




The use of variational inference for lifetime data with spatial correlations

Yueyao Wang^{a,b}, Yili Hong^c, Laura Freeman^c, and Xinwei Deng^c 

^aSchool of Statistics and Mathematics, Zhejiang Gongshang University, Hangzhou, China; ^bCollaborative Innovation Center of Statistical Data Engineering Technology & Application, Zhejiang Gongshang University, Hangzhou, China; ^cDepartment of Statistics, Virginia Tech, Blacksburg, Virginia

ABSTRACT

Lifetime data with spatial correlations are often collected for analysis in modern engineering, clinical, and medical applications. For such spatial lifetime data, statistical models usually account for the spatial dependence through spatial random effects, such as the cumulative exposure model and the proportional hazards model. For these models, the Bayesian estimation is commonly used for model inference, but often encounters computational challenges when the number of spatial locations is large. The conventional Markov Chain Monte Carlo (MCMC) methods for sampling the posterior can be time-consuming. In this case-study article, we investigate the capability of variational inference (VI) for the model inference on spatial lifetime data, aiming for a good balance between the estimation accuracy and computational efficiency. Specifically, the VI methods with different divergence metrics are investigated for the spatial lifetime models. In the case study, the Titan graphics processing unit (GPU) lifetime data and the pine tree lifetime data are used to examine the VI methods in terms of their computational advantage and estimation accuracy.

KEYWORDS

α -divergence; Accelerated failure time model; Bayesian inference; Hamiltonian Monte Carlo; Proportional hazards model; Spatial random effects


1. Problem description

Spatially correlated lifetime data are commonly observed in engineering, medical, and biostatistics fields, when the data are collected over spatial locations, such as clinical sites and geographical regions. For example, Ostrouchov et al. (2020) provide a detailed analysis of Cray XK7 Titan supercomputer graphics processing unit (GPU) lifetime data. Inside the Titan supercomputer, GPU nodes are regularly arranged in the server room and present a spatial pattern. For such data, there exists unobserved heterogeneity in observations among different locations. Besides, the lifetime data in adjacent locations may be correlated, because closer locations can share similar environmental characteristics. Thus, modeling the spatial correlation is crucial for understanding the time-to-event data across locations.

Existing approaches incorporate the spatial correlation *via* extending lifetime models with spatially correlated random effects. The spatial random effects are assumed to follow a certain distribution, and the correlations between random effects are described by the variance-covariance matrix of the distribution. For

example, Henderson, Shimakura, and Gorst (2002) investigate possible spatial variation in the survival of adult Leukemia patients based on Bayesian hierarchical proportional hazards (PH) and proportional odds models. The study of Banerjee, Wall, and Carlin (2003) employs the PH model with gamma frailties to explain the pattern of infant mortality among counties while accounting for important covariates and spatial correlations. Diva, Dey, and Banerjee (2008) study the joint modeling of multiple cancers while accounting for spatial correlations. Zhao, Hanson, and Carlin (2009) illustrate a flexible spatial frailty lifetime modeling with an analysis of spatially oriented breast cancer data. Pan et al. (2014) study the interval-censored spatial survival under a proportional hazards frailty model. Li et al. (2015) present a semi-parametric PH model for the lifetime of pine trees across 182 sites in the United States. Min et al. (2025) use mixture distributions within the lifetime regression model to study the potential effect of spatial locations on the Titan GPU lifetime with competing risks. In the above studies, the Bayesian framework is often adopted for the model inference. Because the number of spatial locations can be large with over thousands of locations,

CONTACT Xinwei Deng  xdeng@vt.edu  Department of Statistics, Virginia Tech, Blacksburg, VA 24061.

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/00224065.2026.2633083>.

© 2026 American Society for Quality

the posterior distribution is often intractable. Markov chain Monte Carlo (MCMC) approaches are used in many studies to generate samples from the posterior. Recently, Hamiltonian Monte Carlo (HMC), as an efficient MCMC method in dealing with high-dimensional distributions, receives great attention. Various algorithms, such as No-U-Turn Sampler (NUTS), are deployed for HMC Hoffman and Gelman (2014).

With a massive amount of data and complex models, the computational cost of HMC can, however, limit its applications. Thus, we are motivated to investigate an approximation inference approach, namely variational inference (VI), for spatial lifetime model inferences. VI is an optimization-based approach to approximate intractable distributions. VI approximates the complicated posterior with a simple variational distribution by minimizing the metrics that quantify the closeness between two distributions. Commonly used metrics include the Kullback-Leibler (KL) divergence, the α -divergence, and the χ -divergence (Wan, Li, and Hovakimyan 2020). Blei, Kucukelbir, and McAuliffe (2017) provide a thorough review of VI. Minimizing the KL divergence between the posterior and a variational distribution that belongs to a mean-field family is a standard VI procedure in many studies. But the parameter independence assumption of the mean-field family limits the flexibility of the variational distribution, and the KL divergence can lead to the underestimation of the posterior variance. Structured VI that allows dependencies between parameters is one direction to introduce flexibility to variational distributions. For example, Hoffman and Blei (2015) allow arbitrary dependencies between global and local hidden variables for better parameter estimations. Kucukelbir et al. (2017) consider Gaussian variational families with non-diagonal covariance. Ranganath, Tran, and Blei (2016) propose hierarchical Gaussian process priors to allow the correlation between parameters through a more flexible and structured approximation distribution. The other direction is to use different divergence metrics, such as the α -divergence as in Hernandez-Lobato et al. (2016) to allow flexible behavior of the variational distribution.

Some articles apply VI to lifetime data inference. Kim and Pavlovic (2018) apply VI to a Gaussian process survival analysis model, which is an extension of the Cox PH model. Boškosi et al. (2021) use variational Bayes to estimate parameters in an unemployment problem. Jung and Gerstung (2021) propose a population-level inference in electronic health records with stochastic VI. Xiu, Tao, and Henao (2020) study

the variational inference capability of model fitting with individual survival analysis. However, these works did not consider the VI for lifetime data with spatial correlations. Thus, it is worth investigating the capability and computational efficiency of VI in spatially correlated lifetime data and conducting the performance comparison between HMC and variational inference.

In this article, motivated by the importance of computational efficiency, we focus on investigating the performance of VI for spatial lifetime models regarding computing time and inference capability. The rest of this article is organized as follows. Section 2 provides a detailed description of the Titan GPU dataset and the pine tree lifetime dataset. Section 3 introduces the models used for spatially correlated lifetime data as well as a general VI inference procedure. The numerical results are presented in Section 4. Finally, Section 5 summarizes the contribution and conclusion of this work.

2. Data collection and processing

2.1. Titan GPU lifetime data

The original Titan GPU lifetime dataset is publicly available in Ostrouchov et al. (2020), where the serial numbers, locations, failure type, and life spans are extracted from inventory runs and failure event log records from 2012 to 2019. During the service time of the Titan supercomputer, a GPU node can fail multiple times due to different reasons. In this study, since our goal is to investigate the computational efficiency of VI, we only take into account the first error time of each GPU node. The modified version of the dataset as described in Min et al. (2025) is adopted. In particular, one main failure type, off the bus (OTB), is focused on. GPU nodes that failed due to OTB are recorded as failed, and other failure modes, including censored nodes, are recorded as censored.

In the Titan supercomputer, GPU nodes are placed in 8×25 cabinets. There are three cages within each cabinet, eight slots within each cage, and four nodes within each slot. Figure 1 shows the physical structure of the Titan supercomputer. Since Ostrouchov et al. (2020) found that GPU failure is related to the heat dissipation in the room, we consider using the row and column of cabinets as the spatial location of each GPU node, as it can serve as a proxy for the regional temperature. The cage, slot, and node positions are considered as covariates. Table 1 gives an illustration of the information contained in our dataset. Figure 2 presents the failure proportion of GPU nodes at each

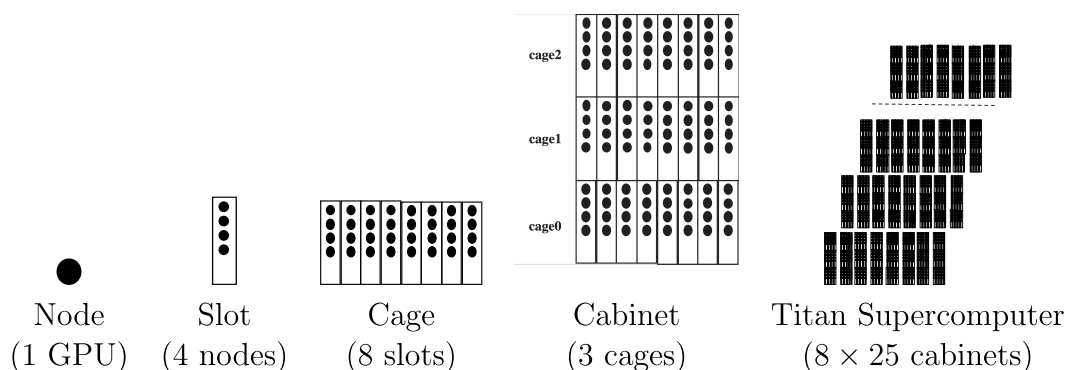


Figure 1. The physical organization of Titan supercomputer.

Table 1. Illustration of Titan GPU lifetime data structure.

SN	Column	Row	Cage	Slot	Node	Failure Type	Failure Time (years)
323512026070	11	4	0	7	2	sensor	5.91
323512026071	18	6	1	3	0	sensor	3.68
323512026072	24	1	0	4	2	sensor	0.44
323512026108	9	1	0	0	2	OTB	4.00
323512026161	12	4	2	6	2	OTB	3.15

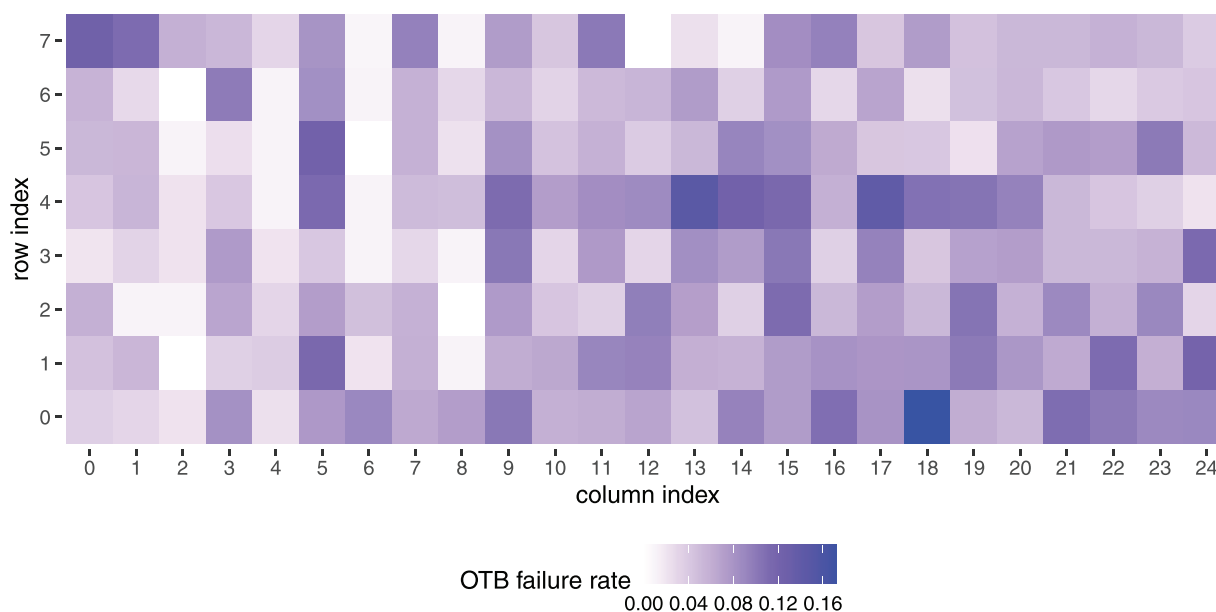


Figure 2. The OTB failure proportion of GPU nodes in different cabinet locations of Titan supercomputer.

cabinet location. It is clear that heterogeneity exists in failure rates at different locations and tends to be similar in adjacent cabinets compared to those far away.

2.2. Pine tree lifetime data

The pine tree dataset is from Li et al. (2015). This dataset tracks the lifetime of trees under varying living conditions and treatments. A total of 42,525 pine trees

across 182 sites are collected throughout the Piedmont and Atlantic Coastal Plain regions. During each tree's lifetime, variables such as total height (TH), diameter at breast height (DBH), and crown class are recorded every three years up to seven times. Other time-invariant variables, such as physical region and thinning treatment (control, light thinning, and heavy thinning) applied to the tree are also collected. In this study, due to the computational limitation, a subset of 60 randomly selected sites with 14,044 trees is used

for analysis. Table 2 presents the exploratory variables and summary statistics of pine trees in the 60 selected sites. The event of interest is the tree mortality. A tree is recorded as censored if it survives till the seventh follow-up period. Figure 3 presents the mortality rate of pine trees across 60 selected sites. To avoid overlapping markers in the visualization, we applied small uniform random displacements to the original coordinates.

3. Methodologies

In Section 3.1, we first introduce the two models, the cumulative exposure model (CEM) and proportional hazards model, that are most appropriate for

Table 2. The description and summary statistics of explanatory variables of 60 selected sites in the pine tree dataset.

Variable	Description	N (%)
SN	A serial number assigned to each tree	14044 (100%)
Thinning Intensity	1-control	3370 (24.00%)
	2-light thinning	5674 (40.40%)
	3-heavy thinning	5000 (35.60%)
Physical Region	1-Coastal Plain	7448 (53.00%)
	2-Piedmont	5261 (37.50%)
	3-other types	1335 (9.50%)
Crown Class	1-dominant	5543 (39.50%)
	2-codominant	9603 (68.40%)
	3-intermediate	3823 (27.20%)
	4-suppressed	1302 (9.30%)
Variable	Description	Mean (SD)
DBH	Diameter at breast height (cm)	7.39 (2.19)
TH	Total height of the tree (meters)	49.36 (13.86)

analyzing the spatially correlated lifetime datasets described in Section 2. Then, in Section 3.2, we present the Bayesian framework for these models and detail the variational inference approach, which incorporates two types of divergences, along with the parameter estimation procedure.

3.1. Modeling spatially correlated lifetime data

We first introduce some general notations for both models. Suppose there are m distinct locations, s_1, \dots, s_m . Let t_{ij} be the time-to-event for the j th unit in the i th location s_i , where $i = 1, \dots, m, j = 1, \dots, n_i$. Here n_i is the number of units in location s_i , and the total number of units is $n = \sum_{i=1}^m n_i$. Let δ_{ij} be the corresponding censoring indicator. Set $\delta_{ij} = 1$ if the unit failed and otherwise $\delta_{ij} = 0$. Denote $\mathbf{x}_{ij}(t)$ to be the p -dimensional vector of explanatory variables at time t at location i . Here, we consider a general notation of the covariates vector $\mathbf{x}_{ij}(t)$, which is a function of time t . For a time-invariant variable, the corresponding elements of $\mathbf{x}_{ij}(t)$ can be set to constants. For categorical variables, a dummy variable coding can be implemented. The time-to-event data collection can be denoted by $\mathcal{D} = \{t_{ij}, \delta_{ij}, \mathbf{x}_{ij}(t) : t \leq t_{ij}, i = 1, \dots, m, j = 1, \dots, n_i\}$.

3.1.1. Cumulative exposure model (CEM)

For the GPU dataset, the problem of interest is to model the time-to-event of GPU nodes using their physical architecture covariates and the spatial

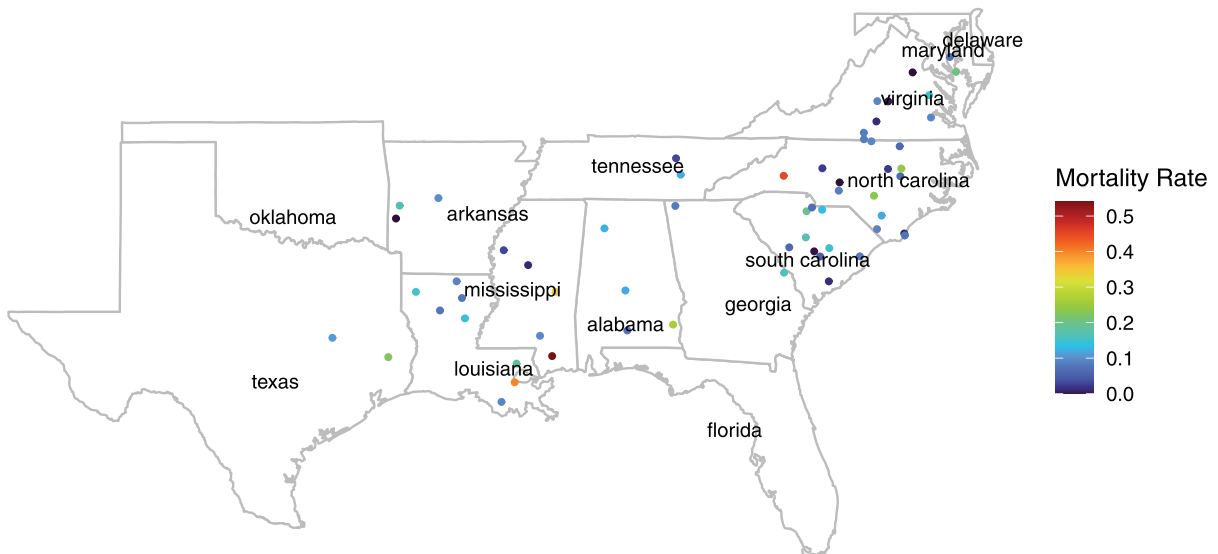


Figure 3. The mortality rate of pine trees across 60 selected U.S. sites. To avoid the overlapping between points, jitters with a uniformly distributed random noise are added to the original site location.

locations. Accelerated failure time (AFT) model, as a special case of the CEM, is in an additive form that can easily incorporate multiple covariates and random effects. Newby (1988) and Meeker, Escobar, and Pascual (2022) demonstrate that AFT can properly model the mechanism of failure. The model form of AFT assumes that the mechanism of failure does not change with the covariates, but is simply accelerated or slowed down. The model's applicability in engineering is further supported by studies such as Hong and Meeker (2010), Kundu, Darpe, and Kulkarni (2019), and Min et al. (2025). Therefore, in this study, the AFT model is considered for the GPU dataset.

For presentation convenience, we first describe the general model, CEM, which can handle both time-varying and time-invariant covariates. We then demonstrate how the CEM can be simplified into an AFT model. The CEM describes the cumulative damage level of the product by a certain time t given time-varying covariates $\mathbf{x}(t)$ (Hong and Meeker 2013). Let $\boldsymbol{\beta}$ be the coefficient vector of covariates. The cumulative exposure $u_{ij}(t)$ of unit j in location i by time t is defined as

$$u_{ij}(t) = \int_0^t \exp\left[-\mathbf{x}_{ij}(s)^\top \boldsymbol{\beta}\right] ds.$$

The unit fails at time T when the amount of cumulative exposure reaches a random threshold U . Let T_{ij} be the random time-to-event of the j th unit in the i th location. Usually, the log cumulative exposure $\log[u_{ij}(T_{ij})]$ is assumed to follow a location-scale distribution. When considering the spatial random effect, the model can be written as

$$\log[u_{ij}(T_{ij})] = \mu + \gamma_i + \sigma\epsilon_{ij}, \quad [1]$$

where γ_i is the spatial random parameter that represents the random effect of the i th location, and ϵ_{ij} is independent and identically distributed and follows a standard location-scale distribution. Here, μ and σ are the corresponding location and scale parameters. Let $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_m)^\top$ be the spatial random parameter vector. The random effect vector $\boldsymbol{\gamma}$ is assumed to follow a multivariate normal distribution $\boldsymbol{\gamma} \sim \text{MVN}(\boldsymbol{\theta}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} = s_\gamma^2 \boldsymbol{\Omega}$. Here, s_γ^2 describes the overall spatial variability and $\boldsymbol{\Omega} = (\rho_{i,i'})_{m \times m}$ is the correlation matrix. The (i, i') th element of the correlation matrix $\rho_{i,i'}$ describes the spatial correlation between the random effect of location s_i and $s_{i'}$. Commonly used correlation functions include exponential, Gaussian, and Matérn correlations. In this article, we use the exponential correlation function $\rho_{i,i'} = \exp[-d(s_i, s_{i'})/\nu]$,

$\nu > 0$, where $d(s_i, s_{i'})$ is the Euclidean distance between locations s_i and $s_{i'}$.

Denote the model parameter vector as $\boldsymbol{\theta} = (\mu, \boldsymbol{\beta}^\top, \boldsymbol{\gamma}^\top, \sigma, s_\gamma^2, \nu)^\top$. Given the distributional assumptions, the likelihood of the model parameters is

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}|\mathcal{D}) &= L(\boldsymbol{\gamma}, s_\gamma^2, \nu) L(\mu, \boldsymbol{\beta}, \sigma|\mathcal{D}, \boldsymbol{\gamma}) \quad [2] \\ &= (2\pi)^{-m/2} |\boldsymbol{\Sigma}|^{-1/2} \exp\left(-\frac{1}{2} \boldsymbol{\gamma}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma}\right) \prod_{ij} \left[\frac{1}{\sigma t_{ij}} \phi(z_{ij})\right]^{\delta_{ij}} \\ &\quad [1 - \Phi(z_{ij})]^{1-\delta_{ij}}, \end{aligned}$$

where

$$z_{ij} = \frac{\log(u_{ij}) - \gamma_i}{\sigma}.$$

Here, $u_{ij} = u(t_{ij})$ is the cumulative exposure level at the time-to-event t_{ij} of j th unit in the i th location, and $\phi(\cdot)$ and $\Phi(\cdot)$ are the probability density function (pdf) and cumulative distribution function (cdf) of standard location-scale distribution, respectively.

In the GPU dataset, the covariates are time-invariant. Therefore, the AFT model, a special case of CEM, is considered for analyzing the dataset. When the covariates are time-invariant, the cumulative exposure becomes

$$u(t) = \int_0^t \exp\left(-\mathbf{x}_{ij}^\top \boldsymbol{\beta}\right) ds = \exp\left(-\mathbf{x}_{ij}^\top \boldsymbol{\beta}\right) t.$$

The CEM model in Eq. [1] can be simplified to an AFT model:

$$\log[u_{ij}(T_{ij})] = -\mathbf{x}_{ij}^\top \boldsymbol{\beta} + \log(T_{ij}) = \mu + \gamma_i + \sigma\epsilon_{ij}, \quad [3]$$

which is equivalent to

$$\log(T_{ij}) = \mu + \mathbf{x}_{ij}^\top \boldsymbol{\beta} + \gamma_i + \sigma\epsilon_{ij}.$$

The z_{ij} term in the likelihood function Eq. [2] becomes $z_{ij} = [\log(t_{ij}) - \mathbf{x}_{ij}^\top \boldsymbol{\beta} - \gamma_i]/\sigma$.

3.1.2. Proportional hazards model with time-dependent covariates

For the pine tree survival data, how the survival rate of pine trees is impacted by the thinning treatment, environmental factors across different regions is of interest. The PH model Cox (1972) provides a suitable framework for analyzing such lifetime data. This model is well-suited for the pine tree dataset because it accommodates time-varying covariates (e.g., DBH and TH) as well as random effects. Additionally, it allows for direct estimation of pine tree survival rates. The PH model assumes that the covariates as well as random effects are multiplicatively related to the hazard. The hazard function of the j th unit in the i th location is modeled as

$$h_{ij}(t) = h_0(t) \exp \left[\mathbf{x}_{ij}(t)^\top \boldsymbol{\beta} + \gamma_i \right], \quad [4]$$

where $h_0(t)$ is the baseline hazard function, $\boldsymbol{\beta}$ is the coefficient vector for covariates, and γ_i is the spatial random effect at location s_i . Then the cumulative hazard function can be expressed as

$$H_{ij}(t) = \int_0^t h_{ij}(s) ds = \exp(\gamma_i) \int_0^t h_0(s) \exp \left[\mathbf{x}_{ij}(s)^\top \boldsymbol{\beta} \right] ds.$$

The survival function of the j th unit in the i th location can be obtained by $S_{ij}(t) = \exp[-H_{ij}(t)]$. In this article, to accommodate the Bayesian framework, we consider a parametric form of the baseline hazard $h_0(t; \boldsymbol{\theta}_h)$, where $\boldsymbol{\theta}_h$ is the parameter vector. Commonly used hazard functions, like Weibull or log-normal hazard function can be considered. Then the model parameter can be written as $\boldsymbol{\theta} = (\boldsymbol{\beta}^\top, \boldsymbol{\gamma}^\top, \boldsymbol{\theta}_h^\top, s_\gamma^2, \nu)^\top$. The likelihood of the model parameter is

$$\begin{aligned} L(\boldsymbol{\theta}|\mathcal{D}) &= L(\boldsymbol{\gamma}, s_\gamma^2, \nu) L(\boldsymbol{\beta}, \boldsymbol{\theta}_h | \mathcal{D}, \boldsymbol{\gamma}) \\ &= (2\pi)^{-m/2} |\boldsymbol{\Sigma}|^{-1/2} \exp \left(-\frac{1}{2} \boldsymbol{\gamma}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\gamma} \right) \\ &\quad \prod_{ij} [h_{ij}(t_{ij})]^{\delta_{ij}} S_{ij}(t_{ij}). \end{aligned}$$

3.2. Variational inference for spatially correlated lifetime data

For statistical inferences of the model parameter $\boldsymbol{\theta}$, priors are assigned to $\boldsymbol{\theta}$ under the Bayesian framework. In the above models, since the scale parameter s_γ^2 and length-scale parameter ν in the spatial covariance matrix are positive, inverse-gamma priors are considered in this study. Non-informative priors are considered for the covariate parameter vector $\boldsymbol{\beta}$. The parameters μ and $\sigma_l = \log(\sigma)$ in the spatial AFT model and $\boldsymbol{\theta}_h$ in the spatial PH model are also assumed to follow non-informative priors. The Bayesian framework for the spatial lifetime model can be expressed as:

$$\begin{aligned} \boldsymbol{\gamma} &\propto \text{MVN}(\mathbf{0}, s_\gamma^2 \boldsymbol{\Omega}), \\ s_\gamma^2 &\propto \text{IGAM}(a_\sigma, b_\sigma), \\ \nu &\propto \text{IGAM}(a_\nu, b_\nu), \\ \boldsymbol{\beta}_p &\propto \mathbf{I}_p. \end{aligned}$$

For spatial AFT:

$$\begin{aligned} \log [u_{ij}(T_{ij})] &= \mu + \gamma_i + \sigma \epsilon_{ij}, \\ \mu &\propto 1, \quad \sigma_l \propto 1. \end{aligned}$$

For spatial PH:

$$\begin{aligned} h_{ij}(t) &= h_0(t; \boldsymbol{\theta}_h) \exp \left[\mathbf{x}_{ij}(t)^\top \boldsymbol{\beta} + \gamma_i \right], \\ \boldsymbol{\theta}_h &\propto 1. \end{aligned}$$

Here, IGAM stands for inverse-gamma distribution. The parameter a_σ, b_σ and a_ν, b_ν are the shape and scale parameters for s_γ^2 and ν , respectively.

Denote the priors for $\boldsymbol{\theta}$ by $p(\boldsymbol{\theta})$. The posterior distribution is proportional to the joint density function of the data collection \mathcal{D} and model parameters $\boldsymbol{\theta}$, which is

$$\begin{aligned} p(\boldsymbol{\theta}|\mathcal{D}) &\propto p(\mathcal{D}|\boldsymbol{\theta})p(\boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{\theta}|\mathcal{D})p(\boldsymbol{\theta}) \\ &= \mathcal{L}(\boldsymbol{\theta}|\mathcal{D})p(s_\gamma^2)p(\nu). \end{aligned} \quad [5]$$

Instead of using MCMC methods to draw samples from the posterior, in this study, we focus on how to use VI for model inference. The key idea of VI is to use a variational distribution $q(\boldsymbol{\theta}|\boldsymbol{\eta})$ to approximate the exact posterior. Here $\boldsymbol{\eta}$ is the parameter vector in the variational probability distribution. One would like to find $\boldsymbol{\eta}$ that allows the variational distribution $q(\boldsymbol{\theta}|\boldsymbol{\eta})$ to be as close to the exact posterior $p(\boldsymbol{\theta}|\mathcal{D})$ as possible for better approximation. Usually, divergence metrics are used to evaluate the distance between $p(\boldsymbol{\theta}|\mathcal{D})$ and $q(\boldsymbol{\theta}|\boldsymbol{\eta})$, and the optimal $\boldsymbol{\eta}$ is obtained by minimizing the distribution distance. Given the estimated $\boldsymbol{\eta}$, one can sample model parameters from the variational distribution $q(\boldsymbol{\theta}|\boldsymbol{\eta})$ and conduct inferences as in MCMC methods. Note that VI reformulates the Bayesian inference problem into an optimization problem. So the computational cost of VI depends on the complexity of the optimization problem. As it is often the case, the variational distribution $q(\boldsymbol{\theta}|\boldsymbol{\eta})$ is chosen to be relatively simple to gain the computational advantages. While this simplification enables faster computation compared to MCMC, it may introduce approximation bias since the variational family may not fully capture the true posterior. Nevertheless, for spatially correlated lifetime data, VI remains valuable for efficient explorations of complex models.

3.2.1. Two types of divergence

One important aspect of VI is the choice of the divergence metric used to measure the distance between the true posterior $p(\boldsymbol{\theta}|\mathcal{D})$ and the approximate posterior $q(\boldsymbol{\theta}|\boldsymbol{\eta})$. KL divergence is the most commonly used metric in VI estimation procedure. However, for some complex model and data structures, the KL divergence may not be flexible enough to capture the true posterior's characteristics. Recent advances in VI have

extended the KL divergence to richer divergence families, such as Rényi's α -divergence, to better approximate the exact posterior distribution. Therefore, in this study, we are interested in the inference performance of both KL divergence and α -divergence.

The KL divergence between $p(\theta|\mathcal{D})$ and $q(\theta|\eta)$ is defined as:

$$\begin{aligned} \text{KL}[q(\theta|\eta)||p(\theta|\mathcal{D})] &= \int q(\theta|\eta) \log \left[\frac{q(\theta|\eta)}{p(\theta|\mathcal{D})} \right] d\theta \quad [6] \\ &= \mathbf{E}_{q(\theta|\eta)} \{ \log [q(\theta|\eta)] \} \\ &\quad - \mathbf{E}_{q(\theta|\eta)} \{ \log ([p(\mathcal{D}, \theta)]) \} + \log [p(\mathcal{D})]. \end{aligned}$$

Because $\log[p(\mathcal{D})]$ is intractable, we cannot directly minimize KL divergence. An equivalent optimization problem is to maximize the evidence lower bound (ELBO):

$$\begin{aligned} \mathcal{L}_{\text{VI}} &= \log [p(\mathcal{D})] - \text{KL}(q(\theta|\eta)||p(\theta|\mathcal{D})) \\ &= \mathbf{E}_{q(\theta|\eta)} \left\{ \log \left[\frac{p(\theta, \mathcal{D})}{q(\theta|\eta)} \right] \right\}. \quad [7] \end{aligned}$$

One characteristic of the KL divergence is that it is an asymmetric measurement. That means $\text{KL}[q(\theta|\eta)||p(\theta|\mathcal{D})] \neq \text{KL}[p(\theta|\mathcal{D})||q(\theta|\eta)]$. The choice of KL divergence direction affects the behavior of the variational distribution.

Rényi's α -divergence is a more general metric that measures the closeness of $p(\theta|\mathcal{D})$ and $q(\theta|\eta)$. It is defined as

$$D[q(\theta|\eta)||p(\theta|\mathcal{D})] = \frac{1}{\alpha - 1} \log \left[\int q(\theta|\eta)^\alpha p(\theta|\mathcal{D})^{1-\alpha} d\theta \right],$$

where $\alpha > 0, \alpha \neq 1$. When $\alpha \rightarrow 1$, the α -divergence approaches the KL divergence. So the α -divergence can be considered as the root family of KL divergence.

Similar to the KL divergence, we can not compute the α -divergence directly. Alternatively, we can maximize the variational Rényi (VR) bound, which is defined as

$$\begin{aligned} \mathcal{L}_\alpha &= \log [p(\mathcal{D})] - D[q(\theta|\eta)||p(\theta|\mathcal{D})] \\ &= \frac{1}{1 - \alpha} \log \left\{ \mathbf{E}_{q(\theta|\eta)} \left[\left(\frac{p(\theta, \mathcal{D})}{q(\theta|\eta)} \right)^{1-\alpha} \right] \right\}. \quad [8] \end{aligned}$$

Unlike the original Rényi's α -divergence's positive support, the VR lower bound can be extended to $\alpha \leq 0$. Besides, the objective \mathcal{L}_α is continuous and non-increasing on α (Li and Turner 2016). When $\alpha \rightarrow -\infty$, it encourages mass-covering behavior. That means θ is encouraged to explore a wider range of possible values. When $\alpha \rightarrow \infty$, it encourages the zero-

forcing behavior. That means the approximation tends to return an over-confident θ estimate with a small variance. By adjusting the α value, we can encourage flexible variational distribution.

3.2.2. Black-Box VI parameter estimation procedure

For some specific models with convenient conditional distributions, the VI with KL divergence parameter estimation can be obtained *via* a closed-form coordinate ascent algorithm (Ranganath, Gerrish, and Blei 2014). For generic models and α -divergence, the VR bound \mathcal{L}_α is often intractable due to the model complexity and high-dimensional integration over θ . Therefore, we adopt the black-box VI idea as presented in Hernandez-Lobato et al. (2016) and Ranganath, Gerrish, and Blei (2014), and use a Monte Carlo (MC) method to evaluate the VR bound. Since KL divergence is a special case of α -divergence with $\alpha \rightarrow 1$, in the following, we focus on the VR bound \mathcal{L}_α . Specifically, we randomly sample from the variational distribution $\theta_k \sim q(\theta|\eta), k = 1, \dots, h$, and \mathcal{L}_α in Eq. 8 can be approximated by using the sample mean to substitute the expectation:

$$\hat{\mathcal{L}}_\alpha = \frac{1}{1 - \alpha} \log \left\{ \frac{1}{h} \left[\sum_{k=1}^h \left(\frac{p(\theta_k, \mathcal{D})}{q(\theta_k|\eta)} \right)^{1-\alpha} \right] \right\}. \quad [9]$$

Note that Eq. [9] is a general expression that works for both KL divergence and α -divergence. After obtaining the approximated VR bound, $\hat{\mathcal{L}}_\alpha$ can be optimized and updated iteratively to obtain variational parameter estimates η^* . We optimize $\hat{\mathcal{L}}_\alpha$ using the Adam algorithm, which is an extension to the stochastic gradient descent method (Kingma and Ba 2014). The derivation for the gradient of the VR bound is shown below:

$$\begin{aligned} \frac{\partial \mathcal{L}_\alpha}{\partial \eta} &= \frac{\partial \mathcal{L}_\alpha}{\partial \log [q(\theta|\eta)]} \frac{\partial \log [q(\theta|\eta)]}{\partial \eta} \\ &= \left(\frac{1}{1 - \alpha} \frac{\partial}{\partial \log [q(\theta|\eta)]} \log \left\{ \int q(\theta|\eta) \left[\frac{p(\theta, \mathcal{D})}{q(\theta|\eta)} \right]^{1-\alpha} d\theta \right\} \right) \\ &\quad \frac{\partial \log [q(\theta|\eta)]}{\partial \eta} \\ &= \frac{\alpha}{1 - \alpha} \frac{\int q(\theta|\eta) \left[\frac{p(\theta, \mathcal{D})}{q(\theta|\eta)} \right]^{1-\alpha} \frac{\partial \log [q(\theta|\eta)]}{\partial \eta} d\theta}{\int q(\theta|\eta) \left[\frac{p(\theta, \mathcal{D})}{q(\theta|\eta)} \right]^{1-\alpha} d\theta}. \end{aligned}$$

Algorithm 1 describes a general approach to obtain η^* . After getting the variational parameter estimate η^* , we can sample model parameters from $\theta^l \sim q(\theta|\eta^*), l = 1, \dots, L$, and use the sample mean $\sum_{l=1}^L \theta^l / L$ as the point estimates. The $100(1 - \alpha)\%$

credible intervals (CIs) can be obtained from the $\alpha/2$ and $(1 - \alpha/2)$ quantiles.

Algorithm 1. Obtain parameter estimation through black-box VI

Require: $q(\theta|\eta)$: the variational distribution and τ : the stopping threshold.

Require: η_0 : initial variational parameter vector.

$r \leftarrow 1$ (initialize iteration number);

$\eta \leftarrow \eta_0$ (initialize variational parameter vector);

stop = 0 (initialize the stopping indicator);

1: **while** stop = 0 **do**

2: Take h samples from the variational distribution $\theta_k \sim q(\theta|\eta_{r-1}), k = 1, \dots, h$ and compute a stochastic estimate of $\hat{\mathcal{L}}_\alpha$ as in Eq. [9].

3: Take a gradient descent step in Adam algorithm to obtain the updated parameter estimate η_r ;

4: **if** $r > 2$ & $(|\hat{\mathcal{L}}_\alpha^r - \hat{\mathcal{L}}_\alpha^{r-1}|/|\hat{\mathcal{L}}_\alpha^r| < \tau$ &

$|\hat{\mathcal{L}}_\alpha^{r-1} - \hat{\mathcal{L}}_\alpha^{r-2}|/|\hat{\mathcal{L}}_\alpha^{r-1}| < \tau)$ **then**

5: stop = 1

6: **end if**

7: $r \leftarrow r + 1$

8: **end while**

9: **return** $\eta^* \leftarrow \eta_r$

4. Numerical studies

In this section, we investigate the inference performance of VI using both KL divergence and α -divergence on the GPU and pine tree datasets. Section 4.1 first introduces the benchmark method, HMC, used for comparison. Then the computational efficiency and inference capability of VI compared to HMC are evaluated with two case studies in Sections 4.2 and 4.3, respectively.

4.1. The benchmark method: HMC

As mentioned in Section 1, MCMC is the most often used method for drawing samples from the posterior distribution (Li and Meeker 2014). Among various MCMC techniques, HMC utilizes Hamiltonian dynamics to rapidly traverse the parameter space and enhance MCMC mixing (Neal 2011). As an improvement on HMC, NUTS eliminates the manual parameter tuning step and is extensively applied in Bayesian model inference, including Min et al. (2025). Thus, NUTS in Stan (Carpenter et al. 2017) is considered as a benchmark inference method in this article. When implementing the log-likelihood functions, positive model parameters are transformed to the range of

$(-\infty, \infty)$ and the random effect design matrix is transformed to avoid co-linearity as presented in Min et al. (2025). The chain length and tree depth parameters can be determined by running short chains and checking convergence first.

In the case studies, we compare HMC, VI with KL divergence, and VI with α -divergence using the following three metrics: the Cox-Snell residual probability plot, the negative log-likelihood (NLL) value, and the computing time. The first two metrics are used to evaluate the inference capability of VI compared to HMC, while the computing time is used to evaluate the computational efficiency.

4.2. Case study 1: The Titan GPU lifetime data

4.2.1. Bayesian AFT model for GPU dataset

In the GPU dataset, we mainly focused on a more “global” room temperature between cabinets that causes the spatial correlations. Within each cabinet, the temperature is assumed to be relatively stable compared to variations observed across cabinets that are farther apart. Although researchers from the Oak Ridge National Laboratory found that GPU located at lower cage level may have relatively higher survival probability due to the cooling airflow from bottom to top (Ostrouchov et al. 2020), in this study, instead of directly modeling the impact of within-cabinet correlation, we modeled the impact of temperature across cage levels as a fixed effect. That is, we considered the cage level as a categorical variable in the model. Similarly, slot and node are also considered as categorical variables. Consider a dummy variable encoding, then there are 12 dummy variable covariates. Denote covariate as $\mathbf{x} = (1, \mathbf{x}_{\text{cage}}^\top, \mathbf{x}_{\text{slot}}^\top, \mathbf{x}_{\text{node}}^\top)^\top$, where \mathbf{x}_{cage} represents the vector of dummy variable encoding of cage variable and similar for \mathbf{x}_{slot} and \mathbf{x}_{node} . The spatial AFT model of the j th GPU in the i th location can be written as

$$\log(T_{ij}) = \mathbf{x}_{ij}^\top \boldsymbol{\beta} + \gamma_i + \sigma \epsilon_{ij}.$$

For the GPU data, we assume the error term ϵ_{ij} follows the standard smallest extreme value (SEV) distribution. According to Section 3.1.1, the spatial random effect γ is assumed to follow a multivariate normal distribution $\gamma \sim \text{MVN}(\mathbf{0}, \Sigma)$, where $\Sigma = s_\gamma^2 \Omega$ and $(\Omega)_{i\ell} = (\exp[-d(s_i, s_\ell)]/\nu)$. Since the exact physical locations of each cabinet in the server room are unavailable, we use the row and column indices as proxies for spatial location coordinates. To have a stable estimation of ν , we standardized the distance along the row and column direction, respectively, into

the range [0,1] and then computed the Euclidean distance as the distance between two cabinets. This approach may not provide a perfect approximation of the true spatial positions; nevertheless, it offers a practical way to leverage the available dataset to construct spatial structures.

Denote the model parameter as $\theta = (\boldsymbol{\beta}^\top, \boldsymbol{\gamma}^\top, \sigma_l, s_\gamma^2, \nu)^\top$. Then the log posterior can be derived as

$$\begin{aligned} \log [p(\theta|\mathcal{D})] &= \text{Constant} \\ &+ \sum_{ij} (\delta_i \{-\log(t_{ij}) - \sigma_l + \log[\phi(z_{ij})]\} + (1 - \delta_i) \log[1 - \Phi(z_{ij})]) \\ &- \frac{1}{2} \left[\log(s_\gamma^2 \Omega) + \frac{\boldsymbol{\gamma}^\top \Omega^{-1} \boldsymbol{\gamma}}{s_\gamma^2} \right] - \frac{1}{2} (a_\sigma + 1) \log(s_\gamma^2) \\ &- \frac{b_\sigma}{s_\gamma^2} - \frac{1}{2} (a_\nu + 1) \log(\nu) - \frac{b_\nu}{\nu}. \end{aligned}$$

In the GPU dataset, for the hyper-parameters inside the inverse-gamma prior, we set a weak prior for s_γ^2 by letting $a_\sigma = 1$, $b_\sigma = 1$, and a strong prior for ν by setting $a_\nu = 13$, $b_\nu = 0.1$.

For the variation inference of model parameters, we assume the following variational distributional assumptions:

$$\begin{aligned} \boldsymbol{\beta} &\sim \text{MVN}(\boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta), \text{ where } \boldsymbol{\Sigma}_\beta = \text{diag}(\boldsymbol{\sigma}_\beta^2), \\ \boldsymbol{\gamma} &\sim \text{MVN}(\boldsymbol{\mu}_\gamma, \boldsymbol{\Sigma}_\gamma), \text{ where } \boldsymbol{\Sigma}_\gamma = \text{diag}(\boldsymbol{\sigma}_\gamma^2), \\ \sigma_l &\sim \text{N}(\mu_\sigma, \sigma_\sigma^2), \\ s_\gamma^2 &\propto \text{IGAM}(c_\sigma, d_\sigma), \\ \nu &\propto \text{IGAM}(c_\nu, d_\nu), \end{aligned}$$

where $\boldsymbol{\Sigma}_\beta$ and $\boldsymbol{\Sigma}_\gamma$ are diagonal matrices with diagonal elements being $\boldsymbol{\sigma}_\beta^2$ and $\boldsymbol{\sigma}_\gamma^2$. The variational distributions are independent of each other. In the above framework, c_σ, d_σ and c_ν, d_ν are the scale and shape parameters in inverse-gamma distribution for s_γ^2 and ν , respectively. Let $\boldsymbol{\sigma}_{l\beta} = \log(\boldsymbol{\sigma}_\beta)$, $\boldsymbol{\sigma}_{l\gamma} = \log(\boldsymbol{\sigma}_\gamma)$, and $\sigma_{l\sigma} = \log(\sigma_\sigma)$. Then the variational parameter vector can be denoted as $\boldsymbol{\eta} = (\boldsymbol{\mu}_\beta^\top, \boldsymbol{\sigma}_{l\beta}^\top, \boldsymbol{\mu}_\gamma^\top, \boldsymbol{\sigma}_{l\gamma}^\top, \mu_\sigma, \sigma_{l\sigma}, c_\sigma, d_\sigma, c_\nu, d_\nu)^\top$. The log density of the joint variational distribution is

$$\begin{aligned} \log [q(\theta|\boldsymbol{\eta})] &= \text{Constant} - \frac{1}{2} \sum_j \left[2\sigma_{l\beta,j} + \frac{(\beta_j - \mu_{\beta,j})^2}{\exp(2\sigma_{l\beta,j})} \right] \\ &- \frac{1}{2} \sum_j \left[2\sigma_{l\gamma,j} + \frac{(\gamma_j - \mu_{\gamma,j})^2}{\exp(2\sigma_{l\gamma,j})} \right] - \frac{1}{2} \left[2\sigma_{l\sigma} + \frac{(\sigma_l - \mu_\sigma)^2}{\exp(2\sigma_{l\sigma})} \right] \\ &+ c_\sigma \log(d_\sigma) - \log[\Gamma(c_\sigma)] - (c_\sigma + 1) \log(s_\gamma^2) - \frac{d_\sigma}{s_\gamma^2} \\ &+ c_\nu \log(d_\nu) - \log[\Gamma(c_\nu)] - (c_\nu + 1) \log(\nu) - \frac{d_\nu}{\nu}, \end{aligned}$$

where $\mu_{\beta,j}$ is the j th element of mean vector $\boldsymbol{\mu}_\beta$ and $\sigma_{l\beta,j}$ is the j th element of the log diagonal variance vector $\boldsymbol{\sigma}_\beta$ for $\boldsymbol{\beta}$. Similarly for $\mu_{\gamma,j}$ and $\sigma_{l\gamma,j}$. Also, $\Gamma(\cdot)$ denotes the gamma function. Plugging in $\log [p(\theta, \mathcal{D})]$ and $\log [q(\theta|\boldsymbol{\eta})]$ in Eqs. [7] and [8], we can obtain the objective functions \mathcal{L}_α and \mathcal{L}_{VI} in variational inference. The procedures in Section 3.2.2 can be used to obtain estimates $\boldsymbol{\eta}^*$. Then we can sample θ from $q(\theta|\boldsymbol{\eta}^*)$ and conduct inferences for the spatial AFT model.

4.2.2. Analysis results

For the GPU data, we compare HMC, VI with KL divergence, and VI with α -divergence with $\alpha = 0.8$. To understand how well the posterior samples from each inference approach fit the model, we use the residual plot to conduct diagnostics. The censored Cox-Snell residual of the spatial AFT model is an extension of the standardized residual, which is defined as

$$\hat{c}_{ij} = \exp \left[\frac{\log(t_{ij}) - \mathbf{x}_{ij}^\top \hat{\boldsymbol{\beta}} - \gamma_i}{\hat{\sigma}} \right].$$

For a censored unit, the corresponding residual is also censored. With the SEV distribution assumption, the residuals should approximately follow a Weibull distribution. The probability plots of censored Cox-Snell residuals of three inference methods are presented in Figure 4. We can see that for α -divergence, its probability plot of residual is almost aligned with the standard Weibull distribution, while the KL divergence has some deviation at the tail of residuals.

Table 3 summarizes the NLL values and computational times across methods. VI with α -divergence has an NLL value very close to HMC while using the fastest computation time. Parameter estimates and standard errors (SE) in Table 4 show that α -divergence has estimates closer to HMC benchmarks than KL divergence does. While the Slot effect estimates show slightly greater discrepancies, the absolute value of this effect is small in HMC results, keeping the bias within acceptable limits. From both α -divergence and HMC estimates, we can see that when other variables are fixed, cage 0-1 accelerate a unit's failure compared to the baseline unit with cage 2, while node 0-2 decelerates a unit's failure compared to node 3. The impact of a GPU's slot position is relatively small to its failure.

The following results in Figures 5 and 6 present the estimated result obtained using α -divergence. Figure 5(a) shows the estimated spatial correlation

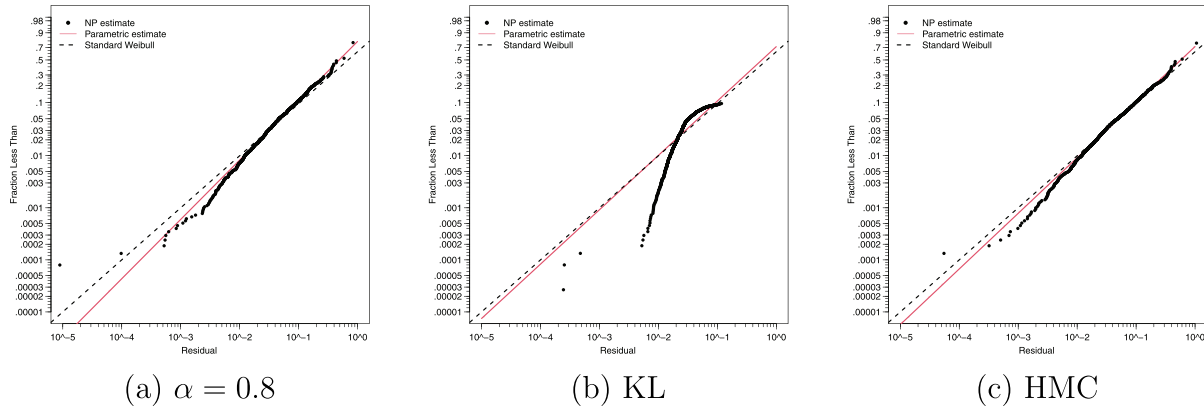


Figure 4. Weibull probability plot of residuals for α -divergence, KL divergence, and HMC with the GPU data.

Table 3. The negative log likelihood and computing time of α -divergence, KL divergence, and HMC inferences with the GPU data.

	$\alpha = 0.8$	KL	HMC
NLL	4481.04	5368.06	4470.13
Time (minutes)	42.7	55.23	86.53

Table 4. The parameter estimates and SE of three inference methods for the GPU data.

Parameter	α -divergence		KL divergence		HMC	
	Estimate	SE	Estimate	SE	Estimate	SE
Constant	2.0235	0.0166	2.5989	0.0279	1.9306	0.0278
Cage 0	0.6922	0.0488	0.1127	0.0483	0.6783	0.0169
Cage 1	0.2755	0.0208	-0.0027	0.0466	0.2734	0.0146
Slot 0	0.0702	0.0284	0.0234	0.0503	0.0412	0.0242
Slot 1	0.0439	0.0337	0.0087	0.0445	0.0285	0.0236
Slot 2	0.0606	0.0348	-0.0221	0.0507	0.0520	0.0240
Slot 3	0.0675	0.0271	0.0147	0.0484	0.0398	0.0237
Slot 4	0.0811	0.0262	0.0330	0.0486	0.0709	0.0251
Slot 5	0.0141	0.0259	-0.0069	0.0487	0.0119	0.0233
Slot 6	-0.0065	0.0289	-0.0340	0.0496	-0.0083	0.0232
Node 0	-0.3049	0.0216	-0.0169	0.0489	-0.2895	0.0192
Node 1	-0.3235	0.0231	-0.0187	0.0447	-0.3102	0.0192
Node 2	-0.0870	0.0279	0.0169	0.0494	-0.0666	0.0201
σ	0.2289	0.0056	0.3866	0.0095	0.2038	0.0053
s^2_γ	0.0940	0.0318	0.0457	0.0136	0.0488	0.0160
ν	0.1337	0.0575	1.8981	3.0627	0.2734	0.2426

from VI with α -divergence as the distance increases. The standardized distance range in the GPU dataset is [0,1.414], with units of “standardized grid cells” for parameter ν . One “standardized grid cell” can be difficult to transform back to the original grid cell. For easier interpretation, Figure 5(b) shows the standardized distances from the cabinet at Row 0, Column 0 to other cabinets. The spatial correlation curve indicates correlations decay to 0.2 at a distance of 0.21, indicating weak correlation. This suggests that cabinets separated by standardized distances greater than 0.21 can be considered weakly correlated or uncorrelated. In original grid cell terms, this corresponds to

cabinets separated by more than one row or five columns. Figure 6 presents the heat map of the estimated random effect obtained using the VI with α -divergence at each location. The heat map shows that the positions in the bottom right corner tend to decelerate the GPU’s failure.

It is also interesting to examine whether the number of spatial locations impacts the computing time for each of the three inference methods. A series of subsets with the number of locations $m = 50, 75, 100, 150$ are randomly sampled from the original dataset without replacements. The process is repeated 10 times and each of the three methods is applied to the subsets. The boxplot of computing time versus the number of locations is presented in Figure 7. We can see that the computing time difference between HMC and variational approaches increases as the number of locations increases. The computational advantage of α -divergence is more obvious when the number of locations is large.

4.3. Case study 2: The pine tree lifetime data

4.3.1. Bayesian PH model for tree data

As introduced in Section 3.1.2, we consider the spatial PH model with time-dependent covariates for the pine tree data. For the PH model in Eq. [4], the baseline hazard function $h_0(t)$ needs to be specified. In this study, we consider the Weibull hazard with functional form $h_0(t) = at^b$ as the baseline function, where a and b are parameters. Let $\mathbf{x}(t) = [\text{DBH}(t), \text{TH}(t), \mathbf{x}_{\text{Crown}}(t), \mathbf{x}_{\text{PhyReg}}, \mathbf{x}_{\text{Trt}}]$, where $\mathbf{x}_{\text{Crown}}(t)$ represent the dummy variable encoding of the crown class at time t , similar for $\mathbf{x}_{\text{PhyReg}}, \mathbf{x}_{\text{Trt}}$. Then the PH model can be expressed as

$$h_{ij}(t) = at^b \exp \left[\mathbf{x}_{ij}(t)^\top \boldsymbol{\beta} + \gamma_i \right].$$

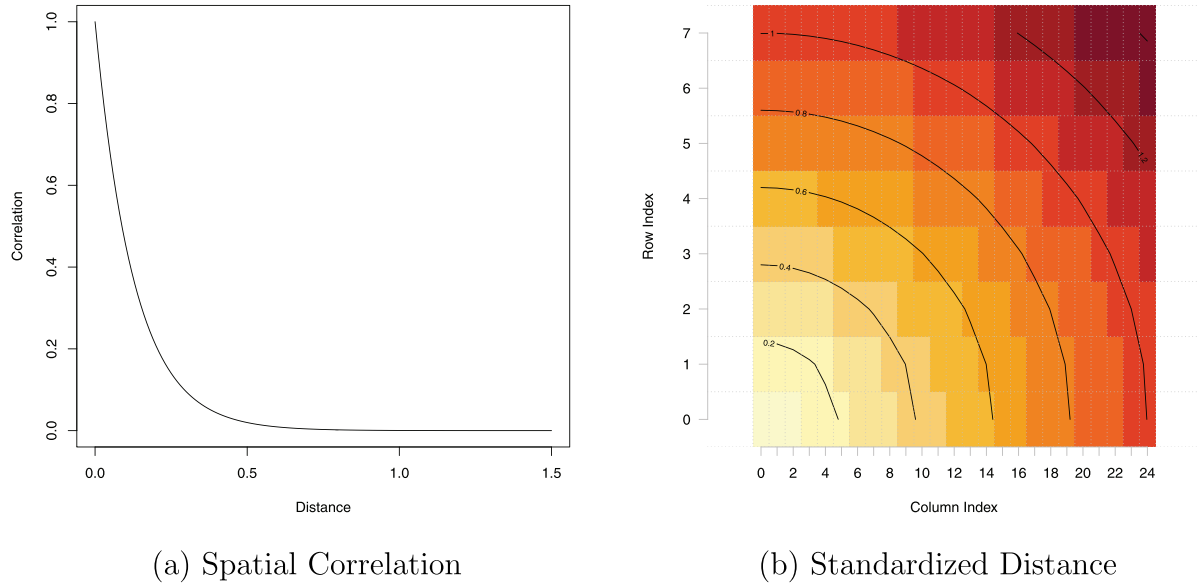


Figure 5. The estimated spatial correlation versus distance from VI with α -divergence, $\alpha = 0.8$ in the AFT model for the GPU data (a) and the “standardized grid cell” distance from the cabinet at Row 0, Column 0 to other cabinets (b).

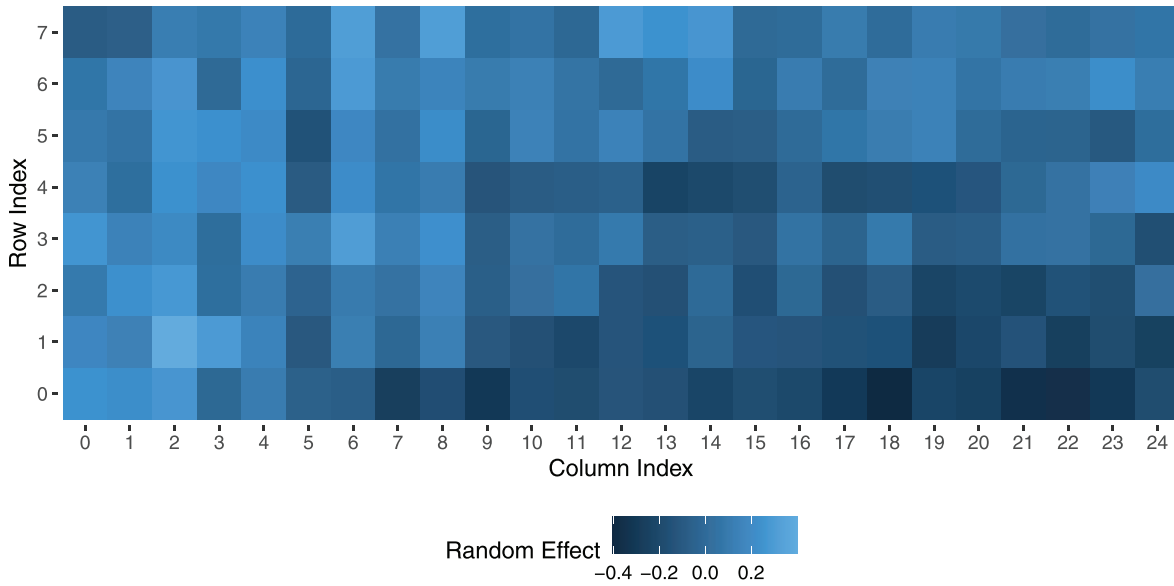


Figure 6. The heat map of estimated random effects from VI with α -divergence, $\alpha = 0.8$ in the AFT model for the GPU data.

Similarly, the MVN distribution is considered for the spatial random effect vector γ . In the pine tree dataset, to have stable estimates of the length-scale parameter ν , we used the great circle distance and standardized it into the range $[0,1]$ by dividing the maximum great circle distance among all location pairs.

Denote the model parameter as $\theta = (a_l, b_l, \beta, \gamma, s_\gamma^2, \nu)$, where $a_l = \log(a), b_l = \log(b)$. The log posterior can be derived as

$$\begin{aligned} \log [p(\theta|\mathcal{D})] &= \text{Constant} \\ &+ \sum_{ij} \left\{ \delta_{ij} \left[\gamma_i + a_l + b_l \log(t) + \mathbf{x}_{ij}(t_{ij})^\top \beta \right] - H_{ij}(t_{ij}) \right\} \\ &- \frac{1}{2} \left[\log(s_\gamma^2 \Omega) + \frac{\gamma^\top \Omega^{-1} \gamma}{s_\gamma^2} \right] \\ &- \frac{1}{2} (a_\sigma + 1) \log(s_\gamma^2) - \frac{b_\sigma}{s_\gamma^2} - \frac{1}{2} (a_\nu + 1) \log(\nu) - \frac{b_\nu}{\nu}. \end{aligned}$$




Figure 7. The computing time of three inference methods versus the number of locations with the GPU data.

In the pine tree dataset, for the hyper-parameters inside the inverse-gamma prior, we set a weak prior for s_γ^2 by letting $a_\sigma = 1$, $b_\sigma = 1$ and a strong prior for ν by setting $a_\nu = 5$, $b_\nu = 0.1$. We assume the following variational distributional assumptions:

$$\begin{aligned} \boldsymbol{\beta} &\sim \text{MVN}(\boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta), \text{ where } \boldsymbol{\Sigma}_\beta = \text{diag}(\boldsymbol{\sigma}_\beta^2), \\ \boldsymbol{\gamma} &\sim \text{MVN}(\boldsymbol{\mu}_\gamma, \boldsymbol{\Sigma}_\gamma), \text{ where } \boldsymbol{\Sigma}_\gamma = \text{diag}(\boldsymbol{\sigma}_\gamma^2), \\ a_l &\sim \text{N}(\mu_a, \sigma_a^2), \\ b_l &\sim \text{N}(\mu_b, \sigma_b^2), \\ s_\gamma^2 &\propto \text{IGAM}(c_\sigma, d_\sigma), \\ \nu &\propto \text{IGAM}(c_\nu, d_\nu), \end{aligned}$$

where $\boldsymbol{\Sigma}_\beta$ and $\boldsymbol{\Sigma}_\gamma$ are diagonal matrices with diagonal elements being $\boldsymbol{\sigma}_\beta^2$ and $\boldsymbol{\sigma}_\gamma^2$. The variational distributions are independent to each other. Let $\boldsymbol{\sigma}_{l\beta} = \log(\boldsymbol{\sigma}_\beta)$, $\boldsymbol{\sigma}_{l\gamma} = \log(\boldsymbol{\sigma}_\gamma)$, $\sigma_{la} = \log(\sigma_a)$, and $\sigma_{lb} = \log(\sigma_b)$. Then the variational parameter vector can be denoted as $\boldsymbol{\eta} = (\boldsymbol{\mu}_\beta^\top, \boldsymbol{\sigma}_{l\beta}^\top, \boldsymbol{\mu}_\gamma^\top, \boldsymbol{\sigma}_{l\gamma}^\top, \mu_a, \sigma_{la}, \mu_b, \sigma_{lb}, c_\sigma, d_\sigma, c_\nu, d_\nu)^\top$. The log density of joint variational distribution is

$$\begin{aligned} \log [q(\boldsymbol{\theta}|\boldsymbol{\eta})] &= \text{Constant} - \frac{1}{2} \sum_j \left[2\sigma_{l\beta,j} + \frac{(\beta_j - \mu_{\beta,j})^2}{\exp(2\sigma_{l\beta,j})} \right] \\ &- \frac{1}{2} \sum_j \left[2\sigma_{l\gamma,j} + \frac{(\gamma_j - \mu_{\gamma,j})^2}{\exp(2\sigma_{l\gamma,j})} \right] \\ &- \frac{1}{2} \left[2\sigma_{la} + \frac{(a_l - \mu_a)^2}{\exp(2\sigma_{la})} \right] - \frac{1}{2} \left[2\sigma_{lb} + \frac{(b_l - \mu_b)^2}{\exp(2\sigma_{lb})} \right] \\ &+ c_\sigma \log(d_\sigma) - \log[\Gamma(c_\sigma)] - (c_\sigma + 1) \log(s_\gamma^2) - \frac{d_\sigma}{s_\gamma^2} \\ &+ c_\nu \log(d_\nu) - \log[\Gamma(c_\nu)] - (c_\nu + 1) \log(\nu) - \frac{d_\nu}{\nu}. \end{aligned}$$

Similar to Section 4.2, we can apply VI based on the unnormalized posterior $\log [p(\boldsymbol{\theta}, \mathcal{D})]$ and variational distribution $\log [q(\boldsymbol{\theta}|\boldsymbol{\eta})]$ and obtain the estimated $\hat{\boldsymbol{\theta}}$.

4.3.2. Analysis results

For a PH model with time-dependent covariates, the Cox-Snell residual is defined as

$$\hat{\epsilon}_{ij} = \int_0^{T_{ij}} \hat{H}_0(t) \exp \left[\mathbf{x}_{ij}(t)^\top \hat{\boldsymbol{\beta}} + \hat{\gamma}_i \right] dt,$$

where $\hat{H}_0(t)$ is the estimated cumulative baseline hazard rate by plugging in \hat{a} and \hat{b} . If the model is correct, then $\hat{\epsilon}_{ij}$ approximately follows an exponential distribution with $\lambda = 1$ and censoring, which is a special case of the Weibull distribution (Klein and Moeschberger 2003). Therefore, we can use Weibull probability plot of residuals to check three inference methods posterior samples.

Figure 8 shows the probability plot of residuals for α -divergence, KL divergence, and HMC. While there exists curvature at the tail of α -divergence residuals, it is close to the HMC method. The KL divergence residuals deviate far more from the standard Weibull distribution. Table 5 summarizes the NLL and computational time of three methods. The NLL value obtained using α -divergence is very close to HMC while taking much less computing time. Table 6 shows the parameter estimates and SE of three inference methods. Compared to the KL divergence, the point estimates of VI with α -divergence are closer to HMC, except for the Physical Region effect

estimation, while the relative impact of different Physical Regions remains consistent between α -divergence and HMC. From both HMC and α -divergence estimates, it is seen that if other covariates are fixed, control and light thinning treatment reduce the survival rates compared to heavy thinning. Pine trees in the Piedmont region tend to have higher survival rates compared to the Coastal Plain region. Trees with larger DBH and TH tend to have longer lifetime.

Since the statistical inference of the α -divergence is very close to HMC and it takes less computing time, its inference results are presented in the following. Figure 9(a) shows the estimated spatial correlation versus the standardized great circle distance. The curve shows that the spatial correlation decays to 0.15 at a distance 0.02. Since the distance matrix is

Table 5. The negative log likelihood and computing time of α -divergence, KL divergence, and HMC inferences with pine tree data.

	$\alpha = 0.8$	KL	HMC
NLL	11149.40	13712.82	11147.80
Time (hours)	3.19	5.25	7.79

standardized into range $[0, 1]$ by dividing the maximum great circle distance between site pairs (1,983.41 km), it means that locations within a range of 39.67 km ($0.02 \times 1,983.41$ km) have correlations greater than or equal to 0.15. Locations that are over 47.6 km have correlations smaller than 0.1, which can be considered as no correlations. Figure 9(b) shows the histogram of estimated random effects, which are in the range of $(-2, 2)$, indicating large spatial variability in survival rate. Figure 10 shows the estimated spatial random effect at each site on the map. The map shows that moving from Virginia to South Carolina, dots tend to change from light blue/green to orange and brown, indicating the spatial random effect increases and the survival rate decreases. There are also several sites in Texas and Louisiana that have large estimated random effects, indicating locally reduced survival rates.

5. Conclusion and discussion

In this study, driven by the computational demands for spatial lifetime models with a large number of

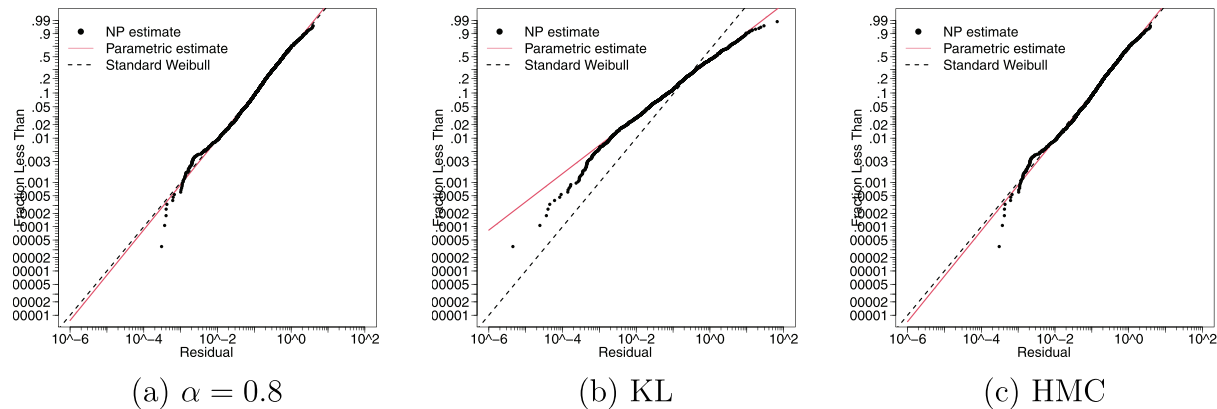



Figure 8. Weibull probability plot of residuals for α -divergence, KL divergence, and HMC with the pine tree data.

Table 6. The parameter estimates and SE of three inference methods for the pine tree data.

Parameter	α -divergence		KL divergence		HMC	
	Estimate	SE	Estimate	SE	Estimate	SE
Treatment (Control)	1.0620	0.0458	0.7204	0.0817	1.0648	0.0767
Treatment (Light)	0.2364	0.0532	0.0728	0.0842	0.2365	0.0800
Physical region (Coastal Plain)	0.0273	0.0501	-0.1729	0.1007	0.5873	0.7354
Physical region (Piedmont)	-0.3027	0.0556	-0.2214	0.1088	0.1517	0.7535
Crown class (dominant)	-2.1722	0.0764	-2.5599	0.1101	-2.1696	0.1163
Crown class (codominant)	-2.1563	0.0591	-2.6741	0.0833	-2.1554	0.0840
Crown class (intermediate)	-1.1116	0.0495	-1.3718	0.0599	-1.1052	0.0660
DBH	-0.2040	0.0086	-0.4066	0.0223	-0.1929	0.0287
TH	-0.0211	0.0012	-0.0789	0.0028	-0.0222	0.0040
a	0.0444	0.0023	2.1949	0.1054	0.0243	0.0240
b	0.9157	0.0137	0.9487	0.0484	0.9178	0.0154
s^2_γ	1.4298	0.2348	4.9903	1.0327	1.3903	0.3621
ν	0.0102	0.0028	0.0425	0.1024	0.0103	0.0029



(a) Spatial Correlation (b) Random Effects

Figure 9. Estimated spatial correlation versus distance and histogram of random effects from VI with α -divergence, $\alpha = 0.8$ for pine tree data.

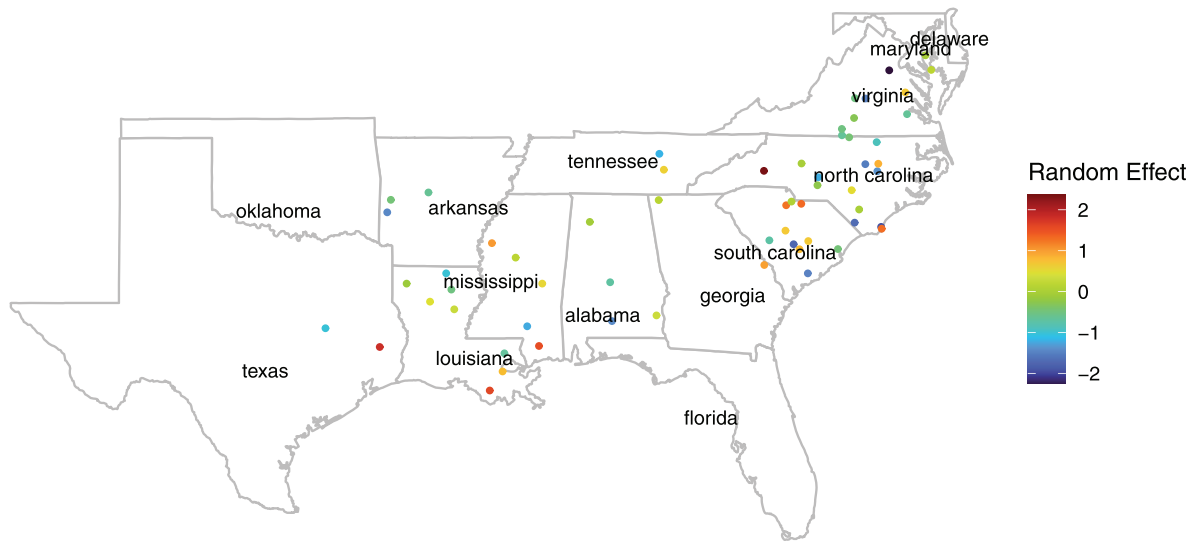


Figure 10. Estimated random effects obtained from VI with α -divergence, $\alpha = 0.8$ at the 60 selected locations for the pine tree data. To avoid the overlapping between points, the same jitters are added as in Figure 3.

locations, we focus on the capability and computational efficiency of VI in analyzing two real spatially correlated lifetime datasets. The classic MCMC inference approach often requires extensive efforts on parameter tuning to achieve convergence. The VI method, in contrast, reformulates the inference challenge to an optimization problem. Once the optimized

variational parameter is obtained, the model inference can be completed immediately.

There are several contributions in this work. First, we consider a unified VI framework that comprises a stochastic optimization algorithm and MC approximation for the variational bound. To enhance generality, the Rényi's α -divergence, which is the root divergence

of the classical KL divergence, is adopted in the framework. Second, the performance of VI is studied in two spatial lifetime models, the PH model and the CEM, both incorporating time-varying covariates, aiming to demonstrate the versatility of VI across different types of spatial lifetime models. Third, two real datasets, Titan GPU lifetime data and pine tree survival data are used to study the VI inference framework and show the performance of VI with α -divergence is comparable to HMC with less computational cost.

Case studies with the Titan GPU lifetime data and pine tree lifetime data show that VI with α -divergence has comparable inference performance as HMC. VI with α -divergence has very close NLL values to HMC while speeding up the computing time by around 2 times in both datasets. The parameter estimation and uncertainty comparisons in two datasets indicate that the KL divergence produces a different posterior from the other two methods. Besides, KL divergence takes a longer time to meet the convergence criteria. A possible explanation is that the parameter space explored by KL divergence is more restricted than α -divergence, potentially leading to slower convergence and local optimal solutions. In contrast, α -divergence encourages a more flexible variational distribution, resulting in better inference performance. In summary, based on these two applications, we find α -divergence with $\alpha < 1$ has comparable performance as HMC but with better computational efficiency.

There are a few aspects to consider when adopting the VI framework for analyzing spatially correlated lifetime data. The first aspect is the choice of the number of samples h in the MC approximation in Eq. [9]. Choosing the number of samples used in MC is crucial for the approximation quality. Since the proposed VI algorithm relies on the gradient-based optimization, both the variational bound (the objective function) and its gradient need to be well approximated. To select a reasonable h , one wants the number of samples h to provide relatively stable estimations of the variational bound and its gradient. While larger h generally improves estimation stability, it also increases computational cost. For the GPU and pine tree datasets in this study, we evaluated the tradeoff between computational time and approximation stability across different h values. Based on the variance of the variational bound and its gradient, we selected $h = 500$ for the GPU dataset and $h = 800$ for the tree dataset. Additional details and results are available in the [Supplementary Section 1](#). In practice,

one needs to manually select h by balancing computational efficiency with approximation quality.

The selection of α values in the α -divergence is another important aspect to be considered. While cross-validation (CV) would be the ideal approach for selecting α , CV is a challenging task with spatially correlated lifetime data. Randomly splitting the dataset into k folds can result in spatial dependencies between the response in the training, testing, and validation sets, potentially leading to overfitting and impacting the choice of α -value. In the case studies, we consider a few α values with $\alpha < 1$ and observe similar inference performances with both datasets. So, in this article, $\alpha = 0.8$ is considered. [Supplementary Section 2](#) presents detailed inference results using VI with $\alpha \in [0.3, 0.9]$ (with increments of 0.1) for the GPU dataset. The consistent performance across this range suggests that choosing any α value smaller than 1 and not too close to 1 (e.g., $\alpha \leq 0.9$) can yield good inference results. However, developing appropriate CV methods for selecting α in spatially structured datasets remains an interesting direction for future research.

The choice of variational distribution also needs to be considered. In this study, to simplify the optimization in the VI algorithm, we adopt the commonly used mean-field approach, that is, assuming the variational distributions are independent of each other. Several works (e.g., Bernardi, Bianchi, and Bianco 2024, Quiroz, Nott, and Kohn 2023) in the literature also considered independent or conditionally independent variational distributions in spatial-temporal models (e.g., vector autoregressive and state-space models) to balance computational efficiency and estimation accuracy. However, the independence assumption can introduce some estimation bias, especially for the spatial random effects. We evaluate the impact of this choice using the two datasets in this study, and the results are available in [Supplementary Section 3](#). The results show that for data with stronger spatial correlations, independent variational distributions can lead to relatively larger bias in estimating random effects. Nevertheless, residual plots and negative log-likelihood comparisons between α -divergence and HMC methods reveal that this assumption does not significantly compromise the α -divergence's overall inference capability. While adopting a variational distribution that incorporates spatial dependency structures for the random effects could improve estimation accuracy, it would also introduce computational challenges and additional complexities in the VI estimation process. This represents a tradeoff between the accuracy of random effects estimation and the

efficiency of optimization. In practice, researchers can choose the variational distribution of the random effects based on the dataset size, model complexity, and the spatial correlation ranges.

There are some future research directions. First, regarding α -value selection in α -divergence, while our current approach uses empirical testing of different values, developing more systematic methods for α selection would be beneficial. Second, although the main focus of this study is to demonstrate the inference capability of variational inference in spatial survival models, it will be interesting to examine prediction performance for HMC and the VI methods on the testing set. For spatially correlated lifetime data, both research directions require a proper CV design or training-testing set split, where methods reviewed in Otto, Fassò, and Maranzano (2024) can be applicable. Works by Jacquez (1999), Cressie and Kornak (2003), and Santi et al. (2021) also provide insights regarding the desired prior information of spatial dependency when splitting the data into different folds. Additionally, this study focuses on the Bayesian analysis framework. It will be interesting to consider frequentist approaches, like the generalized spatial mixed effect models, as alternatives. Comparing these methods' inference capabilities and computational efficiency with the current VI framework would provide valuable insights. Finally, it may be interesting to consider a mixture variational distribution or other flexible structures to capture the possible dependencies among parameters as well as the multi-mode behavior of posterior distributions.

Acknowledgment

The work by Wang was partially supported by the Summit Advancement Disciplines of Zhejiang Province (Zhejiang Gongshang University–Statistics).

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

The work by Wang was partially supported by the Summit Advancement Disciplines of Zhejiang Province (Zhejiang Gongshang University–Statistics) under Grant 1020JYN6524001G-034.

ORCID

Xinwei Deng  <http://orcid.org/0000-0002-1560-2405>

Data availability statement

The GPU lifetime data in this study are openly available at <https://doi.ccs.ornl.gov/dataset/4aa54c30-3d51-5443-b839-88a130e4c713>, and the pine tree lifetime data are openly available at <https://doi.org/10.1080/01621459.2014.995793?scroll=top>.

References

- Banerjee, S., M. M. Wall, and B. P. Carlin. 2003. Frailty modeling for spatially correlated survival data, with application to infant mortality in Minnesota. *Biostatistics* 4 (1):123–42. doi: [10.1093/biostatistics/4.1.123](https://doi.org/10.1093/biostatistics/4.1.123).
- Bernardi, M., D. Bianchi, and N. Bianco. 2024. Variational inference for large Bayesian vector autoregressions. *Journal of Business & Economic Statistics* 42 (3):1066–82. doi: [10.1080/07350015.2023.2290716](https://doi.org/10.1080/07350015.2023.2290716).
- Blei, D. M., A. Kucukelbir, and J. D. McAuliffe. 2017. Variational inference: A review for statisticians. *Journal of the American Statistical Association* 112 (518):859–77. doi: [10.1080/01621459.2017.1285773](https://doi.org/10.1080/01621459.2017.1285773).
- Boškoski, P., M. Perne, M. Rameša, and B. M. Boshkoska. 2021. Variational Bayes survival analysis for unemployment modelling. *Knowledge-Based Systems* 229:107335. doi: [10.1016/j.knsys.2021.107335](https://doi.org/10.1016/j.knsys.2021.107335).
- Carpenter, B., A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. 2017. Stan: A probabilistic programming language. *Journal of Statistical Software* 76:1. doi: [10.18637/jss.v076.i01](https://doi.org/10.18637/jss.v076.i01).
- Cox, D. R. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society Series B* 34 (2):187–202. doi: [10.1111/j.2517-6161.1972.tb00899.x](https://doi.org/10.1111/j.2517-6161.1972.tb00899.x).
- Cressie, N., and J. Kornak. 2003. Spatial statistics in the presence of location error with an application to remote sensing of the environment. *Statistical Science* 18 (4):436–56. doi: [10.1214/ss/1081443228](https://doi.org/10.1214/ss/1081443228).
- Diva, U., D. K. Dey, and S. Banerjee. 2008. Parametric models for spatially correlated survival data for individuals with multiple cancers. *Statistics in Medicine* 27 (12): 2127–44. doi: [10.1002/sim.3141](https://doi.org/10.1002/sim.3141).
- Henderson, R., S. Shimakura, and D. Gorst. 2002. Modeling spatial variation in Leukemia survival data. *Journal of the American Statistical Association* 97 (460):965–72. doi: [10.1198/016214502388618753](https://doi.org/10.1198/016214502388618753).
- Hernandez-Lobato, J., Y. Li, M. Rowland, T. Bui, D. Hernández-Lobato, and R. Turner. 2016. Black-box alpha divergence minimization. In *International Conference on Machine Learning*, 1511–20.
- Hoffman, M. D., and D. M. Blei. 2015. Stochastic structured variational inference. Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics. *Proceedings of Machine Learning Research* 38: 361–369.
- Hoffman, M. D., and A. Gelman. 2014. The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15 (1):1593–623.
- Hong, Y., and W. Q. Meeker. 2010. Field-failure and warranty prediction based on auxiliary use-rate information.

- Technometrics* 52 (2):148–59. doi: [10.1198/TECH.2010.09097](https://doi.org/10.1198/TECH.2010.09097).
- Hong, Y., and W. Q. Meeker. 2013. Field-failure predictions based on failure-time data with dynamic covariate information. *Technometrics* 55 (2):135–49. doi: [10.1080/00401706.2013.765324](https://doi.org/10.1080/00401706.2013.765324).
- Jacquez, G. M. 1999. Spatial statistics when locations are uncertain. *Annals of Gis* 5 (2):77–87. doi: [10.1080/10824009909480517](https://doi.org/10.1080/10824009909480517).
- Jung, A. W., and M. Gerstung. 2021. Bayesian Cox regression for population-scale inference in electronic health records. arXiv Preprint *arXiv:2106.10057*.
- Kim, M., and V. Pavlovic. 2018. Variational inference for Gaussian process models for survival analysis. In 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018, 435–445.
- Kingma, D. P., and J. Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klein, J. P., and M. L. Moeschberger. 2003. *Survival analysis: Techniques for censored and truncated data*. Berlin: Springer.
- Kucukelbir, A., D. Tran, R. Ranganath, A. Gelman, and D. M. Blei. 2017. Automatic differentiation variational inference. *Journal of Machine Learning Research* 18 (14): 1–45.
- Kundu, P., A. K. Darpe, and M. S. Kulkarni. 2019. Weibull accelerated failure time regression model for remaining useful life prediction of bearing working under multiple operating conditions. *Mechanical Systems and Signal Processing* 134:106302. doi: [10.1016/j.ymssp.2019.106302](https://doi.org/10.1016/j.ymssp.2019.106302).
- Li, J., Y. Hong, R. Thapa, and H. E. Burkhart. 2015. Survival analysis of Loblolly pine trees with spatially correlated random effects. *Journal of the American Statistical Association* 110 (510):486–502. doi: [10.1080/01621459.2014.995793](https://doi.org/10.1080/01621459.2014.995793).
- Li, M., and W. Q. Meeker. 2014. Application of Bayesian methods in reliability data analyses. *Journal of Quality Technology* 46 (1):1–23. doi: [10.1080/00224065.2014.11917951](https://doi.org/10.1080/00224065.2014.11917951).
- Li, Y., and R. E. Turner. 2016. Rényi divergence variational inference. *Advances in Neural Information Processing Systems* 29:1–9.
- Meeker, W. Q., L. A. Escobar, and F. G. Pascual. 2022. *Statistical methods for reliability data*, 2nd ed. New York: John Wiley & Sons.
- Min, J., Y. Hong, W. Meeker, and O. George. 2025. A spatially correlated competing risks time-to-event model for supercomputer GPU failure data. *Technometrics* 67 (3): 531–45. doi: [10.1080/00401706.2025.2475783](https://doi.org/10.1080/00401706.2025.2475783).
- Neal, R. M. 2011. *Handbook of Markov Chain Monte Carlo, chapter 5, MCMC using Hamiltonian dynamics*. New York: CRC Press.
- Newby, M. 1988. Accelerated failure time models for reliability data analysis. *Reliability Engineering & System Safety* 20 (3):187–97. doi: [10.1016/0951-8320\(88\)90114-7](https://doi.org/10.1016/0951-8320(88)90114-7).
- Ostrouchov, G., D. Maxwell, R. A. Ashraf, C. Engelmann, M. Shankar, and J. H. Rogers. (2020). GPU lifetimes on Titan supercomputer: Survival analysis and reliability. In SC20: International Conference for High Performance Computing, Networking, Storage and Analysis, 1–14. doi: [10.1109/SC41405.2020.00045](https://doi.org/10.1109/SC41405.2020.00045).
- Otto, P., A. Fassò, and P. Maranzano. 2024. A review of regularised estimation methods and cross-validation in spatiotemporal statistics. *Statistics Surveys* 18:299–340. doi: [10.1214/24-ss150](https://doi.org/10.1214/24-ss150).
- Pan, C., B. Cai, L. Wang, and X. Lin. 2014. Bayesian semi-parametric model for spatially correlated interval-censored survival data. *Computational Statistics & Data Analysis* 74:198–208. doi: [10.1016/j.csda.2013.11.016](https://doi.org/10.1016/j.csda.2013.11.016).
- Quiroz, M., D. J. Nott, and R. Kohn. 2023. Gaussian variational approximations for high-dimensional state space models. *Bayesian Analysis* 18 (3):989–1016. doi: [10.1214/22-BA1332](https://doi.org/10.1214/22-BA1332).
- Ranganath, R., S. Gerrish, and D. Blei. 2014. Black box variational inference. *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, PMLR* 33:814–822.
- Ranganath, R., D. Tran, and D. Blei. 2016. Hierarchical variational models. In *International Conference on Machine Learning*, 324–33.
- Santi, F., M. M. Dickson, G. Espa, E. Taufer, and A. Mazzitelli. 2021. Handling spatial dependence under unknown unit locations. *Spatial Economic Analysis* 16 (2):194–216. doi: [10.1080/17421772.2020.1769171](https://doi.org/10.1080/17421772.2020.1769171).
- Wan, N., D. Li, and N. Hovakimyan. 2020. f -divergence variational inference. *Advances in Neural Information Processing Systems* 33:17370–9.
- Xiu, Z., C. Tao, and R. Henao. 2020. Variational learning of individual survival distributions. In Proceedings of the ACM Conference on Health, Inference, and Learning, 10–8. ACM. doi: [10.1145/3368555.3384454](https://doi.org/10.1145/3368555.3384454).
- Zhao, L., T. E. Hanson, and B. P. Carlin. 2009. Mixtures of Polya trees for flexible spatial frailty survival modelling. *Biometrika* 96 (2):263–76. doi: [10.1093/biomet/asp014](https://doi.org/10.1093/biomet/asp014).