

SPATIAL VARIATIONS IN THE AVERAGE RAINFALL–ALTITUDE RELATIONSHIP IN GREAT BRITAIN: AN APPROACH USING GEOGRAPHICALLY WEIGHTED REGRESSION

C. BRUNSDON^a, J. McCLATCHEY^b and D.J. UNWIN^{c,*}

^a *Department of Geography, University of Newcastle, Newcastle upon Tyne, UK*

^b *School of Environmental Science, Nene College, Northampton, UK*

^c *School of Geography, Birkbeck College, University of London, London, UK*

Received 7 October 1999

Revised 26 September 2000

Accepted 28 September 2000

ABSTRACT

The relationship between annual rainfall totals and gauge elevation over Great Britain is re-examined using the recently developed technique of geographically weighted regression (GWR). This enables the spatial drift of regression parameters to be identified, estimated and mapped. It is shown that the rate of increase of precipitation with height, or height coefficient, varies from around 4.5 mm/m in the northwest to almost zero in the southeast. There is a particularly rapid change in this value across the English Midlands. The predicted sea level precipitation varies from 1250 mm to less than 600 mm in much the same way. Copyright © 2001 Royal Meteorological Society.

KEY WORDS: altitude; precipitation; regression modelling; Great Britain

1. INTRODUCTION: THE PROBLEM

Although rainfall (precipitation) is gauged at points over typically sparse networks, for many environmental studies it is useful to interpolate these records as a continuous surface, represented by isohyets. Such estimates are nowadays often used in geographical information systems relating them to other physical or human factors. Often, they are integrated over space and time to give river catchment totals for use in hydrology and water resource planning and they are implicit in most maps of rainfall totals.

It has long been known that a statistical relationship exists between rainfall totals integrated or averaged over periods of time from the month upwards and the station elevation (see Salter, 1918; for an extended review, see Barry, 1992). In this paper, we study the problem of rainfall interpolation making use of a new method, known as geographically weighted regression (GWR, Brunsdon *et al.*, 1996), that enables spatial and temporal non-stationarity in the statistical relationship between rainfall totals and altitude to be addressed directly. Although computationally intensive, the method is conceptually simple, and undoubtedly has further utility in climatology. In the remainder of this paper, Section 2 outlines previous work in the field. Section 3 shows some results using scientific visualization in a highly interactive computing environment. Section 4 outlines the method of GWR, with results for annual totals over Great Britain given in Sections 5 and 6. Section 7 provides a conclusion.

* Correspondence to: School of Geography, Birkbeck College, University of London, London, UK; e-mail: d.unwin@bbk.ac.uk

2. APPROACHES TO STUDYING RAINFALL DISTRIBUTION

Orographic enhancement of rainfall arises from a number of mechanisms which serve to increase both the chances of rainfall over hills and the rate when it is occurring (Sawyer, 1956; Pedgley, 1970). These mechanisms are well-known, and it is not the purpose of this paper to elaborate on them. In the mid-latitudes, most studies have found that a simple linear model seems to fit the observed data well, at least for extended time periods and over relatively small areas. Investigations of this orographic enhancement for individual storms or by modelling have proved to be difficult (Dore *et al.*, 1982).

The simple linear model, usually fitted by ordinary least squares (OLS) methods, is

$$P = c_0 + c_1(H) + \varepsilon(\text{mm}) \quad (1)$$

P = rainfall (mm)

c_0 = rainfall at sea level (mm)

c_1 = dimensionless rate of increase in rainfall with altitude, or height coefficient (mm/m)

H = station altitude (m)

ε = an error term

Unusually, for a climatological regression, this is dimensionally balanced and the estimated parameters have direct physical interpretations as the predicted rainfall at sea level (c_0) and the rate of increase in rainfall with altitude or height coefficient (c_1). For some 6500 stations over Great Britain, Bleasdale and Chan (1972) found an overall relation:

$$P_{\text{est}} = 714 + 2.42(H) \dots (\text{mm})$$

It should be noted that lack of data from stations at altitude makes both the linear specification and the parameter estimation uncertain as one moves to high values of (H). There are very few rain gauges in Great Britain at altitudes above 250 m (Taylor, 1976).

There is evidence in the Bleasdale and Chan study to suggest that the model used was incorrectly specified. Evidence for this is provided by the clear pattern of spatial auto-correlation in the regression residuals. Values > 600 mm were found in a large block over the west of mainland Scotland with values < 600 mm concentrated over Cairngorm and the east. The line of exact fit, with residuals close to zero, closely followed the main watershed through Scotland, northern England and Wales. It is of interest to note that in his original paper of 1918, Salter gave estimates of the rate of increase with height (c_1) equivalent to 1.5–2.0 mm/m for southeast England and 2.0–3.0 mm/m for west Wales.

A number of authors have attempted to improve the fit of this basic model by resort to a variety of methods. One approach is to fit other functions, as, for example, in Ballentyne's use of a curvilinear function to model the relationship on An Teallach, a mountain in the northwest Highlands of Scotland (Ballentyne, 1983). Alternatively, several authors have found that taking a local average height around each gauge site can improve the fit of the relation. Examples include Chaun and Lockwood's (1974) use of the mean height over 8 km around each station and, in more extreme form, Hill *et al.*'s (1981) use of the average height over 4 km grid squares centred 1.5 km to the southwest of the gauge site. Recently, using data from a 17 gauge network ranging in altitude from 8 to 385 m in Lochaber in the western Highlands of Scotland, Harrison and Clark (1998) have shown that even in an area where altitude effects are usually held to be pronounced, it is possible to sample spatially and temporally such that the rainfall–altitude relationship disappears or can even be negative. Directional effects and gauge exposure in individual rainfall events have been modelled in a number of ways (Weston and Roy, 1994). Many studies have used multiple linear regression to incorporate several effects in the same model (see Basist *et al.*, 1994). Possibly the most careful analysis using a combination of methods is that of Spackman (1993), who used monthly totals to provide a series of rainfall maps. In his study, use was made of some 32 explanatory variables, including measures of the geographic position, local topography, exposure and land use around the gauge site and model calibration was by iterative least squares. In all these studies, it is debatable whether the additional effort in data collection and model complexity is really justified by the increase in model fit to the observed data.

There are good climatological reasons for expecting the rainfall–altitude relationship to vary spatially. It is clear that over Great Britain the orographic effect is most pronounced in warm fronts and the warm sector of depressions but is not important in cold frontal rain events. It follows that spatial variation in the synoptic mix of rain-producing events will give variations in the relationship, and such variations have been well-documented (Shaw, 1962; Smithson, 1969a,b; Atkinson and Smithson, 1976). We can expect spatial variations in both the intercept term, c_0 , and the height coefficient, c_1 . Spatial variation in c_0 relates to what is often called the ‘rain shadow’ effect. Similar variation in c_1 is evidence that the rate of increase with altitude cannot be modelled by a single relation valid over a wide area.

Most of the studies cited above have recognized this spatial non-stationarity and have attempted to incorporate it into their analysis, often by use of some form of spatial trend surface (Unwin, 1969). Alternatively, individual regressions have been calculated for arbitrary (but sensible) regions, as, for example, in McClatchey’s (1996) study of the rainfall/height relation across Scotland, in which he showed that the rainfall gradients are steeper in the northwest and west and lowest in the central and southeast Highlands. Much of the most recent work has tended to attack this problem using methods from geostatistics using Matheron’s theory of regionalized variables to obtain statistically optimum interpolations (Bastin *et al.*, 1984; Bigg, 1991; Goovaerts, 1999).

3. VISUALIZING THE RAINFALL-ALTITUDE RELATION

Figure 1 presents a plot of the rainfall/height relation over Great Britain created using a system realized in the X-LISPstat environment (Brunsdon, 1998). The data plotted are for the mean annual rainfall at some 10925 gauges over Great Britain. These are all the gauges for which the UK Meteorological Office has published an annual average total for the period 1961–1990. Over 70% of the gauges had records of 10 years or longer, and over 90% had records of 5 years or more. These records were checked to determine whether there was any difference in the distributions of gauges with altitude between the stations with 10 or more years of observations, and those with fewer than 10 years, but there was no significant difference (*t*-test) between the two sets. Although these data originated in the UK Meteorological Office, they were kindly supplied by the Climate Research Unit, University of East




Figure 1. Plot of the rainfall/height relationship for all stations across Great Britain

Anglia, where they have been assembled and made available for academic research. *The authors would be the first to point out that generating such a comprehensive and carefully quality-controlled data set as this forms a major input to the research results we report here, and we are happy to acknowledge our gratitude for this work in rather more than the usual footnote.*

Overall, there is a weak linear relationship that hardly differs from that originally found by Bleasdale and Chan (1972):

$$P_{\text{est}} = 730.83 + 2.28(H) \dots (\text{mm})$$

With a correlation coefficient of just 0.582, it is clear that much of the variance remains unexplained. This might be a result of true random variation, of the assumption that the parameters are spatially stationary over the country, or of a failure to introduce other variables into the model specification.

To examine spatial variation in the parameters, an extension of the method of using pre-defined blocks of data is to use what is known as object hot linking in a highly interactive computing environment. This enables a scatter plot of all the observations to be linked to a map, such that any subset of data highlighted on the map is also highlighted on the scatter plot and vice versa.

Figure 2 illustrates how we can use this method to explore the relationship a little further. On it, we have used the mouse to 'lasso' a block of gauges in the west of the area. These are shown in the darker symbolism, and instantaneously highlighted in the same way on the scatter plot. Any rainfall/height regression for these highlighted stations would show very different estimates of the coefficients c_0 and c_1 , and the correlation than those when all the gauges are considered. Specifically, we would expect the regression for this subset to have higher values for both the intercept, c_0 , and the height coefficient, c_1 , than is the case overall. For this western subset, there is also considerable scatter around the regression, indicative of a relatively poor fit that presumably results from a greater variety of local effects on the rainfall totals than is found in the east.

Figure 3 shows an equivalent pair of displays on which gauges in the east have been highlighted. The plot shows that the same model form must have a lower intercept and rate constant, but will have a much better fit. We conclude that the regression (Equation (1)) cannot possibly be considered stationary. Continued exploration by this simple, but highly effective, visualization technique leads to a number of simple hypotheses. There is spatial variation in c_0 and c_1 and most probably also in the spread (fit, R^2). As a result of these visualizations, we expect:

- (a) That the estimated c_0 will decrease systematically from north and west to east and south across England.
- (b) That the estimated c_1 will be highest in the north and west and decline in the same direction.

4. GEOGRAPHICALLY WEIGHTED REGRESSION

In this section, we suggest and exemplify a direct approach to handling this problem of spatial non-stationarity based around the idea of a local statistic (Unwin, 1996). As outlined in Section 2, Equation (1) is the simplest form that we can reasonably suggest for the rainfall/height relationship. In geographically weighted regression (GWR, Brunsdon *et al.*, 1996), we consider the stability of the coefficients c_0 and c_1 over space. As our review of the literature and the visualizations presented above show, in the case of orographic enhancement it is incorrect to hold that the same linear relationship is appropriate in all places. Retaining the same linear model, we can allow its two parameters, the intercept constant and the height coefficient to change, or 'drift', over space. That is, if (x, y) is a coordinate pair on a cartesian projection such as the UK National Grid, the simple linear model of Equation (1) can be expanded to

$$P = c_0(x, y) + c_1(x, y) H + \varepsilon \tag{2}$$




Figure 2. Object linking a subset of stations in the west of Great Britain to the overall plot

This revised model allows the coefficients to vary as continuous functions over space, so that each may be thought of as a three-dimensional surface over the geographical study area rather than as a single, fixed, real number. Typically, one wishes to estimate the values of $c_0(x, y)$ and $c_1(x, y)$ for points on a regular grid. This is achieved by considering data for places in the vicinity of a given location (x, y) . For example, if one drew a circle of some radius, r , around a point (x, y) , and calibrated an ordinary least squares regression model on the basis of observations that fell within this circle, then the c_0 and c_1 obtained could be thought of as an estimate of the local coefficients at the point (x, y) . By evaluating the coefficients for each (x, y) over a regular grid, it is possible to obtain a set of estimates of spatially varying parameters without specifying a functional form for that spatial variation. In this sense, this technique 'lets the data speak for itself' when providing estimates for each $c_0(x, y)$ and $c_1(x, y)$. This is the underlying concept of GWR. The approach is similar to the drift analysis of regression parameters developed by econometricians (Casetti and Can, 1999), as well as the ideas of kernel and nearest neighbour regression developed by statisticians (Cleveland, 1979; Cleveland and Devlin, 1988). In all these approaches, the philosophy is not to seek and calibrate some form of invariant 'law', but instead, it holds

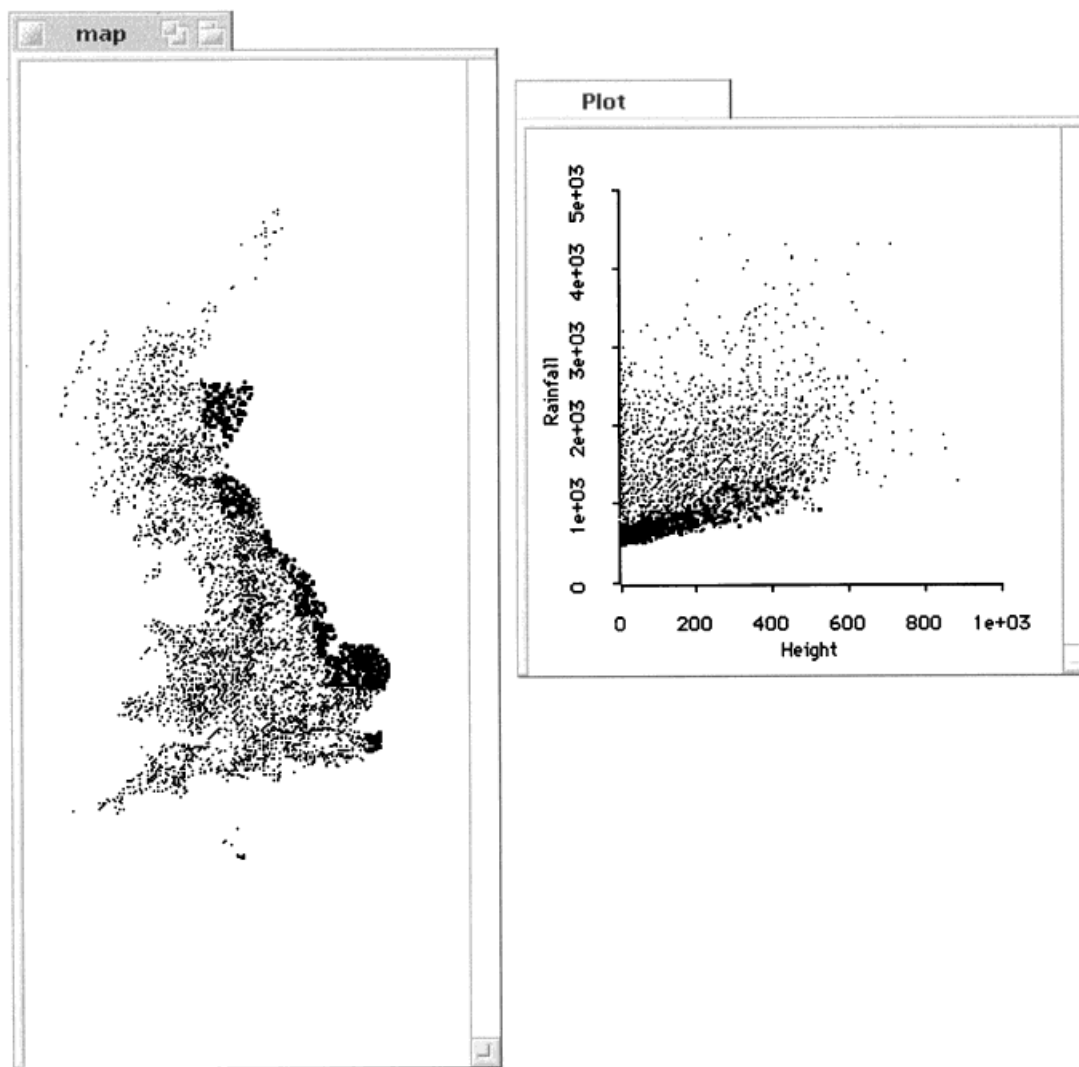


Figure 3. Object linking a subset of stations in the east of Great Britain to the overall plot

that looking for interesting relationships across some space or other is often a more scientifically fruitful line of enquiry.

An initial consideration of the technique outlined above raises issues to do with the notion of the 'circle of inclusion' of observations around each point. This has been specified to have radius, r , but what value should r take? If r is very large, the data included in each and every point-based estimate will include most of the study area, so the estimated coefficients will all be very close to those obtained by applying a global model. If it is very small, few observations will be included in the calibration, which will lead to estimates with very large standard errors, giving high variability in the estimates. Another issue that may be raised is the binary nature of the inclusion of observations in the regression model calibration. Any observation whose distance from (x, y) falls just below r will be included in the model, whereas one whose distance just exceeds this will be excluded. It does not seem natural that the spatial association between the variables should end so abruptly.

In the next few paragraphs, we address the second of these issues. To overcome the binary inclusion/exclusion problem, an approach using weighting is taken. The regression model in Equation (2) centred around each (x, y) is thought of as a weighted OLS regression on the whole data set, with

observations in the circle of inclusion weighted as unity and other observations as zero. This, for any given (x, y) the weight given to the k -th observation is:

$$\alpha_k = 1 \quad \text{if } d_k(x, y) < r, \quad 0 \text{ otherwise} \tag{3}$$

in which $d_k(x, y)$ is the straight-line distance between the location of observation k and the point (x, y) . But there is no reason to restrict the weighting to a step function. It is also possible to relate $d_k(x, y)$ to α_k with a continuous function. For example, a Gaussian distance decay weighting is achieved by

$$\alpha_k = \exp(-d_k(x, y)^2/2r^2) \tag{4}$$

Here, the value of the weight decays gradually with distance to the extent that, at a distance three times the bandwidth ($d_k(x, y) = 3r$), the weighting is virtually zero. A variety of other functions might be used, and in the literature these are referred to as *kernel functions*, or *kernels*, and are denoted by the symbol K , as in ($k = K(d_k(x, y))$). In each case, the constant bandwidth, r , provides control over the range of the ‘circle of influence’ of the geographical data, as it does in the basic step-function example, but, as we have just seen, the weighting decays with distance, rather than dropping suddenly to zero when a certain distance is reached. Increasing r increases the smoothness of the estimates of the functions $c_0(x, y)$ and $c_1(x, y)$. The choice of K controls the shape with which the weighting decays. Although experimentally the choice of K appears less important than the choice of r , desirable features of K are:

- $K(0) = 1$
- $\lim_{d \rightarrow \infty} K(d) = 0$
- K is a positive monotone decreasing function for positive arguments.

Note that if $K(\cdot)$ is a continuous function, then the kernel estimates for $c_0(x, y)$ and $c_1(x, y)$ will themselves be spatially continuous.

Whilst the above approach overcomes the binary inclusion/exclusion problem, it does not address the ‘circle radius’ problem, which is equivalent to choosing r in, for example, Equation (4). In the standard text on this type of work, Silverman (1986) suggests that there are occasions when subjective choice lends itself well to the problem in hand, but there are many situations in which the prior knowledge needed to make an informed choice does not exist. In such situations, some form of automatic, data-driven choice of r is more appropriate. One method suggested here is that of *least squares cross validation*. By choosing (x_k, y_k) as the location of observation k , one can obtain estimates for $c_0(x, y)$ and $c_1(x, y)$. Using the observed height at point k , H_k , we obtain a predicted value for the rainfall as

$$P_{k(\text{est})} = c_0(x_k, y_k) + c_1(x_k, y_k)H_k \tag{5}$$

The sum of squared errors between predicted rainfall totals and those observed may then be written as

$$SS(h) = \sum_k (P_k - P_{k(\text{est})}(r))^2 \tag{6}$$

The estimated value of P depends on the choice of r , and it is sensible to choose r , such that Equation (6) is minimized. However, attempting to implement this techniques immediately throws up a difficulty since as $r \rightarrow 0$, the ‘circle of influence’ around the k th point shrinks and so downweights the influence of all observations except the k th. As this happens, $P_{k(\text{est})}(r) \rightarrow P_k$. This implies that an unmodified least-squares automatic choice of r will always suggest a zero value, or possibly give a computational error. This problem can be avoided if, for each k , a GWR estimate of P_k is obtained omitting the k th observation from the model. This is equivalent to replacing the kernel function K , by a modified function K^* , such that

$$K^*(0) = 0$$

$$K^*(d) = K(d), \quad \text{if } d \neq 0$$

If the modified GWR estimate of P_k is denoted $P_k^*(h)$, then the cross-validated sum of squared errors is given by

$$\text{CVSS}(h) = \sum_k (P_k - P_k^*(r))^2 \quad (7)$$

Finding an r to minimize Equation (7) provides a method for choosing r automatically that does not suffer from the problems encountered when working with Equation (6). However, it should be noted that although this method is 'automatic', it is not necessarily optimal. In the related area of generalized additive models, Hastie and Tibshirani (1990) find that, in some cases, cross-validation methods tend to undersmooth. When this occurs, the plot of r against $\text{CVSS}(r)$ tends to show multiple minima. As a working rule, it is suggested that r should be plotted against $\text{CVSS}(r)$. If only a single minima is observed, then the corresponding value of r should be chosen. If tell-tale multiple minima occur, then the globally minimizing value of r should be inflated by a small safety factor to compensate for any possible undersmoothing.

There is a brief primer on the GWR technique at the website <http://www.ncl.ac.uk/~ngeog/GWR/> to which interested readers are referred for additional information and appropriate FORTRAN code.

5. SPATIAL VARIATIONS IN THE RAINFALL-ALTITUDE RELATIONSHIP

A GWR analysis was carried out on the rainfall data according to Equation (2) using the cross-validation approach outlined above. The software used was written by the first author of this paper and is not the same as that at the website mentioned above. A plot of r against $\text{CVSS}(r)$ (not shown) suggests an optimal value for r of about a kilometre. However, as the plot shows some evidence of multiple minima, it is probable that this value is too low and, in response, a value 2 km was used to derive estimates for $c_0(x, y)$ and $c_1(x, y)$. This implies that, for each estimation, we are considering gauges in a constant area of around 113 km², with radius 6 km around each evaluation point. Although this choice of r is to an extent arbitrary, it seems a reasonable scale of analysis given the overall number and areal density of available rain gauge sites. The real test of its suitability lies in the results obtained. It will be seen that, although the choice gives what we judge to be 'good' results over most of England and Wales, it may be that it is not optimal for much of the Highlands of Scotland. In these areas, the number of gauges used in each determination can be very low indeed, such that, in conventional statistical terms, the standard errors of the estimates are high. In fact, an examination of the inter-gauge spacing for gauges in Scotland shows that typically there are six to eight gauges within a 6 km radius of each other, but there are areas where the density is even lower. Clearly, we should not have much confidence in regressions based on three or four points and for this application it is possible that less equivocal results would have been obtained using a bandwidth that was allowed to vary spatially as a function of the gauge density in the manner developed for kernel density estimation (Brunsdon, 1995).

Despite this problem, the results, where the regressions are evaluated over a 1 × 1 km grid, show some clear patterns that cannot be artefacts of the sampling. The results for the constant bandwidth $r = 2$ km are presented as Figures 4 and 5. Figure 4 maps the results of a GWR for the height coefficient, $c_1(x, y)$ in mm/m contoured at intervals of 0.5 mm/m. The map not only confirms our first hypothesis of spatial non-stationarity in this parameter, it provides some numerical calibration for it and displays a number of interesting features. There appears to be a steep increase in the height coefficient moving from the southeast of England through the Midlands to the north and northwest, with the same feature apparent moving from east to west across southern England. Interestingly, simple OLS regression models on subsets of the data for the southeast and south Midlands, the Midlands and north Midlands give height coefficients of 0.95, 1.9 and 2.1, respectively (the corresponding R^2 values being 0.48, 0.81 and 0.85). The GWR model appears to have correctly identified this change which would not necessarily be apparent using sub-sets of the data. A similar increase in the height coefficient is also found using OLS regression on subsets of the data for southern England working westwards across the area of rapid increase shown in Figure 4. This feature of the precipitation climatology of England has not as far as we are aware been described before and presumably reflects the greater proportion of rainfall from convective events to the south and east of this zone.




Figure 4. Contour map for the height coefficient

The height coefficient of around 3.0 mm/m shown for Cheshire and Lancashire in northwest England is unexpected. A simple OLS regression model using gauge totals in that area gives a value of 2.1 ($R^2=0.62$), but if the subset area is expanded westwards to include Snowdonia and the Lake District, it increases to 3.1 ($R^2=0.45$). The choice of subset area is, therefore, crucial, and this emphasizes that the GWR is identifying areas of interest that may not be apparent using subsets of the data. Maximum values in excess of 4.5 mm/m occur over Snowdonia in northwest Wales. It is of interest to note that, in an earlier study using 46 gauges from Snowdonia, Unwin (1969) found that the linear model

$$P_{\text{est}} = 1064 + 4.57(H) \dots (\text{mm})$$

fitted the data well ($R^2 = 0.66$) with a height coefficient close to that given by the GWR.

The Highlands of Scotland do, however, appear to have presented a problem for the GWR model. The largest GWR height coefficients identified for the Highlands are close to 4.5 mm/m, and are centred on Lochaber. Yet Harrison and Clark (1998) noted that there did not appear to be a strong relationship between altitude and precipitation in this area. There is a paucity of higher level stations in this area, and the orientation of the Great Glen from southwest to northeast could lead to higher low level precipitation totals than might otherwise be expected. It is also of note that if the Ben Nevis Observatory annual total (Thom, 1974) is used in conjunction with the totals at station around Fort William, the height coefficient




Figure 5. Contour map for the intercept constant

is close to 1.8 mm/m. This value is almost certainly too low, and there is likely to be a considerable underestimate of winter precipitation on the summit of Ben Nevis. Although these factors may explain why there is a higher than expected height coefficient in Lochaber, it does not explain why larger coefficients have not also been identified by the GWR to the northwest and to the south of that area (as noted in McClatchey, 1996).

The height coefficients in the eastern Highlands of Scotland of around 3.0 mm/m are much larger than expected when compared with OLS regressions on subsets of the data (McClatchey, 1996) which are in the range of 1.5–1.8 mm/m. Again, this may be that the GWR in adjusting across the whole country has led to an overestimate of the height coefficients in these areas, but it is worth noting that the high level gauges in the eastern Highlands are the ones known to be most prone to underestimating winter precipitation. This arises as the proportion of precipitation falling as snow is greater in the east than in the west. It may well be that the GWR estimates are closer to the 'true' value than OLS regression estimates based on subsets of the available data.

The difficulties with the GWR estimates of the height coefficient in the Highlands are also likely to be a result of the directional nature of precipitation events. The contribution to precipitation in the Highlands is greatest from the southwest sector, but this contribution decreases from 50% in the west to 35% in the east (Weston and Fowler, 1991). In addition, orographic enhancement of precipitation is not independent of direction (Weston and Roy, 1994) and, as specified, the GWR estimates are not adjusted for such downstream changes.

6. SPATIAL VARIATIONS IN THE ESTIMATED SEA LEVEL RAINFALL

Figure 5 shows the estimated intercept constant, c_0 , for Great Britain, contoured at 50 mm intervals from 600 to 1250 mm. It can be seen that the gradient anticipated in our second hypothesis is present. As might be expected, the intercept noted in the global regression, at 730.83 mm, lies along the demarcation of a 'lowland' and a 'highland' Great Britain, with the higher values as anticipated to the west of this line and the lower values, below 600 mm, to the east. Over Snowdonia, it suggests an intercept of around 750 mm, which is considerably less than the 1064 mm found by Unwin (1969).

As an obvious test of the GWR method we can compare this map with observed totals for selections of gauges close to sea level. Overall, the GWR intercept coefficients plotted in Figure 5 appear to provide a good match for precipitation totals at low lying gauges for England and much of Wales, and a fairly good match in eastern Scotland. The match appears to be less good in the west and northwest Highlands. While there are coastal stations in that area at which average annual precipitation totals are close to the highest values shown, they are further north than indicated, and most low lying coastal stations have higher totals than suggested by the intercept values. This difference between the GWR intercept coefficients and the observed totals may be a combination of the rapid change in average annual precipitation totals from west to east across Scotland, and the influence of the totals recorded on the Western Isles (included in the data but not shown in Figure 4 or Figure 5). The Western Isles totals are lower than those on the west coast of Scotland.

7. CONCLUSION

The GWR estimates of the relationship between the average annual rainfall and altitude across Great Britain clearly illustrate the non-stationary nature of that relationship which should perhaps be written:

$$P = (< 600 \text{ to } > 1250) + (0.0-5.0)*(H) + \varepsilon(\text{mm})$$

In this, we have given values for the range in both intercept and height constant. In addition, the approach has highlighted areas of rapid change in the height coefficient across lowland England that are not readily apparent using local regressions and that will repay further investigation. In comparison with previous approaches, the GWR model has the advantage of being much simpler than the more complex of the multiple regression approaches (Spackman, 1993; Basist *et al.*, 1994), and has the considerable advantage that any unusual features observed in the results are not a function of some complicated and possibly untraceable relationship between derived parameters in a multiple regression model, but can be examined in the way that we have done, by 'drilling down' to look at the local detail. Finally, the GWR estimates in conjunction with a digital elevation matrix of the country would enable the rapid production of objective rainfall maps of considerable use to ecologists and hydrologists.

What we have not yet examined is any temporal variation in these relationships through the year and between individual years. Further work is also being undertaken to assess the problem associated with downstream changes in the proportion of rainfall from different directions and any changes in the rainfall/altitude relationships. It should be noted that GWR also provides a useful method for incorporating the varying relationship between rainfall and altitude across the country into any predicted change in rainfall amounts. The climate of Great Britain is predicted to change over the next 50-100 years as a result of the global climate change resulting from increasing concentrations of carbon dioxide and other greenhouse gases. At present, predicted changes in precipitation in Great Britain are less reliable than the predictions for temperature, yet it is the precipitation changes that are likely to give rise to the greatest economic and social problems. Given the strong relationship between rainfall and altitude, it will be crucial that it is included in assessing the impact of climate change.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the efforts of the Climatic Research Unit at UEA in preparing the data and to the Meteorological Office (Glasgow Weather Centre) and BADC for the provision of additional information relating to rain gauge stations in the UK.

REFERENCES

- Atkinson BW, Smithson PA. 1976. Rainfall. In *The Climate of the British Isles*, Chandler TJ, Gregory S (eds). Longman: London; 129–182, Chapter 6.
- Ballentyne LK. 1983. Precipitation gradients in Wester Ross, north-west Scotland. *Weather* **38**: 379–387.
- Barry RG. 1992. *Mountain Weather and Climate* (2nd edn). Methuen: London.
- Basist A, Bell GD, Meentemeyer V. 1994. Statistical relationships between topography and precipitation patterns. *Journal of Climate* **7**: 1305–1315.
- Bastin G, Lorent B, Duque C, Gevers M. 1984. Optimal estimation of the average rainfall and optimal selection of raingauge locations. *Water Resources Research* **20**: 463–470.
- Bigg GR. 1991. Kriging and intraregional rainfall variability in England. *International Journal of Climatology* **11**: 663–675.
- Bleasdale A, Chan YK. 1972. Orographic influences on the distribution of precipitation. In *Distribution of Precipitation in Mountainous Areas*. World Meteorological Organisation: Geneva; 322–333.
- Brunsdon C. 1995. Estimating probability surfaces for geographical point data: an adaptive technique. *Computers and Geosciences* **21**: 877–894.
- Brunsdon C. 1998. Exploratory spatial data analysis and local indicators of spatial association with XLISP-STAT. *The Statistician* **47**: 471–484.
- Brunsdon C, Fotheringham S, Charlton M. 1996. Geographically weighted regression—modelling spatial non-stationarity. *Geographical Analysis* **28**: 281–289.
- Casetti E, Can A. 1999. The econometric estimation and testing of DARP models. *Journal of Geographical Systems* **1**: 91–106.
- Chun GK, Lockwood JG. 1974. An assessment of topographical controls on the distribution of rainfall in the central Pennines. *Meteorological Magazine* **103**: 275–287.
- Cleveland WS. 1979. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association* **74**: 829–836.
- Cleveland WS, Devlin SJ. 1988. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association* **83**: 596–610.
- Dore AJ, Choularton TW, Brown R, Blackall RM. 1982. Orographic enhancement in the mountains of the Lake District and Snowdonia. *Atmospheric Environment* **26A**: 357–371.
- Goovaerts P. 1999. Using elevation to aid the geostatistical mapping of rainfall erosivity. *Catena* **34**: 227–242.
- Harrison J, Clark G. 1998. The Lochaber raingauge network. *Weather* **53**: 352–357.
- Hastie TJ, Tibshirani RJ. 1990. *Generalised Linear Models*. Chapman and Hall: London.
- Hill FF, Browning KA, Bader MJ. 1981. Orographic rainfall over low hills and associated corrections to radar measurements. *Quarterly Journal of the Royal Meteorological Society* **107**: 642–670.
- McClatchey J. 1996. Spatial and altitudinal gradients of precipitation in Scotland. In *Hydrologie dans les pays celtique*, Merot P, Jogoret A (eds). INRA: Paris; 45–51.
- Pedgley DE. 1970. Heavy rainfalls over Snowdonia. *Weather* **25**: 340–350.
- Salter M de CS. 1918. The relation of rainfall to configuration. *British Rainfall* **58**: 40–56.
- Sawyer JS. 1956. The physical and dynamical problems of orographic rain. *Weather* **11**: 375–381.
- Shaw EM. 1962. An analysis of the origins of precipitation in northern England, 1956–1960. *Quarterly Journal of the Royal Meteorological Society* **88**: 539–547.
- Silverman BW. 1986. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall: London.
- Smithson PA. 1969a. Effects of altitude on rainfall in Scotland. *Weather* **24**: 370–376.
- Smithson PA. 1969b. Regional variations in the synoptic origins of rainfall across Scotland. *Scottish Geographical Magazine* **85**: 182–195.
- Spackman E. 1993. Calculation and Mapping of Rainfall Averages for 1961–1990. Bracknell: Meteorological Office, Observation Provision, Climate Studies Group, notes in summary of a talk given to the British Hydrological Society, 15 December 1993 at the University of Salford, England (cited with the permission of UK Meteorological Office).
- Taylor JA. 1976. Upland climates. In *The Climate of the British Isles*, Chandler TJ, Gregory S (eds). Longman: London; 264–287.
- Thom AS. 1974. Meteorological report. In *Aonach Mor. A planning Report on the Prospect of Winter Sport Development at Fort William*, Parnell BK, Thom AS (eds). Glasgow School of Art: Glasgow; 82–109.
- Unwin DJ. 1969. The areal extension of rainfall records. *Journal of Hydrology* **7**: 404–414.
- Unwin DJ. 1996. GIS, spatial analysis and spatial statistics. *Progress in Human Geography* **20**: 540–551.
- Weston KJ, Fowler D. 1991. The importance of orography in spatial patterns of rainfall acidity in Scotland. *Atmospheric Environment* **25A**: 1517–1522.
- Weston KJ, Roy MG. 1994. The directional-dependence of the enhancement of rainfall over complex topography. *Meteorological Applications* **1**: 267–275.