



CENTRE FOR **STOCHASTIC GEOMETRY**
AND ADVANCED **BIOIMAGING**

Lecture Notes

Analysing spatial point patterns in R

Adrian Baddeley, CSIRO and University of
Western Australia

MISCELLANEA

Analysing spatial point patterns in R

Adrian Baddeley
CSIRO and University of Western Australia

`Adrian.Baddeley@csiro.au`
`adrian@maths.uwa.edu.au`

Workshop Notes
Version 4.1
December 2010

Copyright ©CSIRO 2010

Abstract

This is a detailed set of notes for a workshop on *Analysing spatial point patterns in R*, presented by the author in Australia and New Zealand since 2006.

The goal of the workshop is to equip researchers with a range of practical techniques for the statistical analysis of spatial point patterns. Some of the techniques are well established in the applications literature, while some are very recent developments. The workshop is based on **spatstat**, a contributed library for the statistical package R, which is free open source software.

Topics covered include: statistical formulation and methodological issues; data input and handling; R concepts such as classes and methods; exploratory data analysis; nonparametric intensity and risk estimates; goodness-of-fit testing for Complete Spatial Randomness; maximum likelihood inference for Poisson processes; spatial logistic regression; model validation for Poisson processes; exploratory analysis of dependence; distance methods and summary functions such as Ripley's K function; simulation techniques; non-Poisson point process models; fitting models using summary statistics; LISA and local analysis; inhomogeneous K -functions; Gibbs point process models; fitting Gibbs models; simulating Gibbs models; validating Gibbs models; multitype and marked point patterns; exploratory analysis of multitype and marked point patterns; multitype Poisson process models and maximum likelihood inference; multitype Gibbs process models and maximum pseudolikelihood; line segment patterns, 3-dimensional point patterns, multidimensional space-time point patterns, replicated point patterns, and stochastic geometry methods.

These notes require R **version 2.10.0** or later, and **spatstat version 1.21-2** or later.

Acknowledgements

The author gratefully acknowledges countless comments and suggestions from workshop participants and colleagues, and the support of CSIRO MATHEMATICS INFORMATICS AND STATISTICS, THE UNIVERSITY OF WESTERN AUSTRALIA, THE STATISTICAL SOCIETY OF AUSTRALIA, THE NEW ZEALAND STATISTICAL ASSOCIATION, and THE UNIVERSITY OF WAIKATO.

Copyright ©CSIRO Australia 2010

All rights are reserved. Permission to reproduce individual copies of this document for personal use is granted. Redistribution in any other form is prohibited.

The information contained in this document is based on a number of technical, circumstantial or otherwise specified assumptions and parameters. The user must make its own analysis and assessment of the suitability of the information or material contained in or generated from this document. To the extent permitted by law, CSIRO excludes all liability to any party for any expenses, losses, damages and costs arising directly or indirectly from using this document.

Contents

PART I. OVERVIEW	5
1 Introduction	6
2 Statistical formulation	13
3 The R system	18
4 Introduction to spatstat	20
PART II. DATA TYPES & DATA ENTRY	31
5 Objects, classes and methods in R	32
6 Entering point pattern data into spatstat	38
7 Converting from GIS formats	45
8 Windows in spatstat	46
9 Manipulating point patterns	53
10 Pixel images in spatstat	63
11 Tessellations	71
PART III. INTENSITY	77
12 Exploring intensity	78
13 Dependence of intensity on a covariate	82
PART IV. POISSON MODELS	87
14 Tests of Complete Spatial Randomness	88
15 Maximum likelihood for Poisson processes	95
16 Checking a fitted Poisson model	106
17 Spatial logistic regression	112
PART V. INTERACTION	113
18 Exploring dependence between points	114
19 Distance methods for point patterns	115
20 Simulation envelopes and goodness-of-fit tests	132
21 Spatial bootstrap methods	139

22 Simple models of non-Poisson patterns	139
23 Model-fitting using summary statistics	144
24 Exploring local features	148
25 Adjusting for inhomogeneity	149
PART VI. GIBBS MODELS	155
26 Gibbs models	156
27 Fitting Gibbs models	162
28 Validation of fitted Gibbs models	171
PART VII. MARKED POINT PATTERNS	177
29 Marked point patterns	178
30 Handling marked point pattern data	181
31 Exploratory tools for multitype point patterns	187
32 Exploratory tools for marked point patterns	200
33 Multitype Poisson models	204
34 Gibbs models for multitype point patterns	210
PART VIII. HIGHER DIMENSIONS AND OTHER SPATIAL DATA	215
35 Line segment data	216
36 Point patterns in 3D	218
37 Point patterns in multi-dimensional space-time	219
38 Replicated data and hyperframes	221
39 Stochastic geometry	222
40 Further information on spatstat	224
Bibliography	225
Index	228

PART I. OVERVIEW


The first part of the workshop is a quick overview of spatial statistics for point patterns, and a very quick introduction to the software.

1 Introduction

1.1 Types of data

1.1.1 Points

A **point pattern** dataset gives the locations of objects/events occurring in a study region.




The points could represent trees, animal nests, earthquake epicentres, petty crimes, domiciles of new cases of influenza, galaxies, etc.

The points might be situated in a region of the two-dimensional (2D) plane, or on the Earth's surface, or a 3D volume, etc. They could be points in space-time (e.g. earthquake epicentre location and time).


The **spatstat** package was originally implemented for 2D point patterns. However it is being extended progressively to 3D, space-time, and multi-dimensional space-time point patterns (see Sections 36–37).

1.1.2 Marks

The points may have extra information called **marks** attached to them. The mark represents an “attribute” of the point. The mark variable could be *categorical*, e.g. species or disease status:



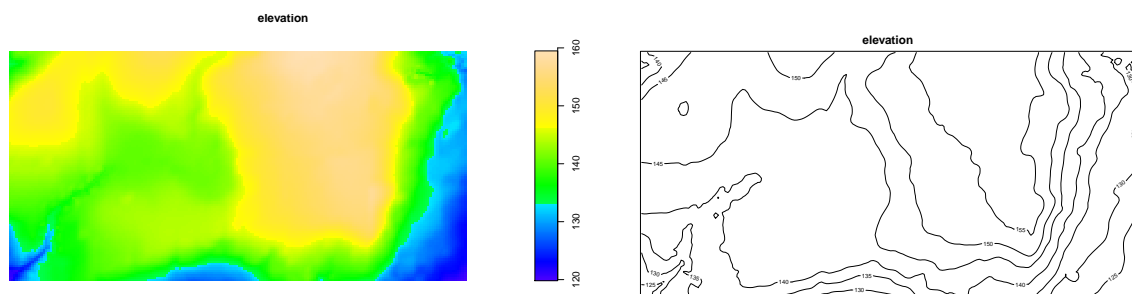
The mark variable could be *continuous*, e.g. tree diameter:



The mark could be multivariate (for example, a tree could be marked by its species and its diameter) or even more complicated.

1.1.3 Covariates

Our dataset may also include **covariates** — any data that we treat as explanatory, rather than as part of the ‘response’. Covariate data may be of any kind. One type of covariate is a *spatial function* $Z(u)$ defined at all spatial locations u , e.g. terrain altitude. Such functions can be displayed as a pixel image or a contour plot:




Another common type of covariate data is a *spatial pattern* such as another point pattern, or a line segment pattern, e.g. a map of geological faults:



1.2 Typical scientific questions


1.2.1 Intensity

‘Intensity’ is the average density of points (expected number of points per unit area). It measures the ‘abundance’ or ‘frequency’ of the events recorded by the points. Intensity may be constant (‘uniform’ or ‘homogeneous’) or may vary from location to location (‘non-uniform’ or ‘inhomogeneous’).




1.2.2 Interaction

‘Interpoint interaction’ is stochastic dependence between the points in a point pattern. Usually we expect dependence to be strongest between points that are close to one another.



Example 1 (Japanese pines) *Locations of 65 saplings of Japanese pine in a 5.7×5.7 metre square sampling region in a natural stand.*

Main question: is the spacing between saplings greater than would be expected for a random pattern? (reflecting competition for resources)




1.2.3 Covariate effects

For a point pattern dataset with covariate data, we typically want to


- investigate whether the intensity depends on the covariates
- allow for covariate effects on intensity before studying interaction between points

Example 2 (Tropical rainforest data) *Locations of 3605 trees in a tropical rainforest, with supplementary grid map of elevation (altitude).*

Main questions: (1) does tree density depend on slope? (2) after accounting for variation in tree density due to slope, is there evidence of clustering of trees?

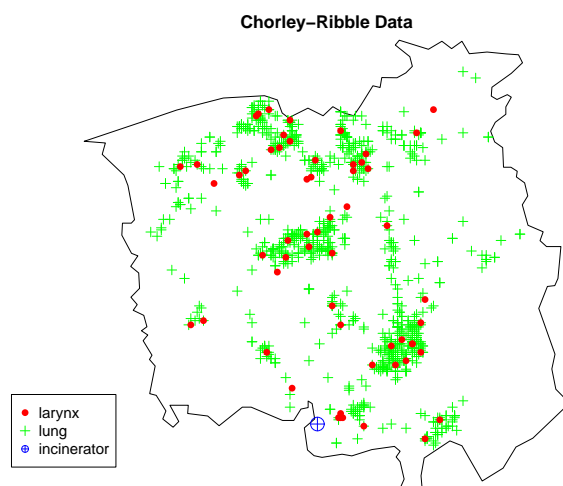


Example 3 (Queensland copper data) *A intensive mineralogical survey yields a map of copper deposits (essentially pointlike at this scale) and geological faults (straight lines). The faults can easily be observed from satellites, but the copper deposits are hard to find. The main question is whether the faults are ‘predictive’ for copper deposits (e.g. copper less/more likely to be found near faults).*



Example 4 (Chorley-Ribble data) *An apparent cluster of cases of cancer of the larynx occurred near a disused industrial incinerator. The area health authority mapped the domicile locations of all cases (58) of cancer of the larynx and, for control purposes, a random sample of cases (978) of lung cancer.*

Main question: after allowing for spatial variation in density of the susceptible population (for which the lung cancer cases are a surrogate), is there evidence of raised incidence of laryngeal cancer near the incinerator?




1.2.4 Segregation of points with different marks

In a marked point pattern, we need to investigate whether points with different mark values are ‘segregated’ (found in different parts of the study region).


Example 5 (Lansing Woods) *In a 20-acre study region in Lansing Woods, Michigan, the locations of 2251 trees and the botanical classification of each tree were recorded.*

Main question: is the study region divided into domains where a single tree species dominates, or are the different species randomly interspersed?



Example 6 (Longleaf Pines) *In a forest of Longleaf Pine trees in Georgia, USA, the locations of 584 trees were recorded along with their diameter at breast height (dbh), a convenient surrogate measure of size and age.*

Main question: explain any spatial variation in the density and age of trees.




1.2.5 Dependence between points of different types

In a point pattern dataset with **categorical** marks, (aka *multitype point pattern*), dependence between the different types may be formulated either as


- interaction between the sub-pattern of points of type i and the sub-pattern of points of type j ; or
- dependence between the mark values of points at two specified locations.

Example 7 (Amacrine cells) *The retina is a flat sheet containing several layers of cells. Amacrine cells occupy two adjacent layers, the ‘on’ and ‘off’ layers. In a microscope field of view, the locations of all amacrine cells were mapped, and classified into ‘on’ and ‘off’.*

Main question: is there evidence that the ‘on’ and ‘off’ layers grew independently of one another?



Example 8 (Ants' nests) *The nests of two species of ants in a plot in Greece were mapped. Auxiliary information records a field/scrub boundary, and the position of a walking track. Main question: does species A intentionally place its nests close to species B?*



1.3 Overview of statistical methods

Statistical methods for spatial point patterns have a quirky history. Although there is a highly-developed branch of probability theory for *point processes*, the corresponding statistical methodology is relatively underdeveloped. Until recently, practical techniques for analysing spatial point patterns were often developed in application areas (notably forestry, ecology, geology, geography and astronomy) rather than in statistical science. Techniques include:

- **summary statistics:** the applied literature is dominated by *ad hoc* methods based on evaluating a summary statistic (e.g. average distance from a point to its nearest neighbour) with very little statistical theory to support them.
- **comparison to Poisson process:** in the applied literature, hypothesis tests are invoked chiefly to decide whether the point pattern is 'completely random' (a uniform Poisson point process) whether or not this is scientifically relevant.
- **modelling:** only in the last decade has it finally become possible to formulate and fit realistic models to point pattern data. There's still a lot of work to be done e.g. in algorithms, model choice, goodness-of-fit.

We'll cover both classical and modern methods. Useful textbooks include [24, 30, 35, 44, 61, 51]. An important recent survey is [50].

2 Statistical formulation

2.1 Point processes

In this workshop, the observed point pattern \mathbf{x} will be treated as a realisation of a random **point process** \mathbf{X} in two-dimensional space. A point process is simply a random set of points; the *number* of points is random, as well as the locations of the points. Our goal is usually to estimate parameters of the distribution of \mathbf{X} .

2.2 Should I treat the data as a point process?

Treating the point pattern as a point process effectively assumes that the pattern is *random* (the locations of the points, and the number of points, are random) and that the pattern is the *observation* or '*response*' of interest. A realisation of a point process is an unordered set of points, so the points do not have a serial order (unless there are marks attached).

Example 9 *A silicon wafer is inspected for defects in the crystal surface, and the locations of all defects are recorded.*

This can be analysed as a point process in two dimensions, assuming the defects are point-like. We're interested in the intensity of defects, spacing between defects, etc.

Example 10 *Earthquake aftershocks in Japan are detected and their latitude, longitude and time of occurrence are recorded.*

This can be analysed as a point process in space-time (where space is the two-dimensional plane or the Earth's surface). If the occurrence times are ignored, it becomes a spatial point process.

Example 11 *The locations of petty crimes that occurred in the past week are plotted on a street map of Chicago.*

This can be analysed as a point process. We're interested in the intensity (propensity for crimes to occur), any spatial variation in intensity, clusters of crimes, etc. One issue here is whether the recorded crime locations can be anywhere in two dimensional space, or whether they are actually restricted to locations on the streets (making them a point process on a 1-dimensional network).

Example 12 *A tiger shark is captured, tagged with a satellite transmitter, and released. Over the next month its location is reported daily. These points are plotted on a map.*

It is probably *not* appropriate to analyse these data as a spatial point process. At the very least, the time of each observation should be included. They could be treated as a space-time point process, except that it's a strange process, as it consists of exactly one point at each instant of time. These data should really be treated as a sparse sample of a continuous trajectory, and analysed using other methods [which, alas, are fairly underdeveloped.] See the R package `trip`.

Example 13 *A herd of deer is photographed from the air at noon each day for 10 days. Each photograph is processed to produce a point pattern of individual deer locations on a map.*

Each day produces a point pattern that could be analysed as a realisation of a point process. However, the observations on successive days are dependent (e.g. constant herd size, systematic foraging behaviour). Assuming individual deer cannot be identified from day to day, this is effectively a ‘repeated measures’ dataset where each response is a point pattern. Methods for this problem are in their infancy.

Example 14 *In a designed controlled experiment, silicon wafers are produced under various conditions. Each wafer is inspected for defects in the crystal surface, and the locations of all defects are recorded as a point pattern.*

This is a designed experiment in which the response is a point pattern. Methods for this problem are in their infancy. There are some methods for *replicated* spatial point patterns [15, 19, 36, 37, 42] that apply when each experimental group contains several point patterns.

Example 15 *The points are not the original data, but were obtained after processing the data. For example,*

- *the original dataset is a pattern of small blobs, and the points are the blob centres;*
- *the original dataset is a collection of line segments, and the points are the endpoints, crossing points, midpoints etc;*
- *the original dataset is a space-filling tessellation of biological cells, and the points are the centres of the cells.*

This is a grey area. Point process methodology can be applied, and may be more powerful or more flexible than existing methodology for the unprocessed data. However the origin of the point pattern may lead to artefacts (for example the centres of biological cells never lie very close together, because cells have nonzero size) which must be taken into account in the analysis.


For more discussion about these topics, see [3].

2.3 Assumptions about the data

The “**standard model**” assumes that the point process \mathbf{X} extends throughout 2-D space, but is observed only inside a region W , the “*sampling window*”. Our data consist of an unordered set

$$\mathbf{x} = \{x_1, \dots, x_n\}, \quad x_i \in W, \quad n \geq 0$$

of points x_i in W . The window W is fixed and known. Usually our goal is inference about parameters of \mathbf{X} .




Data are often supplied without information about the sampling window W . **It is important to know the window W** , since we need to know where points were *not* observed. Even something as simple as estimating the density of points depends on the window. It would be wrong, or at least different, to analyze a point pattern dataset by “guessing” the appropriate window. An analogy may be drawn with the difference between sequential experiments and experiments in which the sample size is fixed *a priori*.

For the same reason, it is not sufficient to observe the values of covariates at the data points only. In order to investigate the dependence of the point process on the covariate, we need to have at least some observations of the covariate at other (“non-data”) locations.

It’s implicitly assumed that all points of \mathbf{X} within W have been mapped without omission.

Most models we use will assume that random points *could* have been observed at any location in the window W , without further constraint. (Examples where this does not apply: GPS locations of cars will usually lie along roads; certain cells lie only inside certain tissues).

When thinking about methodological issues it’s often useful to think about the discretised version of a point process. Suppose the window W is chopped into a large number of tiny ‘pixels’. Each pixel is assigned the value $I = 1$ if it contains a point of \mathbf{X} , and $I = 0$ otherwise. This array of 0’s and 1’s constitutes the data that must be modelled. Thus we need to know where points did *not* occur, as well as where they *did* occur.



To investigate the dependence of these indicators on a covariate, we need to observe the covariate value at some locations where $I = 0$, and not only at locations where $I = 1$.

2.4 Marks and covariates

The main differences between marks and covariates are that

- marks are associated with data points;
- marks are part of the ‘response’ (the point pattern) while covariates are ‘explanatory’.

2.4.1 Marks

A mark variable may be interpreted as an additional coordinate for the point: for example a point process of earthquake epicentre locations (longitude, latitude), with marks giving the occurrence time of each earthquake, can alternatively be viewed as a point process in space-time with coordinates (longitude, latitude, time).

A marked point process of points in space S with marks belonging to a set M is mathematically defined as a point process in the cartesian product $S \times M$. The space M of possible marks may be ‘anything’. In current applications, typically the mark is either a categorical variable (so that the points are grouped into ‘types’) or a real number. Multivariate marks consisting of several such variables are also common.

A marked point pattern is an unordered set

$$\mathbf{y} = \{(x_1, m_1), \dots, (x_n, m_n)\}, \quad x_i \in W, \quad m_i \in M$$

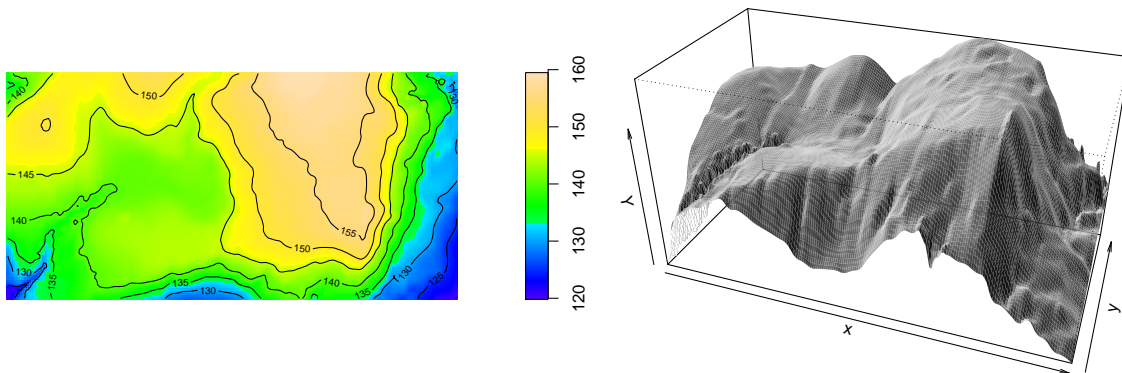
where x_i are the locations and m_i are the corresponding marks.

Marked point patterns are discussed in detail in section 29.

2.4.2 Covariates

Any kind of data may be recruited as an explanatory variable (covariate).


A ‘spatial function’, ‘spatial covariate’ or ‘geostatistical covariate’ is a function $Z(u)$ observable (potentially) at every spatial location $u \in W$. Values of $Z(u)$ may be available for a fine grid of locations u , for example, a terrain elevation map:



The values of a spatial function $Z(u)$ may only be observable at some scattered sampling locations u . An example is the measurement of soil pH at a few sampling locations. In this case, the value of the covariate Z must be observed for all points x_i of the point pattern \mathbf{x} , and must also be observed at some other ‘non-data’ or ‘background’ locations $u \in W$ with $u \notin \mathbf{x}$. You might have to interpolate the observations.

Alternatively, the covariate information may consist of another spatial pattern, such as a point pattern or a line segment pattern. The way in which this covariate information enters the analysis or statistical model depends very much on the context and the choice of model. Typically the covariate pattern would be used to define a surrogate spatial function Z , for

example, $Z(u)$ may be the distance from u to the nearest line segment. Here is a line segment dataset representing the locations of geological faults, and its distance function Z :



3 The R system

We will be using the statistical package R.

3.1 How to obtain R

R is free software with an open-source licence. You can download it from r-project.org and it should be easy to install on any computer (see the instructions at the website).

Books and online tutorials are available to help you learn to use R.

3.2 How commands are printed in the notes

You can run an R session using either a point-and-click interface or a line-by-line command interpreter. In these notes, R commands are printed as they would appear when typed at the command line. So a typical series of R commands looks like this:

```
> pi/2
> sin(pi/2)
> x <- sqrt(2)
> x
```

Note that you are not meant to type the `>` symbol; this is just the prompt for command input in R. To type the first command, just type `pi/2`.

In these notes we will sometimes also print the response that R gives to a set of commands. In the example above, it would look like this:

```
> pi/2

[1] 1.570796
```

```
> sin(pi/2)

[1] 1
```

```
> x <- sqrt(2)
> x

[1] 1.414214
```

If the input is too long, R will break it into several lines, and print the character `+` to indicate that the input continues from the previous line. (You don't type the `+`). Also if you type an expression involving brackets and hit Return before all the open brackets have been closed, then R will print a `+` indicating that it expects you to finish the expression.

```
> folderol <- 1.2
> sin(folderol * folderol * folderol * folderol * folderol * folderol *
+     folderol * folderol * folderol * folderol)

[1] -0.09132148
```

3.3 Contributed libraries for R

In addition to the basic R system, the R website also offers many add-on modules ('libraries' or 'packages') contributed by users. These can be downloaded from the R archive site cran.r-project.org (under 'Contributed Packages').

Packages that may be useful for analysing spatial data are listed under the Spatial Task View (follow the links to *Task Views — Spatial* on cran.r-project.org). For spatial point pattern data, the useful packages include:

<code>adehabitat</code>	habitat selection analysis
<code>ads</code>	spatial point pattern analysis
<code>ash</code>	(includes functions for hexagonal binning)
<code>aspace</code>	'centrographic' analysis of point patterns
<code>DCluster</code>	detecting clusters in spatial count data
<code>ecespa</code>	spatial point pattern analysis
<code>fields</code>	curve and function fitting
<code>geoR</code>	model-based geostatistical methods
<code>geoRglm</code>	model-based geostatistical methods
<code>GeoXp</code>	interactive spatial exploratory data analysis
<code>maptools</code>	geographical information systems
<code>MarkedPointProcess</code>	nonparametric analysis of marked spatial point processes
<code>RArcInfo</code>	interface to ArcInfo system and data format
<code>rgdal</code>	interface to GDAL geographical data analysis
<code>SGCS</code>	spatial graph techniques for detecting clusters
<code>sp</code>	base library for some spatial data analysis packages
<code>sparr</code>	analysis of spatially varying relative risk
<code>spatgraphs</code>	graphs constructed from spatial point patterns
<code>spatialCovariance</code>	spatial covariance for data on grids
<code>spatialkernel</code>	interpolation and segregation of point patterns
<code>spatialsegregation</code>	segregation of multitype point patterns
<code>spatstat</code>	Spatial point pattern analysis and modelling
<code>spBayes</code>	Gaussian spatial process MCMC (grid data)
<code>spdep</code>	spatial statistics for variables observed at fixed sites
<code>spgwr</code>	geographically weighted regression
<code>splanCS</code>	spatial and space-time point pattern analysis
<code>spsurvey</code>	spatial survey methods
<code>trip</code>	spatial trip data formats
<code>tripEstimation</code>	analysis of spatial trip data

To make use of a package, you need to:

1. download the package code (once only) *without unpacking*;
2. 'install' the package code on your system (once only);
3. 'load' the package into your current R session using the command `library` (each time you start a new R session).

The installation step is performed automatically using R, not by manually unpacking the code. Installation is usually a very easy process.

Instructions on how to install a package are given at `cran.r-project.org`. If you are running Windows, first start an R session. Then try the pull-down menu item **Packages — Install packages**. If this menu item is available, then you will be able to download and install any desired packages by simply selecting the package name from the pulldown list. If this menu item is not available (for internet security reasons), you can manually download packages by going to the CRAN website under **Contributed packages --- Windows binaries** and downloading the desired zip files of Windows binary files. To perform step 2, start an R session and use the menu item **Packages — Install from local zip files** to install.

If you are running Linux, step 1 is performed manually by going to the CRAN website under **Contributed Packages** and downloading the tar file `packagename.tar.gz`. Step 2 is performed by issuing the command `R CMD INSTALL packagename.tar.gz`.

4 Introduction to spatstat

4.1 The spatstat package

Spatstat is a contributed R package for analysing spatial data, written by Adrian Baddeley and Rolf Turner. Current versions of **spatstat** deal mainly with **spatial point patterns** in two dimensions. The package supports

- creation, manipulation and plotting of point patterns
- exploratory data analysis
- simulation of point process models
- parametric model-fitting
- hypothesis tests
- residual plots, model diagnostics

Spatstat is one of the largest contributed packages available for R, containing over 1000 user-level functions and a 750-page manual. It has its own web domain, `www.spatstat.org`, offering information about the package.

Spatstat can be downloaded from `cran.r-project.org` (under ‘Contributed packages — **spatstat**’). To install **spatstat** you will also need to download the package **deldir** (some other packages are also recommended but not compulsory).

4.2 Please acknowledge spatstat

If you use **spatstat** for research that leads to publications, it would be much appreciated if you could acknowledge **spatstat** in your publications, preferably citing [10]. Citations help us to justify the expenditure of time and effort on maintaining and developing the package.

4.3 Getting started

Here is a quick demonstration of **spatstat** in action. You can follow the demonstration by typing the commands into R.

To begin any analysis using **spatstat**, first start the R system, and type

```
> library(spatstat)
```

The response will be something like this:

```
deldir 0.0-12
```

```
    Please note: The process for determining duplicated points
    has changed from that used in version 0.0-9 (and previously).
```

```
spatstat 1.21-2
```

```
Type help(spatstat) for an overview of spatstat
```

```
    latest.news() for news on latest version
```

```
    licence.polygons() for licence information on polygon calculations
```

The printout shows that, before loading `spatstat`, the system has loaded the package `deldir` that is required by `spatstat`. Then it loads `spatstat`, showing the version number of the package.

For a list of the commands available in `spatstat`, type

```
> help(spatstat)
```

To get information on a particular *command*, type `help(command)`.

To gain an impression of what is available in `spatstat`, you can run the package demonstration by typing `demo(spatstat)`.

4.4 Licence

The `spatstat` package is free open source software, under the GNU Public Licence.

However, some of the facilities in `spatstat` depend on a polygon geometry package called `gpclib`, and this has a restricted licence, that forbids commercial use. For details, type `licence.polygons()`.

By default, the `gpclib` package is disabled when you start `spatstat`. If you are doing non-commercial work, please enable the polygon clipping library by typing

```
> spatstat.options(gpclib = TRUE)
```

4.5 Inspecting data

For our first demonstration, we'll use one of the standard point pattern datasets that is installed with the package. The 'Swedish Pines' dataset represent the positions of 71 trees in a forest plot 9.6 by 10.0 metres.


```
> data(swedishpines)
```

To avoid typing 'swedishpines' all the time, let us copy the data to another dataset with a shorter name:

```
> X <- swedishpines
```

You can immediately plot the point pattern by typing

```
> plot(X)
```

Simply typing the name of the dataset gives you some basic information:

```
> X
```

```
planar point pattern: 71 points
window: rectangle = [0, 96] x [0, 100] units (one unit = 0.1 metres)
```

Let's study the intensity (density of points) in this point pattern. For a few basic summary statistics, type

```
> summary(X)
```


```
Planar point pattern: 71 points
Average intensity 0.0074 points per square unit (one unit = 0.1 metres)
```

```
Window: rectangle = [0, 96]x[0, 100]units
Window area = 9600 square units
Unit of length: 0.1 metres
```

The coordinates are expressed in decimetres (0.1 metre), so the average intensity is 0.0074 trees per square decimetre or 0.74 trees per square metre.

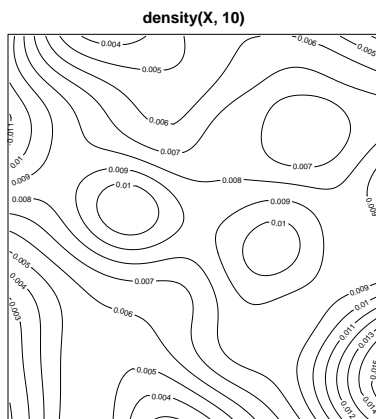
To get an impression of local spatial variations in intensity, we can plot a kernel estimate of intensity:

```
> plot(density(X, 10))
```



where 10 is my chosen value for the standard deviation of the Gaussian smoothing kernel: it is 10 decimetres, i.e. one metre. If you prefer a contour plot,

```
> contour(density(X, 10), axes = FALSE)
```



The contours are labelled in density units of “trees per square decimetre”.

4.6 Exploratory data analysis


Spatstat is designed to support all the standard types of exploratory data analysis for point patterns.

One common example is *quadrat counting*. The study region is divided into rectangles (‘quadrats’) of equal size, and the number of points in each rectangle is counted.

```
> Q <- quadratcount(X, nx = 4, ny = 3)
> Q
```


	x			
y	[0,24]	(24,48]	(48,72]	(72,96]
(66.7,100]	7	3	6	5
(33.3,66.7]	5	9	7	7
[0,33.3]	4	3	6	9

```
> plot(X)
> plot(Q, add = TRUE, cex = 2)
```




Another common example is *Ripley's K function*. I'll explain more about the K function later. For now, we'll just demonstrate how easy it is to compute and plot it. To compute the K function for a point pattern X , type `Kest(X)`. This returns an object which can be plotted.

```
> K <- Kest(X)
> plot(K)
```



In this plot, the empirical K function (solid lines) deviates from the theoretical expected value assuming the points are completely random (dashed lines). To test whether this deviation is statistically significant, the standard approach is to use a Monte Carlo test based on envelopes of the K function obtained from simulated point patterns. In **spatstat** this is done with the `envelope` function:

```
> E <- envelope(X, Kest, nsim = 39)
> plot(E)
```



4.7 Models

The main strength of **spatstat** is that it supports **statistical models** of point patterns. Models can be fitted to point pattern data; the fitted models can be used to summarise the data or make predictions; the fitted models can be simulated (i.e. a random pattern can be generated according to the model); and there are facilities for model selection, for testing whether a term in the model is required (like analysis of variance), and for model criticism (like residuals, regression diagnostics, and goodness-of-fit tests).

Participants in this workshop often say “*I’m not interested in modelling my data; I only want to analyse it.*” However, any kind of data analysis or data manipulation is equivalent to imposing assumptions. We can’t say something is ‘statistically significant’ unless we assume a model, because the p -value is the probability according to a model. The purpose of statistical modelling is to make these assumptions or hypotheses explicit. By doing so, we are able to determine the best and most powerful way to analyse data, we can subject the assumptions to criticism, and we are more aware of the potential pitfalls of analysis. In statistical usage, a model is always tentative; it is assumed for the sake of argument; we might even want it to be wrong. In the famous words of George Box: “All models are wrong, but some are useful.” If you only want to do data analysis without statistical models, your results will be less informative and more vulnerable to critique.

A statistical model for a point pattern is technically termed a *point process* model. Think of a point process as a black box that generates a random spatial point pattern according to some rules. To fit a point process model to a point pattern dataset in **spatstat**, use the function **ppm** (point process model). This is analogous to the standard functions in R for fitting linear models (**lm**), generalized linear models (**glm**) and so on.

```
> data(swedishpines)
> X <- swedishpines
> fit <- ppm(X, ~1, Strauss(9))
> fit
```

Stationary Strauss process

First order term:

```
beta
0.04378316
```

```

Interaction: Strauss process
interaction distance:      9
Fitted interaction parameter gamma:      0.2904

```

Relevant coefficients:


```

Interaction
-1.236324

```


We have fitted a model called the “Strauss point process” to these data. We can generate a simulated realisation of this model:

```
> plot(simulate(fit))
```



We can perform a goodness-of-fit test for this fitted model:

```
> plot(envelope(fit, Kest, nsim = 39))
```



This plot suggests good agreement between the model and the data.

There are many, many other facilities for point process models in **spatstat**, described throughout these notes (mainly in Sections 15–16, 23.1, 27–28 and 34).

4.8 Multitype point patterns

A marked point pattern in which the marks are a categorical variable is usually called a *multitype* point pattern. The ‘types’ are the different values or levels of the mark variable.

Here is the famous Lansing Woods dataset recording the positions of 2251 trees of 6 different species (hickories, maples, red oaks, white oaks, black oaks and miscellaneous trees).

```
> data(lansing)
> lansing
```

```
marked planar point pattern: 2251 points
multitype, with levels = blackoak      hickory      maple      misc      redoak
window: rectangle = [0, 1] x [0, 1] units (one unit = 924 feet)
```

```
> summary(lansing)
```

```
Marked planar point pattern: 2251 points
Average intensity 2250 points per square unit (one unit = 924 feet)
```

```
*Pattern contains duplicated points*
```

```
Multitype:
```

	frequency	proportion	intensity
blackoak	135	0.0600	135
hickory	703	0.3120	703
maple	514	0.2280	514
misc	105	0.0466	105
redoak	346	0.1540	346
whiteoak	448	0.1990	448


```
Window: rectangle = [0, 1]x[0, 1]units
```

```
Window area = 1 square unit
```

```
Unit of length: 924 feet
```

```
> plot(lansing)
```

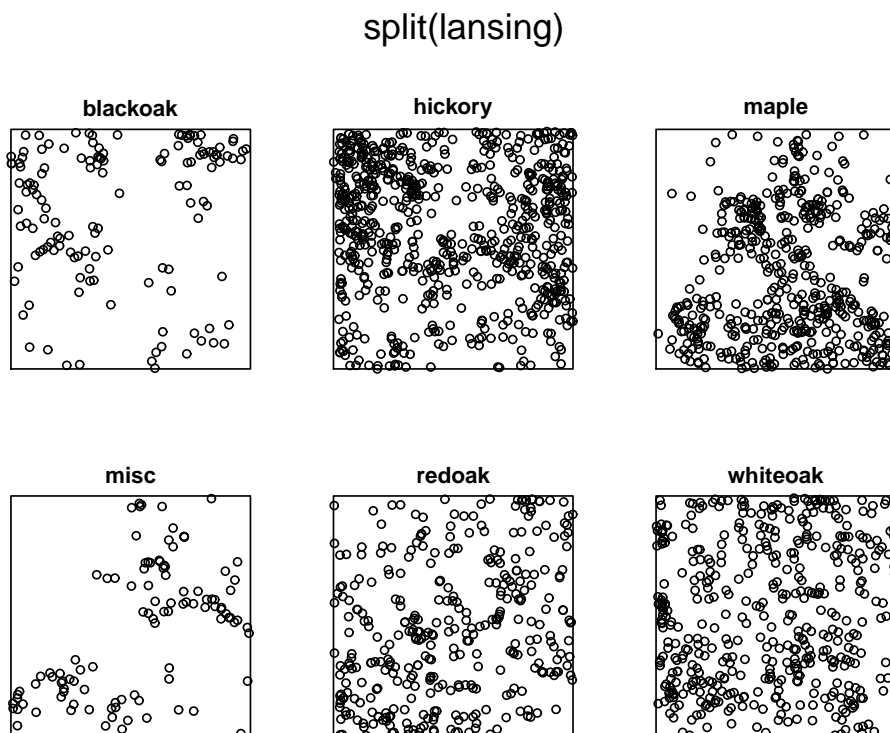
blackoak	hickory	maple	misc	redoak	whiteoak
1	2	3	4	5	6



In this plot, each type of point (i.e. each species of tree) is represented by a different plot symbol. The last line of output above explains the encoding: black oak is coded as symbol 1 (open circle) and so on.

An alternative way to plot these data is to split them into 6 point patterns, each pattern containing the trees of one species. This is done using `split`:


```
> plot(split(lansing))
```



The result of `split(lansing)` is a list of point patterns. The names of the list entries are the names of the types (in this case "blackoak", "hickory", etc). To extract one of these patterns,

e.g. the hickories,





























```
> hick <- split(lansing)$hickory  
> plot(hick)
```






It's also possible to do exploratory analysis and model-fitting for multitype point patterns.

4.9 Installed datasets

For reference, here is a list of the standard point pattern datasets that are supplied with the current installation of **spatstat**:

name	description	marks	covariates	window
amacrine	Hughes' rabbit amacrine cells	2 types	.	
anemones	Upton-Fingleton sea anemones	diameter	.	
ants	Harkness-Isham ant nests	2 species	2 zones	
bei	Tropical rainforest trees	.	topography	
betacells	Wässle et al. cat retinal ganglia	2 types	.	
bramblecanes	Bramble Canes	3 ages	.	
bronzefilter	Bronze particles	diameter	.	
cells	Crick-Ripley biological cells	.	.	
chorley	Chorley-South Ribble cancers	case/control	.	
copper	Queensland copper deposits	.	fault lines	
demopat	artificial data	2 types	.	
finpines	Finnish Pines	diam, height	.	
hamster	Aherne's hamster tumour data	2 types	.	
humberside	Humberside child leukaemia	case/control	.	
japanesepines	Japanese Pines	.	.	
lansing	Lansing Woods	6 species	.	
longleaf	Longleaf Pine trees	diameter	.	
murchison	Murchison gold deposits	.	faults rock type	
nbfires	New Brunswick fires	several	.	
nztrees	Mark-Esler-Ripley NZ trees	.	.	
ponderosa	Getis-Franklin Ponderosa pines	.	.	
redwood	Strauss-Ripley redwood saplings	.	.	
redwoodfull	Strauss redwood map (full set)	.	2 zones	
shapley	Shapley galaxy concentration	several	.	
simdat	Simulated point pattern	.	.	
spruces	Spruce trees in Saxony	diameter	.	
swedishpines	Strand-Ripley Swedish pines	.	.	
urkiola	Urkiola Woods, Spain	2 species	.	

The shape of the window containing the point pattern is indicated by the symbols  (rectangle),  (convex polygon) and  (irregular polygon).

There are also the following datasets which are not 2D point patterns:

name	description	format
heather	Diggle's heather data	binary image (three versions)
osteo	osteocyte lacunae	replicated 3D point patterns with covariates
residualspaper	data from [12]	it's complicated

To flick through a nice display of all these datasets, type **demo(data)**.

To access one of these datasets, type **data(name)** where **name** is the name listed above. To see information about the dataset, type **help(name)**. To plot the dataset, type **plot(name)**.

PART II. DATA TYPES & DATA ENTRY

In Part II of the workshop, we look at the different types of spatial data in **spatstat** (point patterns, windows, pixel images, etc). We explain how to read data into the package and manipulate these data types.

5 Objects, classes and methods in R

The tutorial examples above have used some of the ‘object-oriented’ features of R. It is very useful to know a little about how these work.

5.1 Classes in R

R is an ‘object-oriented’ language. A dataset with some kind of structure on it (e.g. a contingency table, a time series, a point pattern) is treated as a single ‘object’.

For example, R includes a dataset `sunspots` which is a time series containing monthly sunspot counts from 1749 to 1983. This dataset can be manipulated as if it were a single object:

```
> plot(sunspots)
> summary(sunspots)
> X <- sunspots
```

Each object in R is identified as belonging to a particular type or **class** depending on its structure. For example, the `sunspots` dataset is a time series:

```
> class(sunspots)

[1] "ts"
```

Standard operations, such as printing, plotting, or calculating the sample mean, are defined separately for each class of object.

For example, typing `plot(sunspots)` invokes the generic command `plot`. Now `sunspots` is an object of class `"ts"` representing a time series, and there is a special “**method**” for plotting time series, called `plot.ts`. So the system executes `plot.ts(sunspots)`. It is said that the `plot` command is “dispatched” to the method `plot.ts`. The plot method for time series produces a display that is sensible for time series, with axes properly annotated.

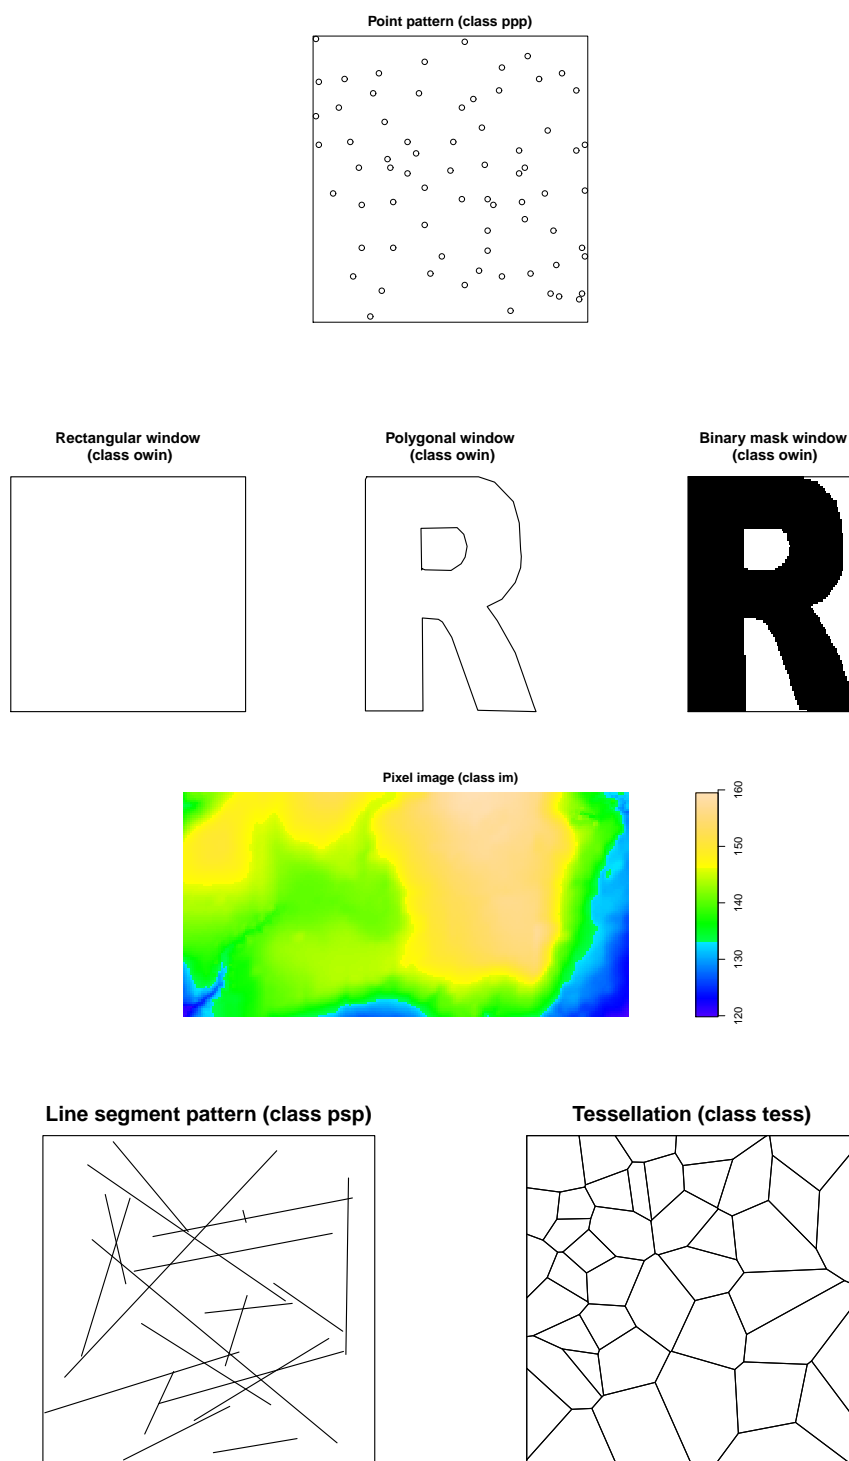
Tip: to find out how to modify the plot for an object of class `"foo"`, consult `help(plot.foo)` rather than `help(plot)`.

5.2 Classes in spatstat

To handle point pattern datasets and related data, the `spatstat` package defines the following important classes of objects:

- `ppp`: planar point pattern
- `owin`: spatial region (‘observation window’)
- `im`: pixel image
- `psp`: pattern of line segments
- `tess`: tessellation


(there are also other classes for specialised use, such as `pp3` for three-dimensional point patterns, `ppx` for multidimensional space-time point patterns, and `hyperframe`).



Most of the functionality in **spatstat** works on such objects. To use this functionality, you'll need to read your raw data into R and then convert it into an object of the appropriate format.

In particular **spatstat** has methods for **plot**, **print** and **summary** for each of these classes. For example, the **plot** method for point patterns, **plot.ppp**, ensures that the x and y scales are equal, and does various other things that are sensible when plotting a spatial point pattern rather than just a list of (x, y) pairs.

```
> data(humberside)
> plot(humberside)
```



Exercise 1 Find out how to modify the command `plot(swedishpines)` so that the title reads “Swedish Pines data” and the points are represented by plus-signs instead of circles.

When you type `print(swedishpines)` or just `swedishpines`, this invokes the generic command `print`, which dispatches to the method `print.ppp`, which prints some sensible information about the point pattern `swedishpines` at the terminal.

```
> swedishpines
```

```
planar point pattern: 71 points
window: rectangle = [0, 96] x [0, 100] units (one unit = 0.1 metres)
```

The generic command `summary` is meant to provide basic summary statistics for a dataset. When you type `summary(swedishpines)` this is dispatched to the method `summary.ppp`, which computes a sensible set of summary statistics for a point pattern, and prints them at the terminal.


```
> summary(swedishpines)
```

```
Planar point pattern: 71 points
Average intensity 0.0074 points per square unit (one unit = 0.1 metres)

Window: rectangle = [0, 96]x[0, 100]units
Window area = 9600 square units
Unit of length: 0.1 metres
```

The command `density` is also generic. It is normally used to compute a kernel density estimate of a probability distribution from a vector of numbers. (This “default method” is called `density.default`.) But there is also a method for point patterns, so that when you type `density(swedishpines)`, this is dispatched to `density.ppp` which computes a two-dimensional kernel estimate of the intensity function.

```
> plot(density(swedishpines, sigma = 10))
```



To see a list of all methods available in R for a particular generic function such as `plot`:

```
> methods(plot)
```

To see a list of all methods that are available for a particular class such as `ppp`:

```
> methods(class = "ppp")
```

[1] <code>affine.ppp</code>	<code>as.data.frame.ppp</code>	<code>as.im.ppp</code>	<code>as.owin.ppp</code>
[5] <code>as.ppp.ppp</code>	<code>bermantest.ppp</code>	<code>by.ppp</code>	<code>closing.ppp</code>
[9] <code>coords<-.ppp</code>	<code>coords.ppp</code>	<code>crossdist.ppp</code>	<code>cut.ppp</code>
[13] <code>density.ppp</code>	<code>dilation.ppp</code>	<code>distfun.ppp</code>	<code>distmap.ppp</code>
[17] <code>duplicated.ppp</code>	<code>envelope.ppp</code>	<code>erosion.ppp</code>	<code>identify.ppp</code>
[21] <code>is.empty.ppp</code>	<code>is.marked.ppp</code>	<code>is.multitype.ppp</code>	<code>kstest.ppp</code>
[25] <code>markformat.ppp</code>	<code>marks<-.ppp</code>	<code>marks.ppp</code>	<code>nnclean.ppp</code>
[29] <code>nn-dist.ppp</code>	<code>nnwhich.ppp</code>	<code>npoints.ppp</code>	<code>opening.ppp</code>
[33] <code>pairdist.ppp</code>	<code>pcf.ppp</code>	<code>periodify.ppp</code>	<code>pixellate.ppp</code>
[37] <code>plot.ppp</code>	<code>[<-.ppp</code>	<code>[.ppp</code>	<code>print.ppp</code>
[41] <code>quadratcount.ppp</code>	<code>quadrat.test.ppp</code>	<code>rebound.ppp</code>	<code>rescale.ppp</code>
[45] <code>rotate.ppp</code>	<code>rshift.ppp</code>	<code>sharpen.ppp</code>	<code>shift.ppp</code>
[49] <code>split<-.ppp</code>	<code>split.ppp</code>	<code>summary.ppp</code>	<code>unique.ppp</code>
[53] <code>unitname<-.ppp</code>	<code>unitname.ppp</code>	<code>unmark.ppp</code>	

5.3 Return values

5.3.1 The return value of a function

Every function in R returns a value. The return value may be ‘null’, or a single number, a list, or any kind of object. When you type an R expression on the command line, the result of evaluating the expression is printed.

```
> 1 + 1
```

```
[1] 2
```

```
> sin(pi/3)
```

```
[1] 0.8660254
```


Just to confuse matters, the result of a function may be tagged as *'invisible'* so that it is not printed.

```
> data(cells)
> plot(cells)
```

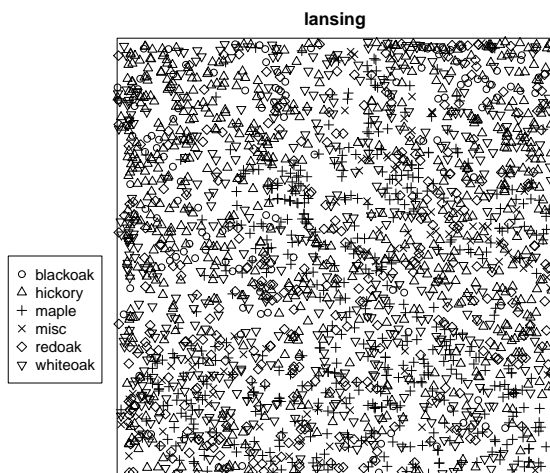
There's still a return value from the function, which can be captured by assigning the result to a variable:

```
> a <- plot(cells)
> a
```

NULL

Tip: Many plotting commands return a value which is useful if you want to annotate the plot. In `spatstat` the function `plot.ppp` plots a point pattern and returns information about the encoding of the marks. After plotting a multitype pattern, to make a nice legend for the plot, save the result of the `plot` call and pass it to the `legend` command:

```
> data(lansing)
> a <- plot(lansing)
> legend(-0.25, 0.5, names(a), pch = a)
```



Tip: To find out the format of the output returned by a particular function `fun`, type `help(fun)` and read the section headed *'Value'*.

5.3.2 Returning an object

A function which performs a complicated analysis of your data will typically return an object belonging to a special class. This is a convenient way to handle calculations that yield large or complicated output. It enables you to store the result for later use, and provides methods for handling the result.

Many of the functions in `spatstat` return an object of a special class. For example, the value returned by `density.ppp` is a pixel image (an object of class "im"). This is effectively a large matrix, giving the values of the kernel estimate of intensity at each point in a fine regular grid of locations.

```
> Z <- density(swedishpines, sigma = 10)
> Z
```

```
real-valued pixel image
100 x 100 pixel array (ny, nx)
enclosing rectangle: [0, 96] x [0, 100] units (one unit = 0.1 metres)
```

The class of pixel images in `spatstat` has methods for `print`, `summary`, `plot` and so on.

```
> summary(Z)

real-valued pixel image
100 x 100 pixel array (ny, nx)
enclosing rectangle: [0, 96] x [0, 100] units
dimensions of each pixel: 0.96 x 1 units
(one unit = 0.1 metres)
Image is defined on the full rectangular grid
Frame area = 9600 square units
Pixel values :
  range = [0.00188947243195949,0.0155470858797917]
  integral = 71.3036909843861
  mean = 0.00742746781087355
```

Another example is the command `Kest` which estimates Ripley's K -function. The value returned by `Kest` is an object of class "fv" ('function value table') containing the estimated values of $K(r)$, obtained using several different estimators, for a range of r values. This class has methods for `print`, `plot` and so on.

```
> u <- Kest(swedishpines)
> u
```

```
Function value object (class fv)
for the function r -> K(r)
Entries:
id      label      description
--      -
r       r          distance argument r
theo    K[pois](r)   theoretical Poisson K(r)
border  K[bord](r)    border-corrected estimate of K(r)
trans   K[trans](r)  translation-corrected estimate of K(r)
iso     K[iso](r)    Ripley isotropic correction estimate of K(r)
-----
```

```
Default plot formula:
. ~ r
```


```
<environment: 0x5f821a0>
```

```
Recommended range of argument r: [0, 24]
```

```
Available range of argument r: [0, 24]
```

```
Unit of length: 0.1 metres
```

```
> plot(u)
```



6 Entering point pattern data into spatstat

To analyse your *own* point pattern data in **spatstat**, you'll need to read the data into R and convert them into an object of class "ppp".

This section explains how to handle 'raw' data in a text file. Section 7 explains how to handle data files in other formats (such as ESRI shapefiles).

6.1 Reading raw data into R

It's good practice to keep a copy of your original data in a text file (where it is not dependent on changes to software, data formats etc). The data can then be loaded into R using standard operations.

Two common formats for the data are

- a comma-separated values (**csv**) file, generated by many spreadsheet packages. To read data from a **csv** file into R, use the command `read.csv`.
- a table format file. The data are arranged in rows and columns, one row for each spatial point, something like this:

Easting	Northing	Diameter
176.111	32.105	10.4
175.989	31.979	7.6
....

The first line of the file is an (optional) *header*. To read these data into R, use the command `read.table`.

To read these datafiles, type either

```
> mydata <- read.csv("mydatafile.csv")
> mydata <- read.table("mydatafile.txt", header = TRUE)
```

In either of these cases, the resulting object `mydata` is a “data frame” in R. You can print the data frame by typing its name.

Individual columns of the data frame can be extracted using `$`. For example

```
> east <- mydata$Easting
```

extracts the `Easting` column of data, and saves it as a vector of numbers, called `east`.

```
> east

[1] 176.111 175.989 176.786 176.394 176.501 175.480 175.041 175.909 176.955
[10] 175.232 176.842 176.752 176.166 175.778 175.176 175.124 175.853 175.866
```

You can also use `scan()` to read a stream of numbers that you type at the keyboard, or `scan(file="filename")` to read a stream of numbers from a file.

6.2 Creating a ppp object

Here is a simple recipe to create a point pattern object from raw data in R.

1. store the x and y coordinates for the points in two vectors `x` and `y`.
2. create two vectors `xrange`, `yrange` of length 2 giving the x and y dimensions of a rectangle that contains all the points.
3. create the point pattern object by

```
> ppp(x, y, xrange, yrange)
```

The value returned by the function `ppp` is an object of class “`ppp`” representing a point pattern inside a rectangle.

If the natural window for the point pattern is **not** a rectangle, then you need to use a command like

```
> ppp(x, y, window = W)
```

where `W` is a window object. See Section 8.5 for details on how to do this.

For example, the following code reads raw data from a text file in table format, and creates a point pattern:

```
> mydata <- read.table("mydatafile.txt", header = TRUE)
> east <- mydata$Easting
> north <- mydata$Northing
> X <- ppp(east, north, c(174, 178), c(29, 33))
```

A slicker way to do the same thing is:

```
> mydata <- read.table("mydatafile.txt", header = TRUE)
> attach(mydata)
> X <- ppp(Easting, Northing, c(174, 178), c(29, 33))
```

6.3 Marks

Recall that a ‘mark’ is an additional attribute of each point in a point pattern. For example, in addition to recording the locations of trees in a forest, we could also record the species, diameter and height of each tree, a chemical analysis of the leaves of each tree, etc.

Suppose `x` and `y` are vectors containing the coordinates of the point locations, as before. If there are marks attached to the points, store the corresponding marks in a vector `m` (with one entry for each point) or in a matrix or data frame `m` (with one row for each point and one column for each mark variable). Recall that a ‘mark’ is an additional attribute of each point in a point pattern. For example, in addition to recording the locations of trees in a forest, we could also record the species, diameter and height of each tree, a chemical analysis of the leaves of each tree, etc.

Suppose `x` and `y` are vectors containing the coordinates of the point locations, as before. If there are marks attached to the points, store the corresponding marks in a vector `m` (with one entry for each point) or in a matrix or data frame `m` (with one row for each point and one column for each mark variable). Then create the marked point pattern by

```
> ppp(x, y, xrange, yrange, marks = m)
```

For example, the following code reads raw data from a text file in table format, and creates a point pattern with a column of numeric marks containing the tree diameters:

```
> mydata <- read.table("mydatafile.txt", header = TRUE)
> attach(mydata)
> X <- ppp(Easting, Northing, c(174, 178), c(29, 33), marks = Diameter)
```

An even slicker way to do this is to convert the data frame directly into a point pattern using the conversion operator `as.ppp`:

```
> mydata <- read.table("mydatafile.txt", header = TRUE)
> X <- as.ppp(mydata, owin(c(174, 178), c(29, 33)))
```


The handling of marks in `spatstat` depends on what type of data they are. Mark values may belong to any of the atomic data types: numeric, integer, character, logical, or complex. They may also be a `factor`, representing categorical values (see below).

Use the commands `is.numeric`, `is.integer`, ..., `is.factor` to test whether your data have the intended type. In a matrix, all the entries have the same atomic type, so you can ask `is.numeric(m)` where `m` is a matrix. A data frame is like a matrix except that different columns may have different types. You can ask `is.numeric(m[,2])` to find out if column 2 in the data frame `m` contains numeric data.

Here is the typical output from `plot.ppp` when the marks are numeric:

```
> data(longleaf)
> plot(longleaf)
```

	0	20	40	60	80
0.000000	1.722522	3.445045	5.167567	6.890090	



The last line of output is the return value from `plot(longleaf)`, which indicates the scale used to plot the marks. The mark value 20 was plotted as a circle of radius 1.72.

6.4 Categorical marks

When the mark is a categorical variable, we have a *multitype point pattern*. The ‘types’ are the different levels of the mark variable. **The mark values should be stored as a ‘factor’ in R.**

Here’s an example of an installed dataset with categorical marks:

```
> data(demopat)
> demopat

marked planar point pattern: 112 points
multitype, with levels = A      B
window: polygonal boundary
enclosing rectangle: [525, 10575] x [450, 7125] furlongs
```

The output (from the `spatstat` function `print.ppp`) indicates that this is a multitype point pattern. Here is the vector of marks:

```
> marks(demopat)

 [1] A B B A B B B A A A B A A B B A A A B B A A A A B B B A A B B B B B A A B
[38] A A B B A A B B B B A B B B B B B B A A A B A B A B B B B B A B B A A B B
[75] B B B A B B A A B A B B B A B A B B B B B A A B A B B B B B A A A B A B B
[112] A
Levels: A B
```

This output (from the base R system) indicates that `marks(demopat)` is a factor with levels A and B.

```
> m <- marks(demopat)
> is.factor(m)

[1] TRUE
```


If the marks are intended to be a categorical variable, ensure that `m` is stored as a **'factor'**.

Here is the typical output from `plot.ppp` when the marks are a factor:

```
> plot(demopat)
```

```
A B
```

```
1 2
```



The last line of output indicates how the marks were plotted: the mark **A** was plotted as symbol 1 (circle) and mark **B** was plotted as symbol 2 (triangle).

Notice that the factor levels are sorted alphabetically (by default). This is one of the common slip-ups with factors in R. To stipulate a different ordering of the levels, do something like

```
> levels(marks(demopat)) <- c("B", "A")
```

Tip: whenever you create a factor `x`, check that the factor levels are as you intended, using `levels(x)`.

Other ways of adding marks to a point pattern will be described in Section 32.

6.5 Multivariate marks

The marks attached to a point pattern may be multivariate (i.e. several variables are attached to each point). For example, the `finpines` point pattern is marked by tree diameter and tree height.

```
> data(finpines)
```

```
> finpines
```

```
marked planar point pattern: 126 points
```


```
Mark variables: diameter, height
```

```
window: rectangle = [-5, 5] x [-8, 2] metres
```

To create such a point pattern, the mark data should be supplied as a data frame. It's important to check that each column of data has the intended type — especially for columns that are intended to be factors.

When a point pattern with multivariate marks is plotted, only one of the columns of marks will be displayed. By default, the first column is selected. You can select another column using the argument `which.marks`.

```
> par(mfrow = c(1, 2))
> plot(finpines)
> plot(finpines, which.marks = "height")
> par(mfrow = c(1, 1))
```



6.6 Checking data

It is prudent to check for quirks in the data.

- Print out the coordinate values and marks to check for errors in data entry, and to determine whether the coordinates have been rounded.
- *Duplicated points* are surprisingly common in data files (i.e. where two records in the file refer to the same (x, y) location). Once you have entered the coordinates into R as a two-column matrix or a data frame *D* say, you can check for duplication using the command `any(duplicated(D))`. If your data are already in the form of a point pattern *X*, you can also type `any(duplicated(X))` to detect duplication. To remove duplicated points, type `Y <- unique(X)`.
- Plotting the point pattern is always wise. Look for unexpected patterns, and points that lie outside the window.
- On a plot of a point pattern *X*, you can identify an individual point by typing `plot(X); identify(X)` then clicking on the point.

The function `ppp` automatically checks for duplicated points, and for points that lie outside the specified window.

6.7 Units

A point pattern *X* may include information about the units of length in which the x and y coordinates are recorded. This information is optional; it merely enables the package to print better reports and to annotate the axes in plots.

If the x and y coordinates in the point pattern *P* were recorded in metres, type

```
> unitname(P) <- c("metre", "metres")
```


at least in Australia or New Zealand. The two strings are the singular and plural forms of the unit. In Scandinavia and Germany you would type

```
> unitname(P) <- "meter"
```

The measurement unit can also be given as some multiple of a standard unit. If, for example, one unit for the x and y coordinates equals 42 centimetres, type

```
> unitname(P) <- list("cm", "cm", 42)
```

Beware that the `unitname` applies only to the coordinates, and not to the marks, of a point pattern.

Altering the `unitname` in an existing dataset is usually not sensible; it simply alters the name of the unit, without changing the entries in the x and y vectors. If you want to convert to different units (e.g. from metres to kilometres or from imperial to metric units), use the command `rescale` as described in Section 9.2.4. If you want to actually change the coordinates by a linear transformation, producing a dataset that is not equivalent to the original one, use `affine`.

6.8 Other ways to make point patterns

To create a point pattern object we can either

- create one from raw data using the function `ppp`
- convert data from other formats (including other packages) using `as.ppp`
- point-and-click on a graphics device using `clickppp`
- read data from a file using `scanpp`
- transform an existing point pattern using a variety of tools
- generate a random pattern using one of the simulation routines
- use one of the standard point pattern datasets supplied with the package.

The package help file `help(spatstat)` lists all the available options.

Note that it is a standard naming convention in R that, for a class `"foo"`, there should be a ‘creator’ function `foo` that creates objects of this class from raw numerical data, and a ‘converter’ function `as.foo` that converts data from other formats into objects of class `"foo"`. We adhere to this convention in `spatstat`:

Class	Creator	Converter
"ppp"	ppp	as.ppp
"owin"	owin	as.owin
"im"	im	as.im

More alternatives for using `ppp` will be covered in Section 8.5.

7 Converting from GIS formats


There is a wide variety of software packages for handling spatial data, especially Geographical Information Systems (GIS). These packages use many different formats to represent spatial data. Typically **spatstat** does not support these formats: this would not be good software design.

Specialised R packages exist for handling different spatial data file formats. The most useful ones are **rgdal**, **shapefiles** and **maptools**. These packages will make it possible for you to read your data from a file into an R session. The **rgdal** package has the most functionality, but can sometimes be difficult to install, as it requires installation of an external library on your system. The packages **shapefiles** and **maptools** have no such difficulty.

The package **sp** provides generic support for spatial data types in R. It enables you to convert between different representations of your data in R.

The usual procedure for converting spatial data is:

1. read your data file into R using a package designed specifically for that file format (e.g. **shapefiles** for ESRI shapefiles), converting it into an R dataset;
2. convert this R dataset into a generic format used by **sp**;
3. convert the generic **sp** format to the required **spatstat** format, using **sp**.



This procedure has to be followed separately for different types of spatial data. Point patterns, windows and pixel images are handled slightly differently. If your point pattern locations are supplied as an ESRI shapefile **mypoints.shp**, then the commands would be

```
> S <- readShapePoints("myfile.shp")
> SP <- as(S, "SpatialPoints")
> P <- as(SP, "ppp")
```

The result is a point pattern (object of class "ppp") in **spatstat**, but you then need to assign the correct window to it.


For further details, see the package vignette *Handling shapefiles in the spatstat package* which is available by typing **vignette("shapefiles")**. For further information on handling GIS formats see [24].

8 Windows in spatstat

Many commands in **spatstat** require us to specify a window, study region or domain. It will be handy to know more about windows in **spatstat**.

An object of class "owin" ("observation window") represents a region or window in two-dimensional space. The window may be

- a rectangle;
- a polygon or polygons, with polygonal holes; or
- an irregular shape represented by a binary pixel image mask.



Objects of this class are created by the function `owin`. There are methods for printing and plotting windows, and numerous geometrical operations.

8.1 Making windows by hand

8.1.1 Rectangular window

To create a rectangular window, type

```
> owin(xrange, yrange)
```

where `xrange`, `yrange` are vectors of length 2 giving the x and y dimensions, respectively, of the rectangle.

```
> owin(c(0, 3), c(1, 2))
```

```
window: rectangle = [0, 3] x [1, 2] units
```

For a square window you can also use `square`:

```
> square(5)
```

```
window: rectangle = [0, 5] x [0, 5] units
```

8.1.2 Circular window

For a circular window use `disc`:

```
> W <- disc(radius = 3, centre = c(0, 0))
```

Currently a circular window is represented as a polygon with a large number of edges.

8.1.3 Polygonal window

Spatstat supports polygonal windows of arbitrary shape and topology. That is, the boundary of the window may consist of one or more closed polygonal curves, which do not intersect themselves or each other. The window may have ‘holes’. Type

```
> owin(poly = p)
```


or

```
> owin(poly = p, xrange, yrange)
```

to create a polygonal window. The argument `poly=p` indicates that the window is polygonal and its boundary is given by the dataset `p`. Note we must use the “name=value” syntax to give the argument `poly`. The arguments `xrange` and `yrange` are optional here; if they are absent, the x and y dimensions of the bounding rectangle will be computed from the polygon.

If the window boundary is a single polygon, then `p` should be a matrix or data frame with two columns, or a list with components `x` and `y`, giving the coordinates of the vertices of the window boundary, **traversed anticlockwise**. For example, the triangle with corners (0,0), (1,0) and (0,1) is created by


```
> Z <- owin(poly = list(x = c(0, 1, 0), y = c(0, 0, 1)))
> plot(Z)
```



Note that polygons should **not** be closed, i.e. the last vertex should **not** equal the first vertex. The same convention is used in the standard plotting function `polygon()`.

If the window boundary consists of several separate polygons, then `p` should be a list, each of whose components `p[[i]]` is a matrix or data frame or a list with components `x` and `y` describing one of the polygons. The vertices of each polygon should be traversed **anticlockwise for external boundaries** and **clockwise for internal boundaries (holes)**. For example, the following creates a triangle with a square hole.


```
> Z <- owin(poly = list(list(x = c(0, 8, 0), y = c(0, 0, 8)), list(x = c(2,
+      2, 3, 3), y = c(2, 3, 3, 2))))
> plot(Z)
```



Notice that the first boundary polygon is traversed anticlockwise and the second clockwise, because it is a hole.

It is often useful to plot a polygonal window with line shading:

```
> plot(Z, hatch = TRUE)
```



8.1.4 Binary mask

A window may be defined by a discrete pixel approximation. Type

```
owin(mask=m, xrange, yrange)
```


to create the window object. Here `m` should be a matrix with logical entries; it will be interpreted as a binary pixel image whose entries are `TRUE` where the corresponding pixel belongs to the window.

The rectangle with dimensions `xrange`, `yrange` is divided into equal rectangular pixels. The correspondence between matrix indices `m[i, j]` and cartesian coordinates is slightly idiosyncratic: the **rows** of `m` correspond to the y coordinate, and the columns to the x coordinate. The entry `m[i, j]` is `TRUE` if the point `(xx[j], yy[i])` (sic) belongs to the window, where `xx`, `yy` are vectors of pixel coordinates equally spaced over `xrange` and `yrange` respectively. The length of `xx` is `ncol(m)` while the length of `yy` is `nrow(m)`.

In some GIS applications the study region will be given as a binary pixel image. A safe strategy is to dump the data from the GIS system to a text file, and read the text file into R using `scan`. Then reformat it as a matrix, and use `owin` to create the window object.

To convert a rectangle or polygonal window to a binary mask, use `as.mask`.

```
> Z <- owin(poly = list(x = c(0, 1, 0), y = c(0, 0, 1)))
> W <- as.mask(Z)
> plot(W)
```



8.2 Converting from GIS formats

If your window (spatial region) is supplied as an “ESRI shapefile” with a name like `myfile.shp`, then type the following:

```
> library(maptools)
> S <- readShapePoly("myfile.shp")
> library(sp)
> SP <- as(S, "SpatialPolygons")
> W <- as(SP, "owin")
```

The `readShapePoly` command reads the file `myfile.shp` and returns an object `S` of class `"SpatialPolygonsDataFrame"`. The next command converts this to an object of class `"SpatialPolygons"` and the last command converts this in turn into a window (object of class `"owin"`) in `spatstat`.

For further information on handling GIS formats see [24] or `vignette("shapefiles")`.

8.3 Functions that make a window

Some functions make a window object. They include

<code>as.owin</code>	Convert other data to a window object
<code>as.polygonal</code>	Convert a window to a polygonal window
<code>as.mask</code>	Convert a window to a binary image mask window
<code>disc</code>	Create a circular window
<code>clickpoly</code>	The user draws a polygon on the screen
<code>bounding.box</code>	Bounding box of a window
<code>bounding.box.xy</code>	Bounding box of a point pattern
<code>convexhull.xy</code>	Convex hull of a point pattern
<code>ripras</code>	Ripley-Rasson estimator of window, given only the points
<code>trim.rectangle</code>	Cut off side(s) of a rectangle
<code>levelset</code>	Level set of a pixel image
<code>solutionset</code>	Solution of an equation involving pixel image(s)
<code>tiles</code>	List of the tiles in a tessellation.

For example, the dataset `bei.extra$elev` is a pixel image containing altitude (elevation) values for a study region. To find the subset where altitude exceeds 145,


```
> elev <- bei.extra$elev
> W <- levelset(elev, 145, ">")
> plot(W)
```



The result *W* is a window.

The accompanying dataset `bei.extra$grad` is a pixel image of the slope (gradient) of the terrain. To find the subset where altitude is below 140 and slope exceeds 0.1,

```
> grad <- bei.extra$grad
> V <- solutionset(elev <= 140 & grad > 0.1)
> plot(V)
```



8.4 Operations on windows

Basic methods for the class "owin" include

<code>print.owin</code>	print short description of a window
<code>summary.owin</code>	print detailed summary of a window
<code>plot.owin</code>	plot a window

Numerous geometrical operations are implemented for window objects. They include:

<code>as.polygonal</code>	Convert a window to a polygonal window
<code>as.mask</code>	Convert a window to a binary image mask window
<code>as.rectangle</code>	Extract the bounding rectangle of a window
<code>area.owin</code>	compute window's area
<code>diameter</code>	compute window's diameter
<code>perimeter</code>	compute window's perimeter length
<code>intersect.owin</code>	intersection of two windows
<code>union.owin</code>	union of two windows
<code>setminus.owin</code>	set difference of two windows
<code>is.subset.owin</code>	determine whether one window contains another
<code>complement.owin</code>	swap inside and outside
<code>bounding.box</code>	Find a tight bounding box for the window
<code>convexhull</code>	Convex hull of a window
<code>is.convex</code>	Test whether a window is convex
<code>rotate</code>	rotate window
<code>shift</code>	translate window
<code>affine</code>	apply affine transformation
<code>rescale</code>	change scale and adjust units
<code>as.mask</code>	convert window to binary image mask
<code>pixellate.owin</code>	convert window to pixel image
<code>dilate.owin</code>	morphological dilation
<code>dilation.owin</code>	morphological dilation
<code>erode.owin</code>	morphological erosion
<code>erosion.owin</code>	morphological erosion
<code>opening.owin</code>	morphological opening
<code>closing.owin</code>	morphological closing
<code>dilated.areas</code>	compute areas of dilated windows
<code>eroded.areas</code>	compute areas of eroded windows
<code>border</code>	create a border region around a window
<code>inside.owin</code>	determine whether a point is inside a window
<code>distmap.owin</code>	distance transform image
<code>distfun.owin</code>	distance transform function
<code>centroid.owin</code>	compute centroid (centre of mass) of window
<code>incircle</code>	find largest circle inside window
<code>simplify.owin</code>	Approximate window by a polygon
<code>deltametric</code>	Measure discrepancy between two windows

8.5 Creating a point pattern in any window

As we saw in Section 6.2, the function `ppp()` will create a point pattern (an object of class "ppp") from raw numerical data in R.

Suppose the x, y coordinates of the points of the pattern are contained in vectors `x` and `y` of equal length. Then

```
ppp(x, y, other.arguments)
```

will create the point pattern. The 'other arguments' must determine a window for the pattern, in one of two ways:

- the other arguments can be passed to `owin` to determine a window:

<code>ppp(x, y, xrange, yrange)</code>	point pattern in rectangle
<code>ppp(x, y, poly=p)</code>	point pattern in polygonal window
<code>ppp(x, y, poly=p, xrange, yrange)</code>	point pattern in polygonal window
<code>ppp(x, y, mask=m, xrange, yrange)</code>	point pattern in binary mask window

- if `W` is a window object (class "owin") then

```
> ppp(x, y, window = W)
```

will create the point pattern.

You may already have a window `W` (an object of class "owin") ready to hand, and now want to create a pattern of points in this window. For example you may want to put a new point pattern inside the window of an existing point pattern `X`; the window is accessed as `X$window`, so type

```
ppp(x, y, window=X$window)
```

9 Manipulating point patterns

Before proceeding, we need to know more about how to manipulate and interrogate point pattern data.

9.1 Point pattern objects

A point pattern is represented in `spatstat` by an object of the class `"ppp"`. This contains the coordinates of the points, optional ‘mark’ values attached to the points, and a description of the study region or spatial ‘window’.

9.1.1 Internal Format

WARNING: It is strongly advisable **NOT** to directly access or modify the internal components of an object.

It is a “beginner’s mistake” to modify the internal components of a structured object such as a point pattern (object of class `"ppp"`). The internal structure of objects in a package can change from one version of the package to another. It is much safer to use operators defined in the package to extract and modify information.

However, in the spirit of “open source”, here is a description of the internal format.

A point pattern object `P` has the following components:

- `P$n` is the number of points (which may be zero).
- `P$x` is a numeric vector containing the x coordinates of the points. Its length equals `P$n` (and may be zero).
- `P$y` is a numeric vector containing the y coordinates of the points. Its length also equals `P$n`.
- `P$marks` contains the marks. It is either `NULL`, or a vector of length `P$n` containing the mark values, or a data frame with `P$n` rows containing the mark values. The entries of `P$marks` may be of any atomic type (character, numeric, integer, logical, complex) or factor.
- `P$window` is an object of class `"owin"` (“observation window”) determining the study region or spatial ‘window’.

It is possible to extract these components individually; for example, to make a histogram of the x coordinates you could just type `hist(P$x)`. However, **do not assign values to these components directly**, or you may create inconsistencies in the data which cause `spatstat` to crash.

To extract or manipulate the data in a point pattern object, use the functions provided in the package. Important ones are:

<code>npoints(X)</code>	number of points in <code>X</code>
<code>marks(X)</code>	marks of <code>X</code>
<code>coords(X)</code>	coordinates of points in <code>X</code>
<code>as.owin(X)</code>	window of <code>X</code>
<code>as.data.frame(X)</code>	coordinates and marks of <code>X</code>
<code>marks(X) <- value</code>	change the marks of <code>X</code>
<code>coords(X) <- value</code>	change the coordinates of <code>X</code>
<code>X[W]</code>	change the window of <code>X</code>

9.1.2 A point pattern needs a window

Note especially that, when you create a new point pattern object, you need to specify the spatial region or window in which the pattern was observed. In `spatstat`, the observation window is an integral part of the point pattern. A point pattern dataset consists of knowledge about where points were *not* observed, as well as the locations where they *were* observed. Even something as simple as estimating the intensity of the pattern depends on the window of observation. It would be wrong, or at least different, to analyze a point pattern dataset by “guessing” the appropriate window (e.g. by computing the convex hull of the points). An analogy may be drawn with the difference between sequential experiments and experiments in which the sample size is fixed *a priori*.

Often, the window of observation is a rectangle, so this requirement just means that we have to specify the x and y dimensions of the rectangle when we create the point pattern. Windows with a more complicated shape can easily be represented in `spatstat`, as described below.

For situations where the window is really unknown, `spatstat` provides the function `ripras` to compute the Ripley-Rasson estimator of the window, given only the point locations.

9.1.3 Order of points

Although a point pattern should (in principle) be treated as an unordered set, the coordinates are obviously stored in a particular order, and can be addressed using that order.

```
> data(longleaf)
> as.data.frame(longleaf)[1:5, ]
```

	x	y	marks
1	200.0	8.8	32.9
2	199.3	10.0	53.5
3	193.6	22.4	68.0
4	167.7	35.6	17.7
5	183.9	45.4	36.9

If the marks are a categorical variable, then `marks(P)` is a **factor**.

```
> data(chorley)
> as.data.frame(chorley)[55:60, ]
```

	x	y	marks
55	355.6	413.9	larynx
56	355.5	413.9	larynx
57	355.7	413.9	larynx
58	355.6	414.1	larynx
59	359.0	417.3	lung
60	353.1	426.9	lung

```
> type <- marks(chorley)
> is.factor(type)
```

```
[1] TRUE
```

```
> levels(type)
```

```
[1] "larynx" "lung"
```

```
> table(type)
```

```
type
larynx  lung
    58   978
```

9.2 Operations on ppp objects

Directly manipulating the entries inside an object is not safe. It is also unnecessary, because these manipulations can be performed using functions or operators.

For point patterns (objects of class "ppp") there are the following operations.

9.2.1 Extracting and altering data

<code>npoints(X)</code>	number of points in <code>X</code>
<code>marks(X)</code>	marks of <code>X</code>
<code>coords(X)</code>	coordinates of points in <code>X</code>
<code>as.owin(X)</code>	window of <code>X</code>
<code>as.rectangle(X)</code>	bounding rectangle of <code>X</code>
<code>as.data.frame(X)</code>	coordinates and marks of <code>X</code>
<code>marks(X) <- value</code>	change the marks of <code>X</code>
<code>coords(X) <- value</code>	change the coordinates of <code>X</code>
<code>X[W]</code>	change the window of <code>X</code>

9.2.2 Extracting a subset of a point pattern

Recall that in R the subset operator is `[]`. If `x` is a vector of numbers, then `x[s]` extracts an element or subset of `x`. The subset index `s` can be

- a positive integer: `x[3]` means the third element of `x`;
- a vector of positive integers indicating which elements to extract: `x[c(2,4,6)]` extracts the 2nd, 4th and 6th elements of `x`;
- a vector of negative integers indicating which elements *not* to extract: `x[-1]` means all elements of `x` except the first one;
- a vector of logical values, of the same length as `x`, with each `TRUE` entry of `s` indicating that the corresponding entry of `x` should be extracted, and `FALSE` indicating that it should not be extracted. For example `x[x > 3.1]` extracts those elements of `x` which are greater than 3.1.

To extract a subset of a point pattern in `spatstat`, we also use the subset operator `[]`. If `X` is a point pattern then `X[s]` is also a point pattern, consisting of those points of `X` selected by the subset index `s`, where `s` can be any of the three types listed above, (Recall that the points in a point pattern object are stored in a particular order; this is the order in which they are indexed by `s`.)

```
> data(bei)
> bei
```

```
planar point pattern: 3604 points
window: rectangle = [0, 1000] x [0, 500] metres
```

```
> bei[1:10]
```

```
planar point pattern: 10 points
window: rectangle = [0, 1000] x [0, 500] metres
```

It is also possible to extract the subset defined by a spatial region. If **X** is a point pattern and **W** is a spatial window (object of class "owin") then **X[W]** is the point pattern consisting of all points of **X** that lie inside **W**.

```
> W <- owin(c(100, 800), c(100, 400))
> W
```

```
window: rectangle = [100, 800] x [100, 400] units
```

```
> bei[W]
```

```
planar point pattern: 918 points
window: rectangle = [100, 800] x [100, 400] units
```

Tip: You may need to put quotes around the subset operator in some contexts. The generic subset operator is `[]` but the help file is summoned by typing `help("[]")`. The subset method for point patterns is called `[],ppp` but the help file is summoned by typing `help("[],ppp")`.

The command `split.ppp` allows you to divide a point pattern into sub-patterns, and the command `by.ppp` allows you to perform an operation on each sub-pattern.

9.2.3 Fiddling with marks

To extract the marks from a point pattern, use `marks`:

```
> m <- marks(X)
```

To add or change marks, use `marks<-`

```
> marks(X) <- whatever
```

To delete marks from a point pattern, assign the marks to `NULL`:

```
> marks(X) <- NULL
```

For convenience, you can also perform these operations inside an expression, using the function `unmark` to remove marks and the binary operator `%mark%` to add marks:

```
> data(redwood)
> radii <- rexp(redwood$n, rate = 10)
> X <- redwood %mark% radii
> X
```

```
marked planar point pattern: 62 points
marks are numeric, of type double
window: rectangle = [0, 1] x [-1, 0] units
```

```
> unmark(X)
```

```
planar point pattern: 62 points
window: rectangle = [0, 1] x [-1, 0] units
```

For a point pattern with real-valued marks, the method `cut.ppp` for the generic function `cut` will divide the range of mark values into several discrete bands, yielding a point pattern with categorical marks:

```
> Y <- cut(X, breaks = 3)
> Y <- cut(X, breaks = c(0, 1, 10, Inf))
> Y
```

```
marked planar point pattern: 62 points
multitype, with levels = (0,1]      (1,10]      (10,Inf]
window: rectangle = [0, 1] x [-1, 0] units
```

9.2.4 Changing scales and units

A scalar dilation can be applied using `affine`. For example, the Swedish Pines data were recorded in decimetres. To convert the coordinates to metres, we could type

```
> data(swedishpines)
> X <- affine(swedishpines, mat = diag(c(1/10, 1/10)))
> unitname(X) <- c("metre", "metres")
> X
```

```
planar point pattern: 71 points
window: rectangle = [0, 9.6] x [0, 10] metres
```

The command `rescale` performs the same function:

```
> data(swedishpines)
> X <- rescale(swedishpines, 10)
> X
```

```
planar point pattern: 71 points
window: rectangle = [0, 9.6] x [0, 10] metres
```

Beware that this does not change the marks in the point pattern. If your marks represent tree diameter and you want to rescale them as well, this must be done by hand.

9.2.5 Geometrical transformations

The commands `rotate`, `shift` and `affine` apply two-dimensional rotation, vector shifts, and affine transformations, respectively.

9.2.6 Random perturbations of a point pattern

It is sometimes useful to randomise the data, for example for hypothesis testing. The command `rshift` will apply the same random shift to each point, while `rjitter` will apply a different random shift to each point. The command `quadratresample` performs a block resampling procedure in which the window is divided into rectangles and these rectangles are randomly resampled.


9.3 Example

We will use one of the standard point pattern datasets that is installed with the package. The NZ trees dataset represent the positions of 86 trees in a forest plot 153 by 95 feet.

```
> data(nztrees)
> nztrees


planar point pattern: 86 points
window: rectangle = [0, 153] x [0, 95] feet

> plot(nztrees)
```



To get an impression of local spatial variations in intensity, we plot a kernel density estimate of intensity.

```
> contour(density(nztrees, 10), axes = FALSE)
```



The density surface has a steep slope at the top right-hand corner of the study region. Looking at the plot of the point pattern itself, we can see a cluster of trees at the top right.

You may also notice a line of trees at the right-hand edge of the study region. It looks as though the study region may have included some trees that were planted as a boundary or avenue. This sticks out like a sore thumb if we plot the x coordinates of the trees:

```
> hist(nztrees$x, nclass = 25)
```



We might want to exclude the right-hand boundary from the study region, to focus on the pattern of the remaining trees. Let's say we decide to trim a 5-foot margin off the right-hand side.

First we create the new, trimmed study region:

```
> chopped <- owin(c(0, 148), c(0, 95))
```

or more slickly,

```
> win <- nztrees$window
> chopped <- trim.rectangle(win, xmargin = c(0, 5), ymargin = 0)
> chopped
```

```
window: rectangle = [0, 148] x [0, 95] feet
```

(Notice that `chopped` is not a point pattern, but simply a rectangle in the plane.)

Then, using the subset operator `[.ppp]`, we simply extract the subset of the original point pattern that lies inside the new window:

```
> nzchop <- nztrees[chopped]
```

We can now study the 'chopped' point pattern:

```
> summary(nzchop)
```

```
Planar point pattern: 78 points
```

```
Average intensity 0.00555 points per square foot
```


```
Window: rectangle = [0, 148]x[0, 95]feet
```

```
Window area = 14060 square feet
```

```
Unit of length: 1 foot
```

```
> plot(density(nzchop, 10))
```

```
> plot(nzchop, add = TRUE)
```

Removing the right margin seems to have produced a much more uniform pattern.

9.4 Splitting and combining point patterns

Sometimes it is useful to split a point pattern dataset into several sub-patterns, and perform some calculations on each sub-pattern.

9.4.1 Splitting a point pattern into sub-patterns

The powerful R command `split` has a method for point patterns. This enables the user to divide a point pattern into sub-patterns using any suitable criterion.

- If `X` is a marked point pattern, and the marks are a factor, then `split(X)` separates the data points into different point patterns according to their mark value.
- If `Z` is a pixel image with factor values, then `split(X,Z)` separates the data points into different point patterns according to the pixel value of `Z` at each point.
- If `Z` is a tessellation, then `split(X,Z)` separates the point pattern `X` into sub-patterns delineated by the tiles of `Z`.

In each case the result is a list of point patterns. You can then use the R command `lapply` to perform any desired operation on each element of the list. For example, to apply adaptive estimation of intensity to each species of tree in the Lansing Woods data,

```
> data(lansing)
> V <- split(lansing)
> A <- lapply(V, adaptive.density)
> plot(as.listof(A))
```

A neater way to operate on sub-patterns is to use `by.ppp`, a method for the R function `by`. The call `by(X, INDICES=Z, FUN=f)` is essentially equivalent to `lapply(split(X,Z), f)`. It splits the dataset `X` into sub-patterns according to `Z`, then applies the function `f` to each sub-pattern. So to apply adaptive estimation of intensity to each species of tree in the Lansing Woods data,

```
> data(lansing)
> A <- by(lansing, FUN = adaptive.density)
> plot(A)
```

9.4.2 Combining point patterns

Any number of point patterns can be combined to make a single pattern, using `superimpose`.

```
> X <- runifpoint(20)
> Y <- runifpoint(10)
> superimpose(X, Y)
```

```
planar point pattern: 30 points
window: rectangle = [0, 1] x [0, 1] units
```

The argument `W`, if given, specifies the window for the combined point pattern.

```
> superimpose(X, Y, W = square(2))
```

```
planar point pattern: 30 points
window: rectangle = [0, 2] x [0, 2] units
```

To attach a separate mark to each component pattern, use argument names:

```
> superimpose(Hooray = X, Boo = Y)
```

```
marked planar point pattern: 30 points
multitype, with levels = Hooray      Boo
window: rectangle = [0, 1] x [0, 1] units
```

9.5 List of operations on point patterns

Here's a summary of basic operations available for a point pattern `X`.

<code>X</code>	print basic info
<code>print(X)</code>	print basic info
<code>summary(X)</code>	print detailed summary
<code>npoints(X)</code>	number of points
<code>coords(X)</code>	extract coordinates of points
<code>coords(X)<-value</code>	assign new coordinates to points
<code>marks(X)</code>	extract marks from point pattern
<code>marks(X)<-value</code>	assign new marks to point pattern
<code>unmark(X)</code>	remove marks
<code>marks(X)<-NULL</code>	remove marks
<code>as.owin(X)</code>	extract window of point pattern
<code>X[subset]</code>	subset of point pattern
<code>plot(X)</code>	plot a point pattern
<code>superimpose(X1, X2,...)</code>	combine several point patterns
<code>duplicated(X)</code>	detect duplicated points
<code>unique(X)</code>	remove duplicated points
<code>identify(X)</code>	point-and-click to identify individual points
<code>cut(X, ...)</code>	classify points into groups
<code>split(X)</code>	divide pattern into sub-patterns
<code>by(X, ...)</code>	apply function to sub-patterns
<code>discretise(X)</code>	discretise coordinate values
<code>pixellate(X)</code>	approximate point pattern by pixel image
<code>as.im(X)</code>	approximate point pattern by pixel image
<code>rotate(X, ...)</code>	rotate entire point pattern and window
<code>shift(X, ...)</code>	shift entire point pattern and window
<code>affine(X, ...)</code>	affine transformation
<code>density(X)</code>	kernel smoothed intensity estimate
<code>smooth.ppp(X)</code>	spatial interpolation of mark values
<code>convexhull(X)</code>	convex hull of point pattern
<code>delaunay(X)</code>	Delaunay triangulation of point pattern
<code>dirichlet(X)</code>	Dirichlet-Voronoi tessellation based on point pattern
<code>periodify(X)</code>	make several translated copies of point pattern
<code>rlabel(X)</code>	random re-labelling of multitype point pattern
<code>rshift(X)</code>	random shifting of points

10 Pixel images in spatstat

An object of class "im" represents a pixel image. It specifies a rectangular grid of locations ("pixels") in two dimensional space, and a numerical value for each pixel. The pixel values can be real numbers, integers, complex numbers, single characters or strings, logical values or categorical values. A pixel's value can also be NA, meaning that it is not defined at that location.

A pixel image represents a spatial function $Z(u)$ in many different contexts. It may contain experimental data (such as a map of terrain elevation) or computed values (such as a kernel estimate of point process intensity) or it may be directly obtained from a camera (such as a satellite image).

10.1 Creating a pixel image

10.1.1 Creating an image from raw data

To create a pixel image from raw data, use `im`:


```
> im(mat, xcol, yrow)
```

where `mat` is a matrix containing the pixel values. The pixel values could have been generated by hand, or read from a file.

The correspondence between matrix indices `mat[i,j]` and cartesian coordinates is slightly idiosyncratic: the **rows** of `m` correspond to the y coordinate, and the columns to the x coordinate.

The argument `xcol` is a vector of equally-spaced x coordinate values corresponding to the **columns** of `mat`, and `yrow` is a vector of equally-spaced y coordinate values corresponding to the **rows** of `mat`. These vectors determine the spatial position of the pixel grid. The length of `xcol` is `ncol(mat)` while the length of `yrow` is `nrow(mat)`. If `mat` is not a matrix, it will be converted into a matrix with `nrow(mat) = length(yrow)` and `ncol(mat) = length(xcol)`.

```
> vec <- seq(-5, 5, length = 1200) + rnorm(1200)
> mat <- matrix(vec, nrow = 30, ncol = 40)
> noisy <- im(mat, xcol = seq(0, 4, length = 40), yrow = seq(0,
+   3, length = 30))
> plot(noisy)
```



10.1.2 Factor valued images

For some strange reason, R does not allow matrices with categorical (factor) values, and many operations that create factors in R will convert a matrix to a vector.

```
> cutvec <- cut(mat, 3)
> is.factor(cutvec)
```

```
[1] TRUE
```

```
> is.matrix(cutvec)
```

```
[1] FALSE
```

Although `mat` was a matrix, `cutvec` is a vector, with factor values.

To create a pixel image with categorical values, leave the pixel values as a vector, and let the `im` reshape it:

```
> cutnoise <- im(cutvec, xcol = seq(0, 1, length = 40), yrow = seq(0,
+      1, length = 30))
> cutnoise
```

factor-valued pixel image


factor levels:

```
[1] "(-7.33,-2.43]" "(-2.43,2.47]" "(2.47,7.37]"
```

30 x 40 pixel array (ny, nx)

enclosing rectangle: [-0.012821, 1.0128] x [-0.017241, 1.0172] units

```
> plot(cutnoise)
```



[Another alternative is to create an integer-valued matrix, and assign a `levels` attribute to it. This will be interpreted as a matrix with categorical values.]

10.1.3 Converting a function to an image

The command `as.im` will convert other types of data to a pixel image.

A function $f(x,y)$ can be converted into a pixel image. This makes it easy to create a pixel image in which the pixel values are defined by an algebraic formula in the x and y coordinates.

```
> f <- function(x, y) {
+   x^2 + y^2
+ }
> w <- owin(c(-1, 1), c(-1, 1))
> Z <- as.im(f, w)
```

The second argument of `as.im` is a window object (class "owin") specifying the domain of the image.

10.1.4 Functions that return a pixel image

Functions that return an object of class "im" include:

<code>as.im</code>	converts other data to a pixel image
<code>density.ppp</code>	kernel smoothing of point pattern
<code>density.psp</code>	kernel smoothing of line segment pattern
<code>pixellate.ppp</code>	approximate point pattern by pixel image
<code>pixellate.psp</code>	approximate line segment pattern by pixel image
<code>pixellate.owin</code>	approximate window by pixel image
<code>distmap.owin</code>	distance function of window
<code>distmap.ppp</code>	distance function of point pattern
<code>distmap.psp</code>	distance function of line segment pattern
<code>setcov</code>	geometric covariance function of a window
<code>connected</code>	identify connected components of a window
<code>predict.ppm</code>	fitted intensity of a point process model
<code>[.im</code>	subset of an image (or look up pixel values)
<code>shift.im</code>	vector shift of image domain
<code>rescale.im</code>	rescaling of image domain
<code>eval.im</code>	evaluate any expression involving images
<code>cut.im</code>	convert numeric image to factor image
<code>split.im</code>	divide pixel image into sub-images
<code>by.im</code>	apply function to subsets of pixel image
<code>interp.im</code>	spatial interpolation of image
<code>blur</code>	spatial blurring and extrapolation of image

10.2 Inspecting an image

10.2.1 Basic information

For basic information about an image `Z`, use the following:

<code>Z</code>	print basic information
<code>print(Z)</code>	print basic information
<code>summary(Z)</code>	print detailed information
<code>dim(Z)</code>	pixel raster dimensions (y, x)
<code>nrow(Z)</code>	number of rows (y coordinate)
<code>ncol(Z)</code>	number of columns (x coordinate)
<code>range(Z)</code>	range of pixel values
<code>max(Z)</code>	maximum of pixel values
<code>min(Z)</code>	minimum of pixel values
<code>mean(Z)</code>	mean of pixel values
<code>median(Z)</code>	median of pixel values
<code>quantile(Z, ...)</code>	quantiles of pixel values
<code>sum(Z)</code>	sum of pixel values
<code>integral.im(Z)</code>	sum of pixel values times pixel area

To compute other numerical summaries of pixel values that are not on this list, you can extract the pixel values using `as.matrix(Z)` then apply the desired operation.


10.2.2 Plotting an image

Methods for plotting an image object include:

<code>plot.im</code>	display as colour image
<code>contour.im</code>	contour plot
<code>persp.im</code>	perspective plot of surface

These are methods for generic functions, so you would type `plot(Z)`, `contour(Z)` or `persp(Z)` to display a pixel image `Z`.


```
> opa <- par(mfrow = c(1, 3))
> data(redwood)
> D <- density(redwood)
> plot(D)
> persp(D)
> contour(D)
> par(opa)
```



For `plot.im`, note that the default colour map for image plots in R has only 12 colours and can convey a misleading impression of the gradation of pixel values in the image. Use the argument `col` to control the colour map.


```
> opa <- par(mfrow = c(1, 2))
> plot(D)
```

```
> plot(D, col = grey(seq(1, 0, length = 512)))
> par(opa)
```



In the example above, the argument `col` was a vector of colour data. The range of pixel values in the image `Z` was mapped to these colours. Unfortunately this means that if we plot two images `Z1`, `Z2` using the same `col` vector, the interpretation of the colours will be different! To avoid this, set the argument `col` to be an object of the special class "`colourmap`", created by the function `colourmap`. An object of this class specifies a mapping between numerical values and colours.


```
> mymap <- colourmap(grey(seq(1, 0, length = 512)), range = c(0,
+ 140))
> plot(D, col = mymap)
```



See `help(colourtools)` for tools that manipulate colours.

For `persp.im`, see also the help for `persp.default` for the names of various arguments to control the appearance of the plot. For example, the viewing direction is controlled by the angles `theta` and `phi`.

```
> persp(density(redwood), theta = 30)
```

Similarly for `contour.im`, consult also the help file for `contour.default` to control the appearance of the contours.

For some inspiring examples of perspective and contour plots with beautiful colour schemes and shading, see the R graphics demonstration by typing `demo(graphics)`.

10.2.3 Exploratory analysis

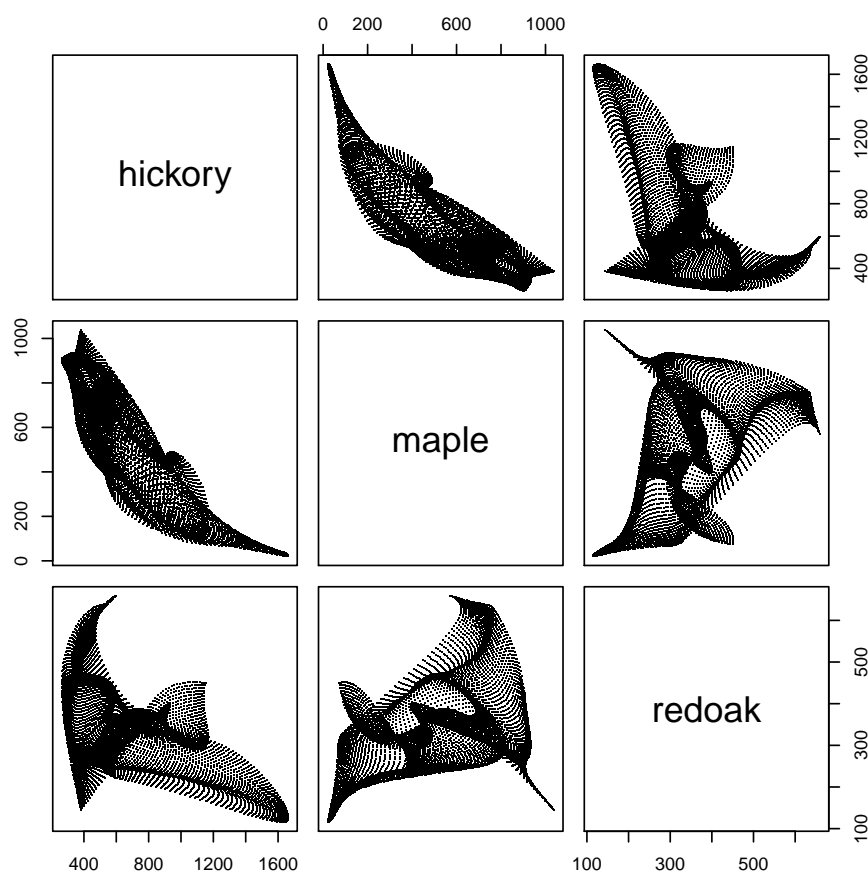
To inspect an image, the following are useful.

<code>as.matrix</code>	extract matrix of pixel values from image
<code>cut.im</code>	convert numeric image to factor image
<code>hist.im</code>	histogram of pixel values
<code>ecdf.im</code>	cumulative distribution function of pixel values

For an image `Z` with any type of values, `plot(cut(Z, 3))` will divide the pixel values into 3 bands, and display the image with the 3 bands rendered in 3 different colours.

To study the relation between two or more images, it's useful to display the `pairs` plot, a scatterplot of the corresponding pixel values of each image. See `pairs.im`.

```
> data(lansing)
> pairs(density(split(lansing)[c(2, 3, 5)]))
```



This command divided the Lansing Woods point pattern dataset into 6 sub-patterns of different tree species, extracted the 3 most common species, computed the kernel smoothed intensity estimate for each species, and then displayed scatterplots of the intensity estimates for each pair of species. The plot suggests that hickory and maple trees are strongly segregated from one another (since a high density of hickories is strongly associated with a low density of maples).

10.3 Manipulating images

10.3.1 Subsets of an image

The subset operator `[]` has a method for pixel images, `$.im`:

```
> X[S]
> X[S, drop = TRUE]
```

The subset to be extracted is determined by the index argument `S`.

- If `S` is a point pattern, or a `list(x,y)`, then the values of the pixel image `X` at these points are extracted, and returned as a vector.
- If `S` is a window (an object of class "owin"), the values of the image inside this window are extracted. The result is a pixel image if possible, and a numeric vector otherwise (see `help("$.im")` for details).

- If **S** is a pixel image with logical values, it is interpreted as a window (with TRUE inside the window).

The logical argument **drop** determines whether pixel values that are undefined are omitted (**drop = TRUE**) or returned as the value NA (**drop=FALSE**).

See `help("[.im]")` for full details.


The subset operator can be used to look up the value of a pixel image at a single point:

```
> data(bei)
> elev <- bei.extra$elev
> elev[list(x = 142, y = 356)]
```

```
[1] 147.08
```

or to display a subregion:

```
> S <- owin(c(200, 300), c(100, 200))
> plot(elev[S])
```



This can even be performed interactively, using the R function `locator` to click on a point in the window:

```
> elev[locator(1)]
```

10.3.2 Computation with images

The handy function `eval.im` allows us to perform pixel-by-pixel calculations on an image or on several compatible images.

If **Z** is a pixel image, to take the logarithm of each pixel value,

```
> logZ <- eval.im(log(Z))
```

If **A** and **B** are two pixel images with compatible grids of pixels (i.e. having the same numbers of pixels and the same coordinate locations), then to find the sum of the corresponding pixel values,

```
> C <- eval.im(A + B)
```

The expressions may involve constants and functions as well, so long as the expression is ‘parallelised’.


```
> W <- eval.im(sin(pi * Z))
> V <- eval.im(Z > 3)
> U <- eval.im(ifelse(Z > 3, 42, Z))
```

Other functions which manipulate images include the following:

<code>shift.im</code>	vector shift of an image
<code>cut.im</code>	convert numeric image to factor image
<code>split.im</code>	divide pixel image into sub-images
<code>by.im</code>	apply function to subsets of pixel image
<code>interp.im</code>	spatially interpolate an image
<code>levelset</code>	threshold an image (produces a window)
<code>solutionset</code>	find the region where a statement is true (produces a window)

11 Tessellations

A “tessellation” is a division of space into non-overlapping regions (“tiles”).




Tessellations have several uses in **spatstat**. The tessellation may be ‘real’, for example, a continent divided into states or provinces. The tessellation may be completely artificial, for example, the rectangular quadrats which we use in quadrat counting. Or the tessellation may be computed from other data, for example, the Dirichlet tessellation defined by a set of points.

11.1 Creating a tessellation

An object of class “**tess**” represents a tessellation. Currently **spatstat** supports three kinds of tessellations:

- **rectangular tessellations** in which the tiles are rectangles with sides parallel to the coordinate axes;
- **tile lists**, tessellations consisting of a list of windows, usually polygonal windows;
- **pixellated tessellations**, in which space is divided into pixels and each tile occupies a subset of the pixel grid.



All three types of tessellation can be created by the command `tess`.

To create a rectangular tessellation:

```
> tess(xgrid = xg, ygrid = yg)
```

where `xg` and `yg` are vectors of coordinates of vertical and horizontal lines determining a grid of rectangles. Alternatively, if you want to divide a rectangular window `W` into rectangles of equal size, you can type

```
> quadrats(W, nx, ny)
```

where `nx,ny` are the numbers of rectangles in the *x* and *y* directions, respectively. A common use of this command is to create quadrats for a quadrat-counting method.

To create a tessellation from a list of windows,

```
> tess(tiles = z)
```

where `z` is a list of objects of class "owin". The windows should not be overlapping; currently `spatstat` does not check this. This command is commonly used when the study region is divided into administrative regions (states, départements, postcodes, counties) and the boundaries of each sub-region are provided by GIS data files.

To create a tessellation from a pixel image,

```
> tess(image = Z)
```

where `Z` is a pixel image with factor values. Each level of the factor represents a different tile of the tessellation. The pixels that have a particular value of the factor constitute a tile. This command is often used to separate the landcover types in a landcover image (a pixel image in which each pixel is labelled by the type of vegetation or land use at that location) into different regions.


The command `as.tess` can also be used to convert other types of data to a tessellation.

11.2 Computed tessellations

There are two commands which compute a tessellation from a point pattern.


The command `dirichlet(X)` computes the *Dirichlet tessellation* or *Voronoi tessellation* of the point pattern `X`. The tile associated with a given point of the pattern `X` is the region of space which is closer to that point than to any other point of `X`. The Dirichlet tiles are polygons. The command `dirichlet(X)` computes these polygons and intersects them with the window of `X`.

```
> X <- runifpoint(42)
> plot(dirichlet(X))
```



The command `delaunay(X)` computes the *Delaunay triangulation* of the point pattern `X`. Strictly speaking this is not a tessellation but a network or graph, formed by joining some of the points of `X` by straight lines. Two points of `X` are joined if their Dirichlet tiles share a common edge. The resulting network forms a set of non-overlapping triangles. These triangles cover the *convex hull* of `X` rather than the entire window of `X`.

```
> plot(delaunay(X))
```



11.3 Operations involving a tessellation

There are methods for `print`, `plot` and `[` for tessellations.

Use the command `tiles` to extract a list of the tiles in a tessellation. The result is a list of windows ("owin" objects). This can be handy if, for example, you want to compute some characteristic of the tiles in a tessellation, such as their areas or diameters:

```
> X <- runifpoint(10)
> V <- dirichlet(X)
> U <- tiles(V)
> unlist(lapply(U, area.owin))
```

Tile 1	Tile 2	Tile 3	Tile 4	Tile 5	Tile 6	Tile 7
0.17096279	0.09645337	0.02587679	0.18646932	0.15352600	0.07760811	0.03357145
Tile 8	Tile 9	Tile 10				
0.01904539	0.15038824	0.08609854				


Tessellations can be used to classify the points of a point pattern, in `split.ppp`, `cut.ppp` and `by.ppp`. If `X` is a point pattern and `V` is a tessellation, then

- `cut(X,V)` attaches marks to the points of `X` identifying which tile of `V` each point falls into;
- `split(X,V)` divides the point pattern into sub-patterns according to the tiles of `V`, and returns a list of the sub-patterns;
- `by(X,V,FUN)` divides the point pattern into sub-patterns according to the tiles of `V`, applies the function `FUN` to each sub-pattern, and returns the results as a list.


```
> par(mfrow = c(1, 3))
> X <- runifpoint(100)
> plot(X)
> Z <- dirichlet(runifpoint(16))
> plot(Z)
> plot(cut(X, Z))
```

```
1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16
1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16
```

```
> par(mfrow = c(1, 1))
```




```
> plot(split(X, Z))
```



If we plot two tessellations on the same spatial domain, what we see is another tessellation. The “*intersection*” (or “*overlay*” or “*common refinement*”) of two tessellations X and Y is the tessellation whose tiles are the intersections between tiles of X and tiles of Y . The command `intersect.tess` computes the intersection of two tessellations.

```
> opa <- par(mfrow = c(1, 3))
> plot(X)
> plot(Y)
> plot(intersect.tess(X, Y))
> par(opa)
```



Other operations for tessellations include:

<code>bdist.tess</code>	compute distance from tile to boundary line
<code>chop.tess</code>	divide tessellation along a line
<code>rpoislinetess</code>	generate tessellation based on random lines

PART III. INTENSITY

Finally we can start working on statistical methods for analysing point pattern data. Part III of the workshop discusses how to investigate the intensity of a point pattern, and its dependence on covariates.

12 Exploring intensity

When we analyse numerical data, we often begin by taking the sample mean. The analogue of the mean or expected value of a random variable is the *intensity* of a point process.

‘Intensity’ is the average density of points (expected number of points per unit area). Intensity may be constant (‘uniform’ or ‘homogeneous’) or may vary from location to location (‘inhomogeneous’). Investigation of the intensity should be one of the first steps in analysing a point pattern.

12.1 Uniform intensity

12.1.1 Theory

If the point process \mathbf{X} is homogeneous, then for any sub-region B of two-dimensional space, the expected number of points in B is proportional to the area of B :

$$\mathbb{E}[N(\mathbf{X} \cap B)] = \lambda \text{area}(B)$$

and the constant of proportionality λ is the intensity. Intensity units are numbers per unit area (length⁻²). If we know that a point process is homogeneous, then the empirical density of points,

$$\bar{\lambda} = \frac{n(\mathbf{x})}{\text{area}(W)}$$

is an unbiased estimator of the true intensity λ .

12.1.2 Implementation in spatstat

To compute the estimator $\bar{\lambda}$ in `spatstat`, use `summary.ppp`:

```
> data(swedishpines)
> summary(swedishpines)
```

```
Planar point pattern: 71 points
Average intensity 0.0074 points per square unit (one unit = 0.1 metres)
```

```
Window: rectangle = [0, 96]x[0, 100]units
Window area = 9600 square units
Unit of length: 0.1 metres
```

The estimated intensity is $\bar{\lambda} = 0.0074$ points per square unit. To extract this intensity value, type

```
> lamb <- summary(swedishpines)$intensity
> lamb
```

```
[1] 0.007395833
```

The units are decimetres, so this is 0.74 points per square metre.

12.2 Inhomogeneous intensity

12.2.1 Theory

In general the intensity of a point process will vary from place to place. Assume that the expected number of points falling in a small region of area du around a location u is equal to $\lambda(u) du$. Then $\lambda(u)$ is the “*intensity function*” of the process, satisfying

$$\mathbb{E}[N(\mathbf{X} \cap B)] = \int_B \lambda(u) du$$

for all regions B .

More generally there could be singular concentrations of intensity (e.g. earthquake epicentres may be concentrated along a fault line) so that an intensity function does not exist. Then we speak of the “*intensity measure*” Λ defined by

$$\Lambda(B) = \mathbb{E}[N(\mathbf{X} \cap B)]$$

for each $B \subset \mathbb{R}^2$, assuming the expectation is finite.

If it is suspected that the intensity may be inhomogeneous, the intensity function or intensity measure can be estimated nonparametrically by techniques such as quadrat counting and kernel smoothing.

In quadrat counting, the window W is divided into subregions (‘quadrats’) B_1, \dots, B_m of equal area. We count the numbers of points falling in each quadrat, $n_j = n(\mathbf{x} \cap B_j)$ for $j = 1, \dots, m$. These are unbiased estimators of the corresponding intensity measure values $\Lambda(B_j)$.

The usual *kernel estimator* of the intensity function is

$$\tilde{\lambda}(u) = e(u) \sum_{i=1}^n \kappa(u - x_i), \quad (1)$$

where $\kappa(u)$ is the kernel (an arbitrary probability density) and

$$e(u)^{-1} = \int_W \kappa(u - v) dv \quad (2)$$

is an edge effect bias correction. Clearly $\tilde{\lambda}(u)$ is an unbiased estimator of

$$\lambda^*(u) = e(u) \int_W \kappa(u - v) \lambda(v) dv,$$

a smoothed version of the true intensity function $\lambda(u)$. The choice of smoothing kernel κ involves a tradeoff between bias and variance.

Intensity can also be estimated using parametric methods, as we explain in Section 15.

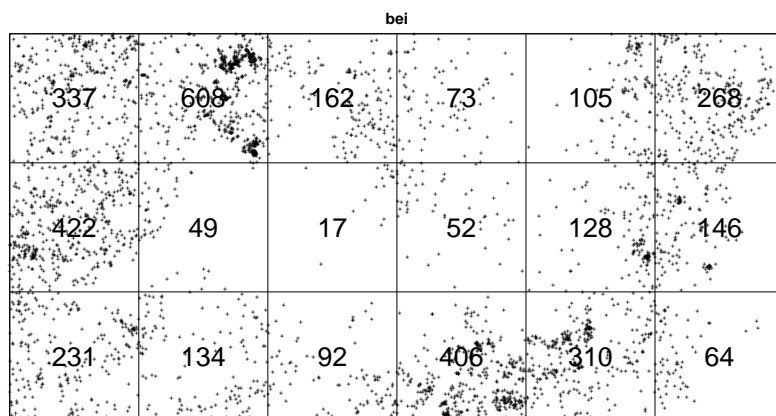
12.2.2 Implementation in spatstat

Quadrat counting is performed in `spatstat` by the function `quadratcount`.

```
> data(bei)
> quadratcount(bei, nx = 4, ny = 2)
```

		x			
y		[0,250]	(250,500]	(500,750]	(750,1e+03]
	(250,500]	666	677	130	481
	[0,250]	544	165	643	298


```
> Q <- quadratcount(bei, nx = 6, ny = 3)
> plot(bei, cex = 0.5, pch = "+")
> plot(Q, add = TRUE, cex = 2)
```



The value returned by `quadratcount` is an object belonging to the special class "quadratcount". We have used the plot method for this class to get the display above.


Kernel density (or *intensity*) estimation using an isotropic Gaussian kernel is implemented in `spatstat` by the function `density.ppp`, a method for the generic command `density`.

```
> den <- density(bei, sigma = 70)
> plot(den)
> plot(bei, add = TRUE, cex = 0.5)
```




The value returned by `density.ppp` is a pixel image (object of class "im"). This class has methods for `print`, `summary`, `plot`, `contour` (contour plots), `persp` (perspective plots) and so on.

```
> persp(den)
```




```
> contour(den, axes = FALSE)
```



Alternatively, there is an **adaptive estimator of intensity** which uses a fraction f of the data to construct a Dirichlet tessellation, then forms an intensity estimate that is constant in each tile of the tessellation:

```
> aden <- adaptive.density(bei, f = 0.01, nrep = 10)
> plot(aden, main = "Adaptive intensity")
> plot(bei, add = TRUE, cex = 0.5)
```




The value returned by `adaptive.density` is also a pixel image (object of class "im").

13 Dependence of intensity on a covariate

13.1 Spatial covariates

Often we want to know whether the intensity of points depends on the values of a covariate. For example, the tropical rainforest point pattern dataset `bei` comes with an extra set of covariate data `bei.extra`, which contains a pixel image of terrain elevation `bei.extra$elev` and a pixel image of terrain slope `bei.extra$grad`. It is of interest to determine whether the trees prefer steep or flat terrain, and whether they prefer a particular altitude.

```
> data(bei)
> slope <- bei.extra$grad
> par(mfrow = c(1, 2))
> plot(bei)
> plot(slope)
> par(mfrow = c(1, 1))
```




13.2 Quadrats determined by a covariate

In quadrat counting methods, any choice of quadrats is permissible. From a theoretical viewpoint, the quadrats do not have to be rectangles of equal area, and could be regions of any shape.

Quadrat counting is more useful if we choose the quadrats in a meaningful way. One way to do this is to define the quadrats using covariate information.

For the tropical rainforest data `bei`, it might be useful to split the study region into several sub-regions according to the terrain slope.

```
> data(bei)
> Z <- bei.extra$grad
> b <- quantile(Z, probs = (0:4)/4)
> Zcut <- cut(Z, breaks = b, labels = 1:4)
> V <- tess(image = Zcut)
> plot(V)
> plot(bei, add = TRUE, pch = "+")
```



The call to `quantile` gave us the quartiles of the slope values, so the four tiles in the tessellation `V` have equal area (ignoring discretisation effects). In other words, we have divided the study region into four zones of equal area according to the terrain slope.


We can now use this tessellation to study the point pattern `bei`. We could invoke the commands `split`, `cut` or `by` to divide the points according to this tessellation and manipulate the sub-patterns.

The command `quadratcount` also works with tessellations:

```
> qb <- quadratcount(bei, tess = V)
> qb

tile
  1    2    3    4
271  984 1028 1321

> plot(qb)
```



The text annotations show the number of trees in each region. Since the four regions have equal area, the counts should be approximately equal if there is a uniform density of trees. Obviously they are not equal; there appears to be a strong preference for steeper slopes.

13.3 Relative distribution estimate


Let us assume that the intensity of the point process is a *function* of the covariate Z . At any spatial location u , let $\lambda(u)$ be the intensity of the point process, and $Z(u)$ the value of the covariate. Then we are assuming

$$\lambda(u) = \rho(Z(u))$$

where ρ is a function that we want to investigate, telling us how the intensity of points depends on the value of the covariate.

Kernel smoothing can be used to estimate the function ρ , using methods of relative distribution or relative risk, as explained in [11, 5].


```
> plot(rhohat(bei, slope))
```




The plot is an estimate of the intensity $\rho(z)$ as a function of terrain slope z . It indicates that the *Beilschmiedia* trees are relatively unlikely to be found on flat terrain (where the slope is less than 0.05) compared to steeper slopes.

Additional capabilities will be added into `spatstat` in the near future.

13.4 Distance map

The dataset `copper` gives the locations of copper deposits in a survey region, and also the location of geological lineaments (which are mostly geological faults). It is conjectured that copper is more likely to be deposited close to a fault.

```
> data(copper)
> X <- rotate(copper$SouthPoints, pi/2)
> L <- rotate(copper$SouthLines, pi/2)
> plot(X, pch = 16, main = "copper data")
> plot(L, add = TRUE)
```




To apply the methods described above, the covariate information contained in the map of geological faults L must be converted into a covariate that is a function $Z(u)$ of spatial location u . A natural choice is the *distance function*

$$Z(u) = \text{distance from } u \text{ to } L$$

This can be computed by the command `distmap`, which returns a pixel image containing the values $Z(u)$ at a fine grid of pixels u .

```
> Z <- distmap(L)
> plot(L, lwd = 2, main = "")
> contour(Z, add = TRUE)
```



Having created this covariate image we can now apply the other techniques such as relative distributions.

```
> plot(rhohat(X, Z), xlab = "Distance to nearest fault")
```



A slightly more sophisticated version of `distmap` is the command `distfun`. Whereas `distmap` returns a pixel image at a certain spatial resolution, `distfun` returns a function with arguments (x,y) that can be evaluated at any spatial location.

```
> f <- distfun(L)
> f
```

```
Distance function for line segment pattern
planar line segment pattern: 90 line segments
window: rectangle = [-158.233, -0.19] x [-0.335, 35] km
```

```
> f(-42, 10)
```

```
[1] 2.387029
```

In most commands in `spatstat` where a pixel image is required, a `distfun` can be used in its place. This increases the precision of many calculations. It is usually advisable to call `distfun`, unless you really need a pixel image.

PART IV. POISSON MODELS

Part IV of the workshop discusses how to assess whether a pattern is completely random, and how to model a random pattern with a spatial trend.

14 Tests of Complete Spatial Randomness

The basic ‘reference’ or ‘benchmark’ model of a random point pattern is the *uniform Poisson point process* in the plane with intensity λ , sometimes called *Complete Spatial Randomness (CSR)*. Its basic properties are

- the number of points falling in any region A has a Poisson distribution with mean $\lambda \times \text{area}(A)$
- given that there are n points inside region A , the locations of these points are i.i.d. and uniformly distributed inside A
- the contents of two disjoint regions A and B are independent.

For historical reasons, many researchers are focussed on establishing that their data do *not* conform to this model. The logic is that, if a point pattern is completely random, then there is nothing “interesting” happening (because the points are completely unpredictable, and have no trend or association with anything else). Statisticians would say that the uniform Poisson process often serves as the ‘null model’ in a statistical analysis.

14.1 Definition

The *homogeneous Poisson process* of intensity $\lambda > 0$ has the properties


- (PP1): the number $N(\mathbf{X} \cap B)$ of points falling in any region B is a Poisson random variable;
- (PP2): the expected number of points falling in B is $\mathbb{E}[N(\mathbf{X} \cap B)] = \lambda \cdot \text{area}(B)$;
- (PP3): if B_1, B_2 are disjoint sets then $N(\mathbf{X} \cap B_1)$ and $N(\mathbf{X} \cap B_2)$ are independent random variables;
- (PP4): given that $N(\mathbf{X} \cap B) = n$, the n points are independent and uniformly distributed in B .

The list is redundant; (PP2) and (PP3) are sufficient.

This process is often called “*Complete Spatial Randomness (CSR)*” especially in biological science. Under CSR, points are independent of each other and have the same propensity to be found at any location.

It is easy to simulate the Poisson process directly by following the properties (PP1)–(PP4). In **spatstat**, use the command **rpoispp** (by convention, random data generators have names beginning with **r**).

```
> plot(rpoispp(100))
```




Conceptually, if we discretise a homogeneous Poisson process into infinitesimal pixels, the indicators I are independent and identically distributed, with success probability $\mathbb{P}\{I = 1\} = \lambda dA$ where dA is the infinitesimal area of a pixel.

To develop some intuition about completely random patterns, it's useful to repeat the command `plot(rpoispp(100))` several times (use the up-arrow key to recall the previous command line) so that you see several replicates of the Poisson process. In particular you will notice that the points in a homogeneous Poisson process are not 'uniformly spread': there are empty gaps and clusters of points.

The command `rpoispp` has arguments `lambda` (the intensity) and `win` (the window in which to simulate). The default window is the unit square.

```
> data(letterR)
> plot(rpoispp(100, win = letterR))
```



If you want to simulate a Poisson process *conditionally* on a fixed number of points, use the command `runifpoint`.

```
> runifpoint(100, win = letterR)

planar point pattern: 100 points
window: polygonal boundary
enclosing rectangle: [2.017, 3.93] x [0.645, 3.278] units
```

14.2 Quadrat counting tests for CSR

In classical literature, the homogeneous Poisson process (CSR) is usually taken as the appropriate 'null' model for a point pattern. Our basic task in analysing a point pattern is to find evidence against CSR.

A classical test for the null hypothesis of CSR is the χ^2 test based on quadrat counts. As explained earlier, the window W is divided into subregions ('quadrats') B_1, \dots, B_m of equal area. We count the numbers of points falling in each quadrat, $n_j = n(\mathbf{x} \cap B_j)$ for $j = 1, \dots, m$. Under the null hypothesis of CSR, the n_j are i.i.d. Poisson random variables with the same expected value. The Pearson χ^2 goodness-of-fit test can be used.

```
> quadrat.test(nzchop, nx = 3, ny = 2)
```

Chi-squared test of CSR using quadrat counts

```
data:  nzchop
X-squared = 5.0769, df = 5, p-value = 0.4066
```

The value returned by `quadrat.test` is an object of class `"htest"` (the standard R class for hypothesis tests). Printing the object (as shown above) gives comprehensible output about the outcome of the test. Inspecting the p -value, we see that the test does not reject the null hypothesis of CSR for the (chopped) New Zealand trees data.


The return value `quadrat.test` also belongs to the special class `"quadrat.test"`. Plotting the object will display the quadrats, annotated by their observed and expected counts and the Pearson residuals (observed counts n_j at top left; expected count at top right; Pearson residuals at bottom).

```
> M <- quadrat.test(nzchop, nx = 3, ny = 2)
> M
```

Chi-squared test of CSR using quadrat counts

```
data:  nzchop
X-squared = 5.0769, df = 5, p-value = 0.4066
```

```
> plot(nzchop)
> plot(M, add = TRUE, cex = 2)
```



The p -value can also be extracted by

```
> M$p.value
[1] 0.4065648
```

14.3 Critique

Since this kind of technique is often used in the applied literature, a few comments are appropriate.

The main critique of the quadrat test approach is the lack of information. This is a goodness-of-fit test in which the alternative hypothesis H_1 is simply the negation of H_0 , that is, the alternative is that “the process is not a homogeneous Poisson process”. A point process may fail to satisfy properties (PP1)–(PP4) either because it violates (PP2) by having non-uniform

intensity, or because it violates (PP3)–(PP4) by exhibiting dependence between points. There are too many types of departure from H_0 .

The usual justification for the classical χ^2 goodness-of-fit test is to assume that the counts are independent, and derive a test of the null hypothesis that all counts have the same expected value. Invoking it here is slightly naive, since the independence of counts is also open to question here.

Indeed we can also turn things around and view the χ^2 test as a test of the Poisson distributional properties (PP2)–(PP3) assuming that the intensity is uniform. The Pearson χ^2 test statistic

$$X^2 = \frac{\sum_j (n_j - n/m)^2}{n/m}$$

(where $n = \sum_j n_j$ is the total number of points) coincides, up to a constant factor, with the sample variance-to-mean ratio of the counts n_j , which is often interpreted as a measure of over/under-dispersion of the counts n_j assuming they have constant mean.

The power of the quadrat test depends on the size of quadrats, and falls to zero for quadrats which are either very large or very small. The power also depends on the alternative hypothesis, in particular on the ‘spatial scale’ of any departures from the assumptions of constant intensity and independence of points. The choice of quadrat size carries an implicit assumption about the spatial scale.

14.4 Kolmogorov-Smirnov test of CSR

Typically a more powerful test of CSR is the Kolmogorov-Smirnov test in which we compare the observed and expected distributions of the values of some function T .

We specify a real-valued function $T(x, y)$ defined at all locations (x, y) in the window. We evaluate this function at each of the data points. Then we compare this empirical distribution of values of T with the predicted distribution of values of T under CSR, using the classical Kolmogorov-Smirnov test.

In **spatstat** the spatial Kolmogorov-Smirnov test is performed by **kstest**. This function is generic. The method for point patterns, **kstest.ppp**, performs the Kolmogorov-Smirnov test for CSR.

If **X** is the data point pattern, then

```
> kstest(X, covariate)
```

performs the test, where **covariate** is the spatial covariate that will be used. Here **covariate** can be a pixel image, a **function(x,y)** in the R language, or one of the strings **"x"** or **"y"** indicating one of the Cartesian coordinates.

For example, let's consider the **nzchop** data and choose the function T to be the x coordinate, $T(x, y) = x$. This means we are simply comparing the observed and expected distributions of the x coordinate.


```
> kstest(nzchop, "x")
```

Spatial Kolmogorov-Smirnov test of CSR

```
data: covariate x evaluated at points of nzchop
      and transformed to uniform distribution under CSR
D = 0.0719, p-value = 0.7882
alternative hypothesis: two-sided
```


The result of `kstest` is an object of class `"htest"` (the standard R class for hypothesis tests) and also of class `"kstest"` so that it can be printed and plotted. The print method (demonstrated above) reports information about the hypothesis test such as the p -value. The plot method displays the observed and expected distribution functions.

```
> KS <- kstest(nzchop, "x")
> plot(KS)
> pval <- KS$p.value
```



Sometimes this test generates a warning message about tied values. Typically this occurs because the coordinates in the dataset have been rounded to the nearest integer, so that there are tied observations.

14.5 Using covariate data

We are often interested in testing whether the point pattern intensity depends on a covariate. For example, our preliminary analysis of the tropical rainforest pattern `bei` in Section 13.2 suggested that the density of trees depends on terrain slope. To test this formally we can divide the region into irregular quadrats according to the terrain slope, and apply the χ^2 test. The command `quadrat.test` accepts a tessellation and uses the tiles of the tessellation as the quadrats:

```
> data(bei)
> Z <- bei.extra$grad
> b <- quantile(Z, probs = (0:4)/4)
> Zcut <- cut(Z, breaks = b, labels = 1:4)
> V <- tess(image = Zcut)
> quadrat.test(bei, tess = V)
```

Chi-squared test of CSR using quadrat counts

```
data:  bei
X-squared = 661.8402, df = 3, p-value < 2.2e-16
```


Because of the large counts in these regions, we can probably ignore concerns about independence, and conclude that the trees are not uniform in their intensity.

A more powerful test (if that were needed!) is the Kolmogorov-Smirnov test using the slope covariate:

```
> KS <- kstest(bei, Z)
> plot(KS)
> KS
```

Spatial Kolmogorov-Smirnov test of CSR

```
data: covariate Z evaluated at points of bei
      and transformed to uniform distribution under CSR
D = 0.1948, p-value < 2.2e-16
alternative hypothesis: two-sided
```



The Kolmogorov-Smirnov test would typically be preferred if the covariate Z has continuously-varying numerical values. If the covariate is a factor or discrete variable, then the Kolmogorov-Smirnov test is ineffective because of tied values, and the χ^2 test based on quadrat counts would be preferable.


14.6 Berman's tests

Berman [20] proposed two tests for the dependence of a point process on a spatial covariate. These tests are optimal against a certain class of alternatives. They are performed by the command `bermantest` which is analogous to `kstest`.

```
> B <- berrmantest(bei, Z)
> plot(B)
> B
```

Berman Z1 test of CSR

```
data: covariate Z evaluated at points of bei
Z1 = 10.844, p-value < 2.2e-16
alternative hypothesis: two-sided
```



Two vertical lines show the mean values of these distributions. If the model is correct, the two curves should be close; the test is based on comparing the two vertical lines.

When the covariate Z is the distance to a spatial pattern, another useful diagnostic is Foxall's *J*-function [40], available using `Jfox`.

15 Maximum likelihood for Poisson processes

If we are willing to assume (tentatively) that the points are independent, then we can apply some decent statistical methods to the investigation of the intensity.

15.1 Inhomogeneous Poisson process

The *inhomogeneous* Poisson process with intensity function $\lambda(u)$, $u \in \mathbb{R}^2$, is a modification of the homogeneous Poisson process, in which properties (PP2) and (PP4) above are replaced by


(PP2'): the number $N(\mathbf{X} \cap B)$ of points falling in a region B has expectation

$$\mathbb{E}[N(\mathbf{X} \cap B)] = \int_B \lambda(u) \, du.$$

(PP4'): given that $N(\mathbf{X} \cap B) = n$, the n points are independent and identically distributed, with common probability density $f(u) = \lambda(u)/I$, where $I = \int_B \lambda(u) \, du$.

This process can also be simulated using `rpoispp` using the same properties. The intensity argument `lambda` can be a constant, a `function(x,y)` giving the values of the intensity function at coordinates x, y , or a pixel image containing the intensity values at a grid of locations.

```
> lambda <- function(x, y) {
+   100 * (x + y)
+ }
> plot(rpoispp(lambda))
```



If we discretise an inhomogeneous Poisson process, the indicators I are independent, but have unequal success probabilities, $\mathbb{P}\{I(u) = 1\} = \lambda(u) \, dA$.

The inhomogeneous Poisson process is a plausible model for point patterns under several scenarios. One is **random thinning**: suppose that a homogeneous Poisson process of intensity β is generated, and that each point is either deleted or retained, independently of other points. Suppose the probability of retaining a point at the location u is $p(u)$. Then the resulting process of retained points is inhomogeneous Poisson, with intensity $\lambda(u) = \beta p(u)$.

Consider, for example, a model of plant propagation which assumes that seeds are randomly dispersed according to a Poisson process, and seeds randomly germinate or do not germinate, independently of each other, with a germination probability that depends on the local soil conditions. The resulting pattern of plants is an inhomogeneous Poisson process.

15.2 Likelihood methods

The log-likelihood for the homogeneous Poisson process with intensity λ is

$$\log L(\lambda; \mathbf{x}) = n(\mathbf{x}) \log \lambda - \lambda \text{area}(W) \quad (3)$$

where $n(\mathbf{x})$ is the number of points in the dataset \mathbf{x} . The maximum likelihood estimator of λ is

$$\hat{\lambda} = \frac{n(\mathbf{x})}{\text{area}(W)}$$

which is also an unbiased estimator. The variance of $\hat{\lambda}$ is $\text{var}[\hat{\lambda}] = \lambda/\text{area}(W)$.

Consider an inhomogeneous Poisson process with intensity function $\lambda_\theta(u)$ depending on a parameter θ . The log-likelihood for θ is

$$\log L(\theta; \mathbf{x}) = \sum_{i=1}^n \log \lambda_\theta(x_i) - \int_W \lambda_\theta(u) \, du \quad (4)$$

This is a well-behaved likelihood; for example if $\log \lambda_\theta(u)$ is linear in θ , then the log-likelihood is concave, so there is a unique MLE. However, the MLE $\hat{\theta}$ is not analytically tractable, so it must be computed using numerical algorithms such as Newton's method.

The usual asymptotic theory of maximum likelihood applies: under suitable large sample conditions, the MLE of θ is asymptotically normal. If we wish to test CSR, the likelihood ratio test statistic

$$R = 2 \log \frac{L(\hat{\theta})}{L(\hat{\lambda})}$$

is asymptotically χ^2 under CSR, and this gives an asymptotically optimal test of CSR against the alternative of an inhomogeneous Poisson process with intensity $\lambda_\theta(u)$.

15.3 Fitting Poisson processes in spatstat

Mark Berman and Rolf Turner [21] (see also [47, 25, 48]) developed a clever computational device for finding the MLE of θ by exploiting a formal similarity between the Poisson log-likelihood (4) and that of a loglinear Poisson regression.

The Berman-Turner algorithm is implemented in **spatstat**. The intensity function $\lambda_\theta(u)$ must be loglinear in the parameter θ :

$$\log \lambda_\theta(u) = \theta \cdot S(u) \quad (5)$$

where $S(u)$ is a real-valued or vector-valued function of location u . The form of S is arbitrary so this is not much of a restriction. In practice $S(u)$ could be a function of the spatial coordinates of u , or an observed covariate, or a mixture of both. Assuming (5), the log-likelihood (4) is a convex function of θ , so maximum likelihood is well-behaved.

15.3.1 Model-fitting function

The fitting function is called **ppm** ('point process model') and is very closely analogous to the model fitting functions in R such as **lm** and **glm**. The statistic $S(u)$ is specified by an R language formula, like the formulas used to specify the systematic relationship in a linear model or generalised linear model. The basic syntax is:

```
> ppm(X, ~trend)
```

where X is the point pattern dataset, and `~trend` is an R formula with no left-hand side. This should be viewed as a model with log link, so *the formula `~trend` specifies the form of the logarithm of the intensity function.*

To fit the homogeneous Poisson model:

```
> ppm(bei, ~1)
```

Stationary Poisson process

Uniform intensity: 0.007208

To fit an inhomogeneous Poisson model with an intensity that is log-linear in the cartesian coordinates, i.e. $\lambda_\theta((x, y)) = \exp(\theta_0 + \theta_1 x + \theta_2 y)$,

```
> ppm(bei, ~x + y)
```

Nonstationary Poisson process

Trend formula: `~x + y`

Fitted coefficients for trend formula:

(Intercept)	x	y
-4.7245290274	-0.0008031288	0.0006496090

Here `x` and `y` are reserved names that always refer to the cartesian coordinates. In the output, the ‘fitted coefficients’ are the maximum likelihood estimates of $\theta_0, \theta_1, \theta_2$, the coefficients of the ‘linear predictor’. The fitted intensity function is

$$\lambda_\theta((x, y)) = \exp(-4.724529 - 0.000803x + 0.00065y).$$

To fit an inhomogeneous Poisson model with an intensity that is log-quadratic in the cartesian coordinates, i.e. such that $\log \lambda_\theta((x, y))$ is a quadratic in x and y :

```
> ppm(bei, ~polynom(x, y, 2))
```

Nonstationary Poisson process

Trend formula: `~polynom(x, y, 2)`

Fitted coefficients for trend formula:

(Intercept)	polynom(x, y, 2)[x]	polynom(x, y, 2)[y]
-4.275762e+00	-1.609187e-03	-4.895166e-03
polynom(x, y, 2)[x^2]	polynom(x, y, 2)[x.y]	polynom(x, y, 2)[y^2]
1.625968e-06	-2.836387e-06	1.331331e-05

Essentially any kind of model formula can be used, involving the reserved names `x` and `y` and any covariates (as we explain later).

To fit a model with constant but unequal intensities on each side of the vertical line $x = 500$, the explanatory variable $S(u)$ should be a factor with two levels, `Left` and `Right` say, taking the value `Left` when the location u is to the left of the line $x = 500$.

```
> side <- function(z) factor(ifelse(z < 500, "left", "right"))
> ppm(bei, ~side(x))
```

Nonstationary Poisson process

Trend formula: `~side(x)`

Fitted coefficients for trend formula:

```
(Intercept) side(x)right
-4.8026460   -0.2792705
```

When factors are involved, the interpretation of the coefficients depends on which ‘contrasts’ are in force. By default the ‘treatment contrasts’ are assumed. This means that the treatment effect is taken to be zero for the first level of the factor, and the estimated treatment effects for other levels are effectively estimates of the difference from the first level. In this case “left” comes alphabetically before “right”, so by default, the first level is “left”. The fitted model is

$$\lambda_{\theta}((x, y)) = \begin{cases} \exp(-4.8026) & \text{if } x < 500 \\ \exp(-4.8026 + (-0.2793)) & \text{if } x \geq 500 \end{cases}$$


Rather than relying on such interpretations, it is prudent to use the command `predict` to compute predicted values of the model, as explained in Section 15.4 below.

15.3.2 Models involving spatial covariates

It is also possible to fit an inhomogeneous Poisson process model with an intensity function that depends on an observed covariate. Let $Z(u)$ be a covariate that has been measured at every location u in the study window. Then $Z(u)$, or any transformation of it, can serve as the statistic $S(u)$ in the parametric form (5) for the intensity function.

The point pattern dataset `bei` is supplied with accompanying covariate data `bei.extra`. The covariates are the elevation (altitude) and the slope of the terrain at each location in the window, given as two pixel images `bei.extra$elev` and `bei.extra$grad`.

```
> data(bei)
> grad <- bei.extra$grad
> plot(grad)
```



To fit the inhomogeneous Poisson model with intensity which is a loglinear function of slope, i.e.

$$\lambda(u) = \exp(\beta_0 + \beta_1 Z(u)) \quad (6)$$

where β_0, β_1 are parameters and $Z(u)$ is the slope at location u , we type

```
> ppm(bei, ~slope, covariates = list(slope = grad))
```

Nonstationary Poisson process

Trend formula: ~slope

Fitted coefficients for trend formula:

```
(Intercept)      slope
-5.390553      5.022021
```

In the call to `ppm`, the argument `covariates` should be a list of `name=value` pairs that provide the covariate data for the model. Every variable name that appears in the model formula should match one of the `names` in this list. Each `value` should be either a pixel image, a `function(x,y)` in the R language, a window of class `"owin"`, or a single number.

The printout includes the fitted coefficients β_0, β_1 so the fitted model is

$$\lambda(u) = \exp(-5.390553 + 5.022021 Z(u)). \quad (7)$$

It might be more appropriate to fit the inhomogeneous Poisson model with intensity that is *proportional* to slope,

$$\lambda(u) = \beta Z(u) \quad (8)$$

where again $Z(u)$ is the slope at u . Equivalently

$$\log \lambda(u) = \log \beta + \log Z(u). \quad (9)$$

There is no coefficient in front of the term $\log Z(u)$ in (9), so this term is an ‘offset’. To fit this model,

```
> ppm(bei, ~offset(log(slope)), covariates = list(slope = grad))
```

Nonstationary Poisson process

Trend formula: ~offset(log(slope))

Fitted coefficients for trend formula:

```
(Intercept)
-2.427127
```

The fitted coefficient is the constant $\log \beta$ appearing in (9), so converting back to the form (8), the fitted model is

$$\lambda(u) = e^{-2.427127} Z(u) = 0.0883 Z(u).$$

15.4 Fitted models

The value returned by the model-fitting function `ppm` is an object of class `"ppm"` that represents the fitted model. This is analogous to the fitting of linear models (`lm`), generalised linear models (`glm`) and so on.

15.4.1 Standard operations

The following standard operations on fitted models in R can be applied to point process models (i.e. these operations have methods for the class "ppm"):

<code>print</code>	print basic information
<code>summary</code>	print detailed summary information
<code>plot</code>	plot the fitted intensity
<code>predict</code>	compute the fitted intensity
<code>fitted</code>	compute the fitted intensity at data points
<code>update</code>	re-fit the model
<code>coef</code>	extract the fitted coefficient vector $\hat{\theta}$
<code>vcov</code>	variance-covariance matrix of $\hat{\theta}$
<code>anova</code>	analysis of deviance
<code>logLik</code>	log-likelihood value
<code>formula</code>	extract the model formula
<code>terms</code>	extract the terms in the model
<code>model.matrix</code>	compute the design matrix

For information on these methods, see `print.ppm`, `summary.ppm`, `plot.ppm` etc. The following commands also work on "ppm" objects:

<code>step</code>	stepwise model selection
<code>drop1</code>	one step model deletion
<code>AIC</code>	Akaike Information Criterion

```
> fit <- ppm(bei, ~x + y)
> fit
```


Nonstationary Poisson process

Trend formula: `~x + y`

Fitted coefficients for trend formula:

(Intercept)	x	y
-4.7245290274	-0.0008031288	0.0006496090

```
> plot(fit, how = "image", se = FALSE)
```



```
> predict(fit, type = "trend")
```

real-valued pixel image

50 x 50 pixel array (ny, nx)

enclosing rectangle: [0, 1000] x [0, 500] metres

```
> predict(fit, type = "cif", ngrid = 256)

real-valued pixel image
256 x 256 pixel array (ny, nx)
enclosing rectangle: [0, 1000] x [0, 500] metres

> coef(fit)

      (Intercept)           x           y
-4.7245290274 -0.0008031288  0.0006496090

> vcov(fit)

      (Intercept)           x           y
(Intercept)  1.854091e-03 -1.491267e-06 -3.528289e-06
x           -1.491267e-06  3.437842e-09  1.208410e-14
y           -3.528289e-06  1.208410e-14  1.338955e-08

> sqrt(diag(vcov(fit)))

      (Intercept)           x           y
4.305915e-02  5.863311e-05  1.157132e-04

> round(vcov(fit, what = "corr"), 2)

      (Intercept)           x           y
(Intercept)      1.00 -0.59 -0.71
x              -0.59  1.00  0.00
y              -0.71  0.00  1.00
```

This is the fitted model with intensity function


$$\lambda_{\theta}((x, y)) = \exp(\theta_0 + \theta_1 x + \theta_2 y) \quad (10)$$

with the following estimates:

i	θ_i	$\text{var}(\hat{\theta}_i)$	standard deviation
0	-4.724529	0.001854091	0.04305915
1	-0.0008031288	3.437842e-09	5.863311e-05
2	0.000649609	1.338955e-08	0.0001157132

It is also possible to compute the standard error of the fitted intensity $\lambda_{\theta}(u)$ at each location u , as a pixel image. Use `predict(fit, type="se")` or `plot(fit, se=TRUE)`.

```
> SE <- predict(fit, type = "se")
> plot(SE, main = "standard error of fitted intensity")
```



If the model formula involves transformations of the original covariates, then `model.matrix(fit)` gives the design matrix whose columns contain these transformed covariates, and `model.images(fit)` gives a list of pixel images of these transformed covariates.


```
> fit <- ppm(bei, ~sqrt(slope) + x, covariates = list(slope = grad))
> mo <- model.images(fit)
> mo
```

```
(Intercept) :
real-valued pixel image
100 x 100 pixel array (ny, nx)
enclosing rectangle: [0, 1000] x [0, 500] metres
```

```
sqrt(slope) :
real-valued pixel image
100 x 100 pixel array (ny, nx)
enclosing rectangle: [0, 1000] x [0, 500] metres
```


```
x :
real-valued pixel image
100 x 100 pixel array (ny, nx)
enclosing rectangle: [0, 1000] x [0, 500] metres
```

```
> plot(mo[[2]])
```



It is also possible to plot the ‘effect’ of a single covariate in the model. The command `effectfun` computes the intensity of the fitted model as a function of one of its covariates. This is chiefly useful if the model only has one covariate.

```
> fit <- ppm(bei, ~slope, covariates = list(slope = grad))
> plot(effectfun(fit, "slope"))
```



15.4.2 Model selection

Analysis of deviance for nested Poisson point process models is implemented in `spatstat` as `anova.ppm`. The first model should be a sub-model of the second.

```
> fit <- ppm(bei, ~slope, covariates = list(slope = grad))
> fitnull <- update(fit, ~1)
> anova(fitnull, fit, test = "Chi")
```

Analysis of Deviance Table

Model 1: .mpl.Y ~ 1

Model 2: .mpl.Y ~ slope

	Resid. Df	Resid. Dev	Df	Deviance	P(> Chi)
1	20507	18728			
2	20506	18346	1	382.25	< 2.2e-16 ***

Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

This effectively performs the **likelihood ratio test** of the null hypothesis of a homogeneous Poisson process (CSR) against the alternative of an inhomogeneous Poisson process with intensity that is a loglinear function of the slope covariate (6). The p -value is extremely small, indicating rejection of CSR in favour of the alternative. (Please ignore the columns `Resid.Df` and `Resid.Dev` which are artefacts of the discretisation. Only the deviance difference and the difference in degrees of freedom are valid.)

Note that standard Analysis of Deviance requires the null hypothesis to be a sub-model of the alternative. Unfortunately the model (8), in which intensity is proportional to slope, does *not* include the homogeneous Poisson process as a special case, so we cannot use analysis of deviance to test the null hypothesis of homogeneous Poisson against the alternative of an inhomogeneous Poisson with intensity (8).

One possibility here is to use the Akaike Information Criterion **AIC** for model selection.

```
> fitprop <- ppm(bei, ~offset(log(slope)), covariates = list(slope = grad))
> fitnull <- ppm(bei, ~1)
> AIC(fitprop)
```

```
[1] 42496.65
```

```
> AIC(fitnull)
```

```
[1] 42763.92
```

The smaller AIC favours the model (8) with intensity is proportional to slope.

Automatic model selection can be performed using `step`. By default, this performs stepwise deletion. Starting from the fitted model, the procedure considers each term in the model, and determines whether the term should be deleted (according to AIC). The deletion giving the biggest improvement in AIC is carried out. This is applied recursively until no more terms can be deleted.

```
> X <- rpoispp(100)
> fit <- ppm(X, ~x + y)
> step(fit)
```

```
Start:  AIC=-580.96
```

```
~x + y
```

	Df	AIC
- x	1	-582.29
- y	1	-581.91
<none>		-580.96

```
Step:  AIC=-582.29
```

```
~y
```

	Df	AIC
- y	1	-583.25
<none>		-582.29

```
Step:  AIC=-583.25
```

```
~1
```


Stationary Poisson process

```
Uniform intensity:      85
```

15.5 Simulating the fitted model

A fitted Poisson model can be simulated automatically using the function `rmh` or `simulate.ppm`.

```
> X <- rmh(fitprop)
> plot(X, main = "realisation of fitted model")
```



It is possible to perform conditional simulation (conditional on the number of points, on the configuration of points in a particular subregion, or on the presence of certain points). See `rmhcontrol` for details.

16 Checking a fitted Poisson model

After fitting a point process model to a point pattern dataset, we should check that the model is a good fit ('goodness-of-fit'), and that each component assumption of the model was appropriate ('validation'). This section presents some techniques available for checking a fitted Poisson model.

Model checking can be either 'formal' or 'informal'. Formal techniques are based on detailed probabilistic assumptions about the data, and allow us to make probabilistic statements about the outcome. They include hypothesis tests (χ^2 tests, goodness-of-fit tests, Monte Carlo tests) and Bayesian model selection.

In contrast, 'informal' tools do not impose assumptions on the data and their interpretation depends on human judgement. A typical example is the residual, defined for each observation by $(\text{residual}) = (\text{observed}) - (\text{fitted})$. If the model is a good fit, then the residuals should be 'noise', centred around zero.

16.1 Goodness-of-fit

A goodness-of-fit test is a formal test of the null hypothesis that the model is true, against the very general alternative that the model is not true.

The χ^2 goodness-of-fit test based on quadrat counts can be applied to a fitted Poisson model, homogeneous or inhomogeneous. Under the null hypothesis, the quadrat counts are independent Poisson variables with different mean values, and the means are estimated by the fitted model.

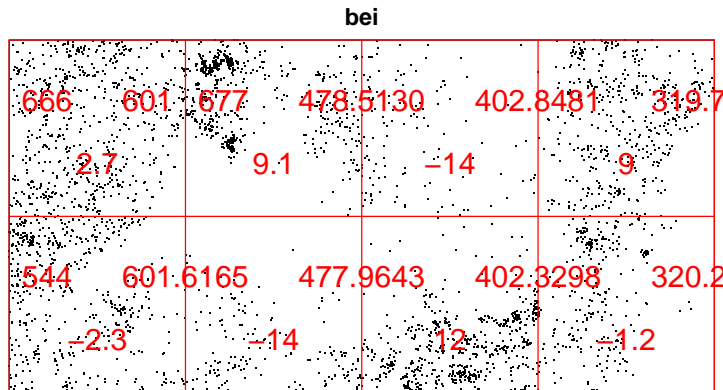
```
> data(bei)
> fit <- ppm(bei, ~x)
> M <- quadrat.test(fit, nx = 4, ny = 2)
> M
```

Chi-squared test of fitted Poisson model fit using quadrat counts

```
data: data from fit
X-squared = 711.5036, df = 6, p-value < 2.2e-16
```

If (as in this case) the formal goodness-of-fit test rejects the fitted model, we would then like to gain an informal impression of the type of departure from the model (i.e. in what way the data appear to depart from the predictions of the model) so that we may formulate a better model. To do this we can inspect the residual counts.

```
> plot(bei, pch = ".")
> plot(M, add = TRUE, cex = 1.5, col = "red")
```



The plot displays, for each quadrat, the observed number of points (top left), the predicted number of points according to the model (top right), and the Pearson residual (bottom) defined by

$$\text{Pearson residual} = \frac{(\text{observed}) - (\text{expected})}{\sqrt{\text{expected}}}$$

If the original data were Poisson, this transformation approximately standardises the residuals so that they have mean zero and variance 1 when the model is true. A Pearson residual of -14 is a gross departure from the fitted model.

The Kolmogorov-Smirnov test can also be applied to a fitted Poisson model, with homogeneous or inhomogeneous intensity.

```
> kstest(fit, "y")
```

Spatial Kolmogorov-Smirnov test of inhomogeneous Poisson process

```
data: covariate y evaluated at points of bei
      and transformed to uniform distribution under fit
D = 0.1068, p-value < 2.2e-16
alternative hypothesis: two-sided
```

This uses the method `kstest.ppm` for the generic function `kstest`.

16.2 Validation using residuals

16.2.1 Residuals

Residuals from the fitted model are an important diagnostic tool in other areas of applied statistics, but in spatial statistics they have only recently been developed ([52, 60], [58, pp. 49–50], [12]).

For a fitted Poisson process model, with fitted intensity $\hat{\lambda}(u)$, the predicted number of points falling in any region B is $\int_B \hat{\lambda}(u) du$. Hence the **residual** in each region $B \subset \mathbb{R}^2$ is defined [12] to be the *observed minus predicted* number of points falling in B : [12]

$$R(B) = n(\mathbf{x} \cap B) - \int_B \hat{\lambda}(u) du \quad (11)$$

where \mathbf{x} is the observed point pattern, $n(\mathbf{x} \cap B)$ the number of points of \mathbf{x} in the region B , and $\hat{\lambda}(u)$ is the intensity of the *fitted* model.

These residuals are closely related to the residuals for quadrat counts that were used above. Taking the set B to be one of our quadrats, the ‘observed’ quadrat count is $n(\mathbf{x} \cap B)$. The ‘expected’ quadrat count is $\hat{\lambda} \text{area}(B)$ if the model is CSR, or more generally $\int_B \hat{\lambda}(u) du$ if the model is an inhomogeneous Poisson process. Hence the ‘raw residual’ is **observed** – **expected** = $n(\mathbf{x} \cap B) - \int_B \hat{\lambda}(u) du$.


16.2.2 Residual measure

Equation (11) defines the total residual for any region B , large or small.

Intuitively the residuals can be visualised as an electric charge, with unit positive charge at each data point, and a diffuse negative charge at all other locations u , with density $\hat{\lambda}(u)$. If the model is true, then these charges should approximately cancel.

If we’d like to visualise this electric charge, one way is to plot the observed points and the fitted intensity function together:

```
> data(bei)
> fit <- ppm(bei, ~x + y)
> plot(predict(fit))
> plot(bei, add = TRUE, pch = "+")
```




Each data point should be visualised as a charge of +1, while the colour image indicates a negative charge density. If the model is true then these positive and negative charges should even out to zero.

16.2.3 Smoothed residuals

A more useful way to visualise the residuals is to smooth them.

```
> data(bei)
> fitx <- ppm(bei, ~x)
> diagnose.ppm(fitx, which = "smooth")
```



This is an image plot of the ‘smoothed residual field’

$$s(u) = \hat{\lambda}(u) - \lambda^\dagger(u) \quad (12)$$

where $\hat{\lambda}(u)$ is the nonparametric, kernel estimate of the intensity,

$$\hat{\lambda}(u) = e(u) \sum_{i=1}^{n(\mathbf{x})} \kappa(u - x_i)$$

while $\lambda^\dagger(u)$ is a correspondingly-smoothed version of the parametric estimate of the intensity according to the fitted model,

$$\lambda^\dagger(u) = e(u) \int_W \kappa(u - v) \lambda_{\hat{\theta}}(v) \, dv.$$

Here κ is the smoothing kernel and $e(u)$ is the edge correction (2) on page 79. The difference (12) should be approximately zero if the model is true.

In this example the smoothed residual image contains a visible trend, suggesting that the model is inappropriate.


16.2.4 Lurking variable plot

If there is a spatial covariate $Z(u)$ that plays an important role in the analysis, it may be useful to display a *lurking variable plot* of the residuals against Z . This is a plot of $C(z) = R(B(z))$ against z , where

$$B(z) = \{u \in W : Z(u) \leq z\}$$

is the region of space where the covariate value is less than or equal to z .

```
> grad <- bei.extra$grad
> lurking(fitx, grad, type = "raw")
```




Note that the lurking variable plot typically starts and ends at the horizontal axis, since (for any model with an intercept term) the total residual for the entire window W must equal zero. This is analogous to the fact that the residuals in linear regression sum to zero.

The plot also shows approximate 5% significance bands for the cumulative residual $C(x)$ or $C(y)$, obtained from the asymptotic variance under the model.

This plot indicates that the model is grossly inadequate; the fitted intensity function fails to capture the dependence of intensity on slope.

It can be helpful to display the derivative $C'(z)$, which often indicates which values of z are associated with a lack of fit.

```
> lurking(fitx, grad, type = "raw", cumulative = FALSE)
```



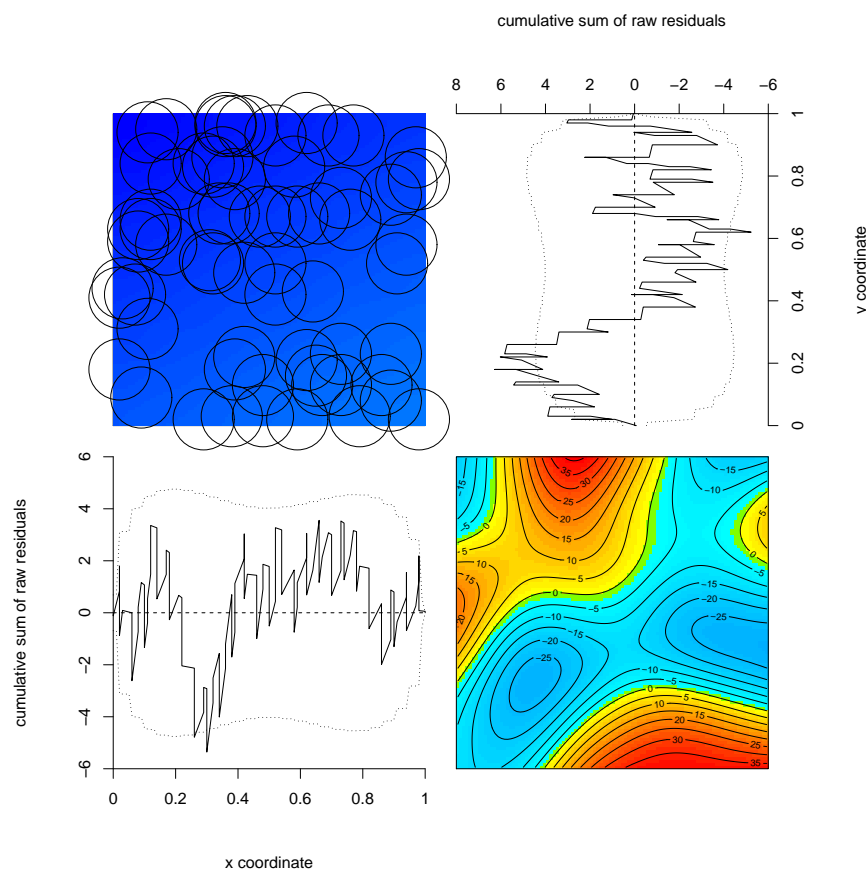
The derivative is estimated using a smoothing spline and you may need to tweak the smoothing parameters (argument `splineargs`) to get a useful plot. Also the package currently does not plot significance bands for $C'(z)$.

Additional techniques described in [5] will soon be added to `spatstat`.

16.2.5 Four-panel plot

If there are no spatial covariates, use the command `diagnose.ppm` to plot the residuals:

```
> data(japanesepines)
> fit <- ppm(japanesepines, ~x + y)
> diagnose.ppm(fit)
```



This combination of four plots has proved to be a useful quick indication of departure from the trend in the model.

The bottom right panel is an image of the smoothed residual field.

The top left panel is a direct representation of the residual ‘charge’, with circles representing the data points (positive residuals) and a colour scheme representing the fitted intensity (negative residuals). However, it is often difficult to interpret.

The two other panels are lurking variables against one of the cartesian coordinates. For example, the bottom left panel is a lurking variable plot for the x -coordinate. Imagine a vertical line which sweeps from left to right across the window. The progressive total residual to the left of the line is plotted against the position of the line.

In this example, the lurking variable plot for the y coordinate suggests a lack of fit at about $y = 0.15$, and the image of the smoothed residual field suggests an excess of positive residuals at about $x = 0.8, y = 0.15$, both indicating that the model *underestimates* the true intensity of points in this vicinity.


16.2.6 Caveats

The residual plots described above are useful for detecting misspecification of the *trend* in a fitted Poisson process model. They are not very useful for checking the *independence* assumption, that is, for checking the properties (PP3)–(PP4) of a Poisson process listed on page 88.

Effective diagnostics of independence or dependence between points include the K -function (section 19.4) and a Q–Q plot of the residuals (section 28.2.3).

17 Spatial logistic regression

A popular technique for analysing point pattern data in Geographical Information Systems is *spatial logistic regression* [62, 1]. The spatial domain is divided into a fine grid of pixels; each pixel is assigned the value $y = 1$ if it contains at least one data point, and $y = 0$ otherwise.



Then logistic regression is used to model the presence probability $p = P(Y = 1)$ as a function of a covariate x in the form

$$\log \frac{p}{1-p} = \beta_0 + \beta_1 x$$

where β_0, β_1 are parameters to be estimated. Similarly for multiple covariates and so on.

In fact, spatial logistic regression is very close to a Poisson point process model [4]. However, `spatstat` provides facilities for performing spatial logistic regression, for comparison purposes. Models are fitted using the command `slrm`.

```
> data(copper)
> X <- rotate(copper$SouthPoints, pi/2)
> L <- rotate(copper$SouthLines, pi/2)
> D <- distfun(L)
> fit <- slrm(X ~ D)
> fit
```

Fitted spatial logistic regression model

Formula: X ~ D

Fitted coefficients:

```
(Intercept)      D
-4.69015189  0.05211539
```

`slrm` produces a “fitted spatial logistic regression model” object of class “`slrm`”. Methods for this class include `print`, `plot`, `predict`, `coef`, `fitted`, `update`, `terms`, `formula`, `anova` and `logLik`. You can also use `step` for model selection.


PART V. INTERACTION

Part V of the workshop explains how to investigate dependence between the points in a point pattern.


18 Exploring dependence between points

Suppose that a point pattern appears to have constant intensity, and we wish to assess whether the pattern is Poisson. The alternative is that the points are dependent (they exhibit ‘interaction’).


Classical writers suggested a simple trichotomy between ‘independence’ (the Poisson process), ‘regularity’ (where points tend to avoid each other), and ‘clustering’ (where points tend to be close together). [The concept of ‘clustering’ does not imply that the points are organised into identifiable ‘clusters’; merely that they are closer together than expected for a Poisson process.]



One simple diagnostic for dependence between points is a *Morishita plot*. The spatial domain is divided into quadrats, and the χ^2 statistic based on quadrat counts is computed. The quadrats are repeatedly subdivided. The Morishita plot shows the χ^2 statistic against the linear size of the quadrats. It is computed by the command `mipLOT`.



A more sophisticated option is the *Fry plot*. This is a scatterplot of the vector differences $x_j - x_i$ between all pairs of distinct points in the pattern. Suppose you have printed the point pattern on a sheet of paper. Take a sheet of tracing paper, and mark a red dot in the middle. Now place the tracing paper over the point pattern, and move it until the red dot coincides with one of the data points. Now copy all the other data points onto the tracing paper. Repeat this for every data point, and you have the Fry plot. It is computed by the command `fryplot`.



While these graphical diagnostics can be very useful for spotting features of the point pattern, they involve subjective interpretation.

19 Distance methods for point patterns

19.1 Distances


The main classical techniques for investigating interpoint interaction are *distance methods*, based on measuring the distances between points. Specifically we may consider

- **pairwise distances** $s_{ij} = \|x_i - x_j\|$ between all distinct pairs of points x_i and x_j ($i \neq j$) in the pattern;
- **nearest neighbour distances** $t_i = \min_{j \neq i} s_{ij}$, the distance from each point x_i to its nearest neighbour;
- **empty space distances** $d(u) = \min_i \|u - x_i\|$, the distance from a fixed reference location u in the window to the nearest data point.

If you need to compute these directly, they are available in **spatstat** using the functions **pairdist**, **ndist** and **distmap** respectively. If **X** is a point pattern object,


- **pairdist(X)** returns the matrix of pairwise distances.
- **ndist(X)** returns the vector of nearest neighbour distances.
- **distmap(X)** returns a pixel image whose pixel values are the empty space distances to the pattern **X** measured from every pixel.

```
> data(cells)
> emp <- distmap(cells)
> plot(emp, main = "Empty space distances")
> plot(cells, add = TRUE)
```

Tip: Quite a useful exploratory tool is the *Stienen diagram* obtained by drawing a circle around each data point of diameter equal to its nearest neighbour distance:

```
> plot(X %mark% (nndist(X)/2), markscale = 1, main = "Stienen diagram")
```



In order to develop formal statistical analysis, we typically use the empirical cumulative distribution function of these distances. This is explained in the following subsections.

19.2 Empty space distances

It's easiest to start by explaining the analysis of the empty space distances.

The distance


$$d(u, \mathbf{x}) = \min\{\|u - x_i\| : x_i \in \mathbf{x}\}$$

from a fixed location $u \in \mathbb{R}^2$ to the nearest point in a point pattern \mathbf{x} , is called the 'empty space distance' or 'void distance'. It can be computed for all locations u on a fine grid, using the `spatstat` function `distmap` as we saw above.

19.2.1 Edge effects

It is not easy to interpret a histogram of the empty space distances. The empirical distribution of the empty space distances depends on the geometry of the window W as well as on characteristics of the point process \mathbf{X} .

Another viewpoint is that the window introduces a sampling bias. Recall that under the ‘standard model’ (Section 2.3) the point process \mathbf{X} extends throughout 2-D space, but is observed only inside W . This leads to bias in the distance measurements. Confining observations to a window W implies that the observed distance $d(u, \mathbf{x}) = d(u, \mathbf{X} \cap W)$ to the nearest data point inside W , may be greater than the true distance $d(u, \mathbf{X})$ to the nearest point of the complete point process \mathbf{X} .



19.2.2 Empty space function F

Ignoring the edge problems for a moment, let us focus on the entire point process \mathbf{X} .

Assuming \mathbf{X} is *stationary* (statistically invariant under translations), we can define the cumulative distribution function of the empty space distance

$$F(r) = \mathbb{P}\{d(u, \mathbf{X}) \leq r\} \quad (13)$$

where u is an arbitrary reference location. If the process is stationary then this definition does not depend on u .

The empirical distribution function of the observed empty space distances on a grid of locations u_j , $j = 1, \dots, m$,

$$F^*(r) = \frac{1}{m} \sum_j \mathbf{1}\{d(u_j, \mathbf{x}) \leq r\} \quad (14)$$

is a negatively biased estimator of $F(r)$, for reasons explained above.

Corrections for this ‘edge effect bias’ are required. Many edge corrections are available. Typically they are weighted versions of the ecdf,

$$\hat{F}(r) = \sum_j e(u_j, r) \mathbf{1}\{d(u_j, \mathbf{x}) \leq r\} \quad (15)$$

where $e(u, r)$ is an edge correction weight designed so that $\hat{F}(r)$ is unbiased. These corrections are effectively forms of the Horvitz-Thompson estimator of survey sampling fame.

The edge effect problem can also be regarded as a form of censoring (analogous to right-censoring in survival data), as first pointed out by CSIRO researcher Geoff Laslett [46]. A counterpart of the Kaplan-Meier estimator is available. For further information see [13].

Thus, *assuming that the point process is homogeneous*, we are able to compute an unbiased and reasonably accurate estimate of the empty space function F defined by (13).

To interpret this estimate, a useful benchmark is the Poisson process. Notice that $d(u, \mathbf{X}) > r$ if and only if there are no points of X in the disc $b(u, r)$ of radius r centred on u . For a homogeneous Poisson process of intensity λ , the number of points falling in $b(u, r)$ is Poisson with mean $\mu = \lambda \text{area}(b(u, r)) = \lambda \pi r^2$, so the probability that there are no points in this region is $\exp(-\mu) = \exp(-\lambda \pi r^2)$. Hence for a Poisson process

$$F_{\text{pois}}(r) = 1 - \exp(-\lambda \pi r^2). \quad (16)$$

Typically we compare $\hat{F}(r)$ with the value of $F_{\text{pois}}(r)$ obtained by plugging in the estimated intensity $\hat{\lambda} = n(\mathbf{x})/\text{area}(W)$. Values $\hat{F}(r) > F_{\text{pois}}(r)$ suggest that empty space distances in the point pattern are shorter than for a Poisson process, suggesting a regularly space pattern; while values $\hat{F}(r) < F_{\text{pois}}(r)$ suggest a clustered pattern.

19.2.3 Implementation in spatstat

The function **Fest** computes estimates of $F(r)$ using several edge corrections, and the benchmark value for the Poisson process.

```
> data(cells)
> plot(cells)
> Fc <- Fest(cells)
> Fc
```

Function value object (class fv)

for the function $r \rightarrow F(r)$

Entries:

id	label	description
---	-----	-----
r	r	distance argument r
theo	F[pois](r)	theoretical Poisson F(r)
cs	F[cs](r)	Chiu-Stoyan estimate of F(r)
rs	F[bord](r)	border corrected estimate of F(r)
km	F[km](r)	Kaplan-Meier estimate of F(r)
hazard	hazard(r)	Kaplan-Meier estimate of hazard function lambda(r)


Default plot formula:

. ~ r

<environment: 0x6a251e0>

Recommended range of argument r: [0, 0.085]

Available range of argument r: [0, 0.2975]



Tip: Don't use `F` as a variable name! It's a reserved word — an abbreviation for `FALSE`.


The value returned by `Fest` is an object of class `"fv"` ("function value table"). This is effectively a data frame with some extra information. The printout for `Fc` indicates that the columns in the data frame are named `r`, `theo`, `cs`, `rs`, `km` and `hazard`. The first column `r` contains a sequence of values of the function argument r . The next column `theo` contains the corresponding values of $F(r)$ for a homogeneous Poisson process. The columns `cs`, `rs` and `km` contain different estimates of the empty space function F , namely the Chiu-Stoyan estimator, the 'reduced sample' estimator, and the Baddeley-Gill Kaplan-Meier estimator, respectively. The column `hazard` contains an estimate of the hazard rate of F , i.e. $h(r) = (d/dr) \log(1 - F(r))$, a by-product of the Kaplan-Meier estimate.

If you don't want to compute all these estimates (for example, for the sake of efficiency), you can use the argument `correction` to specify which estimate or estimates are required.

```
> Fest(cells, correction = "km")
> Fest(cells, correction = c("km", "cs"))
```

```
> par(pty = "s")
> plot(Fest(cells))
```


	lty	col	key	label	meaning
km	1	1	km	$F[\text{km}](r)$	Kaplan-Meier estimate of $F(r)$
rs	2	2	rs	$F[\text{bord}](r)$	border corrected estimate of $F(r)$
cs	3	3	cs	$F[\text{cs}](r)$	Chiu-Stoyan estimate of $F(r)$
theo	4	4	theo	$F[\text{pois}](r)$	theoretical Poisson $F(r)$



This is a call to `plot.fv`. The printed output is the return value from `plot.fv`, which explains the encoding of the different function estimates using the R graphics parameters `lty` (line type) and `col` (line colour).


You'll notice that, by default, the hazard rate `hazard` was not plotted. The choice of estimates to be plotted, and the style in which they are plotted, are controlled by the second argument to `plot.fv`, which should be an R language formula involving the identifier names `r`, `theo`, `cs`, `rs`, `km` and `hazard`. To plot the hazard rate against r ,

```
> plot(Fest(cells), hazard ~ r, main = "Hazard rate of F")
```



To plot only the Kaplan-Meier and Chiu-Stoyan estimators with the theoretical curve,


```
> plot(Fest(cells), cbind(km, cs, theo) ~ r)
```



Notice the use of `cbind` to specify several different graphs on the same plot.

To plot the estimates of $F(r)$ against the Poisson value, in the style of a P-P plot:

```
> plot(Fest(cells), cbind(km, rs, theo) ~ theo)
```




(including `theo` on the left side here gives us the diagonal line).

The symbol `.` stands for ‘all recommended estimates of the function’. So an abbreviation for the last command is

```
> plot(Fest(cells), . ~ theo)
```


Transformations can be applied to these function values. For example, to subtract the theoretical Poisson value from the estimates,

```
> plot(Fest(cells), . - theo ~ r)
```



To apply Fisher's variance stabilising transformation $\phi(\hat{F}(t)) = \sin^{-1}(\sqrt{\hat{F}(t)})$,

```
> plot(Fest(cells), asin(sqrt(.)) ~ r)
```



19.3 Nearest neighbour distances

For other types of distances we encounter similar problems. For the nearest neighbour distances $t_i = \min_{j \neq i} \|x_i - x_j\|$, again it is not easy to interpret a histogram of the observed distances. The empirical distribution of the nearest neighbour distances depends on the geometry of the window W as well as on characteristics of the point process \mathbf{X} . Confining observations to a window W implies that the observed nearest-neighbour distances are larger, in general, than the 'true' nearest neighbour distances of points in the entire point process \mathbf{X} . Corrections for this edge effect bias are required.

19.3.1 G function

Assuming the point process \mathbf{X} is stationary, we can define the cumulative distribution function of the nearest-neighbour distance for a typical point in the pattern,

$$G(r) = \mathbb{P}\{d(u, \mathbf{X} \setminus \{u\}) \leq r \mid u \in \mathbf{X}\} \quad (17)$$

where u is an arbitrary location, and $d(u, \mathbf{X} \setminus \{u\})$ is the shortest distance from u to the point pattern \mathbf{X} excluding u itself. If the process is stationary then this definition does not depend on u .

The empirical distribution function of the observed nearest-neighbour distances

$$G^*(r) = \frac{1}{n(\mathbf{x})} \sum_i \mathbf{1}\{t_i \leq r\} \quad (18)$$

is a negatively biased estimator of $G(r)$, for reasons we explained above. Many edge corrections are available. Typically they are weighted versions of the ecdf,

$$\hat{G}(r) = \sum_i e(x_i, r) \mathbf{1}\{t_i \leq r\} \quad (19)$$

where $e(x_i, r)$ is an edge correction weight designed so that $\hat{G}(r)$ is approximately unbiased. A counterpart of the Kaplan-Meier estimator is also available.

For a homogeneous Poisson point process of intensity λ , the nearest-neighbour distance distribution function is known to be

$$G_{\text{pois}}(r) = 1 - \exp(-\lambda\pi r^2). \quad (20)$$

This is identical to the empty space function for the Poisson process. Intuitively, because points of the Poisson process are independent of each other, the knowledge that u is a point of \mathbf{X} does not affect any other points of the process, hence G is equivalent to F .

Interpretation of $\hat{G}(r)$ is the reverse of $\hat{F}(r)$. Values $\hat{G}(r) > G_{\text{pois}}(r)$ suggest that nearest neighbour distances in the point pattern are shorter than for a Poisson process, suggesting a clustered pattern; while values $\hat{G}(r) < G_{\text{pois}}(r)$ suggest a regular (inhibited) pattern.

The function `Gest` computes estimates of $G(r)$ using several edge corrections, and the benchmark value for the Poisson process.

```
> Gc <- Gest(cells)
> Gc
```

```
Function value object (class fv)
for the function r -> G(r)
```

```
Entries:
```

id	label	description
---	-----	-----
r	r	distance argument r
theo	G[pois](r)	theoretical Poisson G(r)
han	G[han](r)	Hanisch estimate of G(r)
rs	G[bord](r)	border corrected estimate of G(r)
km	G[km](r)	Kaplan-Meier estimate of G(r)
hazard	hazard(r)	Kaplan-Meier estimate of hazard function lambda(r)

```
Default plot formula:
```

```
. ~ r
```


```
<environment: 0x6ab90d8>
```

```
Recommended range of argument r: [0, 0.15]
```

```
Available range of argument r: [0, 0.29539]
```



```
> par(pty = "s")
> plot(Gest(cells))
```



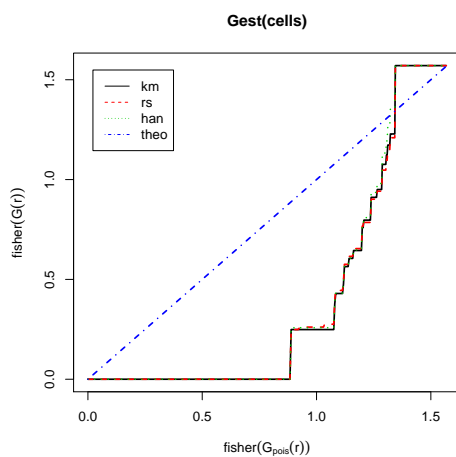
The estimate of $G(r)$ suggests strongly that the pattern is regular. Indeed, $\hat{G}(r)$ is zero for $r \leq 0.07$ which indicates that there are no nearest-neighbour distances shorter than 0.07.

Common ways of plotting \hat{G} include:

$\hat{G}(r)$ and $G_{\text{pois}}(r)$ plotted against r	<code>plot(Gest(X))</code>
$\hat{G}(r) - G_{\text{pois}}(r)$ plotted against r	<code>plot(Gest(X), . - theo ~ r)</code>
$\hat{G}(r)$ plotted against $G_{\text{pois}}(r)$ in P-P style	<code>plot(Gest(X), . ~ theo)</code>

and Fisher's variance-stabilising transformation $\phi(G(t)) = \sin^{-1}(\sqrt{G(t)})$ applied to the P-P plot:

```
> fisher <- function(x) {
+   asin(sqrt(x))
+ }
> plot(Gest(cells), fisher(.) ~ fisher(theo))
```



19.4 Pairwise distances and the K function

The observed pairwise distances $s_{ij} = \|x_i - x_j\|$ in the data pattern \mathbf{x} constitute a biased sample of pairwise distances in the point process, with a bias in favour of smaller distances. For example, we can never observe a pairwise distance greater than the diameter of the window.

Ripley [54] defined the K -function for a stationary point process so that $\lambda K(r)$ is the expected number of other points of the process within a distance r of a typical point of the process. Formally

$$K(r) = \frac{1}{\lambda} \mathbb{E}[n(\mathbf{X} \cap b(u, r) \setminus \{u\}) \mid u \in \mathbf{X}]. \quad (21)$$

For a homogeneous Poisson process, intuitively, the knowledge that u is a point of \mathbf{X} does not affect the other points of the process, so that $\mathbf{X} \setminus \{u\}$ is conditionally a Poisson process. The expected number of points falling in $b(u, r)$ is $\lambda \pi r^2$. Thus for a homogeneous Poisson process

$$K_{\text{pois}}(r) = \pi r^2 \quad (22)$$

regardless of the intensity.

Numerous estimators of K have been proposed. Most of them are weighted and renormalised empirical distribution functions of the pairwise distances, of the general form

$$\hat{K}(r) = \frac{1}{\hat{\lambda}^2 \text{area}(W)} \sum_i \sum_{j \neq i} \mathbf{1}\{\|x_i - x_j\| \leq r\} e(x_i, x_j; r) \quad (23)$$

where $e(u, v, r)$ is an edge correction weight. The choice of estimator does not seem to be very important, as long as *some* edge correction is applied.

Again we usually compare the estimate $\hat{K}(r)$ with the Poisson K function. Values $\hat{K}(r) > \pi r^2$ suggest clustering, while $\hat{K}(r) < \pi r^2$ suggests a regular pattern.

In **spatstat** the function **Kest** computes several estimates of the K -function.

```
> Gc <- Kest(cells)
> Gc
```

Function value object (class fv)

for the function $r \rightarrow K(r)$

Entries:

id	label	description
--	----	-----
r	r	distance argument r
theo	K[pois](r)	theoretical Poisson K(r)
border	K[bord](r)	border-corrected estimate of K(r)
trans	K[trans](r)	translation-corrected estimate of K(r)
iso	K[iso](r)	Ripley isotropic correction estimate of K(r)

Default plot formula:


$\cdot \sim r$

<environment: 0x60d97a8>

Recommended range of argument r : [0, 0.25]

Available range of argument r : [0, 0.25]

```
> par(pty = "s")
> plot(Kest(cells))
```



In this case, the interpretation of all three summary statistics F , G and K is the same: emphatic evidence of a regular pattern. It is not always the case that these three summaries give equivalent messages.


A commonly-used transformation of K is the L -function

$$L(r) = \sqrt{\frac{K(r)}{\pi}}$$

which transforms the Poisson K function to the straight line $L_{\text{pois}}(r) = r$, making visual assessment of the graph much easier. The square root transformation also approximately stabilises the variance of the estimator, making it easier to assess deviations.

To compute the estimated L function, use `Lest`.

```
> L <- Lest(cells)
> plot(L, main = "L function")
```




Another related summary function is the *pair correlation function*

$$g(r) = \frac{K'(r)}{2\pi r}$$

where $K'(r)$ is the derivative of K . The pair correlation is in some ways easier to interpret than either K or L , although it is more difficult to estimate. Roughly speaking, the pair correlation $g(r)$ is the probability of observing a pair of points separated by a distance r , divided by the corresponding probability for a Poisson process. This is a non-centred correlation which may take any nonnegative value. The value $g(r) = 1$ corresponds to complete randomness; for the Poisson process the pair correlation is $g_{\text{pois}}(r) \equiv 1$. For other processes, values $g(r) > 1$ suggest clustering or attraction at distance r , while values $g(r) < 1$ suggest inhibition or regularity.

To compute the estimated pair correlation function, use `pcf`.

```
> plot(pcf(cells))
```



Here we have used the method `pcf.ppp`. This computes a standard kernel estimate which performs well except at very small values of r . So it is prudent not to read too much into the behaviour of the `pcf` close to $r = 0$.

If you want to try another algebraic transformation of a summary function, the transformation can be computed using `eval.fv`. You can also plot algebraic transformations of a summary function using the ‘plotting formula’ argument to `plot.fv`. For example, if we have already computed the K function, we can plot the L function by

```
> K <- Kest(cells)
> plot(K, sqrt(./pi) ~ r)
```

and compute the L function using `eval.fv`:

```
> K <- Kest(cells)
> L <- eval.fv(sqrt(K/pi))
```

If you have already computed the K function and wish to derive the pair correlation, there is another algorithm `pcf.fv` that calculates $g(r) = K'(r)/(2\pi r)$ by numerical differentiation.

```
> K <- Kest(cells)
> g <- pcf(K)
```

19.5 J function

A useful combination of F and G is the J function [64]

$$J(r) = \frac{1 - G(r)}{1 - F(r)} \quad (24)$$

defined for all $r \geq 0$ such that $F(r) < 1$. For a homogeneous Poisson process, $F_{\text{pois}} = G_{\text{pois}}$, so that

$$J_{\text{pois}}(r) \equiv 1. \quad (25)$$

Values $J(r) > 1$ suggest regularity, and $J(r) < 1$ suggest clustering.

An appealing property of the J function is that the superposition $\mathbf{X}_{\bullet} = \mathbf{X}_1 \cup \mathbf{X}_2$ of two *independent* point processes $\mathbf{X}_1, \mathbf{X}_2$ has J -function

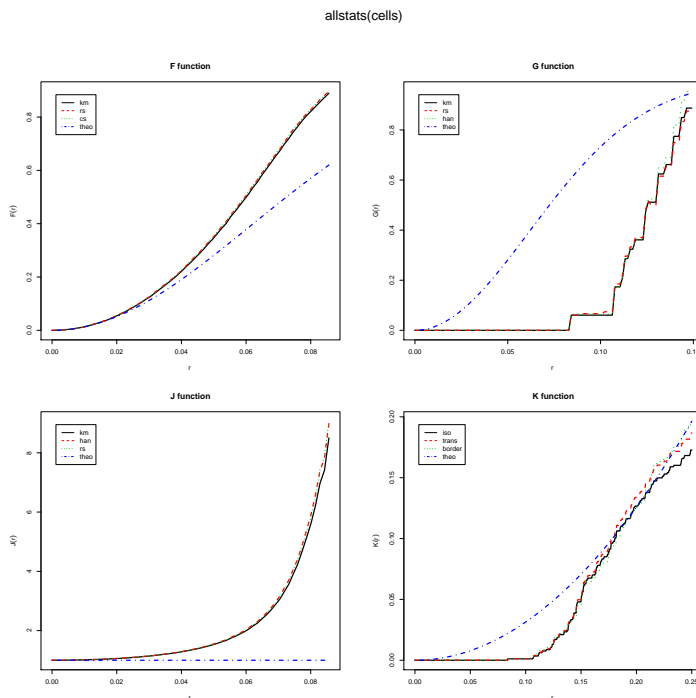
$$J(t) = \frac{\lambda_1}{\lambda_1 + \lambda_2} J_1(t) + \frac{\lambda_2}{\lambda_1 + \lambda_2} J_2(t)$$

where J_1, J_2 are the J -functions of $\mathbf{X}_1, \mathbf{X}_2$ respectively and λ_1, λ_2 are their intensities.

The J function is computed by `Jest`.

The convenient function `allstats` efficiently computes the F , G , J and K functions for a dataset. They can be plotted automatically.

```
> plot(allstats(cells))
```



19.6 Manipulating and plotting summary functions

As explained above, the summary function commands `Fest`, `Gest`, `Kest`, `Lest`, `pcf` etc. return a function value table (an object of class `"fv"`). This is a data frame (i.e. it also belongs to the class `"data.frame"`) with some extra information. One column of the data frame contains

values of the distance argument r , while the other columns contain different estimates of the value of the function, or the theoretical value of the function under CSR.

The following operations are defined on this class:

<code>print.fv</code>	print a summary description
<code>plot.fv</code>	plot the function estimates
<code>as.data.frame</code>	strip extra information (returns a data frame)
<code>\$</code>	extract one column (returns a numeric vector)
<code>[,fv</code>	extract subset (returns an "fv" object)
<code>with.fv</code>	perform calculations with specific columns of data frame
<code>eval.fv</code>	perform calculation on all columns of data frame
<code>cbind.fv</code>	combine several "fv" objects
<code>collapse.fv</code>	combine several redundant "fv" objects
<code>bind.fv</code>	combine an "fv" object and a data frame
<code>smooth.fv</code>	apply smoothing to function values
<code>stieltjes</code>	compute Stieltjes integral with respect to function

To make life easier, there are several options for manipulating the function values.

To manipulate or combine one or more columns of the data frame, it is typically easiest to use the command `with.fv`, which is a method for the generic command `with`. For example:

```
> data(redwood)
> K <- Kest(redwood)
> y <- with(K, iso - theo)
> x <- with(K, r)
```

In this case, the results `x` and `y` are numeric vectors, where `x` contains the values of the distance argument r , and `y` contains the difference between the columns `iso` (isotropic correction estimate) and `theo` (theoretical value for CSR) for the K -function estimate of the redwood seedlings data. For this to work, we have to know that `K` contains columns named `r`, `iso` and `theo`.

The general syntax is `with(X, expr)` where `X` is an "fv" object and `expr` can be any expression involving the names of columns of `X`. The expression can include functions, so long as they are capable of operating on numeric vectors. The expression can also involve the symbol `.` representing "all recommended estimates of the function". Thus:

```
> L <- with(K, sqrt(./pi))
```

computes the estimates of $L(r) = \sqrt{K(r)}$ by all the available edge correction methods. In this case, the result `L` is an "fv" object. You can also get a single numeric result, for example

```
> with(Kest(redwood), max(abs(iso - theo)))
```

```
[1] 0.04945199
```

To plot a transformed function, you can also use the `plot` method. Its second argument is a formula in the R language. The left side of the formula represents what curve or curves will be plotted on the y axis, and the right side determines the x variable for the plot. Thus:

```
> plot(K, sqrt(./pi) ~ r)
```

plots the estimates of $L(r) = \sqrt{K(r)}$, by all the available edge correction methods, against r . The symbol $.$ again signifies “all recommended estimates of the function”. The left hand side of the formula may use the command `cbind` to indicate that several different curves should be plotted. For example, to plot only two curves, giving the isotropic correction estimate and the theoretical value of $K(r)$:

```
> plot(K, cbind(iso, theo) ~ r)
```

The right-hand side can be any expression that evaluates to a numeric vector, and the left hand side is any expression that evaluates to a vector or matrix, of compatible dimensions.

To manipulate or combine one or more “fv” objects, use `eval.fv`. Its argument is an expression containing the names of “fv” objects. For example

```
> K <- Kest(redwood)
> L <- eval.fv(sqrt(K/pi))
```

This can be used to perform computations involving several “fv” objects provided they are compatible (they must have the same vector of r values).

```
> K1 <- Kest(redwood)
> K2 <- Kest(runifpoint(redwood$n, redwood$window))
> DK <- eval.fv(K1 - K2)
```

If these facilities are not sufficient, then direct access to the function values is also possible. A single column of the data frame can be extracted using the `$` operator in the usual way. The object can also be converted to a data frame using `as.data.frame` and the entries extracted in any desired fashion.

19.7 Caveats


The use of summary functions for analysing point patterns has become established across wide areas of applied science, following Ripley’s influential paper [54] and many subsequent textbooks [30, 33, 35, 63, 56, 57, 61] until quite recently.

There is a tendency to apply them uncritically and exclusively. It’s important to remember that

1. the functions F , G and K are defined and estimated under the *assumption that the point process is stationary (homogeneous)*.
2. these summary functions *do not completely characterise the process*.
3. if the process is not stationary, deviations between the empirical and theoretical functions (e.g. \hat{K} and K_{pois}) are not necessarily evidence of interpoint interaction, since they may also be attributable to variations in intensity.

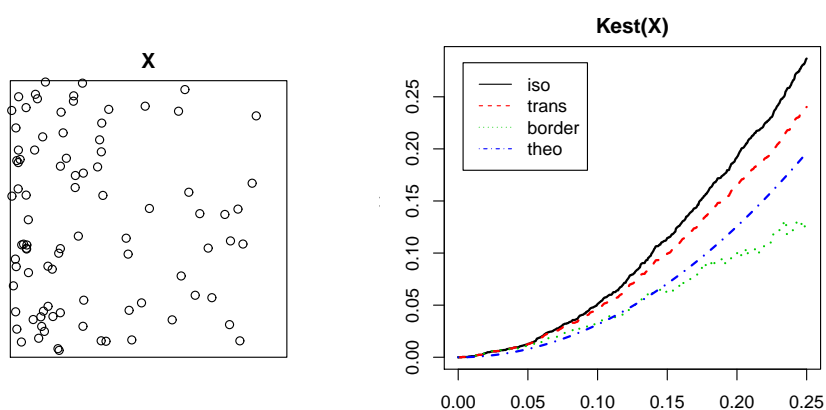
For an example of caveat 2, here is a point process constructed by Baddeley and Silverman [16] which has the same K function as the homogeneous Poisson process:

```
> par(mfrow = c(1, 2))
> X <- rcell(nx = 15)
> plot(X)
> plot(Kest(X))
```



For an example of caveat 3, we generate an inhomogeneous Poisson pattern and apply the ordinary K function estimator. The result appears to show clustering, but this is an artefact of the spatial inhomogeneity.

```
> par(mfrow = c(1, 2))
> X <- rpoispp(function(x, y) {
+   300 * exp(-3 * x)
+ })
> plot(X)
> plot(Kest(X))
```



20 Simulation envelopes and goodness-of-fit tests

Although summary statistics such as the K -function are intended primarily for exploratory purposes, it is also possible to use them as a basis for statistical inference.


20.1 Envelopes and Monte Carlo tests

20.1.1 Motivation


In Section 19 we examined plots of the K -function to judge whether a point pattern dataset is completely random. The K -function estimated from the point pattern data, $\hat{K}(r)$, was compared graphically with the theoretical K -function for a completely random pattern, $K_{\text{pois}}(r) = \pi r^2$. In the toy examples, large discrepancies between \hat{K} and K_{pois} were observed, indicating that the toy examples were not completely random patterns.

However, because of random variability, we will never obtain perfect agreement between \hat{K} and K_{pois} , even with a completely random pattern. Try typing `plot(Kest(rpoispp(50)))` a few times to get an idea of the inherent variability.

The following plot shows the K -function estimated from the `cells` dataset (thick line), and also the K -functions of 20 simulated realisations of CSR with the same intensity (thin lines).



The next plot shows the upper and lower *envelopes* of the simulated K -functions, that is, the maximum and minimum values of $\hat{K}(r)$ for each value of r . The region between the envelopes is shaded.



Clearly, the K -function estimated from the `cells` data lies outside the typical range of values of the K -function for a completely random pattern.

To conclude formally that there is a ‘significant’ difference between \hat{K} and K_{pois} , we use the language of hypothesis testing. Our *null hypothesis* is that the data point pattern is a realisation of complete spatial randomness. The *alternative hypothesis* is that the data pattern is a realisation of another, unspecified point process.

20.1.2 Monte Carlo tests

A *Monte Carlo test* is a test based on simulations from the null hypothesis. The principle was originated independently by Barnard [18] and Dwass [39]. It was applied in spatial statistics by Ripley [54, 56] and Besag [22, 23]. See also [41]. Monte Carlo tests are a special case of *randomisation tests* which are commonly used in nonparametric statistics.

Suppose the reference curve is the theoretical K function for CSR. Generate M independent simulations of CSR inside the study region W . Compute the estimated K functions for each of these realisations, say $\hat{K}^{(j)}(r)$ for $j = 1, \dots, M$. Obtain the pointwise upper and lower envelopes of these simulated curves,

$$\begin{aligned} L(r) &= \min_j \hat{K}^{(j)}(r) \\ U(r) &= \max_j \hat{K}^{(j)}(r). \end{aligned}$$

For any fixed value of r , consider the probability that $\hat{K}(r)$ lies outside the envelope $[L(r), U(r)]$ for the simulated curves. If the data came from a uniform Poisson process, then $\hat{K}(r)$ and $\hat{K}^{(1)}(r), \dots, \hat{K}^{(M)}(r)$ are statistically equivalent and independent, so this probability is equal to $2/(M+1)$ by symmetry. That is, the test which rejects the null hypothesis of a uniform Poisson process when $\hat{K}(r)$ lies outside $[L(r), U(r)]$, has exact significance level $\alpha = 2/(M+1)$. Instead of the pointwise maximum and minimum, one could use the pointwise order statistics (the pointwise k th largest and k smallest values) giving a test of exact size $\alpha = 2k/(M+1)$.

20.1.3 Envelopes in spatstat

In `spatstat` the function `envelope` computes the pointwise envelopes.

```
> data(cells)
> E <- envelope(cells, Kest, nsim = 39, rank = 1)

> E
```

```
Pointwise critical envelopes for K(r)
Edge correction: iso
Obtained from 39 simulations of CSR
Significance level of pointwise Monte Carlo test: 2/40 = 0.05
Data: cells
Function value object (class fv)
for the function r -> K(r)
Entries:
id      label      description
--      -
r       r          distance argument r
```

obs	obs(r)	observed value of $K(r)$ for data pattern
theo	theo(r)	theoretical value of $K(r)$ for CSR
lo	lo(r)	lower pointwise envelope of $K(r)$ from simulations
hi	hi(r)	upper pointwise envelope of $K(r)$ from simulations

Default plot formula:


`. ~ r`

<environment: 0x5d76db8>

Recommended range of argument r : [0, 0.25]

Available range of argument r : [0, 0.25]

```
> plot(E, main = "pointwise envelopes")
```



For example if r had been fixed at $r = 0.10$ we would have rejected the null hypothesis of CSR at the 5% level. The value $M = 39$ is the smallest to yield a two-sided test with significance level 5%.

Tip: A common and dangerous mistake is to misinterpret the simulation envelopes as “confidence intervals” around \hat{K} . They cannot be interpreted as a measure of accuracy of the estimated K function! They are the critical values for a test of the hypothesis that $K(r) = \pi r^2$. **They assume that the pattern is completely random.** [See Section 21 for ways of making confidence intervals for $K(r)$.]

The value returned by `envelope` is an object of class “fv” that can be manipulated in the usual way: you can plot it, transform it, extract columns, and so on (see Section 19.6 on page 128).

20.1.4 Simultaneous Monte Carlo test

Note that the theory of the Monte Carlo test, as presented above, requires that r be fixed in advance. If we plot the envelope and check whether the empirical K function ever wanders

outside the envelope, this is equivalent to choosing the value of r in a data-dependent way, and the true significance level is higher (less ‘significant’).

To avoid this problem we can construct *simultaneous critical bands* which have the property that, under H_0 , the probability that \hat{K} ever wanders outside the critical bands is exactly 5%.


One simple way to achieve this is to compute, for each estimate $\hat{K}(r)$, its maximum deviation from the Poisson K function, $D = \max_r |\hat{K}(r) - K_{\text{pois}}(r)|$. This is computed for each of the M simulated datasets, and the maximum value D_{max} obtained. Then the upper and lower limits are

$$\begin{aligned} L(r) &= \pi r^2 - D_{\text{max}} \\ U(r) &= \pi r^2 + D_{\text{max}}. \end{aligned}$$

The estimated K function for the data transgresses these limits if and only if the D -value for the data exceeds D_{max} . Under H_0 this occurs with probability $1/(M + 1)$. Thus, a test of size 5% is obtained by taking $M = 19$.

```
> E <- envelope(cells, Kest, nsim = 19, rank = 1, global = TRUE)
```


```
> plot(E, main = "global envelopes")
```



A more powerful test is obtained if we (approximately) stabilise the variance, by using the L function in place of K .

```
> E <- envelope(cells, Lest, nsim = 19, rank = 1, global = TRUE)
```

```
> plot(E, main = "global envelopes of L(r)")
```



20.1.5 Envelopes for any fitted model


In the explanation above, we assumed that the null hypothesis was CSR (complete spatial randomness, a uniform Poisson process). In fact the Monte Carlo testing rationale can be applied to any point process model serving as a null hypothesis. We simply have to generate simulated realisations from the null hypothesis, and compute the summary function for each simulated realisation.

To simulate from a fitted point process model (object of class "ppm"), call the `envelope` function, giving the fitted model as the first argument of `envelope`. Then the simulated patterns will be generated according to this fitted model. The original data point pattern, to which the model was fitted, is stored in the fitted model object; the original data are extracted and the summary function for the data is also computed.

The following code fits an inhomogeneous Poisson process to the Beilschmiedia pattern, then generates simulation envelopes of the L function by simulating from the fitted inhomogeneous Poisson model.

```
> data(bei)
> fit <- ppm(bei, ~elev + grad, covariates = bei.extra)
> E <- envelope(fit, Lest, nsim = 19, global = TRUE, correction = "border")

> plot(E, main = "envelope for inhomogeneous Poisson")
```



20.1.6 Envelopes based on any simulation procedure

Envelopes can also be computed using any user-specified procedure to generate the simulated realisations. This allows us to perform randomisation tests, for example.

The simulation procedure should be encoded as an R expression, which will be evaluated each time a simulation is required. For example if we type

```
> sim <- expression(rpoispp(100))
```


then each time the expression `sim` is evaluated, it will yield a different random outcome of the Poisson process with intensity 100 in the unit square.

This expression should be passed to the `envelope` function as the argument `simulate`.

The following code generates simulation envelopes for the L function based on simulations of CSR which have the same number of points as the data pattern.

```
> data(cells)
> n <- npoints(cells)
> W <- as.owin(cells)
> e <- expression(runifpoint(n, W))
> E <- envelope(cells, Lest, nsim = 19, global = TRUE, simulate = e)

> plot(E, main = "envelope with fixed n")
```



20.1.7 Envelopes based on a set of point patterns

Envelopes can also be computed from a user-supplied list of point patterns, instead of the simulated point patterns generated by a chosen simulation procedure. The argument `simulate` can be a list of point patterns:

```
> data(cells)
> Xlist <- list()
> for (i in 1:99) Xlist[[i]] <- runifpoint(42)
> envelope(cells, Kest, nsim = 99, simulate = Xlist)
```

The argument `simulate` can also be an `envelope` object. This improves efficiency and consistency if, for example, we are going to calculate the envelopes of several different summary statistics.

```
> data(cells)
> EK <- envelope(cells, Kest, nsim = 99, savepatterns = TRUE)
> Ep <- envelope(cells, pcf, nsim = 99, simulate = EK)
```

In the first call to `envelope`, the argument `savepatterns=TRUE` indicates that we want to save the simulated point patterns. These are stored in the object `EK`. Then in the second call to `envelope`, the simulated patterns are extracted from `EK` and used to compute the envelopes of the pair correlation function.

20.1.8 Envelopes based on sample mean & variance

Envelopes can be constructed using the sample mean and sample variance of the simulations. By default the envelope is the sample mean plus or minus 2 times the sample standard deviation. This is useful for understanding the variability of the summary function. Be aware that these envelopes do not have the same significance interpretation.


```
> envelope(cells, Kest, nsim = 100, VARIANCE = TRUE)
```

21 Spatial bootstrap methods

Spatstat includes some elementary bootstrap methods for estimating the variance and the distribution of a summary statistic. These can be used to construct **confidence intervals** for the true value of $K(r)$, for example.

The function `varblock` divides the point pattern's window into several quadrats. The specified summary statistic is applied to each of the corresponding sub-patterns. Then the pointwise sample mean, sample variance and sample standard deviation of these summary statistics are computed. The two-standard-deviation confidence intervals are computed. This is an elementary bootstrap estimate of the sampling variance of the summary statistic [35, eq. (4.21), p. 52].

```
> data(finpines)
> Kci <- varblock(finpines, Kest, nx = 3, ny = 3)
> plot(Kci, iso ~ r, shade = c("loiso", "hiiso"), main = "Confidence Interval")
```



The function `quadratresample` generates a randomly resampled version of the data point pattern. The spatial domain is divided into several rectangular quadrats of equal shape and size; the sub-patterns in each quadrat are extracted into a list; and the list is randomly permuted or resampled. The resampled sub-patterns are replaced in the original domains.


```
> X <- unmark(finpines)
> Xlist <- quadratresample(X, nx = 4, ny = 4, nsamples = 100)
> E <- envelope(X, Kest, simulate = Xlist, nsim = 100, VARIANCE = TRUE)
> plot(E)
```

22 Simple models of non-Poisson patterns

A point process that is not Poisson can be said to exhibit ‘interaction’ or dependence between the points. It’s time to introduce some models for such processes. This section covers simple models that are derived from the Poisson process, and still retain some of the tractable features of the Poisson model.

22.1 Poisson cluster processes


In a Poisson cluster process, we begin with a Poisson process \mathbf{Y} of ‘parent’ points. Each ‘parent’ point $y_i \in \mathbf{Y}$ then gives rise to a finite set Z_i of ‘offspring’ points according to some stochastic mechanism. The set comprising all the offspring points forms a point process \mathbf{X} . Only \mathbf{X} is observed.



An example is the *Matérn cluster process* in which the parent points come from a homogeneous Poisson process with intensity κ , and each parent has a Poisson (μ) number of offspring, independently and uniformly distributed in a disc of radius r centred around the parent.

The Matérn cluster process can be generated in **spatstat** using the command `rMatClust`. [By convention, random data generators in R always have names beginning with **r**.]

```
> plot(rMatClust(kappa = 10, r = 0.1, mu = 5))
```



Other Poisson cluster processes implemented in **spatstat** are


- **rThomas**: the *Thomas process*, in which each cluster consists of a Poisson(μ) number of random points, each having an isotropic Gaussian $N(0, \sigma^2 I)$ displacement from its parent.
- **rGaussPoisson**: the *Gauss-Poisson process* in which each cluster is either a single point or a pair of points.
- **rNeymanScott**: the general *Neyman-Scott* cluster process in which the cluster mechanism is arbitrary.

22.2 Cox processes

A Cox point process is effectively a Poisson process with a random intensity function. Let $\Lambda(u)$ be a random function with non-negative values, defined at all locations $u \in \mathbb{R}^2$. Conditional on Λ , let \mathbf{X} be a Poisson process with intensity function Λ . Then \mathbf{X} is a Cox process.

A trivial example is the “mixed Poisson” process in which we generate a random variable Λ and, conditional on Λ , generate a uniform Poisson process with intensity Λ . Following are three different realisations of this process:

```
> par(mfrow = c(1, 3))
> for (i in 1:3) {
+   lambda <- rexp(1, 1/100)
+   X <- rpoispp(lambda)
+   plot(X)
+ }
> par(mfrow = c(1, 1))
```



Moments of Cox processes are tractable (in terms of the moments of Λ). The intensity function of \mathbf{X} is $\lambda(u) = \mathbb{E}[\Lambda(u)]$.

A Cox model is the analogue of a ‘random effects’ model. It is always overdispersed relative to a Poisson process (i.e. the variance of the number of points falling in a region, is greater than the mean). Cox processes are the most convenient models for clustered point patterns. A particularly interesting and useful class is that of *log-Gaussian Cox processes (LGCP)* in which $\log \Lambda(u)$ is a Gaussian random function [51, 50].

The Matérn Cluster process and the Thomas process are both Cox processes.


Currently there are no functions in **spatstat** for generating the general Cox process, but if you have a way of generating realisations of a random function Λ of interest, then you can use **rpoispp** to generate the Cox process. The intensity argument **lambda** of **rpoispp** can be a **function(x,y)** or a pixel image.

22.3 Thinned processes

‘*Thinning*’ means deleting some of the points from a point pattern. Under ‘*independent thinning*’ the fate of each point is independent of other points. When independent thinning is applied to a Poisson process, the resulting process of retained points is Poisson. To get a non-Poisson process we need some kind of *dependent thinning* mechanism.


In *Matérn’s Model I*, a homogeneous Poisson process \mathbf{Y} is first generated. Any point in \mathbf{Y} that lies closer than a distance r from the nearest other point of \mathbf{Y} , is deleted. Thus, pairs of close neighbours annihilate each other.

```
> plot(rMaternI(70, 0.05))
```



In *Matérn's Model II*, the points of the homogeneous Poisson process \mathbf{Y} are marked by 'arrival times' t_i which are independent and uniformly distributed in $[0, 1]$. Any point in \mathbf{Y} that lies closer than distance r from another point that has an earlier arrival time, is deleted.

```
> plot(rMaternII(70, 0.05))
```




22.4 Sequential models

In a sequential model, we start with an empty window, and the points are placed into the window one-at-a-time, according to some criterion.

In Simple Sequential Inhibition, each new point is generated uniformly in the window and independently of preceding points. If the new point lies closer than r units from an existing point, then it is rejected and another random point is generated. The process terminates when no further points can be added.

```
> plot(rSSI(0.05, 200))
```



Sequential point processes are the easiest way to generate highly ordered patterns with high intensity.

23 Model-fitting using summary statistics

Summary statistics can also be used to fit models to data.

In the ‘method of moments’ we estimate a parameter θ by solving

$$\mathbb{E}_\theta[S(\mathbf{X})] = S(\mathbf{x})$$

where $S(\mathbf{x})$ is the observed value of a statistic S for our data \mathbf{x} , and the left side is the theoretical mean of S for the model governed by parameter θ .

The analogue for point process models is to fit the model by matching a summary statistic such as the K function to its theoretical value under the model.

23.1 Fitting a cluster process

In a precious few cases, the K function of a point process is known exactly, as an analytic expression in terms of the model parameters. These happy cases include many Neyman-Scott cluster processes. For example, the K -function of the Thomas process (Section 22.1, page 140) with parameters $\theta = (\kappa, \mu, \sigma)$ is

$$K_\theta(r) = \pi r^2 + \frac{1}{\kappa} \left(1 - \exp\left(-\frac{r^2}{4\sigma^2}\right)\right). \quad (26)$$

We can use this to fit a Thomas model to data. We determine the values of the parameters $\theta = (\kappa, \mu, \sigma)$ to achieve the best match between $K_\theta(r)$ and the estimated K -function of the data, $\hat{K}(r)$. The best match is determined by minimising the discrepancy between the two functions over some range $[a, b]$:

$$D(\theta) = \int_a^b \left| \hat{K}(r)^q - K_\theta(r)^q \right|^p dr \quad (27)$$

where $0 \leq a < b$, and where $p, q > 0$ are indices. This method was originally advocated by Peter Diggle and collaborators, and is now known as the *method of minimum contrast*. See [35].

The command `kppm` fits cluster point process models by the method of minimum contrast. To fit the Thomas model to the redwood data:

```
> data(redwood)
> fit <- kppm(redwood, ~1, "Thomas")
```


The first argument to `kppm` is a point pattern dataset. The second argument is a formula (with no left hand side) describing the log intensity of the model; the formula `~1` indicates a stationary process (see section 25.3 for nonstationary models). The third argument is the name of the cluster mechanism; currently the only options are `"Thomas"` and `"MatClust"`.

The fitted model, `fit`, is an object of class `kppm`. There are methods for printing and plotting objects of this class.

```
> fit
```

```
Stationary cluster point process model
Fitted to point pattern dataset redwood
Fitted using the K-function
Cluster model: Thomas process
Fitted parameters:
      kappa      sigma      mu
23.55389789  0.04704965  2.63226071
```

```
> plot(fit)
```



The plot shows the theoretical K function of the fitted Thomas process (`fit`), three non-parametric estimates of the K function (`iso`, `trans`, `border`) and the Poisson K function (`theo`).

At present, the only cluster process models that can be fitted using `kppm` are the Thomas process and the Matérn cluster process. To fit the Matérn cluster process to the redwood data,

```
> fitM <- kppm(redwood, ~1, "MatClust")
```


A fitted model returned by `kppm` can be simulated immediately:

```
> plot(simulate(fit, nsim = 4))
```

The command `simulate` is generic; here we have used the method `simulate.kppm`.

Simulation envelopes are also available, using the method `envelope.kppm`.

```
> plot(envelope(fit, Lest, nsim = 39))
```



23.2 Fitting cluster models using the pair correlation

Minimum contrast estimation can be applied to any summary statistic. In particular we can use the pair correlation function instead of the K -function.

```
> fitp <- kppm(redwood, ~1, "Thomas", statistic = "pcf")
> fitp
```

```
Stationary cluster point process model
Fitted to point pattern dataset redwood
Fitted using the pair correlation function
Cluster model: Thomas process
Fitted parameters:
      kappa      sigma      mu
25.30537171 0.03968295 2.45007268
```

23.3 Fitting other models with known K function

Apart from cluster processes, there are certain other point process models for which the K -function is known as a function of the model parameters. Minimum contrast methods are also available for these models.

One special case is the *log-Gaussian Cox processes* described in detail in [51]. To fit a log-Gaussian Cox process with exponential covariance function to the redwood data:

```
> fit <- lgcp.estK(redwood, c(sigma2 = 0.1, alpha = 1))
> fit
```

```
Minimum contrast fit (object of class minconfit)
Model: log-Gaussian Cox process
Fitted by matching theoretical K function to Kest(redwood)
Parameters fitted by minimum contrast ($par):
      sigma2      alpha
```

```

1.0485493 0.0997963
Derived parameters of log-Gaussian Cox process ($modelpar):
      sigma2      alpha      mu
1.0485493 0.0997963 3.6028597
Converged successfully after 145 iterations.
Domain of integration: [ 0 , 0.25 ]
Exponents: p= 2, q= 0.25

```

The second argument to `lgcp.estK` gives initial values for the model parameters σ^2 and α .

The result of `lgcp.estK` is an object of class `minconfit` (representing a ‘minimum contrast fit’). There are methods for printing and plotting the fit. Simulation of these models has not yet been implemented in `spatstat`.

Estimation can also be based on the pair correlation function:

```

> fit <- lgcp.estpcf(redwood, c(sigma2 = 0.1, alpha = 1))
> fit

```

```

Minimum contrast fit (object of class minconfit)
Model: log-Gaussian Cox process
Fitted by matching theoretical g function to pcf(redwood)
Parameters fitted by minimum contrast ($par):
      sigma2      alpha
1.30244010 0.07011464
Derived parameters of log-Gaussian Cox process ($modelpar):
      sigma2      alpha      mu
1.30244010 0.07011464 3.47591433
Converged successfully after 83 iterations.
Domain of integration: [ 0.0004883 , 0.25 ]
Exponents: p= 2, q= 0.25

```

23.4 Generic algorithm for minimum contrast

The command `mincontrast` is a generic fitting algorithm for the method of minimum contrast. It can be used in any context where the theoretical function can be computed exactly from the model parameters. A basic call to `mincontrast` is:

```

> mincontrast(observed, theoretical, starpar)

```


where `observed` is an object of class `"fv"` containing the summary function calculated from the data; `theoretical` is a function which returns the theoretical value of the summary function for a given parameter value; and `startpar` is a vector of initial values of the model parameters. For details, see the help file for `mincontrast`.

For the vast majority of point process models, the true K function $K_\theta(r)$ is not known analytically in terms of the parameter θ . In principle we could use Monte Carlo simulation to determine an approximation to $K_\theta(r)$, for any given θ , by generating a large number of simulated realisations of the process with parameter θ , computing the estimated K -function for each realisation, and taking the pointwise sample average. It's possible to do this in `spatstat` using the generic algorithm `mincontrast`. Details are not given here as it is rather fiddly at present, and will change soon.

24 Exploring local features

The `shapley` dataset is an example of a point pattern which is clearly not homogeneous. The data comes from a radioastronomical survey of galaxies in the Shapley Galaxy Concentration: each point is a galaxy in the distant universe. There are very dense concentrations of galaxies in some parts of the survey area.

```
> data(shapley)
> X <- unmark(shapley)
> plot(X, pch = ".", main = "")
```



Exploratory techniques for investigating localised features in a point pattern include LISA (Local Indicators of Spatial Association), nearest-neighbour cleaning, and data sharpening.


In *LISA methods*, a summary statistic is separated into contributions from each of the data points. For example the K function is expressed as a sum of the *local K functions* of each of the data points. These local functions are then compared, and classified into several groups of functions, perhaps using principal component analysis [2, 29, 28].

The `spatstat` functions `localK`, `localL`, `localpcf` compute local versions of the K -function, L -function and pair correlation function, respectively.

Nearest-neighbour cleaning [26] groups the points into two classes — ‘feature’ and ‘noise’ — on the basis of their nearest-neighbour distances. It is quick and often very useful.


```
> Z <- nnclean(X, k = 17)
> plot(Z, chars = c(".", "+"), main = "nearest neighbour cleaning")
```

```
noise feature
" ."      "+"
```



In *data sharpening* [27] the points effectively exert a force of attraction on each other, and are allowed to move in the direction of the resultant force. This tends to enhance tight linear concentrations of points.

```
> Y <- sharpen(X, sigma = 0.5, edgecorrect = TRUE)
> plot(Y, pch = ".", main = "data sharpened")
```



25 Adjusting for inhomogeneity

If a point pattern is known or suspected to be spatially inhomogeneous, then our statistical analysis of the pattern should take account of this inhomogeneity.

25.1 Inhomogeneous K function

There is a modification of the K function that applies to inhomogeneous processes [7]. If $\lambda(u)$ is the true intensity function of the point process \mathbf{X} , then the idea is that each point x_i will be weighted by $w_i = 1/\lambda(x_i)$.

The *inhomogeneous K -function* is defined as

$$K_{\text{inhom}}(r) = \mathbb{E} \left[\sum_{x_j \in \mathbf{X}} \frac{1}{\lambda(x_j)} \mathbf{1} \{0 < \|u - x_j\| \leq r\} \mid u \in \mathbf{X} \right] \quad (28)$$

assuming that this does not depend on location u . Thus, $\lambda(u)K(r)$ is the expected total ‘weight’ of all random points within a distance r of the point u , where the ‘weight’ of a point x_i is $1/\lambda(x_i)$.

If the process is actually homogeneous, then $\lambda(u)$ is constant and $K_{\text{inhom}}(r)$ reduces to the usual K function (21).

It turns out that, for an inhomogeneous Poisson process with intensity function $\lambda(u)$, the inhomogeneous K function is

$$K_{\text{inhom, pois}}(r) = \pi r^2 \quad (29)$$

exactly as for the homogeneous case.

The standard estimators of K can be extended to the inhomogeneous K function:

$$\hat{K}_{\text{inhom}}(r) = \frac{1}{D} \sum_i \sum_{j \neq i} \frac{\mathbf{1} \{\|x_i - x_j\| \leq r\}}{\hat{\lambda}(x_i)\hat{\lambda}(x_j)} e(x_i, x_j; r) \quad (30)$$

where $e(u, v, r)$ is an edge correction weight as before, and $\hat{\lambda}(u)$ is an estimate of the intensity function $\lambda(u)$.

The denominator D in (30) is either the area of the window $D_1 = \text{area}(W)$, or


$$D_2 = \sum_i \frac{1}{\hat{\lambda}(x_i)}$$

which is an unbiased estimator of $\text{area}(W)$ if the intensity is correctly estimated. The denominator D_2 is often preferred on grounds of statistical performance, because it introduces a data-dependent normalisation.

The inhomogeneous K function is computed by the command `Kinhom(X, lambda)` where X is the point pattern and `lambda` is the estimated intensity function. Here `lambda` may be a pixel image, a `function(x,y)` in the R language, a numeric vector giving the values $\hat{\lambda}(x_i)$ at the data points x_i only, or it may be omitted (and will then be estimated from X). By default, the data-dependent denominator D_2 is used.

There remains the question of how to estimate the intensity function $\lambda(u)$. It is usually advisable to obtain the intensity estimate $\hat{\lambda}(u)$ by fitting a parametric model, to avoid overfitting. Here is an example for the tropical rainforest data, using the covariate data to suggest a model for the intensity.

```
> data(bei)
> fit <- ppm(bei, ~elev + grad, covariates = bei.extra)
> lam <- predict(fit, locations = bei)
> Ki <- Kinhom(bei, lam)
> plot(Ki, main = "Inhomogeneous K function")
```



The plot suggests that, even after accounting for dependence on altitude and slope, the trees still appear to be clustered.

The intensity function $\lambda(u)$ could also be estimated by kernel smoothing the point pattern data. However, notice that the estimator (30) of the inhomogeneous K function depends on the estimated intensity values at the *data points*, $\hat{\lambda}(x_i)$. These are positively biased estimates of the true values $\lambda(x_i)$. In order to avoid bias, the value $\hat{\lambda}(x_i)$ should be estimated by kernel smoothing of the point pattern with the point x_i *deleted*. This “leave-one-out” estimator is implemented in `Kinhom` and is invoked when the argument `lambda` is not given:

```
> Ki2 <- Kinhom(bei)
> plot(Ki2, main = "Kinhom using leave-one-out")
```

	lty	col	key	label	
bord.modif	1	1	bord.modif	K[bordm](r)	
border	2	2	border	K[bord](r)	
theo	3	3	theo	K[pois](r)	
					meaning
bord.modif	modified		border-corrected	estimate of K(r)	
border			border-corrected	estimate of K(r)	
theo			theoretical	Poisson K(r)	

(the smoothing parameter σ can also be controlled.)

The inhomogeneous analogue of the L -function is defined by

$$\hat{L}_{\text{inhom}}(r) = \sqrt{\frac{1}{\pi} \hat{K}_{\text{inhom}}(r)}$$

This can be computed using `Linhom`. For an inhomogeneous Poisson process, $L_{\text{inhom}}(r) \equiv r$.

The inhomogeneous analogue of the pair correlation function can be defined, similarly to the homogeneous case, as

$$g_{\text{inhom}}(r) = \frac{K'_{\text{inhom}}(r)}{2\pi r}.$$

It has the same interpretation, namely, that $g_{\text{inhom}}(r)$ is the probability of observing a pair of points at certain locations separated by a distance r , divided by the corresponding probability for a Poisson process of the same (inhomogeneous) intensity.

The inhomogeneous pair correlation function is computed by `pcfinhom`:

```
> g <- pcfinhom(bei)
```

Incidentally there is also a function `pcf.fv` which will convert any K -function into a pair correlation function by numerical differentiation. Thus the following is an alternative:

```
> g <- pcf(Kinhom(bei))
```

25.2 Inhomogeneous cluster models

The inhomogeneous Poisson process was described in Section 15.1. We can also introduce spatial inhomogeneity into any of the non-Poisson models described in Section 22.

In the case of Poisson cluster processes (Section 22.1) we can introduce inhomogeneity in either the parent process or the offspring processes.

To make the *parents* inhomogeneous, we simply generate the parent points from an inhomogeneous Poisson process with some intensity function $\kappa(u)$.


To make the *clusters* inhomogeneous, we use a clever construction by Waagepetersen [65]. For a parent point at location (x_0, y_0) , the offspring are generated from a Poisson process with intensity $\beta(x, y) = \mu(x, y)f(x - x_0, y - y_0)$, where $f(u, v)$ is either the Gaussian probability density (for the Thomas process) or the uniform probability density in a disc (for the Matérn cluster process), and $\mu(x, y)$ is the *reference* or *modulating* intensity. The number of offspring from a given parent (x_0, y_0) is a Poisson random variable with mean

$$B(x_0, y_0) = \int \beta(x, y) dx dy = \int f(x - x_0, y - y_0) \mu(x, y) dx dy.$$

The simulation algorithms `rMatClust` and `rThomas` allow both these options. If the parent intensity parameter `kappa` is given as a `function(x,y)` or a pixel image, then the parents are

Poisson with inhomogeneous intensity **kappa**. If the offspring mean parameter **mu** is given as a `function(x,y)` or a pixel image, then this determines an inhomogeneous reference density for the clusters.

```
> Z <- as.im(function(x, y) {
+     6 * exp(2 * x - 1)
+ }, owin())
> plot(rMatClust(10, 0.05, Z))
```



25.3 Fitting inhomogeneous models by minimum contrast

Minimum contrast methods can be applied to inhomogeneous point process models.

In principle we could fit any model (homogeneous or inhomogeneous) by the method of minimum contrast using any summary statistic. However, the method works best when we have an exact formula for the true value of the summary function for the model, expressed as a function of the parameters of the model.

Waagepetersen [65] pointed out that, if we take a Thomas process or Matérn cluster process (or in general a Neyman-Scott process) with **homogeneous** parent intensity κ and **inhomogeneous** cluster reference density $\mu(u)$, then the overall intensity of the process is

$$\lambda(u) = \kappa \mu(u)$$

and the *inhomogeneous* K -function is the same as it would be if μ were constant.

Thus, we can fit a Thomas process or Matérn cluster process with inhomogeneous clusters as follows:

1. estimate the inhomogeneous intensity $\lambda(u)$ of the process.
2. derive an estimate of the inhomogeneous K -function.
3. use the method of minimum contrast to estimate the parent intensity κ and the cluster scale parameter (Gaussian standard deviation or disc radius), exactly as we would in the homogeneous case.

The command `kppm` performs this algorithm using a parametric model for the trend:

```
> data(bei)
> fit <- kppm(bei, ~elev + grad, "Thomas", covariates = bei.extra)
> fit
```

```
Inhomogeneous cluster point process model
Fitted to point pattern dataset bei
Fitted using the inhomogeneous K-function
Trend formula:~elev + grad
```

```
Fitted coefficients for trend formula:
```

```
(Intercept)      elev      grad
-8.55862210  0.02140987  5.84104065
```

```
Cluster model: Thomas process
```

```
Fitted parameters:
```

```
      kappa      sigma
0.0004290453 5.4110425537
```


In this example, `kppm` first estimates the intensity by fitting the model `ppm(bei, ~elev+grad, covariates = bei.extra)`. Then `predict.ppm` is used to compute the predicted intensity at the data points, and this is passed to `Kinhom` to calculate the inhomogeneous K function. The parameters of the Thomas process are estimated from the inhomogeneous K function by minimum contrast.

The result of `kppm` can be printed, plotted, simulated and “enveloped” as before.

25.4 Local scaling


Locally-scaled point processes and summary functions [53] provide an alternative to the concept of locally-weighted K functions (Section 25.1). In essence, the point process is assumed to be equivalent, in small regions, to a rescaled version of a ‘template’ process, where the template process is stationary, and the rescaling factor can vary from place to place.

```
> data(bronzefilter)
> X <- unmark(bronzefilter)
> plot(X)
```



Spatstat provides the commands `Kscaled` and `Lscaled` which compute the locally-scaled K and L functions. Their syntax is similar to `Kinhom`.

```
> fit <- ppm(X, ~x)
> lam <- predict(fit)
> plot(Kscaled(X, lam), xlim = c(0, 2))
```



PART VI. GIBBS MODELS

Part VI of the workshop explains Gibbs point process models and how to apply them.

26 Gibbs models

One way to construct a statistical model (in any field of statistics) is to write down its probability density. Advantages of doing this are:

- the functional form of the density reflects its probabilistic properties.
- terms or factors in the density often have an interpretation as ‘components’ of the model.
- it is easy to introduce terms that represent the dependence of the model on covariates, etc.

This approach is useful provided the density *can* be written down, and provided the density is tractable.

Spatial point process models that are constructed by writing down their probability densities are called ‘**Gibbs processes**’. Good references on Gibbs point processes are [63, 31].

26.1 Probability densities

It is possible to define probability densities for spatial point processes that live inside a bounded window W .

The probability density will be a function $f(\mathbf{x})$ defined for each finite configuration $\mathbf{x} = \{x_1, \dots, x_n\}$ of points $x_i \in W$ for any $n \geq 0$. Notice that the number of points n is not fixed, and may be zero. Apart from this peculiarity, probability densities for point processes behave much like probability densities in more familiar contexts.

That’s all you need to know for applications. **If you’re interested in the mathematical technicalities, read on; otherwise, skip to section 26.2.**

A point process \mathbf{X} inside W is defined to have probability density f if and only if, for any nonnegative integrable function h ,

$$\mathbb{E}[h(\mathbf{X})] = e^{-|W|}h(\emptyset)f(\emptyset) + e^{-|W|} \sum_{n=1}^{\infty} \frac{1}{n!} \int_W \cdots \int_W h(\{x_1, \dots, x_n\})f(\{x_1, \dots, x_n\}) dx_1 \cdots dx_n \quad (31)$$

where $|W|$ denotes the area of W .

In particular, the probability that \mathbf{X} contains exactly n points is

$$p_n = \mathbb{P}\{n(\mathbf{X}) = n\} = \frac{e^{-|W|}}{n!} \int_W \cdots \int_W f(\{x_1, \dots, x_n\}) dx_1 \cdots dx_n$$

for $n \geq 1$ and $p_0 = \mathbb{P}\{n(\mathbf{X}) = 0\} = e^{-|W|}f(\emptyset)$. Given that there are exactly n points, the conditional joint density of the locations x_1, \dots, x_n is $f(\{x_1, \dots, x_n\})/p_n$.

26.2 Poisson processes

The uniform Poisson process with intensity 1 has probability density $f(\mathbf{x}) \equiv 1$.

The uniform Poisson process in W with intensity λ has probability density

$$f(\mathbf{x}) = \alpha \lambda^{n(\mathbf{x})} \quad (32)$$

where $n(\mathbf{x})$ is the number of points in the configuration \mathbf{x} , and the constant α is

$$\alpha = e^{(1-\lambda)|W|}.$$

The inhomogeneous Poisson process in W with intensity function $\lambda(u)$ has probability density

$$f(\mathbf{x}) = \alpha \prod_{i=1}^n \lambda(x_i). \quad (33)$$

where the constant α is

$$\alpha = \exp \left[\int_W (1 - \lambda(u)) \, du \right].$$

The densities (32) and (33) are products of terms associated with individual points x_i . This reflects the conditional independence property (PP4) of the Poisson process.

26.3 Pairwise interaction models

In order to construct spatial point processes which exhibit interpoint interaction (stochastic dependence between points), we need to introduce terms in the density that depend on more than one point. The simplest are *pairwise interaction models*, which have probability densities of the form

$$f(\mathbf{x}) = \alpha \left[\prod_{i=1}^{n(\mathbf{x})} b(x_i) \right] \left[\prod_{i < j} c(x_i, x_j) \right] \quad (34)$$

where α is a normalising constant, $b(u)$, $u \in W$ is the ‘first order’ term, and $c(u, v)$, $u, v \in W$ is the ‘second order’ or ‘pairwise interaction’ term. The pairwise interaction term introduces dependence between points. The interaction function must be symmetric, $c(u, v) = c(v, u)$. In principle we are free to choose any functions b and c , provided the resulting density is integrable (the right side of (31) should be finite when $h \equiv 1$).

26.3.1 Hard core process


If we take $b(u) \equiv \beta$ and

$$c(u, v) = \begin{cases} 1 & \text{if } \|u - v\| > r \\ 0 & \text{if } \|u - v\| \leq r \end{cases} \quad (35)$$

where $\|u - v\|$ denotes the distance between u and v , and $r > 0$ is a fixed distance, then the density becomes

$$f(\mathbf{x}) = \begin{cases} \alpha \beta^{n(\mathbf{x})} & \text{if } \|x_i - x_j\| > r \text{ for all } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

This is the density of the Poisson process of intensity β in W conditioned on the event that no two points of the pattern lie closer than r units apart. It is known as the (classical) *hard core process*.



26.3.2 Strauss process

Generalising the hard core process, suppose we take $b(u) \equiv \beta$ and


$$c(u, v) = \begin{cases} 1 & \text{if } \|u - v\| > r \\ \gamma & \text{if } \|u - v\| \leq r \end{cases} \quad (36)$$

where γ is a parameter. Then the density becomes

$$f(\mathbf{x}) = \alpha \beta^{n(\mathbf{x})} \gamma^{s(\mathbf{x})} \quad (37)$$

where $s(\mathbf{x})$ is the number of pairs of distinct points in \mathbf{x} that lie closer than r units apart.

The parameter γ controls the ‘strength’ of interaction between points. If $\gamma = 1$ the model reduces to a Poisson process with intensity β . If $\gamma = 0$ the model is a hard core process. For values $0 < \gamma < 1$, the process exhibits inhibition (negative association) between points.



For $\gamma > 1$, the density (37) is not integrable. Hence the Strauss process is defined only for $0 \leq \gamma \leq 1$ and is a model for inhibition between points. This is typical of most Gibbs models.

26.3.3 Other pairwise interaction models

Other pairwise interactions that are considered in **spatstat** include the *Strauss-hard core* interaction (with hard core distance $h > 0$ and interaction distance $r > h$)

$$c(u, v) = \begin{cases} 0 & \text{if } \|u - v\| \leq h \\ \gamma & \text{if } h < \|u - v\| \leq r \\ 1 & \text{if } \|u - v\| > r \end{cases},$$

the *soft-core* interaction (with scale $\sigma > 0$ and index $0 < \kappa < 1$)

$$c(u, v) = \left(\frac{\sigma}{\|u - v\|} \right)^{2/\kappa},$$


the *Diggle-Gates-Stibbard* interaction (with interaction range ρ)

$$c(u, v) = \begin{cases} \sin \left(\frac{\pi \|u - v\|}{2\rho} \right)^2 & \text{if } \|u - v\| \leq \rho \\ 1 & \text{if } \|u - v\| > \rho \end{cases},$$


the *Diggle-Gratton* interaction (with hard core distance δ , interaction distance ρ and index κ)

$$c(u, v) = \begin{cases} 0 & \text{if } \|u - v\| \leq \delta \\ \left(\frac{\|u - v\| - \delta}{\rho - \delta} \right)^\kappa & \text{if } \delta < \|u - v\| \leq \rho \\ 1 & \text{if } \|u - v\| > \rho \end{cases},$$

and the general *piecewise constant* interaction in which $c(\|u - v\|)$ is a step function of $\|u - v\|$.



Piecewise constant interaction



26.4 Higher-order interactions

There are some useful Gibbs point process models which exhibit interactions of higher order, that is, in which the probability density has contributions from m -tuples of points, where $m > 2$.


One example is the *area-interaction* or Widom-Rowlinson process [17] with probability density

$$f(\mathbf{x}) = \alpha \beta^{n(\mathbf{x})} \gamma^{-A(\mathbf{x})} \quad (38)$$

where α is the normalising constant, $\beta > 0$ is an intensity parameter, and $\gamma > 0$ is an interaction parameter. Here $A(\mathbf{x})$ denotes the area of the region obtained by drawing a disc of radius r centred at each point x_i , and taking the union of these discs. The value $\gamma = 1$ again corresponds to a Poisson process, while $\gamma < 1$ produces a regular process and $\gamma > 1$ a clustered process. This process has interactions of all orders. It can be used as a model for moderate regularity or clustering.

26.5 Conditional intensity

The main tool for analysing a Gibbs point process is its *conditional intensity* $\lambda(u, \mathbf{X})$. Intuitively this determines the conditional probability of finding a point of the process at the location u given complete information about the rest of the process. For formal definitions see [31]. Informally, the conditional probability of finding a point of the process inside an infinitesimal neighbourhood of the location u , given the complete point pattern at all other locations, is $\lambda(u, \mathbf{X}) du$.



For point processes in a bounded window, the conditional intensity at a location u given the configuration \mathbf{x} is related to the probability density f by

$$\lambda(u, \mathbf{x}) = \frac{f(\mathbf{x} \cup \{u\})}{f(\mathbf{x})} \quad (39)$$

(for $u \notin \mathbf{x}$), the ratio of the probability densities for the configuration \mathbf{x} with and without the point u added.

The homogeneous Poisson process with intensity λ has conditional intensity

$$\lambda(u, \mathbf{x}) = \lambda$$

while the inhomogeneous Poisson process with intensity function $\lambda(u)$ has conditional intensity

$$\lambda(u, \mathbf{x}) = \lambda(u)$$

. The conditional intensity for a Poisson process does not depend on the configuration \mathbf{x} , because the points of a Poisson process are independent.

For the general pairwise interaction process (34) the conditional intensity is

$$\lambda(u, \mathbf{x}) = b(u) \prod_{i=1}^{n(\mathbf{x})} c(u, x_i). \quad (40)$$

For the hard core process,


$$\lambda(u, \mathbf{x}) = \begin{cases} \beta & \text{if } \|u - x_i\| > r \text{ for all } i \\ 0 & \text{otherwise} \end{cases} \quad (41)$$

which has the nice interpretation that a point u is either ‘permitted’ or ‘not permitted’ depending on whether it satisfies the hard core requirement.

For the Strauss process

$$\lambda(u, \mathbf{x}) = \beta \gamma^{t(u, \mathbf{x})} \quad (42)$$

where $t(u, \mathbf{x}) = s(\mathbf{x} \cup \{u\}) - s(\mathbf{x})$ is the number of points of \mathbf{x} that lie within a distance r of the location u . For $\gamma < 1$, this has the interpretation that a random point is less likely to occur at the location u if there are many points in the neighbourhood.



For the area-interaction process,

$$\lambda(u, \mathbf{x}) = \beta \gamma^{-B(u, \mathbf{x})} \quad (43)$$

where $B(u, \mathbf{x}) = A(\mathbf{x} \cup \{u\}) - A(\mathbf{x})$ is the area of that part of the disc of radius r centred on u that is not covered by discs of radius r centred at the other points $x_i \in \mathbf{x}$. If the points represent trees or plants, we may imagine that each tree takes nutrients and water from the soil inside a circle of radius r . Then we may interpret $B(u, \mathbf{x})$ as the area of the ‘unclaimed zone’ where a new plant at location u would be able to draw nutrients and water without competition from other plants. For $\gamma < 1$ we can interpret (43) as saying that a random point is less likely to occur when the unclaimed area is small.

The conditional intensity of a point process determines the probability density, through (39). Hence we can use the conditional intensity to define a point process. The conditional intensity is the preferred modelling tool for Gibbs processes: it has a direct interpretation, and it is easier to handle than the probability density.

26.6 Simulating Gibbs models

Gibbs models can be simulated by Markov chain Monte Carlo algorithms. Indeed, MCMC algorithms were invented to simulate Gibbs processes [49, 55].

In brief, these algorithms simulate a Markov chain whose states are point patterns. The chain is designed so that its equilibrium distribution is the distribution of the point process we want to simulate. If the chain were run for an infinite time, the state would converge in distribution to the desired point process. In practice the chain is run for a long finite time. Further details are beyond the scope of this workshop; consult [51, 50] for more information.

Currently **spatstat** offers the function **rmh** which simulates Gibbs processes using the Metropolis-Hastings algorithm.

```
> rmh(model, start, control)
```

- **model** determines the point process model to be simulated (see **help(rmhmodel)**).
- **start** determines the initial state of the Markov chain (see **help(rmhstart)**).
- **control** specifies control parameters for running the Markov chain, such as the number of iteration steps (see **help(rmhcontrol)**).

In the simplest uses of `rmh`, the three arguments are lists of parameter values. To generate a simulated realisation of the Strauss process with parameters $\beta = 2, \gamma = 0.7, r = 0.7$ in a square of side 10,

```
> mo <- list(cif = "strauss", par = c(beta = 2, gamma = 0.2, r = 0.7),
+           w = square(10))
> X <- rmh(model = mo, start = list(n.start = 42), control = list(nrep = 1e+06))
```

The other arguments specify a random initial state of 42 points, and that the algorithm shall be run for a million iterations.

27 Fitting Gibbs models

27.1 Maximum pseudolikelihood

Maximum likelihood estimation is intractable for most point process models. At the very least it requires Monte Carlo simulation to evaluate the likelihood (or the score and the Fisher information).

A workable alternative, at least for investigative purposes, is to maximise the log *pseudolikelihood*

$$\log \text{PL}(\theta; \mathbf{x}) = \sum_i \log \lambda(x_i; \mathbf{x}) - \int_W \lambda(u, \mathbf{x}) \, du. \quad (44)$$

You may recognise this as being very similar to the likelihood (4) of the Poisson process. In general it is not a likelihood, but the analogue of the score equation

$$\frac{\partial}{\partial \theta} \log \text{PL}(\theta) = 0$$

is an unbiased estimating equation. Thus the maximum pseudolikelihood estimator is asymptotically unbiased, consistent and asymptotically normal under appropriate conditions.

The main advantage of maximum pseudolikelihood is that, at least for popular Gibbs models, the conditional intensity $\lambda(u, \mathbf{x})$ is easily computable, so that the pseudolikelihood is easy to compute and to maximise. The main disadvantage is the bias and inefficiency of maximum pseudolikelihood in small samples.

More computationally-intensive estimation procedures typically use the maximum pseudolikelihood estimate as their initial guess. We are implementing such procedures in `spatstat` as well.

27.2 Fitting Gibbs models in spatstat

We have already met the function `ppm` for fitting Poisson point process models. In fact this function will fit a wide class of Gibbs models.

`ppm` contains an implementation of the algorithm of Baddeley and Turner [9] for maximum pseudolikelihood (which extends the Berman-Turner device for Poisson processes to a general Gibbs process). The conditional intensity of the model, $\lambda_\theta(u, \mathbf{x})$, must be loglinear in the parameters θ :

$$\log \lambda_\theta(u, \mathbf{x}) = \theta \cdot S(u, \mathbf{x}), \quad (45)$$

generalising (5), where $S(u, \mathbf{x})$ is a real-valued or vector-valued function of location u and configuration \mathbf{x} . Parameters θ appearing in the loglinear form (45) are called ‘regular’ parameters, and

all other parameters are ‘irregular’ parameters. For example, the Strauss process conditional intensity (42) can be recast as

$$\log \lambda(u, \mathbf{x}) = \log \beta + (\log \gamma) t(u, \mathbf{x})$$

so that $\theta = (\log \beta, \log \gamma)$ are regular parameters, but the interaction distance r is an irregular parameter (technically called a ‘bloody nuisance parameter’).

In **spatstat** we split the conditional intensity into first-order and higher-order terms:

$$\log \lambda_\theta(u, \mathbf{x}) = \eta \cdot S(u) + \varphi \cdot V(u, \mathbf{x}). \quad (46)$$

The ‘first order term’ $S(u)$ describes spatial inhomogeneity and/or covariate effects. The ‘higher order term’ $V(u, \mathbf{x})$ describes interpoint interaction.

The model with conditional intensity (46) is fitted by calling **ppm** in the form

```
ppm(X, ~ terms, V)
```

The first argument **X** is the point pattern dataset. The second argument **~terms** is a model formula, specifying the first order term $S(u)$ in (46), in the manner described in Section 15. Thus the first order term $S(u)$ in (46) may take very general forms.

The third argument **V** is an object of the special class **"interact"** which describes the interpoint interaction term $V(u, \mathbf{x})$ in (46). It may be compared to the ‘family’ argument which determines the distribution of the responses in a linear model or generalised linear model. Only a limited number of canned interactions are available in **spatstat**, because they must be constructed carefully to ensure that the point process exists.

To fit the Strauss process to the cells data using **ppm**,

```
> data(cells)
> ppm(cells, ~1, Strauss(r = 0.1))
```

Stationary Strauss process

First order term:

```
beta
762.6005
```

Interaction: Strauss process

```
interaction distance:      0.1
Fitted interaction parameter gamma:      0.008
```

Relevant coefficients:

```
Interaction
-4.825006
```

Here **Strauss** is a special function that creates an ‘interaction’ object (class **"interact"**) describing the interaction structure of the Strauss process. Notice that we had to specify the value of the irregular parameter r (more about that later).

To fit the inhomogeneous Strauss process with conditional intensity

$$\lambda(u, \mathbf{x}) = b(u) \gamma^{t(u, \mathbf{x})}$$

where, say, $b(u)$ is loglinear in the Cartesian coordinates,

$$\log b((x, y)) = \beta_0 + \beta_1 x + \beta_2 y$$

we simply type


```
> ppm(cells, ~x + y, Strauss(r = 0.1))
```

Nonstationary Strauss process

Trend formula: $\sim x + y$

Fitted coefficients for trend formula:

(Intercept)	x	y
6.2922384	0.5269869	0.1576416

Interaction: Strauss process

interaction distance: 0.1

Fitted interaction parameter gamma: 0.0082

Relevant coefficients:

Interaction
-4.805565

To fit an inhomogeneous Strauss process with log-quadratic first order term,

```
> ppm(cells, ~polynom(x, y, 2), Strauss(r = 0.1))
```

Nonstationary Strauss process

Trend formula: $\sim \text{polynom}(x, y, 2)$

Fitted coefficients for trend formula:

(Intercept)	polynom(x, y, 2)[x]	polynom(x, y, 2)[y]
5.9747220	-0.9375707	3.4732733
polynom(x, y, 2)[x^2]	polynom(x, y, 2)[x.y]	polynom(x, y, 2)[y^2]
1.4970947	-0.1838987	-3.3696109

Interaction: Strauss process

interaction distance: 0.1

Fitted interaction parameter gamma: 0.0081

Relevant coefficients:

Interaction
-4.812711

27.3 Interpoint interactions

Instead of `Strauss` we may use any of the following functions to create an interaction:

<code>AreaInter()</code>	area-interaction process
<code>BadGey()</code>	hybrid Geyer saturation process
<code>DiggleGratton()</code>	Diggle-Gratton potential
<code>DiggleGatesStibbard()</code>	Diggle-Gates-Stibbard potential
<code>Fiksel()</code>	Fiksel double exponential potential
<code>Geyer()</code>	Geyer's saturation process
<code>Hardcore()</code>	hard core process
<code>LennardJones()</code>	Lennard-Jones potential
<code>Ord()</code>	Ord model, user-supplied potential
<code>OrdThresh()</code>	Ord process, threshold potential
<code>PairPiece()</code>	pairwise interaction, piecewise constant
<code>Pairwise()</code>	pairwise interaction, user-supplied potential
<code>Poisson()</code>	the Poisson point process (the default)
<code>Saturated()</code>	general saturated model, user-supplied potential
<code>SatPiece()</code>	multiscale saturation process
<code>Softcore()</code>	pairwise interaction, soft core potential
<code>Strauss()</code>	the Strauss process
<code>StraussHard()</code>	the Strauss/hard core point process

(There are three additional ones for multitype point processes, described in section 34.3.2.)

The area-interaction model and the Geyer saturation model are quite handy, as they can be used to model both clustering and regularity.

```
> data(redwood)
> ppm(redwood, ~1, Geyer(r = 0.07, sat = 2))
```

Stationary Geyer saturation process

First order term:

```
beta
12.39488
```

Interaction: Geyer saturation process

```
interaction distance:      0.07
saturation parameter:      2
Fitted interaction parameter gamma:      2.9004
```

Relevant coefficients:

```
Interaction
1.064845
```

```
> ppm(redwood, ~1, AreaInter(r = 0.03))
```

Stationary Area-interaction process

First order term:

```
beta
```

36.53100

Interaction: Area-interaction process
 disc radius: 0.03
 Fitted interaction parameter eta: 15.7515

Relevant coefficients:
 Interaction
 2.756935

The printout for the area-interaction model uses the “scale-free” parameter `eta` defined by

$$\eta = \gamma^{\pi r^2}$$

where γ and r are the parameters appearing in the definition (43). Values of η greater than 1 suggest clustering.

For more detailed explanation of modelling, see [11].

27.4 Fitted point process models

The result of the `ppm` call is an object of class “`ppm`” (‘point process model’). This is very closely analogous to a fitted linear model (`lm`) or fitted generalised linear model (`glm`).

Standard R operations that are defined for fitted point process models (i.e. that have methods for the class “`ppm`”) include:

<code>print</code>	print basic information
<code>summary</code>	print detailed summary information
<code>plot</code>	plot the fitted (conditional) intensity
<code>predict</code>	fitted (conditional) intensity
<code>fitted</code>	fitted (conditional) intensity at data points
<code>update</code>	re-fit the model
<code>coef</code>	extract the fitted coefficient vector $\hat{\theta}$
<code>vcov</code>	variance-covariance matrix of $\hat{\theta}$
<code>anova</code>	analysis of deviance
<code>logLik</code>	evaluate log-pseudolikelihood
<code>model.matrix</code>	extract design matrix
<code>formula</code>	extract trend formula of model
<code>terms</code>	extract terms in model formula

(the methods for `anova` and `vcov` are only available for Poisson models). The following functions are also available:

<code>step</code>	stepwise model selection
<code>drop1</code>	one step backward in model selection
<code>model.images</code>	compute images of canonical covariates in model
<code>effectfun</code>	fitted intensity as function of one covariate

Plotting a fitted model generates a series of image and contour plots of


- the fitted first order term $\exp(\hat{\eta} \cdot S(u))$
- the fitted conditional intensity $\lambda_{\hat{\theta}}(u, \mathbf{x})$ evaluated for the data pattern \mathbf{x}

For Poisson models, the two plots are equivalent, and give the fitted intensity function.

```

> fit <- ppm(cells, ~polynom(x, y, 2), Strauss(r = 0.1))
> par(mfrow = c(1, 2))
> plot(fit, how = "image", ngrid = 256)

```




For non-Poisson models, it is also possible to extract and plot the interpoint interaction function, using `fitin`.

```

> model <- ppm(X, ~1, PairPiece(seq(10, 100, by = 10)))
> f <- fitin(model)
> plot(f)


```



27.5 Simulation from fitted models


A fitted Gibbs model can also be simulated automatically using `rmh`.

```
> fit <- ppm(swedishpines, ~1, Strauss(r = 7))
> Xsim <- rmh(fit)
> plot(Xsim, main = "Simulation from fitted Strauss model")
```



The `envelope` command will also generate simulation envelopes for a fitted model.

```
> plot(envelope(fit, nsim = 39))
```



27.6 Dealing with nuisance parameters

Irregular parameters, such as the interaction radius r in the Strauss process, cannot be estimated directly using `ppm`. Indeed the statistical theory for estimating such parameters is unclear.

For some special cases, a maximum likelihood estimator of the nuisance parameter is available. For example, for the ‘hard core process’ (Strauss process with interaction parameter $\gamma = 0$) with interaction radius r , the maximum likelihood estimator is the minimum nearest-neighbour distance. Thus the following is a reasonable approach to the `cells` dataset:

```
> rhat <- min(nndist(cells))
> rhat <- rhat * 0.99999
> ppm(cells, ~1, Strauss(r = rhat))
```

Stationary Strauss process

First order term:

beta
301.0949

Interaction: Strauss process

interaction distance: 0.0836293018068393

Fitted interaction parameter gamma: 0

Relevant coefficients:

Interaction
-20.77031

The analogue of profile likelihood, *profile pseudolikelihood*, provides a general solution which may or may not perform well. If $\theta = (\phi, \eta)$ where ϕ denotes the nuisance parameters and η the regular parameters, define the profile log pseudolikelihood by

$$\text{PPL}(\phi, \mathbf{x}) = \max_{\eta} \log \text{PL}((\phi, \eta); \mathbf{x}).$$

The right hand side can be computed, for each fixed value of ϕ , by the algorithm `ppm`. Then we just have to maximise $\text{PPL}(\phi)$ over ϕ . This is done by the command `profilepl`:

```
> data(simdat)
> df <- data.frame(r = seq(0.05, 2, by = 0.025))
> pfit <- profilepl(df, Strauss, simdat, ~1)
```

```
> pfit
```

Profile log pseudolikelihood values

for model: ppm(simdat, ~1, interaction = Strauss)

fitted with rbord= 2


Interaction: Strauss

with irregular parameter r in [0.05, 2]

Optimum value of irregular parameter: r = 0.275

The result is an object of class `profilepl` containing the profile log pseudolikelihood function, the optimised value of the irregular parameter r , and the final fitted model. To plot the profile log pseudolikelihood,

```
> plot(pfit)
```



To extract the final fitted model,

```
> pfit$fit
```

Stationary Strauss process

First order term:

```
beta
2.583110
```

Interaction: Strauss process

```
interaction distance:      0.275
Fitted interaction parameter gamma:      0.5631
```

Relevant coefficients:

```
Interaction
-0.5743608
```

There is a `summary` method for these objects as well.

27.7 Improvements over maximum pseudolikelihood

Maximum pseudolikelihood is quick and dirty. There are statistically more efficient alternatives, but they are computationally intensive.

Currently we have implemented the easiest of these alternatives, the Huang-Ogata [43] one-step approximation to maximum likelihood. Starting from the maximum pseudolikelihood estimate $\hat{\theta}_{PL}$, we simulate M independent realisations of the model with parameters $\hat{\theta}_{PL}$, evaluate the canonical sufficient statistics, and use them to form estimates of the score and Fisher information at $\theta = \hat{\theta}_{PL}$. Then we take one Newton-Raphson step, updating the value of θ . The rationale is that the log-likelihood is approximately quadratic in a neighbourhood of the maximum pseudolikelihood estimator, so that one Newton-Raphson step is almost enough.

To use the Huang-Ogata method instead of maximum pseudolikelihood, add the argument `method="ho"`.

```
> fit <- ppm(simdat, ~1, Strauss(r = 0.275), method = "ho")
```

```
> fit
```

Stationary Strauss process

First order term:

```
beta
2.42845
```

Interaction: Strauss process

interaction distance: 0.275

Fitted interaction parameter gamma: 0.5277

Relevant coefficients:

Interaction

-0.6392568

```
> vcov(fit)
```

```
          [,1]      [,2]
[1,] 0.01000399 -0.01255788
[2,] -0.01255788 0.04019740
```


For models fitted by Huang-Ogata, the variance-covariance matrix returned by `vcov` is computed from the simulations.

28 Validation of fitted Gibbs models

Goodness-of-fit testing and model validation for Poisson models were described in Section 16. Checking a fitted Gibbs point process model is more difficult. There is little theory available to support goodness-of-fit tests and the like.

As an example, consider the following data:

```
> data(residualspaper)
> X <- residualspaper$Fig4b
> plot(X)
```



We fit a Strauss process model with a log-quadratic intensity term:

```
> fit <- ppm(X, ~polynom(x, y, 2), Strauss(0.05), correction = "isotropic")
```


The question is how to confirm or validate this model.

28.1 Goodness-of-fit testing for Gibbs processes

For a fitted Gibbs process, no theory is available to support the χ^2 goodness-of-fit test or the Kolmogorov-Smirnov test. The predicted mean number of points in a given region is not known in closed form for a Gibbs process. Thus, the appropriate test statistic for a χ^2 test is not even available in closed form, let alone the null distribution of this statistic.


Instead, goodness-of-fit for fitted Gibbs models often relies on the summary functions K and G . The command `envelope` will accept as its first argument a fitted Gibbs model, and will simulate from this model to determine the critical envelope.

```
> plot(envelope(fit, Lest, nsim = 19, global = TRUE))
```



Let's subtract the theoretical Poisson value $L(r) = r$ to get a more readable plot:

```
> plot(envelope(fit, Lest, nsim = 19, global = TRUE), . - r ~ r)
```



This is fairly consistent with a Strauss process.

28.2 Residuals for Gibbs processes

28.2.1 Definition

Residuals for a general Gibbs model were defined only recently [12, 6]. The total residual in a region $B \subset \mathbb{R}^2$ is defined as

$$R(B) = n(\mathbf{x} \cap B) - \int_B \hat{\lambda}(u, \mathbf{x}) \, du \quad (47)$$

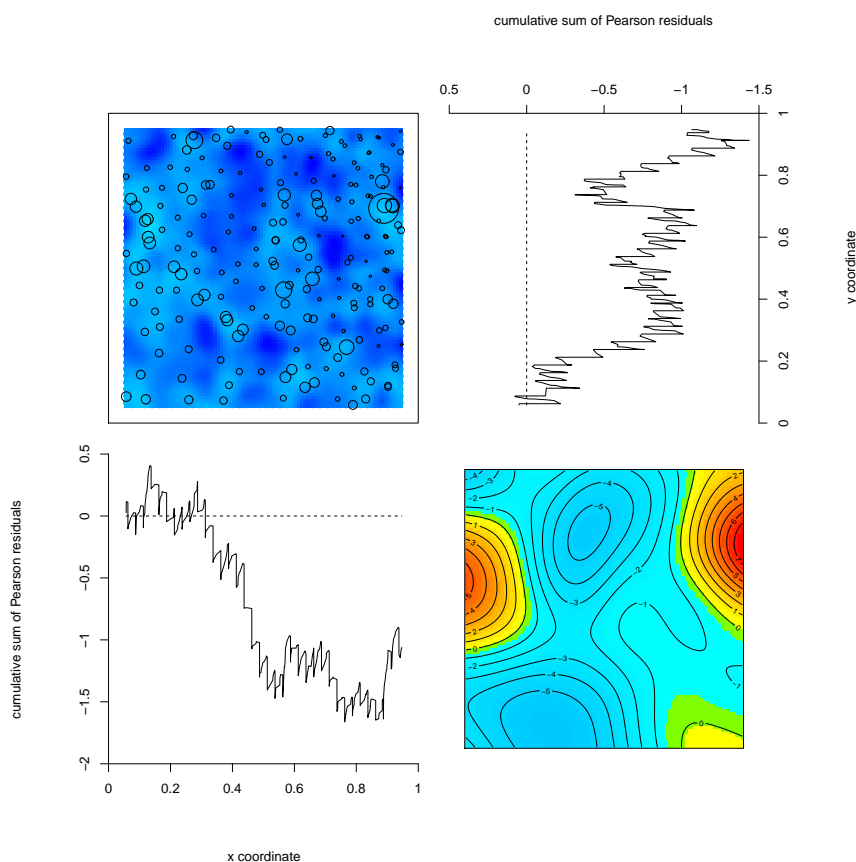
where again $n(\mathbf{x} \cap B)$ is the observed number of points in the region B , and $\hat{\lambda}(u, \mathbf{x})$ is the **conditional** intensity of the fitted model, *evaluated for the data point pattern* \mathbf{x} . If the fitted model is correct, the residuals have mean zero.

This definition is similar to the definition of residuals for Poisson processes (Section 16.2) except that the intensity $\hat{\lambda}(u)$ of the fitted Poisson process has been replaced by the *conditional* intensity $\hat{\lambda}(u, \mathbf{x})$ of the fitted Gibbs process evaluated for the data point pattern \mathbf{x} .

28.2.2 Residual plots

Residuals for Gibbs processes can be plotted using the same techniques as in Section 16.2. Here is the four-panel plot:

```
> diagnose.ppm(fit, type = "Pearson")
```



At the time of writing, `spatstat` does not yet display 2σ significance bands for the lurking variable plots when the fitted model is not Poisson. The interpretation of the lurking variable plots is a little more difficult without the significance bands. One tends to place a little more emphasis on the smoothed residual field. The Pearson residuals should be *approximately* standardised, so that values which are much greater than 2 (in absolute value) suggest a lack of fit.


The four-panel plot above suggests that the model is a reasonable fit.

28.2.3 Q–Q plots

As we noted in Section 16.2.6, the four-panel residual plot and the lurking variable plot are useful for detecting misspecification of the *trend* in a fitted model. They are not very useful for checking misspecification of the *interaction* in a fitted model.

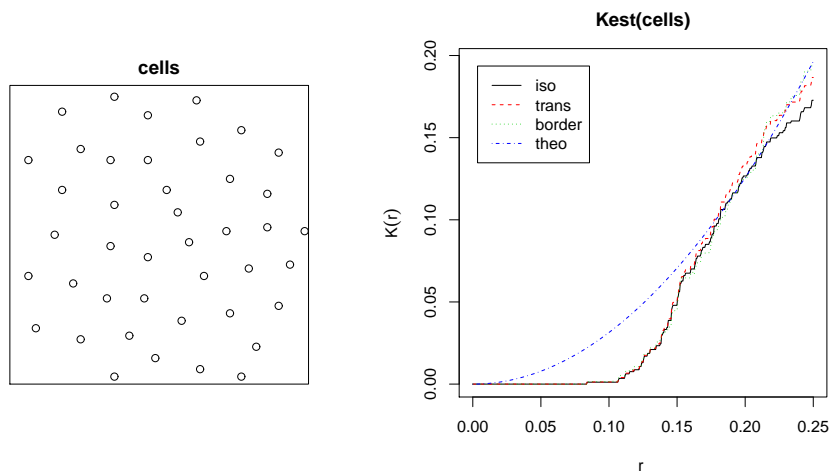
An extreme example is provided by the `cells` dataset. The residual plots for a uniform Poisson process fitted to the cells data suggest that this is a good model:

```
> data(cells)
> fitPois <- ppm(cells, ~1)
> diagnose.ppm(fitPois)
```




However, the K -function shows that the `cells` dataset is clearly not a Poisson pattern, but has strong inhibition:

```
> par(mfrow = c(1, 2))
> plot(cells)
> plot(Kest(cells))
> par(mfrow = c(1, 1))
```



Interaction between points in a point process corresponds roughly to the distribution of the responses in loglinear regression. To validate the interaction terms in a point process model, we should plot the distribution of the residuals. The appropriate tool is a *Q-Q plot*.


```
> qqplot.ppm(fitPois, nsim = 39)
```



This shows a Q-Q plot of the smoothed residuals for a uniform Poisson model fitted to the `cells` data, with pointwise 5% critical envelopes from simulations of the fitted model. This indicates that the uniform Poisson model is grossly inappropriate for the `cells` data.

Returning to the model we fitted at the start of this chapter:

```
> qqplot.ppm(fit, nsim = 39)
```



This shows a Q–Q plot of the smoothed residuals, with pointwise 5% critical envelopes from simulations of the fitted model. This suggests that the Strauss model is reasonable.

These validation techniques generalise and unify many existing exploratory methods. For particular models of interpoint interaction, the Q–Q plot is closely related to the summary functions F , G and K . See [12].

28.3 New methods

Several new diagnostic tools for validating the interaction term in a point process model will be published soon [8]. These methods will then be released in **spatstat**.

PART VII. MARKED POINT PATTERNS

Part VII of the workshop deals with marked point patterns.

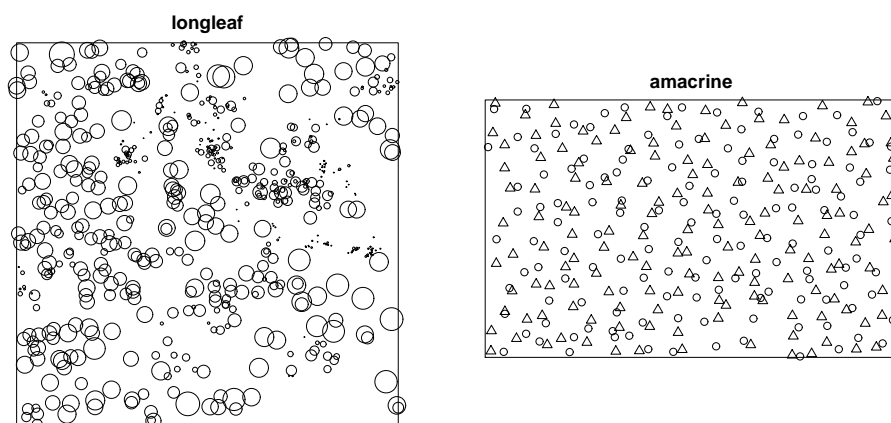
29 Marked point patterns

29.1 Marked point patterns

Each point in a spatial point pattern may carry additional information called a ‘mark’. For example, points which are classified into two or more different types (on/off, case/control, species, colour, etc) may be regarded as marked points, with a mark which identifies which type they are. Data recording the locations and heights of trees in a forest can be regarded as a marked point pattern where the mark attached to a tree’s location is the tree height.

Many of the functions in **spatstat** handle marked point patterns in which the mark attached to each point is either

- a **continuous variate** or “real number”. An example is the Longleaf Pines dataset (**longleaf**) in which each tree is marked with its diameter at breast height. The **marks** component must be a **numeric** vector such that **marks[i]** is the mark value associated with the *i*th point. We say the point pattern has *continuous marks*.
- a **categorical variate**. An example is the Amacrine Cells dataset (**amacrine**) in which each cell is identified as either “on” or “off”. Such point patterns may be regarded as consisting of points of different “types”. The **marks** component must be a **factor** such that **marks[i]** is the label or type of the *i*th point. We call this a *multitype point pattern* and the levels of the factor are the possible types.



Note that, in some other packages, a point pattern dataset consisting of points of two different types (A and B say) is represented by two datasets, one representing the points of type A and another containing the points of type B. In **spatstat** we take a different approach, in which all the points are collected together in one point pattern, and the points are then labelled by the type to which they belong. An advantage of this approach is that it is easy to deal with multitype point patterns with more than 2 types. For example the classic Lansing Woods dataset represents the positions of trees of 6 different species. This is available in **spatstat** as a single dataset, a marked point pattern, with the marks having 6 levels.

29.2 Formulation

A mark variable may be interpreted as an additional coordinate for the point: for example a point process of earthquake epicentre locations (longitude, latitude), with marks giving the occurrence time of each earthquake, can alternatively be viewed as a point process in space-time with coordinates (longitude, latitude, time).

A marked point process of points in space S with marks belonging to a set M is mathematically defined as a point process in the cartesian product $S \times M$. The space M of possible marks may be ‘anything’. In current applications, typically the mark is either a categorical variable (so that the points are grouped into ‘types’) or a real number. Multivariate marks consisting of several such variables are also common.

A marked point pattern is an unordered set

$$\mathbf{y} = \{(x_1, m_1), \dots, (x_n, m_n)\}, \quad x_i \in W, \quad m_i \in M$$

where x_i are the locations and m_i are the corresponding marks.

29.3 Methodological issues

29.3.1 Should the data be treated as a marked point process?

In a marked point process the points are random. Treating the data as a point process is inappropriate if the locations are fixed, or if the locations are not part of the ‘response’.

Example 16 *Today’s maximum temperatures at 25 Australian cities are displayed on a map.*

This is not a point process in any useful sense. The cities are fixed locations. The temperatures are observations of a spatial variable at a fixed set of locations. See the R packages `sp`, `spdep`, `spgwr` for suitable methods.

Example 17 *A mineral exploration dataset records the map coordinates where 15 core samples were drilled, and for each core sample, the assayed concentration of iron in the sample.*

This typically should *not* be treated as a point process. The core sample locations were chosen by a geologist, and are part of the experimental design. The main interest is in the iron concentration at these locations. This should probably be analysed as a geostatistical dataset. See the R packages `geoR`, `geoRglm` for suitable methods.

29.3.2 Joint vs. conditional analysis

There are more choices for analysis (and more traps) when marks are present. Schematically, if we write X for the points and M for the marks, then a statistical model for the marked point pattern could be formulated in several ways:

- $[X] [M|X]$ — ‘conditional on locations’ — points X are first generated according to a spatial point process, then marks M are ‘assigned’ to the points by a random mechanism $[M|X]$;
- $[M] [X|M]$ — ‘conditional on marks’ or ‘split by marks’ — marks M are first generated according to some random mechanism $[M]$, then they are placed at certain locations X by point process(es) $[X|M]$;
- $[X, M]$ — ‘joint’ — marked points are generated according to a marked point process.

These approaches typically lead to different stochastic models and have different inferential interpretations. Correspondingly, there are different null hypotheses that can be tested:

- *random labelling*: given the locations X , the marks are conditionally independent and identically distributed;

- *independence of components*: the sub-processes \mathbf{X}_m of points of each mark m , are independent point processes;
- *complete spatial randomness and independence (CSRI)*: the locations \mathbf{X} are a uniform Poisson point process, and the marks are independent and identically distributed. (This implies both random labelling and independence of components).

These null hypotheses are not equivalent.

The properties of random labelling and independence of components are not equivalent. For example, take a point process \mathbf{X} where nearest neighbour distances are always larger than a threshold r , and attach random marks to the points. The resulting marked point process cannot be generated using the independence construction, because if points with different marks are independent, they can come arbitrarily close to one another.

Example 18 (Ant nests data) *Two species of ants build nests in a desert. We want to investigate ecological interaction between the species, and between different nests of the same species. The locations of all nests are mapped, and marked by the species.*

These data can be analysed as a marked point process consisting of two different types of points. The ‘mark’ attached to each point is its species (a categorical variable). The most natural kind of modelling and analysis is either joint $[X, M]$ or split by species $[M] [X|M]$. We could also treat one of the species as a covariate and analyse the other species conditional on it.

Example 19 *Trees in an orchard are examined and their disease status (infected/not infected) is recorded. We are interested in the spatial characteristics of the disease, such as contagion between neighbouring trees.*

These data probably should *not* be treated as a point process. The response is ‘disease status’. We can think of disease status as a label applied to the trees after their locations have been determined. Since we are interested in the spatial correlation of disease status, the tree locations are effectively fixed covariate values. It would probably be best to treat these data as a discrete random field (of disease status values) observed at a finite known set of sites (the trees).

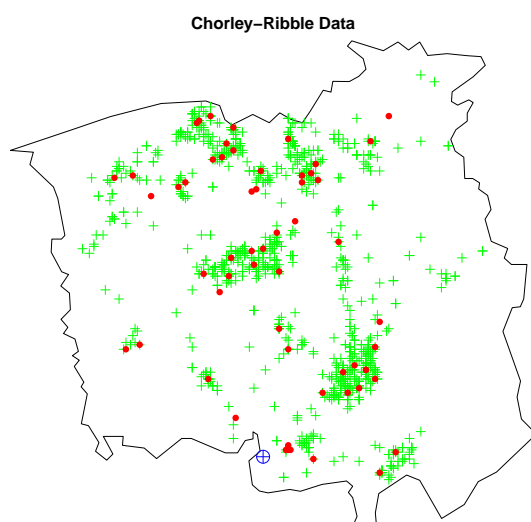
29.3.3 Grey areas

There are some ‘grey areas’ which permit several alternative choices of analysis. It could be appropriate either to analyse the locations and marks jointly (denoted $[X, M]$), or to analyse the marks conditional on the locations ($[M|X]$) or to analyse the locations given the marks ($[X|M]$).

One grey area occurs when the locations are random, but may be ancillary for the parameters of interest.

Example 20 *Case-control study of cancer [34, 38]. The domicile locations of all new cases of a rare cancer are mapped. To allow for spatial variation in the density of the susceptible population, domicile locations are recorded for a random sample of (matched) controls.*

This can be analysed either as a marked point pattern (where the mark is the case/control label) or, by conditioning on locations, as a random field of case/control values attached to the known domicile locations.



For further discussion of these issues, see [3].

30 Handling marked point pattern data

This section explains how to create a marked point pattern dataset in **spatstat**, and how to manipulate it.

30.1 Creating marked point pattern datasets

The marks attached to a point pattern may be stored in a vector (with one entry for each point) or in a matrix or data frame (with one row for each point and one column for each mark variable). The mark values can be of any atomic type: numeric, integer, character, factor, logical or complex.

A marked point pattern dataset can be created using any of the following tools:

ppp	create point pattern dataset
as.ppp	convert other data to point pattern
superimpose	combine several point patterns into a marked point pattern
marks	extract marks from a point pattern
marks<-	attach marks to a point pattern
%mark%	attach marks to a point pattern
unmark	delete marks from a point pattern
scanpp	read point pattern data from text file
clickppp	create a pattern using point-and-click on the screen

The command **ppp** can be used to create a marked point pattern dataset from raw data. The syntax is

```
> ppp(x, y, ..., marks = m)
```

where **x** and **y** are vectors of equal length containing the (x, y) coordinates, **m** is either a vector of the same length as **x** containing the mark values for each point, or a matrix or data frame with **nrow(m) = length(x)** containing the multivariate mark for each point, and **...** are arguments that determine the window for the point pattern.

Tip: If the marks are a vector and are intended to be a categorical variable (representing the types in a multitype point pattern),

- ensure that `m` is stored as a **factor** in R.
- when the point pattern `X` has been created, check that it is multitype using `is.multitype(X)`.
- check that the factor levels are as you intended, using `levels(m)` or `levels(marks(X))` where `X` is the marked point pattern. If the factor levels are character strings, they will be sorted into alphabetical order by default.
- be careful when performing equality/inequality comparisons involving a factor. Particular danger occurs when the factor levels are strings that represent integers.

The command `as.ppp` will convert data in another format (for example, a matrix or data frame) to a point pattern object of class "ppp". The third and subsequent columns of a matrix or data frame will be interpreted as containing the marks.


```
> mydata <- data.frame(x = runif(10), y = runif(10), m = sample(letters[1:3],
+ 10, replace = TRUE))
> as.ppp(mydata, square(1))
```

```
marked planar point pattern: 10 points
multitype, with levels = a      b      c
window: rectangle = [0, 1] x [0, 1] units
```

If point pattern data are stored in a text file, the command `scanpp` will read the data and create a point pattern object of class "ppp". The argument `multitype=TRUE` will ensure that the mark values are interpreted as a factor.

```
> X <- scanpp("myfile.txt", window = square(1), multitype = TRUE)
```

The command `superimpose` combines several point patterns within the same window. It can be used to create a multitype point pattern, if you have already created separate point patterns containing the points of each type. Suppose `X1` and `X2` are unmarked point patterns. Then `superimpose(A=X1, B=X2)` will create a multitype point pattern by attaching the mark `A` to each point of `X1`, attaching the mark `B` to each point of `X2`, and combining the points.



Marks can be attached to an existing point pattern `X` using the function `marks<-` as in

```
> marks(X) <- m
```

or using the binary operator `%mark%`,

```
> Y <- X %mark% m
```

These are convenient when you want to assign new marks to a dataset that are computed using another variable, or perhaps to randomise the marks in a dataset.

A multitype point pattern can also be created interactively using `clickppp`, using the argument `types` to specify the possible types.

30.2 Inspecting a marked point pattern

Basic tools for inspecting a marked point pattern include the `print`, `plot` and `summary` methods.

```
> data(amacrine)
> amacrine

marked planar point pattern: 294 points
multitype, with levels = off      on
window: rectangle = [0, 1.6012085] x [0, 1] units (one unit = 662 microns)


> summary(amacrine)

Marked planar point pattern: 294 points
Average intensity 184 points per square unit (one unit = 662 microns)
Multitype:
      frequency proportion intensity
off      142      0.483      88.7
on       152      0.517      94.9

Window: rectangle = [0, 1.6012085]x[0, 1]units
Window area = 1.60121 square units
Unit of length: 662 microns

> plot(amacrine)

off  on
 1   2
```



If the marks are a data frame (i.e. if there are several columns of marks), the first column of marks will be plotted by default. To change this, use the argument `which.marks` to specify another column.


```

> data(shapley)
> par(mfrow = c(1, 2))
> plot(shapley, which.marks = "V", maxsize = 0.5, main = "Recession velocity")
      -20000      0      20000      40000      60000      80000
-0.1303000  0.0000000  0.1303000  0.2605999  0.3908999  0.5211998

> plot(shapley, which.marks = "Mag", maxsize = 0.1, main = "Magnitude")
      0      5      10      15      20      25
0.0000000 0.0223914 0.0447828 0.0671742 0.0895656 0.1119570

> par(mfrow = c(1, 1))

```



You can also convert a marked point pattern into a data frame for closer inspection of the coordinates and mark values:

```

> as.data.frame(amacrine)
      x      y marks
1 0.0224 0.0243   on
2 0.0243 0.1028   on
3 0.1626 0.1477   on
.....

```

The marks can be extracted using the function `marks`:

```

> data(longleaf)
> m <- marks(longleaf)

```

Beware the possibility that two points with different marks may occupy the same spatial location. This is not currently detected by `ppp` since, for a marked point pattern, the function `duplicated.ppp` regards two points as identical only when their coordinates **and** mark values are identical. To detect duplication of the spatial locations, use `duplicated(unmark(X))`.

Further tools are presented in the next section.

30.3 Manipulating data

30.3.1 Manipulating marks

The following tools can manipulate the marks in a point pattern:

```

marks      extract marks
marks<-    attach marks to a point pattern
%mark%     attach marks to a point pattern
unmark     remove marks from point pattern

```

For example, the Lansing Woods data are tree locations marked by diameter at breast height (dbh) in centimetres. To convert the marks from diameters to circular areas,

```
> data(lansing)
> d <- marks(lansing)
> a <- (pi/4) * d^2
> marks(lansing) <- a
```


30.3.2 Separating points of different types

A *multitype* point pattern can be separated into the sub-patterns of points of each type, using the `split` command.

```
> data(amacrine)
> Y <- split(amacrine)
```

In fact `split` is a generic function and the commands above invoke the `split` method for the class of point patterns, `split.ppp`. The result `Y` is a list of point patterns, with names that correspond to the type labels. This list also belongs to the class "`splitppp`" which can be plotted automatically:

```
> plot(split(amacrine))
```



If the marks are a data frame, you will need to specify the splitting/grouping factor explicitly. For example the `nbfires` dataset records the location of forest fires, marked by 9 different variables. To split the fire locations by the cause of the fire,

```
> data(nbfires)
> Y <- split(nbfires, "cause")
```

30.3.3 Cutting the numerical scale into bands

For a point pattern with *numeric* marks, the marks can be converted to a factor, using a method for the generic function `cut`. The user specifies a series of cut-points on the numerical scale; all mark values between two cut-points are given the same label.

For example, the Longleaf Pines data are the locations of trees marked with their diameter at breast height, `dbh`, in centimetres. By convention we define “adult” trees to be those with `dbh` greater than 30 centimetres. To obtain the bivariate point pattern of adult and juvenile trees,

```
> data(longleaf)
> longleaf
```

```
marked planar point pattern: 584 points
marks are numeric, of type double
window: rectangle = [0, 200] x [0, 200] metres
```

```
> X <- cut(longleaf, breaks = c(0, 30, 80), labels = c("juvenile",
+ "adult"))
> X
```

```
marked planar point pattern: 584 points
multitype, with levels = juvenile      adult
window: rectangle = [0, 200] x [0, 200] metres
```

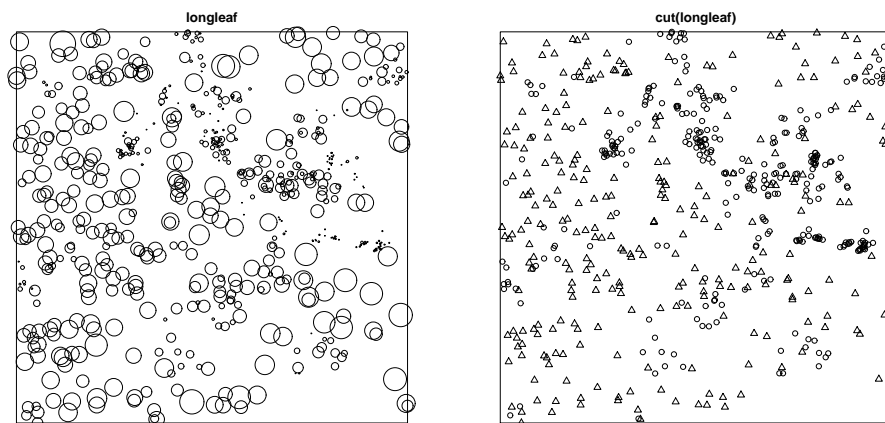
```
> par(mfrow = c(1, 2))
> plot(longleaf)
```

```
      0      20      40      60      80
0.000000 1.722522 3.445045 5.167567 6.890090
```

```
> plot(X, main = "cut(longleaf)")
```

```
juvenile  adult
      1      2
```

```
> par(mfrow = c(1, 1))
```



If the marks are a data frame, use the second argument `z` to specify which column should be used for the cut. For example, to classify New Brunswick fires into three groups by fire size,

```
> data(nbfires)
> Y <- cut(nbfires, "fnl.size", breaks = 4)
```

31 Exploratory tools for multitype point patterns

This section covers some tools for exploratory data analysis of multitype point patterns (i.e. where the marks are categorical).

31.1 Intensity

The Lansing Woods data give the locations of 6 species of trees in a forest in Michigan. Elementary estimates of the frequency distribution of species, and the intensity of each species, are available from `summary.ppp`.

```
> data(lansing)
> summary(lansing)
```

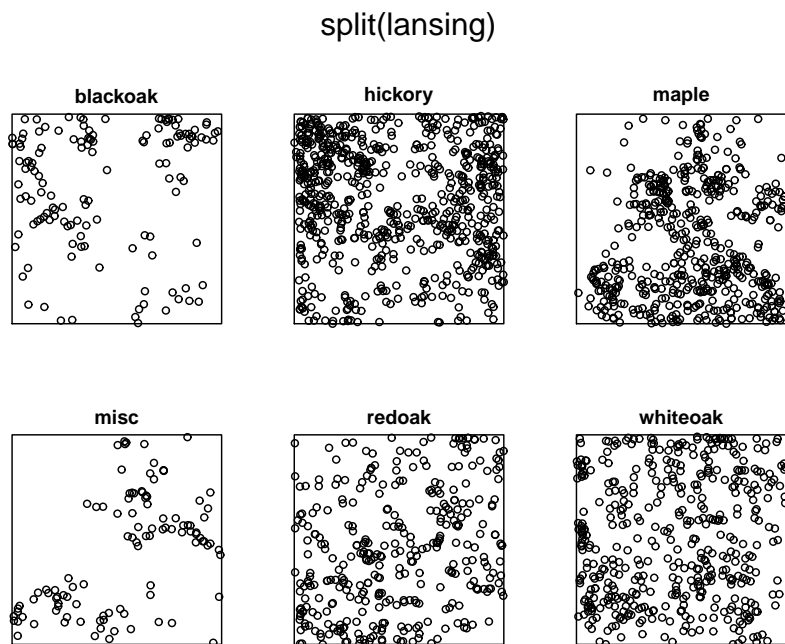
```
Marked planar point pattern: 2251 points
Average intensity 2250 points per square unit (one unit = 924 feet)
```

```
*Pattern contains duplicated points*
Multitype:
      frequency proportion intensity
blackoak      135      0.0600      135
hickory       703      0.3120      703
maple         514      0.2280      514
misc          105      0.0466      105
redoak        346      0.1540      346
whiteoak      448      0.1990      448
```

```
Window: rectangle = [0, 1]x[0, 1]units
Window area = 1 square unit
Unit of length: 924 feet
```


It's sensible to examine the sub-patterns of different types separately, using `split.ppp`.

```
> plot(split(lansing))
```

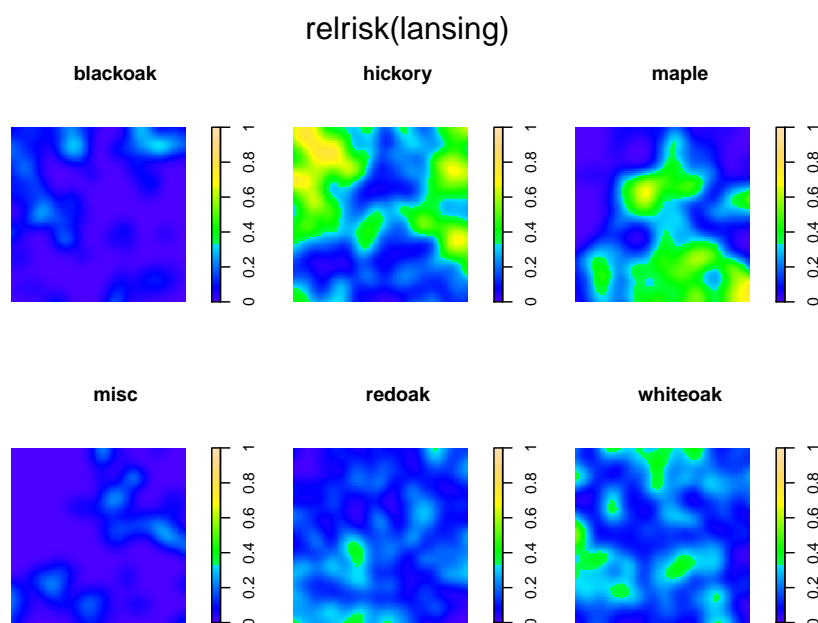
It would be useful to compute and plot a separate estimate of intensity for each type of tree. This is possible using the functions `density.splitppp` and `plot.listof`. They are invoked simply by typing

```
> plot(density(split(lansing)), ribbon = FALSE)
```



The relative proportions of intensity could then be computed by taking ratios of these densities, using `eval.im`. This is done more neatly using the command `relrisk` (which will also select the bandwidth automatically by cross-validation if it is not specified).

```
> plot(relrisk(lansing), zlim = c(0, 1))
```



In the context of the Lansing Woods data, these tools are searching for evidence of *segregation* between the different species of trees. Segregation occurs when the intensities of different species are negatively associated.

Parametric estimates of intensity can be obtained using `ppm`, fitting a Poisson model with an intensity function that may depend on location and/or on the marks. See below.

Evidence for segregation in a multitype point pattern can also be assessed using the Kolmogorov-Smirnov test `kstest` and other specialised tests.

31.2 Simple summaries of neighbouring marks

We are often interested in the marks that are attached to the close neighbours of a typical point.

For a multitype point pattern, the function `marktable` compiles a contingency table of the marks of all points within a given radius of each data point:

```
> data(amacrine)
> M <- marktable(amacrine, R = 0.1)
> M[1:10, ]
```

	mark	
point	off	on
1	1	1
2	2	2
3	4	3
4	3	1
5	4	1
6	2	3
7	3	2
8	1	1
9	3	1
10	3	2

More general summaries of the marks of neighbours can be obtained using the function `markstat`. For example, to compute the average diameter of the 5 closest neighbours of each tree in the Longleaf Pines dataset,

```
> md <- markstat(longleaf, mean, N = 5)
> md[1:10]

[1] 43.40 43.40 48.58 21.70 48.38 53.32 40.28 29.82 24.92 21.70
```

31.3 Distance methods and summary functions

If \mathbf{X} and \mathbf{Y} are two point pattern objects, then

- `crossdist(X,Y)` computes the matrix of distances from each point of \mathbf{X} to each point of \mathbf{Y} ;
- `nncross(X,Y)` finds, for each point in \mathbf{X} , the nearest point of \mathbf{Y} and the distance to this point.

The summary functions F , G , J and K (and other functions derived from K , such as L and the pair correlation function) have been extended to multitype point patterns, using such distances.

31.3.1 A pair of types

Assume the multitype point process \mathbf{X} is stationary. Let \mathbf{X}_j denote the sub-pattern of points of type j , with intensity λ_j . Then for any pair of types i and j ,

- $F_j(r)$ is the empty space function for \mathbf{X}_j .
- $G_{ij}(r)$ is the distribution function of the distance from a point of type i to the nearest point of type j
- $K_{ij}(r)$ is $1/\lambda_j$ times the expected number of points of type j within a distance r of a typical point of type i .
- $L_{ij}(r)$ is the corresponding L -function

$$L_{ij}(r) = \sqrt{\frac{K_{ij}(r)}{\pi}}.$$

- $g_{ij}(r)$ is the corresponding analogue of the pair correlation function

$$g_{ij}(r) = \frac{K'_{ij}(r)}{2\pi r}$$

where $K'_{ij}(r)$ is the derivative of K_{ij} .


- J_{ij} is defined as

$$J_{ij}(r) = \frac{1 - G_{ij}(r)}{1 - F_j(r)}.$$

The functions G_{ij} , K_{ij} , L_{ij} , g_{ij} , J_{ij} are called “cross-type” or “ i -to- j ” summary functions. They are computed in `spatstat` by `Gcross`, `Kcross`, `Lcross`, `pcfcross` and `Jcross` respectively.

```
> data(amacrine)
> amacrine

> plot(Gcross(amacrine, "on", "off"))
```



The interpretation of the cross-type summary functions is similar, *but not identical*, to that of the original functions F , G , K etc:

- if \mathbf{X}_j is a uniform Poisson process (CSR), then $F_j(r) = 1 - \exp(-\lambda_j \pi r^2)$.
- if \mathbf{X}_j is a uniform Poisson process (CSR) *and is independent of* \mathbf{X}_i , then $G_{ij}(r) = 1 - \exp(-\lambda_j \pi r^2)$.
- if \mathbf{X}_i and \mathbf{X}_j are independent, then $K_{ij}(r) = \pi r^2$, $L_{ij}(r) = r$, $g_{ij}(r) = 1$, $G_{ij}(r) = F_{ij}(r)$ and $J_{ij}(r) = 1$.

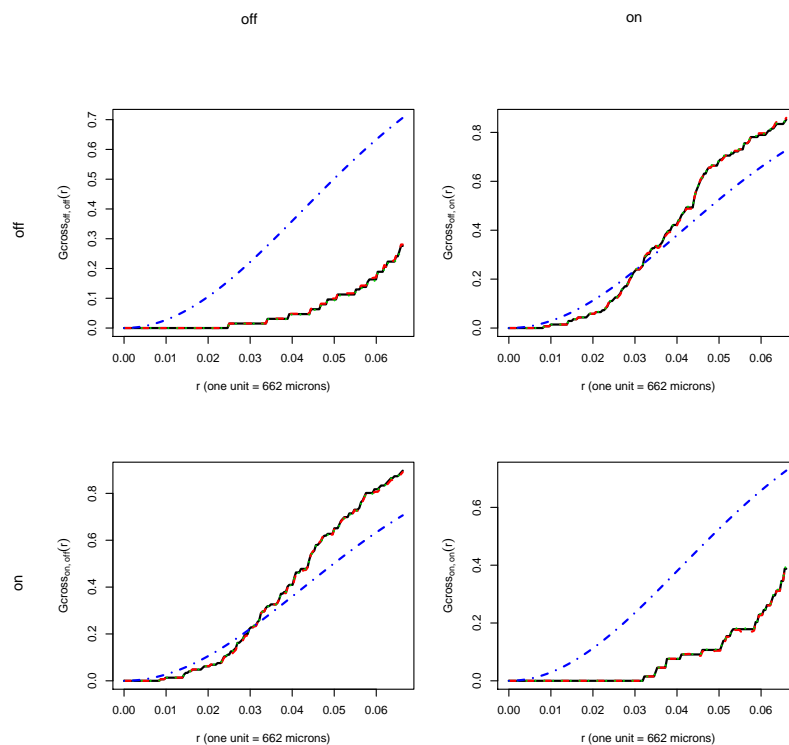
Here ‘independent’ means that the two point processes are probabilistically independent.

31.3.2 All pairs of types

The command `alltypes` enables the user to compute the cross-type summary functions between all pairs of types simultaneously. For example, to compute $G_{ij}(r)$ for all i and j in the amacrine cells data, we would use `alltypes(amacrine, "G")`. The result is automatically displayed as an array of plot panels.

```
> plot(alltypes(amacrine, "G"))
```

array of G functions for amacrine.



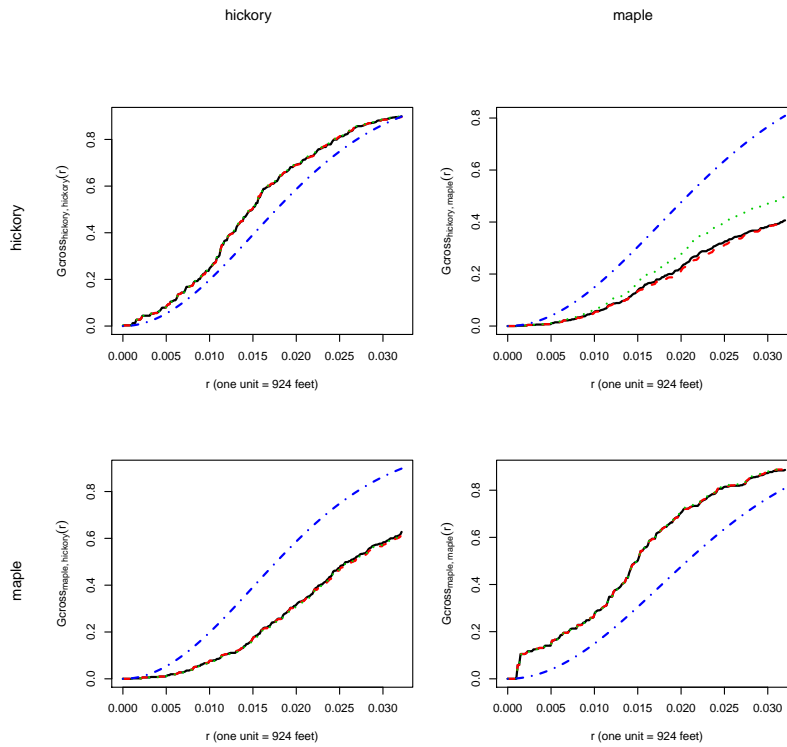
For example the top right panel shows the cumulative distribution function of the distance from an “off” cell to the nearest “on” cell.

The result of `alltypes` is a ‘function array’ (object of class “`fasp`”) which can be indexed by row and column subscripts. If the point pattern has a large number of possible types, you can compute the array of all possible pairwise G functions, then use the subscript operator to inspect a subset of the array.

```
> data(lansing)
> a <- alltypes(lansing, "G")
```

```
> plot(a[2:3, 2:3])
```

array of G functions for lansing.



31.3.3 One type to any type

Also defined are the “*i*-to-any” summaries

- $G_{i\bullet}(r)$, the distribution function of the distance from a point of type i to the nearest other point of any type;
- $K_{i\bullet}(r)$ is $1/\lambda$ times the expected number of points of any type within a distance r of a typical point of type i . Here $\lambda = \sum_j \lambda_j$ is the intensity of the entire process \mathbf{X} .
- $L_{i\bullet}(r)$ is the corresponding L -function

$$L_{i\bullet}(r) = \sqrt{\frac{K_{i\bullet}(r)}{\pi}}.$$

- $g_{i\bullet}(r)$ is the corresponding analogue of the pair correlation function


$$g_{i\bullet}(r) = \frac{K'_{i\bullet}(r)}{2\pi r}.$$

- $J_{i\bullet}$ is defined by

$$J_{i\bullet}(r) = \frac{1 - G_{i\bullet}}{1 - F(r)}$$



These are computed by `Gdot`, `Kdot`, `Ldot`, `pcfdot` and `Jdot` respectively, or using `alltypes`.

```
> plot(Gdot(amacrine, "on"))
```



```
> plot(alltypes(amacrine, "Gdot"))
```

array of Gdot functions for amacrine.




31.3.4 Plotting and manipulating function arrays

A function array (object of class "fasp") can be printed and plotted using methods for this class. It can also be manipulated in various ways.

The `plot` method is similar to `plot.fv` and allows the function values to be transformed:

```
> aG <- alltypes(amacrine, "G")
> fisher <- function(x) asin(sqrt(x))
> plot(aG, fisher(.) ~ fisher(theo))
```

array of G functions for amacrine.



As mentioned above, the function array can be indexed by array subscripts.

```
> data(lansing)
> a <- alltypes(lansing, "G")
> dim(a)
> b <- a[2:3, 2:3]
```

Calculations can be performed on all the functions in the array using `eval.fasp`.

```
> aGfish <- eval.fasp(asin(sqrt(aG)))
```

31.3.5 Mark connection function

The *mark connection function* between types i and j in a stationary multitype point process is

$$p_{ij}(r) = \frac{\lambda_i \lambda_j g_{ij}(r)}{\lambda^2 g(r)}.$$

This can be interpreted as the conditional probability, given that there is a point of the process at a location u and another point of the process at a location v separated by a distance $\|u - v\| = r$, that the first point is of type i and the second point is of type j .

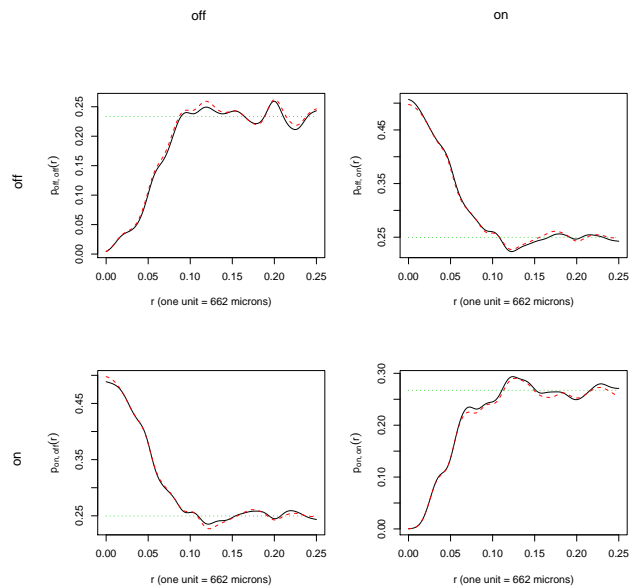
The command `markconnect` computes estimates of the mark connection function.

```
> data(amacrine)
> markconnect(amacrine, "on", "off")
```

We can use `alltypes` to compute the mark connection function p_{ij} for all pairs of types i and j :

```
> plot(alltypes(amacrine, markconnect))
```


array of markconnect functions for amacrine.



31.3.6 Mark equality function


The composite

$$p(r) = \sum_i p_{ii}(r)$$

can be interpreted as the conditional probability, given that there is a point of the process at a location u and another point of the process at a location v separated by a distance $\|u - v\| = r$, that the two points have the *same* type.

This is sometimes called the *mark equality function*. It is a special case of a more general technique of “mark correlation” which we discuss in Section 32.2. To compute the mark equality function of a multitype point pattern, use `markcorr`.

```
> plot(markcorr(amacrine))
```



This plot indicates that nearby points tend to have different types.


31.4 Randomisation tests

Simulation envelopes of summary functions can be used to test various null hypotheses for marked point patterns.

31.4.1 Poisson null

The null hypothesis of a homogeneous Poisson marked point process can be tested by direct simulation, using `envelope` as before. For example, using the cross-type K function as the test statistic,

```
> data(amacrine)
> E <- envelope(amacrine, Kcross, nsim = 39, i = "on", j = "off")
> plot(E, main = "test of marked Poisson model")
```



Notice that the arguments `i` and `j` here do not match any of the formal arguments of `envelope`, so they are passed to `Kcross`. This has the effect of calling `Kcross(X, i="on", j="off")` for each of the simulated point patterns X . Each simulated pattern is generated by the homogeneous Poisson point process with intensities estimated from the dataset `amacrine`.

31.4.2 Independence of components


It's also possible to test other null hypotheses by a randomisation test. We discussed two popular null hypotheses:

- *random labelling*: given the locations X , the marks are conditionally independent and identically distributed;
- *independence of components*: the sub-processes \mathbf{X}_m of points of each mark m , are independent point processes.

In a randomisation test of the independence-of-components hypothesis, the simulated patterns X are generated from the dataset by splitting the data into sub-patterns of points of one type, and randomly shifting these sub-patterns, independently of each other. The shifting is performed by `rshift`:

```
> E <- envelope(amacrine, Kcross, nsim = 39, i = "on", j = "off",
+   simulate = expression(rshift(amacrine, radius = 0.25)))
```

```
> plot(E, main = "test of independent components")
```



The independence-of-components hypothesis seems to be accepted in this example.
Under the independence hypothesis,

$$\begin{aligned} K_{ij}(r) &= \pi r^2 \\ G_{ij}(r) &= F_j(r) \\ J_{ij}(r) &\equiv 1. \end{aligned}$$

while the “*i*-to-any” functions have complicated values. Thus, we would normally use K_{ij} or J_{ij} to construct a test statistic for independence of components.

31.4.3 Random labelling

In a randomisation test of the random labelling null hypothesis, the simulated patterns \mathbf{X} are generated from the dataset by holding the point locations fixed, and randomly resampling the marks, either with replacement (independent random sampling) or without replacement (randomly permuting the marks). The resampling operation is performed by `rlabel`.

Under random labelling,


$$\begin{aligned} J_{i\bullet}(r) &= J(r) \\ K_{i\bullet}(r) &= K(r) \\ G_{i\bullet}(r) &= G(r) \end{aligned}$$

(where G, K, J are the summary functions for the point process without marks) while the other, cross-type functions have complicated values. Thus, we would normally use something like $K_{i\bullet}(r) - K(r)$ to construct a test statistic for random labelling.

To do this, cook up a little function to evaluate $J_{i\bullet}(r) - J(r)$:

```
> Jdif <- function(X, ..., i) {
+   Jidot <- Jdot(X, ..., i = i)
+   J <- Jest(X, ...)
+   dif <- eval.fv(Jidot - J)
+   return(dif)
+ }
> E <- envelope(amacrine, Jdif, nsim = 39, i = "on", simulate = expression(rlabel(amacrine)))
```

```
> plot(E, main = "test of random labelling")
```



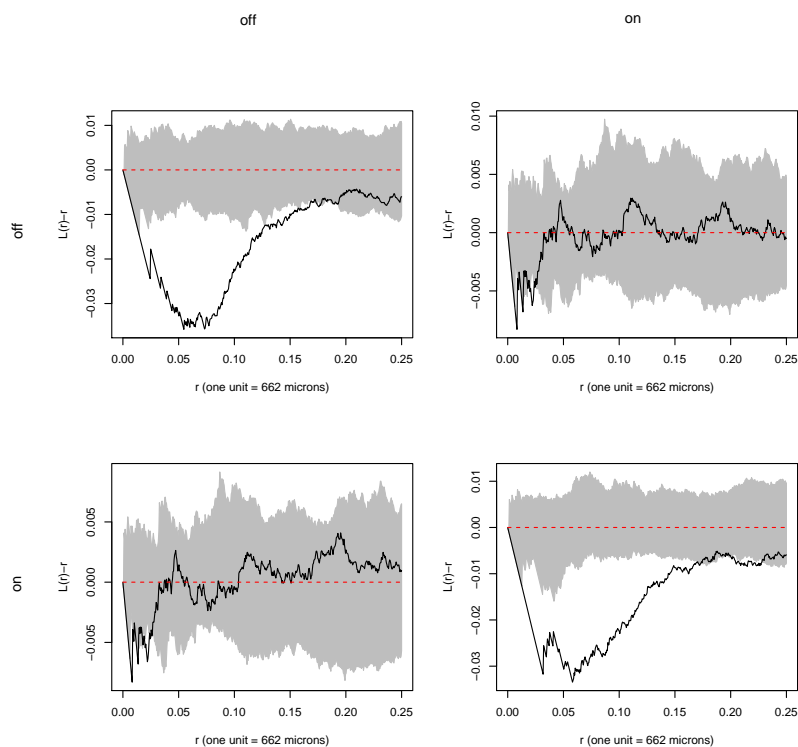
The random labelling hypothesis also seems to be accepted.

31.4.4 Arrays of envelopes

To compute a simulation envelope for the function K_{ij} for each pair of types i and j , use `alltypes` with the argument `envelope=TRUE`.

```
> aE <- alltypes(amacrine, Kcross, nsim = 39, envelope = TRUE)
> plot(aE, sqrt(./pi) - r ~ r, ylab = "L(r)-r")
```

array of envelopes of Kcross functions for amacrine.




31.5 Adjusting for inhomogeneity

Inhomogeneous versions of the multitype K function, L function and pair correlation function also exist.

The command names are `Kdot.inhom`, `Kcross.inhom`, `Ldot.inhom`, `Lcross.inhom`, `pcfdot.inhom` and `pcfcross.inhom`. They require separate estimates of the intensities of the first and second types of points. Again these intensities can be estimated by kernel smoothing.

```
> data(lansing)
> plot(Kcross.inhom(lansing, "whiteoak", "maple", sigma = 0.15))
```




32 Exploratory tools for marked point patterns

32.1 Numeric marks: distribution and trend

For a point pattern with marks that are numeric (real numbers or integers) or logical values, the mark values can be extracted using the `marks` function and inspected using the histogram or kernel density estimate:

```
> data(longleaf)
> hist(marks(longleaf))
```




To assess spatial trend in the marks, one way is to form a kernel regression smoother. The smoothed mark value at location $u \in \mathbb{R}^2$ is

$$\hat{m}(u) = \frac{\sum_i m_i \kappa(u - x_i)}{\sum_i \kappa(u - x_i)}$$

where k is the smoothing kernel, and m_i is the mark value at data point x_i . This is computed by `smooth.ppp`:

```
> plot(smooth.ppp(longleaf))
```

smooth.ppp(longleaf)




The plot shows that there is a region of younger trees in the northeast of the study region. If the marks are a data frame, the result of `smooth.ppp` will be a list of pixel images, one for each mark variable.

You can also use `cut.ppp` followed by `split.ppp` to look for spatial inhomogeneity of the marks:

```
> data(spruces)
> plot(split(cut(spruces, breaks = 3)))
```

split(cut(spruces, breaks = 3))



Other facilities include `markvar` which calculates a smoothed estimate of the local variance of the mark values.

32.2 Mark correlation function

The “mark correlation function” $\rho_f(r)$ of a stationary marked point process \mathbf{Y} is a measure of the dependence between the marks of two points of the process a distance r apart [61]. It is informally defined as

$$\rho_f(r) = \frac{\mathbb{E}[f(M_1, M_2)]}{\mathbb{E}[f(M, M')]}$$

where M_1, M_2 are the marks attached to two points of the process separated by a distance r , while M, M' are independent realisations of the marginal distribution of marks.

Here f is any function $f(m_1, m_2)$ with two arguments which are possible marks of the pattern, and which returns a nonnegative real value. Common choices of f are:

- for continuous real-valued marks, $f(m_1, m_2) = m_1 m_2$;
- for categorical marks (multitype point patterns), $f(m_1, m_2) = \mathbf{1}\{m_1 = m_2\}$;
- for marks taking values in $[0, 2\pi]$, $f(m_1, m_2) = \sin(m_1 - m_2)$.


Note that $\rho_f(r)$ is not a “correlation” in the usual statistical sense. It can take any nonnegative real value. The value 1 suggests “lack of correlation”: under random labelling, $\rho_f(r) \equiv 1$. The interpretation of values larger or smaller than 1 depends on the choice of function f .

The mark correlation function is computed in **spatstat** by **markcorr**. It has the syntax

```
> markcorr(X, f)
```

where **X** is a point pattern and **f** is an R language function.

```
> data(spruces)
> plot(markcorr(spruces))
```



The plot suggests a slight negative association between the sizes of nearby trees. This is somewhat hard to interpret.

The cumulative version of the mark correlation is computed by **markcorrint**.

A concept closely related to the mark correlation, based on the nearest neighbour of each point rather than all neighbours at a given distance, is computed by **nncorr**.

32.3 Reverse conditional moments

Schlather et al [59] defined the functions $E(r)$ and $V(r)$ to be the conditional mean and conditional variance of the mark attached to a typical random point, given that there exists another random point at a distance r away from it. More formally,


$$\begin{aligned} E(r) &= E_{0u}[M(0)] \\ V(r) &= E_{0u}[(M(0) - E(u))^2] \end{aligned}$$

where E_{0u} denotes the conditional expectation given that there are points of the process at the locations 0 and u separated by a distance r , and where $M(0)$ denotes the mark attached to the point 0.

These functions may serve as diagnostics for dependence between the points and the marks. If the points and marks are independent, then $E(r)$ and $V(r)$ should be constant (not depending on r).

These functions are computed using **Emark** and **Vmark**.

```
> plot(Emark(spruces))
```



Looking at the y axis scale we see that a *slight* drop in mean value (by about 10%) occurs near the origin. This plot suggests that a tree with a very close neighbour tends to have a diameter slightly smaller than average. The function **Emark** is bsimpler to interpret than the mark correlation function.

33 Multitype Poisson models

This section covers multitype Poisson process models: basic properties, simulation, and fitting models to data.

33.1 Theory

33.1.1 Complete spatial randomness and independence

A *uniform Poisson marked point process* in \mathbb{R}^2 with marks in \mathcal{M} can be defined in the following equivalent ways.

- randomly marked Poisson process (Poisson $[X]$, iid $[M|X]$): a Poisson point process of locations \mathbf{X} with intensity β is first generated. Then each point x_i is labelled with a random mark m_i , independently of other points, with distribution $\mathbb{P}\{M_i = m\} = p_m$ for $m \in \mathcal{M}$.
- superposition of independent Poisson processes (iid $[M]$, Poisson $[X|M]$): for each possible mark $m \in \mathcal{M}$, a Poisson process \mathbf{X}_m is generated, with intensity β_m . The points of \mathbf{X}_m are tagged with the mark m . Then the processes \mathbf{X}_m with different marks $m \in \mathcal{M}$ are superimposed, to yield a marked point process.
- Poisson marked point process (jointly Poisson $[X, M]$): a Poisson process on $\mathbb{R}^2 \times \mathcal{M}$ is generated, with intensity function $\lambda(u, m) = \beta_m$ at location u and mark m .

These constructions are *equivalent* when $\beta_m = p_m\beta$. See the lovely book by Kingman [45].

Since the established term CSR (‘complete spatial randomness’) is used to refer to the uniform Poisson point process, I propose that the uniform *marked* Poisson point process should be called ‘complete spatial randomness and independence’ (CSRI).

33.1.2 Inhomogeneous Poisson marked point processes

A *inhomogeneous* Poisson marked point process \mathbf{Y} with ‘joint’ intensity $\lambda(u, m)$ for locations u and mark values m is simply defined as an inhomogeneous Poisson point process on $\mathbb{R}^2 \times \mathcal{M}$ with intensity function $\lambda(u, m)$.

Let’s restrict attention to the case of categorical marks, where \mathcal{M} is finite. Then the process \mathbf{Y} has the following properties:

- The locations \mathbf{X} , obtained by removing the marks, constitute an inhomogeneous Poisson process in \mathbb{R}^2 with intensity function


$$\beta(u) = \sum_m \lambda(u, m).$$

- Conditional on the locations \mathbf{X} , the marks attached to the points are independent. For a point x_i the conditional distribution of the mark m_i is $\mathbb{P}\{M_i = m\} = \lambda(x_i, m)/\beta(x_i)$.
- The sub-process \mathbf{X}_m of points with mark m , is an inhomogeneous Poisson point process with intensity $\beta_m(u) = \lambda(u, m)$.
- The sub-processes \mathbf{X}_m of points with different marks m are independent processes.


33.2 Simulation

Realisations of Poisson marked point processes can be generated by 'hand', using `rmppoispp`. The first argument of this command specifies the intensity or intensity function $\lambda(u, m)$. It can be a constant, a vector of constants, or an R function.

```
> par(mfrow = c(1, 2))
> Xunif <- rmppoispp(100, types = c("A", "B"), win = square(1))
> plot(Xunif, main = "CSRI, intensity A=100, B=100")
> Xunif <- rmppoispp(c(100, 20), types = c("A", "B"), win = square(1))
> plot(Xunif, main = "CSRI, intensity A=100, B=20")
> par(mfrow = c(1, 1))
```



```
> X1 <- rmppoispp(function(x, y, m) {
+   300 * exp(-3 * x)
+ }, types = c("A", "B"))
> lamb <- function(x, y, m) {
+   ifelse(m == "A", 300 * exp(-4 * x), 300 * exp(-4 * (1 - x)))
+ }
> X2 <- rmppoispp(lamb, types = c("A", "B"))
> par(mfrow = c(1, 2))
> plot(X1, main = "")
> plot(X2, main = "")
> par(mfrow = c(1, 1))
```



33.3 Fitting Poisson models

Poisson marked point process models may be fitted to point pattern data using `ppm`. Currently the methods are only available for multitype point processes (categorical marks).

33.3.1 Probability densities

Let $W \subset \mathbb{R}^2$ be the study region, and \mathcal{M} the (finite) set of possible marks. Then a marked point pattern is a set

$$\mathbf{y} = \{(x_1, m_1), \dots, (x_n, m_n)\}, \quad x_i \in W, \quad m_i \in \mathcal{M}, \quad n \geq 0$$

of pairs (x_i, m_i) of locations x_i with marks m_i . It can be viewed as a point pattern in the Cartesian product $W \times \mathcal{M}$.

The probability density of a marked point process is a function $f(\mathbf{y})$ defined for all marked point patterns \mathbf{y} including the empty pattern \emptyset .

The process with probability density $f(\mathbf{y}) \equiv 1$ is the uniform Poisson marked point process with intensity 1 **for each mark**. That is, for this model, the sub-process of points with mark $m_i = m$ is a uniform Poisson process with intensity 1. If the marks are removed, we obtain a Poisson point process with intensity equal to $|\mathcal{M}|$, the number of possible types.

The uniform Poisson marked point process with intensity $\lambda(u, m) = \beta_m$ has probability density

$$\begin{aligned} f(\mathbf{y}) &= \exp\left(\sum_{m \in \mathcal{M}} (1 - \beta_m)|W|\right) \prod_{i=1}^{n(\mathbf{y})} \beta_{m_i} \\ &= \exp\left(\sum_{m \in \mathcal{M}} (1 - \beta_m)|W|\right) \prod_{m \in \mathcal{M}} \beta_m^{n_m(\mathbf{y})} \end{aligned}$$

where $n_m(\mathbf{y})$ is the number of points in \mathbf{y} having mark value m .

The inhomogeneous Poisson marked point process with intensity function $\lambda(u, m)$, at location $u \in W$ and mark $m \in \mathcal{M}$, has probability density

$$f(\mathbf{y}) = \exp\left(\sum_{m \in \mathcal{M}} \int_W (1 - \lambda(u, m)) du\right) \prod_{i=1}^{n(\mathbf{y})} \lambda(x_i, m_i). \quad (48)$$

33.3.2 Maximum likelihood

For the multitype Poisson process with intensity function $\lambda(u, m)$ at location $u \in W$ and mark $m \in \mathcal{M}$, the loglikelihood is, up to a constant,

$$\log L = \sum_{i=1}^n \log \lambda(x_i, m_i) - \sum_{m \in \mathcal{M}} \int_W \lambda(u, m) \, du. \quad (49)$$

where m_i is the mark attached to data point x_i . This is formally equivalent to the loglikelihood of a Poisson loglinear regression, so the Berman-Turner algorithm can again be used to maximise the loglikelihood.

33.3.3 Model-fitting in spatstat

Poisson marked point process models are fitted to data using `ppm`.

The trend formula in the call to `ppm` may involve the reserved name `marks` as a variable. This refers to the marks of the points. Since the marks are categorical, `marks` is treated as a `factor` variable for modelling purposes.

To fit the homogeneous multitype Poisson process (CSRI), equation (50), we call

```
> ppm(X, ~marks)
```

The formula `~marks` indicates that the trend depends only on the marks, and not on spatial location; since `marks` is a factor, the trend has a separate constant value for each level of `marks`. This is the model (50).

Note that if we had typed

```
> ppm(X, ~1)
```

this would have fitted the special case of CSRI where the intensities β_m are equal, $\beta_m \equiv \alpha$ say, for all possible marks. That model is only appropriate if we believe that all mark values are equally likely.

For the Lansing Woods data, the minimal model that makes sense is (50), so we call

```
> ppm(lansing, ~marks)
```

Stationary multitype Poisson process

Possible marks:

```
blackoak hickory maple misc redoak whiteoak
```

Trend formula: ~marks

Intensities:

beta_blackoak	beta_hickory	beta_maple	beta_misc	beta_redoak
135	703	514	105	346
beta_whiteoak				
448				

Since `lansing` is a multitype point pattern (its marks are categorical), the variable `marks` in the formula is a factor. The model has one parameter/coefficient for each level of the factor, i.e. one coefficient for each type of point. In other words, this is the homogeneous Poisson marked point process with intensity β_m for points of mark m .

You'll notice that the parameter estimates $\hat{\beta}_m$ coincide with those obtained from `summary.ppp` above. That is a consequence of the fact that the maximum likelihood estimates (obtained by `ppm`) are also the method-of-moments estimates (obtained by `summary.ppp`).

A more complicated example is

```
> ppm(lansing, ~marks + x)
```

Nonstationary multitype Poisson process

Possible marks:

blackoak hickory maple misc redoak whiteoak

Trend formula: ~marks + x

Fitted coefficients for trend formula:

(Intercept)	markshickory	marksmaple	marksmisc	marksredoak
4.94294727	1.65008211	1.33694849	-0.25131442	0.94116400
markswhiteoak	x			
1.19951845	-0.07581624			

This is the marked Poisson process whose intensity function $\lambda((x, y, m))$ at location (x, y) and mark m satisfies

$$\log \lambda((x, y, m)) = \alpha_m + \beta x$$

where $\alpha_1, \dots, \alpha_6$ and β are parameters. The intensity is loglinear in x , with a different intercept for each mark, but the same slope ("parallel loglinear regression"). In the printout above, the fitted slope parameter β is $\hat{\beta} = -0.07581624$. As discussed in Section 15.3 on page 98, the fitted coefficients α_m for the categorical mark are interpreted in the light of the 'contrasts' in force. The default is the treatment contrasts, and the first level of the mark is **blackoak**, so in this case the fitted coefficient for $m=\text{blackoak}$ is 4.942947, while the fitted coefficient for $m=\text{hickory}$ is $4.942947 + 1.650082 = 6.593029$ and so on.

```
> ppm(lansing, ~marks * x)
```

Nonstationary multitype Poisson process

Possible marks:

blackoak hickory maple misc redoak whiteoak

Trend formula: ~marks * x

Fitted coefficients for trend formula:

(Intercept)	markshickory	marksmaple	marksmisc	marksredoak
5.2378062	1.4424915	0.6795604	-0.8482907	0.6916392
markswhiteoak	x	markshickory:x	marksmaple:x	marksmisc:x
1.0901772	-0.7063987	0.4511157	1.3243326	1.2138278
marksredoak:x	markswhiteoak:x			
0.5380413	0.2421379			

The symbol $*$ here is an 'interaction' in the usual sense for linear models. The fitted model is the marked Poisson process with

$$\log \lambda((x, y, m)) = \alpha_m + \beta_m x$$

where $\alpha_1, \dots, \alpha_6$ and β_1, \dots, β_6 are parameters. The intensity is loglinear in x with a different slope and intercept for each mark.

The result of `ppm` is again an object of class "ppm" representing a fitted point process model. To plot the fitted intensity and conditional intensity of the fitted model, use `plot.ppm`. For a multitype point process you will get a separate plot for each possible mark value.

More complicated examples are:

```
> ppm(lansing, ~marks * polynom(x, y, 2))
> ppm(lansing, ~marks * harmonic(x, y, 2))
```

33.3.4 Facilities available

A fitted multitype Poisson process model can be manipulated using any of the methods available for the class `ppm`:

<code>print</code>	print basic information
<code>summary</code>	print detailed summary information
<code>plot</code>	plot the fitted (conditional) intensity
<code>predict</code>	fitted (conditional) intensity
<code>fitted</code>	fitted (conditional) intensity at data points
<code>update</code>	re-fit the model
<code>coef</code>	extract the fitted coefficient vector $\hat{\theta}$
<code>vcov</code>	variance-covariance matrix of $\hat{\theta}$
<code>anova</code>	analysis of deviance
<code>logLik</code>	evaluate log-pseudolikelihood
<code>model.matrix</code>	extract design matrix
<code>formula</code>	extract trend formula of model
<code>terms</code>	extract terms in model formula

The following functions are also available:

<code>step</code>	stepwise model selection
<code>drop1</code>	one step backward in model selection
<code>model.images</code>	compute images of canonical covariates in model
<code>effectfun</code>	fitted intensity as function of one covariate

A fitted multitype Poisson process model can be simulated automatically using `rmh.ppm` or `simulate.ppm`.

34 Gibbs models for multitype point patterns

Gibbs point process models (section 26) are also available for marked point processes, and can be fitted to data using `ppm`. Currently the methods are only implemented for *multitype* point processes (categorical marks), so we restrict attention to this case.

34.1 Gibbs models

Much of the theory of Gibbs models described in Section 26 carries over immediately to *multitype* point processes.

34.1.1 Conditional intensity

The conditional intensity $\lambda(u, \mathbf{X})$ of an (unmarked) point process \mathbf{X} at a location u was defined in section 26.5. Roughly speaking $\lambda(u, \mathbf{x}) du$ is the conditional probability of finding a point near u , given that the rest of the point process \mathbf{X} coincides with \mathbf{x} .

For a marked point process \mathbf{Y} the conditional intensity is a function $\lambda((u, m), \mathbf{Y})$ giving a value at a location u for each possible mark m . For a *finite* set of marks M , we can interpret $\lambda((u, m), \mathbf{y}) du$ as the conditional probability finding a point *with mark m* near u , given the rest of the marked point process.

The conditional intensity is related to the probability density $f(\mathbf{y})$ by

$$\lambda((u, m), \mathbf{y}) = \frac{f(\mathbf{y} \cup \{u\})}{f(\mathbf{y})}$$

for $(u, m) \notin \mathbf{y}$.

For Poisson processes, the conditional intensity $\lambda((u, m), \mathbf{y})$ coincides with the intensity function $\lambda(u, m)$ and does not depend on the configuration \mathbf{y} . For example, the homogeneous Poisson multitype point process or ‘CSRI’ (Section 33.1.1) has conditional intensity

$$\lambda((u, m), \mathbf{y}) = \beta_m \quad (50)$$

where $\beta_m \geq 0$ are constants which can be interpreted in several equivalent ways (section 26.5). The sub-process consisting of points of type *only* is Poisson with intensity β_m . The process obtained by ignoring the types, and combining all the points, is Poisson with intensity $\beta = \sum_m \beta_m$. The marks attached to the points are i.i.d. with distribution $p_m = \beta_m / \beta$.

34.1.2 Pairwise interactions

A *multitype* pairwise interaction process is a Gibbs process with probability density of the form

$$f(\mathbf{y}) = \alpha \left[\prod_{i=1}^{n(\mathbf{y})} b_{m_i}(x_i) \right] \left[\prod_{i < j} c_{m_i, m_j}(x_i, x_j) \right] \quad (51)$$

where $b_m(u), m \in \mathcal{M}$ are functions determining the ‘first order trend’ for points of each type, and $c_{m, m'}(u, v), m, m' \in \mathcal{M}$ are functions determining the interaction between a pair of points of given types m and m' . The interaction functions must be symmetric, $c_{m, m'}(u, v) = c_{m, m'}(v, u)$ and $c_{m, m'} \equiv c_{m', m}$. The conditional intensity is

$$\lambda((u, m); \mathbf{y}) = b_m(u) \prod_{i=1}^{n(\mathbf{y})} c_{m, m_i}(u, x_i). \quad (52)$$

34.1.3 Pairwise interactions not depending on marks

The simplest examples of multitype pairwise interaction processes are those in which the interaction term $c_{m,m'}(u, v)$ does not depend on the marks m, m' . For example, we can take any of the interaction functions $c(u, v)$ described in section 26.3 and use it to construct a marked point process.

Such processes can be constructed equivalently as follows [14]:

- an *unmarked* Gibbs process is generated with first order term $b(u) = \sum_{m \in \mathcal{M}} b_m(u)$ and pairwise interaction $c(u, v)$.
- each point x_i of this unmarked process is labelled with a mark m_i with probability distribution $\mathbb{P}\{m_i = m\} = b_i(x_i)/b(x_i)$ independent of other points.

If additionally the intensity functions are constant, $b_m(u) \equiv \beta_m$, then such a point process has the random labelling property.

34.1.4 Mark-dependent pairwise interactions

Various complex kinds of behaviour can be created by postulating a pairwise interaction that does depend on the marks.

A simple example is the *multitype hard core process* in which $\beta_m(u) \equiv \beta$ and

$$c_{m,m'}(u, v) = \begin{cases} 1 & \text{if } \|u - v\| > r_{m,m'} \\ 0 & \text{if } \|u - v\| \leq r_{m,m'} \end{cases} \quad (53)$$

where $r_{m,m'} = r_{m',m} > 0$ is the hard core distance for type m with type m' . In this process, two points of type m and m' respectively can never come closer than the distance $r_{m,m'}$.

By setting $r_{m,m'} = 0$ for a particular pair of marks m, m' we effectively remove the interaction term between points of these types. If there are only two types, say $\mathcal{M} = \{1, 2\}$, then setting $r_{1,2} = 0$ implies that the sub-processes \mathbf{X}_1 and \mathbf{X}_2 , consisting of points of types 1 and 2 respectively, are independent point processes. In other words the process satisfies the independence-of-components property.

The *multitype Strauss process* has pairwise interaction term

$$c_{m,m'}(u, v) = \begin{cases} 1 & \text{if } \|u - v\| > r_{m,m'} \\ \gamma_{m,m'} & \text{if } \|u - v\| \leq r_{m,m'} \end{cases} \quad (54)$$

where $r_{m,m'} > 0$ are interaction radii as above, and $\gamma_{m,m'} \geq 0$ are interaction parameters.

In contrast to the unmarked Strauss process, which is well-defined only when its interaction parameter γ is between 0 and 1, the multitype Strauss process allows some of the interaction parameters $\gamma_{m,m'}$ to exceed 1 for $m \neq m'$, provided one of the relevant types has a hard core ($\gamma_{m,m} = 0$ or $\gamma_{m',m'} = 0$).

If there are only two types, say $\mathcal{M} = \{1, 2\}$, then setting $\gamma_{1,2} = 1$ implies that the sub-processes \mathbf{X}_1 and \mathbf{X}_2 , consisting of points of types 1 and 2 respectively, are independent Strauss processes.

The *multitype Strauss-hard core process* has pairwise interaction term

$$c_{m,m'}(u, v) = \begin{cases} 0 & \text{if } \|u - v\| < h_{m,m'} \\ \gamma_{m,m'} & \text{if } h_{m,m'} \leq \|u - v\| \leq r_{m,m'} \\ 1 & \text{if } \|u - v\| > r_{m,m'} \end{cases} \quad (55)$$

where $r_{m,m'} > 0$ are interaction distances and $\gamma_{m,m'} \geq 0$ are interaction parameters as above, and $h_{m,m'}$ are hard core distances satisfying $h_{m,m'} = h_{m',m}$ and $0 < h_{m,m'} < r_{m,m'}$.

34.2 Pseudolikelihood for multitype Gibbs processes

Models can be fitted by maximum pseudolikelihood. For a multitype Gibbs point process with conditional intensity $\lambda((u, m); \mathbf{y})$, the log pseudolikelihood is

$$\log \text{PL} = \sum_{i=1}^{n(\mathbf{y})} \log \lambda((x_i, m_i); \mathbf{y}) - \sum_{m \in \mathcal{M}} \int_W \lambda((u, m); \mathbf{y}) \, du. \quad (56)$$

The pseudolikelihood can be maximised using an extension of the Berman-Turner device [9].

34.3 Fitting Gibbs models to multitype data

Marked point process models may be fitted to point pattern data using `ppm`. Currently the methods are only available for multitype point processes (categorical marks).

34.3.1 Interactions not depending on marks

The model-fitting function `ppm` expects an argument `interaction` that specifies the interpoint interaction structure of the point process. The default is ‘no interaction’, corresponding to a Poisson process.

On page 165 there is a list of interpoint interactions for modelling *unmarked* point patterns. These interactions can also be used, without modification, to fit models to *multitype* point patterns.

For example

```
> ppm(lansing, ~marks, Strauss(0.07))
```

fits a multitype version of the Strauss process (section 26.3.2) in which the conditional intensity is

$$\lambda((u, m), \mathbf{y}) = \beta_m \gamma^{t(u, \mathbf{y})}. \quad (57)$$

Here β_m are constants which account for the unequal abundance of the different species of tree. The other quantities are the same as in (42). The interaction between two trees is assumed to be the same for all species, and is controlled by the interaction parameter γ and interaction radius $r = 0.07$. For example, this includes the case $\gamma = 0$ where no two trees (whatever species they belong to) come closer than 0.07 units apart, a ‘multitype hard core process’.

34.3.2 Interactions depending on marks

There are two additional interpoint interactions defined in `spatstat` for multitype point patterns:

MultiStrauss the multitype Strauss process

MultiStraussHard multitype hybrid hard core / Strauss process

In these models, the interaction between two points depends on the types of the points as well as their separation.

In the multitype Strauss process (54), for each pair of types i and j there is an interaction radius r_{ij} and interaction parameter γ_{ij} . In simple terms, each pair of points, with marks i and j say, contributes an interaction term $\gamma_{i,j}$ if the distance between them is less than the interaction distance $r_{i,j}$. These parameters must satisfy $r_{ij} = r_{ji}$ and $\gamma_{ij} = \gamma_{ji}$. The conditional intensity is

$$\lambda((u, i), \mathbf{y}) = \beta_i \prod_j \gamma_{i,j}^{t_{i,j}(u, \mathbf{y})} \quad (58)$$

where $t_{i,j}(u, \mathbf{y})$ is the number of points in \mathbf{y} , with mark equal to j , lying within a distance $r_{i,j}$ of the location u .

To fit the stationary multitype Strauss process to the dataset `betacells`, we must specify the matrix of interaction radii r_{ij} :

```
> data(betacells)
> r <- 30 * matrix(c(1, 2, 2, 1), nrow = 2, ncol = 2)
> ppm(betacells, ~1, MultiStrauss(c("off", "on"), r), rbord = 60)
```

Stationary Multitype Strauss process

Possible marks:

off on

First order terms:

```
      beta_off      beta_on
0.0001373652 0.0001373652
```

Interaction: Pairwise interaction family

Interaction: Multitype Strauss process

2 types of points

Possible types:

```
[1] "off" "on"
```

Interaction radii:

```
      off on
off  30 60
on   60 30
```

Fitted interaction parameters gamma_ij:

```
      off      on
off 0.0000 0.8303
on  0.8303 0.0000
```

Relevant coefficients:

```
markoffxoff markoffxon markonxon
-17.2378706 -0.1860184 -17.2138383
```

To fit a nonstationary multitype Strauss process with log-cubic polynomial trend:

```
> ppm(betacells, ~polynom(x, y, 3), MultiStrauss(c("off", "on"),
+      r), rbord = 60)
```

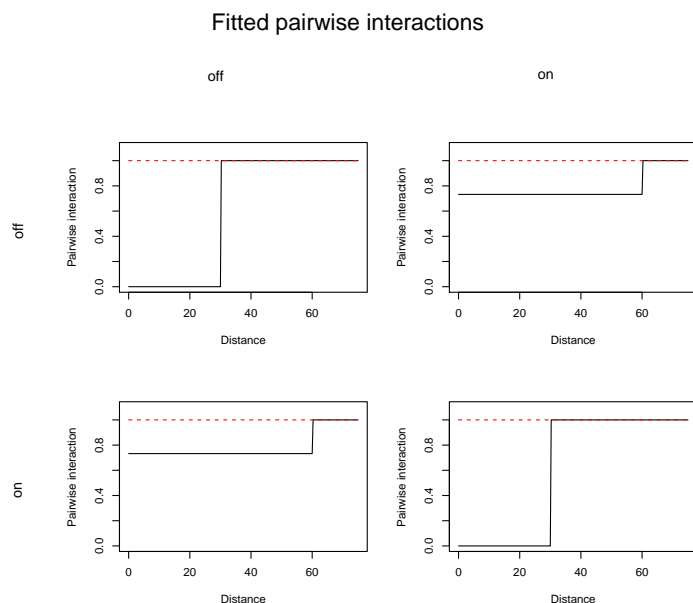
For more detailed explanation and examples of modelling and the interpretation of model formulae for point processes, see [11].

34.3.3 Plotting the fitted interaction

The fitted pairwise interaction in a point process model can be plotted using `fitin`. The value returned by `fitin` is a function array (class "fasp").

```
> model <- ppm(betacells, ~polynom(x, y, 3), MultiStrauss(c("off",
+ "on"), r), rbord = 60)

> plot(fitin(model))
```



34.3.4 Simulating a Gibbs model

A fitted Gibbs point process model can be simulated using `rmh.ppm` or `simulate.ppm`.

```
> rmh(model, verbose = FALSE)
```

```
marked planar point pattern: 139 points
multitype, with levels = off      on
window: rectangle = [28.08, 778.08] x [16.2, 1007.02] microns
```

It's also possible to simulate any specified Gibbs model using `rmh.default` with the model specified 'by hand' using `rmhmodel.default`.

PART VIII. HIGHER DIMENSIONS AND OTHER SPATIAL DATA


Part VIII of the workshop deals with point patterns in 3D, space-time, and multi-dimensional space time. It also covers line segment patterns and stochastic geometry techniques.

35 Line segment data

`spatstat` also has some facilities for handling spatial patterns of *line segments*.

For example, the `copper` dataset in `spatstat` contains a dataset `copper$Lines` that records the locations of geological faults in a survey region.

```
> data(copper)
> L <- copper$Lines
> L <- rotate(L, pi/2)
> plot(L)
```



A spatial pattern of line segments is represented by an object of class "`psp`". It consists of a list of line segments (given by the coordinates of their two endpoints), and a window in which the line segments were observed. The line segments may also carry marks.

Objects of class "`psp`" can be created by the function `psp` or obtained by converting other data using the function `as.psp`.

Capabilities available for this class include:

<code>is.psp</code>	check whether an object is a <code>psp</code>
<code>[.psp</code>	subset operator (also performs clipping)
<code>marks.psp</code>	extract marks
<code>marks<-.psp</code>	assign marks
<code>as.data.frame.psp</code>	extract coordinates and marks
<code>endpoints.psp</code>	extract endpoints of line segments
<code>midpoints.psp</code>	compute midpoints of line segments
<code>unitname.psp</code>	determine units of length
<code>rescale.psp</code>	change units of length
<code>lengths.psp</code>	compute lengths of line segments
<code>angles.psp</code>	compute orientation angles for line segments
<code>rotate.psp</code>	rotate a line segment pattern
<code>shift.psp</code>	shift a line segment pattern
<code>affine.psp</code>	apply affine transformation
<code>crossing.psp</code>	find intersection points between line segments
<code>selfcrossing.psp</code>	find intersection points between line segments
<code>density.psp</code>	kernel-smoothed intensity image
<code>rshift.psp</code>	apply random shift to each line segment
<code>superimposePSP</code>	combine several line segment patterns
<code>identify.psp</code>	point-and-click to identify line segments
<code>as.mask.psp</code>	convert line segments to binary mask
<code>pixellate.psp</code>	measure the length of line in each pixel

There are also the usual methods

```
plot.psp      plot a line segment pattern
print.psp     print information on a line segment pattern
summary.psp   compute summary of a line segment pattern
```

```
> summary(L)
```

```
146 line segments
```

```
Lengths:
```

```
      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.09242  6.61400 12.18000 15.02000 19.95000 65.48000
```

```
Total length: 2192.57251480451 km
```

```
Length per unit area: 0.196937548404655
```

```
Angles (radians):
```

```
      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.008107 0.549500 1.747000 1.378000 2.113000 2.912000
```

```
Window: rectangle = [-158.233, -0.19]x[-0.335, 70.11]km
```


```
Window area = 11133.3 square km
```

```
Unit of length: 1 km
```

```
> plot(distmap(L))
```

```
> plot(L, add = TRUE)
```

distmap(L)



For measuring distances to and from line segments, there are the following facilities:

```
pairedist.psp  distances between line segments
crossdist.psp  distances between two sets of line segments
nndist.psp     closest distances between line segments
nncross        closest distances between points and line segments
nearestsegment closest line segment to each point
project2segment point on line segment closest to specified point
```

It's also possible to generate random points that lie on a set of line segments:

```
runifpointOnLines  fixed number of random points on lines
rpoisppOnLines     Poisson process on lines
```

36 Point patterns in 3D

Basic support for three-dimensional point patterns has recently been added to `spatstat`.

A point pattern in 3D is an object of class "pp3". It is created by the function `pp3`.

```
> X <- pp3(x = runif(100), y = runif(100), z = runif(100), box3(c(0,
+      1)))
> summary(X)
```

Three-dimensional point pattern

100 points

Box: [0, 1] x [0, 1] x [0, 1] units

Volume 1 cubic unit

Average intensity 100 points per cubic unit

The window containing the point pattern is currently required to be a 3D rectangular box, stored as an object of class "box3", created by the function `box3`. Facilities available for `box3` objects include `print`, `summary`, `diameter`, `volume`, `shortside` and `eroded.volumes`.

Facilities available for `pp3` objects include:

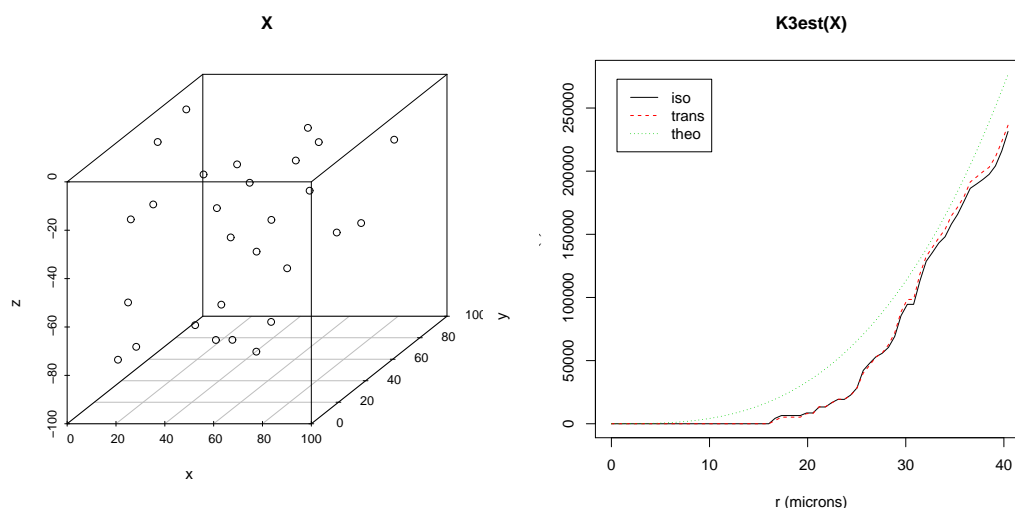
<code>print</code>	print basic information
<code>summary</code>	print detailed summary of data
<code>plot</code>	plot 3D points
<code>as.data.frame</code>	extract coordinates and marks
<code>npoints</code>	number of points
<code>coords</code>	extract coordinates
<code>coords<-</code>	change coordinates
<code>marks</code>	extract marks
<code>marks<-</code>	change marks
<code>unmark</code>	remove marks
<code>unitname</code>	extract name of unit of length
<code>unitname<-</code>	set name of unit of length
<code>pairdist</code>	matrix of distances between all pairs of points
<code>crossdist</code>	distances between all pairs of points in two patterns
<code>nnlist</code>	nearest neighbour distance
<code>nnwhich</code>	identify nearest neighbour
<code>F3est</code>	three-dimensional empty space function
<code>G3est</code>	three-dimensional nearest neighbour function
<code>K3est</code>	three-dimensional K -function
<code>pcf3est</code>	three-dimensional pair correlation function
<code>envelope</code>	simulation envelopes of summary functions
<code>nnclean</code>	Byers-Raftery nearest neighbour cleaning

Random 3D point patterns can be generated by `runifpoint3` and `rpoispp3`. Currently the only 3D point pattern data installed in `spatstat` is the dataset `osteo`.

```
> data(osteo)
> X <- osteo$pts[[36]]
> par(mfrow = c(1, 2))
> plot(X)
> plot(K3est(X))
```

	lty	col	key	label	meaning
iso	1	1	iso	K3[iso](r)	isotropic-corrected estimate of K3(r)
trans	2	2	trans	K3[trans](r)	translation-corrected estimate of K3(r)
theo	3	3	theo	K3[pois](r)	theoretical Poisson K3(r)

```
> par(mfrow = c(1, 1))
```



37 Point patterns in multi-dimensional space-time

Experimental support for multi-dimensional space-time point patterns has recently been added to **spatstat**. These objects belong to the class "ppx" and are created by the function **ppx**. There may be any number of dimensions of space, and any number of dimensions of time.

```
> df <- data.frame(x = runif(100, max = 3), y = runif(100, max = 3),
+   z = runif(100, max = 2), t = runif(100))
> bb <- boxx(c(0, 3), c(0, 3), c(0, 2), c(0, 1))
> X <- ppx(data = df, domain = bb, temporal = "t")
> summary(X)
```

```
Multidimensional point pattern
100 points
3-dimensional space coordinates (x,y,z)
1-dimensional time coordinates (t)
Domain:
  4-dimensional box:
[0, 3] x [0, 3] x [0, 2] x [0, 1] units
```

The point pattern may have marks of any type, stored as a **hyperframe**.

```
> marks(X) <- with(as.hyperframe(df), disc(centre = c(x, y)))
> X
```



```

Multidimensional point pattern
100 points
1 column of marks: marks
Domain:
    4-dimensional box:
[0, 3] x [0, 3] x [0, 2] x [0, 1] units

```

The window containing the point pattern is currently required to be a rectangular box, stored as an object of class "boxx", created by the function `boxx`. Facilities available for `boxx` objects include `print`, `summary`, `diameter`, `volume`, `shortside` and `eroded.volumes`.

Facilities currently available for `ppx` objects include:

<code>print</code>	print basic information
<code>summary</code>	print detailed summary of data
<code>as.data.frame</code>	extract coordinates and marks
<code>npoints</code>	number of points
<code>coords</code>	extract coordinates
<code>coords<-</code>	change coordinates
<code>marks</code>	extract marks
<code>marks<-</code>	change marks
<code>unmark</code>	remove marks
<code>unitname</code>	extract name of unit of length
<code>unitname<-</code>	set name of unit of length
<code>pairdist</code>	matrix of distances between all pairs of points
<code>crossdist</code>	distances between all pairs of points in two patterns
<code>nn-dist</code>	nearest neighbour distance
<code>nnwhich</code>	identify nearest neighbour

Random hyperdimensional point patterns can be generated by `runifpointx` and `rpoisppx`.

There is also code to compute the theoretical distribution of the nearest neighbour distances, i.e. the distance from a reference point to the k -th nearest neighbour in a uniform Poisson point process in m dimensions:

<code>dknn</code>	probability density
<code>pknn</code>	cumulative probability
<code>qknn</code>	quantiles
<code>rknn</code>	random generator

38 Replicated data and hyperframes

Sometimes the data may consist of several point patterns. We may have *replicated* point pattern data obtained by repeating an experiment, such as the locations of cell nuclei in 10 different samples of tissue. These data could be stored in a simple list Z , where the i th entry in the list, $Z[[i]]$, is the i th point pattern object.

A *hyperframe* is like a data frame, except that the entries can be any type of object. For example, a hyperframe can include a column, each of whose entries is a point pattern. The only constraint is that all the entries in a column must be of the same type.

A hyperframe is a convenient way to store replicated point pattern data together with auxiliary data. A hyperframe can store the results of a designed experiment where the “response” is a point pattern and the covariates are other types of data.

The following facilities are available:

<code>hyperframe</code>	create a hyperframe
<code>as.hyperframe</code>	convert other data to a hyperframe
<code>print.hyperframe</code>	print a representation of a hyperframe
<code>dim.hyperframe</code>	dimensions of hyperframe
<code>\$.hyperframe</code>	extract column of hyperframe
<code>[.hyperframe</code>	extract subset of hyperframe
<code>[<-.hyperframe</code>	replace subset of hyperframe
<code>as.data.frame.hyperframe</code>	convert hyperframe to data frame
<code>cbind.hyperframe</code>	combine several hyperframes
<code>rbind.hyperframe</code>	combine several hyperframes
<code>with.hyperframe</code>	compute an expression in each row of hyperframe

The dataset `osteo` installed in `spatstat` contains three-dimensional point patterns recorded in several sampling volumes in each of several bone samples [15]. It is a hyperframe with the following columns:

<code>id</code>	string identifier of bone sample
<code>shortid</code>	last numeral in <code>id</code>
<code>brick</code>	serial number of sampling volume within bone sample
<code>pts</code>	three dimensional point pattern (class <code>pp3</code>)
<code>depth</code>	the depth of the brick in microns

```
> data(osteo)
> osteo[1:5, ]
```

Hyperframe:

```
      id shortid brick  pts depth
1 c77za4      4     1 (pp3)   45
2 c77za4      4     2 (pp3)   60
3 c77za4      4     3 (pp3)   55
4 c77za4      4     4 (pp3)   60
5 c77za4      4     5 (pp3)   85
```

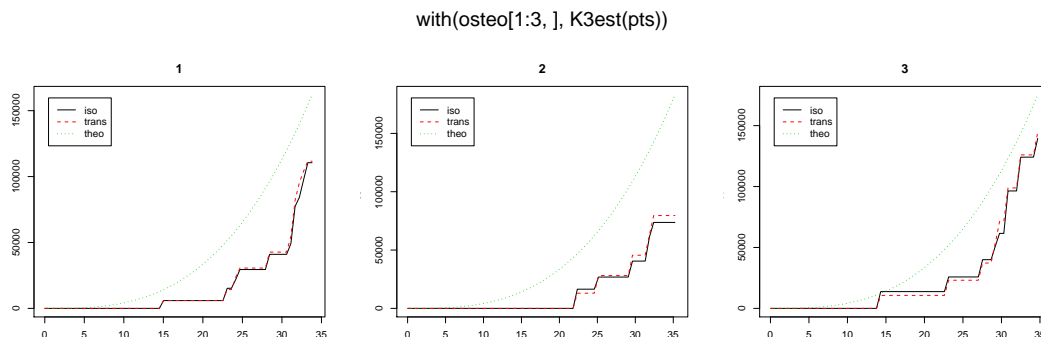
The command `with.hyperframe` can be used to evaluate an expression in each row of the hyperframe. For example, to count the number of points in each of the first 10 patterns:

```
> with(osteo, npoints(pts))[1:10]

 1  2  3  4  5  6  7  8  9 10
13 11 11 12 14 12 16 15 18 16
```

To plot the three-dimensional K -function estimates of the first 3 patterns:

```
> plot(with(osteo[1:3, ], K3est(pts)))
```




39 Stochastic geometry

Spatstat provides some basic facilities for stochastic geometry (model simulation and tools for image analysis).

The dataset **heather** contains Diggle's [32] heather data, a binary pixel mask of the presence/absence of heather (*Calluna vulgaris*) in a survey plot in Sweden. Three versions of the data are given.


```
> data(heather)
> plot(heather)
```



Elementary morphological operations can be applied using **erosion**, **dilation**, **opening** and **closing** along with geometrical operations such as **complement.owin**, **intersect.owin**, **union.owin**, **setminus.owin**. The areas of dilations and erosions can be computed using either **distmap** or **eroded.areas**, **dilated.areas**. Connected components can be identified using **connected**.


The empirical *spherical contact distribution* or empty space function of such a dataset can be computed using **Hest**:

```
> X <- heather$fine
> H <- Hest(X)
> plot(H)
```



Patterns of line segments are represented by objects of the class "psp" described elsewhere. Patterns of *infinite* lines in the two-dimensional plane are represented by objects of the class "infinite" created by the function `infinite`. The command `rpoisline` creates a random tessellation delineated by a Poisson line process. Switzer's random set is generated by `rMosaicSet` and the analogous random field is generated by `rMosaicField`.

```
> X <- rpoislinetess(5)
> Y <- rMosaicSet(X, p = 0.5)
> plot(Y, col = "grey")
```



40 Further information on spatstat

Help files

For information on a particular command in **spatstat**, consult the online help file by typing `help(command)`. The help files are detailed and extensive. The complete manual is over 700 pages long.

For examples of the use of a particular command, read the examples section in the help file, or type `example(command)` to see the examples executed.

Quick reference

Type `help(spatstat)` for a quick-reference overview of all the functions available in the package.

For a demonstration of many of the capabilities of **spatstat**, type `demo(spatstat)`.

For a visual display of all the datasets supplied in **spatstat**, type `demo(data)`.

Website

The website www.spatstat.org contains information on recent updates to the package, frequently-asked questions, bug fixes, literature and other developments.

Updates

Spatstat is updated approximately once a month. Descriptions of the updates are given on the spatstat.org website, and can also be accessed from within the package, by typing `latest.news()` or `news(package="spatstat")`.

Modelling

For more **spatstat** examples on fitting point process models, see [11]. For more discussion on point process modelling strategies, see [3].

Citation

If you use **spatstat** in a research publication, it would be much appreciated if you could cite the paper [10], or mention **spatstat** in the acknowledgements.

In doing so, you will help us to justify the expenditure of time and effort on maintaining and developing the package.

Citation details are also available in the package by typing `citation(package="spatstat")`.

Queries and requests

If you have difficulty in getting the package to do what you want, or if you have a suggestion for additional features that could be added, please contact the package authors, Adrian.Baddeley@csiro.au and r.turner@auckland.ac.nz, or email the R special interest group in spatial and geographical statistics, r-sig-geo@stat.math.ethz.ch.

References

- [1] F.P. Agterberg. Automatic contouring of geological maps to detect target areas for mineral exploration. *Journal of the International Association for Mathematical Geology*, 6:373–395, 1974.
- [2] L. Anselin. Local indicators of spatial association – LISA. *Geographical Analysis*, 27:93–115, 1995.
- [3] A. Baddeley. Modelling strategies. In A.E. Gelfand, P.J. Diggle, M. Fuentes, and P. Guttorp, editors, *Handbook of Spatial Statistics*, chapter 20, pages 339–369. CRC Press, Boca Raton, 2010.
- [4] A. Baddeley, M. Berman, N.I. Fisher, A. Hardegen, R.K. Milne, D. Schuhmacher, R. Shah, and R. Turner. Spatial logistic regression and change-of-support for spatial Poisson point processes. *Electronic Journal of Statistics*, 4:1151–1201, 2010. doi: 10.1214/10-EJS581.
- [5] A. Baddeley, Y.-M. Chang, Y. Song, and R. Turner. Diagnostics for transformation of covariates in spatial Poisson point process models. Submitted for publication.
- [6] A. Baddeley, J. Møller, and A.G. Pakes. Properties of residuals for spatial point processes. *Annals of the Institute of Statistical Mathematics*, 60:627–649, 2008.
- [7] A. Baddeley, J. Møller, and R. Waagepetersen. Non- and semiparametric estimation of interaction in inhomogeneous point patterns. *Statistica Neerlandica*, 54(3):329–350, November 2000.
- [8] A. Baddeley, E. Rubak, and J. Møller. Score, pseudo-score and residual diagnostics for goodness-of-fit of spatial point process models. Submitted for publication.
- [9] A. Baddeley and R. Turner. Practical maximum pseudolikelihood for spatial point patterns (with discussion). *Australian and New Zealand Journal of Statistics*, 42(3):283–322, 2000.
- [10] A. Baddeley and R. Turner. Spatstat: an R package for analyzing spatial point patterns. *Journal of Statistical Software*, 12(6):1–42, 2005. URL: www.jstatsoft.org, ISSN: 1548-7660.
- [11] A. Baddeley and R. Turner. Modelling spatial point patterns in R. In A. Baddeley, P. Gregori, J. Mateu, R. Stoica, and D. Stoyan, editors, *Case Studies in Spatial Point Pattern Modelling*, number 185 in Lecture Notes in Statistics, pages 23–74. Springer-Verlag, New York, 2006. ISBN: 0-387-28311-0.
- [12] A. Baddeley, R. Turner, J. Møller, and M. Hazelton. Residual analysis for spatial point processes (with discussion). *Journal of the Royal Statistical Society, series B*, 67(5):617–666, 2005.
- [13] A.J. Baddeley. Spatial sampling and censoring. In O.E. Barndorff-Nielsen, W.S. Kendall, and M.N.M. van Lieshout, editors, *Stochastic Geometry: Likelihood and Computation*, chapter 2, pages 37–78. Chapman and Hall, London, 1999.
- [14] A.J. Baddeley and J. Møller. Nearest-neighbour Markov point processes and random sets. *International Statistical Review*, 57:89–121, 1989.

-
- [15] A.J. Baddeley, R.A. Moyeed, C.V. Howard, and A. Boyde. Analysis of a three-dimensional point pattern with replication. *Applied Statistics*, 42(4):641–668, 1993.
- [16] A.J. Baddeley and B.W. Silverman. A cautionary example on the use of second-order methods for analyzing point patterns. *Biometrics*, 40:1089–1094, 1984.
- [17] A.J. Baddeley and M.N.M. van Lieshout. Area-interaction point processes. *Annals of the Institute of Statistical Mathematics*, 47:601–619, 1995.
- [18] G. Barnard. Contribution to discussion of “The spectral analysis of point processes” by M.S. Bartlett. *Journal of the Royal Statistical Society, series B*, 25:294, 1963.
- [19] M. Bell and G. Grunwald. Mixed models for the analysis of replicated spatial point patterns. *Biostatistics*, 5:633–648, 2004.
- [20] M. Berman. Testing for spatial association between a point process and another stochastic process. *Applied Statistics*, 35:54–62, 1986.
- [21] M. Berman and T.R. Turner. Approximating point process likelihoods with GLIM. *Applied Statistics*, 41:31–38, 1992.
- [22] J. Besag and P.J. Diggle. Simple Monte Carlo tests for spatial pattern. *Applied Statistics*, 26:327–333, 1977.
- [23] J.E. Besag and P. Clifford. Generalized Monte Carlo significance tests. *Biometrika*, 76:633–642, 1989.
- [24] R. Bivand, E.J. Pebesma, and V. Gómez-Rubio. *Applied spatial data analysis with R*. Springer, 2008.
- [25] D.R. Brillinger. Comparative aspects of the study of ordinary time series and of point processes. In P.R. Krishnaiah, editor, *Developments in Statistics*, pages 33–133. Academic Press, 1978.
- [26] S. Byers and A.E. Raftery. Nearest-neighbour clutter removal for estimating features in spatial point processes. *Journal of the American Statistical Association*, 93:577–584, 1998.
- [27] E. Choi and P. Hall. Nonparametric analysis of earthquake point-process data. In M. de Gunst, C. Klaassen, and A. van der Vaart, editors, *State of the art in probability and statistics: Festschrift for Willem R. van Zwet*, pages 324–344. Institute of Mathematical Statistics, Beachwood, Ohio, 2001.
- [28] N. Cressie and L.B. Collins. Analysis of spatial point patterns using bundles of product density LISA functions. *Journal of Agricultural, Biological and Environmental Statistics*, 6:118–135, 2001.
- [29] N. Cressie and L.B. Collins. Patterns in spatial point locations: local indicators of spatial association in a minefield with clutter. *Naval Research Logistics*, 48:333–347, 2001.
- [30] N.A.C. Cressie. *Statistics for Spatial Data*. John Wiley and Sons, New York, 1991.
- [31] D.J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes*. Springer Verlag, New York, 1988.
-

-
- [32] P.J. Diggle. Binary mosaics and the spatial pattern of heather. *Biometrics*, 37:531–539, 1981.
- [33] P.J. Diggle. *Statistical analysis of spatial point patterns*. Academic Press, London, 1983.
- [34] P.J. Diggle. A point process modelling approach to raised incidence of a rare phenomenon in the vicinity of a prespecified point. *Journal of the Royal Statistical Society, series A*, 153:349–362, 1990.
- [35] P.J. Diggle. *Statistical Analysis of Spatial Point Patterns*. Hodder Arnold, London, second edition, 2003.
- [36] P.J. Diggle, N. Lange, and F. M. Benes. Analysis of variance for replicated spatial point patterns in clinical neuroanatomy. *Journal of the American Statistical Association*, 86:618–625, 1991.
- [37] P.J. Diggle, J. Mateu, and H.E. Clough. A comparison between parametric and non-parametric approaches to the analysis of replicated spatial point patterns. *Advances in Applied Probability (SGSA)*, 32:331–343, 2000.
- [38] P.J. Diggle and B. Rowlingson. A conditional approach to point process modelling of elevated risk. *Journal of the Royal Statistical Society, series A (Statistics in Society)*, 157(3):433–440, 1994.
- [39] M. Dwass. Modified randomization tests for nonparametric hypotheses. *Annals of Mathematical Statistics*, 28:181–187, 1957.
- [40] R. Foxall and A. Baddeley. Nonparametric measures of association between a spatial point process and a random set, with geological applications. *Applied Statistics*, 51(2):165–182, 2002.
- [41] A.C.A. Hope. A simplified Monte Carlo significance test procedure. *Journal of the Royal Statistical Society, series B*, 30:582–598, 1968.
- [42] C.V. Howard, S. Reid, A.J. Baddeley, and A. Boyde. Unbiased estimation of particle density in the tandem-scanning reflected light microscope. *Journal of Microscopy*, 138:203–212, 1985.
- [43] F. Huang and Y. Ogata. Improvements of the maximum pseudo-likelihood estimators in various spatial statistical models. *Journal of Computational and Graphical Statistics*, 8(3):510–530, 1999.
- [44] J. Illian, A. Penttinen, H. Stoyan, and D. Stoyan. *Statistical Analysis and Modelling of Spatial Point Patterns*. John Wiley and Sons, Chichester, 2008.
- [45] J.F.C. Kingman. *Poisson Processes*. Oxford University Press, New York, 1993.
- [46] G.M. Laslett. Censoring and edge effects in areal and line transect sampling of rock joint traces. *Mathematical Geology*, 14:125–140, 1982.
- [47] P.A.W. Lewis. Recent results in the statistical analysis of univariate point processes. In P.A.W. Lewis, editor, *Stochastic point processes*, pages 1–54. Wiley, New York, 1972.
- [48] J.K. Lindsey. *The analysis of stochastic processes using GLIM*. Springer, Berlin, 1992.
-

-
- [49] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.
- [50] J. Møller and R. P. Waagepetersen. Modern spatial point process modelling and inference (with discussion). *Scandinavian Journal of Statistics*, 34:643–711, 2007.
- [51] J. Møller and R.P. Waagepetersen. *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall/CRC, Boca Raton, 2004.
- [52] Y. Ogata. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, 83:9–27, 1988.
- [53] M. Prokešová, U. Hahn, and E.B. Vedel Jensen. Statistics for locally scaled point processes. In A. Baddeley, P. Gregori, J. Mateu, R. Stoica, and D. Stoyan, editors, *Case studies in spatial point process modeling*, number 185 in Lecture Notes in Statistics, pages 99–123. Springer, New York, 2005.
- [54] B.D. Ripley. Modelling spatial patterns (with discussion). *Journal of the Royal Statistical Society, series B*, 39:172–212, 1977.
- [55] B.D. Ripley. Simulating spatial patterns: dependent samples from a multivariate density. *Applied Statistics*, 28:109–112, 1979.
- [56] B.D. Ripley. *Spatial Statistics*. John Wiley and Sons, New York, 1981.
- [57] B.D. Ripley. *Statistical Inference for Spatial Processes*. Cambridge University Press, 1988.
- [58] A. Särkkä. *Pseudo-likelihood approach for pair potential estimation of Gibbs processes*. Number 22 in Jyväskylä Studies in Computer Science, Economics and Statistics. University of Jyväskylä, 1993.
- [59] M. Schlather, P. Riberio, and P.J. Diggle. Detecting dependence between marks and locations of marked point processes. *Journal of Royal Statistical Society Series B*, 66:79–93, 2004.
- [60] D. Stoyan and P. Grabarnik. Second-order characteristics for stochastic structures connected with Gibbs point processes. *Mathematische Nachrichten*, 151:95–100, 1991.
- [61] D. Stoyan and H. Stoyan. *Fractals, Random Shapes and Point Fields*. John Wiley and Sons, Chichester, 1995.
- [62] J.W. Tukey. Discussion of paper by F.P. Agterberg and S.C. Robinson. *Bulletin of the International Statistical Institute*, 44(1):596, 1972. Proceedings, 38th Congress, International Statistical Institute.
- [63] M.N.M. van Lieshout. *Markov Point Processes and their Applications*. Imperial College Press, London, 2000.
- [64] M.N.M. van Lieshout and A.J. Baddeley. A nonparametric measure of spatial interaction in point patterns. *Statistica Neerlandica*, 50:344–361, 1996.
- [65] R. Waagepetersen. An estimating function approach to inference for inhomogeneous Neyman-Scott processes. *Biometrics*, 63:252–258, 2007.
-

Index

- analysis of deviance, 103
- area-interaction process, 159
- binary mask, 33, 48
- circular windows, 46
- classes, 32
 - in R, 32
 - in `spatstat`, 32
- cluster models
 - fitting, 144, 152
 - inhomogeneous, 151
 - fitting, 152
- complete spatial randomness, 88
 - and independence, 179, 204
 - definition, 88
 - Kolmogorov-Smirnov test, 91
 - quadrat counting test, 89
- conditional intensity, 160
 - for marked point processes, 210
- contrasts, 98, 208
- covariate effects, 9
- covariates, 7, 16, 98
 - in `ppm`, 98
- Cox process, 141
- CSRI, 179, 204
 - conditional intensity, 210
 - fitting to data, 207
 - simulating, 205
- data entry, 38
 - checking, 43
 - GIS formats, 45, 49
 - marked point patterns, 181
 - marks, 40
 - point-and-click, 44
- data sharpening, 148
- datasets
 - inspecting, 21
 - provided in `spatstat`, 30
- dispatching, 32
- distance methods, 115
- distances
 - empty space, 115, 116
 - nearest neighbour, 115, 122
 - pairwise, 115, 125
- `distmap`, 115
- edge effects, 116
- empty space distances, 115, 116
- empty space function, 117, 222
- envelopes, 132
 - and Monte Carlo tests, 132
 - for any fitted model, 136
 - for any simulation procedure, 137
 - in `spatstat`, 133
 - of summary functions, 132
- exploratory data analysis, 23
 - for marked point patterns, 200
 - for multitype point patterns, 187
- fitted model, 166
 - goodness-of-fit, 106, 172
 - interpretation of coefficients, 98
 - methods for, 99
 - residuals, 107, 173
 - simulation of, 104
- fitting models
 - by Huang-Ogata method, 170
 - `kppm`, 144, 152
 - maximum pseudolikelihood, 162
 - to marked point patterns, 207, 212
 - via summary statistics, 144
- `fv`, 37
- geometrical transformations, 57
- Gibbs models, 156
 - area-interaction, 159
 - Diggle-Gates-Stibbard, 159
 - Diggle-Gratton, 159
 - fitting, 162
 - by Huang-Ogata method, 170
 - maximum pseudolikelihood, 162
 - `ppm`, 162
 - fitting to marked point patterns, 212
 - goodness-of-fit, 172
 - hard core process, 157
 - in `spatstat`, 165
 - infinite order interaction, 159
 - multitype, 210
 - maximum pseudolikelihood, 212
 - multitype pairwise interaction, 210
 - pairwise interaction, 159
 - residuals, 173
 - simulation, 161

- simulation of fitted model, 168
- soft core, 159
- Strauss process, 158
- Strauss-hard core, 159
- GIS formats, 45, 49
- goodness-of-fit, 106
 - for fitted Gibbs model, 172
 - for Poisson models, 106
- hard core process, 157
 - multitype, 211
- heather data, 222
- Huang-Ogata method, 170
- hyperframe, 221
- hyperframe**, 32
- im**, 32, 63
- images, 63
 - computing with, 70
 - creating, 63
 - from raw data, 63
 - exploratory inspection of, 68
 - extracting subset, 69
 - plotting, 66
 - returned by a function, 65
- independence of components, 179, 197
- intensity
 - function, 79
 - kernel estimator, 79
 - homogeneous, 78
 - inhomogeneous, 79
 - investigation of, 78
 - measure, 79
 - of multitype point process, 187
- interaction, 8, 11
 - distance methods, 115
 - exploratory methods, 114
 - in **spatstat**, 165
 - multitype, 210, 212
 - in **spatstat**, 212
 - plotting a fitted interaction, 214
 - Q–Q plot, 175
 - simple models, 139
 - summary functions, 115
- K* function, 24, 125
 - for multitype point pattern, 190
 - inhomogeneous, 149
- kernel estimator of intensity, 79, 80
- kernel smoothing of marks, 200
- Kolmogorov-Smirnov test
 - of CSR, 91
 - of inhomogeneous Poisson, 107
- kppm**, 144, 152
- line segments, 216
- LISA, 148
- local features, 148
- locally scaled point process, 153
- lurking variable plot, 109
- maptools** package, 45
- mark connection function, 195
- mark correlation function, 202
- marked point patterns
 - cutting marks into bands, 185
 - data entry, 181
 - exploratory data analysis, 200
 - exploring marks, 200
 - inspecting, 183
 - joint and conditional analysis, 179
 - manipulating, 184
 - methodological issues, 179
 - model-fitting, 207, 212
 - probabilistic formulation, 178
 - randomisation tests, 179
 - separating into types, 185
 - summary functions, 190
- marks, 6, 16, 178
 - categorical, 41
 - data entry, 40
 - exploratory data analysis, 200
 - manipulating, 184
 - operations on, 56
 - smoothing, 200
 - spatial trend in, 200
 - versus covariates, 15
- markstat**, 190
- marktable**, 189
- Matern cluster process, 140
- maximum likelihood, 95
- maximum pseudolikelihood, 162, 212
 - for multitype Gibbs models, 212
 - improvements over, 170
- methods, 32
 - default method, 34
 - dispatch, 32
- minimum contrast, 144
- model validation, 106, 171

- models, 25, 224
- Monte Carlo test, 132
 - pointwise, 133
 - simultaneous, 134
- multidimensional point pattern, 219
- multitype hard core process, 211
- multitype point pattern, 10, 11, 27, 41
- multitype point patterns
 - exploratory data analysis, 187
 - separating into types, 185
 - summary functions, 190
- multitype point process
 - intensity, 187
- multitype Strauss process, 211
- nearest neighbour cleaning, 148
- nearest neighbour distances, 115, 122
- nndist**, 115
- nuisance parameters, 168
- owin**, 32, 46
- pairdist**, 115
- pairwise distances, 115, 125
- pairwise interaction process, 157
- point pattern, 6
 - data entry, 38
 - in many dimensions, 219
 - in space-time, 219
 - in three dimensions, 218
 - marked, 178
 - marks, 6, 16
 - multitype, 10, 11
 - needs window, 54
 - point process model for, 13
 - standard model, 14
- point process, 13
- point process models, 25
 - area-interaction, 159
 - Diggle-Gates-Stibbard, 159
 - Diggle-Gratton, 159
 - Gibbs, 156
 - hard core, 157
 - infinite order interaction, 159
 - pairwise interaction, 157, 159
 - soft core, 159
 - Strauss, 158
 - Strauss-hard core, 159
- Poisson cluster processes, 140
- Poisson line process, 223
- Poisson line tessellation, 223
- Poisson models
 - fitting, 96
 - goodness-of-fit, 106
 - homogeneous, 88
 - inhomogeneous, 95
 - log-likelihood, 96
 - marked, 204
 - maximum likelihood, 95
 - residuals, 107
- Poisson point process
 - homogeneous
 - definition, 88
 - simulation, 88
 - inhomogeneous
 - definition, 95
 - fitting, 96
 - likelihood, 96
 - motivation, 95
 - simulation, 95
- Poisson-derived models, 139
- polygonal windows, 33, 47
- pp3**, 32, 218
- ppm**, 99, 166
 - marked Gibbs point process models, 212
 - marked Poisson point process models, 207
 - methods for, 100
- ppp**, 32
 - combining several, 61
 - extracting subset, 55
 - format, 53
 - geometrical transformations, 57
 - in arbitrary window, 51
 - manipulating, 53
 - needs window, 54
 - operations on, 55
 - random perturbations, 58
 - ways to make, 44
- ppx**, 32, 219
- probability density, 156
- profile pseudolikelihood, 168
- pseudolikelihood, 162
 - profile pseudolikelihood, 168
- psp**, 32, 216
- quadrat counting, 23, 79
- quadrat counting test
 - of CSR, 89
- quadrat test

- of inhomogeneous Poisson, 106
- R, 18
 - contributed packages, 19
 - for spatial data formats, 45
 - for spatial statistics, 19
 - where to get, 18
- random labelling, 179, 198
- random perturbations, 58
- random thinning, 95
- randomisation tests, 179, 197
 - for marked point patterns, 197
- rectangular windows, 33, 46
- replicated point patterns, 221
- residuals, 107, 173
 - for fitted Gibbs model, 173
 - for Poisson models, 107
 - lurking variable plot, 109
 - Q–Q plot, 174
 - smoothed residual field, 108
- return value, 35
- rpoispp**, 88, 95
- runifpoint**, 89
- sequential models, 142
- shapefiles, 45
- shapefiles** package, 45
- simulation
 - of fitted Gibbs model, 168
 - of fitted Poisson model, 104
- smoothed residual field, 108
- sp** package, 45
- spatstat**, 20, 224
 - citing, 20, 224
 - getting started, 20
 - help files, 224
 - installing, 20
 - queries and requests, 224
 - updates, 224
 - website, 224
- spherical contact distribution, 222
- split**, 28
- standard model, 14
- stochastic geometry, 222
- Strauss process, 158
 - fitting to data, 163
 - multitype, 211
- summary functions, 115
 - and Monte Carlo tests, 132
 - critique, 130
 - edge effects, 116
 - envelopes, 132
 - F , 117
 - for multitype point patterns, 190
 - G , 122
 - inference using, 132
 - inhomogeneous K , 149
 - J , 128
 - K , 125
 - L , 126
 - mark connection, 195
 - mark correlation, 202
 - model-fitting with, 144
 - pair correlation, 126
 - Switzer's random set, 223
- tess**, 32
- tests
 - χ^2 quadrat counting, 89
 - Kolmogorov-Smirnov, 91, 107
 - Monte Carlo, 132
- thinning, 141
- Thomas process, 140
- three dimensional point pattern, 218
- tips, 32, 36, 42, 56, 116, 119, 134, 181
- treatment contrasts, 98
- unitname**, 43
- units of length, 43
- validation, 106, 171
- windows, 46
 - binary mask, 33, 48
 - circular, 46
 - GIS formats, 49
 - needed in any point pattern, 54
 - operations on, 50
 - polygonal, 33, 47
 - rectangular, 33, 46
 - returned by functions, 49
- χ^2 quadrat counting test, 89