

# BAYESIAN INFERENCE FOR SPATIAL-TEMPORAL NON-GAUSSIAN DATA USING PREDICTIVE STACKING

SOUMYAKANTI PAN<sup>\*</sup>, LU ZHANG<sup>†</sup>, JONATHAN R. BRADLEY<sup>‡</sup>, SUDIPTO BANERJEE<sup>\*</sup>

*<sup>\*</sup>Department of Biostatistics, University of California Los Angeles*

*<sup>†</sup>Division of Biostatistics, Department of Population and Public Health Sciences,  
University of Southern California*

*<sup>‡</sup>Department of Statistics, Florida State University*

**ABSTRACT.** Analysing non-Gaussian spatial-temporal data typically requires introducing spatial dependence in generalised linear models through the link function of an exponential family distribution. However, unlike in Gaussian likelihoods, inference is considerably encumbered by the inability to analytically integrate out the random effects and reduce the dimension of the parameter space. Iterative estimation algorithms struggle to converge due to the presence of weakly identified parameters. We devise an approach that obviates these issues by exploiting generalised conjugate multivariate distribution theory for exponential families, which enables exact sampling from analytically available posterior distributions conditional upon some fixed process parameters. More specifically, we expand upon the Diaconis-Ylvisaker family of conjugate priors to achieve analytically tractable posterior inference for spatially-temporally varying regression models conditional on some kernel parameters. Subsequently, we assimilate inference from these individual posterior distributions over a range of values of these parameters using Bayesian predictive stacking. We evaluate inferential performance on simulated data, compare with fully Bayesian inference using Markov chain Monte Carlo and apply our proposed method to analyse spatially-temporally referenced avian count data from the North American Breeding Bird Survey database.

**Keywords.** Bayesian inference; Generalised linear models; Predictive stacking; Spatial-temporal models.

## 1. INTRODUCTION

Statistical modelling and analysis for spatially oriented non-Gaussian outcomes play a crucial role in diverse scientific applications (see, for example, [De Oliveira et al., 1997](#); [Diggle et al., 1998](#); [Heagerty and Lele, 1998](#); [De Oliveira, 2000](#); [Zhang et al., 2022](#); [Saha et al., 2022](#)). Diverse scientific investigations involve spatially, and possibly temporally, oriented outcomes that are not Gaussian, or even continuous. For example, climate scientists record daily or monthly binary variables at spatial locations indicating whether, or not, rainfall was measurable; ecologists and forest scientists analyse temporal evolution of species counts at locations; and economists study the spatial distribution of the number of insurance claims at different locations over time. While spatial-temporal

---

*Date:* March 4, 2025.

count data are sometimes aggregated into rates or counts over larger regions, statisticians are increasingly encountering count data referenced by space-time coordinates.

For our purposes, we consider a spatial-temporal process as an uncountable set of random variables, say  $\{z(\ell) : \ell \in \mathcal{D}\}$ , over a domain of interest  $\mathcal{D}$ , which is endowed with a probability law specifying the joint distribution for any finite sample from that set. In spatial-temporal settings  $\mathcal{D} = \mathcal{S} \times \mathcal{T}$ , where  $\mathcal{S} \subset \mathbb{R}^d$  is the spatial region,  $\mathcal{T} \subset [0, \infty)$  is the time domain and  $\ell = (s, t)$  is a space-time coordinate with spatial location  $s \in \mathcal{S}$  and time point  $t \in \mathcal{T}$  (see, e.g., [Gneiting and Guttorp, 2010](#), for a review). Following [Diggle et al. \(1998\)](#), we introduce spatial-temporal stochastic processes for non-Gaussian data. Let  $y(\ell)$  be the outcome at  $\ell$  endowed with a probability law from the natural exponential family, which we denote by

$$y(\ell) \sim \text{EF}(x(\ell)^\top \beta + z(\ell); b, \psi_y) \quad (1)$$

for some positive parameter  $b > 0$  and unit log partition function  $\psi_y$  (Section 2.1). Fixed effects regression and spatial dependence, e.g.,  $x(\ell)^\top \beta + z(\ell)$ , is introduced in the natural parameter, where  $x(\ell)$  is a  $p \times 1$  vector of predictors referenced with respect to  $\ell$ ,  $\beta$  is a  $p \times 1$  vector of slopes measuring the trend,  $z(\ell)$  is a zero-centred spatial-temporal process on  $\mathbb{R}^d$  specified by a scale parameter  $\sigma_z$  and a spatial correlation function  $R(\cdot, \cdot; \theta_{\text{sp}})$  with  $\theta_{\text{sp}}$  consisting of spatial-temporal decay and smoothness parameters. This structure is embodied by spatial generalised mixed-effect models ([McCulloch and Searle, 2001](#); [Hughes and Haran, 2013](#)). Bayesian inference for (1) is appealing as it offers fully probabilistic inference for the latent process. However, the presence of  $z(\ell)$  presents challenges for non-Gaussian families as we cannot integrate out  $z$ , which begets a high-dimensional parameter space. Iterative algorithms such as Markov Chain Monte Carlo (MCMC) attempt to sample from the posterior distribution, but convergence is often hampered by high auto-correlations and weakly identified parameters  $\theta_{\text{sp}}$ . Section 1.3 in [Haran \(2011\)](#) offers a detailed overview of MCMC algorithms for estimating such models.

Rather than devising computational algorithms to achieve improved convergence, we develop a framework to conduct Bayesian inference for fixed effects and the latent spatial-temporal process by entirely avoiding convergence issues for MCMC or other iterative algorithms for spatial-temporal non-Gaussian data. Instead, we build a hierarchical model that yields analytically accessible posterior distributions for  $\beta$  and  $z$  subject to fixing weakly identified parameters  $\theta_{\text{sp}}$  and some hyperparameters so that we can draw exact posterior samples for any fixed values of these parameters. This is achieved by availing of recent results developed in [Bradley and Clinch \(2024\)](#) on a new class of analytically accessible conjugate multivariate (CM) and generalised conjugate multivariate (GCM) distributions for spatial models by extending the Diaconis-Ylvisaker family of conjugate priors for exponential families ([Diaconis and Ylvisaker, 1979](#)). Our current manuscript expands significantly over this framework. Not only do we move from spatial to spatial-temporal processes, we further enrich (1) by introducing spatially-temporally varying regression coefficients ([Gelfand et al., 2003](#)). Such models are attractive in the Bayesian paradigm as they capture the impact of explanatory variables or predictors over space and time and offer inference on processes that

are completely unobserved. However, these extensions are not straightforward because of the presence of explanatory variables in the fixed regression component and also in a spatially-temporally varying component.

The literature on spatially-temporally varying coefficient models is rather sparse and has remained purely notional for exponential families, perhaps because they require extensive tuning of MCMC or other iterative algorithms that struggle with weakly identified parameters and high-dimensional random effects. Hence, we devise Bayesian predictive stacking for estimating such models. Stacking (Wolpert, 1992; Breiman, 1996; Clyde and Iversen, 2013) is a model averaging procedure that is widely used in machine learning and has been shown (see, e.g., Le and Clarke, 2017; Yao et al., 2018, 2020, 2021) to be an effective alternative to traditional Bayesian model averaging (Madigan et al., 1996; Hoeting et al., 1999). Let  $\mathcal{M} = \{M_1, \dots, M_L\}$  denote a collection of  $L$  candidate models, where each  $M_l$  corresponds to a fixed set of values of certain parameters ( $\theta_{\text{sp}}$  and, for our subsequent development, a boundary adjustment parameter) that allow us to exactly sample from the posterior distributions  $p_l(\beta, z \mid y)$  for  $l = 1, \dots, L$ . Rather than pursuing full inference using iterative sampling algorithms, we offer inference by “stacking” these  $L$  posterior distributions by solving for the optimal set of weights needed to calculate the weighted average of the  $L$  analytically available posterior distributions. Bayesian model averaging asymptotically chooses a single model in  $\mathcal{M}$  that is closest to an oracle  $M_0$  in terms of Kullback-Leibler divergence and, hence, is appropriate if  $\mathcal{M}$  contains the true data generating model  $M_0$ , but is flawed otherwise. Yao et al. (2018) adapts stacking to assimilate different Bayesian predictive distributions and appears more tenable in practical geostatistics (Zhang et al., 2023).

After briefly discussing CM and GCM distributions with multivariate extensions in Section 2, Section 3 introduces our hierarchical spatially-temporally varying regression model and derives accessible posterior distributions. Section 4 develops predictive stacking with novel results on posterior sampling and predictive inference. Sections 5 and 6 present simulation experiments demonstrating the effectiveness of our method and analyses a spatial-temporal dataset on avian counts. We conclude with some discussion and directions for future work in Section 7.

## 2. CONJUGATE PRIORS FOR EXPONENTIAL FAMILY

**2.1. The Diaconis-Ylvisaker distribution.** Let  $Y$  be distributed from the natural exponential family,  $\text{EF}(\eta; b, \psi)$ , with density

$$p(Y \mid \eta) = \exp\{\eta Y - b\psi(\eta) + c(Y)\}, \quad Y \in \mathcal{Y}, \eta \in \mathcal{H}, \quad (2)$$

where  $\mathcal{Y}$  denotes the support of  $Y$  and  $\mathcal{H} = \{\eta : \psi(\eta) < \infty\}$  denotes the natural parameter space and, therefore, the support of  $\eta$ . The scalar  $b$  may be unknown, while  $\psi(\cdot)$  and  $c(\cdot)$  are known functions. We will be discussing various forms of the unit log partition function  $\psi(\cdot)$  and the scalar  $b$  to denote different distributions from the family. Diaconis and Ylvisaker (1979) provides a proper conjugate prior for  $\eta$  in (2) as  $\text{DY}(\alpha, \kappa; \psi)$ , which has the density

$$p(\eta \mid \alpha, \kappa) \propto \exp\{\alpha\eta - \kappa\psi(\eta)\}, \quad \eta \in \mathcal{H}, \frac{\alpha}{\kappa} \in \mathcal{Y}, \kappa > 0. \quad (3)$$

It is easily seen that the posterior distribution  $\eta \mid Y, \alpha, \kappa \sim \text{DY}(\alpha + Y, \kappa + b; \psi)$ . There are several special cases of the DY distribution other than the Gaussian ( $\psi = \psi_1$ ), log-gamma ( $\psi = \psi_2$ ), and logit-beta ( $\psi = \psi_3$ ) distributions, several of which do not correspond to a member of the exponential family. For example,  $\alpha = 0$ ,  $\psi(t) = \psi_4(t) = \log(1 + t^2/\nu)$ , and  $\kappa = (\nu + 1)/2$  with  $\nu > 0$  results in a Student's  $t$ -distribution with  $\nu$  degrees of freedom.

**2.2. The (Generalised) Conjugate Multivariate distribution.** Bradley et al. (2020) introduced a multivariate version of (3) using linear combinations of mutually independent DY random variables. Let  $\zeta$  be the  $n \times 1$  random vector

$$\zeta = \mu + L\eta, \quad (4)$$

where  $\zeta \in \mathcal{M}$ ,  $\mathcal{M} = \{\zeta : \zeta = \mu + L\eta, \eta \in \mathcal{H}^n\}$ ,  $\mu \in \mathbb{R}^n$  denotes a location vector,  $L$  is an  $n \times n$  lower-triangular matrix with positive diagonal elements and the  $n \times 1$  random vector  $\eta = (\eta_1, \dots, \eta_n)^\top$  consists of  $n$  mutually independent DY random variables,  $\eta_i \sim \text{DY}(\alpha_i, \kappa_i; \psi)$  with  $\kappa_i > 0$  for  $i = 1, \dots, n$ . Define  $\zeta \sim \text{CM}(\mu, L, \alpha, \kappa; \psi)$  with unnormalised density

$$p(\zeta \mid \mu, L, \alpha, \kappa) \propto \exp \left\{ \alpha^\top L^{-1}(\zeta - \mu) - \kappa^\top \psi(L^{-1}(\zeta - \mu)) \right\} \det(L^{-1}) \quad (5)$$

for all  $\zeta \in \mathcal{M}$ , where  $\psi$  operates element-wise on  $L^{-1}(\zeta - \mu)$ ,  $\alpha = (\alpha_1, \dots, \alpha_n)^\top$  and  $\kappa = (\kappa_1, \dots, \kappa_n)^\top$ . If  $\zeta = (\zeta_1^\top, \zeta_2^\top)^\top$  is distributed as  $\text{CM}(\mu, L, \alpha, \kappa; \psi)$ , where  $\zeta_1$  is  $r \times 1$  and  $\zeta_2$  is  $(n - r) \times 1$ , then the conditional distribution of  $\zeta_1$  given  $\zeta_2$  is  $\text{CM}_c(\mu^*, A_1, \alpha, \kappa; \psi)$  with

$$p(\zeta_1 \mid \zeta_2 = c_2, \mu^*, A_1, \alpha, \kappa) \propto \exp \{ \alpha^\top (A_1 \zeta_1 - \mu^*) - \kappa^\top \psi(A_1 \zeta_1 - \mu^*) \}, \quad (6)$$

as the unnormalised density for all  $(\zeta_1^\top, c_2^\top)^\top \in \mathcal{M}$ . Here,  $A_1$  is defined as the  $n \times r$  submatrix of  $L^{-1} = [A_1 : A_2]$ , and  $\mu^* = L^{-1}\mu - A_2 c_2$  for some  $c_2 \in \mathbb{R}^{n-r}$ . The proportionality constant in (6) is strictly positive and finite ensuring that (6) is proper (Bradley et al., 2020).

Bradley and Clinch (2024) generalise the CM distribution by relating  $\zeta$ ,  $\mu$ ,  $L$  and  $\eta$  as in (4), where  $\zeta = (\zeta_1^\top, \dots, \zeta_K^\top)^\top$  and  $\eta = (\eta_1^\top, \dots, \eta_K^\top)^\top$  are  $n \times 1$  with  $n = \sum_{k=1}^K n_k$ , and each element of  $\eta_k$  is independently distributed as  $\eta_{k,i} \sim \text{DY}(\alpha_{k,i}, \kappa_{k,i}; \psi_k)$ ,  $L$  is an  $n \times n$  lower-triangular matrix with positive diagonal elements and  $\mu$  is an  $n \times 1$  location parameter. The density is written analogously to (5) with  $\zeta \in \mathcal{N}$ ,  $\mathcal{N} = \{\zeta : \zeta = \mu + L\eta, \eta_{k,i} \in \mathcal{H}_k, i = 1, \dots, n_k, k = 1, \dots, K\}$  with  $\mathcal{H}_k$  being the parameter space corresponding to the log partition function  $\psi_k$ ,  $\alpha_{k,i}/\kappa_{k,i} \in \mathcal{Y}_k$ ,  $\kappa_{k,i} > 0$  and  $\psi(L^{-1}(\zeta - \mu)) = (\psi_1(J_1 L^{-1}(\zeta - \mu))^\top, \dots, \psi_K(J_K L^{-1}(\zeta - \mu))^\top)^\top$ , where  $J_k = [0 : I_{n_k} : 0]$  is  $n_k \times n$  and each  $\psi_k(\cdot)$  operates element-wise on the vector of arguments,  $\alpha = (\alpha_1^\top, \dots, \alpha_K^\top)^\top$  and  $\kappa = (\kappa_1^\top, \dots, \kappa_K^\top)^\top$  are  $n \times 1$  parameter vectors.

We say  $\zeta$  is distributed as  $\text{GCM}(\mu, L, \alpha, \kappa; \psi)$ . A conditional GCM density up to a normalising constant is obtained analogous to (6) and denoted as  $\text{GCM}_c(\mu^*, A_1, \alpha, \kappa; \psi)$ . We use these distributions for building the hierarchical models in the following sections. In general, we may be unable to sample directly from either the conditional CM or the conditional GCM distributions except for some familiar exceptions (e.g., the conditional distribution of Gaussian is indeed Gaussian). However, it is possible to consider an augmented model with a particular structure that yields a posterior distribution in the GCM family that is easy to sample from (Section 3).

### 3. BAYESIAN HIERARCHICAL MODEL

**3.1. Conjugate spatial-temporal process.** The prevalent specification for  $z(\ell)$  in (1) is a Gaussian process. This is a special case of a “CM” process (CMP) on  $\mathcal{D}$ . We write  $z(\ell) \sim \text{CMP}(\mu_z(\ell), V_z(\cdot, \cdot), \alpha_z, \kappa_z; \psi_z)$  if, for any finite collection  $\mathcal{L} = \{\ell_1, \dots, \ell_n\} \subset \mathcal{D}$ , the  $n \times 1$  vector  $z(\mathcal{L}) = (z(\ell_1), \dots, z(\ell_n))^\top$  is distributed as  $\text{CM}(\mu_z(\mathcal{L}), L_z, \alpha_z, \kappa_z; \psi_z)$ ,  $\mu_z(\mathcal{L}) = (\mu_z(\ell_1), \dots, \mu_z(\ell_n))^\top$  and  $L_z$  is the lower-triangular Cholesky factor of the  $n \times n$  matrix  $V_z(\mathcal{L}; \theta) = (V_z(\ell_i, \ell_j; \theta))$  for  $i, j = 1, \dots, n$  and  $V_z(\ell, \ell'; \theta)$  is a spatial covariance function indexed by parameters  $\theta$ . The CM process endows a CM distribution on  $\{z(\ell) : \ell \in \mathcal{L}\}$ , thereby enabling predictive inference over the entire domain  $\mathcal{D}$  based on partial realisations of the process. The Gaussian process is a CM process with  $\psi_z = \psi_1$ . A crucial extension here is to move from the real-valued process in (1) to a vector-valued process  $z(\ell)$  with each element following a CM process. This vector-valued process is composed of univariate spatial-temporal process that model regression coefficients for certain predictors whose impact is posited to vary over space and time. We elaborate below.

**3.2. Conjugate Bayesian spatially-temporally varying coefficients model.** Let  $\mathcal{L} = \{\ell_1, \dots, \ell_n\}$  be a fixed set of  $n$  distinct space-time coordinates in  $\mathcal{D}$ , where  $y(\mathcal{L}) = (y(\ell_1), \dots, y(\ell_n))^\top \in \mathcal{Y}^n$ , which we simply denote by  $y$ , is the vector of observed outcomes. We introduce spatially-temporally varying coefficients as

$$\begin{aligned} y(\ell_i) \mid \beta, z(\ell_i), \xi_i, \mu_i &\stackrel{\text{ind.}}{\sim} \text{EF} \left( x(\ell_i)^\top \beta + \tilde{x}(\ell_i)^\top z(\ell_i) + \xi_i - \mu_i; b, \psi_y \right), \quad (i = 1, \dots, n) \\ z_j \mid \sigma_{z_j}, L_{z_j}, \alpha_z, \kappa_z &\stackrel{\text{ind.}}{\sim} \text{CM}(\mu_{z_j}, \sigma_{z_j} L_{z_j}(\theta_{\text{sp}j}), \alpha_{z_j}, \kappa_{z_j}; \psi_{z_j}), \quad (j = 1, \dots, r) \\ \beta \mid \sigma_\beta, L_\beta, \alpha_\beta, \kappa_\beta &\sim \text{CM}(\mu_\beta, \sigma_\beta L_\beta, \alpha_\beta, \kappa_\beta; \psi_\beta) \\ \xi \mid \beta, z, \mu, \sigma_\xi, \alpha_\xi, \kappa_\xi &\sim \text{GCM}_c(\tilde{\mu}_\xi, H_\xi, \alpha_\xi, \kappa_\xi; \psi_\xi), \quad \theta \sim \pi(\theta), \end{aligned} \quad (7)$$

where  $x(\ell_i)$  is a  $p \times 1$  vector of predictors,  $\beta$  is the corresponding  $p \times 1$  vector of slopes (fixed effects),  $\tilde{x}(\ell_i)$  is the  $r \times 1$  vector ( $r \leq p$ ) consisting of those predictors in  $x(\ell_i)$  that are posited to have spatially varying regression coefficients, and  $z_j = (z_j(\ell_1), \dots, z_j(\ell_n))^\top$  is  $n \times 1$  with each  $z_j(\ell_i)$  being a spatially-temporally varying coefficient for the predictor  $\tilde{x}_j(\ell_i)$ .

For subsequent developments, we assume that the probability law of the  $nr \times 1$  vector  $z = (z_1^\top, \dots, z_r^\top)^\top$  is derived by assuming that  $z_j(\ell) \stackrel{\text{ind.}}{\sim} \text{CMP}(\mu_{z_j}(\ell), \sigma_{z_j}^2 R_j(\cdot, \cdot; \theta_{\text{sp}j}))$  are a collection of  $r$  independent CM processes with their respective parameters and log partition functions. The spatial covariance matrix for  $z_j$  is given by  $\sigma_{z_j}^2 L_z(\theta_{\text{sp}j}) L_z(\theta_{\text{sp}j})^\top$ , where  $\sigma_{z_j}^2$  is the spatial variance corresponding to process  $z_j(\ell)$  and  $L_z(\theta_{\text{sp}j})$  is the lower-triangular Cholesky factor of an  $n \times n$  spatial correlation matrix  $R_j(\theta_{\text{sp}j})$ . We model the elements of  $R_j(\theta_{\text{sp}j})$  using the spatial-temporal correlation function (see, e.g., [Gneiting and Guttorp, 2010](#), for other examples)

$$R_j(\ell, \ell'; \theta_{\text{sp}j}) = \frac{1}{\phi_{1j}|t - t'|^2 + 1} \exp \left( -\frac{\phi_{2j}\|s - s'\|}{\sqrt{1 + \phi_{1j}|t - t'|^2}} \right), \quad \phi_{1j}, \phi_{2j} > 0, \quad (8)$$

where  $\ell = (s, t)$  and  $\ell' = (s', t')$  are any two distinct space-time coordinates in  $\mathcal{D}$ ,  $\|\cdot\|$  is the Euclidean distance over  $\mathcal{S}$ ,  $\theta_{\text{sp}j} = (\phi_{1j}, \phi_{2j})$ , and  $\phi_{1j}$  and  $\phi_{2j}$  are positive spatial and temporal decay parameters, respectively. For ease of notation, we drop  $\theta_{\text{sp}j}$  in (8) and simply write  $R_j(\cdot, \cdot)$

for the correlation function. The specification for each  $z_j(\ell)$  is completed with CM distribution parameters  $\alpha_{z_j}$ ,  $\kappa_{z_j}$  and  $\psi_{z_j}$ , all specific to the  $j$ th process.

The prior for  $\beta$  is also a CM distribution with parameters defined analogously and the variance-covariance parameter  $\sigma_\beta^2 L_\beta L_\beta^\top$ , where  $L_\beta$  is the Cholesky factor of a fixed correlation matrix  $V_\beta$ . The conditional prior for  $\xi$  is a GCM<sub>c</sub> distribution with the  $2n \times n$  matrix  $H_\xi = [I_n : \sigma_\xi^{-1} I_n]^\top$ , location parameter  $\tilde{\mu}_\xi = ((\mu - X\beta - \tilde{X}z)^\top, \sigma_\xi^{-1} \mu_\xi^\top)^\top$ , where  $\mu = (\mu_1, \dots, \mu_n)^\top$  is  $n \times 1$ , shape and scale parameters  $\alpha_\xi = (\alpha_\epsilon 1_n^\top, 0_n^\top)^\top$  and  $\kappa_\xi = (\kappa_\epsilon 1_n^\top, (1/2)1_n^\top)^\top$ , respectively, for some fixed positive reals  $\alpha_\epsilon$  and  $\kappa_\epsilon$ . The unit log partition function  $\psi_\xi$  is defined as  $\psi_\xi(h) = (\psi_y(h_1)^\top, \psi_1(h_2)^\top)^\top$  for any  $h_1, h_2 \in \mathbb{R}^n$ , where  $\psi_y(\cdot)$  and  $\psi_1(\cdot)$  operate element-wise on their respective arguments. We collect the variance components into a  $(3r + 2)$ -dimensional parameter  $\theta = (\sigma_\beta, \sigma_z, \sigma_\xi, \theta_{\text{sp}})$  where  $\sigma_z = (\sigma_{z_1}, \dots, \sigma_{z_r})$  is  $r$ -dimensional and  $\theta_{\text{sp}} = (\theta_{\text{sp}1}, \dots, \theta_{\text{sp}r})$  is  $2r$ -dimensional and subsequently assign a hyperprior  $\pi$  on  $\theta$ . Gather the fixed and random effects into the  $(n + p + nr) \times 1$  parameter  $\gamma = (\xi^\top, \beta^\top, z^\top)^\top$ . Let  $\tilde{\mu} = (\mu^\top, \mu_\gamma^\top)^\top$  be the  $(2n + p + nr) \times 1$  vector obtained by combining  $\mu$  with the location parameters  $\mu_\gamma = (\mu_\xi^\top, \mu_\beta^\top, \mu_z^\top)^\top$ , where  $\mu_z = (\mu_{z_1}^\top, \dots, \mu_{z_r}^\top)^\top$  is  $nr \times 1$ . Let  $q = -Q^\top \tilde{\mu}$ , where  $Q$  is  $(2n + p + nr) \times n$  and obtained from the decomposition  $QQ^\top = I_{2n+p+nr} - H(H^\top H)^{-1}H^\top$  with the columns of  $Q$  being the  $n$  unit norm orthogonal eigenvectors of  $I_{2n+p+nr} - H(H^\top H)^{-1}H^\top$  corresponding to eigenvalue 1. Hence,  $Q^\top Q = I_n$  and  $H^\top Q = 0$ . Furthermore,  $H = [(I_n : X : \tilde{X})^\top : I_{n+p+nr}]^\top$  is  $(2n + p + nr) \times (n + p + nr)$ ,  $X$  is  $n \times p$  with  $x(\ell_i)^\top$  as its  $i$ th row, and  $\tilde{X} = [\text{diag}(\tilde{x}_1) : \dots : \text{diag}(\tilde{x}_r)]$  is  $n \times nr$  with  $\text{diag}(\tilde{x}_j)$  being an  $n \times n$  diagonal matrix whose  $i$ th diagonal element is  $\tilde{x}_j(\ell_i)$ .

We work with an improper prior on  $q$  given by  $p(q) \propto 1$  which follows from assuming a vague prior on  $\mu$  (see, Proposition A1 and Lemma A1 in Appendix A for technical details on the construction of the improper prior). These specifications yield the posterior distribution

$$p\left((\gamma^\top, q^\top)^\top \mid y, \theta\right) \propto \text{GCM}\left((\gamma^\top, q^\top)^\top \mid 0_{2n+p+nr}, V^*(\theta), \alpha^*, \kappa^*; \psi^*\right), \quad (9)$$

where  $V^*(\theta)^{-1} = L(\theta)^{-1}[H : Q]$ ,  $L(\theta) = \text{blkdiag}(I_n, \sigma_\xi I_n, \sigma_\beta L_\beta, L_z(\theta))$  is  $(3n + p) \times (3n + p)$  block-diagonal with  $\sigma_{z_j} L_{z_j}(\theta_{\text{sp}})$  as the  $j$ th block of  $nr \times nr$  block-diagonal matrix  $L_z(\theta)$ . The shape and scale parameters are  $\alpha^* = ((y + \alpha_\epsilon 1_n)^\top, 0_n^\top, \alpha_\beta^\top, \alpha_z^\top)^\top$  and  $\kappa^* = ((b + \kappa_\epsilon)1_n^\top, (1/2)1_n^\top, \kappa_\beta^\top, \kappa_z^\top)^\top$ . The unit log partition function is  $\psi^*(h) = (\psi_y(h_1)^\top, \psi_1(h_2)^\top, \psi_\beta(h_3)^\top, \psi_z(h_4)^\top)^\top$  for some  $h = (h_1^\top, h_2^\top, h_3^\top, h_4^\top)^\top$  with  $h_1, h_2 \in \mathbb{R}^n$ ,  $h_3 \in \mathbb{R}^p$ ,  $h_4 \in \mathbb{R}^{nr}$ , and,  $\psi_z(h_4) = (\psi_{z_1}(h_{41}), \dots, \psi_{z_r}(h_{4r}))$  for  $h_{4j} \in \mathbb{R}^n$ , where the log partition functions operate element-wise on the arguments (see Theorem A1 in Appendix A). Generalised linear models typically assume that  $q$  is zero, which yields the posterior distribution  $p(\gamma \mid y, \theta) \propto \text{GCM}_c(\gamma \mid \tilde{\mu}, H, \alpha^*, \kappa^*; \psi^*)$ . Recall that we cannot sample directly from a conditional GCM distribution except for some special cases of  $\psi_y$  (e.g., Gaussian). Finally, we remark that  $\mu$  is crucial in producing the posterior distribution within the GCM family and, hence, unlike in traditional generalised linear models, cannot be excluded from (7).

#### 4. PREDICTIVE STACKING

**4.1. Choice of candidate models.** We specifically consider modelling count data typically originating from Poisson ( $\psi_y = \psi_2$ ) or binomial ( $\psi_y = \psi_3$ ) distributions. The choice of the log partition functions  $\psi_\beta(\cdot)$  and  $\psi_{z_j}(\cdot)$  in the prior distributions of  $\beta$  and the  $j$ th process  $z_j$  uniquely determine



its associated DY parameters  $\alpha_\beta$ ,  $\kappa_\beta$ ,  $\alpha_{z_j}$  and  $\kappa_{z_j}$ . For example, in practice, we consider a multivariate Gaussian prior  $\beta \mid \sigma_\beta^2 \sim \mathcal{N}(0, \sigma_\beta^2 I_p)$  which is a special case of (7) with  $\mu_\beta = 0$ ,  $L_\beta = I_p$ ,  $\alpha_\beta = 0_p$ ,  $\kappa_\beta = (1/2)1_p$  and  $\psi_\beta = \psi_1$ . Similarly, for each  $j$ , considering  $\mu_{z_j}(\cdot)$  as the zero function,  $\alpha_{z_j} = 0_n$ ,  $\kappa_{z_j} = (1/2)1_n$  and  $\psi_{z_j} = \psi_1$  implies that  $z_j(\ell) \mid \sigma_{z_j}^2 \sim \text{GP}(0, \sigma_{z_j}^2 R_j(\cdot, \cdot))$  is a zero-centred Gaussian process with the spatial-temporal covariogram (8). For subsequent developments, we proceed with zero-centred multivariate Gaussian priors on  $\beta$  and  $z_j$  for each  $j$ , setting the location parameters  $\mu_\beta$  and  $\mu_{z_j}$  for each  $j$  to be zero. We choose  $\alpha_\epsilon > 0$  (called the boundary adjustment parameter) in the prior of  $\xi$  to ensure that  $\alpha^*$  in (9) does not lie on the boundary of its parameter space. For Poisson data,  $\kappa_\epsilon = 0$  and for binomial data,  $\kappa_\epsilon = 2\alpha_\epsilon$ .

The prior  $\pi(\theta)$  assumes independence for each element of  $\theta$ . While fixing  $\sigma_\xi$  obtained by cross-validation procedures is possible, we place inverse-gamma priors  $\text{IG}(\nu_\beta/2, \nu_\beta/2)$  and  $\text{IG}(\nu_{z_j}/2, \nu_{z_j}/2)$  on  $\sigma_\beta^2$  and  $\sigma_{z_j}^2$  for each  $j$  respectively. We stack over candidate models specified by the different fixed values of  $\Delta = (\alpha_\epsilon, \sigma_\xi, \theta_{\text{sp}})$  while  $\pi(\theta \setminus \Delta)$  remains unchanged. Hence, a hierarchical model (7) is uniquely characterised by the fixed value of  $\Delta$ , denoted by  $M_\Delta$ . We build the collection of candidate models as  $\mathcal{M} = \{M_\Delta : \Delta \in \otimes_{\delta \in \Delta} G_\delta\}$  where  $G_\delta$  is a grid of candidate values of  $\delta \in \Delta$  and  $\otimes$  denotes the Cartesian product.

**4.2. Sampling from posterior.** Sampling from the posterior distribution (9) for a model  $M_\Delta$  first requires sampling  $\beta$  and  $z_j$  for each  $j$  from their respective marginal priors. The Gaussian priors on  $\beta$  and  $z_j$  and inverse-gamma priors on their variances, as specified in Section 4.1, yields the marginal priors of  $\beta$  and  $z$  as central multivariate  $t$ -distributions with degrees of freedom  $\nu_\beta$  and  $\nu_{z_j}$ , respectively, for each  $j$ . This still remains in the DY family. Since  $[H : Q]^{-1} = [H(H^\top H)^{-1} : Q]^\top$ , we compute

$$\gamma^{(r)} = (H^\top H)^{-1} H^\top \tilde{v}, \quad q^{(r)} = Q^\top \tilde{v} \quad (10)$$

to obtain replicates  $(\gamma^{(r)\top}, q^{(r)\top})^\top$  of  $(\gamma^\top, q^\top)^\top$ , where the random vector  $\tilde{v} = L(\theta)v$  and  $v \sim \text{GCM}(0, I_{3n+p}, \alpha^*, \kappa^*; \psi_\gamma)$ . The random vector  $\tilde{v} = (\tilde{v}_\eta^\top, \tilde{v}_\xi^\top, \tilde{v}_\beta^\top, \tilde{v}_z^\top)^\top$  is made up of  $\tilde{v}_\eta$ , where  $\tilde{v}_{\eta,i} \sim \text{DY}(y_i + \alpha_\epsilon, b + \kappa_\epsilon; \psi_y)$  for  $i = 1, \dots, n$  and,  $\tilde{v}_\beta, \tilde{v}_z = (\tilde{v}_{\beta_1}^\top, \dots, \tilde{v}_{\beta_r}^\top)^\top$  are samples from their marginal priors with their scale parameters integrated out and their location parameters set to zero, and  $\tilde{v}_\xi$  is a sample from  $\mathcal{N}(0, \sigma_\xi^2 I_n)$ . The random vector  $\tilde{v}_\eta$  merits special attention as it is the posterior sample of the  $n \times 1$  natural parameter  $\eta$  with  $i$ th element  $\eta(\ell_i) = x(\ell_i)^\top \beta + \tilde{x}(\ell_i)^\top z(\ell_i) + \xi_i - \mu_i$  (Diaconis and Ylvisaker, 1979).

The projection in (10) maps the posterior samples of  $\eta$  to the posterior samples of  $\gamma$  accounting for the effect of its priors. It is instructive to rewrite (10) as

$$\gamma^{(r)} = \left( H_1^\top H_1 + I_{n+p+nr} \right)^{-1} \left( H_1^\top \tilde{v}_\eta + \tilde{v}_\gamma \right), \quad (11)$$

where  $H = [H_1^\top : I_{n+p+nr}]^\top$  with  $H_1 = [I_n : X : \tilde{X}]$  and  $\tilde{v}_\gamma = (\tilde{v}_\xi^\top, \tilde{v}_\beta^\top, \tilde{v}_z^\top)^\top$ . Section B of Appendix B shows that, given  $\tilde{v}$ , computing (11) requires  $O(nr^3 + p^3)$  floating point operations (flops) and storage of  $O(np + nr + nr^2 + p^2)$  (see, Theorem A2 and Algorithm S1 in Appendix B).

We note that CM customarily uses posterior samples of  $\beta$  and  $z$  to draw samples from the posterior predictive distribution  $p(x(\ell)^\top \beta + \tilde{x}(\ell)^\top z(\ell) + \xi - \mu \mid y)$  for any  $\ell \in \mathcal{L}$  by simply considering  $x(\ell)^\top \beta^{(r)} + \tilde{x}(\ell)^\top z^{(r)}(\ell)$ , where  $\{\beta^{(r)}, z^{(r)}(\ell)\}$  are replicates from (9) obtained using (10), which

implicitly estimates  $\mu$  and  $\xi$  to be zero after integrating them out of the posterior distribution (Bradley et al., 2020; Bradley and Clinch, 2024).

**4.3. Prediction.** Given data observed at  $\mathcal{L}$ , let  $\tilde{\mathcal{L}} = \{\tilde{\ell}_1, \dots, \tilde{\ell}_{\tilde{n}}\} \subset \mathcal{D} \setminus \mathcal{L}$  be a collection of  $\tilde{n}$  new space-time coordinates in  $\mathcal{D}$ , where we wish to predict the response and the latent spatial-temporal processes. Let  $\tilde{y}$  and  $\tilde{z}_j$  for each  $j$  be the  $\tilde{n} \times 1$  vectors with  $i$ th elements  $y(\tilde{\ell}_i)$  and  $z_j(\tilde{\ell}_i)$  respectively. For a given model,  $M_\Delta$ , which entails a fixed value of  $\theta_{\text{sp}}$  and some auxiliary model parameters, spatial-temporal predictive inference evaluates the posterior predictive distribution,

$$p(\tilde{y}, \tilde{z} \mid y, M_\Delta) = \int p(\tilde{y} \mid \beta, \tilde{z}) p(\tilde{z} \mid z, M_\Delta) p(\beta, z \mid y, M_\Delta) d\beta dz, \quad (12)$$

where  $\tilde{z} = (\tilde{z}_1^\top, \dots, \tilde{z}_r^\top)^\top$ . Sampling from (12) is facilitated by first drawing  $\{\beta^{(r)}, z^{(r)}\}$  from  $p(\beta, z \mid y, M_\Delta)$  as described in Section 4.2 and then, for each drawn value  $z^{(r)}$  of  $z$ , drawing  $\tilde{z}^{(r)}$  from  $p(\tilde{z} \mid z, M_\Delta)$ . Further, for each posterior sample  $\beta^{(r)}$  and  $\tilde{z}^{(r)}$ , we draw  $\tilde{y}^{(r)}$  from  $p(\tilde{y} \mid \beta, \tilde{z})$  with  $\mu$  in (7) set to 0. This yields samples  $\{\tilde{y}^{(r)}, \tilde{z}^{(r)}\}$  from (12). Under the prior specification of  $z_j$  and  $\sigma_{z_j}^2$  in Section 4.1, the marginal distribution of the  $(n + \tilde{n}) \times 1$  vector  $(z_j^\top, \tilde{z}_j^\top)^\top$  corresponding to the  $j$ th process is a multivariate  $t$ -distribution  $t_{n+\tilde{n}}(\nu_{z_j}, 0_n, \tilde{V}_{z_j})$  with  $\nu_{z_j}$  degrees of freedom, location parameter  $0_n$  and  $(n + \tilde{n}) \times (n + \tilde{n})$  scale matrix  $\tilde{V}_{z_j} = [(R_j : C_j); (C_j^\top : \tilde{R}_j)]$ , where  $\tilde{R}_j$  is the  $\tilde{n} \times \tilde{n}$  correlation matrix for  $\tilde{z}_j$ , and  $C_j = [R_j(\ell, \ell')]$  is  $n \times \tilde{n}$  with  $\ell \in \mathcal{L}, \ell' \in \tilde{\mathcal{L}}$ . This yields

$$\tilde{z}_j \mid z_j, M_\Delta \sim t_{\tilde{n}} \left( \nu_{z_j} + n, C_j^\top R_j^{-1} z_j, \frac{\nu_{z_j} + z_j^\top R_j^{-1} z_j}{\nu_{z_j} + n} (\tilde{R}_j - C_j^\top R_j^{-1} C_j) \right), \quad (j = 1, \dots, r). \quad (13)$$

It is worth noticing that, the conditional scale matrix contains the factor  $(\nu_{z_j} + z_j^\top R_j^{-1} z_j)/(\nu_{z_j} + n)$  which is directly related to the Mahalanobis distance of  $z_j$  implying that the dispersion is enlarged in presence of extreme values of  $z_j$ . The degrees of freedom also increases by a factor  $n$  which means that, the more data we have, the less heavy-tailed  $p(\tilde{z}_j \mid z_j, M_\Delta)$  becomes (Ding, 2016). The analytic tractability as described by (13) provides further motivation behind the choice  $\psi_z = \psi_1$  for practical purposes, since, an efficient way for evaluating the predictive density at an out-of-sample point is particularly crucial in order to effectuate our stacking algorithm.

**4.4. Stacking algorithm.** We collect samples from  $p(\beta, z \mid \chi, \theta_{\text{sp}}, \cdot)$ , where  $\chi = \{y, X, \tilde{X}, \mathcal{L}\}$  denotes the data and  $\cdot$  denotes any other parameters that are required to be conditioned for accessible posterior sampling, for an array of candidate values of these parameters. Following Yao et al. (2018), we devise a stacking algorithm for (7) based on predictive densities. Given a collection of candidate models  $\mathcal{M} = \{M_1, \dots, M_L\}$ , we find the probability distribution in  $\mathcal{C} = \left\{ \sum_{l=1}^L w_l p(\cdot \mid y, M_l) : \sum_{l=1}^L w_l = 1, w_l \geq 0 \right\}$  by solving for the optimal stacking weights  $w = (w_1, \dots, w_L)$  as the solution to the optimisation problem

$$\max_{w_1, \dots, w_L} \frac{1}{n} \sum_{i=1}^n \log \sum_{l=1}^L w_l p(y(\ell_i) \mid y_{-i}, M_l), \quad w_l \geq 0, \sum_{l=1}^L w_l = 1, \quad (14)$$

where  $y_{-i}$  denotes  $(y_{-i}, X_{-i}, \tilde{X}_{-i}, \mathcal{L}_{-i})$ , i.e., the data with the  $i$ th observation deleted.

Solving for  $w$  in (14) involves computing the leave-one-out predictive density  $p(y(\ell_i) \mid y_{-i}, M_l)$  for each  $i$  and  $l$ . For each  $l$ , the exact leave-one-out predictive densities require estimating the model



$n$  times as we need to sample from (9) under  $M_l$ , given data  $(y_{-i}, X_{-i}, \tilde{X}_{-i}, \mathcal{L}_{-i})$ , for  $i = 1, \dots, n$ . Following Vehtari et al. (2017), we apply  $K$ -fold cross validation for faster evaluation of these densities. We randomly permute the data and construct  $K$  blocks using consecutive indices. Let  $\chi_{[k]} = (y_{[k]}, X_{[k]}, \tilde{X}_{[k]}, \mathcal{L}_{[k]})$  be the  $k$ th block of size  $n_k$  and  $\chi_{[-k]} = (y_{[-k]}, X_{[-k]}, \tilde{X}_{[-k]}, \mathcal{L}_{[-k]})$  is its complement of size  $(n - n_k)$  for  $k = 1, \dots, K$ . For each  $k$ , we fit  $M_l$  to  $\chi_{[-k]}$  as we draw  $S$  samples  $\{\beta_{k,l}^s, z_{k,l}^s\}_{s=1}^S$  from  $p((\gamma^\top, q^\top)^\top \mid y_{[-k]}, M_l)$  as given in (9).

Estimating  $M_l$  on  $\chi_{[-k]}$  requires sampling  $z_j$  from its marginal prior for each  $j = 1, \dots, r$ . Each step is dominated by the Cholesky decomposition of the  $(n - n_k) \times (n - n_k)$  matrix  $R_j(\mathcal{L}_{[-k]})$ . Instead of computing the Cholesky factor of  $R_j(\mathcal{L}_{[-k]})$  for every  $\chi_{[-k]}$ , which would cost  $O(Kn^3)$  flops, we execute an efficient block Givens rotation (Golub and Van Loan, 2013, Section 5.1.8) for faster evaluation of Cholesky factors of the  $K$  submatrices of  $R_j(\mathcal{L})$  taking  $O(Kn^3/4)$  operations for all the blocks (see Algorithm S2 in Appendix B). Our algorithm is a block-level variant of Kim et al. (2002). For each sample  $z_{k,l}^s = (z_{1,k,l}^{s\top}, \dots, z_{r,k,l}^{s\top})^\top$  where  $z_{k,l}^s$  is  $(n - n_k)r \times 1$ , we draw posterior samples of  $\tilde{z}_{k,l}^s = (\tilde{z}_{1,k,l}^{s\top}, \dots, \tilde{z}_{r,k,l}^{s\top})^\top$ , the spatially-temporally varying  $n_k r \times 1$  regression coefficients at the  $n_k$  left-out locations,  $\mathcal{L}_{[k]}$  using (13) as detailed in Section 4.3, which we use to evaluate  $p(y(\ell_i) \mid y_{-i}, M_l)$ . In particular, if  $y(\ell_i) \in y_{[k]}$ , then

$$p(y(\ell_i) \mid y_{-i}, M_l) \approx \frac{1}{S} \sum_{s=1}^S \text{EF} \left( y(\ell_i) \mid x(\ell_i)^\top \beta_{k,l}^s + \tilde{x}(\ell_i)^\top \tilde{z}_{k,l}^s(\ell_i); b, \psi_y \right), \quad (15)$$

where  $r \times 1$  vector  $\tilde{z}_{k,l}^s(\ell_i) = (\tilde{z}_{1,k,l}^s(\ell_i), \dots, \tilde{z}_{r,k,l}^s(\ell_i))^\top$  and  $\text{EF}(y_0 \mid \eta_0; b, \psi)$  is the density of  $\text{EF}(\eta_0; b, \psi)$  as defined in Section 2.1 evaluated at  $y_0 \in \mathcal{Y}$ . We repeat these steps for each of the  $L$  models and use (15) to solve the optimisation problem (14) using convex programming. Posterior inference for quantities of interest subsequently proceed from the “stacked posterior”,

$$\tilde{p}(\cdot \mid y) = \sum_{l=1}^L \hat{w}_l p(\cdot \mid y, M_l), \quad (16)$$

where  $\hat{w}_l$  are the optimal weights obtained from (14); see Algorithm S3 in Appendix B.

Unlike MCMC, this stacking algorithm easily distributes independent tasks across multiple computing nodes and accrues substantial computational gains. The predictive stacking algorithm expends  $O(L^g + c^{-1} L K r n^3/4)$  flops for some  $g > 0$ , where  $r$ ,  $L$  and  $K$  are as defined earlier and  $c$  is the number of available processing cores. We use the package CVXR (Fu et al., 2020) in the R statistical computing environment by applying disciplined convex programming (2005 Stanford University Department of Electrical Engineering PhD thesis by M. Grant), while also confirming the convexity of the problem, and find the stacking weights in polynomial time  $O(L^g)$  using an interior-point algorithm. We have used the solvers Mosek (ApS, 2023) and ECOSolveR (Fu and Narasimhan, 2023) for obtaining the stacking weights.

## 5. SIMULATION

**5.1. Simulated data.** We evaluate predictive performance of our proposed methods using two simulated datasets each of sample size  $n$ , with spatial coordinates sampled uniformly inside the unit square  $[0, 1]^2$  and temporal coordinates sampled uniformly within  $[0, 1]$ . The simulated datasets

are based on the natural parameter  $\eta(\ell) = x(\ell)^\top \beta + x(\ell)^\top z(\ell)$ , where the latent Gaussian spatial-temporal processes  $z(\ell) = (z_1(\ell), z_2(\ell))^\top$  are modelled as  $z_j(\ell) \sim \text{GP}(0, \sigma_{z_j}^2 R_j(\cdot, \cdot))$  for  $j = 1, 2$ . In both datasets,  $x(\ell)$  is  $2 \times 1$  consisting of an intercept and one predictor sampled from the standard normal distribution. We take  $\sigma_{z_1}^2 = 0.25$ ,  $\sigma_{z_2}^2 = 0.5$  and,  $(\phi_{11}, \phi_{21}) = (0.5, 2)$  and,  $(\phi_{12}, \phi_{22}) = (1, 4)$ . The first simulation considers responses distributed as  $y(\ell) \sim \text{Poisson}(\exp(\eta(\ell)))$  with  $\beta = (5, -0.5)$ . The second simulation considers binomial data with responses  $y(\ell) \sim \text{Binomial}(m(\ell), \pi(\ell))$  with  $m(\ell)$ , the number of trials sampled independently from a Poisson distribution with mean 20 and probability of success  $\pi(\ell) = \text{ilogit}(\eta(\ell))$  with  $\text{ilogit}(t) = \exp(t)/(1 + \exp(t))$ ,  $t \in \mathbb{R}$  and  $\beta = (1, -0.5)$ . The choice of  $\beta$  in each simulated data is such that the generated data do not contain excessive zeros. The simulation experiments are conducted with  $n$  varying from 200 to 600 with a randomly chosen holdout sample of size  $n_h = 100$  over a set of coordinates  $\mathcal{L}_h$ .

**5.2. Posterior inference.** We stack on the parameters  $\Delta = \{\alpha_\epsilon, \sigma_\xi, \theta_{\text{sp}}\}$ , where  $\theta_{\text{sp}}$  denotes the collection of process parameters corresponding to all  $r$  spatial-temporal processes in the model. We fit the model in (7) with appropriate choices of  $\psi_y$  and  $b$ , i.e., for Poisson data,  $\psi_y(t) = \exp(t)$  and  $b = 1_n$ , and, for the binomial data,  $\psi_y(t) = \log(1 + \exp(t))$  and  $b = (m(\ell_1), \dots, m(\ell_n))^\top$ . For the prior on  $\beta$ , we consider  $\mu_\beta = 0$  and,  $L_\beta = I_2$ . For subsequent inference, we fix hyperparameters  $\nu_\beta = \nu_{z_1} = \nu_{z_2} = 3$  and consider the grid of candidate values of  $\sigma_\xi$  as,  $G_{\sigma_\xi} = \{0.5, 1\}$ . The parameters  $\alpha_\epsilon$  and  $\kappa_\epsilon$  appearing in the GCM<sub>c</sub> prior  $p(\xi \mid \beta, z)$  specify the shape and scale parameters  $\alpha^*$  and  $\kappa^*$  of the posterior distribution (9). If the data is on the boundary of the parameter space with a high frequency (e.g., Poisson data with excessive zeros), then the inference is sensitive to the choice of these parameters. This has been investigated using grid-search with an out-of-sample criterion (Bradley and Clinch, 2024). We formally handle the value of these parameters by stacking on several models with different choices of  $\alpha_\epsilon$  and  $\kappa_\epsilon$ . For the simulated Poisson data and the Binomial data, we choose  $\kappa_\epsilon = 0$  and  $\kappa_\epsilon = 2\alpha_\epsilon$  respectively and fix  $G_{\alpha_\epsilon} = \{0.5, 0.75\}$ . Finally, we choose candidate values of the spatial-temporal decay parameters so that the corresponding effective range is between 20% and 70% of the maximum spatial and temporal inter-coordinate distances (see Banerjee et al., 2003, Chapter 2). Choosing a grid of candidate values for each of these  $2r$  process parameters results in the number candidate models to grow exponentially. In order to negotiate this impracticality, we consider, motivated by the weak identifiability of the process parameters, a special case of (7) which we refer to as the “reduced” model. In the reduced model, we shrink the parameter space of  $\theta_{\text{sp}}$  by considering the process parameters common across all  $r$  spatial processes i.e.,  $\theta_{\text{sp}_1} = \theta_{\text{sp}_2} = (\phi_1, \phi_2)$ . We implement our proposed stacking algorithm on the reduced model. We choose  $G_{\phi_1} = \{0.3, 0.7, 1.2\}$  and  $G_{\phi_2} = \{1.5, 3, 4.5\}$  and, hence, we stack on  $2 \times 2 \times 3 \times 3 = 36$  models. We use the value  $K = 10$  for  $K$ -fold cross validation in order to evaluate the leave-one-out predictive densities (Vehtari and Lampinen, 2002).

We also estimate a fully Bayesian model with prior distributions on the spatial process parameters using MCMC for comparison with predictive stacking. In addition to the same priors for the model parameters, as mentioned above, we assign uniform priors  $\text{U}(0.1, 15)$  on  $\phi_{1j}$  and  $\phi_{2j}$  for  $j = 1, 2$ . Here, it is worth remarking that sampling from the joint posterior distribution, for example using random walk Metropolis steps, suffers considerably from mixing and convergence issues because of

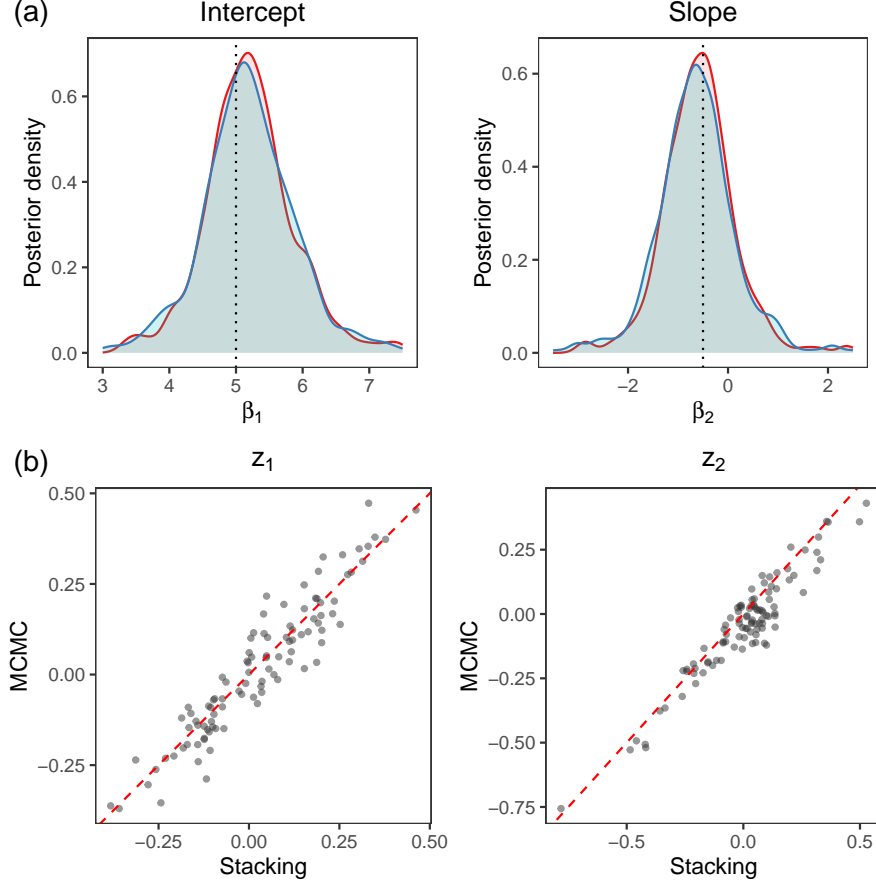


FIGURE 1. Simulated Poisson count data: (a) posterior distributions of the intercept as well as the slope obtained from stacking (blue) and MCMC (red) overlaid dotted vertical line showing their true values; (b) posterior medians of spatial-temporal random effects obtained by stacking and MCMC with the  $y = x$  reference as a red dashed line.

the high-dimensional parameter space and weak identifiability of process parameters. This issue is only partially mitigated using adaptive Metropolis steps (Roberts and Rosenthal, 2009). The Gibbs sampling algorithm involves sampling from the conditional posterior distributions  $p(\theta_{\text{sp}} | \beta, z, \xi, y)$  and  $p(\beta, z, \xi, | \theta_{\text{sp}}, y)$ , where the former involves an adaptive Metropolis update using the R package spBayes (Finley et al., 2015) and, the latter proceeds by sampling from (9) using the projections described in Section 4.2. We evaluate predictive performance of our stacking algorithm with MCMC by computing the mean log-point-wise predictive density at the held out locations, given by  $\text{MLPD} = n_h^{-1} \sum_{\ell \in \mathcal{L}_h} \log \sum_{l=1}^L \hat{w}_l p(y(\ell) | y, M_l)$ .

### 5.3. Results.

5.3.1. *Posterior learning and predictive performance.* Figure 1(a) shows the overlaid posterior densities of the fixed effects obtained from the stacked posterior (blue) and MCMC (red), revealing practically indistinguishable posterior distributions. Figure 1(b) displays high agreement between

Model	Method	Sample size ( $n$ )				
		$n = 100$	$n = 200$	$n = 300$	$n = 400$	$n = 500$
Poisson	Stacking	-7.653	-7.594	-7.579	-7.538	-7.486
	rMCMC	-7.515	-7.398	-7.189	-7.066	-6.957
	MCMC	-6.758	-6.738	-6.736	-6.716	-6.696
Binomial	Stacking	-7.464	-7.449	-7.426	-7.411	-7.398
	rMCMC	-7.428	-7.389	-7.201	-7.182	-7.001
	MCMC	-6.978	-6.813	-6.747	-6.729	-6.711

TABLE 1. Predictive performance of stacking and MCMC. All values correspond to mean log-pointwise predictive density (MLPD) based on 5 replications. Stacking and rMCMC assumes the reduced model with shared process parameters while MCMC assumes separate process parameters.

the posterior medians of the spatial-temporal random effects associated with the intercept and the slope, obtained by our proposed stacking algorithm and MCMC. We also notice that the posterior samples of the spatial-temporal process parameters do not necessarily concentrate around their true values hence, demonstrating their weak identifiability. In addition, we also observe similar patterns in the posterior distributions within the temporal ( $\phi_{11}$  and  $\phi_{12}$ ) and the spatial decay parameters ( $\phi_{21}$  and  $\phi_{22}$ ), which further motivates the assumptions of the reduced model (see, Figure A1 in Appendix C).

Treating the fully Bayesian model with priors on  $\theta_{sp}$  which is fitted using MCMC, we find that the predictive performance of our proposed stacking algorithm gets closer to MCMC as sample size increases (see, Table 1). In Table 1, for Poisson count data, we see that, the difference in mean log-pointwise predictive density of 100 held-out samples between MCMC and our proposed stacking algorithm drops from 11.6% at sample size 100 to 10.5% at sample size 500. Furthermore, we find that the mean log-pointwise predictive density between MCMC on the “full” model (which has  $2r$  process parameters) and MCMC on the reduced model drops from 10.1% at sample size 100 to 3.7% at sample size 500. We observe similar trend in case of the simulated binomial count data.

**5.3.2. Runtime comparison.** Figure 2 reveals that predictive stacking is, on average, about 500 times faster than MCMC further corroborating the efficiency of predictive stacking as an alternative to MCMC. We have implemented our predictive stacking algorithm within the R statistical computing environment. We execute the programs for runtime comparisons on hardware equipped with an 1-Intel(R) Xeon(R) Gold 6140 CPU @ 2.30GHz processor with 36 cores and 1 thread per core, totalling 36 possible threads for parallel computing. The runtime for predictive stacking reported here are based on execution using 6 cores only. We compare the runtime of our proposed algorithm with the adaptive Metropolis-within-Gibbs algorithm that implements full Bayesian inference.




FIGURE 2. Comparison of runtimes of our proposed stacking algorithm (solid), MCMC (dashes) and MCMC on the reduced model (dot-dash). Stacking demonstrates significantly faster execution times across various sample sizes compared to the MCMC algorithms.

## 6. NORTH AMERICAN BREEDING BIRD SURVEY

Terrestrial birds were sampled annually from routes (approximately 40.23 km with point counts every 0.8 km) spread across the United States and Canada as part of the North American Breeding Bird Survey ([Ziolkowski Jr. et al., 2022](#)). We analyse the number of migrant birds observed at different locations across United States with the goal to predict their numbers at an arbitrary location. Our key inferential objective is to evaluate how the impact from nearby vehicles and external noise varies over space and time. For this analysis, we use 2,396 spatial coordinates over the years 2010 through 2019 at which the number of migrant birds were recorded. We aggregated the count data of different species within each spatial unit since very few locations recorded counts for more than one species. We consider two explanatory variables, “car” and “noise”. The variable ‘car’ represents tallies of vehicles passing survey points during each 3-minute count, and the variable ‘noise’ reports unrelated excessive noise at each point from sources other than passing vehicles (for example, from construction work). The presence of excessive noise is defined as noise lasting more than 45 seconds that significantly interferes with the observer’s ability to hear birds at the location during the sampling period. Both these variables are mapped to the same GPS coordinates in the form of latitude and longitude that reference the avian counts.

Statistical analysis of this survey has focused on different species-level analyses with random effects modelling variability within different levels of a factor. For example, [Link and Sauer \(1997\)](#) and [Sauer and Link \(2011\)](#) study route-level population trajectories of a species and observer effects using quasi-likelihood approaches and a Bayesian hierarchical log-linear model respectively. [Sauer](#)

	Median	SD	2.5%	97.5%
(Intercept)	0.891	0.234	0.530	1.207
Car	0.007	0.054	-0.116	0.093
Noise	0.005	0.304	-0.395	0.508
Average Bird Count	15.789	20.677	13.529	31.753

TABLE 2. North American Breeding Bird Survey (2010-19): Posterior summary (95% credible intervals) of global regression coefficients and the average bird count over all spatial-temporal coordinates. SD denotes standard deviation; 2.5% and 97.5% represent quantiles.

and Link (2011) elaborates on the practical difficulties arising from convergence issues compelling a limited number of MCMC iterations for moderately large datasets. Furthermore, their model featured random effects of much lower dimension than we have in our proposed model (7). In addition, none of the previous analyses uses spatial random effects in order to model spatial variability in their log-linear model whereas we utilise available geographic coordinates of the routes to build spatial-temporal processes to account for dependencies in the avian point count data within a rich modelling framework. This framework enables us to account for large scale variation in the mean using explanatory variables (for example the presence of excessive noise from vehicles) so the latent process evinces the residual spatial-temporal association in bird counts that can indicate lurking factors affecting the avian population.

We apply the spatial-temporal Poisson regression model in (7) to analyse the counts of migrant birds at the observed locations with ‘car’ and ‘noise’ as predictors having spatially-temporally varying regression coefficients with the correlation function as given in (8) and with the additional assumption of common process parameters  $\phi_1$  and  $\phi_2$  shared across all the spatial-temporal processes. We implemented our proposed stacking algorithm with  $G_{\phi_2} = \{40, 800, 1000\}$  which considers the effective spatial range to be approximately 20%, 50% and 70% of the maximum inter-site distance,  $G_{\phi_1} = \{0.5, 1, 2\}$ ,  $G_{\sigma_\epsilon} = \{0.5, 1\}$  and  $G_{\alpha_\epsilon} = \{0.5, 0.75\}$ .

Table 2 presents posterior summaries of global regression coefficients (the component that does not vary over space and time) for the intercept, ‘car’ and ‘noise’. We find the global intercept to be significantly larger than zero which contributes approximately 2.5 units to the count with 95% credible interval (1.7, 3.4) in the presence of no passing cars or excessive noise. Neither the number of cars nor levels of excessive noise seem to significantly impact the count in terms of global effects. The last row of Table 2 presents the estimated bird count averaged over all observed spatial-temporal coordinates in the data set. This estimate is obtained from the posterior predictive distribution of the average bird count and provides us with an estimate of the relative influence of the spatially-temporally varying component over the global effects. To be precise, we see that the spatially-temporally varying component of the regression increases the estimate of counts ( $\approx 16$ ) by a factor of 6 times over the 2.5 units contributed by the global effects. This estimate of the average bird count is consistent with what one would obtain from customary generalised linear models with



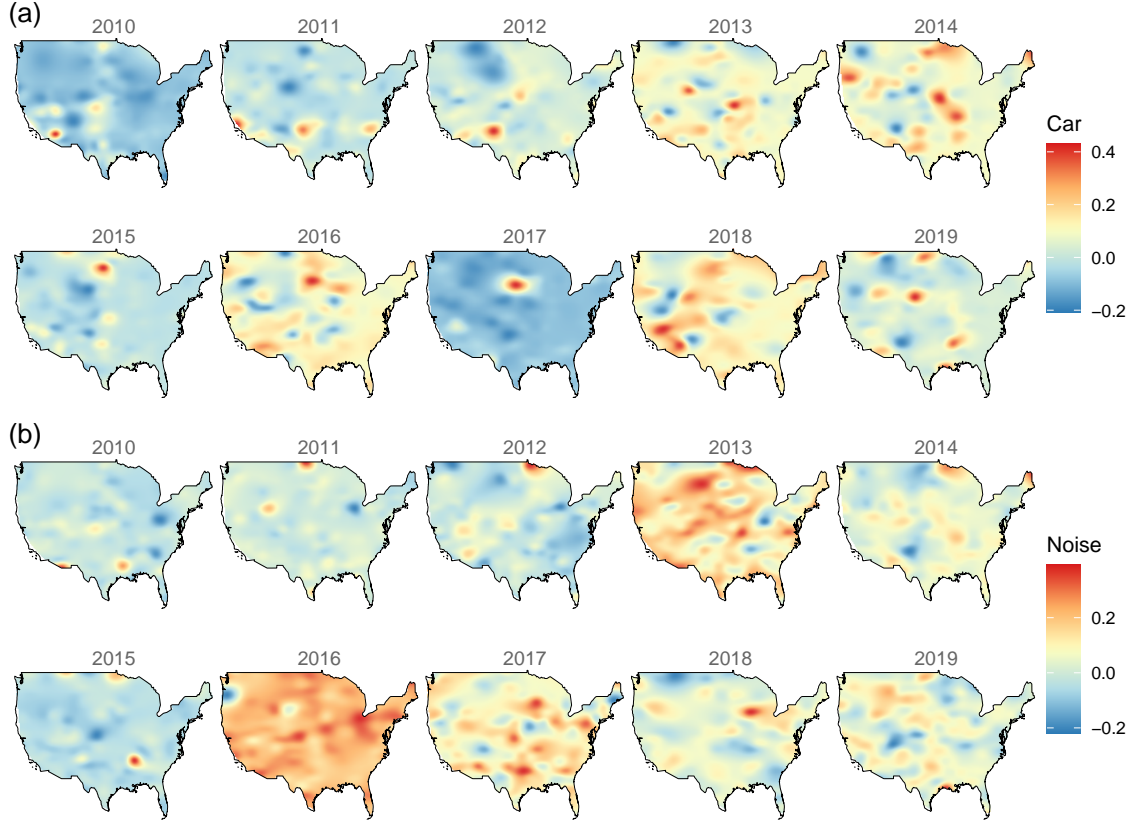


FIGURE 3. North American Breeding Bird Survey (2010-19): Interpolated surfaces of posterior median of spatial-temporal random effects in the slope of the variables (a) ‘car’ and, (b) ‘noise’ obtained by our proposed stacking algorithm.

fixed effects only, but we also offer the spatially and temporally varying impact of predictors that are not captured by the global regression.

Figures 3(a) and (b) reveal how the impact of the predictors ‘car’ and ‘noise’ vary over space and time. For example, in Fig. 3(a) we see significant positive impact (red) of ‘car’ in the Niland and Palo Verde regions of Imperial county, California, (southwest corner of the map) rather consistently between 2010 and 2012. These elevated spatial-temporal coefficients along with higher numbers of cars in these regions (not shown) produce higher than average estimates of bird counts there. We observe a similar pattern in parts of North Dakota and South Dakota between 2015 and 2019. Fig. 3(b) reveals a positive impact (red) of ‘noise’ in the spatial-temporal random effects in northern Minnesota from 2011 to 2015. While this area experiences persistently low levels of ‘noise’ (not shown) the high values of the coefficients produce higher than average estimates of bird counts. In summary, these spatially-temporally varying coefficients represent the local impact of predictors that adjusts the global effects from Table 2. Appendix D presents further details of this analysis. Lastly, we note that while our proposed stacking algorithm delivered inference in 30 minutes using 6 cores, implementing the adaptive Metropolis-within-Gibbs algorithm anticipates at least 15,000

iterations for convergence with each iteration taking 1.3 minutes on average. Hence, our proposed stacking algorithm offers speed-ups of over 650 times over MCMC.

## 7. DISCUSSION

We have devised and demonstrated Bayesian predictive stacking to be an effective tool for estimating spatially-temporally varying regression coefficients and yielding robust predictions for non-Gaussian spatial-temporal data. We develop and exploit analytically accessible distribution theory pertaining to Bayesian analysis of generalised linear mixed models that enables us to directly sample from the posterior distributions. The focus of this article is on effectively combining inference across different closed-form posterior distributions by circumventing inference on weakly identified parameters. Future developments and investigations will consider zero-inflated Bernoulli data for which the current approach will require a boundary adjustment parameter to ensure that the posterior distribution is well-defined. In this context, [Bradley and Clinch \(2024\)](#) explore flexible choices of  $H$  to improve inference for zero-inflated Bernoulli data, although they lose the computational benefits arising from the structured  $H$  in our spatially-temporally varying coefficients model. Exploring predictive stacking for alternate specifications of  $H$  and, in particular, for zero-inflated binary data will comprise future research endeavours. Finally, our method can be adapted to variants of Gaussian process models that scale inference to massive datasets ([Datta et al., 2016](#); [Banerjee, 2017](#); [Heaton et al., 2019](#)) by circumventing the Cholesky decomposition of dense covariance matrices. Other directions can include building more general frameworks to accommodate irregular dispersion and analysis of zero-inflated count data.

## ACKNOWLEDGEMENT

This work used computational and storage services associated with the Hoffman2 Shared Cluster provided by UCLA Office of Advanced Research Computing’s Research Technology Group. Sudipto Banerjee and Soumyakanti Pan were supported by two research grants from the National Institute of Environmental Health Sciences (NIEHS), one grant from the National Institute of General Medical Science (NIGMS) and another from the Division of Mathematical Sciences of the National Science Foundation (NSF-DMS). Jonathan Bradley was supported by NSF-DMS.

## DATA AND CODE AVAILABILITY

Codes and data required to reproduce the results and findings in this article are openly available on GitHub ([https://github.com/SPan-18/spatialGLM\\_stacking](https://github.com/SPan-18/spatialGLM_stacking)). The North American Breeding Bird survey database is openly available at <https://doi.org/10.5066/P97WAZE5>.

## CONFLICTS OF INTEREST

All authors declare that they have no conflicts of interest.

## REFERENCES

- Milton Abramowitz and Irene A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables*. Dover Publications, Inc., New York, 1965.
- MOSEK ApS. *MOSEK Rmosek package 9.3.22*, 2023. URL <https://docs.mosek.com/9.3/rmosek/index.html>.
- Sudipto Banerjee. High-Dimensional Bayesian Geostatistics. *Bayesian Analysis*, 12(2):583 – 614, 2017. doi: 10.1214/17-BA1056R.
- Sudipto Banerjee, Brad P. Carlin, and Alan E. Gelfand. *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC, 2003. doi: 10.1201/9780203487808.
- Jonathan R. Bradley and Madelyn Clinch. Generating independent replicates directly from the posterior distribution for a class of spatial hierarchical models. *Journal of Computational and Graphical Statistics*, 2024. (in press).
- Jonathan R. Bradley, Scott H. Holan, and Christopher K. Wikle. 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, 2020. doi: 10.1080/01621459.2019.1677471.
- Leo Breiman. Stacked regressions. *Machine learning*, 24(1):49–64, 1996.
- Merlise Clyde and Edwin S Iversen. Bayesian model averaging in the m-open framework. *Bayesian theory and applications*, 14(4):483–498, 2013. doi: <http://dx.doi.org/10.1093/acprof:oso/9780199695607.003.0024>.
- Abhirup Datta, Sudipto Banerjee, Andrew O. Finley, and Alan E. Gelfand. Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets. *Journal of the American Statistical Association*, 111(514):800–812, 2016. doi: 10.1080/01621459.2015.1044091. URL <https://doi.org/10.1080/01621459.2015.1044091>. PMID: 29720777.
- Victor De Oliveira. Bayesian prediction of clipped gaussian random fields. *Computational Statistics & Data Analysis*, 34(3):299–314, 2000.
- Victor De Oliveira, Benjamin Kedem, and David A. Short. Bayesian prediction of transformed gaussian random fields. *Journal of the American Statistical Association*, 92(440):1422–1433, 1997. doi: 10.1080/01621459.1997.10473663. URL <https://doi.org/10.1080/01621459.1997.10473663>.
- Persi Diaconis and Donald Ylvisaker. Conjugate Priors for Exponential Families. *The Annals of Statistics*, 7(2):269 – 281, 1979. doi: 10.1214/aos/1176344611.
- P. J. Diggle, J. A. Tawn, and R. A. Moyeed. Model-based geostatistics. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 47(3):299–350, 1998. ISSN 00359254, 14679876. URL <http://www.jstor.org/stable/2986101>.
- Peng Ding. On the conditional distribution of the multivariate t distribution. *The American Statistician*, 70(3):293–295, 2016. doi: 10.1080/00031305.2016.1164756. URL <https://doi.org/10.1080/00031305.2016.1164756>.
- Andrew O. Finley, Sudipto Banerjee, and Alan E. Gelfand. spbayes for large univariate and multivariate point-referenced spatio-temporal data models. *Journal of Statistical Software*, 63(13):

- 1–28, 2015. doi: 10.18637/jss.v063.i13. URL <https://www.jstatsoft.org/index.php/jss/article/view/v063i13>.
- Anqi Fu and Balasubramanian Narasimhan. *ECOSolveR: Embedded Conic Solver in R*, 2023. URL <https://bnaras.github.io/ECOSolveR/>. R package version 0.5.5.
- Anqi Fu, Balasubramanian Narasimhan, and Stephen Boyd. CVXR: An R package for disciplined convex optimization. *Journal of Statistical Software*, 94(14):1–34, 2020. doi: 10.18637/jss.v094.i14.
- Alan E. Gelfand, Hyon-Jung Kim, C. F. Sirmans, and Sudipto Banerjee. Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association*, 98(462):387–396, 2003. doi: 10.1198/016214503000170. URL <https://doi.org/10.1198/016214503000170>.
- T. Gneiting and P. Guttorp. Continuous-parameter spatio-temporal processes. In A.E. Gelfand, P.J. Diggle, M. Fuentes, and P. Guttorp, editors, *Handbook of Spatial Statistics*, Chapman & Hall CRC Handbooks of Modern Statistical Methods, pages 427–436. Taylor and Francis, 2010.
- Gene H. Golub and Charles F. Van Loan. *Matrix Computations - 4th Edition*. Johns Hopkins University Press, Philadelphia, PA, 2013. doi: 10.1137/1.9781421407944. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781421407944>.
- M. Haran. Gaussian random field models for spatial data. In *Markov chain Monte Carlo Handbook* Eds. Brooks, S.P., Gelman, A.E. Jones, G.L. and Meng, X.L., pages 449–478. Chapman and Hall/CRC, 2011.
- Patrick J Heagerty and Subhash R Lele. A composite likelihood approach to binary spatial data. *Journal of the American Statistical Association*, 93(443):1099–1111, 1998.
- Matthew J. Heaton, Abhirup Datta, Andrew O. Finley, Reinhard Furrer, Joseph Guinness, Rajarshi Guhaniyogi, Florian Gerber, Robert B. Gramacy, Dorit Hammerling, Matthias Katzfuss, Finn Lindgren, Douglas W. Nychka, Furong Sun, and Andrew Zammit-Mangion. A case study competition among methods for analyzing large spatial data. *Journal of Agricultural, Biological, and Environmental Statistics*, 24(3):pp. 398–425, 2019. ISSN 10857117, 15372693. URL <https://www.jstor.org/stable/48702926>.
- Jennifer A. Hoeting, David Madigan, Adrian E. Raftery, and Chris T. Volinsky. Bayesian model averaging: a tutorial (with comments by M. Clyde, David Draper and E. I. George, and a rejoinder by the authors. *Statistical Science*, 14(4):382 – 417, 1999. doi: 10.1214/ss/1009212519. URL <https://doi.org/10.1214/ss/1009212519>.
- John Hughes and Murali Haran. 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, January 2013. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2012.01041.x.
- Tae Yoon Kim, Jeong Soo Park, and Dennis D. Cox. Fast algorithm for cross-validation of the best linear unbiased predictor. *Journal of Computational and Graphical Statistics*, 11(4):823–835, 2002. ISSN 10618600. URL <http://www.jstor.org/stable/1391163>.
- Tri Le and Bertrand Clarke. A bayes interpretation of stacking for m-complete and m-open settings. *Bayesian Analysis*, 12(3):807–829, 2017.

- William A. Link and John R. Sauer. Estimation of population trajectories from count data. *Biometrics*, 53(2):488–497, 1997. ISSN 0006341X, 15410420. URL <http://www.jstor.org/stable/2533952>.
- David Madigan, Adrian E Raftery, C Volinsky, and Jennifer Hoeting. Bayesian model averaging. In *Proceedings of the AAAI Workshop on Integrating Multiple Learned Models, Portland, OR*, pages 77–83, 1996.
- Charles E. McCulloch and Shayle R. Searle. *Generalized, linear, and mixed models*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2001. doi: <https://doi.org/10.1002/0471722073>.
- Gareth O. Roberts and Jeffrey S. Rosenthal. Examples of adaptive mcmc. *Journal of Computational and Graphical Statistics*, 18(2):349–367, 2009. doi: 10.1198/jcgs.2009.06134. URL <https://doi.org/10.1198/jcgs.2009.06134>.
- Arkajyoti Saha, Abhirup Datta, and Sudipto Banerjee. Scalable predictions for spatial probit linear mixed models using nearest neighbor gaussian processes. *Journal of Data Science*, 20(4):533–544, 2022. ISSN 1680-743X. doi: 10.6339/22-JDS1073.
- John R. Sauer and William A. Link. Analysis of the north american breeding bird survey using hierarchical models - análisis de los censos de aves reproductivas de norte américa mediante modelos jerárquicos. *The Auk*, 128(1):87–98, 2011. ISSN 00048038, 19384254. URL <http://www.jstor.org/stable/10.1525/auk.2010.09220>.
- Aki Vehtari and Jouko Lampinen. Bayesian Model Assessment and Comparison Using Cross-Validation Predictive Densities. *Neural Computation*, 14(10):2439–2468, 10 2002. ISSN 0899-7667. doi: 10.1162/08997660260293292. URL <https://doi.org/10.1162/08997660260293292>.
- Aki Vehtari, Andrew Gelman, and Jonah Gabry. Practical Bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing*, 27(5):1413–1432, sep 2017. ISSN 0960-3174. doi: 10.1007/s11222-016-9696-4. URL <https://doi.org/10.1007/s11222-016-9696-4>.
- David H Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992.
- Yuling Yao, Aki Vehtari, Daniel Simpson, and Andrew Gelman. Using stacking to average Bayesian predictive distributions (with discussion). *Bayesian Analysis*, 13(3):917–1007, 2018.
- Yuling Yao, Aki Vehtari, and Andrew Gelman. Stacking for non-mixing Bayesian computations: The curse and blessing of multimodal posteriors. *arXiv:2006.12335*, 2020.
- Yuling Yao, Gregor Pirš, Aki Vehtari, and Andrew Gelman. Bayesian hierarchical stacking: Some models are (somewhere) useful. *Bayesian Analysis*, 1(1):1–29, 2021.
- Lu Zhang, Wenpin Tang, and Sudipto Banerjee. Bayesian geostatistics using predictive stacking, 2023.
- Zhongwei Zhang, Reinaldo B Arellano-Valle, Marc G Genton, and Raphaël Huser. Tractable bayes of skew-elliptical link models for correlated binary data. *Biometrics (To appear in)*, 2022.
- David Ziolkowski Jr., Michael Lutmerding, Veronica Aponte, and Marie-Anne Hudson. North American Breeding Bird Survey Dataset 1966 - 2021: U.S. Geological Survey data release, 2022. URL <https://doi.org/10.5066/P97WAZE5>.

## Appendix to “Bayesian inference for spatial-temporal non-Gaussian data using predictive stacking”

### APPENDIX A. DISTRIBUTION THEORY

We outline relevant distribution theory for deriving the posterior distribution of different model parameters in the Bayesian hierarchical model (7) in the main article. Our contribution lies in proving two important results, Proposition A1 and Lemma A1, that justifies constructing the improper prior for the parameter  $q$ , also referred to as the discrepancy parameter (Bradley et al., 2020) and has not been addressed hitherto (Bradley and Clinch, 2024, Theorem 3.1). We assume familiarity with the notations introduced in Section 3 of the main article.

**Proposition A1.** *If  $H = [(I_n : X_{n \times p} : \tilde{X}_{n \times nr})^\top : I_{n+p+nr}]^\top$  and  $Q = [Q_1^\top : Q_2^\top]^\top$ , where  $Q$  is  $(2n + p + nr) \times n$  with  $Q_1$  being  $n \times n$  and  $Q_2$  being  $(n + p + nr) \times n$ , such that the  $n$  columns of  $Q$  are the unit norm orthogonal eigenvectors of the orthogonal projector,  $P_H$ , on the column space of  $H$ , then  $\text{rank}(Q_1) = n$ .*

*Proof.* From our definitions, we note that  $QQ^\top = I_{2n+p+nr} - P_H$ ,  $P_H = H(H^\top H)^{-1}H^\top$ ,  $Q^\top H = 0$  and  $\text{rank}(Q) = n$ . Define  $H_1 = [I_n : X : \tilde{X}]$ . It follows from  $Q^\top H = 0$  that,  $Q_1^\top H_1 + Q_2^\top = 0$  and subsequently,  $Q_2 = -H_1^\top Q_1$ , implying that  $\mathcal{R}(Q_2) \subseteq \mathcal{R}(Q_1)$  where  $\mathcal{R}$  denotes row space of a matrix. Hence the rank of  $Q$  must equal  $\dim(\mathcal{R}(Q_1))$ , the number of independent rows of  $Q_1$ . As  $\text{rank}(Q) = n$ , we have  $\text{rank}(Q_1) = \dim(\mathcal{R}(Q_1)) = n$ .  $\square$

**Lemma A1.** *In the hierarchical model (7), assumption of a vague prior on the parameter  $\mu$  leads to the improper prior on the parameter  $q$ , given by  $p(q) \propto 1$ .*

*Proof.* Recall that  $q = -Q^\top \tilde{\mu}$  where  $\tilde{\mu} = (\mu^\top, \mu_\gamma^\top)^\top$  with  $\mu_\gamma = (\mu_\xi^\top, \mu_\beta^\top, \mu_z^\top)^\top$  and  $Q$  is the  $(2n + p + nr) \times n$  matrix defined in Proposition A1. Partition  $Q = [Q_1^\top : Q_2^\top]^\top$  where  $Q_1$  is  $n \times n$  and  $Q_2$  is  $(n + p + nr) \times n$  to get  $q = -Q_1^\top \mu - Q_2^\top \mu_\gamma$ . Consider the sequence of priors on the discrepancy parameter as  $p_{1,k}(\mu) = \mathcal{N}(\mu \mid 0_n, \tau_k I_n)$  for a real sequence  $\{\tau_k\}_{k \in \mathbb{N}}$  with  $\tau_k > 0, \forall k \in \mathbb{N}$  such that  $\lim_k \tau_k = +\infty$ . By Proposition A1, for any  $k$ , the prior on  $q$  induced by  $p_{1,k}$  has the density  $p_{2,k}(q) = \mathcal{N}(q \mid -Q_2^\top \mu_\gamma, \tau_k Q_1^\top Q_1)$ . Hence, the improper density  $p_2(q) \propto 1$  arises from  $\lim_k p_{2,k}(q) = p_2(q)$  for any  $q \in \mathbb{R}^n$ .  $\square$

Next, we derive the posterior distribution in (9) corresponding to the hierarchical model (7) in the main article. The following result adapts the spatial intercept model of Theorem 3.1 in Bradley and Clinch (2024) to our spatially-temporally varying coefficient model. This extension is not entirely straightforward because of the introduction of possibly different explanatory variables in the fixed regression ( $X$ ) and the spatially-temporally varying component ( $\tilde{X}$ ) in our model.

**Theorem A1.** *The hierarchical model (7), with  $p(q) \propto 1$  leads to*

$$p\left((\gamma^\top, q^\top)^\top \mid y, \theta\right) \propto \text{GCM}\left((\gamma^\top, q^\top)^\top \mid 0_{2n+p+nr}, V^*(\theta), \alpha^*, \kappa^*, \psi^*\right),$$

where  $V^*(\theta)^{-1} = L(\theta)^{-1}[H : Q]$ ,  $(3n + p) \times (3n + p)$  block-diagonal matrix  $L(\theta) = \text{blkdiag}(I_n, \sigma_\xi I_n, \sigma_\beta L_\beta, L_z(\theta))$  with  $\sigma_{z_j} L_{z_j}(\theta_{sp})$  as the  $j$ th block of  $nr \times nr$  block-diagonal matrix  $L_z(\theta)$ . The shape



and scale parameters are  $\alpha^* = ((y + \alpha_\epsilon 1_n)^\top, 0_n^\top, \alpha_\beta^\top, \alpha_z^\top)^\top$  and  $\kappa^* = ((b + \kappa_\epsilon)1_n^\top, (1/2)1_n^\top, \kappa_\beta^\top, \kappa_z^\top)^\top$ . The unit log partition function  $\psi^*$  is  $\psi^*(h) = (\psi_y(h_1)^\top, \psi_1(h_2)^\top, \psi_\beta(h_3)^\top, \psi_z(h_4)^\top)$  for some  $h = (h_1^\top, h_2^\top, h_3^\top, h_4^\top)^\top$  with  $h_1, h_2 \in \mathbb{R}^n$ ,  $h_3 \in \mathbb{R}^p$ ,  $h_4 \in \mathbb{R}^{nr}$ , and,  $\psi_z(h_4) = (\psi_{z_1}(h_{41}), \dots, \psi_{z_r}(h_{4r}))$  for  $h_{4j} \in \mathbb{R}^n$ , where all of the log partition functions operate element-wise on the arguments.

*Proof.* Define the  $(n + p + nr) \times 1$  random vector  $\gamma = (\xi^\top, \beta^\top, z^\top)^\top$ . Then the likelihood is

$$\begin{aligned} p(y | \gamma, \mu) &= \exp \left\{ y^\top (X\beta + \tilde{X}z + \xi - \mu) - b1_n^\top \psi_y(X\beta + \tilde{X}z + \xi - \mu) \right\} \\ &= \exp \left\{ y^\top \left( \begin{bmatrix} I_n & X & \tilde{X} \end{bmatrix} \gamma - \mu \right) - b1_n^\top \psi_y \left( \begin{bmatrix} I_n & X & \tilde{X} \end{bmatrix} \gamma - \mu \right) \right\}. \end{aligned} \quad (\text{A1})$$

According to Section 2.2, the densities of CM priors for  $\beta | \theta$  and  $z | \theta$  are

$$\begin{aligned} p(\beta | \theta) &\propto \exp \left\{ \alpha_\beta^\top L_\beta(\theta)^{-1} (\beta - \mu_\beta) - \kappa_\beta^\top \psi_\beta(L_\beta(\theta)^{-1} (\beta - \mu_\beta)) \right\}, \\ p(z | \theta) &\propto \exp \left\{ \alpha_z^\top L_z(\theta)^{-1} (z - \mu_z) - \kappa_z^\top \psi_z(L_z(\theta)^{-1} (z - \mu_z)) \right\}, \end{aligned} \quad (\text{A2})$$

where  $nr \times nr$  matrix  $L_z(\theta) = \oplus_{j=1}^r \sigma_{z_j} L_{z_j}(\theta_{\text{sp}_j})$ . And the conditional prior  $p(\xi | \beta, z, \theta)$  is

$$\begin{aligned} p(\xi | \beta, z, \mu, \theta) &\propto \exp \left\{ \alpha_\xi^\top \left( \begin{bmatrix} I_n \\ L_\xi(\theta)^{-1} \end{bmatrix} \xi - \begin{bmatrix} \mu - X\beta - \tilde{X}z \\ L_\xi(\theta)^{-1} \mu_\xi \end{bmatrix} \right) \right. \\ &\quad \left. - \kappa_\xi^\top \psi_\xi \left( \begin{bmatrix} I_n \\ L_\xi(\theta)^{-1} \end{bmatrix} \xi - \begin{bmatrix} \mu - X\beta - \tilde{X}z \\ L_\xi(\theta)^{-1} \mu_\xi \end{bmatrix} \right) \right\} \\ &= \exp \left\{ \alpha_\xi^\top \left( \begin{bmatrix} I_n & X & \tilde{X} \\ L_\xi(\theta)^{-1} & 0 & 0 \end{bmatrix} \gamma - \begin{bmatrix} \mu \\ L_\xi(\theta)^{-1} \mu_\xi \end{bmatrix} \right) \right. \\ &\quad \left. - \kappa_\xi^\top \psi_\xi \left( \begin{bmatrix} I_n & X & \tilde{X} \\ L_\xi(\theta)^{-1} & 0 & 0 \end{bmatrix} \gamma - \begin{bmatrix} \mu \\ L_\xi(\theta)^{-1} \mu_\xi \end{bmatrix} \right) \right\}, \end{aligned} \quad (\text{A3})$$

where the last step follows from equality

$$\begin{bmatrix} I_n \\ L_\xi(\theta)^{-1} \end{bmatrix} \xi - \begin{bmatrix} \mu - X\beta - \tilde{X}z \\ L_\xi(\theta)^{-1} \mu_\xi \end{bmatrix} = \begin{bmatrix} I_n & X & \tilde{X} \\ L_\xi(\theta)^{-1} & 0 & 0 \end{bmatrix} \gamma - \begin{bmatrix} \mu \\ L_\xi(\theta)^{-1} \mu_\xi \end{bmatrix}.$$

Hence, combining (A2) and (A3) yields the joint prior density  $p(\gamma | \mu, \theta)$  as

$$\begin{aligned} p(\gamma | \mu, \theta) &\propto \exp \left\{ (\alpha_\xi^\top, \alpha_\beta^\top, \alpha_z^\top) \left( \begin{bmatrix} I_n & X & \tilde{X} \\ L_\xi(\theta)^{-1} & 0 & 0 \\ 0 & L_\beta(\theta)^{-1} & 0 \\ 0 & 0 & L_z(\theta)^{-1} \end{bmatrix} \begin{bmatrix} \xi \\ \beta \\ z \end{bmatrix} - \begin{bmatrix} \mu \\ L_\xi(\theta)^{-1} \mu_\xi \\ L_\beta(\theta)^{-1} \mu_\beta \\ L_z(\theta)^{-1} \mu_z \end{bmatrix} \right) \right. \\ &\quad \left. - (\kappa_\xi^\top, \kappa_\beta^\top, \kappa_z^\top) \psi^* \left( \begin{bmatrix} I_n & X & \tilde{X} \\ L_\xi(\theta)^{-1} & 0 & 0 \\ 0 & L_\beta(\theta)^{-1} & 0 \\ 0 & 0 & L_z(\theta)^{-1} \end{bmatrix} \begin{bmatrix} \xi \\ \beta \\ z \end{bmatrix} - \begin{bmatrix} \mu \\ L_\xi(\theta)^{-1} \mu_\xi \\ L_\beta(\theta)^{-1} \mu_\beta \\ L_z(\theta)^{-1} \mu_z \end{bmatrix} \right) \right\}, \end{aligned} \quad (\text{A4})$$

where  $\psi^*(h) = (\psi_y(h_1)^\top, \psi_1(h_2)^\top, \psi_\beta(h_3)^\top, \psi_z(h_4)^\top)$  for some  $h = (h_1^\top, h_2^\top, h_3^\top, h_4^\top)^\top$  with  $h_1, h_2 \in \mathbb{R}^n$ ,  $h_3 \in \mathbb{R}^p$ ,  $h_4 \in \mathbb{R}^{nr}$ , and,  $\psi_z(h_4) = (\psi_{z_1}(h_{41}), \dots, \psi_{z_r}(h_{4r}))$  for  $h_{4j} \in \mathbb{R}^n$ , where all the log

partition functions operates element-wise on the arguments. Hence, assuming  $p(q) \propto 1$ , we obtain

$$\begin{aligned} p\left((\gamma^\top, q^\top)^\top \mid y, \theta\right) &\propto \exp\left\{(y^\top + \alpha_\epsilon 1_n^\top, 0_n^\top, \alpha_\beta^\top, \alpha_z^\top) L(\theta)^{-1} (H\gamma - \tilde{\mu}) \right. \\ &\quad \left. - ((b + \kappa_\epsilon) 1_n^\top, \frac{1}{2} 1_n^\top, \kappa_\beta^\top, \kappa_z^\top) \psi^* L(\theta)^{-1} (H\gamma - \tilde{\mu})\right\} \\ &= \exp\left\{\alpha^{*\top} L(\theta)^{-1} (H\gamma - \tilde{\mu}) - \kappa^{*\top} \psi^* (L(\theta)^{-1} (H\gamma - \tilde{\mu}))\right\} \end{aligned} \quad (\text{A5})$$

by combining (A1) and (A4). Now, reparameterise by  $q = -Q^\top \tilde{\mu}$ , which, by definition of  $Q$  (see Section 3.2), is equivalent to  $\tilde{\mu} = -Qq$ , and, thus from (A5), we obtain

$$\begin{aligned} p\left((\gamma^\top, q^\top)^\top \mid y, \theta\right) &= \exp\left\{\alpha^{*\top} L(\theta)^{-1} [H : Q] \begin{bmatrix} \gamma \\ q \end{bmatrix} - \kappa^{*\top} \psi^* \left(L(\theta)^{-1} [H : Q] \begin{bmatrix} \gamma \\ q \end{bmatrix}\right)\right\} \\ &\propto \text{GCM}\left((\gamma^\top, q^\top)^\top \mid 0_{3n+p}, V^*(\theta), \alpha^*, \kappa^*; \psi^*\right). \end{aligned} \quad (\text{A6})$$

□

## APPENDIX B. COMPUTATIONAL DETAILS

This section evolves as follows. First, we present Theorem A2 that provides a computationally efficient way to evaluate the projection expression in (11) in the main article and subsequently provide a memory-efficient algorithm implementing Theorem A2. Next, we outline the algorithm for fast Cholesky factor updates required for the cross-validation step, as detailed in Section 4.4 of the main article. Lastly, we summarise our proposed stacking algorithm in Algorithm S3.

Once the Cholesky factors are computed, we are able to obtain the samples,  $\tilde{v}_z$ , from the marginal prior of  $z$ . Given  $\tilde{v}$ , following Section 4.2, it is clear that naive evaluation of  $\gamma^{(r)} = (H^\top H)^{-1} H^\top \tilde{v}$  requires expensive matrix operations on an  $(n + p + nr) \times (n + p + nr)$  matrix  $H^\top H$ . Theorem A2 derives a computationally efficient algorithm to evaluate this projection.

**Theorem A2.** *Given  $\tilde{v}$ , the projection expression in (11) can be computed in  $O(nr^3 + p^3)$  flops and requires a storage of  $O(np + nr + nr^2 + p^2)$ .*

*Proof.* Note that  $H = [(I_n : X : \tilde{X})^\top : I_{n+p+nr}]^\top$  is a  $(2n + p + nr) \times (n + p + nr)$  matrix where  $\tilde{X} = [\text{diag}(\tilde{x}_1) : \dots : \text{diag}(\tilde{x}_r)]$  is  $n \times nr$  with  $\text{diag}(\tilde{x}_j)$  being an  $n \times n$  diagonal matrix whose  $i$ th diagonal element is  $\tilde{x}_j(\ell_i)$ . Without loss of generality, we consider a permutation of the columns of  $\tilde{X}$  and a subsequent suitable permutation on  $z = (z_1^\top, \dots, z_r^\top)^\top$  such that the new  $\tilde{X} = \oplus_{i=1}^n \tilde{x}_i^\top$  and  $z = (z_{(1)}^\top, \dots, z_{(n)}^\top)^\top$  where  $z_{(i)} = (z_1(\ell_i), \dots, z_r(\ell_i))^\top$ . Note that we can always rearrange the obtained vector back to the original ordering by applying the inverse of the permutation. The idea of the proof follows from inversion of a partitioned matrix. Write

$$H^\top H = \begin{bmatrix} 2I_n & X & \tilde{X} \\ X^\top & X^\top X + I_p & X^\top \tilde{X} \\ \tilde{X}^\top & \tilde{X}^\top X & \tilde{X}^\top \tilde{X} + I_{nr} \end{bmatrix} = \begin{bmatrix} A & B^\top \\ B & D \end{bmatrix}, \quad (\text{A7})$$

where  $B = [\tilde{X}^\top : \tilde{X}^\top X]$  and  $D = \tilde{X}^\top \tilde{X} + I_{nr}$ . Note that  $(H^\top H)^{-1} H^\top \tilde{v}$  can be simplified as

$$(H^\top H)^{-1} H^\top \tilde{v} = \begin{bmatrix} S_A^{-1} (v_1 - B^\top D^{-1} v_2) \\ -D^{-1} B S_A^{-1} (v_1 - B^\top D^{-1} v_2) + D^{-1} v_2 \end{bmatrix}, \quad (\text{A8})$$

where  $v_1 = ((\tilde{v}_\eta + \tilde{v}_\xi)^\top, (X^\top \tilde{v}_\eta + \tilde{v}_\beta)^\top)^\top$ ,  $v_2 = (\tilde{X}^\top \tilde{v}_\eta + \tilde{v}_z)^\top$  and  $S_A$  denotes the Schur complement of  $A$  in (A7). Define  $v_3 = v_1 - B^\top D^{-1} v_2$  and rewrite it as  $v_1 - [(\tilde{X} v_4)^\top, (X^\top \tilde{X} v_4)^\top]^\top$  where  $v_4 = (\tilde{X}^\top \tilde{X} + I_{nr})^{-1} v_2$ . Note that  $v_4$  can be evaluated efficiently since  $(\tilde{X}^\top \tilde{X} + I_{nr})^{-1} = \oplus_{i=1}^n (\tilde{x}_i \tilde{x}_i^\top + I_r)^{-1}$  is block diagonal. This involves  $n$  inversions of  $r \times r$  positive definite matrices and hence evaluation of  $v_4$  involves  $O(nr^3)$  operations. Next, we move on to inversion of the Schur complement  $S_A = A - B^\top D^{-1} B$ . Note that  $B^\top D^{-1} B = [I_n : X]^\top \tilde{X} (\tilde{X}^\top \tilde{X} + I_{nr})^{-1} \tilde{X}^\top [I_n : X]$ . We use  $(\tilde{X}^\top \tilde{X} + I_{nr})^{-1} \tilde{X}^\top = \tilde{X}^\top (\tilde{X} \tilde{X}^\top + I_n)^{-1}$  which follows from the equality  $\tilde{X}^\top (\tilde{X} \tilde{X}^\top + I_n) = (\tilde{X}^\top \tilde{X} + I_{nr}) \tilde{X}^\top$ . Hence we have  $\tilde{X} (\tilde{X}^\top \tilde{X} + I_{nr})^{-1} \tilde{X}^\top = \tilde{X} \tilde{X}^\top (\tilde{X} \tilde{X}^\top + I_n)^{-1} = \text{diag}(\|\tilde{x}_i\|^2 / (1 + \|\tilde{x}_i\|^2), i = 1, \dots, n) = \tilde{P}$  (say). Then, writing  $v_3 = (v_{51}^\top, v_{52}^\top)^\top$ ,

$$S_A = \begin{bmatrix} A_* & B_*^\top \\ B_* & D_* \end{bmatrix}, \quad S_A^{-1} v_3 = \begin{bmatrix} A_*^{-1} v_{51} + A_*^{-1} B_*^\top S_{D_*}^{-1} (B_* A_*^{-1} v_{51} - v_{52}) \\ -S_{D_*}^{-1} (B_* A_*^{-1} v_{51} - v_{52}) \end{bmatrix}, \quad (\text{A9})$$

where  $v_{51}$  is  $n \times 1$ ,  $v_{52}$  is  $p \times 1$  and,  $A_* = 2I_n - \tilde{P}$ ,  $B_* = X^\top (I_n - \tilde{P})$ ,  $D_* = X^\top (I_n - \tilde{P}) X + I_p$ . Note that  $A_*^{-1} = (2I_n - \tilde{P})^{-1}$  is again a diagonal matrix with diagonal elements  $(1 + \|\tilde{x}_i\|^2) / (2 + \|\tilde{x}_i\|^2)$  for  $i = 1, \dots, n$ . Hence, we evaluate  $v_6 = B_* A_*^{-1} v_{51} - v_{52} = X^\top \text{diag}(1 / (2 + \|\tilde{x}_i\|^2)) v_{51} - v_{52}$ . Finally, we find the inverse of  $S_{D_*}$ , the  $p \times p$  Schur complement of  $D_*$ , by simplifying as

$$\begin{aligned} S_{D_*} &= D_* - B_* A_*^{-1} B_*^\top \\ &= X^\top (I - \tilde{P}) X + I_p - X^\top (I_n - \tilde{P}) (2I_n - \tilde{P})^{-1} (I_n - \tilde{P}) X \\ &= I_p + X^\top \left[ \text{diag} \left( \frac{1}{1 + \|\tilde{x}_i\|^2} \right) - \text{diag} \left( \frac{1}{1 + \|\tilde{x}_i\|^2} \right) \text{diag} \left( \frac{1 + \|\tilde{x}_i\|^2}{2 + \|\tilde{x}_i\|^2} \right) \text{diag} \left( \frac{1}{1 + \|\tilde{x}_i\|^2} \right) \right] X \\ &= I_p + X^\top \text{diag} \left( \frac{1}{2 + \|\tilde{x}_i\|^2} \right) X. \end{aligned} \quad (\text{A10})$$

Inverting  $S_{D_*}$  takes  $O(p^3)$  operations thus completing the evaluation of  $S_A^{-1} v_3$ . We subsequently find  $D^{-1} B S_A^{-1} v_3$  using the already computed block inverses of the matrix  $D$ .  $\square$

Next, Algorithm S1 provides a memory-efficient step-by-step implementation of Theorem A2. It proceeds by considering the  $n \times nr$  predictor matrix  $\tilde{X}$  with permuted columns as described in Theorem A2. Subsequently  $\tilde{v}_z$  is also permuted adapting to the new  $\tilde{X}$ . This requires memory allocations for one  $nr \times 1$  vector  $v_1$ , two  $n \times 1$  vectors  $v_2$  and  $v_4$ , one  $p \times 1$  vector  $v_3$ , an  $n$ -dimensional list  $\{D_i : i = 1, \dots, n\}$  with each item storing one  $r \times r$  matrix and one  $p \times p$  matrix  $D_*$ , amounting to allocated storage of  $O(n + p + nr + p^2 + nr^2)$ . Throughout Algorithm S1, while  $v_{2,i}$  and  $v_{4,i}$  correspond to the  $i$ th element of the  $n \times 1$  vectors  $v_2$  and  $v_4$ , respectively,  $v_{1,i}$  denotes the  $i$ th  $r \times 1$  block of the  $nr \times 1$  vector  $v_1$  corresponding to the indices running from  $(i-1)r + 1$  to  $ir$  for each  $i = 1, \dots, n$ . The functions **chol** and **trsolve** compute the standard Cholesky decomposition and the solution of a triangular linear system, respectively, and are readily available in numerical linear algebra libraries. Algorithm S1 has been designed to avoid expensive linear algebra, consisting of  $n$  Cholesky decompositions of  $r \times r$  matrices, one Cholesky decomposition of a  $p \times p$  matrix,  $2n$  triangular solvers each of order  $O(r)$ , and one additional  $O(p)$  triangular solve.

Next, we supply a block Givens rotation algorithm (Golub and Van Loan, 2013, Section 5.1.8) for faster model evaluation during cross-validation. Our algorithm is a block-level variant of Kim

---

**Algorithm S1** Computationally efficient algorithm for projection (see, Theorem A2)

---

```

1: Input:  $X, \tilde{X}, \tilde{v}_\eta, \tilde{v}_\xi, \tilde{v}_\beta, \tilde{v}_z$ 
2: Output: Posterior samples of  $\gamma = (\xi^\top, \beta^\top, z^\top)^\top$ 
3: Storage: Pre-allocate  $v_1, v_2, v_3, v_4, \{D_i : i = 1, \dots, n\}, D_*$ 
4: function PROJECTION( $X, \tilde{X}, \tilde{v}_\eta, \tilde{v}_\xi, \tilde{v}_\beta, \tilde{v}_z$ )
5:    $v_1 \leftarrow \tilde{X}^\top \tilde{v}_\eta + \tilde{v}_z$ 
6:   for  $i = 1$  to  $i = n$  do
7:      $D_i \leftarrow \text{chol}(\tilde{x}_i \tilde{x}_i^\top + I_r)$ 
8:      $v_{1,i} \leftarrow \text{trsolve}(D_i, v_{1,i})$ 
9:    $v_2 \leftarrow \tilde{v}_\eta - \tilde{X} v_1$ 
10:   $v_3 \leftarrow X^\top v_2 + \tilde{v}_\beta$ 
11:   $v_2 \leftarrow v_2 + \tilde{v}_\xi$ 
12:  for  $i = 1$  to  $i = n$  do  $v_{4,i} \leftarrow v_{2,i} / (2 + \|\tilde{x}_i\|^2)$ 
13:   $v_3 \leftarrow X^\top v_4 - v_3$ 
14:   $D_* \leftarrow \text{chol}\left(I_p + X^\top \text{diag}\left(\frac{1}{2 + \|\tilde{x}_i\|^2}\right)_{i=1}^n X\right)$ 
15:   $v_3 \leftarrow \text{trsolve}(D_*, v_3)$ 
16:   $v_4 \leftarrow X v_3$ 
17:   $v_2 \leftarrow \text{diag}\left(\frac{1 + \|\tilde{x}_i\|^2}{2 + \|\tilde{x}_i\|^2}\right)_{i=1}^n v_2 + \text{diag}\left(\frac{1}{2 + \|\tilde{x}_i\|^2}\right)_{i=1}^n v_4$ 
18:   $v_4 \leftarrow v_2 - v_4$ 
19:  for  $i = 1$  to  $i = n$  do
20:     $v_{1,i} \leftarrow \text{trsolve}(D_i, v_{4,i} \tilde{x}_i)$ 
21:  return  $(v_2^\top, -v_3^\top, -v_1^\top)^\top$ 

```

---

et al. (2002). Following the setup preceding Algorithm S3, suppose the  $k$ th block  $\chi_{[k]}$  for each  $k$  contains adjacent indices of the data  $\chi$ . Moreover, suppose evaluating a model needs Cholesky decomposition of a  $n \times n$  positive-definite matrix  $R$ , where  $n$  is the number of observations in  $\chi$ . In the context of (7),  $R$  corresponds to the correlation matrix of a spatial-temporal process. Suppose we have evaluated the model on the full data, and we have stored  $L$ , the upper triangular Cholesky factor of  $R$ . The consecutive indices of the partitioned data creates  $K^2$  submatrices partitioning  $R$  and its upper triangular Cholesky factor  $L$  as,

$$R = \begin{bmatrix} R_{11} & R_{21}^\top & \cdots & R_{K1}^\top \\ R_{21} & R_{22} & \cdots & R_{K2}^\top \\ \vdots & \vdots & \ddots & \vdots \\ R_{K1} & R_{K2} & \cdots & R_{KK} \end{bmatrix}, \quad L = \begin{bmatrix} L_{11} & L_{12} & \cdots & L_{1K} \\ 0 & L_{22} & \cdots & L_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & L_{KK} \end{bmatrix}.$$

For each  $k$  such that  $1 < k < K$ , consider the following representation of the matrices  $R, L, R_{-k}$ , the matrix corresponding to  $\chi_{[-k]}$  and its upper triangular Cholesky factor  $L_{-k}$ .

$$R = \begin{bmatrix} R_{11}^k & R_{21}^{k\top} & R_{31}^{k\top} \\ R_{21}^k & R_{22}^k & R_{32}^{k\top} \\ R_{31}^k & R_{32}^k & R_{33}^k \end{bmatrix}, \quad L = \begin{bmatrix} L_{11}^k & L_{12}^k & L_{13}^k \\ 0 & L_{22}^k & L_{23}^k \\ 0 & 0 & L_{33}^k \end{bmatrix}, \quad R_{-k} = \begin{bmatrix} R_{11}^k & R_{31}^{k\top} \\ R_{31}^k & R_{33}^k \end{bmatrix}, \quad L_{-k} = \begin{bmatrix} C_{11}^k & C_{13}^k \\ 0 & C_{33}^k \end{bmatrix},$$

where  $R_{22}^k = R_{kk}$ . Then, Algorithm S2 finds  $L_{-k}$  for each  $k = 1, \dots, K$ . In Algorithm S2, note that, no computation is needed to find  $L_{-K}$  and  $C_{11}^k, C_{13}^k$  for  $1 < k < K$ .

---

**Algorithm S2** Fast Cholesky updates of row-deletion using block Givens rotation.

---

```

1: Input:  $R, L, k$ ; full matrix, its upper-triangular Cholesky factor, block to be deleted
2: Output:  $L_{-k}$ , the Cholesky factor of  $R_{-k}$ , the matrix with  $k$ th block deleted
3: function CHOLESKYCV( $R, L, k$ )
4:   if  $k = 1$  then
5:      $L_{-k} \leftarrow \text{chol}(R_{-k})$ 
6:   else
7:     if  $1 < k < K$  then
8:        $C_{11}^k \leftarrow L_{11}^k$  (pre-computed)
9:        $C_{13}^k \leftarrow L_{13}^k$  (pre-computed)
10:       $C_{33}^k \leftarrow \text{chol}(L_{33}^{k\top} L_{33}^k + L_{23}^{k\top} L_{23}^k)$ 
11:    else
12:       $L_{-k} =$  upper left  $(K - 1) \times (K - 1)$  submatrices of  $L$  (pre-computed)
13:    return  $L_{-k}$ 

```

---

It is worth remarking that the “naive approach” for obtaining  $L_{-k}$  involves Cholesky decomposition of  $R_k$  for each  $k$ , and, hence, requiring  $K^{-2}(K - 1)^3 n^3 / 3$  floating point operations (flops). On the other hand, the time complexity of Algorithm S2 is in the order of

$$\sum_{k=1}^K \left(1 - \frac{k}{K}\right)^3 \frac{n^3}{3} = \frac{n^3}{3K^3} \sum_{k=1}^{K-1} i^3 = \frac{n^3}{3K^3} \frac{K^2(K-1)^2}{4} = \frac{(K-1)^2}{4K} \frac{n^3}{3},$$

where the penultimate step follows from the sum of cubes of natural numbers. Hence, we show that Algorithm S2 is theoretically  $4(K - 1)/K$  times faster than the naive approach. If  $K = 10$  (Vehtari and Lampinen, 2002), then Algorithm S2 theoretically offers approximately 72% efficiency in time complexity over the naive approach. However, it must be noted that modern linear algebra libraries are highly vectorised and are often multithreaded. Hence, it is difficult to accurately translate these purely theoretical flop counts to actual wall clock time. Nevertheless, Algorithm S2 is more efficient than the naive approach.

Lastly, we detail the algorithm for implementing our proposed stacking framework utilising the aforementioned algorithms. Algorithm S3 outlines the steps required for estimating the spatially-temporally varying coefficients model (7) in the main article. The inputs  $X$  and  $\tilde{X}$  in Algorithm S3 correspond to the design matrices appearing in (7),  $\mathcal{L}$  denotes spatial-temporal locations,  $N$  denotes number of posterior samples to be drawn,  $B$  denotes the number of samples to be drawn for evaluating the leave-one-out predictive densities as described in (15),  $\mathcal{M} = \{M_1, \dots, M_L\}$  denotes the collection of  $L$  candidate models and  $K$  denotes the number of folds for cross-validation required for fast evaluation of leave-one-out predictive densities. Following a random permutation of the entire dataset, we construct a partition using consecutive indices of the permuted data. This step significantly accelerates evaluation of Cholesky factors using block Givens rotation (Golub

and Van Loan, 2013) required during the cross-validation step detailed in Algorithm S2. After partitioning the data, which is denoted by  $\chi = (y, X, \tilde{X}, \mathcal{L})$  into  $K$  blocks, each block is denoted by  $\chi_{[k]}$  and the data with the  $k$ th block deleted is denoted by  $\chi_{[-k]}$  for  $k = 1, \dots, K$ . Algorithm S3 is written in a richer and more general context and can be easily adapted to spatial, spatial-temporal, spatially and spatially-temporally varying coefficients models.

---

**Algorithm S3** Predictive stacking algorithm for spatially-temporally varying coefficients GLM

---

- 1: **Input:**  $y, X, \tilde{X}, \mathcal{L}, N, B, \mathcal{M}, K$
  - 2: **Output:** Posterior samples for each model in  $\mathcal{M}$ , optimal stacking weights  $\hat{w}$
  - 3: **function** STVCGLM\_STACKING( $y, X, \tilde{X}, \mathcal{L}, N, B, \mathcal{M}, K$ )
  - 4:   Partition the data  $\chi$  into  $K$  blocks  $\chi_{[1]}, \dots, \chi_{[K]}$ , each containing consecutive indices
  - 5:   **for**  $l = 1$  to  $l = L$  **do**
  - 6:     Fit model  $M_l$  on  $\chi$  and obtain posterior samples  $\{\beta_l^m, z_l^m\}_{m=1}^N$  using (11)
  - 7:     **for**  $k = 1$  to  $k = K$  **do**
  - 8:       Fit model  $M_l$  on  $\chi_{[-k]}$  and obtain posterior samples  $\{\beta_{k,l}^s, z_{k,l}^s\}_{s=1}^S$  using (11)
  - 9:       **for**  $s = 1$  to  $s = S$  **do**
  - 10:          Use (13) to predict the spatial-temporal process  $\tilde{z}_{k,l}^s$  at  $\mathcal{L}_{[k]}$  from  $z_{k,l}^s$
  - 11:          Use (15) to find  $p(y(\ell_i) \mid y_{-i}, M_l), \ell_i \in \mathcal{L}_{[k]}$
  - 12:     Optimise (14) to obtain optimal stacking weights  $\hat{w}$
  - 13:   **return** Optimal model weights  $\hat{w}$ , Posterior samples  $\{\beta_l^m, z_l^m\}_{m=1}^N$  for  $l = 1, \dots, L$
- 

## APPENDIX C. ADDITIONAL DETAILS ON SIMULATION EXPERIMENTS

We present some additional results from the simulation experiments carried out in Section 5 of the main article. Consider the simulated Poisson count data under the spatially-temporally varying coefficient model (7) as described in Section 5.1. Figure A1 presents the posterior distributions of the process parameters that characterise the two spatial-temporal processes  $z_1$  and  $z_2$ , corresponding to the intercept ( $\phi_{11}$  and  $\phi_{21}$ ) and the predictor ( $\phi_{21}$  and  $\phi_{22}$ ) respectively, obtained by MCMC (see Section 5.2). None of the process parameters in Fig. A1 have concentrated around their corresponding true values, illustrating their weak identifiability. We also notice similarity in the posterior distributions of the temporal decay parameters  $\phi_{11}$  and  $\phi_{12}$  for both the processes  $z_1$  and  $z_2$  respectively. The same phenomenon is noticed in the posterior distributions of the spatial decay parameters  $\phi_{21}$  and  $\phi_{22}$ . This between-process similarity in the posterior learning of the process parameters justifies the tenability of the “reduced” model that uses common spatial-temporal process parameters across all the  $r$  processes in (7). This reduces the parameter space by  $r$  folds. The reduced model further effectuates a substantial decrease in the number of candidate models required for the proposed stacking algorithm.

Moreover, in addition to Fig. 1 in the main article, where we plot the posterior medians of the spatial-temporal random effects for the simulated Poisson count data, Figure A2 plots 100 quantiles of the combined posterior samples of spatial-temporal random effects obtained by stacking against quantiles obtained from MCMC. This demonstrates strong agreement between the posterior






FIGURE A1. Histograms of posterior samples of the spatial-temporal process parameters obtained by MCMC on a fully Bayesian model on a simulated Poisson count data of sample 100.

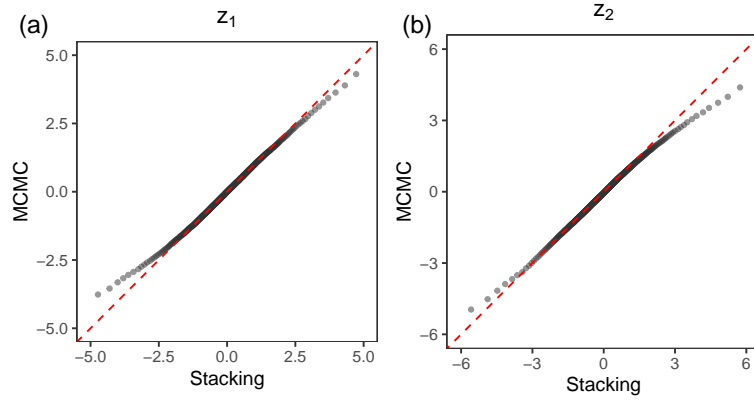


FIGURE A2. Plot of quantiles of the posterior distributions of the spatial-temporal random effects corresponding to (a) intercept and, (b) predictor obtained by stacking against MCMC with the  $y = x$  reference as a red dashed line.

distributions obtained from the competing algorithms with most of the points concentrated along the  $y = x$  line (red dashed).




FIGURE A3. Comparison of interpolated surfaces of (a) the true spatial effects, with posterior median of spatial effects obtained by (b) MCMC, and, (c) our proposed predictive stacking algorithm on a simulated spatial count data.

We also present results from an additional simulation experiment that demonstrates posterior learning of a random field modelling the underlying spatial or spatial-temporal processes. For convenient visualisation of the random field, we consider spatial data instead of spatial-temporal data in a continuous time domain. We simulate a dataset with responses distributed as  $y(s) \sim \text{Poisson}(\exp(x(s)^\top \beta + z(s)))$  with sample size 1000. The locations are sampled uniformly inside the unit square  $[0, 1]^2$ . The explanatory vector  $x(s)$  consists of an intercept and one predictor sampled from the standard normal distribution and the regression coefficients are taken as  $\beta = (5, -0.5)$ . The latent spatial process  $z(s) \sim \text{GP}(0, \sigma_z^2 R(\cdot, \cdot; \phi, \nu))$  is a zero-centred Gaussian process with  $\sigma_z^2 = 0.4$ ,  $\phi = 3.5$  and  $\nu = 0.5$ , where  $R(\cdot, \cdot; \phi, \nu)$  is the Matérn correlation function,

$$R(s, s'; \phi, \nu) = \frac{(\phi |s - s'|)^\nu}{2^{\nu-1} \Gamma(\nu)} K_\nu(\phi |s - s'|), \quad (\text{A11})$$

$|s - s'|$  is the Euclidean distance between  $s$  and  $s'$ , and  $\theta_{\text{sp}} = \{\phi, \nu\}$ . The parameter  $\nu > 0$  controls the smoothness of the realised random field,  $\phi$  is the spatial decay parameter,  $\Gamma(\cdot)$  denotes the gamma function, and  $K_\nu$  is the modified Bessel function of the second kind of order  $\nu$  (Abramowitz and Stegun, 1965, Chapter 10). The model in (7) can be modified for accommodating a spatial regression by considering  $r = 1$  and subsequently  $\tilde{X} = I_n$ . Here,  $\mathcal{S}$  is the domain of interest. We stack on the parameters  $\Delta = \{\alpha_\epsilon, \sigma_\xi, \phi, \nu\}$  with hyperparameters  $\nu_\beta = \nu_z = 3$ . We also implement a fully Bayesian model with uniform priors  $\text{U}(0.5, 10)$  for  $\phi$  and  $\text{U}(0.1, 2)$  for  $\nu$ . We modify our adaptive Metropolis-within-Gibbs algorithm accordingly. Figure A3 compares the posterior distributions of the spatial random effects obtained by predictive stacking and MCMC with its true values. We observe indistinguishable spatial surfaces of the posterior medians of the spatial random effects.

#### APPENDIX D. ADDITIONAL DATA ANALYSIS

We offer some additional data analysis from the North American Breeding Bird Survey to supplement the results in Section 6 of the main article. In the survey, the data are collected annually during the breeding season, primarily in June, along thousands of randomly established roadside survey routes in the United States and Canada. Routes are roughly 24.5 miles (39.2 km) long

with counting locations placed at approximately half-mile (800 m) intervals, for a total of 50 stops. At each stop, a citizen scientist, highly skilled in avian identification, conducts a 3-minute point count recording all birds seen within a quarter-mile (400 m) radius and all birds heard. Routes are sampled once per year. In addition to avian count data, this dataset also contains route location information including country, state and the geographic coordinates of the route start point. The variable ‘car’ records the total number of motorised vehicles passing a particular point count stop during the 3-minute count period. The variable ‘noise’ represents the presence/absence of excessive noise defined as noise from sources other than vehicles passing the survey point (e.g. from streams, construction work, vehicles on nearby roads, etc.) lasting 45 seconds that significantly interferes with the observer’s ability to hear birds at the stop during the 3-minute count period. More information on the survey can be found online at BBS 2022 data release (<https://www.usgs.gov/data/2022-release-north-american-breeding-bird-survey-dataset-1966-2021>).

The main article discusses the global regression coefficients. Figures A4 and A5 display side-by-side plots of the observed point counts and interpolated spatial surfaces of the posterior median of the processes assigned to the intercept and the predictors ‘car’ and ‘noise’, revealing clear spatial-temporal patterns in the effect of the predictors.

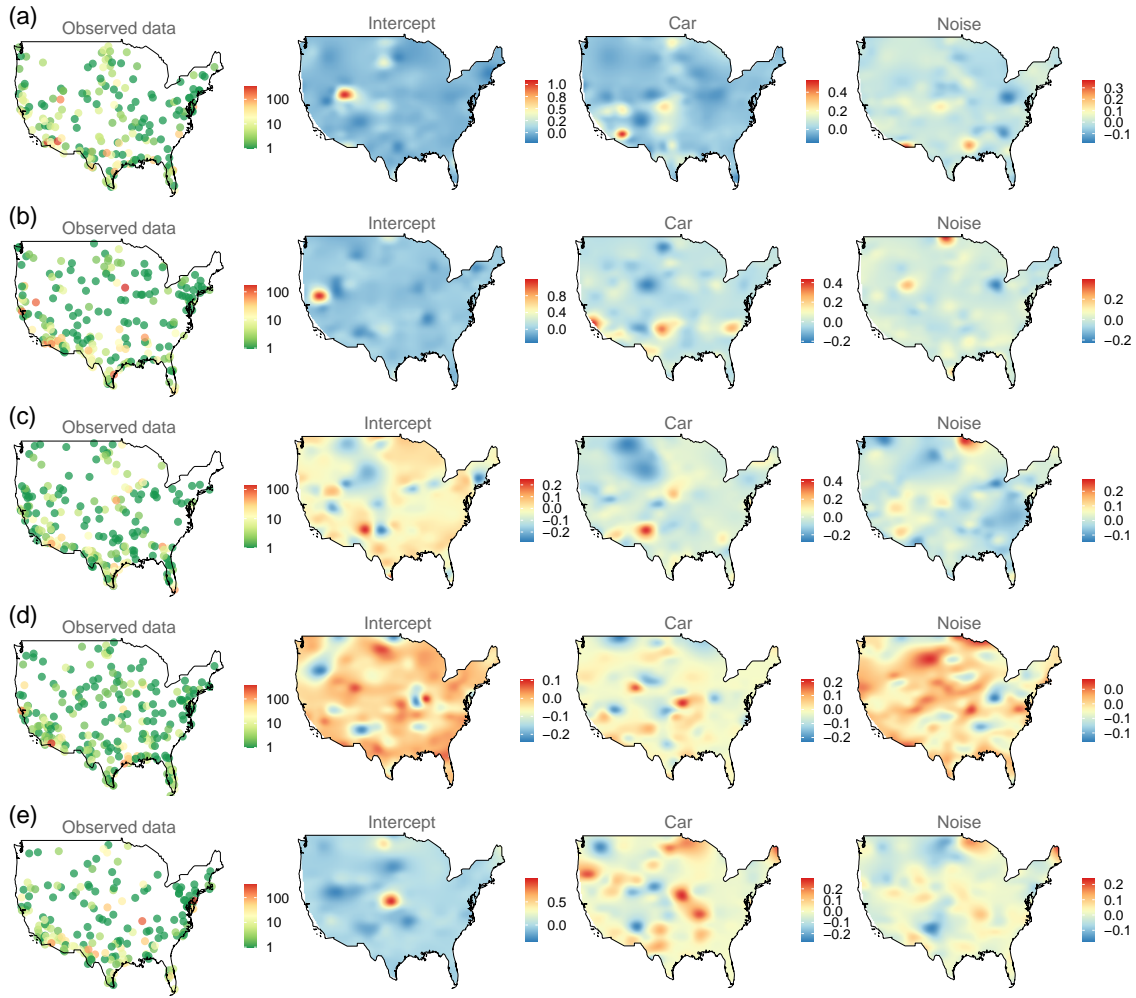


FIGURE A4. (a)–(e) Observed avian point count data and interpolated surfaces of posterior median of spatial-temporal random effects in the intercept as well as slopes of the variables ‘car’ and ‘noise’ for years 2010–2014.

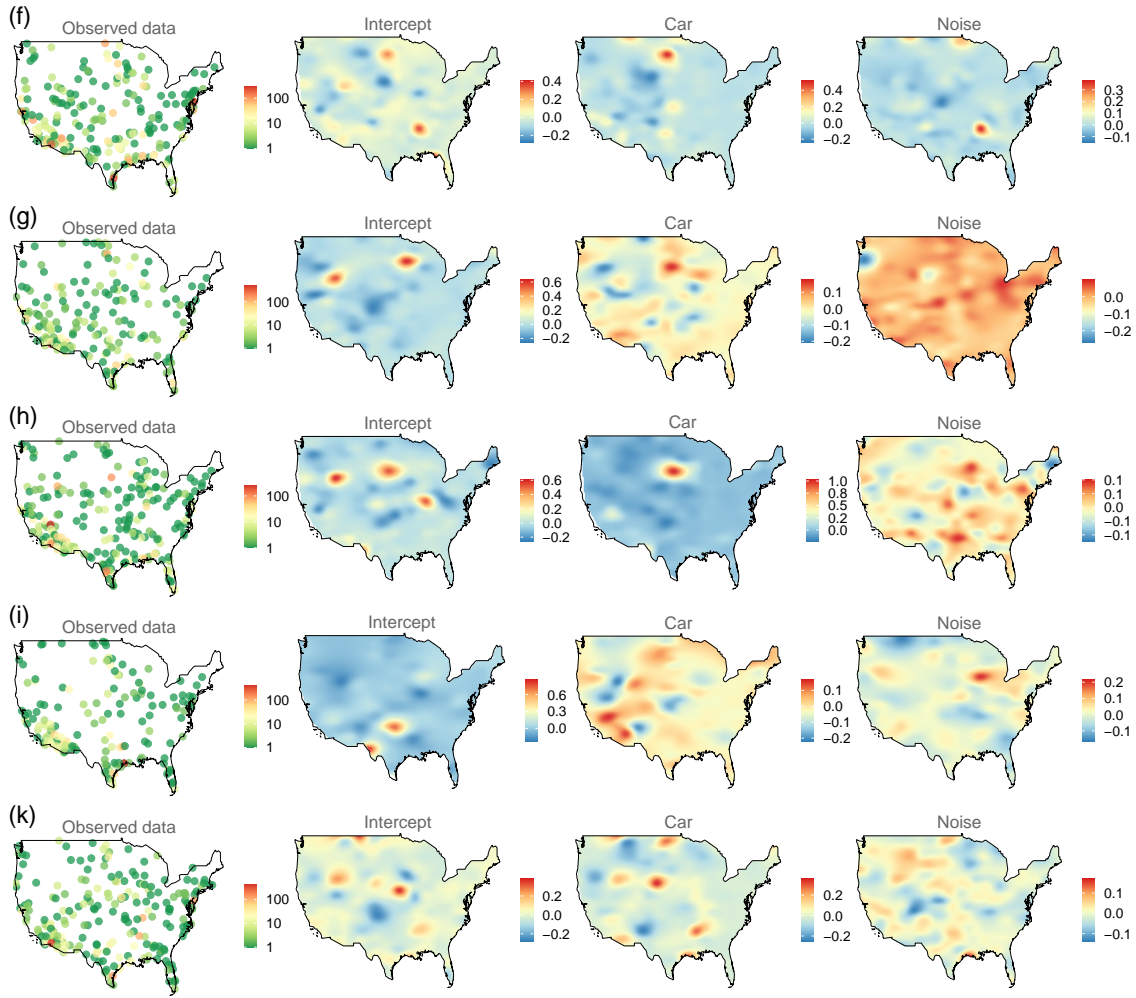


FIGURE A5. (f)–(k) North American Breeding Bird Survey (2015-19): Observed avian point count data and interpolated surfaces of posterior median of spatial-temporal random effects in the intercept as well as slopes of the variables ‘car’ and ‘noise’ for years 2015-2019.