CHAPTER NINE

On the quantification of local variation in biodiversity scaling using wavelets

TIMOTHY H. KEITT
University of Texas at Austin

Introduction

It is obvious to even the most casual naturalist that community composition varies through space and that this variation is related in some way to environmental gradients. Yet the problem of understanding in detail patterns of species turnover in space remains a difficult and unsolved challenge in spatial ecology. How do we relate spatial turnover in community composition to the underlying biological processes regulating the origin and maintenance of biological diversity in landscapes? The basic components are simple. If each species is restricted to a unique set of preferred environmental conditions and the environment varies spatially, then in a deterministic world, spatial variation in community composition simply reflects environmental variation. If rates of environmental change are constant, we should find a highly regular and simple pattern of species turnover. The complicating factors are of course ecology, history and variation in the "texture" of the physical environment. Species unable to track a temporally changing environment will be present in suboptimal conditions and missing from the preferred environment (Vetaas, 2002). The vagaries of extinction–recolonization dynamics imply that some fraction of species which could occur in some particular local environment will be missing solely for historical reasons (Hanski & Gyllenberg, 1997). Species may also be absent from preferred habitats because of competitive exclusion or absence of an obligate mutualist (Caley & Schluter, 1997). Further complicating the picture are lags in species responses to changing conditions simply owing to long life span or because of the length of time required for indirect ecological interactions to propagate through a community (Ives, 1995). In addition, rates of environmental change are highly variable through space and so we may predict that patterns of species turnover also are highly variable. We are led logically to the somewhat uncomfortable conclusion that spatial patterns in biodiversity depend strongly on conditions unique to particular locations. How then can we understand the pattern without being overwhelmed by its inherent complexity?

Community ecologists seem to have approached this question from two opposite ends of a spectrum. A long tradition in ecology has been content to

Scaling Biodiversity, ed. David Storch, Pablo A. Marquet and James H. Brown. Published by Cambridge University Press.

document the enormous variation in communities, fortunately so, as we would have few data to debate were it not for the dedicated efforts of field ecologists. At the other extreme, ecologists have attempted to characterize spatial turnover with a single number, the exponent of a power law relating number of species to sample area. Species–area analysis makes considerable sense in the context of island biogeography where analysis is constrained to species lists from islands of different size (MacArthur & Wilson, 1967). Ocean waters buffer island climate resulting in reduced local environmental variation and the separation of islands means they can reasonably be treated as point samples. Both properties lend themselves to species–area analysis. Application of species–area analysis to regional- to continental-scale landscapes seems far less appropriate, despite heroic efforts to predict the species–area exponent from paired island-like samples embedded in the larger landscape (Harte, Kinzig & Green, 1999). The difficulty is that for reasons stated above, rates of species turnover are likely to be highly variable from location to location, implying that a considerable amount of information is lost by taking a global average of this behavior. A far more compelling approach is to embrace this variation and attempt to explain the pattern in terms of local environmental contexts and the general principles of habitat selection (Pulliam & Danielson, 1991) and isolation-by-distance effects (Slatkin, 1993) that generate and redistribute biodiversity on the landscape. Allowing for localized variation also avoids a tautological conclusion (Klvana, Berteaux & Cazelles, 2004): if scaling patterns are invariant, we can have greater confidence if scale invariance is not a fundamental assumption as is the case when assuming a power-law form of the species–area relationship.

As an initial step towards a functional understanding of biodiversity scaling, it is desirable to identify methods that can capture local variation in species turnover without overwhelming the analyst with the full detail of the raw data. We present here a method of quantifying local turnover in community composition at multiple scales based on the wavelet transform (Daubechies, 1992). The wavelet transform arose in the context of signal processing and was designed precisely to optimize the trade-off between estimating global patterns and adapting to local variation (Mallat, 1999). Wavelets achieve this optimum by measuring local fluctuations in increasingly coarse-grained or locally averaged representations of the original signal (Walker, 1999). In this sense, wavelets are similar to block-quadrat methods developed in vegetation science (Dale, 1999) (see also Gaston *et al.*, this volume). However, unlike more heuristic methods, concepts of pattern and scale have precise mathematical definitions within the wavelet framework facilitating interpretation of results and comparisons among systems. In the remainder of this chapter, we briefly describe the wavelet transform and construct a measure of scale-dependent community dissimilarity in the wavelet domain. We then apply this method to spatial turnover in breeding bird community composition within the State of Texas, USA.

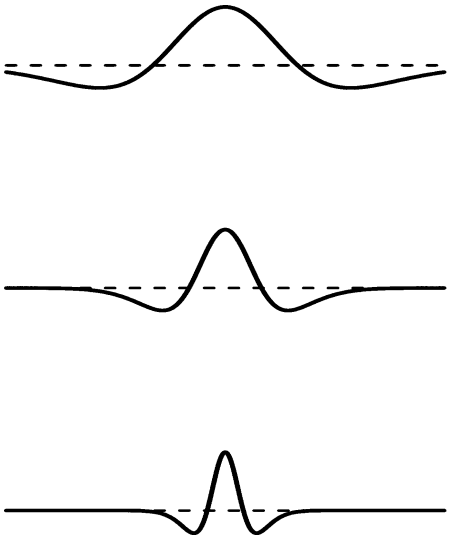

Figure 9.1 The Difference-of-Gaussians wavelet at three different scales. The narrower wavelets extract fine-scale information, whereas the wider wavelets extract broad-scale information.

Wavelets

Wavelets (Daubechies, 1992) were used to extract scale-specific information from the bird community data. Wavelets like the Fourier transform and other harmonic methods allow one to isolate variation at specific frequencies in data. One can imagine the pitch controls on a high fidelity stereo system. If we wish to hear only the low-frequency sounds we turn up the bass and turn down the midrange and treble controls. Wavelets operate in a similar manner – by tuning the width of a chosen wavelet (see Fig. 9.1), we can isolate and analyze separately low-frequency, midrange or high-frequency components of a signal. The Fourier transform is ideal when analyzing a pure stationary sinusoidal signal, such as that produced when a single note is played on a piano, because sine and cosine functions form the basis of the Fourier transform. For signals that have sharp transients – e.g. the sound wave produced by a snare drum – or where the dominant frequency changes over time, the Fourier transform is not ideal because it requires a complex superposition of sine and cosine function such that they cancel in precisely the right way to capture the transient nature of the signal. The interference between sines and cosines fitted to different parts of the signal leads to a large number of extraneous values being returned from the transform. For transient nonstationary patterns, the wavelet transform is preferred because wavelets rapidly decay to zero at longer distances and can thus capture localized behavior without interfering with other wavelets applied to the signal at other points in time or space.

Wavelets come in a wide variety of shapes, some more suited to capturing transient behavior and others more suitable for extracting regular oscillations.

Wavelets can further be classified as either continuous or discrete. For this study, the two-dimensional continuous wavelet transform (Antoine, 1999) was used, defined mathematically by

$$(T^{\mathrm{wav}}f)(a,b,s) = \frac{1}{h(s)} \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \psi\left(\frac{x-a}{s}, \frac{y-b}{s}\right) f(x,y)\mathrm{d}x\mathrm{d}y, \tag{9.1}$$

where $f(x, y)$ represents the spatial pattern of interest and $\psi\left(\frac{x-a}{s}, \frac{y-b}{s}\right)$ is a wavelet with width s centered at a,b. The function $h(s)$ is a normalization that standardizes the variance of the transform across scales. The dominant scale of analysis is governed by the scale parameter s. For the wavelet used in this study, patterns repeating every s distance units were highlighted by the transform. Variation at scales much larger or smaller than s distance units was suppressed by the wavelet filter. We used an adaptive "Difference-of-Gaussians" (DoG) wavelet (Fig. 9.1). Further details are given in the Appendix.

Community dissimilarity in the wavelet domain

Wavelet-based Euclidean distances were used to examine changes in community composition as a function of scale within localized contexts. Given lists of species abundances from two sample areas, recall that the Euclidean distance between the samples is simply the root sum of squared differences in abundance computed species by species. The DoG wavelet used for analysis (see Appendix) was computed by taking the difference between two nested samples, one with a narrow radius and another with a broader radius, both centered on the point of interest. Let $f_i(x, y)$ be the density of the ith species selected from a pool of N species, then

$$D^{\mathrm{wav}}_{a,b}(s) = \sqrt{\sum_{i=1}^{N} [(T^{\mathrm{wav}}f_i)(a,b,s)]^2} \tag{9.2}$$

is "wavelet dissimilarity" computed as the Euclidean distance between the nested samples centered at a,b and highlights community turnover occurring over s distance units (Fig. 9.2).

This measure of community dissimilarity is similar to other methods of measuring beta diversity across scales. An alternative approach, used by for example Condit *et al.* (2002), would be to plot spatial separation between samples versus a measure of distance between community composition and then examine the decay of similarity (or increase of dissimilarity) as a function of increasing spatial separation between samples. In this case, the pairwise differences are accumulated and then averaged within distance bins or fit directly from theory or regression. The power of this approach is in modeling the specific form of the decay curve. What is different is that community similarity is always computed at the sample plot level – the relationship between distance

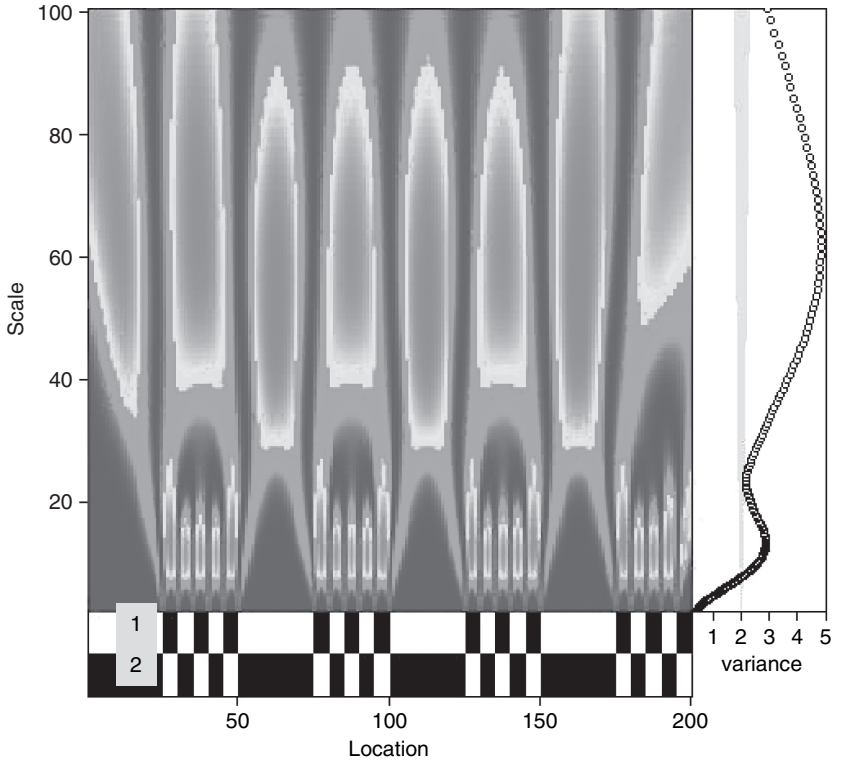

Figure 9.2 Illustration of wavelet dissimilarity applied to simulated data. The 10 species were divided into two communities labeled 1 and 2. The communities alternate along a linear transect according to the black and white sequences at the bottom with white indicating the presence of that community and black the absence. The color plot shows the output of the wavelet transform with blue indicating small values of community dissimilarity and red colors indicating large values. The average behavior is shown to the right as a function of scale. (For color version see Plate 3.)

and similarity is modeled in a second step, typically by averaging within small distance bins. Using wavelets, the averaging and differencing steps are combined into a single operation. To increase scale, the wavelet smoothly averages community composition over increasing regions and then computes differences between these increasingly larger average species pools. The result of computing differences on spatially averaged data is that the measure is quite stable, even at a single location, thus allowing one to look at community turnover as a function of scale at a particular location. With species–area and similarity-by-distance methods, results are generally averaged across the entire landscape. The wavelet method is not unlike breaking a sample transect into increasingly larger blocks and then computing differences between pooled communities in neighboring blocks (so called "block-quadrat" methods in plant ecology). Dale

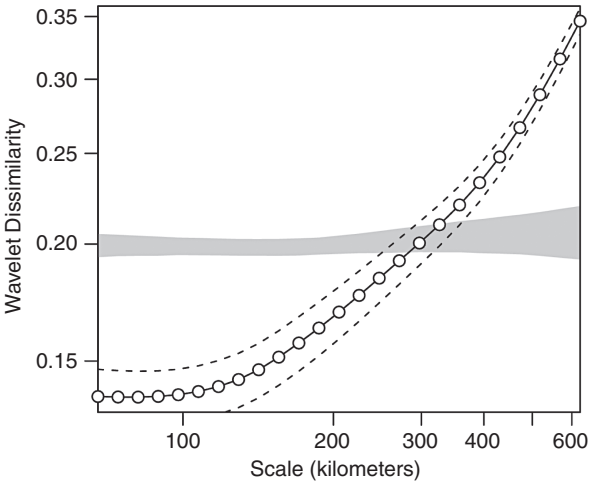

Figure 9.3 Community turnover in BBS data measured at different scales using wavelets. The shaded area bounds 95% of 1000 bootstrap randomizations of the data. Open circles are the spatial average of dissimilarity computed at each route. Dotted lines are ± 1 standard error.

(1999) has pointed out that this procedure is essentially a Haar wavelet transform and so is conceptually aligned to the method applied here. However, the Haar wavelet has certain undesirable properties, particularly in the case of irregular spatial sampling (Daubechies, 1992), hence our preference for the less abrupt DoG wavelet. A further advantage we see of the wavelet method is the natural interpretation of linear regression models on wavelet transformed data (Keitt & Urban, 2005) which would allow, for example, one to test the relationship between spatial turnover in environmental data and spatial turnover in species composition.

Application to community data

Data from the North American Breeding Bird Survey (BBS) were used to study spatial variation in community turnover. The BBS is a census spanning much of North America conducted each spring and was designed to track changes in population sizes of breeding bird species. Once per year, trained observers take point counts at regular intervals along 39-km sampling "routes". Data analyzed here are route totals for each species and locations given are for the route midpoints. Additional details can be found in Peterjohn and Sauer (1993). The full BBS data set was subsetted to include only those routes falling within the State of Texas, USA.

Spatially averaged results of the wavelet dissimilarity analysis are shown (on logarithmic scales) in Fig. 9.3. The pattern of wavelet dissimilarity deviated strongly from a null model where route totals for each species were randomly assigned to route locations. Average dissimilarity was much less than for the null model at scales below 300 km, indicating greatly reduced species turnover

at these scales. The likely explanation is that small regions constitute much less environmental variability and therefore permit fewer community types relative to a purely random sampling of the regional species pool. However, the direction of causality is not necessarily clear. If species distributions are strongly clustered, perhaps because of dispersal limitation or other factors unrelated to environmental turnover, then one still expects to see reduced community turnover as local species pools will be much smaller than the regional pool.

Fortunately, the slope of the dissimilarity curve provides some insight into potential processes regulating diversity. Steeper slopes indicate greater spatial patterning and changes in slope can indicate sharp transitions across ecological boundaries. Notice that the slope of the dissimilarity curve was essentially flat at the finest scales analyzed (50–100 km), albeit still greatly reduced relative to the null model. Again, the overall reduction relative to the null model indicates a sharply reduced local species pool relative to regional diversity. However, the flat slope, paralleling the null model, indicates little or no spatial pattern in community composition across the landscape at the finest scales. This could result from spatially uncorrelated observer errors obscuring the pattern of community turnover, or could result from localized random dispersal events that effectively blur community boundaries. Above 100 km in scale, the curve begins to slope upward, indicating greater patterning and the emergence of distinct community types.

At scales greater than 300 km, dissimilarity greatly exceeds that of the null model. This is likely the result of broad environmental gradients, such as moisture availability and temperature, that are well-known features of the Texas climate. There also appears to be a slight increase in the slope of the dissimilarity curve around 400 km in scale. This is the scale at which one begins to cross a major divide between a wetter, forested east Texas and a dry desert west Texas.

We assessed spatial variation in biodiversity scaling by computing a local community dissimilarity exponent. This was done by computing the median of $\Delta \log D_{a,b}^{\text{wav}}(s)/\Delta \log s$ across 25 scale increments spaced between the minimum and maximum scales analyzed. (Note that smaller exponents should generally correspond to larger species–area exponents, although the relationship is not exact.) These results are plotted in Fig. 9.4. A wide range of scaling exponents was observed, suggesting considerable local variation in biodiversity scaling. No strong spatial patterning was apparent. There did appear to be a band of higher exponents running from the eastern border through central Texas towards the Big Bend region in lower west Texas; however, the pattern is not sufficiently robust to warrant further interpretation. More important appears to be the local variation in dissimilarity scaling. A thorough understanding of the local variation would require detailed environmental data and likely field studies and is beyond the scope of this chapter.

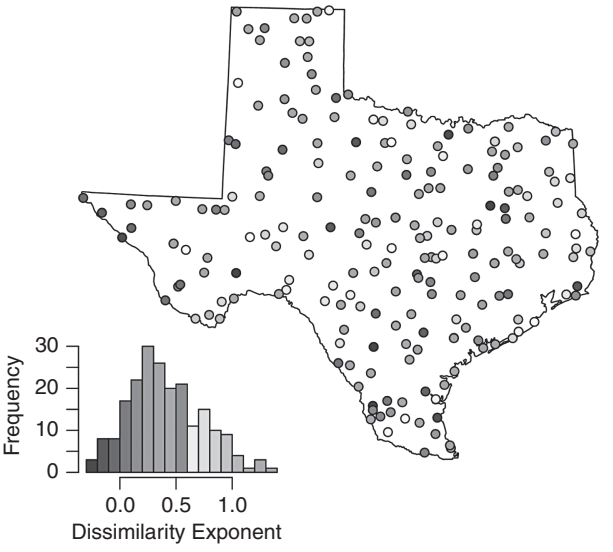

Figure 9.4 Spatial variation in biodiversity scaling across the Texas BBS data set. The inset shows a frequency histogram of values. (For color version see Plate 4.)

Somewhat more insight can be obtained through examination of the pattern of turnover "hot spots" (locations with significantly greater dissimilarity in comparison with the null model at a particular scale) and "cold spots" (locations with significantly reduced dissimilarity in comparison with the null model at a particular scale). These results are plotted in Fig. 9.5.

At the finer scales (75 km–150 km), many routes showed significantly less turnover than the randomized null model ($\alpha = 0.05$). Three areas did, however, show up as turnover hot spots with larger wavelet dissimilarity values than expected at random. One of the hot spots includes several routes in east Texas located just south and east from Dallas, a major metropolitan area. It is difficult to ascertain why these routes are hot spots without field data. However, one can speculate that these routes have been engulfed by suburban sprawl since their initiation and are now sampling a large pool of invasive bird species, such as European Starling (*Sturnus vulgarus*) and House Sparrow (*Passer domesticus*), not present in neighboring survey routes.

Another area includes several routes in the southern panhandle region (northwest quadrant) located between the cities of Midland-Odessa and Lubbock. Again, it seems likely that this is an anthropogenic hot spot driven by extensive irrigation for agriculture. (Green irrigation circles are readily apparent in satellite pictures of this area available at http://maps.google.com/.) The high productivity, increased humidity and anthropogenic disturbance in this area must contrast sharply with neighboring semidesert habitats, possibly leading to high rates of species turnover.

Another hot spot occurred in southern Texas along the Mexican border. This area is well known to ornithologists as a region of high species diversity and rapid

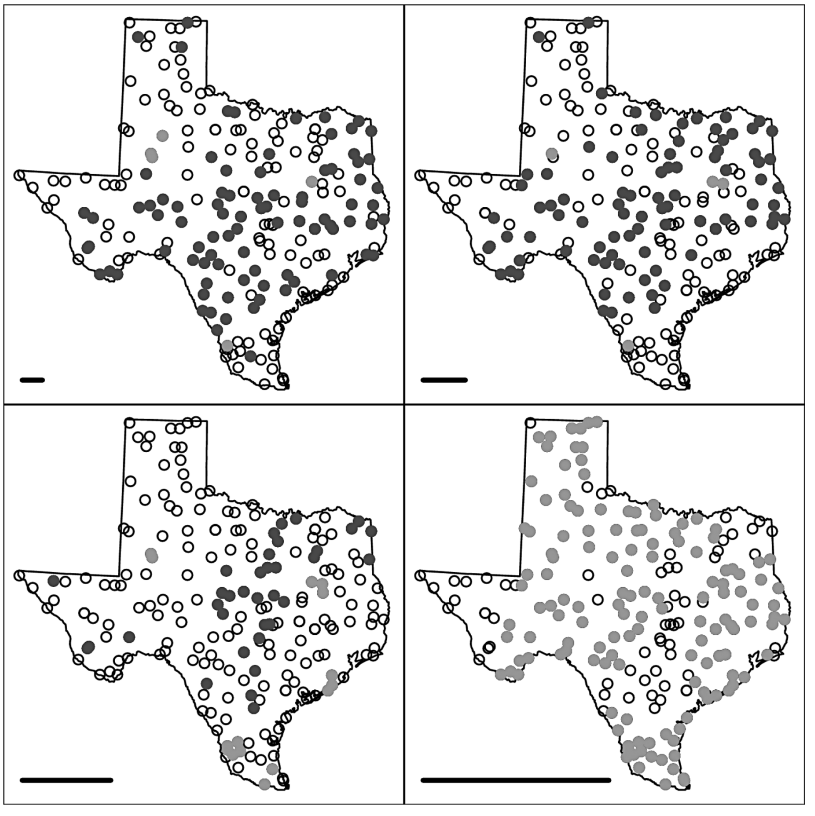

Figure 9.5 Hot and cold spots for community dissimilarity plotted at four scales. Open circles correspond to BBS routes that were not significantly different from the result of 1000 bootstrap randomizations of the data. Blue-filled circles showed significantly ($\alpha = 0.05$) less turnover at that scale and red-filled circles significantly more. The horizontal bar at the lower left of each panel indicates the scale of analysis. (For color version see Plate 5.)

community turnover. A large number of bird species whose primary affinities are tropical reach their northern range limit along lower Rio Grande valley. There is a rapid transition as one moves northward to a more temperate bird fauna and this likely accounts for the high wavelet dissimilarity indices in this region.

At the largest scales (>600 km) most routes showed significantly higher turnover than expected at random. At these distances, local environments range from subtropical to temperate grasslands and from southeastern pine forests to southwestern deserts. It is thus unsurprising that species turnover is quite large at these scales.

Discussion

The current results illustrate the applicability of wavelets to problems in pattern detection and analysis in biology and environmental sciences. The flexibility of

wavelet methods combined with their ability to represent local variation at different scales makes them ideal for probing complex, spatially and temporally varying patterns, such as patterns of species turnover in landscapes.

Despite the fact that scaling patterns can be powerfully characterized using wavelets (Arneodo, Grasseau & Holschneider, 1988; Muzy, Bacry & Arneodo, 1991, 1993; Ivanov *et al.*, 1996), we did not find strong evidence for a single scaling exponent for the biodiversity data analyzed here. We believe this is the result of the many complex interacting factors that influence community composition and species turnover in space. Our view is that quantification of scaling properties is most useful in cases where the pattern of dynamics may be governed by a small number of key variables. For example, Keitt *et al.* (2002) examined time-invariant scaling properties of breeding bird populations across North America. They hypothesized that scaling in the data was a consequence of a few simple factors: the size of a species' geographic range and the pattern of subdivision of the range into independently fluctuating subpopulations (see also Marquet *et al.*, this volume). Is such a result likely in community data? We feel it is unlikely. Compositional turnover in communities across space is likely to be determined by a large array of factors, including evolutionary history. Historical isolating barriers and zones of rapid change in the physical environment should generate spatially varying and scale-dependent patterns in diversity. Instead, it may be more fruitful to attempt to understand why beta diversity patterns are nonstationary and why some areas are turnover hot spots and why hot spots appear and disappear at different scales of analysis.

We see a number of interesting avenues for future studies. It would be quite interesting to incorporate environmental covariates into our analysis of community turnover. Keitt and Urban (2005) discuss wavelet transformation of ordinary linear regression models. One could easily add additional layers of information such as temperature, precipitation, etc. and then model the dependence of observed wavelet dissimilarity measures on local rates of environmental changes (also modeled with wavelets). On a more methodological level, we note that the smoothing functions used here (see Appendix) are equivalent to the intercept of a local linear regression (Silverman, 1986). Some bias reduction near boundaries might be achieved using higher-order local linear estimators (Hastie & Loader, 1993), or perhaps even semiparametric smoothers (Ruppert, Wand & Carroll, 2003) could be used to partial out influences of nuisance covariates.

Acknowledgments

The author acknowledges the kind invitation of the workshop organizers, insights and helpful comments from workshop participants, Keitt Lab members and several anonymous reviewers, and the generous support of the David and Lucile Packard Foundation.

Appendix 9.I

We implemented an adaptive "second generation" wavelet filter conceptually derived from Sweldens' "lifting scheme" (Sweldens, 1998). Second generation wavelets are built upon the idea that variation at a particular scale can be quantified by fitting increasingly smooth functions to input data and then examining changes in the lack-of-fit between successive smoothers. If the input signal contains a great deal of variation at a particular scale then a smooth function fit to the data at that scale will leave behind a large amount of residual variation. If the input signal is quite smooth at a particular scale, then the smoothing function will give a close fit and little residual variation will remain. The magnitude of the variation not explained by a smooth function fitted at a particular scale therefore estimates the intensity of pattern present in the data at that scale. Because second generation wavelets are built piecewise from locally smooth function, they have the advantage of adapting their shape near sampling gaps and boundaries. As a result, artificial boundary adjustments, such as periodic wrapping of the data, are not needed. Most previous applications have involved the discrete wavelet transform in which the scale of analysis jumps by powers of two at each level of the transform. For our purposes, the continuous wavelet transform is far more informative as the scale of analysis can take any value within the limits imposed by the grain and extent of the data.

Our construction of a second-generation continuous wavelet transform is based on the subtraction of local Gaussian-kernel regression estimators producing an adaptive "Difference-of-Gaussians" wavelet (Muraki, 1995). Let the local smoother centered at a,b take the form

$$\eta^{s}_{a,b}(x,y) = \frac{k\left(\frac{x-a}{s}, \frac{y-b}{s}\right)}{\sum\limits_{(u,v)\in\Omega} k\left(\frac{u-a}{s}, \frac{v-b}{s}\right)}, \tag{9.3}$$

where $\Omega = \{(u_1, v_1), (u_2, v_2), (u_3, v_3), \ldots (u_n, v_n)\}$ is the set of sampling points and $k(x,y) = e^{-\left(x^2+y^2\right)/2}$. We can then define the adaptive wavelet filter

$$\psi^{s}_{a,b}(x,y) = \eta^{s}_{a,b}(x,y) - \eta^{\beta s}_{a,b}(x,y), \tag{9.4}$$

where $\beta > 1$. Variance transmitted by the adaptive DoG wavelet at frequencies ω_x, ω_y is asymptotically given by

$$\left(T^{\text{Fourier}}\psi^{s}_{a,b}\right)\left(\omega_x, \omega_y\right) = e^{-s^2\left(\omega_x^2+\omega_y^2\right)/2} - e^{-\beta^2 s^2\left(\omega_x^2+\omega_y^2\right)/2}, \tag{9.5}$$

which has a maximum at

$$\omega_x^2 + \omega_y^2 = \frac{4\ln\beta}{s^2\left(\beta^2-1\right)}. \tag{9.6}$$

For convenience, we chose $\beta = 1.87$ such that the maximum frequency solution was reduced to $\omega_x^2 + \omega_y^2 = s^{-2}$. The result was that for any scaling $x, y \to (x, y)/s$, the dominant scale of analysis was s distance units.

We can then define the adaptive wavelet transform as

$$(T^{\text{wav}} f)(a, b, s) = \frac{1}{h_{a,b}(s)} \sum_{(u,v) \in \Omega} \psi^s_{a,b}(u, v) f(u, v), \tag{9.7}$$

where

$$h_{a,b}(s) = \sqrt{\sum_{(u,v) \in} \left[\psi^s_{a,b}(u, v)\right]^2} \tag{9.8}$$

ensures that wavelet variances are comparable across all locations and scales.

For statistical analysis, species counts were randomly assigned to route locations in repeated Monte Carlo trials. The wavelet transform and any subsequent statistics based on the wavelet coefficients were computed for each trial. Confidence regions were constructed to bound 95% of the Monte Carlo results falling closest to the median of the bootstrap ensemble. Observed values falling outside the confidence region were deemed statistically significant.

References

Antoine, J. P. (1999). The 2-D wavelet transform, physical applications and generalizations. In *Wavelets in Physics*, ed. J. C. Van den Berg, chapter 2, pp. 23–32. Cambridge: Cambridge University Press.

Arneodo, A., Grasseau, G. & Holschneider, M. (1988). Wavelet transform of multifractals. *Physical Review Letters*, **61**, 2281–2284.

Caley, M. J. & Schluter, D. (1997). The relationship between local and regional diversity. *Ecology*, **78**, 70–80.

Condit, R., Pitman, N., Leigh Jr., E. G., *et al.* (2002). Beta-diversity in tropical forest trees. *Science*, **295**, 666–669.

Dale, M. R. T. (1999). *Spatial Pattern Analysis in Plant Ecology*. Cambridge: Cambridge University Press.

Daubechies, I. (1992). *Ten Lectures On Wavelets*. CBMS-NSF Regional Conference Series in Applied Mathematics. Philadelphia: Society for Industrial and Applied Mathematics.

Hanski, I. & Gyllenberg, M. (1997). Uniting two general patterns in the distribution of species. *Science*, **275**, 397–400.

Harte, J., Kinzig, A. & Green, J. (1999). Self-similarity in the distribution and abundance of species. *Science*, **284**, 334–336.

Hastie, T. & Loader, C. (1993). Local regression: automatic kernel carpentry. *Statistical Science*, **8**, 120–129.

Ivanov, P. C., Rosenblum, M. G., Peng, C. K., *et al.* (1996). Scaling behavior of heartbeat intervals obtained by wavelet-based time-series analysis. *Nature*, **383**, 323–327.

Ives, A. R. (1995). Predicting the response of populations to environmental change. *Ecology*, **76**, 926–941.

Keitt, T. H. & Urban, D. L. (2005). Scale-specific inference using wavelets. *Ecology*, **86**, 2497–2504.

Keitt, T. H., Amaral, L. A. N., Buldyrev, S. V. & Stanley, H. E. (2002). Scaling in the growth of geographically subdivided populations: invariant patterns from a continent-wide

biological survey. *Philosophical Transactions of the Royal Society of London, Series B*, **357**, 627–633.

Klvana, I., Berteaux, D. & Cazelles, B. (2004). Porcupine feeding scars and climatic data show ecosystem effects of solar cycle. *American Naturalist*, **164**, 283–297.

MacArthur, R. H. & Wilson, E. O. (1967). *Island Biogeography*. Princeton: Princeton University Press.

Mallat, S. (1999). *A Wavelet Tour of Signal Processing*, 2nd edn. New York: Academic Press.

Muraki, S. (1995). Multiscale volume representation by a DOG wavelet. *IEEE Transactions on Visualization and Computer Graphics*, **1**, 109–116.

Muzy, J. F., Bacry, E. & Arneodo, A. (1991). Wavelets and multifractal formalism for singular signals: application to turbulence data. *Physical Review Letters*, **67**, 3515–3518.

Muzy, J. F., Bacry, E. & Arneodo, A. (1993). Multifractal formalism for fractal signals: the structure-function approach versus the wavelet-transform modulus-maxima method. *Physical Review E*, **47**, 875–884.

Peterjohn, B. G. & Sauer, J. R. (1993). North American breeding bird survey annual summary 1990–1991. *Bird Populations*, **1**, 1–15.

Pulliam, H. R. & Danielson, B. J. (1991). Sources, sinks, and habitat selection: a landscape perspective on population dynamics. *American Naturalist*, **137**, S50–S60.

Ruppert, D., Wand, M. P. & Carroll, R. J. (2003). *Semiparametric Regression*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge: Cambridge University Press.

Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall.

Slatkin, M. (1993). Isolation by distance in equilibrium and non-equilibrium populations. *Evolution*, **47**, 264–279.

Sweldens, W. (1998). The lifting scheme: a construction of second generation wavelets. *SIAM Journal on Mathematical Analysis*, **29**, 511–546.

Vetaas, O. R. (2002). Realized and potential climate niches: a comparison of four rhododendron tree species. *Journal of Biogeography*, **29**, 545–554.

Walker, J. S. (1999). *A Primer on Wavelets and Their Scientific Applications*. New York: Chapman and Hall.