Journal of Transport Geography 31 (2013) 123–131

Contents lists available at SciVerse ScienceDirect

Journal of Transport Geography

journal homepage: www.elsevier.com/locate/jtrangeo

Understanding spatial filtering for analysis of land use-transport data

Yiyi Wang [a], Kara M. Kockelman [b,*], Xiaokun (Cara) Wang [c]

[a] College of Engineering, Montana State University, P.O. Box 173820, Bozeman, MT 59717-3820, USA
[b] Department of Civil, Architectural and Environmental Engineering, The University of Texas at Austin, Austin, TX 78712-1076, USA
[c] Department of Civil and Environmental Engineering, 4031 JEC Building, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180-3590, USA

ARTICLE INFO

Keywords:
Spatial filtering
Land use change
Land valuation
Spatial probit
Spatial autoregressive regression

ABSTRACT

This paper summarizes the literature on spatial filtering (SF) for analysis of spatial data. Given the scarcity of its application in transportation and its fledgling nature, preliminary case studies were conducted using continuous and discrete response data sets, for land values and land use, in comparison with results from spatial autoregressive (SAR) models with distance decay parameters estimated using Bayesian techniques. For both the continuous land value and binary land use cases, the SF approach demonstrates great potential as a worthy competitor to more conventional SAR-based models. In addition to offering high fit statistics, somewhat shorter computing times, and more straightforward computations, the SF approach makes explicit the patterns of spatial dependency in the land value and land use data. By controlling for these spatial relationships, the SF approach yields more reliable marginal effects of policy variables of interest. Model results confirm the important role of transportation access (as quantified using distances to a region's central business district, and various roadway types).

1. Introduction

Spatial relationships typically exist across people and locations in transportation, land use, and demographic data sets. They can be summarized into two types of spatial effects: spatial heterogeneity and spatial autocorrelation (Anselin, 1988). Geographically weighted regression (GWR) is commonly used to characterize spatial heterogeneity, by estimating parameters for each site or observational unit based on all observations within a neighborhood (Fotheringham, 2003; Páez, 2006). A common treatment for spatial autocorrelation is to specify a spatial structure directly, such as a spatial autoregressive (SAR) or spatial error (SEM) model (Anselin, 1988; Anselin and Hudak, 1992; LeSage and Pace, 2009; Ibéas et al., 2012). Work on discrete states of land-use change with such specifications can be found in Chakir and Parent's (2009) spatial multinomial probit model (for cross-sectional data), Munroe et al.'s (2002) series of binary probit and random-effects probit models (using panel techniques), and Wang and Kockelman's (2009a, 2009b, 2009c) dynamic spatial ordered probit model with a temporal component.

Most applications to date rely on specific functional forms (such as SAR and SEM) and arbitrarily pre-determined weight structures to anticipate spatial structure in the data. Several issues can limit the use of specific functional forms in addressing spatial autocorrelation. McMillen (2004) noted how functional misspecifications may lead to spatial autocorrelation and advocated the use of non-parametric methods, to avoid *a priori* assumptions of model form. Computing effort is another important factor to consider, as demonstrated in Wang et al.'s (2011) pursuit of an estimable dynamic spatial multinomial probit specification. Essentially, more complicated models require more complex estimation strategies, such as Bayesian sampling from large-size truncated normals (for latent response variables, in the case of multinomial probit, for example); issues of parameter identification, sample size limitations, and a model's functional flexibility can and do emerge. Another serious challenge relates to computing the log-determinant of a SAR specification: $|I_n - \rho W|$, where I_n is an n by n identity matrix, ρ is the degree of spatial autocorrelation and W is the connectivity or weight matrix. This is especially time-consuming when n is large.

Eigenvector-based spatial filtering, as discussed by Griffith (2007) and Dray et al. (2006), is a relatively new technique for analysis of spatial data sets, and it appears to offer much promise. This technique uses orthogonal and uncorrelated map patterns (represented as eigenvectors obtained from a contiguity- or distance-based weight matrix characterizing the data set's spatial structure) to control for such relationships. Another advantage of this filtering approach is the orthogonality of the eigenvectors, facilitating stepwise variable addition to the model specification. This is also known as "forward regression," as applied in Griffith and Peres-Neto (2006) and Moniruzzaman and Páez (2012). A potential drawback lies in computing the "eigenfunctions" (a general term for the

* Corresponding author. Tel.: +1 512 471 0210.
E-mail addresses: yiyi.wang1984@gmail.com (Y. Wang), kkockelm@mail.utexas.edu (K.M. Kockelman), wangx18@rpi.edu (X. Wang).

0966-6923/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.jtrangeo.2013.06.001

eigenvectors and their associated eigenvalues) for large data sets, which can be a formidable task.[1] However, in comparison with the more widely used SAR and SEM approaches to spatial autocorrelation, such computations may be reasonable. This is especially true for transportation-related contexts, given the relative lack of such research to date, and the highly spatial nature of transport applications.

A major advantage of eigenvector-based spatial filtering, compared to other techniques (such as SAR and geographically weighted regression), is that it explicitly identifies the clustering patterns, which is critical information for land-use forecasting and other forms of spatial-data analysis. Other spatial techniques mostly focus on evaluating the magnitude of spatial effects with an underlying assumption that the spatial effect is universal across the sample. In practice, land developers, transport planners, and policy makers are more interested in the source and pattern of the spatial dependency: If the development of a parcel is influenced by its neighbors, which neighbors are the most influential and to what extent? Where are they located? Do other neighbors also depend on this one, and in which way? Eigenvector-based spatial filtering helps answer such questions. The approach decomposes the spatial dependency into several loadings of synthetic covariates (or eigenvectors) and ranks the clustering effects captured by these eigenvectors. Influential neighbors are identified by examining a visual representation of the selected eigenvectors.

The following sections discuss in some detail the development of spatial filtering techniques. Two sets of case studies, analyzing the effects of travel access on (continuous) land values and (binary) land-development data, are provided for comparison to a special SAR model with a flexible weight matrix. Given the advantages of an eigenfunction-based approach among existing spatial filtering techniques, all comparisons are conducted using this technique and the special SAR models (which enable estimation of direct as well as indirect effects of covariates, in contrast to an SEM specification, whose marginal effects are computed like those in an aspatial model). The land-use-change and land-value data used here tie transportation interests through trip generation and attraction rates, access valuation, right-of-way costs, and so forth.

2. A background on spatial filtering

Apart from the observed covariates, also known as the systematic component, spatial filtering techniques rely on weight-matrix eigenvectors, which serve as synthetic explanatory variables representing the data set's spatial structure. These variables add flexibility to the model and have been called the model's non-parametric component (Tiefelsdorf and Griffith 2007). Different methods for generating these variables lead to two main types of spatial filtering in the literature: eigenfunction-based procedures, as discussed in Griffith (2007) and Dray et al. (2006), using a contiguity- or distance-based weight matrix, and Getis' (1990, 1995) G-statistics-based approach.

The main difference between the two approaches is the manner in which the original variables are decomposed. Getis (1990, 1995) used the difference between observed and expected local spatial statistics to separate spatial from non-spatial effects. In an unusual paper, Getis and Griffith (2002) compared their two approaches using government expenditures per capita across US states. Their distinct filtering methods yielded similar goodness-of-fit statistics, although the z-score of the Moran's I test statistic for residuals switched signs: it was weakly negative in Getis' model and weakly positive in Griffith's model. In addition, parameter estimates from both approaches were similar to those of a SAR model. Importantly, the eigenfunction-based approach was deemed preferable, thanks to its flexibility for application in non-linear model specifications. By contrast, Getis' approach requires that analysts have variables with a natural origin and a linear model specification, thereby limiting its use (Patuelli et al., 2011). In the coming discussions, this paper refers only to the eigenfunction-based approach.

The crux of the eigenfunction-based spatial filtering lies in the linkage between eigenfunctions (i.e., the eigenvectors and corresponding eigenvalues) and spatial autocorrelation. Essentially, the (exogenously specified) W matrix's eigenvectors are used as supplemental covariates in the regression, to "filter" out spatial autocorrelation, thereby allowing for potentially more efficient estimation of primary covariates' parameters (Griffith 1996). Eigenfunction decomposition has been widely used in fields like control theory and imaging, but its usage in spatial analysis is relatively new. A thorough interpretation of eigenfunctions from a regional/spatial perspective can be found in Griffith (1996), who uses a 9-by-9 regular square grid and 3 cases of Canada's urban census tracts to provide natural interpretations of the eigenvectors associated with the largest eigenvalues. As Griffith increased his sample size (n), he observed greater agreement between the eigenvalues of the weight matrix W and the eigenvalues of its transformation matrix $\Omega = (I - 11^T/n)W(I - 11^T/n)$, with their correlations ranging from 0.97 for the small n case to approximately 1.0 for the large n case. Essentially, as n increases, the eigenvalues of Ω will show upper-bound convergence to the eigenvalues of W (Griffith, 1996).

Moreover, the transformation matrix (Ω) is guaranteed to have an eigenvector of purely $1/\sqrt{n}$ values, corresponding to an eigenvalue of unity. One advantage of using Ω instead of W is that Ω's eigenvectors are orthogonal (and thus uncorrelated). In other words, Ω's eigenvectors look a bit like the scaled principal components (see, e.g., Jolliffe, 2002) of the matrix W, though their mathematical derivations are quite different.

2.1. Spatial filtering based on Ω

The extreme (most negative and positive) values of Moran's I for a specific spatial configuration (represented by connectivity matrix C or its row-standardized counterpart W) can be expressed as a function of Ω's eigenvalues, as per Tiefelsdorf and Boots (1995) and De Jong et al. (1984)[2]:

$$\text{Moran's I} = \text{MI} = \frac{n}{1'W1} \cdot \text{eigvalue}(\Omega) \tag{1}$$

In other words, one can compute Moran's I for any set of numerical values (y) observed in any spatial data set of size n, and these are the normalized/scaled eigenvalues of Ω. Moreover, the first eigenvector of matrix Ω (denoted as E_1) is the vector of values yielding the strongest spatial autocorrelation (thus having the largest MI value) in the space W. Thus, it is the most important principal component of the spatial structure, as encoded by W. Ω's second eigenvector (E_2) offers the second largest eigenvalue or MI, and is orthogonal to (and thus uncorrelated with) the first eigenvector (Griffith, 2000).

2.2. Selection of eigenvectors

A key issue for filtered regression applications is the strategy to select meaningful eigenvectors to embody compelling spatial interactions. As noted above, one may set a single MI threshold for all eigenvectors' inclusion (Griffith, 2000; Patuelli et al., 2011). In other words, MIs can be computed for each eigenvector, and only eigenvectors with MIs exceeding a target/threshold value

[1] Fortunately, when the data come from a regular, square grid, computing eigenvalues and vectors can be done directly and quickly, in closed form.

[2] Eq. 1 can be simplified to MI = eigvalue (Ω) for row-standardized weight matrices, W, because $\frac{n}{1W1} = 1$.