

Spatial Disease Mapping Using Directed Acyclic Graph Auto-Regressive (DAGAR) Models

Abhirup Datta^{*¶}, Sudipto Banerjee^{†||}, James S. Hodges[‡], and Leiwen Gao[§]

Abstract. Hierarchical models for regionally aggregated disease incidence data commonly involve region specific latent random effects that are modeled jointly as having a multivariate Gaussian distribution. The covariance or precision matrix incorporates the spatial dependence between the regions. Common choices for the precision matrix include the widely used ICAR model, which is singular, and its nonsingular extension which lacks interpretability. We propose a new parametric model for the precision matrix based on a directed acyclic graph (DAG) representation of the spatial dependence. Our model guarantees positive definiteness and, hence, in addition to being a valid prior for regional spatially correlated random effects, can also directly model the outcome from dependent data like images and networks. Theoretical results establish a link between the parameters in our model and the variance and covariances of the random effects. Simulation studies demonstrate that the improved interpretability of our model reaps benefits in terms of accurately recovering the latent spatial random effects as well as for inference on the spatial covariance parameters. Under modest spatial correlation, our model far outperforms the CAR models, while the performances are similar when the spatial correlation is strong. We also assess sensitivity to the choice of the ordering in the DAG construction using theoretical and empirical results which testify to the robustness of our model. We also present a large-scale public health application demonstrating the competitive performance of the model.

Keywords: areal data, Bayesian inference, directed acyclic graphs, disease mapping, spatial autoregression.

1 Introduction

Epidemiological data for disease rates are often presented as aggregated disease counts over entire geographical regions like states or counties. Such *areal* or *areally-referenced* data are ubiquitous in public health applications. Accurate identification of trends and factors associated with the disease requires accounting for the spatial dependence among the regions. A common approach to analyze areal datasets envisions the geographic domain as an undirected graph with the regions constituting the vertices and an edge between two vertices if the corresponding regions share a geographical border. This creates well defined neighbors for each region which are used to specify the joint or conditional

^{*}Johns Hopkins University, abhidatta@jhu.edu

[†]University of California Los Angeles, sudipto@ucla.edu

[‡]University of Minnesota, hodge003@umn.edu

[§]University of California Los Angeles, gaoleiwen@ucla.edu

[¶]Supported in part by NSF grant DMS-1915803.

^{||}Supported in part by NSF grants DMS 1916349, IIS 1562303, and NIEHS R01ES027027.

distributions of region-specific latent Gaussian random effects in a hierarchical setup. For example, the popular conditional autoregressive (CAR) model (Besag 1974; Clayton and Bernardinelli 1992) incorporates the underlying neighborhood structure in specifying the full conditional distribution for each observation. If w_i denotes the random effect representing the i th region for $i = 1, \dots, k$ and $i \sim j$ indicates that regions i and j are neighbors, then the CAR model specifies the full conditional distributions

$$w_i \mid w_{-i} \sim N \left(\sum_{j \sim i} w_j / n_i, \tau_w n_i \right), \quad (1.1)$$

where w_{-i} denotes the vector of observations leaving out the i th one, n_i denotes the number of neighbors for the i th region and throughout the text we adopt the convention that $N(\alpha, \Delta)$ denotes normal distribution with mean α and precision Δ , both in univariate and multivariate contexts. Hence, in (1.1) above, $\tau_w n_i$ is the conditional precision of $w_i \mid w_{-i}$.

The joint distribution of $w = (w_1, \dots, w_k)^T$ can be derived from (1.1) as $w \sim N(0, \tau_w(D - A))$ where $A = (a_{ij})$ is the adjacency matrix of the neighborhood graph, i.e., $a_{ij} = 1$ if and only if $i \sim j$, and D is a diagonal matrix with n_1, \dots, n_k on the diagonal. As $D - A$ is singular, this construction yields an improper joint distribution of the w_i 's, referred to as the intrinsic or improper CAR (ICAR) model. This impropriety renders the model ineligible for directly modeling the response or for generating data, although both can proceed by using contrasts as demonstrated in Besag and Kooperberg (1995). Also, the distribution can still be used as a prior for latent spatial random effects w and the posterior of w usually remains valid.

The impropriety of the ICAR model can be rectified by generalizing the full conditional mean to $E(w_i \mid w_{-i}) = \rho \sum_{j \sim i} w_j / n_i$ yielding the joint distribution $w \sim N(0, \tau_w(D - \rho A))$ which is proper for a certain range of ρ . Although introduction of ρ imparts more flexibility than the parameter-free improper analogue, it is difficult to interpret ρ as even very high values of ρ induce only modest spatial correlation among the observations (see Banerjee et al. 2014, for a discussion on this). Furthermore, Wall (2004) shows that even negative values of ρ may lead to positive correlation among neighboring regions. Assuncao and Krainski (2009) found that these oddities are a general feature of CAR models.

The second popular approach is the simultaneous autoregressive (SAR) model (Whittle 1954) which proceeds by simultaneously modeling the random effects as

$$w_i = \sum_{j \neq i} b_{ij} w_j + \epsilon_i \text{ for } i = 1, 2, \dots, k \quad (1.2)$$

where $\epsilon_i \stackrel{\text{ind}}{\sim} N(0, \tau_i)$ are errors independent of w . Defining $B = (b_{ij})$ and F to be a diagonal matrix with entries τ_1, \dots, τ_k , the set of equations in (1.2) yields the joint distribution $w \sim N(0, (I - B)F(I - B)^T)$. However, the common choice of defining $b_{ij} = \rho I(i \sim j) / n_i$, where $I(\cdot)$ denotes the indicator function, leads to similar problems with respect to interpretation of the parameter ρ (Wall 2004).

Beyond these two approaches, the inventory of covariance models for areal datasets is very limited. Leroux et al. (2000) and MacNab and Dean (2000) extended the CAR model by accommodating over-dispersion alongside spatial information. They proposed using the precision matrix $\lambda\tau_w(D - A) + (1 - \lambda)\tau_w I$, where $\lambda \in [0, 1]$ controls the degree of dependence among the regions. For a regular graph where all vertices have same number of neighbors d , $D = dI$. In this case, $\lambda\tau_w(D - A) + (1 - \lambda)\tau_w I$ can be rewritten as $\frac{1+(d-1)\lambda}{d}\tau_w(D - \rho^*A)$ where $\rho^* = \frac{d\lambda}{1+(d-1)\lambda}$. Thus, if the numbers of neighbors for the vertices do not vary greatly, this approach is somewhat similar to the proper CAR model and is encumbered by the same interpretability concerns.

For lattice based applications, there is a richer class of parametric intrinsic autoregression models (Besag and Kooperberg 1995; Besag and Higdon 1999). However, all such intrinsic models rely heavily on the lattice structure and cannot be used directly for arbitrary graphs. Applications to irregular areal data can proceed by breaking up the region into a fine lattice, using the intrinsic model on the lattice and aggregating over each area. Besag and Mondal (2005) demonstrated that certain classes of intrinsic autoregressive models can be interpreted as an average of a fine scale Gaussian Process over the entire domain. In disease mapping contexts, where the data are often observed over fixed politically delineated regions, such a latent fine scale spatial process may be difficult to interpret. In this manuscript, we only focus on models that can be formulated directly on the areal units.

We propose a new way of constructing precision matrices for areal models using a directed acyclic graph derived from the original undirected graph. Directed acyclic graphs or *DAGs* have been used in the spatial literature for modeling large spatial datasets (Datta et al. 2016) and for generating image textures (Cressie and Davidson 1998). Instead of modeling the precision matrix directly, we model its Cholesky factor, which for any multivariate Gaussian distribution is determined by the conditional distributions of the w_i 's. We specify these conditional distributions using autoregressive covariance models on a sequence of local trees created from this directed acyclic graph. The resulting Cholesky factor and the precision matrix are sparse. We refer to this model as the directed acyclic graph autoregressive or DAGAR model.

Unlike the ICAR model, our model's covariance matrix is guaranteed to be positive definite. This opens up a new avenue to generate or directly model multivariate Gaussian data with dependence structure derived from a graph. Common examples of such data, besides aggregated regional data, include images or social network data. We establish, both theoretically and empirically, that our model endows ρ with a clear interpretation as a spatial autocorrelation parameter, which, in fact, resolves an important conundrum in the conditional and simultaneous autoregressive models (Wall 2004). Also, the Cholesky factor has the same level of sparsity as the undirected graph ensuring scalability for analyzing very large areal datasets.

Cholesky factors inherit the dependence of directed acyclic graphs on ordering of the regions, thereby making our model order-dependent. As spatial regions generally do not have any natural ordering, to understand the impact of ordering, we propose a novel order free model by averaging over all $k!$ possible orderings. We show that the resulting precision matrix, which is order-free, can be evaluated in closed form and we use it

to present some theoretical results suggesting that the DAGAR precision matrix with a reasonably chosen ordering is often similar to the order-free matrix. The theoretical results complemented by simulation exercises reveal that the choice of ordering does not significantly affect the results. Simulation experiments also show that when the spatial correlation is weak or moderate, the DAGAR model outperforms CAR models in their ability to correctly estimate a latent spatial surface while the performances are similar for data with stronger spatial dependence.

2 Model

2.1 Cholesky Factors

We first review a general approach to modeling Gaussian covariance matrices using sparse Cholesky factors and discuss how this relates to CAR models and general covariance estimation. This helps motivate the subsequent construction of our model in Section 2.2. We assume that the graph of the regions is connected. Disconnected graphs with multiple islands will entail a simple extension with block diagonal covariance structures, where each block represents an island. Let $\mathcal{G} = (V, E)$ denote the connected graph with the regions as vertices V and edges E between neighbors. We denote the i th region simply by i and let $A = (a_{ij})$ denote the adjacency matrix for this undirected graph. To model Cholesky factors we specify distributions of the w_i 's as

$$w_1 = \epsilon_1, w_2 = b_{21}w_1 + \epsilon_2, \dots, w_k = b_{k1}w_1 + \dots + b_{k,k-1}w_{k-1} + \epsilon_k, \quad (2.1)$$

where the ϵ_i 's are independent $N(0, \tau_i)$ errors. Throughout this section and Section 2.2, we assume τ_w , the scale parameter for the w_i 's, is one. $B = (b_{ij})$ in (2.1) is a strictly lower triangular matrix. Let F be the diagonal matrix with τ_1, \dots, τ_k on the diagonal. Then $w \sim N(0, L^T F L)$ where $L = I - B$. Switching from the specification in (1.2) to a strictly lower triangular matrix B is not restrictive because of the following result.

Theorem 1. *Let $w \sim N(0, Q)$ where Q is the (possibly singular) precision matrix. Then there exists a permutation matrix P , a strictly lower triangular matrix B and a diagonal matrix F with non-negative entries such that $PQP^\top = (I - B)^\top F(I - B)$.*

If Q is non-singular, the result is familiar. For rank deficient Q we resort to the algorithm for obtaining the Cholesky factor. All proofs are presented in the Supplement (Datta et al. 2019). Hence, any multivariate normal distribution can be expressed as in (2.1) under certain orderings of the areal units. For low rank distributions this will be equivalent to setting some of the τ_i 's to zero. In fact, switching to the lower triangular B has several advantages. First, L is lower triangular with ones on the diagonal, guaranteeing that $L^T F L$ is positive definite as long as all τ_i 's are positive. Also $\det(L^T F L)$ is simply $\prod_{i=1}^n \tau_i$ and the quadratic form $w^T L^T F L w$ can be expressed as $\tau_1 w_1^2 + \sum_{i=2}^k \tau_i (w_i - \sum_{\{j < i\}} w_j b_{ij})^2$, evaluating which requires $\mathcal{O}(k + s)$ floating point operations (FLOPs) where s is the sparsity, i.e., the number of non-zero entries of B . Hence, if B is sparse, the joint density of w can be evaluated in an extremely scalable manner.

To complete the specification in (2.1), we need to fully specify the matrices B and F . The parameters $\{b_{ij}\}$ and $\{\tau_i\}$ are identifiable up to a marginal precision parameter because the factorization $(I - B)^\top F(I - B)$ is the so called “LDL^T factorization” (a variant of Cholesky decomposition with ones on the diagonal of L), which is unique. If multiple observations have been made for each region, we can estimate B and F without imposing simplifying assumptions. In fact, since a sparse representation of the Cholesky factor $I - B$ is desired, the problem reduces to high-dimensional covariance or precision matrix estimation for which there exists a vast inventory of statistical methods including banding (Wu and Pourahmadi 2003; Bickel and Levina 2008b), tapering (Cai et al. 2010), thresholding (Bickel and Levina 2008a; El Karoui 2008; Rothman et al. 2009) and penalization (Meinshausen and Bühlmann 2006; Friedman et al. 2007; Xue et al. 2012) among others.

On the other hand, for large point-referenced spatial datasets, Datta et al. (2016); Finley et al. (2017); Datta et al. (2016) construct sparse Cholesky factor approximations of the precision matrix from a Matérn covariance function (Stein 1999). These approximations are hence derived from an original joint distribution of the spatial random effects.

However, most areal datasets lack replication that would permit use of fully data-driven learning methods to estimate B and F . Also, there is no analog of Matérn covariance functions on arbitrary graphs from which one can derive sparse Cholesky factors. In fact, our goal here is the opposite, that is to construct a multivariate Gaussian distribution on graphs starting from the sparse Cholesky factor. Consequently, we will make parametric assumptions that will lead to an interpretable covariance model.

2.2 Directed Acyclic Graph Autoregressive Models

To achieve sparsity, we adopt the strategy of defining neighbor sets $N(i)$ such that $b_{ij} = 0$ for all $j \notin N(i)$. The choice and size of the neighbor sets for areal datasets can be predicated upon the underlying neighborhood graph \mathcal{G} . For $i > 1$, we define $N(i) = \{j < i, j \sim i\}$. The constraint $j < i$ is necessary to endow B with a lower triangular structure. This reduces (2.1) to

$$w_1 = \epsilon_1, \quad w_i = \sum_{j \in N(i)} w_j b_{ij} + \epsilon_i, \quad (i = 2, \dots, k) \quad (2.2)$$

This specification is analogous to auto-regressive models for time series. In fact, if w_i denotes the response at time i , $N(i)$ includes all time points less than i up to a lag of r , and $b_{ij} = b_{i-j}$, then (2.2) simply denotes the autoregressive model of order r . In a time-series context, where i and j denote time points, assigning the weights b_{ij} based on the temporal lag seems natural, but for irregular areal datasets, enumeration of the areal units does not have any physical interpretation. In the context of image texture analysis, Cressie and Davidson (1998) used different coefficients for w_j in (2.2) based on the direction of neighbors on a regular lattice, to generate images with a wide range of textures. In general, vertices of irregular graphs based on areal datasets do not share such commonality in terms of spatial orientation of their neighbors. Hence, it is

intuitive to assign equal weights to all the neighbors, i.e., letting $b_{ij} = b_i$ for all $j \in N(i)$. A natural choice would be to let $b_i = 1/n_{<i}$ and $\tau_i \propto n_{<i}$ where $n_{<i} = |N(i)|$ denotes the cardinality of the neighbor set. This specification is similar to (1.1) except that we are only using the directed neighbors $N(i)$ instead of all neighbors. Since $n_{<1} = 0$, this choice of b_{ij} leads to the conundrum of how to specify τ_1 . Either we need to define τ_1 in a manner inconsistent with the definition of τ_i for $i > 1$ or we define $\tau_1 = 0$ which yields an improper distribution for w . We circumvent this using a more general specification described below that includes the degenerate prior with $b_i = 1/n_{<i}$ and $\tau_i \propto n_{<i}$ as a limiting case.

Let d_{ij} denote the length of the shortest path on \mathcal{G} between nodes i and j . If \mathcal{G} is a tree, i.e., an acyclic graph, then for any $0 \leq \rho < 1$, the matrix with elements $\rho^{d_{ij}}$ is positive definite and can be used to model the covariance of w . This extends the autoregressive model of lag 1 (AR(1) model) for time series to any tree graph (Basseville et al. 2006). However, graphs corresponding to areal datasets are rarely acyclic and for loopy graphs such results generally do not hold. A spanning tree of a graph is a subgraph that is a tree and includes all the vertices of the original graph. Spanning trees have been used to iteratively approximate parametric covariance matrices over loopy graphs (Sudderth 2002). Borrowing these ideas, a potential solution would be to define the covariance matrix for w as the AR(1) covariance matrix on a spanning tree of \mathcal{G} . However, for large graphs, strategies for deciding upon the best spanning tree are unclear and computationally expensive. Furthermore, as demonstrated in (Sudderth 2002), ignoring certain edges when pruning \mathcal{G} to a spanning tree can lead to large errors. Instead, we will use local spanning tree embeddings of small subgraphs of \mathcal{G} to construct the lower dimensional conditional densities specified in (2.2). This method will not ignore any edge and yet produces a computationally convenient precision matrix.

Let \mathcal{G}_i be the subgraph of \mathcal{G} comprising vertices $\{i\} \cup N(i)$ and the edges among them. We intend to construct the conditional density $w_i | w_{N(i)}$ using an embedded spanning tree \mathcal{T}_i of \mathcal{G}_i . The natural candidate for \mathcal{T}_i is the tree graph $(\{i\} \cup N(i), \{i \sim j | j \in N(i)\})$ as it contains all edges between i and $N(i)$. We specify the conditional density $w_i | w_{N(i)}$ using the AR(1) model on \mathcal{T}_i with parameter ρ . To be precise, for any $0 \leq \rho < 1$, an auto-regressive AR(1) covariance matrix with parameter ρ on \mathcal{T}_i is given by

$$\begin{pmatrix} 1 & \rho & \rho & \cdots & \rho \\ \rho & 1 & \rho^2 & \cdots & \rho^2 \\ \rho & \rho^2 & 1 & \cdots & \rho^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho & \rho^2 & \cdots & \rho^2 & 1 \end{pmatrix} = \begin{pmatrix} 1 & v_i^\top \\ v_i & \Sigma_i \end{pmatrix}. \quad (2.3)$$

This helps us define $E(w_i | w_{N(i)}) = v_i^T \Sigma_i^{-1} w_{N(i)}$ and $\text{var}(w_i | w_{N(i)}) = 1 - v_i^T \Sigma_i^{-1} v_i$, where v_i is the $n_{<i} \times 1$ vector of covariances between w_i and $w_{N(i)}$, and Σ_i is the $n_{<i} \times n_{<i}$ covariance matrix of $w_{N(i)}$ assuming an AR(1) model on \mathcal{T}_i . From equation (2.3) it is clear that $v_i = \rho \mathbf{1}$ where $\mathbf{1}$ denotes the vector of ones, and Σ_i is the matrix

with one on the diagonals and ρ^2 on the off-diagonals. Equating this with (2.2), we have

$$b_{ij} = \frac{\rho}{1 + (n_{<i} - 1)\rho^2} \quad (i = 2, \dots, k; j \in N(i)), \quad \tau_i = \frac{1 + (n_{<i} - 1)\rho^2}{1 - \rho^2} \quad (i = 1, \dots, k) \tag{2.4}$$

The specifications in (2.4) reveal some desirable intuitive features. First, as discussed earlier, $b_{ij} = b_i$ for all $j \in N(i)$, i.e, we assign equal weights to all the directed neighbors. Also, the conditional precision τ_i for w_i increases with the number of directed neighbors. The formulation of \mathcal{T}_i also ensures that any edge between i and j is incorporated in the conditional specification of w_i or w_j depending on which comes later in the ordering. So, unlike approximating the entire graph with a spanning tree, the local spanning tree approach ensures that no edge of \mathcal{G} is ignored. Furthermore, for any $0 \leq \rho < 1$ all τ_i 's are positive thereby ensuring a proper probability distribution $w \sim N(0, L^T FL)$. The limiting case of $\rho = 1$ is equivalent to the improper prior with $b_i = 1/n_{<i}$ and $\tau_i \propto n_{<i}$.

The constructions in (2.1) and (2.4) assume a specific ordering, which we now generalize to any other ordering. Let $\pi = \{\pi(1), \dots, \pi(n)\}$ be any predetermined ordering of the regions and π^{-1} denote its corresponding inverse permutation. Under this ordering, for any $i \neq \pi(1)$, we define its past observations $w_{<i,\pi}$ as the collection $\{w_j \mid \pi^{-1}(j) < \pi^{-1}(i)\}$ and its set of directed neighbors $N_\pi(i) = \{j \mid i \sim j, \pi^{-1}(j) < \pi^{-1}(i)\}$. Let E_π denote the collection of directed edges from all members of $N_\pi(i)$ to i for every $i \neq \pi(1)$. We now have a directed acyclic graph $\mathcal{D}_\pi = (V, E_\pi)$. Let $n_{\pi(i)} = |N_\pi(i)|$. The generalization of (2.4) based on \mathcal{D}_π is:

$$w_i \mid w_{<i,\pi} \sim N \left(\frac{\rho}{1 + (n_{\pi(i)} - 1)\rho^2} \sum_{j \in N_\pi(i)} w_j, \frac{1 + (n_{\pi(i)} - 1)\rho^2}{1 - \rho^2} \right), \tag{2.5}$$

where for any i that has no directed neighbors under π , $n_{\pi(i)} = 0$ and the conditional mean in (2.5) is zero. If $w_\pi = (w_{\pi(1)}, \dots, w_{\pi(k)})^T$, and L_π and F_π denote the analogous matrices corresponding to this ordering π , then we have

$$w_\pi \sim N(0, L_\pi^T F_\pi L_\pi) \tag{2.6}$$

This completes the specification of a new class of covariance models for areal datasets. Since the construction is predicated upon a directed acyclic graph derived from an original graph \mathcal{G} , we refer to this as the directed acyclic graph autoregressive or DAGAR model. Since L_π is lower triangular with $e = |E_\pi| = |\mathcal{G}|$ non-zero sub-diagonal entries and F_π is diagonal, for any ρ , the determinant of $cov(w_\pi)$ is simply $(1 - \rho^2)^k / \prod_{i=1}^k (1 + (n_{\pi(i)} - 1)\rho^2)$ and the likelihood for the model in (2.6) can be evaluated using $\mathcal{O}(k + e)$ FLOPs. This ensures that our model is scalable and can be used to analyze massive areal datasets.

2.3 Interpretation of ρ

While scalability of the DAGAR model is an important aspect in the analysis of large spatial datasets, our current emphasis is on offering a class of areal models with an interpretable parameterization. In this regard, we resolve the issue of a lack of meaningful

relationship between ρ and spatial correlation in the proper CAR models. We now offer insight about the interpretability of ρ in the DAGAR model for certain special graphs.

Theorem 2. *Let \mathcal{T} denote a tree with vertices $V = \{1, \dots, k\}$ and π denote any ordering such that for any $i \neq \pi(1)$, $n_\pi(i) = 1$. Then the covariance matrix in (2.6) defines the autoregressive Gaussian process on \mathcal{T} , i.e., $(L_\pi^T F_\pi L_\pi)^{-1} = (\rho^{d_{ij}})$ where d_{ij} denotes the shortest path on \mathcal{T} between i and j .*

Theorem 2 shows that if \mathcal{G} is acyclic, then, under certain orderings including breadth-first and pre-order tree traversals, this model is equivalent to the stationary AR(1) model on trees with unit variances and ρ being exactly equal to the correlation between neighboring areas. In fact, for any two vertices separated by a distance d , the correlation is ρ^d . The result also shows that while we require an ordering of the locations to construct the DAGAR model, the resulting matrix in this case is order-free and stationary and simply a function of the underlying undirected graph. We now present a result on interpretation of ρ for an $m \times n$ regular grid graph.

Theorem 3. *Let \mathcal{G} denote the $m \times n$ grid graph with vertices $V = \{(i, j) \mid i = 1, \dots, m; j = 1, \dots, n\}$ and neighbors to the north, south, east, and west, and let π denote any diagonal ordering of the vertices corresponding to non-decreasing or non-increasing order of $i + j$ or $i - j$, then $\text{var}(w_{(i,j)}) = 1$ for all i and j , and for any neighboring pair of vertices (i, j) and (i', j') , $\text{cov}(w_{(i,j)}, w_{(i',j')}) = \rho$.*

Hence, although for a grid, the DAGAR precision matrix is a function of the ordering, for all orderings of Theorem 3, the model yields unit variances and a correlation of ρ for all neighbor pairs. Hence, ρ is still interpretable as the correlation between neighboring units. This result for a grid graph is quite promising as graphs arising from areal data, like the grid graph, are loopy. Also note that when a CAR model is specified as $w \sim N(0, \tau_w(D - \rho A))$, it may seem that $1/\tau_w$ is the marginal variance of each spatial random effect. Unfortunately, this is not true as the specification of the precision matrix for the CAR model effectuates a heteroskedastic distribution. Consequently, τ_w can only be interpreted as a common scale factor for the marginal variances. This is remedied in the DAGAR specification, as we see from Theorems 2 and 3 that the resulting model $N(0, \tau_w Q(\rho))$, where $Q(\rho)$ is the DAGAR precision matrix, is homoskedastic, and hence $1/\tau_w$ is the marginal variance. So, the DAGAR model ensures interpretability for both ρ and τ_w . We will see that this interpretability empowers the DAGAR model to deliver significantly improved inference about the spatial parameters in areal data analysis.

Of course, it is difficult to generalize these theoretical interpretability results for irregular graphs. Hence, we also conducted numerical experiments to corroborate the results in Theorems 2 and 3, and also gain insight into the relationship between ρ and the neighbor-pair correlations for the proper CAR model and the DAGAR model using an irregular graph. So we used three different graphs: a simple path graph with 100 vertices which is analogous to a time-series, a two-dimensional 10×10 lattice or grid graph with edges between vertically or horizontally adjacent vertices, and the state map of the contiguous United States, where two states are said to have an edge if they share a common geographical boundary.

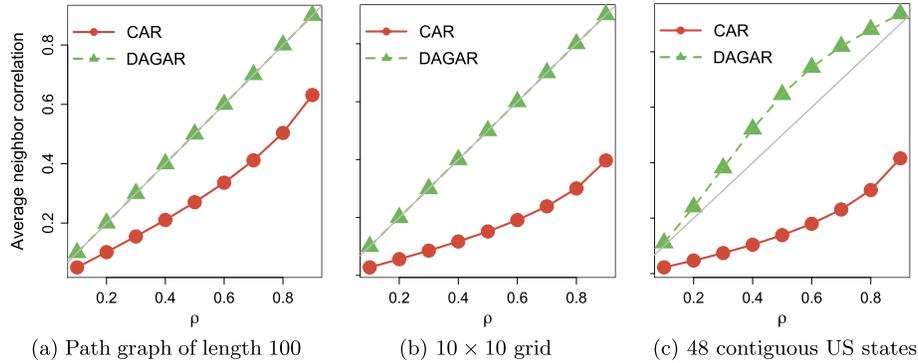


Figure 1: Average neighbor pair correlations as a function of ρ for proper CAR and DAGAR model. The solid gray line represents $x = y$ line.

We generated covariance matrices corresponding to the two models for $\rho \in \{i/10 \mid i = 1, \dots, 9\}$. Figure 1 plots the average neighbor-pair correlation, given by $c(\rho) = \sum_{i \sim j} cov(w_i, w_j) / (2\sqrt{var(w_i)}\sqrt{var(w_j)}) / (\sum n_i)$ as a function of ρ , for proper CAR and DAGAR models. For the path and grid graphs, we find the average neighbor pair correlation $c(\rho)$ for our model is exactly ρ as guaranteed in Theorems 2 and 3. For the highly irregular United States graph, the theoretical results, of course, do not hold. Nevertheless, $c(\rho)$ for the DAGAR model is much closer to ρ than for the proper CAR model. For the CAR model, even when ρ is close to one, $c(\rho)$ is less than 0.4. In fact, for all three graphs, the average neighbor-pair correlation for the proper CAR model remains modest. This seems to be true even for very high values of ρ and is consistent with findings elsewhere (see, e.g., Banerjee et al. 2014).

2.4 Impact of Ordering

Unlike covariance or precision matrices that remain invariant up to a permutation factor under different orderings of the multivariate vector, Cholesky factors depend on the ordering of the observations. Our model in (2.6) assumes a predetermined ordering π . We have already seen that for tree and grid graphs, Theorems 2 and 3 guarantee that the under many different orderings, the DAGAR model retains desirable properties like homoskedasticity and neighbor-pair correlation of ρ .

To understand the impact of ordering beyond the variances and neighbor-pair correlations, we consider an order-free model using a product-of-experts type construction (Hinton 2002). Let P_π denote the permutation matrix corresponding to π , i.e., $P_\pi(x_1, \dots, x_k)^T = (x_{\pi(1)}, \dots, x_{\pi(k)})^T$ for any k -dimensional vector x , and let Q denote the average of the DAGAR precision matrices in (2.6) over all permutations π , i.e.,

$$Q(\rho) = \frac{1}{k!} \sum_{\pi} P_\pi^T L_\pi^T F_\pi L_\pi P_\pi . \tag{2.7}$$

It is clear that $Q = Q(\rho)$ is free of any ordering and is only a function of the undirected graph \mathcal{G} . Also, since it is the average of positive definite matrices, it is also positive definite. We will use the order-free model Q to understand how the DAGAR precision matrices Q_π differ under different choices of the ordering π . In order to do this, we first note that Q can be expressed in closed form.

Theorem 4. *Let $i \approx j$ mean that i and j share at least one common neighbor. There exist functions $f(\rho, r)$ and $g(\rho, r)$ for any positive integer r and $0 \leq \rho < 1$ such that*

$$Q_{ii} = 1 + \frac{n_i \rho^2}{2(1 - \rho^2)} + \frac{\rho^2}{1 - \rho^2} \sum_{j \sim i} f(\rho, n_j)$$

$$Q_{ij} = -\frac{\rho}{1 - \rho^2} I(i \sim j) + \frac{1}{1 - \rho^2} I(i \approx j) \sum_{k \sim N(i) \cap N(j)} g(\rho, n_k).$$

Here $I(\cdot)$ denotes the indicator function. Explicit expressions for $f(\rho, r)$ and $g(\rho, r)$ are provided in the proof of Theorem 4. Let \mathcal{K} denote a set of ‘reasonable’ orderings that one can consider for a given areal dataset. We note that for two orderings π_1 and π_2 in \mathcal{K} , the relative difference $\|Q_{\pi_1}(\rho) - Q_{\pi_2}(\rho)\|_F / \|Q(\rho)\|_F$, where $\|\cdot\|_F$ denotes the Frobenius norm, is bounded above by $2 \max_{\pi \in \mathcal{K}} \|Q_\pi(\rho) - Q(\rho)\|_F / \|Q(\rho)\|_F$. We now investigate the asymptotic behavior of this quantity for the path and grid graphs.

Theorem 5. *Consider the path graph with k nodes, and let π denote the left to right or right to left ordering. Then the relative difference*

$$\lim_{k \rightarrow \infty} \frac{\|Q_\pi(\rho) - Q(\rho)\|_F}{\|Q(\rho)\|_F} = \sqrt{\frac{4\rho^8 + 2\rho^4}{(3 + 6\rho^2 + \rho^4)^2 + 18\rho^2(1 + \rho^2)^2 + 2\rho^4}}. \tag{2.8}$$

Theorem 5 quantifies asymptotically the relative difference between the DAGAR model and the order free version. Figure 2 (a) plots the quantity on the right hand side of (2.8) as a function of $\rho \in [0, 1]$. We see that it is a monotonically increasing function of ρ . For small values of ρ the difference is extremely small (0.02 for $\rho = 0.25$) and even for moderate values of ρ (0.5 – 0.75) the difference is around or less than 10%. Below we also provide the analogous result for the two-dimensional grid graph.

Theorem 6. *Consider a $m \times m$ grid graph and let π denote any of the orderings used in Theorem 3. Then*

$$\lim_{m \rightarrow \infty} \frac{\|Q_\pi(\rho) - Q(\rho)\|_F}{\|Q(\rho)\|_F} = \sqrt{\frac{\rho^4 \left(\frac{s(\rho)}{5} - \frac{2}{1+\rho^2}\right)^2 + 2\left(\frac{1}{3} - \frac{s(\rho)}{30} - \frac{\rho^2}{1+\rho^2}\right)^2 + 12\left(\frac{1}{6} - \frac{s(\rho)}{60}\right)^2}{(1 + \rho^2 + \rho^2 \frac{s(\rho)}{5})^2 + 4\rho^2 + 20\left(\frac{1}{6} - \frac{s(\rho)}{60}\right)^2}} \tag{2.9}$$

where $s(\rho) = \sum_{r=1}^4 \frac{r}{1+(r-1)\rho^2}$.

The limit in Theorem 6 looks more complicated than the analogous result for the path graph. However, things are simplified noting that $1 - s(\rho)/10$ is $O(\rho^2)$ and, hence, so is the numerator in the right hand side of (2.9). Figure 2 (b) plots this quantity as

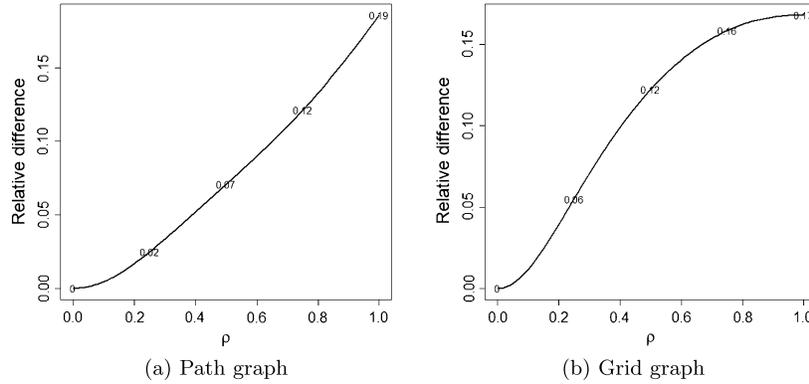


Figure 2: Asymptotic relative difference between the DAGAR model and the order free DAGAR model in terms of Frobenius norm for path graph (left) and grid graph (right). The five numbers on each curve corresponds to the values of the difference at $\rho = 0, 0.25, 0.5, 0.75$ and 1 respectively.

a function of ρ . We see that once again this is monotonic in ρ and the difference is negligible for small ρ . Theorems 5 and 6 show that at least for small ρ , the impact of ordering is negligible, though for larger ρ theoretically the DAGAR precision matrices for different orderings will be somewhat different.

While these results are restricted to the case of special graphs, for an arbitrary areal dataset, one approach would be to use simple intuitive orderings based on the coordinates representing the nodes in some Euclidean embedding of \mathcal{G} . Similar strategies have been used in Cholesky factor based approaches in Datta et al. (2016), Stein et al. (2004) and Vecchia (1988) who observed empirically that the joint distribution seemed to be less sensitive to ordering of the regions. Our own set of simulations, detailed in Section 3.2 will confirm these findings as we observe that results corresponding to different orderings are similar for both regular and irregular graphs.

The order-free model, owing to the availability of the precision matrix in closed form, can be deemed a worthy candidate for analyzing areal data given its liberation from a synthetic ordering. Our simulation analyses detailed in Section 3.1, show that for a wide range of scenarios, performance of the DAGAR model and its order-free version were very similar. However, the order-free model has certain disadvantages. From Theorem 4, it is clear that for $i \neq j$, $Q_{ij} \neq 0$ if and only if either $i \sim j$ or $i \approx j$. Hence, the sparsity of Q is e_2 where e_2 is the number of edges in the second order graph created from \mathcal{G} . As $e_2 > e$, Q is less sparse than the precision matrices for the original DAGAR model in Section 2.2 or the CAR models. Furthermore, unlike the DAGAR precision matrix Q_π , Q does not have a closed form expression for the determinant, which invokes expensive computations. These computational roadblocks limit the possibility of using the order-free model for larger datasets. The results of Theorems 2 and 3 about interpretability of the parameter ρ also do not carry over to the order-free model.

3 Data Analysis

3.1 Data Generated Using an Exponential Gaussian Process

Models for areal datasets are typically used as priors for areal random effects in a hierarchical setup. For example, let y_i denote the response observed at region i and x_i denote the corresponding set of covariates. A spatial generalized linear mixed model framework assumes $h(E(y_i)) = x_i^T \beta + w_i$ where $h(\cdot)$ denotes a suitable link function. Subsequently, a hierarchical areal model is customarily specified as

$$\prod_{i=1}^k pr(y_i | x_i^T \beta + w_i, \theta) \times N(w | 0, \tau_w Q(\rho)) \times pr(\beta, \tau_w, \rho, \theta) \quad (3.1)$$

where $Q(\rho)$ denotes the precision matrix of the areal model, $pr(y_i | x_i^T \beta + w_i, \theta)$ denotes the density corresponding to the link $h(\cdot)$ and $pr(\beta, \tau_w, \rho, \theta)$ is the prior for the parameters. If $h(\cdot)$ is the identity link, e.g., the responses are Gaussian, then we can exploit conjugacy for generating w in a sampler. However, for non-Gaussian responses, we have to sample w and the other parameters using a Metropolis random walk sampler from the joint density in (3.1).

We conducted simulation experiments assessing the performance of the areal data models using the three graph structures described in Section 2.3 — path, grid and US states. For each graph, we embed the vertices on the Euclidean plane and generate the spatial random effect vector w from an Gaussian process, i.e., $w \sim N(0, \tau_w M)$ where M^{-1} is the covariance matrix corresponding to an exponential (Matérn_{1/2}) GP, i.e., $M^{-1} = \exp(-\phi d(i, j))$ with $d(i, j)$ denoting the distance between the embedding of the i^{th} and j^{th} vertex. The path graph has a distance preserving embedding in the Euclidean plane such that $D(i, j) = |i - j|$. We embed the grid graph within a 10×10 grid in the Euclidean plane. Although the resulting distance matrix is not identical to the shortest distance (or geodesic) matrix on the graph, the distance between each neighbor-pair remains one. For the United States graph, we use the centroid of each state to create the distance matrix. We scale the distance matrix so that the average neighbor-pair distance is one. To generate w , we use $\tau_w = 0.25$, $\tau_e = 2.5$ and $\phi = -\log(j/10)$ for $j = 1, \dots, 9$. This implies that for the exponential GP the average neighbor pair correlation $\rho = \exp(-\phi)$ varies between 0.1 and 0.9, thereby covering a wide spectrum of scenarios. Subsequently, we generate the response y comprising independent $y_i = x_i^T \beta + w_i + \epsilon_i$, where x_i is a 2×1 vector comprising two independent standard normal variables, $\beta = (1, 5)^T$, and the ϵ_i 's are independent $N(0, \tau_e)$.

We fitted all models using the hierarchical setup in (3.1) with the six different choices of $Q(\rho)$ — 1) ICAR 2) proper CAR, 3) a scaled ICAR model (as suggested by one reviewer) of Sørbye and Rue (2014) which specifies a prior for τ_w in the ICAR model as $\tau_w \sim \text{Gamma}(2, \sigma_{ref}^2)$ where

$$\sigma_{ref}^2 = \exp \left(\frac{1}{k} \sum_{i=1}^k \log(((D - A)^+))_{ii} \right),$$

$(D - A)^+$ denoting the Moore-Penrose inverse of $D - A$, 4) the dimension-reducing sparse generalized linear mixed model (sparse GLMM or SGLMM, Hughes and Haran 2013), and the 5) ordered (DAGAR) and 6) order-free (DAGAR_{OF}) directed acyclic graph

autoregressive models. To create the directed acyclic graph for our model, we used the ordering π based on the sum of the co-ordinates of the mappings of the vertices. The sparse GLMM was implemented directly using the npSpatial R-package with the method of restricted spatial regression (Hughes and Cui 2018). For the remaining models, the proper CAR, ICAR, scaled ICAR and two DAGAR, the models were fit using Jags and rjags packages. The hierarchical model fit for ICAR is

$$\prod_{i=1}^k N(y_i | x_i^T \beta + w_i, \tau_e) \times N(w | 0, \tau_w(D - A)) \times N(\beta | 0, 0.0001I) \times \text{Gamma}(\tau_w | 2, 1) \times \text{Gamma}(\tau_e | 2, 0.1) . \quad (3.2)$$

The scaled ICAR model was specified same as in (3.2) with replacing the $\text{Gamma}(\tau_w | 2, 1)$ prior with $\text{Gamma}(\tau_w | 2, \sigma_{ref}^2)$. For proper CAR the model fitted is same as in 3.2 but with $D - A$ replaced by $D - \rho A$ and a $\text{Uniform}(0, 1)$ prior for ρ . The ordered and order-free DAGAR model replaced $D - A$ in (3.2) respectively with $Q_\pi(\rho)$ and $Q(\rho)$ along with inclusion of the $\text{Uniform}(0, 1)$ prior for ρ . For each combination of parameter values we used 100 replicate datasets.

Figure 3 plots the mean square error (MSE) between the true and estimated w averaged over 100 replicates for each scenario. We first observe that the mean square errors for the ICAR-based models (ICAR, scaled ICAR and sparse GLMM) are significantly higher than the other three models for all three graphs. The scaled ICAR was the best among these three, producing substantially lower MSEs than the original ICAR and the sparse GLMM. The sparse GLMM also produced lower MSEs than ICAR for path and USA graphs, but had slightly higher MSE for the grid graph. However, the three models involving the additional ρ parameters, i.e., the proper CAR and the two DAGAR models, consistently produced lower MSEs. For the path graph there is no significant difference in MSE among these three models. However, for the grid and USA graphs when ρ is small, the DAGAR models yielded substantially lower errors. We also noticed that, in terms of MSE, there was very little difference between the performance of the ordered DAGAR model (2.6) and its order-free analogue (2.7) for most of the scenarios. This result is encouraging as the scalability of the former is a pragmatic solution for analyzing very large areal datasets or networks.

Next, we consider estimation of ρ , as ensuring interpretability of ρ is the motivation driving the construction of the DAGAR model. As the data was generated using an exponential GP, ρ is the unit-distance spatial correlation. We look at the estimates and credible bands of ρ in Figure 4 for the three models involving ρ (the ICAR-based models do not involve ρ and, hence, are not included). We see that for all three graphs, estimates for ρ from the proper CAR model are considerably higher than the true value. The bias is especially stark when the true ρ is small. The DAGAR models generally perform much better in this respect with much less estimation bias, particularly for small ρ . For larger ρ , the order-free model performs slightly better with the ordered model demonstrating some downwards bias. The 95% confidence bands for both the DAGAR models cover the true value of ρ in most scenarios for all three graphs, while the bands for the proper CAR clearly miss many of the true ρ values.

Finally, in Figure 5 we plot the coverage probabilities (CP) defined as the mean coverage for a parameter by the 95% confidence intervals over the 100 replications. The

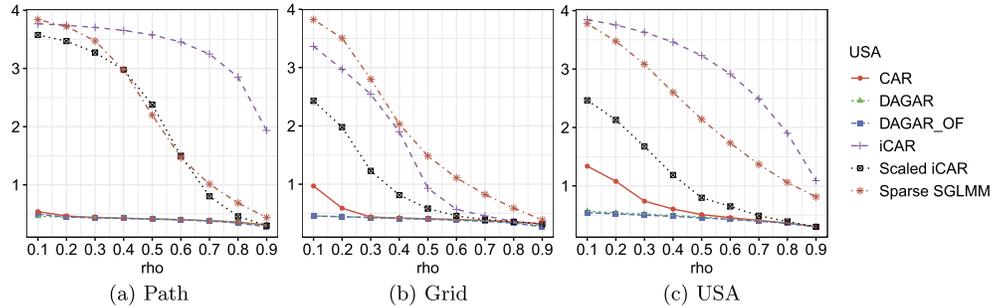


Figure 3: MSE as a function of the true ρ (x-axis) for the simulation data analysis using data generated from an exponential GP.

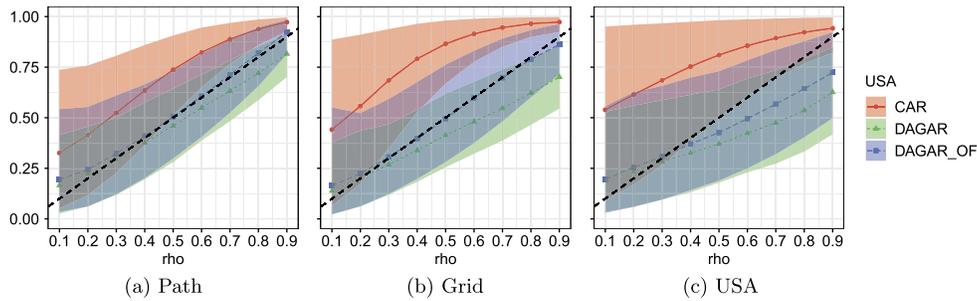


Figure 4: Estimate and confidence bands of ρ as a function of the true ρ (x-axis) for the simulation data analysis using data generated from an exponential GP.

regression coefficients β_1 , β_2 and the error variance $\sigma_e^2 = 1/\tau_e$ are common to all the models, and hence we compare the CPs for all six models for these parameters. For ρ , we only compare the three models using ρ . We do not compare $\sigma_w^2 = 1/\tau_w$, as it has different interpretation for different models. For example, it is the homoskedastic spatial variance for the ordered DAGAR model, but simply a scale factor for the heteroskedastic spatial variances in the CAR models. We see that for ρ , the two DAGAR models offer significantly improved coverage over the proper CAR model. This trend was already reflected in the confidence bands in Figure 4 and is confirmed here. For smaller ρ the coverages of the two DAGAR models are nearly identical to the nominal level of 95%. While the coverages decline for larger values of ρ , they are still almost uniformly and substantially better than the coverage provided by the proper CAR model.

For the regression coefficients, β_1 and β_2 , we see that all models except the sparse GLMM offered satisfactory coverage, close to 95%. This is not surprising as estimates of regression coefficients are typically robust to variance misspecification. The under coverage of the sparse GLMM is also expected as it tries to adjust for spatial confound-

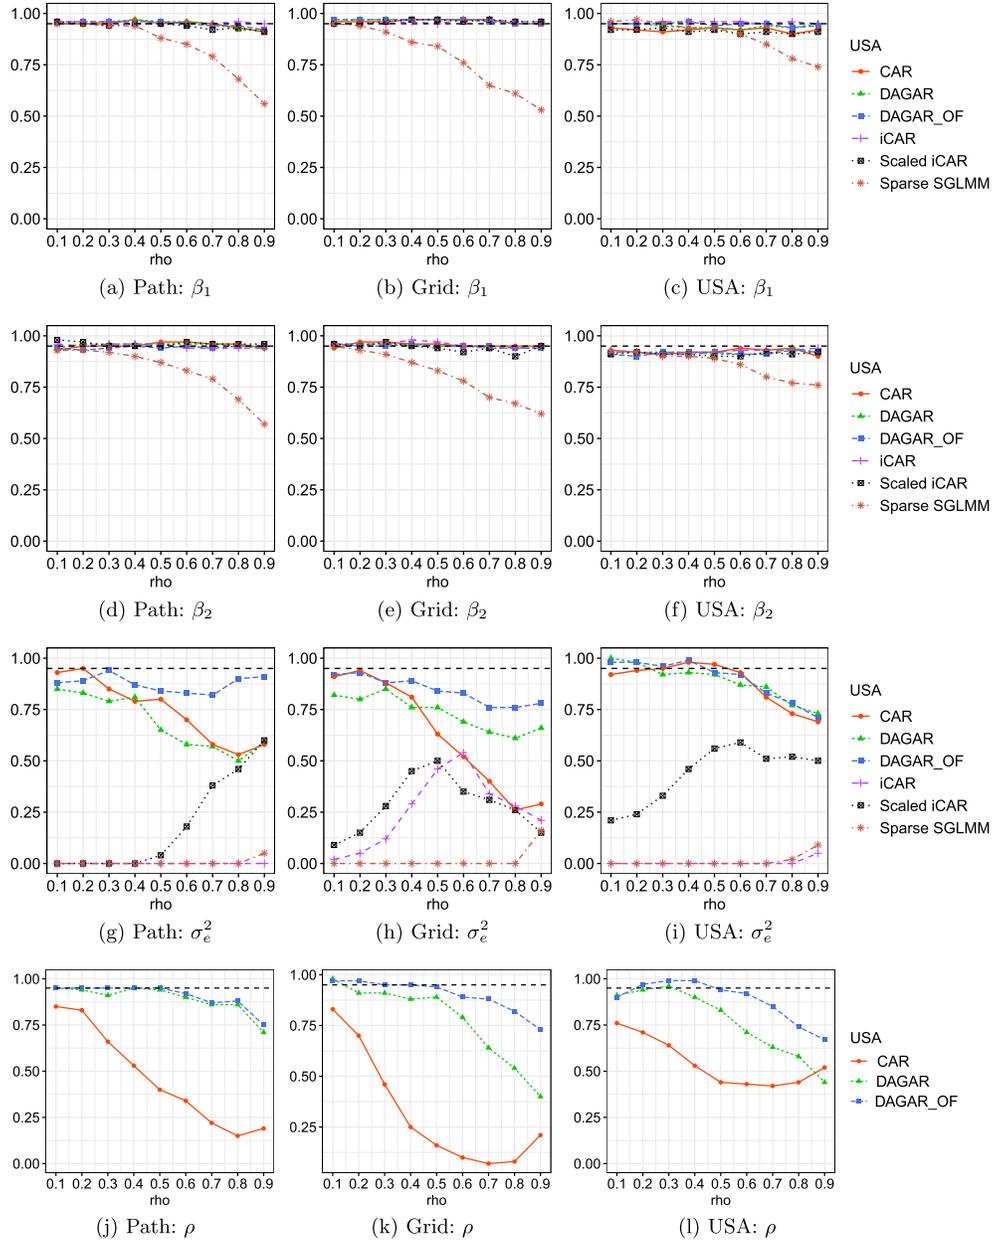


Figure 5: Coverage probabilities of the parameters as a function of the true ρ (x-axis) for the simulation data analysis using data generated from an exponential GP.

ing based on an underlying model assumption, which can lead to worse estimates if that assumption is violated. If we generated data in a way such that the eigenvectors corresponding to non-zero eigenvalues of the covariance matrix are uncorrelated with the covariates, then it is likely that the sparse GLMM will produce the most accurate estimates as the DAGAR or the other CAR models do not adjust for spatial confounding. However, our focus in this manuscript is not on spatial confounding and the data generation scheme we used is extremely common for geo-spatial settings. The dimension reduction approach used by sparse GLMM with the ICAR model as their baseline, can possibly also be adopted for the DAGAR models, to yield versions that guard against confounding. However, care has to be taken avoid eigen decompositions of the covariance matrix at every step of the Markov Chain Monte Carlo (MCMC), as the DAGAR models, unlike the ICAR, involve ρ whose value will be updated at every iteration. We identify this as one of the future research directions.

Turning to the error variance $\sigma_e^2 = 1/\tau_e$, we first note that the ICAR-based models performs surprisingly poorly offering almost zero coverage for small ρ for all the three graphs. Only the scaled ICAR offers somewhat improved coverage for larger values of ρ . The CPs from the DAGAR models are once again close to 95% for smaller values of ρ but decline for larger ρ . The proper CAR generally offers coverage worse than the DAGAR model and better than the ICAR-based models.

Reliable estimation and inference for spatial covariance parameters is a notoriously difficult problem. These results for the coverage probabilities of the spatial parameters clearly demonstrate the value of our interpretable model in delivering more accurate inference about the parameters for areal data. The results present strong evidence for the superiority of the DAGAR model both in terms of effectively recovering the latent spatial surface and the ability to assess hypotheses related to parameters describing the spatial structure in the data.

3.2 Analyses Using Different Orderings

The DAGAR model used in the analyses in Section 3.1 for the USA graph was constructed by ordering the nodes (states) from the southwest to the northeast. In this section, we repeat the analysis for the DAGAR model using three other orderings which start at southeast, northwest and northeast respectively and go approximately diagonally to the opposite end of the map.

Figure 6 plots the average MSE (left) and the estimates and confidence bands for ρ over 100 replicated datasets for the DAGAR model using these three orderings and the original ordering used in Section 3.1. We see that the ordering has little impact on the results as the MSE as well as the estimates and confidence bands for ρ for the four different orderings are nearly indistinguishable.

3.3 Data Generated Using CAR and DAGAR

In Sections 3.1 and S1.1, we generated the data using Gaussian Processes to ensure that the data generating mechanism is different from all the models we are fitting to

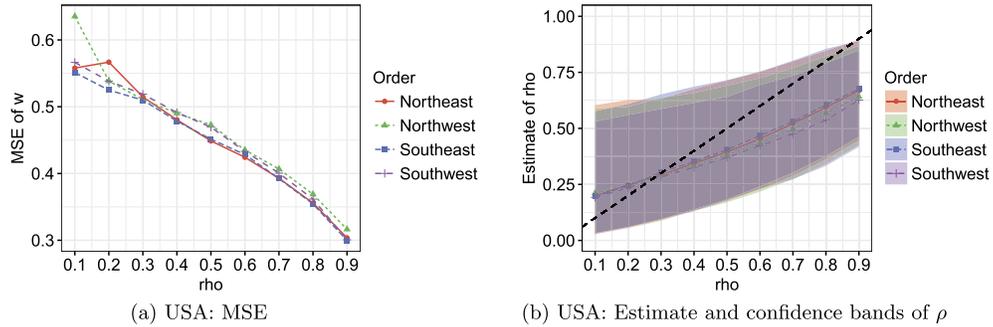


Figure 6: MSE (left) and estimates and confidence bands of ρ (right) as a function of the true ρ (x-axis) for four different orderings of the DAGAR model.

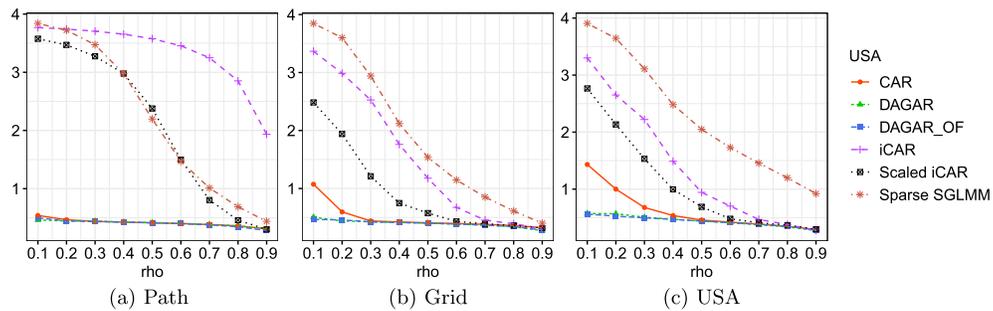


Figure 7: MSE as a function of the true ρ (x-axis) for the simulation data analysis using data generated from a DAGAR model.

the data (except for the path graph for data generated using an exponential GP). In this Section, we considered simulation schemes where the data were generated using the DAGAR or the proper CAR model for all three graphs. All the parameter choices were kept the same as in Section 3.1 and 100 replicates were used for each setting.

Figures 7 and 8 plot the average MSE when the data is generated using a DAGAR model. We see that when the data is generated using a DAGAR covariance, the DAGAR model substantially outperforms all the ICAR-based models with significantly lower MSE for all three graphs and all values of ρ . For the path graph, the proper CAR also produces MSEs similar to the DAGAR model, whereas for grid and USA graphs for smaller values of ρ , the MSE is higher than those of the DAGAR models. The trends in MSE are broadly similar to what was observed in Sections 3.1 and S1.1. When data is generated using a proper CAR, Figure 8 reveals that the DAGAR models, along with the proper CAR model (which is the true model), once again produce MSEs substantially lower than the ICAR-based models for all 3 graphs.

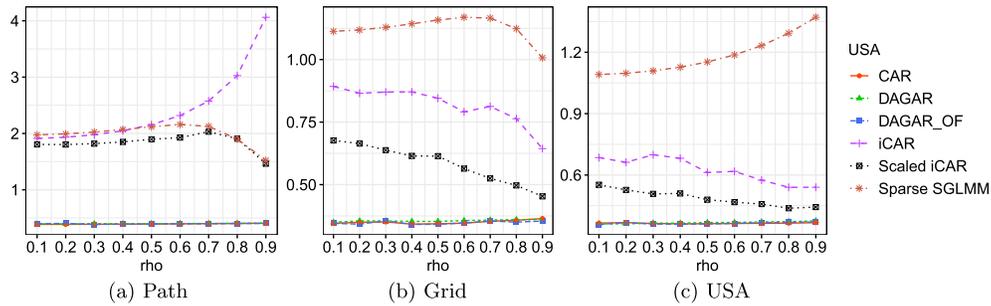


Figure 8: MSE as a function of the true ρ (x-axis) for the simulation data analysis using data generated from a proper CAR model.

3.4 Additional Analyses

The Supplemental file contains additional analyses for a) simulated data generated with spatial random effects coming from a smoother Matérn_{3/2} GP instead of the exponential GP (Section S1.1), and b) non-Gaussian (Poisson) areal data (Section S1.2). Overall, the findings from these analyses concur with the results presented here. Across all the simulation scenarios, we found the performance of the DAGAR model to be remarkably robust, uniformly producing the lowest MSEs, accurately estimating the regression coefficients, error variances, and most remarkably the spatial correlation ρ , which is considered to be notoriously difficult to estimate. With the exception of data generated on a path graph using an exponential GP, all the other settings effectively correspond to misspecified models, and the estimates of the regression coefficients from the DAGAR model were quite robust to this. The proper CAR model also performed quite well, often producing MSEs close to those from the DAGAR models except for cases when the true spatial correlation was weak, i.e., ρ was small. However, both estimation and inference for ρ from the proper CAR model was generally much less accurate than the DAGAR model. We note that across all scenarios, the ICAR-based models generally performed quite poorly, producing large MSE, with the scaled ICAR being the best in this class. Finally, the order-free DAGAR model produced results very similar to the ordered DAGAR model for all scenarios. However, it was much slower, especially due to determinant calculations. While using state-of-the-art sparse matrix algorithms would definitely help scale reduce the computing times for the order-free model, that is not the focus of the current manuscript and hence, we do not consider the order-free model for the real data analyses.

4 County-Level US Infant Mortality Data

We now analyze a large areal data using the DAGAR model. The dataset consists of counts of infant births B_i and deaths D_i for each of 3071 US counties. County-specific covariates, which possibly affect infant death rates, were available and include number

of births with low weight (low_i), percentages of black residents ($black_i$) and Hispanic residents ($Hisp_i$), a Gini index measuring income disparity ($gini_i$), social affluence (aff_i), and a measure of residential stability ($stab_i$). The dataset is publicly available in the ‘ngspatial’ package in R and was analyzed in Hughes and Haran (2013) where more description of the data is available.

We analyzed this dataset using Poisson spatial regression model where each D_i is modeled as an independent Poisson random variable with mean $B_i \exp(\alpha + \beta_1 low_i + \beta_2 black_i + \beta_3 Hisp_i + \beta_4 gini_i + \beta_5 aff_i + \beta_6 stab_i + w_i)$ where the w_i ’s are the spatial random effects. We compared the ICAR and the DAGAR priors for the w_i ’s. ICAR model was implemented in BUGS using the function *car.normal*. The unnormalized weights associated with each pair of areas were set to be 1 to make sure it corresponds to the default ICAR model. The *car.normal* distribution is parameterized to include a sum-to-zero constraint on the random effects to make sure that the intercept α is identifiable. We assigned α an improper uniform prior using the *dflat()* distribution. For remaining parameters, we assign independent $N(0, 10^{-6})$ priors to β_i ’s and *Gamma*(2, 1) to τ_w . We set the same priors for DAGAR model except for the intercept for which the prior for α is the same as β_i ’s with $N(0, 10^{-6})$. We could not implement the proper CAR model for such a large dataset. However, we also add the results of the sparse spatial GLMM model.

In addition to presenting the parameter estimates and confidence intervals, we also use model comparison metrics to evaluate the three covariance models. We used the Deviance Information Criterion (DIC, Spiegelhalter et al. 2002) to compare the posterior distributions. Table 1 presents the results for the three models. Among the seven regression coefficients, estimates for six of them were similar across the three models, with each of the credible intervals yielding the same inference. The exception to this was β_4 , whose estimates differed substantially between the sparse GLMM and the other two models. The sparse GLMM was the only one yielding a credible interval that does not cover zero.

We do not know if the difference for β_4 was due to the sparse GLMM accounting for spatial confounding, as this can only be answered depending on what we believe the

	DAGAR	ICAR	sparse GLMM
α	-5.623 (-5.944, -5.353)	-5.641 (-5.871, -5.413)	-5.430 (-5.616, -5.246)
β_1	7.803 (6.438, 9.172)	7.716 (3.924, 9.166)	8.777 (7.540, 10.032)
β_2	0.00376 (0.00208, 0.00543)	0.00364 (0.00182, 0.00915)	0.00423 (0.00288, 0.00556)
β_3	-0.00347 (-0.00501, -0.00189)	-0.00286 (-0.00859, -0.00262)	-0.00379 (-0.00488, -0.00270)
β_4	-0.0616 (-0.570, 0.480)	0.103 (-0.425, 0.631)	-0.555 (-0.977, -0.125)
β_5	-0.0770 (-0.0911, -0.0632)	-0.0778 (-0.0935, -0.0616)	-0.0757 (-0.0877, -0.0638)
β_6	-0.0413 (-0.0590, -0.0234)	-0.0448 (-0.0643, -0.0249)	-0.0285 (-0.0433, -0.0138)
τ_w	7.544 (3.615, 12.866)	32.080 (14.11, 39.87)	9.450 (3.870, 16.459)
ρ	0.987 (0.974, 0.995)		
DIC	10145.8	9902.0	10110.0

Table 1: Parameter estimates (posterior medians) and model comparison metrics for the US infant mortality data. The numbers inside braces indicates the lower and upper bounds for the 95% credible intervals.

true data generation process was. We have seen consistently in the simulation analyses using the usual data generation paradigm, how the sparse GLMM offered higher MSEs and poor inference on the regression coefficients. Also, while the credible intervals for the sparse GLMM were significant for all 7 regression coefficients (compared to 6 for the other two models), it also produced a slightly higher DIC than the ICAR model despite being a dimension reduction approach with a fewer number of parameters.

DAGAR was the only model to accommodate ρ , and the estimate and confidence intervals suggest strong spatial correlation. We have seen consistently from the simulation exercises that when the underlying spatial correlation is strong, the DAGAR model performs similarly to the CAR models. This is consistent with what we observe here. In fact, the deviance information criterion of the three models is within approximately 1% of each other, demonstrating the competitive performance of the DAGAR model even for large datasets in a non-Gaussian setup under strong spatial dependence. Moreover, through the estimation of ρ , DAGAR provides additional insight about the spatial dependence that is not offered by the ICAR model or the sparse GLMM model.

5 Discussion

The existing repertoire of covariance models used for analyzing areal datasets is extremely limited. In this manuscript, we have developed an alternative parametric model for areal datasets that promises to be a significant addition to this inventory. The parametric DAGAR models we have proposed in Section 2 are novel and offer a greater degree of interpretability than the CAR models, and will be scalable for large datasets. We observe that when spatial dependence is weak or modest, the DAGAR model excels over both variants of the CAR model, while the results are similar when there is strong spatial correlation. Since, the magnitude of the underlying spatial correlation is unknown a priori in most real life applications, we believe the DAGAR model will be a useful alternative to the CAR models. While the ordering of the locations for the DAGAR model is artificial, the theoretical results and extensive simulations strongly suggest that substantive inference from the DAGAR model is expected to be robust to the ordering.

Analyzing prevalence of many diseases simultaneously in a multivariate setup is becoming increasingly important to accommodate the correlations among different disease prevalences. Many of the popular approaches rely on Cholesky factors of conditionally autoregressive precision matrices (Gelfand and Vounatsou 2003; Martinez-Beneito 2013; Martinez-Beneito et al. 2017) which can be computationally prohibitive for large k . Our ordered model lends itself naturally to these settings due to the readily available Cholesky decomposition and, hence, promises to broaden the inventory of multivariate disease mapping models. The ordered model also offers a coherent way of modeling on arbitrary graphs or networks of growing size, i.e., if a new point is added to the graph, the nested distributions remain same, unlike any of the other three models considered here.

Supplementary Material

Supplementary Material for ‘Spatial disease mapping using directed acyclic graph autoregressive (DAGAR) models’ (DOI: [10.1214/19-BA1177SUPP](https://doi.org/10.1214/19-BA1177SUPP); .pdf).

References

- Assuncao, R. and Krainski, E. (2009). “Neighborhood dependence in Bayesian spatial models.” *Biometrical Journal*, 51: 851–869. MR2751717. doi: <https://doi.org/10.1002/bimj.200900056>. 1222
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2014). *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton, FL: Chapman & Hall/CRC, second edition. MR3362184. 1222, 1229
- Basseville, M., Benveniste, A., Chou, K. C., Golden, S. A., Nikoukhah, R., and Willsky, A. S. (2006). “Modeling and Estimation of Multiresolution Stochastic Processes.” *IEEE Transactions on Information Theory*, 38(2): 766–784. URL <http://dx.doi.org/10.1109/18.119735> MR1267630. doi: <https://doi.org/10.1109/81.260214>. 1226
- Besag, J. (1974). “Spatial interaction and statistical analysis of lattice systems.” *Journal of the Royal Statistical Society, Series B*, 36: 192–225. MR0373208. 1222
- Besag, J. and Higdon, D. (1999). “Bayesian analysis of agricultural field experiments.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(4): 691–746. MR1722238. doi: <https://doi.org/10.1111/1467-9868.00201>. 1223
- Besag, J. and Kooperberg, C. (1995). “On conditional and intrinsic autoregressions.” *Biometrika*, 82: 733–746. MR1380811. doi: <https://doi.org/10.1093/biomet/82.4.733>. 1222, 1223
- Besag, J. and Mondal, D. (2005). “First-order intrinsic autoregressions and the de Wijs process.” *Biometrika*, 92(4): 909–920. MR2234194. doi: <https://doi.org/10.1093/biomet/92.4.909>. 1223
- Bickel, P. J. and Levina, E. (2008a). “Covariance regularization by thresholding.” *The Annals of Statistics*, 36(6): 2577–2604. MR2485008. doi: <https://doi.org/10.1214/08-AOS600>. 1225
- Bickel, P. J. and Levina, E. (2008b). “Regularized estimation of large covariance matrices.” *The Annals of Statistics*, 36(1): 199–227. MR2387969. doi: <https://doi.org/10.1214/009053607000000758>. 1225
- Cai, T. T., Zhang, C.-H., and Zhou, H. H. (2010). “Optimal rates of convergence for covariance matrix estimation.” *The Annals of Statistics*, 38(4): 2118–2144. MR2676885. doi: <https://doi.org/10.1214/09-AOS752>. 1225
- Clayton, D. G. and Bernardinelli, L. (1992). “Bayesian Methods for Mapping Disease Risk.” In Elliott, P., Cuzick, J., English, D., and Stern, R. (eds.), *Geographical and*

- Environmental Epidemiology: Methods for Small-Area Studies*, 205–220. Oxford University Press. 1222
- Cressie, N. and Davidson, J. L. (1998). “Image analysis with partially ordered Markov models.” *Computational Statistics and Data Analysis*, 29(1): 1 – 26. URL <http://www.sciencedirect.com/science/article/pii/S0167947398000528> MR1665674. doi: [https://doi.org/10.1016/S0167-9473\(98\)00052-8](https://doi.org/10.1016/S0167-9473(98)00052-8). 1223, 1225
- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016). “Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets.” *Journal of the American Statistical Association*, 111(514): 800–812. MR3538706. doi: <https://doi.org/10.1080/01621459.2015.1044091>. 1223, 1225, 1231
- Datta, A., Banerjee, S., Finley, A. O., Hamm, N. A. S., and Schaap, M. (2016). “Nonseparable dynamic nearest neighbor Gaussian process models for large spatio-temporal data with an application to particulate matter analysis.” *The Annals of Applied Statistics*, 10(3): 1286–1316. MR3553225. doi: <https://doi.org/10.1214/16-AOAS931>. 1225
- Datta, A., Banerjee, S., Hodges, J. S., and Gao, L. (2019). “Supplement to “Spatial disease mapping using directed acyclic graph auto-regressive (DAGAR) models”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/19-BA1177SUPP>. 1224
- El Karoui, N. (2008). “Operator norm consistent estimation of large-dimensional sparse covariance matrices.” *The Annals of Statistics*, 36(6): 2717–2756. MR2485011. doi: <https://doi.org/10.1214/07-AOS559>. 1225
- Finley, A. O., Datta, A., Cook, B. C., Morton, D. C., Andersen, H. E., and Banerjee, S. (2017). “Applying Nearest Neighbor Gaussian Processes to Massive Spatial Data Sets: Forest Canopy Height Prediction Across Tanana Valley Alaska.” <https://arxiv.org/pdf/1702.00434.pdf>. 1225
- Friedman, J., Hastie, T., and Tibshirani, R. (2007). “Sparse inverse covariance estimation with the graphical lasso.” *Biostatistics*, 9: 432–441. 1225
- Gelfand, A. E. and Vounatsou, P. (2003). “Proper multivariate conditional autoregressive models for spatial data analysis.” *Biostatistics*, 4(1): 11. URL <http://dx.doi.org/10.1093/biostatistics/4.1.11> 1240
- Hinton, G. E. (2002). “Training products of experts by minimizing contrastive divergence.” *Neural Computation*, 14: 1711–1800. 1229
- Hughes, J. and Cui, X. (2018). *ngspatial: Fitting the Centered Autologistic and Sparse Spatial Generalized Linear Mixed Models for Areal Data*. Denver, CO. R package version 1.2-1. 1233
- Hughes, J. and Haran, M. (2013). “Dimension reduction and alleviation of confounding for spatial generalized linear mixed models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(1): 139–159. MR3008275. doi: <https://doi.org/10.1111/j.1467-9868.2012.01041.x>. 1232, 1239
- Leroux, B. G., Lei, X., and Breslow, N. (2000). “Estimation of Disease Rates in Small

- Areas: A new Mixed Model for Spatial Dependence.” In Halloran, M. E. and Berry, D. (eds.), *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, 179–191. New York, NY: Springer New York. MR1731684. doi: https://doi.org/10.1007/978-1-4612-1284-3_4. 1223
- MacNab, Y. and Dean, C. (2000). “Parametric bootstrap and penalized quasi-likelihood inference in conditional autoregressive models.” *Statistics in Medicine*, 19: 15–30. 1223
- Martinez-Beneito, M. A. (2013). “A general modelling framework for multivariate disease mapping.” *Biometrika*, 100(3): 539. MR3094436. doi: <https://doi.org/10.1093/biomet/ast023>. 1240
- Martinez-Beneito, M. A., Botella-Rocamora, P., and Banerjee, S. (2017). “Towards a Multidimensional Approach to Bayesian Disease Mapping.” *Bayesian Analysis*, 12(1): 239–259. MR3597574. doi: <https://doi.org/10.1214/16-BA995>. 1240
- Meinshausen, N. and Bühlmann, P. (2006). “High-dimensional graphs and variable selection with the Lasso.” *The Annals of Statistics*, 34(3): 1436–1462. MR2278363. doi: <https://doi.org/10.1214/009053606000000281>. 1225
- Rothman, A. J., Levina, E., and Zhu, J. (2009). “Generalized Thresholding of Large Covariance Matrices.” *Journal of the American Statistical Association*, 104(485): 177–186. MR2504372. doi: <https://doi.org/10.1198/jasa.2009.0101>. 1225
- Sørbye, S. H. and Rue, H. (2014). “Scaling intrinsic Gaussian Markov random field priors in spatial modelling.” *Spatial Statistics*, 8: 39–51. MR3326820. doi: <https://doi.org/10.1016/j.spasta.2013.06.004>. 1232
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). “Bayesian Measures of Model Complexity and Fit.” *Journal of the Royal Statistical Society, Series B*, 64: 583–639. MR1979380. doi: <https://doi.org/10.1111/1467-9868.00353>. 1239
- Stein, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. New York, NY: Springer, first edition. MR1697409. doi: <https://doi.org/10.1007/978-1-4612-1494-6>. 1225
- Stein, M. L., Chi, Z., and Welty, L. J. (2004). “Approximating Likelihoods for Large Spatial Data Sets.” *Journal of the Royal Statistical Society, Series B*, 66: 275–296. MR2062376. doi: <https://doi.org/10.1046/j.1369-7412.2003.05512.x>. 1231
- Sudderth, E. B. (2002). “Embedded Trees: Estimation of Gaussian Processes on Graphs with Cycles.” <http://cs.brown.edu/~sudderth/papers/sudderthMasters.pdf>. 1226
- Vecchia, A. V. (1988). “Estimation and Model Identification for Continuous Spatial Processes.” *Journal of the Royal Statistical Society, Series B*, 50: 297–312. MR0964183. 1231
- Wall, M. (2004). “A close look at the spatial structure implied by the CAR and SAR

- models." *Journal of Statistical Planning and Inference*, 121: 311–324. MR2038824. doi: [https://doi.org/10.1016/S0378-3758\(03\)00111-3](https://doi.org/10.1016/S0378-3758(03)00111-3). 1222, 1223
- Whittle, P. (1954). "On Stationary Processes in the Plane." *Biometrika*, 41(3/4): 434–449. URL <http://www.jstor.org/stable/2332724> MR0067450. doi: <https://doi.org/10.1093/biomet/41.3-4.434>. 1222
- Wu, W. and Pourahmadi, M. (2003). "Nonparametric Estimation of Large Covariance Matrices of Longitudinal Data." *Biometrika*, 90(4): 831–844. URL <http://www.jstor.org/stable/30042091> MR2024760. doi: <https://doi.org/10.1093/biomet/90.4.831>. 1225
- Xue, L., Ma, S., and Zou, H. (2012). "Positive-Definite ℓ_1 -Penalized Estimation of Large Covariance Matrices." *Journal of the American Statistical Association*, 107(500): 1480–1491. MR3036409. doi: <https://doi.org/10.1080/01621459.2012.725386>. 1225