

GeoXp: an R package for exploratory spatial data analysis

T. Laurent^{*}, A. Ruiz-Gazen[†] and C. Thomas-Agnan[‡]

Toulouse School of Economics

November 5, 2009

We present **GeoXp**, an R package implementing interactive graphics for exploratory spatial data analysis. We use data bases coming from the **spdep** package to illustrate the use of these exploratory techniques based on the coupling between a statistical graph and a map. Besides elementary plots like boxplots, histograms or simple scatterplots, **GeoXp** also couples maps with Moran scatterplots, variogram clouds, Lorenz curves, etc. In order to make the most of the multidimensionality of the data, **GeoXp** includes dimension reduction techniques such as principal components analysis and cluster analysis whose results are also linked to the map.

KEY WORDS: Exploratory analysis, spatial econometrics, spatial statistics, interactive graphics, brushing and linking, dimension reduction.

1 Introduction

Exploratory analysis of georeferenced data must take into account their spatial nature. The aims of exploratory spatial data analysis include describing geographical distributions, identifying spatial outliers, discovering trends or heterogeneity, regimes of spatial association, validating models. Geographic Information Systems (GIS) are very elaborate cartographic tools but their statistical analysis capabilities are generally limited. When they include statistical techniques, they are very basic tools from descriptive statistics (boxplots, histograms, barcharts, etc.) but none of the state of the art are tools specific to spatial data from geostatistics or spatial econometrics. Openshaw (1994), Anselin (1994) and Anselin (1998) attempt to define the type of exploratory data analysis techniques that GIS should try to incorporate. Anselin (1994) advocates the integration in the GIS of local measures of spatial association, spatial lag plots, spatial lag scatterplots, Moran scatterplots as well as variogram clouds and pocket plots. Wilhelm (1998) and Unwin et al (1998) also argue for the use of local measures of

^{*}thibault.laurent@univ-tlse1.fr

[†]ruiz@cict.fr

[‡]cthomas@cict.fr

spatial association.

The use of the coupling between a map and a statistical graph such as a histogram, a boxplot or a scattermatrix has already been advocated in the literature (see detailed references below). The coupling is the fact that the selection of a zone on the map results in the automatic highlighting of the corresponding points on the statistical graph or reversely the selection of a portion of the graph results in the automatic highlighting of the corresponding points on the map.

Haslett and Bradley (1991) link histograms, double histograms, scatterplot matrices, and varioclouds (see section 4) with the maps using the PASCAL language. Anselin and Bao (1995) implement the methods advocated in Anselin (1994) linking ArcView and Space-Stat. Brundson (1998) implements the scatterplot matrix, the neighbour plot and the angle plot (see section 4) plus some spatial smoothing of maps for trend detection with the XLISP-STAT language. Haining et al (1998) and Wise et al. (2001) develop SAGE, a software system held in the ARC-INFO GIS, with very similar capabilities as those quoted above. Let us mention also the linkage of ArcView and XGobi by Cook et al. (1996) and Symanzik et al. (2000) and the cartographic data visualizer (cdv) of Dykes (1998) based on the Tcl/Tk language. The ArcGIS Geostatistical Analyst extension¹ includes extensive kriging capabilities and exploratory tools but is mainly oriented towards geostatistics and requires the expensive ArcGIS software. Mondrian², written in JAVA, features interactive descriptive tools such as mosaic plots, scatter plots, bar charts, histograms and parallel coordinates plots.

MANET (Unwin et al (1996)), preceded by SPIDER (Haslett et al (1990)) and REGARD, also contains a number of interactive descriptive tools with a central objective of dealing with missing values, but does not contain any tool from spatial statistics. GeoDatm is a free specialized software for spatial data analysis developed by Anselin (2003) and combines maps with statistical graphs dynamically. It offers many functionalities for exploratory data analysis and spatial regression and its main strength is extensive mapping with full linking and brushing possibilities. In contrast (see Anselin and Syabri and Kho Y (2006)), GeoDa is a “closed box” which does not benefit from the tremendous expansion of the R project and has to be considered as an introductory tool to spatial data analysis.

Wise et al. (2001)³ evaluate and compare cdv, MANET, SAGE and SpaceStat.

Lesage and Pace (2004) develop C/C++ code to export polygons and data information from ArcView shapefiles into Matlab and a GUI interface as well as mapping functions to link a map with a histogram and a Moran scatterplot with the possibility of zooming (see also LeSage (1998)).

The need for a more adaptable, comprehensive and unified tool motivated us to start the development of a set of statistical routines adapted to the exploration of georeferenced data called **GeoXp**. It is mainly an exploratory tool for researchers and experienced users in spatial statistics, spatial econometrics, geography, ecology, epidemiology, etc.

¹<http://www.esri.com/software/arcgis/extensions/geostatistical/index.html>

²<http://stats.math.uni-augsburg.de/Mondrian/>

³The visualization of area-based spatial data, In: Case Studies of Visualization in the Social Sciences online. (ed.) D. Unwin, P. Fisher. Advisory Group on Computer Graphics (AGOCG). Available online via <http://www.agocg.ac.uk/reports/visual/casestud/contents.htm>

GeoXp is a stand-alone (free-standing) package independent of a GIS and this is certainly an advantage. Its functions allow coupling between statistical plots and elementary maps as defined before. The routines are user friendly. The user does not need to write a lot of **R** code except for loading the data and calling a function in the command window: after entering some parameters as arguments of the function (usual inputs are at least the name of the variables concerned by the graph), the user needs to execute it. He is then asked to perform the selections by mouse clicking.

The quality of the cartographic display is not a priority for the exploration itself and this is why the emphasis in **GeoXp** is rather in the implementation of spatial statistics tools as numerous and up to date as possible. The final map for a publication can always be produced by a more sophisticated mapping tool if necessary.

GeoXp is based on **R**: this choice of language is motivated by the flexibility of **R** and the existence of many statistical packages developed in this language. The flexibility and adaptability of **GeoXp** comes from the fact that **R** is an open source software and thus the user who is familiar with **R** can customize **GeoXp** with its own routines and benefit from the large amount of modelling tools available in this environment. **GeoXp** includes spatial econometrics as well as geostatistics tools. The advantage over approaches linking a computer engine for statistical computations and a cartographic device such as ArcView is that **GeoXp** is not specific to an operating system and it avoids file transfers. Some unique features are present such as linking a map with a Lorenz curve (see section 2) or with generalized principal components analysis (see section 6). **GeoXp** offers also some rare tools such as the angle plot of Brundson (1998) (see section 4) and the neighbor plot (see section 5).

As far as timing performances are concerned, we ran some tests on an Optiplex GX745 2 duo 2.13GHz under Windows Vista and using the version 1.2 of **GeoXp**. With a function like `histomap`, the time required to make a selection is under 1 second for a data set of size less than 5 000. With a data set of size 10000, the time required is about 1.5 seconds and for size 50 000, it is about 6.5 seconds. For functions which involve selections on couples of points like for example the `moranplotmap` function, the call takes about 19 seconds for size 1000 (resp. 3mn50s for size 2500). However, beyond a data set of size 4000, an allocation memory problem arises and we should be able to improve on this in the next version of **GeoXp**.

Section 2 describes the basic functionalities of **GeoXp** illustrated through an example. In section 3, we present briefly descriptive functions which link simple univariate or bivariate graphs to maps. In section 4, we focus on geostatistics functions such as the `variocloud` and the drift plot while, in section 5, we describe econometrics functions such as the Moran plot and the neighbor plot. The multivariate functions such as generalized principal component analysis are presented in section 6.

For this paper, we have chosen to illustrate only a selection of the different routines and the reader will find a comprehensive list of the **GeoXp** functions in the annex and more illustrations on the web site⁴.

⁴<http://gremaq.univ-tlse1.fr/stat/Chrisweb/SiteGeoXp/Index.htm>

2 Description of the basic functionalities

2.1 General principles

The **GeoXp** functions apply to the analysis of any data set of variables measured at geographical sites or on geographical zones such as cities, counties, countries, etc. called basic spatial units. For each site (for each zone), the data set must contain the cartesian coordinates of the site (respectively of the centroid of the zone). Variables can be continuous or categorical. In the case of geographical zones, one may use additionally the coordinates of polygonal spatial contours to improve the map quality and to help identifying locations.

As far as format is concerned, any format that can be imported in R can be used as long as it contains the geographical coordinates. For example one can import a shapefile format from ArcView using the function `readShapePoly` of the R package `maptools` or the function `readOGR` of the R package `rgdal`, and a MIF/MID format from MapInfo using the function `readOGR` of the R package `rgdal`. The geographical contours have their own format (coordinates of vertices separated from one unit to the next by the missing value symbol NA). The **GeoXp** functions `polylist2list` and `spdf2list` allow to convert respectively the `maptools`, `sp` and `rgdal` formats into the format of **GeoXp** contours. The names of the main **GeoXp** functions reflect their functionality and always end with “map” (example: `moranplotmap`, `scattermap`). A call to a **GeoXp** function generally opens three windows: two graphical R windows for the statistical graph and the map respectively, and one Tk window for the menu. The user then selects on the menu the graph on which he wants to select points first. This graph then becomes active and the selection by mouse clicking begins.

For selecting the points either on a statistical graph or on a map, the user can choose between selecting individual points (centroids) or selecting points inside a given polygon. For the selection on the statistical graph, there are several cases. In the case of a histogram or a bar plot, it is possible to select several non necessarily contiguous bars. In the case of a density plot, one can select one or several intervals on the x-axis by mouse clicking or by specifying its endpoints. In the case of a `boxplot`, one can select outliers or inter-quartiles ranges. In the case of the Lorenz curve, it is possible to select either a given percentage of spatial units on the first axis or a given threshold value of the variable.

The selection of an already selected unit delete its selection. Upon exit, each function returns a zero-one selection vector (one for selected units and zero for the rest) allowing further analysis of the selection’s characteristics.

Selected units are marked with a different color or alternatively with a different symbol and an option allows a chosen label to be printed too. Polygons representing the boundaries of the spatial units can be added easily if available. Names of the variables can be specified for use in the graph axes labels.

As in cartographic devices, proportional symbol maps can be produced by adding bubbles to the map, proportional to a given variable with a legend for their size.

For most functions, an additional statistical graph can be added in a given choice list, using variables specified in an option. This additional graph is only interactive in one direction though: selections made on the first graph or the map will appear on this

additional graph but one cannot select from the additional graph.

2.2 Example

Using the data set `eire` from **spdep**, FIG. 1 displays a scatterplot of percentage in value terms of gross agricultural output of each county consumed by itself versus arterial road network accessibility in 1961, with a conditional quartile curve. An additional graph shows the bar plot of the variable “Pale”. A selection on the map of the eastern districts, in red on the map, shows on the scatterplot that these districts have low values for the two numeric variables. Moreover, the barplot shows that they belong mainly to the pale one. To underline the simplicity of the code, you will find below the code used to load the data set and produce these plots for version 1.3 of **GeoXp**.

```
data(eire)
eire.contours<-polylist2list(eire.polys.utm)
scattermap(eire.coords.utm$V1,eire.coords.utm$V2,eire.df$ROADACC,
eire.df$OWNCONS, listvar=eire.df,listnomvar=names(eire.df),
carte=eire.contours,xlab="ROADACC",ylab="OWNCONS",
reg.line=TRUE,quantiles=0.95,cex.lab=0.8)
```

3 Descriptive functions

The descriptive functions are called `barmap`, `boxplotmap`, `histomap`, `densitymap`, `histobarmap`, `dbلهistomap`, `dbلهdensitymap`, `polyboxplotmap`, `ginimap` and `scattermap`.

In the case of a simple histogram, the selection of some bars of this histogram will show the corresponding zones on the map, which is just a more elaborate variant of the previous tool as in Haslett et al (1990). In the other direction, a selection of a subregion of the map produces the subhistogram of the distribution of the variable in this subregion. Since the goal is then to compare the distribution of the variable on the whole map to its subdistribution on the selected zone, it is not optimal to use histograms based on counts as most packages do, so we have introduced an alternative function allowing the user to produce two kernel density estimators instead of two histograms. The user can choose the bandwidth or use a default option for this choice. He can also change the initial bandwidth selection with a ruler displayed in the Tk window, resulting in an automatic updating of the graphs. For discrete variables, it is also possible to link a bar plot to the map.

When the statistical graph is a simple boxplot, only the selection on the boxplot is implemented and allows the user to display the zones corresponding to lower or upper quartiles as well as to outliers (as in Haining et al (1998)). The same information is conveyed by choropleth maps in a GIS.

For a couple of variables, a double histogram or a double kernel density estimator can be graphed and linked to the map. Selection is then possible on the map as well as on one of the histograms or density graphs. On FIG. 2, the `dbلهdensitymap` function displays the arterial road network accessibility in 1961 and the percentage in value



Figure 1: Scatterplot of percentage in value terms of gross agricultural output of each county consumed by itself versus arterial road network accessibility in 1961 and barplot of the variable pale

terms of gross agricultural output of each county consumed by itself. A selection of the counties with more than 6300 and less than 3300 arterial road network accessibility on the first density produces on the second plot the graph of the corresponding subdensity for the percentage in value terms of gross agricultural output. The subdensity looks different.

```
data(eire)
eire.contours<-polylist2list(eire.polys.utm)
obs<-dbledensitymap(eire.coords.utm$V1,eire.coords.utm$V2,
eire.df$ROADACC,eire.df$OWNCONS,xlab=c("arterial road networkaccessibility",
"percentage in value terms of gross agricultural output"),
carte=eire.contours,kernel="normal",listvar=eire.df,
listnomvar=names(eire.df))
```

A simple scatterplot of a couple of variables can also be linked to the map and selection is again possible in both directions as in Brundson (1998). A kernel smoother can be added to the scatterplot for convenience with a flexible choice of bandwidth. An option allows the user to overlay conditional quantile estimates instead of the kernel smoother which estimates the conditional mean, thus allowing a more precise exploration of the cloud when one is interested for example in the extreme rather than the average behaviour. The possibility of linking the map with a Lorenz curve allows the study of the geographical component of the concentration or inequality measured by the Gini index (see Gastwirth (1972)). The Lorenz curve is a scatterplot of the relative mass of a given variable X due to the sites with a value of X less than or equal to x versus the relative frequency of such sites. The Gini index (area between the Lorenz curve and the diagonal of the unit square) measures the inequality in the distribution of X .

The selection of a given frequency F on the frequency axis results in the printing of the corresponding relative mass G on the other axis, the corresponding quantile (value x such that the cumulative distribution function of X at x equals to F) as well as the selection of the corresponding points on the map (spatial units such that X is less than or equal to x). For example FIG. 3 shows the Lorenz curve of the total personal income and the bar plot of variable pale. The Lorenz curve, which is away from the diagonal with a Gini index of 0.47, shows that a small number of counties concentrate a large number of income.

```
data(eire)
eire.contours<-polylist2list(eire.polys.utm)
obs<-ginimap(eire.coords.utm$V1,eire.coords.utm$V2,eire.df$INCOME,
listvar=eire.df,listnomvar=names(eire.df),carte=eire.contours,col="orange",
label=as.character(1:26),cex.lab=0.7,xlab="Income",ylab="Total Income")
```





Figure 2: Densities of the arterial road network accessibility and the percentage in value terms of gross agricultural output. Selection corresponds to the dashed curve on the last plot.

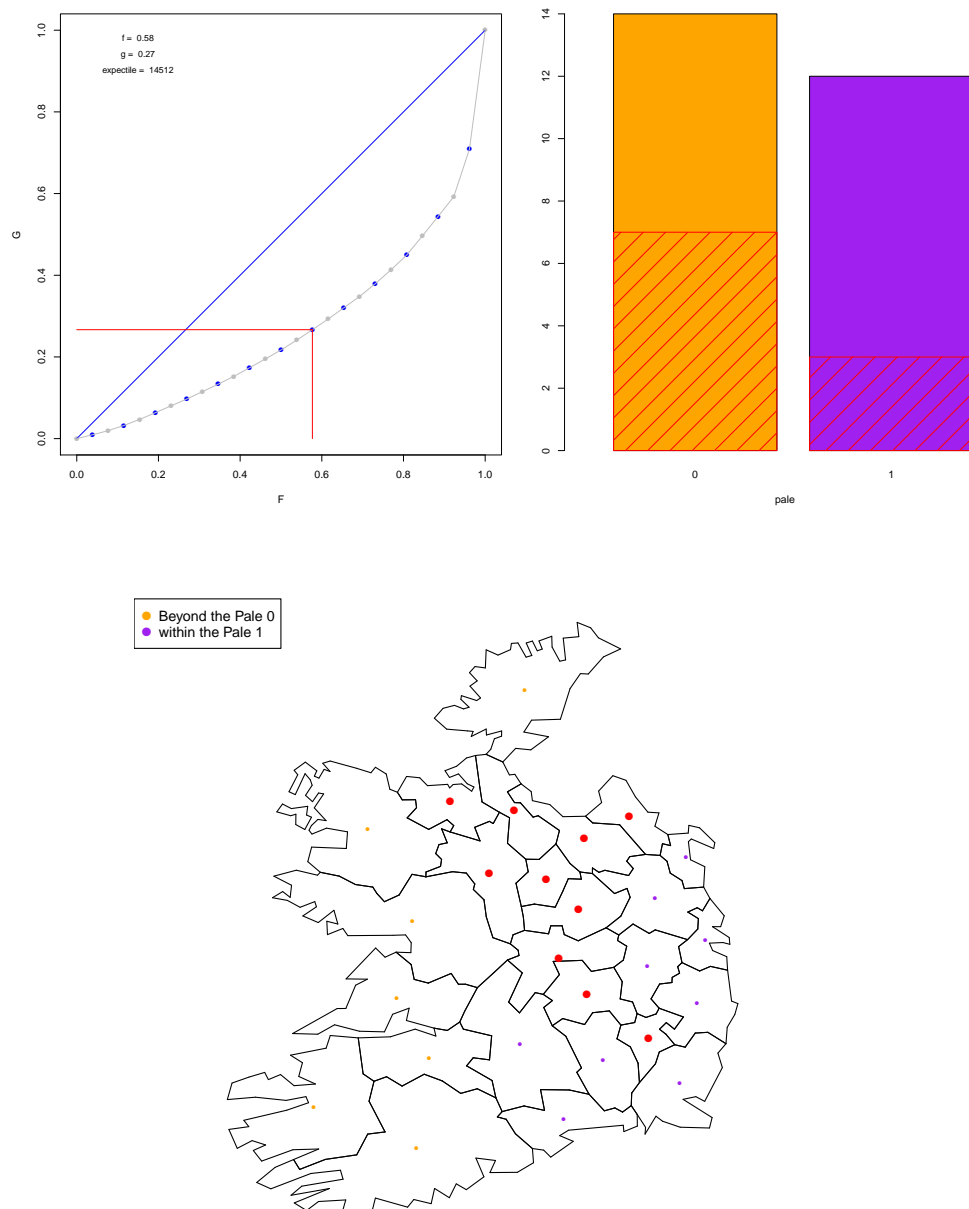


Figure 3: Lorenz curve and Gini index for the number of students: selection of the first 38% of counties sorted by increasing number of income.

4 Geostatistic functions

The geostatistic functions are called `angleplotmap`, `driftmap` and `variocloudmap`. As in Cressie (1993), in order to examine trends in one variable, **GeoXp** creates a grid of a given fineness and for each square of the grid computes the mean of the variable for all basic units intersecting the square. It is then easy to produce row and column means and medians, and plot the row means and medians to the right of the map as well as the column means and medians below the map. No selection is possible here but the study of the variation of the row means with longitude and column means with latitude brings out the north-south and east-west trends if present. An option allows the user to rotate the map by a given angle and thus study trends in any direction. Discrepancies between means and corresponding medians detect the presence of outliers in a given row or column. Generally, the user may have no prior idea of the directions of the main trends. It is then interesting to use an angle plot prior to the trend graphic (see Brundson (1998)) that may reveal unknown spatial heterogeneity. The angle plot implemented here is a scatterplot of the square root of the absolute differences between the values of the variable at two given zones as a function of the bearing of a line joining the centroids of the two zones (in radians or degrees).

On FIG. 4, the selection of the couples of sites with a bearing of $3\pi/4$ radians and with large absolute differences in the topsoil copper concentration (see `help(meuse)`) reveals a disparity between the sites near the river and the sites far to the river. It is interesting to train oneself in the interpretation of angle plots by applying them to deterministic trends such as latitude and longitude.

```
data(eire)
data(meuse)
data(meuse.riv)
obs<-angleplotmap(meuse$x,meuse$y,meuse$copper,col="green",
quantiles=0.9,listvar=meuse,listnomvar=names(meuse),
xlab="topsoil copper concentration, ppm. ",pch=7,carte=meuse.riv)
```

The variogram cloud is another tool inspired by geostatistics to study autocorrelation (Chauvet (1982)). It is a simple scatterplot of the half square of the difference between the value of the variable at two locations against the distance between these points. As in Haslett and Bradley (1991), outliers may be mapped by highlighting those points on this graph which have a high value of the second coordinate. An option allows the user to overlay an empirical variogram or a smooth of this scatterplot thus estimating the variogram function (with the possibility of a robust alternative (Cressie (1993))). This option is important to represent the bulk of the cloud since, because of the high number of couples of positions with a low value on the vertical axis, it is often desirable to combine this with another option allowing to represent only those couples with a value above a chosen threshold percentile (conditional on the value of the horizontal coordinate). Finally, another option allows to concentrate on couples of points in a given direction (with a tolerance) and overlay a directional variogram. On FIG. 5, one can see that high differences in the topsoil lead concentration appear





Figure 4: Angle plot for topsoil copper concentration: selection of large absolute differences for a given angle.

between sites located near the river (with a high concentration) and sites distant to the river. A threshold of 75 percent has been chosen for representing the points.

```
data(meuse)
data(meuse.riv)
obs<-variocloudmap(meuse$x,meuse$y,meuse$zinc,
quantiles=0.75,listvar=meuse,listnomvar=names(meuse),
xlim=c(0,2000),ylim=c(0,500000),pch=2,ylab='semivariance',
carte=meuse.riv[c(21:65,110:153),])
```

5 Econometric functions

The econometric functions are called `moranplotmap` and `neighbourmap`. These two functions use neighborhood or weight matrices of several types, which can be constructed using the auxiliary function `makeneighborsw` or using other similar functions in the R package **spdep** of Bivand⁵. The function `makeneighborsw` creates a weighting scheme based on a given number of nearest neighbors or a given distance threshold or using these two criteria together. The function `normW` performs row standardizing of these matrices if necessary.

To examine spatial autocorrelation, given a spatial binary weight matrix (Bavaud (1998)) containing information about the neighboring relationships of the basic spatial units, one can simply make a scatterplot of the value of the variable on each unit versus the value of the same variable on the neighboring units (neighbor plot). Points far away from the diagonal on this plot identify local outliers and selection is again possible on the plot as well as on the map. When a point is selected on the map, its neighbors are shown connected by lines to this point. We illustrate this method with the 49 neighbourhoods in Columbus on the variable `crime` (see `help(columbus)`). We draw on FIG. 6 a neighbor plot with a weight matrix based on 4 nearest neighbors. This graph shows some amount of spatial autocorrelation with points not too far from the diagonal and we notice the asymmetry due to the corresponding asymmetry of the weight matrix. The selection of the counties with the smallest value of crime reveals that their neighbors have higher values of crime.

This tool is also interesting for investigating a chosen spatial weight matrix as is shown on FIG. 7. For the same 4 nearest neighbors matrix, the couples of points with a large difference in longitude are selected. Large distances between neighbors may arise for some weight matrices (for example those based on a Delaunay triangulation) and this type of graph points out these inappropriate neighbors.

```
columbus <- readShapePoly(system.file("etc/shapes/columbus.shp",
package="spdep")) [1])
```

```
list.OLD <- spdf2list(columbus)
```

⁵see the site <http://cran.r-project.org/web/views/Spatial.html>





Figure 5: Variocloud plot for topsoil zinc concentration, ppm: selection of large absolute differences for a given little distance. The empirical variogram is represented by a continuous line, the robust version is represented by a dashed line

```

contours.OLD <- list.OLD$poly
x.OLD <- list.OLD$X
y.OLD <- list.OLD$Y

W <- makeneighborsw(x.OLD , y.OLD,method="neighbor",4)
obs<-neighbourmap(x.OLD , y.OLD, columbus@data$CRIME, W,
carte=contours.OLD,criteria=(columbus@data$CRIME>mean(columbus@data$CRIME)),
listvar=columbus@data, listnomvar=names(columbus@data),
xlab='crime',id=TRUE, col="grey", pch=4)

columbus <- readShapePoly(system.file("etc/shapes/columbus.shp",
package="spdep")[1])

list.OLD <- spdf2list(columbus)
contours.OLD <- list.OLD$poly
x.OLD <- list.OLD$X
y.OLD <- list.OLD$Y

W <- makeneighborsw(x.OLD , y.OLD,method="neighbor",4)
obs<-neighbourmap(x.OLD , y.OLD, x.OLD, W,carte=contours.OLD,
criteria=(columbus@data$CRIME>mean(columbus@data$CRIME)),
listvar=columbus@data, listnomvar=names(columbus@data),
xlab='longitude',id=TRUE, col="grey", pch=4)

```

A simple scatterplot linked to the map has potentials for more advanced investigations if one applies it to transformations of the raw variables. For example, for a centered variable X and for a given weight matrix W , the classical Moran scatterplot (Anselin (1995)) is the scatterplot of the spatial lag variable WX against X . The function `moranplotmap` of **GeoXp** links this scatterplot to the map and exhibits the regression line whose slope is the Moran index indicating the strength and nature of the spatial autocorrelation. But the observation of the cloud itself conveys more information about changes in spatial autocorrelation regimes and also outliers (see Anselin (1995) for details). The selection of each quadrant on the plot exhibits zones of positive and negative autocorrelation on the map. An option allows the computation of the local Moran statistic for the selected points. The p-value of the Moran gaussian test for spatial autocorrelation is displayed by default and the p-value of the permutation test based on a chosen number of simulations can also be obtained.

FIG. 8 displays the Moran scatterplot of the House sales price in Baltimore in 1978 (see `help(baltimore)`). The Moran index of 0.51 has a p-value of 0.0001 for the gaussian and the permutation tests (with 500 permutations). The selection of the first quadrant, corresponding to districts with a House sales price higher than average as well as their neighbors, shows that these are mainly located in the North.

```

data(baltimore)
W <- makeneighborsw(baltimore$X,baltimore$Y,method="neighbor",4)
moranplotmap(baltimore$X,baltimore$Y,baltimore$PRICE,W,

```





Figure 6: Neighbor plot for crime: selection of small values of crime.





Figure 7: Neighbor plot for longitude: selection of large differences in longitude.


```
label=baltimore$STATION,criteria=(baltimore$PRICE>mean(baltimore$PRICE)),
listvar=baltimore,listnomvar=names(baltimore),flower=TRUE,
locmoran=TRUE,cex.lab=0.7,pch=16)
```

6 Multivariate functions

GeoXp includes the possibility of linking the results of a clustering algorithm (k-means from the **R** function `kmeans` or hierarchical clustering from the **R** function `hclust`) to the map. We suggest using a preliminary dimension reduction technique such as principal components analysis to produce bivariate plots of relevant linear combinations of the variables linked to the map. Exploratory analysis becomes rapidly cumbersome with large numbers of variables hence it is essential to use devices that select interesting projections of the data. The multivariate functions are called `clustermmap` and `pcamap`. The function `pcamap` implements the generalized principal components analysis (PCA) as it is described in Caussinus et al. (2003). Note that using the link between map and scatterplot, users can rapidly customize **GeoXp** to any other dimension reduction method.

In the case of usual PCA, which is a specific case of generalized PCA, one can do a scatterplot of the projection of the cloud for any couple of factorial axes and one can link it to the map. If outliers or groups appear on one of these plots, it is interesting to locate them on the map and explore their relative spatial position. Reversely, the positions on the scatterplot of a selected subregion of the map may provide information about its specificities with respect to the principal axes. The interpretation of the principal axes is guided by the representation of the variables on a separate non interactive plot. In the case of standardized PCA, correlations between the original variables and the principal components are plotted inside a correlations circle.

FIG. 9 illustrates this method on the following set of seven variables: housing value, household income, residential burglaries and vehicle thefts per thousand households in the neighborhood, open space in neighborhood, percentage housing units without plumbing, distance to CBD. The right top plot of FIG. 9 shows that the first axis is positively correlated to residential burglaries and vehicle thefts per thousand households in the neighborhood. The second axis is negatively correlated to the open space in neighborhood. The labels on the axes indicate the percentages of inertia associated with each principal axis (which is nearly 66% for the first two axes of this example) while the percentages for the variables indicate their quality of representation on the principal plane. Three counties have been selected on the extreme right bottom part of the scatterplot. The quality of their representation on the factorial plane is given on the scatterplot and is high (more than 76% of their norm is accounted for by the first two principal coordinates). They differ from other counties because they have low level of crime and low level of income and are located in the center.

```
example(columbus)
coords <- coordinates(columbus)
cont<-spdf2list(columbus)$poly
```

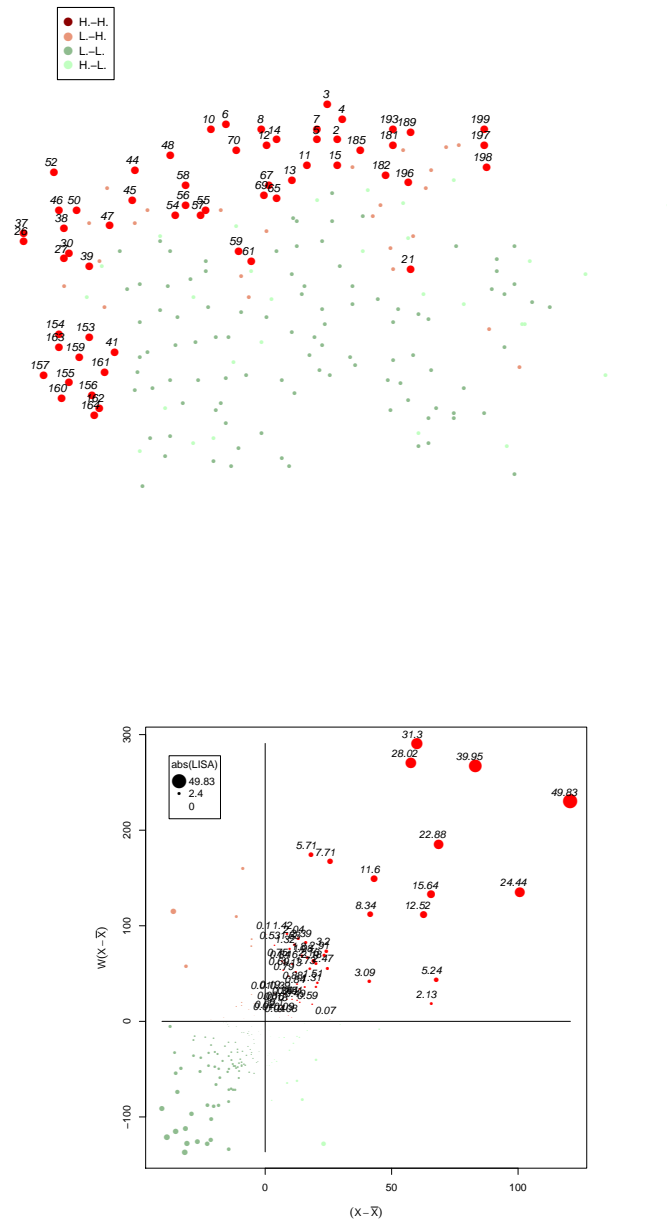


Figure 8: Moran scatterplot of the House sales price in Boston : selection of first quadrant

```

obs<-pcamap(coords[,1], coords[,2],columbus@data[,6:12],
label=columbus@data$NEIGNO, carte=cont,listvar=columbus@data,
qualproj=TRUE,listnomvar=names(columbus@data),
name.dataset=names(columbus@data[,6:12]),
cex.lab=0.7, pch=6, col="purple")

colmbs> columbus <- readShapePoly(system.file("etc/shapes/columbus.shp",
colmbs+ package="spdep")[1])

colmbs> col.gal.nb <- read.gal(system.file("etc/weights/columbus.gal",
colmbs+ package="spdep")[1])

```

7 Conclusion

The project **GeoXp** started before 2000 and has known many different versions. A 1998 Matlab version (working with Matlab 6) still is on the site of the econometrics toolbox of LeSage and contains tools which have not yet been translated to R. It is now an R package downloadable from CRAN. For applications oriented purposes, this set of routines has also been translated into C++ in the context of a contract with the Midi-Pyrénées region council. There are a lot of new tools that we plan to include in **GeoXp** such as a weighted version of ginimap, a Moran scatter plot for residuals of a OLS model, a 3-d version of the scattermap, a micromap display (see Symanzik et al. (2000)), an Apleplotmap based on Li et al (2007), etc. More structural changes will involve in the near future the use of R classes for a larger compatibility of formats with the **sp** R package and also of sparse matrices for handling large weight matrices.

Acknowledgments: We thank the team of students who participated in the writing of the several versions of **GeoXp**. We thank as well our faculty colleagues from the university of Toulouse I for their research assistance and many colleagues for their remarks (E. Malin, I. Heba, J. LeSage, R. Bivand, J. Symanzik., etc.)

8 Annex

GeoXp includes two types of functions: the main functions and the auxiliary ones. The auxiliary functions are just routines called by the main functions. The list of main functions can be found below.

- `angleplotmap`: links a map and an angle plot (only the angle plot is active).
- `barmap`: links a map and a bar plot.
- `boxplotmap`: links a map and a box and whiskers plot.
- `clustermmap`: links a map and a bar map of a clustering variable (kmeans method).
- `dbledensitymap`: links a map and two density estimators.





Figure 9: Principal components analysis: selection of the counties on the right part of the first principal plane.

- `dblehemistomap`: links a map and two histograms.
- `densitymap`: links a map and a density estimator.
- `driftmap`: this function is meant for detecting trends (non interactive)
- `ginimap`: links a map and a Gini plot (Lorenz curve).
- `histobarmap`: links a map to an histogram and a bar plot.
- `histomap`: links a map and an histogram.
- `moranplotmap`: links a map and a Moran scatterplot.
- `neighbourmap`: links a map and a neighbor plot (scatterplot of variable against variable for the neighboring sites)
- `pcamap`: links a map and a scatterplot of principal axes of Principal Components Analysis
- `polyboxplotmap`: links a map and a box and whiskers plot.
- `scattermap`: links a map and a two-dimensional scatterplot.
- `variocloudmap`: links a map and a variogram cloud (only the variogram cloud is active).

References

- Anselin, L., 1994. Exploratory spatial data analysis and geographic information systems. in M. Painho ed., Luxembourg: Eurostat:45-54.
- Anselin, L., 1995. Local indicators of spatial association-LISA. *Geographical Analysis*, 27: 93-115.
- Anselin, L and Bao, S (1995). Exploratory spatial data analysis linking SpaceStat and Arcview. Regional Research Institute, West Virginia University.
- Anselin, L (1998). Exploratory spatial data analysis in a geocomputational environment. in Paul Longley, Sue Brooks, Bill Macmillan and Rachel McDonnell (eds), New York: Wiley.
- Anselin, L (2003). GeoDa 0.9. User's guide. Urbana-Champaign, IL: Spatial Analysis Laboratory (SAL), Department of Agricultural and Consumer Economics, University of Illinois.
- Anselin, L, Syabri I, Kho Y (2006). GeoDa: An Introduction to Spatial Data Analysis. *Geographical Analysis*, 38 (1), 522.
- Bavaud, F (1998). Models for spatial weights: a systematic look. *Geographical Analysis*, 30 (2): 153-171.

- Bessy-Pietri P., Sicamois Y, (2001). Le Zonage en Aires Urbaines en 1999. 4 millions d'habitants en plus dans les aires urbaines. INSEE-Première, 765 : 1-4.
- Brundson, C (1998). Exploratory spatial data analysis and local indicators of spatial association with XLISP-STAT. *The Statistician*, 47: 471-484.
- Caussinus, H., Fekri, M., Hakam, S., Ruiz-Gazen A. (2003). A monitoring display of multivariate outliers. *Computational Statistics and Data Analysis*, 44: 237-252.
- Chauvet P. (1982). The variogram cloud. *Proceedings of the 17th APCOM International Symposium*, Golden, Colorado
- Cook, D., Masure, JJ., Symanzik, J. and Cressie, N. (1996). Dynamic graphics in a GIS: exploring and analysing multivariate spatial data using linked software. *Computational Statistics*, 11: 467-480.
- Cressie, N (1993). *Statistics for spatial data*. Wiley.
- Dykes, J (1998). Cartographic visualization: exploratory spatial data analysis with local indicators of spatial association using Tcl/Tk and cdv. *The Statistician*, 47:485-497.
- Gastwirth, J L, (1972). The Estimation of the Lorenz Curve and Gini Index. *The Review of Economics and Statistics*, MIT Press. vol. 54(3), pages 306-16, August.
- Haining, R., Wise S., and Ma, J. (1998). Exploratory spatial data analysis in a geographic information system environment. *The Statistician*, 47: 457-469..
- Haslett, J., Wills, G. and Unwin, AR. (1990). SPIDER - an interactive statistical tool for the analysis of spatially distributed data. *International Journal of Geographical Information Systems*, 4(3), 285-296.
- Haslett, J., Bradley, R., Craig, P., Unwin A. and Wills G. (1991). Dynamic graphics for exploring spatial data with application to locating global and local anomalies. *The American Statistician*, 45: 234-242.
- LeSage J. and Pace K., (2004). Arc Mat, a Toolbox for Using ArcView Shape Files for Spatial Econometrics and Statistics. in *Geographic Information Science, Proceedings of the Third International Conference*, Max J. Egenhofer, Christian Freksa and Harvey J. Miller (eds.), pp. 179-190.
- LeSage J. (1998). *Spatial Econometrics*. <http://www.econ.utoledo.edu/faculty/lesage/lesage.html>.
- Li H., Calder C. A., and Cressie N., (2007). Beyond Moran's I: Testing for spatial dependence based on the SAR model. *Geographical Analysis*, 39, 357-375.
- Openshaw, S (1994). What is a gisable spatial analysis. in M. Painho (ed.), *Luxembourg: Eurostat*: 36-44.

- Symanzik J., Cook D., Lewin-Koh N., Majure J.J. and Megretskaja I. (2000). Linking Arcview and XGobi: Insight behind the front end. *Journal of Computational and Graphical Statistics*, 9, pp. 470-490.
- Unwin A.R., Hawkins, G., Hofman H. and Siegl, B. (1996). Interactive graphics for data sets with missing values - MANET. *J. Comput. Graph. Statist.* 5: 113-122.
- Unwin, A, Unwin, D (1998). Exploratory spatial data analysis with local statistics. *The Statistician*, 47: 415-421.
- Wilhelm, A., Steck, R. (1998). Exploring spatial data with interactive graphics and local statistics. *The Statistician*, 47: 423-430.
- Wise S, Haining R and Ma J (2001). Providing spatial data analysis functionality for the GIS user: the SAGE project. *International Journal of Geographical Information Science*, 15 (3), pp. 239-254.