

---

**Technical Note:**  
**Co-kriging with the gstat package**  
**of the R environment for statistical computing**

---

*D G Rossiter*  
*Cornell University, Section of Soil & Crop Sciences*  
南京师范大学地理学学院

July 14, 2018

**Contents**

<b>1 Introduction</b>	<b>1</b>
1.1 When to use co-kriging? . . . . .	1
1.2 R and gstat . . . . .	2
1.3 Example Data Set . . . . .	2
<b>2 Target and co-variables</b>	<b>3</b>
2.1 The target variable . . . . .	3
2.2 Selecting the co-variables . . . . .	5
<b>3 Simulation of under-sampling</b>	<b>6</b>
<b>4 Modelling the spatial structure of the target variable</b>	<b>8</b>
<b>5 Ordinary Kriging interpolation of the target variable</b>	<b>13</b>
5.1 Evaluation . . . . .	15
5.2 Cross-validation . . . . .	17
5.3 Evaluating errors in absolute terms . . . . .	19
<b>6 Modelling a bivariate co-regionalisation</b>	<b>20</b>
6.1 Modelling the co-variable . . . . .	21
6.2 Building a data structure to model co-regionalisation . . . . .	23
6.3 Fitting a linear model of co-regionalisation . . . . .	26

---

Version 3.0 Copyright © 2005–2018 D G Rossiter. All rights reserved. Reproduction and dissemination of the work as a whole (not parts) freely permitted if this original copyright notice is included. Sale or placement on a web site where payment must be made to access this document is strictly prohibited. To adapt or translate please contact the author ([d.g.rossiter@cornell.edu](mailto:d.g.rossiter@cornell.edu)).

6.4	Comparing models of regionalisation and co-regionalisation	28
<b>7</b>	<b>Co-kriging with one co-variable</b>	<b>31</b>
7.1	Evaluation	32
7.2	Cross-validation	33
<b>8</b>	<b>Comparison of co-kriging and ordinary kriging</b>	<b>35</b>
8.1	Predictions	35
8.2	Prediction errors	36
8.3	Validations	38
8.4	Cross-validations	39
<b>9</b>	<b>Co-kriging with a different co-variable</b>	<b>40</b>
9.1	Modelling the co-variable Zn	40
9.2	Building a data structure to model co-regionalisation of Pb with Zn	43
9.3	Fitting a linear model of co-regionalisation for Zn and Pb	44
9.4	Co-kriging Pb with Zn as a co-variable	46
9.5	Validation	47
9.6	Cross-validation	48
9.7	Comparison of co-kriging Pb with Zn and ordinary kriging of Pb	50
9.7.1	Predictions	51
9.7.2	Prediction variances	52
9.7.3	Validations	52
9.7.4	Validations	54
9.7.5	Cross-validations	55
<b>10</b>	<b>Comparison of the three interpolations</b>	<b>56</b>
10.1	Predictions	56
10.2	Prediction errors	57
10.3	Evaluations	58
10.4	Cross-validations	64
<b>11</b>	<b>Conclusion</b>	<b>64</b>
	<b>References</b>	<b>67</b>
<b>A</b>	<b>Producing the figures for this document</b>	<b>67</b>
A.1	Plotting functions	68

## 1 Introduction

This technical note shows how to perform co-kriging using the `gstat` geostatistical package [12] of the R environment for statistical computing and visualisation [3]. It does not present the theory of co-regionalisation or co-kriging; this is dealt with in many texts [e.g. 2, 4, 9, 19]. The main aim is to show various R techniques for data manipulation, calculation and graphics; emphasis is on the `gstat` geostatistical package, the `sp` spatial data package [11], and the `lattice` graphics package [15].

**Note:** Since the first version of this tutorial in 2005, the `ggplot2` ‘Grammar of Graphics’ package was developed, along with many packages that depend on it. These provide other ways to produce statistical graphics and maps. Probably these would be more attractive, but I do not have time to re-do the tutorial.

It also shows how to evaluate the success of a kriging interpolation by evaluation (“validation”) from an independent data set and by cross-validation from the sample data set.

These notes are designed so that you can *cut* the code shown and *paste* it directly into the R console; you should then see the output of executing the code, the same as you see in the document. Of course, you are encouraged to edit the code and experiment.

### 1.1 When to use co-kriging?

Co-kriging allows samples of an *auxiliary* variable (also called the *co-variable*), besides the *target* value of interest, to be used when predicting the target value at unsampled locations. The co-variable may be measured at the same points as the target (*co-located* samples), at other points, or both.

A common application of co-kriging is when the co-variable is cheaper to measure, and so has been more densely sampled, than the target variable. In this example we show this typical situation, where there is a small sample set where both variables have been measured (co-located measurements) and a larger sample set where only the co-variable has been measured. This is the application we will examine in this tutorial.

**Note:** Another application is when several correlated variables are measured at the same locations, and using the spatial and feature-space correlation can lead to more precise predictions for all of the variables.

Co-kriging requires that both target and co-variable be, individually, spatially autocorrelated, and in addition that they be spatially cross-correlated.

If the target and co-variable samples are co-located, and in addition the values of the co-variable are known at each prediction location (e.g. an interpolation grid), more appropriate techniques are Kriging with External Drift (KED) or Regression Kriging (RK); these only require a significant feature-space relation between target and co-variable.

## 1.2 R and gstat

R is set up as a *base system* with many *contributed libraries*, among which are several for geostatistics, including `spatial` [17], `geoR` [13], and `gstat` [12]. The latter, `gstat`, is the subject of these notes. Each package has its own way of approaching geostatistics. This is one of the strengths of R: it is open to any practising statisticians who can write code to express their computational methods. Recently, several of these practitioners have together developed a common package `sp` to facilitate spatial statistics [11]; `gstat` depends on this package and loads it as necessary.

As with all R packages, `gstat` has on-line help: `library(help=gstat)` to get a list of methods in the package, `?variogram` to get help for the `variogram` method, and similarly for the other methods. The user's manual [10] refers to the stand-alone version, not the R package, but explains the algorithms in detail.

The `gstat` library uses structures from the `sp` “classes and methods for spatial data” package.

## 1.3 Example Data Set

This example uses the Meuse soil pollution dataset [14] and an interpolation grid of the study area in the south of the Netherlands, both provided with the `sp` R library as sample data sets. A brief explanation of the data fields is given in the on-line help (`?meuse`, `?meuse.grid`) once this library is loaded.

---

**TASK 1** : Load the `sp` and `gstat` libraries, the dataset of sample points, and the interpolation grid. Examine the structure of the point set and grid. •

```
library(sp)
library(gstat)
?meuse
data(meuse)
str(meuse)

'data.frame': 155 obs. of 14 variables:
 $ x      : num  181072 181025 181165 181298 181307 ...
 $ y      : num  333611 333558 333537 333484 333330 ...
 $ cadmium: num   11.7  8.6  6.5  2.6  2.8  3  3.2  2.8  2.4  1.6 ...
 $ copper  : num   85  81  68  81  48  61  31  29  37  24 ...
 $ lead   : num  299 277 199 116 117 137 132 150 133 80 ...
 $ zinc   : num  1022 1141 640 257 269 ...
 $ elev   : num   7.91 6.98 7.8 7.66 7.48 ...
 $ dist   : num  0.00136 0.01222 0.10303 0.19009 0.27709 ...
 $ om     : num   13.6 14 13 8 8.7 7.8 9.2 9.5 10.6 6.3 ...
 $ ffreq  : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 1 ...
 $ soil   : Factor w/ 3 levels "1","2","3": 1 1 1 2 2 2 2 1 1 2 ...
 $ lime   : Factor w/ 2 levels "0","1": 2 2 2 1 1 1 1 1 1 1 ...
 $ landuse: Factor w/ 15 levels "Aa","Ab","Ag",...: 4 4 4 11 4 11 4 2 2 15 ...
 $ dist.m : num   50 30 150 270 380 470 240 120 240 420 ...

?meuse.grid
data(meuse.grid)
```

```
str(meuse.grid)

'data.frame': 3103 obs. of 7 variables:
 $ x      : num  181180 181140 181180 181220 181100 ...
 $ y      : num  333740 333700 333700 333700 333660 ...
 $ part.a : num  1 1 1 1 1 1 1 1 1 ...
 $ part.b : num  0 0 0 0 0 0 0 0 0 ...
 $ dist   : num  0 0 0.0122 0.0435 0 ...
 $ soil   : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 ...
 $ ffreq  : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 ...
```

The data frame of sample points contains 155 observations of twelve variables<sup>1</sup>, each with its coördinates in the Dutch national grid. The prediction grid has 3103 locations, spaced every 40 m in both the E and N grid directions, covering the irregularly-shaped study area. The support (geographic area represented by the soil samples)  $(10 \text{ m})^2 \cdot \pi \approx 314 \text{ m}^2$  [14, §2.2] horizontal, 10 cm vertical. Punctual kriging is thus a prediction of the average concentration over this area, at the centre of each grid cell, averaged over 10 cm depth slice. The ‘point’ support is about 20% of the area of each cell of the interpolation grid, 1600m<sup>2</sup>.

## 2 Target and co-variables

In a mapping exercise we are often interested in one or more *target* variables, i.e. those to be mapped. When co-kriging we also need to select one or more *co-variables* to assist in the mapping of the target variables.

When mapping with co-kriging it is possible to use the same set of variables as target and co-variables; however this is typically used in the case of more or less equal sampling density. In these notes we consider the case of under-sampling of a target variable, using additional samples from co-variables.

### 2.1 The target variable

We select lead (chemical symbol Pb) as the *target variable*, i.e. the one we want to map. This metal is a serious human health hazard. It can be inhaled as dust from disturbed soil or taken up by plants and ingested. The critical value for Pb in agricultural soils, according to the Berlin Digital Environmental Atlas<sup>2</sup>, is 600 mg kg<sup>-1</sup> for agricultural fields: above this level grain crops can not be grown for human consumption. Above 300 mg kg<sup>-1</sup> crops must be tested; above 100 mg kg<sup>-1</sup> consumption of leafy vegetables should be avoided. Levels below 200 mg kg<sup>-1</sup> are required for sports fields or parks where soil may become bare from overuse. Natural levels in unpolluted soils are on the order of 30 mg kg<sup>-1</sup>.

---

**TASK 2 :** Display histograms of the target variable and its log<sub>10</sub> trans-


<sup>1</sup> although there are a few missing values of organic matter

<sup>2</sup> <http://www.stadtentwicklung.berlin.de/umwelt/umweltatlas/ed103103.htm>

form. What proportion of the samples are above the five thresholds?

We will use the `lattice` graphics package throughout this case study; this is an R implementation by Sarkar [15] of the Trellis framework for data visualization developed at Bell Labs, based on the ideas in Cleveland [1]. It must be loaded before use. The `histogram` method is part of this package, not the base `graphics` package<sup>3</sup>.

```
require("lattice")
h1 <- histogram(~ lead, meuse, xlab="Pb", col="thistle3", nint=12)
h2 <- histogram(~ log10(lead), meuse, xlab="log10(Pb)", col="thistle3", nint=12)
print(h1, split = c(1,1,2,1), more=T)
print(h2, split = c(2,1,2,1), more=F)
rm(h1, h2)
```



```
# proportions higher than various thresholds
ph <- function(level) {
  round(100 * sum(meuse$lead > level)/length(meuse$lead), 1)
}
p <- NULL; lvls <- c(600, 300, 200, 100, 30)
for (l in lvls) p <- c(p, ph(l))
# display a table of thresholds and proportions
(data.frame(cbind(level=lvls, percent.higher=p)))

  level percent.higher
1    600             0.6
2    300             8.4
3    200            27.1
4    100            54.8
5     30           100.0

rm(ph, l, lvls, p)
```

The `nint=` optional parameter is used to specify the number of histogram bins; if it is omitted, `lattice` will compute it as  $(\log_2 n) + 1$ , in this case 8.

Not many samples are too polluted for grain crops, but over a quarter are too polluted for use as playgrounds and half are too polluted for garden allotments. All are above background levels.

Because of the right-skew and wide numerical range of the Pb values we will work with the log-transformed target variable; to allow easy interpretation of the results we will use base-10 logarithms.

<sup>3</sup>The `hist` method in the base `graphics` package gives similar output.

## 2.2 Selecting the co-variables

A *co-variable* for use in co-kriging is correlated to the target variable in both feature and geographic space, and has usually been more densely sampled than the target variable.

Candidates for co-variables must have:

1. a *feature-space correlation* with the target variable;
2. a *spatial structure* (i.e. be modelled as a regional variable);
3. a *spatial co-variance* with the target variable.

There are two main ways to select a co-variable:

1. *theoretically*, from knowledge of the spatial process that caused the observed spatial (co-)distribution;
2. *empirically*, by examining the feature-space correlations (scatter-plots) and then the spatial co-variance (cross-correlograms or cross-variograms).

We usually have some idea of the possible co-variables from theory (i.e. an *a priori* hypothesis) based on our knowledge of the process that (we think) gave rise to the data. The modelling process reveals if our hypothesis is supported by this particular data set. The second option, exploring the data (a sort of “data mining”) is appropriate when there is no theory.

In this tutorial we will compare two possible co-variables based on theory, which turn out to have different strengths of relation with the target variable.

First, we select **organic matter content** (abbreviation “OM”). It is cheaper than the heavy metal to measure in the laboratory, and may bind heavy metals in the soil [6, 16]; thus there may be a positive correlation in feature space between this co-variable and target variable. However, it is mostly formed in place in the soil and not transported with flood waters, as is Pb, so it is unclear if there will prove to be a spatial co-variance with Pb in the study area.

Second, we select **zinc content** (chemical symbol Zn). Many of the industrial processes that result in lead pollution also produce Zn, and it is transported similarly, both as dust and sediment. It is generally measured along with lead at no additional cost (i.e. the same soil sample serves for both) so this example is somewhat artificial, but it will illustrate the effect of a highly-correlated (in both feature and geographic spaces) co-variable.


---

**TASK 3** : Display histograms of the co-variables and their log10 transforms. •

```

h1 <- histogram(~ om, meuse, xlab="OM", col="lightblue", nint=12)
h2 <- histogram(~ zinc, meuse, xlab="Zn", col="red4", nint=12)
h3 <- histogram(~ log10(om), meuse, xlab="log10(OM)", col="lightblue", nint=12)
h4 <- histogram(~ log10(zinc), meuse, xlab="log10(Zn)", col="red4", nint=12)
print(h1, split = c(1,1,2,2), more=T)
print(h3, split = c(2,1,2,2), more=T)
print(h2, split = c(1,2,2,2), more=T)
print(h4, split = c(2,2,2,2), more=F)
rm(h1, h2, h3, h4)

```



The co-variable OM is symmetric with a fairly narrow range, and does not need to be transformed for feature-space correlation; however we will work with the transformed value to have comparable numbers for modelling co-regionalisation (§6). The co-variable Zn is skewed much like Pb, and so is log-transformed.

### 3 Simulation of under-sampling

To explore the value of co-kriging, we will simulate the situation where the target variable is under-sampled compared to the co-variable.

---

**TASK 4:** Make a subset of the observations of the target variable (lead) and the co-variables (organic matter and zinc), using every third sample from the full data set. •

```

meuse.pb <- meuse[ seq(1, length(meuse$lead), by=3),
                  c("x", "y", "lead", "om", "zinc")]
str(meuse.pb)

'data.frame': 52 obs. of 5 variables:
 $ x : num 181072 181298 181165 181232 180874 ...
 $ y : num 333611 333484 333370 333168 333339 ...
 $ lead: num 299 116 132 80 285 240 207 123 75 67 ...
 $ om : num 13.6 8 9.2 6.3 15.4 16.2 13.7 7.2 6.9 4.4 ...
 $ zinc: num 1022 257 346 183 1096 ...

```



Note the use of the `seq` method to select every third row, and the `c` method to select only the relevant columns. Another possibility is the `sample` method to select a random sample of rows; however each sample is different, so you would not be able to compare your results with mine.

We can see which observations were selected:

```
rownames(meuse.pb)

 [1] "1"  "4"  "7"  "10" "13" "16" "19" "22" "25" "28" "31"
 [12] "34" "38" "41" "44" "47" "50" "53" "56" "59" "62" "65"
 [23] "69" "79" "82" "85" "88" "123" "70" "92" "95" "98" "101"
 [34] "104" "108" "111" "114" "117" "120" "124" "127" "130" "133" "136"
 [45] "137" "141" "144" "147" "150" "153" "156" "159"
```

Since we will work with the  $\log_{10}$  transformed variables, for convenience we compute these once and add them to the subsample data frame.

---

**TASK 5 :** Add fields with the  $\log_{10}$  transformed target and co-variables to the data frame of the subsample. •

We name the new fields `ltpb` (stands for “log-ten-lead”), `ltom` (stands for “log-ten-organic-matter”), and `ltzn`:

```
meuse.pb <- cbind(meuse.pb,
                  ltpb = log10(meuse.pb$lead),
                  ltom = log10(meuse.pb$om),
                  ltzn = log10(meuse.pb$zinc))

str(meuse.pb)

'data.frame': 52 obs. of 8 variables:
 $ x : num 181072 181298 181165 181232 180874 ...
 $ y : num 333611 333484 333370 333168 333339 ...
 $ lead: num 299 116 132 80 285 240 207 123 75 67 ...
 $ om : num 13.6 8 9.2 6.3 15.4 16.2 13.7 7.2 6.9 4.4 ...
 $ zinc: num 1022 257 346 183 1096 ...
 $ ltpb: num 2.48 2.06 2.12 1.9 2.45 ...
 $ ltom: num 1.134 0.903 0.964 0.799 1.188 ...
 $ ltzn: num 3.01 2.41 2.54 2.26 3.04 ...
```

The `cbind` method is used to add columns to a data frame. The column names may be specified on the left side of a `=` sign for each added column; otherwise they are named by the full expression, e.g. `log10(meuse.pb$zinc)` instead of `ltzn`.

We have a set of 103 points at which lead was measured, but which we didn’t use in the subset, either for modelling or interpolation. This can be used as a *evaluation dataset*<sup>4</sup> to assess the performance of the interpolation; see §5.1. Here we set up the evaluation data set.

---

**TASK 6 :** Make a data frame of the lead values at the extra points that were not included in the subsample and compare the descriptive statistics of the three sets: sample, extra, and full. •

<sup>4</sup> often called the “validation” dataset, but see [7, 8]

Since we already have the subsample in frame `meuse.pb`, we can extract the samples we didn't choose for the subsample by comparing the sample with the full sample, using the `setdiff` method on the row names.

```
meuse.extra <- meuse[setdiff(rownames(meuse), rownames(meuse.pb)),
  c("x", "y", "lead")]
meuse.extra <- cbind(meuse.extra, ltpb = log10(meuse.extra$lead))
str(meuse.extra)

'data.frame': 103 obs. of 4 variables:
 $ x : num 181025 181165 181307 181390 181027 ...
 $ y : num 333558 333537 333330 333260 333363 ...
 $ lead: num 277 199 117 137 150 133 86 97 183 130 ...
 $ ltpb: num 2.44 2.3 2.07 2.14 2.18 ...

summary(log10(meuse$lead)); sd(log10(meuse$lead))

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
1.568  1.860   2.090   2.088   2.316   2.816
[1] 0.2894356

summary(meuse.pb$ltpb); sd(meuse.pb$ltpb)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
1.681  1.831   2.139   2.117   2.386   2.683
[1] 0.3080808

summary(meuse.extra$ltpb); sd(meuse.extra$ltpb)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
1.568  1.863   2.068   2.073   2.293   2.816
[1] 0.2799144
```

The subsample has very similar statistics to the full sample. As expected, the range is a bit narrower but the standard deviation a bit higher; by chance the median and mean are a bit larger. The evaluation set has almost identical statistics to the full set.

#### 4 Modelling the spatial structure of the target variable

We now begin *spatial* analysis. The first step is to convert the data sets to explicitly-spatial, using the `sp` package.

As loaded from the `sp` package, the sample data `meuse` is a `data.frame`. Notice that the coordinates of each point are listed as fields. However, `sp` has defined some classes to make the spatial nature of the data explicit. It also provides a `coordinates` method to set the spatial coordinates and thereby create explicit spatial data.

---

**TASK 7:** Convert the full and subset data frames, and the interpolation grid, to `sp` classes. •

```
class(meuse)

[1] "data.frame"

coordinates(meuse) <- ~ x + y
# alternate command format: coordinates(meuse) <- c("x", "y")
coordinates(meuse.pb) <- ~ x + y
coordinates(meuse.extra) <- ~ x + y
coordinates(meuse.grid) <- ~ x + y
```

```

class(meuse)

[1] "SpatialPointsDataFrame"
attr(,"package")
[1] "sp"

```

The `coordinates` method is used to inform `sp` about the spatial objects. The objects are now spatially-explicit; so for example they have a known bounding box, projection (here not defined), and attributes:

```

summary(meuse.pb)

Object of class SpatialPointsDataFrame
Coordinates:
  min      max
x 178786 181298
y 329822 333611
Is projected: NA
proj4string : [NA]
Number of points: 52
Data attributes:
  lead      om      zinc      ltpb
Min.   : 48.00  Min.   : 1.900  Min.   : 117.0  Min.   :1.681
1st Qu.: 67.75  1st Qu.: 5.450  1st Qu.: 190.0  1st Qu.:1.831
Median :138.00  Median : 7.000  Median : 401.0  Median :2.139
Mean   :167.12  Mean   : 7.782  Mean   : 516.6  Mean   :2.117
3rd Qu.:243.25  3rd Qu.: 9.150  3rd Qu.: 737.8  3rd Qu.:2.386
Max.   :482.00  Max.   :17.000  Max.   :1672.0  Max.   :2.683
      NA's :1
  ltom      ltzn
Min.   :0.2788  Min.   :2.068
1st Qu.:0.7362  1st Qu.:2.279
Median :0.8451  Median :2.603
Mean   :0.8470  Mean   :2.586
3rd Qu.:0.9614  3rd Qu.:2.868
Max.   :1.2304  Max.   :3.223
NA's   :1

```

Many R methods are not yet aware of `sp` classes; for these the original data frame can be recovered when necessary with the `as.data.frame` method:

```

str(as.data.frame(meuse))

'data.frame': 155 obs. of  14 variables:
 $ x      : num  181072 181025 181165 181298 181307 ...
 $ y      : num  333611 333558 333537 333484 333330 ...
 $ cadmium: num  11.7 8.6 6.5 2.6 2.8 3 3.2 2.8 2.4 1.6 ...
 $ copper  : num  85 81 68 81 48 61 31 29 37 24 ...
 $ lead   : num  299 277 199 116 117 137 132 150 133 80 ...
 $ zinc   : num  1022 1141 640 257 269 ...
 $ elev   : num  7.91 6.98 7.8 7.66 7.48 ...
 $ dist   : num  0.00136 0.01222 0.10303 0.19009 0.27709 ...
 $ om     : num  13.6 14 13 8 8.7 7.8 9.2 9.5 10.6 6.3 ...
 $ ffreq  : Factor w/ 3 levels "1","2","3": 1 1 1 1 1 1 1 1 1 1 ...
 $ soil   : Factor w/ 3 levels "1","2","3": 1 1 1 2 2 2 2 1 1 2 ...
 $ lime   : Factor w/ 2 levels "0","1": 2 2 2 1 1 1 1 1 1 1 ...
 $ landuse: Factor w/ 15 levels "Aa","Ab","Ag",...: 4 4 4 11 4 11 4 2 2 15 ...
 $ dist.m : num  50 30 150 270 380 470 240 120 240 420 ...

```

**TASK 8:** Display a postplot of the subsample superimposed on the full sample and compare their geographic distribution. •

```


xyplot(y ~ x, as.data.frame(meuse), asp="iso",
  panel = function(x, ...) {

```

```

panel.points(coordinates(meuse),
             cex=1.8*(log10(meuse$lead) - 1.3),
             pch=1, col="blue");
panel.points(coordinates(meuse.pb),
             cex=1.8*(meuse.pb$tpb - 1.3),
             pch=20, col="red");
panel.grid(h=-1, v=-1, col="darkgrey")
})

```



**Note:** The `xylplot` method is part of the `lattice` graphics package. It allows the use of *panel functions* to superimpose different kinds of plots, in this case two point sets and the grid lines on top of the main x-y (scatter) plot. Each `panel.*` method can have optional graphics parameters; in this case we used `cex=` to specify maximum point size, `pch` to specify the printing character shape, and `col` to specify the character colour. Colours may be given by an RGB code or by name, see `?colors` and `?col2rgb`. The `aspect="iso"` argument scales the map so that the same distance in the N-S and E-W directions is represented by the same distance on the map.

The subsample has a similar but sparser geographic distribution to the full sample.

The first step in any kriging prediction is to model the regionalised target variable. In this case the sample set is small for variogram modelling; it is certainly too small to model anisotropy.

---

**TASK 9:** Compute the omnidirectional spatial variogram of  $\log_{10}\text{Pb}$  and fit a variogram model to it. •

`Gstat` provides methods for each of these steps:


1. The `variogram` method computes the experimental variogram as a variogram object;
2. The `vgm` method creates a variogram model object;
3. The `fit.variogram` method adjusts a variogram model object to a variogram object.

First we examine the variogram cloud to see all the point-pairs; then we compute the experimental variogram. Then we specify a variogram model and its initial parameters (determined by eye looking at the experimental variogram); finally we fit it.

The `variogram` method has optional arguments to specify maximum distance, bin widths, directions, directional tolerances, and whether the robust Cressie estimator should be used. In this case we use the default estimator but specify the bin widths and cutoff.


Plot the variogram cloud:

```
plot(v.ltpb.c <- variogram(ltpb ~ 1, data=meuse.pb, cutoff=1800, cloud=T))
```




Plot the empirical variogram, also showing the number of point-pairs in each bin:

```
plot(v.ltpb <- variogram(ltpb ~ 1, data=meuse.pb, cutoff=1800, width=200), pl=T)
```



Estimate the variogram model form and parameters by eye, visualize on the empirical variogram:

```
m.ltpb <- vgm(0.08, "Sph", 800, 0.03)
plot(v.ltpb, pl=T, model=m.ltpb)
```




Fit the model parameters by weighted least-squares and plot on the empirical variogram:

```
(m.ltpb.f <- fit.variogram(v.ltpb, m.ltpb))

  model      psill    range
1  Nug 0.02395032  0.0000
2   Sph 0.08879807 876.1765

plot(v.ltpb, pl=T, model=m.ltpb.f)
rm(v.ltpb.c, v.ltpb, m.ltpb)
```



The initial variogram parameters (Spherical form, total sill 1.1, range 800, nugget 0.03) are estimated by eye. The form is suggested not only by this experimental variogram but also by previous experience with soil chemical values. The automatic fit did not change the initial parameter estimates much. For a small sample, the variogram has a good structure, and was successfully modelled.

## 5 Ordinary Kriging interpolation of the target variable

First, we predict Pb without the co-variables, from the small subset. This simulates the situation where we have a small sample of the expensive-to-measure target variable.

---

**TASK 10 :** Use the fitted model of regionalization to interpolate the  $\log_{10}\text{Pb}$  content with OK on the prediction grid. Summarize the prediction and error; we will use this as a baseline to evaluate the benefits of co-kriging. •

Gstat provides the `krige` method for univariate kriging. We must supply a *formula* for the spatial dependence (here, `ltpb ~ 1` for a stationary field), the object containing the sample data and from which the coordinates can be extracted (here, `meuse.pb`), an object with locations to be predicted (here, `meuse.grid`) and named coordinates as in the sample data object, and a model of spatial dependence (variogram) (here, `m.ltpb.f`):

```
# interpolate
k.o <- krige(ltpb ~1, locations=meuse.pb, newdata=meuse.grid, model=m.ltpb.f)

[using ordinary kriging]

# summary statistics
summary(k.o)

Object of class SpatialPointsDataFrame
Coordinates:
      min      max
x 178460 181540
y 329620 333740
```

```
Is projected: NA
proj4string : [NA]
Number of points: 3103
Data attributes:
  var1.pred      var1.var
Min.   :1.707    Min.   :0.03569
1st Qu.:1.897    1st Qu.:0.05195
Median :2.027    Median :0.06078
Mean   :2.060    Mean   :0.06380
3rd Qu.:2.231    3rd Qu.:0.07379
Max.   :2.530    Max.   :0.11188
```

The kriged object has both the predictions (field `var1.pred`) and variances of the predictions (field `var1.var`).

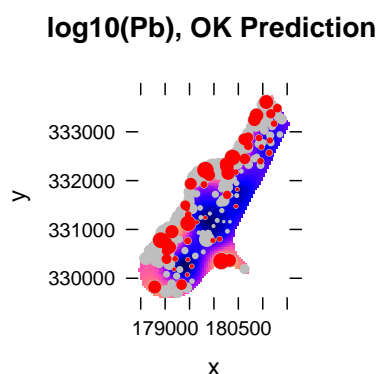
---

**TASK 11** : Display a map of the predictions and their errors. •

A grid map can be displayed with the `levelplot` and `contourplot` functions of the `lattice` package. Here we use panel functions to superimpose the sample points, as well as the extra points not included in the sample. For the predictions, we draw a `postplot` (size of symbols proportional to value); for the errors only the positions, since kriging error depends only on point configuration. The production of these plots has been encapsulated in several plotting functions which make use of `levelplot`; they are presented in §A.1 and also provided as source code in file `ck_plotfns.R`, which must be loaded at this point:

```
source("ck_plotfns.R")
plot.kresults(k.o, "var1", meuse, meuse.pb, "lead", "log10(Pb), OK")
```





The OK map shows the main features of the spatial distribution; however it is clear that some of the sample points that were not included in the sub-sample are poorly predicted. (See for example the extra point near (179750E, 330800N); the prediction is low, but the large green point shows that the value is large.) The OK prediction error map shows the expected low errors near the sample points. There are extra points near most high-error areas; using information on the co-variable from these points should improve the prediction and lower its error.

## 5.1 Evaluation

The `meuse.extra` object has the coordinates and lead levels of the 103 points at which lead was measured, but which we didn't use in the subset, either for modelling or interpolation. This can be used as a *evaluation dataset* (often called the *validation dataset*) to assess the performance of the interpolation.

The procedure is to predict (by kriging) at the points in this evaluation set `meuse.extra`, using the sample points only. Then we compare the predictions with the measured values. We work with the log10-transformed values.

---

**TASK 12** : Interpolate at the extra points, using OK from the sample points, and determine the prediction bias and precision. •

```

# predict at the extra points
k <- krige(ltpb ~ 1, meuse.pb, meuse.extra, m.ltpb.f)

[using ordinary kriging]

# compute and summarize evaluation errors
summary(k)

Object of class SpatialPointsDataFrame
Coordinates:
      min      max
x 178605 181390
y 329714 333558
Is projected: NA
proj4string : [NA]
Number of points: 103
Data attributes:
      var1.pred      var1.var
Min.   :1.744   Min.   :0.04016
1st Qu.:1.920   1st Qu.:0.05255
Median :2.108   Median :0.05797
Mean   :2.098   Mean   :0.06064
3rd Qu.:2.263   3rd Qu.:0.06648
Max.   :2.457   Max.   :0.10325

diff <- k$var1.pred - meuse.extra$ltpb
summary(diff)

      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.66294 -0.06474  0.02155  0.02502  0.12695  0.37164

sqrt(mean(sum(diff^2))) # RMSE (precision)

[1] 1.684435

mean(diff) # mean error (bias)

[1] 0.02502092

median(meuse.extra$ltpb) # median error

[1] 2.068186

```

There is a slight positive bias: the median and mean predictions at the extra points by OK is higher than the true value, as is the mean error: 0.025. The overall precision is good: RMSE: 1.684.

---


**TASK 13** : Display a postplot of the errors. •

We can use the `bubble` method of the `gstat` package to plot spatial points. It requires two arguments: the spatial object (here, a kriging interpolation) and the field in the object to plot, in this case the difference. Positive values are plotted in green and negative in red, with the size of the “bubble” proportional to the distance from zero.

This has also been included in the `plot.valids` function listed in §A.1; this also plots the cross-validation (see next §§). Since we haven’t yet computed the cross-validation, we just look at the evaluation residuals. The `bubble` method requires a spatial object, so we first have to convert the differences, assigning them the coordinates from the evaluation points:

```
diff.df <- as.data.frame(diff)
coordinates(diff.df) <- coordinates(meuse.extra)
bubble(diff.df, zcol="diff", pch=1,
        main="OK evaluation errors at undersampled points, log10(Pb)")
```

### OK evaluation errors at undersampled points, log10(Pb)



This Figure shows the evaluation errors at the extra points. There are several large relative errors, from -0.66 to 0.37, compared to the median value for the extra points, 2.07. Note the very large under-prediction at the extra point near (179750E, 330800N) which we noticed in the OK map.

## 5.2 Cross-validation

We can also evaluate the success of OK by *cross-validation* at the 52 points in the subset. Each point is held out in turn, and the prediction at that point is made from the remaining 51 points, using the common variogram model.

Diagnostic measures are the ME (bias), RMSE (precision), and Mean Squared Deviation Ratio (MSDR) of the residuals to the prediction errors; this should be 1 because the residuals from cross-validation should equal to prediction errors at each point that was held out. For the univariate case (OK) we can use the `krige.cv` cross-validation method of the `gstat`

package.

---

**TASK 14 :** Cross-validate the OK predictions of  $\log_{10}\text{Pb}$  and compute the diagnostic measures. •

```
cv.o <- krige.cv(ltpb ~ 1, meuse.pb, model=m.ltpb.f, nfold=nrow(meuse.pb),  
               verbose=FALSE)
```

```
summary(cv.o)
```

```
Object of class SpatialPointsDataFrame
```

```
Coordinates:
```

```
      min      max
```

```
x 178786 181298
```

```
y 329822 333611
```

```
Is projected: NA
```

```
proj4string : [NA]
```

```
Number of points: 52
```

```
Data attributes:
```

var1.pred	var1.var	observed
Min. :1.783	Min. :0.04124	Min. :1.681
1st Qu.:1.957	1st Qu.:0.05594	1st Qu.:1.831
Median :2.123	Median :0.06213	Median :2.139
Mean :2.114	Mean :0.06430	Mean :2.117
3rd Qu.:2.269	3rd Qu.:0.07295	3rd Qu.:2.386
Max. :2.494	Max. :0.11311	Max. :2.683

residual	zscore	fold
Min. :-0.464118	Min. :-1.867852	Min. : 1.00
1st Qu.:-0.154049	1st Qu.:-0.574692	1st Qu.:13.75
Median :-0.055698	Median :-0.226918	Median :26.50
Mean :0.003671	Mean :0.007067	Mean :26.50
3rd Qu.:0.166942	3rd Qu.:0.699510	3rd Qu.:39.25
Max. :0.658260	Max. :2.683341	Max. :52.00

```
res <- as.data.frame(cv.o)$residual
```

```
sqrt(mean(res^2))
```

```
[1] 0.2369119
```

```
mean(res)
```

```
[1] 0.003670487
```

```
mean(res^2/as.data.frame(cv.o)$var1.var)
```

```
[1] 0.8845749
```

The `nfold` argument specifies the number of times to do the cross-validation, in this case once for each point.

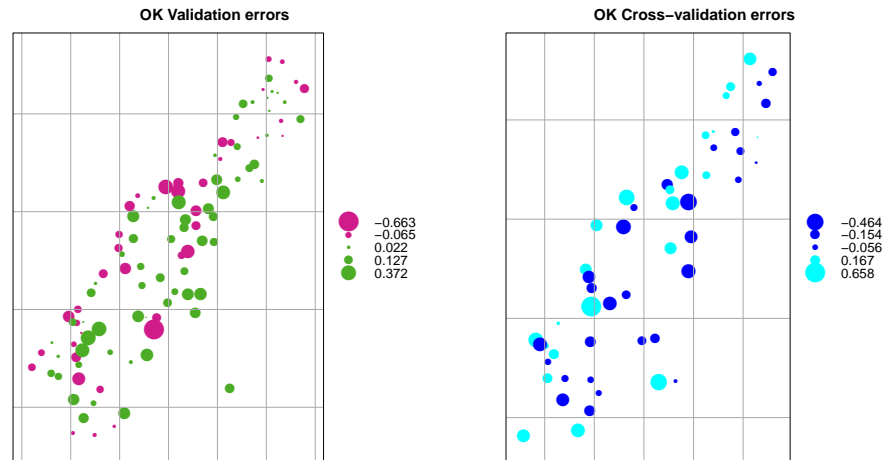
The RMSE is 0.2369, the mean error is almost zero, and the Mean Squared Deviation Ratio (MSDR) of the predictions vs. the sample is a bit lower, 0.8846, than the ideal 1, meaning that the predictions are somewhat less variable than reality; this is to be expected, as kriging is a smoothing estimator.

---

**TASK 15 :** Display a postplot of the errors. •

This is included in the `plot.valids` function, which also shows the evaluation errors from the previous § on the left side plot:

```
plot.valids(k, "var1", meuse.extra, "ltpb", cv.o, "OK")
```



This Figure (right) shows the cross-validation errors at the sample points. There are some very large errors, from -0.464 to 0.658.

### 5.3 Evaluating errors in absolute terms

The maximum absolute residual for both evaluations was  $0.66 \log_{10}\text{Pb}$ ; one might naïvely back-transform this as  $10^{0.66} = 4.6 \text{ mg kg}^{-1} \text{ Pb}$ , an order of magnitude smaller than the smallest data values. However, this conversion is not correct, because the error is *in addition to* the predicted value, which is of another order of magnitude. For example, the mean predictions are about  $2.1 \log_{10}\text{Pb}$ , i.e. about  $10^{2.1} = 125.9 \text{ mg kg}^{-1} \text{ Pb}$ ; at this level an error of  $0.66 \log_{10}\text{Pb}$  corresponds to about  $10^{2.1+0.66} - 10^{2.1} = 450 \text{ mg kg}^{-1} \text{ Pb}$ .

So, to evaluate the severity of the errors, we must compare original Pb values with predicted Pb values, *not* log-transformed values (although these were used for interpolation).

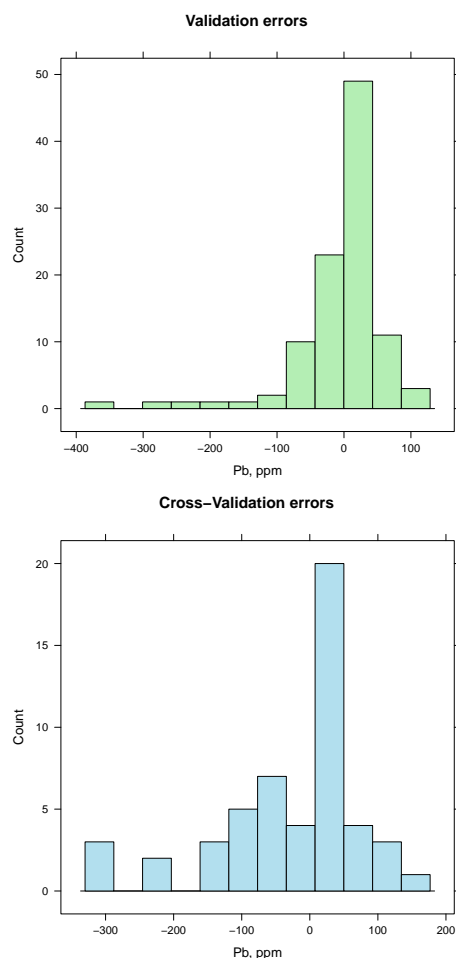
---

**TASK 16 :** Compute and summarize the absolute evaluation and cross-validation errors in units of  $\text{mg kg}^{-1} \text{ Pb}$ . •

```
diff <- 10^(k$var1.pred) - 10^(meuse.extra$ltpb)
summary(diff)
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-367.687 -23.345   6.655  -7.198  27.193  109.603

histogram(diff, col="darkseagreen2", nint=12, type="count",
  main="Validation errors", xlab="Pb, ppm")
diff <- 10^(cv.o$var1.pred) - 10^(meuse.pb$ltpb)
summary(diff)
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-311.36 -63.35  12.31  -24.63  34.42  157.98

histogram(diff, col="lightblue2", nint=12, type="count",
  main="Cross-Validation errors", xlab="Pb, ppm")
```



This Figure shows the histograms of the absolute evaluation and cross-validation errors. There are some serious errors; in particular the high negative residuals are points with high Pb which are seriously under-predicted; recall that  $100 \text{ mg kg}^{-1}$  is the level at which agricultural crops must be tested.

In summary, OK with this small dataset did not perform so well; this is the motivation for trying to improve the predictions with co-kriging.

## 6 Modelling a bivariate co-regionalisation

In order to perform co-kriging, we must first *model the spatial structure* of a co-variable and its covariance with the target variable. This is called a *co-regionalisation*. It is an extension of the theory of a single regionalised variable used for ordinary kriging. This is the same sequence we followed for OK: we first modelled the spatial structure of the target variable (§4) and then used this model to interpolate by OK (§5).

In this section we use the organic matter content as the co-variable. This is cheap to sample and determine in the laboratory, and in practice would be a good choice if it turns out to give better predictions. We will see that this isn't the case, so in a later section (§9) a different

co-variable, zinc content, is used.

## 6.1 Modelling the co-variable

First we examine the feature-space correlation.

---

**TASK 17 :** Display a scatterplot of the co-variable (organic matter) vs. the target variable (Pb) and compute the correlation coefficient. Describe the feature-space relation. •

For this task we can only use the points from the subsample, since we are simulating the situation where the target variable is under-sampled.

```
xyplot(ltpb ~ ltom, data=meuse.pb@data, pch=20, cex=1.2,
       col="blue", ylab="log10(Pb)", xlab="log10(OM)")
with(meuse.pb@data, cor(ltom, ltpb))


[1] NA

sum(is.na(meuse.pb$om))

[1] 1

with(meuse.pb@data, cor(ltom, ltpb, use = "complete"))

[1] 0.61277
```



The `cor` method does not, by default, compute the correlation coefficient if there are any missing values in either vector. Here there is one missing value of organic matter in the sub-sample (detected by the `is.na` method). The `use = "complete"` optional argument to `cor` specifies that only complete observations (with values for both variables) be used to compute the correlation.

The scatterplot and numerical correlation both show that there is a positive feature-space relation between the target and co-variable. However, the relation only explains 37.55% of the total variability (the square of the correlation coefficient); this can be seen in the “cloud”-like form of the scatterplot. We will see if this is too weak for a successful co-kriging.

**TASK 18 :** Model the omnidirectional spatial variogram of the  $\log_{10}$ -transformed covariable. We do not intend to interpolate it; however we need to know its spatial structure for co-kriging. Compare thstructure with that of the target variable. •

First we have to make a data frame without the missing values of organic matter, then we can compute the variogram. The subset method is convenient for this; its second argument is the condition (in this case, to omit the missing values) and its third argument is the list of fields to include. All available samples are used for this model, not just the subset. This simulates the realistic situation where the target variable is under-sampled but the co-variable has a full sample.

```

# all valid covariable observations, with coordinates
meuse.co <- subset(as.data.frame(meuse), !is.na(om), c(x, y, om))
# add log10-transformed variables for convenience
meuse.co <- cbind(meuse.co, ltom = log10(meuse.co$om))
str(meuse.co)

'data.frame': 153 obs. of 4 variables:
 $ x : num 181072 181025 181165 181298 181307 ...
 $ y : num 333611 333558 333537 333484 333330 ...
 $ om : num 13.6 14 13 8 8.7 7.8 9.2 9.5 10.6 6.3 ...
 $ ltom: num 1.134 1.146 1.114 0.903 0.94 ...

# convert to spatial object
coordinates(meuse.co) <- ~ x + y
# experimental variogram
v.ltom <- variogram(ltom ~ 1, meuse.co, cutoff=1800)
plot(v.ltom, pl=T)
# model by eye
m.ltom <- vgm(.035, "Sph", 800, .015)
# fit
(m.ltom.f <- fit.variogram(v.ltom, m.ltom))

  model      psill    range
1  Nug 0.01067190  0.0000
2  Sph 0.03859883 641.6166

plot(v.ltom, pl=T, model=m.ltom.f)
# compare variogram structure to target variable
m.ltom.f$range[2]; m.ltpb.f$range[2]

[1] 641.6166
[1] 876.1765

round(m.ltom.f$psill[1]/sum(m.ltom.f$psill),2)



[1] 0.22

round(m.ltpb.f$psill[1]/sum(m.ltpb.f$psill),2)

[1] 0.21

```





The structure for  $\log_{10}\text{OM}$  is quite similar to that for  $\log_{10}\text{Pb}$ . A spherical model could be fitted in both cases, with an almost identical nugget effect (about 1/5 of the total sill). However, the range for OM is only about 3/4 of that for Pb; it has a more local structure. The linear model of co-regionalisation requires the ranges of the target and co-variable be the same; here we see they are not too different.

## 6.2 Building a data structure to model co-regionalisation

Now comes the hardest part of co-kriging: modelling the co-regionalisation. We have to fit models to both the direct and cross-variograms simultaneously, and these models must lead to a positive definite cokriging system. The easiest way to ensure this is to fit a *linear model of co-regionalisation*: all models (direct and cross) have the same shape and range, but may have different partial sills and nuggets.

It is possible to model any number of variables together; in this case we have three (one target and two co-variables). The modelling and subsequent co-kriging in fact treat all the variables equally: they are all modelled (direct and cross) and predicted.

We will explore the simplest case: one target variable (here, Pb) and one

co-variable (here, OM). In a later section (§9.2) we will use the other co-variable (Zn).

---

**TASK 19 :** Build a `gstat` structure containing the two sample sets: subsample for Pb and full sample for OM. •

Up till now we've worked with data frames, variograms, and variogram models as separate objects, each with its own class. This is sufficient for univariate geostatistics. For multivariate analysis we must work with the more general `gstat` object; this allows all the flexibility of the stand-alone `gstat` program but within the R environment.

First, we build `gstat` objects to specify the two experimental variograms that we computed above. The object has named subframes, and is added to as we go. In the ordinary kriging case we used a 'wrapper' method, `variogram`, for the simple case of a single variogram. Here we have to explicitly build the `gstat` model piece-by-piece, using the `gstat` method. This allows us to build complicated structures both for modelling and interpolation. It works on an object of class `gstat` that is a *list of frames*; each frame has a separate model or interpolation result.

We fill the first frame within the `gstat` object with the  $\log_{10}$ Pb sample observations, only at the subsample points, and the second for the  $\log_{10}$ OM sample observations, at all the points where it was measured.

The first command declares that there is no existing object, whereas the second command begins by naming the existing object, so that `gstat` will add onto it. Note also how each command uses the `id` argument to name the frame within the object. Here we have one frame for Pb (`id="ltpb"`) and one for OM (`id="ltom"`).

```
(g <- gstat(NULL, id = "ltpb", form = ltpb ~ 1, data=meuse.pb))
data:
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6

(g <- gstat(g, id = "ltom", form = ltom ~ 1, data=meuse.co))
data:
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6
ltom : formula = ltom`~`1 ; data dim = 153 x 2
```

The `gstat` object now has two frames, both of which are data observation with their associated formulas and locations. From these, both direct and cross-variograms can be computed and plotted.

---

**TASK 20 :** Compute and display the two direct variograms and one cross-variogram. •

```
v.cross <- variogram(g)
str(v.cross)


Classes 'gstatVariogram' and 'data.frame': 46 obs. of 6 variables:
 $ np      : num  38 194 277 291 357 360 367 360 402 381 ...
 $ dist    : num  76.4 163.1 264.9 373.1 477.4 ...
```

```

$ gamma : num 0.0285 0.0548 0.055 0.0581 0.0755 ...
$ dir.hor: num 0 0 0 0 0 0 0 0 0 ...
$ dir.ver: num 0 0 0 0 0 0 0 0 0 ...
$ id : Factor w/ 3 levels "ltpb.ltom","ltom",...: 1 1 1 1 1 1 1 1 1 1 ...
- attr(*, "direct")= 'data.frame': 3 obs. of 2 variables:
..$ id : Factor w/ 3 levels "ltom","ltpb",...: 3 1 2
..$ is.direct: logi FALSE TRUE TRUE
- attr(*, "boundaries")= num 0 101 202 303 404 ...
- attr(*, "pseudo")= num 0
- attr(*, "what")= chr "semivariance"

```

```
plot(v.cross, pl=T)
```



The variogram object has a field `id`, which identifies the variogram (direct or cross, using the same identifiers as the `gstat` object used to specify the variogram) to which the semi-variance refers; we can also see whether the variogram estimate is from a direct or cross-variogram. Note in the larger number of point pairs for each lag for the co-variable compared to the target variable; the cross-variogram is intermediate between these.

`Gstat` does not compute the cross-variogram as strictly defined, but rather the *pseudo-cross-variogram* [19, p. 217-8]. The cross-variogram only uses co-located points, whereas the pseudo-cross-variogram uses all point pairs which can be formed from the target

and co-variables. This gives many more point pairs with which to estimate the co-regionalisation.

We have already computed the variograms for  $\log_{10}\text{Pb}$  (§4) and  $\log_{10}\text{OM}$  (§6.1) and determined that their structures and ranges were similar. The cross-variogram in this case also has a similar structure to the direct variograms; the nugget effect appears to be about 1/5 of the total sill. The range seems comparable to the range of the direct variogram of the target variable (about 850 m). Since the linear model of coregionalization requires equal ranges, this is good.

### 6.3 Fitting a linear model of co-regionalisation

The next step is to add variogram models to the `gstat` object and then fit these.

---

**TASK 21 :** Add variogram models to the `gstat` object and fit a them using the linear model of co-regionalisation. •

We use the fitted model for the target value as a starting point for all three variogram models. This is because we will use the linear model of co-regionalisation, which requires a single range and structure. By filling all the frames with one model (using the `fill.all = T` argument), these conditions are automatically met.

```
(g <- gstat(g, id = "ltpb", model = m.ltpb.f, fill.all=T))
```

```
data:
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6
ltom : formula = ltom`~`1 ; data dim = 153 x 2
variograms:
```

	model	psill	range
ltpb[1]	Nug	0.02395032	0.0000
ltpb[2]	Sph	0.08879807	876.1765
ltom[1]	Nug	0.02395032	0.0000
ltom[2]	Sph	0.08879807	876.1765
ltpb.ltom[1]	Nug	0.02395032	0.0000
ltpb.ltom[2]	Sph	0.08879807	876.1765

The `gstat` object now contains both the data and the models:

Models can be specified individually for each direct and cross-variogram; however, it is difficult to ensure that the resulting CK system is positive-definite. For example:

```
g <- gstat(g, id=c("ltpb","ltom"), model=vgm(0.02,"Sph",800,0.08))
```

could be used to specify the cross-variogram. Note how the two variables whose cross-variogram is to be specified are named in a list.

Now we fit all three variograms together, ensuring they lead to a positive definite co-kriging system. For this we use the `fit.lmc` method (“fit linear model of co-regionalization”). This takes the initial estimate,

fits all the variograms, and then each of the partial sills is adjusted (by least squares) to the closest value that will result in a positive definite matrices.

There are several variogram fitting methods in `gstat`; the default method `fit.method=7` gives more weight to nearby point-pairs, but can not handle co-located points, as in the cross-variogram (dividing by zero is not possible). So we have to specify another method; a good choice is `fit.method=6`: unweighted (ordinary least squares). Another possibility is `fit.method=1`: weights proportional to the number of point-pairs in each bin.

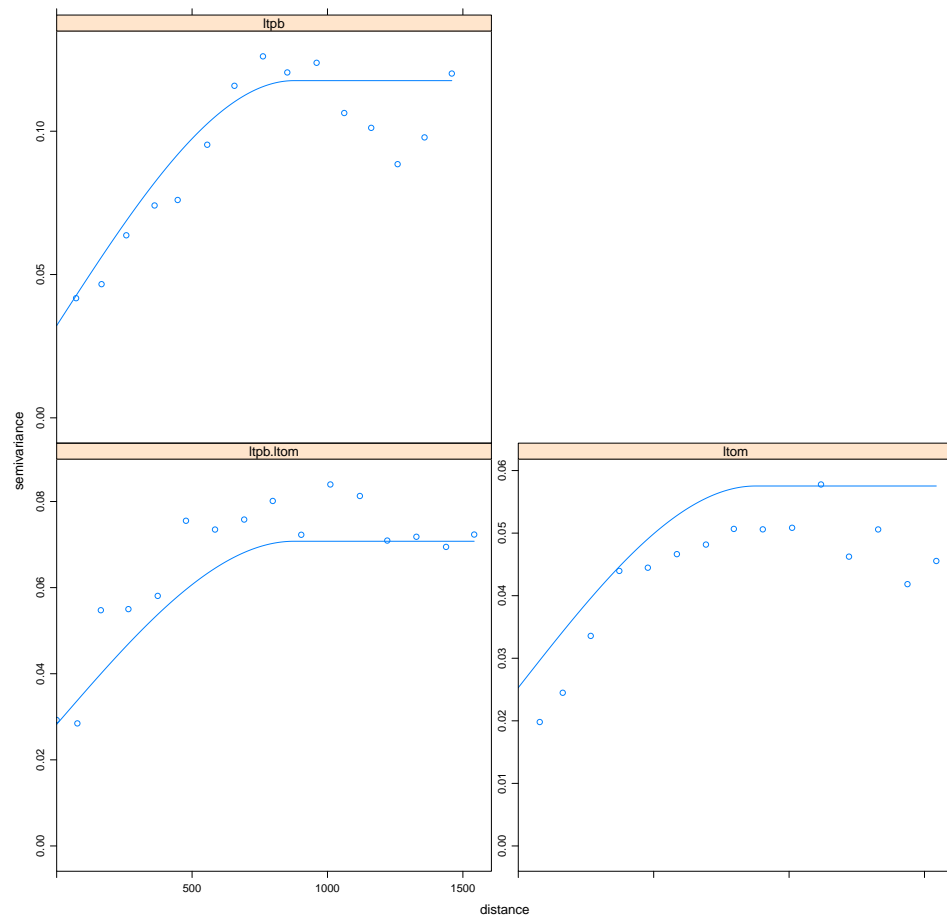
**Note:** The `correct.diagonal` optional argument is often required in order to avoid numerical instability. This is explained in the help for `fit.lmc`:

“`correct.diagonal`: multiplicative correction factor to be applied to partial sills of direct variograms only; the default value, 1.0, does not correct. If you encounter problems with singular covariance matrices during cokriging or cosimulation, you may want to try to increase this to e.g. 1.01 ... The argument `correct.diagonal` was introduced by experience: by zeroing the negative eigenvalues for fitting positive definite partial sill matrices, apparently still perfect correlation may result, leading to singular cokriging/cosimulation matrices. If someone knows of a more elegant way to get around this, please let me know.”

```
(g <- fit.lmc(v.cross, g, fit.method=6, correct.diagonal=1.01))

data:
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6
ltom : formula = ltom`~`1 ; data dim = 153 x 2
variograms:
      model      psill   range
ltpb[1]      Nug 0.03218388  0.0000
ltpb[2]      Sph 0.08546483 876.1765
ltom[1]      Nug 0.02529692  0.0000
ltom[2]      Sph 0.03223500 876.1765
ltpb.ltom[1] Nug 0.02825085  0.0000
ltpb.ltom[2] Sph 0.04251836 876.1765

plot(variogram(g), model=g$model)
```



The model fits fairly well, although the sill for  $\log_{10}OM$  seems a bit high and its range too long. Note that the ranges are not adjusted by `fit.lmc`, however all the partial sills (both for the spherical model and for the nugget) of the co-variable and cross-variogram were adjusted. So, the fitted variograms have the same range but different sills and nuggets.

#### 6.4 Comparing models of regionalisation and co-regionalisation

We can compare the two direct variograms (for subset Pb, §4, and full set OM, §6.1) with the parameters that we determined for each one separately, to see how much fitting the linear model of co-regionalisation adjusted the fit to each variable separately.

First, we have to examine the structure of the model objects (the `gstat` object `g` and the variogram model objects `m.ltom.f` and `m.ltpb.f`) to see where the parameters are stored; then we can compute their differences.

We begin with the single variogram model for OM:

```
str(m.ltom.f)
```

```
Classes 'variogramModel' and 'data.frame': 2 obs. of 9 variables:
 $ model: Factor w/ 20 levels "Nug","Exp","Sph",...: 1 3
 $ psill: num 0.0107 0.0386
```

```

$ range: num 0 642
$ kappa: num 0 0.5
$ ang1 : num 0 0
$ ang2 : num 0 0
$ ang3 : num 0 0
$ anis1: num 1 1
$ anis2: num 1 1
- attr(*, "singular")= logi FALSE
- attr(*, "SSErr")= num 6.69e-08
- attr(*, "call")= language fit.variogram(object = v.ltom, model = m.ltom)

```

This object is of class `variogramModel` and has several fields of interest, including the ranges `range` and partial sills `psill` for the different components of the variogram model. In this case there are two (as shown in the `model` field): the first is a spherical model with partial sill 0.0373 and range 653, and the second is a nugget mode with partial sill 0.0121 (and by definition range 0).

We now examine the structure of the `gstat` object; this is so complex that we look at it a level at a time:

```

str(g, max.level = 1)

List of 3
 $ data :List of 2
 $ model:List of 5
 $ call : language gstat(g = g, id = "ltom", formula = ltom ~ 1, data = meuse.co)
- attr(*, "class")= chr [1:2] "gstat" "list"

```

At the first level this is a list of three items: the data frames (here, for the two variables), the models (here, two direct and three cross-variograms and their parameters), and the most recent call.

Now we look at the structure of the list components:

```

str(g$data, max.level = 1)

List of 2
 $ ltpb:List of 15
 $ ltom:List of 15

str(g$model, max.level = 1)

List of 5
 $ ltpb      :Classes 'variogramModel' and 'data.frame': 2 obs. of  9 variables:
 .. attr(*, "singular")= logi FALSE
 .. attr(*, "SSErr")= num 0.00171
 .. attr(*, "call")= language fit.variogram(object = x, model = m, fit.ranges = fit.ranges,
 $ ltpb.ltom:Classes 'variogramModel' and 'data.frame': 2 obs. of  9 variables:
 .. attr(*, "singular")= logi FALSE
 .. attr(*, "SSErr")= num 0.00142
 .. attr(*, "call")= language fit.variogram(object = x, model = m, fit.ranges = fit.ranges,
 $ ltom      :Classes 'variogramModel' and 'data.frame': 2 obs. of  9 variables:
 .. attr(*, "singular")= logi FALSE
 .. attr(*, "SSErr")= num 0.000233
 .. attr(*, "call")= language fit.variogram(object = x, model = m, fit.ranges = fit.ranges,
 $ ltom.NA   :Classes 'variogramModel' and 'data.frame': 2 obs. of  9 variables:
 .. attr(*, "singular")= logi FALSE
 .. attr(*, "SSErr")= num 1.11e-07
 .. attr(*, "call")= language fit.variogram(object = v.ltpb, model = m.ltpb)
 $ ltom.ltom:Classes 'variogramModel' and 'data.frame': 2 obs. of  9 variables:
 .. attr(*, "singular")= logi FALSE
 .. attr(*, "SSErr")= num 1.11e-07
 .. attr(*, "call")= language fit.variogram(object = v.ltpb, model = m.ltpb)

```

The data has two sub-lists and the models five. There are two cross-variograms we didn't ask for explicitly: `ltom.NA` for the missing values of OM, and `ltom.ltom` for the OM values measured at different sampling locations (subset vs. full set).

```
str(g$data$ltpb, max.level = 1)

List of 15
 $ formula      :Class 'formula' language ltpb ~ 1
 .. ..- attr(*, ".Environment")=<environment: R_GlobalEnv>
 $ data         :Formal class 'SpatialPointsDataFrame' [package "sp"] with 5 slots
 $ has.intercept: int 1
 $ beta        : num(0)
 $ nmax        : num Inf
 $ nmin        : num 0
 $ omax        : num 0
 $ maxdist     : num Inf
 $ force       : logi FALSE
 $ dummy       : logi FALSE
 $ vfn         : int 1
 $ weights     : NULL
 $ degree      : num 0
 $ vdist       : logi FALSE
 $ lambda      : num 1
```

The data list item is itself a list with 11 items. One is the data object used to estimate the variogram for the model list components which use this variable. Another is the formula used for modelling. The others control the modelling for this variable, e.g. is there a trend, should a limiting distance be used, are their weights etc.; all these are explained in the help `?gstat`.

```
str(g$model$ltpb, max.level = 1)

Classes 'variogramModel' and 'data.frame': 2 obs. of 9 variables:
 $ model: Factor w/ 20 levels "Nug","Exp","Sph",...: 1 3
 $ psi11: num 0.0322 0.0855
 $ range: num 0 876
 $ kappa: num 0 0.5
 $ ang1 : num 0 0
 $ ang2 : num 0 0
 $ ang3 : num 0 0
 $ anis1: num 1 1
 $ anis2: num 1 1
 - attr(*, "singular")= logi FALSE
 - attr(*, "SSErr")= num 0.00171
 - attr(*, "call")= language fit.variogram(object = x, model = m, fit.ranges = fit.ranges,
```

The structure for the model within the `gstat` object is exactly that same as for the single variogram; it is of the same class: `variogramModel`.

Now that we know the structure, we can compare the models:

```
g$model$ltom$psi11 - m.ltom.f$psi11

[1] 0.014625015 -0.006363833

sum(g$model$ltom$psi11) - sum(m.ltom.f$psi11)

[1] 0.008261181

sum(g$model$ltom$psi11)

[1] 0.05753192

g$model$ltpb$psi11 - m.ltpb.f$psi11
```



```
[1] 0.008233553 -0.003333242
sum(g$model$ltpb$psill) - sum(m.ltpb.f$psill)
[1] 0.004900311
sum(g$model$ltpb$psill)
[1] 0.1176487
```

The total sills 0.058 for  $\log_{10}\text{OM}$  and 0.118 for  $\log_{10}\text{Pb}$ , were hardly affected (in both cases raised slightly). The partial sills were adjusted even less. Thus the common starting point for the linear model of co-regionalisation was satisfactory. Note that the total sill for Pb is about double that for OM.

## 7 Co-kriging with one co-variable

Now that we have modelled the co-regionalisation, we can use the extra samples of OM to improve (we hope!) the predictions of the target variable.

---

**TASK 22 :** Predict  $\log_{10}\text{Pb}$  on the interpolation grid using the modelled co-regionalisation. Summarize the predictions and their errors. •

The wrapper method `krige` that was used for OK (§5) can only be used for univariate kriging; here we must use the `predict` method. This takes a `gstat` object as the first argument and the prediction points data frame as the second argument,.

```
# interpolate
k.c <- predict(g, meuse.grid)

Linear Model of Coregionalization found. Good.
[using ordinary cokriging]

str(k.c)

Formal class 'SpatialPointsDataFrame' [package "sp"] with 5 slots
..@ data      : 'data.frame': 3103 obs. of  5 variables:
.. ..$ ltpb.pred  : num [1:3103] 2.32 2.35 2.32 2.3 2.38 ...
.. ..$ ltpb.var   : num [1:3103] 0.0764 0.0669 0.07 0.0733 0.0564 ...
.. ..$ ltom.pred  : num [1:3103] 1.01 1.04 1.02 1.01 1.06 ...
.. ..$ ltom.var   : num [1:3103] 0.0444 0.0412 0.042 0.0431 0.0378 ...
.. ..$ cov.ltpb.ltom: num [1:3103] 0.0524 0.0481 0.0493 0.0507 0.0434 ...
..@ coords.nrs : int [1:2] 1 2
..@ coords     : num [1:3103, 1:2] 181180 181140 181180 181220 181100 ...
.. .. attr(*, "dimnames")=List of 2
.. .. ..$ : chr [1:3103] "1" "2" "3" "4" ...
.. .. ..$ : chr [1:2] "x" "y"
..@ bbox       : num [1:2, 1:2] 178460 329620 181540 333740
.. .. attr(*, "dimnames")=List of 2
.. .. ..$ : chr [1:2] "x" "y"
.. .. ..$ : chr [1:2] "min" "max"
..@ proj4string: Formal class 'CRS' [package "sp"] with 1 slot
.. .. ..@ projargs: chr NA

# summarize predictions and their errors
summary(k.c$ltpb.pred); summary(k.c$ltpb.var)
```


Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
1.624	1.835	2.016	2.043	2.254	2.611
Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.04166	0.05317	0.05891	0.06190	0.06789	0.10503

Both variables and their covariance are all predicted at the same time. We are interested in the target variable, but the linear model also predicts the co-variable, in both cases using all sample information.

---

**TASK 23 :** Display the predictions and their errors as maps. •

```
plot.kresults(k.c, "ltpb", meuse, meuse.pb, "lead", "CK of Pb with OM covariable, ")
```




---

There are still some obvious discrepancies between the predictions and extra Pb points, although not so glaring as with OK.

## 7.1 Evaluation

We have a set of 103 points at which lead was measured, but which we didn't use in the subset, either for modelling or interpolation. We did use these points for CK prediction, but only their OM values. Thus, as in §5.1, the lead values at these points can be used as a *evaluation dataset* to assess the performance of the interpolation. We already have the data frame of extra points `meuse.extra` from §5.1.

---

**TASK 24 :** Interpolate at the extra points, using CK from the sample points, and determine the prediction bias and precision. •

```
# predict at the extra points
k <- predict(g, meuse.extra)

Linear Model of Coregionalization found. Good.
```

```
[using ordinary cokriging]

# compute and summarize prediction errors
diff <- k$ltpb.pred - meuse.extra$ltpb
summary(diff)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.74736 -0.06402  0.03529  0.01388  0.10208  0.52287

sqrt(sum(diff^2)/length(diff)) # RMS error (precision)

[1] 0.2118304

sum(diff)/length(diff) # mean error (bias)

[1] 0.01388193
```


There is a slight positive bias: the median and mean predictions at the extra points by CK is higher than the true value, as is the mean error (+0.015). The overall precision is good (RMSE = 0.222). There are several large relative errors (from  $-0.81$  to  $+0.55$ ) compared to the median value for the extra points of 2.07. This evaluation will be compared with the OK evaluation (§5.1) in §8.3, below.

---

**TASK 25 :** Display the evaluation errors as a bubble plot. •

```
diff <- as.data.frame(diff)
coordinates(diff) <- coordinates(meuse.extra)
bubble(diff, zcol="diff", pch=1,
        main="CK evaluation errors at undersampled points, log10(Pb)")
```

CK evaluation errors at undersampled points, log10(Pb)



This Figure shows the evaluation errors at the extra points; this is comparable to the same figure for the OK predictions.

## 7.2 Cross-validation

We can also evaluate the success of CK from its cross-validation, as we did for OK (§5.2). As with OK, the cross-validation is at the 52 points with data for Pb in the subset, since these were used for the interpolation. Each point is held out in turn, and the prediction at that point is made from the remaining 51 points. However, for CK cross-validation,

all 153 samples of the co-variable are used to predict the target variable; these points are not held out. For almost all points in the subset there is a measurement of the co-variable (OM) at that point; this should give a decided advantage to CK if the feature-space correlation is good and there is a good spatial cross-correlation.

For the univariate case (OK) we were able to use the `krige.cv` cross-validation method, but for the more general case of CK we must use the `gstat.cv` cross-validation method that can be applied to arbitrary `gstat` objects. We do not have to specify the model, since this is already stored as part of the `gstat` object.

---

**TASK 26** : Cross-validate the CK predictions of  $\log_{10}\text{Pb}$  and compute the diagnostic measures. •

```
cv.c <- gstat.cv(g)

str(cv.c)

Formal class 'SpatialPointsDataFrame' [package "sp"] with 5 slots
..@ data      : 'data.frame': 52 obs. of  6 variables:
.. ..$ ltpb.pred: num [1:52] 2.4 2.12 2.16 1.88 2.42 ...
.. ..$ ltpb.var : num [1:52] 0.01742 0.01417 0.01011 0.01481 0.00903 ...
.. ..$ observed : num [1:52] 2.48 2.06 2.12 1.9 2.45 ...
.. ..$ residual : num [1:52] 0.0712 -0.0542 -0.0428 0.027 0.0319 ...
.. ..$ zscore   : num [1:52] 0.54 -0.456 -0.425 0.222 0.336 ...
.. ..$ fold     : int [1:52] 1 2 3 4 5 6 7 8 9 10 ...
..@ coords.nrs : num(0)
..@ coords     : num [1:52, 1:2] 181072 181298 181165 181232 180874 ...
.. ..- attr(*, "dimnames")=List of 2
.. .. ..$ : NULL
.. .. ..$ : chr [1:2] "x" "y"
..@ bbox       : num [1:2, 1:2] 178786 329822 181298 333611
.. ..- attr(*, "dimnames")=List of 2
.. .. ..$ : chr [1:2] "x" "y"
.. .. ..$ : chr [1:2] "min" "max"
..@ proj4string:Formal class 'CRS' [package "sp"] with 1 slot
.. .. ..@ projargs: chr NA

summary(cv.c$residual)

      Min.   1st Qu.   Median     Mean   3rd Qu.    Max.
-0.504049 -0.103375  0.014853  0.004864  0.162211  0.468696

sqrt(mean(cv.c$residual^2))

[1] 0.2037365

mean(cv.c$residual)

[1] 0.004864414

mean(cv.c$residual^2/cv.c$ltpb.var)

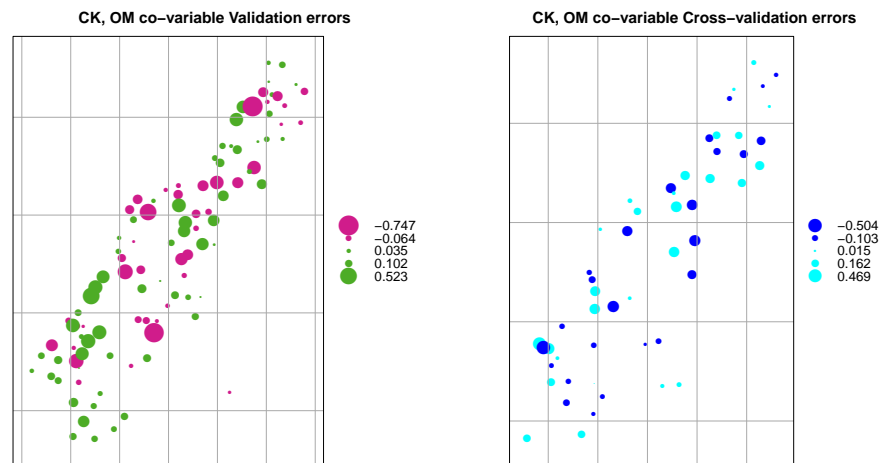
[1] 4.913226
```

The RMSE is 0.209, the mean error is almost zero, but the Mean Squared Deviation Ratio (MSDR) of the prediction residuals vs. the kriging error is very high (12.6) meaning that the actual residuals are an order of magnitude greater than what CK predicts; this is very poor performance, and shows that the CK prediction errors are too low.

---

**TASK 27** : Display a postplot of the errors. •

```
plot.valids(k, "ltpb", meuse.extra, "ltpb", cv.c, "CK, OM co-variable")
```



This Figure (right) shows the cross-validation errors at the sample points; this is comparable to the same Figure for the OK predictions. There are some very large errors, from  $-0.52$  to  $+0.51$ ; this is similar to the OK result.

## 8 Comparison of co-kriging and ordinary kriging

We have now used both OK and CK to predict the lead distribution across the study area. In this section we compare the predictions and their errors, and investigate the reasons for any discrepancies.

### 8.1 Predictions

First we compare the predictions of the target variable from OK and CK.

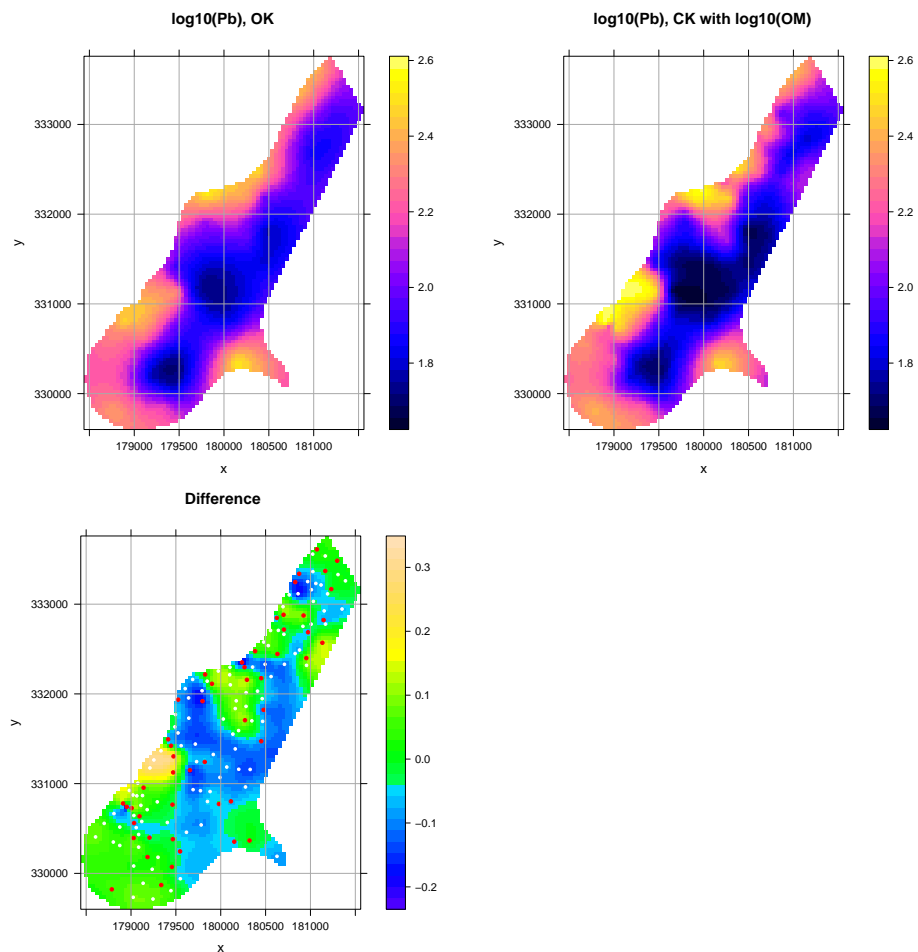
---

**TASK 28** : Determine the differences between co-kriging (with OM) and OK predictions of  $\log_{10}\text{Pb}$ , and display these as maps. •

These can be compared graphically on the same scale by use of the range function on the list of kriging predictions to find the extremes and then by dividing this range into a number (here, 32) of break points. The differences between predictions are computed, placed into a temporary data frame, and displayed with `levelplot`, using a different colour scheme than for the predictions. This has been encapsulated in the `plot.cf` function which is listed in §A.1 and also provided as source code in file `ck_plotfns.R`.

```
plot.cf(k.o, k.c, "var1", "ltpb", meuse, meuse.pb, "pred",  
        "log10(Pb), OK", "log10(Pb), CK with log10(OM)")
```

```
[1] "diff <- k2.o$ltpb.pred- k1.o$var1.pred"
```



There are some regions of the map where the predictions are quite different; these are in the undersampled regions where the OM sample provides new information. Now we compare them numerically:

```
summary(k.o$var1.pred); summary(k.c$1tpb.pred)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
1.707  1.897   2.027   2.060  2.231   2.530
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
1.624  1.835   2.016   2.043  2.254   2.611

diff <- k.c$1tpb.pred - k.o$var1.pred
summary(diff); rm(diff)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.19914 -0.07768 -0.02929 -0.01725  0.03894  0.31329
```

The CK predictions are on average slightly lower than the OK predictions; there are some fairly large differences, both positive and negative.

## 8.2 Prediction errors

Second, we compare the prediction errors for the target variable from OK and CK.

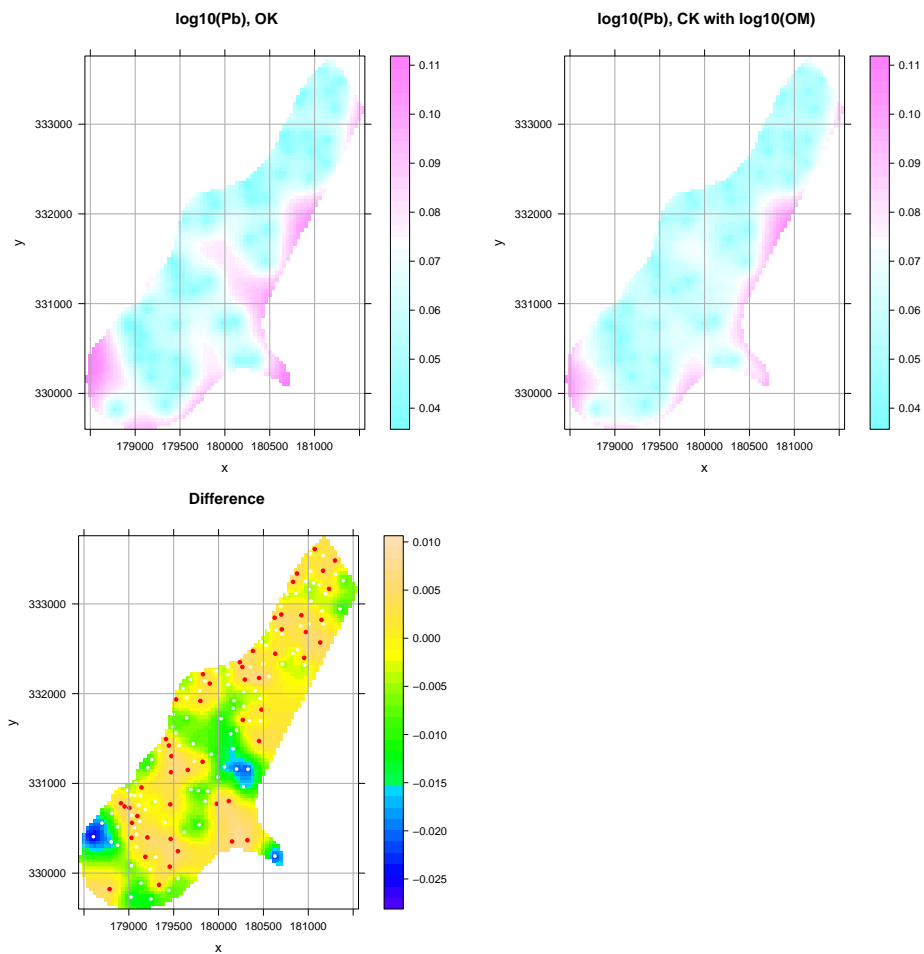
A major aim of using co-kriging is to reduce the prediction variances. With so much new information (103 extra points, plus the 52 collocated

points) and a fairly good model of co-regionalisation, we expect lower prediction errors.

**TASK 29 :** Determine the differences between the prediction errors of CK with OM and OK and display these as maps. •

We compare these graphically on the same scale; in addition we show the differences with a different colour scheme.

```
plot.cf(k.o, k.c, "var1", "ltpb", meuse, meuse.pb, "var",
        "log10(Pb), OK", "log10(Pb), CK with log10(OM)")
[1] "diff <- k2.o$ltpb.var- k1.o$var1.var"
```



Clearly, the cokriging errors are in general larger, especially near the sample points of the subset. This is because the co-kriging predictions are based not only on the target variable but also the co-variable at these points as well as their covariance. Both of these have a significant nugget effect.

However, some areas of the map have much improved predictions; these are the areas far from the subsample points and relatively nearer the extra samples with the co-variable.

Now we compare the prediction errors numerically:

```
summary(k.o$var1.var); summary(k.c$1tpb.var)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.03569 0.05195 0.06078 0.06380 0.07379 0.11188
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.04166 0.05317 0.05891 0.06190 0.06789 0.10503

diff.var <- (k.c$1tpb.var - k.o$var1.var)
summary(diff.var)

  Min.    1st Qu.    Median      Mean   3rd Qu.    Max.
-0.0257395 -0.0052398 -0.0005685 -0.0019007  0.0024603  0.0082648

round(100*sum(diff.var < 0)/length(diff.var))

[1] 54

rm(diff.var)
```

The OK prediction variances are higher than the CK prediction variances over 67% of the map.

### 8.3 Validations

We used the extra sample points to evaluate the performance of OK (§5.1) and CK (§7.1). Here we compare these two validations.

	ME	RMSE	MinE	MaxE
OK	0.0250	0.1660	-0.663	+0.372
CK with OM	0.0148	0.2226	-0.811	+0.555

The OK validation is superior in all measures except bias: the mean error of CK is closer to zero than that of OK. However, the overall precision is greater for OK, and the extreme errors are smaller.

The following Figure shows where the prediction differences are located. There are many large changes in predictions between the two validations.

```
kv.ok <- krige(1tpb ~ 1, meuse.pb, meuse.extra, m.1tpb.f)


[using ordinary kriging]

kv.ck <- predict(g, meuse.extra, debug.level=0)
kv.diff <- (kv.ck$1tpb.pred - kv.ok$var1.pred)
summary(kv.diff)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.77403 -0.12973  0.01907 -0.01114  0.10997  0.44596

d <- SpatialPointsDataFrame(coordinates(kv.ok),
                             data=as.data.frame(kv.diff))
bubble(d, col = c("plum", "seagreen"),
        panel = function(x, ...) {
          panel.xyplot(x, ...);
          panel.grid(h=-1, v=-1, col="darkgrey")}
rm(kv.ok, kv.ck, kv.diff, d)
```





## 8.4 Cross-validations

We summarize the cross-validation performance of OK (§5.2) and CK (§7.2):

	ME	RMSE	MSDR
OK	0.0037	0.2369	00.88
CK with OM	0.0059	0.2093	12.62

The cross-validation of  $\log_{10}\text{Pb}$  at the 52 subsample points shows that CK had slightly lower RMSE; however the MSDR for CK was very high, meaning that the actual residuals are an order of magnitude greater than what CK predicts. By contrast, the MSDR for OK was close to the ideal 1. Both had low bias, OK being a bit lower.


```
cv.diff <- data.frame(diff = cv.c$residual -
  cv.o$residual,
  better = (abs(cv.c$residual) < abs(cv.o$residual)))
summary(cv.diff$diff)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.404650 -0.109298  0.004691  0.001194  0.129863  0.454612

summary(cv.diff$better)

  Mode  FALSE  TRUE
Logical  22    30

d <- SpatialPointsDataFrame(coordinates(cv.c),
  data=as.data.frame(cv.diff$diff))
bubble(d,
  col=c("lightblue", "red4"),
  panel = function(x, ...) {
    panel.xyplot(x, ...);
    panel.grid(h=-1, v=-1, col="darkgrey")})
rm(cv.diff)
```



Slightly more than half (56%) of the individual cross-validation predictions are better with CK than OK. The Figure shows the difference between the cross-validation residuals for the two methods. There are large differences between the residuals, showing the large effect of CK on predictions at the sample points.

## 9 Co-kriging with a different co-variable

The use of organic matter as a co-variable gave disappointing results (§8.3, §8.4). Now we try the other co-variable (Zn), to see if it is a better co-variate. We follow the same procedure as explained in detail in §6 (modelling the co-regionalisation) and §7 (co-kriging with one co-variable), so here we present the code and results without explanation, concentrating on the discussion.

### 9.1 Modelling the co-variable Zn

As with the other co-variable (§6.1), we first we examine the feature-space correlation.


---

**TASK 30** : Display a scatterplot of the co-variable (Zn) vs. the target variable (Pb) and compute the correlation coefficient. Describe the feature-space relation. •

Note that we should only use the points from the subsample, since we are simulating the situation where the target variable is under-sampled.

```
xypplot(ltpb ~ ltz, data=meuse.pb@data, pch=20, cex=1.2,
        col="red4", ylab="log10(Pb)", xlab="log10(Zn)")
with(meuse.pb@data, cor(ltz, ltpb))
```

[1] 0.9737985



The scatterplot and numerical correlation both show that there is a very strong positive feature-space relation between the target and co-variable. However, there is a big difference between the co-variables: in the case of OM (§6.1) the relation only explained about 38% of the total variability, whereas in this case the relation explains  $0,9738^2 = 0.9483$  or about 95% of the total variability; this should give better results.

This close relation is evidence that the processes by which soils in this area are polluted with the two metals are similar.

---

**TASK 31 :** Model the omnidirectional spatial variogram of the log10-transformed covariable. Compare the structure with that of the target variable. •

We use all the available samples for the co-variable, not just the subset.

```
# all valid covariable observations, with coordinates
meuse.co <- as.data.frame(meuse)[, c("x", "y", "zinc")]
# add log10-transformed variables for convenience
meuse.co <- cbind(meuse.co, ltzn = log10(meuse.co$zinc))
coordinates(meuse.co) <- ~ x + y
str(meuse.co)

Formal class 'SpatialPointsDataFrame' [package "sp"] with 5 slots
..@ data      : 'data.frame': 155 obs. of 2 variables:
.. ..$ zinc: num [1:155] 1022 1141 640 257 269 ...
.. ..$ ltzn: num [1:155] 3.01 3.06 2.81 2.41 2.43 ...
..@ coords.nrs : int [1:2] 1 2
..@ coords     : num [1:155, 1:2] 181072 181025 181165 181298 181307 ...
.. .. attr(*, "dimnames")=List of 2
.. .. ..$ : chr [1:155] "1" "2" "3" "4" ...
.. .. ..$ : chr [1:2] "x" "y"
..@ bbox      : num [1:2, 1:2] 178605 329714 181390 333611
.. .. attr(*, "dimnames")=List of 2
.. .. ..$ : chr [1:2] "x" "y"
.. .. ..$ : chr [1:2] "min" "max"
..@ proj4string: Formal class 'CRS' [package "sp"] with 1 slot
.. .. ..@ projargs: chr NA

# experimental variogram
v.ltzn <- variogram(ltzn ~ 1, meuse.co, cutoff=1800)
```

```

plot(v.ltz, pl=T)
# model by eye
m.ltz <- vgm(.11, "Sph", 1000, .02)
# fit
(m.ltz.f <- fit.variogram(v.ltz, m.ltz))

  model    psill    range
1  Nug 0.01552141  0.0000
2  Sph 0.10615089 976.6231

plot(v.ltz, pl=T, model=m.ltz.f)
# compare variogram structure to target variable
m.ltz.f$range[2]; m.ltpb.f$range[2]

[1] 976.6231
[1] 876.1765



round(m.ltz.f$psill[1]/sum(m.ltz.f$psill),2)

[1] 0.13

round(m.ltpb.f$psill[1]/sum(m.ltpb.f$psill),2)

[1] 0.21

```



The structure for  $\log_{10}\text{Zn}$  is quite similar to that for  $\log_{10}\text{Pb}$ . A spherical model could be fitted in both cases, but with a much lower nugget effect (about 1/8 of the total sill) and a somewhat longer range (976 vs. 876 m)

for  $\log_{10}\text{Zn}$ . Thus Zn is more structured than Pb.

## 9.2 Building a data structure to model co-regionalisation of Pb with Zn

As with the other co-variable (§6), the next step is to model its co-regionalisation with the target variable.

---

**TASK 32 :** Build a `gstat` structure containing the two sample sets: subsample for Pb and full sample for Zn. •

```
(g2 <- gstat(NULL, id = "ltpb", form = ltpb ~ 1, data = meuse.pb))
```

```
data:
```

```
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6
```

```
(g2 <- gstat(g2, id = "ltzn", form = ltzn ~ 1, data = meuse.co))
```

```
data:
```

```
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6
```


```
ltzn : formula = ltzn`~`1 ; data dim = 155 x 2
```

---

**TASK 33 :** Compute and display the two direct variograms and one cross-variogram. •

```
v.cross <- variogram(g2)
```

```
plot(v.cross, pl=T)
```



### 9.3 Fitting a linear model of co-regionalisation for Zn and Pb

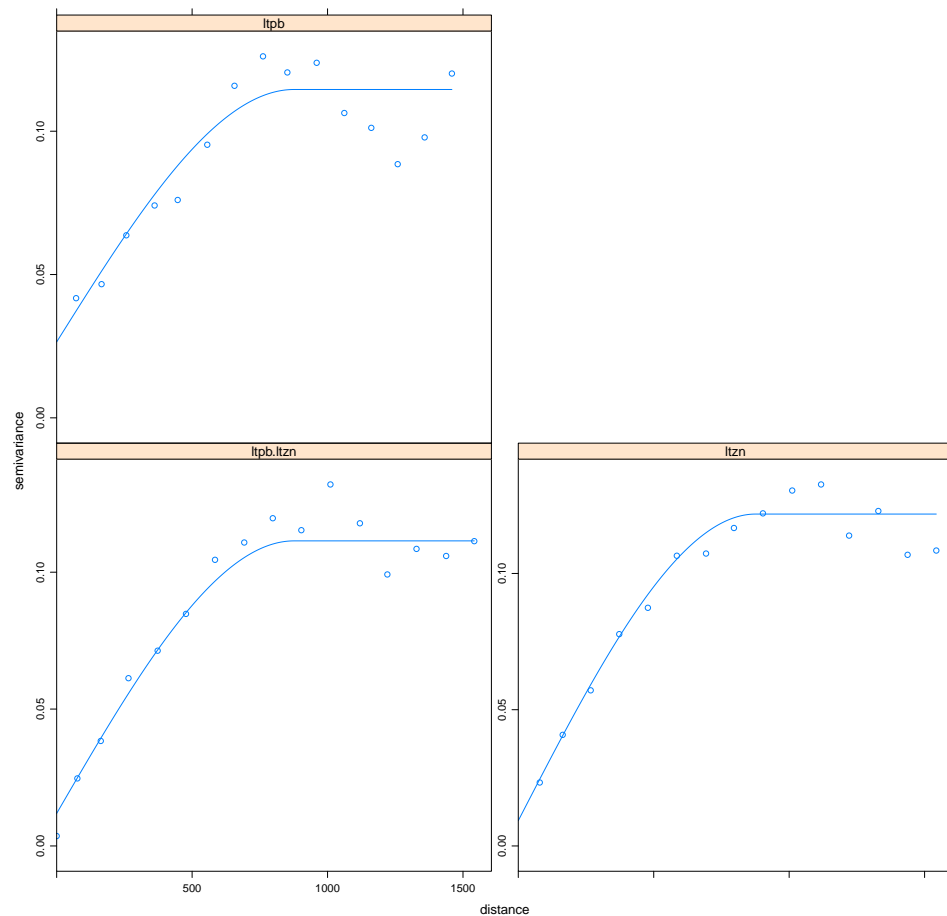
**TASK 34 :** Add variogram models to the `gstat` object and fit a linear model of co-regionalisation to them. •

```
g2 <- gstat(g2, id = "ltpb", model = m.ltpb.f, fill.all=T)
(g2 <- fit.lmc(v.cross, g2, fit.method=6))
```

```
data:
ltpb : formula = ltpb`~`1 ; data dim = 52 x 6
ltzn : formula = ltzn`~`1 ; data dim = 155 x 2
variograms:
```

	model	psill	range
ltpb[1]	Nug	0.026487524	0.0000
ltpb[2]	Sph	0.088075546	876.1765
ltzn[1]	Nug	0.009280069	0.0000
ltzn[2]	Sph	0.112520566	876.1765
ltpb.ltzn[1]	Nug	0.011893222	0.0000
ltpb.ltzn[2]	Sph	0.099550542	876.1765

```
plot(variogram(g2), model=g2$model)
```



The model fits fairly well, although the sill for the cross-variogram seems a bit low. The enforced common range (876 m) does not seem unreasonable.

**TASK 35 :** Compare these fits to the fitted models for the direct variograms. •

```
# zinc
g2$model$ltn$psill - m.ltn.f$psill
[1] -0.006241337 0.006369676
sum(g2$model$ltn$psill) - sum(m.ltn.f$psill)
[1] 0.0001283387
sum(g2$model$ltn$psill)
[1] 0.1218006

# lead
g2$model$ltpb$psill - m.ltpb.f$psill
[1] 0.0025372006 -0.0007225288
sum(g2$model$ltpb$psill) - sum(m.ltpb.f$psill)
```

```
[1] 0.001814672
sum(g2$model$1tpb$psill)
[1] 0.1145631
```

The total sills, 0.124 for  $\log_{10}\text{Zn}$  and 0.117 for  $\log_{10}\text{Pb}$ , were hardly affected (in both cases raised slightly). The partial sills were adjusted even less. Thus the common starting point for the linear model of co-regionalisation was satisfactory. The partial and total sills are quite similar between the two models.

---

**TASK 36 :** Compare the fitted variogram for Pb in the two co-regionalisations; this one with Zn and the previous with OM (§6). •

```
g2$model$1tpb$psill - g$model$1tpb.$psill
[1] -0.00176333  0.04555719
sum(g2$model$1tpb$psill) - sum(g$model$1tpb$psill)
[1] -0.003085639
```

The models are somewhat different: the model co-fitted with Zn has slightly higher partial sill of the spherical model, nugget, and therefore total sill. This shows that the linear model of co-regionalisation depends on the variables included in the model.

## 9.4 Co-kriging Pb with Zn as a co-variable

---

**TASK 37 :** Predict  $\log_{10}\text{Pb}$  on the interpolation grid using the modelled co-regionalisation with  $\log_{10}\text{Zn}$ . Summarize the predictions and their errors. •

```
k.c2 <- predict(g2, meuse.grid, debug.level=0)
summary(k.c2$1tpb.pred); summary(k.c2$1tpb.var)
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
1.663  1.843   1.970   2.022  2.201   2.690
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.03190 0.03975  0.04357  0.04707  0.05108  0.09503
```

The first line of the summary is for  $\log_{10}\text{Pb}$ , the second for the covariable  $\log_{10}\text{Zn}$ . Note that co-kriging predicts both variables together. In this case, the strong correlation might make a better map of both variables, compared to predicting them individually.

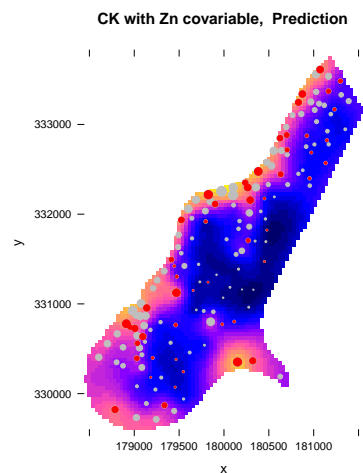
The prediction errors for  $\log_{10}\text{Pb}$  are somewhat lower than for the CK with OM (§7).

---

**TASK 38 :** Display the predictions and their errors as maps. •

```
plot.kresults(k.c2, "1tpb", meuse, meuse.pb, "lead", "CK with Zn covariable, ")
```





This figure shows the results. There are no glaring discrepancies between the predictions and extra Pb points, unlike the CK predictions using OM as co-variable and OK. In particular, the extra point near (179750E, 330800N) is well-predicted; this is because of the high Zn value at that point in the full data set.

## 9.5 Validation

**TASK 39:** Interpolate at the extra (validation) points, using CK from the sample points, and determine the actual prediction bias and precision. •

```
# predict at the extra points
k <- predict(g2, meuse.extra, debug.level=0)
# compute and summarize prediction errors
diff <- k$ltpb.pred - meuse.extra$ltpb
summary(diff)

      Min.   1st Qu.   Median     Mean   3rd Qu.     Max.
-0.132012 -0.054154 -0.007258  0.005356  0.050025  0.200444

sqrt(sum(diff^2)/length(diff)) # RMS error (precision)

[1] 0.07547095


sum(diff)/length(diff) # mean error (bias)

[1] 0.005355924
```

There is a very slight positive bias: the mean predictions at the extra points by CK is higher than the true value, as is the mean error (+0.005); however the median is slightly lower (-0.006). The overall precision is very good (RMSE = 0.075). There are no large relative errors (only from -0.14 to +0.19) compared to the median value for the extra points of 2.07.

**TASK 40:** Display the validation errors as a bubble plot. •

```
plot.valids(k, "ltpb", meuse.extra, "ltpb", cv.o, "CK, Zn co-variable")
```



This Figure (left) shows the prediction errors at the validation points; this is much better than the validation of CK with OM as covariable (note the different scales).

## 9.6 Cross-validation

**TASK 41 :** Cross-validate the CK predictions of  $\log_{10}\text{Pb}$  and compute the diagnostic measures. •

```
cv.c2 <- gstat.cv(g2, verbose=FALSE)

Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
```



```
[using ordinary cokriging]

summary(cv.c2$residual)

      Min.      1st Qu.      Median      Mean      3rd Qu.      Max.
-1.865e-01 -4.148e-02  9.354e-03 -1.105e-05  4.441e-02  2.427e-01

sqrt(mean(cv.c2$residual^2))

[1] 0.07213579

mean(cv.c2$residual)

[1] -1.104855e-05

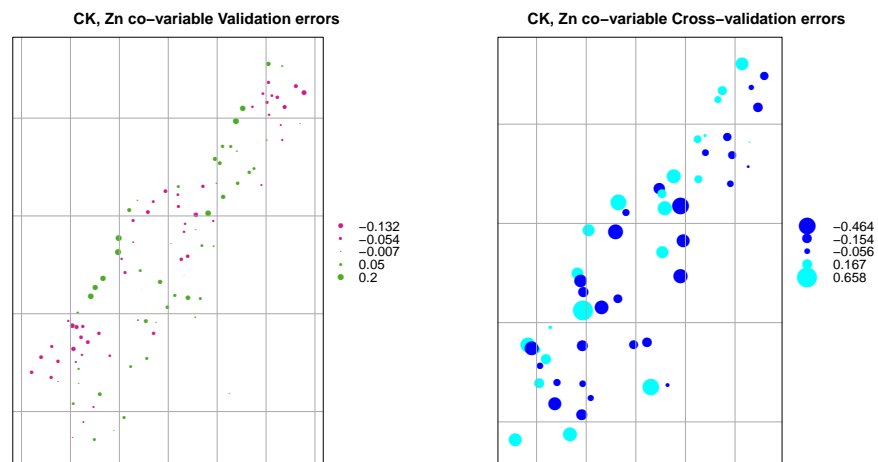
mean(cv.c2$residual^2/cv.c2$1tpb.var)

[1] 0.4144632
```

The RMSE is 0.07, which is much lower than when using OM as the co-variable (0.21, §7.2). The mean error is again almost zero, but the Mean Squared Deviation Ratio (MSDR) lower than the ideal 1 (0.31), meaning that the CK prediction errors are less variable than reality, as expected with a smooth predictor.

**TASK 42 :** Display a postplot of the errors.

```
plot.valids(k, "1tpb", meuse.extra, "1tpb", cv.o, "CK, Zn co-variable")
```



This Figure (right) shows the cross-validation errors at the sample points; this is comparable to the OK predictions and the CK predictions with OM as the co-variable. The maximum errors, from  $-0.18$  to  $+0.23$  are about half those for OK or for CK with co-variable OM.

## 9.7 Comparison of co-kriging Pb with Zn and ordinary kriging of Pb

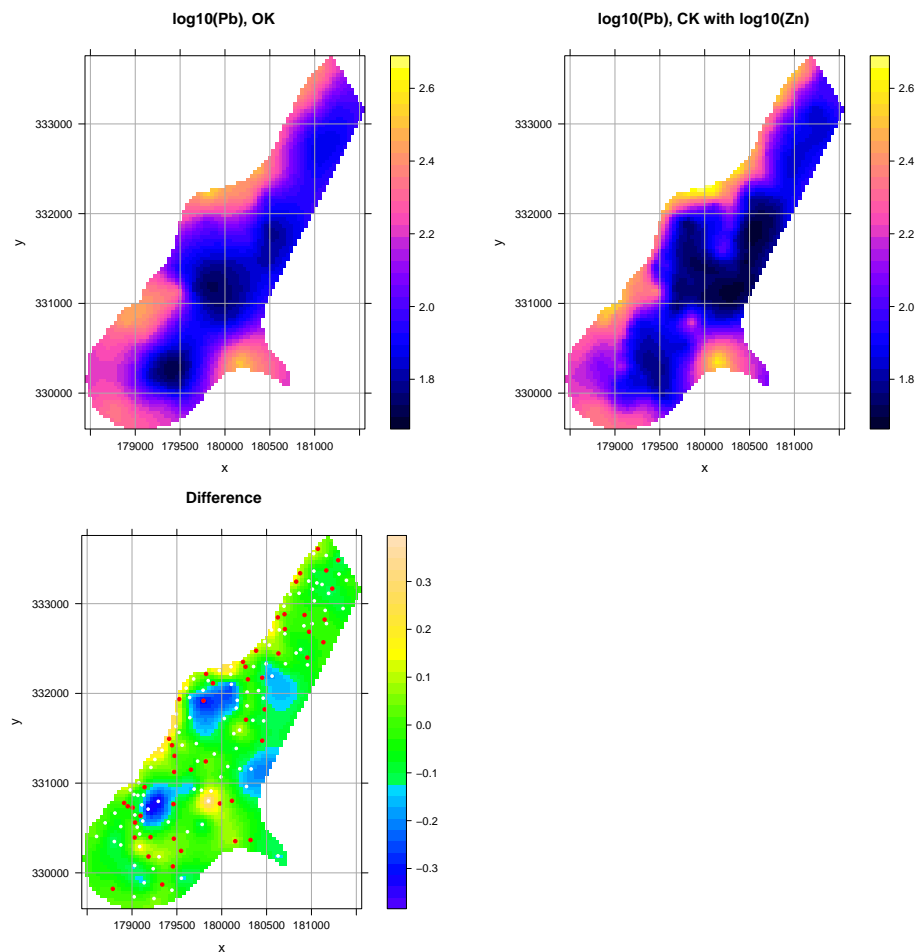
In this section we compare the CK and OK predictions and their errors, and investigate the reasons for any discrepancies for CK with co-variable Zn, just as we did for CK with co-variable OM (§8).

## 9.7.1 Predictions

**TASK 43 :** Determine the differences between co-kriging (with Zn) and OK predictions of  $\log_{10}\text{Pb}$ , and display these as maps. •

```
plot.cf(k.o, k.c2, "var1", "ltpb", meuse, meuse.pb, "pred",
        "log10(Pb), OK", "log10(Pb), CK with log10(Zn)")
```

```
[1] "diff <- k2.o$ltpb.pred- k1.o$var1.pred"
```



There are some regions of the map where the predictions are quite different; these are in the undersampled regions where the Zn measurement provides new information. The CK map shows finer detail. This is similar to CK using OM as the co-variable (§8.1). Now we compare them numerically:

```
summary(k.o$var1.pred); summary(k.c2$ltpb.pred)
```

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
k.o\$var1.pred	1.707	1.897	2.027	2.060	2.231	2.530
k.c2\$ltpb.pred	1.663	1.843	1.970	2.022	2.201	2.690

```
diff <- k.c2$ltpb.pred - k.o$var1.pred
summary(diff); rm(diff)
```

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
diff	-0.100	-0.050	0.000	0.000	0.050	0.100

```
-0.33624 -0.09076 -0.03722 -0.03807 0.01247 0.34837
```

The CK predictions are on average slightly lower than the OK predictions; there some substantial differences, both positive and negative, although half the differences are small, between  $-0.089$  and  $+0.012$ ; this is a narrower range than for CK with OM as the co-variable (§8.1).

### 9.7.2 Prediction variances

---

**TASK 44 :** Determine the differences between the prediction errors of CK with Zn and OK and display these as maps. •

```
summary(k.o$var1.var); summary(k.c2$1tpb.var)
```

```
   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.03569 0.05195 0.06078 0.06380 0.07379 0.11188
   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.03190 0.03975 0.04357 0.04707 0.05108 0.09503
```

```
diff.var <- (k.c2$1tpb.var - k.o$var1.var)
summary(diff.var)
```

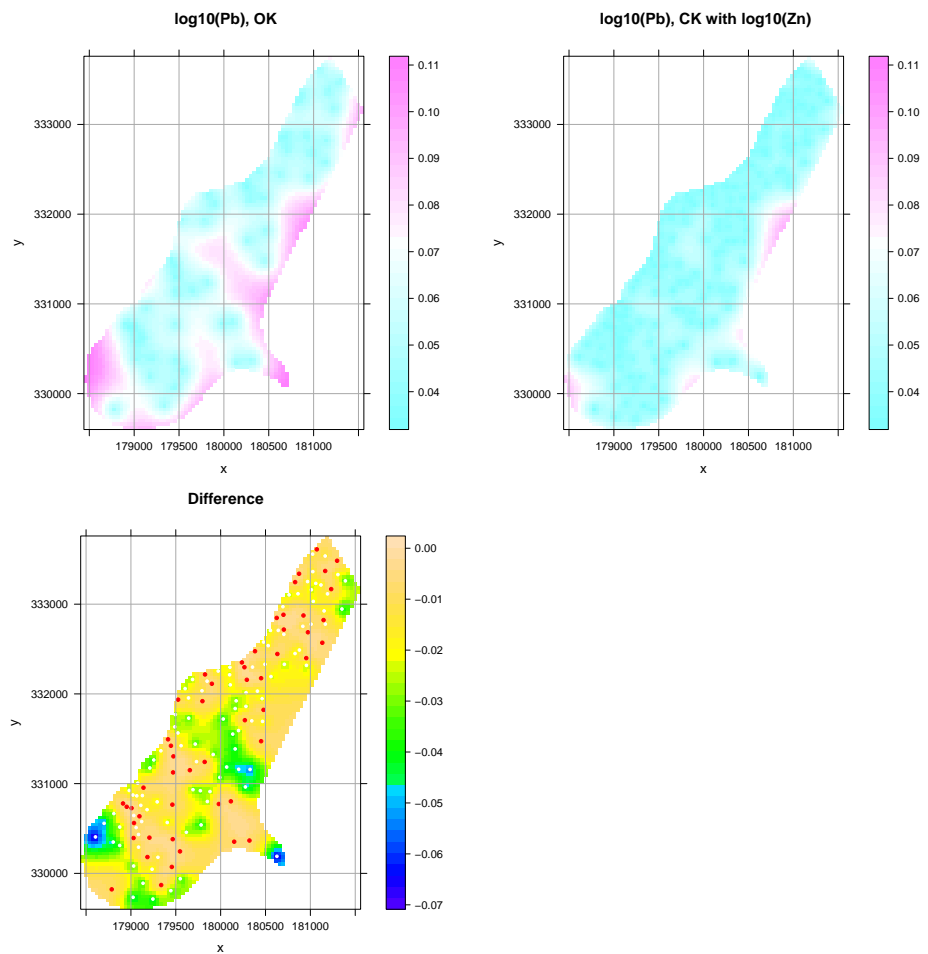
```
   Min.   1st Qu.   Median     Mean   3rd Qu.     Max.
-0.066369 -0.022628 -0.013948 -0.016724 -0.008424 -0.002091
```

The CK prediction variances are all lower than the OK prediction variances. This is a much better result than when using OM as the covariable (§8.2).

### 9.7.3 Validations

```
plot.cf(k.o, k.c2, "var1", "1tpb", meuse, meuse.pb, "var",
        "log10(Pb), OK", "log10(Pb), CK with log10(Zn)")
```

```
[1] "diff <- k2.o$1tpb.var- k1.o$var1.var"
```



The cokriging errors are in general larger, especially near the sample points of the subset. However, some areas of the map have much improved predictions; these are the areas far from the subsample points and relatively nearer the extra samples with the co-variable. This is the same pattern we saw when using CK with OM as the co-variable (§8.2), but here the differences are less marked.

Now we compare the prediction errors numerically:

```
summary(k.o$var1.var); summary(k.c2$1tpb.var)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.03569 0.05195 0.06078 0.06380 0.07379 0.11188
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.03190 0.03975 0.04357 0.04707 0.05108 0.09503

diff.var <- (k.c2$1tpb.var - k.o$var1.var)
summary(diff.var)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.066369 -0.022628 -0.013948 -0.016724 -0.008424 -0.002091
```

The CK prediction variances are all lower than the OK prediction variances. This is a much better result than when using OM as the covariable (§8.2).

## 9.7.4 Validations

We used the extra sample points to evaluate the performance of OK (§5.1) and CK with Zn as co-variable (§9.5). Here we compare these two validations; this is the same as we did in §8.3 for CK with OM as co-variable.

	ME	RMSE	MinE	MaxE
OK	0.0250	0.1660	-0.663	+0.372
CK with Zn	0.0043	0.0775	-0.137	+0.228

The CK validation is well superior in all measures; this is in contrast to the result for CK with OM as the co-variable. This Figure shows where the prediction differences are located. There are many large changes in predictions between the two validations, but not as extreme as when OM was used as the co-variable.

```
kv.ok <- krige(ltpb ~ 1, meuse.pb, meuse.extra, m.ltpb.f)

[using ordinary kriging]

kv.ck <- predict(g2, meuse.extra, debug.level=0)
kv.diff <- data.frame(x = coordinates(kv.ck)[,"x"],
                     y = coordinates(kv.ck)[,"y"],
                     diff = kv.ck$ltpb.pred - kv.ok$var1.pred)
summary(kv.diff$diff)

   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.45342 -0.10540 -0.04435 -0.01966  0.07801  0.59876

coordinates(kv.diff) <- ~ x + y
bubble(kv.diff, z="diff",
       main="CK - OK predictions, validation points, co-variable Zn",
       col = c("plum", "seagreen"),
       panel = function(x, ...) {
         panel.xyplot(x, ...);
         panel.grid(h=-1, v=-1, col="darkgrey")}
rm(kv.ok, kv.ck, kv.diff)
```



CK – OK predictions, validation points, co-variable Zn



### 9.7.5 Cross-validations

We summarize the cross-validation performance of OK (§5.2) and CK with Zn as co-variable (§9.6); this is the same as we did in §8.4 for CK with OM as co-variable.

	ME	RMSE	MSDR
OK	+0.0037	0.2369	0.88
CK with Zn	-0.0035	0.0807	0.32

CK with Zn had much lower overall error than OK. The bias of both cross-validations was very low. The MSDR for CK with Zn was quite a bit further from the ideal 1 than OK, evidence of over-smooth predictions.

```
cv.diff <- data.frame(x = coordinates(cv.c2)[,"x"],
  y = coordinates(cv.c2)[,"y"],
  diff = cv.c2$residual -
  cv.o$residual, better = (abs(cv.c2$residual) < abs(cv.o$residual)));
str(cv.diff);

'data.frame': 52 obs. of 4 variables:
 $ x      : num  181072 181298 181165 181232 180874 ...
 $ y      : num  333611 333484 333370 333168 333339 ...
 $ diff   : num  -0.2964 0.2401 0.0934 0.2366 -0.2133 ...
 $ better : logi  TRUE FALSE TRUE TRUE TRUE FALSE ...


summary(cv.diff$better);

  Mode  FALSE   TRUE
logical    7    45

coordinates(cv.diff) <- ~ x + y
bubble(cv.diff, z="diff",
  main="CK - OK cross-validation residuals, co-variable Zn",
  col = c("lightblue", "red4"),
  panel = function(x, ...) {
    panel.xyplot(x, ...);
    panel.grid(h=-1, v=-1, col="darkgrey")})
```

```
rm(cv.diff)
```

CK – OK cross-validation residuals, co-variable Zn



About 85% of the individual cross-validation predictions are better with CK than OK. This Figure shows the difference between the cross-validation residuals for the two methods. There are large differences between the residuals; this shows that CK with co-variable Zn had a large effect on predictions at the sample points. This effect is stronger than that for CK with co-variable OM.

## 10 Comparison of the three interpolations

Here we compare all three interpolations side-by-side.

### 10.1 Predictions

**TASK 45 :** Compare the predictions of  $\log_{10}\text{Pb}$  for OK (§5), CK with OM as co-variable (§7), and CK with Zn as co-variable (§9.4). •

```
# common colour scale
range <- range(k.o$var1.pred, k.c$1tpb.pred, k.c2$1tpb.pred)
breaks <- seq(range[1], range[2], length=32)
p1 <- levelplot(var1.pred ~ x+y, as.data.frame(k.o), aspect="iso",
  col.regions=bpy.colors(64), at=breaks,
  main="log10(Pb), OK predictions",
  panel = function(x, ...) {
    panel.levelplot(x, ...);
    panel.grid(h=-1, v=-1, col="darkgrey")
  })
p2 <- levelplot(1tpb.pred ~ x+y, as.data.frame(k.c), aspect="iso",
  col.regions=bpy.colors(64), at=breaks,
  main="log10(Pb), CK predictions, OM covariable",
  panel = function(x, ...) {
    panel.levelplot(x, ...);
    panel.grid(h=-1, v=-1, col="darkgrey")
  })
```

```

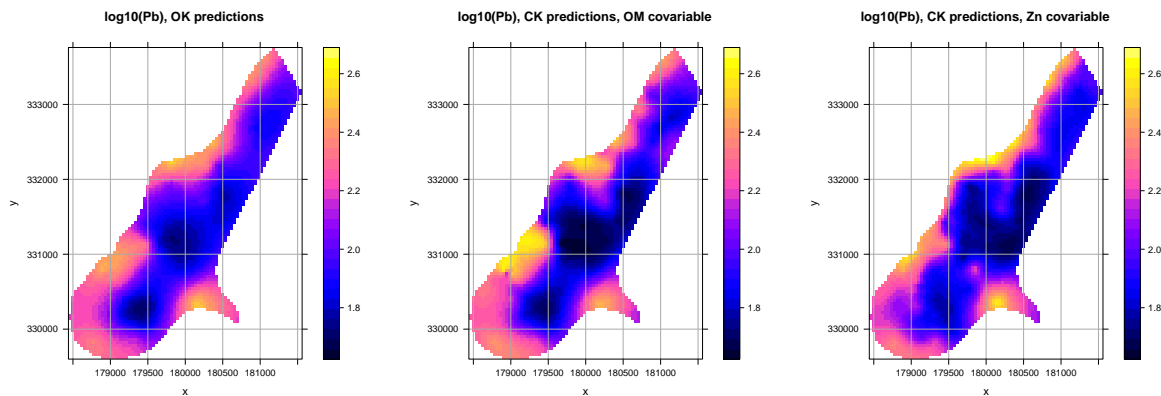
})
p3 <- levelplot(ltpb.pred ~ x+y, as.data.frame(k.c2), aspect="iso",
  col.regions=bpy.colors(64), at=breaks,
  main="log10(Pb), CK predictions, Zn covariable",
  panel = function(x, ...) {
    panel.levelplot(x, ...);
    panel.grid(h=-1, v=-1, col="darkgrey")
  })

```

```

print(p1, split=c(1,1,3,1), more=T)
print(p2, split=c(2,1,3,1), more=T)
print(p3, split=c(3,1,3,1), more=F)
rm(range, breaks)

```



This Figure (left) shows large differences in predictions. The most intricate pattern is for CK with co-variable Zn; this respects the full sample points much better than the other co-variable, because both the feature-space correlation and the spatial co-regionalisation is much stronger for Zn than OM.

## 10.2 Prediction errors

**TASK 46**: Compare the kriging predictions errors of  $\log_{10}\text{Pb}$  for OK (§5), CK with OM as co-variable (§7), and CK with Zn as co-variable (§9.4). •

```

# common colour scale
range <- range(k.o$var1.var, k.c$ltpb.var, k.c2$ltpb.var)
breaks <- seq(range[1], range[2], length=32)
p1 <- levelplot(var1.var ~ x+y, as.data.frame(k.o), aspect="iso",
  col.regions=cm.colors(64), at=breaks, colorkey=T,
  main="log10(Pb), OK errors",
  panel = function(x, ...) {
    panel.levelplot(x, ...);
    panel.points(coordinates(meuse), col="green", pch=20, cex=.6);
    panel.points(coordinates(meuse.pb), col="red", pch=20, cex=.8);
    panel.grid(h=-1, v=-1, col="darkgrey")
  })
p2 <- levelplot(ltpb.var ~ x+y, as.data.frame(k.c), aspect="iso",
  col.regions=cm.colors(64), at=breaks, colorkey=T,
  main="log10(Pb), CK errors, OM covariable",
  panel = function(x, ...) {
    panel.levelplot(x, ...);
    panel.points(coordinates(meuse), col="green", pch=20, cex=.6);
    panel.points(coordinates(meuse.pb), col="red", pch=20, cex=.8);
  })

```

```

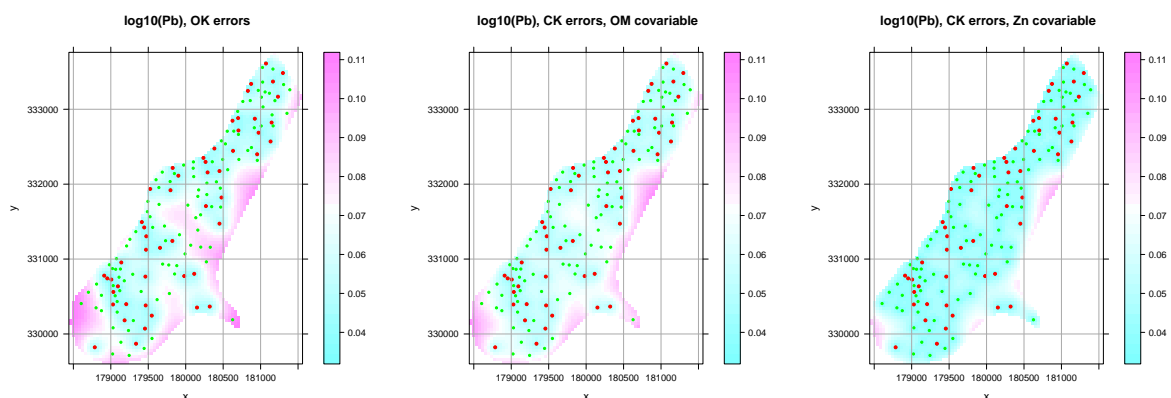
    panel.grid(h=-1, v=-1, col="darkgrey")
  })
p3 <- levelplot(ltpb.var ~ x+y, as.data.frame(k.c2), aspect="iso",
  col.regions=cm.colors(64), at=breaks, colorkey=T,
  main="log10(Pb), CK errors, Zn covariable",
  panel = function(x, ...) {
    panel.levelplot(x, ...);
    panel.points(coordinates(meuse), col="green", pch=20, cex=.6);
    panel.points(coordinates(meuse.pb), col="red", pch=20, cex=.8);
    panel.grid(h=-1, v=-1, col="darkgrey")
  })
rm(range, breaks)

```

```

print(p1, split=c(1,1,3,1), more=T)
print(p2, split=c(2,1,3,1), more=T)
print(p3, split=c(3,1,3,1), more=F)

```



This Figure (right) shows that the OK prediction errors are the smallest of the three near the sample points, whereas CK with co-variable Zn gives the lowest errors averaged across the map.

### 10.3 Evaluations

**TASK 47 :** Compare the evaluations for OK (§5.1), CK with OM as co-variable (§7.1), and CK with Zn as co-variable (§9.5). •

	ME	RMSE	MinE	MaxE
OK	0.0250	0.1660	-0.663	+0.372
CK with OM	0.0148	0.2226	-0.811	+0.555
CK with Zn	0.0043	0.0775	-0.137	+0.228

The CK with Zn as co-variable is far superior to OK in all measures; in turn this is better than CK with OM as co-variable. We see that CK with Zn substantially narrows the range of errors, lowers the RMSE, and has almost no bias at these extra points. This is a very good result.

**TASK 48 :** Plot all three evaluations and cross-validations. •

```

# compute all the evaluations and x-validations, and differences
k.ok <- krige(ltpb ~ 1, meuse.pb, meuse.extra, m.ltpb.f)

[using ordinary kriging]

diff.ok <- k.ok$var1.pred - meuse.extra$ltpb
k.ck <- predict(g, meuse.extra)

Linear Model of Coregionalization found. Good.
[using ordinary cokriging]

diff.ck <- k.ck$ltpb.pred - meuse.extra$ltpb
k.ck.zn <- predict(g2, meuse.extra)

Linear Model of Coregionalization found. Good.
[using ordinary cokriging]

diff.ck.zn <- k.ck.zn$ltpb.pred - meuse.extra$ltpb
cv.o <- krige.cv(ltpb ~ 1, meuse.pb, model=m.ltpb.f, nfold=nrow(meuse.pb),
                verbose=FALSE)
cv.c <- gstat.cv(g, verbose=FALSE)

Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]

```





```

Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]
Linear Model of Coregionalization found. Good.
[using ordinary cokriging]

# value that corresponds to cex=3
extreme <- max(abs(range(diff.ok, diff.ck, diff.ck.zn,
cv.o$residual, cv.c$residual, cv.c2$residual)))
# max. cex in each case proportional to its largest

# display all plots
d <- SpatialPointsDataFrame(coordinates(k.ok), data=as.data.frame(diff.ok))
b.ok <- bubble(d,
main="OK evaluation errors, log10(Pb)",
maxsize = 3 * (max(abs(range(diff.ok)))/extreme,
panel = function(x, ...) {
panel.xyplot(x, ...);
panel.grid(h=-1, v=-1, col="darkgrey")})
d <- SpatialPointsDataFrame(coordinates(k.ck), data=as.data.frame(diff.ck))
b.ck <- bubble(d,
main="CK evaluation errors, log10(Pb), co-variable log10(OM)",
maxsize = 3 * (max(abs(range(diff.ck)))/extreme,
panel = function(x, ...) {
panel.xyplot(x, ...);
panel.grid(h=-1, v=-1, col="darkgrey")})
d <- SpatialPointsDataFrame(coordinates(k.ck.zn), data=as.data.frame(diff.ck.zn))
b.ck.zn <- bubble(d,
main="CK evaluation errors, log10(Pb), co-variable log10(Zn)",
maxsize = 3 * (max(abs(range(diff.ck.zn)))/extreme,
panel = function(x, ...) {
panel.xyplot(x, ...);
panel.grid(h=-1, v=-1, col="darkgrey")})
d <- SpatialPointsDataFrame(coordinates(cv.o), data=as.data.frame(cv.o$residual))
b.x.ok <- bubble(d,
main="OK cross-validation errors, log10(Pb)",
maxsize = 3 * (max(abs(range(diff.ok)))/extreme, col=c(4,5),
panel = function(x, ...) {
panel.xyplot(x, ...);
panel.grid(h=-1, v=-1, col="darkgrey")})
d <- SpatialPointsDataFrame(coordinates(cv.c), data=as.data.frame(cv.c$residual))
b.x.ck <- bubble(d,
main="CK cross-validation errors, log10(Pb), co-variable log10(OM)",
maxsize = 3 * (max(abs(range(diff.ck)))/extreme, col=c(4,5),
panel = function(x, ...) {
panel.xyplot(x, ...);
panel.grid(h=-1, v=-1, col="darkgrey")})
d <- SpatialPointsDataFrame(coordinates(cv.c2), data=as.data.frame(cv.c2$residual))

```




```
b.x.ck.zn <- bubble(d,
  main="CK cross-validation errors, log10(Pb), co-variable log10(Zn)",
  maxsize = 3 * (max(abs(range(diff.ck.zn))))/extreme, col=c(4,5),
  panel = function(x, ...) {
    panel.xyplot(x, ...);
    panel.grid(h=-1, v=-1, col="darkgrey")})
rm(k.ok, k.ck, k.ck.zn)
rm(diff.ok, diff.ck, diff.ck.zn, extreme)
```

```


print(b.ok, split=c(1,1,2,3), more=T)
print(b.ck, split=c(1,2,2,3), more=T)
print(b.ck.zn, split=c(1,3,2,3), more=T)
print(b.x.ok, split=c(2,1,2,3), more=T)
print(b.x.ck, split=c(2,2,2,3), more=T)
print(b.x.ck.zn, split=c(2,3,2,3), more=F)
rm(b.ok, b.ck, b.ck.zn, b.x.ok, b.x.ck, b.x.ck.zn)

```


OK evaluation errors, log<sub>10</sub>(Pb)



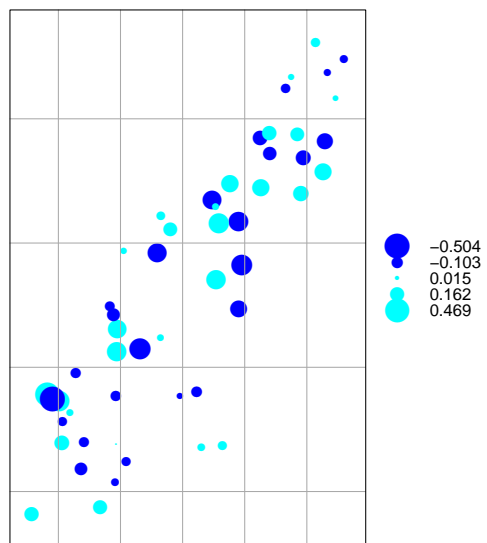
OK cross-validation errors, log<sub>10</sub>(Pb)



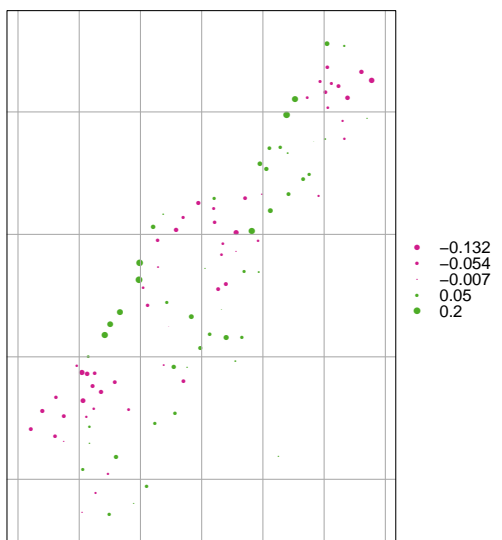
CK evaluation errors, log<sub>10</sub>(Pb), co-variable log<sub>10</sub>(OM)




CK cross-validation errors, log<sub>10</sub>(Pb), co-variable log<sub>10</sub>(OM)



CK evaluation errors, log<sub>10</sub>(Pb), co-variable log<sub>10</sub>(Zn)



CK cross-validation errors, log<sub>10</sub>(Pb), co-variable log<sub>10</sub>(Zn)



This figure (left) shows bubble plots of the evaluation errors; this is a summary of previous evaluation error plots, but with a single symbol size range for all three figures. The superiority of CK with Zn as co-variable is clear.

## 10.4 Cross-validations

**TASK 49 :** Compare the cross-validations for OK (§5.2), CK with OM as co-variable (§7.2), and CK with Zn as co-variable (§9.6). •

	ME	RMSE	MSDR
OK	+0.0037	0.2369	00.88
CK with OM	+0.0059	0.2093	12.62
CK with Zn	-0.0035	0.0807	00.32

CK with Zn as co-variable performs much better on every measure than CK with OM as co-variable. CK with Zn is better than OK except that it underestimates the prediction error at the cross-validation points.

This figure (right) shows bubble plots of the cross-validation errors; this is a summary of the previous cross-validation figures (right), but with a single symbol size range for all three figures. The superiority of CK with Zn as co-variable is clear.

## 11 Conclusion

This technical note has shown how to perform co-kriging using the `gstat` package of the R environment for statistical computing and visualisation, and how to evaluate its success by validation and cross-validation.

In the particular case studied here, the use of organic matter as the co-variable for co-kriging the target variable did not in general improve predictions compared to univariate ordinary kriging. This was despite the fairly feature-space good correlation between the two variables and the fairly good linear model of co-regionalisation that could be fitted to the two variables. This shows that the more sophisticated technique is not always better. The underlying reasons for the poor relative performance of co-kriging in this case are (1) that the feature-space correlation between OM and Pb is too weak (§6.1), explaining only about 38% of the variability; (2) the spatial structures of the two variables were somewhat different, so that the linear model of co-regionalization was not completely appropriate.

Using Zn as the co-variable, with a feature-space correlation that explained about 95% of the variability and also had a good model of co-regionalisation with the target variable, did result in a much-improved map, especially away from the sample points of the (undersampled) target variable (§9).

There is much more to co-regionalisation and co-kriging than is explained in these notes, for example indicator co-kriging, multivariate co-kriging, non-linear models of co-regionalisation, and universal co-kriging. The reader is encouraged to consult relevant texts [e.g. 2, 4, 9, 18, 19] and then experiment with `gstat`. Modelling the co-regionalisation is also tricky and should be approached with caution [e.g. 5]; remember that all kriging approaches are model-based: bad models lead to bad predictions.

## References

- [1] WS Cleveland. *Visualizing data*. AT&T Bell Laboratories; Hobart Press, Murray Hill, N.J., 1993. 4
- [2] P Goovaerts. *Geostatistics for natural resources evaluation*. Applied Geostatistics. Oxford University Press, New York; Oxford, 1997. 1, 64
- [3] R Ihaka and R Gentleman. R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, 5(3):299–314, 1996. 1
- [4] EH Isaaks and RM Srivastava. *An introduction to applied geostatistics*. Oxford University Press, New York, 1990. 1, 64
- [5] RM Lark. Two robust estimators of the cross-variogram for multivariate geostatistical analysis of soil properties. *European Journal of Soil Science*, 54(1):187–201, 2003. 64
- [6] MB McBride. *Environmental chemistry of soils*. Oxford University Press, New York, 1994. 5
- [7] N Oreskes. Evaluation (not validation) of quantitative models. *Environmental Health Perspectives*, 106(Suppl 6):1453–1460, Dec 1998. 7
- [8] N. Oreskes, K. Shrader-Frechette, and K. Belitz. Verification, validation, and confirmation of numerical models in the earth sciences. *Science*, 263:641–646, Feb 1994. URL <http://dx.doi.org/10.1126/science.263.5147.641>. 7
- [9] A Papritz and A Stein. Spatial prediction by linear kriging. In A Stein, F van der Meer, and B G F Gorte, editors, *Spatial statistics for remote sensing*, pages 83–113. Kluwer Academic, Dordrecht, 1999. 1, 64
- [10] EJ Pebesma. *Gstat User's Manual*. Dept. of Physical Geography, Utrecht University, Utrecht, version 2.3.3 edition, 2001. 2
- [11] EJ Pebesma and RS Bivand. Classes and methods for spatial data in R. *R News*, 5(2):9–13, November 2005. URL <http://CRAN.R-project.org/doc/Rnews/>. 1, 2
- [12] EJ Pebesma and CG Wesseling. Gstat: a program for geostatistical modelling, prediction and simulation. *Computers & Geosciences*, 24(1):17–31, 1998. 1, 2
- [13] PJ Ribeiro, Jr. and PJ Diggle. geoR: A package for geostatistical analysis. *R News*, 1(2):14–18, June 2001. 2
- [14] MGJ Rikken and RPG van Rijn. Soil pollution with heavy metals – an inquiry into spatial variation, cost of mapping and the risk evaluation of copper, cadmium, lead and zinc in the floodplains of the Meuse west of Stein, the Netherlands. Doctoraalveldwerkverslag, Dept. of Physical Geography, Utrecht University, 1993. 2, 3

- [15] D Sarkar. Lattice. *R News*, 2(2):19–23, June 2002. 1, 4
- [16] P Schachtschabel, H-P Blume, G Brümmer, KH Hartge, and U Schwertmann. *Lehrbuch der Bodenkunde (Scheffer, Schachtschabel)*. Enke, Stuttgart, 14 edition, 1998. 5
- [17] WN Venables and BD Ripley. *Modern applied statistics with S*. Springer-Verlag, New York, fourth edition, 2002. 2
- [18] H Wackernagel. *Multivariate geostatistics : an introduction with applications*. Springer, Berlin ; New York, 2nd edition, 1998. 64
- [19] R Webster and MA Oliver. *Geostatistics for environmental scientists*. Wiley & Sons, Chichester, 2001. 1, 25, 64

## A Producing the figures for this document

Many figures in this document combine several plots into one. This can easily be accomplished with the `lattice` graphics package. The basic idea is to create plots with any `lattice` method, including `levelplot()`, `xypplot()`, and `histogram()`, but instead of displaying them directly, *saving* them (using the assignment operator `<-`) as *trellis objects*; this is the object type created by `lattice` graphics methods<sup>5</sup>. They are then printed with the `print()` method for `trellis` objects, using the optional `split=` argument to split the display and position each saved plot within it.

The kriged maps in the tutorial were created using the `plot.kresults` function shows in the next section §A.1; see the code listing.

First, two `trellis` objects (`plot.kr` and `plot.kre`) are created, using the `levelplot()` method. Then they are placed into a one-row, two-column split display: The `split = c(1,1,2,1)` optional argument says to place this figure in position (1,1) within a matrix with dimensions (2,1) in (x, y) i.e. (columns, rows). The `more=T` argument indicates that more saved objects will be added to the current plot; `more=F` indicates that the plot is finished.

Plots can be saved for later printing or incorporation into documents by writing them to a graphics “device” that represents a file. Here is an example of how to write a PDF file.

---

### R code:

```
# open the graphics device
pdf("graph/fig_pb_ok.pdf", h=6, w=12)
# put plotting commands here
# ...
# close the graphics device
dev.off()
```

---

<sup>5</sup> The reason for this name is historical: S-PLUS has a ‘trellis’ graphics environment, on which the R package ‘lattice’ is based.

The optional `h=` and `w=` arguments specify the dimensions of the plot, in inches<sup>6</sup>. In this case the plots are side-by-side so the plot is wider than tall.

## A.1 Plotting functions

**Note:** These were written before `sp` package was developed; this package includes the `spplot` method which would probably make this code simpler and easier to understand.

These functions should be loaded into your workspace as follows:

---

### R code:

```
source("ck_plotfns.R")
```

---

The code is reproduced here for your reference; however, rather than attempt to cut-and-paste from here, it is better to download file `ck_plotfns.R` from the author's website where this tutorial is also found<sup>7</sup>.

```
### R source code, file ck_plotfns.R
### author: DG Rossiter, d.g.rossiter@cornell.edu
###
### to accompany Technical Note "Co-kriging with the gstat package
### of the R environment for statistical computing"
### plot Kriging results: predictions and variances
### with postplot of sample/subsample points on predictions
##   colour: bpy, grey, red
##   and locations of sample/subsample points on prediction variance
##   colour: cm, green, red
## arguments
##   kr.o      SpatialPointsDataFrame from kriging
##   coordinates must be named x, y
##   var1      prefix for field names *.pred, *.var, e.g. "var1"
##   samp.pts  SpatialPointsDataFrame sample points
##   subsamp.pts SpatialPointsDataFrame subsample points
##   f         field name (quoted) or number for point data values
##   title
plot.kresults <- function(kr.o, var1, samp.pts, subsamp.pts,
                          f=1, title="") {
  to.eval <- paste("plot.kr <- levelplot(",
                  paste(var1,"pred",sep="."),
                  " ~ x+y, as.data.frame(kr.o),
                  aspect='iso',
                  col.regions=bpy.colors(64), cut=32,
                  main=paste(title, 'Prediction'),
                  panel = function(x, ...) {
                    panel.levelplot(x, ...);
                    panel.points(coordinates(samp.pts), col='grey',
                                         pch=20,
                                         # log scale, but still show minimum
                                         cex=1.6 * (log10(samp.pts[[f]]) -
                                                    0.9 * min(log10(samp.pts[[f]])))));
                    panel.points(coordinates(subsamp.pts), col='red',
```

---

<sup>6</sup> 1" = 2.54 cm exactly

<sup>7</sup> [http://www.css.cornell.edu/faculty/dgr2/pubs/list.html#pubs\\_m\\_R](http://www.css.cornell.edu/faculty/dgr2/pubs/list.html#pubs_m_R)

```

        pch=20,
        cex=1.6 * (log10(subsamp.pts[[f]]) -
        0.9 * min(log10(subsamp.pts[[f]])));
        panel.grid(h=-1, v=-1, col='')
    })" );
eval(parse(text=to.eval));
to.eval <- paste("plot.kr.e <- levelplot(",
    paste(var1,"var",sep="."),
    "~ x+y, as.data.frame(kr.o),
    aspect='iso',
    col.regions=cm.colors(64), cut=32,
    main=paste(title, 'Prediction variance'),
    panel = function(x, ...) {
        panel.levelplot(x, ...);
        panel.points(coordinates(samp.pts), col='green',
            pch=20, cex=.6);
        panel.points(coordinates(subsamp.pts), col='red',
            pch=20, cex=.8); # subsample points larger
        panel.grid(h=-1, v=-1, col='darkgrey')
    })" );
eval(parse(text=to.eval));
print(plot.kr, split = c(1,1,2,1), more=T);
print(plot.kr.e, split = c(2,1,2,1), more=F)
}

### plot Kriging validation and cross-validation errors
##     colours: validation: bubble() default: palette()[2:3]
##     colours: x-valid: palette()[4:5]
## arguments
## kv.o     SpatialPointsDataFrame from kriging to validation points
## var1     prefix for kriging field name *.pred, e.g. "var1"
## valid.pts SpatialPointsDataFrame with validation points
## f       field name (quoted) or number for point data values
## cv.o    SpatialPointsDataFrame from x-validation kriging
## title
plot.valids <- function(kv.o, var1, valid.pts, f, cv.o, title="") {
    # validation errors
    to.eval <- paste("diff <- kv.o$", paste(var1,"pred",sep="."),
        " - valid.pts[[f]]")
    eval(parse(text=to.eval))
    extreme <- max(abs(range(diff, as.data.frame(cv.o)$residual)))
    d <- SpatialPointsDataFrame(kv.o, data=as.data.frame(diff))
    b1 <- bubble(d,
        main=paste(title,"Validation errors"),
        maxsize = 3 * (max(abs(range(diff))))/extreme,
        panel = function(x, ...) {
            panel.xyplot(x, ...);
            panel.grid(h=-1, v=-1, col="darkgrey")}
    )
    b2 <- bubble(cv.o, z="residual",
        main=paste(title,"Cross-validation errors"), col=c(4,5),
        maxsize = 3 * (max(abs(range(cv.o$residual))))/extreme,
        panel = function(x, ...) {
            panel.xyplot(x, ...);
            panel.grid(h=-1, v=-1, col="darkgrey")}
    )
    print(b1, split=c(1, 1, 2, 1), more=T)
    print(b2, split=c(2, 1, 2, 1), more=F)
}

```



```

### compare Kriging results: predictions or prediction variance
###   with locations of sample/subsample points on predictions
###   colour: bpy, grey, red
###   and locations of sample/subsample points on prediction variance
###   colour: cm, green, red
## arguments
##   k1.o, k2.o   SpatialPointsDataFrame's from kriging
##   coordinates must be named x, y
##   var1.1, var1.2 prefix for field names *.pred, *.var, e.g. "var1"
##   in the two objects
##   samp.pts   SpatialPointsDataFrame sample points
##   subsamp.pts SpatialPointsDataFrame subsample points
##   type      what to compare: "pred" or "var"
##   title.1, title.2 titles for the two kriging objects
plot.cf <- function(k1.o, k2.o, var1.1, var1.2, samp.pts, subsamp.pts,
  type="pred", title.1="", title.2="") {
  # common scale
  to.eval <- paste("range <- range(k1.o$", paste(var1.1, type, sep="."),
    ", k2.o$", paste(var1.2, type, sep="."), ")", sep="")
  eval(parse(text=to.eval)); # makes range
  breaks <- seq(range[1], range[2], length=32)
  to.eval <- paste("plot.k1 <- levelplot(",
    paste(var1.1, type, sep="."),
    " ~ x+y, as.data.frame(k1.o), aspect='iso', col.regions=",
    ifelse(type=="pred", 'bpy.colors', 'cm.colors'),
    "(64), at=breaks, main=title.1,
    panel = function(x, ...) {
      panel.levelplot(x, ...);
      panel.grid(h=-1, v=-1, col='darkgrey')}}")
  );
  eval(parse(text=to.eval)); # makes plot.k1
  to.eval <- paste("plot.k2 <- levelplot(",
    paste(var1.2, type, sep="."),
    " ~ x+y, as.data.frame(k2.o), aspect='iso', col.regions=",
    ifelse(type=="pred", 'bpy.colors', 'cm.colors'),
    "(64), at=breaks, main=title.2,
    panel = function(x, ...) {
      panel.levelplot(x, ...);
      panel.grid(h=-1, v=-1, col='darkgrey')}}")
  );
  eval(parse(text=to.eval)); # makes plot.k2
  to.eval <- paste("diff <- k2.o$", paste(var1.2, type, sep="."),
    "- k1.o$", paste(var1.1, type, sep="."), sep="")
  print(to.eval)
  eval(parse(text=to.eval)); # makes diff
  tmp <- data.frame(x = coordinates(k1.o)[,"x"],
    y = coordinates(k2.o)[,"y"], diff)
  plot.diff <-
    levelplot(diff ~ x+y, tmp, aspect="iso",
      col.regions=topo.colors(64), cut=32,
      main="Difference",
      panel = function(x, ...) {
        panel.levelplot(x, ...);
        panel.points(coordinates(samp.pts), col="white",
          pch=20, cex=.6);
        panel.points(coordinates(subsamp.pts), col="red",
          pch=20, cex=.8);
        panel.grid(h=-1, v=-1, col="darkgrey")
      })
  # display the plots in one figure

```

```
print(plot.k1, split = c(1,1,2,2), more=T)
print(plot.k2, split = c(2,1,2,2), more=T)
print(plot.diff, split = c(1,2,2,2), more=F)
}
```