

Mapping the Results of Geographically Weighted Regression

Jeremy Mennis

Department of Geography and Urban Studies, Temple University, 1115 West Berks Street, 309 Gladfelter Hall, Philadelphia, PA 19066, USA.

Email: jmennis@temple.edu

Geographically weighted regression (GWR) is a local spatial statistical technique for exploring spatial nonstationarity. Previous approaches to mapping the results of GWR have primarily employed an equal step classification and sequential no-hue colour scheme for choropleth mapping of parameter estimates. This cartographic approach may hinder the exploration of spatial nonstationarity by inadequately illustrating the spatial distribution of the sign, magnitude, and significance of the influence of each explanatory variable on the dependent variable. Approaches for improving mapping of the results of GWR are illustrated using a case study analysis of population density–median home value relationships in Philadelphia, Pennsylvania, USA. These approaches employ data classification schemes informed by the (nonspatial) data distribution, diverging colour schemes, and bivariate choropleth mapping.

INTRODUCTION

Local forms of spatial analysis have recently gained in prominence. For example, local adaptations have been developed for conventional summary statistics (Brunsdon *et al.*, 2002) as well as for the analysis of spatial dependency in both quantitative (Anselin, 1995; Ord and Getis, 1995) and categorical data (Boots, 2003). Because local spatial statistics often generate georeferenced data, maps and other graphics are typically used to present, and aid in the interpretation of, local spatial statistical results. And because these local statistics are generally exploratory, as opposed to confirmatory, in nature, they have much in common theoretically with recent research in cartography focusing on the use of maps and statistical graphics for data exploration (e.g. MacEachren and Ganter, 1990; Andrienko *et al.*, 2001; Carr *et al.*, 2005). Few cartographers, however, have explicitly addressed the adaptation of conventional mapping techniques for local spatial statistics.

Geographically weighted regression (GWR) is a local spatial statistical technique used to analyze spatial nonstationarity, defined as when the measurement of relationships among variables differs from location to location (Fotheringham *et al.*, 2002). Unlike conventional regression, which produces a single regression equation to summarize global relationships among the explanatory and dependent variables, GWR generates spatial data that express the spatial variation in the relationships among variables. Maps generated from these data play a key role in exploring and interpreting spatial nonstationarity.

A number of recent publications have demonstrated the analytical utility of GWR for investigating a variety of topical areas, including climatology (Brunsdon *et al.*, 2001), urban poverty (Longley and Tobon, 2004), environmental justice (Mennis and Jordan, 2005), and the ecological inference problem (Calvo and Escobar, 2003). However, a standard approach for mapping the results of GWR has not yet been developed. This may be due to the relatively recent development of the technique itself, but is also likely a result of the complications in displaying the results of GWR. Note that each GWR analysis can produce a voluminous amount of spatial data, including multiple georeferenced variables. Some of these variables can be considered ratio data while other variables can be interpreted as nominal. Numeric variables may be highly skewed and range over positive and negative values.

The purpose of this research is to review previous approaches to mapping the results of GWR and suggest methods to improve upon them. I focus on GWR as applied to the analysis of areal data, as opposed to data taken as samples of a continuous surface, as the vast majority of GWR research has been applied to socio-economic data aggregated to census or other spatial units. As a case study, a number of mapping approaches are used to interpret the results of a GWR analysis of median home value in Philadelphia, Pennsylvania, USA using 2000 US Bureau of the Census tract level data.

GEOGRAPHICALLY WEIGHTED REGRESSION

Because readers may not be familiar with the details of GWR, a brief explanation of it is offered here. The conventional regression equation can be expressed as

$$\hat{y}_i = \beta_0 + \sum_k \beta_k x_{ik} + \varepsilon_i \quad (1)$$

where \hat{y}_i is the estimated value of the dependent variable for observation i , β_0 is the intercept, β_k is the parameter estimate for variable k , x_{ik} is the value of the k^{th} variable for i , and ε_i is the error term. Instead of calibrating a single regression equation, GWR generates a separate regression equation for each observation. Each equation is calibrated using a different weighting of the observations contained in the data set. Each GWR equation may be expressed as

$$\hat{y}_i = \beta_0(u_i, v_i) + \sum_k \beta_k(u_i, v_i) x_{ik} + \varepsilon_i \quad (2)$$

where (u_i, v_i) captures the coordinate location of i (Fotheringham *et al.*, 1998). The assumption is that observations nearby one another have a greater influence on one another's parameter estimates than observations farther apart. The weight assigned to each observation is based on a distance decay function centred on observation i . In the case of areal data, the distance between observations is calculated as the distance between polygon centroids.

The distance decay function, which may take a variety of forms, is modified by a bandwidth setting at which distance the weight rapidly approaches zero. The bandwidth may be manually chosen by the analyst or optimized using an algorithm that seeks to minimize a cross-validation score, given as

$$CV = \sum_{i=1}^n (y_i - \hat{y}_{i \neq i})^2 \quad (3)$$

where n is the number of observations, and observation i is omitted from the calculation so that in areas of sparse observations the model is not calibrated solely on i . Alternatively, the bandwidth may be chosen by minimizing the Akaike Information Criteria (AIC) score, give as

$$AIC_c = 2n \log_e(\hat{\sigma}) + n \log_e(2\pi) + n \left\{ \frac{n + \text{tr}(\mathbf{S})}{n - 2 - \text{tr}(\mathbf{S})} \right\} \quad (4)$$

where $\text{tr}(\mathbf{S})$ is the trace of the hat matrix. The AIC method has the advantage of taking into account the fact that the degrees of freedom may vary among models centred on different observations. In addition, the user may choose a fixed bandwidth that is used for every observation or a variable bandwidth that expands in areas of sparse observations and shrinks in areas of dense observations (Charlton *et al.*, no date).

Because the regression equation is calibrated independently for each observation, a separate parameter estimate, t-value, and goodness-of-fit is calculated for each observation. These values can thus be mapped, allowing the analyst to visually interpret the spatial distribution of the nature and strength of the relationships among explanatory and

dependent variables. For more information on the theory and practical application of GWR the reader is referred to (Fotheringham *et al.*, 2002)

CHALLENGES TO MAPPING THE RESULTS OF GWR

A survey of research incorporating GWR reveals that maps play a central role in interpreting GWR results. However, there are a number of issues that have led these maps to obscure the GWR results as much as illuminate them. One issue is that the spatial distribution of the parameter estimates must be presented in concert with the distribution of significance, as indicated by a t-value, in order to yield meaningful interpretation of the results. Some researchers have chosen to map only the parameter estimates and not associated t-values (Fotheringham *et al.*, 1998; Huang and Leung, 2002; Lee, 2004), which can be very misleading as it may visually emphasize the areas of highest (or lowest, if the relationship is primarily negative) parameter estimation, regardless of the significance of the estimate. Thus, one may get the impression that the areas with the highest parameter estimates exhibit the strongest relationship between the explanatory and dependent variables, when those estimates may not, in fact, be significant. Clearly, maps of the spatial distribution of the parameter estimates must be accompanied by associated t-value data if spatial nonstationarity is to be interpreted effectively by the map reader.

A second issue concerns data classification. The equal step approach, where the data range is divided into classes of equal extent (Dent, 1999), appears to be the most common data classification technique for mapping the distribution of parameter estimates and t-values generated from GWR (e.g. Longley and Tobon, 2004). It should be noted, however, except in cases where exogenous classification criteria are used, the choice of data classification scheme for quantitative data is typically informed by the non-spatial data distribution (Evans, 1977; Dent, 1999). The equal step classification is most appropriate for uniformly distributed data, which in the case of GWR-generated parameter estimates would occur when the frequencies of the estimates were approximately the same over the range of the estimates. While possible, this is certainly unlikely. Other classification schemes are likely to be more appropriate, such as the use of standard deviation classification for normally distributed data, or the use of optimal methods for maximizing within-class homogeneity (e.g. Coulson, 1987; Cromley, 1996).

In addition, the data classification for t-values should account for certain exogenous criteria that are of importance to the variable being mapped (Evans, 1977), namely the threshold values that distinguish parameter estimates that are significant from those that are not. When a class interval extends across a significance threshold to encompass both significant and not significant t-values within one class, as it may be using an equal step classification scheme, it becomes impossible to visually distinguish significant parameter estimates from those that are not significant on the map.

A third issue is the choice of colour scheme. Many GWR researchers have employed a sequential no-hue colour



Figure 1. Important neighbourhoods of Philadelphia, Pennsylvania in the context of the case study, overlain with tract boundaries

scheme, which assigns a series of class intervals increasing shades of grey (Brewer, 1994) for choropleth mapping of both parameter estimates and t-values (Fotheringham *et al.*, 1998; Longley and Tobon, 2004; Lee, 2004). Such a colour scheme gives the impression of a gradation of increasing influence (i.e. from a lighter to darker shade of grey) of the explanatory variable on the dependent variable. In cases where the parameter estimates are all of the same sign, the sequential approach may be appropriate. However, this colour scheme is problematic in cases where the parameter estimate is positive in some locations and negative in others (which is not an unusual occurrence, e.g. Huang and Leung, 2002; Lee, 2004; Mennis and Jordan, 2005), as it ignores the fact that the sign of the parameter estimate indicates an importance difference in the nature of the relationship of the explanatory with the dependent variable. In this case, a diverging colour scheme (Brewer, 1994; 1996), which indicates the magnitude of departure from a midpoint value (i.e. zero in the case of distinguishing positive from negative relationships), is most appropriate.

A fourth issue is the sheer number of individual maps required to report both the parameter estimates and t-values for each explanatory variable. This is problematic in terms of cost of map production (e.g. physical space in a journal publication) and the cognitive effort in map comprehension required from the map reader.

Table 1. Descriptive statistics

Variable	Minimum	Maximum	Mean	Standard deviation
Home value (US dollars)	9 999	843 800	75 860	70 362
Population density (people km ⁻²)	120	21 168	6 618	3 853

Table 2. Conventional regression of home value

Independent variable	Coefficient	t-value
Constant	-106 524.30***	-14.87
Population density	-4.63***	-4.96

*** Significance <0.005, $N = 357$, Adjusted $R^2 = 0.062$.

Choropleth mapping has been extended to two variables simultaneously, as in a bivariate choropleth map (Olson, 1975). Combining parameter estimates and t-values in a single choropleth map would reduce the volume of maps necessary for exploring the results of GWR.

CASE STUDY: GWR OF HOME VALUE IN PHILADELPHIA, PA

The case study concerns the GWR of median owner-occupied home value (US dollars) in Philadelphia, Pennsylvania, USA using population density (people km⁻²) as the explanatory variable. These 2000 data were acquired from the US Bureau of the Census at the tract level. Note that the purpose of the case study is not to demonstrate anything novel about home values in Philadelphia per se, but rather to show and compare different strategies for mapping the results of GWR. The focus is on maps of parameter estimates and t-values as these are the most commonly reported maps in research using GWR. The use of only one explanatory variable in the case study keeps the volume of GWR results to a manageable level while generating interesting patterns of spatial nonstationarity that can be used to illustrate the benefits and pitfalls of various mapping strategies. Of the 381 tracts in Philadelphia, 24 were removed from the analysis because they represented very sparsely populated or unpopulated areas (i.e. parks, airports, and industrial land uses), leaving 357 tracts for use in the analysis. A map of Philadelphia neighbourhoods relevant to the case study is presented in Figure 1. Descriptive statistics and choropleth maps of the variables used in the analysis are presented in Table 1 and Figure 2, respectively.

The results of a conventional linear regression of home value are reported in Table 2. The model indicates that population density is negatively and significantly related to home value; as home values increase, population density decreases. Note, however, that the model is poorly specified, explaining only approximately 6% of the variation in home value. Reasons for this poor specification will be made clear in the GWR.

The data were entered into the GWR software using a variable bandwidth setting that minimizes the AIC. The variable bandwidth approach was chosen to account for the spatial variation in the size of the tracts, and hence the density of tract centroids. As noted above, the most

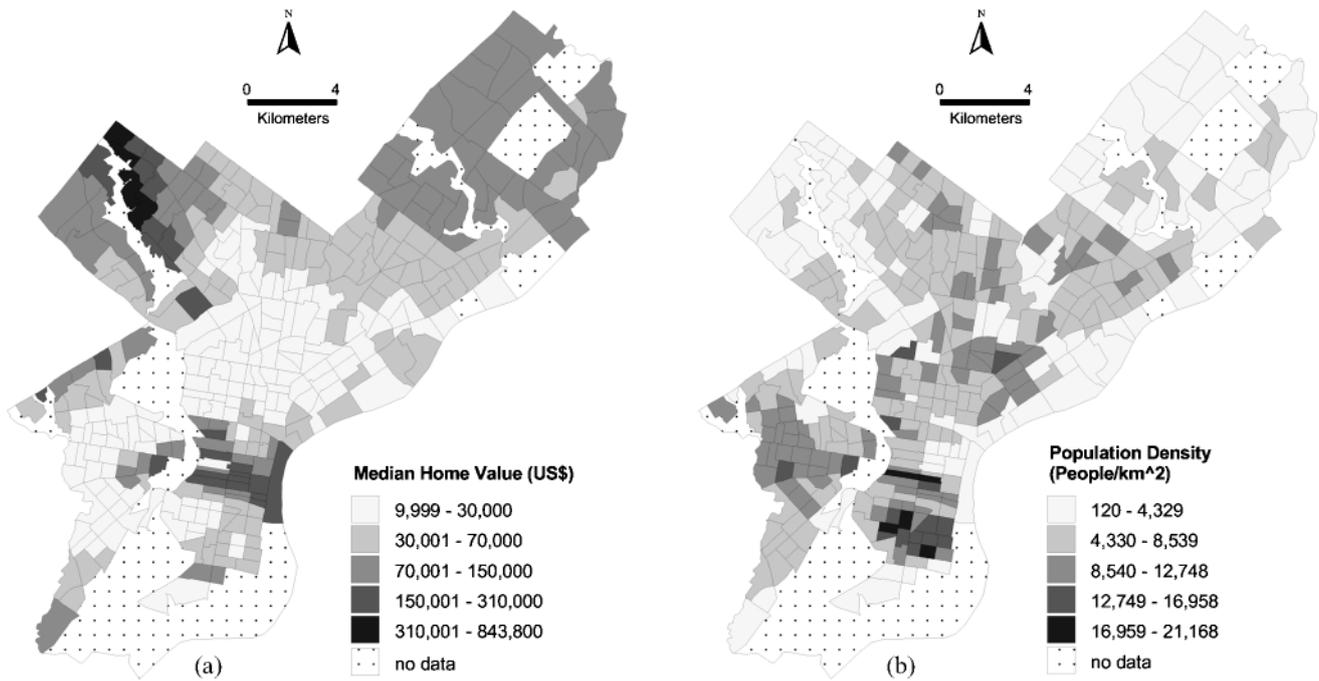


Figure 2. Choropleth maps of *a* median home value and *b* population density by census tract in Philadelphia, PA

common approach to presenting the results of GWR is to generate choropleth maps of the parameter estimates using a sequential no-hue colour scheme and an equal-step classification. Figure 3a presents such a map of the population density parameter estimate. One can immediately see that this map is problematic, as the imposition of this colour scheme and classification ignore relevant variations in the data that should be brought to the attention of the viewer. First, the sequential colour scheme

suggests that the influence of population density on home value increases monotonically. In fact, in some tracts this relationship is negative and in others it is positive. Perhaps even more troubling is that the majority of the mapped area is occupied by a single class that includes both positive and negative parameter estimates (i.e. the class interval -7 to 12). Thus, it is impossible to tell within which areas the population density–home value relationship is positive versus negative. Finally, because no information on the

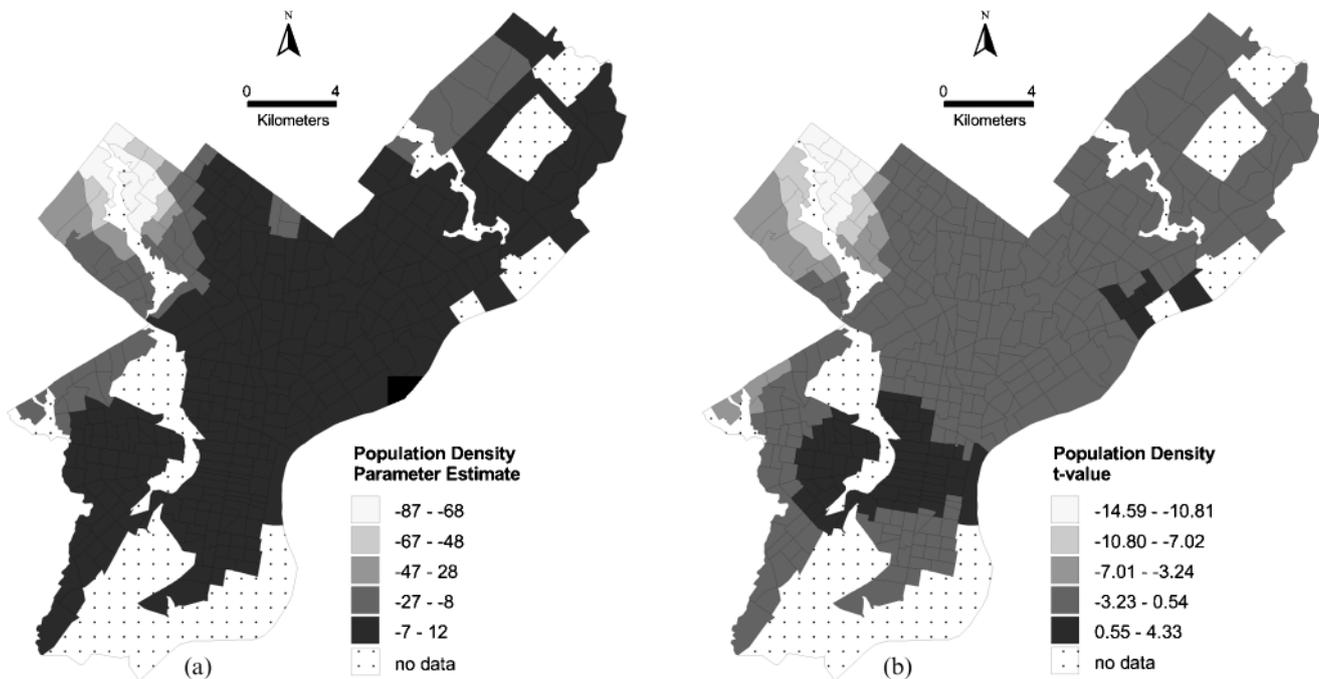


Figure 3. Choropleth maps of *a* parameter estimates and *b* t-values by census tract for the GWR of median home value using an equal step data classification and a sequential no-hue colour scheme for each map

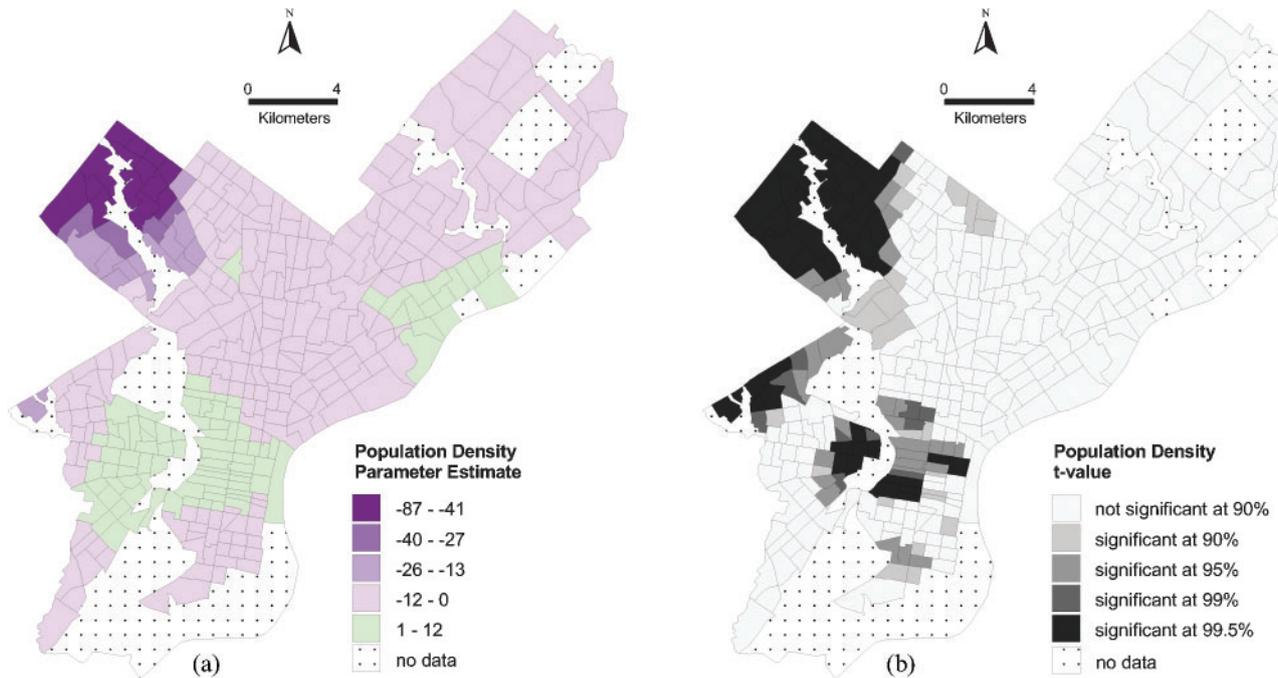


Figure 4. Choropleth maps of a parameter estimates and b t -values by census tract for the GWR of median home value. In the parameter estimate map, a modified standard deviation data classification and a diverging colour scheme is used whereas in the t -value map, an exogenous data classification based on commonly accepted significance thresholds and a sequential no-hue colour scheme is used

distribution of t -values is provided, one cannot detect the areas in which the relationship between explanatory and dependent variables is significant. This last problem can be amended simply by creating a map of t -values (Figure 3a), presented here also using the conventional sequential no-hue colour scheme and equal step classification, though similar problems regarding classification and choice of colour scheme apply.

Figure 4a presents a map that addresses the classification and colour scheme problems present in the choropleth map of parameter estimates presented in Figure 3a. In Figure 4a, the classification is based generally on a standard deviation classification scheme, as the data approach a normal distribution. In addition, manual adjustments to the statistically-derived data classification scheme are made to facilitate map interpretation (Monmonier, 1982). The class breaks were shifted to distinguish positive from negative parameter estimates, and, because the range of negative parameter estimates is greater than the range of positive parameter estimates, the interval boundaries were set to allow the direct comparison of positive and negative parameter estimates of equivalent magnitude. Thus, of five classes, only one contains all the tracts with positive parameter estimates. A diverging colour scheme was also employed to differentiate negative from positive parameter estimates by hue, while expressing increasing magnitudes of the estimates using a combination of saturation and value. Unlike Figure 3a, Figure 4a clearly shows that the areas of positive relationship between population density and home value are largely limited to the greater Center City and University City neighbourhoods, as well as nearby Frankford. A negative population density–home value relationship of equal magnitude is evident in the remainder of the city, with the exception of the Roxborough and

Chestnut Hill neighbourhoods, within which stronger negative relationships occur.

Figure 4b presents a map that addresses the classification and colour scheme problems present in Figure 3b. Figure 4b has a classification scheme based on commonly used significance thresholds: 90, 95, 99, and 99.5%. A sequential colour scheme is used to represent different levels of significance. Unlike in Figure 3b, Figure 4b clearly indicates that in the majority of Philadelphia the relationship between population density and home value is, in fact, not significant at the 90% confidence level. It is significant primarily in University City, western Center City, Girard Estates, and a number of neighbourhoods in the north-western part of the city. Clearly, this significance information is key to interpreting Figure 4a, as Figure 4a appears to suggest an equivalency between Center City and Frankford in the relationship of population density with home value. Figure 4b, however, clearly shows that in Frankford the relationship between the two variables is not significant at the 90% confidence level and, within those areas where the relationship between the variables is significant, the magnitude of the significance varies. Some parts of those areas show a significant relationship at the 99.5% confidence level (e.g. Chestnut Hill and Roxborough), while others only meet the 90% confidence level threshold (e.g. East Falls and West Oak Lane).

The maps presented in Figure 4 are a marked improvement over those presented in Figure 3, as they allow for a much more accurate assessment of which areas have positive and negative relationships of the explanatory variable with the dependent variable, the magnitude of those relationships, and the significance of those relationships. However, given a regression with many explanatory variables, as opposed to just the one used in this case study, many maps

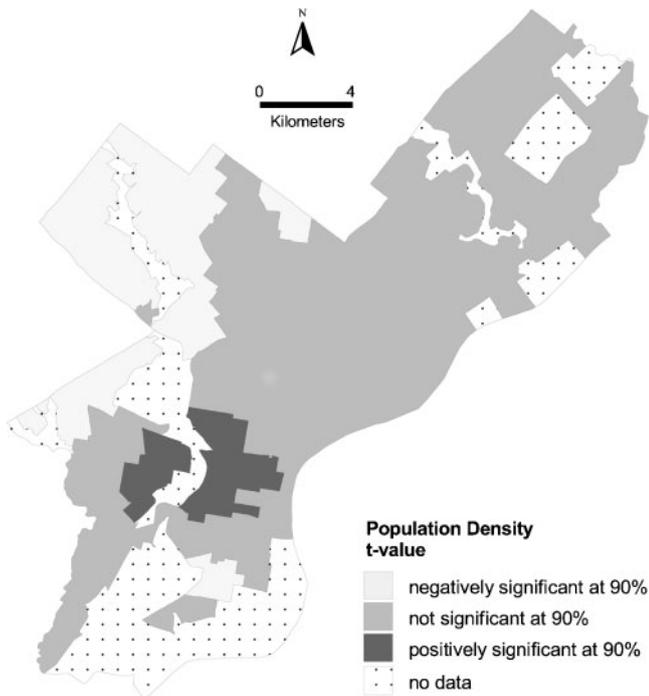


Figure 5. An area-class map of positively and negatively significant and not significant t-values, for the GWR of median home value

are required to communicate this information, as each explanatory variable demands two separate maps – one for the parameter estimate and one for the t-value. Figure 5 offers a potential solution to this problem by encoding certain key characteristics of Figures 4a and 4b in a single area-class map. Here, tracts are classified according to their

relationship between the explanatory and dependent variable, characterized as positively significant, negatively significant, and not significant (at the 90% confidence level). These classes are treated as nominal data and assigned varying lightness levels of grey in the map in a qualitative colour scheme that is intended to differentiate among classes without implying rank or quantity (Brewer, 1994). Note that the linework of the tract boundaries has been removed to reduce the visual complexity of the map. The advantage of this mapping approach is that one can easily see qualitative differences among areas in the sign of the relationship between the explanatory and dependent variable, as well as distinguish between areas exhibiting a significant versus not significant relationship. Another advantage is that a grey-scale, as opposed to colour, map may be used. Of course, the disadvantage of this mapping approach is that potentially interesting patterns may not be observed regarding the magnitude of the relationship between the explanatory and dependent variable as contained in the actual parameter estimate values, as well as in the magnitude of the significance.

Bringing colour back into the map allows for a compromise between Figures 4a and 5 as contained in a single map, presented in Figure 6a. Here, a map showing the parameter estimates in a manner similar to that of 3a is used, except that a significance threshold (at 90% confidence level) is used to mask out all those areas in which the relationship between the explanatory and dependent variables is not significant. Here, it is implied that distinguishing between positive and negative parameter estimates (and associated t-values) in these areas is unnecessary. These areas are given a neutral grey tone and their linework for the tract boundaries is removed, the

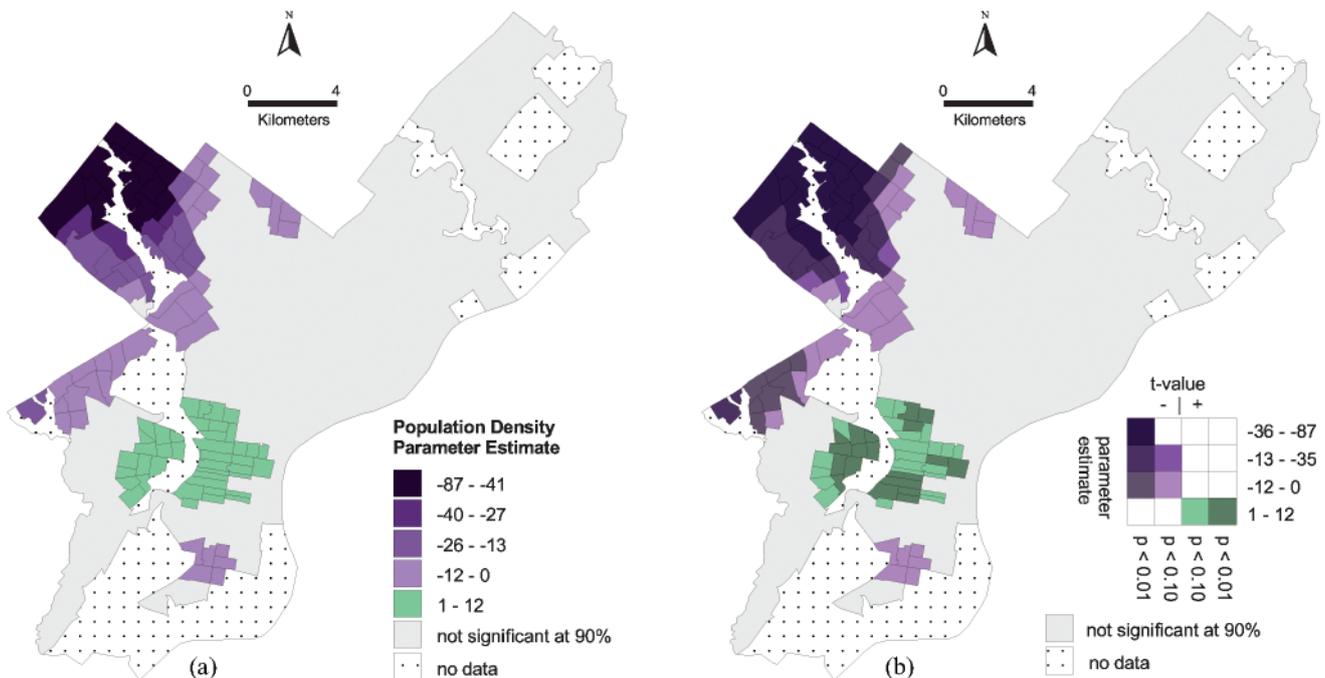


Figure 6. Choropleth maps simultaneously displaying both the magnitude and significance of the parameter estimate by census tract: *a* a mask is applied to those tracts with a t-value with a significance less than 90%; *b* both the parameter estimate and associated significance are incorporated in a bivariate data classification and colour scheme

assumption being that these areas are of less interest to an analyst than those areas that are significant.

Figure 6a can also be modified by using a bivariate colour scheme to simultaneously depict both the magnitude of the parameter estimate and the magnitude of the significance. In Figure 6b, a 4×4 class colour matrix is used to depict various combinations of parameter estimate and significance. A diverging colour scheme using two different hues is used to map the parameter estimate values, as in Figure 6a, because they range from positive to negative values. A sequential scheme using saturation is used to map significance, where increased saturation indicates higher significance, because the sign of the relationship is already captured by the hue in the vertical axis of the matrix. Thus, the map may be considered to use a diverging-sequential, bivariate colour scheme.

Because colours are only assigned to tracts with a significant relationship between the explanatory and dependent variables (at greater than or equal to 90% confidence), the matrix's class intervals are not continuous along the horizontal axis. All tracts that do not exhibit a significant relationship between population density and home value (i.e. fall within the vertical class partition in the centre of the matrix) are assigned a neutral grey colour. Note also that the matrix is sparsely populated (i.e. there are a number of 'empty' cells) because the t-value and parameter estimate always share the same sign.

DISCUSSION AND CONCLUSION

Although the purpose of the case study concerns cartographic methodology and not the substantive topic of home values in Philadelphia, it is worth taking a moment to discuss the substantive results as a means to evaluate the various mapping approaches. First, the reason that the conventional regression was not specified properly is explained, at least in part, by the spatial nonstationarity indicated by the GWR. Clearly, a linear regression model that is global in nature will not be able to accurately characterize the relationship between explanatory and dependent variables when the relationship is positive in some portions of the study region and negative in others, as Figure 4a indicates. The negative relationship between population density and home value is perhaps one that could be expected; expensive homes are likely to occur in sparsely populated areas where single-family homes sit on large lots. This is indeed the case in certain Philadelphia neighbourhoods at the urban periphery, such as Roxborough, Chestnut Hill, and Overbrook, as Figures 4, 5, and 6 show.

The positive relationship between population density and home value exhibited in University City and western Center City is probably related to their historic roots as centres of wealth, high-end commercial activity, and higher education within the city core. Both neighbourhoods have maintained densely populated residential areas even as many nearby working-class neighbourhoods in North, South, and West Philadelphia have lost population in recent years. Population decline is associated with housing abandonment and marginal home appreciation (or even decline), thus

creating the local positive relationship between population density and home value for University City and western Center City that can now be observed in Figures 4, 5, and 6.

This research demonstrates that the conventional approach of using an equal step classification and sequential no-hue colour scheme for choropleth mapping of GWR-generated parameter estimates is clearly inadequate. As Figure 3a shows, such a map is not only uninformative but can be downright misleading, even when paired with another map of t-values as an indicator of significance. Adjustments to the data classification and colour scheme to improve the cartographic representation of the sign, magnitude, and significance of parameter estimates, as in Figure 4, offer an improvement in interpreting the GWR results, but two maps are required for the representation of each explanatory variable.

The advantage of Figure 5 is that, because it is an area-class map with only three classes, it appears relatively uncluttered and is therefore easy to visually interpret. Yet it effectively communicates the basic pattern of spatial nonstationarity as captured by the GWR. On the downside, however, it does not show the spatial distribution of the magnitude of the parameter estimates. The maps contained in Figure 6 are unique in that they convey spatial information on both the magnitude and significance of the parameter estimates in a single map. Because Figure 6a employs a simple significance threshold, whereas Figure 6b maps the distribution of significance, Figure 6b contains more information. For example, Figure 6b clearly shows that some tracts in western Center City have a much higher significance than others, a pattern that cannot be observed in Figure 6a. And one can see that in Overbrook population density has a highly significant, negative relationship with home value, though the influence of the explanatory variable on the dependent variable is relatively marginal compared with its influence in other areas, such as Chestnut Hill.

However, the bivariate colour scheme used in Figure 6b can be difficult to visually interpret, particularly given the fact that additional colour assignments are needed for representing observations which are classified as not significant or which have no data. And while knowing the spatial distribution of significance values is certainly important, significance is typically treated as a threshold. For these reasons, I advocate the mapping approach taken in Figure 6a as a good rule-of-thumb for mapping the results of GWR. Or, an analyst may choose to use a map like that presented in Figure 5, if this reduced level of information communication is deemed sufficient.

It is worth noting that while the case study focuses on mapping the parameter estimate and t-value for GWR using a single explanatory variable, most GWR applications will have multiple explanatory variables. In such a situation, GWR may be used to interpret maps of parameter estimates and/or t-values to determine within which region(s) specific explanatory variables are particularly influential. Such an analysis demands a comparison of choropleth maps in a series, for which design criteria may differ from that used for a single map (Brewer and Pickle, 2002) Mennis and Jordan (2005) facilitate such a comparison by using

area-class maps like that presented in Figure 5, thus supporting map comparison by standardizing maps according to a significance threshold applied uniformly to all explanatory variables. However, if choropleth mapping of parameter estimates is used to indicate the magnitude of influence of each explanatory variable, each parameter estimate must be standardized before being mapped (i.e. the standardized β). Likewise, standardization of the data classification and colour scheme across all maps in the series will facilitate map comparison, even if some maps contain data for only a subset of the classification range (Brewer and Pickle, 2002). It is also worth noting that not all parameter estimates and attached significance values necessarily need to be mapped in order to generate an effective visualization of the overall quality and most relevant characteristics of a GWR model.

A software package devoted to automated mapping of GWR results would be a useful tool for assisting researchers in developing informative and useful maps for exploring spatial nonstationarity. Such a software package could ingest the output from GWR analysis and offer automated intelligent rules for cartographic display, based on the data classification, colour scheme, and bivariate mapping approaches described above. In addition, a software package whose purpose is to support the exploration of the results of GWR ought to include characteristics that have been developed for exploratory data analysis in other cartographic contexts, such as the use of small multiples for the visualization of many variables (Pickle *et al.*, 1996), dynamically linked maps and other graphical displays (MacEachren *et al.*, 1999), and modes of interactivity (Crampton, 2002). For example, consider the significance threshold of 90% confidence used in Figure 6a to mask out tracts in which the relationship between population density and home value is considered not significant. A slider bar or other interactive device could facilitate the exploration of the effect of changing the threshold significance value on the interpretation of spatial nonstationarity. Interactive devices for dynamically altering class breaks for parameter estimates and/or significance values would be useful in exploring the maps presented Figures 4 and 6, as well as in transforming the t-values to nominal data in Figure 5.

It would be useful to provide choropleth maps of the explanatory and dependent variables, linked to the choropleth maps of the analogous parameter estimates and t-values so that panning, zooming, selection and other interactions in one map would be effective in all maps. In addition, dynamically linking statistical graphics, such as scatter plots and parallel coordinate plots (e.g. Gahegan *et al.*, 2002), to the maps of parameter estimates and significance would facilitate the exploration of the multivariate 'signatures' associated with regions of homogeneity regarding the relationship between explanatory and dependent variables.

ACKNOWLEDGEMENTS

The choice of colour schemes used in this research were informed by ColorBrewer, an online mapping tool for choosing colour schemes for choropleth maps (Harrower

and Brewer, 2003) and *Mapping Census 2000: The Geography of US Diversity* (Brewer and Suchan, 2001).

REFERENCES

- Andrienko, N., Andrienko, G., Savinov, A., Voss, H., and Wettschereck, D. (2001). 'Exploratory analysis of spatial data using interactive maps and data mining', *Cartography and Geographic Information Science*, 28, 151–165.
- Anselin, L. (1995). 'Local indicators of spatial association – LISA', *Geographical Analysis*, 27, 93–115.
- Boots, B. (2003). 'Developing local measures of spatial association for categorical data', *Journal of Geographical Systems*, 5, 139–160.
- Brewer, C. (1994). 'Color use guidelines for mapping and visualization', in *Visualization in Modern Cartography*, ed. by MacEachren, A. and Taylor, D.R.F., p. 123–147, Elsevier, New York.
- Brewer, C. A. (1996). 'Guidelines for selecting colors for diverging schemes on maps', *The Cartographic Journal*, 33, 79–86.
- Brewer, C. A. and Pickle, L. (2002). 'Evaluation of methods for classifying epidemiological data on choropleth maps in a series', *Annals of the Association of American Geographers*, 92, 662–681.
- Brewer, C. A. and Suchan, T. A. (2001). *Mapping Census 2000: The Geography of US Diversity*. US Census Bureau Special Report, Series CENSR/01-1. US Government Printing Office. Washington DC.
- Brunsdon, C., Fotheringham, A. S. and Charlton, M. E. (2002). 'Geographically weighted summary statistics: a framework for localized exploratory data analysis', *Computers, Environment and Urban Systems*, 501–524.
- Brunsdon, C., McClatchey, J. and Unwin, D. (2001). 'Spatial variations in the average rainfall–altitude relationships in Great Britain: an approach using geographically weighted regression', *International Journal of Climatology*, 21, 455–466.
- Calvo, C. and Escolar, M. (2003). 'The local voter: a geographically weighted approach to ecological inference', *American Journal of Political Science*, 47, 189–204.
- Carr, D. B., White, D., and MacEachren, A. M. (2005). 'Conditioned choropleth maps and hypothesis generation', *Annals of the Association of American Geographers*, 95, 32–53.
- Charlton, M., Fotheringham, S. and Brunsdon, C. (no date). Geographically Weighted Regression Version 2.x, User's Manual and Installation Guide.
- Coulson, M. R. C. (1987). 'In the matter of class intervals for choropleth maps: with particular reference to the work of George Jenks', *Cartographica*, 24, 16–39.
- Crampton, J.W. (2002). 'Interactivity types in geographic visualization', *Cartography and Geographic Information Science*, 29, 85–98.
- Cromley, R. G. (1996). 'A comparison of optimal classification strategies for choropleth displays of spatially aggregated data', *International Journal of Geographical Information Science*, 10, 405–424.
- Dent, B. D. (1999). *Cartography: Thematic Map Design, Fifth Edition*, WCB/McGraw Hill, Boston.
- Evans, I. A. (1977). 'Selection of class intervals', *Transactions of the Institute of British Geographers, New Series*, 2, 98–124.
- Fotheringham, A. S., Brunsdon, C. and Charlton, M. E. (1998). 'Geographically weighted regression: a natural evolution of the expansion method for spatial data analysis', *Environment and Planning A*, 30, 1905–1927.
- Fotheringham, A. S., Brunsdon, C., and Charlton, M. E. (2002). *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships*, Wiley, Chichester.
- Gahegan, M., Takatsuka, M., Wheeler, M. and Hardisty, F. (2002). 'Introducing GeoVISTA Studio: an integrated suite of visualization and computational methods for exploration and knowledge construction in geography', *Computers, Environment and Urban Systems*, 26, 267–292.

- Harrower, M. A. and Brewer, C. A. (2003). 'ColorBrewer.org: an online tool for selecting colour schemes for maps', **The Cartographic Journal**, 40, 27–37.
- Huang, Y. and Leung, Y. (2002). 'Analyzing regional industrialization in Jiangsu province using geographically weighted regression'. **Journal of Geographical Systems**, 4, 233–249.
- Lee, S.-I. (2004). 'Spatial data analysis for the US regional income convergence, 1969–1999: a critical appraisal of b-convergence', **Journal of the Korean Geographical Society**, 39.
- Longley, P. A. and Tobon, C. (2004). 'Spatial dependence and heterogeneity in patterns of hardship: an intra-urban analysis', **Annals of the Association of American Geographers**, 94, 503–519.
- MacEachren, A. M. and Ganter, J. H. (1990). 'A pattern identification approach to cartographic visualization', **Cartographica**, 27, 64–81.
- MacEachren, A. M., Wachowicz, M., Edsall, R., Haug, D., and Masters, R. (1999). 'Constructing knowledge from multivariate spatiotemporal data: integrating geographical visualization with knowledge discovery in database methods', **International Journal of Geographical Information Science**, 13, 311–334.
- Mennis, J. and Jordan, L. (2005). 'The distribution of environmental equity: exploring spatial nonstationarity in multivariate models of air toxic releases', **Annals of the Association of American Geographers**, 95, 249–268.
- Monmonier, M.S. (1982). 'Flat laxity, optimization, and rounding in the selection of class intervals', **Cartographica**, 19, 16–26.
- Olson, J. (1975). 'Spectrally encoded two-variable maps', **Annals of the Association of American Geographers**, 71, 259–276.
- Ord, J. K. and Getis, A. (1995). 'Local spatial autocorrelation statistics: distributional issues and an application'. **Geographical Analysis**, 27, 286–306.
- Pickle, L. W., Mingle, M., Jones, G. K., and White, A. A. (1996). **Atlas of United States Mortality**, US National Center for Health Statistics, Hyattsville, Maryland, USA.