Geographical Analysis ISSN 0016-7363

"The Problem of Spatial Autocorrelation" and Local Spatial Statistics

A. Stewart Fotheringham

National Centre for Geocomputation, National University of Ireland,
Maynooth, Maynooth, Ireland

This article examines the relationship between spatial dependency and spatial heterogeneity, two properties unique to spatial data. The property of spatial dependence has led to a large body of research into spatial autocorrelation and also, largely independently, into geostatistics. The property of spatial heterogeneity has led to a growing awareness of the limitation of global statistics and the value of local statistics and local statistical models. The article concludes with a discussion of how the two properties can be accommodated within the same modelling framework.

Introduction

The challenge presented in this article is to address the question: "What impacts has the article by Cliff and Ord (1969) about spatial autocorrelation had on the development of local spatial statistics?" Given that the basis of spatial autocorrelation is *spatial dependency*, and the basis for local statistics is *spatial nonstationarity* and these two properties of spatial processes are not necessarily linked, the tempting answer is "Very little." The spatial autocorrelation measure developed in Cliff and Ord's article is, like the concept on which it is based, a global one. Local variants of spatial autocorrelation statistics were not developed until much later (e.g., Anselin (1995), paralleling local formulations of other statistics that possibly owed more to articles such as that by Casetti (1972), which promulgated the concept of spatial nonstationarity, more than this article by Cliff and Ord. In Cliff and Ord's article, for example, no mention is made of spatial nonstationarity, spatial heterogeneity, or spatially varying relationships.[1] However, the real answer to the preceding question is more complex, and the development of local statistical models owes much to the popularization of the concept of spatial dependency, which is the central theme of Cliff and Ord's (1969) article. Indeed, the two properties of spatial dependency and spatial heterogeneity are shown here to have a rather complex interaction: spatial dependency is the cornerstone of the solution to modeling spatial nonstationarity,

Correspondence: A. Stewart Fotheringham, Director, National Centre for Geocomputation, John Hume Building, National University of Ireland Maynooth, Maynooth, Co. Kildare, Ireland
e-mail: stewart.fotheringham@nuim.ie

and local statistical models can solve the problem of residual spatial dependency often encountered in global models.

Spatial dependency and spatial autocorrelation

Cliff and Ord were by no means the first to consider spatial autocorrelation, or even to derive statistics to measure it (see Moran 1950; Geary 1954; Whittle 1954, for example). However, Cliff and Ord (1969) did propose a more general measure of spatial autocorrelation, and placed this formulation on a firm statistical foundation so that it became the basis of much of the subsequent research in this area. The essence of spatial autocorrelation is that of spatial dependency: the situation whereby observations drawn from different locations are not independent of each other, with observations at locations nearer to each other being more similar (positive spatial autocorrelation) or less similar (negative spatial autocorrelation) than observations at locations farther apart.[2] Therefore, a key element in measuring and understanding spatial autocorrelation is a description of the relationship between the degree of similarity between observations and the distance separating them. This relationship is the same key concept underlying geostatistics, an area of research that has evolved largely independently from the work about spatial autocorrelation, and has developed its own vocabulary (Goovaerts 1997). The description of spatial dependency then can be used to define a spatial weighting function that relates the similarity of observations to the distance they are apart. Spatial autocorrelation is a measure of the strength and direction of this relationship.

However, despite being measures of spatial dependency, typically spatial autocorrelation statistics have been viewed as global descriptors of data. For example, the degree of spatial dependency in house prices across a city might be described by calculating a spatial autocorrelation statistic and interpreting its value. We would expect that two houses with identical attributes would have more similar prices if they were located close to each other than if they were located far apart. Of course, we do not know, and the spatial autocorrelation statistic does not indicate, to what extent spatial dependency in this case results from intrinsic variations in the perceptions of areas on the part of buyers, or to the extent to which people base the price of their houses on the sales prices of nearby houses. In either case, the measure of spatial autocorrelation is a global one, meaning that it measures the average spatial dependency across a region. Interesting spatial variations may well exist in the degree of spatial dependency in the data that are completely obscured by the calculation of one globally average measure. Only through the computation of local measures of spatial autocorrelation, such as those proposed by Anselin (1995), can we measure any spatial variation in the level of spatial dependency of an attribute.

Local statistics and local statistical models

Local spatial statistics are spatial disaggregations of more familiar global statistics. To make the distinction obvious, consider the mean temperature measured across the United States at a single point in time. This is a global statistic—it summarizes what is happening across the country in a single value—and of course it disguises a great deal of local variation. The local variant of this statistic could be a locally weighted mean of neighboring values, or even the individual temperature readings themselves across the United States, of which the global statistic is the mean. Typically, local statistics are represented by a surface of some description, whereas the global statistic is single-valued.

With data the distinction between global and local is easy to see; but, what is meant by "local" becomes less obvious in the context of something such as a regression model, in which we estimate relationships through calibration. The classic example of a regression model calibrated with spatial data produces global estimates of relationships between a dependent variable and each of a set of independent variables, ceteris paribus. That is, the single-valued estimates of each relationship are global averages of unknown sets of local relationships. A typical assumption, more for convenience rather than being based upon any logical reasoning, is that the relationships underlying the global average are constant across space—the assumption of spatial stationarity. However, this assumption is increasingly being questioned, and local variants of a regression, such as geographically weighted regression (Fotheringham, Charlton, and Brunsdon 2002), have been developed to allow local parameters to be estimated—in essence generating surfaces of local relationships that can be mapped and which are equivalent to the surface of US temperatures referred to above. Moreover, a large literature has developed about local spatial statistics and local spatial statistical models (Fotheringham 1997; Lloyd 2007).

How spatial dependency informs about spatial nonstationarity

Local statistics are calculated and local statistical models are calibrated by weighting observations according to their distance from a focal point such that the observations closer to the focal point are weighted more heavily in the calculation than are observations farther away. This procedure means that the statistic calculated is unique to a focal point, and that each time the location of the focal point changes, a new local statistic is calculated, generating a set of such values that can be displayed as a surface to show the extent of any spatial variation. In effect, this set of calculations creates a whole new geography—that of spatially varying relationships. Therefore, the calculation of local statistics and the calibration of local statistical models depend on a spatial weighting function by which observations at locations nearer to a focal point are given more weight than are observations farther away. That is, the construction of local statistics depends on spatial dependency—the property that observations closer to each other in geographic space are more likely to be similar than are observations farther apart.

The use of spatial dependency in this manner to create local models is not an arbitrary one. Consider a variable Y (of which y_i is a realization at location i) regressed on a set of variables $X_1, X_2, \ldots, X_n$ (of which $x_{1i}, x_{2i} \ldots, x_{ni}$ are realizations at location i), where the relationships between Y and each X are spatial nonstationary. The ideal situation is to have repeated observations of $Y, X_1, X_2, \ldots, X_n$ at each location i, and to use these measurement replicates to calibrate parameter estimates for each i. Unfortunately, this scenario usually is not possible, resulting in data being drawn from other locations for estimation purposes (the exchangeability assumption). But these pseudoreplicates introduce bias into parameter estimates if the processes being modeled are nonstationary. In order to reduce this bias, data need to be weighted according to their proximity to i, with data from nearby locations being weighted more heavily than data from locations farther away—hence utilizing the property of spatial dependency to calibrate local models.

How local models inform about spatial dependency

In a global regression model calibrated with spatial data, common practice is to examine the residuals of such a model for possible spatial autocorrelation. Spatial dependency among residuals raises suspicions about any inference based on the estimated parameters from the calibration of a model. Examination of the residuals is sometimes carried out visually by mapping them and examining the spatial distribution for any obvious clustering. Alternatively, it can be undertaken with various analytical graphical methods, or by computing an autocorrelation statistic and examining its significance. If the residuals from a regression model are found to exhibit spatial autocorrelation, various spatial regression models can be used that have been proposed to handle this problem (Schabenberger and Gotway 2005; Lloyd 2007). However, if spatial autocorrelation detected in a set of residuals is caused by spatial nonstationarity, then an alternative, and arguably more informative solution, involves the calibration of local rather than global models.[3] To illustrate this situation, suppose a global model (kept very simple, for illustrative purposes) of the form

$$y_i = \alpha + \beta x_i \tag{1}$$

where α is the global intercept and β the global regression slope parameter is calibrated using data generated from a process that is spatially varying—that is, where the true model is

$$y_i = \alpha + \beta_i x_i \tag{2}$$

where β_i is the local regression slope parameter, varies spatially and exhibits spatial dependence,[4]—then the estimate of β from incorrectly calibrating equation (1) is a weighted average of the β_i values.

Assuming there are no other sources of error in the model specification (this is not necessary for the argument), then the estimate of y_i from calibrating the global model given by equation (1) is $\alpha^*+\beta^* x_i$, where * denotes an estimated value. If equation (2) represents the true process, and assuming that β_i is positive for all i, then for locations where $\beta_i > \beta^*$, the global model underestimates y_i and the corresponding residual is positive. For locations where $\beta_i < \beta^*$, the global model overestimates y_i and the corresponding residual is negative. Where $\beta_i = \beta^*$, the global model correctly estimates y_i and the corresponding residual is zero. Hence, if β_i exhibits spatial dependency (of note is that this is the specification used in hierarchical Bayesian spatial models), so will the residuals of a global model calibrated with data generated by a spatially nonstationary process. Consequently, calibrating local models removes this problem of spatial dependency in residuals, as demonstrated by Fotheringham, Charlton, and Brunsdon (2002).

Combing spatial dependency and spatial nonstationarity

While spatial regression models incorporate the effects of various types of spatial dependency, and local models can account for spatial nonstationarity, modeling both spatial dependence and spatial nonstationarity within the same framework is possible. This end can be achieved, for example, in a geographically weighted version of a spatial regression model. For instance, a geographically weighted spatial lag model yields local measures of spatial autocorrelation through the estimated parameters on the spatially lagged *y* variable. Therefore, such models account for local spatial dependence as well as provide information about the degree of spatial nonstationarity in the processes being modelled.

Summary

Cliff and Ord's (1969) article, despite being exclusively concerned with a global statistic, helped popularize the concept of spatial autocorrelation as a measure of spatial dependency, the latter being a key element in the development of local statistics and local spatial models that are used to assess spatial nonstationarity. Indeed, seeing how the two concepts can be combined is easy. Local variants of spatial regression models, such as spatial lag and spatial error models, can be formulated so that both spatial nonstationarity and spatial dependency are modeled simultaneously. Such a model allows not only the relationships within the model, but also the degree of spatial dependency in either the dependent variable or the residuals, to vary spatially. After 40 years, the weakness in the calculation of any spatial autocorrelation statistic remains the often subjective nature of the definition of a spatial weights matrix, representing the scale of the spatial dependency being measured. Local spatial models overcome this problem to some extent by determining an optimal spatial weights matrix based on model goodness of fit. What therefore is surprising is that more use is not made of a geographically weighted version of a spatial lag model, which would allow more objective estimation of local spatial autocorrelation statistics through estimates of local parameters on the lagged dependent variable term. Thus the essence of Cliff and Ord's (1969) article continues to have an impact on our thinking about modeling spatial processes.

Notes

1 Interestingly, the article preceding that of Cliff and Ord in the same issue of *London Papers in Regional Science* by Granger (1969), subsequently a Nobel Prize winner in Economics, does raise the issue of spatial nonstationarity.

2 Positive spatial autocorrelation is much more frequently encountered than negative spatial autocorrelation, and indeed some of what has been measured as negative spatial autocorrelation possibly reflects the scale at which observations are made being much greater than the scale at which processes operate; henceforth in this paper, all references to spatial autocorrelation refer to positive spatial autocorrelation.

3 Of course, not all residual spatial autocorrelation is caused by applying a global model to a spatially nonstationary process; other causes of spatially autocorrelated residuals exist. What is unknown is the extent to which commonly observed residual autocorrelation is caused by the incorrect application of global models.

4 For the sake of clarity here, we allow only the slope parameter to vary spatially, although this makes no difference to the argument.

References

Anselin, L. (1995). ''Local Indicators of Spatial Association—LISA.'' *Geographical Analysis* 27, 93–115.

Casetti, E. (1972). ''Generating Models by the Expansion Method: Applications to Geographic Research.'' *Geographical Analysis* 4, 81–91.

Cliff, A. D., and J. K. Ord. (1969). ''The Problem of Spatial Autocorrelation.'' In *London Papers in Regional Science*, 25–55, edited by A. Scott. London: Pion.

Fotheringham, A. S. (1997). ''Trends in Quantitative Methods I: Stressing the Local.'' *Progress in Human Geography* 21, 88–96.

Fotheringham, A. S., M. Charlton, and C. Brunsdon. (2002). *The Geography of Parameter Space: The Analysis of Spatially Varying Relationships*. Chichester, UK: Wiley.

Geary, R. C. (1954). ''The Contiguity Ratio and Statistical Mapping.'' *The Incorporated Statistician* 5, 115–45.

Goovaerts, P. (1997). *Geostatistics for Natural Resources Evaluation*. New York: Oxford University Press.

Granger, C. W. J. (1969). ''Spatial Data and Time Series Analysis.'' In *London Papers in Regional Science*, 1–24, edited by A. Scott. London: Pion.

Lloyd, D. (2007). *Local Models for Spatial Analysis*. Boca Raton, FL: CRC Press.

Moran, P. A. P. (1950). ''Notes on Continuous Stochastic Phenomena.'' *Biometrika* 37, 17–23.

Schabenberger, O., and C. A. Gotway. (2005). *Statistical Methods for Spatial Data Analysis*. Boca Raton, FL: Chapman and Hall/CRC.

Whittle, P. (1954). ''On Stationary Processes in the Plane.'' *Biometrika* 41, 434–49.