

Customising spatial data classes and methods*

Edzer Pebesma[†]

Feb 2008

Contents

1	Programming with classes and methods	2
1.1	S3-style classes and methods	5
1.2	S4-style classes and methods	6
2	Animal track data in package trip	7
2.1	Generic and constructor functions	8
2.2	Methods for trip objects	8
3	Multi-point data: SpatialMultiPoints	10
4	Spatio-temporal grids	14
5	Analysing spatial Monte Carlo simulations	18

Although the classes defined in the **sp** package cover many needs, they do not go far beyond the most typical GIS data models. In applied research, it often happens that customised classes would suit the actual data coming from the instruments better. Since **S4** classes have mechanisms for inheritance, it may be attractive to build on the **sp** classes, so as to utilise their methods where appropriate. Here, we will demonstrate a range of different settings in which **sp** classes can be extended. Naturally, this is only useful for researchers with specific and clear needs, so our goal is to show how (relatively) easy it may be to prototype classes extending **sp** classes for specific purposes.

*This vignette formed pp. 127–148 of the first edition of Bivand, R. S., Pebesma, E. and Gómez-Rubio V. (2008) Applied Spatial Data Analysis with R, Springer-Verlag, New York. It was retired from the second edition (2013) to accommodate material on other topics, and is made available in this form with the understanding of the publishers. It has been updated to the 2013 state of the software, e.g. using **over**.

[†]Institute for Geoinformatics, University of Muenster, Weseler Strasse 253, 48151 Münster, Germany. edzer.pebesma@uni-muenster.de

1 Programming with classes and methods

This section will explain the elementary basics of programming with classes and methods in R. The S language (implemented in R and S-PLUS™) contains two mechanisms for creating classes and methods: the traditional S3 system and the more recent S4 system (see Section 2.2 in [Bivand et al. \(2008\)](#), in which classes were described for the useR — here they are described for the developer). This chapter is not a full introduction to R programming (see [Braun and Murdoch, 2007](#), for more details), but it will try to give some feel of how the `Spatial` classes in package `sp` can be extended to be used for wider classes of problems. For full details, the interested reader is referred to e.g. [Venables and Ripley \(2000\)](#) and [Chambers \(1998\)](#), the latter being a reference for new-style S4 classes and methods. Example code is for example found in the source code for package `sp`, available from CRAN.

Suppose we define `myfun` as

```
> myfun <- function(x) {  
+   x + 2  
+ }
```

then, calling it with the numbers 1, 2 and 3 results in

```
> myfun(1:3)  
[1] 3 4 5
```

or alternatively using a named argument:

```
> myfun(x = 1:3)  
[1] 3 4 5
```

The return value of the function is the last expression evaluated. Often, we want to wrap existing functions, such as a plot function:

```
> plotXplus2Yminus3 <- function(x, y, ...) {  
+   plot(x = x + 2, y = y - 3, ...)  
+ }
```

In this case, the `...` is used to pass information to the `plot` function without explicitly anticipating what it will be: named arguments `x` and `y`, or the first two arguments if they are unnamed are processed, remaining arguments are passed on. The `plot` function is a generic method, with an instance that depends on the class of its first (S3) or first n arguments (S4). The available instances of `plot` are shown for S3-type methods by

```
> methods("plot")  
[1] plot,ANY,ANY-method  
[2] plot,Extent,missing-method  
[3] plot,Intervals,ANY-method  
[4] plot,Intervals,missing-method  
[5] plot,Intervals_full,ANY-method  
[6] plot,Intervals_full,missing-method  
[7] plot,Raster,ANY-method  
[8] plot,Raster,Raster-method  
[9] plot,STF,missing-method  
[10] plot,STFDF,missing-method  
[11] plot,STI,missing-method  
[12] plot,STS,missing-method
```

- [13] plot, STSDF, missing-method
- [14] plot, STT, missing-method
- [15] plot, Satellite, ANY-method
- [16] plot, SpatExtent, missing-method
- [17] plot, SpatGraticule, missing-method
- [18] plot, SpatNetwork, missing-method
- [19] plot, SpatRaster, SpatRaster-method
- [20] plot, SpatRaster, character-method
- [21] plot, SpatRaster, missing-method
- [22] plot, SpatRaster, numeric-method
- [23] plot, SpatVector, character-method
- [24] plot, SpatVector, missing-method
- [25] plot, SpatVector, numeric-method
- [26] plot, SpatVectorCollection, missing-method
- [27] plot, SpatVectorCollection, numeric-method
- [28] plot, SpatVectorProxy, missing-method
- [29] plot, Spatial, missing-method
- [30] plot, SpatialGrid, missing-method
- [31] plot, SpatialGridDataFrame, missing-method
- [32] plot, SpatialLines, missing-method
- [33] plot, SpatialMultiPoints, missing-method
- [34] plot, SpatialPixels, missing-method
- [35] plot, SpatialPixelsDataFrame, missing-method
- [36] plot, SpatialPoints, missing-method
- [37] plot, SpatialPolygons, missing-method
- [38] plot, profile.mle, missing-method
- [39] plot.HoltWinters*
- [40] plot.Intervals*
- [41] plot.Intervals_full*
- [42] plot.R6*
- [43] plot.SOM*
- [44] plot.StVariogram*
- [45] plot.TukeyHSD*
- [46] plot.acf*
- [47] plot.bclust*
- [48] plot.classIntervals*
- [49] plot.data.frame*
- [50] plot.decomposed.ts*
- [51] plot.default
- [52] plot.dendrogram*
- [53] plot.density*
- [54] plot.ecdf
- [55] plot.factor*
- [56] plot.formula*
- [57] plot.freelway*
- [58] plot.function
- [59] plot.gstatVariogram*
- [60] plot.hclust*
- [61] plot.histogram*
- [62] plot.ica*
- [63] plot.isoreg*
- [64] plot.lm*
- [65] plot.medpolish*
- [66] plot.mlm*
- [67] plot.pal_continuous*
- [68] plot.pal_discrete*
- [69] plot.pointPairs*

```

[70] plot.ppr*
[71] plot.prcomp*
[72] plot.princomp*
[73] plot.profile*
[74] plot.profile.nls*
[75] plot.raster*
[76] plot.replot_xts*
[77] plot.sf*
[78] plot.sfc_CIRCULARSTRING*
[79] plot.sfc_GEOMETRY*
[80] plot.sfc_GEOMETRYCOLLECTION*
[81] plot.sfc_LINESTRING*
[82] plot.sfc_MULTILINESTRING*
[83] plot.sfc_MULTIPPOINT*
[84] plot.sfc_MULTIPOLYGON*
[85] plot.sfc_POINT*
[86] plot.sfc_POLYGON*
[87] plot.sfg*
[88] plot.shingle*
[89] plot.somgrid*
[90] plot.spec*
[91] plot.stepfun
[92] plot.stft*
[93] plot.stl*
[94] plot.svm*
[95] plot.table*
[96] plot.transform*
[97] plot.trellis*
[98] plot.ts
[99] plot.tskernel*
[100] plot.tune*
[101] plot.units*
[102] plot.variogramCloud*
[103] plot.variogramMap*
[104] plot.variogramModel*
[105] plot.xts*
[106] plot.zoo*
see '?methods' for accessing help and source code

```

and for S4-type methods by

```

> library(sp)
> showMethods("plot")

```

```

Function: plot (package base)
x="ANY", y="ANY"
x="Extent", y="missing"
x="Intervals", y="ANY"
x="Intervals", y="missing"
x="Intervals_full", y="ANY"
x="Intervals_full", y="missing"
x="Raster", y="ANY"
x="Raster", y="Raster"
x="STF", y="missing"
x="STFDF", y="missing"
x="STI", y="missing"
x="STS", y="missing"
x="STSDF", y="missing"
x="STT", y="missing"

```

```

x="Satellite", y="ANY"
x="SpatExtent", y="missing"
x="SpatGraticule", y="missing"
x="SpatNetwork", y="missing"
x="SpatRaster", y="SpatRaster"
x="SpatRaster", y="character"
x="SpatRaster", y="missing"
x="SpatRaster", y="numeric"
x="SpatVector", y="character"
x="SpatVector", y="missing"
x="SpatVector", y="numeric"
x="SpatVectorCollection", y="missing"
x="SpatVectorCollection", y="numeric"
x="SpatVectorProxy", y="missing"
x="Spatial", y="missing"
x="SpatialGrid", y="missing"
x="SpatialGridDataFrame", y="missing"
x="SpatialLines", y="missing"
x="SpatialMultiPoints", y="missing"
x="SpatialPixels", y="missing"
x="SpatialPixelsDataFrame", y="missing"
x="SpatialPoints", y="missing"
x="SpatialPolygons", y="missing"
x="profile.mle", y="missing"

```

where we first loaded `sp` to make sure there are some `S4` plot methods to show.

1.1 S3-style classes and methods

Building S3-style classes is simple. Suppose we want to build an object of class `foo`:

```

> x <- rnorm(10)
> class(x) <- "foo"
> x

[1] 0.9074673 2.0500511 -1.6567538 -1.8064411 0.4462380 -1.4624058 0.7551423
[8] -0.6608567 1.0524294 0.1700278
attr(,"class")
[1] "foo"

```

If we plot this object, e.g., by `plot(x)` we get the same plot as when we would not have set the class to `foo`. If we know, however, that objects of class `foo` need to be plotted without symbols but with connected lines, we can write a plot method for this class:

```

> plot.foo <- function(x, y, ...) {
+   plot.default(x, type = "l", ...)
+ }

```

after which `plot(x)` will call this particular method, rather than a default plot method.

Class inheritance is obtained in S3 when an object is given multiple classes, as in

```

> class(x) <- c("foo", "bar")
> plot(x)

```

For this plot, first function `plot.foo` will be looked for, and if not found the second option `plot.bar` will be looked for. If none of them is found, the default `plot.default` will be used.

The S3 class mechanism is simple and powerful. Much of R works with it, including key functions such as `lm`.

```
> data(meuse)
> class(meuse)

[1] "data.frame"

> class(lm(log(zinc) ~ sqrt(dist), meuse))

[1] "lm"
```

There is, however, no checking that a class with a particular name does indeed contain the elements that a certain method for it expects. It also has design flaws, as method specification by dot separation is ambiguous in case of names such as `as.data.frame`, where one cannot tell whether it means that the method `as.data` acts on objects of class `frame`, or the method `as` acts on objects of class `data.frame`, or none of them (the answer is: none). For such reasons, S4-style classes and methods were designed.

1.2 S4-style classes and methods

S4-style classes are formally defined, using `setClass`. As an example, somewhat simplified versions of classes `CRS` and `Spatial` in `sp` are

```
> setClass("CRS", representation(proj4args = "character"))
> setClass("Spatial", representation(bbox = "matrix",
+   proj4string = "CRS"), validity <- function(object) {
+   bb <- bbox(object)
+   if (!is.matrix(bb))
+     return("bbox should be a matrix")
+   n <- dimensions(object)
+   if (n < 2)
+     return("spatial.dimension should be 2 or more")
+   if (any(is.na(bb)))
+     return("bbox should never contain NA values")
+   if (any(!is.finite(bb)))
+     return("bbox should never contain infinite values")
+   if (any(bb[, "max"] < bb[, "min"]))
+     return("invalid bbox: max < min")
+   TRUE
+ })
```

The command `setClass` defines a class name as a formal class, gives the names of the class elements (called slots), and their type—type checking will happen upon construction of an instance of the class. Further checking, e.g., on valid dimensions and data ranges can be done in the `validity` function. Here, the validity function retrieves the bounding box using the generic `bbox` method. Generics, if not defined in the base R system, e.g.,

```
> isGeneric("show")

[1] TRUE
```

can be defined with `setGeneric`. Defining a specific instance of a generic is done by `setMethod`:

```
> setGeneric("bbox", function(obj) standardGeneric("bbox"))
> setMethod("bbox", signature = "Spatial", function(obj) obj@bbox)
```

where the signature tells the class of the first (or first n) arguments. Here, the `@` operator is used to access the `bbox` slot in an `S4` object, not to be confused with the `$` operator to access list elements.

We will now illustrate this mechanism by providing a few examples of classes, building on those available in package `sp`.

2 Animal track data in package `trip`

CRAN Package `trip`, written by Michael Sumner ([Kirkwood et al., 2006](#); [Page et al., 2006](#)), provides a class for animal tracking data. Animal tracking data consist of sets of (x, y, t) stamps, grouped by an identifier pointing to an individual animal, sensor or perhaps isolated period of monitoring. A strategy for this (slightly simplified from that of `trip`) is to extend the `SpatialPointsDataFrame` class by a length 2 character vector carrying the names of the time column and the trip identifier column in the `SpatialPointsDataFrame` attribute table.

Package `trip` does a lot of work to read and analyse tracking data from data formats typical for tracking data (Argos DAT), removing duplicate observations and validating the objects, e.g., checking that time stamps increase and movement speeds are realistic. We ignore this and stick to the bare bones.

We now define a class called `trip` that extends `SpatialPointsDataFrame`:

```
> library(sp)
> setClass("trip", representation("SpatialPointsDataFrame",
+   TOR.columns = "character"), validity <- function(object) {
+   if (length(object@TOR.columns) != 2)
+     stop("Time/id column names must be of length 2")
+   if (!all(object@TOR.columns %in% names(object@data)))
+     stop("Time/id columns must be present in attribute table")
+   TRUE
+ })
> showClass("trip")
```

```
Class "trip" [in ".GlobalEnv"]
```

```
Slots:
```

```
Name: TOR.columns      data coords.nrs      coords      bbox
Class: character data.frame      numeric      matrix      matrix
```

```
Name: proj4string
Class:      CRS
```

```
Extends:
```

```
Class "SpatialPointsDataFrame", directly
Class "SpatialPoints", by class "SpatialPointsDataFrame", distance 2
Class "Spatial", by class "SpatialPointsDataFrame", distance 3
Class "SpatialVector", by class "SpatialPointsDataFrame", distance 3
```

that checks, upon creation of objects, that indeed two variable names are passed and that these names refer to variables present in the attribute table.

2.1 Generic and constructor functions

It would be nice to have a constructor function, just like `data.frame` or `SpatialPoints`, so we now create it and set it as the generic function to be called in case the first argument is of class `SpatialPointsDataFrame`.

```
> trip.default <- function(obj, TORnames) {
+   if (!is(obj, "SpatialPointsDataFrame"))
+     stop("trip only supports SpatialPointsDataFrame")
+   if (is.numeric(TORnames))
+     TORnames <- names(obj)[TORnames]
+   new("trip", obj, TOR.columns = TORnames)
+ }
> if (!isGeneric("trip")) setGeneric("trip", function(obj,
+   TORnames) standardGeneric("trip"))

[1] "trip"

> setMethod("trip", signature(obj = "SpatialPointsDataFrame",
+   TORnames = "ANY"), trip.default)
```

We can now try it out, with turtle data:

```
> turtle <- read.csv(system.file("external/seamap105_mod.csv",
+   package = "sp"))
> timestamp <- as.POSIXlt(strptime(as.character(turtle$obs_date),
+   "%m/%d/%Y %H:%M:%S"), "GMT")
> turtle <- data.frame(turtle, timestamp = timestamp)
> turtle$lon <- ifelse(turtle$lon < 0, turtle$lon + 360,
+   turtle$lon)
> turtle <- turtle[order(turtle$timestamp), ]
> coordinates(turtle) <- c("lon", "lat")
> proj4string(turtle) <- CRS("+proj=longlat +ellps=WGS84")
> turtle$id <- c(rep(1, 200), rep(2, nrow(coordinates(turtle)) -
+   200))
> turtle_trip <- trip(turtle, c("timestamp", "id"))
> summary(turtle_trip)
```

Object of class trip

Coordinates:

	min	max
--	-----	-----

lon	140.923	245.763
-----	---------	---------

lat	21.574	39.843
-----	--------	--------

Is projected: FALSE

proj4string : [+proj=longlat +ellps=WGS84 +no_defs]

Number of points: 394

Data attributes:

	id	obs_date	timestamp
Min.	:1.000	Length :394	Min. :1996-08-11 01:15:00
1st Qu.:	1.000	N.unique :394	1st Qu.:1996-10-30 00:10:04
Median :	1.000	N.blank : 0	Median :1997-01-24 23:31:01
Mean :	1.492	Min.nchar: 19	Mean :1997-01-26 06:56:46
3rd Qu.:	2.000	Max.nchar: 19	3rd Qu.:1997-04-10 12:26:20
Max.	:2.000		Max. :1997-08-13 20:19:46

2.2 Methods for trip objects

The summary method here is not defined for `trip`, but is the default summary inherited from class `Spatial`. As can be seen, nothing special about the trip

features is mentioned, such as what the time points are and what the identifiers. We could alter this by writing a class-specific summary method

```
> summary.trip <- function(object, ...) {
+   cat("Object of class \"trip\"\\nTime column: ")
+   print(object@TOR.columns[1])
+   cat("Identifier column: ")
+   print(object@TOR.columns[2])
+   print(summary(as(object, "Spatial")))
+   print(summary(object@data))
+ }
> setMethod("summary", "trip", summary.trip)
> summary(turtle_trip)

Object of class "trip"
Time column: [1] "timestamp"
Identifier column: [1] "id"
Object of class Spatial
Coordinates:
      min      max
lon 140.923 245.763
lat  21.574  39.843
Is projected: FALSE
proj4string : [+proj=longlat +ellps=WGS84 +no_defs]
      id      obs_date      timestamp
Min.   :1.000   Length :394   Min.     :1996-08-11 01:15:00
1st Qu.:1.000   N.unique :394   1st Qu.:1996-10-30 00:10:04
Median :1.000   N.blank  : 0   Median :1997-01-24 23:31:01
Mean   :1.492   Min.nchar: 19   Mean   :1997-01-26 06:56:46
3rd Qu.:2.000   Max.nchar: 19   3rd Qu.:1997-04-10 12:26:20
Max.   :2.000           Max.   :1997-08-13 20:19:46
```

As `trip` extends `SpatialPointsDataFrame`, subsetting using `"["` and column selection or replacement using `"[["` or `"$"` all work, as these are inherited. Creating invalid trip objects can be prohibited by adding checks to the validity function in the class definition, e.g., will not work because the time and/or id column are not present any more.

A custom plot method for trip could be written, for example using colour to denote a change in identifier:

```
> setGeneric("lines", function(x, ...) standardGeneric("lines"))
[1] "lines"
> setMethod("lines", signature(x = "trip"), function(x,
+   ..., col = NULL) {
+   tor <- x@TOR.columns
+   if (is.null(col)) {
+     l <- length(unique(x[[tor[2]]]))
+     col <- hsv(seq(0, 0.5, length = l))
+   }
+   coords <- coordinates(x)
+   lx <- split(1:nrow(coords), x[[tor[2]]])
+   for (i in 1:length(lx)) lines(coords[lx[[i]], ],
+     col = col[i], ...)
+ })
```

Here, the `col` argument is added to the function header so that a reasonable default can be overridden, e.g., for black/white plotting.

3 Multi-point data: SpatialMultiPoints

One of the feature types of the OpenGeospatial Consortium (OGC) simple feature specification that has not been implemented in **sp** is the **MultiPoint** object. In a **MultiPoint** object, each feature refers to a *set of* points. The **sp** classes **SpatialPointsDataFrame** only provide reference to a single point. Instead of building a new class up from scratch, we'll try to re-use code and build a class **SpatialMultiPoint** from the **SpatialLines** class. After all, lines are just sets of ordered points.

In fact, the **SpatialLines** class implements the **MultiLineString** simple feature, where each feature can refer to multiple lines. A special case is formed if each feature only has a single line:

```
> setClass("SpatialMultiPoints", representation("SpatialLines"),
+   validity <- function(object) {
+     if (any(unlist(lapply(object@lines,
+       function(x) length(x@Lines))) !=
+       1))
+       stop("Only Lines objects with one Line element")
+     TRUE
+   })
> SpatialMultiPoints <- function(object) new("SpatialMultiPoints",
+   object)
```

As an example, we can create an instance of this class for two **MultiPoint** features each having three locations:

```
> n <- 5
> set.seed(1)
> x1 <- cbind(rnorm(n), rnorm(n, 0, 0.25))
> x2 <- cbind(rnorm(n), rnorm(n, 0, 0.25))
> x3 <- cbind(rnorm(n), rnorm(n, 0, 0.25))
> L1 <- Lines(list(Line(x1)), ID = "mp1")
> L2 <- Lines(list(Line(x2)), ID = "mp2")
> L3 <- Lines(list(Line(x3)), ID = "mp3")
> s <- SpatialLines(list(L1, L2, L3))
> smp <- SpatialMultiPoints(s)
```

If we now plot object **smp**, we get the same plot as when we plot **s**, showing the two lines. The plot method for a **SpatialLines** object is not suitable, so we write a new one:

```
> plot.SpatialMultiPoints <- function(x, ..., pch = 1:length(x@lines),
+   col = 1, cex = 1) {
+   n <- length(x@lines)
+   if (length(pch) < n)
+     pch <- rep(pch, length.out = n)
+   if (length(col) < n)
+     col <- rep(col, length.out = n)
+   if (length(cex) < n)
+     cex <- rep(cex, length.out = n)
+   plot(as(x, "Spatial"), ...)
+   for (i in 1:n) points(x@lines[[i]]@Lines[[1]]@coords,
+     pch = pch[i], col = col[i], cex = cex[i])
+ }
> setMethod("plot", signature(x = "SpatialMultiPoints",
+   y = "missing"), function(x, y, ...) plot.SpatialMultiPoints(x,
+   ...))
```

Here we chose to pass any named ... arguments to the plot method for a `Spatial` object. This function sets up the axes and controls the margins, aspect ratio, etc. All arguments that need to be passed to `points` (`pch` for symbol type, `cex` for symbol size and `col` for symbol colour) need explicit naming and sensible defaults, as they are passed explicitly to the consecutive calls to `points`. According to the documentation of `points`, in addition to `pch`, `cex` and `col`, the arguments `bg` and `lwd` (symbol fill colour and symbol line width) would need a similar treatment to make this plot method completely transparent with the base `plot` method—something an end user would hope for.

Having `pch`, `cex` and `col` arrays the length of the number of `MultiPoints` *sets* rather than the number of points to be plotted is useful for two reasons. First, the whole point of `MultiPoints` object is to distinguish *sets* of points. Second, when we extend this class to `SpatialMultiPointsDataFrame`, e.g., by

```
> cName <- "SpatialMultiPointsDataFrame"
> setClass(cName, representation("SpatialLinesDataFrame"),
+   validity <- function(object) {
+     lst <- lapply(object@lines, function(x) length(x@Lines))
+     if (any(unlist(lst) != 1))
+       stop("Only Lines objects with single Line")
+     TRUE
+   })
> SpatialMultiPointsDataFrame <- function(object) {
+   new("SpatialMultiPointsDataFrame", object)
+ }
```

then we can pass symbol characteristics by (functions of) columns in the attribute table:

```
> df <- data.frame(x1 = 1:3, x2 = c(1, 4, 2), row.names = c("mp1",
+   "mp2", "mp3"))
> smp_df <- SpatialMultiPointsDataFrame(SpatialLinesDataFrame(smp,
+   df))
> setMethod("plot", signature(x = "SpatialMultiPointsDataFrame",
+   y = "missing"), function(x, y, ...) plot.SpatialMultiPoints(x,
+   ...))
> grys <- c("grey10", "grey40", "grey80")
> plot(smp_df, col = grys[smp_df[["x1"]]], pch = smp_df[["x2"]],
+   cex = 2, axes = TRUE)
```

for which the plot is shown in Figure 1.

Hexagonal grids are like square grids, where grid points are centres of matching hexagons, rather than squares. Package `sp` has no classes for hexagonal grids, but does have some useful functions for generating and plotting them. This could be used to build a class. Much of this code in `sp` is based on postings to the R-sig-geo mailing list by Tim Keitt, used with permission.

The spatial sampling method `spsample` has a method for sampling points on a hexagonal grid:

```
> data(meuse.grid)
> gridded(meuse.grid) = ~x + y
> xx <- spsample(meuse.grid, type = "hexagonal", cellsize = 200)
> class(xx)

[1] "SpatialPoints"
attr(,"package")
[1] "sp"
```




Figure 1: Plot of the `SpatialMultiPointsDataFrame` object.

gives the points shown in the left side of Figure 2. Note that an alternative hexagonal representation is obtained by rotating this grid 90 degrees; we will not further consider that here.

```
> HexPts <- spsample(meuse.grid, type = "hexagonal", cellsize = 200)
> spplot(meuse.grid["dist"], sp.layout = list("sp.points",
+       HexPts, col = 1))
> HexPols <- HexPoints2SpatialPolygons(HexPts)
> df <- over(HexPols, meuse.grid)
> HexPolsDf <- SpatialPolygonsDataFrame(HexPols, df, match.ID = FALSE)
> spplot(HexPolsDf["dist"])
```

for which the plots are shown in Figure 2.

We can now generate and plot hexagonal grids, but need to deal with two representations: as points and as polygons, and both representations do not tell by themselves that they represent a hexagonal grid.

For designing a hexagonal grid class we will extend `SpatialPoints`, assuming that computation of the polygons can be done when needed without a prohibitive overhead.

```
> setClass("SpatialHexGrid", representation("SpatialPoints",
+       dx = "numeric"), validity <- function(object) {
+   if (object@dx <= 0)
+     stop("dx should be positive")
+   TRUE
+ })
> setClass("SpatialHexGridDataFrame",
+       representation("SpatialPointsDataFrame",
+       dx = "numeric"), validity <- function(object) {
+   if (object@dx <= 0)
+     stop("dx should be positive")
+   TRUE
+ })
```

Note that these class definitions do not check that instances actually do form valid hexagonal grids; a more robust implementation could provide a test that




Figure 2: Hexagonal points (left) and polygons (right).

distances between points with equal y coordinate are separated by a multiple of dx , that the y -separations are correct and so on.

It might make sense to adapt the generic `spsample` method in package `sp` to return `SpatialHexGrid` objects; we can also add `plot` and `spsample` methods for them. Method `over` should work with a `SpatialHexGrid` as its first argument, by inheriting from `SpatialPoints`. Let us first see how to create the new classes. Without a constructor function we can use

```
> HexPts <- spsample(meuse.grid, type = "hexagonal", cellsize = 200)
> Hex <- new("SpatialHexGrid", HexPts, dx = 200)
> df <- over(Hex, meuse.grid)
> spdf <- SpatialPointsDataFrame(HexPts, df)
> HexDf <- new("SpatialHexGridDataFrame", spdf, dx = 200)
```

Because of the route taken to define both `HexGrid` classes, it is not obvious that the second extends the first. We can tell the S4 system this by `setIs`:

```
> is(HexDf, "SpatialHexGrid")
[1] FALSE

> setIs("SpatialHexGridDataFrame", "SpatialHexGrid")
> is(HexDf, "SpatialHexGrid")
[1] TRUE
```

to make sure that methods for `SpatialHexGrid` objects work as well for objects of class `SpatialHexGridDataFrame`.

When adding methods, several of them will need conversion to the polygon representation, so it makes sense to add the conversion function such that e.g. `as(x, "SpatialPolygons")` will work:

```

> setAs("SpatialHexGrid", "SpatialPolygons",
+       function(from) HexPoints2SpatialPolygons(from,
+       from@dx))
> setAs("SpatialHexGridDataFrame", "SpatialPolygonsDataFrame",
+       function(from) SpatialPolygonsDataFrame(as(obj,
+       "SpatialPolygons"), obj@data,
+       match.ID = FALSE))

```

We can now add `plot`, `splot`, `spsample` and `over` methods for these classes:

```

> setMethod("plot", signature(x = "SpatialHexGrid", y = "missing"),
+       function(x, y, ...) plot(as(x, "SpatialPolygons"),
+       ...))
> setMethod("splot", signature(obj = "SpatialHexGridDataFrame"),
+       function(obj, ...) splot(SpatialPolygonsDataFrame(as(obj,
+       "SpatialPolygons"), obj@data, match.ID = FALSE),
+       ...))
> setMethod("spsample", "SpatialHexGrid", function(x, n,
+       type, ...) spsample(as(x, "SpatialPolygons"), n = n,
+       type = type, ...))
> setMethod("over", c("SpatialHexGrid", "SpatialPoints"),
+       function(x, y, ...) over(as(x, "SpatialPolygons"),
+       y))

```

After this, the following will work:

```

> splot(meuse.grid["dist"], sp.layout = list("sp.points",
+       Hex, col = 1))
> splot(HexDf["dist"])

```

Coercion to a data frame is done by

```

> as(HexDf, "data.frame")

```

Another detail not mentioned is that the bounding box of the hexgrid objects only match the grid centre points, not the hexgrid cells:

```

> bbox(Hex)
      min      max
x 178580.2 181380.2
y 329646.2 333629.9
> bbox(as(Hex, "SpatialPolygons"))
      min      max
x 178480.2 181480.2
y 329530.8 333745.4

```

One solution for this is to correct for this in a constructor function, and check for it in the validity test. Explicit coercion functions to the points representation would have to set the bounding box back to the points ranges. Another solution is to write a `bbox` method for the hexgrid classes, taking the risk that someone still looks at the incorrect `bbox` slot.

4 Spatio-temporal grids

Spatio-temporal data can be represented in different ways. One simple option is when observations (or model-results, or predictions) are given on a regular space-time grid.

Objects of class or extending `SpatialPoints`, `SpatialPixels` and `SpatialGrid` do not have the constraint that they represent a two-dimensional space; they may have arbitrary dimension; an example for a three-dimensional grid is

```
> n <- 10
> x <- data.frame(expand.grid(x1 = 1:n, x2 = 1:n, x3 = 1:n),
+               z = rnorm(n^3))
> coordinates(x) <- ~x1 + x2 + x3
> gridded(x) <- TRUE
> fullgrid(x) <- TRUE
> summary(x)
```

```
Object of class SpatialGridDataFrame
Coordinates:
      min max
x1 0.5 10.5
x2 0.5 10.5
x3 0.5 10.5
Is projected: NA
proj4string : [NA]
Grid attributes:
      cellcentre.offset cellsize cells.dim
x1              1          1          10
x2              1          1          10
x3              1          1          10
Data attributes:
      z
Min.   :-3.00805
1st Qu.: -0.70751
Median :-0.04317
Mean   :-0.02348
3rd Qu.: 0.68843
Max.    : 3.81028
```

We might assume here that the third dimension, `x3`, represents time. If we are happy with time somehow represented by a real number (in double precision), then we are done. A simple representation is that of decimal year, with e.g. 1980.5 meaning the 183rd day of 1980, or e.g. relative time in seconds after the start of some event.

When we want to use the `POSIXct` or `POSIXlt` representations, we need to do some more work to see the readable version. We will now devise a simple three-dimensional space-time grid with the `POSIXct` representation.

```
> setClass("SpatialTimeGrid", "SpatialGrid",
+   validity <- function(object) {
+     stopifnot(dimensions(object) ==
+               3)
+     TRUE
+   })
```

Along the same line, we can extend the `SpatialGridDataFrame` for space-time:

```
> setClass("SpatialTimeGridDataFrame", "SpatialGridDataFrame",
+   validity <- function(object) {
+     stopifnot(dimensions(object) == 3)
+     TRUE
+   })
```

```
> setIs("SpatialTimeGridDataFrame", "SpatialTimeGrid")
> x <- new("SpatialTimeGridDataFrame", x)
```

A crude summary for this class could be written along these lines:

```
> summary.SpatialTimeGridDataFrame <- function(object,
+   ...) {
+   cat("Object of class SpatialTimeGridDataFrame\n")
+   x <- gridparameters(object)
+   t0 <- ISOdate(1970, 1, 1, 0, 0, 0)
+   t1 <- t0 + x[3, 1]
+   cat(paste("first time step:", t1, "\n"))
+   t2 <- t0 + x[3, 1] + (x[3, 3] - 1) * x[3, 2]
+   cat(paste("last time step: ", t2, "\n"))
+   cat(paste("time step:      ", x[3, 2], "\n"))
+   summary(as(object, "SpatialGridDataFrame"))
+ }

> setMethod("summary", "SpatialTimeGridDataFrame",
+   summary.SpatialTimeGridDataFrame)
> summary(x)
```

```
Object of class SpatialTimeGridDataFrame
first time step: 1970-01-01 00:00:01
last time step: 1970-01-01 00:00:10
time step:      1
Object of class SpatialGridDataFrame
Coordinates:
  min max
x1 0.5 10.5
x2 0.5 10.5
x3 0.5 10.5
Is projected: NA
proj4string : [NA]
Grid attributes:
  cellcentre.offset cellsize cells.dim
x1              1          1         10
x2              1          1         10
x3              1          1         10
Data attributes:
  z
Min.  :-3.00805
1st Qu.: -0.70751
Median :-0.04317
Mean   :-0.02348
3rd Qu.: 0.68843
Max.   : 3.81028
```

Next, suppose we need a subsetting method that selects on the time. When the first subset argument is allowed to be a time range, this is done by

```
> subs.SpatialTimeGridDataFrame <- function(x, i, j, ...,
+   drop = FALSE) {
+   t <- coordinates(x)[, 3] + ISOdate(1970, 1, 1, 0,
+   0, 0)
+   if (missing(j))
+     j <- TRUE
+   sel <- t %in% i
+   if (!any(sel))
+     stop("selection results in empty set")
+   fullgrid(x) <- FALSE
```

```

+   if (length(i) > 1) {
+     x <- x[i = sel, j = j, ...]
+     fullgrid(x) <- TRUE
+     as(x, "SpatialTimeGridDataFrame")
+   }
+   else {
+     gridded(x) <- FALSE
+     x <- x[i = sel, j = j, ...]
+     cc <- coordinates(x)[, 1:2]
+     p4s <- CRS(proj4string(x))
+     SpatialPixelsDataFrame(cc, x@data, proj4string = p4s)
+   }
+ }
> setMethod("[", c("SpatialTimeGridDataFrame", "POSIXct",
+   "ANY"), subs.SpatialTimeGridDataFrame)
> t1 <- as.POSIXct("1970-01-01 0:00:03", tz = "GMT")
> t2 <- as.POSIXct("1970-01-01 0:00:05", tz = "GMT")
> summary(x[c(t1, t2)])

```

```

Object of class SpatialTimeGridDataFrame
first time step: 1970-01-01 00:00:01
last time step: 1970-01-01 00:00:10
time step: 1
Object of class SpatialGridDataFrame
Coordinates:
  min max
x1 0.5 10.5
x2 0.5 10.5
x3 0.5 10.5
Is projected: NA
proj4string : [NA]
Grid attributes:
  cellcentre.offset cellsize cells.dim
x1              1         1         10
x2              1         1         10
x3              1         1         10
Data attributes:
  z
Min.   :-3.008049
1st Qu.: -0.649691
Median :  0.008089
Mean   :  0.018708
3rd Qu.:  0.679773
Max.   :  3.810277
NAs    :800

```

```
> summary(x[t1])
```

```

Object of class SpatialPixelsDataFrame
Coordinates:
  min max
x1 0.5 10.5
x2 0.5 10.5
Is projected: NA
proj4string : [NA]
Number of points: 100
Grid attributes:
  cellcentre.offset cellsize cells.dim
x1              1         1         10

```

```

x2          1          1          10
Data attributes:
  z
Min.   :-2.40310
1st Qu.: -0.42295
Median :  0.13533
Mean   :  0.09981
3rd Qu.:  0.77787
Max.   :  2.64917

```

The reason to only convert back to `SpatialTimeGridDataFrame` when multiple time steps are present is that the time step (“cell size” in time direction) cannot be found when there is only a single step. In that case, the current selection method returns an object of class `SpatialPixelsDataFrame` for that time slice.

Plotting a set of slices could be done using `levelplot`, or writing another `spplot` method:

```

> spplot.stgdf <- function(obj, zcol = 1, ..., format = NULL) {
+   if (length(zcol) != 1)
+     stop("can only plot a single attribute")
+   if (is.null(format))
+     format <- "%Y-%m-%d %H:%M:%S"
+   cc <- coordinates(obj)
+   df <- unstack(data.frame(obj[[zcol]], cc[, 3]))
+   ns <- format(coordinatevalues(getGridTopology(obj))[[3]] +
+     ISOdate(1970, 1, 1, 0, 0, 0), format = format)
+   cc2d <- cc[cc[, 3] == min(cc[, 3]), 1:2]
+   obj <- SpatialPixelsDataFrame(cc2d, df)
+   spplot(obj, names.attr = ns, ...)
+ }
> setMethod("spplot", "SpatialTimeGridDataFrame", spplot.stgdf)

```

Now, the result of

```

> library(lattice)
> trellis.par.set(canonical.theme(color = FALSE))
> spplot(x, format = "%H:%M:%S", as.table = TRUE, cuts = 6,
+   col.regions = grey.colors(7, 0.55, 0.95, 2.2))

```

is shown in Figure 3. The `format` argument passed controls the way time is printed; one can refer to the help of

```
> `?`(as.character.POSIXt)
```

for more details about the `format` argument.

5 Analysing spatial Monte Carlo simulations

Quite often, spatial statistical analysis results in a large number of spatial realisations or a random field, using some Monte Carlo simulation approach. Regardless whether individual values refer to points, lines, polygons or grid cells, we would like to write some methods or functions that aggregate over these simulations, to get summary statistics such as the mean value, quantiles, or cumulative distributions values. Such aggregation can take place in two ways. Either we aggregate over the probability space, and compute summary statistics for each geographical feature over the set of realisations (i.e., the rows of




Figure 3: `spplot` for an object of class `SpatialTimeGridDataFrame`, filled with random numbers.

the attribute table), or for each realisation we aggregate over the complete geographical layer or a subset of it (i.e., aggregate over the columns of the attribute table).

Let us first generate, as an example, a set of 100 conditional Gaussian simulations for the zinc variable in the meuse data set:

```
> data(meuse)
> coordinates(meuse) <- ~x + y
> if (require(gstat, quietly = TRUE)) {
+   v <- vgm(0.5, "Sph", 800, 0.05)
+   sim <- krige(log(zinc) ~ 1, meuse, meuse.grid, v,
+     nsim = 100, nmax = 30)
+   sim@data <- exp(sim@data)
+ }
```

drawing 100 GLS realisations of beta...
[using conditional Gaussian simulation]

where the last statement back-transforms the simulations from the log scale to the observation scale. A quantile method for Spatial object attributes can be written as

```
> quantile.Spatial <- function(x, ..., byLayer = FALSE) {
+   stopifnot("data" %in% slotNames(x))
+   apply(x@data, ifelse(byLayer, 2, 1), quantile, ...)
+ }
```

after which we can find the sample lower and upper 95% confidence limits by

```
> if (require(gstat, quietly = TRUE)) {
+   sim$lower <- quantile.Spatial(sim[1:100], probs = 0.025)
```

```
+   sim$supper <- quantile.Spatial(sim[1:100], probs = 0.975)
+ }
```

To get the sample distribution of the areal median, we can aggregate over layers:

```
> if (require(gstat, quietly = TRUE)) {
+   medians <- quantile.Spatial(sim[1:100], probs = 0.5,
+     byLayer = TRUE)
+ }
> hist(medians)
```

It should be noted that in these particular cases, the quantities computed by simulations could have been obtained faster and exact by working analytically with ordinary (block) kriging and the normal distribution (Section 8.7.2 in [Bivand et al. \(2008\)](#)).

A statistic that cannot be obtained analytically is the sample distribution of the area fraction that exceeds a threshold. Suppose that 500 is a crucial threshold, and we want to summarise the sampling distribution of the area fraction where 500 is exceeded:

```
> fractionBelow <- function(x, q, byLayer = FALSE) {
+   stopifnot(is(x, "Spatial") || !("data" %in%
+     slotNames(x)))
+   apply(x@data < q, ifelse(byLayer,
+     2, 1), function(r) sum(r)/length(r))
+ }
> if (require(gstat, quietly = TRUE)) {
+   over500 <- 1 - fractionBelow(sim[1:100], 200, byLayer = TRUE)
+   summary(over500)
+   quantile(over500, c(0.025, 0.975))
+ }
      2.5%      97.5%
0.6398244 0.7381888
```

For space-time data, we could write methods that aggregate over space, over time or over space and time.

This chapter has sketched developments beyond the base **sp** classes and methods used otherwise in this book. Although we think that the base classes cater for many standard kinds of spatial data analysis, it is clear that specific research problems will call for specific solutions, and that the R environment provides the high-level abstractions needed to help busy researchers get their work done.

References

- Roger S. Bivand, Edzer J. Pebesma and Virgilio Gomez-Rubio (2008). *Applied spatial data analysis with R* Springer, NY
- Braun, W. J. and Murdoch, D. J. (2007). *A first course in statistical programming with R*. Cambridge University Press, Cambridge.
- Chambers, J.M. (1998). *Programming with Data*. Springer, New York.

- Kirkwood, R., Lynch, M., Gales, N., Dann, P., and Sumner, M. (2006). At-sea movements and habitat use of adult male Australian fur seals (*Arctocephalus pusillus doriferus*). *Canadian Journal of Zoology*, 84:1781–1788.
- Page, B., McKenzie, J., Sumner, M., Coyne, M., and Goldsworthy, S. (2006). Spatial separation of foraging habitats among New Zealand fur seals. *Marine Ecology Progress Series*, 323:263–279.
- Venables, W. N. and Ripley, B. D. (2000). *S Programming*. Springer, New York.