Published in final edited form as:

Stat Med. 2008 May 30; 27(12): 2127–2144. doi:10.1002/sim.3141.

Parametric models for spatially correlated survival data for individuals with multiple cancers

Ulysses Diva¹, Dipak K. Dey², and Sudipto Banerjee^{3,*},†,‡

¹ Global Biometric Sciences, Bristol-Myers Squibb Company, Wallingford, CT, U.S.A

² Department of Statistics, University of Connecticut, Storrs, CT, U.S.A

³ Department of Biostatistics, University of Minnesota, Minneapolis, MN, U.S.A

SUMMARY

Incorporating spatial variation could potentially enhance information coming from survival data. In addition, simultaneous (joint) modeling of time-to-event data from different diseases, such as cancers, from the same patient could provide useful insights as to how these diseases behave together. This paper proposes Bayesian hierarchical survival models for capturing spatial correlations within the proportional hazards (PH) and proportional odds (PO) frameworks. Parametric (Weibull for the PH and log-logistic for the PO) models were used for the baseline distribution while spatial correlation is introduced in the form of county–cancer-level frailties. We illustrate with data from the Surveillance Epidemiology and End Results database of the National Cancer Institute on patients in Iowa diagnosed with multiple gastrointestinal cancers. Model checking and comparison among competing models were performed and some implementation issues were presented. We recommend the use of the spatial PH model for this data set.

Keywords

Bayesian hierarchical models; frailty models; Markov chain Monte Carlo (MCMC); proportional hazards; proportional odds; spatial association; survival modeling

1. INTRODUCTION

The importance of accounting for underlying spatial correlation in modeling data with geographic information has been recognized in several areas (see, e.g. [1–4]). Accounting for spatial correlation could provide insights that would have been overlooked otherwise. On the other hand, the failure to include such information could potentially lead to spurious or misleading results.

In biostatistics and epidemiology, there has been a growing interest in modeling survival data accounting for spatial associations. Li and Ryan [5] and Banerjee *et al.* [6] addressed this problem under a proportional hazards (PH) framework from classical and Bayesian perspectives, respectively. Semi-parametric approaches were proposed by Banerjee and Dey [7] using a proportional odds (PO) structure on the survival times, and by Li and Lin [8] using normal transformations of the survival times that marginally follow the PH model.

*Correspondence to: Sudipto Banerjee, School of Public Health, University of Minnesota, A460 Mayo Building, MMC 303, 420 Delaware Street, SE, Minneapolis, MN 55455, U.S.A.

†E-mail: sudiptob@biostat.umn.edu

‡Associate Professor of Biostatistics.

As opposed to modeling disease incidence and mortality, survival modeling provides a slightly different perspective with regard to the nature of the disease. It focuses upon how many are expected to survive after a certain period of time, how fast is the rate of failure, and what drives shortened or prolonged survival; all of these may be influenced by several factors such as gender, race, age, type of cancer, treatment obtained, and access to healthcare facilities.

Among the widely investigated diseases are the different types of cancers. The SEER (Surveillance Epidemiology and End Results) database of the National Cancer Institute provides a rich amount of data on patients from different registries possibly suffering from multiple types of cancers. Consequently, survival models for *multiple cancers* are sought for practical reasons, such as to assess survival from multiple primary cancers simultaneously and to adjust the survival rates from a specific primary cancer in the presence of other primary cancers (see, e.g. [9,10]). Recently, Carlin and Banerjee [11] implemented spatial survival models, where each patient developed possibly several cancers. They based their modeling upon the first primary cancer incorporating the effects of subsequent cancers (if any) on that individual as binary regressors indicating the presence or absence, but ignored the time-to-event information available for any subsequent primary cancers.

On the other hand, Diva *et al.* [12] (henceforth, DBD) proposed semi-parametric PH models that directly incorporated information on survival times from subsequent primary cancers. For modeling the spatial associations, both studies employed multivariate conditionally autoregressive (MCAR) model, which is, in fact, a multi-layered Markov random field (MRF) (see, e.g. [13,14]). In this approach, each type of cancer is given its own spatial distribution over an underlying map, but are also spatially correlated among themselves, so that the cancer effects nested within space have a joint distribution that follows a multi-layered MRF or, more specifically, an MCAR distribution. While we adopt the multivariate approach employed by DBD to handle survival data from multiple cancers, here we consider parametric specifications of the PH and PO models. In general, parametric specifications are simpler to implement and faster to execute, can produce better model performance when the specified parametric form is close to the observed pattern, and can provide a better framework to distinguish between models. Although the nature of the data sets used in DBD and this paper is similar, here we actually used different search criteria, which correspond to a different target population and lead to slightly different model interpretations. For purposes of comparing the performance of the parametric and semi-parametric implementations, we also applied our model to the DBD data although the results are not presented in detail.

We present our modeling approach in Section 2. Technical and numerical aspects of the Bayesian implementation of the models are described in Section 3. We illustrate and evaluate our proposed models in Section 4 with data on cancers of six (6) gastrointestinal (GI) organs. Finally, Section 5 concludes the paper with a summary and some recommendations.

2. PARAMETRIC MODELS FOR MULTIPLE CANCERS

2.1. Data on patients with multiple cancers

Let us consider a relatively simple version of the SEER data, where we extracted survival information on patients suffering from multiple cancers. Along with survival information, patient-specific and patient-cancer combination-specific covariates are also available. Patient-specific information includes gender, race, marital status, number of primary cancers diagnosed, and their county of residence. Patient-cancer-specific information relates to age at diagnosis, the stage of the cancer (*in situ*, local, distant or regional), the type (radiation and/or surgery) of treatment that each patient underwent for each type of the cancer, and so on. As an example, consider a patient, say with ID 00171, who first contracted cancer of the small intestine (the first primary cancer) in 1986, developed colon cancer in 1991, pancreatic cancer

in 1995, and eventually died in 1998. We present the data for this particular patient in Table I, incorporating information about the disease progression. We extract information about this patient from the SEER database, where each entry corresponding to the patient would list the type of cancer, the age at diagnosis, the stage of the cancer, the status at end point (dead or alive), and the survival time (or censorship time) in months. Note that the status is labeled as 'Dead' for all entries, as it corresponds to the status at the end point.

The succeeding subsections describe approaches in modeling survival time with the data setup described above. While there are other viable approaches that may be used we decided to concentrate on the parametric PH and PO models for various reasons. Firstly, it should be noted that the two models can be considered direct opposites in the assumptions that they employ [15]. Secondly, it is straightforward to incorporate spatial and non-spatial correlations in the corresponding regression models. Although the use of a parametric specification in the baseline distributions could greatly limit the flexibility of the resulting models, it lends ease of interpretation for the model parameters and reduces the computational burden considerably. Finally, and perhaps most importantly, exploratory data analysis techniques described in Section 4 suggest that the parametric models may be appropriate. The Bayesian modeling framework will be used for its flexibility and facility in handling complex hierarchical models.

2.2. PH model

The PH model [16] is one of the most commonly used survival models that incorporates covariate effects. It is given by

$$h(t|\mathbf{x})=h_0(t)e^{\mathbf{x}^T\boldsymbol{\beta}} \quad (1)$$

where $h(t)$ denotes the hazard function, $\mathbf{x}$ are the covariates, and $\boldsymbol{\beta}$ is the corresponding vector of regression coefficients. Also, $h_0(t)$ denotes a baseline hazard function, which is often modeled semi-parametrically (see, e.g. [12,17–19]). However, in this paper we specified a Weibull(λ_1, λ_2) baseline hazard function, which gives $h_0(t) = \lambda_1\lambda_2t^{\lambda_1-1}$. The parametric setup is equivalent to specifying a Weibull regression model for the logarithm of the survival time [20]. That is

$$\log(t)=\mu+\mathbf{x}^T\boldsymbol{\beta}^*+\gamma Y \quad (2)$$

with $\mu=-(1/\lambda_1)\log(\lambda_2)$, $\gamma=1/\lambda_1$, $\boldsymbol{\beta}^* = -(1/\lambda_1)\boldsymbol{\beta}$, and Y follows the extreme value distribution with density $f_Y(y) = \exp(y - e^y)$, $-\infty < y < \infty$. This gives a natural interpretation of λ_1 and λ_2 to be the scale and intercept parameters, respectively. The above setup also has an accelerated failure-time interpretation, with acceleration factor given by $e^{\mathbf{x}^T\boldsymbol{\beta}^*}$.

The term $e^{\mathbf{x}^T\boldsymbol{\beta}}$ in (1) is interpreted as the multiplicative change in the hazard function from baseline due to the covariates. As with any regression setting care must be given in setting up the covariates so that the value $\mathbf{x}=\mathbf{0}$ is meaningful. The PH model assumes that the hazard ratio (relative risk) between two levels of a covariate, $\mathbf{X}$, is constant over time, i.e. $h(t|\mathbf{X})/h(t|\mathbf{X}^*) = e^{(\mathbf{X}-\mathbf{X}^*)^T\boldsymbol{\beta}}$, while the odds ratio approaches either 0 or $+\infty$.

2.3. PO model

Another model that could be used for multiple cancer data incorporating covariate effects is the PO model, popularized by Bennett [21]. The model is given by

$$\frac{S(t|\mathbf{x})}{1 - S(t|\mathbf{x})} = \frac{S_0(t)}{1 - S_0(t)} e^{\mathbf{x}^T \boldsymbol{\beta}} \quad (3)$$

where $S(t|\mathbf{x})$ is the survival function evaluated at t conditional on the observed covariates, $\mathbf{x}$, and $\boldsymbol{\beta}$ is the corresponding vector of regression coefficients. The baseline survival function, $S_0(t)$, is oftentimes modeled semi-parametrically (see, e.g. [7,22]). However, in this paper, we modeled the baseline distribution as log-logistic(λ_1, λ_2) so that $S_0(t) = 1/(1+\lambda_2 t^{\lambda_1})$. Similar to the PH model with Weibull baseline, this formulation of the PO model has a corresponding linear regression model in log-time given by

$$\log(t) = \mu + \mathbf{x}^T \boldsymbol{\beta}^* + \gamma \tilde{Y} \quad (4)$$

with $\mu = -(1/\lambda_1)\log(\lambda_2)$, $\gamma = 1/\lambda_1$, $\boldsymbol{\beta}^* = -(1/\lambda_1)\boldsymbol{\beta}$, $\tilde{Y}$ follows the logistic distribution with and density $f_{\tilde{Y}}(y) = e^{-y}/[1+e^{-y}]^2$, $-\infty < y < \infty$ [20].

The PO model may be considered to be on the opposite end of the spectrum as the PH model since it assumes that the hazard ratio approaches unity over time, i.e. that the covariate effects on the hazards disappear over time. In addition, although the estimates of the parameters may be similar between the two models, the interpretation of the regression component significantly differs. The term $e^{\mathbf{x}^T \boldsymbol{\beta}}$ in the PO model is interpreted as the change in the odds of survival (or failure, depending on the parametrization), given the observed covariates or risk factors.

2.4. Capturing correlation using frailties

Suppose that there are I counties and in the i th county we observe n_i patients suffering from at least two of K cancer types of interest. We uniquely identify each of the $N = \sum_{i=1}^I n_i$ patients through the ordered pair (i, j) , as the j th patient ($j = 1, 2, \dots, n_i$) from the i th county. Let t_{ijk} denote the survival time, the time to death, or end of study (censoring) for the (i, j) th patient from diagnosis with the k th type of cancer. Since not all patients develop all the K types of cancer, we list the possible cancers as $\{1, 2, \dots, K\}$ and form the subset $C_{(i,j)} \subseteq \{1, 2, \dots, K\}$ as the indices of the cancers developed by the (i, j) th patient. Thus, if patient 00171 in our example above is the 10th patient from the seventh county and if small intestine, colon, and pancreas are cancer types 2, 3, and 6, respectively, in $\{1, 2, \dots, K\}$ then $C_{(7,10)} = \{2, 3, 6\}$. Hence, when referring to t_{ijk} , $k \in C_{(i,j)}$ unless otherwise noted. Clearly, these survival times will be correlated for similar county-patient-cancer combinations. We capture these correlations by introducing appropriate frailties.

Let $u_{(i,j)}$ denote the frailty for the (i, j) th patient, v_k be the frailty for the k th cancer type, and φ_{ik} be the frailty for the k th cancer type, nested within the i th county. For purposes of discussion, let us assume the patient frailties to be independent and normally distributed (zero-centered) with county-specific variances, i.e. $u_{(i,j)} \stackrel{\text{indep}}{\sim} N(0, \sigma_i^2)$. However, in the implementation of the models, we did not include the patient frailties for this particular data set as convergence problems were encountered, probably due to the scarcity and imbalance of information at the patient-county level. Let us then collect the K cancer frailties into a vector $\mathbf{v}$ and assume a joint zero-centered normal distribution with unknown covariance matrix $\boldsymbol{\Lambda}$. That is, $\mathbf{v} = (v_1, \dots, v_K)^T \sim N(\mathbf{0}, \boldsymbol{\Lambda})$, where $\boldsymbol{\Lambda}$ is a $K \times K$ covariance matrix.

The spatial frailties, $\{\varphi_{ik}\}$, are assigned an MCAR(α, Λ) distribution by first forming the $I \times K$ vector $\Phi = (\varphi_1^T, \dots, \varphi_I^T)^T$, where each $\Phi_i = (\varphi_{i1}, \dots, \varphi_{iK})^T$ is the $K \times 1$ vector of the spatial frailties for the K cancers within the i th county, and assigning a joint zero-centered normal distribution with a covariance matrix that accounts for the spatial structure of the region. That is, by $\{\varphi_{ik}\} \sim \text{MCAR}(\alpha, \Lambda)$, we mean

$$\Phi = (\varphi_1^T, \dots, \varphi_I^T)^T \sim N(\mathbf{0}, \sum_w (\alpha) \otimes \Lambda)$$

where $\Sigma_W(\alpha) = (\text{Diag}(m_i) - \alpha \mathbf{W})^{-1}$, m_i is the number of neighbors of the i th county, and $\mathbf{W}$ is the adjacency matrix of the *graph* representing our region. It should be noted that only α , the spatial smoothness parameter, is random in $\Sigma_W(\alpha)$. Carlin and Banerjee [11] have shown that as long as $\alpha \in (0, 1)$, we have a proper MCAR distribution. In this setup, we are essentially using the same covariance matrix for the cancer main effects and the spatial frailties within a particular county. This is accomplished by using a Kronecker product of spatial structure and ‘cancer dispersion’, an approach for modeling interactions that has been suggested earlier in non-spatial contexts by Clayton [23]. Several generalizations may be used to further enrich the models, such as using a different cancer covariance matrix $\mathbf{T}$ resulting in MCAR($\alpha, \mathbf{T}$), using different spatial smoothness parameters for each cancer type resulting in MCAR($(\alpha_1, \dots, \alpha_K), \Lambda$) [11, 14], and having $\mathbf{v}_i \sim N(\mathbf{0}, \Lambda_i)$ (not identically distributed) where each county has its own cancer–covariance pattern. However, the above generalizations of the MCAR(α, Λ) require substantial information on multiple cancer patients, which, unfortunately, is rare in practice. For instance, in the SEER database there is considerable imbalance of information at the county–cancer level, with some counties having more cases of a particular cancer and no observed case of the other cancers, rendering many of the more general models unidentifiable. Therefore, we restrict our subsequent attention to the MCAR(α, Λ) specification and compare with non-spatial models. For more information about MCAR models and possible generalizations the reader is referred to [11, 13, 14, 24].

3. NUMERICAL IMPLEMENTATION

3.1. PH model

Considering the kind of data described in Section 2.1, a parametric PH model with Weibull (λ_1, λ_2) baseline distribution incorporating covariate effects, main effects for patients and cancers and nested effects for cancers within counties, may be expressed as

$$h(t_{ijk} | \mathbf{x}_{ijk}, u_{(i,j)}, v_k, \varphi_{ik}) = \lambda_1 \lambda_2 t_{ijk}^{\lambda_1 - 1} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}} \quad (5)$$

for $i = 1, 2, \dots, I$, $j = 1, 2, \dots, n_i$, and $k = 1, 2, \dots, K$. Using the above specification and letting $\delta_{(i,j)}$ be a death indicator (status at end point) for the (i, j) th patient, we obtain the likelihood

$$\begin{aligned} L(\boldsymbol{\beta}, \lambda_1, \lambda_2, \{u_{(i,j)}\}, \{v_k\}, \{\varphi_{ik}\}; \{t_{ijk}\}, \{\delta_{(i,j)}\}, \mathbf{x}_{ijk}) \\ = \prod_{i=1}^I \prod_{j=1}^{n_i} \prod_{k \in C(i,j)} [S(t_{ijk} | \mathbf{x}_{ijk}, u_{(i,j)}, v_k, \varphi_{ik}, \boldsymbol{\beta}, \lambda_1, \lambda_2)] [\lambda_1 \lambda_2 t_{ijk}^{\lambda_1 - 1} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}}]^{\delta_{(i,j)}} \end{aligned} \quad (6)$$

where $S(t_{ijk}|\mathbf{x}_{ijk}, u_{(i,j)}, v_k, \varphi_{ik}, \boldsymbol{\beta}, \lambda_1, \lambda_2) = \exp[-\lambda_2 t_{ijk}^{\lambda_1} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}}]$ is the survival function evaluated at t_{ijk} conditional on the parameters and observed covariates.

To complete the Bayesian specification of the PH model, we need to assign prior distributions for the parameters. We assign vague normal, say $\boldsymbol{\beta} \sim N(\mathbf{0}, 10^5 \mathbb{I}_{p \times p})$, or improper flat prior, $f(\boldsymbol{\beta}) \propto 1$, for the regression coefficients, a $\text{Gamma}(s_0, s_1)$ prior for the Weibull scale parameter λ_1 , and a log-normal(μ_0, τ^2) for the Weibull shape parameter λ_2 , where s_0, s_1, μ_0 , and τ^2 are hyper-parameters specified to obtain vague prior distributions. We also need to specify priors for the parameters of the frailty distributions presented in Section 2.4. We assign inverted-Gamma, $\text{IG}(a_i, b_i)$, priors for the σ_i^2 's, an inverted-Wishart prior, $\text{IW}(r_0, \Lambda_0)$, for the covariance matrix $\boldsymbol{\lambda}$, and $\text{U}(0, 1)$ or a Beta prior for the spatial smoothness parameter α , where the hyper-parameters are chosen to obtain vague prior distributions.

The parameters of the baseline hazard function, λ_1 and λ_2 , and the regression parameters, $\boldsymbol{\beta}$ (with either flat or vague Gaussian priors), all have full conditional distributions that are log-concave and can be updated using the adaptive rejection sampling (ARS) algorithm [25] or a Metropolis step. The same applies to the patient frailties $\{u_{(i,j)}\}$, the cancer frailties $\{v_k\}$, and the spatial frailties $\{\varphi_{ik}\}$. The full conditional distributions for the σ_i^2 's, are conjugate inverted-Gamma distributions that are directly updated. The cancer covariance parameter, $\boldsymbol{\lambda}$, has a conjugate inverted-Wishart distribution, following from a lemma proved in DBD and stated in the Appendix. Finally, the spatial smoothness parameter, α , is updated using a *slice sampler* [26,27], which amounts to a rejection sampler from the prior of α .

Comparing with a non-spatial PH model with $\varphi_{ik} = 0, i = 1, \dots, I$ and $k = 1, \dots, K$, we obtain essentially the same likelihood but without the spatial components $\{\varphi_{ik}\}$ and α . Consequently, the full conditional distribution of $\boldsymbol{\Lambda}$ now excludes the contribution from the spatial frailties. The full conditional distributions are presented in the Appendix.

3.2. PO model

Analogous to the PH model presented in the previous subsection, a parametric PO model with log-logistic(λ_1, λ_2) baseline distribution incorporating covariate effects, main effects for patients and cancers and nested effects for cancers within counties, may be expressed as

$$\frac{S(t_{ijk}|\mathbf{x}_{ijk}, u_{(i,j)}, v_k, \varphi_{ik})}{1 - S(t_{ijk}|\mathbf{x}_{ijk}, u_{(i,j)}, v_k, \varphi_{ik})} = \frac{1}{\lambda_2 t_{ijk}^{\lambda_1}} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}} \quad (7)$$

for $i = 1, 2, \dots, I, j = 1, 2, \dots, n_i$, and $k = 1, 2, \dots, K$. Using the above specification with $\delta_{(i,j)}$ as in Section 3.2, the likelihood contribution of the (i, j) th patient is given by

$$L(\boldsymbol{\beta}, \lambda_1, \lambda_2 \{u_{(i,j)}\}, \{v_k\}, \{\varphi_{ik}\}; \{t_{ijk}\}, \delta_{(i,j)}, \mathbf{x}_{ijk}) \\ = \prod_{i=1}^I \prod_{j=1}^{n_i} \prod_{k \in C(i,j)} \left[\frac{(\lambda_1 \lambda_2 t_{ijk}^{\lambda_1 - 1} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}})^{\delta_{(i,j)}}}{(1 + \lambda_2 t_{ijk}^{\lambda_1} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}})^{1 + \delta_{(i,j)}}} \right] \quad (8)$$

with the parameters defined in previous sections.

We next assign prior distributions to complete the hierarchical specification. As in the PH setting we assign vague multivariate normal or improper flat prior for the regression coefficients, a $\text{Gamma}(s_0, s_1)$ prior for the scale parameter λ_1 and a log-normal(μ_0, τ^2) for the

shape parameter λ_2 of the log-logistic baseline distribution, inverted-Gamma, $IG(a_i, b_i)$, priors for the σ_i^2 's, an inverted-Wishart prior, $IW(r_0, \Lambda_0)$, for the covariance matrix Λ , and $U(0, 1)$ or a Beta prior for the spatial smoothness parameter α , where the hyper-parameters are chosen to obtain vague prior distributions.

Although the likelihoods arising from PH differ from PO, their full conditional distributions behave very similarly in terms of functional characteristics (i.e. log-concavity, conjugacy, etc.) and hence were updated using the same approaches described in Section 3.1. The details regarding the full conditional distributions are presented in the Appendix.

4. ILLUSTRATION

Data were obtained from the 1973–2001 SEER Public Use Incidence Database [28] on patients from the state of Iowa diagnosed with at least two of the following GI cancers: stomach, small intestine, colon, rectum, gall bladder, and pancreas. The above sites were identified in the SEER database by the variable *site recode*. We have collected a total of 7740 cases (stomach: 196 cases; small intestine: 144 cases; colon: 5702 cases; rectum: 1435 cases; gall bladder: 52 cases; and pancreas: 211 cases) from 3666 patients representing all 99 counties in the state. Survival time in months (from diagnosis to end of study or death) and the status at endpoint were recorded. For purposes of this investigation, death was defined based on *all-cause mortality* for lack of a more definitive source of information. All autopsy (diagnosed after death) cases were discarded as they do not contribute any information to either the hazard rate or the survival odds but would only complicate the numerical implementation. In addition to the survival information, we also looked at the gender, number of primaries (nprimes dichotomized as ‘= 2(Mode)’ or ‘>2’), radiation and/or surgery received for each site (0 for neither, 1 otherwise), age at diagnosis in years, and the stage of the cancer (*in situ* and unstaged (Stage I), local (Stage II), regional (Stage III), distant (Stage IV)). In this data set we have 49.6 per cent of the patients to be females and 31.6 per cent having more than two primary cancers. Of the total number of cases, 43.9, 33, and 10.5 per cent were Stages II, III, and IV, respectively, and only 19.6 per cent received site-specific radiation and/or surgery.

Graphical approaches were used to evaluate the appropriateness of the modeling framework we are proposing [20]. The Kaplan–Meier estimates of the survival function for each cancer site are plotted in Figure 1, while appropriate functions of the estimated cumulative hazards ($\log(H)$ for Weibull and $\log(\exp(H)-1)$ for log-logistic distributions) are plotted against the natural logarithm of time in Figure 2. Both plots suggest that the survival distributions of the different cancers may be coming from the same family, differing only in some specific parameters. In particular, the survival distributions of the different cancers may be different only in terms of the shape parameters under either Weibull or log-logistic distributions. The parallel lines in Figure 2(a) suggest that the Weibull distribution may be appropriate for modeling the survival distribution of the different cancers. A similar observation could be made for the log-logistic distribution, Figure 2(b), but some curvature is notable at larger survival times indicating that this model may not fit as well as the Weibull model.

We then fitted the $MCAR(\alpha, \Lambda)$ and the non-spatial PH and PO models to the data. As mentioned in Section 2.4, the models were implemented without the patient-specific frailties due to convergence issues. Henceforth, we adopt the naming convention presented in Table II. Appropriate Gibbs samplers were set up to obtain samples from the posterior distributions. Metropolis–Hastings algorithms with normal candidates were used to update the posterior distributions of the regression parameters, the frailties, and the parameters of the baseline distribution. The cancer variance–covariance matrix Λ was updated using direct draws from the appropriate inverted-Wishart distribution. The results of the model-fitting exercise with 10

000 post-burnin samples were summarized using the Bayesian output analysis program [29] and are presented in Table III.

It can be observed from Table III that the parameter estimates are consistent across the different models, which should give additional support to the modeling framework. For the PH models, the regression terms contribute to the increase or decrease in the hazard rate. The negative coefficients for gender suggest that the relative risk of dying is significantly lower in women than in men. This finding is consistent with the estimated mortality rates presented in the SEER cancer statistics review [30] as well as in the results of DBD. The number of primary cancers has a significant negative coefficient, which indicates a lowering of hazard rate as more primary cancers are observed. The same phenomenon has been observed in DBD, where they noted that this variable is *post hoc* in the sense that one does not know how many primaries the patient will develop in the beginning, hence could not be taken as a prognostic factor. An interesting finding is the increased risk for those who received site-specific radiation and/or surgery (positive regression coefficient). Upon further investigation we found out that more advanced stage (Stages III and IV) cancer cases tend to receive radiation and/or surgery (21.07 per cent, s.e.= 0.01 per cent) compared with early stage (Stages I and II) cancer cases (18.44 per cent, s.e.= 0.01 per cent). In addition, for early-stage cancer cases the risk of dying given a site-specific radiation/surgery was performed is 37.10 per cent (s.e.= 0.02 per cent) compared with 56.48 per cent (s.e.= 0.02 per cent) for cases in the advanced stages. This suggests that those receiving radiation and/or surgery already have a higher risk of dying to start with. It is also worth mentioning that the relative risk of failure increases with age as indicated by the positive regression coefficient.

For the PO model, similar interpretations of the regression parameter estimates can be made although the regression terms now affect the odds of failure. From Model 3, for instance, the -0.44 coefficient for gender suggests that the odds of failure at any time is significantly lower for females than males by a factor of 0.64 ($= \text{failure odds ratio}$). A similar finding is observed for a number of primary cancers with a failure odds ratio (>2 vs $= 2$) of 0.74 . An increase in odds of failure is noted for cases who underwent site-specific radiation/surgery compared with those who did not. An increase in the odds of failure with older age (1.05 times for each unit increase in age) is observed. Note that the estimate does not depend on the unit of age while in the PH model a unit change results in the re-scaling of the estimate.

We employed two approaches to evaluate the performance of the proposed models, the first of which is the conditional predictive ordinate (CPO) with the average log-marginal pseudo-likelihood (ALMPL). Roughly speaking, the CPO statistic is a Bayesian criterion that is used to evaluate how well the model predicts a particular observation given the rest of the data. Following Gelfand and Dey [31], the CPO statistic for the (i, jk) th case is given by

$$\text{CPO}_{ijk} = f(D_{ijk} | \mathbf{x}_{ijk}, \mathbf{D}^{(-ijk)}) = \left(\int \frac{1}{f(D_{ijk} | \mathbf{x}_{ijk}, \boldsymbol{\Theta})} \pi(\boldsymbol{\Theta} | \mathbf{D}) d\boldsymbol{\Theta} \right)^{-1} \quad (9)$$

where $D_{ijk} = (i, jk)$ th observation, $\mathbf{x}_{ijk}$ is the corresponding covariate vector, $\mathbf{D}^{(-ijk)}$ is the data without the (i, jk) th observation, $\mathbf{D}$ is the complete data, and $\boldsymbol{\Theta}$ is the collection of parameters. The four models under consideration have very similar expressions for the CPO statistic, except for the presence of additional terms in the MCAR(α, Λ) models and the actual expressions of the likelihood contribution. The CPO expression for each model is as follows:

$$\text{CPO}_{ijk}^m = \left(\int \frac{1}{f(D_{ijk} | \mathbf{x}_{ijk}, \mathbf{Z}_m, \boldsymbol{\Theta}_m)} g(\mathbf{Z}_m | \mathbf{D}, \boldsymbol{\Theta}_m) \pi(\boldsymbol{\Theta}_m | \mathbf{D}) d\mathbf{Z}_m d\boldsymbol{\Theta}_m \right)^{-1}, \quad m=1, \dots, 4$$

where $D_{ijk} = \{t_{ijk}, \delta_{ij}; k \in C_{ii,j}\}$, $f(D_{ijk}|\mathbf{x}_{ijk}, \mathbf{Z}_m, \boldsymbol{\Theta}_m)$ is the corresponding likelihood contribution, $g(\mathbf{Z}_m|\mathbf{D}, \boldsymbol{\Theta}_m)$ is the full conditional distribution of the frailties, and $\pi(\boldsymbol{\Theta}_m|\mathbf{D})$ is the posterior distribution of the parameters. For the MCAR(α, Λ) models ($m = 1, 3$), $\mathbf{Z}_m = \{v_k, \varphi_{ik}; k \in C_{(i,j)}\}$ and $\boldsymbol{\Theta}_m = (\beta, \lambda_1, \lambda_2, \Lambda, \alpha)$. For the non-spatial models ($m = 2, 4$), $\mathbf{Z}_m = \{v_k; k \in C_{(i,j)}\}$ and $\boldsymbol{\Theta}_m = (\beta, \lambda_1, \lambda_2, \Lambda)$. The integrals given above cannot be obtained in closed form and hence were evaluated using Monte Carlo integration. For each model, the ALMPL (over the total number of cases) was used as an overall measure of model fit and was calculated as

$$\text{ALMPL}^m = \frac{\sum_{ijk} \log \text{CPO}_{ijk}^m}{n}$$

where n is the total number of cases. Larger values of ALMPL indicate better model fit. The calculated ALMPL values are notably close for the four models ($-3.99, -3.88, -3.92$, and -3.90 for Model 1, Model 2, Model 3, and Model 4, respectively). A similar finding could be noted from the boxplots of the $\{\log(\text{CPO}_{ijk})\}$ (Figure 3). The boxplots are exhibiting very similar distributions except for a couple of extreme values in Model 1. It appears that the CPO statistic was not capable of distinguishing between the proposed parametric models, which leads us to consider another method for evaluating the fit of the model.

One of the classical approaches in evaluating the goodness of fit of a parametric survival model is using the Cox–Snell residuals. The Cox–Snell residuals are essentially the cumulative hazard rates based on the fitted model and are expected to follow a standard exponential distribution if the model is adequate [20]. However, in this particular application, we need to take into account the presence of the frailty terms. Hence, we define conditional (on the realized frailties) Cox–Snell residuals as

$$r_{ijk}^m = H(t_{ijk}|\mathbf{x}_{ijk}, \mathbf{Z}_m, \boldsymbol{\Theta}_m), \quad m = 1, \dots, 4$$

where H is the cumulative hazard function evaluated at t_{ijk} given the covariates, frailties, and the model parameters. It can be argued that the observed frailty terms, $\mathbf{Z}_m$, merely contribute constant terms to the shape parameter and hence we would expect the same properties as the classical Cox–Snell residuals. The quantile plots of the conditional Cox–Snell residuals against the unit exponential distribution (not shown) suggest that of the four models, the PH model with the MCAR(α, Λ) spatial frailties is the most adequate. The other three models have distributions with much lighter tails than the unit exponential indicating a lack of fit. However, since some of the survival times are censored, there is potential censoring of the Cox–Snell residuals as well. A more appropriate presentation is the plot of the distribution functions based on the Cox–Snell residuals vs the distribution function of a standard exponential distribution (Figure 4). The residuals from the four models appear to have lighter left tails than that of the exponential. Further investigation shows that a higher percentage of the residuals from cancer types with sparse data (stomach, small intestine, gall bladder, and pancreas) fall on the left tails of the distribution functions than colon and rectal cancers. In addition to having lighter left tails, models based on the log-logistic distribution (Models 3 and 4) also have lighter right tails. This suggests that the PH models (Models 1 and 2) might be more appropriate.

On the basis of the results of the model evaluation and comparison techniques, we suggest the use of Model 1, the parametric PH model with the MCAR(α, Λ) spatial frailties, to be the final model. The spatial frailties for this model are summarized in the choropleth maps (Plate 1). We could observe from the map that there is a clustering effect among adjacent counties. It

should also be noted that the frailty values are not to be interpreted in absolute terms as they simply correspond to either a decrease or an increase in relative risks in a particular county compared with the rest. This is due to the zero-mean constraint imposed for identifiability. The spatial frailties together with some additional data, such as hospital/healthcare facility locations and general socioeconomic status at the county level, could augment the information available to policymakers and researchers in assessing which factors not available in the SEER database potentially contribute to elevating/lowering of hazard rates.

The *a priori* cancer correlation matrix for the recommended model is presented in Table IV. The coefficients range from insignificant to moderate strength. The strongest correlation is between colon and stomach cancers (0.38) while small intestine is weakly correlated with all the other cancers except colon. Negative correlation with the other cancers is generally observed for pancreatic cancer. Positive coefficients between two cancers indicate that their hazard rates go up or down together while negative coefficients point to opposing trends in the hazard rates.

5. SUMMARY AND FUTURE WORK

This paper proposed parametric models for spatially correlated multivariate cancer survival data. Using data on 6 GI cancers in the state of Iowa collected from the SEER database, we illustrated that there are significant correlations between the hazard rates among the cancers. There is also evidence of spatial association across the different counties. We also note that for this particular data set, females have relatively lower risks of dying compared with males and that more primary cancers result in lower hazard rate for a specific cancer. Higher (older) age at diagnosis increases the risk of failure, which may be attributed to the use of *all-cause* mortality as endpoint. We also found that cases who received site-specific radiation and/or surgery had higher risks compared with those who did not. This could be attributed to the fact that most of the radiation/surgery cases were already at elevated risk by being in more advanced stages.

For this particular data set, the use of parametric models is both practical and appropriate. The proposed parametric models lend themselves to straightforward implementation and ease of interpretation of the parameters. Semi-parametric approaches could also be used for modeling the baseline distributions as in DBD. Preliminary investigations comparing parametric and semi-parametric implementations of the PH model with the $\text{MCAR}(\alpha, \Lambda)$ structure were conducted using the colon and rectal cancer data set and recommended hazard model presented in DBD. Although results are not presented here, we have found that the parametric model performed better (ALMPL = -4.03) than its semi-parametric counterpart (ALMPL = -6.03). However, a more extensive evaluation is necessary to determine under which circumstances it would be more appropriate to use the parametric or semi-parametric implementation.

As mentioned in Section 2.4, the $\text{MCAR}(\alpha, \Lambda)$ structure may be generalized in several ways: accounting for different spatial and non-spatial cancer variance-covariance matrices, $\text{MCAR}(\alpha, \mathbf{T})$; accounting for different spatial smoothness parameters for each type of cancer, $\text{MCAR}((\alpha_1, \dots, \alpha_K), \Lambda)$; and accounting for different cancer variance-covariance patterns for each county, $\text{MCAR}(\alpha, (\Lambda_1, \dots, \Lambda_I))$. These extensions of the $\text{MCAR}(\alpha, \Lambda)$ model were not implemented as they would require a more robust data set as a result of complexity. It is also worth mentioning that convergence problems were encountered when introducing patient-specific frailties. This is likely due to the fact that most patients in the data set were diagnosed with only two of the 6 GI cancers studied.

A potential issue with the multivariate modeling approach presented is the identifiability of the dependence structure (not the model) due to the different cancers censoring each other's

survival times. A suggested remedy is to use age at diagnosis as the endpoint instead of survival times. However, while both endpoints are time to event in nature and are clinically relevant, the resulting models would have very different interpretations. The question of which endpoint is more appropriate and relevant is beyond the scope of this paper. Finally, while the MCAR structure provides a useful framework for modeling the dependencies, it is possible that using a different dependence structure could give different results and model interpretation. We therefore recommend that caution be taken when interpreting the observed correlations.

References

1. Turechek WW, Madden LV. A generalized linear modeling approach for characterizing disease incidence in spatial hierarchy. *Phytopathology* 2002;93:458–466. [PubMed: 18944361]
2. Ramsay T, Burnett R, Krewski D. Exploring bias in a generalized additive model for spatial air pollution data. *Environmental Health Perspectives* 2003;111:1283–1288. [PubMed: 12896847]
3. Lichstein JW, Simons TR, Shriner SA, Franzreb KE. Spatial autocorrelation and autoregressive models in ecology. *Ecological Monographs* 2002;72(3):445–463.
4. Biggeri A, Marchi M, Lagazio C, Martuzzi M, Böhning D. Nonparametric maximum likelihood estimators for disease mapping. *Statistics in Medicine* 2000;19:2539–2554. [PubMed: 10960870]
5. Li Y, Ryan L. Modeling spatial survival data using semiparametric frailty models. *Biometrics* 2002;58:287–297. [PubMed: 12071401]
6. Banerjee S, Wall M, Carlin BP. Frailty modelling for spatially correlated survival data with application to infant mortality in Minnesota. *Biostatistics* 2003;4:123–142. [PubMed: 12925334]
7. Banerjee S, Dey DK. Semiparametric proportional odds model for spatially correlated survival data. *Lifetime Data Analysis* 2005;11:175–191. [PubMed: 15938545]
8. Li Y, Lin X. Semiparametric normal transformation models for spatially correlated survival data. *Journal of the American Statistical Association* 2006;101:591–603.
9. Heinavaara, S. PhD Thesis, University of Helsinki, Statistical Research Reports, No. 18. The Finnish Statistical Society; 2003. Modelling survival of patients with multiple cancers.
10. Sankila R, Hakulinen T. Survival of patients with colorectal carcinoma: effect of prior breast cancer. *Journal of the National Cancer Institute* 1998;90:6365.
11. Carlin, BP.; Banerjee, S. Hierarchical multivariate car models for spatio-temporally correlated survival data. In: Bernardo, JM.; Bayarri, MJ.; Berger, JO.; Dawid, AP.; Heckerman, D.; Smith, AF.; West, M., editors. *Bayesian Statistics. Vol. 7*. Oxford University Press; Oxford: 2003. p. 45-64.
12. Diva UA, Banerjee S, Dey DK. Modeling spatially correlated survival data for individuals with multiple cancers. *Statistical Modeling* 2007;7(2):1–23.
13. Mardia KV. Multidimensional multivariate Gaussian Markov random fields with application to image processing. *Journal of Multivariate Analysis* 1998;24:45–64.
14. Gelfand AE, Vounatsou P. Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* 2002;4:11–25. [PubMed: 12925327]
15. Zucker DM, Yang S. Inference for a family of survival models encompassing the proportional hazards and proportional odds model. *Statistics in Medicine* 2005;25:995–1014. [PubMed: 16220492]
16. Cox DR. Regression models with life tables. *Journal of the Royal Statistical Society* 1972;34:187–220.
17. Gelfand AE, Ghosh SK, Christiansen C, Soumerai SB, McLaughlin TJ. Proportional hazards models: a latent competing risk approach. *Applied Statistics* 2000;49:385–397.
18. Goetghebuer E, Ryan L. Semiparametric regression analysis of interval-censored data. *Biometrics* 2000;56:1139–1144. [PubMed: 11129472]
19. Sinha D, Chen MH, Ghosh SK. Bayesian analysis and model selection for interval-censored survival data. *Biometrics* 1999;55:585–590. [PubMed: 11318218]
20. Klein, JP.; Moeschberger, ML. *Survival Analysis: Techniques for Censored and Truncated Data*. Springer; New York: 1997. p. 373-400.
21. Bennett S. Analysis of survival data by the proportional odds model. *Statistics in Medicine* 1983;2:273–277. [PubMed: 6648142]

22. Yang S, Prentice RL. Semiparametric inference in the proportional odds regression model. *Journal of the American Statistical Association* 1999;94:125–136.
23. Clayton DG. Some approaches to the analysis of recurrent event data. *Statistics in Medical Research* 1995;3:244–262.
24. Jin X, Carlin BP, Banerjee S. Generalized hierarchical multivariate car models for areal data. *Biometrics* 2005;61:950–961. [PubMed: 16401268]
25. Gilks WR, Wild P. Adaptive rejection sampling for Gibbs sampling. *Journal of the Royal Statistical Society, Series C* 1992;41:337–348.
26. Agarwal D, Gelfand AE. Slice Gibbs sampling for simulation based fitting of spatial models. *Statistics and Computing* 2005;15:61–69.
27. Neal R. Slice sampling. *Annals of Statistics* 2003;31(3):705–767.
28. *Surveillance, Epidemiology, and End Results (SEER) Program (www.seer.cancer.gov) SEER*Stat Database: Incidence—SEER 11 Regs + AK Public-use, November 2003 Sub (1973–2001 varying)*. National Cancer Institute, DCCPS, Surveillance Research Program, Cancer Statistics Branch, released April 2004, based on the November 2003 submission.
29. Smith, BA. Bayesian Output Analysis Program (BOA) Version 1.1 User's Manual. [08 August 2005]. <http://www.public-health.uiowa.edu/boa/>
30. Ries, LAG.; Eisner, MP.; Kosary, CL.; Hankey, BF.; Miller, BA.; Clegg, L.; Mariotto, A.; Feuer, EJ.; Edwards, BK. SEER Cancer Statistics Review. National Cancer Institute; Bethesda, MD: 1975–2002. Available from: <http://seer.cancer.gov/csr/19752002/> based on November 2004 SEER data submission, posted to the SEER Web site 2005
31. Gelfand AE, Dey DK. Bayesian model choice: asymptotics and exact calculations. *Journal of Royal Statistical Society, Series B* 1994;56:501–514.

Appendix

APPENDIX

1. Lemma

Let A and B be $I \times I$ and $K \times K$ matrices and let $\mathbf{x}$ be any vector of length $I \times K$. Then, there exist a $K \times K$ matrix C and an $I \times I$ matrix D such that $\mathbf{x}^T(A \otimes B)\mathbf{x} = \text{tr}(C B) = \text{tr}(D A)$.

Proof

See DBD.

2. Full conditional distributions for the PH models

In the succeeding expressions we let $H(t_{ijk}; \mathbf{x}_{ijk}, \Theta)$ be the cumulative hazard function with parameters $\Theta = \{\lambda_1, \lambda_2, \beta, \mathbf{v}, \mathbf{u}, \Phi\}$ and $\mathbf{D} = \{t_{ijk}, \delta_{(i,j)}, \mathbf{x}_{ijk}\}$ be the observed data. For notational convenience, let $\Sigma_{i,j,k}$ denote $\sum_{i=1}^I \sum_{j=1}^{n_i} \sum_{k \in C(i,j)}$ and $\Pi_{i,j,k}$ denote $\prod_{i=1}^I \prod_{j=1}^{n_i} \prod_{k \in C(i,j)}$ unless otherwise specified.

a. For λ_1 with $\text{Gamma}(s_0, s_1)$ prior, the full conditional distribution is given by

$$f(\lambda_1 | s_0, s_1, \Theta, \mathbf{D}) \propto \lambda_1^{\sum_{i,j,k} \delta_{(i,j)} + s_0 - 1} \left[\prod_{i,j,k} t_{ijk}^{\delta_{(i,j)}} \right]^{\lambda_1} \exp \left[-\frac{\lambda_1}{s_1} - \sum_{i,j,k} H(t_{ijk}; \mathbf{x}_{ijk}, \Theta) \right]$$

which can be updated using an ARS or a Metropolis step.

b. For λ_2 with $\text{log-normal}(\mu_0, \tau_0^2)$ prior, the full conditional distribution is given by

$$f(\lambda_2|\mu_0, \tau_0^2, \Theta, \mathbf{D}) \propto \lambda_2^{\sum_{i,j,k} \delta_{(i,j)} - 1} \exp \left[-\frac{(\log \lambda_2 - \mu_0)^2}{2\tau_0^2} - \sum_{i,j,k} H(t_{ijk}; \mathbf{x}_{ijk}, \Theta) \right]$$

which is updated using an ARS or a Metropolis step.

- c. For β with $\text{MVN}(\mathbf{0}, \sigma_\beta^2 \mathbf{I})$ prior, the full conditional distribution is given by

$$f(\beta|\sigma_\beta^2, \Theta, \mathbf{D}) \propto \exp \left[-\frac{\beta' \beta}{2\sigma_\beta^2} + \beta' \sum_{i,j,k} \delta_{(i,j)} \mathbf{x}_{ijk} - \sum_{i,j,k} H(t_{ijk}; \mathbf{x}_{ijk}, \Theta) \right]$$

which is updated using an ARS or a Metropolis step.

- d. For $u_{(i,j)}$ with independent $N(0, \sigma_i^2)$ prior, the full conditional distribution is given by

$$f(u_{(i,j)}|\sigma_i^2, \Theta, \mathbf{D}) \propto \exp \left[-\frac{u_{(i,j)}^2}{2\sigma_i^2} + \sum_{k \in C(i,j)} \delta_{(i,j)} u_{(i,j)} - \sum_{i,j,k} H(t_{ijk}; \mathbf{x}_{ijk}, \Theta) \right]$$

which is updated using an ARS or a Metropolis step.

- e. For σ_i^2 with $\text{IG}(a_i, b_i)$ prior, the full conditional distribution is given by

$$\sigma_i^2 | \mathbf{u}, a_i, b_i \sim \text{IG}(a_i + n_i, b_i + \frac{1}{2} \mathbf{u}_i^T \mathbf{u}_i) \text{ for } i=1, 2, \dots, I$$

where n_i is the number of patients in the i th county and $\mathbf{u}_i = [u_{(i,j)}]_{j=1}^{n_i}$. This is updated using direct draws from its density.

- f. For $\mathbf{v}$ with $\text{MVN}(\mathbf{0}, \Lambda)$ prior, the full conditional distribution is given by

$$f(\mathbf{v}|\Lambda, \Theta, \mathbf{D}) \propto \exp \left[-\frac{\mathbf{v}' \Lambda^{-1} \mathbf{v}}{2} + \sum_{i,j,k} \delta_{(i,j)} v_k - \sum_{i,j,k} H(t_{ijk}; \mathbf{x}_{ijk}, \Theta) \right]$$

which is updated using an ARS or a Metropolis step.

- g. For Φ with $\text{MCAR}(\alpha, \Lambda)$ prior, the full conditional distribution is given by

$$f(\Phi|\alpha, \Lambda, \Theta, \mathbf{D}) \propto \exp \left[-\frac{\Phi' (\sum_w (\alpha) \otimes \Lambda)^{-1} \Phi}{2} + \sum_{i,j,k} \delta_{(i,j)} \varphi_{ik} - \sum_{i,j,k} H(t_{ijk}; \mathbf{x}_{ijk}, \Theta) \right]$$

which is updated using an ARS or a Metropolis step.

- h. For Λ with $\text{IW}(r_0, \Lambda_0)$ prior, the full conditional distribution follows from the lemma above and is given by

$$\Lambda|r_0, \Lambda_0, \mathbf{v}, \Phi, \alpha \sim \text{IW}(r_0 + I, \Lambda_0^*)$$

with $\Lambda_0^* = [\Lambda_0^{-1} + \mathbf{v}\mathbf{v}' + \sum_{i=1}^I \sum_{j=1}^{n_i} (\sum_w^{-1}(\alpha))_{ii'} \varphi_i \varphi_{i'}']^{-1}$ which is updated using direct draws from its density. The derivations are omitted for brevity and the reader is referred to DBD for details.

- i. For α with $U(0,1)$ prior, the full conditional distribution is given by

$$f(\alpha|\Phi, \Lambda) \propto f(\alpha) \times \frac{1}{|\sum_w(\alpha)|^{K/2}} \exp \left[-\frac{\Phi'(\sum_w^{-1}(\alpha) \otimes \Lambda^{-1})\Phi}{2} \right]$$

which is updated using a *slice sampler* outlined below:

- i. Within the current state of the Gibbs sampler, let

$$f(\Phi; \alpha, \Lambda) = \frac{1}{|\sum_w(\alpha)|^{K/2}} e^{-\Phi'(\sum_w^{-1}(\alpha) \otimes \Lambda^{-1})\Phi/2}$$

- ii. Generate $U \sim U(0, f(\Phi; \alpha, \Lambda))$.
 iii. Generate $\alpha^* \sim f(\alpha)$.
 iv. If $U > f(\Phi; \alpha^*, \Lambda)$ set $\alpha = \alpha^*$. Otherwise, repeat Step (i).

This procedure is equivalent to a rejection sampling from the prior of α .

- j. For the non-spatial PH model, the $\{\varphi_{ik}\}$'s and α will disappear and we would need to make the following modification to the full conditional distribution of Λ :

$$\Lambda|r_0, \Lambda_0, \mathbf{v}, \sim \text{IW}(r_0 + I, \Lambda_0^*)$$

with $\Lambda_0^* = [\Lambda_0^{-1} + \mathbf{v}\mathbf{v}']^{-1}$ which is updated using direct draws from its density.

3. Full conditional distributions for the PO models

In the succeeding expressions we let $G(t_{ijk}; \mathbf{x}_{ijk}, \Theta) = 1 + \lambda_2 t_{ijk}^{\lambda_1} e^{\mathbf{x}_{ijk}^T \boldsymbol{\beta} + u_{(i,j)} + v_k + \varphi_{ik}}$ with parameters $\Theta = \{\lambda_1, \lambda_2, \boldsymbol{\beta}, \mathbf{v}, \mathbf{u}, \Phi\}$ and the observed data $\mathbf{D} = \{t_{ijk}, \delta_{(i,j)}, \mathbf{x}_{ijk}\}$. For notational convenience, let $\sum_{i,j,k}$ denote $\sum_{i=1}^I \sum_{j=1}^{n_i} \sum_{k \in C(i,j)}$ and $\prod_{i,j,k}$ denote $\prod_{i=1}^I \prod_{j=1}^{n_i} \prod_{k \in C(i,j)}$ unless otherwise specified. The same expressions for the full conditional distributions were obtained for Λ and α as in the PH model and hence were omitted.

- a. For λ_1 with $\text{Gamma}(s_0, s_1)$ prior, the full conditional distribution is given by

$$f(\lambda_1|s_0, s_1, \Theta, \mathbf{D}) \propto \lambda_1^{\sum_{i,j,k} \delta_{(i,j)} + s_0 - 1} e^{-\lambda_1 (\sum_{i,j,k} \delta_{(i,j)} \log t_{ijk} + 1/s_1)} \prod_{i,j,k} G(t_{ijk}; \mathbf{x}_{ijk}, \Theta)^{-(\delta_{(i,j)} + 1)}$$

which can be updated using a Metropolis step.

- b.** For λ_2 with $\log\text{-normal}(\mu_0, \tau_0^2)$ prior, the full conditional distribution is given by

$$f(\lambda_2|\mu_0, \tau_0^2, \boldsymbol{\Theta}, \mathbf{D}) \propto \lambda_2^{\sum_{i,j,k} \delta_{(i,j)} - 1} \exp \left[-\frac{(\log \lambda_2 - \mu_0)^2}{2\tau_0^2} \right] \prod_{i,j,k} G(t_{ijk}; \mathbf{x}_{ijk}, \boldsymbol{\Theta})^{-(\delta_{(i,j)} + 1)}$$

which is updated using a Metropolis step.

- c.** For $\boldsymbol{\beta}$ with $\text{MVN}(\mathbf{0}, \sigma_\beta^2 \mathbf{I})$ prior, the full conditional distribution is given by

$$f(\boldsymbol{\beta}|\sigma_\beta^2, \boldsymbol{\Theta}, \mathbf{D}) \propto \exp \left[-\frac{\boldsymbol{\beta}' \boldsymbol{\beta}}{2\sigma_\beta^2} + \boldsymbol{\beta}' \sum_{i,j,k} \delta_{(i,j)} \mathbf{x}_{ijk} \right] \prod_{i,j,k} G(t_{ijk}; \mathbf{x}_{ijk}, \boldsymbol{\Theta})^{-(\delta_{(i,j)} + 1)}$$

which is updated using a Metropolis step.

- d.** For $u_{(i,j)}$ with independent $N(0, \sigma_i^2)$ prior, the full conditional distribution is given by

$$f(u_{(i,j)}|\sigma_i^2, \boldsymbol{\Theta}, \mathbf{D}) \propto \exp \left[-\frac{u_{(i,j)}^2}{2\sigma_i^2} + \sum_{k \in C(i,j)} \delta_{(i,j)} u_{(i,j)} \right] \prod_{i,j,k} G(t_{ijk}; \mathbf{x}_{ijk}, \boldsymbol{\Theta})^{-(\delta_{(i,j)} + 1)}$$

which is updated using a Metropolis step.

- e.** For σ_i^2 with $I G(a_i, b_i)$ prior, the full conditional distribution is given by

$$\sigma_i^2 | \mathbf{u}, a_i, b_i \sim \text{IG}(a_i + n_i, b_i + \frac{1}{2} \mathbf{u}_i^T \mathbf{u}_i) \text{ for } i=1, 2, \dots, I$$

where n_i is the number of patients in the i th county and $\mathbf{u}_i = [u_{(i,j)}]_{j=1}^{n_i}$. This is updated using direct draws from its density.

- f.** For $\boldsymbol{\nu}$ with $\text{MVN}(\mathbf{0}, \boldsymbol{\Lambda})$ prior, the full conditional distribution is given by

$$f(\boldsymbol{\nu}|\boldsymbol{\Lambda}, \boldsymbol{\Theta}, \mathbf{D}) \propto \exp \left[-\frac{\boldsymbol{\nu}' \boldsymbol{\Lambda}^{-1} \boldsymbol{\nu}}{2} + \sum_{i,j,k} \delta_{(i,j)} \nu_k \right] \prod_{i,j,k} G(t_{ijk}; \mathbf{x}_{ijk}, \boldsymbol{\Theta})^{-(\delta_{(i,j)} + 1)}$$

which is updated using a Metropolis step.

- g.** For $\boldsymbol{\Phi}$ with $\text{MCAR}(\alpha \boldsymbol{\Lambda})$ prior, the full conditional distribution is given by

$$f(\boldsymbol{\Phi}|\alpha, \boldsymbol{\Lambda}, \boldsymbol{\Theta}, \mathbf{D}) \propto \exp \left[-\frac{\boldsymbol{\Phi}' (\sum_w (\alpha) \otimes \boldsymbol{\Lambda})^{-1} \boldsymbol{\Phi}}{2} + \sum_{i,j,k} \delta_{(i,j)} \varphi_{ik} \right] \prod_{i,j,k} G(t_{ijk}; \mathbf{x}_{ijk}, \boldsymbol{\Theta})^{-(\delta_{(i,j)} + 1)}$$

which is updated using a Metropolis step.

Figure 1.
Kaplan–Meier plots of the survival times for the six GI cancers studied in the state of Iowa.

Figure 2.

Exploratory data analysis plots to check appropriateness of the Weibull ($\log(H)$) and log-logistic ($\log(\exp(H)-1)$) distributions based on Kaplan–Meier estimates of the cumulative hazard function. The x -axes for both graphs represent the natural logarithm of the survival time. Linear trends provide support to the use of the specified distribution: (a) Weibull distribution; and (b) Log-logistic distribution.

Figure 3. Boxplots of the log(CPO) for each model. Larger values of the log(CPO) indicate more support for the model.

Figure 4.

Diagnostics plots to check for adequacy of the $\text{MCAR}(\alpha, \Lambda)$ and non-spatial Weibull and log-logistic distributions. Clockwise from top-left: Model 1, Model 2, Model 4, and Model 2. For all graphs, the y-axes represent the Kaplan–Meier estimate of distribution function of the Cox–Snell residuals, while the x-axes represent the distribution function of the standard exponential. Linear trends suggest adequacy of the model.

Plate 1.

Plots of the spatial frailties for each type of GI cancer studied based on the proportional hazards model with Weibull(λ_1, λ_2) baseline distribution (Model 1). color change from blue to red indicates an increase from negative (lower risk) to positive frailties (higher risk). The same cutoff points were used to facilitate comparison between cancer types.

Table 1
Illustrative example of data for a patient with multiple cancers.

County	Patient ID	Gender	No. of primaries	Cancer type	Stage	Age at diagnosis	Radiation/surgery	Survival time*	Status
7	00171	Female	3	Small intestine	Distant	57	No	143	Dead
7	00171	Female	3	Colon	Local	63	No	78	Dead
7	00171	Female	3	Pancreas	Distant	67	Yes	25	Dead

* Survival time is given in months after diagnosis.

Table II

Descriptions of the four parametric models considered.

Model	Model description
Model 1	Proportional hazards model with MCAR(α, Λ)
Model 2	Non-spatial proportional hazards model
Model 3	Proportional odds model with MCAR(α, Λ)
Model 4	Non-spatial proportional odds model

Table III

Parameter estimates and 95 per cent highest posterior density (HPD) intervals under various models.

Parameter	Model 1	Model 2	Model 3	Model 4
α	0.39 (0.00, 0.85)	—	0.38 (0.00, 0.85)	—
λ_1	0.90 (0.88, 0.92)	0.91 (0.89, 0.93)	1.21 (1.18, 1.24)	1.20 (1.17, 1.23)
λ_2	0.01 (0.01, 0.01)	0.02 (0.02, 0.02)	0.01 (0.01, 0.01)	0.01 (0.01, 0.01)
β_{Female}	-0.33 (-0.38, -0.27)	-0.34 (-0.40, -0.28)	-0.44 (-0.52, -0.35)	-0.43 (-0.52, -0.35)
$\beta_{\text{Primaries}}$	-0.15 (-0.21, -0.09)	-0.18 (-0.24, -0.12)	-0.30 (-0.39, -0.22)	-0.31 (-0.40, -0.23)
$\beta_{\text{Rad/surg}}$	0.13 (0.05, 0.21)	0.19 (0.11, 0.27)	0.30 (0.18, 0.42)	0.31 (0.20, 0.43)
$\beta_{\text{Age Diagnosed}}$	0.04 (0.04, 0.05)	0.04 (0.04, 0.05)	0.05 (0.05, 0.06)	0.05 (0.05, 0.06)
$\beta_{\text{Stage II}}$	-0.27 (-0.36, -0.19)	-0.35 (-0.43, -0.27)	-0.73 (-0.86, -0.60)	-0.74 (-0.87, -0.61)
$\beta_{\text{Stage III}}$	0.08 (-0.01, 0.17)	-0.05 (-0.13, 0.04)	-0.22 (-0.35, -0.09)	-0.23 (-0.37, -0.11)
$\beta_{\text{Stage IV}}$	1.39 (1.28, 1.49)	1.25 (1.15, 1.36)	1.90 (1.73, 2.07)	1.87 (1.70, 2.04)

Table IV
Correlation coefficients (95 per cent highest posterior density (HPD) intervals) of the cancer frailties based on the proportional hazards model with Weibull(λ_1, λ_2) baseline distribution (Model 1).

Type	Stomach	Small intestine	Colon	Rectum	Gall bladder
Small intestine	0.17 (0.11,0.24)	—	—	—	—
Colon	0.38 (0.26, 0.50)	−0.02 (−0.08,0.04)	—	—	—
Rectum	−0.01 (−0.03,0.01)	0.15 (0.09, 0.22)	−0.01 (−0.04,0.02)	—	—
Gall bladder	−0.09 (−0.13, −0.05)	0.19 (0.11, 0.27)	−0.12 (−0.17, −0.07)	0.00 (−0.03,0.05)	—
Pancreas	−0.16 (−0.22, −0.10)	−0.22 (−0.29, −0.14)	−0.04 (−0.08,0.01)	−0.13 (−0.17, −0.08)	−0.11 (−0.16, −0.06)