

Bayesian Variable Selection in Double Generalized Linear Tweedie Spatial Process Models

ARITRA HALDER^{*,1}, SHARIQ MOHAMMED¹, AND DIPAK K. DEY

Abstract

Double generalized linear models provide a flexible framework for modeling data by allowing the mean and the dispersion to vary across observations. Common members of the exponential dispersion family including the Gaussian, Poisson, compound Poisson-gamma (CP-g), Gamma and inverse-Gaussian are known to admit such models. The lack of their use can be attributed to ambiguities that exist in model specification under a large number of covariates and complications that arise when data display complex spatial dependence. In this work we consider a hierarchical specification for the CP-g model with a spatial random effect. The spatial effect is targeted at performing uncertainty quantification by modeling dependence within the data arising from location based indexing of the response. We focus on a Gaussian process specification for the spatial effect. Simultaneously, we tackle the problem of model specification for such models using Bayesian variable selection. It is effected through a continuous spike and slab prior on the model parameters, specifically the fixed effects. The novelty of our contribution lies in the Bayesian frameworks developed for such models. We perform various synthetic experiments to showcase the accuracy of our frameworks. They are then applied to analyze automobile insurance premiums in Connecticut, for the year of 2008.

KEYWORDS AND PHRASES: Bayesian Modeling, Gaussian Process, Hierarchical Spatial Process Models, Spike and Slab Priors, Tweedie Double Generalized Linear Models.

1. INTRODUCTION

Spatial processes have occupied center stage in statistical theory and applications for the last few decades. Their voracious use can largely be explained by geographically tagged data becoming increasingly commonplace in modern applications. Such data are often composed of complex variables which are no longer amenable to a Gaussian assumption. For example, spatially indexed counts [see for e.g., 61, 9, 2, 39], proportions [see, for e.g., 23, 26, 21, 20], time to event or, survival outcomes [see, for e.g. 4, 44, 67] are some frequently occurring variables where spatial processes have proved invaluable in performing uncertainty quantification. The purpose being to quantify unobserved dependence introduced within the variable of interest due to varying location. The cornerstone for such studies is a spatially indexed process variable of interest, often termed as a response process and denoted by, $y(\mathbf{s})$. This is accompanied by covariate information $\mathbf{X}(\mathbf{s}) = [\mathbf{x}_1(\mathbf{s}), \mathbf{x}_2(\mathbf{s}), \dots, \mathbf{x}_p(\mathbf{s})]$. Here $\mathbf{s} \in \mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L\}$ is the spatial indexing and, \mathcal{S} is a finite set of indices or, locations over which the response variable and covariates are observed. The investigator often encounters observation-level covariates that account for response specific characteristics when learning such processes. It becomes important to understand which of these covari-

ates are important contributors to variation in the data. From a model parsimony standpoint, model choice becomes an important issue to the investigator. In statistical theory, this problem is often addressed by performing shrinkage or, variable selection on the model coefficients. Moreover, performing spatial uncertainty quantification produces accurate inference for model coefficients, which also raises the concern regarding a more “honest” subset of covariates within $\mathbf{X}(\mathbf{s})$ that primarily determine the variation in $y(\mathbf{s})$. The crown jewel of our contribution is Bayesian methodology for performing spatial uncertainty quantification and model choice simultaneously.

Spatial process modeling generally requires a hierarchical specification of an unobserved random effect within the model [14]. Maintaining a hierarchical approach allows for exploration of the effect of covariates, $\mathbf{X}(\mathbf{s})$, and the random effect $\mathbf{w}(\mathbf{s})$ jointly on the response process $y(\mathbf{s})$. Particularly, considering generalized linear spatial process modeling, it is assumed that $y(\mathbf{s}) \mid \boldsymbol{\beta}, \mathbf{w}(\mathbf{s})$ arise in a *conditionally* independent fashion from a member of the exponential family with mean $\mu(\mathbf{s})$ such that $g(\mu(\mathbf{s})) = \mathbf{X}(\mathbf{s})\boldsymbol{\beta} + \mathbf{F}(\mathbf{s})\mathbf{w}(\mathbf{s})$, where g is a monotonic link function, $\boldsymbol{\beta}$ are model coefficients, or fixed effects and $\mathbf{F}(\mathbf{s})$ is a spatial incidence matrix. In contrast to Gaussian response processes where a direct hierarchical specification on the response is feasible, modeling a non-Gaussian spatial process leverages the generalized linear model framework to employ a latent process

*Corresponding author.

¹Equal contribution.

specification [see for e.g., 64, 17, 16, 65, 5]. This is facilitated by the existence of a valid joint probability distribution, $\pi(y(\mathbf{s}_1), y(\mathbf{s}_2), \dots, y(\mathbf{s}_L) \mid \boldsymbol{\beta}, \boldsymbol{\theta}_{pr})$, where $\boldsymbol{\theta}_{pr}$ denotes process parameters required for specification of $\mathbf{w}(\mathbf{s})$ [see discussion in section 6.2, 5]. This hierarchical specification gives us a natural way to perform variable selection by incorporating shrinkage priors into the hierarchical prior formulation. There are several choices of shrinkage priors that differ in their prior specification such as: the discrete spike-and-slab prior [see, for e.g. 46, 27], other priors based on the Gaussian-gamma family in linear Gaussian models [see, for e.g. 51, 7], the continuous spike-and-slab prior [see, for e.g. 33], the Bayesian counterparts of LASSO and elastic net [see, for e.g. 50, 41], mixtures of g -priors [see, for e.g. 42], the horseshoe priors and its variants [see, for e.g. 13], among several others. We use the continuous spike-and-slab prior to effect shrinkage on model coefficients.

We focus on a subset of probability distributions within the exponential family, termed as the exponential dispersion family [see, for e.g., 34, 35, 36, 37]. It allows the dispersion along with the mean to vary across observations, suppressing the need for having a constant dispersion across observations. We focus on a particular member of the family, the Tweedie compound Poisson-gamma (CP-g), more commonly referred to as the Tweedie (probability) distribution [see, 58]. The corresponding random variable is constructively defined as a Poisson sum of independently distributed Gamma random variables. Allowing for a varying dispersion across observations enables exploration of the effect of covariates $\mathbf{X}(\mathbf{s})$ on the mean and the dispersion separately, by employing two separate generalized linear models (GLMs). This gives rise to the double generalized linear model (DGLM) [see, for e.g., 54, 59, 56]. Hierarchical frameworks for specifications of DGLMs were first developed in [40]. Although not mandatory, it is customary to use the same covariates, $\mathbf{X}(\mathbf{s})$ for both the mean and dispersion GLMs to avoid ambiguities in model specification. Previous work [see for e.g., 29, 30] uses this approach and considers developing inference for DGLMs under a frequentist framework. Inference on spatial effects is obtained through penalizing the graph Laplacian. In this paper we adopt a Bayesian discourse by supplementing the DGLM framework with the continuous version of the spike-and-slab prior to effect shrinkage and thereby achieve better model specification. We integrate the spike and slab prior into our hierarchical prior formulation for both mean and dispersion models. We show that these priors provide a natural way of incorporating sparsity into the model, while offering straightforward posterior sampling in the context of our spatial DGLMs.

The scale for spatial indexing is assumed to be point-referenced. For example, latitude-longitude or easting-northing. Generally, $\mathbf{s} \in \mathbb{R}^2$. Specification of a neighborhood structure or, proximity is naturally important when attempting to quantifying the behavior of response in locations that are *near* each other. We select the Euclidean distance between locations. This results in a Gaussian process

[see, for e.g., 60] prior on the spatial process, $\mathbf{w}(\mathbf{s})$. Other choices exist for specifying such spatial process, $\mathbf{w}(\mathbf{s})$, for e.g. log-Gamma [see 10, 11] etc. We particularly focus on a Gaussian process specification on $\mathbf{w}(\mathbf{s}) \sim GP(\mathbf{0}, K(\cdot))$, where K is a covariance function. For arbitrary locations, \mathbf{s} and \mathbf{s}' , dependence between $y(\mathbf{s})$ and $y(\mathbf{s}')$ is specified through $K(\mathbf{s}, \mathbf{s}')$, which governs the covariance between $w(\mathbf{s})$ and $w(\mathbf{s}')$. For point-referenced data, the Matérn family [see, for e.g., 45] provides the most generic and widely adopted covariance specification.

Next, we address Bayesian model specification. In the absence of such concerns for the hierarchical process models discussed above, prior specification follows the generic framework, $[\text{data} \mid \text{process}, \boldsymbol{\theta}] \times [\text{process} \mid \tilde{\boldsymbol{\theta}}] \times [\boldsymbol{\theta}_m] \times [\boldsymbol{\theta}_{pr}]$. Here, $\tilde{\boldsymbol{\theta}} = \{\boldsymbol{\theta}_m, \boldsymbol{\theta}_{pr}\}$ denote model parameters [see for e.g. 8, 24, 5, Chapter 6, p. 125]. In particular, $\boldsymbol{\theta}_{pr}$ constitute parameters instrumental in specification of the process, while $\boldsymbol{\theta}_m$ are other model parameters. We adopt a proper prior on $\boldsymbol{\theta}_{pr}$ to avoid the risk of generating improper posteriors [see, for e.g., 6]. Building a Bayesian variable selection framework that facilitates model specification for $\boldsymbol{\theta}_m$ requires an additional layer of hierarchical prior specification, appending the latter framework with variable selection parameters, $\boldsymbol{\theta}_{vs}$ and, thereby producing

$$[\text{data} \mid \text{process}, \boldsymbol{\theta}] \times [\text{process} \mid \boldsymbol{\theta}] \times [\boldsymbol{\theta}_{pr}] \times [\boldsymbol{\theta}_m] \times [\boldsymbol{\theta}_{vs}], \quad (1.1)$$

where $\boldsymbol{\theta} = \{\tilde{\boldsymbol{\theta}}, \boldsymbol{\theta}_{vs}\}$. We resort to Markov Chain Monte Carlo (MCMC) sampling [see, for e.g. 12, 28] for performing posterior inference on $\boldsymbol{\theta}$. The novelty of our approach lies in the simple Bayesian computation devised—employing only Gibbs sampling updates for $\boldsymbol{\theta}_{vs}$. To the best of our knowledge *hierarchical Bayesian frameworks* for fitting (a) Tweedie DGLMs, (b) spatial Tweedie DGLMs with (or without) variable selection, do not exist in the statistical literature. Evidently, proposed methodology in (1.1) remedies that.

The ensuing developments in the paper are organized as follows: In Section 2 we detail the proposed statistical framework outlining Tweedie distributions—the likelihood and parameterization, model formulation and the hierarchical prior specification. Section 3 provides comprehensive synthetic experiments that document the efficacy of our proposed statistical framework for Bayesian variable selection in spatial DGLMs. Section 5 considers application of the developed framework to automobile insurance premiums in Connecticut, USA during 2008. Additional synthetic experiments capturing various performance aspects for the models are provided in the Supplementary Material.

2. STATISTICAL FRAMEWORK

The Tweedie distribution produces observations composed of exact zeros with a continuous Gamma tail. Their ability to model mixed data types featuring exact zeros and continuous measurements jointly makes them suitable for

modeling response arising from a variety of domains. Some of the current applications include, actuarial science [see 55, 62, 29, 30], ecology [see for e.g., 57], public health [63], environment [see for e.g., 38], ecology [53], gene expression studies [43]. As evidenced by these applications, the presence of unobserved dependence between observations, affecting the quality of inference, is not unlikely. In the following subsections we provide more details on Tweedie distributions, followed by the model formulation and hierarchical prior specification.

2.1 The Exponential Dispersion Family: Tweedie Distributions

The Tweedie family of distributions belong to the exponential dispersion (ED) family of models whose probability density/mass function has the generic form,

$$\pi(y | \theta, \phi) = a(y, \phi) \exp \{ \phi^{-1}(y\theta - \kappa(\theta)) \}, \quad (2.1)$$

where θ is the natural or canonical parameter, ϕ is the dispersion parameter, and $\kappa(\theta)$ is the cumulant function. Characterizing the Tweedie family is an index parameter ξ , varying values of which produce different members of the family. For e.g., the CP-g is obtained with $\xi \in (1, 2)$, for $\xi = 1$ we obtain a Poisson and $\xi = 2$ produces a Gamma distribution, for $\xi \in (0, 1)$ they do not exist [for further details see Table 1, 29]. We are particularly interested in the CP-g distributions in this work. In general, for the ED family we have the mean, $\mu = E(y) = \kappa'(\theta)$ and the variance, $Var(y) = \phi\kappa''(\theta)$. For the CP-g we have $\kappa(\theta) = (2 - \xi)^{-1}\{(1 - \xi)\theta\}^{2-\xi/1-\xi}$. Using the relation, $\kappa'(\theta) = \mu$, some straightforward algebra yields, $\kappa(\theta) = (2 - \xi)^{-1}\mu^{2-\xi}$ and $\kappa''(\theta) = \mu^\xi$, implying $Var(y) = \phi\mu^\xi$ and denoting $\alpha = (1 - \xi)^{-1}(2 - \xi)$ we have,

$$a(y, \phi) = 1 \cdot I(y = 0) + y^{-1} \sum_{j=1}^{\infty} \left[\frac{y^{-\alpha}(\xi - 1)^\alpha}{\phi^{1-\alpha}(2 - \xi)} \right]^j \frac{1}{j! \Gamma(-j\alpha)} \times I(y > 0).$$

Evidently, $\pi(0 | \theta, \phi) = \exp\{-\phi^{-1}\kappa(\theta)\}$.

We introduce some notation. Let $y_{ij}(\mathbf{s}_i)$ denote the j -th response at the i -th location $\mathbf{s}_i \in \mathcal{S}$, where $j = 1, 2, \dots, n_i$ and $i = 1, 2, \dots, L$ with $\sum_{j=1}^L n_i = N$. Together we denote, $\mathbf{y} = \mathbf{y}(\mathbf{s}) = \{\{y_{ij}(\mathbf{s}_i)\}_{j=1}^{n_i}\}_{i=1}^L$ as the $N \times 1$ response. Similarly, $\boldsymbol{\mu} = \boldsymbol{\mu}(\mathbf{s}) = \{\{\mu_{ij}(\mathbf{s}_i)\}_{j=1}^{n_i}\}_{i=1}^L$ and $\boldsymbol{\phi} = \{\{\phi_{ij}\}_{j=1}^{n_i}\}_{i=1}^L$ denotes the mean and dispersion vectors respectively. If $\mathbf{y} | \boldsymbol{\mu}, \boldsymbol{\phi}, \xi$ arises independently from a CP-g distribution, then the likelihood is given by

$$\pi(\mathbf{y} | \boldsymbol{\mu}, \boldsymbol{\phi}, \xi) = \prod_{i=1}^L \prod_{j=1}^{n_i} a_{ij}(y_{ij}(\mathbf{s}_i) | \phi_{ij}) \times \exp \left[\phi_{ij}^{-1} \left(\frac{y_{ij}(\mathbf{s}_i)\mu_{ij}(\mathbf{s}_i)^{1-\xi}}{1 - \xi} - \frac{\mu_{ij}(\mathbf{s}_i)^{2-\xi}}{2 - \xi} \right) \right]. \quad (2.2)$$

Working with the likelihood, $\pi(\cdot)$, when devising computation, evaluating the infinite series representation of $a(y, \phi)$

is required. The two commonly used methods are—saddle-point approximation [see for e.g. 49, 55, 18, 66] and Fourier inversion [see, for e.g. 19]. The saddle-point approximation to (2.1) uses a deviance function based representation where, $\tilde{\pi}(y | \mu, \phi) = b(y, \phi) \exp\{-(2\phi)^{-1}d(y | \mu)\}$. For CP-g distributions, the deviance function is $d(y | \mu) = d(y | \mu, \xi) = 2\{(y^{2-\xi} - y\mu^{1-\xi})(1 - \xi)^{-1} - (y^{2-\xi} - \mu^{2-\xi}) \times (2 - \xi)^{-1}\}$, and $b(y | \phi, \xi) = (2\pi\phi y^\xi)^{-1/2}I(y > 0) + 1 \cdot I(y = 0) \approx a(y, \phi) \exp\{\phi^{-1}y^{2-\xi}(1 - \xi)^{-1}(2 - \xi)^{-1}\}$. We performed experiments which showed that the saddle-point approximation performs well when fewer zeros are present in the data. Under higher proportions of zeros its performance was sub-optimal. However, albeit its computationally intensive nature, in all scenarios the Fourier inversion based method had stable performance. Hence, in this paper we use the evaluation of $a(y, \phi)$ that is based on Fourier inversion. The adopted Bayesian approach requires MCMC computation that relies on accurate likelihood evaluations. Hence, we emphasize the importance of choosing the appropriate likelihood function for application purposes. We denote the likelihood in (2.2) as $Tw(\boldsymbol{\mu}, \boldsymbol{\phi}, \xi)$. Tweedie distributions are the only members of the ED family that possess a scale invariance property [see, for e.g., 37, Theorem 4.1]. This suggests for $c_{ij} > 0$, $\mathbf{y}^*(\mathbf{s}) = \{c_{ij}y_{ij}(\mathbf{s}_i)\} \sim Tw(\mathbf{c}^T \boldsymbol{\mu}, \mathbf{c}^{2-\xi T} \boldsymbol{\phi}, \xi)$ allowing observations with different scales of measurement to be modeled jointly.

2.2 Model Formulation

Formulating DGLMs with spatial effects theoretically involves specification of a spatial random effect in both mean and dispersion models. In such a scenario complex dependencies can be specified to account for varied degrees of uncertainty quantification. In the simplest case the corresponding spatial random effects for the mean and dispersion models are independent Gaussian processes. More complex scenarios can feature dependent Gaussian processes, where the dependence arises from a cross-covariance matrix. Spatial random effects in the mean model are readily interpretable—risk faced (adjustment to mean premium paid in our case) owing to location. However, spatial random effects in the dispersion model are not readily interpretable. Subsequently, we choose to include spatial random effects only in the mean model for this work.

Let $\mathbf{x}_{ij}(\mathbf{s}_i)$ denote a $p \times 1$ vector of observed covariates for the mean model and $\boldsymbol{\beta}$ be the corresponding a $p \times 1$ vector of coefficients. $\mathbf{f}_{ij}(\mathbf{s}_i)$ denotes a $L \times 1$ vector specifying the location and $w(\mathbf{s}_i)$ is the spatial effect at location \mathbf{s}_i with $\mathbf{w} = \mathbf{w}(\mathbf{s}) = (w(\mathbf{s}_1), w(\mathbf{s}_2), \dots, w(\mathbf{s}_L))^T$ denoting the $L \times 1$ vector of spatial effects; $\mathbf{z}_{ij}(\mathbf{s}_i)$ denotes a $q \times 1$ vector of known covariates for the dispersion model and $\boldsymbol{\gamma}$ is the corresponding $q \times 1$ vector of coefficients. A Bayesian hierarchical double generalized linear model (DGLM) using a non-canonical logarithmic link function is specified as

$$\begin{aligned} \log \mu_{ij}(\mathbf{s}_i) &= \mathbf{x}_{ij}^T(\mathbf{s}_i)\boldsymbol{\beta} + \mathbf{f}_{ij}(\mathbf{s}_i)^T \mathbf{w}(\mathbf{s}_i), \\ \log \phi_{ij} &= \mathbf{z}_{ij}^T \boldsymbol{\gamma}, \end{aligned} \quad (2.3)$$

which implies $\mu_{ij}(\mathbf{s}_i) = \mu_{ij}(\boldsymbol{\beta}, \mathbf{w}) = \exp(\mathbf{x}_{ij}^T(\mathbf{s}_i)\boldsymbol{\beta} + \mathbf{f}_{ij}^T(\mathbf{s}_i)\mathbf{w}(\mathbf{s}_i))$ and $\phi_{ij} = \phi_{ij}(\boldsymbol{\gamma}) = \exp(\mathbf{z}_{ij}^T\boldsymbol{\gamma})$.

2.3 Hierarchical Prior Specification

In this section, we first present the hierarchical prior formulation for the model and process parameters, $\tilde{\boldsymbol{\theta}}$, followed by the prior formulation for the variable selection parameters $\boldsymbol{\theta}_{vs}$. Prior specification for model and process parameters are as follows:

$$\begin{aligned} \text{Model Parameters: } & \xi \sim U(a_\xi, b_\xi); \\ & \boldsymbol{\beta} \sim N_p(\mathbf{0}_p, \boldsymbol{\lambda}_\beta^T \mathbf{I}_p); \boldsymbol{\gamma} \sim N_q(\mathbf{0}_q, \boldsymbol{\lambda}_\gamma^T \mathbf{I}_q), \\ \text{Process Parameters: } & \phi_s \sim U(a_{\phi_s}, b_{\phi_s}); \\ & \sigma^{-2} \sim \text{Gamma}(a_\sigma, b_\sigma); \nu \sim U(a_\nu, b_\nu), \\ \text{Process: } & \mathbf{w} \sim N_L(\mathbf{0}_L, \sigma^2 \mathbf{R}(\phi_s)), \end{aligned} \quad (2.4)$$

where $\mathbf{0}_m$ is the $m \times 1$ zero vector and \mathbf{I}_m is the $m \times m$ identity matrix; $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\mathbf{Z} \in \mathbb{R}^{n \times q}$ are design matrices corresponding to the mean and dispersion models, with coefficients $\boldsymbol{\beta} \in \mathbb{R}^p$ and $\boldsymbol{\gamma} \in \mathbb{R}^q$, respectively; $\mathbf{F} \in \mathbb{R}^{n \times L}$ is the spatial incidence matrix, $\mathbf{w} \in \mathbb{R}^L$ is the spatial effect, and $\mathbf{R}(\phi_s) = \sigma^2(\phi \|\Delta\|)^\nu K_\nu(\phi \|\Delta\|)$, where K_ν is the modified Bessel function of order ν [1], is the Matérn covariance kernel. Here $\{\|\Delta\|\}_{ii'} = \|\mathbf{s}_i - \mathbf{s}_{i'}\|_2$, is the Euclidean distance between locations \mathbf{s}_i and $\mathbf{s}_{i'}$. $U(\cdot \mid a, b)$ denotes the uniform distribution, $N_m(\cdot \mid \mathbf{0}, \Sigma)$ is the m -dimensional Gaussian with zero mean and covariance matrix Σ , and $\text{Gamma}(\cdot \mid a, b)$ is the Gamma distribution with shape-rate parameterization. Note that the priors on $\boldsymbol{\lambda}_\beta = (\lambda_{\beta,1}, \dots, \lambda_{\beta,p})$ and $\boldsymbol{\lambda}_\gamma = (\lambda_{\gamma,1}, \dots, \lambda_{\gamma,q})$ are a part of the variable selection priors and are discussed next. Referring to the framework in (1.1), the resulting posterior from (2.4) establishes the [data | process, $\boldsymbol{\theta}$] \times [process | $\tilde{\boldsymbol{\theta}}$] step. Conditional posteriors for $\tilde{\boldsymbol{\theta}}$ are outlined in Section S1 of Supplementary Materials.

For the continuous spike-and-slab prior formulation, $\boldsymbol{\theta}_{vs} = \{\zeta_\beta, \zeta_\gamma, \sigma_\beta^2, \sigma_\gamma^2, \alpha_\beta, \alpha_\gamma\}$ [see, for e.g., 33]. Note that we have separate prior formulations for mean and dispersion models. Let $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_p)^T$ and $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots, \gamma_q)^T$ be the model coefficients corresponding to the mean and the dispersion models. Let us define $\lambda_{\beta,u} = \zeta_{\beta,u} \sigma_{\beta,u}^2$ and $\lambda_{\gamma,v} = \zeta_{\gamma,v} \sigma_{\gamma,v}^2$ for $u = 1, \dots, p$ and $v = 1, \dots, q$, respectively. We consider the following prior formulation:

$$\pi(\boldsymbol{\theta}_{vs}) = \begin{cases} \zeta_{\beta,u} \stackrel{iid}{\sim} (1 - \alpha_\beta) \delta_{\nu_0}(\cdot) + \alpha_\beta \delta_1(\cdot), \\ \zeta_{\gamma,v} \stackrel{iid}{\sim} (1 - \alpha_\gamma) \delta_{\nu_0}(\cdot) + \alpha_\gamma \delta_1(\cdot), \\ \alpha_\beta \sim U(0, 1), \alpha_\gamma \sim U(0, 1), \\ \sigma_{\beta,u}^{-2} \stackrel{iid}{\sim} \text{Gamma}(a_{\sigma_\beta}, b_{\sigma_\beta}), \\ \sigma_{\gamma,v}^{-2} \stackrel{iid}{\sim} \text{Gamma}(a_{\sigma_\gamma}, b_{\sigma_\gamma}), \end{cases} \quad (2.5)$$

where β_u and γ_v have normal priors with mean 0 and variance $\zeta_{\beta,u} \sigma_{\beta,u}^2$ and $\zeta_{\gamma,v} \sigma_{\gamma,v}^2$, respectively. Here, $\delta_c(\cdot)$ denotes

the discrete measure at c ; hence, $\zeta_{\beta,u}$ and $\zeta_{\gamma,v}$ are indicators taking values 1 or ν_0 (small number close to 0) based on the selection of their corresponding covariates. The probabilities of these indicators taking the value 1 is given by α_β and α_γ respectively. We place a uniform prior on these selection probabilities and an inverse-Gamma prior on the parameters $\sigma_{\beta,u}^2$ and $\sigma_{\gamma,v}^2$. The choice of the shape and rate parameters ($a_{\sigma_\beta}, b_{\sigma_\beta}; a_{\sigma_\gamma}, b_{\sigma_\gamma}$) of these inverse-Gamma priors induces a continuous bimodal distributions on $\zeta_{\beta,u} \sigma_{\beta,u}^2$ and $\zeta_{\gamma,v} \sigma_{\gamma,v}^2$ with a spike at ν_0 and a right continuous tail. Combining the priors in (2.4) and (2.5) completes the hierarchical prior formulation for parameters $\boldsymbol{\theta}$ as defined in (1.1). Evidently, the above prior formulation allows for sufficient flexibility regarding variations in implementation. For instance, a hierarchical Bayesian framework for a simple DGLM can be obtained by omitting the process specification and variable selection. Analogously, DGLMs featuring variable selection or, spatial effects are obtained by omitting respective components from the prior specification outlined previously.

2.4 Bayesian Estimation and Inference

In its full capacity (a model with spatial effects and variable selection) the model structure with prior specifications in (2.4) and (2.5) contains $3p + L + 3q + 4$ parameters. Depending on the dimensions of $\mathbf{X}(\mathbf{s})$ and \mathbf{Z} , and the number of locations L , posterior inference can be a sufficiently daunting task. Traditional Metropolis–Hasting (M–H) random walk strategies are sub-optimal, involving costly pilot runs to determine viable initial starting points and unreasonably long chains while performing MCMC sampling. To avoid such issues, we use an adaptive rejection sampling while leveraging the log-concavity of the posteriors to perform effective inference that is not plagued by the above described issues [for more details, see, for e.g., 28]. In the following, we describe (i) briefly, our adaptive rejection MCMC sampling approach (more details are provided in the Supplementary Materials), (ii) the identifiability issues on the overall intercept that arise due to inclusion of a spatial effect and a strategy to address this, and (iii) a false discovery rate (FDR)–based approach for performing variable selection.

The joint posterior $\pi(\boldsymbol{\theta} \mid \mathbf{y})$ generated as a result of the hierarchical priors in Eq. 2.4 is sampled using a hybrid sampling strategy that includes M–H random walk and the Metropolis-Adjusted Langevin Algorithm (MALA) [52, 28]. We consider MALA updates for the model parameters $\{\boldsymbol{\beta}, \mathbf{w}\}$ for the mean model. The dispersion model coefficients are sampled depending on the choice of likelihood, i.e., $\boldsymbol{\gamma}$ is sampled using a MALA if a saddle-point approximation of the likelihood is considered, otherwise $\boldsymbol{\gamma}$ is sampled using a MALA with a numerical approximation to the derivative of the conditional posterior for $\boldsymbol{\gamma}$ or using a M–H random walk. The parameters $\{\xi, \phi_s\}$ are updated using a M–H random walk. All the other remaining parameters are sampled using Gibbs sampling. In particular, we employ block updates for $\boldsymbol{\beta}_{\mathbf{w}} = \{\boldsymbol{\beta}, \mathbf{w}\}$ and $\boldsymbol{\gamma}$. Proposal variances

feature adaptive scaling such that the optimal acceptance rate ($\approx 58\%$) to capture Langevin dynamics is achieved upon convergence [see, 12, 28]. Proposal variances in the M-H updates also feature adaptive scaling such that the optimal acceptance rate ($\approx 33\%$) for random walks is achieved upon convergence. We outline the full sampling algorithm at the end of Section S1 of the Supplement. For the hierarchical DGLM in (2.3), the specification of a spatial effect translates to fitting a random intercept mean model. Consequently, having an additional overall intercept β_0 in the model renders it unidentifiable [see, 25, 22]. Hence, β_0 is not estimable, although $\beta_0 + \mathbf{w}$ is estimable. β_0 is estimated through hierarchical centering of the posterior for \mathbf{w} [see, for e.g., 22].

The MCMC samples of β and γ explore their conditional posterior distributions and point estimates for these model parameters can be obtained using *maximum a-posteriori* (MAP) estimates or, the posterior means. Although we obtain point estimates, these estimates do not yield exact zero values since we have considered a continuous spike-and-slab prior with a spike at ν_0 (a small positive number). Additionally, these estimates do not make use of all the MCMC samples. We use a Bayesian model averaging-based strategy that leverages all the MCMC samples to build inference [see, for e.g., 32]. Specifically, we use a FDR-based strategy—combining Bayesian model averaging with point estimates [see for e.g., 48, 47]. Let $\beta_u^{(m)}$ for $m = 1, \dots, M$ denote the MCMC samples (after burn-in and thinning) for the coefficients of the mean model. We compute $p_u = \frac{1}{M} \sum_m I(|\beta_u^{(m)}| \leq c)$, where $I(\cdot)$ is the indicator function; these probabilities p_u can be interpreted as local FDR [see, for e.g., 48]. The probability $(1 - p_u)$ can be interpreted as the probability that covariate u is significantly associated with the response. We use the p_u s to decide on which covariates to select while controlling the FDR at level α . Explicitly, we infer that the covariate u has a *non-zero* coefficient if $p_u < \kappa_\alpha$ for some threshold $\kappa_\alpha \in (0, 1)$. We compute the threshold κ_α as follows: We first sort the probabilities p_u and denote the sorted probabilities as $p_{(u)}$ for $u = 1, \dots, p$. We then assign $\kappa_\alpha = p_{(u^*)}$, where $u^* = \max\{\tilde{u} \mid \frac{1}{u} \sum_{u=1}^{\tilde{u}} p_{(u)} \leq \alpha\}$. This approach caps our false discoveries of selected variables at $100\alpha\%$. We employ the same approach using the MCMC samples of γ to select significant coefficients for the dispersion models.

Posterior inference on β , γ is performed using MAP (point) estimates along with posterior mean, median, stan-

dard deviation and highest posterior density (HPD) intervals. For \mathbf{w} we employ the mean, median, standard deviation and HPD intervals to perform posterior inference. Next, we demonstrate some synthetic experiments that document the performance of our proposed models using the discussed metrics. The computation has been performed in the R statistical environment. The required subroutines can be accessed via an open-source repository at: <https://github.com/arh926/sptwdglm>.

3. SYNTHETIC EXPERIMENTS

We begin with an observation—the spatial heterogeneity that our models aim to quantify is not observed in real life. Hence, it is imperative to document the accuracy of estimating such effects through synthetic experiments. Settings used are outlined—we consider varying proportion of zeros (15%, 30%, 60%, 80% and 95%) under which the quality of posterior inference for θ is assessed. Proportion of zeros can be interpreted as an inverse signal-to-noise ratio for the synthetic response. For the sake of brevity, we only show the results for synthetic experiments pertaining to Bayesian variable selection in the presence of spatial effects. Additional simulations can be found in the online Supplement. To construct the synthetic data we consider three scenarios pertaining to *model structure*, (a) there is no overlap (i.e. selected β 's and γ 's do not intersect) (b) 50% overlap (in the union of all selected variables across the mean and dispersion models) (c) 100% overlap between mean and dispersion model specification. We use 10 covariates including an intercept, where the columns of the synthetic design matrices \mathbf{X} and \mathbf{Z} are hierarchically centered and scaled, independently sampled Gaussian variables with mean 0 and variance 1. Naturally, specification of the true β , \mathbf{w} and γ parameters determine the proportion of zeros in the synthetic response. Table 2 contains the parameter specifications used. The true value of the index parameter, $\xi = 1.5$. In an attempt to produce a synthetic setup that resembles reality we simulate, $\mathbf{w} \sim N(5(\sin(3\pi s_1) + \cos(3\pi s_2)), 1)$ (see Figure 1, second row). The alternative route would be to fix values of σ^2 and ϕ_s and generate a realization $\mathbf{w} \sim N_L(\mathbf{0}_L, \sigma^2 \mathbf{R}(\phi_s))$ (see Figure 1, first row). Under each setting we consider $M = 10$ replications. Within each replication we fit all the proposed modeling frameworks as shown in Table 1. The hyper-parameter settings used while specifying priors for the models are, $a_\xi = 1$, $b_\xi = 2$, $a_\sigma = a_{\sigma_\beta} = a_{\sigma_\gamma} = 2$,

Table 1. Proposed Bayesian Hierarchical Double Generalized Linear Modeling Frameworks.

Models	Frameworks	Specification (θ)	Number of Parameters
M1	DGLM	θ_m	$p + q + 1$
M2	DGLM + Variable Selection	θ_m, θ_{vs}	$3p + 3q + 1$
M3	DGLM + Spatial Effect	θ_m, θ_{pr}	$p + q + L + 4$
M4	DGLM + Spatial Effect + Variable Selection	$\theta_m, \theta_{vs}, \theta_{pr}$	$3p + 3q + L + 4$

Table 2. Parameter settings used to obtain varying proportion of zeros in the synthetic data.

Proportion of 0s	$\mu_{\beta_{-0}} (\sigma_{\beta_{-0}})$	γ_0	$\mu_{\gamma_{-0}} (\sigma_{\gamma_{-0}})$
15%	0.50 (0.1)	-1.50	0.50 (0.1)
30%	0.50 (0.1)	0.70	0.50 (0.1)
60%	0.50 (0.1)	2.50	0.50 (0.1)
80%	1.00 (0.1)	4.50	0.50 (0.1)
95%	1.00 (0.1)	7.00	0.50 (0.1)

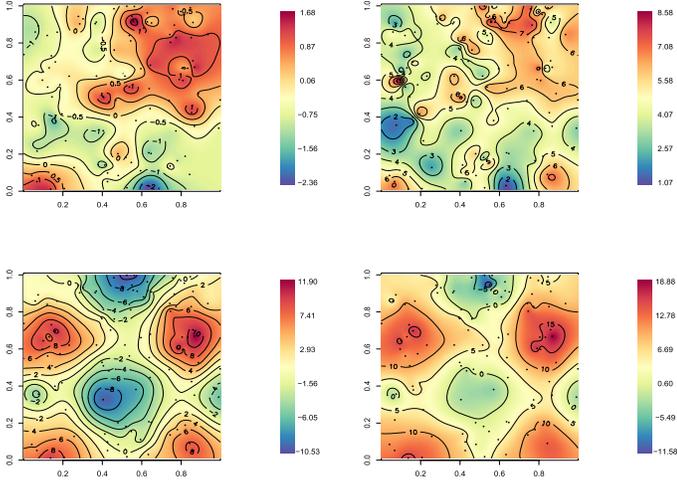


Figure 1: Plots showing synthetic spatial patterns, pattern 1 (top, left column) and pattern 2 (bottom, left column) and corresponding logarithm of aggregated synthetic response (right column).

$b_\sigma = b_{\sigma_\beta} = b_{\sigma_\gamma} = 1$ (producing inverse-Gamma priors with mean 1 and infinite variance), $a_{\phi_s} = 0$, $b_{\phi_s} = 30$, $\sigma_\beta^{-2} = \sigma_\gamma^{-2} = 10^{-6}$ and $\nu_0 = 5 \times 10^{-4}$, producing a vague and non-informative hierarchical prior. We maintain an FDR of 5% for all settings while performing model selection. The sample size varies from $N = \{2 \times 10^3, 5 \times 10^3, 1 \times 10^4\}$ and the number of locations are $L = 1 \times 10^2$. Across replications, false positive rate (FPR) and true positive rate (TPR) are computed to measure accuracy of our model selection procedure. To record the quality of estimation, we compute the mean squared error (MSE), for e.g. $MSE(\beta) = \frac{1}{p} \sum_{i_\beta=1}^p (\beta_{i_\beta} - \hat{\beta}_{i_\beta})^2$, which can be computed similarly for the other parameters. We also compute average coverage probabilities, for e.g. considering these probabilities for β we define $CP(\beta) = \frac{1}{M} \sum_{m=1}^M I(\beta_{true} \in (l_m(\beta), u_m(\beta)))$, where $l_m(\beta)$ and $u_m(\beta)$ are the lower and upper 95% HPD intervals respectively for β in replication m ; we obtain coverage probabilities for \mathbf{w} and γ similarly. The results obtained under the above settings are shown in Table 3. The first column is named configuration (abbreviated as config.) with entries denoting the proportion of overlap between selected coefficients in the mean and dispersion models, which is indicative of model structure. This is estimated by ob-

serving the overlap between selected variables following the model fit (for models M2 and M4). No variable selection is performed for models M1 and M3.

From the results shown, we see that models M1 and M2 perform poorly. Estimates, $\hat{\beta}$ remain fairly unaffected as compared to $\hat{\gamma}$ and $\hat{\xi}$, where all of the variation not quantified, yet present in the synthetic data spills over to corrupt and compromise the quality of estimates. This also does not produce reliable results pertaining to model structure recovery for M1 and M2. However, significant improvements show up with M3 and M4. Particularly, under higher proportion of zeros in the synthetic data (low signal to noise ratio) the performance of M4 remains stable with respect to model structure recovery and estimation of parameters (refer to Table 3), thereby producing *robust inference* among the models in comparison. As an example within our simulation setting, under 95% of 0s in the data and under low sample sizes, for example 2000 or, 5000, the estimates of model coefficients and spatial effects in M3 and M4 are adversely affected by locations having fewer non-zero observations. This observation addresses the concern around specifying DGLMs without spatial random effects in a scenario where the data displays spatial variation. The results demonstrate expected gains when our model in its full capacity is used instead of an usual DGLM.

We use the MCMC algorithm featuring MALA updates for β, \mathbf{w} and γ . Chain lengths are set to 1×10^4 , with the initial 5,000 samples as burn-in and thin the rest by selecting every 10-th sample which reduces any remaining auto-correlation and produces 500 independent posterior samples for each setting. The posterior estimate, $\hat{\theta}$ is then obtained using the produced samples by computing the median or a MAP estimate as applicable for the model. Coverage probabilities for model M4 remained sufficiently high (≈ 1) across all settings; only declining marginally for \mathbf{w} (remaining above 90%) under high proportions of zeros (low signal to noise ratio) in the data. We performed additional synthetic experiments to showcase (a) the performance of M3 with respect to the quality of estimation for spatial effects and (b) the performance of M2. They are detailed in the Supplementary materials—we briefly outline its contents in the next section.

4. SUPPLEMENTARY ANALYSIS

The online Supplement to this paper contains details on the derivations of the posteriors essential for constructing MCMC subroutines. They are outlined in Section S1. Section S2 features additional simulation experiments that supplement those outlined previously in Section 3. It documents performance of M2, shown in Table S1, contains results of experiments for scenarios featuring spatial covariates, shown in Table S2, and varying spatial patterns as seen in Figure 1, shown in Tables S3 and S4. Convergence diagnostics are shown for selected model parameters (index parameter ξ)

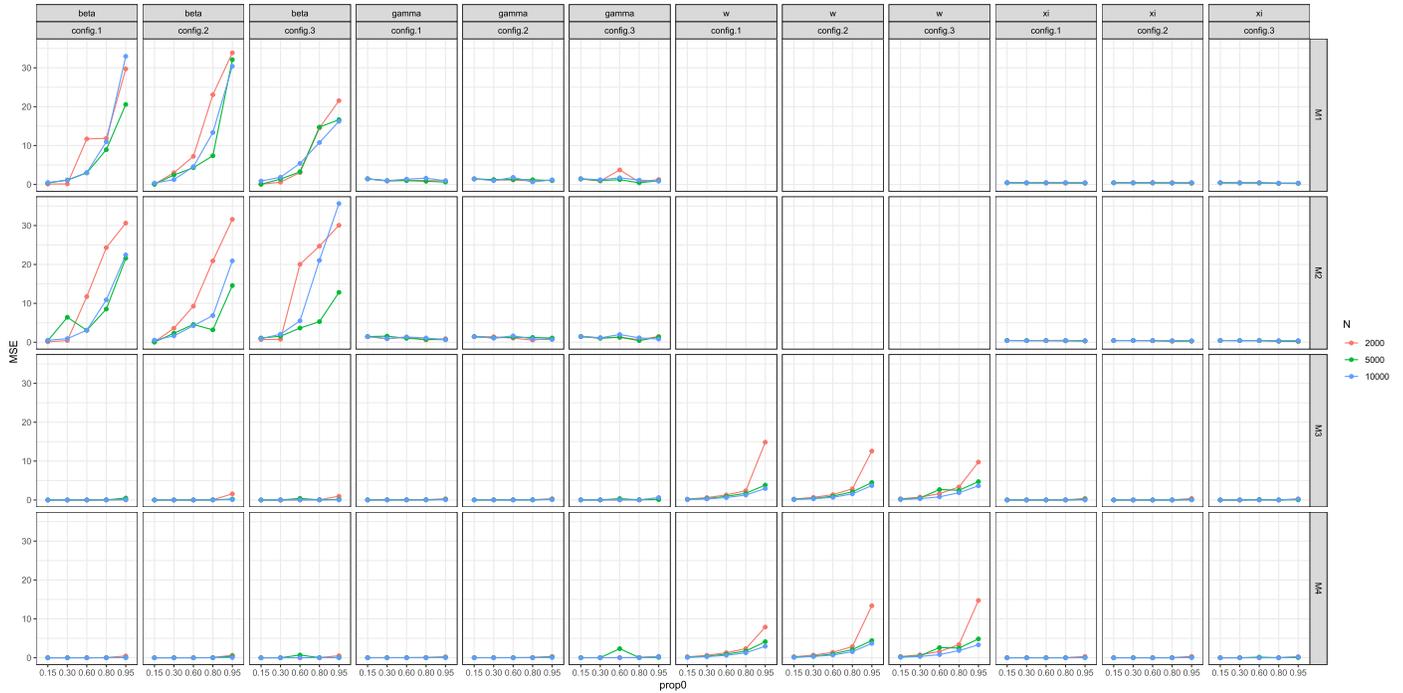


Figure 2: Results for synthetic experiments showing model performance, where MSE is plotted against proportion of zeros in the synthetic response which is tabulated across models (M1, M2, M3, M4) vs. $\{\theta_m, \theta_{pr}\} \times$ configuration.

in Section S2.2. Contents of the R-package are described in Section S2.3. Finally, results for models M1 and M3 pertaining to the real data analysis described in the next section appear in Section S3, Tables S5, S6, S7 and S8.

5. AUTOMOBILE INSURANCE PREMIUMS, CONNECTICUT, 2008

We analyze automobile insurance premiums for collision coverage in the state of Connecticut, for the year 2008. The data is obtained from a comprehensive repository named Highway Loss Data Institute (HLDI) maintained by the independent non-profit, Insurance Institute for Highway Safety (IIHS) (<https://www.iihs.org/>) working towards reducing losses arising from motor vehicle collisions in North America. We briefly describe the variables contained in HLDI data. It records covariate information at three levels for an insured individual. They are as follows,

- a. *Individual Level:* (i) accident and model year of the vehicle, (ii) age, gender, marital status.
- b. *Policy level:* (i) policy payments, measured in United States dollars, (ii) exposure—measured in policy years, for e.g. 0.5 indicates a coverage period of 6 months or, half a year, measured in years, (iii) policy risk—having two levels, which is assigned by the insurance company based on provided information by the individual, (iv) deductible limit—with 8 categories.
- c. *Spatial:* 5-digit zip code.

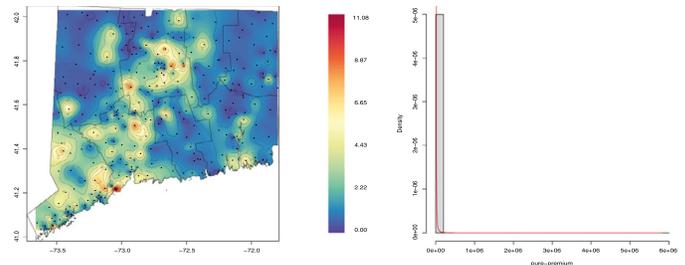


Figure 3: (left) Spatial plot of zip-code level aggregated pure-premium $\times 10^{-6}$ for the state of Connecticut, 2008 (right) histogram for the pure-premium overlaid with a probability density estimate (in red).

Derived variables like age categories, vehicle age in years and interactions like gender \times marital status are computed and used as covariates in the model. For the state of Connecticut, 1,513,655 (≈ 1.5 million) data records were obtained in the year 2008, at 281 zip-codes. Zip-codes are areal units, we consider the latitude-longitude corresponding to the centroid of each zip code as the point reference counterpart unit for our application purposes. Distance between two zip-codes is then specified as the Euclidean distance between their centroids. The proportion of zeros in the payments is 95.73%. From an insurer’s perspective, policy rate-making is the problem of assigning policy-premium to a new customer’s policy based on their covariate information (for in-

Table 3. Table showing results of synthetic experiments for model selection for models M2 and M4. Corresponding standard deviations are shown in brackets below.

N	True Overlap	Prop. of 0s	M2			M4		
			Overlap	FPR	TPR	Overlap	FPR	TPR
5000	0.00	0.15	0.03 (0.06)	0.04 (0.06)	0.69 (0.07)	0.00 (0.00)	0.00 (0.00)	0.90 (0.00)
		0.30	0.05 (0.09)	0.05 (0.09)	0.89 (0.10)	0.00 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.60	0.01 (0.04)	0.01 (0.04)	0.98 (0.04)	0.00 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.80	0.04 (0.06)	0.04 (0.06)	0.99 (0.03)	0.00 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.95	0.10 (0.20)	0.06 (0.10)	0.89 (0.20)	0.08 (0.10)	0.04 (0.11)	0.95 (0.08)
		0.15	0.24 (0.16)	0.01 (0.05)	0.67 (0.09)	0.50 (0.00)	0.00 (0.00)	0.90 (0.03)
	0.50	0.30	0.49 (0.08)	0.01 (0.05)	0.94 (0.06)	0.50 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.60	0.51 (0.03)	0.00 (0.00)	0.93 (0.06)	0.50 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.80	0.56 (0.09)	0.07 (0.08)	0.97 (0.04)	0.49 (0.02)	0.01 (0.05)	1.00 (0.00)
		0.95	0.65 (0.31)	0.20 (0.17)	0.85 (0.22)	0.48 (0.10)	0.04 (0.15)	0.92 (0.18)
		0.15	0.40 (0.25)	0.03 (0.05)	0.66 (0.07)	1.00 (0.00)	0.00 (0.00)	0.90 (0.00)
		0.30	0.90 (0.14)	0.04 (0.08)	0.86 (0.08)	1.00 (0.00)	0.00 (0.00)	1.00 (0.00)
	1.00	0.60	1.00 (0.00)	0.00 (0.00)	0.96 (0.05)	1.00 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.80	1.00 (0.00)	0.00 (0.00)	1.00 (0.00)	1.00 (0.00)	0.00 (0.00)	1.00 (0.00)
		0.95	0.35 (0.21)	0.15 (0.16)	0.85 (0.22)	0.89 (0.21)	0.10 (0.11)	0.89 (0.10)

stance, individual level, policy and residence zip-code). We achieve this via out-sample prediction. To that end, we consider a 60-40 split for the data, the split is performed by using stratified sampling without replacement over zip-codes, such that the same 281 zip-codes are also available in the out-sample. The training data then contains $N_{tr} = 908,741$ observations with $N_{pr} = 604,914$ observations kept in reserve for prediction constituting the out-sample data.

We denote payments towards a policy made by individual i , residing in zip-code j as y_{ij} with an exposure of t_{ij} . We assume that the policy-premium defined as $y_{ij}^* = \frac{y_{ij}}{t_{ij}} \sim Tw(\mu_{ij}, \phi_{ij}, \xi)$, which implies $y_{ij} \sim Tw(t_{ij}\mu_{ij}, t_{ij}^{2-\xi}\phi_{ij}, \xi)$ using the scale invariance property. The following hierarchical DGLM is then specified,

$$\begin{aligned} \log \mu_{ij}(\mathbf{s}_i) &= -\log t_{ij} + \mathbf{x}_{ij}^T(\mathbf{s}_i)\boldsymbol{\beta} + \mathbf{f}_{ij}(\mathbf{s}_i)^T \mathbf{w}(\mathbf{s}_i), \\ \log \phi_{ij} &= -(2 - \xi) \log t_{ij} + \mathbf{z}_{ij}^T \boldsymbol{\gamma}, \end{aligned} \tag{5.1}$$

when considered with a spatial specification, where the terms $-\log t_{ij}$ and $-(2 - \xi) \log t_{ij}$ act as offsets for the

respective mean and dispersion models. Given the covariates described at the beginning $p = q = 29$, producing a $N_{tr} \times (p - 1)$ design matrix for the mean model and a $N_{tr} \times q$ design matrix for the dispersion model. The model in (5.1) specifies model M3 from Table 1, M4 is obtained by specifying $\pi(\boldsymbol{\theta}_{vs})$ from (2.5) on $\boldsymbol{\theta}_m$. M1 is obtained by setting $f_{ij}(\mathbf{s}_i) = \mathbf{0}$ and M2 is obtained by specifying $\pi(\boldsymbol{\theta}_{vs})$ on $\boldsymbol{\theta}_m$ for the resulting model. For M1 and M2 we include an intercept in the mean model. We fit models M1-4 on the training data. Model selection is performed using FDR based variable selection on the posterior MCMC samples obtained from fitting models M2 and M4, controlling for FDR at 1%. The performance of M1-4 is assessed using the Akaike Information Criteria (AIC) [3]. While specifying (2.4) and (2.5) the hyper-parameter settings used are, $\sigma_\beta^2 = \sigma_\gamma^2 = 10^6$, $a_{\sigma_\beta} = a_\sigma = a_{\sigma_\gamma} = 2$ and $b_{\sigma_\beta} = b_\sigma = b_{\sigma_\gamma} = 1$ (generating inverse-gamma priors with mean 1 and infinite variance), $a_{\phi_s} = 0$, $b_{\phi_s} = 60$. We maintain $a_\xi = 1$ and $b_\xi = 2$ for all models. For ease of implementation, the fractal parameter, ν is fixed at 0.5, producing the exponential covariance kernel. We consider 1×10^5 MCMC iterations for gener-

Table 4. Estimated coefficients for fixed effects corresponding to model M2. We show the MAP, median, mean, standard deviation and HPD intervals. The variables listed are a result of FDR-based variable selection at 1%.

	Parameters	Levels	MAP	Median	Mean	SD	Lower HPD	Upper HPD
β	(Intercept)	–	5.815	5.965	5.955	0.151	5.706	6.227
	age.car	1	0.113	0.112	0.110	0.033	0.044	0.173
		2	0.234	0.239	0.240	0.030	0.186	0.304
		5	0.124	0.122	0.121	0.033	0.062	0.189
	risk	S	–0.213	–0.217	–0.217	0.023	–0.258	–0.168
	agec	3	–0.399	–0.399	–0.400	0.030	–0.454	–0.341
		4	–0.518	–0.515	–0.515	0.035	–0.580	–0.444
		5	–0.712	–0.710	–0.708	0.065	–0.832	–0.575
	gender	F	0.493	0.502	0.513	0.067	0.400	0.648
		M	0.608	0.614	0.619	0.053	0.521	0.731
	marital	M	–0.376	–0.392	–0.403	0.067	–0.577	–0.302
	deductible	B	0.898	1.051	1.104	0.244	0.740	1.525
		E	0.616	0.482	0.478	0.158	0.192	0.741
		F	0.580	0.428	0.435	0.154	0.135	0.656
		G	0.287	0.348	0.358	0.159	0.073	0.633
γ	(Intercept)	–	7.354	7.346	7.345	0.048	7.249	7.435
	age.car	0	–0.820	–0.811	–0.811	0.053	–0.911	–0.714
		1	–1.024	–1.017	–1.016	0.051	–1.115	–0.922
		2	–0.888	–0.874	–0.873	0.052	–0.969	–0.776
		3	–0.864	–0.854	–0.851	0.052	–0.953	–0.760
		4	–0.843	–0.838	–0.838	0.052	–0.936	–0.743
		5	–0.788	–0.781	–0.780	0.052	–0.877	–0.682
		6	–0.770	–0.765	–0.765	0.053	–0.862	–0.665
	risk	7	–0.730	–0.723	–0.723	0.053	–0.829	–0.627
		S	0.114	0.114	0.114	0.011	0.093	0.135
	agec	5	–0.289	–0.293	–0.293	0.029	–0.345	–0.234
	deductible	B	–0.971	–0.963	–0.951	0.099	–1.136	–0.759
		C	–0.519	–0.522	–0.517	0.049	–0.614	–0.414
		D	–0.565	–0.568	–0.564	0.044	–0.657	–0.470
		E	–0.506	–0.501	–0.495	0.042	–0.568	–0.394
F		–0.348	–0.346	–0.340	0.039	–0.410	–0.242	
H		0.411	0.419	0.426	0.097	0.251	0.618	
genderMarital	A	–0.113	–0.106	–0.097	0.043	–0.168	–0.001	
ξ	–	–	1.673	1.673	1.673	0.001	1.671	1.676

ating samples from respective posteriors with burn-in diagnosed at 5×10^4 and include every 20-th sample to compute posterior estimates as our thinning strategy. Convergence is assessed through inspecting trace plots. Proposal variances were scaled in an adaptive fashion to provide optimal acceptance rates of 58% (MALA) and 33% (MH). Predictive performance for models M1–4 is judged based on square root deviance on the out-sample data. The results are shown in Table 6. Optimal values are marked in bold. We show the results for estimated model coefficients featuring Bayesian variable selection (models M2 and M4) in Tables 4 and 5. Results for M1 and M3 are postponed to Tables S5, S6, S7 and S8 in the Supplementary Materials. Posterior estimates for the spatial effects in models M3 and M4 are showed in

Figure 4. Zip-codes with significant effects are color coded appropriately.

Comparing the models we observe that M4 produces the most optimal model fit criteria among the models considered. This extends to out-sample performance when predicting policy premiums. Plots produced for spatial effects in models M3 and M4 are mean adjusted. Since specification of M3 and M4 differ only in presence/absence of the hierarchical Bayesian variable selection component, the produced spatial effects mimic each other after adjusting for the mean. Comparing results in Tables 4 and 5 we observe that including the spatial effect results in more categories for vehicle age, driver age and deductible being selected. Overall, we observe that the findings remain consistent with our earlier

Table 5. Estimated coefficients for fixed effects corresponding to model M4. We show the MAP, median, mean, standard deviation and HPDs. The variables listed are a result of FDR-based variable selection at 1%.

	Parameters	Levels	MAP	Median	Mean	SD	Lower HPD	Upper HPD	
β	(Intercept)	–	5.215	5.213	5.216	0.135	5.165	5.327	
	age.car	0	0.252	0.251	0.251	0.063	0.123	0.377	
		1	0.329	0.317	0.318	0.058	0.217	0.431	
		2	0.424	0.423	0.425	0.061	0.318	0.549	
		3	0.209	0.193	0.195	0.060	0.084	0.318	
		4	0.236	0.252	0.255	0.062	0.149	0.387	
		5	0.297	0.302	0.307	0.059	0.191	0.428	
		6	0.182	0.209	0.214	0.061	0.113	0.343	
	7	0.155	0.154	0.157	0.063	0.049	0.286		
	risk	S	–0.190	–0.184	–0.183	0.023	–0.226	–0.137	
	agec	2	–0.121	–0.126	–0.128	0.030	–0.186	–0.068	
		3	–0.461	–0.467	–0.470	0.031	–0.531	–0.412	
		4	–0.608	–0.617	–0.619	0.034	–0.685	–0.554	
		5	–0.721	–0.699	–0.696	0.065	–0.808	–0.561	
	gender	F	0.535	0.540	0.541	0.066	0.407	0.679	
		M	0.670	0.705	0.720	0.100	0.535	0.919	
	deductible	B	1.337	1.383	1.408	0.216	1.036	1.866	
		C	0.209	0.256	0.303	0.152	0.073	0.624	
		D	0.603	0.627	0.672	0.161	0.397	0.982	
		E	0.800	0.825	0.864	0.157	0.625	1.189	
F		0.753	0.774	0.819	0.155	0.592	1.154		
G		0.756	0.786	0.825	0.158	0.580	1.169		
H		0.820	0.829	0.824	0.190	0.463	1.149		
genderMarital		B	–0.368	–0.236	–0.169	0.233	–0.457	0.311	
γ	(Intercept)	–	6.423	6.429	6.415	0.073	6.310	6.504	
	age.car	0	–0.790	–0.809	–0.811	0.040	–0.891	–0.731	
		1	–1.006	–1.025	–1.028	0.039	–1.099	–0.954	
		2	–0.875	–0.885	–0.889	0.039	–0.972	–0.819	
		3	–0.842	–0.861	–0.863	0.039	–0.950	–0.799	
		4	–0.831	–0.844	–0.848	0.039	–0.923	–0.777	
		5	–0.781	–0.792	–0.795	0.039	–0.879	–0.728	
		6	–0.762	–0.766	–0.769	0.039	–0.849	–0.703	
	7	–0.725	–0.732	–0.735	0.039	–0.814	–0.662		
	risk	S	0.110	0.110	0.110	0.012	0.087	0.132	
	agec	5	–0.262	–0.262	–0.261	0.030	–0.316	–0.201	
	deductible	B	–1.023	–1.026	–1.029	0.158	–1.328	–0.721	
		C	–0.534	–0.553	–0.571	0.069	–0.700	–0.461	
		D	–0.591	–0.611	–0.634	0.067	–0.760	–0.534	
		E	–0.533	–0.552	–0.575	0.066	–0.702	–0.478	
		F	–0.377	–0.390	–0.416	0.065	–0.543	–0.335	
		H	0.352	0.361	0.369	0.107	0.170	0.591	
		ξ	–	1.667	1.667	1.667	0.001	1.665	1.670

research [see, 29] but within a more robust proposed model choice framework. This is evident when comparing model estimates between Table 5 and Table S7 and S8 in the Supplement. The marital status and interaction between gender and marital status is not selected. We conclude by observing that the spatial effects are significantly positive in major cities in Connecticut, indicating a higher spatial risk, as op-

posed to sparsely populated regions showing significantly lower risk.

6. DISCUSSION

Double generalized linear models have not seen much use after their inception by [40]. Hindrances presented by ambiguities existing around model specification/choice have

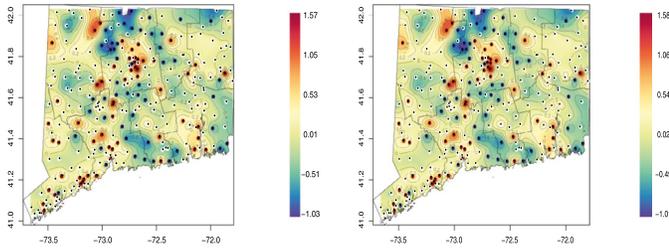


Figure 4: Spatial plot showing the posterior estimates for 281 zipcodes from (left) M3, and (right) M4 in Connecticut. The locations are color coded based on significance, with white indicating a location with 0 in its HPD interval, blue (red) indicating HPD interval with both endpoints negative (positive).

Table 6. Table showing AIC and out-sample square root deviance for models M1–M4.

	M1	M2	M3	M4
AIC	1340060	1117733	1117114	1115363
$\sqrt{\text{Deviance}}$	5565.209	5509.549	5507.926	5441.230

been addressed in this paper. We propose Bayesian modeling frameworks that perform model selection using continuous spike and slab priors for hierarchical double generalized linear Tweedie spatial process models. Leveraging Langevin dynamics we are able to successfully produce practical implementations for the proposed frameworks which would otherwise remain unachievable with standard MCMC techniques. The proposed algorithms are available as a publicly accessible package for the R–statistical environment. Although the formulation considers the CP-g densities, evidently such modeling could be effected under any probabilistic framework that allows for varying dispersion. The application offers some key insights into the actuarial domain. It is generally believed that marital status and gender play a key role. However, the model inference suggests otherwise, with marital status not being selected as a significant feature.

Future work is aimed at extending this framework in multiple directions. Firstly, with the advent of modern Bayesian variable selection priors—for example, the Bayesian Lasso, the Horseshoe prior etc., a comparative model selection performance remains to be seen when considered within hierarchical DGLM formulations. Secondly, with the emerging techniques for handling large spatial and spatiotemporal data [see, for e.g., 31] the DGLM framework could be extended to model spatially or, spatio-temporally indexed observations over massive geographical domains. With respect to our application, this would allow us to investigate properties of the premium surface over much larger domains, for instance a country-wide study. Finally, extending these models to a spatiotemporal setting could be achieved using spatiotemporal covariance kernels that are commonly used.

Depending on the nature of spatial and temporal interaction, we can have separable and non-separable kernels at our disposal [see, 15, and references therein]. Bayesian variable selection could then be effected to examine resulting changes in model specification upon inclusion of random effects that address spatiotemporal variation in the data.

SUPPLEMENTARY MATERIAL

Supplementary Material containing further details as described in Section 4 is available online. The R–package is available for installation and deployment at: <https://github.com/arh926/sptwdglm>.

ACKNOWLEDGEMENTS

The authors would like to thank the Editor and an anonymous Reviewer for suggestions that improved the paper. They would also like to thank Brien Aronov and Keith Holler for their insight during early stages of this work.

Aritra Halder would like to thank Sudipto Banerjee for discussions that improved the writing and presentation significantly. He extends his gratitude to Sallie Keller and Stephanie S. Shipp from the Biocomplexity Institute and Initiative, University of Virginia for their valuable support.

FUNDING

Shariq Mohammed was supported by institutional research funds from Boston University (BU) School of Public Health and Rafik B. Hariri Institute for Computing and Computational Science & Engineering at BU.

Accepted 31 May 2023

REFERENCES

- [1] ABRAMOWITZ, M., STEGUN, I. A. and ROMER, R. H. (1988). *Handbook of mathematical functions with formulas, graphs, and mathematical tables*. American Association of Physics Teachers. MR0415962
- [2] AGARWAL, D. K., GELFAND, A. E. and CITRON-POUSTY, S. (2002). Zero-inflated models with application to spatial count data. *Environmental and Ecological statistics* 9(4) 341–355. <https://doi.org/10.1023/A:1020910605990>. MR1951713
- [3] AKAIKE, H. (1998). Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike* 199–213 Springer. MR1486823
- [4] BANERJEE, S. and CARLIN, B. P. (2004). Parametric spatial cure rate models for interval-censored time-to-relapse data. *Biometrics* 60(1) 268–275. <https://doi.org/10.1111/j.0006-341X.2004.00032.x>. MR2044123
- [5] BANERJEE, S., CARLIN, B. P. and GELFAND, A. E. (2014). Hierarchical Modeling and Analysis for Spatial Data. MR3362184
- [6] BERGER, J. O., DE OLIVEIRA, V. and SANSÓ, B. (2001). Objective Bayesian analysis of spatially correlated data. *Journal of the American Statistical Association* 96(456) 1361–1374. <https://doi.org/10.1198/016214501753382282>. MR1946582
- [7] BERGER, J. O., PERICCHI, L. R., GHOSH, J., SAMANTA, T., DE SANTIS, F., BERGER, J. and PERICCHI, L. (2001). Objective Bayesian methods for model selection: Introduction and comparison. *Lecture Notes-Monograph Series* 135–207. <https://doi.org/10.1214/lmns/1215540968>. MR2000753

- [8] BERLINER, M. (2000). Hierarchical Bayesian modeling in the environmental sciences. *ASA Advances in Statistical Analysis* **2**(84) 141–153. <https://doi.org/10.1214/06-BA130>. MR2282211
- [9] BEST, N. G., ICKSTADT, K. and WOLPERT, R. L. (2000). Spatial Poisson regression for health and exposure data measured at disparate resolutions. *Journal of the American statistical association* **95**(452) 1076–1088. <https://doi.org/10.2307/2669744>. MR1821716
- [10] BRADLEY, J. R., HOLAN, S. H. and WIKLE, C. K. (2018). Computationally efficient multivariate spatio-temporal models for high-dimensional count-valued data (with discussion). *Bayesian Analysis* **13**(1) 253–310. <https://doi.org/10.1214/17-BA1069>. MR3773410
- [11] BRADLEY, J. R., HOLAN, S. H. and WIKLE, C. K. (2020). Bayesian hierarchical models with conjugate full-conditional distributions for dependent data from the natural exponential family. *Journal of the American Statistical Association* **115**(532) 2037–2052. <https://doi.org/10.1080/01621459.2019.1677471>. MR4189775
- [12] CARLIN, B. P. and LOUIS, T. A. (2008) *Bayesian methods for data analysis*. CRC press. MR2442364
- [13] CARVALHO, C. M., POLSON, N. G. and SCOTT, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika* **97**(2) 465–480. <https://doi.org/10.1093/biomet/asq017>. MR2650751
- [14] CLARK, J. and GELFAND, A. (2006). A future for models and data in ecology. *Trends in Ecology and Evolution* **21** 375–380.
- [15] CRESSIE, N. (2015) *Statistics for spatial data*. John Wiley & Sons. MR3559472
- [16] DEY, D. K., GHOSH, S. K. and MALLICK, B. K. (2000) *Generalized linear models: A Bayesian perspective*. CRC Press. MR1893779
- [17] DIGGLE, P. J., TAWN, J. A. and MOYEED, R. A. (1998). Model-based geostatistics. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **47**(3) 299–350. <https://doi.org/10.1111/1467-9876.00113>. MR1626544
- [18] DUNN, P. K. and SMYTH, G. K. (2005). Series evaluation of Tweedie exponential dispersion model densities. *Statistics and Computing* **15**(4) 267–280. <https://doi.org/10.1007/s11222-005-4070-y>. MR2205390
- [19] DUNN, P. K. and SMYTH, G. K. (2008). Evaluation of Tweedie exponential dispersion model densities by Fourier inversion. *Statistics and Computing* **18**(1) 73–86. <https://doi.org/10.1007/s11222-007-9039-6>. MR2416440
- [20] EIDSVIK, J., FINLEY, A. O., BANERJEE, S. and RUE, H. (2012). Approximate Bayesian inference for large spatial datasets using predictive process models. *Computational Statistics & Data Analysis* **56**(6) 1362–1380. <https://doi.org/10.1016/j.csda.2011.10.022>. MR2892347
- [21] FINLEY, A. O., BANERJEE, S. and McROBERTS, R. E. (2009). Hierarchical spatial models for predicting tree species assemblages across large domains. *The annals of applied statistics* **3**(3) 1052. <https://doi.org/10.1214/09-AOAS250>. MR2750386
- [22] GELFAND, A. E., SAHU, S. K. and CARLIN, B. P. (1996). Efficient Parametrizations for Generalized Linear Mixed Models. *Bayesian Statistics 5: Proceedings of the Fifth Valencia International Meeting* 165–180. MR1425405
- [23] GELFAND, A. E. (2000). Modeling and Inference for Point-Referenced Binary Spatial Data. *Generalized linear models: a Bayesian perspective* 373. MR1893801
- [24] GELFAND, A. E. and BANERJEE, S. (2017). Bayesian modeling and analysis of geostatistical data. *Annual review of statistics and its application* **4** 245–266.
- [25] GELFAND, A. E., SAHU, S. K. and CARLIN, B. P. (1995). Efficient parametrizations for normal linear mixed models. *Biometrika* **82**(3) 479–488. <https://doi.org/10.1093/biomet/82.3.479>. MR1366275
- [26] GELFAND, A. E., SCHMIDT, A. M., WU, S., SILANDER JR, J. A., LATIMER, A. and REBELO, A. G. (2005). Modelling species diversity through species level hierarchical modelling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **54**(1) 1–20. <https://doi.org/10.1111/j.1467-9876.2005.00466.x>. MR2134594
- [27] GEORGE, E. I. and McCULLOCH, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association* **88**(423) 881–889.
- [28] GIROLAMI, M. and CALDERHEAD, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(2) 123–214. <https://doi.org/10.1111/j.1467-9868.2010.00765.x>. MR2814492
- [29] HALDER, A., MOHAMMED, S., CHEN, K. and DEY, D. K. (2021). Spatial Tweedie exponential dispersion models: an application to insurance rate-making. *Scandinavian Actuarial Journal* **2021**(10) 1017–1036. <https://doi.org/10.1080/03461238.2021.1921017>. MR4345874
- [30] HALDER, A., MOHAMMED, S., CHEN, K. and DEY, D. K. (2022). Spatial Risk Estimation in Tweedie Double Generalized Linear Models. *Proceedings of International E-Conference on Mathematical and Statistical Sciences: A Selcuk Meeting* **2022** 62.
- [31] HEATON, M. J., DATTA, A., FINLEY, A. O., FURRER, R., GUINNESS, J., GUHANIYOGI, R., GERBER, F., GRAMACY, R. B., HAMMERLING, D., KATZFUSS, M. et al. (2019). A Case Study Competition Among Methods For Analyzing Large Spatial Data. *Journal of Agricultural, Biological and Environmental Statistics* **24**(3) 398–425.
- [32] HOETING, J. A., MADIGAN, D., RAFTERY, A. E. and VOLINSKY, C. T. (1999). Bayesian model averaging: a tutorial (with comments by M. Clyde, David Draper and EI George, and a rejoinder by the authors. *Statistical science* **14**(4) 382–417. <https://doi.org/10.1214/ss/1009212519>. MR1765176
- [33] ISHWARAN, H. and RAO, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics* **33**(2) 730–773. <https://doi.org/10.1214/009053604000001147>. MR2163158
- [34] JØRGENSEN, B. (1986). Some properties of exponential dispersion models. *Scandinavian Journal of Statistics* 187–197. MR0873073
- [35] JØRGENSEN, B. (1987). Exponential dispersion models. *Journal of the Royal Statistical Society: Series B (Methodological)* **49**(2) 127–145. MR0905186
- [36] JØRGENSEN, B. (1992). Exponential dispersion models and extensions: A review. *International Statistical Review/Revue Internationale de Statistique* 5–20.
- [37] JØRGENSEN, B. (1997) *The theory of dispersion models*. CRC Press. MR1462891
- [38] KOKONENDJI, C. C., BONAT, W. H. and ABID, R. (2021). Tweedie regression models and its geometric sums for (semi-) continuous data. *Wiley Interdisciplinary Reviews: Computational Statistics* **13**(1) 1496. <https://doi.org/10.1002/wics.1496>. MR4186771
- [39] LAWSON, A. B. (2018) *Bayesian disease mapping: hierarchical modeling in spatial epidemiology*. Chapman and Hall/CRC. MR2484272
- [40] LEE, Y. and NELDER, J. A. (2006). Double hierarchical generalized linear models (with discussion). *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **55**(2) 139–185. <https://doi.org/10.1111/j.1467-9876.2006.00538.x>. MR2226543
- [41] LI, Q. and LIN, N. (2010). The Bayesian elastic net. *Bayesian analysis* **5**(1) 151–170. <https://doi.org/10.1214/10-BA506>. MR2596439
- [42] LIANG, F., PAULO, R., MOLINA, G., CLYDE, M. A. and BERGER, J. O. (2008). Mixtures of g-priors for Bayesian variable selection. *Journal of the American Statistical Association* **103**(481) 410–423. <https://doi.org/10.1198/016214507000001337>. MR2420243
- [43] MALLICK, H., CHATTERJEE, S., CHOWDHURY, S., CHATTERJEE, S., RAHNAVARD, A. and HICKS, S. C. (2022). Differential expression of single-cell RNA-seq data using Tweedie models. *Statistics in medicine* **41**(18) 3492–3510. <https://doi.org/10.1002/sim.9430>. MR4453460
- [44] MARTINO, S., AKERKAR, R. and RUE, H. (2011). Approximate Bayesian inference for survival models. *Scandinavian Journal of Statistics* **38**(3) 514–528. <https://doi.org/10.1111/j.1467-9469>

- 2010.00715.x. MR2833844
- [45] MATÉRN, B. (2013) *Spatial variation* **36**. Springer Science & Business Media. <https://doi.org/10.1007/978-1-4615-7892-5>. MR0867886
- [46] MITCHELL, T. J. and BEAUCHAMP, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American statistical association* **83**(404) 1023–1032. MR0997578
- [47] MOHAMMED, S., BHARATH, K., KURTEK, S., RAO, A. and BALADANDAYUTHAPANI, V. (2021). RADIOHEAD: Radiogenomic analysis incorporating tumor heterogeneity in imaging through densities. *The Annals of Applied Statistics* **15**(4) 1808–1830. <https://doi.org/10.1214/21-aos1458>. MR4355077
- [48] MORRIS, J. S., BROWN, P. J., HERRICK, R. C., BAGGERLY, K. A. and COOMBES, K. R. (2008). Bayesian analysis of mass spectrometry proteomic data using wavelet-based functional mixed models. *Biometrics* **64**(2) 479–489. <https://doi.org/10.1111/j.1541-0420.2007.00895.x>. MR2432418
- [49] NELDER, J. A. and PREGIBON, D. (1987). An extended quasi-likelihood function. *Biometrika* **74**(2) 221–232. <https://doi.org/10.1093/biomet/74.2.221>. MR0903123
- [50] PARK, T. and CASELLA, G. (2008). The Bayesian lasso. *Journal of the American Statistical Association* **103**(482) 681–686. <https://doi.org/10.1198/016214508000000337>. MR2524001
- [51] RAFTERY, A. E., MADIGAN, D. and HOETING, J. A. (1997). Bayesian model averaging for linear regression models. *Journal of the American Statistical Association* **92**(437) 179–191. <https://doi.org/10.2307/2291462>. MR1436107
- [52] ROBERTS, G. O. and STRAMER, O. (2002). Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and computing in applied probability* **4**(4) 337–357. <https://doi.org/10.1023/A:1023562417138>. MR2002247
- [53] SHONO, H. (2008). Application of the Tweedie distribution to zero-catch data in CPUE analysis. *Fisheries Research* **93**(1-2) 154–162.
- [54] SMYTH, G. K. (1989). Generalized linear models with varying dispersion. *Journal of the Royal Statistical Society: Series B (Methodological)* **51**(1) 47–60. MR0984992
- [55] SMYTH, G. K. and JØRGENSEN, B. (2002). Fitting Tweedie’s compound Poisson model to insurance claims data: dispersion modelling. *ASTIN Bulletin: The Journal of the IAA* **32**(1) 143–157. <https://doi.org/10.2143/AST.32.1.1020>. MR1930491
- [56] SMYTH, G. K. and VERBYLA, A. P. (1999). Adjusted likelihood methods for modelling dispersion in generalized linear models. *Environmetrics: The official journal of the International Environmetrics Society* **10**(6) 695–709.
- [57] SWALLOW, B., BUCKLAND, S. T., KING, R. and TOMS, M. P. (2016). Bayesian hierarchical modelling of continuous non-negative longitudinal data with a spike at zero: An application to a study of birds visiting gardens in winter. *Biometrical Journal* **58**(2) 357–371. <https://doi.org/10.1002/bimj.201400081>. MR3499119
- [58] TWEEDIE, M. C. et al. (1984). An index which distinguishes between some important exponential families. In *Statistics: Applications and new directions: Proc. Indian statistical institute golden Jubilee International conference* **579** 604. MR0786162
- [59] VERBYLA, A. P. (1993). Modelling variance heterogeneity: residual maximum likelihood and diagnostics. *Journal of the Royal Statistical Society: Series B (Methodological)* **55**(2) 493–508. MR1224412
- [60] WILLIAMS, C. K. and RASMUSSEN, C. E. (2006) *Gaussian processes for machine learning* **2**. MIT press Cambridge, MA. MR2514435
- [61] WOLPERT, R. L. and ICKSTADT, K. (1998). Poisson/gamma random field models for spatial statistics. *Biometrika* **85**(2) 251–267. <https://doi.org/10.1093/biomet/85.2.251>. MR1649114
- [62] YANG, Y., QIAN, W. and ZOU, H. (2018). Insurance premium prediction via gradient tree-boosted Tweedie compound Poisson models. *Journal of Business & Economic Statistics* **36**(3) 456–470. <https://doi.org/10.1080/07350015.2016.1200981>. MR3828973
- [63] YE, T., LACHOS, V. H., WANG, X. and DEY, D. K. (2021). Comparisons of zero-augmented continuous regression models from a Bayesian perspective. *Statistics in Medicine* **40**(5) 1073–1100. <https://doi.org/10.1002/sim.8795>. MR4384363
- [64] ZEGER, S. L. and KARIM, M. R. (1991). Generalized linear models with random effects; a Gibbs sampling approach. *Journal of the American statistical association* **86**(413) 79–86. MR1137101
- [65] ZHANG, H. (2002). On estimation and prediction for spatial generalized linear mixed models. *Biometrics* **58**(1) 129–136. <https://doi.org/10.1111/j.0006-341X.2002.00129.x>. MR1891051
- [66] ZHANG, Y. (2013). Likelihood-based and Bayesian methods for Tweedie compound Poisson linear mixed models. *Statistics and Computing* **23**(6) 743–757. <https://doi.org/10.1007/s11222-012-9343-7>. MR3247830
- [67] ZHOU, H. and HANSON, T. (2015). Bayesian spatial survival models. *Nonparametric Bayesian Inference in Biostatistics* 215–246. MR3411022

Aritra Halder. Department of Biostatistics, Dornsife School of Public Health, Drexel University, Philadelphia, PA, USA. E-mail address: aritra.halder@drexel.edu

Shariq Mohammed. Department of Biostatistics, Boston University School of Public Health; Rafik B. Hariri Institute for Computing and Computational Science & Engineering, Boston University, Boston, MA, USA. E-mail address: shariqm@bu.edu

Dipak K. Dey. Department of Statistics, University of Connecticut, Storrs, CT USA. E-mail address: dipak.dey@uconn.edu