Tutorial: Spatial Point Pattern Analysis

D G Rossiter Cornell University

March 27, 2024

Contents

1	Examining some point patterns						
2	First-order properties: the G function						
3	Kernel density estimation						
4	Second-order properties: the K function 4.1 The L function: a linearized K function 4.2 * Modifying the window	16 18 19					
5	The F function; non-rectangular windows 21						
6	Marked point patterns6.1Categorical marks6.2Interaction between point patterns: the cross-K function6.3Combining point patterns6.4Continuous marks	26 28 30 32					
7	Models of spatial processes7.1Null model7.2Trend surface7.3Strauss process7.4Covariates	34 36 36 39 41					
8	Spatial prediction	45					

Version 6.0 Copyright © 2012–7, 2020–4 Cornell University All rights reserved. Reproduction and dissemination of the work as a whole (not parts) freely permitted if this original copyright notice is included. Sale or placement on a web site where payment must be made to access this document is strictly prohibited. To adapt or translate please contact the author (http://www.css.cornell.edu/faculty/dgr2/).

9 Spatio-temporal analysis	49	
10 Further reading	59	
11 Answers	59	
A Preparing data for point pattern analysisA.1 Shapefiles and other common geospatial formatsA.2 Text files	64 64 66	
References		
Index of R concepts		

冰冻三尺,非一日之寒 "One day of cold weather is not enough to freeze ice three feet thick!" – Chinese proverb

This tutorial gives an overview of **spatial point-pattern analysis**, and some practical experience with such analysis using the R Environment for Statistical Computing.

Spatial point-pattern analysis considers the distribution of one or more sets of **points** in some **bounded region** as possibly being **the result of some stochastic process** which produces a **finite number of "events" or "occurrences"**. Examples are a forest plot with the locations of individual trees, a microscope slide with the locations of individual cell centres, and a municipal boundary with point pollution sources. These may be viewed as pure point-patterns (just the locations) but sometimes attributes are included in the analysis; for example, the size or species of a tree in the forest plot.

After completing this tutorial you should be able to:

- 1. Display and examine spatial point patterns (§1);
- 2. Quantify deviations from **Complete Spatial Randomness** with the *G* function (§2);
- 3. Analyze single spatial point patterns with the F, G, K and L functions;
- 4. Compute a kernel density of an inhomogeneous point pattern (§3);
- 5. Analyze interactions between two point patterns with the *K* and *L* functions (§4);
- 6. Analyze **marked** point patterns (§6), i.e., where there is some attribute recorded at each point;
- 7. Model spatial point patterns as realizations of a **spatial data generating process** (sDGP) including both trend and interaction components, and evaluate model success (§7);
- 8. Predict over an areas based on a fitted model $(\S8)$;
- Do some simple spatio-temporal analysis of an evolving point pattern (§9).

In addition there is an Appendix (A) on how to prepare a dataset for pointpattern analysis.

The theory behind point-pattern analysis is comprehensively presented in the text of Diggle [6]. An accessible and less technical introduction is Boots and Getis [3]. Bivand et al. [2, Ch. 7] present worked examples in the context of R processing. Here we work through some of the main ideas only. Some of the code here is adapted from that chapter. Illian et al. [7] present a computational framework for fitting complex spatial point process models using a recently-developed methodology known as INLA.

Point-pattern analysis is based on theories of point processes. A modern review article is by Møller and Waagepetersen [8].

We will consider two kinds of properties of point patterns:

- **First-order**: considering points as **individuals**, no interaction, over the whole region. An example is the **spatial density**, which is taken as the indication of the **intensity** of the process that gave rise to the PPA;
- Second-order: considering interactions between marked sets of points, e.g., their tendency to cluster or repel. An example is bird nests of two different species in the same area.

In addition, we will attempt (§7) to model the presumed spatial data generating process.

Note: The code in these exercises was tested with Knitr [12] on R version 4.2.3 (2023-03-15), sf package 1.0.15, splancs package 2.1.44, stpp package 2.0.7, terra package 1.7.71, and spatstat package 3.0.7 running on Mac OS X Sonoma 14.3.1. The text and graphical output you see here was automatically generated and incorporated into IATEX by running the code through R and its packages. Then the IATEX document was compiled into the PDF version you are now reading. Your output may be slightly different on different versions and on different platforms.

1 Examining some point patterns

Supplementary reading:

• Bivand et al. [2, §7.2]: R packages relevant for spatial point-pattern analysis.

We use the same examples as Bivand et al. [2, Ch. 7]: location of cell centres in a microscope slide ("Cells"); locations of Japanese black pine saplings ("Japanese"); and locations of saplings of California redwood trees ("Redwood").

Task 1 : Load the Japanese pines example dataset japanesepines and summarize it.

These are examples in the **spatstat** package and are provided in a suitable format, namely as objects of R class **ppp**, a "planar point pattern".

require(spatstat) data(japanesepines) class(japanesepines) ## [1] "ppp" str(japanesepines) ## List of 5 :List of 4 ## \$ window ..\$ type : chr "rectangle" ## ..\$ xrange: num [1:2] 0 1 ## ## ..\$ yrange: num [1:2] 0 1 ## ..\$ units :List of 3\$ singular : chr "metre" ## ##\$ plural : chr "metres"\$ multiplier: num 5.7 ##- attr(*, "class")= chr "unitname" ## ## ..- attr(*, "class")= chr "owin"

```
## $ n
               : int 65
## $ x
               : num [1:65] 0.09 0.29 0.38 0.39 0.48 0.59 0.65 0.67 0.73 0.79 ...
## $ y
              : num [1:65] 0.09 0.02 0.03 0.18 0.03 0.02 0.16 0.13 0.13 0.03 ...
## $ markformat: chr "none"
   - attr(*, "class")= chr "ppp"
##
summary(japanesepines)
## Planar point pattern: 65 points
## Average intensity 65 points per square unit (one unit = 5.7 metres)
##
## Coordinates are given to 2 decimal places
## i.e. rounded to the nearest multiple of 0.01 units
## (one unit = 5.7 metres)
##
## Window: rectangle = [0, 1] x [0, 1] units
## Window area = 1 square unit
## Unit of length: 5.7 metres
```

Note: Notice that the ppp class has a structure that, according to the authors of the spatstat package, facilitates point-pattern analysis. Of course it has the coordinates (fields x and y) and the number of points (field n), but it also defines a window (field window) as a list of four characteristics: the shape, the limiting coordinates, and the units of measure. This field is of class owin.

Q1 : How many trees are represented by this point pattern? Jump to A1
•

Q2: What is the area covered by this point pattern? Jump to $A2 \bullet$

Task 2 : Plot the locations of the trees.

The generic plot method specializes to plot.ppp for an object of class ppp: plot.ppp(japanesepines, main = "Locations of Japanese pine trees", axes = T) grid()

Locations of Japanese pine trees



Q3 : Does this pattern look completely random, clustered, or regular? $\begin{array}{c} Jump \ to \ A3 \bullet \end{array}$

 Task 3 : Load the other two example datasets redwoodfull and cells;

 view the spatial distribution of all three on one graph.

 data(redwoodfull)

 data(cells)

```
par(mfrow = c(2, 2))
plot.ppp(japanesepines, main = "Japanese Pines", axes = T)
plot.ppp(redwoodfull, main = "Redwoods", axes = T)
plot.ppp(cells, main = "Biological cells", axes = T)
par(mfrow = c(1, 1))
```



 $\mathbf{Q4}$: Are there differences in the point patterns? Which pattern(s) look completely random, clustered, or regular? Jump to $A4 \bullet$

2 First-order properties: the G function

Supplementary reading:

• Bivand et al. [2, §7.3]: Preliminary Analysis of a Point Pattern

A suitable null hypothesis for the locations of a set of points is that there is no pattern, i.e., the points are distributed at random. This is known as **Complete Spatial Randomness** (abbreviation CSR). Assuming CSR we can compute several expected distributions of the points:

• The **G** function: the distribution of the distances from an arbitrary observed **point** to its nearest neighbour; expressed as the **cumulative distribution function** G(r) of the proportion of points that have at least one neighbour within a distance r. Formally:

$$d_i = \min_j \{ d_{ij}, \forall j \neq i \in S \}, i = 1, \dots, n$$

$$\tag{1}$$

$$\widehat{G}(r) = \frac{\{\#d_i : d_i \le r, \forall i\}}{n}$$
(2)

The **F** function: the distribution of the distances from an arbitrary **lo**-• **cation** in the plane to its nearest observation; this is sometimes called the empty space function: it measures the average empty space between observed points. This is examined in §5, below.

A homogeneous Poisson process (HPP) is a process on the "landscape" by which points (occurrences, events ...) are produced at specific locations, and in which the points are independently and uniformly distributed over a given region.

hus the location of one point does *not* affect the location of other points; another point can be anywhere. There are no clusters and no "empty" subregions, except by chance.

A homogeneous Poisson process with known **intensity** (a first-order property), conventionally called λ , will produce a CSR pattern, with the statistical properties:

$$G(r) = F(r) = 1 - \exp\{-\lambda \pi r^2\}$$
(3)

rs

where r is the distance, and λ is the mean number of points per unit area, in the given distance units; it is also called the **intensity** of the process.

In this section we examine the G function, also known as the **nearest**neighbour-distance distribution function. It is easily interpretable as the distance one must travel from any observation, to find at least one other observation. It is "short-sighted", since it only considers the nearest neighbour, and so gives no information about behaviour at long distances. It is a **point-related** function, i.e., computed from each point in the pattern. As such it gives no information about empty space; for that see the F function $(\S5)$, which is a **location-related** function

Task 4 : Compute and display the empirical vs. theoretical *G* function for the Japanese pines.

The Gest function of the spatstat.explore package computes this function on an object of class ppp, so we use the japanesepines object originally loaded as the sample **ppp** object.

```
G <- Gest(japanesepines)</pre>
class(G)
## [1] "fv"
                 "data.frame"
summarv(G)
##
                       theo
                                      han
        r
## Min. :0.00000 Min. :0.0000
                                  Min. :0.0000
                                                 Min. :0.0000
## 1st Qu.:0.05936
                  1st Qu.:0.5130
                                  1st Qu.:0.5187
                                                 1st Qu.:0.4583
## Median :0.11872 Median :0.9438 Median :0.9704
                                                 Median :0.9697
  Mean :0.11872
                                  Mean :0.7354
##
                  Mean :0.7383
                                                 Mean :0.7218
   3rd Qu.:0.17808
                   3rd Qu.:0.9985
                                  3rd Qu.:1.0000
                                                  3rd Qu.:1.0000
##
## Max. :0.23744 Max. :1.0000 Max. :1.0000
                                                       :1.0000
                                                 Max.
##
        km
                    hazard
                                     theohaz
                  Min. : 0.00
   Min. :0.0000
##
                                  Min. : 0.00
## 1st Qu.:0.4785 1st Qu.: 0.00
                                  1st Qu.:24.24
## Median :0.9580 Median : 0.00
                                  Median :48.49
##
   Mean :0.7227 Mean : 13.33
                                  Mean
                                       :48.49
```

##	3rd (u.:1.0000	3rd Qu.	: 0.00	3rd Qu.	:72.73
##	Max.	:1.0000	Max.	:1494.63	Max.	:96.97

The returned object has fields for the distance (\mathbf{r}) , the theoretical value of G(r) (theo), and four variants of an empirical estimate of G. These differ in how edge effects are accounted for; see ?Gest for an explanation of the options.

Note: The "edge effect" occurs because any points produced by the spatial process but outside the window can not be observed. If not accounted for there is bias, because points near the edge may well have a nearer neighbour outside the window, but since that is not observed, we can not compute the distance to it. So, we presume there are such unobserved points, and they are "similarly" distributed as the ones we observe; statisticians have proposed various ways to implement this "similarity".

The plot method specializes to plot.fv, to handle the returned object of class fv; this is just a convenient structure for this sort of plot. plot(G, main = "G-function, Japanese pines")



This graph can be interpreted as follows:

- *r*-axis: the distance away from an arbitrary point;
- G(r)-axis: the G function, namely, he proportion of points that have at least one neighbour within a distance r.

The default graph does not quite show the whole empirical function, i.e., where G(r) = 1, so we re-draw and specify the radius limits explicitly: plot(G, xlim = c(0, 0.16), main = "G-function, Japanese pines")

G-function, Japanese pines



Note that at 0.16 the expected value of G(r) is: 1 - exp(-japanesepines\$n * pi * (0.16)^2)

[1] 0.9946337

This is of course an asymptotic function, so never reaches 1; the empirical function does.

Q5:

(a) What is the distance across the unit square at which at least 95% of the points have at least one neighbour within that distance, according to the "reduced sample" (**rs**), also called "border", method of edge correction?

(b) What proportion of points have at least one neighbour within 0.05 units? $Jump \ to \ A5 \ \bullet$

To answer this questions, we can make a visual estimate from the graph, or look inside the object; field **r** is the radius and **rs** is the reduced-sample estimate of G(r). The selection function which picks out the entries meeting the required condition, and the min "minimum" function finds the first one in the list. For example, we can find the radius at which at least 95% of the points have at least one point within that threshold $gr[min(which(G\rs \ge 0.95))]$

```
## [1] 0.116867
G$rs[min(which(G$r >= 0.05))]
## [1] 0.3877551
```

Q6: How closely do the four variants of the empirical *G* function match the theoretical *G* function for this intensity? Note this last is marked $G_{\text{pois}}(r)$

on the graph; the estimates are marked $(\hat{G})_{\text{method}}(r)$. Jump to $A6 \bullet$

Another way to look at this is as expected vs. actual, which should be a 1:1 relation.

Task 5 : Plot the actual reduced-sample estimate against the theoretical value of G(r).

```
## 1 : example of random pattern
plot(G, cbind(rs,theo) ~ theo,
    main="G-function, Japanese pines")
```





Here both axes are values of G(r); the radius r is not shown.

Q7: How closely does the empirical match the theoretical? Jump to $A7 \bullet$

Task 6 : Compute the *G* function and plot it for the other two datasets. \bullet

First the Redwood trees:

```
## 2 : example of clustered pattern
G.clust <- Gest(redwoodfull)
par(mfrow=c(1,2))
plot(G.clust, cbind(rs,theo) ~ r,
    main="G-function, Redwood trees")
plot(G.clust, cbind(rs,theo) ~ theo,
    main="G-function, Redwood trees")
par(mfrow=c(1,1))</pre>
```



 $\mathbf{Q8}$: Describe and interpret the differences between this G function and that for the Japanese pines. Jump to $A8 \bullet$

Now the cells:

```
## 3 : example of dispersed / regular pattern
G.disp <- Gest(cells)
par(mfrow=c(1,2))
plot(G.disp, cbind(rs,theo) ~ r,
    main="G-function, cells")
plot(G.disp, cbind(rs,theo) ~ theo,
    main="G-function, cells")
par(mfrow=c(1,1))</pre>
```



Q9 : Describe and interpret the differences between this G function and that for the Japanese pines. $Jump \text{ to } A9 \bullet$

So far we have formed subjective opinions about which patterns are consistent with CSR. There is no formal test; the way this is assessed is to form an **envelope** of how the empirical G function would look, under the null hypothesis of CSR and with the observed homogeneous intensity. This envelope is computed by a large number of **simulations** of the Poisson point process

that gives rise to CSR. Each realization will be a bit different, because it's a discrete simulation, not the theoretical curve. Then the envelope is plotted along with the actual empirical G function, and we can see if it is contained inside, and if not, at what points it diverges.

Q10 : Would simulations be more or less variable as the intensity of the process increases? Jump to $A10 \bullet$

Task 7 : Compute an envelope for the G function of the Japanese pines, and plot it along with the observed empirical G function.

The envelope function of the spatstat.explore package computes this for an object of class ppp. Arguments include the object, the function name (here, Gest), the radii at which to compute point-wise envelopes, the rank of the envelope value among the simulated values, and the number of simulations.

Note: The **nrank** argument controls the width of the envelope. A low value, such as the default **nrank=1**, uses the extreme values (minimum and maximum) of the simulated distributions as the envelope; this is the widest and most conservative with respect to rejecting the null hypothesis of CSR. Using a higher rank corresponds roughly to setting the one-sided α for a t-test. In this case we have 99 simulations; so using **nrank=2** is excluding the single maximum and minimum at each point, thereby narrowing the envelope. This is roughly equivalent to one-sided $\alpha = (1 - (2/99)) = 0.9798$, i.e., two-sided $\alpha \approx 96\%$ confidence level considering both minimum and maximum.

We try to simulate out to the radius where the empirical G(r) = 1. Also, we use the **set.seed** function to initialize the random number generator, so your results match ours – you would not do this in an actual analysis, you would want the randomness. If we did not specify a large enough radius in the previous step, the *G* function may not reach 1, so we specify a somewhat smaller ending radius.

Note: The argument to set.seed is arbitrary, it has no meaning.

Japanese pines, G-function envelope



Q11 : Is the empirical G function within the envelope throughout? Can we reject the null hypothesis of CSR? $Jump \text{ to } A11 \bullet$

Task 8 :Compute and plot the G function envelopes for the other two
datasets.

The maximum radius at which a point is encountered varies considerably; here we show each one with its own maximum:



Q12 : Describe and interpret the differences between the envelopes for theseG functions and that for the Japanese pines.Jump to A12 •

3 Kernel density estimation

Supplementary reading:

• Bivand et al. [2, §7.4.1–3]: Statistical Analysis of Spatial Point Processes

An important question is whether the stochastic point-process is **homogeneous** (the same across the whole area) or **inhomogeneous**. Equivalently, is the **intensity** of the process the same everywhere? If not, we assume an **inhomogeneous Poisson process** (IPP), so that a single intensity λ describing the Poisson process is replaced by a spatial function $\lambda(x)$, where x is the spatial position. The question then arises: what is the spatial function of the density?

Note: The assumption of homogeneity is often not realistic, from what we know about the process. For example, environmental factors favour some sub-areas over others for occurrence of a given tree species¹, so that the point-process by which a species is located can not be assumed to be homogeneous. However, the other assumption of the Poisson process still holds: given an intensity, events are independent and uniformly distributed.

Naturally, this depends on the scale at which we examine it. At broad scales, all points are taken together and the process is by definition homogeneous; at fine scales, very small neighbourhoods are considered and random fluctuations can lead to different intensity estimates. In other words, a small bandwidth will lead to a "spiky" map, whereas a large bandwidth will lead to a smooth map – which one best represents the (in)homogeneity of the point process? This is the **bandwidth** problem.

 $^{^{1}}$ E.g., the eastern hemlock (*Tsuga canadensis*) grows by preference in moist, shallow soils on hillsides, whereas beech (*Fagus* spp.) prefers well-drained hilltop positions

The next question is the shape of the **smoothing kernel**. This is a function of the two coordinates, providing a smooth 2D surface with highest probability of finding a point at the centre and smoothly decreasing probability away from it. The default used by **density.ppp** is the **Gaussian** (normal distribution) kernel, but many applications prefer the **quartic** kernel:

$$\kappa(u) = \begin{cases} \frac{3}{\pi} (1 - ||u||^2)^2 & \text{if } u \in (-1, 1) \\ 0 & \text{otherwise} \end{cases}$$
(4)

where $||u||^2 = u_1^2 + u_2^2$, i.e., the squared norm, centred on the point to be estimated. Thus there is an inverse-square decrease in density outward from a point, to zero outside the unit circle; the "unit" is set by the bandwidth h, hence its importance:

$$u = ||x - x_i||/h \tag{5}$$

Although the kernel density is conceptually spatially continuous, in practice it is computed at many points over a fine grid and displayed as a raster "image".

Task 9 : Plot the kernel density of the Redwoods dataset, at various bandwidths.

The density.ppp function of the spatstat.explore package computes the density surface directly from the point pattern. The sigma argument specifies a function to compute an "optimum" bandwidth. We compare three of these functions, and pick the "Diggle" for subsequent analysis: print(bw.o <- bw.diggle(redwoodfull))

```
## sigma
## 0.01981409
print(bw.CvL(redwoodfull))
## sigma
## 0.09340959
print(bw.ppl(redwoodfull))
## sigma
## 0.0392317
```

Quite some difference.

Now make plots with the "optimum" bandwidth and some multiples of it. The density.ppp of the spatstat.explore package computes a density surface, of class im. To plot im objects, the plot generic method specializes to the plot.im function of the spatstat.geom package. Contour lines can be added with the contour function.

```
d1 <- density.ppp(redwoodfull, sigma = bw.o, kernel = "quartic")
class(d1)
## [1] "im"
d05 <- density.ppp(redwoodfull, sigma = bw.o, kernel = "quartic", adjust = 0.5)
d4 <- density.ppp(redwoodfull, sigma = bw.o, kernel = "quartic", adjust = 4)</pre>
```

```
d2 <- density.ppp(redwoodfull, sigma = bw.o, kernel = "quartic", adjust = 2)
```

```
par(mfrow = c(2, 2))
plot.im(d05, main = paste("Bandwidth=", round(bw.o * 0.4, 4), " (optimum*.5)",
    sep = "", collapse = ""))
contour(d05, add = TRUE)
plot(d1, main = paste("Bandwidth=", round(bw.o, 4), " (optimum)", sep = "",
    collapse = ""))
contour(d1, add = TRUE)
plot(d2, main = paste("Bandwidth=", round(bw.o * 1.5, 4), " (optimum*1.5)",
    sep = "", collapse = ""))
contour(d2, add = TRUE)
plot(d4, main = paste("Bandwidth=", round(bw.o * 2, 4), " (optimum*2)",
    sep = "", collapse = ""))
contour(d4, add = TRUE)
plot(d4, main = paste("Bandwidth=", round(bw.o * 2, 4), " (optimum*2)",
    sep = "", collapse = ""))
contour(d1, add = TRUE)
plot(d4, main = paste("Bandwidth=", round(bw.o * 2, 4), " (optimum*2)",
    sep = "", collapse = ""))
contour(d1, add = TRUE)
par(mfrow = c(1, 1))
```





Q13 : What happens to the kernel density (intensity) statistics as the bandwidth increases? (Note the numbers in the four scale legends.) Jump to A13 \bullet

Q14: Describe the trend in kernel density as the bandwidth increases. Does the "optimum" found by the Diggle algorithm seems to give the "best" view of the varying intensity of the point process that is assumed to be causing the Redwood point pattern? (Hint: compare with the point pattern). Jump

to A14 \bullet

4 Second-order properties: the K function

Supplementary reading:

• Bivand et al. [2, §7.4.5]: Second-order properties

A second-order property of a point process refers to the interactions between points². Examples are clustering (attraction) or competition (repulsion), with obvious ecological interest. Below (§6.2) we examine interactions between two types of points; here we consider one type of points, but at any distances. The F and G functions are "short-sighted", they only consider nearest neighbours to an arbitrary point or location, respectively. Here we consider any radius from an arbitrary point.

Ripley [9] proposed a K function for quantifying second-order properties for a HPP. It counts the number of points within a given distance of a point, and is defined as:

$$K(s) = \lambda^{-1} E[N_0(s)] \tag{6}$$

where E[.] is the expectation and $N_0(s)$ is the number of points found within radius s of point x_0 (an arbitrary point of the point pattern). This expectation is computed as:

$$\hat{K}(s) = (n(n-1))^{-1} |A| \sum_{i=i}^{n} \sum_{j \neq i} w_{ij}^{-1} \cdot x_j : d(x_i, x_j) \le s$$
(7)

where the **radius** $d(x_i, x_j)$ is the distance between two points, and the **weights** w_{ij} are the proportion of the area inside region A, of size |A|, of the circle centred on the target point x_i . The term $(n(n-1))^{-1}$ normalizes for the total number of point-pairs.

For the HPP, we have $K(s) = \pi s^2$, i.e., the area of a circle with radius s. If K(s) is greater, this indicates clustering, i.e., more points than expected with the radius; the inverse indicates a regular (dispersed) process.

Task 10: Compute and graph the K function for the three example datasets.

We use the Kest function of the spatstat.explore package; this works with point-patterns of class ppp.

```
Kjap <- Kest(japanesepines)
Kred <- Kest(redwoodfull)
Kcells <- Kest(cells)
```

 $^{^2}$ Recall: the **first-order** property refers to properties at a single point, e.g., the intensity of the process

par(mfrow=c(1,3))
plot(Kjap, main="Japanese pines")
plot(Kred, main="Redwoods")
plot(Kcells, main="Cells")
par(mfrow=c(1,1))



As with the G function, there are several border corrections; see help(Kest) for details.

We can also compute a simulation envelope for the K function, in the same manner as for the G function.

Task 11: Compute and graph simulation envelopes for the K function, for
the three example datasets.

We choose to display these to a radius that is 1/3 across the diagonal of the unit bounding box.

```
r <- seq(0, sqrt(2)/6, by = 0.005)
envjap <- envelope(japanesepines, fun=Kest, r=r, nrank=2, nsim=99, verbose=F)
envred <- envelope(redwoodfull, fun=Kest, r=r, nrank=2, nsim=99, verbose=F)
envcells <- envelope(cells, fun=Kest, r=r, nrank=2, nsim=99, verbose=F)</pre>
```

```
par(mfrow = c(1, 3))
plot(envjap, main = "Japanese pines, K-function envelope")
plot(envred, main = "Redwood trees, K-function envelope")
plot(envcells, main = "Cells, K-function envelope")
par(mfrow = c(1, 1))
```



These envelopes confirm the interpretations.

4.1 The *L* function: a linearized *K* function

A linear version of K may be easier to interpret; therefore Besag proposed a function $L(r) = \sqrt{K(r)/\pi}$. All this does is linearize the expected value, so it appears on the plot as a straight line,

Task 12 : Compute and plot the *L* function and their envelopes for the three patterns. \bullet

```
r <- seq(0, sqrt(2)/6, by = 0.005)
envjap <- envelope(japanesepines, fun = Lest, r = r, nrank = 2, nsim = 99,
        verbose = F)
envred <- envelope(redwoodfull, fun = Lest, r = r, nrank = 2, nsim = 99,
        verbose = F)
envcells <- envelope(cells, fun = Lest, r = r, nrank = 2, nsim = 99, verbose = F)</pre>
```

```
par(mfrow = c(1, 3))
plot(envjap, main = "Japanese pines, L-function envelope")
plot(envred, main = "Redwood trees, L-function envelope")
plot(envcells, main = "Cells, L-function envelope")
par(mfrow = c(1, 1))
```



4.2 * Modifying the window

Recall that an object of class ppp includes a window of class owin, that defines the window in which the point-pattern is evaluated. Suppose we want to only evaluate the point-pattern in some smaller area. To do this, we can create a new window with the owin function, and then use it to extract just those points that fall in the window.

Task 13 :Create a window covering the upper left-hand (NW) quadrant of
the Japanese pines point-pattern. Extract just the Japanese pines points in
this window and plot them.

We use the **owin** function to specify the new window, and then the **inside.owin** logical function to determine which points are in the new window. We then use the logical vector to select the points in the new window, and save these as a new point pattern.

```
str(japanesepines, max.level = 1)
## List of 5
##
   $ window
                :List of 4
    ..- attr(*, "class")= chr "owin"
##
   $ n
##
                : int 65
                : num [1:65] 0.09 0.29 0.38 0.39 0.48 0.59 0.65 0.67 0.73 0.79 ...
##
   $ x
                : num [1:65] 0.09 0.02 0.03 0.18 0.03 0.02 0.16 0.13 0.13 0.03 ...
##
   $у
##
    $ markformat: chr "none"
   - attr(*, "class")= chr "ppp"
##
print(japanesepines$window)
## window: rectangle = [0, 1] x [0, 1] units (one unit = 5.7 metres)
(window.nw <- owin(xrange = c(0, 0.5), yrange = c(0.5, 1)))
## window: rectangle = [0, 0.5] x [0.5, 1] units
table(is.in <- inside.owin(japanesepines, w = window.nw))</pre>
```

```
##
## FALSE TRUE
## 43 22
japanesepines.nw <- japanesepines[is.in]</pre>
```

We see only 22 of the 65 Japanese pines are in this window.

However, this selection does not change the window size. We can see this with the Windowunction of the spatstat.geom package:

Window(japanesepines.nw)

window: rectangle = [0, 1] x [0, 1] units (one unit = 5.7 metres)

If we want to reduce the window size of the new point pattern, we again use the Window function to set the window size:

```
Window(japanesepines.nw) <- window.nw
Window(japanesepines.nw)
```

window: rectangle = [0, 0.5] x [0.5, 1] units (one unit = 5.7 metres)

Now we can plot the reduced window and its points: plot(japanesepines.nw, main = "Japanese pines, NW quadrant")



Japanese pines, NW quadrant

Task 14: Compute and plot the G and L functions for this subset; compare with these functions for the full set. Limit the plot radius to 0.10, i.e., the close-range part of the function

```
par(mfrow = c(2, 2))
G <- Gest(japanesepines.nw)
plot(G, main = "G-function, Japanese pines, NW quadrant", xlim = c(0, 0.1))
G <- Gest(japanesepines)
plot(G, main = "G-function, Japanese pines, all", xlim = c(0, 0.1))
L <- Lest(japanesepines.nw)
plot(L, main = "L-function, Japanese pines, NW quadrant", xlim = c(0, 0.1))
L <- Lest(japanesepines)
plot(L, main = "L-function, Japanese pines, all", xlim = c(0, 0.1))</pre>
```

par(mfrow = c(1, 1))



It's clear that the smaller window, with fewer points, results in a more irregular function. The G-function is considerably different: for the quadrant we see some clustering around r = 0.05; this is not seen in the full point-pattern.

5 The F function; non-rectangular windows

The G function explored in §2 is a **point-related** function, i.e., computed from each point in the pattern. Another way to examine point distribution is with a **location-related** function, i.e., computed from any location, whether or not it is a point. This provides information about empty space. Such a function developed from the theory of Poisson processes is the F "empty space" function, which we examine in this section.

However, there is a complication. The above examples used rectangular windows "filled" with the point pattern. The implicit assumption (which we now make explicit) is that the data-generating process (i.e., process by which the points were placed) operates over the whole window, and in some border area outside the window. The process may not be homogeneous, as we saw in the kernel density estimation (§3), but it does "fill" the window

– there is a probability that any location in the window could have a point. However, if the rectangular window includes areas that were not observed or not part of the study, there will be "white space" which appears as part of the pattern, but is not. In that case some statistics will be misleading, in particular, the "empty space" F function, and any plots will include areas that are not interesting.

Another issue is that we may be given a point-pattern that is clearly only filling part of a map, and we want to extract that area. An example is the point-pattern of trees from which we want to derive the boundary of a forest. The spatstat.geom package has several useful functions for that purpose.

We illustrate the process of specifying a window with the **meuse** example dataset provided with the **sp** package. This is a set of observation points in the river Maas (Meuse) floodplain near the village of Stein, Limburg province, Netherlands.

Task 15: Load the meuse example dataset, restrict it to just the Pb content and flooding frequency attributes, convert to a spatial object of class sf, make an equivalent point-pattern object of class ppp, and plot as a marked point-pattern, marked by the flooding frequency (field ffreq).

Here we use the as.ppp method to convert from sf to ppp. The meuse data set is first converted to an sf object with the st_as_sf method, specifying also the coördinate fields. The flood frequency is specified as the marks of the point pattern with the marks function.

```
data(meuse, package = "sp")
meuse <- meuse[,c("x","y","ffreq")]</pre>
require(sf)
meuse.sf <- st as sf(meuse, coords = 1:2)</pre>
meuse.ppp <- as.ppp(meuse.sf)</pre>
str(meuse.ppp)
## List of 6
               :List of 4
## $ window
    ..$ type : chr "rectangle"
##
##
    ..$ xrange: num [1:2] 178605 181390
    ..$ yrange: num [1:2] 329714 333611
##
##
    ..$ units :List of 3
##
    ....$ singular : chr "unit"
                      : chr "units"
##
    .. ..$ plural
##
    .. ..$ multiplier: num 1
    ....- attr(*, "class")= chr "unitname"
##
     ..- attr(*, "class")= chr "owin"
##
## $ n : int 155
## $ x : num [1:155] 181072 181025 181165 181298 181307 ...
## $ y
                : num [1:155] 333611 333558 333537 333484 333330 ...
## $ markformat: chr "vector"
               : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 ...
## $ marks
## - attr(*, "class")= chr "ppp"
marks(meuse.ppp) <- meuse$ffreq</pre>
tmp <- plot(meuse.ppp, use.marks=TRUE,</pre>
     cols=c("red","orange","green"),
     chars=16, which.marks="marks",
     main="Meuse floodplain flood frequency class",
     axes=T)
grid()
legend("left", pch=16,
```





We can see that the point-pattern only partially fills the bounding rectangle. By default, the type conversion to class **ppp** defines the window as the rectangular bounding box of the point-pattern; we can see this as the **window** field of the **ppp** object:

meuse.ppp\$window

```
## window: rectangle = [178605, 181390] x [329714, 333611] units
```

Task 16 : Compute a bounding window and replace the rectangular bound-ary with it.

We compute the window as the Ripley and Rasson [10] estimate of the spatial domain. This is a clever way of expanding the convex hull (which contains the outermost points) consistent with the intensity of the pattern.

We use the **ripras** "Ripley-Rasson forest edge" function to compute the window; we plot this along with the convex hull computed by the **convexhull** function.

Note: Both ripras and convexhull return an object of class owin; this includes the boundary in field bdry as a list of coördinate vectors.

To plot a point pattern we use the plot.ppp function, which is automatically called by the generic plot method for an object of class ppp.

```
meuse.ppp.r <- meuse.ppp</pre>
(meuse.ppp.r$window <- ripras(meuse.ppp))</pre>
## window: polygonal boundary
## enclosing rectangle: [178543.73, 181443.23] x [329644.9, 333702.2]
## units
tmp <- plot(meuse.ppp.r, use.marks=TRUE,</pre>
     cols=c("red","orange","green"),
chars=16, which.marks="ffreq",
     main="Meuse floodplain flood frequency class",
     boundary=2, axes=T)
grid()
legend("left", pch=16,
       col=c("red","orange","green"),
       legend=c("Annually","2-5 Years", "> 5 Years"))
ch <- convexhull(meuse.ppp)</pre>
lines(ch$bdry[[1]]$x, ch$bdry[[1]]$y, lty=2)
legend("bottomright", lty=1:2,
      legend=c("Ripley-Rasson", "convex hull"))
```

Meuse floodplain flood frequency class



Q16 : Describe the polygonal window.

Jump to $A16 \bullet$

Q17 : How did changing the boundary affect the mean intensity of the point process? Jump to $A17 \bullet$

The intensity function of the spatstat.geom package computes the intensity from the number of points and the window area: str(meuse.ppp)

List of 6

```
## $ window :List of 4
##
    ..$ type : chr "rectangle"
    ..$ xrange: num [1:2] 178605 181390
##
##
     ..$ yrange: num [1:2] 329714 333611
    ..$ units :List of 3
##
##
    ....$ singular : chr "unit"
##
     .. ..$ plural
                     : chr "units"
##
    .. ..$ multiplier: num 1
    ... - attr(*, "class")= chr "unitname"
##
##
     ..- attr(*, "class")= chr "owin"
## $ n : int 155
## $ x : num [1:155] 181072 181025 181165 181298 181307 ...

##
   $у
               : num [1:155] 333611 333558 333537 333484 333330 ...
## $ markformat: chr "vector"
## $ marks : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 ...
##
   - attr(*, "class")= chr "ppp"
intensity(meuse.ppp)
##
             1
                         2
                                       3
## 7.739692e-06 4.422681e-06 2.119201e-06
intensity(meuse.ppp.r)
##
                         2
            1
                                       3
## 1.428895e-05 8.165116e-06 3.912452e-06
round(intensity(meuse.ppp.r)/intensity(meuse.ppp), 2)
    1 2
##
              3
## 1.85 1.85 1.85
```

Clearly, this is not a perfect boundary of the area from which the points were taken; we know from the documentation that it is bounded by a large meander of the river Maas (Meuse), and is limited on the east side by a canal (the Julianakanaal) and steep cliff, so ideally we'd have a bounding polygon of the actual study area. Absent this, the Ripley-Rasson method at least restricts the area.

Note: See A.2 for how to import a polygonal boundary in ESRI shapefile format, and use it for the window.

A major effect of reducing the window to the actual area sampled is to properly estimate the "empty space" function, i.e., average distance from an arbitrary location in the window to the nearest point (event).

Task 17: Compute and plot the "empty space" function F for the Meusepoint-pattern in the rectangular and polygonal windows.

The Fest function of the spatstat.explore package computes this function on an object of class ppp. We specify the same x-axes to compare the functions side-by-side:

```
par(mfrow = c(1, 2))
plot(Fest(meuse.ppp), main = "rectangular window", xlim = c(0, 550))
plot(Fest(meuse.ppp.r), main = "polygonal window", xlim = c(0, 550))
par(mfrow = c(1, 2))
```



As with the G function, there are several cumulative functions; $F_{\text{pois}}(r)$ is the theoretical distribution in the case of CSR; $F_{\text{bord}}(r)$ is the same corrected for border effects; the estimates are marked $\hat{F}_{\text{method}}(r)$ for Kaplan-Meier ("km"), Chiu-Stoyan ("cs"). Note the different x-axis scales of the two plots.

Q18 : What proportion of space is expected to have at least one point within 100 m for the rectangular and polygonal windows? Jump to $A18 \bullet$

Q19: Describe the agreement (or lack thereof) of the observed $\hat{F}_{\rm km}(r)$ with the theoretical for CSR $F_{\rm pois}(r)$. Jump to A19 •

6 Marked point patterns

A marked point pattern is one where each point has some attribute; this can be a continuous value (e.g., tree size) ($\S6.4$) or a categorical attribute (e.g., tree species, tree size class) ($\S6.1$).

6.1 Categorical marks

We first examine a point-pattern marked with a **categorical** attribute: the "forest fires" dataset clmfires, supplied as an example in the spatstat package. This is a record of forest fires (1998-2007) in the Castilla-La Mancha region (E). For each fire there are four types of marks, i.e., attributes: cause, date, day of year, and size.

Task 18: Load the clmfires dataset and display the locations of the fires, along with their cause.

Castilla-La Mancha forest fires



Q20 : What do the marks represent? How many classes are there? Jump to $A20 \bullet$

Task 19 :Split the marked point-pattern and show the pattern for eachmark separately.•

The split.ppp method of the spatstat.geom package specializes the generic split method. Similarly, the plot.splitppp method of the spatstat.geom package specializes the generic plot method

Here we don't need to plot any mark types, since we've already split on the cause. So the use.marks argument to plot.splitppp is set to FALSE. clmfires.split <- split(clmfires)

```
str(clmfires.split, max.level=1)
## List of 4
## $ lightning :List of 6
## ..- attr(*, "class")= chr "ppp"
## $ accident :List of 6
## ..- attr(*, "class")= chr "ppp"
## $ intentional:List of 6
## ..- attr(*, "class")= chr "ppp"
## $ other :List of 6
```

```
## ..- attr(*, "class")= chr "ppp"
## - attr(*, "class")= chr [1:4] "splitppp" "ppplist" "solist" "list"
## - attr(*, "fsplit")= Factor w/ 4 levels "lightning","accident",..: 3 1 1 1 4 4 2 4 2 2 ...
## - attr(*, "fgroup")= Factor w/ 4 levels "lightning","accident",..: 3 1 1 1 4 4 2 4 2 2 ...
plot(clmfires.split, use.marks=FALSE,
    main="Castilla-La Mancha forest fires",
    pch=21, bg=2)
```

Castilla-La Mancha forest fires



Q21 : Do the fires with different causes appear to have different point patterns? Jump to $A21 \bullet$

We could compare the separate patterns with the usual G, F, J, K or L functions and compare the function plots visually.

6.2 Interaction between point patterns: the cross-K function

Another question can be raised when there are several patterns covering the same area: what is the **interaction** between them? That is, do occurrences of one mark "attract" or "repel" those of other marks?

Note: The quotes for "attract" and "repel" remind us that we need meta-statistical information to propose the causes of observed interactions.

We call such a process a **multi-type** process, that is, we assume that there may be some interaction between the types. In the current example, we may expect that an area burned with one kind of fire would not be susceptible to another kind of fire, because the necessary fuel would have been removed by the first fire. We assume that the multi-type process is stationary across the area.

The appropriate statistic to investigate this is a so-called "cross" pointpattern function, from mark type *i* to mark type *j* or vice-versa. For example, a crossed *K* function of a stationary multi-type point process with intensity λ_j of point type *j* is defined so that $\lambda_j K_{ij}(r)$ is the expected number of **additional** random points of type *j* within a distance *r* of a typical point of type *i*.

The $K_{ij}(r)$ function plotted over a range of distances is used to form hypotheses about the multi-type point pattern. If the two point-processes are independent, the expected value $K_{ij}(r) = \pi r^2$, that is, the number of additional points just depends on the area of the circle centred on a source point. If the empirical K_{ij} function is above the theoretical function πr^2 (i.e., a parabola), there are more points of type j near to the source points of type i than expected; this suggest dependence between the processes. If the empirical function is below the theoretical, this suggests repulsion or avoidance.

Task 20 :Compute and plot the cross-K function for the relation betweenintentional and lightning-induced fires.

The Kcross function computes Ripley's K, as for the univariate case (function Kest), but the measure is the number of neighbours of another pattern within a radius of a given point³. However, this function works on an "multi-type point-pattern" object, which is a point-pattern with a single mark. The clmfires object has four kinds of marks:

str(clmfires\$marks)

```
## 'data.frame': 8488 obs. of 4 variables:
## $ cause : Factor w/ 4 levels "lightning","accident",..: 3 1 1 1 4 4 2 4 2 2 ...
## $ burnt.area : num 0.4 0 0.4 0 1.05 3 0.1 0.02 0.4 2.85 ...
## $ date : Date, format: "1998-01-07" ...
## $ julian.date: num 6 6 6 6 6 7 7 7 8 8 ...
```

So, we make an object with just a single mark, i.e., the causes: clmfires.cause <- clmfires

```
is.multitype(clmfires.cause)
## [1] FALSE
```

```
clmfires.cause$marks <- clmfires$marks$cause
is.multitype(clmfires.cause)</pre>
```

[1] TRUE

Now we can compute the cross-K function, using Kcross. We specify the

 $^{^3}$ There are similar analogues of the $L,\,G,\,{\rm and}\;J$ functions, but not the F "empty space" function.

'translation' edge correction, suitable for complex geometries such as the province boundaries:

Kcross.il <- Kcross(clmfires.cause, "intentional", "lightning", correction = "translate")
plot(Kcross.il)
grid()</pre>



Q22 : Is there evidence for interaction between the processes that producethe intentional and accidental fires?Jump to $A22 \bullet$

6.3 Combining point patterns

We may have two or more unmarked point patterns which represent different types of points, which we want to combine into a marked point pattern. For example, the **redwoodfull** and **japanesepines** point patterns represent two kinds of trees as unmarked point patterns. If we suppose these two patterns are from the same area⁴ we may ask what is the relation between them. We can discover this with the cross-K function, if we can combine them into a single multi-type marked point pattern.

Task 21 : Combine the redwoodfull and japanesepines point patterns into a single multi-type marked point pattern.

The superimpose function of the spatstat.geom package superimpose several point patterns. These can optionally be supplied with marks applied to all points in each pattern.

```
two.trees <- superimpose(rw = redwoodfull, jp = japanesepines)
str(two.trees)
## List of 6
## $ window :List of 4</pre>
```

 $^{^{4}}$ which is not true, but allows us to illustrate the techniques

```
..$ type : chr "rectangle"
##
##
    ..$ xrange: num [1:2] 1e-09 1e+00
     ..$ yrange: num [1:2] 1e-09 1e+00
##
##
     ..$ units :List of 3
##
    ....$ singular : chr "metre"
##
    ....$ plural : chr "metres"
##
     .. ..$ multiplier: num 5.7
    ....- attr(*, "class")= chr "unitname"
##
    ..- attr(*, "class")= chr "owin"
##
##
   $ n
               : int 260
##
   $ x
               : num [1:260] 0.931 0.939 0.935 0.98 0.787 ...
               : num [1:260] 0.818 0.764 0.722 0.665 0.661 ...
##
   $ y
##
    $ markformat: chr "vector"
              : Factor w/ 2 levels "rw","jp": 1 1 1 1 1 1 1 1 1 ...
##
   $ marks
   - attr(*, "class")= chr "ppp"
##
plot(two.trees, main = "Superimposed point patterns", cols = c("green",
   "blue"))
```

œ Δ Δ Δ 2 \wedge Δ_{\triangle} Δ Δ ◬ 0 Δ ×, 980 C $\Delta \Delta$ æд Δο Ø 8 ⊘ rw 8 Δo Δ 0 œ Δ jp **^8** °0 0 Ø 0 ଚ 0 <u>م</u>8ّ ₿ ΟΔ 0 P Λ 0 Δ Δ ଡ 49₀₀₀₀ 0 _ **6**0 ° 0 0 0 °∆° S cΔ Δo Δ

Superimposed point patterns

In this case the windows had the same extent; by default a union of the windows of the superimposed point patterns is used. The optional W "Window" argument provides several additional ways to specify a window.

```
Task 22 : Compute the crossed K-function for these two tree species.

• Kcross.jp.rw <- Kcross(two.trees)
plot(Kcross.jp.rw)
grid()
</pre>
```



Q23 : Does there appear to be any interaction between the point processes that produced these two tree species? Is this surprising? Jump to $A23 \bullet$

6.4 Continuous marks

The marks on a point-pattern may be continuous variables rather than categories. An example is the longleaf dataset, which shows the locations and diameters at breast height (DBH) of 584 longleaf pines (*Pinus palustris*) in a 200 x 200 metre region in southern Georgia (USA)

Task 23 : Load and display this point pattern. Show the mature trees witha red symbol, and saplings with a green symbol.

The summary function summarizes the dataset. To just see the window size, use the Window function (note the capital "W").

```
summary(longleaf)
## Marked planar point pattern: 584 points
## Average intensity 0.0146 points per square metre
##
## Coordinates are given to 1 decimal place
## i.e. rounded to the nearest multiple of 0.1 metres
##
## marks are numeric, of type 'double'
## Summary:
##
     Min. 1st Qu. Median
                             Mean 3rd Qu.
                                             Max.
##
                           26.84 42.12
     2.00
            9.10
                    26.15
                                            75.90
##
## Window: rectangle = [0, 200] x [0, 200] metres
## Window area = 40000 square metres
## Unit of length: 1 metre
```

```
Window(longleaf)
```

data(longleaf)



Longleaf pines, location and DBH

Q24 : Do the trees appear to be clustered, regularly-spaced, or randomly placed? Do trees of similar size appear to be clustered? What appears to be the relation between mature trees and saplings? $Jump \text{ to } A24 \bullet$

Task 24 : Compute and plot the K function for the longleaf pines.

Again the estimated K function is computed with Kest.

K.long <- Kest(longleaf)
plot(K.long, main = "K function, longleaf pines")</pre>

K function, longleaf pines



The trees clearly show some clustering at all distances to 50 m.

7 Models of spatial processes

Supplementary reading:

• Bivand et al. [2, §7.4.4]: Likelihood of an inhomogeneous Poisson process

The observed point pattern is presumably the **realization** of some **point process**, i.e., a spatial data generating process (sDGP) by which points, also called "events", are placed on the landscape. These could be completely random with some intensity (a **homogeneous** Poisson process), random but with varying intensity across the region (an **inhomogeneous** Poisson process), a process depending on inter-point interactions, depending on a regional trend, depending on environmental covariables, or any combination. If we can fit a model to the observed process we can (1) infer the sDGP which generated it; (2) map the results of the process.

The result of such models is a **conditional intensity** $\lambda(u, \mathbf{x})$, a function of the location u and the observed point pattern \mathbf{x} . The units are the number of points per unit area. In practice we compute this over some "small" cell.

Baddeley and Turner [1] explain how to fit stochastic models to observed point patterns with the versatile ppm function of the spatstat.model package. The issue of modelling is quite deep and you are encouraged to read this paper before building your own models; here we only show some possibilities.

The general formula for models that can be fit with ppm is [1, Eqn. (4)]:

$$\lambda(u, \mathbf{x}) = \exp\left(\psi^T B(u) + \phi^T C(u, \mathbf{x})\right)$$
(8)

where the two components are:
- 1. the **spatial trend** B(u) which depends only on location; this could also include covariates at these locations;
- 2. the stochastic interactions $C(u, \mathbf{x})$, i.e., the dependence between the points of the point process.

The analyst specifies the forms of B and C and ppm estimates the coefficients (ψ, ϕ) .

We continue with the Castilla-La Mancha forest fires example of §6. In this modelling exercise we subset the whole dataset to just one kind of fire; it should be easier to interpret the model results.

Task 25 : Restrict the dataset to intentional fires.

Again use the split.ppp function, and then select one of the subsets. clmfires.i <- split(clmfires, "cause")\$intentional

```
plot(clmfires.i, chars = 21, cex = 0.5, bg = 2, axes = T, main = "Castilla-La Mancha intentional fores
    use.marks = FALSE)
grid()
```





Task 26 : Remove the marks from the point pattern.

This is necessary because ppm is not yet implemented for marked patterns. Here we remove the marks for the entire pattern; it would also be possible to split the pattern according to a mark using the split.ppp function, remove the marks from each of the sub-patterns and analyze each one.

7.1 Null model

Task 27: Model the forest fire incidence as a Poisson process, i.e., complete spatial randomness (CSR). This is a "null" model because it is the simplest hypothesis about how points are placed.

This is the simplest conditional intensity: $\lambda(u, \mathbf{x}) = \beta$, where β is a single intensity of a (presumed) homogeneous Poisson process. In the terminology of Eqn. (8), both *B* (trend) and *C* (interaction) are absent.

We specify this to ppm by setting the trend argument to ~1, i.e., the process only has a mean intensity. We also specify the type of interaction between points by setting the interaction argument to $NULL^5$.

```
print(m.pois <- ppm(clmfires.i, trend = ~1, interaction = NULL))</pre>
```

```
## Stationary Poisson process
## Fitted to point pattern dataset 'clmfires.i'
## Intensity: 0.02250655
##
               Estimate
                             S.E. CI95.lo CI95.hi Ztest
                                                                   Zval
## log(lambda) -3.793949 0.02366243 -3.840326 -3.747571 *** -160.3364
class(m.pois)
## [1] "ppm"
exp(coef(m.pois))
## log(lambda)
## 0.02250655
intensity(clmfires.i)
## [1] 0.02250655
(clmfires.i$n/summary(clmfires.i)$window$areas)
## [1] 0.02250655
```

A model fitted by ppm is of class ppm.

This is not a very interesting model, since we could get the same result simply from the average intensity, using the **intensity** function or even direct computation from the number of points and the window area. The only complication is that **ppm** works with the logarithm of the parameters, in this case just the Poisson intensity β . Notice however the standard error and confidence intervals that are provided with the model summary.

7.2 Trend surface

A more complex model is $\lambda(u, \mathbf{x}) = \beta(u)$, where $\beta(u)$ is a variable intensity, dependent on the location u, of a (presumed) inhomogeneous Poisson process. This is termed a **trend**, which may be a function of the coördinates or

⁵ These are both defaults and so don't have to be explicitly specified.

of covariables. In the terminology of Eqn. (8), B (trend) is defined but C (interaction) are absent. The form of B is a trend surface, i.e., a polynomial function of the coördinates.

Task 28: Model the intentional forest fire incidence as first- and secondorder regional trend plus a Poisson process, i.e., complete spatial randomness (CSR) after accounting for a trend.

We specify the trend to ppm by setting the trend argument to a formula; for example x+y for a first-order trend: the intensity changes linearly along some plane to be computed. Here we use the polynom function to specify both first- and second-order trend surface:

```
(m.ts1 <- ppm(clmfires.i, trend = ~polynom(x, y, 1), interaction = NULL))</pre>
## Nonstationary Poisson process
## Fitted to point pattern dataset 'clmfires.i'
##
## Log intensity: ~x + y
##
## Fitted trend coefficients:
## (Intercept)
                          х
## -3.287880898 -0.005732725 0.002884705
##
##
                                            CI95.lo
                  Estimate
                                  S.E.
                                                        CI95.hi Ztest
## (Intercept) -3.287880898 0.0731704116 -3.431292270 -3.144469527 ***
## x -0.005732725 0.0002707268 -0.006263340 -0.005202111
                                                                   ***
## y
             0.002884705 0.0003026018 0.002291616 0.003477793
##
                    Zval
## (Intercept) -44.934569
## x
             -21.175316
                9.533004
## y
(m.ts2 <- ppm(clmfires.i, trend = ~polynom(x, y, 2), interaction = NULL))</pre>
## Nonstationary Poisson process
## Fitted to point pattern dataset 'clmfires.i'
##
## Log intensity: -x + y + I(x^2) + I(x * y) + I(y^2)
##
## Fitted trend coefficients:
                                                  I(x^2)
## (Intercept)
                                                              I(x * v)
                           x
                                         v
## -5.268868e+00 1.202084e-02 9.535875e-03 -3.409373e-05 -2.969984e-05
##
         I(y^2)
## -2.229839e-06
##
##
                  Estimate
                                   S.E.
                                              CI95.lo
                                                           CI95.hi
## (Intercept) -5.268868e+00 2.541951e-01 -5.767081e+00 -4.770655e+00
## x 1.202084e-02 1.686762e-03 8.714846e-03 1.532683e-02
## y
              9.535875e-03 1.615694e-03 6.369173e-03 1.270258e-02
             -3.409373e-05 3.582964e-06 -4.111621e-05 -2.707125e-05
## I(x^2)
## I(x - J.
## I(y^2) -2.2200
Ztest
## I(x * y) -2.969984e-05 4.670127e-06 -3.885312e-05 -2.054656e-05
             -2.229839e-06 3.526177e-06 -9.141018e-06 4.681341e-06
                          Zval
## (Intercept) *** -20.7276522
               *** 7.1265770
*** 5.9020307
## x
## y
## I(x^2)
            *** -9.5155088
## I(x * y) *** -6.3595372
```

I(y^2) -0.6323672

Now we see the fitted coefficients β ; the intercept is the overall log-intensity at (0,0) (the lower-left corner of the pattern) and the coefficients show the change intensity in the x and y directions, along with their standard errors and confidence intervals; recall these are logarithms, so we convert to original units to interpret them. Here we see an increase in intensity towards the WNW, almost equal in both axes. This accords with our visual estimate.

Task 29 : Plot the trend surfaces.

The plot.ppm function calls predict.ppm (see below, \$8) to compute the spatial trend and conditional intensity of the fitted point process model on a grid, and then displays the result; by default the grid is 40 by 40 pixels filling the bounding box.

We first visualize these by a colour ramp 2.5D plot; the **how** argument specifies the type of plot:

2nd-order trend

```
par(mfrow = c(1, 2))
plot.ppm(m.ts1, ngrid = c(80, 80), how = "image", superimpose = F, trend = T,
    se = F, pause = F, main = "1st-order trend")
plot.ppm(m.ts2, ngrid = c(80, 80), how = "image", superimpose = F, trend = T,
    se = F, pause = F, main = "2nd-order trend")
par(mfrow = c(1, 1))
```

1st-order trend



We can also see the trends as perspective plots:

```
par(mfrow = c(1, 2))
plot.ppm(m.ts1, ngrid = c(80, 80), how = "persp", theta = -30, phi = 30,
    trend = T, se = F, pause = F, main = "1st-order trend")
plot.ppm(m.ts2, ngrid = c(80, 80), how = "persp", theta = -30, phi = 30,
    trend = T, se = F, pause = F, main = "2nd-order trend")
par(mfrow = c(1, 1))
```



The anova.ppm function performs analysis of deviance for two or more fitted models with Poisson interaction terms, i.e., independence:

```
anova.ppm(m.ts2, m.ts1, m.pois)
## Analysis of Deviance Table
##
## Model 1: ~x + y + I(x^2) + I(x * y) + I(y^2) Poisson
## Model 2: ~x + y Poisson
## Model 3: ~1 Poisson
## Model 3: ~1 Poisson
## 1 6
## 2 3 -3 -126.47
## 3 1 -2 -504.87
```

Here we see that the trend surfaces uses more degrees of freedom but both reduce the residual deviance slightly as the model becomes increasingly complex.

7.3 Strauss process

So far we've treated the observations as independent (a Poisson process), possibly influenced by a regional trend in overall intensity. Another possibility is that fires are not independent. In the terminology of Eqn. (8), *B* (trend) is absent but *C* (interaction) is present; the analyst must define the form of *C*. One way to model that is as a **Strauss process**: $\lambda(u, x) = \beta \gamma^{t(u, x)}$,

where β is the overall homogeneous intensity and γ is an **interaction** parameter $0 \leq \gamma \leq 1$ and $t(u, \mathbf{x})$ is the number of points of the pattern \mathbf{x} closer than the **interaction radius** r of the location u. This has an interesting interpretation: $\gamma = 0 \Longrightarrow \lambda = 0$, that is, within the radius there is no chance of finding another point, perhaps because of the intrinsic size of a "point". With $\gamma < 1$ the chance of a second point is reduced, at $\gamma = 1$ this is equivalent to a Poisson process (no effect one way or the other).

Note: A Strauss process is an example of a so-called **Gibbs process**, derived from physics to model repulsion; they include an intensity and an interaction function.

Task 31 : Fit a homogeneous Strauss process model to the forest fire incidence and plot the resulting surface.

Here we add an interaction argument to the ppm function. See ?ppm for the choices. We choose StraussHard. Then we must choose an interaction radius; the parameter $0 \le \gamma \le 1$ is estimated by ppm.

We investigate this with the K function. Above (§4) we saw that this measures the number of points within a given radius of a given point, as a function of radius.

plot(Kest(clmfires.i, r = seq(0, 10, by = 0.2)))



Kest(clmfires.i, r = seq(0, 10, by = 0.2))

There seems to be an inflection point around 4 km, so we pick this as an interaction radius.

```
## Interaction distance: 4
## Hard core distance: 0.003995263
## Fitted interaction parameter gamma: 1.0781255
##
## Relevant coefficients:
## Interaction
## 0.07522386
##
## For standard errors, type coef(summary(x))
exp(coef(m.strauss.4))
## (Intercept) Interaction
## 0.01768717 1.07812547
```

There are two fitted parameters: the overall log intensity β and the strengthof-interaction parameter γ . This latter was fit as 1.0781, i.e., $\gamma > 1$, so there is on average clustering.

Task 32 : Compare the likelihood of the several fitted models.

•

The logLik function shows the log-likelihood of the fitted parameters: data.frame(

```
model=c("Poisson","1st order trend", "2nd order trend", "Strauss/hard core"),
likelihood=c(logLik(m.pois, warn=F),
logLik(m.ts1, warn=F),
logLik(m.ts2, warn=F),
logLik(m.strauss.4, warn=F)))
## model likelihood
## 1 Poisson -8561.993
## 2 1st order trend -8306.779
## 3 2nd order trend -8243.545
## 4 Strauss/hard core -6776.067
```

Clearly the Strauss model is superior: there is local clustering, not CSR, and this is a better fit to the observations than either trend surface.

7.4 Covariates

The observed point pattern may well depend on environmental factors. For example, density of trees in a forest may depend on soil type, elevation, temperature or rainfall. Baddeley and Turner [1, §7] explain how how to include covariates, such as environmental factors, in the model. The clmfires dataset is accompanied by a raster dataset clmfires.extra, a list of two objects of class im, also defined by spatstat, which is a matrix of images (i.e., a layer stack). One of the objects in the list is clmcov200, a 200 x 200 pixels grid in the same coördinate system as clmfires, showing four possible covariates that might affect fire incidence: elevation, orientation (aspect), slope and landuse. The anova.ppm function uses this image to extract the value of the covariable(s) at the locations of observed events (points).

Loading clmfires also loaded the covariate images as object clmfires.extra: names(clmfires.extra)

```
## [1] "clmcov100" "clmcov200"
names(clmfires.extra$clmcov200)
```

<pre>## [1] "elevation" "orientation" "slope" "landuse"</pre>									
name	<pre>names(clmfires.extra\$clmcov200\$landuse)</pre>								
## ##	[1] [8]	"v" "yrow"	"dim" "type	"xrange" "units"	"yrange"	"xstep	" "ystep"	"xcol"	
name	<pre>names(clmfires.extra\$clmcov200\$landuse)</pre>								
## ##	[1] [8]	"v" "yrow"	"dim" "type	"xrange" " "units"	"yrange"	"xstep	" "ystep"	"xcol"	
<pre>levels(clmfires.extra\$clmcov200\$landuse\$v)</pre>									
## ## ##	[1] [5] [9]	"urban" "conifer' "scrub"	1	"farm" "mixedforest" "artifgreen"	"meadow" ' "grassla	' and"	"densefores "bush"	st"	

Task 33 : Display the covariate images.

The generic plot method specializes to plot.im for objects of class im. plot(clmfires.extra%clmcov200, main = "200 m grid covariates")



200 m grid covariates

We can get a better view of the landuse classes by considering them as an SpatRaster object as defined by the terra package:

```
require(terra)
clmfires.lu.grid <- rast(clmfires.extra$clmcov200$landuse)
class(clmfires.lu.grid)
## [1] "SpatRaster"
## attr(,"package")
## [1] "terra"</pre>
```

plot(clmfires.lu.grid, col = hcl.colors(11, palette = "viridis"))



An interesting question is whether different land uses have different forest fire incidences.

Task 34: Model the incidence of intentional forest fire as a function of the
"landuse" covariate, without any interaction process. Compare the model fit
with the null model.

In the terminology of Eqn. (8), B (trend) depends on the covariates (not the coördinates as in the trend surface), and C (interaction) is absent. The analyst must define the form of B, here, a linear model of the covariate. levels(clmfires.extra\$clmcov200\$landuse)

## ## ##	[1] "urban" [5] "conifer" [9] "scrub"	"farm" "mixedforest" "artifgreen"	"meadow" "grassland"	"denseforest" "bush"					
(m.	<pre>(m.lu <- ppm(clmfires.i, ~ landuse - 1,</pre>								
##	## Nonstationary Poisson process								
##	## Fitted to point pattern dataset 'clmfires.i'								
##	##								
##	## Log intensity: ~landuse - 1								
##	t#								
##	# Fitted trend coefficients:								
##	landuseurban	landuse	efarm land	lusemeadow					
##	-4.266410	-3.7:	L1950	-4.172085					

##	landusedense forest	landuse	landuseconifer landusemixedforest				
##	-4.000644	-3	3.716348	-4.211720			
##	landusegrassland	land	lusebush	landusescrub			
##	-4.003044	-3	3.967172	-3.62677	73		
##	landuseartifgreen						
##	-15.302585						
##							
##		Estimate	S.E.	CI95.lo	CI95.hi		
##	landuseurban	-4.266410	0.16439899	-4.588626	-3.944194		
##	landusefarm	-3.711950	0.03251280	-3.775674	-3.648226		
##	landusemeadow	-4.172085	0.13736056	-4.441307	-3.902863		
##	landusedenseforest	-4.000644	0.12700013	-4.249560	-3.751729		
##	landuseconifer	-3.716348	0.07669650	-3.866671	-3.566026		
##	landusemixedforest	-4.211720	0.19611614	-4.596100	-3.827339		
##	landusegrassland	-4.003044	0.11704115	-4.232440	-3.773647		
##	landusebush	-3.967172	0.07832604	-4.120688	-3.813656		
##	landusescrub	-3.626773	0.06250000	-3.749270	-3.504275		
##	landuseartifgreen	-15.302585	251.20102020	-507.647538	477.042367		
##		Ztest	Zval				
##	landuseurban	*** -25.	.95155663				
##	landusefarm	*** -114.	.16887274				
##	landusemeadow	*** -30.	.37323862				
##	landusedenseforest	*** -31.	.50110610				
##	landuseconifer	*** -48.	45525351				
##	landusemixedforest	*** -21.	.47564168				
##	landusegrassland	*** -34.	.20202124				
##	landusebush	*** -50.	. 64945926				
##	landusescrub	*** -58.	.02836198				
##	landuseartifgreen	-0.	.06091769				
<pre>data.frame(model=c("Poisson", "Landuse"),</pre>							
	<pre>likelihood=c(logLik(m.pois),</pre>						
	<pre>logLik(m.lu, warn=F)))</pre>						
## ##	model likelihoo 1 Poisson -8561.99	od 93					
##	2 Landuse -8533.13	31					

Q26 : Does the landuse explain some of the intentional fire pattern? Whichland use classes are more prone to fire?Jump to $A26 \bullet$

Task 35: Model the incidence of intentional forest fire as a function of the "landuse" covariate, *also* taking into account a presumed Strauss interaction process. Compare the model fit with the null, landuse-only and interaction-only models.

In the terminology of Eqn. (8), B (trend) depends on the covariates (not the coördinates as in the trend surface), and C (interaction) is present. The analyst must define the forms of both B (here, a linear model of the covariate) and C (here the Strauss process).

```
## -4.4280748 0.3394140 0.2921413
## landusedenseforest landuseconifer landusemixedforest
                       0.3072505
landusebush
##
          0.4687907
                                              0.2188922
##
    landusegrassland
                                            landusescrub
##
     0.3457191
                           0.4270869
                                             0.7187600
## landuseartifgreen
##
        -10.8745103
##
## Interaction distance: 4
## Hard core distance: 0.003995263
## Fitted interaction parameter gamma: 1.0789022
##
## Relevant coefficients:
## Interaction
## 0.07594402
##
## For standard errors, type coef(summary(x))
data.frame(
   model=c("Poisson", "Landuse", "Strauss", "Landuse + Strauss"),
   likelihood=c(logLik(m.pois),
       logLik(m.lu, warn=F),
       logLik(m.strauss.4, warn=F),
       logLik(m.lu.strauss.4, warn=F)))
##
               model likelihood
## 1
            Poisson -8561.993
## 2
             Landuse -8533.131
## 3
              Strauss -6776.067
## 4 Landuse + Strauss -6757.784
```

Q27 : Which model is most likely, given the observations? What do you conclude about the origin of intentionally-set fires? Jump to $A27 \bullet$

8 Spatial prediction

The models fit with ppm can be used to predict the point-pattern intensity over a study area. Obviously, they can not predict individual events (e.g., new forest fires) but they can predict the conditional intensity $\lambda(u, x)$ of an occurrence at each location over a grid. The predict.ppm function (called just as predict on an object of class ppm, i.e., a "point pattern model") evaluates the intensity at each grid location.

Task 36 : Predict the conditional intensity of intentional fires using thelanduse as predictor and plot it.

```
pred.lu <- predict(m.lu, covariates = clmfires.extra$clmcov200)
summary(pred.lu)
## real-valued pixel image
## 128 x 128 pixel array (ny, nx)
## enclosing rectangle: [4.131124, 391.3795] x [18.565, 385.189]
## kilometres
## dimensions of each pixel: 3.03 x 2.86425 kilometres
## finage is defined on a subset of the rectangular grid
## Subset area = 79462.0730449286 square kilometres
## Subset area fraction = 0.56
## Pixel values (inside window):
## range = [2.260329e-07, 0.0266019]
## integral = 1787.036</pre>
```

```
## mean = 0.02248917
image(pred.lu, main = "Fire intensity based on land use")
```



Fire intensity based on land use

Notice the default grid 128 x 128 pixels, and the automatic calculation of the size of each grid cell. This can be changed with the optional ngrid argument. Intensity at any set of locations can be requested with the optional locations argument. See ?predict.ppm for details.

Another model was the trend surface; we already saw that prediction in the previous section.

The best model of the previous section was landuse + Strauss (interaction process).

Task 37 :Predict the conditional intensity of intentional fires using the
landuse + Strauss process as predictors. Compare the range and mean with
the land use-only model.

```
## kilometres
## dimensions of each pixel: 3.03 x 2.86425 kilometres
## Image is defined on a subset of the rectangular grid
## Subset area = 79462.0730449286 square kilometres
## Subset area fraction = 0.56
## Pixel values (inside window):
## range = [2.260329e-07, 0.0244943]
## integral = 1402.767
## mean = 0.01765328
range(pred.lu.strauss)
## [1] 2.260329e-07 2.449430e-02
range(pred.lu)
## [1] 2.260329e-07 2.660190e-02
mean(pred.lu.strauss)
## [1] 0.01765328
mean(pred.lu)
## [1] 0.02248917
```

Q29 : Which model has the higher predicted mean intensity and wider range? Jump to $A29 \bullet$

Task 38 : Plot the two predicted maps with the same stretch.

•

land use

land use + Strauss process





9 Spatio-temporal analysis

Point patterns can evolve over time. In this short section we introduce one way to analyze these: by comparing **time slices** of a point pattern where each point is associated with a **time stamp**, i.e., time of observation.

The objective is to analyze point-patterns which may change over time, for example:

- locations of live trees in a forest plot (some die, some new ones grow);
- locations of crime or disease incidences; these occur at known times.

There is a rich literature on spatio-temporal point process models, see Diggle [5] and Taylor et al. [11]. Here we only show some visualizations and simple analysis, without any attempt to build models.

We ask several questions about the point pattern:

- 1. Does the structure of the point-pattern change over time?
 - Evaluate with intensity, kernel density, G, F, K, L functions.
- 2. Does the point-pattern at one time affect the pattern at a later time?
 - Evaluate with the crossed K function.

And of course the aim is to interpret the answers in terms of the **process** that produced the spatio-temporal point pattern.

We use an example of occurrences of foot-and-mouth disease of cattle from North Cumbria (England), fmd, in the stpp "Spatio-temporal Point Patterns" package.

Task 39: Load the foot-and-mouth disease temporal point-pattern dataset, and the study area boundary northcumbria. Summarize the dataset. • library("stpp")

```
data("fmd")
data("northcumbria")
summary(fmd)
##
                                  ReportedDay
        Х
                        Y
   Min. :295580 Min. :494470 Min. : 28.00
##
## 1st Qu.:327742 1st Qu.:534362
                                 1st Qu.: 51.00
## Median :340625 Median :544235 Median : 60.50
##
  Mean :340190
                  Mean :542980
                                  Mean : 71.83
##
  3rd Qu.:352670
                  3rd Qu.:553052
                                  3rd Qu.: 76.00
## Max. :384530 Max. :575320 Max. :198.00
dim(fmd)
## [1] 648
           3
```

Task 40 : Examine the dataset description. help(fmd)

•

Q30: What are the three fields? How many cases of foot-and-mouth disease were reported? Jump to $A30 \bullet$

```
Task 41 : Display a histogram of the occurrences over time.
hist(fmd[,"ReportedDay"], xlab="reported day",
    main="Cases of Foot-and-mouth disease", breaks=16)
rug(fmd[,"ReportedDay"])
```





Q31 : Describe the temporal pattern of the epidemic. Jump to $A31 \bullet$

Task 42 : Convert the point-pattern to an object of R class stpp "spatio-
temporal point pattern".

The function as.3dpoints performs this conversion:

```
class(fmd)
## [1] "matrix" "array"
fmd <- as.3dpoints(fmd)
class(fmd)
## [1] "stpp"</pre>
```

Task 43 : Plot the occurrence locations, with an indication of the data ofoccurrence.

We show the occurrence by the size of the symbol, stretched from mark.cexmin to mark.cexmin:



For reference, here is the study area from Google Maps. Northern Cumbria county does not include Windermere and further south.



Q32 : Describe the overall spatial point pattern, not considering time of occurrence. $Jump \text{ to } A32 \bullet$

Q33 : Describe the evolution of the spatial point pattern over time. Jump to $A33 \bullet$

Now for some analysis. We will compare the G and F functions for **time-slices** of the point-pattern, and examine their interaction with the crossed K function. The time slices discretize the continuous evolution of the epidemic. In practice these would be set by the epidemiologist according to the presumed process; for convenience we chose 50-day time slices.

Note: You can experiment with different time slices.

Task 44 : Slice the data set into 50-day intervals; report the number of cases in each slice.

Slicing is with the [] selection operator and various logical operators, including < and <=, to form logical conditions.

```
## [1] 648
fmd.1 <- as.3dpoints(fmd[fmd[, 3] <= 50, ])
fmd.2 <- as.3dpoints(fmd[(fmd[, 3] > 50) & (fmd[, 3] <= 100), ])
fmd.3 <- as.3dpoints(fmd[(fmd[, 3] > 100) & (fmd[, 3] <= 150), ])
fmd.4 <- as.3dpoints(fmd[fmd[, 3] > 150, ])
dim(fmd.1)[1]
## [1] 156
```

dim(fmd)[1]

```
dim(fmd.2)[1]
## [1] 404
dim(fmd.3)[1]
## [1] 40
dim(fmd.4)[1]
## [1] 48
```



Jump to $A34 \bullet$

```
Task 45 : Plot each slice's point pattern.
plot(fmd.1, s.region = northcumbria, pch = 21, col = "blue", bg = "red",
    mark = T, mark.col = 0, mark.cexmin = 1, mark.cexmax = 1)
title("Days 0-50")
grid()
```



Days 0-50



Days 51-100



plot(fmd.3, s.region = northcumbria, pch = 21, col = "blue", bg = "red", mark = T, mark.col = 0, mark.cexmin = 1, mark.cexmax = 1) title("Days 101-150") grid()



Days 101-150

plot(fmd.4, s.region = northcumbria, pch = 21, col = "blue", bg = "red", mark = T, mark.col = 0, mark.cexmin = 1, mark.cexmax = 1)

```
title("Days 151-200")
grid()
```



We see an obvious difference in location and clustering.

Task 46 : Compute an owin "point-pattern window" object, in order to compute intensity, G, F and K functions.

w <- owin(poly = list(x = northcumbria[, 1], y = northcumbria[, 2]))

Task 47 : Compute the point-pattern intensity within the window, expressed as cases per km². At the same time, make a class ppp object from the point-pattern.

```
1/(intensity(fmd.1.ppp <- ppp(fmd.1[, 1], fmd.1[, 2], window = w)) * 10^6)
## [1] 35.61729
1/(intensity(fmd.2.ppp <- ppp(fmd.2[, 1], fmd.2[, 2], window = w)) * 10^6)
## [1] 13.75321
1/(intensity(fmd.3.ppp <- ppp(fmd.3[, 1], fmd.3[, 2], window = w)) * 10^6)
## [1] 138.9074
1/(intensity(fmd.4.ppp <- ppp(fmd.4[, 1], fmd.4[, 2], window = w)) * 10^6)
## [1] 115.7562</pre>
```

As shown by the histogram, there is a big difference in intensity between the time slices.

Task 48 : Compare the F "empty space" functions for the four time slices.

```
par(mfrow = c(2, 2))
plot(Fest(fmd.1.ppp), main = "Days 0-50")
plot(Fest(fmd.2.ppp), main = "Days 51-100")
plot(Fest(fmd.3.ppp), main = "Days 101-150")
plot(Fest(fmd.4.ppp), main = "Days 151-200")
par(mfrow = c(1, 1))
```



Q35 : Describe the evolution of the F function over time. Jump to $A35 \bullet$

 $Task \ 49: \ Compare \ the \ G \ ``closest \ point'' \ functions \ for \ the \ four \ time \ slices.$

```
par(mfrow = c(2, 2))
plot(Gest(fmd.1.ppp), main = "Days 0-50")
plot(Gest(fmd.2.ppp), main = "Days 51-100")
plot(Gest(fmd.3.ppp), main = "Days 101-150")
plot(Gest(fmd.4.ppp), main = "Days 151-200")
par(mfrow = c(1, 1))
```



Q36 : Describe the evolution of the G function over time. Jump to $A36 \bullet$

Now we want to evaluate the relation between time slices with the crossed K function. This will reveal if there is any interaction (attraction, dispersion, independence) between patterns. This can then be interpreted by the epidemiologist.

To compute the crossed K function the pattern must be marked.

Task 50 : Combine the four time slices into one marked point pattern. Plotthe marked point pattern.

We do this with the superimpose function, and name the four slices.

```
fmd.all.ppp <- superimpose(Q1 = fmd.1.ppp, Q2 = fmd.2.ppp, Q3 = fmd.3.ppp,
    Q4 = fmd.4.ppp)
plot(fmd.all.ppp, main = "2001 Foot-and-mouth disease, 50-day intervals",
    cex = 0.9, pch = 21, col = 1, bg = 2:5)
```



2001 Foot-and-mouth disease, 50-day intervals

 ${\bf Task}\; {\bf 51}: \ \ {\rm Compute} \ {\rm and} \ {\rm display} \ {\rm the \ crossed} \ {\rm K} \ {\rm function} \ {\rm for \ each \ time \ step}.$

There are three of these.

Kcross.1.2	<-	<pre>Kcross(fmd.all.ppp,</pre>	"Q1",	"Q2")
Kcross.2.3	<-	<pre>Kcross(fmd.all.ppp,</pre>	"Q2",	"Q3")
Kcross.3.4	<-	<pre>Kcross(fmd.all.ppp,</pre>	"Q3",	"Q4")

```
par(mfrow = c(1, 3))
plot(Kcross.1.2, main = "0-50 vs. 51-100")
plot(Kcross.2.3, main = "51-100 vs. 101-150")
plot(Kcross.3.4, main = "101-150 vs. 151-200")
par(mfrow = c(1, 1))
```



Q37 : What is the relation between the point-patterns in successive time slices? Jump to $A37 \bullet$

10 Further reading

Point-pattern analysis is based on theories of point processes. A modern review article is by Møller and Waagepetersen [8]. The text of Diggle [4] presents a detailed explanation and many worked examples of the concepts presented here, and many more. Bivand et al. [2, Ch. 7] presents worked examples of some of these questions; in particular §7.5 presents some applications in spatial epidemiology. Illian et al. [7] present a computational framework for fitting complex spatial point process models using a recently-developed methodology known as INLA. Spatio-temporal point pattern modelling is covered by Diggle [5] and Taylor et al. [11].

11 Answers

A1 : 65 trees.

Return to Q1 \bullet

A2: A square of 5.7 m x 5.7 m; the units have been normalized to $[0 \dots 1]$. Return to $Q2 \bullet$

A4: Yes, they are quite different. The Japanese pines appear to be completely

randomly distributed; the Redwoods clustered, the cells more or less regular (anyway, dispersed). Return to Q4

A5:

(a) The distance at which at least 95% of the points have a neighbour is 0.117 units.

(b) the proportion of points with a neighbour within 0.05 units is 0.388. Return to $Q5 \bullet$

A6 : They all match well, with very little difference between them. Only at the
furthest distance is the empirical function somewhat lower (fewer neighbours than
expected) than expected by CSR.Return to $Q6 \bullet$

A7: They match fairly well, again the discrepancy around G(r) = 0.8 where there are fewer points with first nearest neighbour in that range than expected. Return to Q7•

A8 : For the Redwoods dataset the empirical G function is well above (greater than) the theoretical after a radius of about 0.01 to about 0.05. This indicates strong clustering: nearest neighbours are found at closer distances than expected by CSR. An interesting feature is that there are no neighbours until 0.01 - a very short-range repulsion probably due to the size of an individual tree, making it impossible for two trees to be closer than the size of one tree. Return to $Q8 \bullet$

A10 : Low intensities lead to more sampling error, so high intensities would have
narrower envelopes.Return to Q10 •

A11: With this simulation the empirical G function is almost completely within the envelope throughout, so we can not reject the null hypothesis of CSR. The anomaly near r = 0.1 is clear, indeed at one point of this simulation the empirical value is below the lower limit of the 96% confidence envelope. Return to Q11 •

A12 : The empirical G function for both are well outside the simulation envelopefor much of the radius range. The Redwood trees are almost surely clustered andthe cells almost surely dispersed.Return to Q12 •

A13 : As the bandwidth increases, the maximum intensity decreases: the "hot

spots" are not as "hot" (dense). The minimum intensity in all cases is zero, meaning there are some regions with effectively no probability of a point occurrence. The mean intensity is almost the same; theoretically it should be the same but there are variable edge effects depending on bandwidth. The quartiles show a clear trend: first quartile and median increasing with bandwidth, third quartile decreasing. The intensities are concentrated below the mean at wider bandwidths. This is especially clear with the maximum intensity, which decreases as the counts are averaged across increasingly larger areas. Return to Q13 •

A14 : As the bandwidth increases, the density becomes more uniform . By fourtimes the optimum the density is almost homogeneous, at 50% of the optimum thedensity there are spurious patches. The optimum seems to give a good representa-tion.Return to Q14

•

A15: The K-function for the Japanese pines is quite close to the theoretical for CSR, although slightly below (dispersed) for radiuses around 0.15. The K function for the Redwood trees is consistently above (clustered) at all radii but especially near zero-separation, except for the very close range. The K function for the cells is well below the theoretical for separations to about 0.15; after this is conforms to CSR, meaning that after the initial dispersion to that radius, the number of points within the radius is as expected by CSR. This shows that the dispersion is not on a regular grid. Return to Q15 •

A16 : The convex hull has expanded slightly outward, consistent with average interpoint spacing. The point at the extreme SE controls the SSE and E boundaries, adding a large amount of unsampled area to the polygon.Return to Q16 •

A17 : The removal of extraneous "white space" increases the intensity by 185, 185, 185 %. Return to Q17 •

A18 : For the rectangular window about 0.4; for the polygonal ≈ 0.55 . This isbecause of the smaller area of the rectangle with the same number of points, i.e.,higher average intensity.Return to Q18 •

A19: In the rectangular window the observed proportion matches the theoretical under CSR up to about 80 m, after which the observed is much lower than the theoretical, i.e., much of the area is further from a point than expected under CSR. This is because of the large areas without any points in the rectangle. In the polygonal window the theoretical (under CSR) and actual match well till about 120 m, after which the observed is slightly lower than the theoretical, by about 0.1. Thus there is less area far from the nearest point; this indicates some larger areas of empty space compared to CSR; in this case these are the areas in the SSE and E controlled by the south-easternmost point and not part of the study area. Return to Q19 •

A20: The marks are the cause of each forest fire; there are four classes (causes):

A21 : Yes; for example, the fires caused by lightning are heavily clustered in
the east, with a large space in the centre with almost no fires; by contrast, the
intentional fires are more prevalent in the NW.Return to Q21 •

A22 : Up to about 15 km distances the expected and observed numbers of additional fires from the second process are almost the same; however, beyond that, they are consistently fewer than would be expected by chance, indicating dispersion of one process "caused by" the other. In this case the apparent dispersion may be an artefact of non-stationary intensities of both processes. Return to Q22 •

A23 : The actual cross-K function is very close to the theoretical cross-K functionfrom two unrelated processes. This makes sense because these are two independentpatterns that we superimposed just to show that operation.Return to Q23 •

A24 : The trees appear to be clustered; there are some areas with no trees. Insome sections trees of similar size cluster together but there are also very small treesnear very large (see centre E).Return to Q24 •

A25 : Yes, it appears that there is trend from higher intensity in the NNW tolower in the SSE.Return to $Q25 \bullet$

A26 : The landuse is a somewhat more likely explanation for the observed pattern
of fires than the null model. Scrubland, coniferous forests, and farmland have higher
intensities of the Poisson process. Urban land has less.Return to $Q26 \bullet$

A27 : The model with both land use and Strauss process (interaction) is the
most likely; the model with just land use is quite poor. The combined model is a
bit better than the interaction-only model. There is definitely interaction between
points, i.e., clustering within the 4 km radius. Fires are more likely on scrubland,
conifer forest, and dense forest,Return to Q27 •

A28 : The intensities follow the land use classes (compare with the figure of §7.4).The lowest intensities are in the "urban" and "artificial green" areas (the "coldspot" in the lower-right), the highest in mixed forests.Return to Q28 •

A29: The landuse-only model has both a higher mean and wider range. It is not adjusted to account for inter-point interaction, which reduces the intensity. The Strauss interaction coefficient was 0.76 > 0, indicating local clustering, accounting for some of the intensity within the most susceptible land uses. Return to Q29 •

A30: The three fields are x, y, and x. These are the east and north coördinates in an unspecified CRS, and the reported time of occurrence of a case of foot-and-

mouth disease, in days from an unspecified 0 (maybe day of year 2001). Return to $Q30 \bullet$

A31 : There was a gradual start to the epidemic, then a very strong peak, and along tail with a few additional cases.Return to Q31

A32 : The cases form clear clusters, with a few scattered cases outside of these(e.g., in the SW and NE). The pattern within the clusters seems random. Notethat the blank areas with no cases are probably because no cattle is raised there –the large central area is the Lake District national park, and the eastern edge arethe North Pennine mountains.Return to Q32 •

A33 : The earliest cases (smallest symbols) are in the centre and NW, then there
are cases more towards the NW and W, and finally the most recent cases (largest
symbols) are concentrated in the SE.Return to Q33 •

 A34 :
 0-50 days: 156 cases; 51-100 days: 404 cases; 101-150 days: 40 cases; 151-200 days: 48 cases.

 Return to $Q34 \bullet$

A35:

These all show a longer distance from an arbitrary location in the study area to the nearest case than would be expected by chance. The pattern changes: increasingly strong in the last time-slice, since most of the cases are found only in the SE of the study area. Return to Q35 •

A36: These all show strong clustering: the observed distance to nearest neighbouris well above the theoretical line. Notice the different distance (\mathbf{r}) scales. Theclustering is strongest for the 151-200 day slice: almost all points have a neighbourwithin 2.2 km, whereas for the 101-150 time slice this is not reached until about8 km.Return to Q36 •

A37 : There is a clear repulsion influence of the first slice (0-50 days) on the second (51-100), and the third (101-150) on the fourth (151-200). That is, the nearest point in one pattern is further than expected by chance from the point in the other pattern. The second and third time-slice patterns are almost independent. Return to Q37 •

A Preparing data for point pattern analysis

As shown in §5, objects of class ppp can be converted from "Simple Features" objects of class sf, either with or without attributes which can be used as marks. This uses the as.ppp function.

So, all that is required is to import point data (possibly with attributes) into an **sf** object. There are two common methods:

- Directly from ESRI shapefiles of points, other common geospatial formats such as GPKG ("Geopackage") or GeoJSON, using the st_read function of the sf package (§A.1). See st_drivers for a list of drivers available on your system.
- 2. From data frames imported with the read.table function or its variants such as read.csv for comma-separated values (CSV) files (§A.2).

For records stored in Excel spreadsheets, see the R data import/export FAQ^{6} . The easiest way to import Excel spreadsheets is to first export the sheet from Excel as a CSV file, and follow option (2) below (§A.2).

A.1 Shapefiles and other common geospatial formats

A shapefile is a specification for geospatial data interchange among ESRI and other information systems, and is one of the native formats used by ArcGIS. It consists of three files with the same name and different file extensions: (1) shp for the geometry; (2) shx for the spatial index; (3) dbf for the attribute table. We illustrate how to read a shapefile for of the sample datasets provided with the maptools package, using st_read function to read it.

```
baltim <- sf::st_read(system.file("shapes/baltim.shp", package = "maptools"))</pre>
## Reading layer `baltim' from data source
    `/Library/Frameworks/R.framework/Versions/4.2-arm64/Resources/library/maptools/shapes/baltim.shp
##
##
    using driver `ESRI Shapefile'
## Simple feature collection with 211 features and 17 fields
## Geometry type: POINT
## Dimension:
                XY
## Bounding box: xmin: 860 ymin: 505.5 xmax: 987.5 ymax: 581
## CRS:
                NA
str(baltim)
## Classes 'sf' and 'data.frame': 211 obs. of 18 variables:
## $ STATION : int 1 2 3 4 5 6 7 8 9 10 ...
## $ PRICE : num 47 113 165 104.3 62.5 ...
## $ NROOM : num 4777766867...
             : num 0 1 1 1 1 1 1 1 1 1
  $ DWELL
##
## $ NBATH
           : num 1 2.5 2.5 2.5 1.5 2.5 2.5 1.5 1 2.5 ...
## $ PATIO : num 0 1 1 1 1 1 1 1 1 ...
   $ FIREPL : num 0 1 1 1 1 1 0 1 1 ...
##
             : num 0101001011...
## $ AC
## $ BMENT : num 2 2 3 2 2 3 3 0 3 3 ...
##
   $ NSTOR : num 3 2 2 2 2 3 1 3 2 2 ...
             : num 0222012002..
## $ GAR
## $ AGE
             : num 148 9 23 5 19 20 20 22 22 4 ...
##
   $ CITCOU : num 0 1 1 1 1 1 1 1 1 ...
## $ LOTSZ : num 5.7 279.5 70.6 174.6 107.8 ...
```

```
<sup>6</sup> http://cran.r-project.org/doc/manuals/R-data.html#Reading-Excel-spreadsheets
```

Note: Packages are stored in the directory found with the .libPaths function. The system.file function expands its argument with this, to give a full path and file name.

This is a data.frame as well as an sf obhjec, since it has attributes; these appear to be information on house sales in Baltimore (USA). It does not have a defined coördinate reference system; if the source shapefile has one, it is imported. This is then easily converted to a ppp object with as.ppp. The data frame of attributes is then added as marks with the marks function.

```
require(spatstat)
baltim.ppp <- as.ppp(baltim)</pre>
str(baltim.ppp)
## List of 6
    $ window :List of 4
..$ type : chr "rectangle"
## $ window
##
    ..$ xrange: num [1:2] 860 988
##
##
    ..$ yrange: num [1:2] 506 581
##
    ..$ units :List of 3
##
    ....$ singular : chr "unit"
    ....$ plural : chr "units"
##
##
    .. ..$ multiplier: num 1
    ....- attr(*, "class")= chr "unitname"
##
##
    ..- attr(*, "class")= chr "owin"
##
   $ n
              : int 211
## $ x
               : num [1:211] 907 922 920 923 918 900 918 907 918 897 ...
## $ y
             : num [1:211] 534 574 581 578 574 577 576 576 562 576 ...
##
   $ markformat: chr "vector"
              : int [1:211] 1 2 3 4 5 6 7 8 9 10 ...
##
   $ marks
   - attr(*, "class")= chr "ppp"
##
window(baltim.ppp)
## Marked planar point pattern: 6 points
## marks are numeric, of storage type 'integer'
## window: rectangle = [860, 987.5] x [505.5, 581] units
marks(baltim.ppp) <- st_drop_geometry(baltim)</pre>
summary(marks(baltim.ppp))
                                      NROOM
                                                      DWELL
##
      STATION
                      PRICE
## Min. : 1.0 Min. : 3.50 Min. : 3.000
                                                  Min. :0.0000
## 1st Qu.: 53.5 1st Qu.: 30.95 1st Qu.: 5.000 1st Qu.:0.0000
## Median :106.0 Median : 40.00
                                  Median : 5.000
                                                  Median :1.0000
##
   Mean :106.0
                  Mean : 44.31
                                   Mean : 5.199
                                                   Mean :0.5355
## 3rd Qu.:158.5 3rd Qu.: 53.75
                                  3rd Qu.: 6.000
                                                   3rd Qu.:1.0000
## Max. :211.0 Max. :165.00
                                  Max. :10.000 Max. :1.0000
      NBATH
##
                     PATIO
                                     FIREPL
                                                        AC
## Min. :1.000 Min. :0.0000
                                  Min. :0.0000
                                                   Min. :0.0000
## 1st Qu.:1.000 1st Qu.:0.0000
                                  1st Qu.:0.0000 1st Qu.:0.0000
## Median :1.500
                  Median :0.0000
                                   Median :0.0000
                                                   Median :0.0000
```

Mean :0.2417

3rd Qu.:2.000 3rd Qu.:0.0000 3rd Qu.:0.0000 3rd Qu.:0.0000

Mean :0.2417

Mean :1.573 Mean :0.1469

```
##
         :5.000
                   Max.
                          :1.0000
                                           :1.0000
                                                     Max.
                                                            :1.0000
   Max.
                                    Max.
       BMENT
##
                       NSTOR
                                        GAR.
                                                         AGE
##
         :0.000
                   Min. :1.000
                                         :0.0000
   Min.
                                   Min.
                                                    Min.
                                                          : 0.0
   1st Qu.:2.000
                   1st Qu.:2.000
                                   1st Qu.:0.0000
                                                    1st Qu.: 20.0
##
##
   Median :2.000
                   Median :2.000
                                   Median :0.0000
                                                    Median : 25.0
##
   Mean :1.981
                   Mean :1.905
                                   Mean :0.2512
                                                    Mean : 30.1
##
   3rd Qu.:3.000
                   3rd Qu.:2.000
                                   3rd Qu.:0.0000
                                                    3rd Qu.: 40.0
                   Max. :3.000
                                   Max.
##
         :3.000
                                          :3.0000
                                                    Max. :148.0
    Max.
##
       CITCOU
                        LOTSZ
                                          SQFT
                                                           Х
##
   Min. :0.0000
                    Min. : 5.70
                                     Min. : 5.76
                                                     Min.
                                                           :860.0
   1st Qu.:0.0000
                    1st Qu.: 20.76
                                     1st Qu.:11.02
                                                     1st Qu.:889.0
##
                    Median : 56.25
##
                                     Median :13.44
                                                     Median :910.0
   Median :1.0000
##
   Mean :0.6066
                    Mean : 72.28
                                     Mean :16.43
                                                     Mean :911.6
##
   3rd Qu.:1.0000
                    3rd Qu.: 84.32
                                     3rd Qu.:19.94
                                                     3rd Qu.:933.5
                    Max. :400.37
                                     Max. :47.61
                                                     Max. :987.5
##
   Max.
         :1.0000
##
         Y
##
         :505.5
   Min.
##
   1st Qu.:528.8
##
   Median :544.5
   Mean :544.2
##
##
   3rd Qu.:559.0
##
   Max. :581.0
#
op <- par(no.readonly = TRUE)</pre>
par(mar=rep(0.5, 4))
plot(baltim.ppp, which.marks = "PRICE", axes=T,
    main="Baltimore house sale prices, k$")
par(op)
```



The bounding box is set to the exact limits of the point data set.

Note: The par function retrieves and sets base graphics parameters. Here we reduce the default margins, which are specified with the mar argument to par.

A.2 Text files

A text file to be converted to a point pattern typically has one header line giving the variable names, usually with the two coördinates as the first two columns. The following lines are one record per point, with a number of fields. For a point pattern, at least two fields are needed, i.e., the coördinates. Others (the attributes) are optional. In a CSV file, fields within each record are separated by commas, and text is quoted. But this is only one possible format; the many arguments to read.table (see its help) allow almost any text file format to be read into a data frame.

We use as an example the Meuse dataset meuse of the sp package. We write it to a text file using write.csv, examine its structure, and show how to import it using read.csv.

First the export: data(meuse, package = "sp") class(meuse) ## [1] "data.frame" write.csv(meuse, file = "tmp.csv", row.names = FALSE)

We examine its structure with file.show, but you can also view in any plain-text editor.

file.show("tmp.csv")

```
"x","y","cadmium","copper","lead","zinc","elev","dist","om","ffreq","soil","lime","landuse","dist.m"
181072,333611,11.7,85,299,1022,7.909,0.00135803,13.6,"1","1","1","Ah",50
181025,333558,8.6,81,277,1141,6.983,0.0122243,14,"1","1","1","Ah",30
181165,333537,6.5,68,199,640,7.8,0.103029,13,"1","1","1","Ah",150
181298,333484,2.6,81,116,257,7.655,0.190094,8,"1","2","0","Ga",270
...
180627,330190,2.7,27,124,375,8.261,0.0122243,5.5,"3","3","0","W",40
```

Then the import; although here we just duplicate what we already had in the example data frame, for your own data this would be the point at which you bring your data into R.

Notice that the CSV file has no information on data types; read.table guesses but is not always right. Here it can not determine that ffreq, soil and lime are classes, because they are coded as integer labels. So they must be converted explicitly with as.factor.

\$ zinc : int 1022 1141 640 257 269 281 346 406 347 183 ...
\$ elev : num 7.91 6.98 7.8 7.66 7.48 ...
\$ dist : num 0.00136 0.01222 0.10303 0.19009 0.27709 ...
\$ om : num 13.6 14 13 8 8.7 7.8 9.2 9.5 10.6 6.3 ...
\$ ffreq : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 1 ...
\$ soil : Factor w/ 3 levels "1","2","3": 1 1 1 2 2 2 2 1 1 2 ...
\$ lime : Factor w/ 2 levels "0","1": 2 2 2 1 1 1 1 1 1 1 ...
\$ dist.m : int 50 30 150 270 380 470 240 120 240 420 ...

Convert to a sf object using the st_as_sfethod, specifying the coordinate columns and the CRS (see ?meuse):

```
require(sf)
pp.sf <- st_as_sf(pp, coords = c("x", "y"), crs = 28992)
class(pp.sf)
## [1] "sf" "data.frame"</pre>
```

summary(pp.sf)

##	cadm:	ium	copp	er	le	ad	zi	nc
##	Min.	: 0.200	Min. :	14.00	Min.	: 37.0	Min.	: 113.0
##	1st Qu.	: 0.800	1st Qu.:	23.00	1st Qu.	: 72.5	1st Qu.	: 198.0
##	Median	: 2.100	Median :	31.00	Median	:123.0	Median	: 326.0
##	Mean	: 3.246	Mean :	40.32	Mean	:153.4	Mean	: 469.7
##	3rd Qu.	: 3.850	3rd Qu.:	49.50	3rd Qu.	:207.0	3rd Qu.	: 674.5
##	Max.	:18.100	Max. :	128.00	Max.	:654.0	Max.	:1839.0
##								
##	ele	ev	dis	t		om	ffreq	soil
##	Min.	: 5.180	Min. :	0.00000	Min.	: 1.000	1:84	1:97
##	1st Qu.	: 7.546	1st Qu.:0	0.07569	1st Qu	.: 5.300	2:48	2:46
##	Median	: 8.180	Median :	0.21184	Median	: 6.900	3:23	3:12
##	Mean	: 8.165	Mean :	0.24002	Mean	: 7.478		
##	3rd Qu.	: 8.955	3rd Qu.:	0.36407	3rd Qu	.: 9.000		
##	Max.	:10.520	Max. :	0.88039	Max.	:17.000		
##					NA's	:2		
##	lime	landuse	9	dis	st.m		geom	etry
##	0:111	Length:15	55	Min.	: 10.0	POINT		:155
##	1: 44	Class :ch	naracter	1st Qu	.: 80.0	epsg:2	28992	: 0
##		Mode :ch	naracter	Median	: 270.0	+proj=	=ster	: 0
##				Mean	: 290.3			
##				3rd Qu	.: 450.0			
##				Max.	:1000.0			
##								

Finally, this can be converted to a point pattern with the as.ppp function: pp.m <- as.ppp(pp.sf) class(pp.m)

```
## [1] "ppp"
```

summary(pp.m)

```
## Marked planar point pattern: 155 points
## Average intensity 1.428157e-05 points per square unit
##
## Coordinates are integers
## i.e. rounded to the nearest unit
##
## marks are numeric, of type 'double'
## Summary:
##
    Min. 1st Qu. Median
                           Mean 3rd Qu.
                                           Max.
    0.200 0.800 2.100 3.246 3.850 18.100
##
##
## Window: rectangle = [178605, 181390] x [329714, 333611] units
##
                      (2785 x 3897 units)
## Window area = 10853100 square units
```

```
window(pp.m)
```

```
## Marked planar point pattern: 6 points
## marks are numeric, of storage type 'double'
## window: rectangle = [178605, 181390] x [329714, 333611] units
```

This automatically has marks from the first column in the dataframe. To change the marks to some other factor:

```
summary(marks(pp.m))
##
     Min. 1st Qu. Median
                              Mean 3rd Qu.
                                               Max.
    0.200 0.800
                    2,100
                             3.246 3.850 18.100
##
marks(pp.m) <- pp.sf$soil</pre>
table(marks(pp.m))
##
## 1 2 3
## 97 46 12
op <- par(no.readonly = TRUE)</pre>
par(mar=rep(0.5, 4))
plot(pp.m, cols=c("red","blue","green"),
    pch=20, main="Meuse soil types", axes=T)
par(op)
```



The window boundary can be limited as explained in §5. If a bounding polygon is available, it can be imported and used as the window. I have prepared a crude boundary for this area, by polygonizing the interpolation grid meuse.grid supplied in the sp package, taking the union of the grid cell polygons, and writing the result as a shapefile meuseBoundary. This should have been supplied with this exercise, as a compressed folder with the four files which together make up the ESRI shapefile format.

The st_read function of the sf package can read ESRI shapefiles in to an sf object. The dsn "data source name" argument to st_read for a shapefile is the folder name in which the shapefile is located; in the code below it is

given as ".", i.e., the current working directory⁷; you can change this as you wish. The **layer** "layer name" argument is the name of the shapefile, without extension.

```
meuseBoundary <- sf::st_read(dsn = ".", layer = "meuseBoundary")
class(meuseBoundary)
## Reading layer `meuseBoundary' from data source</pre>
```

```
## `/Users/rossiter/Library/Mobile Documents/com~apple~CloudDocs/Documents/data/edu/PointPatternAnal
## using driver `ESRI Shapefile'
## Simple feature collection with 1 feature and 1 field
## Geometry type: POLYGON
## Dimension: XY
## Bounding box: xmin: 178440 ymin: 329600 xmax: 181560 ymax: 333760
## Projected CRS: Stereographic
## [1] "sf" "data.frame"
```

The polygon shapefile import creates a **sf** object. This can then be converted to an **owin** "observation window" object with the **as.owin** function of the **spatstat.geom** package:

```
class(meuseBoundary)
## [1] "sf" "data.frame"
(meuseBoundary.win <- as.owin(meuseBoundary))
## window: polygonal boundary
## enclosing rectangle: [178440, 181560] x [329600, 333760] units</pre>
```

Finally, this window can be substituted for the original rectangular window by direct assignment to the window field of the ppp object:

pp.m\$window <- meuseBoundary.win</pre>

When this is plotted we see the study area window:

```
op <- par(no.readonly = TRUE)
par(mar=rep(0.5, 4))
plot(pp.m, cols=c("red","blue","green"),
        pch=20, main="Meuse soil types", axes=T)
par(op)</pre>
```

 $^{^7}$ You can see what this is with the ${\tt getwd}$ function.


Challenge: Repeat the analysis of §5 (the F function) with this polygonal window, and compare the results with those from the rectangular bounding box and the window found by the Ripley-Rasson method.

Clean up from this section; this includes removing the temporary file with the unlink function⁸:

unlink("tmp.csv")
rm(meuse, pp, pp.sf, pp.m, meuseBoundary, meuseBoundary.win, op)

 $^{^{8}}$ This name for file deletion is inherited from the Unix operating system.

References

- A. Baddeley and R. Turner. Modelling spatial point patterns in R. In A. Baddeley, P. Gregori, J. Mateu, R. Stoica, D. Stoyan, J. Berger, S. Fienberg, J. Gani, K. Krickeberg, I. Olkin, and B. Singer, editors, *Case Studies in Spatial Point Process Modeling*, volume 185 of *Lecture Notes in Statistics*, pages 23–74. Springer New York, 2006. ISBN 978-0-387-31144-9. 34, 41
- R. S. Bivand, E. J. Pebesma, and V. Gómez-Rubio. Applied Spatial Data Analysis with R. UseR! Springer, 2008. http://www.asdar-book.org/.
 1, 2, 5, 13, 16, 34, 59
- [3] B. N. Boots and A. Getis. *Point pattern analysis*. Number v. 8 in Scientific geography series. Sage Publications, Newbury Park, Calif, 1988. ISBN 0803922450. 1
- [4] P. Diggle. Statistical analysis of spatial point patterns. Arnold, London, 2nd edition, 2003. ISBN 0122158504. 59
- P. J. Diggle. Statistical analysis of spatial and spatio-temporal point patterns. CRC Press, Boca Raton, 3rd edition, 2013. ISBN 978-1-4665-6023-9. 49, 59
- [6] Peter Diggle. Statistical Analysis of Spatial and Spatio-Temporal Point Patterns. CRC Press, Boca Raton, third edition. edition, 2014. ISBN 978-1-4665-6024-6. doi: 10.1201/b15326.
- [7] Janine B. Illian, Sigrunn H. Sorbye, and Havard Rue. A toolbox for fitting complex spatial point process models using integrated nested laplace approximation (inla). Annals of Applied Statistics, 6(4):1499– 1530, Dec 2012. doi: 10.1214/11-AOAS530. 1, 59
- [8] Jesper Møller and Rasmus P. Waagepetersen. Modern statistics for spatial point processes. *Scandinavian Journal of Statistics*, 34(4):643– 684, Dec 2007. doi: 10.1111/j.1467-9469.2007.00569.x. 1, 59
- B. D. Ripley. Modelling spatial patterns. Journal of the Royal Statistical Society. Series B (Methodological), 39(2):172–212, January 1977. 16
- [10] B. D. Ripley and J. P. Rasson. Finding the edge of a Poisson forest. Journal of Applied Probability, 14(3):483–491, September 1977. doi: 10.2307/3213451. 23
- [11] Benjamin M. Taylor, Tilman M. Davies, Barry S. Rowlingson, and Peter J. Diggle. lgcp: an R package for inference with spatial and spatio-temporal log-Gaussian Cox processes. *Journal of Statistical Soft*ware, 52(4), 2013. ISSN 1548-7660. doi: 10.18637/jss.v052.i04. URL http://www.jstatsoft.org/v52/i04/. 49, 59
- [12] Yihui Xie. knitr: Elegant, flexible and fast dynamic report generation with R, 2011. URL http://yihui.name/knitr/. Accessed 21-Mar-2024. 2

Index of Commands

.libPaths, 65 < operator, 52 \leq operator, 52 [] operator, 52 anova.ppm (spatstat.model package), 39, 41 as.3dpoints (stpp package), 50 as.factor, 67 as.owin (spatstat.geom package), 70 as.ppp (spatstat.geom package), 22, 64, 65.68 cells dataset, 4 clmfires dataset, 26, 41 clmfires.extra dataset, 41 contour, 14 convexhull (spatstat.geom package), 23 data.frame class, 65 density.ppp (spatstat.explore package), 14 dsn argument (st_read function), 69 envelope (spatstat.explore package), 11 Fest (spatstat.explore package), 25 file.show, 67 fmd dataset (stpp package), 49 fv class, 7 Gest (spatstat.explore package), 6 getwd, 70 how argument (plot.ppp function), 38 im class, 14, 41, 42 inside.owin (spatstat.geom package), 19 intensity (spatstat.geom package), 24, 36 interaction argument (ppm function), 36, 40 japanesepines dataset, 2, 30 Kcross (spatstat.explore package), 29 Kest (spatstat.explore package), 16, 29, 33 layer argument (st_read function), 70 locations argument (predict.ppm function), 46

logLik (stats package), 41 longleaf dataset, 32 maptools package, 64 mar argument (par function), 66 mark.cexmin argument (plot.stpp function), marks (spatstat.geom package), 22, 65 meuse dataset, 22, 67 meuse.grid dataset, 69 min, 8 ngrid argument (predict.ppm function), 46 northcumbria dataset (stpp package), 49 nrank argument (envelope function), 11 owin (spatstat.geom package), 19 owin class, 3, 19, 23, 55, 70 par, 66 plot, 3, 7, 14, 23, 27, 42 plot.fv (spatstat.explore package), 7 plot.im (spatstat.geom package), 14, 42 plot.ppm (spatstat.model package), 38 plot.ppp (spatstat.geom package), 3, 23 plot.splitppp (spatstat.geom package), 27polynom (spatstat.model package), 37 ppm (spatstat.model package), 34-37, 40, 45ppm class, 36, 45 ppp class, 2, 3, 6, 11, 16, 19, 22, 23, 25, 55, 64, 65, 70 predict.ppm (spatstat.model package), 38, 45read.csv, 64, 67 read.table, 64, 67redwoodfull dataset, 4, 30 ripras (spatstat.geom package), 23 set.seed, 11 sf class, 22, 64, 65, 68-70 sf package, 2, 64, 69 sigma function argument, 14 sp package, 22, 67, 69 SpatRaster class, 42

spatstat package, 2, 3, 26, 41

```
spatstat.explore package, 6, 11, 14, 16,
        25
spatstat.geom package, 14, 20, 22, 24, 27,
       30, 70
spatstat.model package, 34
splancs package, 2
split, 27
split.ppp (spatstat.geom package), 27, 35
st_as_sf (m package), 68
st_as_sf (sf package), 22
st_drivers (sf package), 64
st_read (sf package), 64, 69
stpp class, 50
stpp package, 2, 49
summary (spatstat.geom package), 32
superimpose (spatstat.geom package), 30,
       57
system.file, 65
terra package, 2, 42
trend argument (ppm function), 36, 37
unlink, 71
use.marks argument (plot.ppp function),
        27
W argument (superimpose function), 31
which, 8
Window (f package), 20
Window (spatstat.geom package), 20, 32
write.csv, 67
```