

SMOOTHED ANOVA WITH SPATIAL EFFECTS AS A COMPETITOR TO MCAR IN MULTIVARIATE SPATIAL SMOOTHING

BY YUFEN ZHANG, JAMES S. HODGES AND SUDIPTO BANERJEE¹

Novartis Pharmaceuticals, University of Minnesota and University of Minnesota

Rapid developments in geographical information systems (GIS) continue to generate interest in analyzing complex spatial datasets. One area of activity is in creating smoothed disease maps to describe the geographic variation of disease and generate hypotheses for apparent differences in risk. With multiple diseases, a multivariate conditionally autoregressive (MCAR) model is often used to smooth across space while accounting for associations between the diseases. The MCAR, however, imposes complex covariance structures that are difficult to interpret and estimate. This article develops a much simpler alternative approach building upon the techniques of smoothed ANOVA (SANOVA). Instead of simply shrinking effects without any structure, here we use SANOVA to smooth spatial random effects by taking advantage of the spatial structure. We extend SANOVA to cases in which one factor is a spatial lattice, which is smoothed using a CAR model, and a second factor is, for example, type of cancer. Datasets routinely lack enough information to identify the additional structure of MCAR. SANOVA offers a simpler and more intelligible structure than the MCAR while performing as well. We demonstrate our approach with simulation studies designed to compare SANOVA with different design matrices versus MCAR with different priors. Subsequently a cancer-surveillance dataset, describing incidence of 3-cancers in Minnesota's 87 counties, is analyzed using both approaches, showing the competitiveness of the SANOVA approach.

1. Introduction. Statistical modeling and analysis of spatially referenced data receive considerable interest due to the increasing availability of geographical information systems (GIS) and spatial databases. For data aggregated over geographic regions such as counties, census tracts or ZIP codes (often called *areal* data), with individual identifiers and precise locations removed, inferential objectives focus on models for spatial clustering and variation. Such models are often used in epidemiology and public health to understand geographical patterns in disease incidence and morbidity. Recent reviews of methods for such data include Lawson et al. (1999), Elliott et al. (2000), Waller and Gotway (2004) and Rue and Held (2005). Traditionally such data have been modeled using conditionally specified probability models that shrink or smooth spatial effects by borrowing strength

Received August 2008; revised December 2008.

¹Supported in part by NIH Grant 1-R01-CA95995.

Key words and phrases. Analysis of variance, Bayesian inference, conditionally autoregressive model, hierarchical model, smoothing.

from neighboring regions. Perhaps the most pervasive model is the conditionally autoregressive (CAR) family pioneered by Besag (1974), which has been widely investigated and applied to spatial epidemiological data [Wakefield (2007) gives an excellent review]. Recently the CAR has been extended to multivariate responses, building on multivariate conditional autoregressive (MCAR) models described by Mardia (1988). Gelfand and Vonatsou (2003) and Carlin and Banerjee (2003) discussed their use in hierarchical models, while Kim, Sun and Tsutakawa (2001) presented a different “twofold CAR” model for counts of two diseases in each areal unit. Other extensions allowing flexible modeling of cross-correlations include Sain and Cressie (2002), Jin, Carlin and Banerjee (2005) and Jin, Banerjee and Carlin (2007). The MCAR can be viewed as a conditionally specified probability model for interactions between space and an attribute of interest. For instance, in disease mapping interest often lies in modeling geographical patterns in disease rates or counts of several diseases. The MCAR acknowledges dependence between the diseases as well as dependence across space. However, practical difficulties arise from MCAR’s elaborate dependence structure: most interaction effects will be weakly identified by the data, so the dependence structure is poorly identified. In hierarchical models [e.g., Gelfand and Vonatsou (2003), Jin, Carlin and Banerjee (2005, 2007)], strong prior distributions may improve identifiability, but this is not uncontroversial, as inferences are sensitive to the prior and perhaps unreliable without genuine prior information. This article proposes a much simpler and more interpretable alternative to the MCAR, modeling multivariate spatial effects using smoothed analysis of variance (SANOVA) as developed by Hodges, Carlin and Fan (2007), henceforth HCSC. Unlike an ANOVA that is used to identify some interaction effects to retain and others to remove, SANOVA mostly retains effects that are large, mostly removes those that are small, and partially retains middling effects. (Loosely speaking, “large,” “middling” and “small” describe the size of the unsmoothed effects compared to their standard errors.) To accommodate rich dependence structures, MCAR introduces weakly identifiable parameters that complicate estimation. SANOVA, on the other hand, focuses instead on smoothing interactions to yield more stable and reliable results. Our intended contribution is to show how SANOVA can solve the multiple disease mapping problem while avoiding the dauntingly complex covariance structures imposed by MCAR and its generalizations. We demonstrate that SANOVA produces inference that is largely indistinguishable from MCAR, yet SANOVA is simpler, more explicit, easier to put priors on and easier to estimate. The rest of the article is as follows. Section 2 reviews SANOVA and MCAR, identifying SANOVA as a special case of MCAR. Section 3 is a “tournament” of simulation experiments comparing SANOVA with MCAR for normal and Poisson data, while Section 4 analyzes data describing the number of deaths from lung, larynx and esophagus cancer in Minnesota between 1990 and 2000. A summary and discussion of future research in Section 5 concludes the paper. Zhang, Hodges and Banerjee (2009) (Appendices) gives computational and technical details.

2. The competitors.

2.1. *Smoothing spatial effects using SANOVA.*

2.1.1. *SANOVA for balanced, single-error-term models [HCSC (2007)].* Consider a balanced, single-error-term analysis of variance, with M_1 degrees of freedom for main effects and M_2 degrees of freedom for interactions. Specify this ANOVA as a linear model: let A_1 denote columns in the design matrix for main effects, and A_2 denote columns in the design matrix for interactions. Assume the design has c cells and n observations per cell, giving cn observations in total. To simplify later calculations, normalize the columns of A_1 and A_2 so $A_1' A_1 = I_{M_1}$ and $A_2' A_2 = I_{M_2}$. (Note: HCSC normalized columns differently, fixing $A_1' A_1 = cn I_{M_1}$ and $A_2' A_2 = cn I_{M_2}$.) Then write the ANOVA as

$$(1) \quad \mathbf{y} = [A_1|A_2] \begin{bmatrix} \Theta_1 \\ \Theta_2 \end{bmatrix} + \epsilon = A_1 \Theta_1 + A_2 \Theta_2 + \epsilon,$$

where $\epsilon \sim N(0, \frac{1}{\eta_0} I)$ with η_0 being a precision, \mathbf{y} is $cn \times 1$, A_1 is $cn \times M_1$, A_2 is $cn \times M_2$, Θ_1 is $M_1 \times 1$, Θ_2 is $M_2 \times 1$, and ϵ is $cn \times 1$. This ANOVA is smoothed by further modeling Θ . HCSC emphasized smoothing interactions, although main effects can be smoothed by exactly the same means. Following HCSC, we add constraints (or a prior) on Θ_2 as $\theta_{M_1+j} \sim N(0, 1/\eta_j)$ for $j = 1, \dots, M_2$, written as

$$(2) \quad \mathbf{0}_{M_2} = [0_{M_2 \times M_1} | I_{M_2}] \begin{bmatrix} \Theta_1 \\ \Theta_2 \end{bmatrix} + \delta,$$

where $\delta \sim N(0, \text{diag}(\frac{1}{\eta_j}))$, in the manner of Lee and Nelder (1996) and Hodges (1998). Combining (1) and (2), express this hierarchical model as a linear model:

$$(3) \quad \begin{bmatrix} \mathbf{y} \\ \mathbf{0}_{M_2} \end{bmatrix} = \begin{bmatrix} A_1 & A_2 \\ 0_{M_2 \times M_1} & I_{M_2} \end{bmatrix} \begin{bmatrix} \Theta_1 \\ \Theta_2 \end{bmatrix} + \begin{bmatrix} \epsilon \\ \delta \end{bmatrix}.$$

More compactly, write

$$(4) \quad \mathbf{Y} = X \Theta + \mathbf{e},$$

where \mathbf{Y} has dimension $(cn + M_2) \times 1$ and \mathbf{e} 's covariance Γ is block diagonal with blocks $\Gamma_1 = \frac{1}{\eta_0} I_{cn}$ for the data cases (rows of X corresponding to the observation \mathbf{y}) and $\Gamma_2 = \text{diag}(1/\eta_1, \dots, 1/\eta_{M_2})$ for the constraint cases (rows of X with error term δ). For convenience, define the matrix $X_D = [A_1|A_2]$, the data-case part of X . This development can be done using the mixed linear model (MLM) formulation traditionally written as $\mathbf{y} = X\beta + Z\mathbf{u} + \epsilon$, where our (1) supplies this equation and $\mathbf{u} = \Theta_2 \sim N(\mathbf{0}, \Gamma_2)$. The development to follow can also be done using the MLM formulation at the price of slightly greater complexity, so we omit it. HCSC developed SANOVA for exchangeable priors on groups formed from components of Θ_2 . The next section develops the extension to spatial smoothing.

2.1.2. *What is CAR?* Suppose a map has N regions, each with an unknown quantity of interest ϕ_i , $i = 1, \dots, N$. A conditionally autoregressive (CAR) model specifies the full conditional distribution of each ϕ_i as

$$(5) \quad \phi_i \mid \phi_j, j \neq i, \sim N\left(\frac{\alpha}{m_i} \sum_{i \sim j} \phi_j, \frac{1}{\tau m_i}\right), \quad i, j = 1, \dots, N,$$

where $i \sim j$ denotes that region j is a *neighbor* of region i (typically defined as spatially adjacent), and m_i is the number of region i 's neighbors. Equation (5) reduces to the well-known intrinsic conditionally autoregressive (ICAR) model [Besag, York and Mollié (1991)] if $\alpha = 1$ or an independence model if $\alpha = 0$. The ICAR model induces “local” smoothing by borrowing strength from neighbors, while the independence model assumes spatial independence and induces “global” smoothing. The CAR prior’s smoothing parameter α also controls the strength of spatial dependence among regions, though it has long been appreciated that a fairly large α may be required to induce large spatial correlation; see Wall (2004) for recent discussion and examples. It is well known [e.g., Besag (1974)] that the conditional specifications in (5) lead to a valid joint distribution for $\boldsymbol{\phi} = (\phi_1, \dots, \phi_N)'$ expressed in terms of the map’s neighborhood structure. If Q is an $N \times N$ matrix such that $Q_{ii} = m_i$, $Q_{ij} = -\alpha$ whenever $i \sim j$ and $Q_{ij} = 0$ otherwise, then the intrinsic CAR model [Besag, York and Mollié (1991)] has density

$$(6) \quad p(\boldsymbol{\phi} \mid \tau) \sim \tau^{N^*/2} \exp\left(-\frac{\tau}{2} \boldsymbol{\phi}' Q \boldsymbol{\phi}\right), \quad \text{with}$$

$$N^* = \begin{cases} N, & \text{if } \alpha \in (0, 1), \\ N - G, & \text{if } \alpha = 1. \end{cases}$$

In (6) τ is the spatial precision parameter, τQ is the precision matrix in this multivariate normal distribution and G is the number of “islands” (disconnected parts) in the spatial map [Hodges, Carlin and Fan (2003)]. When $\alpha \in (0, 1)$, (6) is a proper multivariate normal distribution. When $\alpha = 1$, Q is singular with $Q\mathbf{1} = \mathbf{0}$; Q has rank $N - G$ in a map with G islands, therefore, the exponent on τ becomes $(N - G)/2$. In hierarchical models, the CAR model is usually used as a prior on spatial random effects. For instance, let Y_i be the observed number of cases of a disease in region i , $i = 1, \dots, N$, and E_i be the expected number of cases in region i . Here the Y_i are treated as random variables, while the E_i are treated as fixed and known, often simply proportional to the number of persons at risk in region i . For rare diseases, a Poisson approximation to a binomial sampling distribution for disease counts is often used, so a commonly used likelihood for mapping a single disease is

$$(7) \quad Y_i \stackrel{\text{ind}}{\sim} \text{Poisson}(E_i e^{\mu_i}), \quad i = 1, \dots, N,$$

where $\mu_i = \mathbf{x}_i' \boldsymbol{\beta} + \phi_i$. The \mathbf{x}_i are explanatory, region-specific regressors with coefficients $\boldsymbol{\beta}$ and the parameter μ_i is the log-relative risk describing departures of

observed from expected counts, that is, from E_i . The hierarchy’s next level is specified by assigning the CAR distribution to ϕ and a hyper-prior to the spatial precision parameter τ . In the hierarchical setup, the improper ICAR with $\alpha = 1$ gives proper posterior distributions for spatial effects. In practice, Markov chain Monte Carlo (MCMC) algorithms are designed for estimating posteriors from such models and the appropriate number of linear constraints on the ϕ suffices to ensure sampling from proper posterior distributions [Banerjee, Carlin and Gelfand (2004), pages 163–164, give details].

2.1.3. *How does CAR fit into SANOVA?* To use CAR in SANOVA, the key is re-expressing the improper CAR, that is, (6) with $\alpha = 1$. Let Q have spectral decomposition $Q = VDV'$, where V is an orthogonal matrix with columns containing Q ’s eigenvectors and D is diagonal with nonnegative diagonal entries. D has G zero diagonal entries, one of which corresponds to the eigenvector $\frac{1}{\sqrt{N}}\mathbf{1}_N$, by convention the N th (right-most) column in V . Define a new parameter $\Theta = V'\phi$, so Θ has dimension N and precision matrix τD . Giving an N -vector Θ a normal prior with mean zero and precision τD is equivalent to giving $\phi = V\Theta$ a CAR prior with precision τQ . Θ consists of $\Theta_N = \frac{1}{\sqrt{N}}\mathbf{1}'_N\phi = \sqrt{N}\bar{\phi}$, the scaled average of the ϕ_i , along with $N - 1$ contrasts in ϕ , which are orthogonal to $\frac{1}{\sqrt{N}}\mathbf{1}_N$ by construction. Thus, the CAR prior is informative (has positive precision) only for contrasts in ϕ , while putting zero precision on $\Theta_{GM} = \Theta_N = \frac{1}{\sqrt{N}}\mathbf{1}'_N\phi$, the overall level, and on $G - 1$ orthogonal contrasts in the levels of the G islands. In other words, the CAR model can be thought of as a prior distribution on the contrasts rather than individual effects (hence the need for the sum-to-zero constraint). A related result, discussed in Besag, York and Mollié (1995), shows the CAR to be a member of a family of “pairwise difference” priors. This reparameterization allows the CAR model to fit into the ANOVA framework, with Θ_{GM} corresponding to the ANOVA’s grand mean and the rest of Θ , Θ_{Reg} , corresponding to $V^{(-)'}\phi$, where $V^{(-)}$ is V excluding the column $\frac{1}{\sqrt{N}}\mathbf{1}_N$, consisting of $N - 1$ orthogonal contrasts among the N regions and giving the $N - 1$ degrees of freedom in the usual ANOVA:

$$\begin{aligned} \phi &= [\phi_1, \phi_2, \dots, \phi_N]' \\ &= V\Theta \\ &= \begin{bmatrix} V^{(-)} & \frac{1}{\sqrt{N}}\mathbf{1}_N \end{bmatrix} \begin{bmatrix} \Theta_{Reg} \\ \Theta_{GM} \end{bmatrix}. \end{aligned}$$

Giving ϕ a CAR prior is equivalent to giving Θ a $N(\mathbf{0}, \tau D)$ prior; the latter are the “constraint cases” in HCSC’s SANOVA structure. The precision $D_{NN} = 0$ for the overall level is equivalent to a flat prior on Θ_{GM} , though Θ_{GM} could alternatively have a normal prior with mean zero and finite variance. If $G > 1$, the CAR prior also puts zero precision on $G - 1$ contrasts in ϕ , which are contrasts in the levels of the G islands [Hodges, Carlin and Fan (2003)].

2.2. *SANOVA as a competitor to MCAR.*

2.2.1. *Multivariate conditionally autoregressive (MCAR) models.* With multiple diseases, we have unknown ϕ_{ij} corresponding to region i and disease j , where $i = 1, \dots, N$ and $j = 1, \dots, n$. Letting Ω be a common precision matrix (i.e., inverse of the covariance matrix) representing correlations between the diseases in a given region, MCAR distributions arise through conditional specifications for $\phi_i = (\phi_{i1}, \dots, \phi_{in})'$:

$$(8) \quad \phi_i | \{\phi_{i'}\}_{i' \neq i} \sim \text{MVN} \left(\frac{\alpha}{m_i} \sum_{i' \sim i} \phi_{i'}, \frac{1}{m_i} \Omega^{-1} \right).$$

These conditional distributions yield a joint distribution for $\phi = (\phi'_1, \dots, \phi'_N)'$:

$$(9) \quad f(\phi | \Omega) \propto \|\Omega\|^{(N-G)/2} \exp \left(-\frac{1}{2} \phi' (Q \otimes \Omega) \phi \right),$$

where Q is defined as in Section 2.1.2 and again (9) is an improper density when $\alpha = 1$. However, as for the univariate CAR, this yields proper posteriors in conjunction with a proper likelihood. The specification above is a “separable” dispersion structure, that is, the covariances between the diseases are invariant across regions. This may seem restrictive, but relaxing this restriction gives even more complex dispersion structures [see Jin, Banerjee and Carlin (2007) and references therein]. As mentioned earlier, our focus is to retain the model’s simplicity without compromising the primary inferential goals. We propose to do this using SANOVA and will compare it with the separable MCAR only.

2.2.2. *SANOVA with Minnesota counties as one factor.* We now describe the SANOVA model using the Minnesota 3-cancer dataset. Consider the Minnesota map with $N = 87$ counties, and suppose each county has counts for $n = 3$ cancers. County i has an n -vector of parameters describing the n cancers, $\phi_i = (\phi_{i1}, \phi_{i2}, \dots, \phi_{in})'$; define the Nn vector ϕ as $\phi = (\phi'_1, \phi'_2, \dots, \phi'_N)'$. For now, we are vague about the specific interpretation of ϕ_{ij} ; the following description applies to any kind of data. Assume the $N \times N$ matrix Q describes neighbor pairs among counties as before. The SANOVA model for this problem is a 2-way ANOVA with factors cancer (“CA,” n levels) and county (“CO,” N levels) and no replication. As in Section 2.1.1, we model ϕ with a saturated linear model and put the grand mean and the main effects in their traditional positions as in ANOVA (matrix dimensions and definitions appear below the equation):

$$(10) \quad \begin{aligned} \phi &= [\phi'_1, \phi'_2, \dots, \phi'_N]' = [A_1 | A_2] \Theta \\ &= \left[\underbrace{\frac{1}{\sqrt{Nn}} \mathbf{1}_{Nn}}_{\substack{\text{Grand mean} \\ Nn \times 1}} \quad \underbrace{\frac{1}{\sqrt{N}} \mathbf{1}_N \otimes H_{CA}}_{\substack{\text{Cancer} \\ \text{main effect} \\ Nn \times (n-1)}} \right] \end{aligned}$$

$$\underbrace{V^{(-)} \otimes \frac{1}{\sqrt{n}} \mathbf{1}_n}_{\substack{\text{County} \\ \text{main effect} \\ Nn \times (N-1)}} \left[\underbrace{V^{(-)} \otimes H_{CA}^{(1)} \cdots V^{(-)} \otimes H_{CA}^{(n-1)}}_{\substack{\text{Cancer} \times \text{County} \\ \text{interaction} \\ Nn \times (N-1)(n-1)}} \right] \\ \times \begin{bmatrix} \Theta_{GM} \\ \Theta_{CA} \\ \Theta_{CO} \\ \Theta_{CO \times CA} \end{bmatrix},$$

where H_{CA} is an $n \times (n - 1)$ matrix whose columns are contrasts among cancers, so $\mathbf{1}'_n H_{CA} = \mathbf{0}'_{n-1}$, and $H'_{CA} H_{CA} = I_{n-1}$; $H_{CA}^{(j)}$ is the j th column of H_{CA} ; and $V^{(-)}$ is V without its N th column $\frac{1}{\sqrt{N}} \mathbf{1}_N$, so it has $N - 1$ columns, each a contrast among counties, that is, $\mathbf{1}'_N V^{(-)} = \mathbf{0}'_{N-1}$, and $V^{(-)'} V^{(-)} = I_{N-1}$. The column labeled “Grand mean” corresponds to the ANOVA’s grand mean and has parameter Θ_{GM} ; the other blocks of columns labeled as main effects and interactions correspond to the analogous ANOVA effects and to their respective parameters Θ_{CA} , Θ_{CO} , $\Theta_{CO \times CA}$. Defining prior distributions on Θ completes the SANOVA specification. We put independent flat priors (normal with large variance) on Θ_{GM} and Θ_{CA} , which are, therefore, not smoothed. This is equivalent to putting a flat prior on each of the n cancer-specific means. To specify the smoothing priors, define $H_{CA}^{(0)} = \frac{1}{\sqrt{n}} \mathbf{1}_n$. Let the county main effect parameter Θ_{CO} have prior $\Theta_{CO} \sim N_{N-1}(\mathbf{0}, \tau_0 D^{(-)})$, where $D^{(-)}$ corresponds to $V^{(-)}$, that is, D without its N th row and column, $\tau_0 > 0$ is unknown and $\tau_0 D^{(-)}$ is a precision matrix. Similarly, let the j th group of columns in the cancer-by-county interaction, $V^{(-)} \otimes H_{CA}^{(j)}$, have prior $\Theta_{CO \times CA}^{(j)} \sim N_{N-1}(\mathbf{0}, \tau_j D^{(-)})$, for $\tau_j > 0$ unknown. Each of the priors on Θ_{CO} and the $\Theta_{CO \times CA}^{(j)}$ is a CAR prior; the overall level of each CAR, with prior precision zero, has been included in the grand mean and cancer main effects.

To compare this to the MCAR model, use SANOVA’s priors on Θ to produce a marginal prior for ϕ comparable to the MCAR’s prior on ϕ (Section 2.2.1); in other words, integrate Θ_{CO} and $\Theta_{CO \times CA}$ out of the foregoing setup. A priori,

$$(11) \quad \left[K \begin{pmatrix} \Theta_{CO} \\ \Theta_{CO \times CA}^{(1)} \\ \vdots \\ \Theta_{CO \times CA}^{(n-1)} \end{pmatrix} \right]$$

has precision $Q \otimes (H_A^{(+)} \text{diag}(\tau_j) H_A^{(+)})'$, where K is the columns of the design matrix for the county main effects and cancer-by-county interactions—the right-most




FIG. 1. Comparing prior precision matrices for ϕ in MCAR and SANOVA.

$n(N - 1)$ columns in equation (10)’s design matrix—and $H_A^{(+)} = (\frac{1}{\sqrt{n}}\mathbf{1}_n | H_{CA})$ is an orthogonal matrix. Appendix A in Zhang, Hodges and Banerjee (2009) gives a proof.

2.2.3. *Comparing SANOVA vs MCAR.* Defining ϕ as in Sections 2.2.1 and 2.2.2, consider the MCAR prior for ϕ , with within-county precision matrix Ω . Let Ω have spectral decomposition $V_\Omega D_\Omega V_\Omega'$, where D_Ω is $n \times n$ diagonal and V_Ω is $n \times n$ orthogonal. Then the prior precision of ϕ is $Q \otimes (V_\Omega D_\Omega V_\Omega')$, where Q is the known neighbor relations matrix and V_Ω and D_Ω are unknown. Comparing MCAR to SANOVA, the prior precision matrices for the vector ϕ are as in Figure 1. SANOVA is clearly a special case of MCAR in which $H_A^{(+)}$ is known. Also, as described so far, $H_A^{(+)}$ has one column proportional to $\mathbf{1}_n$ with the other columns being contrasts, while MCAR avoids this restriction. MCAR is thus more flexible, while SANOVA is simpler, presumably making it better identified and easier to set priors for. MCAR should have its biggest advantage over SANOVA when the “true” V_Ω is not like $H_A^{(+)}$ for any specification of the smoothing precisions τ_j . However, because data sets often have modest information about higher-level variances, it may be that using the wrong $H_A^{(+)}$ usually has little effect on the analysis. In other words, SANOVA’s performance may be relatively stable despite having to specify $H_A^{(+)}$, while MCAR may be more sensitive to Ω ’s prior.

2.3. *Setting priors in MCAR and SANOVA.*

2.3.1. *Priors in SANOVA.* For the case of normal errors, based on equations (1) and (10), setting priors for Θ, τ_j, η_0 completes a Bayesian specification. Since τ and η_0 are precision parameters, one possible prior is Gamma; this paper uses a Gamma with mean 1 and variance 10. As mentioned, the grand mean and cancer main effects $\theta_1, \theta_2, \theta_3$ have flat priors with $\pi(\theta) \propto 1$, though they could have proper informative priors. The priors for θ_4, \dots are set according to the SANOVA structure as in Section 2.2.2. We ran chains drawing in the order θ, τ and η_0 [Appendix B in Zhang, Hodges and Banerjee (2009) gives details]. Hodges, Carlin and Fan (2007) also considered priors on the degrees of freedom in the fitted model, some conditioned so the degrees of freedom in the model’s fit were

fixed at a certain degree of smoothness. The present paper emphasizes comparing MCAR and SANOVA, so we do not consider such priors. For the case of Poisson errors, we use a normal prior with mean 0 and variance 10^6 for the grand mean and cancer main effects $\theta_1, \theta_2, \theta_3$. The other θ_i s are given normal CAR priors as discussed in Section 2.2.2. For the prior on the smoothing precisions τ_j , we use Gamma with mean 1 and variance 10. To reduce high posterior correlations among the θ s, we used a transformation during MCMC; Appendix C in Zhang, Hodges and Banerjee (2009) gives details.

2.3.2. *Priors in MCAR.* MCAR models were fitted in WinBUGS. For the normal-error case, we used this model and parameterization:

$$(12) \quad \begin{aligned} Y_{ij} &\sim N\left(\mu_{ij}, \frac{1}{\eta_0}\right), \\ \mu_{ij} &= \beta_j + S_{ij}, \end{aligned}$$

$i = 1, \dots, N; j = 1, \dots, n$, where η_0 has a gamma prior with mean 1 and variance 10 as for SANOVA. To satisfy WinBUGS’s constraint that $\sum_i S_{ij} = 0$, we add cancer-specific intercepts β_j . We give β_j a flat prior and for S , the spatial random effects, we use an intrinsic multivariate CAR prior. Similarly, in the Poisson case

$$(13) \quad \begin{aligned} Y_{ij} &\sim \text{Poisson}(\mu_{ij}), \\ \log(\mu_{ij}) &= \log(E_{ij}) + \beta_j + S_{ij}, \end{aligned}$$

where E_{ij} is an offset. Prior settings for β_j and S_{ij} are as in the normal case. For MCAR priors, the within-county precision matrix Ω needs a prior; a Wishart distribution is an obvious choice. If $\Omega \sim \text{Wishart}(R, \nu)$, then $E(\Omega) = \nu R^{-1}$. We want a “vague” Wishart prior; usually $\nu = n$ is used but little is known about how to specify R . Thus, we considered three different R s, each proportional to the identity matrix. One of these priors sets R ’s diagonal entries to $R_{ii} = 0.002$, close to the setting used in an example in the GeoBUGS manual (oral cavity cancer and lung cancer in West Yorkshire). The other two R s are the identity matrix and 200 times the identity. For the special case $n = 1$, where the Wishart reduces to a Gamma, these Wisharts are $\Gamma(0.5, 0.001)$, $\Gamma(0.5, 0.5)$ and $\Gamma(0.5, 100)$, respectively.

3. Simulation experiment. For this simulation experiment, artificial data were simulated from the model used in SANOVA with a spatial factor, as described in Section 2.2.2. Three different types of Bayesian analysis were applied to the simulated data: SANOVA with the same $H_A^{(+)}$ used to generate the simulated data (called “SANOVA correct”); SANOVA with incorrect $H_A^{(+)}$; and MCAR. SANOVA correct is a theoretical best possible analysis in that it takes as known things that MCAR estimates, that is, it uses additional correct information.

SANOVA correct cannot be used in practice, of course, because the true $H_A^{(+)}$ is not known. MCAR vs SANOVA with incorrect $H_A^{(+)}$ is the comparison relevant to practice, and comparing them to SANOVA correct shows how much each method pays for its “deficiency” relative to SANOVA correct. Obviously it is not enough to test the SANOVA model using only data generated from a similar SANOVA model. To avoid needless computing and facilitate comparisons, instead of generating data from an MCAR model and fitting a SANOVA model as specified above, we use a trick that is equivalent to this. Section 3.1.2 gives the details.

3.1. *Design of the simulation experiment.* We simulated both normally-distributed and Poisson-distributed data. For both types of data, we considered two different true sets of smoothing parameters $\mathbf{r} = \boldsymbol{\tau}/\eta_0$ or $\boldsymbol{\tau}$ (Table 1). For the normal data, we considered $\boldsymbol{\tau}/\eta_0$, since this ratio determines smoothing in normal models, and we also considered two error precisions η_0 (Table 1).

3.1.1. *Generating the simulated data sets.* To generate data from the SANOVA model, we need to define the true $H_A^{(+)}$. Let

$$HA_1 = \begin{pmatrix} 1 & -2 & 0 \\ 1 & 1 & -1 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{3}} & 0 & 0 \\ 0 & \frac{1}{\sqrt{6}} & 0 \\ 0 & 0 & \frac{1}{\sqrt{2}} \end{pmatrix}.$$

We used HA_1 as the correct $H_A^{(+)}$; its columns are scaled to have length 1. Given $V^{(-)}$ and with $H_A^{(+)}$ known, one draw of Θ and ϵ produces a draw of $X_D\Theta + \epsilon$, therefore a draw of \mathbf{y} . In the simulation, we let the grand mean and main effects, which are not smoothed, have true value 5. Each observation is simulated from a 3×20 factorial design, where 3 is the number of cancers and 20 is the number of counties. We used the 20 counties in the right lower corner of Minnesota’s map, with their actual neighbor relations. Thus, the dimension of each artificial data set is 60. The simulation experiment is a repeated-measures design, in which a “subject” s in the design is a draw of $(\boldsymbol{\delta}^{(s)}, \boldsymbol{\gamma}^{(s)})$, referring

TABLE 1
Experimental conditions in the simulation experiments

Error distribution	η_0	$(\tau_0/\eta_0, \tau_1/\eta_0, \tau_2/\eta_0) (\tau_0, \tau_1, \tau_2)$	Data name
Normal	1	(100, 100, 0.1)	Data1
	1	(0.1, 100, 0.1)	Data2
	10	(100, 100, 0.1)	Data3
	10	(0.1, 100, 0.1)	Data4
Poisson	NA	(100, 100, 0.1)	Data5
	NA	(0.1, 100, 0.1)	Data6

to equation (3), where $\delta_{1-3}^{(s)} = 5$ and $\delta_{4-60}^{(s)} \sim N_{57}(\mathbf{0}, I_3 \otimes D^{(-)})$ specify Θ and $\boldsymbol{y}^{(s)} \sim N_{60}(\mathbf{0}, I_{60})$ gives ϵ . For the normal-errors case, 100 such “subjects” were generated. Given a design cell in the simulation experiment with $\boldsymbol{\tau} = (a, b, c)$ and $\eta_0 = d$, the artificial data set for subject s is $\boldsymbol{y}^{(s)} = X_D \text{diag}(\mathbf{1}'_3, \frac{1}{\sqrt{a}}\mathbf{1}'_{19}, \frac{1}{\sqrt{b}}\mathbf{1}'_{19}, \frac{1}{\sqrt{c}}\mathbf{1}'_{19})\delta^{(s)} + \frac{1}{\sqrt{d}}\boldsymbol{y}^{(s)}$. All factors of the simulation experiment were applied to each of the 100 “subjects.” For the normally-distributed data, the simulation experiment had these factors: (a) the true $(\tau_0/\eta_0, \tau_1/\eta_0, \tau_2/\eta_0)$: (100, 100, 0.1) or (0.1, 100, 0.1); (b) the true error precision η_0 : 1 or 10; and (c) six statistical methods, described below in Section 3.1.2. Each design cell described in Table 1 thus had 100 simulated data sets. Similarly, for the Poisson-data experiment, another 100 “subjects” were generated, but now there is no $\boldsymbol{y}^{(s)}$. Thus, each “subject” s is a vector $\delta^{(s)}$, where $\delta^{(s)}$ is as described above. For the design cell with $\boldsymbol{\tau} = (100, 100, 0.1)$, the artificial data for subject s is $\boldsymbol{y}^{(s)} \sim \text{Poisson}(\boldsymbol{\mu}^{(s)})$, where $\log(\boldsymbol{\mu}^{(s)}) = \log(\mathbf{E}) + X_D \text{diag}(\mathbf{1}'_3, \frac{1}{10}\mathbf{1}'_{19}, \frac{1}{10}\mathbf{1}'_{19}, \frac{1}{\sqrt{0.1}}\mathbf{1}'_{19})\delta^{(s)}$. In the simulation experiment, we use “internal standardization” of the Minnesota 3-cancer data to supply the expected numbers of cancers E_{ij} . Among the 20 extracted counties, Hennepin county has the largest average population over 11 years, about 1.1 million; its cancer counts are 5294, 119 and 439 for lung, larynx and esophagus respectively. Faribault county has the smallest average population, 16,501, with cancer counts 110, 7 and 13 respectively. The E_{ij} have ranges 80 to 5275, 2 to 113 and 7 to 449 for lung, larynx and esophagus cancer respectively. For the Poisson data, the simulation experiment had these factors: (a) the true τ_0, τ_1, τ_2 : (100, 100, 0.1) or (0.1, 100, 0.1); and (b) six statistical methods described below in Section 3.1.2. Again, each of the two design cells in Table 1 had 100 simulated data sets.

3.1.2. *The six methods (procedures).* For each simulated data set, we did a Bayesian analysis for each of six different models described in Table 2. The six models are: SANOVA with the correct $H_A^{(+)}$, HA_1 ; SANOVA with a somewhat incorrect $H_A^{(+)}$, HA_2 given below; a variant SANOVA with a very incorrect $H_A^{(+)}$,

TABLE 2
The six statistical methods used in the simulation experiment

Procedure	Prior
SANOVA with correct $H_A^{(+)}$	$\eta_0, \tau_j \sim \Gamma(0.1, 0.1)$ for $j = 0, 1, 2$
SANOVA with incorrect $H_A^{(+)}$	$\eta_0, \tau_j \sim \Gamma(0.1, 0.1)$ for $j = 0, 1, 2$
Variant SANOVA with H_{AM}	$\eta_0, \tau_j \sim \Gamma(0.1, 0.1)$ for $j = 0, 1, 2$
MCAR	$\Omega \sim \text{Wishart}(R, 3), R = 0.002I_3$
MCAR	$\Omega \sim \text{Wishart}(R, 3), R = I_3$
MCAR	$\Omega \sim \text{Wishart}(R, 3), R = 200I_3$

H_{AM} given below; MCAR with $R_{ii} = 0.002$; MCAR with $R_{ii} = 1$; and MCAR with $R_{ii} = 200$ (see Section 2.3.2). HA_2 and H_{AM} are

$$HA_2 = \begin{pmatrix} 1 & 1 & 1 \\ 1 & -2 & 0 \\ 1 & 1 & -1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{3}} & 0 & 0 \\ 0 & \frac{1}{\sqrt{6}} & 0 \\ 0 & 0 & \frac{1}{\sqrt{2}} \end{pmatrix},$$

$$H_{AM} = \begin{pmatrix} 0.56 & -0.64 & -0.52 \\ -0.53 & -0.77 & 0.36 \\ -0.63 & 0.07 & -0.77 \end{pmatrix}.$$

The incorrect HA_2 has the same first column (grand mean) as the correct HA_1 , so it differs from the correct HA_1 , though less than it might. As noted above, we need to see how the SANOVA model performs for data generated from an MCAR model in which V_Ω from Figure 1 does not have a column proportional to $\mathbf{1}_n$. To do this without needless computing, we used a trick: we used the data generated from a SANOVA model with HA_1 and fit the variant SANOVA mentioned above, in which $H_A^{(+)}$ is replaced by the orthogonal matrix H_{AM} with no column proportional to $\mathbf{1}_n$, chosen to be very different from HA_1 . For normal errors (Data1 through Data4), this is precisely equivalent to fitting a SANOVA with $H_A^{(+)} = HA_1$ to data generated from an MCAR model with $V_\Omega = BHA_1$, for $B = HA_1 H_{AM}^{-1}$, that is,

$$(14) \quad V_\Omega = \begin{bmatrix} 0.43 & -0.74 & -0.52 \\ -0.13 & -0.63 & 0.77 \\ -0.89 & -0.26 & -0.37 \end{bmatrix}$$

(to 2 decimal places). For Poisson errors (Data5, Data6), the equivalence is no longer precise but the divergence of fitted SANOVA [using $H_A^{(+)} = H_{AM}$] and generated data [using $H_A^{(+)} = HA_1$] is quite similar. Finally, we considered three priors for MCAR because little is known about how to set this prior and we did not want to hobble MCAR with an ill-chosen prior. For the SANOVA and variant SANOVA analyses, we gave τ_j a $\Gamma(0.1, 0.1)$ prior with mean 1 and variance 10 for both the normal data and the Poisson data.

3.2. *Outcome measures.* To compare the six different methods for normal and Poisson data, we consider three criteria. The first is average mean squared error (AMSE). For each of the 60 $(X_D \Theta)_{ij}$, the mean squared error is defined as the average squared error over the 100 simulated data sets. AMSE for each design cell in the simulation experiment is defined as the average of mean squared error over the 60 $(X_D \Theta)_{ij}$. Thus, for the design cell labeled DataK in Table 1, define

$$(15) \quad \widehat{AMSE}_K = \frac{1}{L} \sum_{d=1}^L \sum_{i=1}^N \sum_{j=1}^n [(X_D \widehat{\Theta})_{ij}^d - (X_D \Theta)_{ij}^d]^2 / Nn,$$

where $L = 100, N = 20, n = 3, K = 1, \dots, 4$ for Normal, $K = 5, 6$ for Poisson, Θ is the true value and $\hat{\Theta}$ is the posterior median of Θ . For each design cell (K), the Monte Carlo standard error for AMSE is $(100)^{-0.5}$ times the standard deviation, across Data K 's 100 simulated data sets, of $\sum_{i=1}^N \sum_{j=1}^n [(X_D \hat{\Theta})_{ij}^d - (X_D \Theta)_{ij}^d]^2 / Nn$. The second criterion is the bias of $X_D \Theta$. For each of Data K 's 100 simulated data sets, first compute posterior medians of $(X_D \Theta)_{1,1}, \dots, (X_D \Theta)_{20,3}$, then average each of those posterior medians across the 100 simulated data sets. From this average, subtract the true $(X_D \Theta)_{ij}$ s to give the estimated bias for each of the 60 $(X_D \Theta)_{ij}$ s. MBIAS is defined as the 2.5th, 50th and 97.5th percentiles of the 60 estimated biases. More explicitly, for design cell Data K , MBIAS is

$$(16) \quad \widehat{MBIAS}_K = 2.5\text{th}, 50\text{th}, 97.5\text{th percentiles of} \\ \left(\frac{1}{L} \sum_{d=1}^L (X_D \hat{\Theta}^d - X_D \Theta^d) \right).$$

Finally, the coverage rate of Bayesian 95% equal-tailed posterior intervals, ‘‘PI rate,’’ is the average coverage rate for the 60 individual $(X_D \Theta)_{ij}$ s.

3.3. *Markov chain Monte Carlo specifics.* While the MCAR models were implemented in WinBUGS, our SANOVA implementations were coded in R and run on Unix. The different architectures do not permit a fair comparison between the run times of SANOVA and MCAR. However, the SANOVA models have lower computational complexity than the MCAR models: MCAR demands a spectral decomposition in every iteration, while SANOVA does not. For each of our models, we ran three parallel MCMC chains for 10,000 iterations. The CODA package in R (www.r-project.org) was used to diagnose convergence by monitoring mixing using Gelman–Rubin diagnostics and autocorrelations [e.g., Gelman et al. (2004), Section 11.6]. Sufficient mixing was seen within 500 iterations for the SANOVA models, while 200 iterations typically revealed the same for the MCAR models; we retained 8000×3 samples for the posterior analysis.

3.4. *Results.* Table 3 and Figures 2 and 3 show the simulation experiment’s results. Table 3 shows AMSE; for all methods and design cells, the standard Monte Carlo errors of AMSE are small, less than 0.07, 0.005 and 0.025 for Data1/Data2, Data3/Data4 and Data5/Data6 respectively. Figure 2 shows MBIAS, where the middle symbols represent the median bias and the line segments represent the 2.5th and 97.5th percentiles. Figure 3 shows coverage of the 95% posterior intervals. Denote SANOVA with the correct $H_A^{(+)}$ (HA_1) as ‘‘SANOVA correct,’’ SANOVA with HA_2 as ‘‘SANOVA incorrect,’’ the variant SANOVA with H_{AM} as ‘‘SANOVA variant,’’ MCAR with $R_{ij} = 0.002$ as ‘‘MCAR0.002’’ and so on.

TABLE 3
AMSE for simulated normal and Poisson data

Model	Normal-error model				Poisson-error model	
	Data1	Data2	Data3	Data4	Data5	Data6
SANOVA with HA_1	0.34	0.60	0.04	0.06	0.02	0.04
SANOVA with HA_2	0.47	0.84	0.05	0.07	0.02	0.14
SANOVA with HA_{AM}	0.48	0.74	0.05	0.06	0.03	0.11
MCAR with $R_{ii} = 0.002$	0.66	1.88	0.04	0.13	0.02	0.04
MCAR with $R_{ii} = 1$	0.36	0.84	0.04	0.06	0.02	0.06
MCAR with $R_{ii} = 200$	0.93	0.92	0.09	0.09	0.24	0.36

3.4.1. *As expected, SANOVA with correct $H_A^{(+)}$ performs best.* For normal data, SANOVA correct has the smallest AMSE for all true η_0 and τ (Table 3). The advantage is larger in Data1 and Data2 where the error precision η_0 is 1 than in Data3 and Data4 where η_0 is 10 (i.e., error variation is smaller). For Poisson data, SANOVA correct also has the smallest AMSE. Considering MBIAS (Figure 2), SANOVA correct has the narrowest MBIAS intervals for all cases. In Figure 3, the posterior coverage for SANOVA correct is nearly nominal. As expected, then, SANOVA correct performs best among the six methods.

3.4.2. *SANOVA with incorrect HA_2 and H_{AM} perform very well.* Table 3 shows that, for normal data, both SANOVA incorrect and SANOVA variant have smaller AMSEs than MCAR200 and MCAR0.002, and AMSEs at worst close to MCAR1's. For Poisson data, Table 3 shows that MCAR0.002 and MCAR1

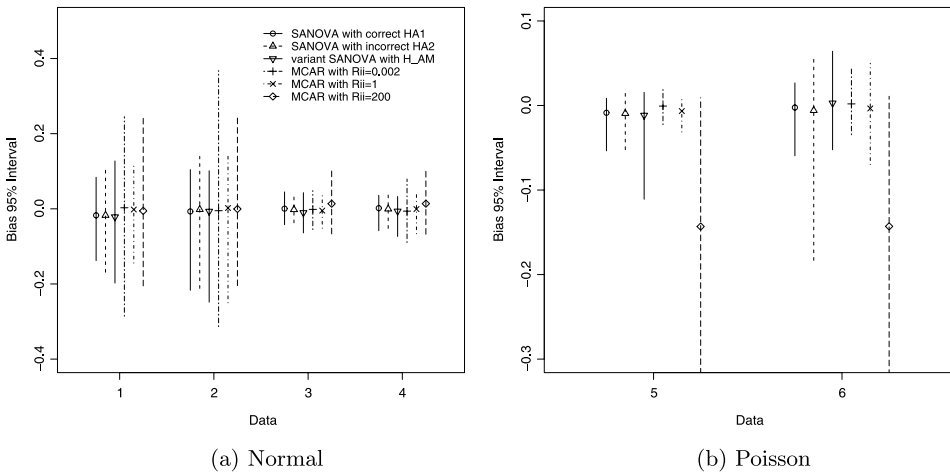


FIG. 2. MBIAS for simulated normal and Poisson data.

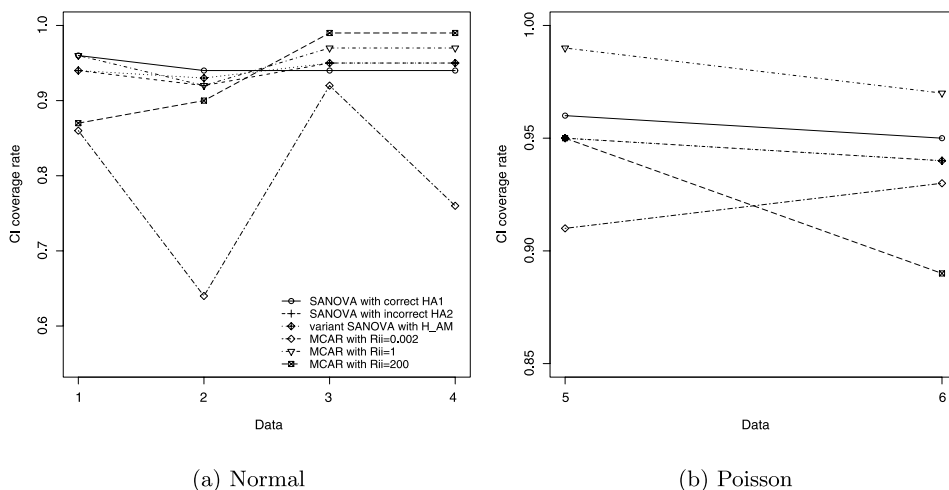


FIG. 3. PI rate for simulated normal and Poisson data.

do somewhat better than SANOVA incorrect and variant SANOVA. Considering MBIAS in normal data [Figure 2(a)], the width of the 95% MBIAS intervals for SANOVA incorrect are the same as or smaller than for all three MCAR procedures. Similarly, SANOVA variant has MBIAS intervals better than MCAR0.002 and MCAR200 and almost as good as MCAR1. Figure 2(b) for Poisson data shows SANOVA correct, MCAR0.002 and MCAR1 have similar MBIAS intervals. SANOVA variant in Data5 and SANOVA incorrect in Data6 show the worst performance for MBIAS apart from MCAR200, whose MBIAS interval is much the widest. Figure 3(a) shows that, for normal data, interval coverage for SANOVA incorrect and SANOVA variant is very close to nominal. It appears that the specific value of $H_A^{(+)}$ has little effect on PI coverage rate for the cases considered here. Apart from MCAR200 for Data1/Data2 and MCAR0.002 for Data1 through Data4, which show low coverage, all the other methods have coverage rates greater than 90% for normal data, most close to 95%. For Data3 and Data4, PI rates for MCAR200 reach above 99%. For Poisson data, the PI rates for SANOVA incorrect and SANOVA variant are close to nominal and better than MCAR0.002 and MCAR200. In particular, all SANOVAs have the closest to nominal coverage rates for both normal and Poisson data, which again shows the stability of SANOVA under different $H_A^{(+)}$ settings.

3.4.3. MCAR is sensitive to the prior on Ω . To fairly compare SANOVA and MCAR, we considered MCAR under three different prior settings. For normal data, MCAR1 has the smallest AMSEs and narrowest MBIAS intervals among the MCARs considered, while MCAR0.002 has the largest and widest, respectively. For Poisson data, however, MCAR0.002 has the best AMSE and MBIAS among

the MCARs. MCAR200 performs poorly for both Normal and Poisson. The coverage rates in Figure 3 show similar comparisons. These results imply that the prior matters for MCAR: no single prior was always best. By comparison, SANOVA seems more robust, at least for the cases considered.

3.5. Summary. As expected, SANOVA correct had the best performance because it uses more correct information. For normal data, SANOVA incorrect and SANOVA variant had similar AMSEs, better than two of the three MCARs for the data sets considered. For Poisson data, SANOVA incorrect and SANOVA variant had AMSEs as good as those of MCAR0.002 and MCAR1 for Data5 and somewhat worse for Data6, while showing nearly nominal coverage rates in all cases and less tendency to bias than MCAR in most cases. Replacing the $\Gamma(0.1, 0.1)$ prior for τ with $\Gamma(0.001, 0.001)$ left AMSE and MBIAS almost unchanged and coverage rates a bit worse (data not shown). MCAR, on the other hand, seems more sensitive to the prior on Ω . MCAR0.002 tends to smooth more than MCAR1, more so in normal models where the prior is more influential than in Poisson models. (The latter is true because data give more information about means than variances, and the Poisson model's error variance is the same as its mean, while the normal model's is not.) For the normal data, MCAR0.002's tendency to extra shrinkage appears to make it oversmooth and perform poorly for Data2 and Data4, where the truth is least smooth. For the Poisson data, MCAR0.002 and MCAR1 give results similar to each other and somewhat better than the SANOVAs except for interval coverage. Therefore, SANOVA, with stable results under different $H_A^{(+)}$ and with parameters that are easier to understand and interpret, may be a good competitor to MCAR in multivariate spatial smoothing.

4. Example: Minnesota 3-cancer data. Researchers in different fields have illustrated that accounting for spatial correlation could provide insights that would have been overlooked otherwise, while failure to account for spatial association could potentially lead to spurious and sometimes misleading results [see, e.g., Turechek and Madden (2002), Ramsay, Burnett and Krewski (2003), Lichstein et al. (2002)]. Among the widely investigated diseases are the different types of cancers. We applied SANOVA and MCAR to a cancer-surveillance data set describing total incidence counts of 3 cancers (lung, larynx, esophagus) in Minnesota's 87 counties for the years 1990 to 2000 inclusive. Minnesota's geography and history make it plausible that disease incidence would show spatial association. Three major North American land forms meet in Minnesota: the Canadian Shield to the north, the Great Plains to the west, and the eastern mixed forest to the southeast. Each of these regions is distinctive in both its terrain and its predominant economic activity: mining and outdoors tourism in the mountainous north, highly mechanized crop cultivation in the west, and dairy farming in the southeast. The different regions were also settled by somewhat different groups of in-migrants,

for example, disproportionately many Scandinavians in the north. These factors imply spatial association in occupational hazards as well as culture, weather, and access to health care especially in the thinly-populated north, which might be expected to produce spatial association in diseases. With multiple cancers one obvious option is to fit a separate univariate model for each cancer. But diseases may share the same spatially distributed risk factors, or the presence of one disease might encourage or inhibit the presence of another in a region, for example, larynx and esophagus cancer have been shown to be closely related spatially [Baron et al. (1993)]. Thus, we may need to account for dependence among the different cancers while maintaining spatial dependence between sites. Although the data set has counts broken out by age groups, for the present purpose we ignore age standardization and just consider total counts for each cancer. Age standardization would affect only the expected cancer counts E_{ij} , while other covariates could be added to either SANOVA or MCAR as unsmoothed fixed effects (i.e., in the A_1 design matrix). Given the population and disease count of each county, we estimated the expected disease count for each cancer in each county using the Poisson model. Denote the 87×3 counts as $y_{1,1}, \dots, y_{87,3}$; then the model is

$$(17) \quad \begin{aligned} y_{ij} | \mu_{ij} &\sim \text{Poisson}(\mu_{ij}), \\ \log(\mu_{ij}) &= \log(E_{ij}) + (X_D \Theta)_{ij}, \end{aligned}$$

where $X_D \Theta$ is the SANOVA structure and Θ has priors as in Section 2.2.2. For disease j in county i , $E_{ij} = P_i \frac{\sum_i O_{ij}}{\sum_i P_i}$, where O_{ij} is the disease count for county i and disease j and P_i is county i 's population. For the SANOVA design matrix, we consider HA_1 and HA_2 from the simulation experiment, though now neither is known to be correct. We also consider a variant SANOVA analysis using $H_A^{(+)}$ estimated from the MCAR1 model, to test the stability of the SANOVA results. Appendix D in Zhang, Hodges and Banerjee (2009) describes the latter analysis. Figures 4 to 6 show the data and results for MCAR1 and SANOVA with HA_1 . In each figure, the upper left plot shows the observed y_{ij}/E_{ij} ; the two lower plots show the posterior median of μ_{ij}/E_{ij} for MCAR1 and SANOVA with HA_1 . Lung cancer counts tended to be high and thus were not smoothed much by any method, while counts of the other cancers were much lower and thus smoothed considerably more (see also Figure 7). Since SANOVA with HA_1 , HA_2 and estimated $H_A^{(+)}$ gave very similar results, only those for HA_1 are shown. Results for MCAR0.002 are similar to those for MCAR1, so they are omitted. As expected, MCAR200 shows the least shrinkage among the three MCARs and gives some odd μ_{ij}/E_{ij} . To compare models, we calculated the Deviance Information Criterion [DIC; Spiegelhalter et al. (2002)]. To define DIC, define the deviance $D(\theta) = -2 \log f(\mathbf{y}|\theta) + 2 \log h(\mathbf{y})$, where θ is the parameter vector in the likelihood and $h(\mathbf{y})$ is a function of the data. Since h does not affect model comparison, we set $\log h(\mathbf{y})$ to 0. Let $\bar{\theta}$ be the posterior mean of θ and \bar{D} the posterior expectation of $D(\theta)$. Then define $p_D = \overline{D(\theta)} - D(\bar{\theta})$ to be a measure of model

complexity and define $DIC = \bar{D} + p_D$. Table 4 shows \bar{D} , p_D and DIC for nine analyses, SANOVA with 3 different $H_A^{(+)}$, MCAR with 3 different priors for Ω , and 3 fits of univariate CAR models to the individual diseases, discussed below. Considering \bar{D} , the three SANOVAs and MCAR1 are similar; Figures 4 to 6 show the fits are indeed similar. Figure 7 reinforces this point, showing that MCAR1 and SANOVA with HA_1 induce similar smoothing for the three cancers. SANOVA with $H_A^{(+)}$ estimated from MCAR has the smallest \bar{D} (1458), though its model complexity penalty ($p_D = 103$) is higher than MCAR0.002's ($p_D = 79$). Despite having the second worst fit (\bar{D}), MCAR0.002 has the best DIC, and the three SANOVAs have DICs much closer to MCAR0.002's than to the other MCARs. Generally, all SANOVA models have similar \bar{D} (≈ 1460) and DIC (≈ 1562), while MCAR results are sensitive to Ω 's prior, consistent with the simulation experiment. For comparison, we fit separate univariate CAR models to the three diseases considering three different priors for the smoothing precision, $\tau \sim \text{Gamma}(a, a)$ for $a = 0.001, 1$ and 1000 . For each prior, we added up \bar{D} , p_D and DIC for three diseases (see Table 4). With $a = 0.001$ and 1 , we obtained \bar{D} 's (1461 and 1453 respectively) competitive with SANOVA, MCAR0.002 and MCAR1 but with considerably greater complexity penalties (141 and 149 respectively) and thus DICs slightly larger than 1600. For $a = 1000$, we obtained an even lower \bar{D} (1432), but an increased penalty (180) resulted in a poorer DIC score. Figure 7 shows fitted values for CAR1, which were smoothed like MCAR1 and SANOVA for lung and esophagus cancers but smoothed rather more for larynx cancer. Overall, these results reflect some gain in performance from accounting for the space-cancer interactions/associations.

To further examine the smoothing under SANOVA, Figure 8 shows separate maps for the county main effect and interactions from the SANOVA fit with HA_1 . The upper left plot is the cancer main effect, the mean of the three cancers; the lower left plot is the comparison of lung versus average of larynx and esophagus; the lower right plot is the comparison of larynx versus esophagus. All values are on the same scale as y_{ij}/E_{ij} in Figures 4 to 6 and use the same legend. The two interaction contrasts are smoothed much more than the county main effect, agreeing with previous research that larynx and esophagus cancer are closely related spatially [Baron et al. (1993)]. To see whether the interactions are necessary, we fit a SANOVA model (using HA_1) without the interactions. As expected, model complexity decreased ($p_D = 77$), while \bar{D} increased slightly, so DIC became 1558, a bit better than SANOVA with interactions. Now consider the posterior of the MCAR's precision matrix Ω . The posterior mean of Ω is much larger for MCAR0.002 than MCAR200; the diagonal elements are larger by 4 to 5 orders of magnitude. This may explain the poor coverage for MCAR0.002 in the simulation. Further, consider the correlation matrix arising from the inverse of Ω 's posterior mean. As the diagonals of R change from 0.002 to 200, the correlation between any two cancers decreases and the complexity penalty p_D increases. By




FIG. 4. Lung cancer data and fitted values.

comparison, the three SANOVAs have similar model fits and complexity penalties, leading to similar DICs. So again, in this sense SANOVA shows greater stability.

5. Discussion and future work. We used SANOVA to do spatial smoothing and compared it with the much more complex MCAR model. For the cases considered here, we found SANOVA with spatial smoothing to be an excellent competitor to MCAR. It yielded essentially indistinguishable inference, while being easier to fit and interpret. In the SANOVA model, $H_A^{(+)}$ is assumed known. For most of the SANOVA models considered, $H_A^{(+)}$'s first column was fixed to




FIG. 5. Larynx cancer data and fitted values.

represent the average over diseases, while other columns were orthogonal to the first column. Alternatively, $H_A^{(+)}$ could be treated as unknown and estimated as part of the analysis. With this extension, SANOVA with spatial effects is a reparameterization of the MCAR model and gains the MCAR model’s flexibility at the price of increased complexity. This extension would be nontrivial, involving sampling from the space of orthogonal matrices while avoiding identification problems arising from, for example, permuting columns of $H_A^{(+)}$. Other covariates can be added to a spatial SANOVA. Although (10) is a saturated model, spatial smoothing “leaves room” for other covariates. Such models would suffer from collinearity of the CAR random effects and the fixed effects, as discussed by Reich, Hodges and




FIG. 6. *Esophagus cancer data and fitted values.*

Zadnik (2006), who gave a variant analysis that avoids the collinearity. For data sets with spatial and temporal aspects, for example, the 11 years in the Minnesota 3-cancer data, interest may lie in the counts' spatial pattern and in their changes over time. By adding a time effect, SANOVA can be extended to a spatiotemporal model. Besides spatial and temporal main effects, their interactions can also be included and smoothed. There are many modeling choices; the simplest model is an additive model without space-time interactions, where the spatial effect has a CAR model and the time effect a random walk, which is a simple CAR. But many other choices are possible. We have examined intrinsic CAR models, where $Q_{ij} = -1$ if region i and region j are connected. SANOVA with spatial smoothing could be extended to more general CAR models. Banerjee, Carlin and Gelfand




FIG. 7. Comparing data and fitted values for each cancer. The “Data” panel shows the density for y_{ij}/E_{ij} , while the other three panels show the posterior median of μ_{ij}/E_{ij} for univariate CAR, SANOVA and MCAR.

(2004) replaced Q with the matrix $D_w - \rho W$, where D_w is diagonal with the same diagonal as Q and $W_{ij} = 1$ if region i is connected with region j , otherwise $W_{ij} = 0$. Setting $\rho = 1$ gives the intrinsic CAR model considered in this paper. For known ρ , the SANOVA model described here is easily extended by replacing Q in Section 2 with $D_w - \rho W$. However, for unknown ρ , our method cannot be adjusted so easily, because updating ρ in the MCMC would force V and the design matrix to be updated as well, but this would change the definition of the parameter Θ . Therefore, a different approach is needed for unknown ρ . A different extension of SANOVA would be to survival models for areal spatial data [e.g., Li and Ryan (2002), Banerjee, Wall and Carlin (2003), Diva, Dey and Banerjee (2008)]. If the regions are considered strata, then random effects corresponding to nearby regions might be similar. In other words, we can embed the SANOVA

TABLE 4
Model comparison using DIC

Model	\bar{D}	p_D	DIC
SANOVA with HA_1	1461	103	1564
SANOVA with HA_2	1463	102	1565
SANOVA with H_A estimated from MCAR1	1458	103	1561
MCAR0.002	1476	79	1555
MCAR1	1459	132	1591
MCAR200	1559	356	1915
CAR0.001	1461	141	1602
CAR1	1453	149	1602
CAR1000	1432	180	1612

structure in a spatial frailty model. For example, the Cox model with SANOVA structure for subject j in stratum i is

$$(18) \quad h(t_{ij}, X_{ij}) = h_0(t_{ij}) \exp(X_{ij}\beta),$$

where X is the design matrix, which may include a spatial effect, a temporal effect, their interactions and other covariates. Banerjee, Carlin and Gelfand (2004) noted that in the CAR model, considering both spatial and nonspatial frailties, the frailties are identified only because of the prior, so the choice of priors for precisions is very important. Besides the above extensions, HCSC introduced tools for normal SANOVA models that can be extended to nonnormal SANOVA models. For example, HCSC defined the degrees of freedom in a fitted model as a function of the smoothing precisions. This can be used as a measure of the fit’s complexity, or a prior can be placed on the degrees of freedom as a way of inducing a prior on the unknowns in the variance structure. The latter is under development and will be presented soon.

Acknowledgments. The authors thank the referees, the Associate Editor and the Editor for valuable comments and suggestions.

SUPPLEMENTARY MATERIAL

Appendices, data and code (DOI: [10.1214/09-AOAS267SUPP](https://doi.org/10.1214/09-AOAS267SUPP); .zip). Our supplementary material includes four sections as appendices. In Appendix A we present a derivation of the precision matrix of (11). Details of our MCMC algorithms can be found in Appendix B. Appendix C discusses the mean transformation for the Poisson case, while Appendix D discusses the estimation of the $H_A^{(+)}$ from MCAR1. In addition, we provide a compressed folder containing the data set for our 3-cancer Minnesota example as well as an R code example to implement the SANOVA models.




FIG. 8. *SANOVA with HA_1* : (a) county main effect; (b) cancer \times county interaction 1 for larynx; (c) cancer \times county interaction 2 for esophagus.

REFERENCES

- BANERJEE, S., CARLIN, B. P. and GELFAND, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall/CRC Press, Boca Raton, FL.
- BANERJEE, S., WALL, M. M. and CARLIN, B. P. (2003). Frailty modeling for spatially correlated survival data, with application to infant mortality in Minnesota. *Biostatistics* **4** 123–142.
- BARON, A. E., FRANCESCHI, S., BARRA, S., TALAMINI, R. and LA VECCHIA, C. (1993). Comparison of the joint effect of alcohol and smoking on the risk of cancer across sites in the upper aerodigestive tract. *Cancer Epidemiol. Biomarkers Prev.* **2** 519–523.
- BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. Roy. Statist. Soc. Ser. B* **36** 192–236. MR0373208

- BESAG, J., GREEN, P., HIGDON, D. and Mengersen, K. (1995). Bayesian computation and stochastic systems (with discussion). *Statist. Science* **10** 3–66. [MR1349818](#)
- BESAG, J., YORK, J. C. and MOLLIE, A. (1991). Bayesian image restoration, with two applications in spatial statistics (with discussion). *Ann. Inst. of Statist. Math.* **43** 1–59. [MR1105822](#)
- CARLIN, B. P. and BANERJEE, S. (2003). Hierarchical multivariate CAR models for spatio-temporally correlated survival data. In *Bayesian Statistics 7* (J. M. Bernardo et al., eds.) 45–64. Oxford Univ. Press, Oxford. [MR2003166](#)
- ELLIOTT, P., WAKEFIELD, J. C., BEST, N. G. and BRIGGS, D. J. (2000). *Spatial Epidemiology: Methods and Applications*. Oxford Univ. Press, Oxford.
- GELFAND, A. E. and VOUNATSOU, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* **4** 11–25.
- GELMAN, A., CARLIN, J. B., STERN, H. S. and RUBIN, D. B. (2004). *Bayesian Data Analysis*, 2nd ed. Chapman and Hall/CRC Press, Boca Raton, FL.
- HODGES, J. S. (1998). Some algebra and geometry for hierarchical models. *J. Roy. Statist. Soc. Ser. B* **60** 497–536. [MR1625954](#)
- HODGES, J. S., CARLIN, B. P. and FAN, Q. (2003). On the precision of the conditionally autoregressive prior in spatial models. *Biometrics* **59** 317–322. [MR1987398](#)
- HODGES, J. S., CUI, Y., SARGENT, D. J. and CARLIN, B. P. (2007). Smoothing balanced single-error-term analysis of variance. *Technometrics* **49** 12–25. [MR2345448](#)
- JIN, X., BANERJEE, S. and CARLIN, B. P. (2007). Order-free coregionalized lattice models with application to multiple disease mapping. *J. Roy. Statist. Soc. Ser. B* **69** 817–838. [MR2368572](#)
- JIN, X., CARLIN, B. P. and BANERJEE, S. (2005). Generalized hierarchical multivariate CAR models for areal data. *Biometrics* **61** 950–961. [MR2216188](#)
- KIM, H., SUN, D. and TSUTAKAWA, R. K. (2001). A bivariate Bayes method for improving the estimates of mortality rates with a twofold conditional autoregressive model. *J. Amer. Statist. Assoc.* **96** 1506–1521. [MR1946594](#)
- LAWSON, A. B., BIGGERI, A. B., BOHNING, D., LESAFFRE, E., VIEL, J. F. and BERTOLLINI, R. (1999). *Disease Mapping and Risk Assessment for Public Health*. Wiley, New York.
- LEE, Y. and NELDER, J. A. (1996). Hierarchical generalized linear models (with discussion). *J. Roy. Statist. Soc. Ser. B* **58** 619–673. [MR1410182](#)
- LI, Y. and RYAN, L. (2002). Modeling spatial survival data using semi-parametric frailty models. *Biometrics* **58** 287–297. [MR1908168](#)
- LICHSTEIN, J. W., SIMONS, T. R., SHRINER, S. A. and FRANZREB, K. E. (2002). Spatial autocorrelation and autoregressive models in ecology. *Ecological Monographs* **72** 445–463.
- MARDIA, K. V. (1988). Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *J. Multivar. Anal.* **24** 265–284. [MR0926357](#)
- RAMSAY, T., BURNETT, R. and KREWSKI, D. (2003). Exploring bias in a generalized additive model for spatial air pollution data. *Environ. Health Perspect.* **111** 1283–1288.
- REICH, B. J., HODGES, J. S. and ZADNIK, V. (2006). Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models. *Biometrics* **62** 1197–1206. [MR2307445](#)
- RUE, H. and HELD, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC, Boca Raton. [MR2130347](#)
- SAIN, S. R. and CRESSIE, N. (2002). Multivariate lattice models for spatial environmental data. In *Proceedings of ASA Section on Statistics and the Environment* 2820–2825. Amer. Statist. Assoc., Alexandria, VA.
- SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit (with discussion). *J. Roy. Statist. Soc. Ser. B* **64** 583–639.
- TURECHEK, W. W. and MADDEN, L. V. (2002). A generalized linear modeling approach for characterizing disease incidence in spatial hierarchy. *Phytopathology* **93** 458–466.
- WAKEFIELD, J. (2007). Disease mapping and spatial regression with count data. *Biostatistics* **8** 158–183.

- WALL, M. M. (2004). A close look at the spatial structure implied by the CAR and SAR models. *J. Statist. Plann. Inference* **121** 311–324. [MR2038824](#)
- WALLER, L. A. and GOTWAY, C. A. (2004). *Applied Spatial Statistics for Public Health Data*. Wiley, New York. [MR2075123](#)
- ZHANG, Y., HODGES, J. S. and BANERJEE, S. (2009). Supplement to “Smoothed ANOVA with spatial effects as a competitor to MCAR in multivariate spatial smoothing.” [DOI:10.1214/09-AOAS267SUPP](#).

Y. ZHANG
NOVARTIS PHARMACEUTICALS
EAST HANOVER, NEW JERSEY 07936
USA

J. S. HODGES
DIVISION OF BIOSTATISTICS
SCHOOL OF PUBLIC HEALTH
UNIVERSITY OF MINNESOTA
MINNEAPOLIS, MINNESOTA 55455
USA

S. BANERJEE
DIVISION OF BIOSTATISTICS
SCHOOL OF PUBLIC HEALTH
UNIVERSITY OF MINNESOTA
MINNEAPOLIS, MINNESOTA 55455
USA
E-MAIL: sudiptop@biostat.umn.edu