

Hierarchical and Joint Site-Edge Methods for Medicare Hospice Service Region Boundary Analysis

Haijun Ma, Bradley P. Carlin, and Sudipto Banerjee*

MMC 303, School of Public Health, University of Minnesota, Minneapolis, Minnesota 55455-0392, U.S.A.

*email: sudiptob@biostat.umn.edu

SUMMARY. Hospice service offers a convenient and ethically preferable health-care option for terminally ill patients. However, this option is unavailable to patients in remote areas not served by any hospice system. In this article, we seek to determine the service areas of two particular cancer hospice systems in northeastern Minnesota based only on death counts abstracted from Medicare billing records. The problem is one of spatial boundary analysis, a field that appears statistically underdeveloped for irregular areal (lattice) data, even though most publicly available human health data are of this type. In this article, we suggest a variety of hierarchical models for areal boundary analysis that hierarchically or jointly parameterize *both* the areas and the edge segments. This leads to conceptually appealing solutions for our data that remain computationally feasible. While our approaches parallel similar developments in statistical image restoration using Markov random fields, important differences arise due to the irregular nature of our lattices, the sparseness and high variability of our data, the existence of important covariate information, and most importantly, our desire for full posterior inference on the boundary. Our results successfully delineate service areas for our two Minnesota hospice systems that sometimes conflict with the hospices' self-reported service areas. We also obtain boundaries for the spatial residuals from our fits, separating regions that differ for reasons yet unaccounted for by our model.

KEY WORDS: Areal data; Conditionally autoregressive (CAR) model; Health services research; Ising model; Wombling.

1. Introduction

1.1 Minnesota Medicare Hospice Utilization Data

Access to hospice care is an important issue for more than half a million people who die of cancer and other terminal diseases each year in the United States. Hospices provide palliative care for the terminally ill, with services that typically include supportive medical, social, emotional, and spiritual care that is usually offered in the patient's home. In the United States, Medicare is by far the largest hospice service payer, accounting for over three-fourths of hospice patients and expenditures. The number of Medicare beneficiaries receiving hospice care has grown rapidly, more than doubling from 1992 to 1998 (U.S. GAO, 2000).

Studies of hospice-use rates report significant geographic variations, with rural beneficiaries less likely to use a hospice (U.S. GAO, 2000). This suggests that patients may not have equal access to hospice services, perhaps due to the longer travel distances required of hospice workers in reach rural areas. Another possible reason for access disparity is that small hospices located in rural areas often have higher per diem costs, yet receive lower Medicare payments due to a wage index adjustment (Virnig and Kind, 2003). As hospice care is largely provided in patients' homes, geographic variations in hospice usage may reflect more serious underlying problems in hospice development and delivery.

Our particular interest lies in identifying unserved areas in the state of Minnesota. Our data consist of ZIP code area-level Medicare beneficiary death counts from 2000 to 2002, as well as the number of these deaths among patients served

by each hospice, both based on Medicare billing records. The use of ZIP code areas as our zonal system has inherent problems (Grubisec, 2008), including that it evolves over time at the whim of the US Postal Service. Here, we use the ZIP code grid from the final year (2002) as obtained from ESRI. While there likely were small changes to this grid over our time period, we simply treat all three years' of data as arising from the 2002 grid. In what follows, we focus on the two hospice systems headquartered in the city of Duluth that serve rural northeast and north-central Minnesota, St. Luke's and St. Mary's/Duluth Clinic (SMDC). Figures 1a and c give raw data maps for St. Luke's, while those for SMDC appear in Figures 1b and d. The first row of the figure maps the numbers of hospice deaths during the three-year period by ZIP code for the two hospice systems, while the second row maps the internally standardized mortality ratios, i.e., actual hospice death count divided by expected deaths (taken as proportional to the total Medicare death count) in each ZIP code area. Using either definition of "service," St. Luke's service area appears much smaller and more tightly clustered than SMDCs.

Determining the "service area" for each hospice system based only on the ZIP code area-specific hospice and total death counts is not as easy as simply drawing boundaries that separate ZIP code areas with zero counts from those with nonzero counts, since a patient's actual and billing addresses may not coincide. Some patients still listed in the database have relocated (say, to a grown child's residence) out of the service area; conversely, some very low-population ZIP code areas may lie within the service area despite having

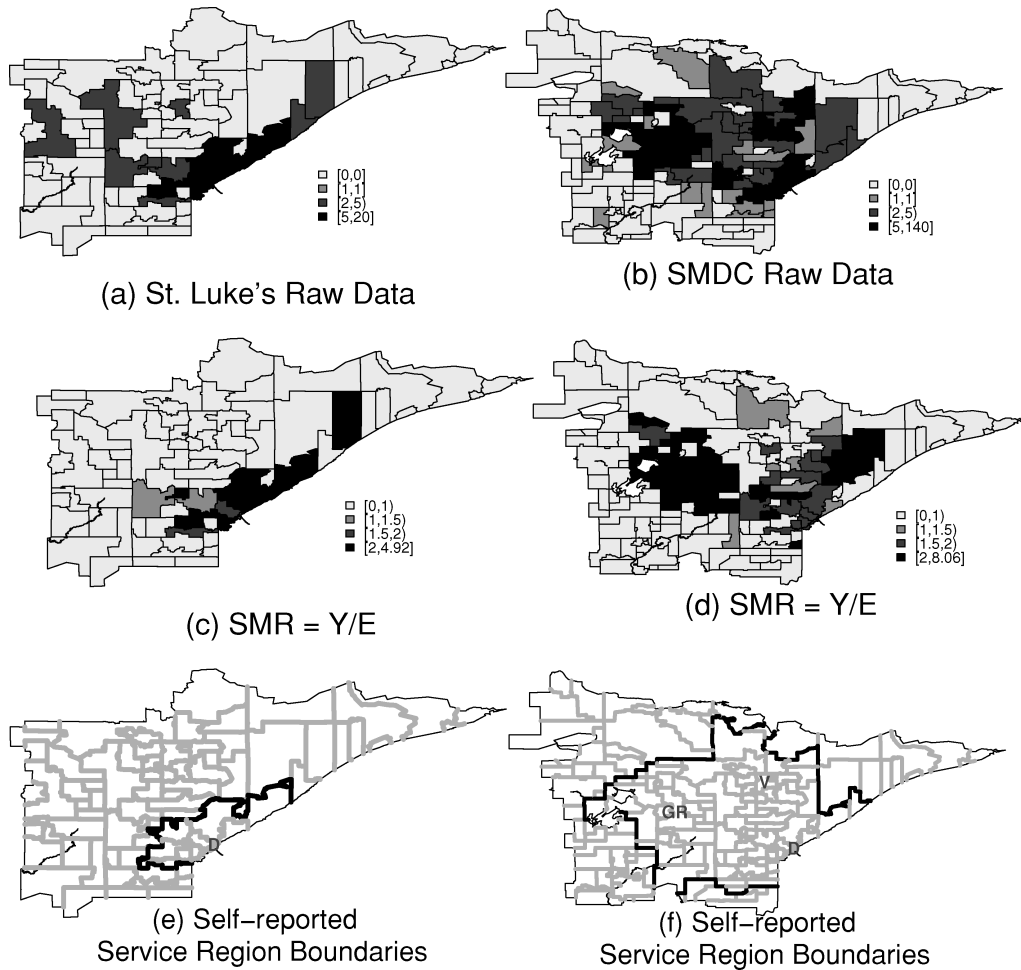


Figure 1. St. Luke's and SMDC hospice system usage data, northeastern Minnesota ZIP codes: (a) St. Luke's hospice death counts; (b) SMDC hospice death counts; (c) St. Luke's internally standardized mortality ratios; (d) SMDC internally standardized mortality ratios; (e) St. Luke's self-reported service area boundaries; (f) SMDC self-reported service area boundaries. In (e) and (f), hospice home bases are marked (D = Duluth, GR = Grand Rapids, V = Virginia). Note that the self-reported service area for St. Luke's is entirely contained by the self-reported service area for SMDC.

zero observed counts over the period. Traditional approaches to this problem include gravity-type flow models with distance decay functions that can be cast as log-linear models, and ad hoc methods using interpoint distances to delineate service areas. We initially analyzed our data with a similar spatial classification model using logistic regression, but found it unhelpful in identifying service area boundaries (see Web Appendix B).

While calling every hospice in the country is infeasible, one might wonder if this would at least provide a “gold standard” to help validate a statistical model in a few cases. To check this, the two hospices were contacted and lists of ZIP code areas that each said it served were obtained. These results are shown in Figures 1e and f for St. Luke's and SMDC, respectively. The former self-reported service area appears congruent with the observed data for St. Luke's, a smaller hospice focusing on ZIP codes in or just northeast of Duluth (indicated with a “D” in panels (e) and (f)). But this is not the case for SMDC, a fast-developing hospice system with home base in Duluth and two satellite bases in Grand Rapids (“GR”,

ZIP code 55744) and Virginia (“V”, ZIP code 55792), both north of Duluth. Comparing the SMDC self-report to the actual hospice death counts in Figure 1b, it does not appear that its service region extended quite as far south or west as claimed during the three years covered by our data. That is, while not all zero-count regions are truly “unserved,” large groups of zero-count regions far from the hospice bases and lying on the periphery of the map seem very unlikely to be served.

1.2 Spatial Smoothing and Boundary Analysis

Identifying boundaries over a geographic surface is an important topic in spatial statistics. The area is often referred to as *boundary analysis* or *wombling*, a name that (like “kriging”) pays homage to an early important paper in the area (Womble, 1951). Boundary analysis is concerned with finding edges across which areal units are significantly different. In public health, this is useful for detecting regions of significantly different disease mortality or incidence, thus improving decision-making regarding disease prevention and

control, allocation of resources, and so on. Boundary analysis is also common in disease ecology; see e.g. Wheeler and Waller (2008), who use spatially varying effects in wombling with river and mountain range barriers. Subsequent work modified the original wombling algorithm and developed algorithms for image or areal polygon data (e.g. Jacquez et al., 2000).

Spatial modeling of image data on regular lattices (e.g. pixels) often uses Markov random field (MRF) models (e.g. Geman and Geman, 1984; Besag, 1986; Geman and McClure, 1987; Jeng and Woods, 1991; Cressie, 1993, Sec 7.4; Rue and Held, 2005). Here, local spatial dependence between the observed image characteristics is captured by a *neighborhood structure*, where a pixel is independent of the rest given the values of its neighbors. Different neighborhood structures are possible, but all propose higher statistical dependence between data values from areas that are spatially closer, thus inducing local smoothing. However, this leads to a new problem: when real discontinuities (boundaries) exist between neighboring pixels, MRF models often lead to oversmoothing, blurring these edges. In order to preserve edges while still smoothing the image, a hidden edge process can be included in the model (e.g., Figueiredo and Leitao, 1997). However, the main purpose still lied in segmentation or noise reduction at areal level. Helterbrand et al. (1994) identify closed object boundaries that are precisely one-pixel wide using Bayesian statistical methods implemented via Markov chain Monte Carlo (MCMC) computational methods (e.g. Banerjee et al., 2004, Sec. 4.3).

MRF models implemented via Bayes–MCMC have been extensively applied to spatial mapping and smoothing problems. Although the boundary analysis problem for public health data bears similarities to the edge-preserving problem in image processing (Geman and Reynolds, 1992; Aykroyd, 1998), many significant differences exist. First, areal public health data are typically aggregated over an *irregular* polygonal lattice, such as a collection of counties or ZIP codes. Second, there are usually far fewer of these areas than the number of pixels that would arise in a typical image restoration problem, so we have far less data. Third, the areal units (polygons) are often quite different in size, shape, and number of neighbors, leading, for example, to different degrees of smoothing in urban and rural regions, as well as near the external boundary of the study region. Clever weighting schemes (e.g., Lawson et al., 1999) may offer a solution here, but they are typically ad hoc. Fourth, areal public health data often come with a number of spatially referenced covariates that need to be considered in the modeling; image denoising or reconstruction problems usually do not involve any additional variables. Finally, image data typically feature large homogeneous regions where neighboring pixels take *exactly* the same value (color), making edges relatively easy to spot (e.g., Lowell, 1997). By contrast, public health data are usually generously endowed with random noise, obscuring any true boundaries in the underlying risk or exposure surface. Indeed, many of the boundaries we seek will not be sharp, but instead correspond to more gradual changes over the spatial domain, so that traditional spatial smoothing methods are more appropriate. Thus, few of the image-restoration methods appearing in the statistical literature are directly applicable to the public health areal data setting.

Since our data are summaries over areal units, drawing lines across the regions (e.g. Banerjee and Gelfand, 2006) does not appear sensible. Identified boundaries in the surface should instead follow the existing geopolitical borders that generate the data and are a collection of segments dually indexed by ij , corresponding to the two adjacent regions i and j the segment separates. Deterministic methods include polygon wombling methods (e.g. Jacquez et al., 2000), as well as the algorithm implemented in the **BoundarySeer** software (<http://www.terraseer.com/products/boundaryseer.html>). Such methods ignore the stochastic nature of the data, and typically separate the problems of boundary detection and areal smoothing. Boundary segments are simply those that produce the k largest discrepancies between neighboring regions in a raw or smoothed areal map for some user-selected k . No stochastic measure of uncertainty can be associated with the identified boundaries, though see Jordan et al. (2005) for a discussion of boundary uncertainty derived from fuzzy-set principles.

Lu and Carlin (2005) proposed a fully model-based hierarchical Bayesian wombling model that accounted for stochasticity. This approach uses MRF models to account for spatial structure, but suffers from oversmoothing and also fails to produce the long series of connected boundary segments we desire.

Section 2 reviews Bayesian MRF technology and applies it to both the area and edge domains, including the case of random spatial neighborhoods. Section 3 then introduces a variety of hierarchical MRF areal wombling approaches that model the area and edge effects using various distributions and conditioning orders. Section 4 applies several of our proposed areal wombling methods to the northeastern Minnesota Medicare hospice utilization data, comparing the results graphically and analytically. Finally, Section 5 summarizes and discusses topics for future investigation.

2. Boundary Analysis and Edge Processes

2.1 Existing Methods for Areal Boundary Analysis

In wombling using areal data, dissimilarity measures called *boundary likelihood values* (BLVs) are calculated for each pair of adjacent areas based on a chosen metric. Then if a BLV exceeds some threshold c , or is among the top $k\%$ percent of BLVs (for some prechosen c or k), the edge corresponding to this BLV is declared to be part of the boundary, and is labeled a *boundary element* (BE). In *crisp* wombling, the corresponding *boundary membership value* (BMV) for this edge is then assigned to be 1; otherwise, the BMV is 0. In *fuzzy wombling*, BMVs between 0 and 1 are allowed to indicate partial boundary membership. The choice of boundary threshold value is subjective and usually ad hoc.

Lu and Carlin (2005; henceforth abbreviated LC) embedded the BLV idea within a Bayesian hierarchical model framework. These authors used a Poisson log-linear model for the disease count Y_i in region i , namely $Y_i | \beta, \phi_i \sim \text{Poisson}(\mu_i)$, where $\log \mu_i = \log E_i + \mathbf{x}_i' \beta + \phi_i$ for $i = 1, \dots, n$. Here, the E_i are internally standardized expected counts (assumed fixed and known) and the \mathbf{x}_i are known region-specific covariates observed over the n regions. The $\phi = (\phi_1, \dots, \phi_n)'$ are *random effects* that account for extra-Poisson variability in the observed data. A *conditionally autoregressive* (CAR)

distribution (Besag, 1974; see also Subsection 2.2 below) is used as the prior for ϕ to model the local dependency among the random effects.

LC define the theoretical BLVs as $\Delta_{|\eta|,ij} = |\eta_i - \eta_j|$ for all i adjacent to j , where $\eta_i = \mu_i/E_i$ measures the true underlying relative rate. Crisp and fuzzy wombled boundaries are based upon the posterior distribution of the $\Delta_{|\eta|,ij}$. In the crisp case, the border between area i and j is a BE if $E(\Delta_{|\eta|,ij} | \mathbf{y}) > c$, where again c is some prespecified constant believed to be of scientific interest, or simply set to deliver the top $k\%$ of the segments as BEs. For the fuzzy case, following the idea of an *exceedance probability*, the partial boundaries can be quantified by taking the $Pr(\Delta_{|\eta|,ij} > c | \mathbf{y})$ values as the BMVs. In either case, the posterior summaries are routinely obtained via MCMC methods. To construct boundaries based on the absolute (instead of relative) level of the process, we replace η_i by the mean μ_i in the preceding formulae. This is of interest since our wombling must address the presence or absence of hospice service in each area, *not* the relative saturation of hospice use per individual.

2.2 Markov Random Fields and Local Smoothing

The CAR model belongs to the more general class of MRF probability models (Besag, 1974; Rue and Held, 2005). For a given map $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ with \mathcal{E} as the neighborhood structure and \mathcal{S} as the set of sites, a vector ϕ forms an MRF if its joint distribution is determined by the full conditional distributions $p(\phi_i | \phi_j, i \neq j) = p(\phi_i | \phi_j, j \sim i)$, where \sim indicates i and j are neighbors according to \mathcal{E} .


The CAR distribution is an MRF formed from consideration of all *pairs* of neighbors. The joint distribution $CAR(\tau_\phi, W)$ for an $n \times 1$ vector ϕ is defined as proportional to

$$\exp \left\{ -\frac{\tau_\phi}{2} \phi' (D_w - W) \phi \right\}, \tag{1}$$

where τ_ϕ is a positive scale parameter, W is an $n \times n$ proximity matrix, and $D_w = \text{Diag}(w_{i+})$ with $w_{i+} = \sum_j w_{ij}$. Note that this joint distribution is improper since $(D_w - W) \mathbf{1} = \mathbf{0}$.


The most common choice for W is the 0-1 adjacency matrix: $w_{ij} = 1$ if and only if $i \neq j$ and $i \sim j$; otherwise $w_{ij} = 0$. Figure 2a illustrates this neighborhood structure on an idealized regular lattice. The dark square (Region 7) has four neighbors (Regions 3, 6, 8, and 11, shaded light gray). In this case, we have $w_{i+} = m_i$, the number of neighbors for region i , so the conditional distribution has mean $\bar{\phi}_i$, the average of the neighboring ϕ_j 's, and variance inversely proportional to m_i . Figure 2c repeats this exercise on an irregular lattice, St. Luke's map from the first column of Figure 1. Note the dark shaded region (Region 4) is adjacent to regions that are physically small and more densely populated (e.g., Regions 43 and 66), but also regions that are much larger and more thinly populated (e.g., Region 9).

The improper CAR above is sometimes called an *intrinsically autoregressive* (IAR) distribution in spatial statistics, or a *weak membrane*-type Gauss-MRF model (Figueiredo and Leitao, 1997) in image processing. A *proper* CAR can be obtained by setting $p(\phi | W, \rho) \propto \exp \left\{ -\frac{\tau_\phi}{2} \phi' (D_w - \rho W) \phi \right\}$, where ρ is chosen to make $(D_w - \rho W)$ nonsingular (Cressie,



(a) Area Adjacency, Regular Lattice

(b) Edge Adjacency, Regular Lattice



(c) Area Adjacency, Irregular Lattice

(d) Edge Adjacency, Irregular Lattice

Figure 2. Illustration of area and edge domain neighborhood structures: (a) areal neighborhood structure, regular idealized map; (b) edge neighborhood structure, regular idealized map; (c) areal neighborhood structure, irregular St. Luke's hospice map; (d) edge neighborhood structure, irregular St. Luke's hospice map. In panel (b), note that edges (such as (6,10) and (7,11)) can be neighbors even though their index sets have no areal units in common. In panel (d), note that the highlighted edge (4,9) is a neighbor of the very short edge (4,14) and the very long edge (9,14), but *not* edge (4,66), even though the entirety of this short edge is much closer to (4,9) than most of "neighboring" edge (4,14).

1993). Another proper CAR model has a joint distribution proportional to $\exp\{-\frac{\tau_\phi}{2}\phi'(D_\epsilon - W)\phi\}$, where $D_\epsilon = \text{Diag}(w_{i+} + \epsilon)$ for some small $\epsilon > 0$.

2.3 Edge Smoothing and Random Neighborhood Structure

As mentioned above, boundary analysis attempts to find edges across which areal units are significantly different. LC placed a statistical model on data arising from the areal units themselves, with final boundaries arising from these smoothed estimates. Although such methods are sensible for areal rate estimation, they appear less so for boundary analysis, since they do not directly model the edges, and may also smooth over true discontinuities in the surface that we hope to detect.

Recently, Ma et al. (2006) proposed a direct modeling in the edge domain, where the basic data elements are assumed to arise on the edge segments themselves. A CAR model for the edge segments is adopted to favor connected boundaries. For example, in Figure 2b, the thick black boundary corresponding to edge (7,11) has six “neighboring” edges, highlighted as thick gray lines. Thus, edge segments are adjacent if and only if they connect to one another. Note that edges (6,10) and (8,12) are adjacent to edge (7,11) even though these segments have no areal units in common. Figure 2d again repeats this exercise for the irregular St. Luke’s lattice, again with occasionally more obscure results (see figure caption).

In most CAR model implementations, the neighborhood structure is assumed known (i.e., W in (1) is fixed in advance). However, this may be antithetic to the goals of boundary detection. Following an idea from statistical social network analysis (Hoff et al., 2002), Lu et al. (2007) proposed use of a *random* neighborhood structure, wherein the w_{ij} are modeled as unknown parameters to be estimated. Specifically, let

$$w_{ij} | p_{ij} \sim \text{Bernoulli}(p_{ij}) \quad \text{and} \quad \text{logit}(p_{ij}) = \mathbf{z}'_{ij} \boldsymbol{\gamma}, \quad (2)$$

where \mathbf{z}_{ij} is a covariate vector of information relevant to the neighborhood structure for regions i and j . The components of \mathbf{z}_{ij} might capture geographical features (such as whether the two regions are separated by a river or mountain range) or sociodemographic features (such as the absolute difference in the two regions’ racial makeups or median incomes). This approach allows the neighborhood structure to be stochastically shaped by values of the process in each area, as well as other covariates indicating two areas’ inherent “closeness.” We refer to this model as the CARw, the additional w standing for the random weights w_{ij} .

The p_{ij} and w_{ij} in the CARw model measure the similarity between two regions, and thus provide a natural framework for areal wombling. Lu et al. (2007) redefine the BLV as $1 - E(w_{ij} | \mathbf{y})$ and show that these new *variance-based* BLVs offer an alternative to the use of the (mean-based) $\Delta_{\eta, ij}$ boundaries encouraged by LC. Note that the posterior summaries of the p_{ij} would also be natural fuzzy wombling inputs.

3. Hierarchical and Joint Area-Edge Boundary Analysis

3.1 Two-level CAR Model

As mentioned in subsection 2.3, the edge elements in the adjacency matrix can be modeled as random, potentially offering a natural framework for areal wombling. Since we prefer connected boundaries, given that a particular edge segment is

part of the boundary, we would like our model to favor the inclusion of neighboring edge segments in the boundary as well. The standard, 0–1 adjacency-based CAR model appears naturally suited to this task: all we require is a *second* CAR model on the edge space (in addition to the original CAR on the areal unit space) with edge adjacency matrix W^* determined by the regional map as illustrated in Figure 2.

Let us explore this *two-level hierarchical CAR* (CAR2) model in the case of Poisson data. Similar to the approach in LC, we start with

$$Y_i | \boldsymbol{\beta}, \phi_i \stackrel{\text{i.i.d.}}{\sim} \text{Poisson}(\mu_i)$$

$$\text{where } \log(\mu_i) = \log(E_i) + \mathbf{x}'_i \boldsymbol{\beta} + \phi_i, \quad i = 1, \dots, n,$$

$$\text{and } p(\boldsymbol{\phi} | \tau_\phi, W) = C(\tau_\phi, W) \exp\left\{-\frac{\tau_\phi}{2}\boldsymbol{\phi}'(D_w - W)\boldsymbol{\phi}\right\}, \quad (3)$$

where $C(\tau_\phi, W)$ is an unknown normalizing constant, discussed further in Web Appendix D. We then augment the CARw model (2) to

$$w_{ij} | p_{ij} \sim \text{Bernoulli}(p_{ij}) \quad \text{and} \quad \text{logit}(p_{ij}) = \mathbf{z}'_{ij} \boldsymbol{\gamma} + \theta_{ij}, \quad (4)$$

where θ_{ij} is a spatial random effect associated with the edge separating areas i and j . Note that in our random W setting, if two regions i and j are neighbors (i.e., $w_{ij} = 1$) then they must also be adjacent, but the converse need not be true. Because of the symmetry of W , we need only be concerned with its upper triangle.

For consistency with previous notation, we reorder the w_{ij} into the singly indexed vector $\boldsymbol{\xi} = (\xi_1, \dots, \xi_K)'$, where K is the number of regional adjacencies in the map. We also carry out a corresponding reordering of the θ_{ij} into a vector $\boldsymbol{\psi} = (\psi_1, \dots, \psi_K)'$. We then place the second-level CAR as the prior on the edge random effects, i.e.,

$$\boldsymbol{\psi} | \tau_\psi \sim \text{CAR}(\tau_\psi, W^*), \quad (5)$$

so that ψ_k has conditional distribution $N(\bar{\psi}_k, 1/(\tau_\psi w_{k+}^*))$, where $\tau_\psi > 0$ and W^* is the fixed $K \times K$ 0–1 adjacency matrix for $\boldsymbol{\psi}$, determined as in Figure 2b.

Equations (3)–(5) comprise the CAR2 model. Vague conjugate gamma prior distributions for the precision hyperparameters τ_ϕ and τ_ψ , along with normal or flat priors for $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$, complete the hierarchical specification. The posterior distribution of the parameters can be estimated via MCMC techniques; The “inhomogeneous model” of Aykroyd (1998) is a Gaussian-response variant of the CAR2 for image data over a regular grid.

3.2 Joint Modeling of Areal and Edge Effects for Areal Data

A primary issue in implementing the CAR2 (or CARw) method is the determination of good “discrepancy” covariates \mathbf{z}_{ij} . Although $\boldsymbol{\gamma}$ is estimable even under a noninformative prior distribution, these second-level regression coefficients are often hard to estimate. At the same time, p_{ij} (and correspondingly w_{ij}) can be sensitive to the prior specification of $\boldsymbol{\gamma}$. Since the edge parameters enter the model only to specify the variances of the first-level random effects, they may be “too far away from the data” in the hierarchical model. This motivates a model with fewer levels or more direct modeling of edge effects.

As such, in this subsection we now consider “site-edge” (SE) models, where both the areal units (sites) *and* the edges between them contribute random effects to the mean structure. Let $\mathcal{G} = (\mathcal{S}, \mathcal{E})$, where $\mathcal{S} = \{1, \dots, n\}$ is a set of sites/areas, and $\mathcal{E} = \{(i, j) : i \sim j\}$ is a set of edges, where \sim indicates the symmetric “adjacency” relation. Suppose the data follow an exponential family likelihood, and let $\phi = (\phi^S, \phi^E)$ be a vector of site- and edge-level effects, respectively. The general form of our SE model is then $g(\mu_i) = f(\mathbf{x}_i, \beta, \phi)$ and $\phi = (\phi^S, \phi^E) \sim MRF$, where μ_i is the mean of the likelihood and $g(\cdot)$ stands for the canonical link function (e.g., the log link for the Poisson).

Hierarchically modeling both $p(\phi^S | \phi^E)$ and $p(\phi^E)$ as MRFs guarantees that the joint distribution $p(\phi^S, \phi^E)$ is an MRF as well by construction (Jeng and Woods, 1991). Here, we take f to be linear, but this is not required for posterior propriety. To facilitate parameter identification and data information flow while encouraging sensible interaction between the areal and edge random effects, we now propose the hierarchical model,

$$\begin{aligned} Y_i | \beta, \phi_i^S &\sim \text{Poisson}(\mu_i) \\ \text{with } \log \mu_i &= \log E_i + \mathbf{x}_i' \beta + \phi_i^S, \quad i = 1, \dots, n, \\ p(\phi^S | \phi^E, \tau_\phi) &= C(\tau_\phi, \phi^E) \exp \left\{ -\frac{\tau_\phi}{2} \sum_{i \sim j} (1 - \phi_{ij}^E) (\phi_i^S - \phi_j^S)^2 \right\}, \end{aligned} \quad (6)$$

and

$$p(\phi^E) \propto \exp \left\{ -\nu \sum_{ij \sim kl} \phi_{ij}^E \phi_{kl}^E \right\}, \quad (7)$$

where $\phi_i^S \in \mathfrak{R}$ as before, but now $\phi_{ij}^E \in \{0, 1\}$ for all edges $(i, j) \in \mathcal{E}$. The conditional distribution in (6) is IAR, with $(1 - \phi_{ij}^E)$ playing the roles of the w_{ij} in (3). That is, $\phi_{ij}^E = 1$ if edge (i, j) is a BE, and 0 otherwise. Thus, smoothing of neighboring ϕ_i^S and ϕ_j^S is only encouraged if there is no boundary between them. The prior for ϕ^E in (7) is an *Ising* model with tuning parameter ν , often used in image restoration (e.g. Geman and Geman, 1984, p. 725). This prior yields a binary MRF that allows binary variables (the ϕ_{ij}^E 's) to directly borrow strength across their neighbors, avoiding the need for a link-function to introduce continuous spatial effects as in (4). The ν here is interpreted as measuring “binding strength” between the edges; smaller values of ν lead to more connected BEs, hence more separated areal units. In a Bayesian framework, ν can be assigned a prior distribution or take a fixed value; we follow the image-restoration literature and choose the second approach in this article, though we would recommend comparing results under different ν values. We refer to models (6) and (7) as an *SE-Ising* model.

Our use of the improper CAR prior for ϕ^S makes the joint prior $p(\phi^S, \phi^E) \equiv p(\phi^S | \phi^E) p(\phi^E)$ improper regardless of the choice of $p(\phi^E)$, but the joint *posterior* of these parameters will still be proper. To see this, note that $p(\phi^S | \phi^E, \mathbf{y})$ is proper under the usual improper CAR prior, and the discrete support of $p(\phi^E)$ in (7) means it too is proper by construction. Since $p(\phi^S, \phi^E | \mathbf{y}) \propto p(\phi^S | \phi^E, \mathbf{y}) p(\phi^E)$, the joint posterior is proper as well.

While the SE-Ising model is quite sensible for boundary analysis, it does not explicitly encourage long strings of connected boundary segments of the sort that would be needed to separate a hospice service area from an unserved area. As such, we further propose a *penalized* SE-Ising distribution

$$p(\phi^E) \propto \exp \left\{ -\nu \sum_{ij \sim kl} \phi_{ij}^E \phi_{kl}^E + \kappa M \right\}, \quad (8)$$

where M is the number of strings of connected “on” edges ($\phi_{ij}^E = 1$) and $\kappa < 0$ is a second tuning parameter. Adding this additional penalty on edge arrangements that do not favor series of connected boundary segments helps to impose the kind of structure we want on our fitted boundaries; see Web Tables 1 and 2 in Web Appendix A for a brief investigation of how this penalized SE-Ising prior behaves for various ν and κ .

4. Data Analysis

We now apply the LC, CAR2, SE-Ising, and penalized SE-Ising models to the Medicare hospice utilization data. With the latter three models we use edge correction, while for all four methods we use a *thresholding* approach designed to detect a boundary only when differences in the means of adjacent ZIP codes lie on opposite sides of some predetermined minimum service level; see Web Appendix C for full details.

Our analysis considers a single covariate x_i , the inter-centroidal (geodetic) distance from the patient’s ZIP code area to the nearest relevant hospice home base ZIP code area (see Figures 1e and f for locations). Since hospice services are provided in the patient’s home, increasing this distance should decrease the probability of that ZIP code area being served.

We use vague $N(0, 10^4)$ priors for both the intercept β_0 and distance effect β_1 . All of our models for both hospices also employ gamma priors for τ_ϕ having mean 1 and variance 1; this prior still permits significant prior-to-posterior Bayesian learning for this parameter while delivering acceptable MCMC convergence. For the SE-Ising model, we begin by setting the binding strength parameter ν equal to 0.5, and additionally set $\kappa = -3$ in the penalized SE-Ising model, though we revisit these selections in Web Tables 1 and 2. For the CAR2 model, we were unable to identify satisfactory areal discrepancy covariates \mathbf{z}_{ij} at the ZIP code area level, though in a previous, county-based analysis (Ma and Carlin, 2007) we used median income, local business pattern, and health insurance coverage. While median income would likely be useful here as well, the logit in (4) contains only the random effects θ_{ij} , assigned the second-stage zero-centered CAR in (5). For τ_ψ , the precision parameter of this second-level CAR, we use the same gamma prior as that for τ_ϕ . We tried different gamma distributions and even fixed τ_ψ at its MLE based on the self-reported boundaries. Although τ_ψ and ψ estimates are sensitive to prior choices, the lower-level parameter estimates are fairly robust. Finally, our summary displays are based on the more robust posterior medians (not means), acknowledging possible skewness in the posterior densities.

4.1 Model Selection

We use the deviance information criterion (DIC; Spiegelhalter et al., 2002) to select the best model for our data set. Table 1 gives DIC comparisons for several areal wombling models for St. Luke’s and SMDC, respectively. Overall, the

Table 1
 Comparison of p_D and DIC Values for LC, CAR2, SE-Ising and Penalized SE-Ising Areal Wombling Methods, St. Luke's and SMDC Hospice Data

	St. Luke's data			SMDC data		
	$\bar{D}(\theta)$	p_D	DIC = $\bar{D}(\theta) + p_D$	$\bar{D}(\theta)$	p_D	DIC = $\bar{D}(\theta) + p_D$
LC	112.2	23.7	135.9	357.8	70.7	428.5
LC with x	113.2	15.6	128.9	357.8	61.8	419.6
CAR2	133.2	22.7	155.9	482.6	65.3	547.9
CAR2 with x	117.9	14.3	132.2	416.1	44.5	460.6
SE-Ising	117.7	13.2	130.9	361.5	62.7	424.3
SE-Ising with x	117.1	10.2	127.3	363.3	45.2	408.5
Penalized SE-Ising	113.2	20.8	134.0	369.0	58.7	427.7
Penalized SE-Ising with x	113.7	14.1	127.8	355.8	56.5	412.3

SE-Ising models perform best: compared to LC, these enjoy similar fit (\bar{D}) but smaller effective parameter count (p_D). Apparently by incorporating the random edge process, areal effects not separated by boundaries are more homogeneous, thus reducing the overall effective parameter burden. Although both CAR2 and SE-Ising impose random edge structures, the former's additional upper-level Bernoulli and CAR variability seems to weaken identifiability and degrade fit (larger \bar{D}), especially for SMDC. Including the distance to nearest home base covariate improves the DIC (though not the \bar{D}) score of every model. This gain is intuitively sensible for St. Luke's data, with its single home base and geographically condensed service region.

The corresponding effects are significant: for example, using the SE-Ising model, the posterior medians and 95% equal-tail credible sets for β are -2.8 ($-3.75, -1.71$) for St. Luke's and -2.7 ($-4.19, -1.64$) for SMDC. Using the penalized SE-Ising model, the posterior medians and 95% equal-tail credible sets

for β are -2.93 ($-4.26, -1.58$) for St. Luke's and -3.06 ($-5.03, -1.50$) for SMDC. The negative signs indicate that the farther away a ZIP code is from the nearest hospice home base, the less likely it is to be served by that hospice. As such, we include the distance covariate as x in all of our subsequent analyses. We also adopt the penalized version of the SE-Ising model in what follows, in order to take advantage of its ability to better connect boundary segments in our service area determination problem. Web Appendix A presents a sensitivity analysis for the binding strength parameter ν and the penalty parameter κ in (8).

4.2 Service Area Boundaries

Figures 3a and b show μ -based boundary maps for St. Luke's, while Figures 3c and d give them for SMDC. All four panels in these figures are based on absolute posterior medians of $\Delta_{\mu,ij} = \mu_i - \mu_j$. Panels (a) and (c) give results from the LC model, which appears to do a credible job for both hospices.

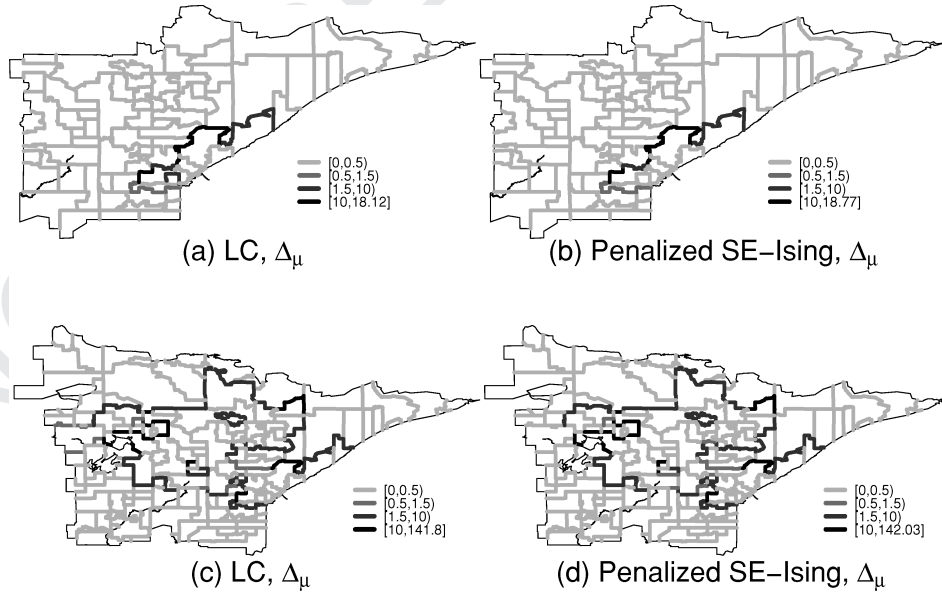


Figure 3. Maps of St. Luke's and SMDC's service area boundaries: (a) St. Luke's service area boundaries given by the LC (usual CAR) model; (b) St. Luke's service area boundaries given by the penalized SE-Ising model; (c) SMDC's service area boundaries given by the LC (usual CAR) model; (d) SMDC's service area boundaries given by the penalized SE-Ising model.

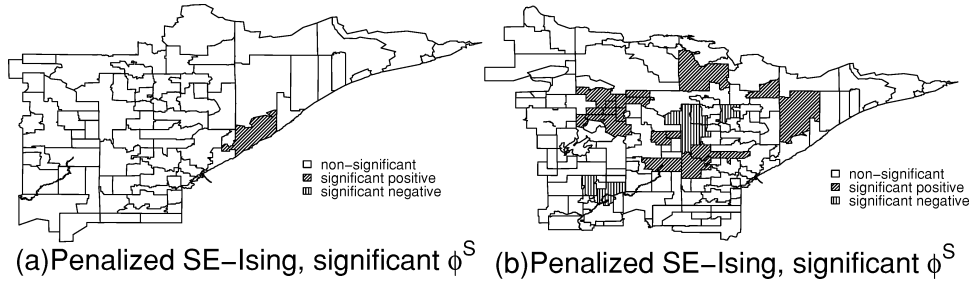


Figure 4. Maps of significant areal residuals ϕ_i^S , penalized SE-Ising model: (a) St. Luke's hospice; (b) SMDC hospice.

However, even in the easier St. Luke's case, the LC map does include some "clutter" (identified boundary segments apparently internal to the service area) near the bottom of the map.

Panels (b) and (d) in Figure 3 give the hierarchically smoothed boundaries from the penalized SE-Ising model. For St. Luke's, the penalized SE-Ising boundaries in Figure 3b are quite satisfactory, showing essentially no internal clutter and offering a good match with the self-reported boundaries in Figure 1e. However, for SMDC the boundaries are less well connected, perhaps owing to the more complex nature of the data that feature a much larger service region and comprise three offices shown in Figure 1f.

The wombled boundaries in Figures 3c and d are quite similar, as the μ_i are fairly well-estimated by any reasonable spatial model. However, none of our SMDC wombled maps provide a very good match to the self-reported boundaries in the south, since the data do not support the claim of service coverage there. This disagreement between our results and the self-report could be the result of reporting lags or migration in and out of service areas, but is more likely due to the focus of some hospices (especially larger ones, like SMDC) on urban patients, at the expense of harder-to-reach rural ones.

4.3 Residual Boundaries

Looking again at our penalized SE-Ising model, (6)–(8), in the presence of any covariates \mathbf{x}_i , the ϕ_i^S can be interpreted as spatial *residuals*, contributing spatially oriented adjustments to the Poisson mean structure beyond what is explained by the covariates. The ϕ_{ij}^E in turn delineate the edges of this spatial residual surface. Thus, boundaries based not on the $\Delta_{\mu,ij}$ but on the $\Delta_{\phi^S,ij} \equiv \phi_i^S - \phi_j^S$ or on the ϕ_{ij}^E themselves would provide information about boundaries separating areas having significantly different *residuals*. This is potentially important secondary information, since such boundaries could suggest that areas to compare as part of a search for spatially varying covariates not yet included in the \mathbf{x}_i vector.

Figure 4 maps the significant ϕ_i^S from our penalized SE-Ising model for St. Luke's and SMDC, respectively, where we define "significant" as those regions having a central 95% posterior CI for ϕ_i^S that excludes 0. The former map is easy to interpret, featuring only a few ZIP code areas with significantly positive residuals northeast of the hospice home base in Duluth. While these ZIP code areas are fairly far from the home base, they are well-connected by road to Duluth and are thus easily served. Meanwhile the SMDC map is once again more interesting but more difficult to interpret. A few areas of negative residuals are seen in the far southwest, with an

even more interesting group in the center of the map near Hibbing, a city of roughly 20,000 people that lies in between the three SMDC bases. Although it lies inside the service region, it apparently serves fewer patients than expected. Conversely, several positive residuals are seen in ZIP code areas near one of the three SMDC home bases that serve even more patients than their small x_i 's would predict.

Turning to the corresponding maps of wombled boundaries, panels (a) and (c) in Figure 5 highlight the posterior medians of the significant $\Delta_{\phi^S,ij}$ (again, those whose posterior 95% CIs do not include 0), while panels (b) and (d) simply shade boundaries according to the posterior means of ϕ^E . Maps in the first row of this figure correspond to St. Luke's and those in the second row correspond to SMDC. Clearly the $\Delta_{\phi^S,ij}$ and ϕ^E maps do not reflect precisely the same information, but one can still spot many similarities between them. For example, a diagonal edge separating urban Duluth from the rest of the service area is clearly visible in both St. Luke's maps, and the southern and western boundaries of Hibbing are outlined in both SMDC maps. Also of interest are the two identified segments (one long, one short) in the northeast corner of Figures 5b and d. These segments separate a ZIP code area consisting primarily of the low-population Superior National Forest from several other, more human-inhabited ZIP code areas. Perhaps an indicator of whether a ZIP code area is a largely uninhabited, protected area should be added as an areal covariate \mathbf{x}_i in (6) and –(7).

5. Discussion and Future Work

We have proposed new areal models to obtain boundaries for Medicare hospice service areas. We used WinBUGS, R, and C++ languages and ArcView shape files to obtain ZIP code boundaries. Our models depart from existing deterministic methods by allowing stochastic uncertainty on the edges and offering full posterior inference on them. The edge-corrected penalized SE-Ising method excels over the CAR2 in encircling the Medicare service areas and exhibits less internal clutter. The LC method emerged as a surprisingly strong competitor, and appears sensible whenever the risk of oversmoothing is low.

Finally, our current models do not account for correlation between the two hospice systems that may be negative, if the competition for hospice patients were a zero-sum game; or positive, if the increasing popularity of hospice care means hospice systems can expand enrollment simultaneously. A key issue here is whether we seek a single set of boundaries (corresponding to the jointly served areas) or multiple sets (one

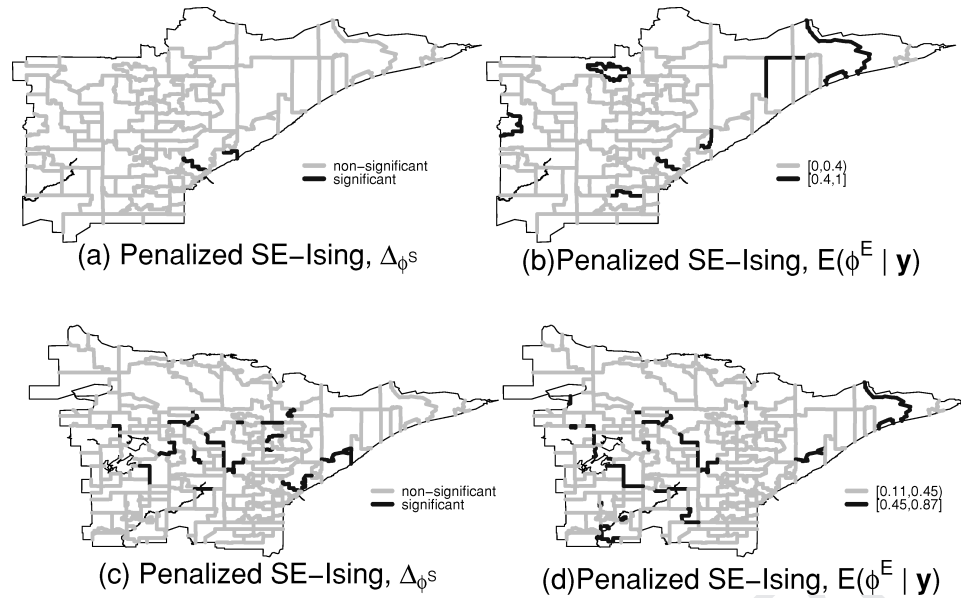


Figure 5. Residual wombled maps based on the penalized SE-Ising model: (a) residual boundaries for St. Luke's based on the $\Delta_{\phi^S, ij}$; (b) residual boundaries for St. Luke's based on $E(\phi^E | \mathbf{y})$; (c) residual boundaries for SMDC based on the $\Delta_{\phi^S, ij}$; (d) residual boundaries for SMDC based on $E(\phi^E | \mathbf{y})$.

for each hospice). In either case, a *multivariate* CAR will be required for a random site or edge effect distribution. Ma and Carlin (2007) apply these ideas to the problem of boundary analysis for multiple areal-level disease data.

6. Supplementary Materials

Web Appendices A–D, Figures and Tables referenced in Sections 1, 3 and 4 are available under the Paper Information link at the *Biometrics* website <http://www.biometrics.tibs.org>.

ACKNOWLEDGEMENTS

The authors thank the referees, the associate editor and the editor for their valuable comments and suggestions. The work of all three authors was supported in part by NIH grants 1–R01–CA112444–03 and 1–R01–CA95955–01.

REFERENCES

- Aykroyd, R. G. (1998). Bayesian estimation for homogeneous and inhomogeneous Gaussian random fields. *IEEE Transactions Pattern Analysis and Machine Intelligence* **20**, 533–539.
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton, FL: Chapman and Hall/CRC Press.
- Banerjee, S. and Gelfand, A. E. (2006). Bayesian wombling: Curvilinear gradient assessment under spatial process models. *Journal of American Statistical Association* **101**, 1487–1501.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of Royal Statistical Society, Series B* **36**, 192–236.
- Cressie, N. A. C. (1993). *Statistics for Spatial Data*, 2nd ed. New York: Wiley.
- Figueiredo, M. A. T. and Leitao, J. M. N. (1997). Unsupervised image restoration and edge location using compound Gauss-Markov random fields and the MDL principle. *IEEE Transactions on Image Processing* **6**, 1089–1102.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6**, 721–742.
- Geman, S. and McClure, D. E. (1987). Statistical methods for tomographic image reconstruction. *Bulletin of the International Statistical Institute* **52**, 5–21.
- Geman, D. and Reynolds, G. (1992). Constrained restoration and the recovery of discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**, 367–383.
- Grubisec, T. H. (2008). Zip codes and spatial analysis: Problems and prospects. *Socio-Economic Planning Sciences* **42**, 129–149.
- Helterbrand, J. D., Cressie, N., and Davidson, J. L. (1994). A statistical approach to identifying closed object boundaries in images. *Advances in Applied Probability* **26**, 831–854.
- Hoff, P. D., Raftery, A. E., and Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of American Statistical Association* **97**, 1090–1098.
- Jacquez, G. M., Maruca, S. L., and Fortin, M.-J. (2000). From fields to objects: A review of geographic boundary analysis. *Journal of Geographical Systems* **2**, 221–241.
- Jeng, F. C., and Woods, J. W. (1991). Compound Gauss-Markov random fields for image estimation. *IEEE Transactions in Signal Processing* **39**, 683–691.
- Jordan, G. J., Fortin, M.-J., and Lertzman, K. P. (2005). Assessing spatial uncertainty associated with forest fire boundary delineation. *Landscape Ecology* **20**, 719–731.
- Lawson, A. B., Biggeri, A., and Dreassi, E. (1999). Edge effects in disease mapping. In *Disease Mapping and Risk Assessment for Public Health*, A. Lawson et al. (eds.), pp. 83–96, Chichester: Wiley.
- Lowell, K. (1997). Effect(s) of the “no-same-color-touching” constraint on the join-count statistic: A simulation study. *Geographical Analysis* **29**, 339–353.
- Lu, H. and Carlin, B. P. (2005). Bayesian areal wombling for geographical boundary analysis. *Geographical Analysis* **37**, 265–285.
- Lu, H., Reilly, C., Banerjee, S., and Carlin, B. P. (2007). Bayesian areal wombling via adjacency modeling. *Environmental and Ecological Statistics* **14**, 433–452.

- Ma, H. and Carlin, B. P. (2007). Bayesian multivariate areal wombling for multiple disease boundary analysis. *Bayesian Analysis* **2**, 281–302.
- Ma, H., Virnig, B., and Carlin, B. P. (2006). Spatial methods in areal administrative data analysis. *Italian Journal of Public Health* **3**, 94–104.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. Boca Raton, FL: Chapman and Hall/CRC Press.
- Spiegelhalter, D. J., Best, N., Carlin, B. P., and van derLinde, A. (2002). Bayesian measures of model complexity and fit (with discussion). *Journal of Royal Statistical Society, Series B* **64**, 583–639.
- U.S. General Accounting Office (2000). Medicare: More beneficiaries use hospice, but for fewer days of care. *GAO/HEHS-00-182*, Washington, DC: GAO.
- Virnig, B. A. and Kind, S. (2003). Using administrative data to study hospice use. In *Symptom Research: Methods and Opportunities, Interactive Clinical Research Textbook*, M. B. Max and J. Lynn (eds.).
- Wheeler, D. C. and Waller, L. A. (2008). Mountains, valleys, and rivers: The transmission of raccoon rabies over a heterogeneous landscape. *Journal of Agricultural Biology and Environmental Statistics* **13**, 388–406.
- Womble, W. H. (1951). Differential systematics. *Science* **114**, 315–322.

Q2

Received August 2008. Revised March 2009.
Accepted March 2009.

Queries

Q1 Author: Please provide the expanded form of ESRI.

Q2 Author: Please provide publisher's name and detail in Vimig and Kind (2003).