



Chapter 13

Spatial and Spatiotemporal Patterns

13.1 Spatiotemporal Patterns

Spatial and spatiotemporal data analysis is of great importance in disease dynamics for a number of reasons such as looking for space-time clustering, hotspot detection, characterizing invasion waves, and quantifying spatial synchrony. Spatial synchrony—the level of correlation in outbreak dynamics at different locations—is of particular significance to acute immunizing infections, because asynchrony may permit regional persistence of infections despite local chains-of-transmission breaking during post-epidemic troughs (Keeling et al., 2004, see Sect. 15.7). Conversely, spatial synchrony can exacerbate the economic and public health burden because the resulting regionalized outbreaks can overwhelm logistical capabilities as was evident in the early part of the 2013–2014 West African Ebola outbreak and the 2020–2021 SARS-CoV-2 pandemic. Spatial statistics is also important in order to correct for the problem of spurious associations between incidence and environmental data because spatial autocorrelation violates the assumption of independence. This is further discussed in Sect. 18.2.

13.2 A Plant-Pathogen Case Study

Pathogenic fungi are generally not very important pathogens of mammals, though a virulent species of *Pseudogymnoascus* emerged in North America in 2007 to cause white-nose syndrome and exert major mortality events of bats (Blehert et al., 2009; Hoyt et al., 2021). In humans they cause ringworm and several opportunistic infections such as aspergillosis and candidiasis that are of minor importance except for in immunocompromised people. In various non-vertebrate animal case studies fungal pathogens have been shown to cause major epizoonoses. For example, *Aspergillus sydowii* has recently decimated Caribbean sea fan corals (Bruno et al.,

This chapter uses the following R package: `ncf`.

2011) and fungal infections frequently slaughter their way through insect populations (Hajek & St. Leger, 1994). Fungi are very common pathogens of plants on which non-systemic pathogens are often called rust. Systemic infections cause devastating disease like Dutch elm disease and chestnut blight. The latter completely altered the nature of North American hardwood forests when emerging during the first decade of the twentieth century (Anagnostakis, 1987).

While a bit idiosyncratic, the spatial dataset from Jennifer Koslow's experiment on a foliar, non-systemic rust fungus (*Coleosporium asterum*) that infects the flat-top goldenrod (*Euthamia graminifolia*) provides useful illustrations of various geo-statistical methods. The *euthamia* data present the severity of rust disease expression (`$score`, from 0 to 10) on host-plants planted within mesocosms (`$plot`) in an old field near Ithaca in New York State. The mesocosms were in a checkerboard grid with locations specified by coordinates `$xloc` and `$yloc`. Each mesocosm contained 3 focal *E. graminifolia* plants. The field also contained naturally occurring *E. graminifolia*, as well as several other hosts of the rust, most notably the Canada goldenrod (*Solidago canadensis*). Two different treatments, species composition (`$comp`, with three levels) and watering treatment (`$water`, with two levels), were applied to the mesocosms in a fully factorial design. Finally, to account for spatial variation the field were divided into four blocks with treatment combinations randomly assigned within each block.

For some of the analyses we need jittered coordinates because the three plants within each plot were not given separate coordinates. Figure 13.1 shows the spatial layout of the study. The vertical lines mark the predefined blocks.

```
data(euthamia)
euthamia$jx = jitter(euthamia$xloc)
euthamia$ jy = jitter(euthamia$yloc)
symbols(y = euthamia$xloc, x = euthamia$yloc,
        circles = euthamia$score, inches = 0.1,
        xlab = "y", ylab = "x")
abline(v = 47.5, col = 2)
abline(v = 97.5, col = 2)
abline(v = 147.5, col = 2)
```




Fig. 13.1: Rust scores from Keslow's experiment

13.3 Spatial Autocorrelation

Spatial statistics is a very rich field. This chapter focuses on a subset of methods that are commonly used in epidemiology involving the notion of spatial autocorrelation. Legendre (1993) is a great introduction to the use of spatial autocorrelation statistics in ecological studies in general. While all the methods discussed—such as Mantel tests, parametric and nonparametric correlation functions, local indicators of spatial association, etc.—come in canned packages (this chapter uses the `ncf` package), it is useful to spend a bit of time on the underlying ideas.

Many geostatistical methods to describe spatial pattern are focused on either spatial variance (Gary's C) or spatial correlation (Moran's I). This chapter largely focuses on the family of correlational methods. The regular (Pearson's) product-moment correlation (ρ) between two random variables, Z_1 and Z_2 , is defined as:

$$\rho_{12} = \frac{(Z_1 - \mu_1)(Z_2 - \mu_2)}{\sigma_1 \sigma_2},$$

where μ 's are expectations and σ 's are standard deviations.¹ The *autocorrelation* has exactly the same definition and is used when the Z 's are measurements of the same quantity (e.g., prevalence, incidence, presence/absence, etc.) at different spatial locations (or different times; Sect. 7.2).

The calculation needs to know (or have an estimate of) the values of the μ 's and σ 's. In the case of single snapshot spatial data the marginal mean and marginal standard deviation is normally used.² For the *euthamia rust* data (Fig. 13.1) these quantities are:

```
n = length(euthamia$score)
# marginal mean:
mu = mean(euthamia$score)
# marginal MLE sd:
sig = sd(euthamia$score) * (n - 1)/n
```

Using the `outer` function that provides all pairwise products of two vectors, the estimated autocorrelation matrix (`rho`) among all 360 plants is then:

```
# rescale Zs
zscale = (euthamia$score - mu)/sig
# autocorrelation matrix
rho = outer(zscale, zscale)
```

Note that while the individual pairwise values are not constrained to be between -1 and 1 , as correlations need to be, the various geostatistical methods discussed in the following involves manipulations of this matrix to normalize values. Most

¹ It is, again, unfortunate that these Greek symbols as used in statistics take a different meaning than their previous usage in epidemiology, but it cannot be helped since the study of epidemics leans on so many different fields of science.

² Note that the geostatistical methods usually use the maximum likelihood estimator of σ rather than the [best linear unbiased \(BLUE\) estimator](#); the denominator is n rather than $n - 1$.

of the methods also require some sort of associated spatial distance matrix. Most commonly used are the Euclidian distance for UTM coordinates or the [great-circle distance](#) for latitude/longitude coordinates. The Euclidean distance matrix among all 360 plants in the euthamia dataset is:

```
dst = as.matrix(dist(euthamia[, c("xloc", "yloc")]))
```




Fig. 13.2: Scatterplot of pairwise ρ against pairwise distance

To understand the different geostatistical methods, consider the plot of the paired autocorrelations as a function of their spatial distance (Fig. 13.2).

```
plot(dst, rho, cex = 0.1, ylab = expression("pairwise " * rho), xlab = "Pairwise distance (m)")
```

With this it is easy to erect a conceptual understanding of many different geostatistical methods.

- **Mantel test:** An overall test for whether there is any significant relationship between the elements in the two matrices. This is essentially a test for significant correlation between ρ and distance.
- **Correlogram:** The most classic tool of testing how autocorrelation depends on distance without assuming any particular function. Hack the distance x -axis into

segments (given by specifying some distance increment) and calculate the average ρ within each distance class.³

- **Parametric correlation functions:** Assume the relationship follows some parametric relationship—such as Exponential, Gaussian, or Spherical functions—and do the appropriate nonlinear regression of ρ on distance. Section 18.2 provides an example of such fitting via the `lme` function of the `nlme` library.
- **Nonparametric correlation function:** Fit a nonparametric regression (usually a smoothing spline or a kernel smoother) to the relationship (Hall & Patil, 1994). This also goes by the name of the spline correlogram (Bjørnstad & Falck, 2001).
- **Local indicators of spatial association (LISA):** A test for hotspots (Anselin, 1995) specifying a neighborhood size and for each location calculates the average ρ with all the other locations that belongs to its neighborhood to find areas of significant above average values.

There are a bunch of other named methods that are variations of these. Several of which are extensions to when there is multiple observations at each location (such as a spatial panel of time series), in which case it is natural to estimate the autocorrelation matrix using the regular correlation matrix. The modified correlogram of Koenig (1999) is the multivariate extension of the correlogram (see also Bjørnstad et al., 1999b). The time-lagged spatial cross-correlation function has been used to study waves of spread (see below and Sect. 12.8). Various other versions allow the spatial correlation function to vary by cardinal direction (so-called anisotropic correlograms) to investigate directional patterns (Bjørnstad et al., 2002b).

13.4 Testing and Confidence Intervals

An important reason why specialized methods are needed for these analyses, despite most being conceptually simple, is because while the n original data points may (or may not) be statistically independent the n^2 numbers in the autocorrelation matrix is obviously very statistically non-independent and the interdependence is very intricate (as nicely discussed and visualized by Rousseeuw and Molenberghs, 1994). None of the usual ways of testing for significance or generating confidence intervals is therefore applicable. Testing is usually done using permutation tests under the null hypothesis of no spatial patterns. The correlogram (or Mantel test, or ...) of the real data should look no different than that of a random re-allocation of observations to spatial coordinates if the null hypothesis is true. Statistical significance is calculated by comparing the observed estimate to the distribution of estimates for, say, 999 different randomized datasets.⁴ If the observed is more extreme than 950 (990) of the randomized data we conclude that there is significant deviation from

³ The variogram is similar to the correlogram but instead of using the autocorrelation similarity measure it uses the semivariance dissimilarity measure: $(Z_i - Z_j)^2/2$.

⁴ This produces a total of 1000 known possible outcomes; the 999 we randomly generated plus the one nature provided.

spatial randomness at a nominal 5%-level (1%-level). For some of the methods it is possible to generate confidence envelopes using bootstrapping (resampling with replacement; Bjørnstad and Falck, 2001). All the above methods are available in the `ncf` package.

```
require(ncf)
```

13.5 Mantel test

We continue using the `euthamia` data as a case study:

```
test1 = mantel.test(M1 = rho, M2 = dst, quiet = TRUE)
test1

## $correlation
## [1] -0.04603662
##
## $p
## [1] 0.000999001
##
## $call
## [1] "mantel.test(M1 = rho, M2 = dst, quiet = TRUE) "
##
## attr(,"class")
## [1] "Mantel"
```

There is a significant negative association between similarity and distance showing that the rust data are not spatially random. The Mantel test is a crude tool but it does reveal that locations near each other tend to be more similar in disease status than those separated by a greater distance. If instead of having two matrixes have spatial coordinates and observations the syntax is:

```
mantel.test(x, y, z, latlon = FALSE)
```

In this case coordinates can either be Euclidian or latitude/longitude if `latlon = TRUE`.

13.6 Correlograms

The correlogram shows how the autocorrelation is a function of distance (Fig. 13.3). The shape of the correlogram can indicate random, diffusive, or clinal patterns. Random patterns show up as a flat non-significant correlogram, diffusive patterns will have significantly positive values at short distances that tapers off to zero, and gradient patterns will have significantly positive values at short distances and signifi-

cantly negative values at long distances.⁵ Legendre and Fortin (1989) provide visual probes for patterns using various characteristics of the correlogram. The illustration using the euthamia data is:

```
test2 = correlog(x = euthamia$xloc, y = euthamia$yloc,
  z = euthamia$score, increment = 10)
plot(test2)
```

The first distance class is significantly positive and the estimated distance to which the local positive value decays to zero (the x-intercept) is 44 meters indicative of significant local similarity. There is further evidence of significantly negative autocorrelation at long distances suggestive of a gradient across the field (Fig. 13.3).

13.7 Nonparametric Spatial Correlation Functions

Finer resolution and confidence intervals can be found using the nonparametric spatial covariance function (Hall & Patil, 1994; Bjørnstad & Falck, 2001):

```
test3 = spline.correlog(x = euthamia$xloc, y = euthamia$yloc,
  z = euthamia$score, quiet = TRUE)

summary(test3)$estimate

##           x           e           y
## estimate 36.53666  5.981457  0.5824953

summary(test3)$quantiles

##           x           e           y
## 0       17.82638  0.005531342 -0.003692403
## 0.025   27.34768  0.409260018  0.126207964
## 0.25    32.96260  3.119520264  0.295493112
## 0.5     36.03922  5.992822095  0.391821835
## 0.75    39.95760  8.451637015  0.517451204
## 0.975   44.19246  12.655055265  0.778859891
## 1       59.30986  15.005569404  1.200445017
```

The `spline.correlogram` returns a bunch of stuff; in fact all the summary statistics I thought might be of relevance in some previous spatial analyses (Bjørnstad & Falck, 2001). These are:

- `estimates`: A vector of benchmark statistics.

⁵ Inhibitory processes such as the Janzen-Connell effect discussed in Sect. 12.4 will produce significantly negative values at short distances that tapers off.

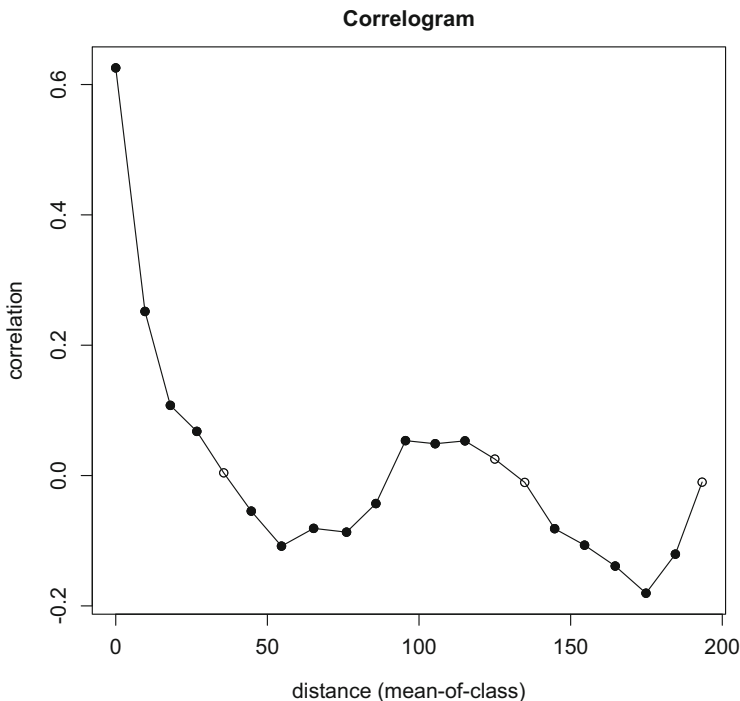


Fig. 13.3: The spatial correlogram for the *euthamia* rust data. Values that significantly deviates from that expected under the null hypothesis of complete spatial randomness are represented by filled black circles

- x : The lowest value at which the function is $= 0$.⁶
- e : The lowest value at which the function is $= 1/e$ (i.e., the spatial scale parameter in the presence of exponential or Gaussian spatial correlation; recall Sect. 12.2).
- y : The extrapolated value at $x = 0$.
- `quantiles`: A matrix summarizing the quantiles in the bootstrap distributions of the benchmark statistics. The 2.5- and 97.5-percentiles represent the 95% confidence interval.

```
plot(test3)
```

⁶ If correlation is initially negative, the distance calculated appears as a negative measure. This may seem a little strange, but some locally inhibitory processes predict significant negative local auto- or cross-correlation (e.g., Seabloom et al., 2005).




Fig. 13.4: The spline correlogram of the *euthamia rust* data. The grey polygon represents the 95% bootstrap confidence envelope

Figure 13.4 shows the estimated correlation function with its bootstrap 95% confidence envelope. The confidence envelope allows comparisons of correlation functions for different datasets to look for significant differences (Bjørnstad et al., 1999a).

13.8 LISA

The previous methods average across all locations to study how similarity depends on distance. Local indicators of spatial association (Anselin, 1995) quantify how similar observations are within neighborhoods of each observation. This can be used to test for significant spatial hot/cold-spots of disease (Fig. 13.5). For this we have to define the radius of the neighborhood. Spatial dependence in the *euthamia* data decay to zero at around 40m (Fig. 13.4), so we use 20 meters.

```
test4 = lisa(x = euthamia$yloc, y = euthamia$xloc,
            z = euthamia$score, neigh = 20)
```

```
plot(test4)
```




Fig. 13.5: LISA analysis of Koslow's rust data (with a 20m neighborhood). Filled red circles are significant spatial hotspots. Squares are cold-spots. The size of the symbols reflects how much the disease-score deviates from the mean

13.9 Cross-Correlations

Janis Antonovics and his colleagues have done roadside surveys of antler smut disease counting number of healthy and diseased wild campions (*Silene alba*) at the Mountain Lake Biological field station for more than 20 years (Antonovics, 2004). The `silene` data contains the mean number of healthy `$hmean` and diseased `$dmean` plants for each road segment, as well as latitude `$lat` and longitude `$lon` (Fig. 13.6).

```
data(silene)
symbols(silene$lon, silene$lat, circles = sqrt(silene$dmean),
        inches = 0.2, xlab = "Longitude", ylab = "Latitude")
```




Fig. 13.6: Burden of antler smut on wild campion at the Mountain Lake field station (Antonovics, 2004)

Most geostatistical methods can be extended to consider spatial *cross*-correlation between different variables. As an example we can use the `silene` dataset to investigate if prevalence is spatially cross-correlated with abundance using the spline cross-correlogram (Fig. 13.7).

```
silene$ab = silene$dmean + silene$hmean
silene$prev = silene$dmean/(silene$dmean + silene$hmean)
```

The square-root transform of the abundance measure helps normalize the variance of the count data. There is significant positive cross-correlation within a 1 km range (95% CI: {0.6, 2.9} km) meaning that where the host tends to be abundant, the pathogen tends to be prevalent.

```
testcc = spline.correlog(x = silene$lon, y = silene$lat,
  z = silene$prev, w = sqrt(silene$ab), latlon = TRUE,
  na.rm = TRUE)
plot(testcc)
```

We can use a spatial cross-correlogram (using 25m distance increments) to study if presence/absence of rust is spatiotemporally cross-correlated between 1994 and 1995 in the `filipendula` dataset discussed in Sect. 12.3.




Fig. 13.7: Spatial cross-correlation of prevalence and abundance in the `silene` dataset

```
data(filip)
testcc2 = correlog(x = filip$X, y = filip$Y, z = filip$y94,
  w = filip$y95, increment = 25)
```

The local inter-year correlation (`corr0`) is 0.75 and the first cross-correlation is significantly positive with a cross-correlogram x-intercept of 148m.⁷

```
testcc2$corr0

## [1] 0.7651124

testcc2$x.intercept

## (Intercept)
##      148.939
```

Locations heavily affected in 1994 were thus also heavily affected in 1995 testifying to the importance of local contagion and/or habitat heterogeneity in infection risk. This is an example of a time-lagged cross-correlogram (Bjørnstad et al., 2002b).

13.10 Gypsy Moth

The gypsy moth (Sect. 12.5) was introduced to the northeastern USA in the late 1860s and has spread at a rate of 10–20 km per year since. The larvae eats leaves of a wide range of trees and shrubs and reaches outbreak densities usually around every 10 years. The outbreaks end through epizootics of the *Lymantria dispar* nuclear polyhedrosis virus and more recently the entomopathogenic fungus *Entomophaga maimaiga* that together kills virtually all larvae following outbreaks. Bjørnstad et al. (2010) used the nonparametric spatial covariance function to study the spatiotemporal patterns in these outbreaks. The `gm` dataset contains UTM coordinates and fraction of forests defoliated each year between 1975 and 2002 in 20×20 km grid cells across northeast USA. The following characterize the patterns of synchrony and time-lagged cross-correlation in the outbreak time series:

```
data(gypsymoth)
sel = apply(gypsymoth$defoliation[,2:28], 1, sum)!=0
#Synchrony:
fit1 = Sncf(gypsymoth$xy[sel, 1]/1000, gypsymoth$xy[sel, 2]/1000,
  gypsymoth$defoliation[sel, ], resamp = 500)
#Lag 1 cross-correlation
fit2 = Sncf(gypsymoth$xy[sel, 1]/1000, gypsymoth$xy[sel, 2]/1000,
  z = gypsymoth$defoliation[sel, 1:27],
  w = gypsymoth$defoliation[sel, 2:28], resamp = 500)
#Lag 2 cross-correlation
fit3 = Sncf(gypsymoth$xy[sel, 1]/1000, gypsymoth$xy[sel, 2]/1000,
  z = gypsymoth$defoliation[sel, 1:26],
  w = gypsymoth$defoliation[sel, 3:28], resamp = 500)
```

⁷ The spline cross-correlogram would give bootstrap confidence intervals on these quantities.

The outbreaks are highly synchronized out to 200 km, with a regional average outbreak correlation of around 0.2. The time-lagged cross-correlogram show significant local cross-correlation at the 1-year lag but not 2-year lag, indicating that outbreaks tend to persist spatially for 2 years before collapsing (Fig. 13.8):

```
par(mfrow = c(1, 3))
plot(fit1, ylim = c(-0.1, 1))
title("Lag 0")
plot(fit2, ylim = c(-0.1, 1))
title("Lag 1")
plot(fit3, ylim = c(-0.1, 1))
title("Lag 2")
```

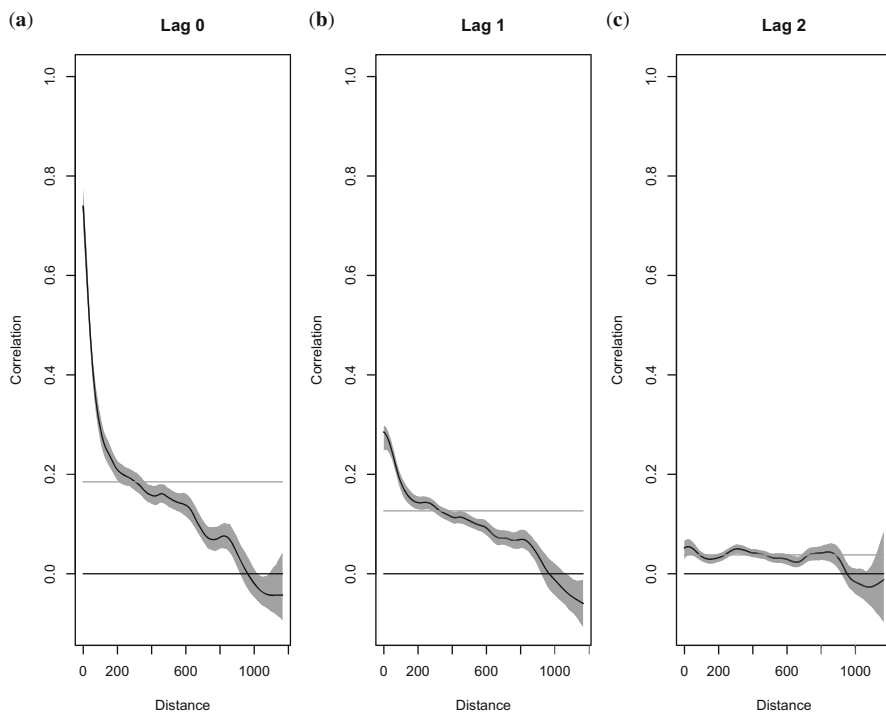


Fig. 13.8: The (a) nonparametric spatial covariance function, (b) lag one and (c) lag two cross-correlation function of gypsy moth outbreak data from northeastern USA between 1975 and 2002