

CHAPTER FIVE

Geometry of species distributions: random clustering and scale invariance

ARNOŠT L. ŠIZLING

Charles University, Prague

DAVID STORCH

Charles University, Prague, and The Santa Fe Institute

Introduction

Spatial biodiversity patterns are tightly related to the patterns of spatial distribution of individual species. It has been recognized that the spatial distribution of individuals is never random nor homogeneous within some well-defined clusters but is aggregated on many spatial scales: individuals form clusters which themselves are aggregated into larger clusters and so on. The most useful way to capture these patterns is with fractal geometry, which treats such patterns as self-similar sets (Kunin, 1998; Halley *et al.*, 2004). Indeed, it has been shown that species spatial distribution is often close to fractal (Virkkala, 1993; Condit *et al.*, 2000; Ulrich & Buszko, 2003) and that the assumption of fractality of species spatial distribution is appropriate for deriving multispecies macroecological patterns, namely the species–area relationship (Harte, Kinzig & Green, 1999; Šizling & Storch, 2004). By contrast, species sometimes reveal distributions that deviate from strict fractality (Hartley *et al.*, 2004; He & Condit, this volume; Lennon *et al.*, this volume). More importantly, although there are several ways in which fractal distributions could emerge (Halley *et al.*, 2004), there is no strong biological reason why species spatial distribution should be exactly fractal, i.e. it is unclear which biological processes should produce fractal distribution.

Here we show that species spatial distributions which are very close to fractal can emerge from random processes leading to aggregation on several spatial scales. These processes have relatively straightforward biological interpretation and the spatial patterns they produce are in many parameters effectively undistinguishable from classical fractals. Moreover, when we compose together many spatial distributions resulting from these simple processes, we obtain relatively realistic species–area relationships, as well as a frequency distribution of occupied areas which is close to the observed distribution of species abundances. We thus propose a null model of the geometry of species distribution which is biologically reliable, realistic, and more general than is the fractal distribution.

Self-similarity and hierarchical aggregation

There is a wide range of possible types of spatially aggregated distributions, and for our purpose it is useful to distinguish self-similarity from fractality. Self-similarity is often taken as synonymous to fractality, but as we will show, fractal distribution can in fact be treated as a rather specific type of self-similar distribution. A geometric object is self-similar, according to Hastings and Sugihara (1993), “if it can be written as a union of rescaled copies of itself”. The problem is the exact nature of the rescaling. Imagine that we have a broad-scale spatial pattern which is patchy on this broad scale of resolution (in ecology these patches can represent, for example, patches of suitable macrohabitat). A self-similar structure emerges by replacing these patches with patterns *similar* to the original patchy pattern. The resulting structures can be very different, depending on the exact meaning of *similarity*, as we will show. Anyway, it is useful to see the emergence of self-similarity as a hierarchical, top-down process, where we obtain the whole structure by a downscaling performed in discrete steps. If we apply an unlimited number of these steps, we would obtain a real, mathematically self-similar set – but in the real world we always have a limited number of these replacements. In the following text we will be dealing only with self-similar structures emerging from a finite number of steps, and will call these steps the *levels of aggregation*. Self-similarity is, in this view, a property that links two subsequent levels of aggregation to one another.

This hierarchical approach to self-similar structures has clear advantages. Most importantly, structures which originate by this process are interpretable in terms of biological processes acting within each level of aggregation. We can imagine, for instance, that the largest patches (i.e. within the zero level of aggregation) correspond to broadly defined macrohabitats characterized by particular climate or elevation. The next level of aggregation can represent habitat patches characterized by particular vegetation, and subsequent levels can be formed for instance by patches of resources. The observed spatial distribution of a species then can form the next level of aggregation, as not all available habitat patches can be occupied, and patterns of occupancy within available habitat patches are given by spatial population (and/or metapopulation) dynamics. The pattern aggregated on many scales of resolution can therefore reflect the nested nature of habitat hierarchy, and also the hierarchical nature of spatial population processes. This hierarchical nature alone obviously does not ensure any regularity concerning the similarity of patterns observed at various levels of aggregation (i.e. self-similarity) or even fractality, but we will show that simple statistical assumptions concerning the nature of these hierarchical processes can lead to surprising regularities.

Note that this hierarchical approach to some extent reconciles two very different approaches to the multiscale nature of ecological systems. One of

these approaches, the hierarchical concept of ecosystems (Allen & Starr, 1982) takes seriously the fact that ecological processes differ in different spatial and temporal scales, and treats ecological systems as hierarchical, where each scale has its own rules of behavior. By contrast, the scaling approach consists of looking for patterns which are actually independent of scale, i.e. for scale invariance. We will show that some types of scale invariance can emerge even if we assume different processes acting on each scale – simply because these processes, albeit different, can have something in common, even if it is something whose nature is purely statistical.

Fractals

Fractals are self-similar structures where the “similarity” is defined in a particular way; classically through particular geometrical projections leading to various structures called “self-similar sets (*sensu stricto*)”, “self-affine sets” and “random fractals”. Similarity is there defined as the proportional lessening of the original pattern in the case of the self-similarity *sensu stricto*, the result of an affine projection of the original pattern in the case of self-affine sets, and the pattern “statistically indistinguishable” from the original pattern in the case of random fractals (Falconer, 1990). Behind all of these definitions is, however, the idea of scale invariance. This idea has been originally related to the problem of the measurement of the length of shoreline (Richardson, 1961) which apparently depends on the resolution of the scale used for this measurement. In planar terminology, the measured size of any real area depends on the resolution of the scale used, because finer resolutions necessarily lead to respecting more and more details. We are therefore looking for a measure of area which is independent of the scale used – i.e. is *scale invariant*. Scale invariance *sensu stricto* is thus a property of the way in which we measure geometrical objects rather than a necessary property of the objects themselves.

When attempting to fulfill the requirement of scale invariance, it is necessary to obey the formal condition that the measured area of the whole must be equal the sum of measured areas of particular patches (see Fig. 5.1). To meet this condition, we have to treat the term *area* more generally than it is in Euclidean space. This can be done as follows.

Let the whole be an i th square patch at any level of aggregation. The Euclidean area of the whole equals L_i^2 where L_i is the length of the edge of this patch; the Euclidean area calculated using the finer information on the structure of the patches at the subsequent level of aggregation equals $\sum_{j=1}^{n_i} l_{i,j}^2$ where n_i is the number of subpatches within the i th patch and $l_{i,j}$ are the lengths of their edges. As we can see, the equality of the areas calculated using coarse and fine scale is




Figure 5.1 Three levels of aggregation of the most common random fractal. While the Euclidean area at the coarsest scale equals L^2 , it equals $l_1^2 + l_2^2 + l_3^2$ for the following level and $\sum l_{i,j}^2$ through all possible combinations of i,j for the bottom level of aggregation.

met only if patches at the subsequent level of aggregation completely cover the whole. This can occur only if the *similar copy of the original pattern* is identical to the whole, and the topological dimension is 2. However, if the similar copy of the original pattern is patchy, and we are still going to meet the requirement of the scale invariance of area, we have to change the dimension of the set. In other words, we have to change the definition of area. The formal condition of self-similarity is then

$$L_i^D = \sum_{j=1}^{n_i} l_{i,j}^D. \tag{5.1}$$

To be accurate - if there is such $D \geq 0$ that

$$\sum_{j=1}^{n_i} \lambda_{i,j}^D = 1 \text{ where } \lambda_{i,j} = l_{i,j}/L_i \tag{5.2}$$

for all i and all levels of aggregation, then we call the set *fractal* and D its *fractal/Hausdorff dimension* (the definition adapted from Falconer, 1990).

It can be seen that a pair “a patch and its replacement” has always $D \geq 0$, and the Eq. (5.2) is thus always met. However, this D could potentially vary from patch to patch within a particular level of aggregation, as well as between subsequent levels. In the case of fractals, however, we have to keep this dimension stable through all the set to meet the condition of scale invariance. This can be realized in several ways. The simplest is the case of proper schoolbook fractals - the $\lambda_{i,j}$ are the same for all i,j , and they are also in the same number for all i . For example, in Fig. 5.2 λ equals to $1/3$ and n is 5. Thus, $\sum_{j=1}^{n_i} \lambda_{i,j}^D = 5(\frac{1}{3})^D = 1$ and the D is necessarily $\frac{\ln 5}{\ln 3} \approx 1.46$.

For self-similar sets as they are usually defined (i.e. where “similar” refers to the proportional lessening of the original pattern), both the number of

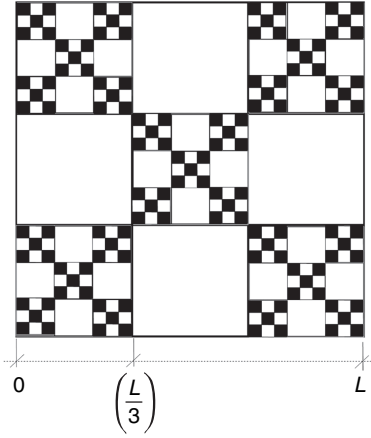


Figure 5.2 Classical regular fractal. In this case, the fractal dimension, D , remains the same for all subsequent levels of aggregation and for all chosen patches. D for the first two subsequent levels of aggregation is the solution of equation $L^D = 5(\frac{L}{3})^D$, for the following pair of levels it is the solution of $5(\frac{L}{3})^D = 25(\frac{L}{9})^D$, etc. The fractal area of this set is thus L^D .

subpatches, n_i , and their proportional sizes, $\lambda_{i,j}$, are constant for all replacements (i.e. for all i). Consequently, D is kept through the whole set and the set is fractal (it is scale invariant). This is also valid for those random fractals that keep the number and sizes of subpatches and where the randomness consists just in the random shuffling of particular patches. For more irregular structures, e.g. the patterns that vary in the number of subpatches for a particular replacement, or in the vector of $\lambda_{i,j}$, the fractality, however, represents very strong and very restrictive condition. Namely, the condition of scale invariance (Eq. 5.2) implies that for each subpatch k of a patch i , the equality

$$\lambda_{i,k} = D \sqrt[1 - \sum_{\substack{j=1 \\ j \neq k}}^{n_i} \lambda_{i,j}^D]{\quad} \quad (5.3)$$

must be met. This means that the area of at least one of the subpatches must depend on the areas of other subpatches and that all subpatches must fall into a particular range of possible sizes. This condition is derived purely from the aforementioned requirement of the independency of the area measurement from the scale used (Eq. 5.1), and thus it is entirely a matter of how we measure areas. There is no reason why this should be met in nature, because there is no reason (in a metaphorical sense) why nature should take care of our problems of measurement. A more realistic view seems to be that there is some other way in which the number and sizes of the subpatches are related. For instance, relative Euclidean area of subpatches, instead of fractal area, could be kept constant - $\sum_{j=1}^{n_i} \lambda_{i,j}^2 = \text{const}$. In the next step, we will try to generalize the idea of fractals with the aim of releasing it from the mentioned restrictive conditions.

One way to generalize the idea of fractals would be to adopt the classical (albeit a bit extreme) interpretation of random fractals which says that

the fractal is almost any spatial structure (Falconer, 1990). According to this definition, a fractal is each set with fractal dimension, D , defined as the root of

$$E\left(\sum_{j=1}^{n_i} \lambda_{i,j}^D\right) = 1, \quad (5.4)$$

where $E(x)$ is the expectation of x , i.e. its mean value. Apparently, there is always such a D that obeys this equation, and thus every set can be considered as (random) fractal. This approach, however, is as useful as the claim that each curve is a line because of the existence of a linear regression for each set of points ($N > 1$), and consequently it cannot say anything relevant about specific properties of particular self-similar sets. Therefore, we have to be more specific in the attempt to generalize fractals in a useful way.

Generalized fractals

As we have shown, self-similarity is a general concept of aggregated geometrical structures which can be applied for any mechanism that forms species spatial distributions. Fractality is, by contrast, the result of the postulate of scale invariance of our measurement, which imposes very special conditions on the process of spatial clustering. Such a process, for example, would have to cause the negative correlation between the number of patches and the proportional area occupied, $\sum_{j=1}^{n_i} \lambda_{i,j}^2$, for all levels of aggregation (the larger is the fractal dimension, the stronger would necessarily be this correlation). Therefore, we have to broaden the concept of fractality to release its narrow and – at least for ecological systems – apparently unrealistic restrictions. We call the broader concept *generalized fractals*, which we define as structures that originate by replacing patches of the original pattern with any pattern formed by any process. The only condition on this process is that it has to remain essentially the same for all levels of aggregation, though it could vary in its parameters. Note that the process can be very broadly defined – it can be either entirely random or constrained in some aspects, as shown below.

To be accurate, we can define generalized fractals as follows. Each spatial structure which can be decomposed into several patches and their subpatches is a generalized fractal if there is a number of subpatches within the i th patch, n_i , and if there is a r_i ($0 < r_i \leq 1$) such that

$$F\left(\lambda_1^2, \lambda_2^2, \dots, \lambda_{n_i}^2\right) = r_i \quad (5.5)$$

for all i , where F is any function of the vector $\vec{\lambda}^2$. Note that for our purposes the general form of F can be replaced with the sum of particular functions, f , which yields

$$\sum_{j=1}^{n_i} f(\lambda_j^2) = r_i. \quad (5.6)$$

In nature, various processes can form spatially aggregated structures, and these processes will in fact probably differ between different levels of aggregation. However, we can expect that some very general statistical regularity should apply universally, across different levels of aggregation, so that real spatial distributions can be modeled using the universal assumptions (say, macro-assumptions) concerning these regularities. There are several possibilities:

1. We can assume scale invariance in the strict sense, which leads to $f(x) = x^{\frac{D}{2}}$ and $r_i = 1$ for all i (which leads to true fractals). We have shown that this requirement is probably biologically meaningless, and thus we have to set up other assumptions that could be more biologically relevant – or at least more easily interpretable.
2. Total proportion of area occupied within any patch is kept within each level and through all levels of aggregation, whereas the number, the sizes of individual patches, and their location are random.
3. Total proportion of occupied area of all subpatches within any patch is kept within each level of aggregation, but it varies randomly between these levels.
4. Everything is random. There is preference neither in number of subpatches nor in their sizes. Apparently, this is a null model for generalized fractals.
5. The probability distribution (but not the parameters of this distribution) that controls the randomness of the proportion of occupied area within patches is independent of the total area considered and is universal for all taxa. This can seem quite a special assumption, but it is actually the most general one, as will be explained in detail later (see also Appendix 5.I). Simply said, this assumes an existence of processes that are actually independent of scale, and that do not depend even on our choice of the “zero” level of aggregation.

Generating generalized fractals – models of more or less random multiscale aggregation

To explore the properties of spatial distributions characterized as generalized fractals, we constructed models of spatial distributions differing in the assumptions mentioned above, and compared their features. Model realizations were in all cases based on the modified percolation model (Falconer, 1990; Lennon *et al.*, this volume), the algorithm being identical in all cases. The models thus differ only in their parameters, and thus they are comparable to one another. The basic pattern was square, i.e. not only the total area considered (the zero level of aggregation) was square-shaped, but also all patches and their subpatches were squares. The models are as follows.

M1 *The fractal model*

In the first step, the fractal dimension, D ($0 < D \leq 2$), was drawn from the distribution of fractal dimensions observed for spatial distributions of bird species within the Czech Republic (see below). Then the square was overlapped with a grid of a randomly chosen number of square-shaped cells, drawn from a regular distribution between 2×2 and 5×5 . In the next step, n squares smaller than one grid cell were randomly located within randomly chosen (yet unoccupied) grid cells, n being drawn from a regular distribution between 1 and the total number of the grid cells (so that each grid cell could be finally either empty or occupied by just one square of random size and random position within the cell). The only condition was that the sum of proportional fractal areas of these squares, $\sum_{i=1}^n \lambda_i^D$, was equal to 1 (see Eq. 5.2). This process was applied repeatedly for each level of aggregation (i.e. for each new square), keeping D for all squares within all levels of aggregation. We stopped the process when reaching the fifth level of aggregation. Note that whereas in an ideal case the distribution of subpatches should be entirely random, here it is quasi-random because their location is constrained by the grid used, i.e. by locations of grid cells. This is the case of all following models as well, and the reason is that otherwise the process would be too excessively time-consuming, as large already located patches would too strongly constrain possible locations of the other patches in the case of large D .

M2 *The model of stable proportion of occupied area among levels*

In the first step, the proportional area occupied, r ($0 < r < 1$), was drawn from regular distribution and then the same procedure as for the fractal model was used. The only difference between these two models was in their parameters. While we used D as an exponent and 1 as the sum of areas in the previous case, we used 2 as the exponent and r as the sum of areas, respectively, in the case of this model. Therefore, this model did not keep the fractal dimension across all levels of aggregation, keeping instead r for each patch across all levels of aggregation and changing the number of subpatches within each patch.

This model can be biologically interpreted as the model of habitat hierarchy where a species occurrence is restricted to some portion of respective level of habitat hierarchy, and this portion remains the same across all levels and all patches. Species can differ in this proportion (as r can vary among species), but each species has some ability to be potentially present in a part of available habitats regardless of the level of habitat hierarchy, and this ability can be determined, for example by its niche width (i.e. specialists would potentially occupy smaller portion of habitat on each level).

M3 The model of stable proportion of area within levels

This model is similar to the previous one except that the proportion of occupied area r is set up separately for each level of aggregation (but remains the same for different patches within one level). This is actually more biologically reasonable than the previous model, as it is probable that each level of aggregation (say, level of habitat hierarchy) is very different and thus the proportion of patches where species occurrence is allowed varies from level to level.

M4 The random proportion model

Here we set up the r separately (and randomly) not only for each level of aggregation, but also for each patch. This model can thus be regarded as an entirely null model, as nothing is kept stable during the process of the emergence of respective spatial structure.

M5 The area- and taxa-invariance model

The main idea behind this model is that the previous model is not as “null” as it looks, because the rules comprising individual patches cannot be equally applied to the basic squares of all sizes (see Appendix 5.I). Therefore, the model M4 in fact assumes that there is some basic level (the zero level of aggregation) which is given, and which cannot itself represent a part of some larger area. Model M5 ensures that the same rules can be applied for basic squares of all possible sizes, i.e. it is independent of the scale we start with. In practice this means that the process is the same as in the previous model (M4), but r is not chosen from the uniform distribution but from the distribution which had been proven to be indeed independent of scale (for details see Appendix 5.I and Fig. 5.8). The parameters of this distribution were generated randomly, but since this process allowed too wide a range of possible patterns, we selected the set of simulations with the distribution of fractal dimensions equal to the observed distribution of D .

Model properties and tests

We explored statistical properties of all the models, choosing properties which are biologically relevant and have been studied previously. These are the ability of the models (1) to predict the scale-occupancy relationship, or more accurately the relationship between area and the probability of species occurrence, (2) to predict the species-area relationship and (3) to predict the distribution of occupied Euclidean area, which we assume to be proportional to the species abundance distribution. We compared these model properties between models M1–M5 and also with respective patterns of species distributions and abundances in central European birds. Species distribution data comprised quadrat-based distributional atlases of breeding birds of the Czech Republic (Št'astný, Bejček & Hudec, 1996) and Europe (Hagemeijer & Blair, 1997), from which we

used squares of 16×16 quadrats (see Šizling & Storch, 2004 for details). Data on abundance distributions were obtained from BirdLife International/EBCC (2000), comprising estimated abundances of breeding birds from 33 European countries excluding Russia and Turkey. These abundance data represent probably the only suitable set for this type of comparison, because they are at the same time large scale and reasonably accurate, containing information on abundances across a large range of magnitudes, from one breeding pair to several million. Abundances were calculated as a geometrical mean of the upper and lower estimate for each country if the lower estimate was not zero; if it was, the abundance was calculated as an arithmetical mean.

The relationship between area and probability of occurrence

The relationship between the area of grid cell and the probability of occurrence within cells of respective area (hereafter p-area relationship) is a very important measure of scaling of spatial distribution. It potentially provides a way to estimate species occupancies or even abundances within fine scales from coarse-scale censuses (Kunin, 1998; Kunin, Hartley & Lennon, 2000; He & Gaston, 2000; He & Condit, this volume; note, however, that these authors were not dealing with the probability of occurrence but with relative occupancy; nevertheless, these measures converge in the limit as area approaches zero). The knowledge of p-area functions is sufficient for the construction of species-area curves simply by summing p-area functions for individual species (Coleman, 1981; Ney-Nifle & Mangel, 1999; Šizling & Storch, 2004). The p-area relationship is generally assumed to be the power law in the case of fractals (i.e. it should be linear in the log-log scale, although see He & Condit, this volume), and for this reason we tested the linearity of the growing part of these relationships for the observed data and models. As a measure of linearity we used the maximal correlation coefficient between $\log(\text{area})$ and $\log(\text{probability of occurrence})$ obtained by rotating axes. This correlation coefficient does not depend on the slope of the regression line and reflects purely the linearity of the relationship. The curvilinearity of this relationship may, however, generally depend on species rarity – the rarer the species, the less linear it is (He & Condit, this volume). For this reason we plotted the maximum correlation against the calculated fractal dimension of respective spatial distribution (assuming that D tightly correlates with both abundance and occupancy) to have comparison across different rarity classes.

All the models have apparently equal ability to predict the shape of the p-area relationship (Fig. 5.3). When using the Nachman's and logistic models instead of the power law (He & Condit, this volume), i.e. using $-\ln(1 - p)$ and $\frac{p}{1-p}$, respectively, instead of the logarithm of probability of occurrence $\ln p$, we obtained smaller correlations with the logarithm of area, i.e. larger curvilinearity, for both models and observed data – but again, all the models were indistinguishable from each other and also from the observed p-area relationships.




Figure 5.3 Maximum correlation coefficient (see the text) of the growing part of the p -area relationship in the log-log scale, plotted against the fractal dimension of the spatial distribution. The correlation coefficient is a measure of the linearity of the p -area relationship in the logarithmic scale (i.e. how close it is to the power law) and the fractal dimension correlates with species occupancy (and abundance). Note that all correlation coefficients are quite high, and correlations for different models fill approximately the same space in the plot. This indicates that all models, as well as the observed data, provide the p -area relationship which is very close to the power law. CR, Czech Republic; CE, central Europe; for details of the data see Šizling & Storch (2004).

Another possibility for comparing the models in terms of their ability to predict the p -area relationship is to compare the p -area relationships produced by the models with an “ideal” power-law relationship. In this case, we compared the probability of occurrence for each area predicted by each model with the probability predicted by the power-law relationship obtained by approximating the modeled relationship with the regression line in the log-log scale up to the point of saturation (i.e. up to the point where the probability of occurrence was 1; see Šizling & Storch, 2004). All the models, including the fractal one (M1), as well as observed spatial distributions, apparently deviate from the power-law p -area relationship (Fig. 5.4) but the mean deviation of all the models is relatively low and – more importantly – approximately the same for all models and observations for all areas. Notably, the highest deviation was revealed by the observed spatial distributions of birds within the Czech Republic, and also by the area- and taxa-invariance model (M5) whose parameters were obtained from these observed distributions. All the other models are effectively




Figure 5.4 Mean (marks) and $CI_{0.05}$ of the mean (whiskers) of deviations, ε , from the finite area model for assemblages of 200 species in the cases of individual models, and for observations of 142 and 193 bird species within the Czech Republic (CR) and central Europe (CE), respectively. The ε was calculated for each area as the difference between the probability of occurrence calculated by the finite area model (i.e. the power law bounded by the point of saturation) and the probability predicted by each model or observation.

indistinguishable from one another in terms of their mean deviation from the power-law p–area relationship.

The species–area relationship

The species–area relationship (hereafter SAR) can be exactly predicted using the knowledge of spatial distribution of individual species, because mean number of species occurring within a plot of area A is exactly equal to the sum of probabilities of occurrence in this area across all species (Coleman, 1981; Ney-Nifle & Mangel, 1999; Lennon, Kunin & Hartley, 2002; Šizling & Storch, 2004). Therefore, we can construct the SAR simply by summing the p–area curves for all species. We constructed 200 realizations (i.e. 200 randomly generated “species”) of each model, and compared resulting SARs with one another and also with SARs observed for birds in the Czech Republic and central Europe (after normalizing the species numbers and areas). We have previously shown that the fractal model (M1) produces SARs which are indistinguishable from those observed, provided that the distribution of fractal dimensions used in the model is the same as the observed distribution of measured fractal dimensions (Šizling & Storch, 2004). Also, it is not surprising that the model M5 gives the SAR which is very close to the observed one, as this model was also based on




Figure 5.5 Species-area curves for individual models (marked thin lines) and observed avifauna within the Czech Republic (bold full line) and central Europe (bold dashed line). The prediction of the fractal model (M1) is not shown, as this model needs a parameterization by some distribution of fractal dimensions, when it has been shown to exactly follow the observed SARs (Šizling & Storch, 2004). Note that the curvilinearity and the associated deviation of modeled species numbers from those observed for large areas could be potentially due to the fact that the construction of generalized fractals was not entirely random (see the section headed “Generalized fractals”), and this quasi-randomness apparently affected large areas more strongly.

an observed distribution of fractal dimensions. Much more surprisingly, the other models which were not dependent on any ad-hoc parameterization produced SARs which were also very close to observed SARs, and individual models predicted SARs very similar to each other (Fig. 5.5). The species-specific properties which must have been adjusted to give exact prediction of the SAR in the case of the fractal model (i.e. fractal dimensions) thus emerged from the other models themselves. Indeed, the distributions of fractal dimensions (calculated by standard box-counting method) predicted by the models closely followed the observed distributions (Fig. 5.6). The shape of the SAR is therefore very well predicted by several models of generalized fractals without any assumptions concerning parameters of species spatial distributions. Note also that the slope of the linearly increasing part of predicted SARs is approximately 0.17, which is very close to the generally observed mean slope of mainland SARs (0.15; Rosenzweig, 1995).

Frequency distribution of the occupied Euclidean area

So far, we have shown that particular generalized fractals (including the scale invariant sets in the original meaning) have similar abilities to predict the p-area relationships for individual species, as well as the species-area relationships.




Figure 5.6 The distribution of fractal dimensions that emerged from models M2–M4, expressed as the rank–dimension relationship (the rank was rescaled to obtain comparable curves for observations and predictions). Fractal dimensions were extracted from 200 independent simulations of species spatial distributions (open diamonds, circles, and closed triangles for M2, M3 and M4, respectively) and the observed distributions of fractal dimensions of bird spatial distributions within the Czech Republic (CR, full squares) and central Europe (CE, open squares). The results of models M1 and M5 are not shown, as in these cases the frequency distribution of fractal dimensions was the input parameter and was set to be equal to the distribution observed for the avifauna of the Czech Republic. Note that the distributions predicted by the models fall between the distributions observed for CR and CE, with the distribution produced by the model M2 falling exactly in the middle.

Consequently, it is no wonder that most published studies showed that species spatial distributions are more or less close to fractal, even if they were not true fractals at all. To show the differences between particular self-similar models, we have had to test the property in which they should differ, and this is the distribution of the Euclidean areas they produce. Each such area corresponds to the area potentially occupied by an individual species and should be in some level of aggregation proportional to species abundance. The strong proportionality would be expected if the spatial requirements of individuals were equal for all species, otherwise the area divided by the mean individual's spatial requirement (i.e. mean home range) should be proportional to abundance.

The distributions of the Euclidean areas occupied, P , obey formulae

$$P_n = L^2 r^n, \quad (\text{M2})(5.7)$$

$$P_n = L^2 \prod_{j=1}^n r_j \quad (\text{M3})(5.8)$$

and

$$P_n = L^2 0.5^{n-1} r, \quad (\text{M4})(5.9)$$

where r is a random number between 0 and 1 (drawn from the regular distribution), n is the number of levels of aggregation and L^2 is the area of the original patch, in the case of M2, M3 and M4, respectively. The area occupied is in these cases independent from any parameter but r (for details see Appendix 5.II). In the cases of the other two models (M1 and M5), the distribution of area occupied depends on an external parameter such as the distribution of fractal dimensions, and this cannot be expressed so easily. We calculated these distributions numerically for various levels of aggregation using the models depicted above.

To compare the distributions of areas occupied for the individual models with observed distributions of species abundances, we extracted the variance and normalized skewness of the distribution of their logarithmically (base 2) transformed values from both the model predictions (for various levels of aggregation) and the abundance distributions of birds in European countries. Only two models fitted the data well: the model of the stable proportion of area within levels (M3) and the area- and taxa-invariance model (M5). In the case of the model M5 it is not too surprising, as this model always provides almost lognormal distribution of areas, fitting even better than the lognormal model to observed distributions (using Kolmogorov–Smirnov statistics; Šizling & Storch, unpublished manuscript). The fit of model M3 is more interesting (Fig. 5.7). Note that we should not expect a perfect fit, as the areas occupied on n th level of aggregation cannot be exactly proportional to the abundances, considering that individuals from different species have different spatial requirements (home ranges). Indeed, when we included an assumption of different home range sizes, assuming an approximately lognormal distribution, the fit of the model M3 improved, and even the model of the stable proportion of occupied area among levels (M2) provided good predictions of species abundance distributions. Therefore the spatial distribution characterized by generalized fractals is likely to lie behind observed species abundance distributions.

Discussion and conclusions

Our results show that spatial structures that are effectively indistinguishable from true fractals can emerge by simple random hierarchical processes taking place on several spatial scales. Such processes can be modeled in several ways, differing in the constraints on the randomness. Quite intriguingly, the models which best fitted the observed data were those where the proportion of area potentially occupied within the lower level of aggregation was kept constant throughout space within

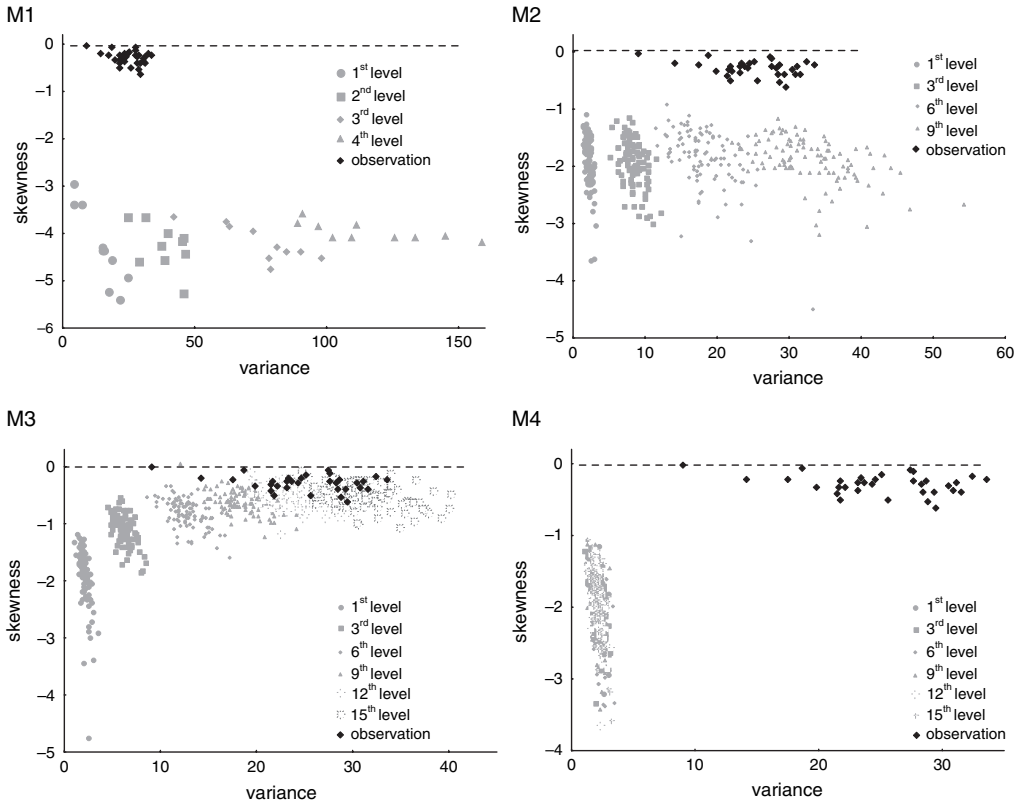


Figure 5.7 Variance and normalized skewness of the frequency distribution of occupied areas for models M1, M2, M3 and M4 (grey marks; the number of species for each assemblage/point was 200; the number of assemblages/points for each setting was 10 in M1 and 100 in M2–M4), and corresponding parameters for logarithmized abundance distributions of avifaunas of European countries (black marks). The parameters for different levels of aggregation are marked with different shapes (see legends inside the figures). Note that with the exception of the model M4, the more levels of aggregation, the higher is the variance of the respective distribution. The frequency distribution of fractal dimensions used for the fractal model (M1) was extracted from the observed central European avifauna data set (see text). Only in the case of model M3 did the parameters of observed and modeled distributions overlap, albeit only for relatively high levels of aggregation.

each level of aggregation (M3) or even among different levels of aggregation (M2), although in the latter case the occupancy/abundance distributions fitted only when accounting for different spatial requirements of individuals, i.e. assuming particular nonequal distribution of home ranges (unpublished simulations). The real pattern of species spatial distribution can in fact lie somewhere between these two models. Anyway, regardless of which model is closer to nature, all the models produce structures that are similar to fractals in most properties, although they are not true fractals, and are in fact much more directly interpretable biologically.

Different levels of modeled spatial aggregation can represent different levels of habitat association (macrohabitat – mesohabitat – resource patches) and/or of spatial population dynamics (e.g. range dynamics – local metapopulation dynamics – behavioral decision-driven changes in occurrence). Various proportions of area selected at each scale that ultimately lead to interspecific differences in the distribution and abundance can then be related to species-specific ecological properties such as body size, niche width, mean dispersal distance or competitive ability. Therefore, although we assume entirely random processes, they can in fact refer to particular biologically relevant mechanisms, and their “randomness” is just a way to treat different processes in one framework – or, in other words, it is a way to ensure that different spatial scales and processes taking place within them can be related to each other in a formal manner.

The finding that our model of random aggregation does not only produce spatial structures which are very close to observed fractal-like distributions, and predicts realistic species–area curves, but also gives the distributions of occupied area which are close to observed species abundance distributions, deserves attention. There are plenty of models of abundance distributions within species assemblages (for reviews see Tokeshi, 1999; Gaston & Blackburn, 2000), some of them based on niche divisions among species (Sugihara, 1980; Tokeshi, 1996), some on nonlinearities in population dynamics (Bell, 2000; Hubbell, 2001). Recently it has become clear that a successful model of the species abundance distribution within larger scales should take space seriously, as abundance always has an inherent spatial dimension. Therefore, good models must be implicitly or explicitly spatial. The most prominent of these models is Hubbell’s neutral theory of biodiversity, according to which the distribution of species abundance is the result of “community drift” and processes of random colonization, extinction, migration and speciation (see also He, 2005; Borda-de-Água *et al.*, this volume). However, this theory assumes that individual species do not differ ecologically, which is apparently an extremely restrictive assumption. Interspecific niche differences are almost surely as important as space, and a successful theory should comprise them as well. There are models that comprise both space and interspecific niche differences (Marquet, Keymer & Cofré, 2003), and our model can be treated as one of them. Arguably, our model is the null model of species abundance distribution, because it does not assume any particular way in which niches are divided among species, neither any sort of interspecific interactions (either in evolutionary or ecological time). It is entirely individualistic (i.e. community-level properties emerge purely from species-level spatial processes) and ecological differences between species are assumed implicitly, as these contribute to the process of hierarchical random aggregation.

Generally, there are two major approaches which attempt to treat most macroecological patterns in species distribution and diversity within one universal framework. One of them is the above-mentioned Hubbell’s neutral theory

of biodiversity and biogeography (Hubbell, 2001; see also Bell, 2001; Chave, 2004). The second one is the HEAP (Hypothesis of Equal Allocation Probabilities) approach which attempts to derive most patterns from basic statistical assumptions (Harte *et al.*, 2005; Harte, this volume). Both approaches have, however, some limitations. Hubbell's theory does not rely only on the restrictive assumption of species per-capita ecological identity mentioned above, but also assumes complete biotic saturation, i.e. that total number of individuals per unit area remains constant, which leads to strong interspecific competition on space, albeit the outcome of this competition is not given a priori. By contrast, Harte's HEAP model represents a rather statistical description, whose biological interpretation is still unclear (see Harte, this volume). Moreover, although HEAP predicts many patterns in species distribution and diversity, it needs the abundance distribution of the whole assemblage as an input. On the contrary, our model of hierarchical random aggregation does not need any assumption concerning interspecific relationships or other community-level constraints, nor any external parameters to predict realistic patterns of species abundance distribution (the abundance distribution being one of its outputs).

It is not yet clear whether our theory is able to predict as wide a range of spatial ecological phenomena as those predicted either by Hubbell's neutral theory or the HEAP. However, since our theory predicts realistic spatial scaling of species distribution and diversity with a minimal set of assumptions and without free parameters, and is at the same time biologically reasonable, it can be treated as a null hypothesis of spatial scaling of ecological patterns.

Acknowledgments

We are grateful to László Hajnal for his mental support in discussions between biologists and mathematicians, and to Stephen Hartley for his useful comments on the manuscript. The study was supported by the Grant Agency of the Academy of Sciences of the CR (KJB6197401), Grant no. CC06073 of the Czech Ministry of Education, and the Research Program CTS MSM0021620845.

Appendix 5.1 The area- and taxa-invariant distribution

As mentioned above, the random proportion model (M4) is not simultaneously applicable for all sizes of patch area, and thus is not "scale invariant" even in the most general sense. Assuming, for example, its validity for an area of $500 \times 500 \text{ km}^2$, we cannot simultaneously assume its validity for any larger or smaller area. The reason is that this model assumes a regular distribution of proportional areas within each patch, r , through all of S simulations (i.e. for an assemblage of S species). However, if we do assume it, for example for the mentioned area $500 \times 500 \text{ km}^2$, we obtain an irregular distribution of r for an area composed from several such areas (Fig. 5.8). The area of $500 \times 500 \text{ km}^2$ is thus in this case privileged and must be considered as basal.

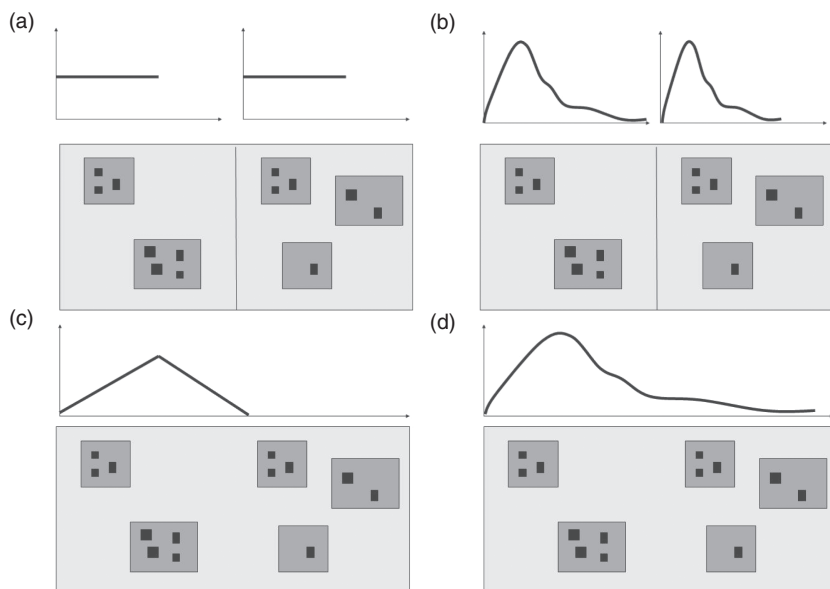


Figure 5.8 Top: the frequency distributions of the proportion of occupied area used for the construction of spatial distributions within the left and the right plots in the case of (a) the random proportion model (M4) and (b) the area- and taxa-invariance model (M5). Below: the resulting distributions of the proportion of occupied area when joining the left and right plots for models M4 (c) and M5 (d). While in the case of the regular distribution (a) the joining of the two areas yields a roof-shaped distribution (c), in the case of the multiexponential distribution (b) the distribution of joined plots remains multiexponential (d).

To build a model that is really independent of “scale” we have to find a distribution that conserves its form when enlarging the basal area (see Fig. 5.8). Since area is always positive, the distribution must allow only positive random values, and also should attain zero, as we never know how many species are absent on the modeled area. The distribution that obeys these conditions is independent of area chosen as the basal area (zero level of aggregation) and of the set of considered species, and we thus call it the area- and taxa-invariant distribution. We found that this distribution can be expressed using the multiexponential form

$$f(x) = \sum_{i=1}^{\infty} c_i (e^{-A_i x} - e^{-\alpha_i x}), \quad (5.10)$$

where A_i and α_i are positive and c_i is any real number (Šizling & Storch, manuscript in preparation). The number of additive terms, $c_i (e^{-A_i x} - e^{-\alpha_i x})$, does not have to be necessarily infinite – it is defined as infinite because it can vary unlimitedly as it arises whenever any two areas are joined. However, three or

four terms are usually enough for quite accurate description of observed data. Additionally, not all additive terms must be necessarily positive (when joining two areas, some of c appears negative), and thus the c_i does not necessarily approach zero as the number of additive terms approaches infinity.

For building the fitting procedure, which fits the multiexponential distribution to the data, we expressed the expectation of the distribution as

$$E(x) = \sum_{i=1}^{\infty} c_i (A_i^{-2} - \alpha_i^{-2}). \quad (5.11)$$

Obviously, the expectation consists of expectations for individual additive terms. Therefore, we can roughly fit this distribution on a sample of observations just by dividing the sample into several subsets of observations (as many as the number of the additive terms we would like to use), and by fitting each of these terms on one of these subsets separately. The only condition is that these subsets should have a similar number of observations.

One term can be fitted by normalizing it (by setting the integral from zero to infinity as equal to one), by which we obtain

$$E_i(x) = A_i^{-1} + \alpha_i^{-1}. \quad (5.12)$$

If $\alpha_i > A_i$, the expected value can vary between A_i^{-1} and $2A_i^{-1}$. Estimating the expectations as a mean, M , we can estimate the parameters as

$$A_i = \frac{2 - \kappa / (1 + \kappa)}{M_i} \text{ and } \alpha_i = (1 + \kappa)A_i \text{ where } \kappa > 0. \quad (5.13)$$

The vector of parameters of proportionality, c_i , was estimated as the best fitted vector from 200 vectors calculated as $c_i = \frac{\text{random}}{E_i(x)}$ where *random* is a random number between 0 and 1. The term “best fitted” refers to the vector for which the Kolmogorov–Smirnov statistics (i.e. the maximal deviation between two cumulative distribution functions) between the analytical form and the fitted data was minimal. The estimation is not too sensitive to the choice of κ but in an attempt to minimize numerical errors we set it as 10^{99} .

Appendix 5.II Calculating the distribution of occupied areas

To set up the general form for the total areas occupied we used the proportional areas, λ^2 , for each patch within each level of aggregation. For two levels of aggregation, for example, we can write

$$P_2 = L^2 \left(\begin{array}{l} \lambda_{i_1=1}^2 \left(\lambda_{i_1=1, i_2=1}^2 + \lambda_{i_1=1, i_2=2}^2 + \cdots + \lambda_{i_1=1, i_2=k_1=1}^2 \right) + \\ + \lambda_{i_1=2}^2 \left(\lambda_{i_1=2, i_2=1}^2 + \lambda_{i_1=2, i_2=2}^2 + \cdots + \lambda_{i_1=2, i_2=k_1=2}^2 \right) + \\ \cdots \\ + \lambda_{i_1=k}^2 \left(\lambda_{i_1=k, i_2=1}^2 + \lambda_{i_1=k, i_2=2}^2 + \cdots + \lambda_{i_1=k, i_2=k_1=k}^2 \right) \end{array} \right) \quad (5.14)$$

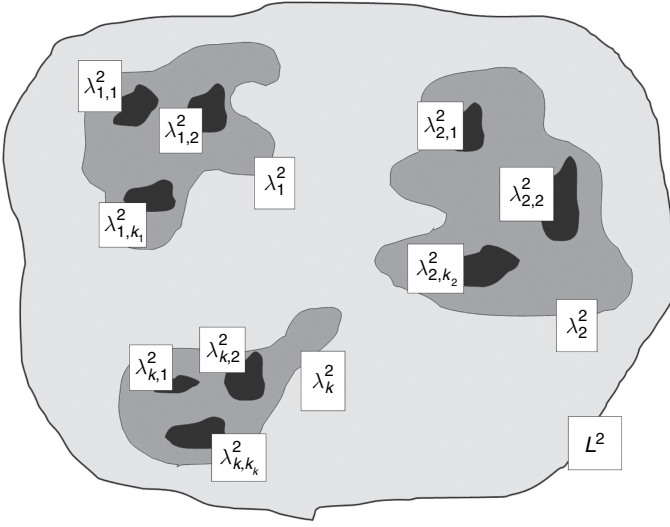


Figure 5.9 Three levels of aggregation with the labels indicating areas of individual patches and subpatches. Using these areas, total occupied area within respective level of aggregation can be calculated using Eq. (5.14). Note that the calculation is actually independent of the exact shape of individual patches, and the L and λ can be imaginary values (because they are the square roots of areas of the patches).

where P_2 is the total area occupied by patches at the second level of aggregation, L^2 is the total area of the area of origin, and indexes i_1 and i_2 refer to the order of a particular patch within the first and second level of aggregation, respectively. So, for instance, $\lambda_{i_1=m, i_2=n}^2$ means the proportional area of n th patch within the m th patch at the second level of aggregation (see Fig. 5.9; note that the indices are simplified in this figure so that $\lambda_{i_1=m, i_2=n}^2$ is replaced with $\lambda_{m,n}^2$, for example).

Making the form (5.14) more compact, we can write

$$P_2 = L^2 \sum_{i_1=1}^k \left(\lambda_{i_1}^2 \sum_{i_2=1}^{k_{i_1}} \lambda_{i_1, i_2}^2 \right) \quad (5.15)$$

in the case of two levels of aggregation, and generally

$$P_n = L^2 \sum_{i_1=1}^k \left(\lambda_{i_1}^2 \sum_{i_2=1}^{k_{i_1}} \left(\lambda_{i_1, i_2}^2 \sum_{i_3=1}^{k_{i_1, i_2}} \left(\lambda_{i_1, i_2, i_3}^2 \cdots \left(\lambda_{i_1, i_2, i_3, \dots, i_{n-1}}^2 \sum_{i_n=1}^{k_{i_1, i_2, i_3, \dots, i_{n-1}}} \lambda_{i_1, \dots, i_n}^2 \right) \cdots \right) \right) \right) \quad (5.16)$$

for n of those levels. This relatively complicated general form can be simplified in the cases of individual models as follows.

M2 The model of stable proportion of occupied area between levels

In this case, the proportional areas occupied, r , were kept constant through the whole set. In general, considering the form Eq. (5.16) and replacing all proportional areas at the lowest level with r , we obtain

$$P_n = L^2 \sum_{i_1=1}^k \left(\lambda_{i_1}^2 \sum_{i_2=1}^{k_{i_1}} \left(\lambda_{i_1, i_2}^2 \sum_{i_3=1}^{k_{i_1, i_2}} \left(\lambda_{i_1, i_2, i_3}^2 \cdots \sum_{i_{n-1}=1}^{k_{i_1, i_2, i_3, \dots, i_{n-2}}} \left(\lambda_{i_1, i_2, i_3, \dots, i_{n-1}}^2 r \right) \right) \cdots \right) \right) \quad (5.17)$$

After simplifying

$$P_n = L^2 r \sum_{i_1=1}^k \left(\lambda_{i_1}^2 \sum_{i_2=1}^{k_{i_1}} \left(\lambda_{i_1, i_2}^2 \sum_{i_3=1}^{k_{i_1, i_2}} \left(\lambda_{i_1, i_2, i_3}^2 \cdots \sum_{i_{n-1}=1}^{k_{i_1, i_2, i_3}} \lambda_{i_1, i_2, i_3, \dots, i_{n-1}}^2 \right) \cdots \right) \right), \quad (5.18)$$

and applying this replacement repeatedly, we obtain Eq. (5.7) presented above, which is the simplest form of the frequency distribution of the total occupied area.

M3 *The model of stable proportion of area within levels*

The only difference between this and the previous model is that here the proportional area changes when stepping down one level of aggregation. Here there is not the only one r for the whole set, but as many r_j as there are levels of aggregation (one r_j for each level). This leads to Eq. (5.8) presented above.

M4 *The random proportion model*

The calculation of the frequency distribution of total area occupied is here a bit more complicated. Let us assume two levels of aggregation again. Replacing each sum of proportional areas occupied at the second level of aggregation with its average value, 0.5, plus its deviation, ε , we obtain

$$P_2 = L^2 \left(\begin{array}{l} \lambda_{i_1=1}^2 (0.5 + \varepsilon_1) + \\ + \lambda_{i_1=2}^2 (0.5 + \varepsilon_2) + \\ \dots \\ + \lambda_{i_1=k}^2 (0.5 + \varepsilon_k) \end{array} \right). \quad (5.19)$$

After simplifying,

$$P_2 = L^2 0.5 (\lambda_{i_1=1}^2 + \lambda_{i_1=2}^2 + \cdots + \lambda_{i_1=k}^2) + L^2 (\lambda_{i_1=1}^2 \varepsilon_1 + \lambda_{i_1=2}^2 \varepsilon_2 + \cdots + \lambda_{i_1=k}^2 \varepsilon_k). \quad (5.20)$$

Obviously, the area is given by the area at the higher level of aggregation and an error term. The error term is a number drawn from a unimodal distribution (nearly Gaussian) which has average value of zero (because all ε_i follow the distribution with the mean of zero).

Replacing the proportional area of the first level with r_1 , we obtain

$$P_2 = L^2 0.5 r_1 + L^2 \varepsilon \quad (5.21)$$

in the case of two levels of aggregation, and using this algorithm repeatedly we obtain

$$P_n = L^2 0.5^{n-1} r_1 + L^2 \varepsilon \quad (5.22)$$

in general. The n is the number of levels of aggregation and r_1 is the proportional area occupied at the first level of aggregation. Since the average value of the error term, $L^2 \varepsilon$, is zero, for larger samples we obtain Eq. (5.9) presented above.

This formula could be criticized, however, for allowing an unrealistic extent of areas. For instance, the area of L^2 is not allowed by this form,

although, according to the model, it is principally possible (all levels of aggregation having r equal to one for all subpatches). The areas larger than $L^{2.0.5^{n-1}}$ are, by contrast, very improbable, so that Eq. (5.9) can be used for nearly any number of species (spatial distributions). This was verified comparing the outputs from a numerical model ($N = 50$; 4 levels of aggregation) with the analytical form derived.

References

- Allen, T. F. H. & Starr, T. (1982). *Hierarchy: Perspectives for Ecological Complexity*. Chicago: University of Chicago Press.
- Bell, G. (2000). The distribution of abundance in neutral communities. *American Naturalist*, **155**, 606–617.
- Bell, G. (2001). Neutral macroecology. *Science*, **293**, 2413–2418.
- BirdLife International/European Bird Census Council (2000). *European Bird Populations: Estimates and Trends*. Cambridge: BirdLife International.
- Chave, J. (2004). Neutral theory and community ecology. *Ecology Letters*, **7**, 241–253.
- Coleman, D. B. (1981). On random placement and species-area relations. *Mathematical Biosciences*, **54**, 191–215.
- Condit, R., Ashton, P. S., Baker, P., et al. (2000). Spatial patterns in the distribution of tropical tree species. *Science*, **288**, 1414–1418.
- Falconer, K. J. (1990). *Fractal Geometry: Mathematical Foundations and Applications*. Chichester: John Wiley.
- Gaston, K. J. & Blackburn, T. M. (2000). *Pattern and Process in Macroecology*. Oxford: Blackwell Science.
- Hagemeijer, W. J. M. & Blair, M. J. (1997). *The EBCC Atlas of European Breeding Birds*. London: T. & A. D. Poyser.
- Halley, J. M., Hartley, S., Kallimanis, S. A., Kunin, W. E., Lennon, J. J. & Sgardelis, S. P. (2004). Uses and abuses of fractals in ecology. *Ecology Letters*, **7**, 254–271.
- Harte, J., Kinzig, A. & Green, J. L. (1999). Self-similarity in the distribution and abundance of species. *Science*, **284**, 334–336.
- Harte, J., Conlisk, E., Ostling, A., Green, J. L. & Smith, A. B. (2005). A theory of spatial structure in ecological communities at multiple spatial scales. *Ecological Monographs*, **75**, 179–197.
- Hartley, S., Kunin, W. E., Lennon, J. J. & Pockock, M. J. O. (2004). Coherence and continuity in the scaling of species' distribution patterns. *Proceedings of the Royal Society of London, Series B*, **271**, 81–88.
- Hastings, H. M. & Sugihara, G. (1993). *Fractals, a User's Guide for the Natural Sciences*. Oxford: Oxford University Press.
- He, F. L. (2005). Deriving a neutral model of species abundance from fundamental mechanisms of population dynamics. *Functional Ecology*, **19**, 187–193.
- He, F. L. & Gaston, K. J. (2000). Estimating species abundance from occurrence. *American Naturalist*, **156**, 553–559.
- Hubbell, S. P. (2001). *The Unified Theory of Biodiversity and Biogeography*. Princeton: Princeton University Press.
- Kunin, W. E. (1998). Extrapolating species abundances across spatial scales. *Science*, **281**, 1513–1515.
- Kunin, W. E., Hartley, S. & Lennon, J. J. (2000). Scaling down: on the challenge of estimating abundance from occurrence patterns. *American Naturalist*, **156**, 560–566.
- Lennon, J. J., Kunin, W. E. & Hartley, S. (2002). Fractal species distributions do not produce power-law species-area relationships. *Oikos*, **97**, 378–386.

- Marquet, P. A., Keymer, J. E. & Cofré, H. (2003). Breaking the stick in space: of niche models, metacommunities and patterns in the relative abundance of species. In *Macroecology: Concepts and Consequences*, ed. T. M. Blackburn & K. J. Gaston, pp. 64–81. Oxford: British Ecological Society and Blackwell Science.
- Ney-Nifle, M. & Mangel, M. (1999). Species-area curves based on geographic range and occupancy. *Journal of Theoretical Biology*, **196**, 327–342.
- Richardson, L. F. (1961). The problem of contiguity: an appendix of statistics of deadly quarrels. *General Systems Yearbook*, **6**, 139–187.
- Rosenzweig, M. L. (1995). *Species Diversity in Space and Time*. Cambridge: Cambridge University Press.
- Šizling, A. L. & Storch, D. (2004). Power-law species-area relationships and self-similar species distributions within finite areas. *Ecology Letters*, **7**, 60–68.
- Št'astný, K., Bejček, V. & Hudec, K. (1996). *Atlas of Breeding Bird Distribution in the Czech Republic 1985–1989*. Jihlava: Nakladatelství a vydavatelství H & H, in Czech.
- Sugihara, G. (1980). Minimal community structure: an explanation of species abundance pattern. *American Naturalist*, **116**, 770–787.
- Tokeshi, M. (1996). Power fraction: a new explanation for species abundance patterns. *Oikos*, **75**, 543–550.
- Tokeshi, M. (1999). *Species Coexistence: Ecological and Evolutionary Perspectives*. Oxford: Blackwell Science.
- Ulrich, W. & Buszko, J. (2003). Self-similarity and the species-area relation of Polish butterflies. *Basic and Applied Ecology*, **4**, 263–270.
- Virkkala, R. (1993). Ranges of northern forest passerines: a fractal analysis. *Oikos*, **67**, 218–226.