HHS Public Access

Author manuscript

Spat Spatiotemporal Epidemiol. Author manuscript; available in PMC 2020 April 14.

Published in final edited form as:

Spat Spatiotemporal Epidemiol. 2019 June ; 29: 97–109. doi:10.1016/j.sste.2019.01.002.

Quantifying geographic regions of excess stillbirth risk in the presence of spatial and spatio-temporal heterogeneity

David Zahrieh[1], **Jacob J. Oleson**[2], **Paul A. Romitti**[3]

[1]Department of Health Sciences Research, Mayo Clinic, Rochester, MN 55905, U.S.A

[2]Department of Biostatistics, The University of Iowa, Iowa City, IA 52242, U.S.A

[3]Department of Epidemiology, The University of Iowa, Iowa City, IA 52242, U.S.A

Abstract

Motivated by population-based geocoded data for Iowa stillbirths and live births delivered during 2005–2011, we sought to identify spatio-temporal variation of stillbirth risk. Our high-quality data consisting of point locations of these delivery events allows use of a Bayesian Poisson point process approach to evaluate the spatial pattern of events. With this large epidemiologic dataset, we implemented the integrated nested Laplace approximation (INLA) to fit the conditional formulation of the point process via a Bayesian hierarchical model and empirically showed that INLA, compared to Markov chain Monte Carlo (MCMC) sampling, is an attractive approach. Furthermore, we modeled the temporal variability in stillbirth to better understand how stillbirths are geographically linked over the seven-year study period and demonstrate the similarity between the conditional formulation of the spatio-temporal model and a log Gaussian Cox process governed by discrete space-time random fields. After controlling for important features of the data, the Bayesian temporal relative risk maps identified areas of increasing and decreasing stillbirth risk over the birth period, which may warrant further public health investigation in the regions identified.

Keywords

Bayesian; Point process; Spatio-temporal heterogeneity; Stillbirth

1. Introduction

In Iowa, a stillbirth is defined as a fetal death with a gestational weight ≥ 350 grams or a gestation age of ≥ 20 weeks. The estimated prevalence for stillbirth in Iowa is about one in 180 pregnancies compared to the U.S. estimate of one in 160 pregnancies (NICHD, 2017). We seek to learn more about where and when stillbirths are occurring by examining both the spatial and temporal patterns in stillbirths. The location of a stillbirth can be represented by the maternal residence at the time of delivery. The Iowa Registry for Congenital and Inherited Disorders (IRCID) has been actively monitoring stillbirth deliveries statewide

Correspondence: David Zahrieh, Department of Health Sciences Research, Mayo Clinic, Rochester, MN 55905, U.S.A. zahrieh.david@mayo.edu.

since 2005 (Romitti, 2015). For the birth period 2005–2011, the maternal residence at delivery for each stillbirth and live birth delivered to an Iowa resident was geocoded and the corresponding date of delivery recorded. The combined data for individual stillbirth events across a specified study period can form a point map of events within a spatial and temporal region. With our rich source of geocoded data, we aim to model the spatial and temporal patterns in stillbirth to learn more about how the spatial and temporal patterns have changed and how risk factors might be related to these changes.

During the seven-year study period, the ratio of the number of stillbirth to live birth events was 0.0044, or 0.44%; pregnancies with multiple fetuses were excluded. Because little is known about the underlying mechanism driving the spatial distribution of stillbirth events and with the relatively precise geocoded data available, a point process modeling approach was initially applied to these population-based surveillance data to quantify excess stillbirth risk (Zahrieh et al., 2018). The point process method allowed for analysis of multiple births by the same mother in independent pregnancies via a maternal contextual effect; this effect was a random factorial effect whereby stillbirth events experienced by the same mother from independent pregnancies were grouped so that the events within a group shared the same random effect. The methodologic approach adopted for our previous analysis identified and quantified several geographic regions of excess stillbirth risk beyond the underlying at-risk population (i.e. maternal residence at delivery of each live birth). Because our starting point focused on modeling the spatial distribution of stillbirth events, we did not take into account the timing of the events. Furthermore, a limitation of the initial methodologic approach was that it could not incorporate covariate information attached to the live births. In the current data analysis, we apply an appealing approach that can analytically address both of these insufficiencies.

To our knowledge, risk factor studies for stillbirth have predominantly been from population-based case-control studies (Flenady et al., 2011, provide a systematic review), with a few studies incorporating minimal geospatial information into data analyses (for example, Hall et al., 2014 and DeFranco et al., 2015). Logistic regression was the corresponding analytic method used to estimate associations in these studies; that is, the outcome events (stillbirth versus live birth) at observed spatial locations were modeled as conditionally independent binomial outcomes, but the spatial and temporal variation inherent in the locations and timing of the events within the study-specific and temporal region were ignored. Statistical methods that take the spatial arrangement of maternal residence and timing of the deliveries into account can potentially provide additional insights into antecedents that contribute to stillbirth occurrence.

In this paper, we use a conditional formulation of the point process via a Bayesian hierarchical model to view the joint realization of stillbirths and live births and, conditional on this realization, examine the probability that the binary label on a point is either a stillbirth or live birth. This spatial dependence addresses the labeling (delivery event), rather than the event locations themselves and simplifies analysis and interpretation compared with modeling maternal residence directly via a point process model. Importantly, with the conditional formulation of the point process we can now incorporate covariate information attached to both stillbirth and live birth. Our approach extends the conditional formulation of

a Bayesian point process model to include both spatial and temporal effects and to study empirically the recovery of spatial and temporal model components in this framework. We note that adding time (i.e. delivery date) to the investigation may be critical with respect to adequate understanding of how maternal residences are geographically linked. Furthermore, for most epidemiologic applications, the relations of individual level outcomes to individual level predictors are examined, and a compelling argument can be made to consider spatial and temporal effects as contextual effects (Lawson, 2012). The analysis of our stillbirth surveillance data falls within this framework. However, to determine the value of adding a temporal component to our model, we fit both spatial and spatio-temporal models to our data and applied a Bayesian model selection criterion to determine the better fitting model. In any event, we aim to apply the conditional formulation to the stillbirth and live birth data to investigate possible risk factors associated with the probability of a stillbirth event in the presence of spatial and spatio-temporal variation. Additionally, after accounting for selected risk factors, as well as a maternal contextual effect, we aim to identify possible geographic regions of high intensity within the spatial and temporal region warranting further investigation.

In developing our analytic approach, we acknowledge that a Markov Chain Monte Carlo (MCMC) approach conventionally is used to estimate posterior quantities for Bayesian models. Approximation to posterior distributions is also available through other techniques; therefore, we chose to use integrated nested Laplace approximation (INLA) in the R INLA package (H. Rue and Chopin (2009), Lindgren et al. (2011), Martins et al. (2013), Simpson et al. (2012a), Simpson et al. (2012b)). We chose to apply INLA, because this approach does not require posterior sampling methods, provides estimates quickly, and is amenable to large data sets. The INLA approximation to the posterior distributions ideally should be similar to the estimated posterior distributions obtained via MCMC sampling. To support the application of INLA in our work, we compare how similar the two approaches are by using the conditional formulation of the point process when the stillbirth event locations arise from a log Gaussian Cox process (LGCP). Conditional on the intensity of the process, the stillbirth event locations are a realization of a Poisson process, and the live birth event locations are an independent realization of a Poisson point process.

Detailed description of our proposed analytic approach begins in Section 2 with a description of our motivating data set and methodologic approach. Herein, we detail the Bayesian approach to estimation of the model. In Section 3, we describe our simulation technique and compare the two Bayesian estimation methods for the conditional formulation of the spatial and spatio-temporal models with respect to the recovery of true parameter values. In Section 4, we describe the application of the methodology to our stillbirth surveillance data, and in Section 5, we conclude our presentation by summarizing our findings, comparing our findings to those initially obtained from the point process approach, and discussing several remaining methodologic issues.

2. Methods

2.1. Spatial Model

Stillbirths for the years 2005 through 2011 are in the form of a set of n event locations $\left\{\mathbf{s}_i = \begin{pmatrix} \text{Latitude}_i \\ \text{Longitude}_i \end{pmatrix} : i = 1, \ldots, n\right\}$ within Iowa, a polygonal region denoted as **W**. This set of events represents the geocoded maternal residence at delivery for each event within **W** and are expected to exhibit substantial spatial variation in intensity. Additionally, some mothers experienced a stillbirth in more than one pregnancy within **W** during the study period. In the initial point process approach to modeling maternal residence, we adopted a multiplicative model (Gatrell et al., 1996) for the intensity function $\lambda(\mathbf{s})$ of the form

$$\lambda(\mathbf{s}) = \rho\lambda_0(\mathbf{s})\lambda_1(\mathbf{s}|\boldsymbol{\theta}). \quad (1)$$

The scaling parameter ρ represented the ratio of the number of stillbirth to live birth events; this quantity, assigned in advance as all occurrences of stillbirths and live births within **W** during the study period were recorded, was incorporated in the analysis. In model (1), $\lambda_0(\mathbf{s})$ represented the background intensity (i.e. the number of pregnancies at risk per unit area in the neighborhood of the location **s**) and $\lambda_1(\mathbf{s}|\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is a vector of parameters, represented the possible increase in risk as a function of **s**.

In this paper, we assumed that the stillbirth event locations $\mathbf{s}_i : i = 1, \ldots, n$ and the live birth event locations $\mathbf{s}_i : i = n + 1, \ldots, N$, where $N = n + m$, the total number of events, were independent realizations of Poisson processes, with their respective intensities governing the processes $\lambda(\mathbf{s})$ and $\lambda_0(\mathbf{s})$. The superposition of the two point processes is also a Poisson point process (Diggle and Rowlingson, 1994). Following the work of Diggle and Rowlingson (1994), we defined a binary random variable Y to take the value 1 or 0 according to whether the i^{th} event in the superposition was an event of the first or the second element of the process. Conditioning on the joint realization of these processes, it is straightforward to algebraically show that the conditional probability of a stillbirth event at any location is $\mathbf{Pr}(y_i = 1) = \frac{\rho \cdot \lambda_1(\mathbf{s}_i|\boldsymbol{\theta})}{1 + \rho \cdot \lambda_1(\mathbf{s}_i|\boldsymbol{\theta})} = p_i$. In deriving the conditional probability, the nuisance background intensity $\lambda_0(\mathbf{s})$ is conveniently eliminated from the model. This formulation leads to a spatial logistic regression model where a linear predictor, including a contextual effect that captures spatially correlated heterogeneity $w(\mathbf{s}_i)$, is assumed within $\lambda_1(\mathbf{s}_i|\boldsymbol{\theta})$. For example, a log linear formulation for $\lambda_1(\mathbf{s}_i|\boldsymbol{\theta})$ leads to a logit link to p_i, i.e. $p_i = \frac{\exp(\eta_i)}{1 + \exp(\eta_i)}$, where $\eta_i = \mathbf{x}^T(\mathbf{s}_i)\boldsymbol{\beta} + w(\mathbf{s}_i) + \gamma_{j_i \in j}$ and $\gamma_{j_i \in j}$, where j denotes the mother, represents a maternal contextual effect. Let $\beta_0 = \log(\rho)$ denote the intercept and serves the role of the scientifically uninteresting constant background rate; that is, it is unnecessary to also include the constant ρ with the above parameterization. The resulting Bernoulli likelihood is then given by

$$L(\theta|\mathbf{s}) = \prod_{i=1}^{N} \left[\frac{\{\exp(\eta_i)\}^{y_i}}{1+\exp(\eta_i)} \right].$$

2.2. Spatio-Temporal Model

Because a stillbirth event was observed with a delivery date, it is possible to extend the conditional formulation by considering spatio-temporal effects. Specifically, we observed within study region **W** and a seven-year time period T, a set of n stillbirth events, with maternal residence at delivery given as $\{\mathbf{s}_i\}$, $i = 1, \ldots, n$, and also time labels $\{t_i\}$, $i = 1, \ldots, n$. Here, the random variables were the spatial location and the timing of delivery. Note that mothers may have experienced multiple independent events, and the spatial location may have changed from one event to the next. In this paper, we readily extended model (1) and considered a multiplicative model for the intensity function $\lambda(\mathbf{s}, t)$ that now represents the variation across space and time of the intensity of stillbirth events, of the form

$$\lambda(\mathbf{s}, t) = \rho\lambda_0(\mathbf{s}, t)\lambda_1(\mathbf{s}, t|\theta). \quad (2)$$

Using a similar argument as with the spatial model, if the m live birth event locations are a realization of a spatio-temporal Poisson point process on $\mathbf{W} \times T$ with $\lambda_0(\mathbf{s}, t)$ and the stillbirth event locations are a realization of an independent spatio-temporal Poisson point process with $\lambda(\mathbf{s}, t)$, then the superposition of the two point processes is also a spatio-temporal Poisson point process provided that the two point processes are separable in space and time (Lawson, 2013). Conditioning on this joint realization, the binary labeling of these $n + m$ events form a set of mutually independent Bernoulli random variables with spatio-temporal dependent probabilities $\mathbf{Pr}(y_i = 1) = \frac{\rho \cdot \lambda_1(\mathbf{s}_i, t_i|\theta)}{1 + \rho \cdot \lambda_1(\mathbf{s}_i, t_i|\theta)}$. The quantity ρ is a constant background rate now in space × time units. Our conditioning converts the statistical model to a linear binary regression model. This conversion avoids the problem of estimating $\lambda_0(\mathbf{s}, t)$, which describes the spatio-temporal at-risk background population, and allows a straightforward extension of the multiplicative decomposition of model (1) to incorporate parameters relating to spatial, temporal, and spatio-temporal components considered germane to the application. Specifically, a log linear formulation for $\lambda_1(\mathbf{s}, t|\boldsymbol{\theta})$ leads to a logit link for p_i, where now we can define $\eta_i = \mathbf{x}^T(\mathbf{s}_i, t_i)\boldsymbol{\beta} + w(\mathbf{s}_i) + g(t_i) + c(\mathbf{s}_i, t_i) + \gamma_{j_i \in j}$ and $\gamma_{j_i \in j}$, where j denotes the mother, represents a maternal contextual effect as before. Here, we consider discrete time labels corresponding to the number of days from January 1, 2005 and we assume a separable covariance structure in space and time, where $w(\mathbf{s}_i)$, $g(t_i)$, and $c(\mathbf{s}_i, t_i)$ represent a spatially correlated term, temporally correlated term, and uncorrelated spatio-temporal interaction term, respectively. This spatio-temporal interaction term allows for overdispersion.

2.3. Bayesian Estimation

A Bayesian hierarchical model was used for fitting models (1) and (2) to our stillbirth surveillance data. The focus in our analysis was to make inference about the coefficients $\boldsymbol{\beta}$

allowing for the unobserved confounders $w(\mathbf{s}_i)$, $g(t_i)$, $c(\mathbf{s}_i, t_i)$, and $\gamma_{j_i \in j}$, as well as obtaining a quantitative description of variation in the local intensity of stillbirth events within the spatial and temporal region. INLA (version 0.0–1468872408) was used to carry out the fitting of the conditional formulation of the spatial and spatio-temporal models to our stillbirth surveillance data and to obtain posterior quantities of the parameters $\theta = \{\beta, \sigma_w^2, \phi, \sigma_g^2, \sigma_c^2, \sigma_\gamma^2\}$. The parameters σ_w^2, σ_c^2, and σ_γ^2 are the respective variances associated with the spatially correlated heterogeneity term **w**(**s**), the uncorrelated space-time component **c**(**s**, *t*), and the maternal contextual effect γ, respectively. With regard to the temporal dependence, in this descriptive analysis a first-order autoregressive time component was assumed because it provided a parsimonious and intuitively appealing mechanism for describing temporal dependence; therefore, the parameters ϕ and σ_g^2 corresponded to the time dependent parameter and the variance associated with the white noise, respectively. The deviance information criterion (DIC) - appropriate for model comparison in complex hierarchical models, such as these spatial and spatio-temporal models applied to the stillbirth data - was used to assess model adequacy and to compare the models. When comparing models using the DIC measure, Spiegelhalter et al. (2002) considered a difference in DIC of 2–3 and greater as meaningful; in our application, a difference in DIC greater than 3 was used to ascertain if the DIC was exhibiting a preference.

2.4. Prior Distributions

We adopted an intrinsic conditional autoregressive (CAR) prior distribution for the spatially correlated heterogeneity **w**(**s**) (Besag and Mollie, 1991) given as

$$w_i(\mathbf{s}_i) | w_j(\mathbf{s}_j), j \neq i, n_{\delta_i}, \sigma_w^2 \sim N\left(\sum_{j \in \delta_i} \frac{w_j(\mathbf{s}_j)}{n_{\delta_i}}, \frac{\sigma_w^2}{n_{\delta_i}} \right),$$

where n_{δ_i} was the total number of first-order neighbors in the j^{th} area (i.e. the regions which share common geographical boundaries with the i^{th} region) and δ_i was the first-order neighborhood of the i^{th} region. Following Lawson (2012), the neighborhood relation assumed between event locations was based on a Dirichlet tessellation for a point process where the tiling of the locations leads to sets of natural neighbors defined by the adjoining edges of the tile. In other words, two locations were defined as neighbors if they shared a common border when the Dirichlet tessellation was used. Rather than apply some arbitrary means for defining neighboring points (e.g. a distance threshold), the Dirichlet tessellation was appealing because it has the remarkable mathematical property that all locations within a tile are closer to the tile point than to any other point (Rogers, 1964). And although a fully specified multivariate normal prior distribution could have been considered for the correlated component **w**(**s**), an intrinsic Gaussian prior distribution was adopted because it provided an attractive means for handling potentially complicated joint spatial dependencies, which were modeled simply through a collection of conditional dependencies, and was computationally advantageous. The prior distribution for each β was set to a normal distribution with mean 0 and variance 1000. The maternal contextual effects γ were specified as spatially

uncorrelated effects so that $\gamma_{j_i \in j} \sim N\left(0, \sigma_\gamma^2\right)$, where i denotes the event and j denotes the mother. A first-order autoregressive model was assumed for the temporally correlated effect $\mathbf{g}(t)$ component, i.e. $\mathbf{g}(t+1) = \boldsymbol{\phi} \cdot \mathbf{g}(t) + \boldsymbol{\eta}(t)$, where $\boldsymbol{\eta}(t)$ are independent and identically distributed $N\left(0, \sigma_g^2\right)$. The uncorrelated space-time component $\mathbf{c}(\mathbf{s}, t)$, which is a residual effect, was given a $N\left(0, \sigma_c^2\right)$ prior distribution.

With INLA, modestly informed priors on the hyperparameters for random effects were shown to be needed in a suite of simulation studies comparing INLA with MCMC sampling results using OpenBUGS in Bayesian disease mapping (Carroll et al., 2015). Therefore, a modestly vague gamma prior (1, 1) was placed on the inverse of the variance components σ_w^2, σ_γ^2, and σ_c^2. Similarly, modestly vague priors were defined for the temporal component. Specificially, a modestly vague loggamma prior (1, 0.30) was placed on the natural logarithm of the marginal precision parameter $\frac{1}{\sigma_m^2}$, where $\sigma_m^2 = \frac{1}{\sigma_g^2 \cdot \left(1 - \phi^2\right)}$; additionally, a modestly vague normal prior $N(0, 0.35)$ was assumed for the $\log\left(\frac{1+\phi}{1-\phi}\right)$, corresponding to a transformation for the time dependent parameter $\boldsymbol{\phi}$. Choosing more informative priors for these hyperparameters than the modestly vague priors detailed here did not significantly impact any of the results; however, and as was reported by Carroll et al. (2015), applying increasingly more vague priors drastically impacted the results with INLA.

3. Simulation Studies

To support the use of INLA, as opposed to MCMC sampling, as a reliable algorithm to estimate posterior quantities for the spatial and spatio-temporal models we are proposing, simulation studies, defined by three general parameterizations of the intensity function $\lambda(\mathbf{s})$, were conducted. The goal of the simulation studies was not to validate the nuanced parameterization for the intensity function $\lambda(\mathbf{s})$ applied to and estimated from our stillbirth data. Rather, our goal for carrying out these simulation studies was to compare the two Bayesian estimation methods INLA and MCMC sampling for the spatial and spatio-temporal models in a slightly more general setting, albeit while maintaining some similarity with the parameterization for $\lambda(\mathbf{s})$ applied to our stillbirth data. Achieving similar results between the two Bayesian estimation methods in such a setting would provide the desired support for the use of INLA in our data application.

The m control points were simulated first, assuming complete spatial randomness. Next, we separately constructed a point pattern of n case events, which proceeded in two steps, beginning with generating $\mathbf{w}(\mathbf{s})$ with zero mean from a conditional specification of a Gaussian Markov random field (GMRF) and then, generating locations given $\mathbf{w}(\mathbf{s})$ along with non-stationary mean $\boldsymbol{\mu}(\mathbf{s})$. This construction is a LGCP. Each realized point pattern was simulated from a LGCP. Conditional on $\lambda(\mathbf{s})$, we have a nonhomogeneous Poisson point process (NHPP). Following Besag and Mollie (1991), the conditional specification of the GMRF was given as

$$w_i(\mathbf{s}_i) \mid w_j(\mathbf{s}_j), j \neq i, n_{\delta_i}, \sigma_w^2 \sim N\left(\alpha \sum_{j \in \delta_i} \frac{w_j(\mathbf{s}_j)}{n_{\delta_i}}, \frac{\sigma_w^2}{n_{\delta_i}}\right),$$

where α was a correlation parameter to ensure a proper stochastic mechanism and n_{δ_i} was the total number of first-order neighbors in the j^{th} area. The first-order neighborhood was found from a Dirichlet tessellation for a point process. We set $\alpha = 0.95$, corresponding to strong spatial correlation.

The intensity function $\lambda(\mathbf{s})$ was considered to reflect etiologic heterogeneity in occurrence of stillbirth, e.g. due to heterogeneous risk of experiencing a stillbirth based on some unobserved risk factors. Conceivably, this heterogeneity could be considered a random quantity that changes; therefore, one reason for simulating a point pattern in this doubly stochastic fashion was to assess the effect of the additional randomness induced by the inclusion of $\mathbf{w}(\mathbf{s})$ on our model that was parameterized to approximate the latent GMRF. This was accomplished by evaluating the reasonableness of the intrinsic CAR prior specification, defined earlier, as an approximation to the latent GMRF, as measured by σ_w^2. An intrinsic CAR prior was used because the expected posterior estimate of the correlation parameter α from a proper weighted CAR prior was consistently approaching one. Following Banerjee et al. (2015), within a Bayesian framework, a prior on α that encourages a consequential amount of spatial correlation would place most of its mass near one anyway. In any event, we also wanted to assess the recovery of the parameter β_1, i.e. the coefficient ascribed to a spatially-referenced covariate driving the point pattern in the spatial model (1) as well as the recovery of all model parameters in the spatio-temporal model (2), separable in space and time.

In our simulation studies, we used the bounded state of Iowa for region $\mathbf{W}$ and $T = [1, 60]$, treated as an indexing set $\{1, 2, \ldots, 60\}$. With our real data reflecting the study time period of 2005 through 2011, 60 discrete time points corresponded to an evolving stillbirth event map every 6 weeks. Furthermore, we considered the setting where the ratio of case to control events was one (i.e. $\rho = 1$ or equivalently $\log(\rho) = \beta_0 = 0$), $\lambda_0(\mathbf{s}) = \lambda_0$, a constant, and $n = 1000$; a realization of size $n = 1000$ corresponds roughly to the complete enumeration of 1,195 stillbirth events observed during the study period. Applying $\rho = 0.004$, as was observed in our stillbirth data, and preserving a point pattern of $n = 1000$ would require $m = 250{,}000$ simulated control points. With our goal to compare INLA and MCMC sampling in a slightly more general setting, we justified setting $\rho = 1$ so that both Bayesian methods of estimation could be implemented easily. For an event location $\mathbf{s}$ and time t, the log linear model parameterizations considered to address our goals were

$$\text{I. } \log \lambda_1(\mathbf{s} \mid \theta) = \beta_0 + w(\mathbf{s}).$$

$$\text{II. } \log \lambda_1(\mathbf{s} \mid \theta) = \beta_0 + \beta_1 x(\mathbf{s}) + w(\mathbf{s}).$$

$$\text{III. } \log \lambda_1(\mathbf{s}, t|\boldsymbol{\theta}) = \beta_0 + \beta_1 x(\mathbf{s}) + w(\mathbf{s}) + g(t).$$

Intensity parameterizations I and II, respectively, allow us to assess the reasonableness of an intrinsic CAR prior specificiation as an approximation to a latent GMRF and to assess the recovery of the true parameter β_1 in the presence of unobserved spatial variation for both Bayesian estimation methods in the spatial-only setting. For both Bayesian estimation methods, to evaluate the recovery of all true parameters used in generating the spatio-temporal point process, separable in space and time, we used intensity parameterization III.

The standardized spatially-referenced covariate **x**(**s**) represented a simulated point-level covariate, which was not time-varying, but exhibited strong spatial correlation as might be found in epidemiologic applications; a Gaussian random variable with mean zero and powered exponential covariance structure was assumed. The covariate was generated using the RandomFields package in R (Schlather et al., 2015) and then standardized. We specified the covariance structure by using the RMpoweredexp command where the variance was set to 1, the range to 1, the nugget to 0, and the power to 1.5. In intensity parameterization III, we assumed a discretized version of the random fields so that any realization of the field $\{(\mathbf{s}_i, t_j)\}$ had a separable correlation structure. Within this assumption, the likelihood remained that of a conditionally modulated Poisson process. The temporal component **g**(t) was a first-order autoregressive model. Unlike the model applied to our stillbirth surveillance data, our carefully controlled simulated experiments obviated the need to include a residual effect corresponding to the uncorrelated space-time residual component c(**s**, t) in model (2). See Appendices A.1 and A.2, respectively, for the simulation algorithm used for the spatial and spatio-temporal models.

3.1. Simulation Results

In each simulated experiment, we generated 100 realizations. For each realization, posterior quantiles were estimated from the Bayesian spatial and spatio-temporal models using INLA (version 0.0–1468872408) and MCMC sampling. The total number of iterations used in MCMC sampling was 750,000 with the first 250,000 treated as burn-in. To decrease autocorrelation, samples were thinned, using only every 50^{th} step in the sampler. The simulation studies were implemented in R using the INLA package and R2OpenBUGS. Tabular summaries were used to display the average measure of error (i.e. the bias) and the corresponding standard deviation.

3.1.1. Intensity Parameterization I: $\log\lambda_1(\mathbf{s}|\boldsymbol{\theta}) = \beta_0 + w(\mathbf{s})$—As part of our goal for conducting simulation studies, we assessed the reasonableness of an intrinsic CAR prior specification as an approximation to a latent GMRF assuming the conditional formulation of the spatial model. We considered modest ($\sigma_w = 0.708$), large ($\sigma_w = 1.225$), and very large ($\sigma_w = 2.5$) unobserved variation for the latent GMRF. Because we assumed the ratio of case to control events was one (i.e. $\rho = 1$), the corresponding true value of β_0 (i.e. $\log(\rho)$) used in simulating the point patterns was zero. For both Bayesian estimation methods, β_0 was well estimated and the average measure of error associated with β_0 was consistently negligible (data not shown). The top portion of Table 1 displays the bias of the estimates $\hat{\sigma}_w$ comparing

INLA with MCMC sampling. Both estimation methods underestimated the true value of σ_w used in simulating the point patterns. Compared to INLA, however, the average measure of error was consistently smaller with MCMC sampling. Although the average measure of error increased for both estimation methods in the presence of incrementally larger unobserved spatial variation, both estimation methods arguably performed reasonably well in the presence of modest unobserved spatial variation. With a point pattern of size $n = 1000$, the respective average runtimes for INLA and MCMC sampling were approximately 4 and 3290 seconds.

3.1.2. Intensity Parameterization II: $\log\lambda_1(s|\theta) = \beta_0 + \beta_1 x(s) + w(s)$—

As part of our goal for conducting simulation studies, we also assessed the recovery of the true parameter β_1 in the presence of modest unobserved spatial variation. The average measure of error for three effect sizes are shown in the bottom portion of Table 1 assuming a point pattern of size $n = 1000$. The average measure of error was similar and arguably negligible for both INLA and MCMC estimation methods.

3.1.3. Intensity Parameterization III: $\log\lambda_1(s, t|\theta) = \beta_0 + \beta_1 x(s) + w(s) + g(t)$—

We assessed the recovery of all true parameters used in generating the point process, separable in space and time, using INLA and MCMC sampling to fit the spatio-temporal model. Table 2 shows the average measure of error for each of the model components when assuming the true value of the time dependent parameter ϕ for the first-order autoregressive distribution was either 0.5 or 0.9; note, $\phi = 1$ corresponded to a random walk and $\phi = 0$ represented no time dependency. For $\phi \in \{0.5,0.9\}$, on average, the bias was negligible and comparable between the two estimation methods for the parameter β_1 and for the standard deviation σ_g associated with the white noise of the first-order autoregressive time series. On average, the time dependent parameter ϕ was somewhat underestimated, and perhaps slightly more so in the presence of stronger temporal dependency, for both estimation methods. Although the average measure of error associated with ϕ was not that different between the two estimation methods, MCMC sampling consistently outperformed INLA. For both estimation methods, the spatial variance component (σ_w) was underestimated by a similar magnitude as seen with Intensity Parameterization I. The average runtimes for INLA and MCMC sampling with a point pattern of size $n = 1000$ were approximately 10 and 7,901 seconds, respectively.

3.2. Simulation Conclusions

In the conditional formulation of the spatial and spatio-temporal models, applicable to a more general setting than with the nuanced parameterization applied to our stillbirth data, INLA gave reasonable results compared to MCMC, particularly for the epidemiologically interesting parameter β_1. Furthermore, results from the simulation study using INLA to estimate the spatio-temporal model suggested that the time dependent parameter ϕ was adequately recovered and not that different from MCMC sampling. If our primary interest was in estimating the unobserved spatial variation, then MCMC would be preferred; furthermore, the researcher could consider less vague priors on the precision hyperparameters with MCMC sampling. Nonetheless, σ_w was negatively biased for both methods which we believe was likely due to the inability of the conditional formulation of

the models to capture the additional randomness associated with the intensity surface. With the focus of our epidemiologic application being on the vector of $\boldsymbol{\beta}$ coefficients, adjusting for the spatially correlated heterogeneity effects $\mathbf{w}(\mathbf{s})$, and because MCMC sampling applied to our large stillbirth data set would be very slow based on our simulation studies, we proceeded with INLA in estimating the spatial and spatio-temporal models in the data application and not MCMC sampling.

4. Data Application

The IRCID began actively monitoring stillbirth deliveries statewide in 2005 (Romitti, 2015). For the years 2005–2011, the maternal residence at delivery (event location) of each stillbirth and live birth delivered to an Iowa resident was geocoded. There were 1,195 stillbirth deliveries from independent pregnancies enumerated by the IRCID and 271,791 live births recorded during the seven-year study period; pregnancies with multiple fetuses were excluded, because of differing risks for birth weight, gestational age, and fetal growth among multiple compared to singleton pregnancies. There were 1,167 mothers who experienced a stillbirth in one pregnancy, and 14 mothers who experienced a stillbirth in two independent pregnancies.

The event locations of all live births in Iowa, during the study period from families who experienced at least one stillbirth during the study period, were included in the set of control event locations. This corresponded to 1,150 control event locations for which all had covariate information. A complete set of covariate information was available for 270,323 control event locations from families who did not experience a stillbirth during the study period. These 270,323 control event locations corresponded to 195,502 unique families. Because of the considerable memory requirements needed to process the neighborhood relation between event locations, defined based on the Dirichlet tesselation, and to fit the spatial-temporal models, 50,000 (25%) control families were randomly selected from the 195,502 unique families. Our analysis included 71,316 events (1,195 stillbirth events, including 1,150 sibling controls; 68,971 remaining controls), which corresponded to 51,181 families. We repeated our analyses on three randomly sampled data sets of similar size, and the results remained robust for each data set (data not shown). The spatial distribution of stillbirth event locations did not differ appreciably from the distribution of the event locations for the at-risk population, namely, maternal residence at delivery for each live birth (Zahrieh et al., 2018).

Our surveillance data were observed with a time label, defined as the number of days from January 1, 2005, and a spatial location, namely, the maternal residence at the time of a stillbirth or live birth event. Figure 1 displays the successive number of stillbirths to the number of live births per month during January 1, 2005 through December 31, 2011. Although there was considerable variability, following a slight increase in the ratio through 2008, the ratio appears to modestly decrease.

We fit several Bayesian spatial and spatio-temporal logistic regression models for the binary outcome (stillbirth, live birth), where the probability was a function of space-only or space and time (Table 3). The DIC expressed a strong preference for the spatio-temporal model

that included the time varying covariate maternal age at the time of a delivery, maternal race/ethnicity, and two zip code tabulation area (ZCTA) covariates: percentage of childbearing women with less than a bachelor's degree and median income, both obtained from the 2007–2011 American Community Survey data. Notably, we observed a strong association between maternal race/ethnicity and urbanized ZCTAs of Iowa, as defined by the 2010 U.S. census (67% of mothers categorized as other race/ethnicity were estimated to live in urban areas compared with 45% of non-Hispanic whites) and based on the DIC, we selected the parsimonious model excluding the urban versus rural indicator. Table 4 displays the estimated posterior quantities obtained from fitting our selected Bayesian spatio-temporal logistic model. With the exception of median income, the expected posterior estimates and the corresponding 95% credible intervals suggested that the parameter estimates were important; in particular, the 95% credible interval for the time dependent parameter $\boldsymbol{\phi}$ excluded zero (mean: 0.8517; 95% credible interval: 0.4604, 0.9938). The estimated runtime to fit the selected spatio-temporal logistic model with INLA was 143,282 seconds.

From our final model that included the point-level covariates maternal age at delivery and maternal race/ethnicity, as well as the regional-level (ZCTA) covariates percent of childbearing women with less than a bachelor's degree and median income, and controlled for the live birth events, the posterior expected estimates for the spatial correlation component, temporal correlation component, the spatio-temporal residual component, and the time series plot indexed by the number of days from January 1, 2005 are displayed in Figure 2 for the 1,195 stillbirth events. The Bayesian stillbirth event map for the spatial correlation component (Figure 2a) suggests that there are peaks in the spatial component in the south-central and south-east areas of Iowa, as well as in a south-easterly direction from Des Moines; we single out this latter area because this observation was evidenced when fitting the point process model. The Bayesian stillbirth event map for the temporal correlation component (Figure 2b) indicates temporal variations with marked changes in several pockets across the state. Based on the Bayesian map of the space-time residual component (Figure 2c), there is considerable extra variation remaining not accounted for by the separable space and time components. Lastly, the time series plot of the temporal correlation component (Figure 2d) does not indicate any apparent cyclical behavior over the study period and mirrors the observed trend depicted in Figure 1.

Identifying where stillbirth prevalence exceeds a certain relative risk threshold over time can be more useful than reporting posterior quantities from fitting the model. Therefore, to assess localized spatio-temporal behavior of the model and the assessment of unusual aggregation of stillbirth events over time, heat-contour maps of relative risk within the time intervals 2005–2006, 2007–2008, and 2009–2011 were investigated. For the 1,195 stillbirth events, Figure 3 displays the marginal posterior expectation results for the spatio-temporal model that included the covariates and controlled for the live births, where the spatio-temporal relative risk was defined as $\lambda_1(\mathbf{s}_i, t_i|\boldsymbol{\theta}) = \exp\left(\mathbf{x}^T(\mathbf{s}_i, t_i)\boldsymbol{\beta} + w(\mathbf{s}_i) + g(t_i) + c(\mathbf{s}_i, t_i) + \gamma_{j_i \in j}\right)$. In other words, the posterior estimates of spatio-temporal relative risk included all model components of risk excluding the intercept. The estimated mean of the posterior distribution of the model parameters was used in obtaining localized relative risk at the three temporal increments over the study

period. The cities with a population >50,000 residents are superimposed and presented as solid green diamonds. Edge effects were present, so care needs to be exercised in interpreting the heat-contour maps along the boundary of the spatial region. Nonetheless, apparent changes were seen with this spatio-temporal presentation of the Bayesian relative risk maps. Notably, the Bayesian relative risk maps indicated increasing risk in the west south-west areas of Iowa and decreasing risk in the north and south-easterly direction from the central counties of Iowa, over the birth period.

5. Discussion

In our epidemiologic setting where it can reasonably be assumed that the stillbirth and live birth event locations arise from independent Poisson point processes, the conditional formulation of the point process model still allowed us to capture the salient features of our stillbirth surveillance data while quantifying localized geographic regions of high relative risk. Moreover, the conditional approach greatly simplified the analysis and interpretation compared with modeling the maternal residence via a point process model. As opposed to the point process modeling approach applied initially, the conditional formulation was easily extended to include temporal effects and allowed for the inclusion of covariate information attached to both stillbirth and live birth; therefore, we were now able to quantify geographic regions of excess stillbirth risk after adjusting for our set of covariates, and both spatial and temporal effects. Although we were no longer modeling the spatial distribution of event locations, we can still, to some extent, assume that the data arose from a LGCP where the intensity of the process is governed by a GMRF. That is, conditional on the intensity, the data are a Poisson point process and then conditional on the locations we showed that we can reasonably account for unobserved spatially correlated heterogeneity assuming an intrinsic CAR specification within a relatively simple Bayesian spatial logistic regression model estimated with INLA.

In our data application, we added the time of a stillbirth event to the event location to facilitate our understanding of how stillbirth events were geographically linked within Iowa during the study period. Although we did not model the spatio-temporal distributions of the event locations directly, we can pragmatically assume that the data arose from a spatio-temporal LGCP where the intensity of the process was governed by discretized space-time random fields. Our general simulation study demonstrated the similarity between the conditional formulation of the spatio-temporal model and a spatio-temporal LGCP. In particular, in the presence of modest spatial variation associated with the GMRF, the conditional formulation of the spatio-temporal model estimated with INLA was sensitive to modest and strong temporal dependence assuming a first-order autoregressive model.

The argument for using the INLA R package to estimate the Bayesian spatial and spatio-temporal models was twofold. First, INLA provided a faster and reasonably accurate alternative to MCMC sampling for posterior parameter estimation. Ideally, results using INLA should be close to the MCMC approach to estimation, which we observed in our simulation studies applicable to a more general setting for the recovery of the parameter β_1, the time dependent parameter ϕ, and the standard deviation σ_g associated with the white noise, but less so for the recovery of the variance σ_w^2 associated with the spatially correlated

heterogeneity effects **w**(**s**). Despite parameterizing the intensity function to represent a form of a LGCP, the underestimation of σ_w^2 was expected because the heterogeneity due to variability in the intensity surface is not captured in the likelihood function for the conditional formulation of the point process. INLA was particularly useful for our application where conventional MCMC sampler programs would be extremely slow. Second, it offered a flexible model specification to capture the salient features of our stillbirth surveillance data and was applicable to a log-linear Gaussian model. The main disadvantage of using INLA was the need to use a less vague prior on the precision hyperparameters. In the current application, the focus of the analysis was on the vector of $\boldsymbol{\beta}$ coefficients, adjusting for the confounding variables **w**(**s**), γ, and **g**(t). If the goal, however, was to model and draw inference on the unobserved spatial variation, the data analyst would need to be prepared to carry out sensitivity analyses.

The parameterization of the modeled excess risk component in the multiplicative formulation of the intensity function flexibly permitted inclusion of model components that captured important and nuanced features of the application, such as a maternal contextual effect. Additionally, the spatio-temporal model allowed us to obtain a quantitative description of variation in local intensity of stillbirth events in space and time. Although our focus was on relative risk estimation rather than cluster detection, localized areas of excess aggregation of stillbirth events over time were quantified based on a host of important features captured by our model parameterization and identified for further investigation. There was some agreement as well as some differences between the results obtained from the conditional approach and the findings previously obtained from the point process approach. The mapped regions of high levels of spatially correlated heterogeneity were qualitatively similar to the mapped regions obtained from the point process model applied to these data, and, notably, neither map indicated a random scatter of areas of high levels. Although maternal age at the time of a stillbirth delivery was not shown to be associated with the spatial distribution of stillbirth after applying the point process model, it was predictive in the conditional approach that modeled the conditional probability of stillbirth; the latter finding was consistent with a recent systematic review and meta-analysis from 14 case-control studies that showed advanced maternal age increases the odds of stillbirth (Lean et al., 2017). Also, after controlling for the at-risk live birth intensity in the point process model, maternal residence in urban locations was strongly associated with the spatial distribution of stillbirth; however, the presence of race/ethnicity in the conditional model obviated the need for the urban/rural factor.

There were several limitations of our methodologic approach. We adopted a CAR prior for the spatially correlated heterogeneity where the neighborhood relation between event locations was based on a Dirichlet tessellation for a point process. This defensible approach for defining neighbors resulted in a connected graph (or collection of nodes and contiguous edges) that required high memory requirements to process; coupled with the rather high memory requirements needed to fit the spatio-temporal model we, therefore, based our inference on a random sample of 50,000 control families or 25% of the 195,502 unique families. However, the analysis was repeated on 3 randomly sampled data sets of similar size and the conclusions were unchanged. As this was a descriptive analysis, edge effects were

ignored. Also, for environmental risk assessment where continuous risk fields may be affecting the at-risk population, a range of appropriate and suspected spatio-temporal covariates are needed to quantitatively describe excess risk. An autoregressive prior distribution was used, assuming a first-order autoregressive model. Although this choice for a prior distribution allowed for a linear (i.e. with respect to the previous value) non-parametric temporal effect, alternative formulations could be considered for the temporal component. Lastly, further work is needed to validate the final model applied to our stillbirth data.

There were also some notable study design limitations. Although the maternal residence at delivery was used to represent exposure to environmental risk, albeit treated as a contextual effect in our model formulation, it ignored the possibility that exposure to environmental risks may occur elsewhere. Also, the residence at delivery may not represent the residence at conception or during early pregnancy. In addition, unavailable fetal and maternal risk factors not included in our model limits the interpretation of our results. Future follow-up of mothers within the spatial and temporal geographic regions of excess risk compared to non-risk may provide insights into these unmeasured environmental and social factors. Lastly, a longer study time period is needed to better characterize spatio-temporal changes in relative risk.

Conclusions

Using a conditional approach to modeling the geocoded stillbirth and live birth data, we quantified and mapped the excess stillbirth risk in the presence of spatial and temporal heterogeneity and after adjusting for covariates attached to both stillbirth and live birth. Our model was fitted with INLA, as opposed to MCMC sampling, with reasonable accuracy, and INLA accommodated our large data set. Furthermore, our use of the conditional formulation was readily extended to include temporal effects. To our knowledge, our study is the first to conduct a formal spatio-temporal analysis of stillbirth surveillance data. Although the temporal correlation component indicated temporal variations with marked changes in several areas across the state, the residual space-time component indicated that there was extra variation remaining not captured by the separable space-time components and covariates available for analysis. Nonetheless, our Bayesian relative risk maps indicated increasing and decreasing risk over the birth period, which may warrant further public health investigation in the geographic regions identified.

Acknowledgements

This research was supported in part through computational resources provided by The Department of Biostatistics at The University of Iowa. This research was also supported by funding from the Iowa Stillbirth Surveillance Project (5U50DD000730), Iowa CBDRP: Birth Defects Study To Evaluate Pregnancy Exposures (5U01DD001035), and The University of Iowa Environmental Health Sciences Research Center (P30 ES005605). The authors gratefully thank Dong Liang, Anthony Rhoads, and Derek Shen for their assistance with data cleaning, data preparation, and data management.

Appendix A.: Simulation Technique

A.1. Spatial Model

First, we simulated m control points assuming complete spatial randomness. The m geographic (longitude, latitude) pairs were generated uniformly in **W** using rejection sampling. Next, we independently generated a point pattern of n case events. Constructing a point pattern proceeded in two steps, beginning with generating **w**(**s**) with mean zero from a conditional specificiation of a Gaussian Markov random field (GMRF) with a first-order neighborhood structure based on a Dirichlet tessellation for a point process and then, generating locations given **w**(**s**) along with non-stationary mean $\boldsymbol{\mu}(\mathbf{s})$. This construction is a log Gaussian Cox process. Conditional on $\lambda(\mathbf{s})$, we have a nonhomogeneous Poisson point process (NHPP). For the given conditional specification of the GMRF and nonstationary mean $\boldsymbol{\mu}(\mathbf{s})$, our algorithm to sample one realization of n case events, therefore, was as follows:

1. For large N, simulate $\mathbf{s} = \{\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_N\}$ assuming complete spatial random ness within **W**
2. Generate **w**(**s**) a GMRF with mean zero
3. Calculate $\lambda(\mathbf{s}) = \exp(\boldsymbol{\mu}(\mathbf{s}) + \mathbf{w}(\mathbf{s}))$
4. Define $\lambda_{max} = MAX(\lambda(\mathbf{s}))$
5. Thin the simulated process as follows
    a. Randomly draw a point **s** with replacement
    b. Generate a random number u from the uniform distribution (0, 1)
    c. If $\frac{\lambda_{sampled}}{\lambda_{max}} > u$, accept the point **s**
    d. Repeat (a)-(c) until the desired number of points ($n = 1000$) are generated

The N geographic (longitude, latitude) pairs were generated uniformly in **W** using rejection sampling. Complete spatial randomness in step 1 of the algorithm was achieved by generating geographic pairs uniformly on the enclosing rectangle [−96.60641, −90.13772] × [40.37634, 43.51041]. The approximate area of the polygonal boundary **W** is 56,025.5 square miles and the approximate area of the enclosing rectangle is 73,029.64 square miles.

The case event locations were a realization of a Poisson point process on **W**, with intensity $\lambda(\mathbf{s})$ and the control event locations were a realization of an independent Poisson point process with constant intensity $\lambda_0(\mathbf{s}) \equiv \lambda_0$. Conditional on the $n + m = 2n$ locations $\mathbf{s}_i$, which we call *events*, the *labels* of these events were a set of mutually independent Bernoulli random variables; that is, we associated with each event location a binary variable (y_i) which labels the event either as a case ($y_i = 1$) or a control ($y_i = 0$). Conditioning on the joint realization of these point processes, the conditional probability of a case at any location is

$$Pr(y_i = 1) = \frac{\lambda(\mathbf{s}_i|\theta)}{1 + \lambda(\mathbf{s}_i|\theta)}.$$

A.2. Spatio-Temporal Model

We assumed a separable covariance function in space and time (i.e. an additive form in spatial and temporal effects) without a space-time residual. First, we simulated m control points assuming complete spatial randomness, and for each location, we uniformly simulated discrete time [1, 60]. Next, we independently generated a point pattern of n case events from a log Gaussian Cox process in space and time using a global majorizing function within a rejection algorithm. Conditional on $\lambda(\mathbf{s}, t)$, we had a spatio-temporal NHPP for the case event locations within the spatial and temporal region.

For the given conditional specification of the Gaussian Markov random field (GMRF) and nonstationary mean $\boldsymbol{\mu}(\mathbf{s})$, as well as a first order autoregressive time series $\mathbf{g}(t)$, our algorithm to sample a realization of n case events was as follows:

1. For large N, simulate $\mathbf{s} = \{\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_N\}$ assuming complete spatial randomness within $\mathbf{W}$
2. For each of the N locations, uniformly simulate discrete time [1, 60] with replacement
3. Generate a first order autoregressive time series $\mathbf{g}(t)$ for $t \in [1, 60]$ and merge with the data set of size N by t
4. Generate $\mathbf{w}(\mathbf{s})$ a GMRF with mean zero
5. Calculate $\lambda(\mathbf{s}) = \exp(\boldsymbol{\mu}(\mathbf{s}) + \mathbf{w}(\mathbf{s}) + \mathbf{g}(t))$
6. Define $\lambda_{max} = MAX(\lambda(\mathbf{s}))$
7. Thin the simulated process as follows
    a. Randomly draw a point $\mathbf{s}$ with replacement
    b. Generate a random number u from the uniform distribution (0, 1)
    c. If $\frac{\lambda_{sampled}}{\lambda_{max}} > u$, accept the point $\mathbf{s}$
    d. Repeat (a)-(c) until the desired number of points ($n = 1000$) are generated

The case event locations were a realization of a spatio-temporal Poisson point process on $\mathbf{W} \times T$, with intensity $\lambda(\mathbf{s}, t)$ and the control event locations were a realization of an independent spatio-temporal Poisson point process with constant intensity $\lambda_0(\mathbf{s}, t) \equiv \lambda_0$. Conditional on the $n + m = 2n$ locations $(\mathbf{s}_i, t_i)$, which we call *events*, the *labels* of these events were a set of mutually independent Bernoulli random variables; that is, we associated with each event location a binary variable (y_i) which labeled the event either as a case ($y_i =$

1) or a control ($y_i = 0$). Conditioning on the joint realization of these point processes, the conditional probability of a case at any event location was

$$Pr(y_i = 1) = \frac{\lambda(\mathbf{s}_i, t_i|\boldsymbol{\theta})}{1 + \lambda(\mathbf{s}_i, t_i|\boldsymbol{\theta})}.$$

References

Banerjee S, Carlin BP, and Gelfand AE (2015). Hierarchhical modeling and analysis for spatial data. Chapman & Hall-CRC, second edition.

Besag J, Y. J and Mollie A (1991). Bayesian image restoration with two applications in spatial statistics. Annals of the Institute of Statistical Mathematics 43, 1–59.

Carroll R, Lawson AB, Faes C, Kirby RS, Aregay M, and Watjou K (2015). Comparing inla and openbugs for hierarchical poisson modeling in disease mapping. Spatial and Spatio-temporal Epidemiology 14–15, 45–54.

DeFranco E, Hall E, Hossain M, Chen A, Haynes EN, Jones D, Ren S, Lu L, and Muglia L (2015). Air pollution and stillbirth risk: Exposure to airborne particulate matter during pregnancy is associated with fetal death. PLoS One 10, e0120594. [PubMed: 25794052]

Diggle P and Rowlingson B (1994). A conditional approach to point process modelling of elevated risk. Journal of the Royal Statistical Society 157, 433–440.

Flenady V, Koopmans L, Middleton P, Froen JF, Smith GC, Gibbons K, Coory M, Gordon A, Ellwood D, McIntyre HD, Fretts R, and Ezzati M (2011). Major risk factors for stillbirth in high-income countries: a systematic review and meta-analysis. The Lancet 377, 1331–1340.

Gatrell AC, Bailey TC, Diggle PJ, and Rowlingson B (1996). Spatial point pattern analysis and its application in geographical epidemiology. The Royal Geographical Society 21, 256–274.

Rue H, M. S and Chopin N (2009). Approximate bayesian inference for latent gaussian models using integrated nested laplace approximations (with discussion). Journal of the Royal Statistical Society, Series B 71, 319–392.

Hall ES, Connolly N, Jones DE, and DeFranco EA (2014). Integrating public health data sets for analysis of maternal airborne environmental exposures and stillbirth. The American Medical Informatics Association Annual Symposium Proceedings 2014, 599–605.

Lawson AB (2012). Bayesian point event modeling in spatial and environmental epidemiology. Statistical Methods in Medical Research 21, 509–529. [PubMed: 23035034]

Lawson AB (2013). Bayesian Disease Mapping Hierarchical Modeling in Spatial Epidemiology. Chapman & Hall-CRC.

Lean SC, Derricott H, Jones RL, and Heazell AEP (2017). Advanced maternal age and adverse pregnancy outcomes: A systematic review and meta-analysis. PLoS One 12, e0186287. [PubMed: 29040334]

Lindgren F, Rue H, and Lindstrom 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 73, 423–498.

Martins TG, Simpson D, Lindgren F, and Rue H (2013). Bayesian computing with inla: new features. Computational Statistics and Data Analysis 67, 68–83.

NICHD (2017). Stillbirth collaborative research network.

Rogers CA (1964). Packing and Covering. London: Cambridge University Press.

Romitti PA (2015). Stillbirth surveillance consortium In BMC Pregnancy and Childbirth, volume 15 Proceedings of the Stillbirth Summit 2014.

Schlather M, Malinowski A, Menck PJ, Oesting M, and Strokorb K (2015). Analysis, simulation and prediction of multivariate random fields with package randomfields. Journal of Statistical Software 63, 1–25.

Simpson D, Lindgren F, and Rue H (2012a). In order to make spatial statistics computationally feasible, we need to forget about the covariance function. Environmetrics 23, 65–74.

Simpson D, Lindgren F, and Rue H (2012b). Think continuous: Markovian gaussian models in spatial statistics. Spatial Statistics 1, 16–29.

Spiegelhalter DJ, Best NG, Carlin BP, and van der Linde A (2002). Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society, Series B 64, 583639.

Zahrieh D, Oleson JJ, and Romitti PA (2018). Bayesian point process modeling to quantify geographic regions of excess stillbirth risk. Geographical Analysis.

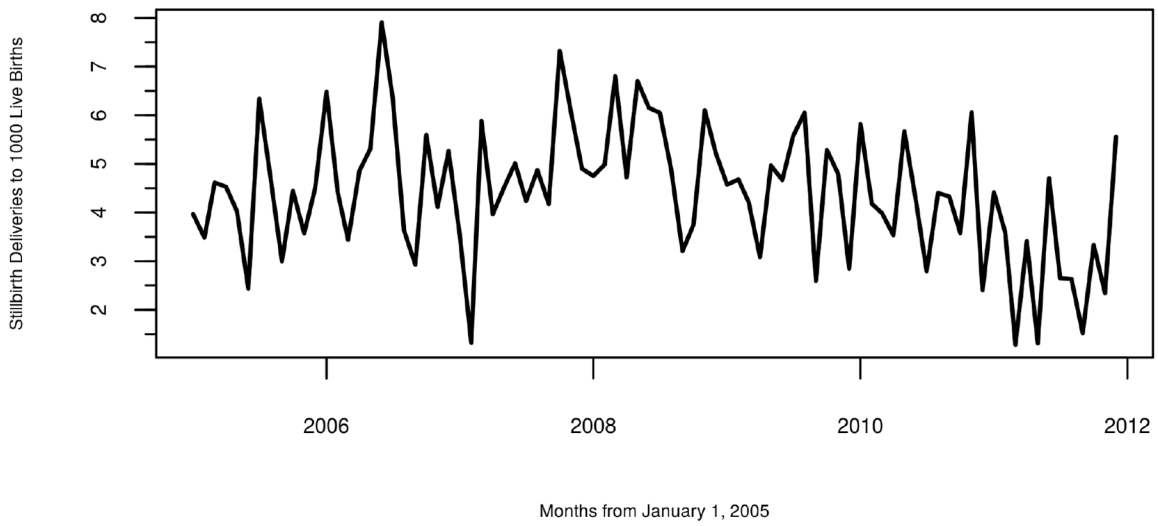

Figure 1:
Ratio of stillbirth deliveries to 1000 live births: Months since January 1, 2005 of successive stillbirth deliveries per 1000 live births between January 1, 2005 and December 31, 2011.

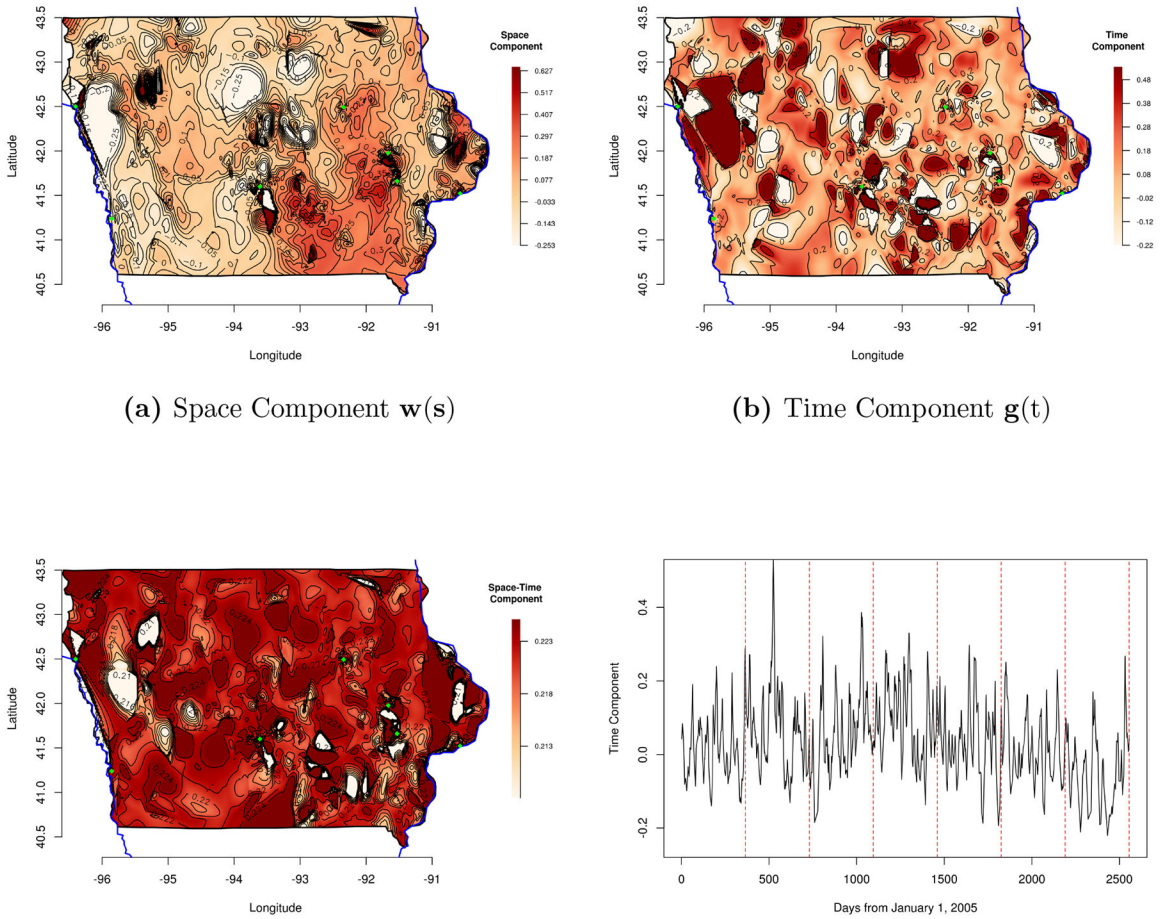

(a) Space Component $\mathbf{w}(\mathbf{s})$ **(b)** Time Component $\mathbf{g}(\mathrm{t})$

(c) Space-Time Residual Component $\mathbf{c}(\mathbf{s}, \mathrm{t})$ **(d)** Time Component $\mathbf{g}(\mathrm{t})$: Time Series

Figure 2:

Bayesian disease maps from the final model: For the 1,195 stillbirth deliveries, the posterior expected estimates from the Bayesian spatio-temporal logistic model fit to the stillbirth surveillance data: four displays correponding to the spatial component (a), a temporal correlation component (b), a space-time residual component (c), and a time component time series plot indexed by the number of days from January 1, 2005 (d). The red dotted lines in (d) indicate one-year intervals. Estimated posterior quantities were obtained from the integrated nested Laplace approximation.

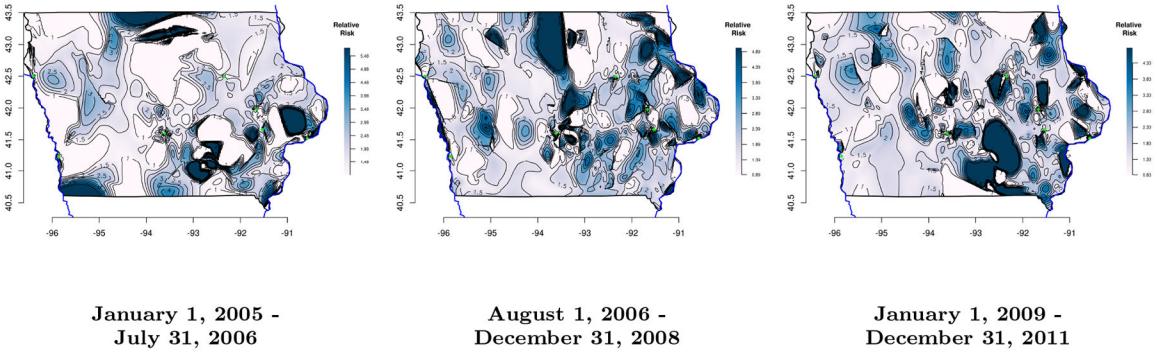

Figure 3:

Bayesian disease maps of relative risk from the final model: Bayesian relative risk maps, presented as heat-contour plots, based on the estimated posterior expectations of all model components from the spatio-temporal Bayesian hierarchical model that controlled for the live births, mapped at the maternal residence for the 1,195 stillbirth deliveries. Estimated posterior quantities were obtained from the integrated nested Laplace approximation.

Table 1:

Average estimated mean of the posterior distribution (standard deviation) and average measure of error associated with the conditional approach to estimation via INLA and MCMC for intensity parameterizations I and II based on 100 realizations of point patterns of size $n = 1000$.

Intensity Parameterization I: log $\lambda_1(s|\theta) = \beta_0 + w(s)$

		True Value (σ_w)			
		0.708	**1.225**	**2.5**	**Runtime**
INLA	**Mean (SD)**	0.5495 (0.0541)	0.6620 (0.0827)	1.1702 (0.1612)	2.8 sec
	Ave$(\hat{\sigma}_w - \sigma_w)$	−0.1585	−0.5630	−1.3300	
MCMC	**Mean (SD)**	0.6290 (0.0809)	0.7707 (0.1400)	1.708 (0.2948)	3290 sec
	Ave$(\hat{\sigma}_w - \sigma_w)$	−0.0790	−0.4543	−0.7920	

Intensity Parameterization II: log $\lambda_1(s|\theta) = \beta_0 + \beta_1 x(s) + w(s)$

		True Value (β_1)		
		−0.50	**−0.25**	**0.00**
INLA	**Mean (SD)**	−0.4896 (0.0650)	−0.2411 (0.0548)	0.0051 (0.0520)
	Ave$\left(\hat{\beta}_1 - \beta_1\right)$	0.0104	0.0089	0.0051
MCMC	**Mean (SD)**	−0.4908 (0.0614)	−0.2352 (0.0622)	−0.0048 (0.0566)
	Ave$\left(\hat{\beta}_1 - \beta_1\right)$	0.0092	0.0148	−0.0048

Note: INLA = integrated nested Laplace approximation; MCMC = Markov chain Monte Carlo; SD = standard deviation. Modest unobserved spatial variation (i.e. $\sigma_w = 0.708$ or $\sigma_w^2 = 0.5$) was assumed for parameterization II. An intrinsic conditional autoregressive prior distribution was adopted for the spatially correlated heterogeneity **w(s)** and a modestly vague gamma prior (1, 1) was placed on the inverse of the variance component σ_w^2. The prior distribution for β_1 was set to a normal distribution with mean 0 and variance 1000.

Table 2:

Average measure of error (standard deviation) associated with the conditional approach to estimation via INLA and MCMC for intensity parameterization III.

Intensity Parameterization III: $\log\lambda_1(s|\theta) = \beta_0 + \beta_1 x(s) + w(s) + g(t)$

True Value				Average Measure of Error (Standard Deviation)			
		ϕ	β_1	$\mathbf{Ave}\left(\hat{\beta}_1 - \beta_1\right)$	$\mathbf{Ave}(\hat{\sigma}_w - \sigma_w)$	$\mathbf{Ave}\left(\hat{\phi} - \phi\right)$	$\mathbf{Ave}\left(\hat{\sigma}_g - \sigma_g\right)$
INLA							
$\sigma_w = 0.708$	$\sigma_g = 0.35$	0.50	−0.50	0.0094 (0.0644)	−0.1796 (0.0433)	−0.0548 (0.2421)	−0.0002 (0.0505)
		0.50	−0.25	0.0075 (0.0565)	−0.1836 (0.0435)	−0.0490 (0.2161)	−0.0030 (0.0600)
		0.50	0.00	−0.0053 (0.0549)	−0.1904 (0.0362)	−0.0607 (0.2395)	−0.0010 (0.0596)
$\sigma_w = 0.708$	$\sigma_g = 0.35$	0.90	−0.50	0.0213 (0.0640)	−0.1740 (0.0481)	−0.0717 (0.1040)	−0.0057 (0.0672)
		0.90	−0.25	0.0054 (0.0576)	−0.1902 (0.0436)	−0.0683 (0.0916)	−0.0095 (0.0646)
		0.90	0.00	0.0046 (0.0591)	−0.1823 (0.0444)	−0.0810 (0.1036)	−0.0118 (0.0692)
MCMC							
$\sigma_w = 0.708$	$\sigma_g = 0.35$	0.50	−0.50	0.0107 (0.0696)	−0.0975 (0.0740)	−0.0098 (0.2428)	−0.0122 (0.0655)
		0.50	−0.25	−0.0013 (0.0614)	−0.1095 (0.0675)	−0.0284 (0.1845)	−0.0031 (0.0676)
		0.50	0.00	−0.0040 (0.0633)	−0.1006 (0.0768)	−0.0160 (0.1963)	−0.0035 (0.0595)
$\sigma_w = 0.708$	$\sigma_g = 0.35$	0.90	−0.50	0.0071 (0.0659)	−0.1068 (0.0625)	−0.0532 (0.0934)	0.0009 (0.0667)
		0.90	−0.25	0.0052 (0.0555)	−0.0959 (0.0652)	−0.0484 (0.0808)	0.0000 (0.0616)
		0.90	0.00	0.0102 (0.0553)	−0.1074 (0.0828)	−0.0478 (0.0758)	0.0015 (0.0652)

Note: INLA = integrated nested Laplace approximation; MCMC = Markov chain Monte Carlo. Results are based on 100 realizations of point patterns of size $n = 1000$. Time $t \in [1, 60]$, was treated as an indexing set {1, 2, …, 60}. An intrinsic conditional autoregressive prior distribution was adopted for the spatially correlated heterogeneity **w**(**s**) and a modestly vague gamma prior (1, 1) was placed on the inverse of the variance component σ_w^2. The prior distribution for β_1 was set to a normal distribution with mean 0 and variance 1000. The standard deviation associated with the white noise of the first-order autoregressive time series was set to $\sigma_g = 0.35$, which corresponded to a marginal standard deviation of $\sigma_m = 0.4041$ and $\sigma_m = 0.8030$ when ϕ equals 0.5 and 0.9, respectively. A modestly vague loggamma prior (1, 0.30) was placed on the natural logarithm of the marginal precision parameter $1/\sigma_m^2$, where $\sigma_m^2 = 1/\left[\sigma_g^2 \cdot \left(1 - \phi^2\right)\right]$ and a modestly vague normal prior $N(0, 0.35)$ was assumed for the log $[(1 + \phi)/(1 - \phi)]$, corresponding to a transformation for the time dependent parameter ϕ.

Table 3:

Model preference using the deviance information criterion (DIC).

Model	Covariates	DIC	Runtime
$w(\mathbf{s}) + \gamma$	None	12068.05	845 sec
$\mathbf{x}^T(\mathbf{s})\boldsymbol{\beta} + w(s) + \gamma$	Included	11991.31	722 sec
$w(\mathbf{s}) + g(t) + c(\mathbf{s},t) + \gamma$	None	11996.36	80958 sec
$\mathbf{x}^T(\mathbf{s},t)\boldsymbol{\beta} + w(\mathbf{s}) + g(t) + c(\mathbf{s},t) + \gamma$	Included + Urban vs. Rural Indicator	11926.05	136289 sec
$\mathbf{x}^T(\mathbf{s},t)\boldsymbol{\beta} + w(\mathbf{s}) + g(t) + c(\mathbf{s},t) + \gamma$	Included	11922.89	143282 sec

Note: The included vector of covariates $\mathbf{x}^T$ comprised maternal age at delivery, an indicator for maternal race/ethnicity, and the ZCTA-level covariates percentage of childbearing women with less than a bachelor's degree and median income; continuous covariates were centered and standardized. The percent of child-bearing women with less than a bachelor's degree and median income were calculated for each ZCTA from the 2007–2011 American Community Survey data. A ZCTA was designated as urban if it intersected with the urbanized areas of Iowa, as defined by the 2010 census (US Census Bureau) or as rural if it did not intersect with urbanized areas; urbanized areas are defined by the census as having a population of >50,000 residents with a density of at least 500 people per square mile. The model parameters $w(\mathbf{s})$, γ, $g(t)$, and $c(\mathbf{s}, t)$ represented spatially correlated heterogeneity, a maternal contextual effect, a temporally correlated term, and an uncorrelated spatio-temporal interaction term, respectively. 50,000 (25%) control families were randomly selected from the 195,502 unique families who did not experience a stillbirth during the study period; therefore, the analysis population included 71,316 events (1,195 stillbirth events, including 1,150 sibling controls; 68,971 remaining controls), which corresponded to 51,181 families. The model shown in the last row of the table was selected as the final model.

Table 4:

Estimated posterior quantities from fitting the final model.

Description of Explanatory Variable		Mean (SD)	95% Credible Interval	
Intercept	β_0	−4.5520 (0.0401)	−4.6312	−4.4738
Point-Level Covariates				
Maternal Age (Years)	β_1	0.0692 (0.0292)	0.0117	0.1265
Other Races/Ethnicities versus non-Hispanic White Indicator	β_2	0.6733 (0.0712)	0.5327	0.8122
Regional-Level (ZCTA) Covariates				
Percent of Childbearing Women with Less than a Bachelor's Degree	β_3	0.0973 (0.0392)	0.0201	0.1739
Median Income	β_4	0.0172 (0.0384)	−0.0578	0.0929
Random Effects				
Spatial Component	$\frac{1}{\sigma_w^2}$	4.4863 (1.4223)	2.4824	7.9881
Temporal Component	ϕ	0.8517 (0.1432)	0.4604	0.9938
	$\frac{1}{\sigma_m^2}$	13.9404 (6.0656)	6.8327	29.7302
Space-Time Residual Component	$\frac{1}{\sigma_c^2}$	4.6876 (1.6492)	2.2746	8.6778
Maternal Contextual Effect	$\frac{1}{\sigma_\gamma^2}$	5.3071 (1.7304)	2.8016	9.5144

Note: All continuous covariates were centered and standardized. Estimates of posterior quantities were obtained from the INLA package. The percent of child-bearing women with less than a Bachelor's degree and median income were calculated for each ZCTA from the 2007–2011 American Community Survey data. Other race/ethnicities included unknown race/ethnicity. Precisions are presented for the random effects. The corresponding mean of the estimated posterior distribution for σ_w, σ_m, σ_c, and σ_γ were 0.4721, 0.2678, 0.4619, and 0.4341, respectively.