

A large-scale spatio-temporal binomial regression model for estimating seroprevalence trends

Stella C. Watson Self¹ | Christopher S. McMahan¹  | Derek A. Brown¹ | Robert B. Lund¹ | Jenna R. Gettings² | Michael J. Yabsley^{2,3}

¹Department of Mathematical Sciences, Clemson University, Clemson, South Carolina

²Southeastern Cooperative Wildlife Disease Study, Department of Population Health, University of Georgia, Athens, Georgia

³Warnell School of Forestry and Natural Resources, University of Georgia, Athens, Georgia

Correspondence

Christopher S. McMahan, Department of Mathematical Sciences, Clemson University, Clemson, SC 29634-0975. Email: mcmaha2@clemson.edu

Funding information

National Science Foundation, Grant/Award Number: DMS-1127914, DMS-1407480, CMMI-1563435 and EEC-1744497; National Institutes of Health, Grant/Award Number: R01 AI121351; Companion Animal Parasite Council (CAPC); Boehringer Ingelheim Vetmedica-CAPC Infectious Disease Postdoctoral Fellowship

Abstract

This paper develops a large-scale Bayesian spatio-temporal binomial regression model to investigate regional trends in antibody prevalence to *Borrelia burgdorferi*, the causative agent of Lyme disease. Our model uses Gaussian predictive processes to estimate the spatially varying trends and a conditional autoregressive scheme to account for spatio-temporal dependence. A novel framework, easily scalable to large spatio-temporal data, is developed. The proposed model is used to analyze about 16 million *B. burgdorferi* antibody Lyme tests performed on canine samples in the conterminous United States over the sixty-month period from January 2012 to December 2016. This analysis identifies areas of increasing canine Lyme disease risk; prevalence of infection is getting worse in endemic regions and increases are also seen in non-endemic regions. Because Lyme disease is zoonotic, affecting both humans and dogs, the analysis also serves to pinpoint areas of increasing human risk.

KEYWORDS

Borrelia burgdorferi, CAR model, chromatic sampling, Gaussian predictive processes, Lyme disease

1 | INTRODUCTION

Lyme disease is a vector-borne disease that impacts both humans and several other mammalian species, with domestic dogs often contracting the infection (Little, Heise, Blagburn, Callister, & Mead, 2010). The disease results via infection by *Borrelia burgdorferi*, a spirochetal bacteria that is transmitted by ticks. In the United States, *B. burgdorferi* is the only known etiologic agent of Lyme disease in dogs; other species can cause Lyme disease in Europe and Asia. Lyme disease incidence in humans is considered emerging, with a growing number of high-incidence counties being reported (Adams et al., 2017). Humans and dogs are infected by the same vectors (Little et al., 2010); hence, the risks of exposure for both are related. In fact, dogs are viewed as sentinels for regional Lyme disease risk in humans (Mead, Goel, & Kugeler, 2011). The effects of Lyme disease are usually more severe in humans, often debilitating.

Dogs are regularly tested for exposure to *B. burgdorferi* as part of annual wellness examinations. Commonly, veterinarians use a serologic test that detects antibodies against the C6 peptide that is present in the blood of exposed animals. The presence of C6 is indicative of an intermediate or late-term infection and is often detectable 3 to 6 weeks

after exposure (Wagner et al., 2012). Among dogs that are infected, only about 5% develop any clinical signs of Lyme disease (Levy & Magnarelli, 1992). The remaining dogs may either clear the infection without developing the disease or be subclinically infected and never show symptoms. The routine Lyme disease testing done on dogs provides an opportunity to measure the proportion of exposed dogs in the relatively healthy canine population that visits veterinary clinics and is tested. This proportion is a prevalence based on serologic (blood) tests and is therefore called a seroprevalence (henceforth shortened to prevalence). In this population, the national prevalence is about 6%, but county Lyme disease prevalence can exceed 40% in some cases.

Monitoring prevalence is useful for many reasons. Prevalence indicate exposure risk within a region, allowing veterinarians to provide effective preventative care and make testing recommendations. Indirectly, prevalence can help identify the range of *Ixodes* spp. tick vectors. This is important because *Ixodes* spp. also transmit other pathogens, including *Anaplasma* spp., *Ehrlichia muris euclairensis*, and *Babesia microti* (Nelder et al., 2016), several of which are also zoonotic. The shared tick vector and extensive testing data for dogs makes them a good sentinel for quantifying human Lyme disease risk. In short, trends in dog prevalence should aid our understanding of Lyme disease risk changes for humans.

The goal of this paper is to identify U.S. regions that are experiencing increasing prevalence in dogs. Our data are 16,571,562 serologic *B. burgdorferi* tests conducted on domestic dogs in the conterminous United States (US) from January 2012 to December 2016, aggregated by county and month. Figure 1 displays the raw prevalence estimates (the proportion of positive tests) after aggregating over all sixty months in this study. Data were reported from 69,876 county-month pairs. To locate where prevalence is increasing, our model needs to have a spatially-varying temporal trend component. To make reliable inferences, the strong positive spatio-temporal dependence of the tests needs to be taken into account. The size of this data set and its large spatio-temporal support (about 3,100 distinct counties or county-equivalent regions) motivates some of our methodological choices.

Gaussian processes (GPs) are popular geostatistical modeling tools due to their flexibility and ability to quantify uncertainty in nonparametric regressions (Neal, 1998; O'Hagan, 1978). Good GP modeling overviews are provided in Cressie (1993), Rasmussen and Williams (2006), Cressie and Wikle (2011), and Gelfand and Schliep (2016). Banerjee, Carlin, and Gelfand (2015) discussed Bayesian aspects of GPs. The objective prior specification for GP models is studied in the work of Berger, de Oliveira, and Sansó (2001). GPs have become standard tools in a wide variety of applications, including oceanography (Jona-Lasinio, Gelfand, & Jona-Lasinio, 2012), water quality analysis (Zhang & El-Shaarawi, 2009), image classification (Morales-Álvarez, Pérez-Suay, Molina, & Camps-Valls, 2017), neuroimaging (Lazar, 2008), and computer experiments (Santner, Williams, & Notz, 2003). GPs have also been previously used to model disease prevalence, including dengue fever (Johnson et al., 2017), malaria (Andrade-Pacheco, Mubangizi, Quinn, & Lawrence, 2015), and influenza (Senanayake, O'Callaghan, & Ramos, 2016). Gelfand, Kim, Sirmans, and Banerjee (2003) allowed GP linear model coefficients to vary smoothly over space, an approach used here to allow for regional prevalence trends.

GP modifications and algorithms for analyzing big spatial data have received significant recent attention, including fixed-rank kriging (Cressie & Johannesson, 2008) and LatticeKrig (Nychka, Bandyopadhyay, Hammerling, Lindgren, & Sain, 2015). Both methods employ basis function expansions of spatial random effects to reduce the dimensions of the

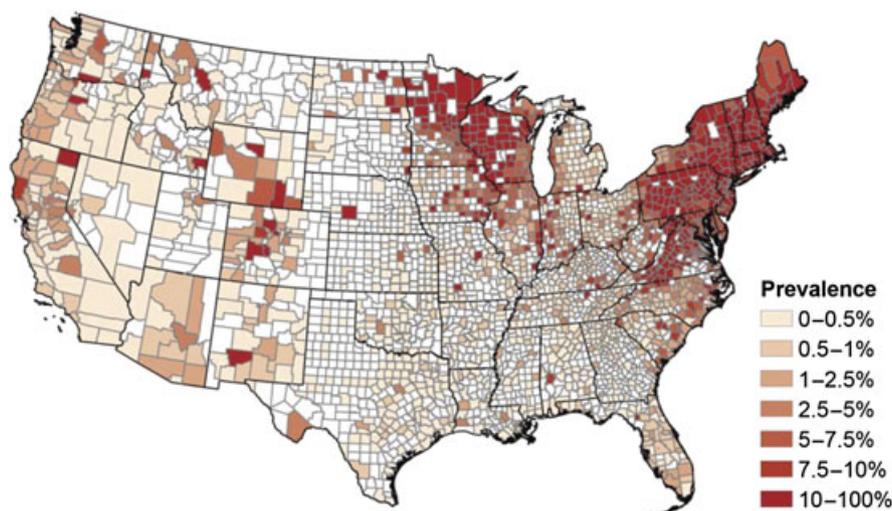


FIGURE 1 Observed seroprevalence of *Borrelia burgdorferi*, aggregated over January 2012 to December 2016. White counties are those that did not report any test results

covariance matrices in the model. Katzfuss (2017) takes a similar approach, applying basis functions to a succession of refined resolutions. Spatial partitioning (e.g., Heaton, Christensen, & Terres, 2017; Sang, Jun, & Huang, 2011) can be used to split regions into smaller, more manageable subregions with computation being accelerated via a conditional independence assumption. Covariance tapering (Furrer, Genton, & Nychka, 2006) uses a covariance structure with compact support to induce sparsity. Nearest-neighbor processes (Datta, Banerjee, Finley, & Gelfand, 2016) achieve computational efficiency by conditioning on a subset of nearby observations. A similar idea is used by Gramacy and Apley (2015) to find the largest number of neighbors that are computationally feasible for prediction, optimally chosen by minimizing a prediction variance. Heaton et al. (2017) provided an overview and comparison of these and other procedures. Our approach involves Gaussian predictive processes (GPPs; Banerjee, Gelfand, Finley, & Sang, 2008) and is discussed further in Section 2.

The most common approach for modeling spatially dependent areal data involves Gaussian Markov random fields (GMRFs; Rue & Held, 2005), with Gaussian conditional autoregressive (CAR) models (Banerjee et al., 2015) being particularly popular. As special cases of Markov random fields (Besag, 1974), GMRFs are collections of jointly distributed Gaussian variables satisfying a Markov dependence structure quantified through a precision matrix. GMRFs are extended to flexible degrees of smoothness in the works of Brezger, Fahrmeir, and Hennerfeind (2007) and Yue and Speckman (2010). Brown, Datta, and Lazar (2017) adjusted the CAR precision matrix to build a unified model for independent and dependent cases and study neighborhood structures other than those based on physical adjacency. GMRF and GP connections are explored in the works of Rue and Tjelmund (2002), Song, Fuentes, and Ghosh (2008), and Lindgren, Rue, and Lindström (2011). CAR models are by now standard in disease mapping problems (e.g., Waller, Carlin, Xia, & Gelfand, 1997).

To achieve our goals, a large-scale spatio-temporal binomial regression model is developed that has both GPP and CAR components. The former is used to capture regionally varying trends by treating the trend coefficient as a non-parametric surface over space, whereas the latter accounts for local heterogeneity. Through data augmentation steps and a novel sampling strategy, a modeling framework is developed that is computationally scalable to large non-Gaussian spatio-temporal data sets. In particular, straightforward Gibbs sampling is facilitated via a data augmentation step involving latent Pólya-gamma variables. To avoid computationally expensive matrix calculations, a chromatic sampling strategy is used in our Gibbs sampler. Our methodology easily handles missing data. The finite sample properties of our approach are studied via simulation before our Lyme disease analysis is conducted.

The remainder of this paper is organized as follows. Section 2 describes the model and our GPP and CAR structures. Section 3 discusses model fitting procedures, emphasizing computational tractability with large spatio-temporal data. Section 4 presents a simulation study supporting our approach and Section 5 analyzes the canine serology data described above. Concluding remarks are offered in Section 6.

2 | MODELING METHODS

Let Y_{st} denote the number of cases (e.g., positive *B. burgdorferi* tests) observed in n_{st} tests taken in region s at time t , for $s = 1, \dots, S$ and $t = 1, \dots, T$. Set $\mathbf{Y}_s = (Y_{s1}, \dots, Y_{sT})'$, $\mathbf{Y} = (\mathbf{Y}'_1, \dots, \mathbf{Y}'_S)' \in \mathbb{R}^{ST}$, $\mathbf{n}_s = (n_{s1}, \dots, n_{sT})'$, and $\mathbf{n} = (\mathbf{n}'_1, \dots, \mathbf{n}'_S)' \in \mathbb{N}^{ST}$. In addition to the disease tests, the covariates Z_{stq} and X_{stp} , for $q = 1, \dots, Q$ and $p = 1, \dots, P$, are assumed available in region s at time t . The Z_{stq} are covariates whose effects are constant over the study area, whereas X_{stp} are covariates whose associated effects vary by region.

To relate the observed test data to the covariates, a Bayesian generalized linear mixed model (Banerjee et al., 2015; Diggle, Tawn, & Moyeed, 1998; McCullagh & Nelder, 1989) is adopted. Our general model is the binomial regression: $Y_{st} | n_{st}, p_{st} \sim \text{Binomial}(n_{st}, p_{st})$ with

$$v_{st} := g^{-1}(p_{st}) = \mathbf{Z}'_{st} \boldsymbol{\delta} + \mathbf{X}'_{st} \boldsymbol{\beta}(\boldsymbol{\ell}_s) + \xi_{st}; \quad s = 1, \dots, S; \quad t = 1, \dots, T, \quad (1)$$

where $g : \mathbb{R} \rightarrow (0, 1)$ is a known link function (e.g., logistic) relating the linear predictor v_{st} to the prevalence p_{st} , $\mathbf{Z}_{st} = (1, Z_{st1}, \dots, Z_{stQ})' \in \mathbb{R}^{Q+1}$, $\mathbf{X}_{st} = (X_{st1}, \dots, X_{stP})' \in \mathbb{R}^P$, $\boldsymbol{\delta} = (\delta_0, \dots, \delta_Q)'$ are global regression coefficients, $\boldsymbol{\beta}(\cdot) = (\beta_1(\cdot), \dots, \beta_P(\cdot))'$ are spatially varying regression coefficients, $\boldsymbol{\ell}_s = (\ell_{s1}, \ell_{s2})'$ is a vector of spatial coordinates (e.g., latitude and longitude) that identifies the centroid of region s , and ξ_{st} is a spatio-temporal random effect. Following the work of Gelfand et al. (2003), the spatially varying regression coefficients are regarded as unknown smooth surfaces over the study region. To model these unknown surfaces while maintaining computational tractability, GPPs are used.

A GP is a stochastic process whose finite dimensional distributions are multivariate normal. A GP $\beta_p(\cdot)$, given a covariance parameter θ_p and denoted

$$\beta_p \mid \theta_p \sim \mathcal{GP}(\mu_p(\cdot), C(\cdot, \cdot; \theta_p)),$$

is uniquely determined by its mean and covariance, $\mu_p(\mathcal{L}_s) := E[\beta_p(\mathcal{L}_s)]$ and $C(\mathcal{L}_s, \mathcal{L}_{s'}; \theta_p) := \text{Cov}(\beta_p(\mathcal{L}_s), \beta_p(\mathcal{L}_{s'})) = \sigma_p^2 \rho_p(\mathcal{L}_s, \mathcal{L}_{s'}; \theta_p)$, where $\rho_p(\cdot, \cdot; \theta_p)$ is a correlation function depending on θ_p . For smoothing and interpolation, a constant mean is often assumed (Bayarri et al., 2007). Our work *a priori* posits that $\mu_p(\cdot) \equiv 0$ for all p . Thus, $\beta_p = (\beta_p(\mathcal{L}_1), \dots, \beta_p(\mathcal{L}_S))'$, $S \in \mathbb{N}$, follows a multivariate normal distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{C}_p = \sigma_p^2 \mathbf{R}_p$, where $(\mathbf{R}_p)_{ss'} = \rho_p(\mathcal{L}_s, \mathcal{L}_{s'}; \theta_p)$. In general, the covariance matrix inversions and factorizations needed to calculate quantities in our posterior distributions are $\mathcal{O}(S^3)$ in computational time. In Markov chain Monte Carlo (MCMC) algorithms, these operations will need to be repeated thousands of times. Thus, as S grows large, GPs quickly become computationally unwieldy.

To reduce the dimension of the problem, our GPP employs a “parent” process based on a strategically chosen set of knots and interpolates to points of interest via kriging. Let $\{\mathcal{L}_1^*, \dots, \mathcal{L}_{S_p^*}^*\}$ denote the knot set with $S_p^* \ll S$. Define $\beta_p^* = (\beta_p(\mathcal{L}_1^*), \dots, \beta_p(\mathcal{L}_{S_p^*}^*))'$ and note that $\beta_p^* \mid \sigma_p^2, \theta_p \stackrel{\text{ind}}{\sim} N(\mathbf{0}, \mathbf{C}_p^*)$, for all p , where $\mathbf{C}_p^* = \sigma_p^2 \mathbf{R}_p^*$ and $(\mathbf{R}_p^*)_{ss'} = \rho_p(\mathcal{L}_s^*, \mathcal{L}_{s'}^*; \theta_p)$. The GPP simply replaces β_p with $\tilde{\beta}_p := E(\beta_p \mid \beta_p^*, \theta_p) = \tilde{\mathbf{R}}_p^* (\mathbf{R}_p^*)^{-1} \beta_p^*$, where $\tilde{\mathbf{R}}_p^*$ is an $S \times S_p^*$ matrix whose (s, s') th element is $\rho_p(\mathcal{L}_s, \mathcal{L}_{s'}^*; \theta_p)$. When S_p^* is not large, $(\mathbf{R}_p^*)^{-1}$ can be quickly computed. For more on GPPs, see the work of Banerjee et al. (2008).

Fully specifying a GPP requires specifying its knot locations. Banerjee et al. (2008) discussed several methods of knot selection, including placing them on a regular grid, selecting them at random from the observation locations, and methods that place more knots in areas with more observations. Finley, Sang, Banerjee, and Gelfand (2009) suggested choosing knot locations to minimize conditional variances at observation locations. Guhaniyogi, Finley, Banerjee, and Gelfand (2011) proposed an adaptive knot selection strategy where knot locations are treated as a point process. Following the work of Eidsvik, Finley, Banerjee, and Rue (2012), our knots are chosen via K -means clustering with S_p^* clusters; that is, using K -means clustering, the S counties are partitioned into S_p^* clusters based on their locations \mathcal{L}_s . The knot locations are taken as the centroids of the S_p^* clusters. For further details on K -means clustering, see the work of Hartigan and Wong (1979).

A variety of ways exist to model spatio-temporal dependence of areal data. Some commonly used methods include intrinsic CAR models, proper CAR models, and the so-called Besag-York-Mollie models for disease mapping problems. These methods can be extended to handle spatio-temporal correlation in many ways. For examples, a separate Besag-York-Mollie model can be fit at each time point (Knorr-Held & Besag, 1998; Waller et al., 1997), or a CAR model can be combined with a random walk in time (Knorr-Held & Besag, 1998), a spline-based temporal structure (MacNab & Dean, 2001) or a local autoregressive model in time (Congdon & Southall, 2005). For a thorough review and comparison of existing spatio-temporal models, see the work of Anderson and Ryan (2017). A county-by-county exploratory analysis of our Lyme disease data suggests that a first-order autoregressive (AR(1)) model is sufficient for handling temporal dependence. Thus, following the work of Rushworth, Lee, and Mitchell (2014) and Lee and Lawson (2014), first-order vector autoregression is used with GMRF errors:

$$\xi_t = \zeta \xi_{t-1} + \phi_t, \quad (2)$$

where $\xi_t = (\xi_{1t}, \dots, \xi_{St})'$, $\zeta \in (-1, 1)$ is a parameter controlling temporal correlation, and $\xi_0 = \mathbf{0}$ is taken as a starting condition. We assume the ϕ_t 's to be independent and identically distributed as a proper intrinsically autoregressive model (Besag & Kooperberg, 1995); that is, $\phi_t \sim N(\mathbf{0}, \tau^2 (\mathbf{D} - \omega \mathbf{W})^{-1})$, where $\tau^2 > 0$ and $\omega \in (0, 1)$ is a so-called “propriety parameter” that ensures that the precision matrix is nonsingular (Banerjee et al., 2015). The neighborhood matrix $\mathbf{W} \in \mathbb{R}^{S \times S}$ is such that $(\mathbf{W})_{ss'}$ is equal to 1 if and only if location s is adjacent to location s' , $s \neq s'$, zero otherwise, and

$$\mathbf{D} = \text{diag} \left(\sum_{j=1}^S (\mathbf{W})_{sj}, s = 1, \dots, S \right).$$

To avoid confounding with the intercept, the standard sum-to-zero constraint

$$\sum_{t=1}^T \sum_{s=1}^S \xi_{st} = 0$$

is imposed.

The proposed model is completed by specifying prior distributions on the regression coefficients and the variance and correlation parameters. In the absence of strong prior information, hyperparameters are chosen to induce vague prior distributions. A Gaussian prior is assumed for the global regression coefficients, and inverse gamma (IG) priors are placed on variance components. A truncated Gaussian prior with support $(-1, 1)$ is specified for ζ . A Beta(α_ω, v_ω) prior is placed on ω to concentrate it close to unity, because previous empirical work has shown that $\omega \approx 1$ is necessary to induce noticeable spatial association (Banerjee et al., 2015). These specifications lead to the following hierarchy:

$$\begin{aligned}
Y_{st} | n_{st}, v_{st} &\stackrel{\text{indep.}}{\sim} \text{Binomial}(n_{st}, p_{st} = g(v_{st})), \quad s = 1, \dots, S; t = 1, \dots, T; \\
\beta_p^* | \sigma_p^2, \theta_p &\stackrel{\text{indep.}}{\sim} \text{N}(\mathbf{0}, \sigma_p^2 \mathbf{R}_p^*(\theta_p)), \quad p = 1, \dots, P; \\
\sigma_p^2 &\stackrel{\text{indep.}}{\sim} \text{IG}(\alpha_{\sigma_p^2}, v_{\sigma_p^2}), \quad p = 1, \dots, P; \\
\theta_p &\stackrel{\text{i.i.d.}}{\sim} \pi(\theta_p), \quad p = 1, \dots, P; \\
\delta &\sim \text{N}(\mathbf{0}, \sigma_\delta^2 \mathbf{I}), \quad \sigma_\delta^2 > 0; \\
\xi_t | \xi_{t-1}, \tau^2, \omega, \zeta &\sim \text{N}(\zeta \xi_{t-1}, \tau^2 (\mathbf{D} - \omega \mathbf{W})^{-1}), \quad t = 1, \dots, T; \\
\tau^2 &\sim \text{IG}(\alpha_{\tau^2}, v_{\tau^2}), \quad \alpha_{\tau^2}, v_{\tau^2} > 0; \\
\omega &\sim \text{Beta}(\alpha_\omega, v_\omega), \quad \alpha_\omega, v_\omega > 0; \\
\zeta &\sim \text{Truncated-Normal}(0, \sigma_\zeta^2, -1, 1), \quad \sigma_\zeta^2 > 0,
\end{aligned} \tag{3}$$

where $v_{st} = \mathbf{Z}'_{st} \boldsymbol{\delta} + \mathbf{X}'_{st} \tilde{\boldsymbol{\beta}}(\boldsymbol{\ell}_s) + \xi_{st}$, $\tilde{\boldsymbol{\beta}}(\boldsymbol{\ell}_s) = (\tilde{\beta}_1(\boldsymbol{\ell}_s), \dots, \tilde{\beta}_P(\boldsymbol{\ell}_s))'$, and $\xi_0 = \mathbf{0}$. Each coefficient in $\tilde{\boldsymbol{\beta}}(\boldsymbol{\ell}_s)$ is obtained from the P predictive processes via $\tilde{\beta}_p = \tilde{\mathbf{R}}_p^*(\mathbf{R}_p^*)^{-1} \beta_p^*$. Appropriate (identical) priors for $\theta_1, \dots, \theta_P$ depend on the correlation function selected in the GPP model.

While the combination of a continuous support GPP and a discrete support GMRF has not been extensively used previously, it is motivated in our application. Because the $\tilde{\beta}_\ell$ coefficients contain trends that are thought to vary smoothly over space, we appeal to GPP models because an explicit covariance function allows for the direct imposition of smoothness assumptions and a meaningful prediction function through kriging—the latter is useful for estimating seroprevalence trends at unobserved locations. The spatio-temporal random effects, on the other hand, are of secondary interest and serve only to smooth the extraregression variability beyond that explained by the predictors. Because they are defined over an areal lattice, a CAR model is a natural choice. In Web Appendix B, we consider replacing the GPP with a CAR model on the regression coefficients and empirically study the results via simulation. We find that the GPP model is able to produce reliable estimates of both “regional” and “local” trends, whereas the CAR model only estimates local trends. Because both models reliably estimate local trends and the GPP model can also accurately estimate regional trends, the GPP model is used. For further discussion of local and regional trends, including the differences between them, see Section 5 and Web Appendix B.

3 | POSTERIOR SAMPLING

3.1 | Data augmentation

We assume conditional independence given the covariate effects and spatio-temporal effects and observe that \mathbf{Y} depends on the regression coefficients and random effects only through $\mathbf{v} = (v_{11}, \dots, v_{1T}, v_{21}, \dots, v_{ST})'$. Hence, the likelihood is

$$f(\mathbf{Y} | \mathbf{v}) \propto \prod_{t=1}^T \prod_{s=1}^S g(v_{st})^{Y_{st}} \{1 - g(v_{st})\}^{n_{st} - Y_{st}}. \tag{4}$$

To develop a posterior sampling algorithm, let $g(\cdot)$ be the logistic link. Other link functions are possible and can be implemented following the work of Albert and Chib (1993) or Gamerman (1997). Metropolis–Hastings steps (Hastings, 1970; Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953) can be used either component-wise or in blocks, but such samplers can be difficult to tune in high dimensions. To facilitate the derivation of a Gibbs sampler for the regression coefficients and spatio-temporal random effects, a data augmentation scheme is used that leads to sampling these parameters from Gaussian full conditional distributions.

Our data augmentation approach follows the work of Polson, Scott, and Windle (2013) and relies on the fact that $\exp(\nu)^a \{1 + \exp(\nu)\}^{-b} = 2^{-b} \exp(\kappa\nu) \int_0^\infty \exp(-\psi\nu^2/2) p(\psi|b, 0) d\psi$, where $a \in \mathbb{R}$, $b \in \mathbb{R}^+$, $\kappa = a - b/2$, and $p(\cdot|b, 0)$ is the probability density function of a Pólya-gamma random variable with parameters b and 0. From these, under the logistic link, (4) can be written as

$$\begin{aligned} f(\mathbf{Y}|\boldsymbol{\nu}) &\propto \prod_{t=1}^T \prod_{s=1}^S \exp(\kappa_{st}\nu_{st}) \int_0^\infty \exp(-\psi_{st}\nu_{st}^2/2) p(\psi_{st}|n_{st}, 0) d\psi_{st} \\ &\propto \prod_{t=1}^T \prod_{s=1}^S \int_0^\infty f_{Y,\psi}(Y_{st}, \psi_{st}|\nu_{st}) d\psi_{st}, \end{aligned}$$

where $\kappa_{st} = Y_{st} - n_{st}/2$ and $f_{Y,\psi}$ is the joint density of (Y_{st}, ψ_{st}) . By introducing the ψ_{st} as latent random variables to be sampled via MCMC, we obtain

$$f_{Y,\psi}(\mathbf{Y}, \boldsymbol{\psi} | \boldsymbol{\nu}) \propto \exp(-\boldsymbol{\nu}' \mathbf{D}_\psi \boldsymbol{\nu} / 2 + \boldsymbol{\kappa}' \boldsymbol{\nu}) \prod_{t=1}^T \prod_{s=1}^S p(\psi_{st}|n_{st}, 0),$$

where $\boldsymbol{\psi} = (\psi_{11}, \dots, \psi_{1T}, \psi_{21}, \dots, \psi_{ST})'$, $\mathbf{D}_\psi = \text{diag}(\boldsymbol{\psi})$, and $\boldsymbol{\kappa} = (\kappa_{11}, \dots, \kappa_{1T}, \kappa_{21}, \dots, \kappa_{ST})'$. Hence, data augmentation yields a Gaussian density in $\boldsymbol{\nu}$ up to a normalizing constant. Consequently, the full conditional distributions for most parameters take a known form and are easy to sample. For specifics, the full conditional distribution of ψ_{st} is Pólya-gamma, $\boldsymbol{\beta}_p^*$ is multivariate normal, $\boldsymbol{\delta}$ is multivariate normal, σ_p^2 is IG, τ^2 is IG, and ζ is truncated normal. Web Appendix A provides additional conditional distributions.

From the data augmentation, a posterior sampling algorithm involving Gibbs steps for the above parameters can be constructed in the usual manner. Metropolis–Hastings steps are used to sample each $\boldsymbol{\theta}_p$ and ω . While the full conditional distribution of $\boldsymbol{\xi}_t$ is multivariate normal, sampling from this high-dimensional distribution is computationally expensive. For more efficient repeated updates of $\boldsymbol{\xi}_t$, the Markov structure of the CAR model is exploited to construct a chromatic sampler that updates conditionally independent blocks of $\boldsymbol{\xi}_t$ in parallel. For further discussion, see the works of Gonzalez, Low, Gretton, and Guestrin (2011) and Brown, McMahan, and Watson (2017).

3.2 | A note on missing data

In our application, data are not reported at all county–month pairs. To account for this, let \mathcal{R} be the set of all ordered pairs (s, t) for which tests are observed. The augmented likelihood is

$$f(\mathbf{Y}(\mathcal{R}), \boldsymbol{\psi}(\mathcal{R}) | \boldsymbol{\nu}(\mathcal{R})) \propto \exp(-\boldsymbol{\nu}(\mathcal{R})' \mathbf{D}_{\boldsymbol{\psi}(\mathcal{R})} \boldsymbol{\nu}(\mathcal{R}) / 2 + \boldsymbol{\kappa}(\mathcal{R})' \boldsymbol{\nu}(\mathcal{R})) \prod_{(s,t) \in \mathcal{R}} p(\psi_{st}|n_{st}, 0),$$

where $\boldsymbol{\nu}(\mathcal{R}) = \mathbf{Z}(\mathcal{R})\boldsymbol{\delta} + \mathbf{X}(\mathcal{R})\tilde{\boldsymbol{b}} + \mathbf{I}(\mathcal{R})\boldsymbol{\xi}$ and the convention that $\mathbf{A}(\mathcal{R})$ is the matrix formed by retaining the rows of \mathbf{A} whose indices are in \mathcal{R} is used. Here, $\mathbf{Z} = (\mathbf{Z}'_1, \dots, \mathbf{Z}'_S)' \in \mathbb{R}^{ST \times (Q+1)}$ with $\mathbf{Z}_s = (\mathbf{Z}_{s1}, \dots, \mathbf{Z}_{sT})'$. Similarly, $\mathbf{X} = \bigoplus_{s=1}^S \mathbf{X}_s \in \mathbb{R}^{ST \times SP}$ with $\mathbf{X}_s = (\mathbf{X}_{s1}, \dots, \mathbf{X}_{sT})'$, \mathbf{I} is the identity matrix, and $\tilde{\boldsymbol{b}} = (\tilde{\boldsymbol{\beta}}'(\mathcal{L}_1), \dots, \tilde{\boldsymbol{\beta}}'(\mathcal{L}_S))' \in \mathbb{R}^{SP}$. Because $\boldsymbol{\xi} \in \mathbb{R}^{ST}$ is the vector of spatial random effects over *all* locations within the study region for all time points, a well-defined full conditional distribution for $\boldsymbol{\xi}$ is obtained, provided that the prior on $\boldsymbol{\xi}$ is proper. This joint density representation permits the imputation of any missing effects via posterior realizations.

4 | A SIMULATION STUDY

This section studies via simulation how well our methods estimate model coefficients and how GPP knot selection influences results. Data were generated on a regularly spaced $S \times S$ grid over 60 time points, where $S = 13$, and then drawing $Y_{st}|n_{st}, p_{st} \stackrel{\text{indep.}}{\sim} \text{Binomial}(n_{st}, p_{st})$ observations, where

$$g^{-1}(p_{st}) = \delta_0 + \tilde{\boldsymbol{\beta}}_1(\mathcal{L}_s)t/60 + \xi_{st}, \quad s = 1, \dots, S; \quad t = 1, \dots, 60,$$

and $g(\cdot)$ is the logistic link. The test counts n_{st} were randomly sampled from a discrete uniform distribution ranging from 100 to 200. The random effects ξ_{st} are generated from the CAR model defined in Section 2 with $\zeta = 0.9$, $\tau^2 = 0.005$, $\omega \in \{0.00, 0.55, 0.90\}$, and a neighborhood matrix \mathbf{W} set so that two areas are neighbors if and only if they share a common edge or corner. The ω values 0.00, 0.55, and 0.90 correspond to no, weak, and strong spatial dependence, respectively.

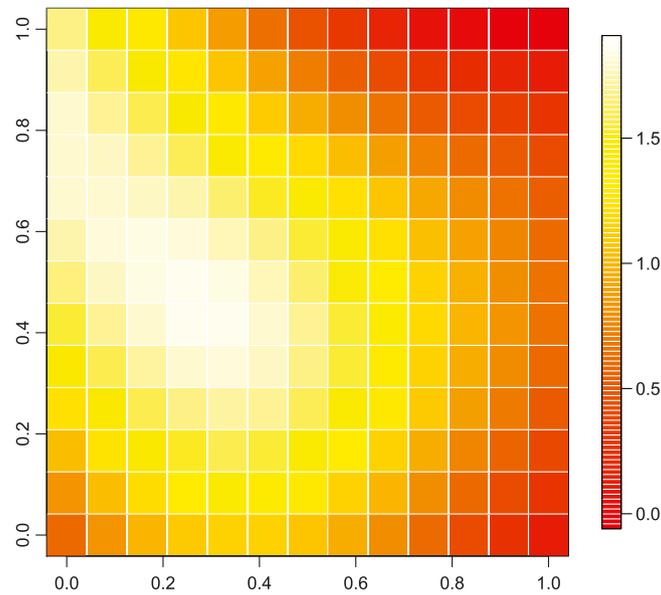


FIGURE 2 The true $\tilde{\beta}_1$ surface used to generate the independent data sets in the simulation example

The true intercept was set to $\delta_0 = -1$ and the surface $\tilde{\beta}_1(\cdot)$ at each study location is generated from the GPP model in (3). Specifically, a realization of the parent process is first simulated on a 5×5 grid of equally spaced knots. The parent process took $\mu_1(\mathcal{L}_s^*) \equiv 1$ and $\rho(\mathcal{L}_s^*, \mathcal{L}_{s'}^*; \theta_1) = \theta_1^{d_{ss'}^2}$, where $d_{ss'}$ is the Euclidean distance between \mathcal{L}_s^* and $\mathcal{L}_{s'}^*$, $\theta_1 = 0.6$, and $\sigma_1^2 = 1.5$. The resulting $\tilde{\beta}_1(\cdot)$ is depicted in Figure 2. Using this surface, 500 independent data sets were generated from the model for each ω .

Our model was fit to each data set using three separate knot set configurations. The first configuration uses the same knots as those generating the true surface, representing an ideal situation. The other two configurations take 4×4 and 7×7 grids of equally spaced knots. For priors in (3), we take $\alpha_{\sigma_1^2} = \nu_{\sigma_1^2} = \alpha_{\tau^2} = \nu_{\tau^2} = 2$, $\sigma_\delta^2 = 1,000$, $\alpha_\omega = 900$, $\nu_\omega = 100$, and $\sigma_\zeta^2 = 10$. In the GPP, the correlation function was taken as $\rho(\mathcal{L}_s, \mathcal{L}_{s'}; \theta_1) = \theta_1^{d_{ss'}^2}$, the same as the true GPP. A Uniform(0,1) prior on θ_1 was used. For each data set, 5,000 MCMC iterates are retained after a burn-in of 5,000 samples. Convergence of the chains was assessed via trace plots and judged to be acceptable.

Figure 3 summarizes our results for the temporal trend parameter $\tilde{\beta}_1(\cdot)$ when $\omega = 0.90$. This includes a spatial depiction of the arithmetic average of the 500 point estimates, as well as empirical biases and mean squared errors. Here, for each data set, a point estimate of $\tilde{\beta}_1(\cdot)$ was obtained as the mean of the 5,000 retained MCMC iterates. Web Figure 1 summarizes results for the other ω values. The methods estimate the spatially varying regression coefficient well for every considered ω ; that is, the mean estimates show little bias and have a relatively small mean squared error. Estimator variability increases near the region's edges—this boundary effect is expected and is common in nonparametric regressions. Figure 3 shows little practical difference for the estimates obtained under the three different knot configurations, suggesting that the methods can recover the true coefficient surface across the entire study region (assuming the model is correct up to choice of knots).

Two additional simulations were conducted. The first simulation examined the performance of the proposed methodology in the presence of missing data. The second study examined the effects of increasing spatial dimension. Results for these are presented in Web Appendix B. The results indicate that the proposed methodology performs well in these more challenging situations.

5 | LYME DISEASE ANALYSIS

5.1 | Background

Our data contain 16,571,562 tests on domestic dogs living throughout the conterminous United States from January 2012 to December 2016. The data were provided by IDEXX Laboratories, Inc. to the Companion Animal Parasite Council

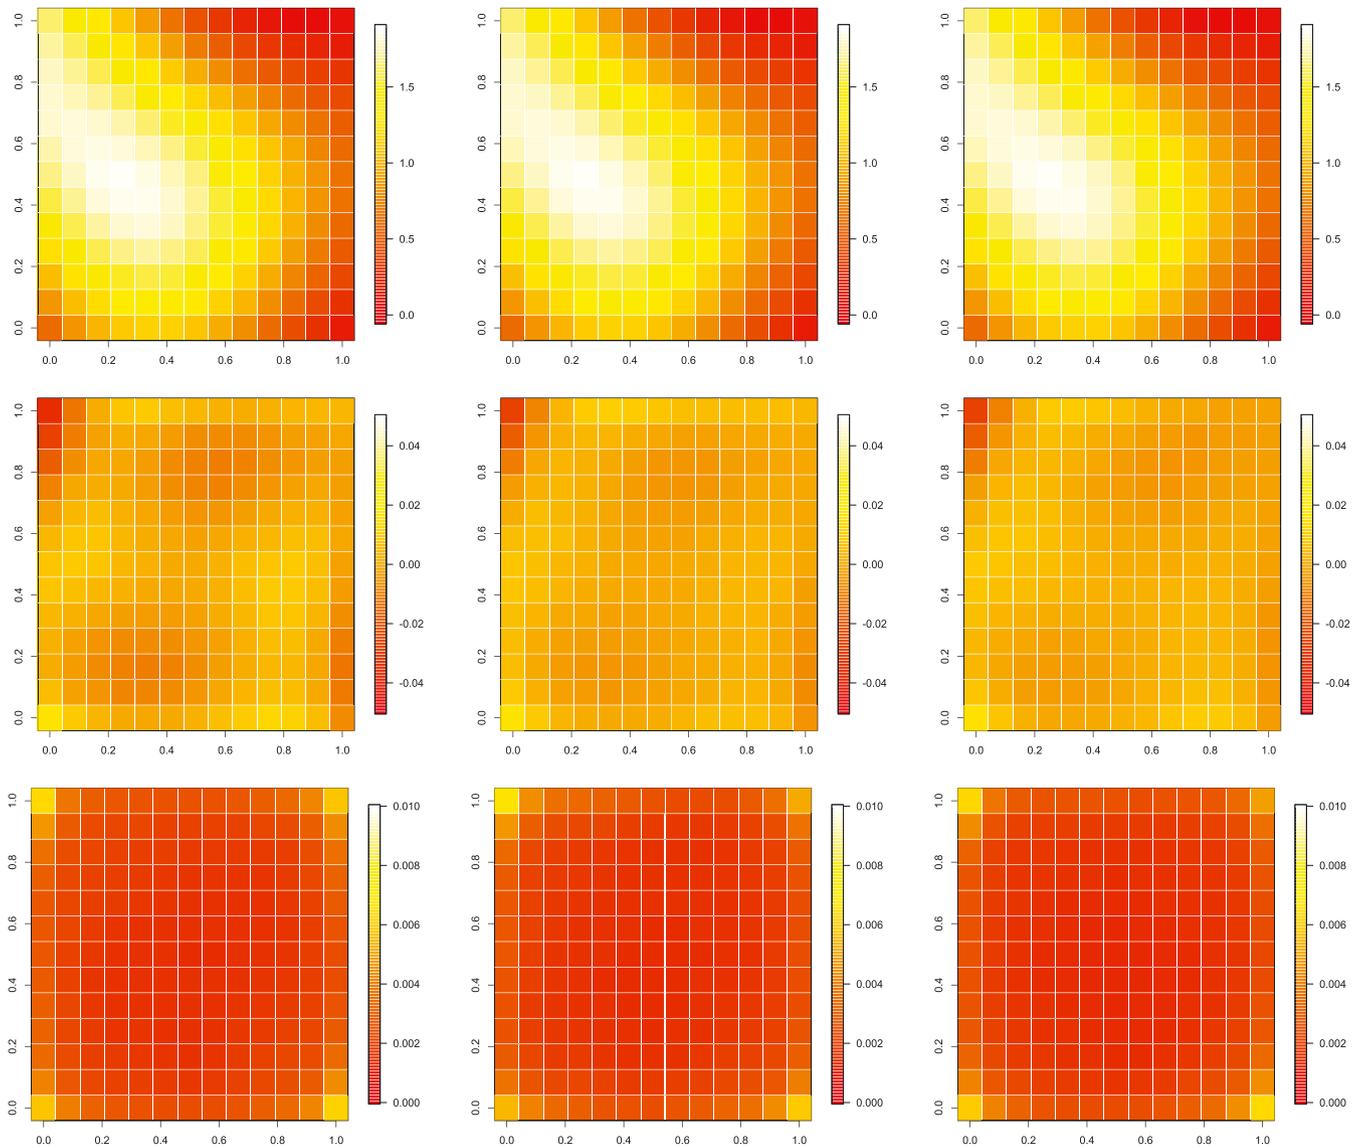


FIGURE 3 Summary of the posterior estimates of $\tilde{\beta}_1$ obtained in the simulation example when $\omega = 0.90$. Presented results include the sample mean of the (top row) posterior estimates, (middle row) empirical bias, and (bottom row) empirical mean squared error. From left to right, the columns correspond to the use of 4×4 , 5×5 , and 7×7 grids of knots

(CAPC), who made them available online at <https://www.capcvet.org>. The data are aggregated by month and county; 69,876 county-month pairs report at least one test.

In general, the spatial distribution of a vector-borne disease is strongly influenced by regional environments and the vector's hosts (e.g., deer populations), leading to correlated data (Legendre, 1993). A strong spatial correlation is seen in these data, as indicated by Figure 1 and a Moran's I statistic of 0.378 (p value ≈ 0). Such data are also positively temporally correlated. Figure 4 displays raw county-level prevalence estimates aggregated over all 12 months in the two years of 2012 and 2016. A comparison of these graphics suggests where a significant increase in prevalence is expected, including Western Pennsylvania, Virginia, West Virginia, Minnesota, and Iowa.

5.2 | Model building and seasonality

As an exploratory step, a county-by-county time series analysis of prevalence was conducted for 672 counties reporting a sufficient amount of data; that is, counties reporting 10 or more positive tests each month. Following the work of Dunsmuir and Scott (2015), a binomial generalized linear model with only a linear time trend (on the logit scale) was fit to each county's time series. The partial autocorrelations of the Pearson residuals were used to assess autoregressive

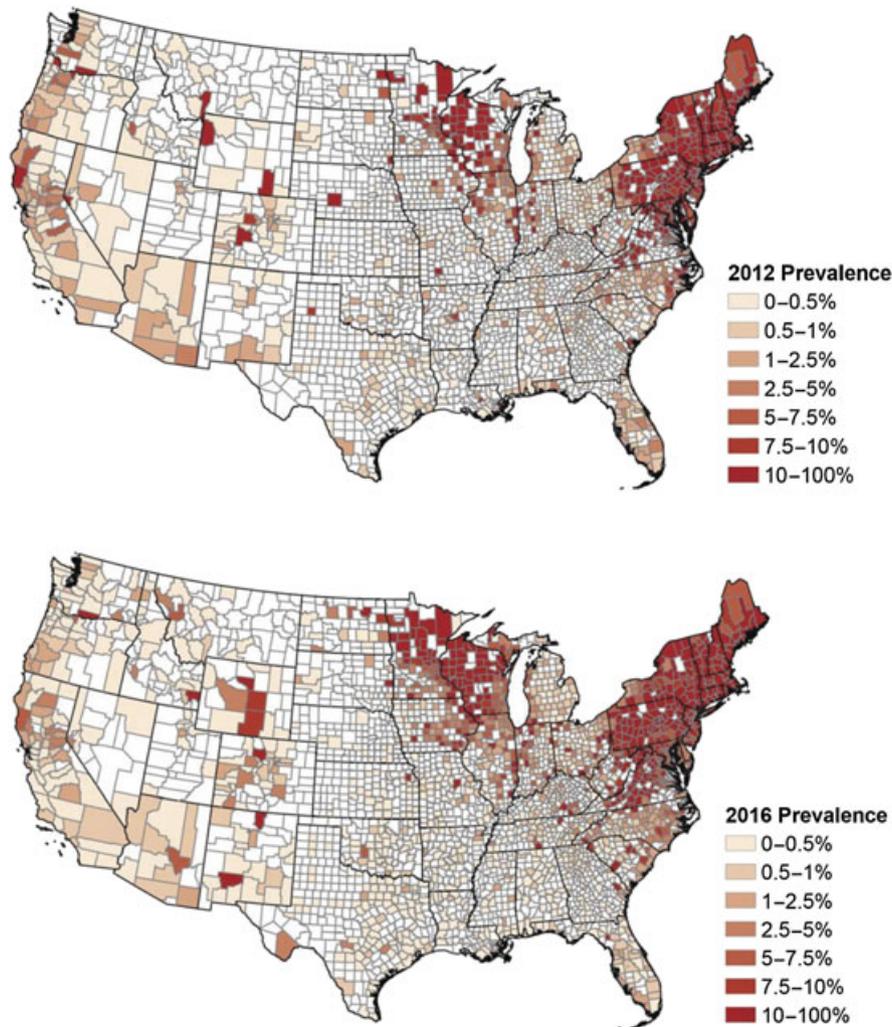


FIGURE 4 Raw reported canine seroprevalences in (top) 2012 and (bottom) 2016. White counties did not report any tests

orders in the usual manner. From this analysis, an AR(1) model was deemed reasonable for most counties. Next, a generalized linear AR(1) model with the mean structure described above was fit to each county. Histograms of the probability integral transformations were constructed and used to assess the suitability of the AR(1) model. These results support an AR(1) model to account for temporal dependence, providing justification for the form taken in (2). For more details on this analysis, see Web Appendix C.

Given the seasonality of tick activity, seasonality could also be present in Lyme disease prevalence. Though no strong evidence of seasonality surfaced in our exploratory analysis, a more thorough investigation of seasonality was conducted by fitting the model

$$v_{st} = \delta_0 + \tilde{\beta}_1(\mathcal{L}_s)I_1(t) + \tilde{\beta}_2(\mathcal{L}_s)I_2(t) + \tilde{\beta}_3(\mathcal{L}_s)I_3(t) + \tilde{\beta}_4(\mathcal{L}_s)t + \xi_{st}, \quad (5)$$

where t denotes time (rescaled to the unit interval) and $I_p(t)$ is a seasonal indicator for $p = 1, 2, 3$. Seasons are defined as follows: Winter (December–February), Spring (March–May), Summer (June–August), and Fall (September–November), where winter is regarded as the baseline. This model allows for spatially varying seasonal effects and spatially varying trend effects. While covariates such as county-level temperatures and precipitations are available, these are not used in this fit because our goal is to quantify trends, not determine the specific drivers of these trends.

Model (5) was fit with the prior specifications and correlation functions described in Section 4. Two specifications for the GPP model were considered, using 50 and 100 knots, respectively. In both cases, knot placement for all GPP models was done by K-means clustering. For sampling, 30,000 MCMC iterates were generated, with the last 10,000 being retained for inference. Convergence of the MCMC chains was assessed using trace plots. We stress the computational scalability of this approach. This model contains four a priori independent coefficient surfaces, each replete with 3,109 spatial locations and 186,540 spatio-temporal random effects.

Two primary findings arise. First, there are no appreciable differences between the estimates using 50 and 100 knots. As both specifications are computationally feasible, all subsequent analyses used 100 knots. Second, there is evidence of seasonality in the location parameters, but these appear constant across space. In particular, Figure 5 depicts 95% credible

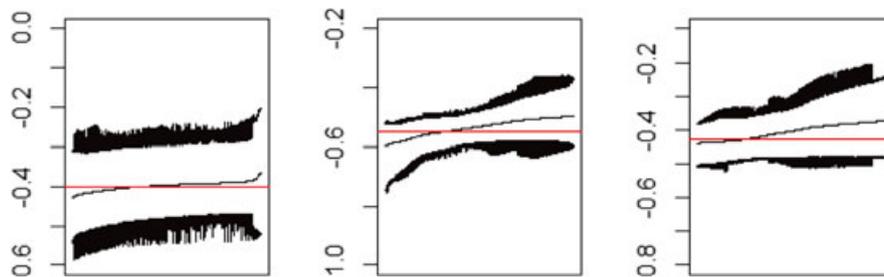


FIGURE 5 A summary of the spatially varying seasonal effects estimate from (5), that is, estimates (ordered from smallest to largest) of (left panel) $\tilde{\beta}_1$, (center panel) $\tilde{\beta}_2$, and (right panel) $\tilde{\beta}_3$. Included are the posterior mean and the upper and lower endpoints of 95% credible intervals. The red horizontal line is included to demonstrate that all of the credible intervals contain the same constant, thus indicating that a constant seasonal effect might be appropriate

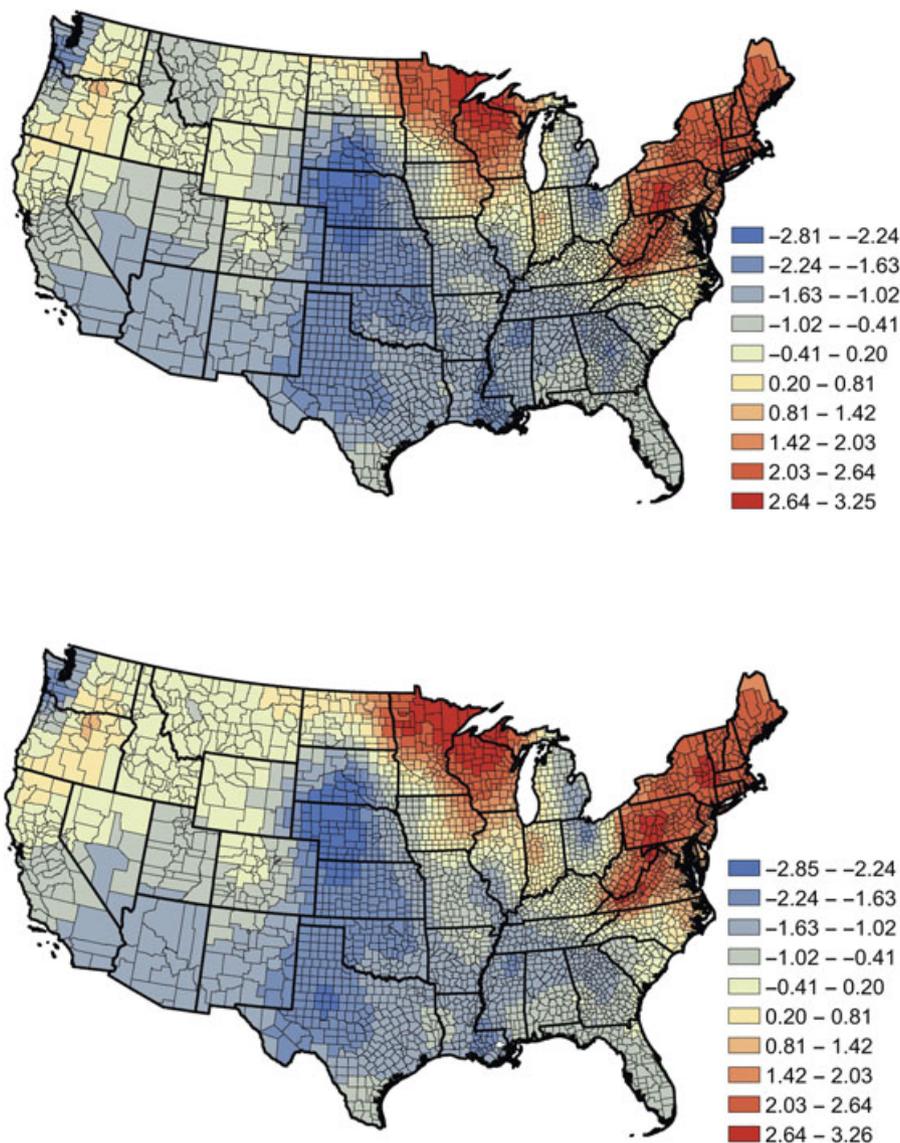


FIGURE 6 Estimate of the regional trend $\tilde{\beta}_1$ from the (top) seasonal model (7) and (bottom) nonseasonal model (8) used to analyze the seroprevalence data

intervals for each county-level seasonal effect. These intervals all contain a common nonzero value, indicating that a spatially constant seasonal effect is reasonable. Thus, the simpler model

$$v_{st} = \delta_0 + \delta_1 I_1(t) + \delta_2 I_2(t) + \delta_3 I_3(t) + \tilde{\beta}_1(\mathcal{L}_s)t + \xi_{st} \quad (6)$$

was considered. Credible intervals at level 95% indicate that the model can be further reduced to

$$v_{st} = \delta_0 + \delta_1 I_1^*(t) + \tilde{\beta}_1(\mathcal{L}_s)t + \xi_{st}, \quad (7)$$

where $I_1^*(t)$ is a seasonal indicator that equals one if t is between March and November, and zero otherwise. Approximate 95% credible intervals for δ_0 and δ_1 are $[-3.95, -3.82]$ and $[-0.20, -0.10]$, respectively.

For further insight, the model in (7) was compared with the nonseasonal model

$$v_{st} = \delta_0 + \tilde{\beta}_1(\mathcal{L}_s)t + \xi_{st}. \quad (8)$$

For this model, an approximate 95% credible interval for δ_0 is $[-4.08, -4.03]$. Figure 6 displays estimates of $\tilde{\beta}_1(\cdot)$ from both models. Very similar large-scale patterns in the estimated trends are seen. In short, while seasonality exists in the location parameters, its effect on trends seems negligible.

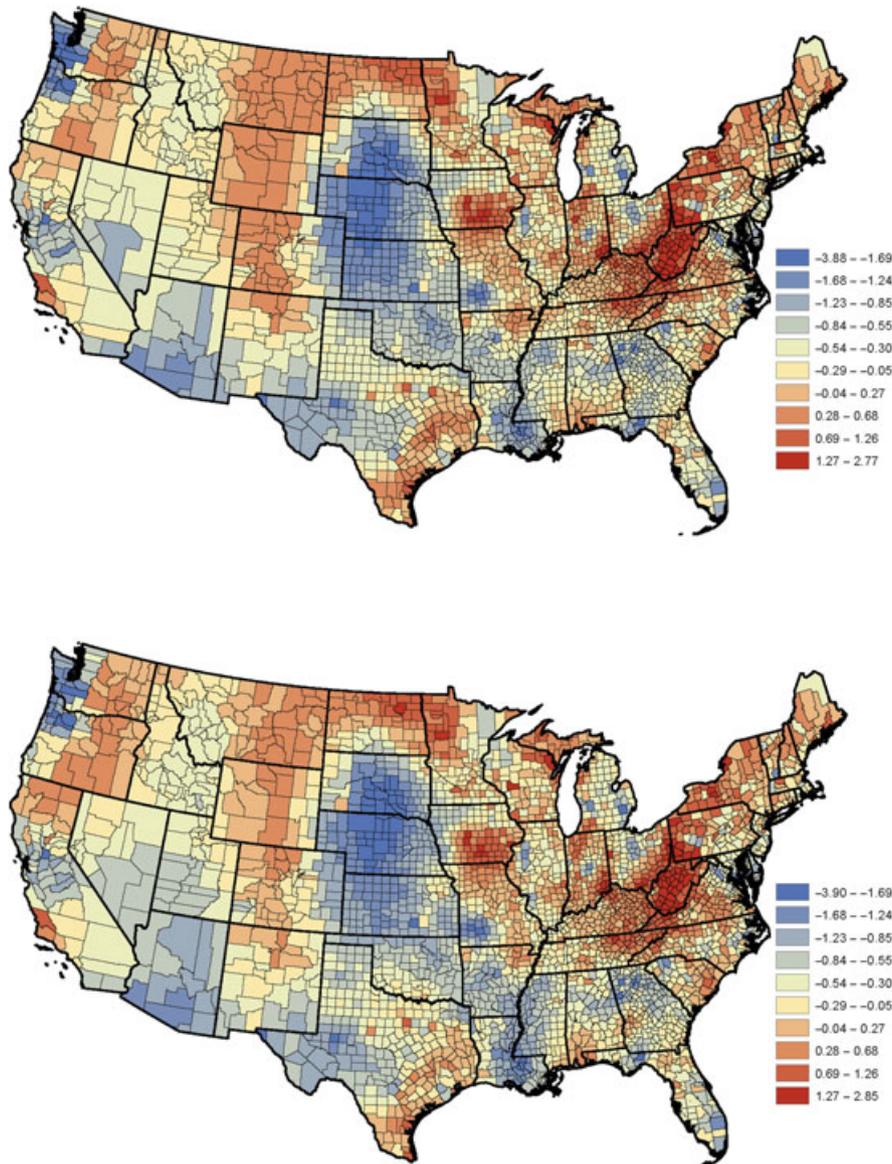


FIGURE 7 County-level trends. The top graphic displays the posterior mean estimate of v_s from model (7), and the bottom, from model (8)

The temporal trend surface $\tilde{\beta}$ represents a regional effect estimated at a particular area by incorporating information from a relatively large swath of surrounding areas. While regional trends are useful for estimating trends in areas with few tests, it may be desirable to separate local effects from regional trends to provide a county-level assessment. Web Appendix B demonstrates how our modeling framework facilitates such a separation. Specifically, county-specific trends are estimated as follows. Let $v_s^{(g)}$ be the least-squares estimate of the county s slope obtained at by fitting a simple linear regression to $\{(t, v_{st}^{(g)}) : t = 1, \dots, T\}$, where $v_{st}^{(g)}$ is the g th posterior draw of v_{st} obtained from the MCMC output. Then, $v_s^{(g)}$ can be regarded as a realization of the linear time trend at county s . Using the $\{v_s^{(g)}\}$ as a random sample from the marginal posterior distribution of v_s , point estimates and inferences can be obtained for county-level trends.

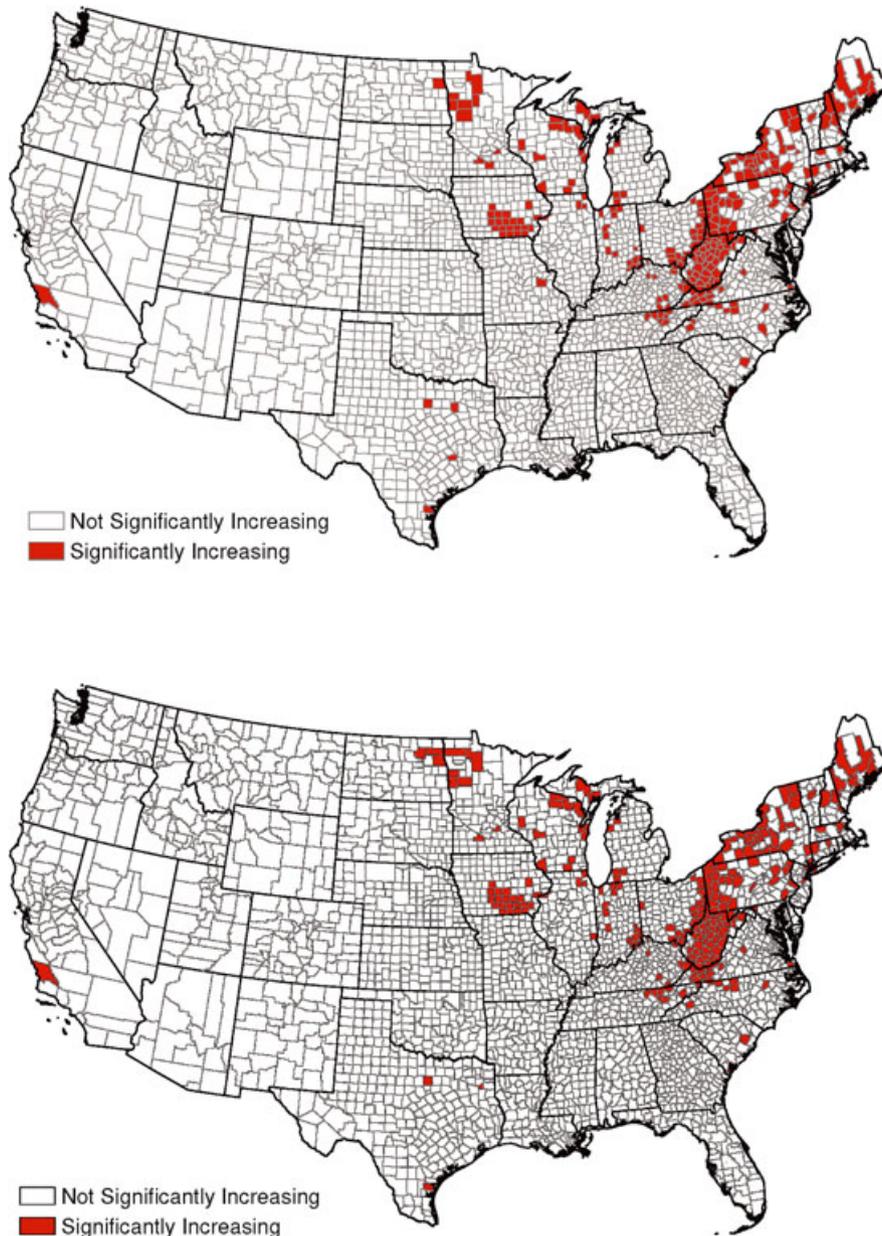


FIGURE 8 Counties where v_s was significantly positive at the 95 % confidence level. The top graphic corresponds to model (7), and bottom, to model (8)

5.3 | Results

Figure 6 displays the estimated posterior mean of the regional temporal trend surface $\tilde{\beta}_1$. The regional rate of change in prevalence between January 2012 and December 2016 is positive in all states that are currently recognized as having high human Lyme disease incidence (Centers for Disease Control and Prevention, 2017), including portions of the Northeast and the Upper Midwest. The rate of increase varies by region, with high-incidence regions generally exhibiting the greatest changes. These regions include Maine, West Virginia, Virginia, and the northern parts of Minnesota and Wisconsin.

Figure 7 displays estimated posterior means of the county-level trends v_s , $s = 1, \dots, 3,109$. Figure 8 shows counties where local trends are significantly positive, assessed using approximate 95% equal-tailed credible intervals. Increasing local trends are seen in much of the Northeast, extending southwards through West Virginia and Virginia, and into North Carolina and Tennessee. This region includes localities where Lyme disease is reportedly increasing. Increasing local trends in parts of Northwestern Minnesota, Northern Wisconsin, and Southeastern Iowa are also apparent. In the Great Lakes region, increasing trends are observed in Eastern Ohio, Indiana, and Western Michigan. In much of eastern New England, where human Lyme disease was first recognized, the prevalence appears stable, albeit high.

6 | DISCUSSION

This paper developed a computationally feasible binomial regression model for large spatio-temporal data that can identify localized trends. Our novel approach combined several recent advances in large-scale spatial modeling and MCMC sampling. The end product is a flexible scalable methodology for modern spatio-temporally referenced count data.

Our proposed approach was used to identify regions of the U.S. experiencing increasing canine Lyme disease risk. Because human and canine risks are similar, such regions are likely also experiencing increasing human exposure. While human Lyme disease data may not be publicly available and, in many regions, scarce due to lack of testing, our canine prevalence data had over 16 million tests. The size of the spatial domain created computational challenges. While monthly and county-level aggregation reduced the size of the response vector from 16,581,562 tests to 69,876 county-month pairs, a binomial response in an MCMC context typically requires sampling via Metropolis-Hastings steps, which can be difficult to tune in extremely high dimensions (over 180,000, in our case). Under the logistic link, a recently proposed Polyá-Gamma data augmentation was used to facilitate direct Gibbs sampling on full conditional distributions. GPPs were used to model smoothly varying high-dimensional coefficients through a low-dimensional representation. Local spatio-temporal heterogeneity was captured by random effects following a time-varying Gaussian CAR distribution. Chromatic sampling was used to facilitate efficient updating of the GMRFs in our MCMC algorithm.

This study was motivated by the rise in Lyme disease cases in the United States (Adams et al., 2017) and, in particular, rising incidence in states not traditionally considered endemic. Our results suggest that (1) canine prevalence is rising in tandem with human cases (Centers for Disease Control and Prevention, 2017; Hendricks & Mark-Carew, 2017; Kugeler, Farley, Forrester, & Mead, 2015), (2) prevalences are increasing most in areas where the pathogen has recently encroached, and (3) prevalence in dogs is rising in states traditionally not considered to be of high incidence for humans (Centers for Disease Control and Prevention, 2017), suggesting that human risk is also increasing in these areas. Several recent studies have recognized increasing risk in traditionally low incidence areas. These areas include Illinois (Herrmann, Dahm, Ruiz, & Brown, 2014), Iowa (Lingren, Rowley, Thompson, & Gilchrist, 2005), North Dakota (Russart, Dougherty, & Vaughan, 2014), Ohio (Wang et al., 2014), and Michigan (Lantos et al., 2017). Significant increases in canine prevalence are also seen in some areas that have not yet reported significant human incidence. Given the proximity of these locations to recognized high-incidence areas, it is reasonable to infer that canine prevalence is more sensitive to changes and can be used as an early warning system to signal changes in human risk. West Virginia, Western Pennsylvania, and Eastern Ohio are such areas and can be viewed as a leading edge of rising prevalence in Lyme disease's westward expansion. This is supported by evidence in increased reports of ticks in these regions (Eisen, Eisen, & Beard, 2016).

Examining local, as opposed to regional, trends shows that some adjacent counties have trends in opposite directions. To fully understand this heterogeneity, further ecological analyses are needed. Possible factors to consider include the presence of urban centers, degree of forestation or other habitat factors, tick populations, reservoir presence and densities, vaccination, and preventative medication use. Medication use is likely driven by socioeconomic factors, whereas the other factors are related to climate or changing habitats.

Our approach made several simplifying assumptions. The link function in the model was treated as known, which might be a strong assumption. As poorly specified link functions can induce bias in the estimates of the covariate effects (Neuhaus, 1999), relaxing this assumption could be fruitful. We also assumed that the spatially varying coefficients followed independent and identically distributed GPs. A more flexible approach would allow these coefficients to be correlated through a multivariate GP (Ver Hoef & Barry, 1998); however, multivariate GPs are more difficult to use and challenges remain in their development (e.g., Fricker, Oakley, & Urban, 2013). The observed prevalence rates suggest that smoothness of the random effects may change by region, suggesting that a heteroskedastic GP might be more appropriate (Binois, Gramacy, & Ludkovski, 2016). Further, GMRFs are known to oversmooth salient features (Smith & Fahrmeir, 2007). However, approximating GPs with GMRFs via stochastic partial differential equations to maintain computational feasibility (Lindgren et al., 2011) could prove promising for our application.

In addition to statistical challenges, future applications of our model include human Lyme disease and heartworm disease, ehrlichiosis, and anaplasmosis in dogs. The ecological, entomological, and environmental implications of the canine prevalence analysis presented here is the subject of ongoing research.

ACKNOWLEDGEMENTS

The authors thank IDEXX Laboratories, Inc. for their data. This work was partially supported by the National Science Foundation Grants DMS-1127914, DMS-1407480, CMMI-1563435, and EEC-1744497, and the National Institutes of Health Grant R01 AI121351. Stella Watson Self, Christopher McMahan, Robert Lund, Jenna Gettings, and Michael Yabsley thank the Companion Animal Parasite Council (CAPC) for the support and initiating this study; Jenna Gettings was further supported by the Boehringer Ingelheim Vetmedica-CAPC Infectious Disease Postdoctoral Fellowship. Comments from two referees greatly improved our presentation.

ORCID

Christopher S. McMahan  <http://orcid.org/0000-0001-5056-9615>

REFERENCES

- Adams, D. A., Thomas, K. R., Jajosky, R. A., Foster, L., Baroi, G., Sharp, P., ... Anderson, W. J. (2017). Summary of notifiable infectious diseases and conditions—United States, 2015. *Morbidity and Mortality Weekly Report*, 64, 1–143.
- Albert, J. H., & Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422), 669–679.
- Anderson, C., & Ryan, L. M. (2017). A comparison of spatio-temporal disease mapping approaches including an application to ischaemic heart disease in New South Wales, Australia. *International Journal of Environmental Research and Public Health*, 14(2), 146.
- Andrade-Pacheco, R., Mubangizi, M., Quinn, J., & Lawrence, N. (2015). Monitoring short term changes in infectious diseases in Uganda with Gaussian processes. In A. Douzal-Chouakria, J. A. Vilar, & P.-F. Marteau (Eds.), *Advanced analysis and learning on temporal data: First ECML PKDD workshop, AALTD 2015, Porto, Portugal, September 11, 2015, revised selected papers* (95–110). Cham, Switzerland: Springer International Publishing Switzerland.
- Banerjee, S., Carlin, B. P., & Gelfand, A. E. (2015). *Hierarchical modeling and analysis for spatial data* (2nd ed.). Boca Raton, FL: Chapman and Hall/CRC.
- Banerjee, S., Gelfand, A. E., Finley, A. O., & Sang, H. (2008). Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Methodological)*, 70(4), 825–848.
- Bayarri, M. J., Berger, J. O., Paulo, R., Sacks, J., Cafeo, J. A., Cavendish, J., ... Tu, J. (2007). A framework for validation of computer models. *Technometrics*, 49(2), 138–154.
- Berger, J. O., de Oliveira, V., & Sansó, B. (2001). Objective Bayesian analysis of spatially correlated data. *Journal of the American Statistical Association*, 96(456), 1361–1374.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2), 192–236.
- Besag, J., & Kooperberg, C. (1995). On conditional and intrinsic autoregressions. *Biometrika*, 82, 733–746.
- Binois, M., Gramacy, R. B., & Ludkovski, M. (2016). Practical heteroskedastic Gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, 1–41.
- Brezger, A., Fahrmeir, L., & Hennerfeind, A. (2007). Adaptive Gaussian Markov random fields with applications in human brain mapping. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 56(3), 327–345.
- Brown, D. A., Datta, G. S., & Lazar, N. A. (2017). A Bayesian generalized CAR model for correlated signal detection. *Statistica Sinica*, 27, 1125–1153.

- Brown, D. A., McMahan, C. S., & Watson, S. C. (2017). Sampling strategies for fast updating of Gaussian Markov random fields. ArXiv preprint arXiv: 1702:05518.
- Centers for Disease Control and Prevention (2017). Data and statistics: Lyme disease.
- Congdon, P., & Southall, H. (2005). Trends in inequality in infant mortality in the north of England, 1921–1973, and their association with urban and social structure. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 168(4), 679–700.
- Cressie, N. (1993). *Statistics for spatial data*. Hoboken, NJ: John Wiley & Sons.
- Cressie, N., & Johannesson, G. (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society: Series B (Methodological)*, 70, 209–226.
- Cressie, N., & Wikle, C. K. (2011). *Statistics for spatio-temporal data*. Hoboken, NJ: John Wiley & Sons.
- Datta, A., Banerjee, S., Finley, A. O., & Gelfand, A. E. (2016). Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111, 800–812.
- Diggle, P. J., Tawn, J. A., & Moyeed, R. A. (1998). Model-based geostatistics. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 47(3), 299–350.
- Dunsmuir, W., & Scott, D. (2015). The glarma package for observation-driven time series regression of counts. *Journal of Statistical Software*, 67(7), 1–36.
- Eidsvik, J., Finley, A. O., Banerjee, S., & Rue, H. (2012). Approximate Bayesian inference for large spatial datasets using predictive process models. *Computational Statistics and Data Analysis*, 56(6), 1362–1380.
- Eisen, R. J., Eisen, L., & Beard, C. B. (2016). County-scale distribution of *Ixodes scapularis* and *Ixodes pacificus* (Acari: Ixodidae) in the continental United States. *Journal of Medical Entomology*, 53, 349–386.
- Finley, A. O., Sang, H., Banerjee, S., & Gelfand, A. E. (2009). Improving the performance of predictive process modeling for large datasets. *Computational Statistics and Data Analysis*, 53(8), 2873–2884.
- Fricker, T. E., Oakley, J. E., & Urban, N. M. (2013). Multivariate Gaussian process emulators with nonseparable covariance structures. *Technometrics*, 55(1), 47–56.
- Furrer, R., Genton, M. G., & Nychka, D. (2006). Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15, 502–523.
- Gamerman, D. (1997). Sampling from the posterior distribution in generalized linear mixed models. *Statistics and Computing*, 7(1), 57–68.
- Gelfand, A. E., Kim, H.-J., Sirmans, C. F., & Banerjee, S. (2003). Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association*, 98(462), 387–396.
- Gelfand, A. E., & Schliep, E. M. (2016). Spatial statistics and Gaussian processes: A beautiful marriage. *Spatial Statistics*, 18(Part A), 86–104.
- Gonzalez, J. E., Low, Y., Gretton, A., & Guestrin, C. (2011). Parallel Gibbs sampling: From colored fields to thin junction trees. In G. Gordon, D. Dunson, & M. Dudik (Eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, Volume 15 of Proceedings of Machine Learning Research*, Fort Lauderdale, FL, pp. 324–332.
- Gramacy, R. B., & Apley, D. W. (2015). Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24, 561–578.
- Guhaniyogi, R., Finley, A. O., Banerjee, S., & Gelfand, A. E. (2011). Adaptive Gaussian predictive process models for large spatial datasets. *Environmetrics*, 22(8), 997–1007.
- Hartigan, J. A., & Wong, M. A. (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1), 100–108.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57, 97–109.
- Heaton, M. J., Christensen, W. F., & Terres, M. A. (2017). Nonstationary Gaussian process models using spatial hierarchical clustering from finite differences. *Technometrics*, 59, 93–101.
- Heaton, M. J., Datta, A., Finley, A., Furrer, R., Guhaniyogi, R., Gerber, F., ... Zammit-Mangion, A. (2017). Methods for analyzing large spatial data: A review and comparison. ArXiv preprint arXiv: 1710.05013.
- Hendricks, B., & Mark-Carew, M. (2017). Using exploratory data analysis to identify and predict patterns of human Lyme disease case clustering within a multistate region, 2010–2014. *Spatial and Spatio-Temporal Epidemiology*, 20, 35–43.
- Herrmann, J. A., Dahm, N. M., Ruiz, M. O., & Brown, W. M. (2014). Temporal and spatial distribution of tick-borne disease cases among humans and canines in Illinois (2000–2009). *Environmental Health Insights*, 8(Supplement 2), 15.
- Johnson, L. R., Gramacy, R. B., Cohen, J., Mordecai, E., Murdock, C., Rohr, J., ... Weikel, D. (2017). Phenomenological forecasting of disease incidence using heteroskedastic Gaussian processes: A dengue case study. *Annals of Applied Statistics*. In Press.
- Jona-Lasinio, G., Gelfand, A., & Jona-Lasinio, M. (2012). Spatial analysis for wave direction data using wrapped Gaussian processes. *Annals of Applied Statistics*, 6(4), 1478–1498.
- Katzfuss, M. (2017). A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association*, 112, 201–214.
- Knorr-Held, L., & Besag, J. (1998). Modelling risk from a disease in time and space. *Statistics in Medicine*, 17(18), 2045–2060.
- Kugeler, K. J., Farley, G. M., Forrester, J. D., & Mead, P. S. (2015). Geographic distribution and expansion of human Lyme disease, United States. *Emerging Infectious Diseases*, 21(8), 1455–1457.
- Lantos, P. M., Tsao, J., Nigrovic, L. E., Auwaerter, P. G., Fowler, V. G., Ruffin, F., ... Hickling, G. (2017). Geographic expansion of Lyme disease in Michigan, 2000–2014. *Open Forum Infectious Diseases* (Vol. 4, No. 1). Oxford University Press.

- Lazar, N. A. (2008). *The statistical analysis of functional MRI data*. New York, NY: Springer Science+Business Media.
- Lee, D., & Lawson, A. (2014). Cluster detection and risk estimation for spatio-temporal health data. ArXiv preprint arXiv: 1408.1191.
- Legendre, P. (1993). Spatial autocorrelation: Trouble or new paradigm? *Ecology*, 74(6), 1659–1673.
- Levy, S. A., & Magnarelli, L. (1992). Relationship between development of antibodies to *Borrelia burgdorferi* in dogs and the subsequent development of limb/joint borreliosis. *Journal of the American Veterinary Medical Association*, 200(3), 344–347.
- Lindgren, F., Rue, H., & Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Methodological)*, 73(4), 423–498.
- Lingren, M., Rowley, W. A., Thompson, C., & Gilchrist, M. (2005). Geographic distribution of ticks (Acari: Ixodidae) in Iowa with emphasis on *Ixodes scapularis* and their infection with *Borrelia burgdorferi*. *Vector-Borne and Zoonotic Diseases*, 5(3), 219–226.
- Little, S. E., Heise, S. R., Blagburn, B. L., Callister, S. M., & Mead, P. S. (2010). Lyme borreliosis in dogs and humans in the USA. *Trends in Parasitology*, 26(4), 213–218.
- MacNab, Y. C., & Dean, C. B. (2001). Autoregressive spatial smoothing and temporal spline smoothing for mapping rates. *Biometrics*, 57, 949–956.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models* (2nd ed.). London, UK: Chapman and Hall/CRC.
- Mead, P., Goel, R., & Kugeler, K. (2011). Canine serology as adjunct to human Lyme disease surveillance. *Emerging Infectious Diseases*, 17(9), 1710–1712.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21, 1087–1092.
- Morales-Álvarez, P., Pérez-Suay, A., Molina, R., & Camps-Valls, G. (2017). Remote sensing image classification with large-scale Gaussian processes. *IEEE Transactions on Geoscience and Remote Sensing*, 56(2), 1103–1114.
- Neal, R. M. (1998). Regression and classification using Gaussian process priors. In J. M. Bernardo, J. O. Berger, & A. P. Dawid (Eds.), *Bayesian statistics 6* (pp. 475–501). New York, NY: Oxford University Press.
- Nelder, M. P., Russell, C. B., Sheehan, N. J., Sander, B., Moore, S., Li, Y., ... Sider, D. (2016). Human pathogens associated with the blacklegged tick *Ixodes scapularis*: A systematic review. *Parasites and Vectors*, 9(1), 265.
- Neuhaus, J. M. (1999). Bias and efficiency loss due to misclassified responses in binary regression. *Biometrika*, 86(4), 843–855.
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., & Sain, S. (2015). A multiresolution Gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics*, 24, 579–599.
- O'Hagan, A. (1978). Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society: Series B (Methodological)*, 40, 1–42.
- Polson, N. G., Scott, J. G., & Windle, J. (2013). Bayesian inference for logistic models using Pólya-Gamma latent variables. *Journal of the American Statistical Association*, 108(504), 1339–1349.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.
- Rue, H., & Held, L. (2005). *Gaussian Markov random fields: Theory and applications*. Boca Raton, FL: Chapman and Hall/CRC.
- Rue, H., & Tjelmeland, H. (2002). Fitting Gaussian Markov random fields to Gaussian fields. *Scandinavian Journal of Statistics*, 29, 31–49.
- Rushworth, A., Lee, D., & Mitchell, R. (2014). A spatio-temporal model for estimating the long-term effects of air pollution on respiratory hospital admissions in greater London. *Spatial and Spatio-Temporal Epidemiology*, 10, 29–38.
- Russart, N. M., Dougherty, M. W., & Vaughan, J. A. (2014). Survey of ticks (Acari: Ixodidae) and tick-borne pathogens in North Dakota. *Journal of Medical Entomology*, 51(5), 1087–1090.
- Sang, H., Jun, M., & Huang, J. Z. (2011). Covariance approximation for large multivariate spatial datasets with an application to multiple climate model errors. *Annals of Applied Statistics*, 5(4), 2519–2548.
- Santner, T. J., Williams, B. J., & Notz, W. I. (2003). *The design and analysis of computer experiments*. New York, NY: Springer-Verlag New York.
- Senanayake, R., O'Callaghan, S., & Ramos, F. (2016). Predicting spatio-temporal propagation of seasonal influenza using variational Gaussian process regression. *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*, 3901–3907.
- Smith, M., & Fahrmeir, L. (2007). Spatial Bayesian variable selection with application to functional magnetic resonance imaging. *Journal of the American Statistical Association*, 102(478), 417–431.
- Song, H.-R., Fuentes, M., & Ghosh, S. (2008). A comparative study of Gaussian geostatistical models and Gaussian Markov random fields. *Journal of Multivariate Analysis*, 99, 1681–1697.
- Ver Hoef, J. M., & Barry, R. P. (1998). Constructing and fitting models for cokriging and multivariable spatial prediction. *Journal of Statistical Planning and Inference*, 69, 275–294.
- Wagner, B., Freer, H., Rollins, A., Garcia-Tapia, D., Erb, H. N., Earnhart, C., ... Meeus, P. (2012). Antibodies to *Borrelia burgdorferi* OspA, OspC, OspF, and C6 antigens as markers for early and late infection in dogs. *Clinical and Vaccine Immunology*, 19(4), 527–535.
- Waller, L. A., Carlin, B. P., Xia, H., & Gelfand, A. E. (1997). Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical Association*, 92(438), 607–617.
- Wang, P., Glowacki, M. N., Hoet, A. E., Needham, G. R., Smith, K. A., Gary, R. E., & Li, X. (2014). Emergence of *Ixodes scapularis* and *Borrelia burgdorferi*, the Lyme disease vector and agent, in Ohio. *Frontiers in Cellular and Infection Microbiology*, 4, 70.
- Yue, Y., & Speckman, P. L. (2010). Nonstationary spatial Gaussian Markov random fields. *Journal of Computational and Graphical Statistics*, 19(1), 96–116.
- Zhang, H., & El-Shaarawi, A. (2009). On spatial skew-Gaussian processes and applications. *Environmetrics*, 21(1), 33–47.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

The supplementary material includes three web appendices. Web Appendix A provides the full conditional distributions required to develop the proposed sampling procedure, Web Appendix B provides additional simulation results, and Web Appendix C provides details about the exploratory data analysis conducted in Section 5.2.

How to cite this article: Self SCW, McMahan CS, Brown DA, Lund RB, Gettings JR, Yabsley MJ. A large-scale spatio-temporal binomial regression model for estimating seroprevalence trends. *Environmetrics*. 2018;29:e2538. <https://doi.org/10.1002/env.2538>