.3.
[2] The number of observations used for each regression is indicated in Figures 6.5 to 6.8.

vations each set has been reduced to. This varies for the case of the structured walk from between 9% and 22% of the original data set, and to as little as between 5% and 10% in the case of the equipaced polygon methods. From Figures 6.5 to 6.8, it is clear that we could impose even harsher constraints on the range of observations used, but although this would probably improve the results still further, relevant scale effects would probably be removed too.

Tables 6.5 and 6.6 show the reestimations of the two models using the four methods applied to each land use and the total of all land uses. There are marginal increases in the performance of the conventional model as comparisons between Tables 6.3 and 6.5 indicate. There is increased consistency between the methods with respect to the dimensions estimated with

Table 6.6. Reestimation of the fractal dimensions for the modified model[1,2]

Method	Residential	Commercial–industrial	Educational	Transport	Open space	All land uses
Structured walk	1.213 –0.006 0.990	1.050 –0.023 0.996	1.070 –0.035 0.986	1.094 –0.006 0.993	1.203 –0.019 0.986	1.203 –0.008 0.991
Equipaced polygon	1.162 –0.011 0.999	0.700 –0.051 0.921	1.030 –0.031 0.996	1.048 –0.007 0.999	0.074 –0.062 0.952	1.117 –0.022 0.980
Hybrid walk	1.237 –0.007 0.994	1.017 –0.031 0.988	1.024 –0.057 0.961	1.093 –0.007 0.994	1.176 –0.026 0.995	1.216 –0.010 0.997
Cell-count	1.261 –0.005 0.992	1.166 –0.017 0.997	1.249 –0.015 0.991	1.110 –0.007 0.996	1.303 –0.012 0.984	1.285 –0.007 0.991

[1] Format of this table is as Table 6.4.
[2] The number of observations used for each regression is indicated in Figures 6.5 to 6.8.

the exception of the equipaced polygon method. This is the most volatile of all the methods with the structured walk being the most consistent in terms of the original estimation and the reestimation. With respect to the ranking of land uses by dimension, an even clearer pattern emerges. Those with the higher fractal dimensions are open space and all land uses, followed by commercial–industrial and residential, with much lower dimensions for educational and transport uses. This bears out the *a priori* analysis even more strongly but it must be noted that the performance of the conventional model is barely adequate.

The modified model results shown in Table 6.6 are even better than those of Table 6.4. The ranking pattern is more variable than that of the conventional model with the commercial–industrial, educational and open space land uses having the highest degree of non-linearity as measured by the dispersion coefficient. The equipaced polygon method, somewhat ironically perhaps, has by no means the worst performance, but it still generates coefficients out of line with the other methods. As with the conventional model, the structured walk provides the most consistent results over each land use, and together with the cell-count method gives the best performance.

It is now worth summarizing all these results with respect to the fractal dimensions produced. In Figure 6.9, an attempt has been made to capture the variations in dimension produced across all methods and land uses in a single diagram. Each of the diagrams shows this variation with respect to the area–perimeter, conventional perimeter–scale and modified perimeter–scale methods, the latter two being shown with respect to their original estimation and reestimation. It is quite clear from these plots that the equipaced polygon method is the most problematic and should be excluded. Yet the structure of these results does show that there are greater differences between land uses than between methods, and this bears out the original hypothesis that such differences can be detected and possibly explained with respect to the processes governing the formation and evolution of different land use activities. We will say more about this in our conclusion but before we explore the variations between land parcels, we have averaged the dimensions produced in the last three sections, and these are shown, together with those of the subsequent section, in Table 6.7.

It is clear that the area–perimeter method produces quite different results from the perimeter–scale methods but that the patterns produced by these latter methods are more robust and consistent with our *a priori* theorizing. For the conventional model, the order of magnitude values of the fractal dimensions vary from $D \approx 1.5$ for open space and all land uses to $D \approx 1.4$ for residential and commercial–industrial to $D \approx 1.3$ for educational and transport. With respect to the modified model, $D \approx 1.2$ (in its limit) for open space, all land use, and residential, while for the other three land uses, $1.0 \leq D \leq 1.1$. This implies that these three – commercial–industrial, transport and educational land uses – present greater non-linearity; that is, their fractal dimensions vary more strongly with scale. These results mask the wide variation in dimension between land use parcels within any land use type, and do not in any way address the equality of fractal dimension over common boundaries between different land uses. In one sense of course, the purpose of this chapter is to ultimately focus on these questions and thus, we will address some of these in the next section.

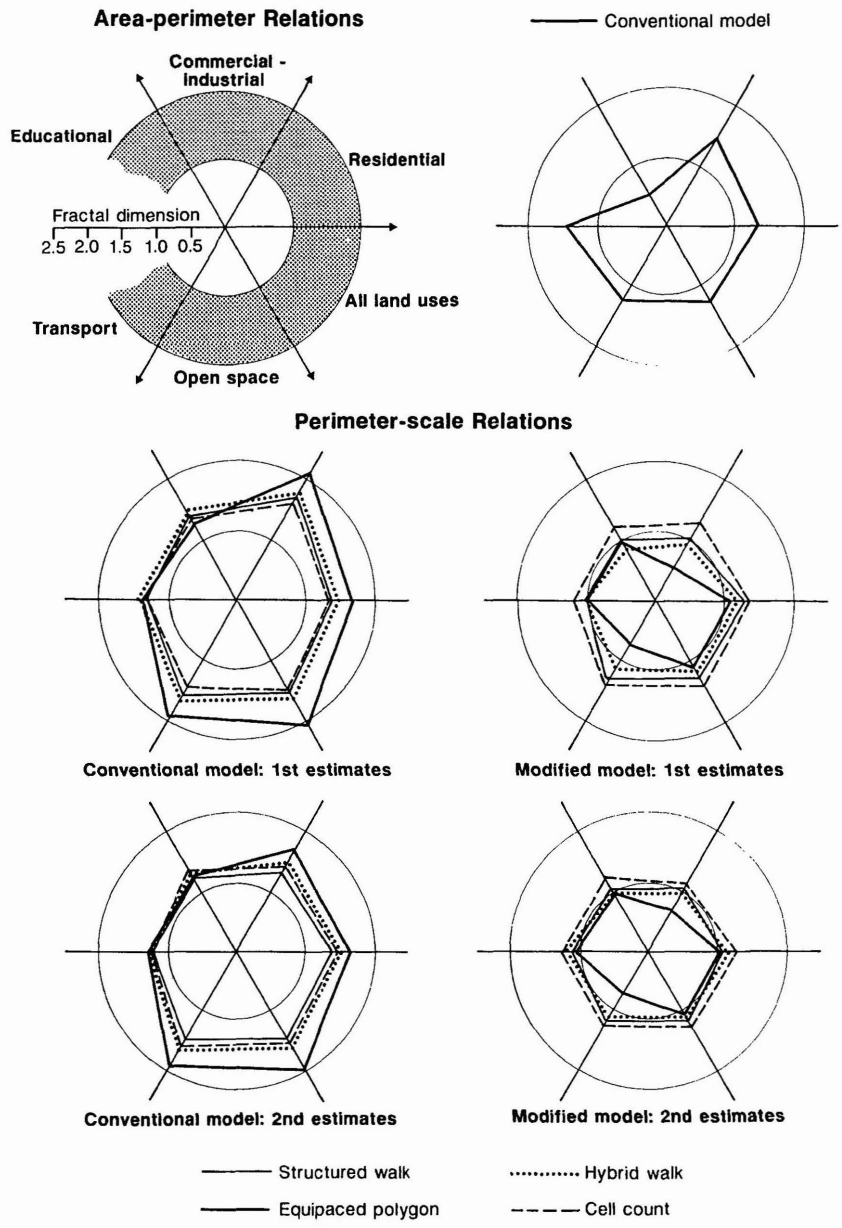

Figure 6.9. Variations in fractal dimensions across models, methods and land uses.

6.6 Fractal Dimensions of Individual Land Parcels

Figure 6.2 shows that there are 72 distinct land use parcels although in the previous analysis, the inner boundaries of some residential land parcels where such boundaries existed as 'holes' in the urban fabric, were added to the aggregate perimeters. There are eight such inner boundaries, all relating to residential land use as shown in Figure 6.2, and in the subsequent

Table 6.7. Fractal dimensions 'averaged' over methods of aggregation

Models/methods[1]	Residential	Commercial–industrial	Educational	Transport	Open space	All land uses
Conventional model: area–perimeter	1.33	1.47	0.56	1.45	1.24	1.29
Conventional model: PS 1st estimates	1.46	1.52	1.41	1.29	1.59	1.58
Conventional model: PS 2nd estimates	1.42	1.43	1.28	1.23	1.51	1.52
Modified model: PS 1st estimates	1.28	1.07	1.06	1.10	1.27	1.29
Modified model: PS 2nd estimates	1.23	1.07	1.10	1.10	1.22	1.23
Conventional model: average land parcels	1.15	1.10	1.09	1.11	1.13	1.13
Modified model: average land parcels	1.08	1.05	1.05	1.06	1.08	1.07

[1] PS: Perimeter–Scale.

analysis, they are treated as separate land parcels, thus augmenting the number of parcels treated to 80. First, all four aggregation methods – the structured and hybrid walks, the equipaced polygon and the cell-count – were applied to each of the 80 parcels, with the number of aggregations structured in geometric form as implied by equations (6.16)–(6.20), but with ω fixed and m varying accordingly.

In the case of the equipaced polygon method, the aggregation of 16 perimeters out of the 80 possible yielded too few observations for any subsequent regression. The other methods produced Richardson plots that were generally more linear than those shown in Figures 6.5 to 6.8, and therefore it was decided to fit the conventional model to all sets of observations generated. The r^2 values ranged from 0.833 to 0.999 in the case of applying the equipaced polygon method, but it was the structured walk that produced the most consistent plots in contrast to the hybrid and cell-count methods which were more volatile across the land parcels. Some methods produced dimensions for individual land parcels outside the range $1 < D < 2$. It was therefore decided to pursue more detailed analysis and model fitting using a narrower range of observations taken from the structured walk method only. In fact, the emphasis in this section is on the variation between land parcels, not on the variation between methods, hence our choice of the most robust method to generate the perimeter–scale data.

The application of both the conventional and modified models is shown in Figure 6.10 with respect to their fractal dimensions and associated r^2 statistics. The results for each land parcel are shown in the arbitrary order of Figure 6.2 according to the way the parcels were digitized but ordered within land use types as given previously. All fractal dimensions for the

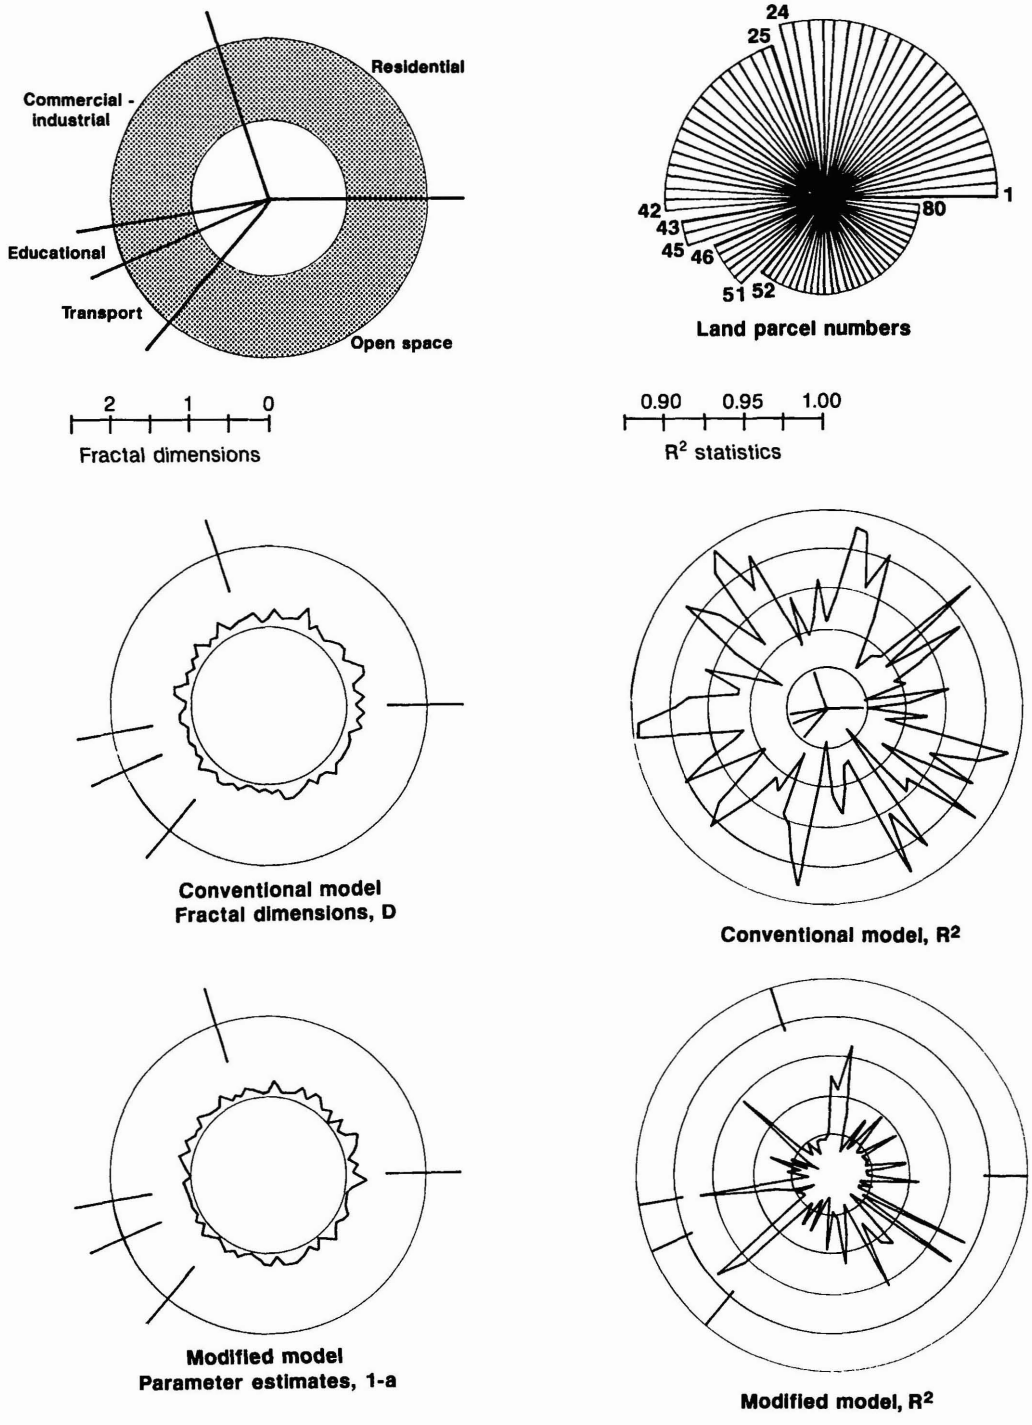

Figure 6.10. Fractal dimensions of individual land use parcels.

conventional model are in the postulated range from $1 < D < 2$. For the largest land parcels, the dimensions appear higher than for the rest, but a regression of the number of digitized points of all parcels on their conventional model dimensions yields an r^2 value of only 0.156. With respect to the five sets of land uses, these r^2 values were more variable, rising to 0.686 in the case of commercial–industrial. But in general, there does not appear to be a strong bias to higher dimensions for those land use parcels with the highest number of perimeter coordinates. Figure 6.10 also shows that the parameters of the modified non-linear model are consistent with the range of $1 < D < 2$, and the level of dispersion reflecting each parcel's non-linearity over scale, is fairly modest in every case.

One way of summarizing these parameters and statistics is by computing means and standard deviations. Table 6.8 presents these results for both models. The parameters and dimensions in this table are in their original form as predicted from the use of equations (6.10) and (6.12), that is, where the coefficient of the conventional model is $g(D) = 1 - D$, and those for the modified model λ and ϕ where $\lambda \rightarrow (1 - D)$ as the dispersion parameter $\phi \rightarrow 0$. Table 6.8 presents the variation in size of the land parcels for each land use and over the whole set in terms of their mean number of coordinates. The distribution of these coordinates with respect to the number of land parcels is skewed with a much greater proportion of parcels below their mean size. In the case of the residential parcels, this distribution is highly skewed, largely because of the existence of the one large parcel which provides the skeletal structure of the town.

The variation in parameter values and performance of the models, however, is much less than the variation in the features of the land parcels themselves. Figure 6.10 makes this apparent, while these results are averaged for each land use over all parcels in Table 6.8. With respect to the conventional model, the r^2 values only range from 0.934 in the case of residential parcels down to 0.919 for transport and the range for each land use over the land parcels is also quite narrow. The fractal dimensions D also show a pattern over the land uses which is consistent with the aggregated perimeter–scale results but is considerably clearer. The ranking of land uses from largest to smallest D is ordered from residential ($D \approx 1.152$), open space (1.132), transport (1.113), commercial–industrial (1.105), and educational (1.091) with an average over all land uses of 1.129. These values are considerably smaller than those shown previously, yet they are more in line with the examples developed in Chapter 5. In fact, the largest residential parcel (see Figure 6.2) is just one of five parcels which has a dimension greater than 1.2. From these results it is clear that the much higher dimensions produced by the aggregated perimeter–scale relations are due to the method of aggregating individual perimeters into strings of coordinates. It would appear that the aggregation picks up arbitrary scale effects which are central to the method itself and not the order or orientation of the individual parcels in the process of forming these composite perimeters.

The modified model results also shown in Table 6.8 have a wider range of variation around their mean estimates than those of the conventional model. In terms of the parameter λ, the residential, open space and all land use parcels have a dimension higher than those of transport, commercial–industrial and educational, in that order, although these values are over a

Table 6.8. 'Average' dimensions and statistics for the individual land parcels[1]

Land use	Zones	No. of coords	Mean coords	σ_c	$1 - D_1$
All land parcels	80	6059	75.7	177.4	–0.129
Residential	24	2989	124.5	317.2	–0.152
Commercial–industrial	18	1030	57.2	33.3	–0.105
Educational	3	109	36.3	6.65	–0.091
Transport	6	510	85.0	65.8	–0.113
Open space	29	1421	49.0	38.9	–0.132

Land use	σ_{D1}	r^2	σ_{r1}	$\lambda = 1 - D_2$	σ_{D2}
All land parcels	0.057	0.927	3.3	–0.071	0.051
Residential	0.063	0.934	3.3	–0.079	0.063
Commercial–Industrial	0.061	0.920	3.1	–0.050	0.042
Educational	0.035	0.921	4.3	–0.047	0.005
Transport	0.040	0.920	2.2	–0.065	0.028
Open space	0.046	0.928	3.6	–0.077	0.048

Land use	ϕ	σ_+	r^2	σ_{r2}
All land parcels	–0.024	0.021	0.969	2.3
All land parcels	–0.029	0.022	0.972	1.8
Residental	–0.018	0.011	0.979	1.6
Commercial–industrial				
Educational	–0.022	0.018	0.952	3.1
Transport	–0.006	0.001	0.960	3.9
Open space	–0.030	0.025	0.966	2.3

[1] The standard deviations are defined as: σ_c of the coordinates, σ_{D1} of the slope parameter $(1 - D_1)$ in the conventional model, σ_{D2} of the parameter λ in the modified model, σ_ϕ of the parameter ϕ in the modified model and σ_r (σ_{r1} and σ_{r2}) of the r^2 fits of the appropriate model to the land parcel data.

narrower range. The values of the dispersion factors also confirm this order and the r^2 estimates, although slightly better than those of the conventional model, are not as high as those produced by the aggregate perimeter–scale relations. Nevertheless, the non-linear model is an improvement over the linear and in general, these results for the individual parcels are better than anticipated.

6.7 The Problem of Measurement

The analysis presented here is suggestive rather than definitive and it reveals some basic problems of observation and measurement which are generic to all empirical science. In the development of fractal geometry, these problems have only just been broached and they will hold the center stage for a long time yet. In terms of developing a morphology of urban land use based on fractal geometry, it would appear that residential and open space land uses have a greater degree of irregularity than commercial–industrial, educational and transport. There is a logic here which we spelt out before we began the analysis in that for land uses which are larger in scale, there is likely to be less effort put into the geometric control of land under development. Yet there remains considerable uncertainty over the processes in operation. We have, however, shown that in general, scale effects vary with scale itself, and this is likely to be the result of multiple processes changing their relative importance through the range of scales. This argument is consistent with our treatment in Chapter 5 of cartographic lines as bounding geographical phenomena which are ostensibly isolated.

We have added to Table 6.7 the results of the last section where the whole range of models and methods applied throughout this chapter are displayed in suitably 'averaged' form. This shows up the arbitrariness of the analysis, with dimensions varying from as large as 1.6 to as low as 1.1. In previous work, the methods themselves have been subject to considerable variation but here despite some association of dimension values with land uses, the main variation concerns the way area, perimeter and scale are defined and measured, and the emphasis on area–perimeter or perimeter–scale relations. Questions of scale are never very clear in much fractal analysis, despite the fact that fractals are defined by scale-invariance. The area–perimeter method assumes that objects of varying sizes show the effects of varying scale itself (Woronow, 1981).

In short, a small residential development will not pick up the aggregate scale effects which can be detected by a large scale development, so runs the logic. However, this will depend on the base level of resolution in the first place, but there has seldom been much discussion of this in the field to date. The method of aggregating perimeters used in the composite perimeter–scale analyses of land uses is also suspect, because of arbitrary scale effects which can be produced, despite careful control over the process of aggregation. Lastly, the individual land parcel analysis using conventional perimeter–scale, not aggregated relations, suffers from its very inability to aggregate parcels, other than by arbitrary statistics such as simple averages.

What is clearly required in future work is a close examination of these approaches in terms of scale effects. We have paid great attention to problems of defining scale limits and ranges here but on reflection, our analysis should probably have employed a much narrower range of scales. However, from the Richardson plots in Figures 6.5 to 6.8, the reader can get some sense of how the dimensions might change if narrower ranges were to be used.

New methods are urgently required which are more robust than those used here, and we have now convinced ourselves that in perimeter–scale analysis, the hybrid walk and equipaced polygon methods should be abandoned in favor of methods such as the structured walk and cell-count whose properties of aggregation are better understood. But the final conclusion to these last two chapters relates to more substantive questions. Although we have tackled both individual and aggregate analysis here, much finer analysis of the fractal dimension of parts of perimeter boundaries is required. Further classification of the fractal shapes of land parcels will not emerge until the common boundary problem is directly broached. This must involve a detailed examination of how such boundaries are formed and how they evolve over time. By explaining the development process, more satisfactory explanations can be given of the way land uses 'stick' to each other to form the whole town. Only by extending the analysis along these lines can conclusive results about the ways in which urban morphologies are structured and evolve, be demonstrated.

Although we will not concern ourselves any further with conceptual problems of physical definition and practical problems of measurement, at least in terms of urban boundaries and edges, we will in fact begin to examine the ways in which entire morphologies of towns evolve, but at a more aggregate level. Our focus will move away from cities composed of edges and boundaries to cities composed of activities, mainly development in general and population in particular, which fill space. In the next two chapters, we will also retreat back to examining single fractal objects as complete cities, but this time with respect to the way they evolve and grow. We will, however, continue to examine the way the land parcels which compose the fabric of the city 'stick' to one another, and once again, we will trace the way the smaller threads of urban development can be woven into complete mosaics whose form is similar across many scales.