PERGAMON

Computers, Environment and Urban Systems
25 (2001) 53–68

Computers, Environment and Urban Systems

www.elsevier.com/locate/compenvurbsys

The comap: exploring spatial pattern via conditional distributions

C. Brunsdon *

Department of Geography, University of Newcastle Upon Tyne, Newcastle Upon Tyne NE1 7RU, UK

Accepted 18 October 2000

Abstract

An exploratory data analysis tool termed the *comap* is introduced. This is essentially a geographical variant of the *coplot*, an exploratory graphical method to investigate the relationship between a pair of variables conditioned on a third variable (and perhaps also a fourth). In the comap, the first pair of variables represent geographical location, and the graphical technique is adapted to reflect this. After the concept of the comap is discussed and an example is given, computational aspects are considered. The paper concludes with a brief discussion. © 2001 Elsevier Science Ltd. All rights reserved.

1. Introduction

Investigating the relationship between variables is a fundamental aspect of data analysis. One common approach is to consider the distribution of one variable conditioned on another one. A typical example of this is bivariate linear regression. Although this technique is frequently viewed as a means of *predicting* some variable Y from another variable X, in actuality one is attempting to find the probability distribution for the variable Y *given* X — the so-called predicted Y values are simply estimates of the conditional mean value of Y for a given X value.

However, in many situations it is not just the mean of Y that is of interest. For instance, knowing something about the standard deviation of Y for a given X is also important, since this gives some idea of the reliability of the mean of Y given X as a predictor of Y. More complex information about the distribution of Y given X is also useful. The relationship between the two variables can be better understood by answering questions such as "Does the variance of Y change for given X values?"

* Tel.: +44-191-2225421; fax: +44-191-222643.
E-mail address: chris.brundsdon@ncl.ac.uk

PII: S0198-9715(00)00042-9

and "Is the distribution of Y multimodal for any given X values?". It is usually helpful to have tentative answers to questions like these before any formal mathematical or statistical modelling is attempted. For this reason, it is generally helpful to investigate such phenomena using exploratory graphical techniques as a first stage of data analysis. Some excellent examples supporting this principle are provided by Cleveland (1993) and Tukey (1977). For the "Y conditioned on X" situation discussed above, basic scatterplots are often powerful tools.

The above reasoning is no less true when applied to geographical information — although here the Y variable may be a two-dimensional map entity rather than a simple continuous variable. For example, Y might be the location of a weather station and X the average annual rainfall. One may then consider conditional distributions of Y, and pose questions such as "given that a rain gauge has recorded an average annual rainfall of Y, where is it likely to be located?". As before, one would need to consider more than a single statement of location — there may be several areas having an annual rainfall of around Y, or a very large area over which Y would be a representative figure. As with the non-geographical situation, issues of variability and multimodality are important in understanding the data. Also, in accordance with the non-geographical situation, the use of exploratory graphical techniques provide an important tool for initial data investigation. However, in the geographical situation it is not obvious what graphical tools could be used. This paper proposes such a tool, based on the concept of the *coplot* (derived from the term COnditional PLOT). The coplot is essentially a multi-panelled plot where each panel (or small graph) is constructed using data selected conditionally on some variable. In this paper, the *comap* is introduced. The name here is derived from the term Conditional MAP.

Since coplots are not used often in geographical, planning or environmental literature, this paper will begin with a review of the coplot together with a practical example of its use. Subsequently, the extension of these ideas to the notion of comap will be discussed, and applied to the earlier example. Following this, practical computational issues will be considered.

2. The coplot: a review

The notion of a coplot utilises the principle of using 'small multiples' of diagrams to highlight differences in pattern, as discussed in Tufte (1990). In particular, coplots show changes in the relationship between the variables Y_1 and Y_2 conditioned on a third variable X. An example is given in Fig. 1.

The data set used here is a subset of mean annual rainfall figures recorded from several rain gauges throughout the UK. Here, 600 of these observations, all from Scotland, were selected. These data were originally compiled by the Climate Research Unit, University of East Anglia, where they were made available for academic research. In this study, four variables were recorded: the easting and northing and height above sea level of each gauge, and the mean annual rainfall. For the remainder of this paper, these variables will be referred to as east, north, height and

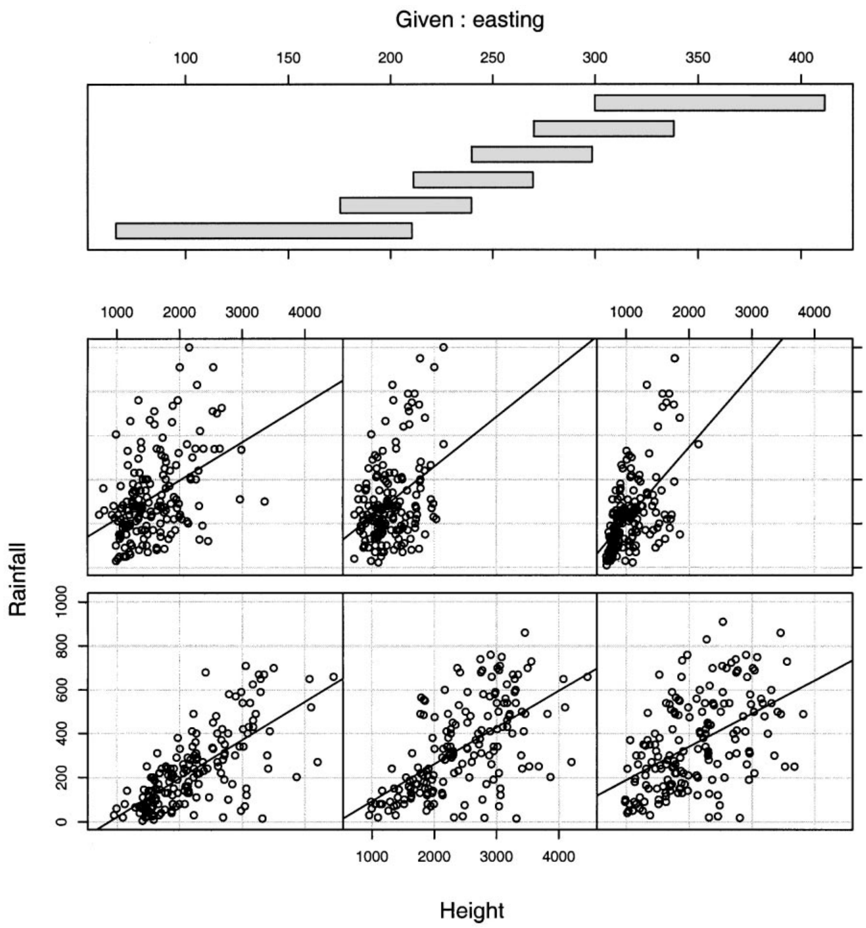

Fig. 1. Coplot of height vs. rainfall given easting of gauge. Reading from left to right, the bottom scatterplots are numbered 1 to 3, and the top are numbered 4 to 6 in order of ascending easting category midpoint.

rainfall, respectively. In relation to the more generic terminology for variables used in earlier discussion, in Fig. 1 we have height = Y_1, rainfall = Y_2 and east = X. At this stage, north will not be considered.

The plot contains seven panels — a broader one at the top and six below this comprising a two-by-three matrix of scatterplots. Each of the six panels in the lower matrix contain only data whose X values fall within a specific range. The ranges are selected so that an identical number of observations appear in each scatterplot, and also so that the ranges of X overlap. The top panel (which Cleveland terms the *given* panel) shows the six ranges graphically. Suppose we number the scatterplots as below:

4	5	6
1	2	3

Then, as the scatterplots progress from 1 to 6, the midpoint of the range of X-values increases. Since in this example X = east, scatterplots 1 to 6 show the change in relationship between height and rainfall as a geographical 'moving window' progresses from east to west. Certain patterns are immediately apparent: most prominently, the slope of the relationship between the two quantities becomes steeper as one moves west. This trend is made easier to see in the scatterplot panels by the addition of least squares regression lines. Other features are also noticeable — the variability of the intercept point of the regression line is perhaps also of interest. These patterns are consistent with the phenomenon of *orographic enhancement* which arises from a number of mechanisms that serve to increase the levels of rain precipitation over and near to hills in some areas (Pedgley, 1970; Sawyer, 1956). The coplot is useful in allowing one to explore exactly how the relationship between elevation and precipitation varies across geographical space. In this instance, prior knowledge of the topography of Scotland lead to the choice of east as the conditioning variable.

This example demonstrates many of the useful features of coplots. Firstly, as a concept they are highly intuitive; provided one is already familiar with the concept of a scatterplot, it is a small step of the imagination to consider a sequence of scatterplots based on 'windows' of a third conditioning variable. Secondly, it is relatively easy to extend the basic coplot. Here, for example, regression lines were added in order to explore the changing relationship between height and rainfall. Other modifications are also possible — in Fig. 2 scatterplots of absolute residual (that is, the positive difference between fitted and actual rainfall values) against the rank of height are used. The replacement of heights with their ranks will be justified at the end of this section. Here we see that in all six scatterplot panels the magnitude of the residuals increases with height — and arguably the effect is more pronounced in the extreme east and extreme west scatterplots. This calls into doubt the assumption in simple linear regression that the error term variance is constant and perhaps suggests that even the *way* it differs from constancy varies geographically. In addition, we note that in all six panels the greatest *level* of variation is about the same, but recall from Fig. 1 that the relationship with height differs — the panels toward the west attain the greatest level of residual variance at lower heights.

Since the ranges of the conditioning variable overlap along the scatterplots, each point will be shown on more than one scatterplot — this is particularly helpful if only a small dataset is being used. In this case partitioning the observations between scatterplots without duplication would result in very few points appearing in individual panels, making trends or patterns difficult to spot.

A final useful property of the coplot is the fact that each scatterplot has the same number of points — or as close to this as the sample size, number of scatterplot panels and degree of overlap allows. Here, each scatterplot has 171 points. In an informal sense, this means that each graph is subject to roughly the same degree of sampling uncertainty, and that no graph is notably more prone to 'freak' results or outliers than any of the others — thus if some panels appear to show more varied results than others this is more likely to be due to genuine differences in the relationship between Y_1 and Y_2 as X varies. This last point is perhaps best demonstrated

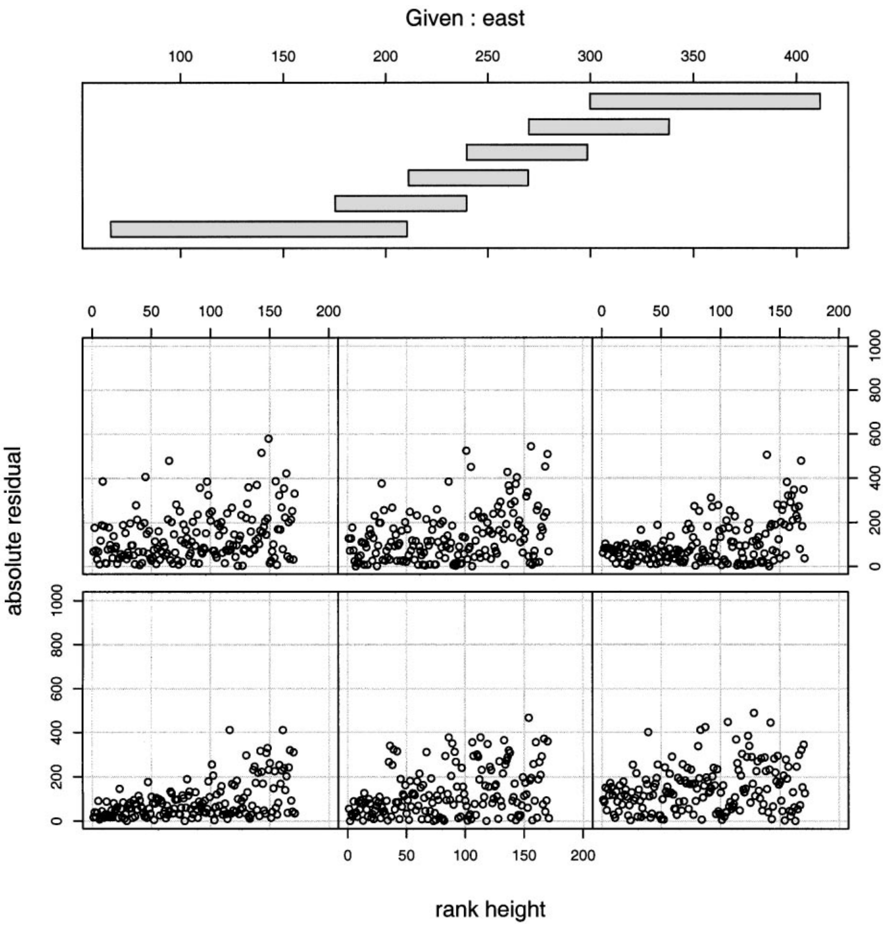

Fig. 2. Coplot of rank height vs. regression residual given easting of gauge. Reading from left to right, the bottom scatterplots are numbered 1 to 3, and the top are numbered 4 to 6 in order of ascending easting category midpoint.

visually: Fig. 3 shows two sets of pairs of points (Y_1, Y_2). The only difference between the two samples is that there are 1000 (Y_1, Y_2) pairs in the left hand panel, but only 100 on the right — those on the right are a random subsample of those on the left. The points on the left appear to be spread over a larger area, and also a seem to have greater proportion of outliers, but in fact these observations are entirely attributable to the difference in sample size, as both were generated from the same population distribution. Some of the patterns observed in scatterplots are an artifact of the number of points used, and keeping the number of points in each panel in the comap the same is a method of controlling for this.

An interesting variant of this problem occurs when using scatterplots to investigate the variance of the Y_2 variable as Y_1 varies (or *vice versa*). If the marginal density of the Y_1 variable changes along the axis, then the effect is that of having

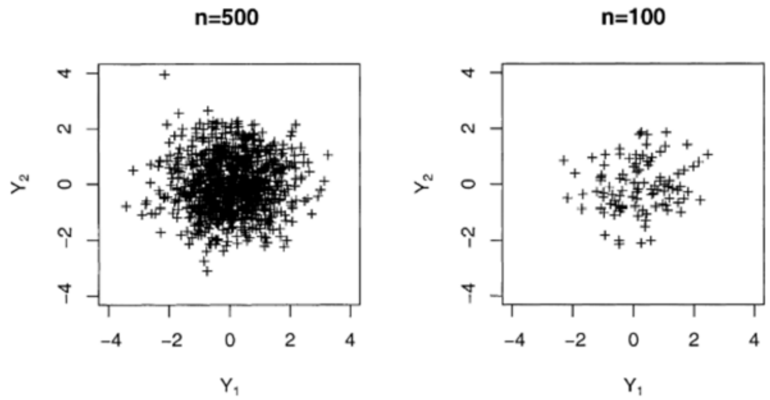

Fig. 3. Effect of sample size on observed scatterplot patterns.

more points in some Y_1-ranges of the plot than in others. Thus, even if the variance of Y_2 is constant given Y_1, it may appear to alter. This would result in a misdiagnosis of heteroskedasticity. A good example of this is given in (Cleveland, 1979); see also Brunsdon, Fotheringham and Charlton (1999). One way of overcoming this problem is to replace Y_1 with the rank of Y_1 in the scatterplot. Since ranks will be spread evenly between 1 and n, the marginal density of the new variable will remain constant — and observed heteroskedasticity becomes much less likely to be a 'false alarm'. This is the reason that ranks were used for the height variable in Fig. 2.

3. The comap: extending the coplot

In this section, ways in which the coplot may be adapted to provide more appropriate exploratory data analysis tools for geographical data sets will be considered. The first, and arguably most obvious shortcoming of the basic technique for geographical data is that only one of the two spatial dimensions is used as a conditioning variable. In Figs. 1 and 2 this is not as limiting as in some cases, as we suspect that there is a strong variation in the relationship between height and rainfall in the east-west direction. However, while the plots seem to bear this out, they shed no light on any other geographical patterns. Also, in a more general setting prior information about the direction of pattern variation may not be available. Fortunately, Cleveland shows how coplots can be extended in an obvious way to incorporate two conditioning variables. In Fig. 4 this is shown, with east and north as the conditioning variables.

The main difference between this plot and the earlier ones is that there are now two 'given' panels — the second along the side of the scatterplots. Both east and north are used to divide the observations between the scatterplots. The bottom row corresponds to observations in the lowest north windows and so on up to the top row, and for the columns a similar correspondence occurs, this one depending on the east windows. Using this plot it is possible to investigate the relationship between height

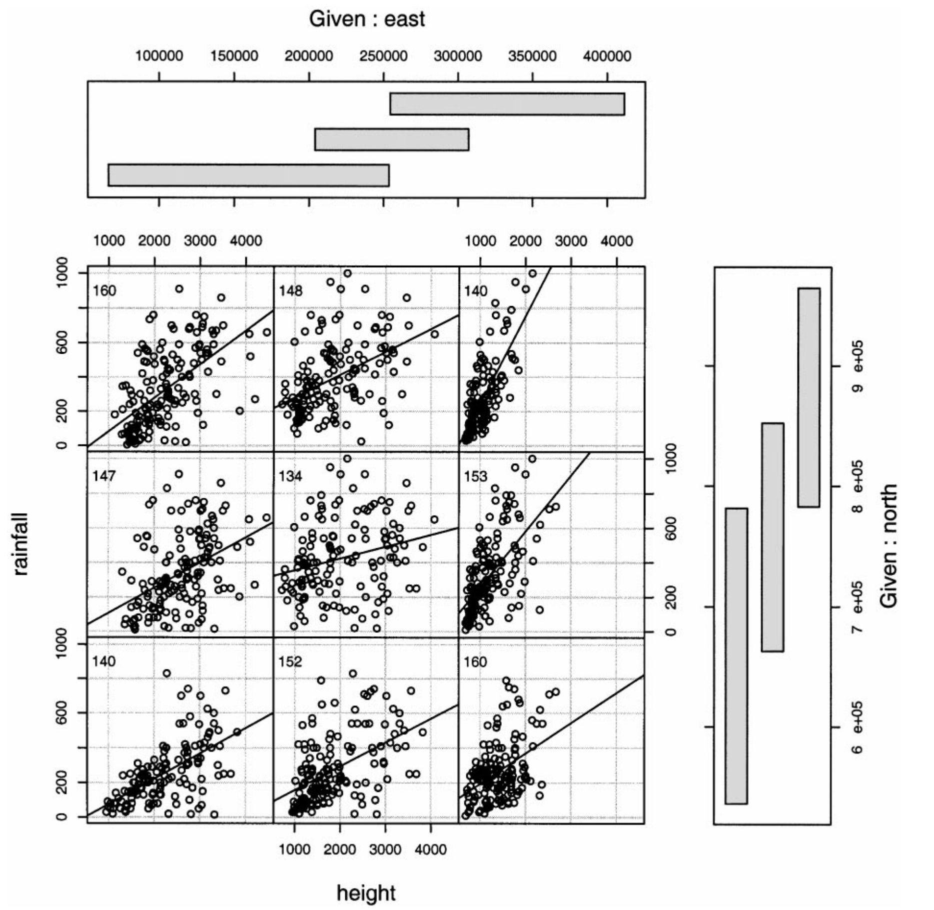

Fig. 4. Coplot using two conditioning variables. Number of observations is shown in top left hand corner of each scatterplot.

and rainfall over two-dimensional geographic space. One feature that now becomes apparent is that the slope of the regression line of rainfall on height seems to get steeper as one heads north, but only on the eastern side of the country.

One problem encountered with the two-conditioning-variable coplot is that it is no longer possible to choose ranges of the given variables that ensure that each scatterplot has exactly the same number of observations. It is possible to solve this problem *marginally*, so that the total number of points in each row of scatterplots is the same, as is the total number of points in each column of scatterplots, but to solve for each individual plot would require different east windows for each column, or different north windows for each row — so that visually scanning a line of plots may not be comparing like with like.

Here, the author proposes a minor addition to the scatterplots. The number of observations in each scatterplot is printed in the top left hand corner. Although this clearly does not remedy the possible problems that differences in numbers of observations create such as those illustrated in Fig. 3, it does provide a quick diagnostic; if one plot is seen to have an unusual pattern it is then possible to check whether this may be due to that plot having an unusually large or small number of observations. It could also be argued that the very presence of these numbers serves as a reminder that observation counts are not guaranteed to be perfectly balanced between the plots. In Fig. 4 it can be seen there although the number of points does vary between scatterplots, this variation is not too great, ranging between 134 and 160 points in a single scatterplot. This is only a variation of under one tenth of the mean number of points in a plot — note that the effect in Fig. 3 was due to the number of points in a pair of plots differing by a factor of five.

Next, note that the conditioning variables are divided into a small number of overlapping classes, whilst the scatterplot variables are depicted in continuous space. This is inevitable since a scatterplot matrix with too many units would be difficult to read. However, for geographers there are situations where we may wish to reverse the roles of the two pairs of variables. In the example here, this would mean that the conditioning variables would be height and rainfall, and the scatterplot variables would be east and north. In essence, the scatterplots are now crude point maps of rain gauge locations. This "rôle reversal" of the variables would show spatial locations in much more detail, at the expense of detailed graphical information for height and rainfall. This may be useful in identifying small-scale local patterns. Obviously, a more direct correspondence to geographical patterns is now established. In Fig. 4, not only are eastings and northings reduced to three broad (albeit overlapping) groupings, but also the geography is distorted — Scotland is deformed into a square!

This type of plot is illustrated in Fig. 5. As well as the exchanging of variable pairs, a number of other modifications have been made. Outlines of the Scottish coast have been added to the scatterplot panels, and the gridlines have been dropped. The gridlines in the original coplot provided graphical reference features when comparing plots. For instance, to see whether the intercept of a regression line is changing between plots, one could compare the point where the line meets the y-axis against the horizontal grid-lines. However, when the points in the scatterplots were geographical locations, it was felt that *geographical* reference features were more useful — hence the inclusion of the coastline. Note that other such features could also be added. In this example relating to orographic enhancement of rainfall, some generalised details of features of mountains could perhaps be added. However, one should perhaps exercise restraint in adding such features. A small number of these may aid interpretation, but too many may serve to clutter the diagram, and obscure patterns in the locations of interest. Tufte's (1990) 'small multiples' work best when the graphical panels are simple. From Fig. 5 it is notable that in the highest altitude category, the greatest degree of rainfall occurs on the western coast, and that in the lowest altitude category there is a shift away from the eastern coast as rainfall levels increase. This is again consistent with orographic enhancement — the gauges in the lowest altitude category recording the most rainfall are closer to those in the higher altitude category recording high rainfall. The advantage of visualising the data in

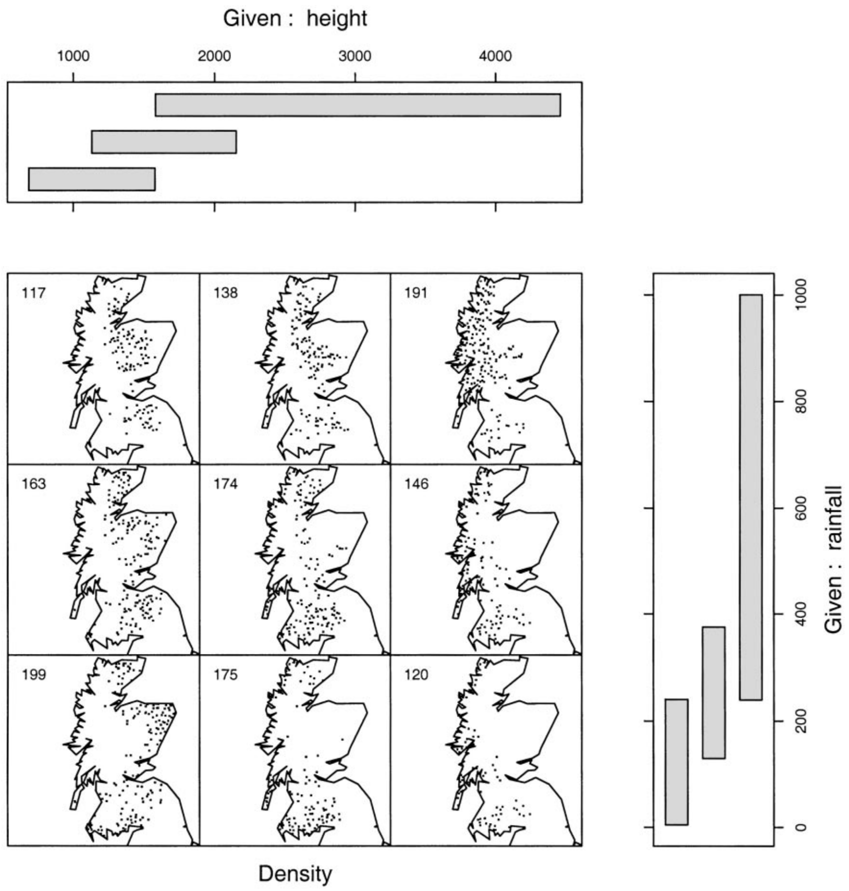

Fig. 5. Comap using two conditioning variables. Each point on a map panel represents an individual rain gauge.

this way, rather than as in Fig. 4 is that a more precise picture of the geographical pattern is communicated.

This line of argument is not intended to be dismissive of the original coplot approach used to generate Fig. 4 — clearly the technique provides very strong graphical evidence that the relationship between height and rainfall is not constant over geographical space. However, it may be helpful to visualise a more precise geographical structure of the variations in the relationship. For this reason, it is recommended that data of the kind being considered here should be visualised in both ways. Using a greater variety of graphical views of a multivariate data collection should increase the likelihood of features in the data being identified.

Another modification in five is that all scatterplots now have an aspect ratio of unity — the scale in the x-direction is the same as the scale in the y-direction. This is a departure from a standard coplot — (Cleveland, 1993) suggests rules for choosing aspect ratios for scatterplots which may not suggest unity even when the x and y

axes relate to quantities having the same physical dimension. There are good reasons for adopting these rules in many situations, but the aim here is to represent geographical patterns faithfully, which in this case neccessitates over-riding the original rules. As observed earlier, once the above alterations have been made to the standard coplot, the scatterplot panels have essentially become maps. Bearing this in mind, the term *comap* is defined to mean a coplot in which the scatterplot panels have been replaced by maps of some sort.

4. Beyond the basic comap: an example using density estimation

A comap need not use point-based representations of spatial data. Like coplots, the basic idea is adaptable — indeed the map panels may be any form of map that could represent the geographical information. As an example, instead of simply plotting the point locations for east and north, a smoothed map of point intensity could be used. The purpose of this would be to highlight larger scale trends in geographical pattern. This could be achieved using *kernel density estimation* techniques (Brunsdon, 1995; Silverman, 1986). In its two dimensional form this method may be used to provide an estimate of a probability density function $f(Y_1, Y_2)$ given a set of n observed values of Y_1 and Y_2, say (y_{1_i}, y_{2_i}) for $i = \{1 \ldots n\}$. The formula used for the estimation is

$$\hat{f}(Y_1; Y_2) = \frac{1}{nh^2} \sum_{i=1}^{n} K\left(\frac{Y_1 - y_{1_i}}{h}\right) K\left(\frac{Y_2 - y_{2_i}}{h}\right) \quad (1)$$

where the hat over f denotes that it is estimated rather than a known true value, K is a univariate probability density function, and h is a constant (whose dimension is length in geographical applications) referred to as the *bandwidth*. In the geographical case an isotropic model is used — it is possible, for instance, to have different bandwidths in the Y_1 and Y_2 directions, but in practice this has not proved worthwhile. $\hat{f}(Y_1, Y_2)$ may be thought of as a surface or set of contours showing where the greatest concentration of points occur. Contour maps of this kind will be used as the map panels in the comaps proposed here. Inspecting Eq. (1) suggests that the choice of K and h will determine the appearance of $\hat{f}$. In fact, the outcome is fairly robust to the choice of K — provided it is smooth, unimodal and symmetrical about the mode — see Bowman and Azzalini (1997), Wand and Jones (1995) or Silverman (1986) for further discussion. Here, K will simply be the normal density function. However, $\hat{f}$ does depend fairly heavily on the bandwidth. Very large values of h tend to smooth out features, whereas very small values tend to give ‘spikey’ surfaces. Neither of these outcomes are desirable, and so some way of choosing h from the data is needed.

There are many such techniques in the statistical literature, but it is argued here that an appropriate method for this application should be based on a conservative choice of h. ‘Conservative’, in this context, means that it is likely that some degree of oversmoothing occurs. As stated above, it is not desirable to massively oversmooth the data, but since the aim here is to detect geographical trends at a larger scale it is helpful to be reasonably confident that any features detected do genuininely reflect some underlying process, and are not a one-off artifact of the particular data sample being visualised. A number of conservative smoothing approaches have been suggested. Bowman and Azzalini (1997) suggest choosing the optimal h under the assumption that (Y_1, Y_2) are normally distributed, optimal being defined in terms of minimising the expected value of the integrated mean square error between f and $\hat{f}$. For the density estimate specified in Eq. (1) with K a normal density function this gives

$$h_{opt} = \sigma n^{-\frac{1}{6}} \quad (2)$$

where σ is the pooled estimate of the standard deviation of Y_1 and Y_2. In the comapping context, σ and n would be supplied on a panel-by-panel basis, so that different bandwidths would be used in each map panel. To some extent, this overcomes the problem of differing numbers of points in each panel — when there are a smaller number of points in a map panel h tends to be larger, and the density surface is spread out more, compensating for the effect demonstrated in Fig. 3.

The standard deviations of Y_1 and Y_2 are assumed to be the same in Eq. (2). Note that in reality, neither this assumption or the assumption of normality are expected to apply. The reason that they are used for conservative smoothing is that they are likely to provide an over-estimate of h. If the data really were bivariate normal, then a relatively large degree of smoothing would be an appropriate choice, as it would tend to yield a univariate $\hat{f}$. However, the requirement for an *optimal* choice of h tends to stop it becoming *too* large — very large values would certainly produce normal-shaped $\hat{f}$'s, but with too much variance. Typically, the outcome of choosing a bandwidth in this way is that one is unlikely to obtain a multimodal (or multi-featured) density estimate if the true distribution is unimodal, but there is some chance that a multimodal true distribution may be oversmoothed giving a unimodal density estimate. This essentially encapsulates the conservative nature of the technique — if a multimodal density estimate is obtained it is unlikely to have arisen spuriously. Terrell (1990) goes on to consider and extend the idea of conservative smoothing in more detail.

This may seem something of a mathematical diversion from what is essentially an exploratory technique, but it may be viewed as some 'under the bonnet' work to ensure that the density visualisation proposed here is not misleading. For such a technique to be used easily, some form of automatic choice of h is needed. Often in density estimation, 'automatic' bandwidth choice does not mean optimal, and so rather than aim for an optimal h, a 'safe' choice is made, in the sense of one that is unlikely to highlight spurious features. The slight oversmoothing can also be justified from a visualisation viewpoint — the idea of the smoothed map is to compliment Fig. 5, and so concentrate on large-scale trends. If a smooth map is seen as being 'risk-averse' in terms of showing spurious features, a point map is 'risk-prone', since it shows every single data point. Perhaps a more wholistic view of the geographical structure in the data may be obtained by viewing both of these maps.

An example of a density-based comap is given in Fig. 6. Here, each map panel is based on kernel density estimates as described above. The maps themselves use five grey-shaded contour bands, clipped around the coast of Scotland. It was decided to omit keys on these maps, since they tended to overload the diagram. Also, the main goal here was to communicate the changing shape of the surfaces between map panels, rather than the exact values of the densities. The comap thus produced tells essentially the same story as Fig. 5, although perhaps the pattern has been made more striking. For the lowest height band, there is a distinct eastward shift in the densities as one progresses up the rainfall bands. Also, there appears to be a corresponding northward shift when one considers the medium height category maps.

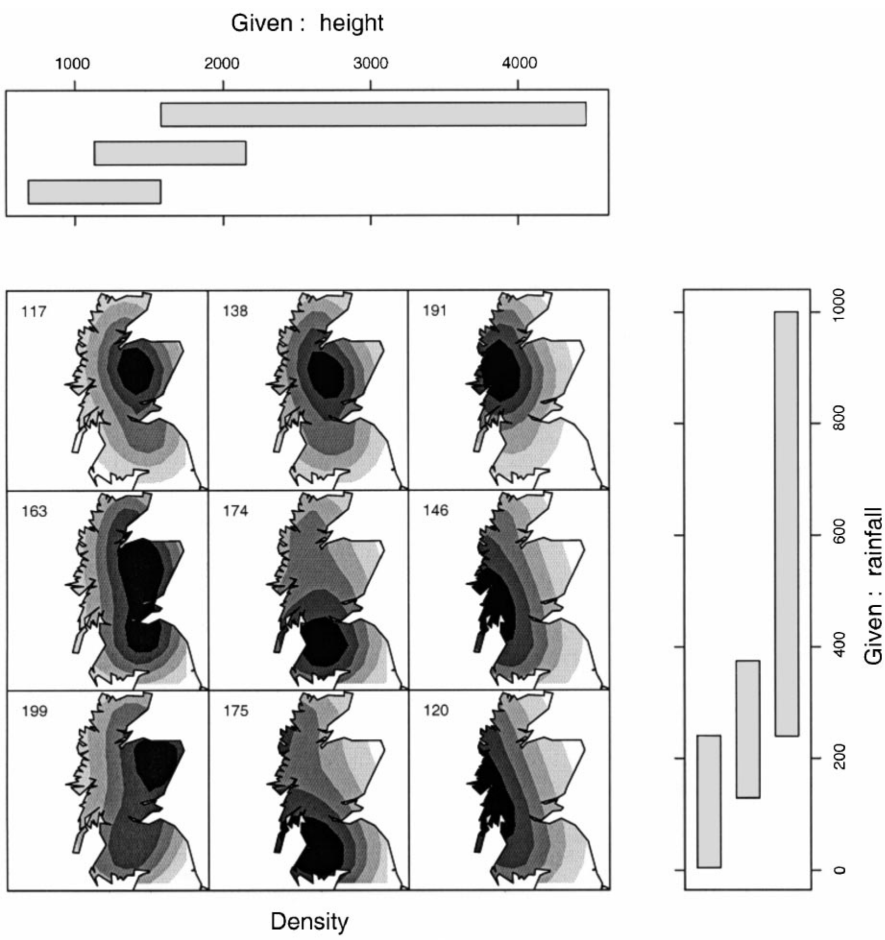

Fig. 6. Density estimation comap.

5. Software issues

All of the coplots and comaps in this paper were produced using the R software package. R consists of a programming language together with a runtime environment with graphical capabilities (Ihaka & Gentleman, 1996), and it is available as Free Software under the terms of the Free Software Foundation's GNU General Public License (Free Software Foundation, 1991). It is similar in functionality to the commercial software package S (Becker, Chambers & Wilks, 1988; Chambers & Hastie, 1992), although its implementation is somewhat different — for discussion of this see Hornik (1999).

R is used here for a number of reasons. Firstly, a coplot function is already provided in the standard distribution. This function allows the basic coplots (with one or two conditioning variables) to be drawn. This function is designed to be flexible — one very useful feature of R (and also of S) is the ability to treat functions as though they were variables, so that R code can be passed as parameters to other R functions. This is the key to adapting the basic coplot. One of the arguments to the coplot function (named panel) provides code which determines how the scatterplot panels are to be drawn. If no value is supplied, it defaults to a simple point-based plot. However, defining your own functions and passing these as the panel argument allows different kinds of plots to be drawn. For instance, Fig. 1 was created by defining a function which plotted a set of (Y_1, Y_2) pairs, computed the regression line for this data set, and then added this to the plot. A similar approach was used to add the counts of numbers of points in the top left hand corners of panels for subsequent coplots and comaps.

A second reason for using R is that the source code is freely accessible and modifiable — this is part of the terms of the GNU Public Licence. This is particularly helpful when adapting coplots into comaps in R. A number of features of comaps could be added using the method described above — but some features of the coplot function could not be adjusted in this way. In particular, setting the aspect ratio of the scatterplot panels to unity and switching off the grid could not be achieved via the panel argument, or indeed any other arguments to coplot. However, since the code for the coplot function is provided (in fact, it is written in the R language) it is possible to take a copy of these, and then define a modified function (here called adjusted.coplot) in which these adjustments (and a few others) have been made.

Once adjusted.coplot is defined, it is a relatively simple task to define functions to produce the comaps shown in this paper. R lends itself to data analysis and statistical data processing, and a two-dimensional kernel density estimation function may be defined in a small number of lines of code. Creating the maps themselves is relatively simple, as R provides primitives to add points, lines, polygons and text to a graphics window. Perhaps the most sophisticated mapping task is to draw filled contour maps of the estimated density functions which have been clipped by the UK boundary — although even this is relatively simple. The clipping effect is actually produced by drawing a set of filled contours over a rectangular area (a function already provided by R), and then masking out the parts of the area that are not on the Scottish mainland. For the whole of the UK, two such 'masking' polygons are shown in Fig. 7. If it is assumed that the background of the maps are white, then drawing these two polygons with solid white shading and solid white borders has the desired clipping effect. The border could then be emphasided by adding the UK border as an unfilled polygon with black lines.

In summary it is relatively easy to produce comaps in R for two reasons — part of the problem has already been solved by the coplot function, and the openness and flexibility of the system allow one to extend and modify this partial solution and go on to finish the job. The R code used here is available on the world-wide web (Brunsdon, 2000).

Fig. 7. Masking polygons to clip maps around the UK coast. The two shaded polygons cover any area in the surrounding rectangle that is not on the UK mainland.

6. Conclusion

In this paper Cleveland's concept of a coplot has been discussed, with particular attention to geographical applications. Following on from this, a new but related concept, the comap, has been defined. Both coplots and comaps offer a relatively simply understood means of visualising and exploring multivariate spatial data, based on Tufte's notion of 'small multiples' of diagrams. In particular, comaps allow the investigation of spatial patterns given the variation of one or two atribute variables for geographical objects. As with many 'small multiple' applications of visualisation, it is best if each of the map panels in a comap are relatively simple, and that they are identical in design — the only characteristics that alter between panels should be those that depend on the individual data samples.

Although the example given in this paper related to rainfall data, the range of potential applications for comaps and coplots is very much broader than this. For example, one could consider house price and floor area as two conditioning variables — in many situations the relationship between housing characteristics and price may vary geographically (Brunsdon et al., 1999). Another useful application may be in visualising point patterns in space and time — here time could be used as a conditioning variable. Indeed, in the latter approach one could then add a second conditioning variable — for example if the space-time point patterns were incidence of some disease, then one could consider age of patient at time of diagnosis. This may be useful in identifying changing trends in the characteristics of some illness in both temporal and spatial terms.

The two examples given above give some idea of the variety of applications when considering geographical *point* patterns, but comaps need not be limited to this one kind of geographic object. Since the way maps are drawn in comap panels may be easily altered or extended, there is no reason why zone-based or arc-based objects could not be viewed in this way. Given this, the comap approach could potentially lead to a very broad class of visualisation techniques for investigating bivariate geographical patterns or space-time patterns.

References

Becker, R. A., Chambers, J. M., & Wilks, A. R. (1988). *The new S language*. London: Chapman and Hall.

Bowman, A., & Azzalini, A. (1997). *Applied smoothing techniques for data analysis: the kernel approach with S-plus illustrations*. Oxford: Oxford University Press.

Brunsdon, C. (1995). Estimating probability surfaces for geographical points data: an adaptive kernel algorithm. *Computers and Geosciences, 21*, 877–894.

Brunsdon, C. (2000). R code for comaps. http://www.staff.ncl.ac.uk/chris.brunsdon/code/code.html.

Brunsdon, C., Fotheringham, A. S., & Charlton, M. (1999). Some notes on parametric significance tests for geographically weighted regression. *Journal of Regional Science, 39*, 497–524.

Chambers, J. M., & Hastie, T. J. (1992). *Statistical models in S*. London: Chapman and Hall.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association, 74*, 829–836.

Cleveland, W. S. (1993). *Visualizing data*. New Jersey: Hobart Press.

Free Software Foundation (1991). GNU General Public Licence. http://www.gnu.org/copyleft/gpl.html.

Hornik, K. (1999). The R FAQ. http://www.ci.tuwien.ac.at/hornik/R/.

Ihaka, R., & Gentleman, R. (1996). R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics, 5*(3), 299–314.

Pedgley, D. (1970). Heavy rainfalls over Snowdonia. *Weather, 25*, 340–350.

Sawyer, J. (1956). The physical and dynamic problems of orographic rain. *Weather, 11*, 375–381.

Silverman, B. W. (1986). *Density estimation for statistics and data analysis*. London: Chapman and Hall.

Terrell, G. (1990). The maximal smoothing principal in density estimation. *Journal of the American Statistical Association, 85*, 470–477.

Tufte, E. R. (1990). *Envisioning information*. Cheshire, Connecticut: Graphics Press.

Tukey, J. (1977). *Exploratory data analysis*. New York: Addison-Wesley.

Wand, M., & Jones, C. (1995). *Kernel smoothing*. London: Chapman and Hall.