Multicollinearity in geographically weighted regression coefficients: Results from a new simulation experiment

Antonio Páez*

Associate Professor
Centre for Spatial Analysis/School of Geography and Earth Sciences
McMaster University

Steven Farber

Graduate Student
Centre for Spatial Analysis/School of Geography and Earth Sciences
McMaster University

David Wheeler

Research Fellow
Cancer Prevention Fellowship Program
National Cancer Institute

* corresponding author e-mail: paezha@mcmaster.ca

Abstract. Multicollinearities among the coefficients obtained from the method of geographically weighted regression have been identified in recent research. This is a serious issue that poses a critical challenge for the utility of the method as a tool to investigate multivariate relationships. The evidence regarding the ability of GWR to retrieve spatially varying processes remains mixed due to partial and inconclusive experiments. The objective of this paper is to provide stronger support to the thesis that multicollinearities are inherent to the method. This objective is accomplished by: 1) Investigating multicollinearity in situations where the underlying process is stationary and non-stationary; and 2) Using advanced visualization to report results for a range of outcomes as opposed to the average of *r* replications as in previous research. Extensive simulation experiments that test two different implementations of GWR provide evidence of spurious multicollinearity between local regression coefficients. This suggests that extreme caution should be exercised when drawing conclusions regarding spatial relationships retrieved using this modeling approach.

Keywords: Geographically weighted regression, multicollinearity, locally linear estimation, simulation, goodness-of-fit, inference

JEL: C13, C21

Introduction

Since its introduction to the geographical and spatial econometric literature in 1996 (Brunsdon et al, 1996; McMillen, 1996), the non-parametric approach termed geographically/locally weighted regression (GWR/LWR) has become a popular tool for the study of geo-referenced data. Geographical weights in spatial analysis are introduced with the object of producing sub-samples of observations that can then be analyzed using linear and non-linear regression analysis (e.g. Atkinson et al, 2003; Paez, 2006). The principle of geographical weighting is simple, and can in fact be applied for spatial segmentation prior to any form of statistical analysis, for example to derive geographically weighted descriptive statistics (Brunsdon et al, 2002), discriminant analysis (Brunsdon et al, 2007), visualization (Dykes and Brunsdon, 2007), etc. The notion of sub-sampling, while innovative in the spatial analysis literature, was not new at the time in statistical analysis, and the key development in the case of the papers by Brunsdon et al. (1996) and McMillen (1996) was to transpose a technique that was more generally applied to variable space, and introduce it for the specific case of analysis in geographical space. Original references in these papers trace a lineage that connects with the work of Cleveland and Devlin (1988), which itself dates back to earlier work on smoothing techniques for histograms (Cleveland, 1979). Interestingly, an earlier reference (Pelto et al, 1968) had in fact proposed the use of weighted regression, using weights that depended on geographical distance between observations, as a way to improve automated interpolation (i.e. contouring) algorithms. With the exception of Farwig (1986), Muller (1996), and Fedorov et al. (1999), this reference appears to have been largely missed by later statistical and spatial analysis research. The (re)discovery that weighted approaches could be used in the case of geographical information turned out to be pivotal, coming as it did on the heels of other influential work in spatial analysis that was beginning to develop local autocorrelation statistics (Getis and Ord, 1993; Ord and Getis, 1995; Anselin, 1995). The ability to work with spatial sub-samples opened the door for local spatial analysis within a multivariate framework, which provided a valuable counterpart to the univariate (i.e. autocorrelation) analysis enabled by the G_i and LISA statistics.

The work on geographically and locally weighted techniques can be fairly characterized as having had an important influence in spatial analysis in the course of the intervening decade. Together, the papers by Brunsdon et al. and McMillen had been cited 137 times by the end of 2008 (according to ISI Web of Science). The Geographical Analysis paper by Brunsdon et al., not being online, is not picked up by Google Scholar, but a related paper (Brunsdon et al, 1998) and the afore mentioned paper by McMillen collect 325 citations there as of June 2009. The pattern of citations (as per ISI) reflects the typical lag before a technique is adopted in the literature (see Figure 1a), and a rapid increase in the number of papers that further develop, discuss, or apply geographically weighted regression can be observed in the last 5 years (Figure 1b). Many of the early papers were published in leading spatial analysis-oriented journals, including Geographical Analysis (e.g. Brunsdon et al, 1996; Atkinson et al, 2003; Páez, 2004), the Journal of Urban Economics (e.g. McMillen, 1996), Environment and Planning A (e.g. Fotheringham et al, 1998; Leung et al, 2000; Páez et al, 2002a; Páez et al, 2002b), and the Journal of Regional Science (e.g. Brunsdon et al, 1999), among others. The approach has since migrated to a number of subject areas, including environmental studies and sciences (e.g. Cho et al, 2007; Tu and Xia, 2008), physical geography (e.g. Propastin and Kappas, 2008), forestry (e.g.

Zhang and Shi, 2004), planning and development (e.g. Yu, 2006; Ali et al, 2007; Bitter et al, 2007), remote sensing (e.g. Foody, 2005), and so forth.

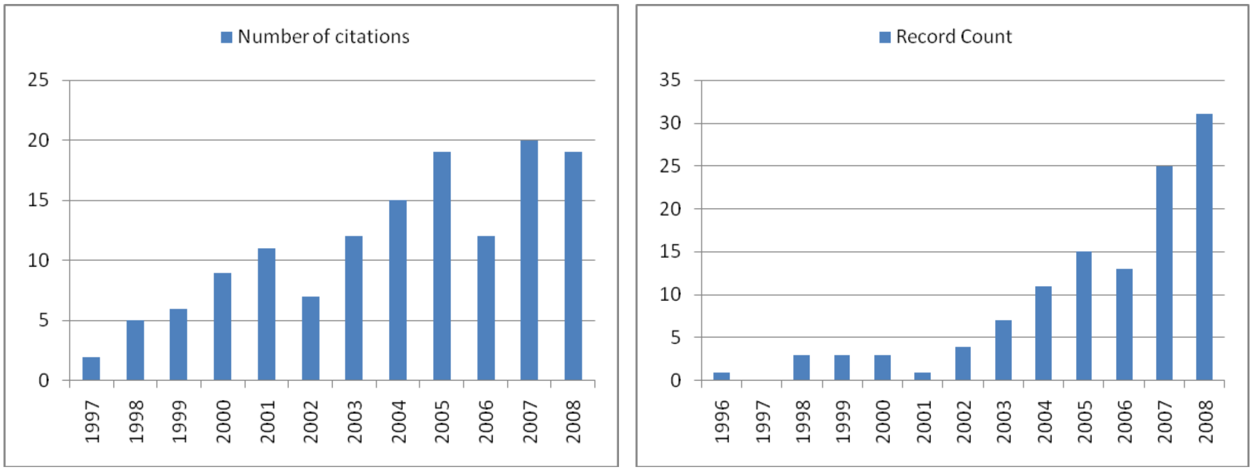

Figure 1. a) ISI cites to Brunsdon et al. (1996) and McMillen (1996) (left panel) b) Geographically weighted regression publications (right panel)

Weighted regression approaches were initially intended for interpolation and curve-fitting applications (e.g. Pelto et al, 1968; Cleveland, 1979). Later work propounded the idea that locally weighted regression could be used in addition to study correlative effects in regression. Cleveland and Devlin (1988), for instance, identify and illustrate three major uses of the local-fitting methodology: as an exploratory graphical tool; as a diagnostic tool to check the adequacy of parametric models; and as a replacement for parametric regression, by using instead the estimated locally weighted surface. Some geographical applications have remained in the realm of curve-fitting and interpolation. McMillen (2003), for example, illustrates how the non-parametric approach can be used to provide a better fit to the data, and to diagnose parametric model misspecification. Páez et al. (2008) use a case study to compare a number of techniques, and find that GWR provides a more accurate interpolator (as measured by an ex-sample validation exercise) relative to the alternative approaches investigated. By and large, however, the emphasis in research so far has been not in curve-fitting or interpolation, but rather on the study of spatially-varying *relationships* (e.g. Fotheringham et al, 2002), and this is the way the method is more widely understood in the literature. Geographically weighted regression papers thus typically include a global model to be used as a benchmark, and maps showing GWR coefficients and their variation over space. These maps tend to motivate fairly extended discussions of the implications of the patterns discovered. For instance, in a paper that investigates the level of regional industrialization in Jiangsu in China (Huang and Leung, 2002), the coefficient associated with fixed capital investment per unit of GDP is interpreted thusly: "The fixed investment...had the smallest effect...in the southern areas" which "did not rely very much on the amount of capital investment". This may be reflective of a regional economy that is "affected by more advanced factors such as technology". In a paper that investigates land use change as a matter of proximity to transportation infrastructure Páez (2006) interprets the variation of the coefficients in this way: "Land use change is less likely to take place in the immediate surroundings of the station; a clear east-west trend suggest that land use change is a more likely outcome to the east of the station". The power of these types of insights in terms of assessing policy and generating knowledge is perhaps one of the reasons why the approach has become popular (see Ali et al, 2007).

Geographically weighted regression has not been without critics. Unfortunately, very large variations in the values of the coefficients are observed, and not infrequently this includes sign reversals that make interpretation difficult. For example, in the analysis of species richness due to Foody (2004), the relationship between species diversity and a vegetation index appears to be positive for most of Africa, but negative for a large patch in the central parts of the continent. Páez (2003), on the other hand, reports a sign reversal for the coefficient of NOx concentration levels in a hedonic price model, suggesting that more pollution associates both with lower *and* higher housing prices depending on location. The issue of sign reversals is noted by Jetz et al. (2005) in a commentary where they raise the possibility that the relationships may in fact be global and that local variability observed may be due to missing variables or interaction terms, or be attributed to the excessive flexibility of the method. In other words, the argument is made that spatially-varying relationships may not be legitimate but rather a consequence of model misspecification or, even worse, an artefact of the method. Jetz et al. (2005) are among the first to note that GWR coefficients display a curious tendency to be correlated with each other (see Figure 2 panels A-D in Foody, 2004). The question of multicollinearity (i.e. excessive correlation) among GWR coefficients is discussed in much more depth in a cotemporaneous paper by Wheeler and Tiefelsdorf (2005), who provide evidence of multicollinearity between the coefficients of a model that investigates bladder cancer mortality data. In the words of these authors (p. 162), these correlations "invalidate any meaningful interpretation...because the regression coefficients are no longer uniquely defined".

The key question at the center of these arguments is whether the variability observed is legitimate, that is, whether it is an accurate representation of an underlying geographical process. Case studies showing that GWR coefficients are frequently correlated with each other may motivate the question, but are not sufficient to settle it because the true underlying processes are unknowable. This is the reason why regression analysis is used in the first place, to try to retrieve a process. If the true process is not known, it is impossible to determine whether the spatially-varying coefficients are reflective of it or instead whether they are somehow unduly introduced by the mechanics of the approach. The inability to resolve this question based on empirical examples is understood by Wheeler and Tiefelsdorf (2005), who set out to conduct controlled numerical experiments to elucidate the point. As we show below, the experimental design devised by these authors is such that the results are in fact inconclusive. At the same time, other evidence (e.g. Paez, 2005; Farber and Páez, 2007; Wang et al, 2008) does not indicate the presence of multicollinearity among GWR coefficients, although presentation of results may have not been sufficient to uncover this effect.

Given: 1) the seriousness of the argument, which may importantly determine whether GWR is more appropriately seen as an exploratory tool or a technique to draw inferences regarding multivariate relationships; and 2) the fact that GWR has increasingly become established as a tool in the investigation of spatially-varying relationships; our objective in this paper is to present evidence to more conclusively settle the question concerning the ability of GWR to retrieve spatially-varying multivariate processes. Since, as we noted above, cases studies are not sufficient to demonstrate the question one way or the other, we approach this objective by means of an extensive set of controlled numerical experiments, whereby the true process is known and the results can then be contrasted for reliability. The experiments are designed in such a way as to overcome the limitations of previous work, in particular the simulations reported by Wheeler and Tiefelsdorf (2005), and results are reported using a more advanced visualization format that is better suited to reveal the extent of variability in retrieved coefficients compared to the (mean) surfaces reported by Páez (2005), Farber and Páez (2007), and Wang et al. (2008). The results of our experiments provide sobering evidence that caution should be exercised when attempting to interpret coefficients obtained using GWR in its current implementation based on existing validation and estimation techniques.

Background

Wheeler and Tiefelsdorf (2005) design a simulation experiment to investigate the question of multicollinearity in GWR regression coefficients. A concern in said paper is that multicollinearity in the explanatory variables is potentially responsible for this effect. In order to control for this potentially confounding effect, variables are selected that are known to be uncorrelated. The experiment is conducted using an irregular tessellation based on the administrative partition of counties in Georgia, US (n=159). In order to obtain variables with zero correlation, the binary connectivity matrix **W** that describes the spatial relationship of counties in Georgia is used. The eigenvectors of the following transformation of matrix **W**:

$$\left(\mathbf{I}-\frac{\boldsymbol{11'}}{n}\right)\frac{n}{\sum_i\sum_j w_{ij}}\mathbf{W}\left(\mathbf{I}-\frac{\boldsymbol{11'}}{n}\right) \quad (1)$$

are known to represent latent map patterns embedded in the geographical system (Griffith, 2003). Since eigenvectors are uncorrelated by design, their use facilitates the controlled variation of the level of correlation between the explanatory variables in the experiment. The data generating process is defined in the following way:

$$\boldsymbol{y}=\beta_0+\beta_1\boldsymbol{x}_1+\beta_2\boldsymbol{x}_2+\boldsymbol{\varepsilon} \quad (2)$$

with:

$$\beta_0=\beta_1=\beta_2=1 \quad (3)$$

Eigenvector 1 ($\mathbf{e_1}$) is selected to represent variable vector $\boldsymbol{x}_1$, whereas eigenvector 61 ($\mathbf{e_{61}}$), with a known level of autocorrelation very close to zero, is standardized and used in lieu of random terms $\boldsymbol{\varepsilon}$. In order to introduce controlled levels of correlation between variables $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ the latter variable is defined in two different ways for successive experiments, using eigenvectors 3 and 4 ($\mathbf{e_3}$ and $\mathbf{e_4}$):

$$\boldsymbol{x}_2=\begin{cases}\sin(\theta)\mathbf{e_3}+\cos(\theta)\mathbf{e_1}\\ \sin(\theta)\mathbf{e_4}+\cos(\theta)\mathbf{e_1}\end{cases} \quad (4)$$

Parameter θ is varied between -0.95 and 0.95 in the experiments to induce levels of correlation for $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ between -0.81 and 0.81. Using the date generated in this way, GWR coefficients are estimated using a bi-square nearest neighbor weighting function and bandwidth sizes S=40, 143, and 159. The results indicate that there are substantial levels of correlation between the local coefficient estimates when the two independent variables are uncorrelated, and that these levels tend to increase as the independent variables display increasing levels of correlation between them.

The experiment is intriguing, but ultimately inconclusive for various reasons. First, unlike many other controlled numerical experiments which use random numbers to introduce a level of stochasticity in the data generation process, there is no randomness involved in this case. There is only one instance of each $\boldsymbol{x}_1$ and $\boldsymbol{\varepsilon}$, and for a given level of θ, only one instance of $\boldsymbol{x}_2$. More critical is the fact that the data generating process is based on a global, not a local formulation of the model. The regression coefficients of the true process are all spatial constants, set to 1. GWR models, moreover, are not typically estimated using an arbitrary bandwidth size S, but rather the bandwidth is calibrated using some sort of cross-validation procedure. The experiment is successful in demonstrating how the vectors of local coefficients are correlated. However, this should not be altogether surprising, considering that the true coefficients in the data generating process are all identical. A more interesting question is whether estimation of the GWR coefficients reverts to the global model of the true data generating process, in which case the question of multicollinearity becomes moot. We replicate the experiment, but instead of using the bandwidths reported in the paper, we calibrate the kernel bandwidth using a cross-validation procedure. The results indicate that in each case, for both definitions of $\boldsymbol{x}_2$ and every level of correlation θ, the cross-validation procedure identifies a bandwidth size of 158 nearest neighbors as appropriate. Since the sample size is n=159, and the cross-validation procedure leaves one observation out at a time, this means that the GWR model should be estimated using all observations in the sample. A typical cross-validation figure is shown in Figure 2, where it can be seen that the tendency is to revert towards the global model.

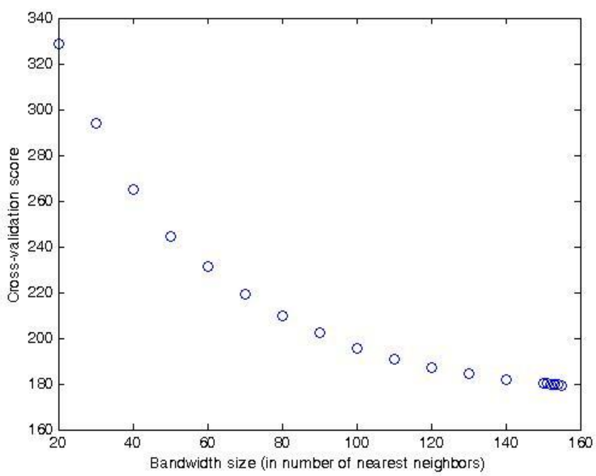

Figure 2. Typical cross-validation plot using the parameters in Wheeler and Tiefelsdorf (2005)

The experiment shows that if GWR is estimated appropriately, the simulation parameters used in Wheeler and Tiefelsdorf (2005) cannot possibly be used to demonstrate the existence of multicollinearity between local coefficients. Other experiments reported in the literature, on the other hand, do not support the thesis that multicollinearity is an artifact of the method, and the level of correspondence between the true coefficient surface used in the data generating process and the surface obtained from averaging the large number of replications typically used when random variables are generated for the simulation, appears to be quite high (see Figures 1-5 in Páez, 2005; Figures 8-12 in Farber and Páez, 2007; and Figures 1-2 and 7-10 in Wang et al, 2008). The operational term here is *average*. The simulations reported in these papers did not focus on multicollinearity and the results were clearly reported without this issue in mind.

However, while the average of R replications appears to closely resemble the true coefficient surface, it is unclear what the situation is at the level of individual replications. The experiments described and reported in the reminder of this paper intend to clarify the ambiguities left unresolved by previous research.

Methods and Experimental Design

We begin our experiments by defining a data generating process, as follows:

$$y_i = \beta_{0i} + \beta_{1i}x_{1i} + \beta_{2i}x_{2i} + \varepsilon_i \qquad (5)$$

The dependent variable y_i is the outcome of a linear combination of the value of spatial variables x_1 and x_2 at location i, possibly spatially-varying coefficients β_{ki} (sub-indexed by i), and a random term ε_i. For comparability, in the experiments we use the same geographical system employed by Wheeler and Tiefelsdorf (i.e. counties in Georgia). Use of an irregular tessellation if further justified by results reported by Farber et al. (2009) indicating that regular tessellations frequently used in experimental spatial statistics are not, from a topological perspective, representative of real geographical systems. The independent variables are generated from random draws of the standard normal distribution, while the vector of error terms is obtained from draws of the normal distribution with mean 0 and standard deviation 0.25. In order to introduce varying levels of correlation between variables $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, we use a similar transformation to that used by Wheeler and Tiefelsdorf (2005):

$$x_2^{\theta} = \sin(\theta)x_1 + \cos(\theta)x_2 \qquad (6)$$

The levels of correlation between the explanatory variables are approximate in this case, and depend to some extent on the initial level of correlation of the random draws, which is usually fairly low, but never exactly zero. Random variables are drawn a total of R=100 times (i.e. replications) for each level of θ between -0.75 and 0.75 in 0.25 increments. After selecting coefficients β as appropriate, variable $\boldsymbol{y}$ is generated using model (5), so that there are 100 unique realizations of the process for each level of correlation between the independent variables.

We devise two variations for our experiments. In the first one, we are interested in the false positive rate, that is, the frequency with which validation of the weighting function fails to collapse into the global model when the true process is indeed global. For this experiment, the coefficients are spatial constants as follows:

$$\beta_{0i} = \beta_0 = \beta_{1i} = \beta_1 = \beta_{2i} = \beta_2 = 1 \qquad (7)$$

For the second experiment, we are interested in the ability of GWR to retrieve a spatially-varying process. Following Wheeler and Tiefelsdorf (2005), we adopt the eigenvectors of the matrix in expression (1). However, instead of using the eigenvectors to represent the independent variables in the data generating process, we use them in place of the spatially-varying coefficients. The advantage of using the eigenvectors for the spatially-varying coefficients is that the coefficients are uncorrelated by design. Three eigenvectors with distinctive patterns are selected as follows: $\boldsymbol{\beta}_0=\mathbf{e_1}$, $\boldsymbol{\beta}_1=\mathbf{e_3}$, $\boldsymbol{\beta}_2=\mathbf{e_4}$, as shown in Figure 3. Please note that the coordinates have been standardized to a unit square.

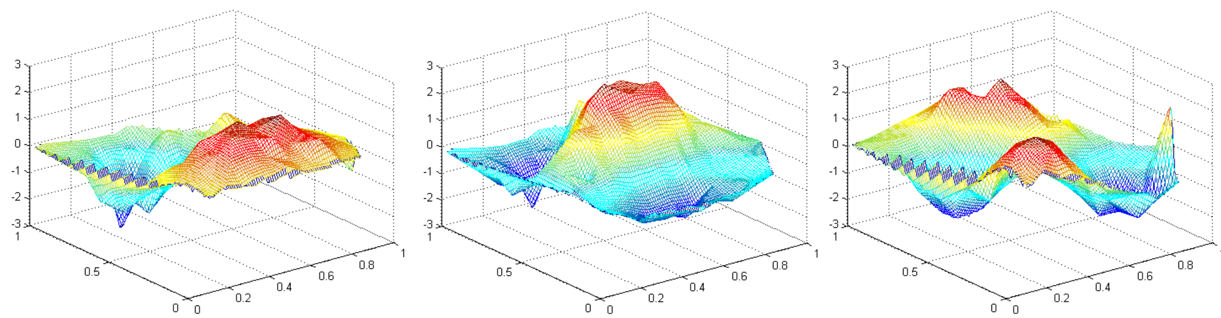

Figure 3. Spatially varying coefficients β_0, β_1, and β_2

Estimation of GWR models involves a number of steps. First, a kernel function is selected for the analysis. For this, we adopt the bi-square nearest neighbor formulation. Next, the kernel bandwidth must be calibrated, something that is typically done by means of a leave one out cross-validation procedure. Two different forms of the cross-validation score are used here. The first one is well known from the literature (see Brunsdon et al, 1996):

$$CVS(s) = \sum_i \left(y_i - \hat{y}_{\neq i}(s)\right)^2 \tag{8}$$

The cross-validation score is the sum of squared differences between the observed value of y at i, and the value predicted by a model estimated using kernel bandwidth s after removing observation i from the sample. A modified version of the cross-validation score is discussed by Farber and Páez (2007) that helps to reduce the dominance of some observations in determining the window size. The score in this case is based on a decomposition of the cross-validation calculations. A partial score for each observation i is calculated and converted to a value between 0 and 1 as follows:

$$CVS_i(s) = \frac{\left(y_i - \hat{y}_{\neq i}(s)\right)^2}{\sum_s \left(y_i - \hat{y}_{\neq i}(s)\right)} \tag{9}$$

and the total score is simply:

$$CVS(s) = \sum_i CVS_i(s) \tag{10}$$

Standardization of the score in this way ensures that all observations exert an identical level of influence across bandwidth sizes, and so the selection of the bandwidth is not influenced by a few observations with atypically large errors. Coefficients estimated using bandwidths selected using this form of the cross-validation score tend to display more compact distributions (i.e. less extreme variability). For further details, please refer to Farber and Páez (2007).

Linked to the issue of validating a bandwidth size is the estimation of the coefficients. We investigate two coefficient estimation techniques. Again, the first one is well known from the literature, and involves employing a diagonal matrix of geographical weights to obtain locally weighted estimates as follows (see Brunsdon et al, 1996):

$$\boldsymbol{\beta}_i = \left(\mathbf{X}'\mathbf{G}_i\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{G}_i\boldsymbol{y} \quad (11)$$

The vector of local estimates is a function, in addition to the data matrix **X** and vector ***y***, of $\mathbf{G}_i$, a of geographic weights where diagonal element g_i is the kernel function implemented with the cross-validated kernel bandwidth. The second implementation of the method is due to Wang et al. (2008), and introduces a local linear approximation based on the Taylor's expansion of the spatially varying coefficient β_k at location i with coordinates (u_i,v_i). Under the assumption that each of the coefficients has second continuous partial derivatives with respect to the coordinates, the coefficient can be locally approximated by the following expression:

$$\beta_k(u,v) \approx \beta_k(u_0,v_0) + \beta_k^{(u)}(u_0,v_0)(u-u_0) + \beta_k^{(v)}(u_0,v_0)(v-v_0) \quad (12)$$

where $\beta_k^{(u)}(u_0,v_0)$ and $\beta_k^{(v)}(u_0,v_0)$ are the partial derivatives of the coefficient with respect to u and v at the location. We consider this method since the experimental results reported by Wang et al. (2008) indicate that this approach outperforms the typical formulation of GWR, i.e., that shown in equation (11).

False Positive Rate

Retrieving a Known Spatial Process

Implications and Concluding Remarks

References

Ali K, Partridge MD, Olfert MR, 2007, "Can geographically weighted regressions improve regional analysis and policy making?" *International Regional Science Review* **30** (3) 300 - 329.

Anselin L, 1995, "Local Indicators of Spatial Association - LISA" *Geographical Analysis* **27** 93 - 115.

Atkinson PM, German SE, Sear DA, Clark MJ, 2003, "Exploring the relations between riverbank erosion and geomorphological controls using geographically weighted logistic regression" *Geographical Analysis* **35** (1) 58 - 82.

Bitter C, Mulligan GF, Dall'erba S, 2007, "Incorporating spatial variation in housing attribute prices: A comparison of geographically weighted regression and the spatial expansion method" *Journal of Geographical Systems* **9** (1) 7 - 27.

Brunsdon C, Fotheringham AS, Charlton M, 1999, "Some notes on parametric significance tests for geographically weighted regression" *Journal of Regional Science* **39** (3) 497 - 524.

Brunsdon C, Fotheringham AS, Charlton M, 2002, "Geographically weighted summary statistics - a framework for localised exploratory data analysis" *Computers, Environment and Urban Systems* **26** (6) 501 - 524.

Brunsdon C, Fotheringham AS, Charlton ME, 1996, "Geographically weighted regression: A method for exploring spatial nonstationarity" *Geographical Analysis* **28** (4) 281 - 298.

Brunsdon C, Fotheringham S, Charlton M, 1998, "Geographically weighted regression - modelling spatial non-stationarity" *Journal of the Royal Statistical Society Series D-the Statistician* **47** 431 - 443.

Brunsdon C, Fotheringham S, Charlton M, 2007, "Geographically weighted discriminant analysis" *Geographical Analysis* **39** (4) 376 - 396.

Cho SH, Chen Z, Yen ST, English BC, 2007, "Spatial variation of output-input elasticities: Evidence from Chinese county-level agricultural production data" *Papers in Regional Science* **86** (1) 139 - 157.

Cleveland WS, 1979, "Robust Locally Weighted Regression and Smoothing of Scatterplots" *Journal of the American Statistical Association* **74** 823 - 836.

Cleveland WS, Devlin SJ, 1988, "Locally Weighted Regression - An Approach to Regression-Analysis by Local Fitting" *Journal of the American Statistical Association* **83** (403) 596 - 610.

Dykes J, Brunsdon C, 2007, "Geographically weighted visualization: Interactive graphics for scale-varying exploratory analysis" *Ieee Transactions on Visualization and Computer Graphics* **13** (6) 1161 - 1168.

Farber S, Páez A, 2007, "A Systematic Investigation of Cross-Validation in GWR Model Estimation: Empirical Analysis and Monte Carlo Simulations" *Journal of Geographical Systems* **9** (4) 371 - 396.

Farber S, Páez A, Volz E, 2009, "Topology and dependency tests in spatial and network autoregressive models" *Geographical Analysis* **41** (2) 158 - 180.

Farwig R, 1986, "Multivariate Interpolation of Arbitrarily Spaced Data by Moving Least-Squares Methods" *Journal of Computational and Applied Mathematics* **16** (1) 79 - 93.

Fedorov VV, Montepiedra G, Nachtsheim CJ, 1999, "Design of experiments for locally weighted regression" *Journal of Statistical Planning and Inference* **81** (2) 363 - 382.

Foody GM, 2004, "Spatial nonstationarity and scale-dependency in the relationship between species richness and environmental determinants for the sub-Saharan endemic avifauna" *Global Ecology and Biogeography* **13** (4) 315 - 320.

Foody GM, 2005, "Mapping the richness and composition of British breeding birds from coarse spatial resolution satellite sensor imagery" *International Journal of Remote Sensing* **26** (18) 3943 - 3956.

Fotheringham AS, Brunsdon C, Charlton M, 2002, *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships* (Wiley, Chichester)

Fotheringham AS, Charlton ME, Brunsdon C, 1998, "Geographically weighted regression: a natural evolution of the expansion method for spatial data analysis" *Environment and Planning A* **30** (11) 1905 - 1927.

Getis A, Ord JK, 1993, "The Analysis of Spatial Association by Use of Distance Statistics (Vol 23, Pg 189, 1991)" *Geographical Analysis* **25** (3) 276 - 276.

Griffith DA, 2003, *Spatial Autocorrelation and Spatial Filtering: Gaining Understanding Through Theory and Scientific Visualization* (Springer-Verlag, Berlin)

Huang Y, Leung Y, 2002, "Analysing regional industrialisation in Jiangsu province using geographically weighted regression" *Journal of Geographical Systems* **4** (2) 233 - 249.

Jetz W, Rahbek C, Lichstein JW, 2005, "Local and global approaches to spatial data analysis in ecology" *Global Ecology and Biogeography* **14** (1) 97 - 98.

Leung Y, Mei CL, Zhang WX, 2000, "Testing for spatial autocorrelation among the residuals of the geographically weighted regression" *Environment and Planning A* **32** (5) 871 - 890.

McMillen DP, 1996, "One hundred fifty years of land values in Chicago: A nonparametric approach" *Journal of Urban Economics* **40** (1) 100 - 124.

Mcmillen DP, 2003, "Spatial autocorrelation or model misspecification?" *International Regional Science Review* **26** (2) 208 - 217.

Muller WG, 1996, "Optimal design for local fitting" *Journal of Statistical Planning and Inference* **55** (3) 389 - 397.

Ord JK, Getis A, 1995, "Local Spatial Autocorrelation Statistics - Distributional Issues and An Application" *Geographical Analysis* **27** (4) 286 - 306.

Páez A, 2003, "Investigating heterogeneity in hedonic price models using geographically weighted regression" Paper presented at the 50th Annual North American Meetings of the Regional Science Association International, Philadephia, 20 - 22 November.

Páez A, 2004, "Anisotropic Variance Functions in Geographically Weighted Regression Models" *Geographical Analysis* **36** (4) 299 - 314.

Paez A, 2005, "Local analysis of spatial relationships: A comparison of GWR and the expansion method" *Computational Science and Its Applications - Iccsa 2005, Pt 3* **3482** 162 - 172.

Páez A, 2005, "Local Analysis of Spatial Relationships: A Comparison of GWR and the Expansion Method" *Lecture Notes in Computer Science* **3482** 162 - 172.

Paez A, 2006, "Exploring contextual variations in land use and transport analysis using a probit model with geographical weights" *Journal of Transport Geography* **14** (3) 167 - 176.

Páez A, Long F, Farber S, 2008, "Moving Window Approaches for Hedonic Price Estimation: An Empirical Comparison of Modeling Techniques" *Urban Studies* **45** (8) 1565 - 1581.

Páez A, Uchida T, Miyamoto K, 2002a, "A general framework for estimation and inference of geographically weighted regression models: 1. Location-specific kernel bandwidths and a test for locational heterogeneity" *Environment and Planning A* **34** (4) 733 - 754.

Páez A, Uchida T, Miyamoto K, 2002b, "A general framework for estimation and inference of geographically weighted regression models: 2. Spatial association and model specification tests" *Environment and Planning A* **34** (5) 883 - 904.

Pelto CR, Elkins TA, Boyd HA, 1968, "Automatic Contouring of Irregularly Spaced Data" *Geophysics* **33** (3) 323 - 343.

Propastin PA, Kappas M, 2008, "Reducing uncertainty in modeling the NDVI-precipitation relationship: A comparative study using global and local regression techniques" *Giscience & Remote Sensing* **45** (1) 47 - 67.

Tu J, Xia ZG, 2008, "Examining spatially varying relationships between land use and water quality using geographically weighted regression I: Model design and evaluation" *Science of the Total Environment* **407** (1) 358 - 378.

Wang N, Mei CL, Yan XD, 2008, "Local linear estimation of spatially varying coefficient models: an improvement on the geographically weighted regression technique" *Environment and Planning A* **40** (4) 986 - 1005.

Wheeler D, Tiefelsdorf M, 2005, "Multicollinearity and correlation among local regression coefficients in geographically weighted regression" *Journal of Geographical Systems* **7** (2) 161 - 187.

Yu DL, 2006, "Spatially varying development mechanisms in the Greater Beijing Area: a geographically weighted regression investigation" *Annals of Regional Science* **40** (1) 173 - 190.

Zhang LJ, Shi HJ, 2004, "Local modeling of tree growth by geographically weighted regression" *Forest Science* **50** (2) 225 - 244.