

Hierarchical Transformed Scale Mixtures for Flexible Modeling of Spatial Extremes on Datasets with Many Locations

Likun Zhang

Department of Statistics, Pennsylvania State University

Benjamin A. Shaby

Department of Statistics, Colorado State University

Jennifer L. Wadsworth

Department of Mathematics and Statistics, Fylde College, Lancaster University

December 1, 2020

Abstract

Flexible spatial models that allow transitions between tail dependence classes have recently appeared in the literature. However, inference for these models is computationally prohibitive, even in moderate dimensions, due to the necessity of repeatedly evaluating the multivariate Gaussian distribution function. In this work, we attempt to achieve truly high-dimensional inference for extremes of spatial processes, while retaining the desirable flexibility in the tail dependence structure, by modifying an established class of models based on scale mixtures Gaussian processes. We show that the desired extremal dependence properties from the original models are preserved under the modification, and demonstrate that the corresponding Bayesian hierarchical model does not involve the expensive computation of the multivariate Gaussian distribution function. We fit our model to exceedances of a high threshold, and perform coverage analyses and cross-model checks to validate its ability to capture different types of tail characteristics. We use a standard adaptive Metropolis algorithm for model fitting, and further accelerate the computation via parallelization and Rcpp.

Lastly, we apply the model to a dataset of a fire threat index on the Great Plains region of the US, which is vulnerable to massively destructive wildfires. We find that the joint tail of the fire threat index exhibits a decaying dependence structure that cannot be captured by limiting extreme value models.

Keywords: Asymptotic dependence, Asymptotic independence, Censored likelihood, Threshold exceedance

1 Introduction

Modeling the dependence structure in the extremes of spatial processes is of great consequence for risk analysis of extreme events. In this paper, we make a slight alteration to the flexible class of randomly scaled transformed Gaussian process models to sidestep computational bottlenecks normally encountered in the likelihood. Our modification enables high-dimensional inference, while preserving submodels that transition smoothly between extremal dependence classes.

Generally, the probability that two spatially-indexed random variables exceed a high level simultaneously varies by their separation distance, and the particular way in which this joint probability decays must be well-represented in models if one hopes to accurately to assess risks posed by spatial extremal phenomena. Good estimation of how the dependence changes both with distance in space and as one moves farther into the joint tail will enable us to accurately calculate exceedance probabilities of areal quantities, predict at un-observed locations, and, secondarily, get a more realistic picture of marginal quantities.

In classical spatial modeling, Gaussian processes have been widely used due to their mathematical simplicity and tractability for larger datasets. However, the Gaussian density function is very light-tailed, and thus has the potential to underestimate probabilities associated with extreme events; furthermore, Gaussian models stipulate that the dependence among rare events at distinct locations will always diminish such that the probability of observing an extreme at one location, conditional upon an extreme at another location, is zero in the limit. This property is termed asymptotic independence, but models that only exhibit this phenomenon may be too inflexible for applications where the true tail dependence structure is uncertain.

Max-stable processes form an important class of models that exhibit the alternative

scenario of asymptotic dependence. They are the natural extension of classical univariate extreme value theory to infinite dimensional settings, and therefore can provide an asymptotically-justified modeling framework for datasets consisting of block-maxima. Counterparts of max-stable processes suitable for threshold exceedances are called generalized Pareto processes (Ferreira and de Haan, 2014; Thibaud and Opitz, 2015). These processes are also asymptotically dependent, but possess the advantage that they bypass many of the computational difficulties of max-stable processes.

Despite the theoretical appeal of limiting max-stable and generalized Pareto processes, there are two main drawbacks to these models: (i) the assumption of asymptotic dependence may be incorrect and (ii) even if the data are asymptotically dependent, they will often not appear to follow such limiting models at sub-asymptotic levels. Both max-stable and generalized Pareto process dependence structures exhibit stability properties, meaning that their dependence structures are invariant to the operations of taking maxima and conditioning upon exceedances of higher thresholds, respectively. If the true data generating process exhibits weakening dependence in the un-observed region of the tail, inference drawn under these models about the far joint tail will over-estimate risk, sometimes substantially.

On account of the limitations of limiting models, it is desirable to find a family of spatial models that can transition between asymptotic dependence and asymptotic independence. In particular, we will be examining a class of marginally transformed Gaussian scale mixture models, which includes those recently proposed by Huser et al. (2017) and Huser and Wadsworth (2019). These models are of great interest due to their appealing theoretical properties and ability to flexibly capture both types of extremal dependence structure. Unfortunately, inference for these models is not feasible for large numbers of observed sites, as calculation of the censored likelihood, the preferred method for fitting joint tail mod-

els, entails integration over high-dimensional multivariate Gaussian distribution functions. To increase scalability, we propose an adaptation of the model by adding an independent measurement error term to each component. By adding this nugget effect, the new model circumvents the lengthy computation of the multivariate normal distribution function. Also it can elegantly avoid the integral of the process below the censoring threshold by considering the uncensored process as latent and drawing from it using Gibbs sampling, allowing for truly high-dimensional inference. Furthermore, we show that the modified models retain all the significant asymptotic properties of the original smooth models despite the presence of the measurement errors, which lays a solid theoretical foundation for correctly capturing the sub-asymptotic dependence behavior.

The article is organized as follows. Section 2 provides a brief literature review on the measures of extremal dependence and hybrid spatial extreme models, and further explains the intractability of the existing censored likelihood approaches for inference on these hybrid models. Section 3 describes our new model that alleviates the computational problems, and studies its extremal dependence properties. Section 4 includes a marginal transformation in the hierarchical model and details the inference using Gibbs sampling. Section 5 presents a simulation study that validates the methodology. We apply our model to a dataset of the Fosberg Fire Weather Index (FFWI) on the Great Plains in Section 6. Section 7 concludes with some discussion. Appendix A provides proofs of all the theoretical results. Appendix B includes supplementary diagnostics for the data application.

2 Spatial Dependence for Extremes

For a stochastic process $\{X(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$, we write $X_j = X(\mathbf{s}_j)$ and so forth for simplicity, where \mathbf{s}_j denotes the j th spatial location. Let F_j denote the marginal distribution of X_j ,

and F_{jk} denote the joint distribution of X_j and X_k . It is useful to summarize the extremal dependence implied by the observed process concisely.

We restrict the scope to the bivariate case, focusing on stationary and isotropic random fields. One example of a bivariate dependence measure is the upper tail dependence coefficient:

$$\chi_u(h) = P(F_j(X_j) > u \mid F_k(X_k) > u) \quad (1)$$

where $h = \|\mathbf{s}_j - \mathbf{s}_k\|$. Joe (1993) defined the upper tail dependence parameter as the limit $\chi(h) = \lim_{u \rightarrow 1} \chi_u(h)$. Asymptotic dependence is attained if and only if $\chi(h) > 0$, while $\chi(h) = 0$ defines asymptotic independence.

For max-stable processes, $\chi_u(h) = 2 - V(1, 1) + O(1 - u)$, $u \rightarrow 1$, where $V(\cdot, \cdot) = -\log F_{jk}(\cdot, \cdot)$ when X has unit Fréchet margins. Max-stable distributions can be associated to a generalized Pareto counterpart, for which $\chi_u(h) \equiv \chi(h) = 2 - V(1, 1)$ for all u above a certain level (Rootzén et al., 2018). The fact that $\chi_u(h)$ does not depend on u is a manifestation of the threshold-stability of generalized Pareto processes. In practice, empirical estimates of (1) from data tend to show $\chi_u(h)$ decreasing with both h and u , meaning that realistic models should also have this property.

When $\chi(h) = 0$, i.e., the case of asymptotic independence, further detail about the behavior of $\chi_u(h)$ is obtained by exploiting the joint tail assumption of Ledford and Tawn (1996):

$$P(F_j(X_j) > u \mid F_k(X_k) > u) = \mathcal{L}(1 - u)(1 - u)^{1/\eta_X(h)-1} \quad (2)$$

where \mathcal{L} is slowly varying at zero, that is, $\lim_{t \rightarrow 0} \mathcal{L}(tx)/\mathcal{L}(t) = 1$ for any $x > 0$, and $\eta_X(h) \in (0, 1]$ is the coefficient of tail dependence of the process X . The pair of variables (X_j, X_k) are asymptotically dependent when $\eta_X(h) = 1$ and $\mathcal{L}(\cdot) \rightarrow 0$. The remaining cases are all asymptotically independent, and the value of $\eta_X(h)$ characterizes the strength of extremal dependence in the upper joint tail. In the case of a Gaussian process, $\eta_X(h) = \{1 + \rho(h)\}/2$,

where $\rho(h)$ is the correlation at lag h . The variables are called positively associated when $\eta_X(h) > 1/2$ and negatively associated when $\eta_X(h) < 1/2$. Near independence corresponds to $\eta_X(h) = 1/2$.

Gaussian processes are asymptotically independent for all correlations $\rho(h) \neq 1$. They might be considered candidates for modeling the joint tail of asymptotically independent phenomena, but as there is no theory to specifically recommend Gaussian processes in this scenario, it is desirable to consider other models as well. As an alternative, [Opitz \(2016\)](#) captures spatial dependence in asymptotically independent processes by construction of Laplace random fields, defined as mixtures of Gaussian processes with a random variance that is exponentially distributed. [Wadsworth and Tawn \(2012\)](#) proposed the class of inverted max-stable processes, for which the tail decay is specified fully by $\eta_X(h)$, although inference is computationally challenging.

In real datasets, it is difficult to conclude definitively whether data exhibit asymptotic independence or asymptotic dependence, and incorrectly assuming an asymptotically independent model can lead to equally severe problems with bias as incorrectly assuming an asymptotically dependent model. Because of this, a recent focus in the literature has been on models that can encompass both scenarios.

2.1 Traversing Asymptotic Independence and Dependence in Spatial Extremes

[Wadsworth and Tawn \(2012\)](#) were the first to introduce hybrid models that combine max-stable and inverted max-stable processes so that asymptotic dependence prevails at short distances, and asymptotic independence at long distances. However, inference for this model is difficult because there are a fairly large number of parameters involved, and the

transition between the dependence classes takes place at the boundary of the parameter space.

Recently, several Gaussian scale mixture models were proposed to allow more flexible transitions between dependence classes. Through multiplying an asymptotically independent Gaussian process by a random effect that governs the extremal dependence, these models can be described by a small number of parameters and have non-trivial asymptotically independent and asymptotically dependent submodels. More precisely, suppose $\{Z(\mathbf{s}), \mathbf{s} \in \mathcal{S}\}$ is a standard isotropic and stationary Gaussian process with covariance function $C_{\boldsymbol{\theta}_C}(h)$ indexed by a parameter vector $\boldsymbol{\theta}_C$, where h is the length of the separation vector, so that $\Sigma_{\boldsymbol{\theta}_C}$ is the covariance matrix of associated finite-dimensional distributions. The class of Gaussian scale mixture models can be constructed as

$$X^*(\mathbf{s}) = R \cdot g(Z(\mathbf{s})), \quad R | \boldsymbol{\theta}_R \sim F_R, \quad (3)$$

where $g(\cdot)$ is a link function, and $R > 0$ is a random scaling factor, from distribution F_R indexed by $\boldsymbol{\theta}_R$, that can be interpreted as a constant random process over spatial domain \mathcal{S} with perfect dependence. Loosely speaking, heavier tailed R , impacting simultaneously the whole domain \mathcal{S} , induces asymptotic dependence in X^* , whereas lighter tailed R induces asymptotic independence. [Engelke et al. \(2019\)](#) provide a fuller description of how extremal dependence of X^* relates to the relative marginal tail heaviness of R and $g(Z)$.

[Morris et al. \(2017a\)](#) uses a space-time model based on skew- t process, where $g(\cdot)$ is a identity function, $R^2 \sim \text{IG}(a/2, b/2)$ is an inverse gamma random variable, and $C_{\boldsymbol{\theta}_C}$ is a Matérn covariance function. On top of the mixture, they added covariate effects and a skew term. Since the inverse gamma distribution is heavy tailed, the skew- t process is asymptotically dependent for $a < \infty$. Asymptotic independence is achieved only when $a \rightarrow \infty$.

Huser et al. (2017) also used an identity link function, but placed few assumptions on the random scale, and provided more general results on the joint tail decay rates of the mixture processes. They showed that a wide class of Weibull-like tail decay in R yields asymptotic independence, while a Pareto-like tail that is regularly varying at infinity gives asymptotic dependence. They also proposed a parametric model that bridges the two asymptotic regimes and provides a simple transition, in which R is a two-parameter distribution

$$F_R(r) = \begin{cases} 1 - \exp\{-\gamma(r^\beta - 1)/\beta\}, & \beta > 0, \\ 1 - r^{-\gamma}, & \beta = 0 \end{cases} \quad (4)$$

where $\gamma > 0$, and the support is $[1, \infty)$. Since $(r^\beta - 1)/\beta$ converges to $\log r$ as β approaches 0, (4) forms a continuous parametric family on β . When $\beta > 0$, (4) constitutes a class of Weibull-type distributions and thus assures asymptotic independence. When $\beta = 0$, the variable R is Pareto distributed and thus gives asymptotic dependence. This shows that the model provides greater flexibility and can transition from asymptotic dependence to independence via adjusting the value of β .

However, the previous two Gaussian scale mixture models both make the transition between the dependence classes at the limit or the boundary of the parameter space. They are also inflexible in their representation of asymptotic dependence structures because there is dominating preference over one dependence class. It may be more desirable to find a model for which the transition takes place in the interior of the parameter space so one we can quantify the uncertainty about the dependence class in a simpler manner. To overcome this, Huser and Wadsworth (2019) proposed a marginally transformed Gaussian scale mixture model, where $g(\cdot)$ transforms a standard Gaussian variable to standard Pareto, and R itself is Pareto distributed:

$$g(z) = \frac{1}{1 - \Phi(z)}, \quad R | \delta \sim \text{Pareto} \left(\frac{1 - \delta}{\delta} \right), \quad \delta \in [0, 1]. \quad (5)$$

Here the type of asymptotic dependence is determined by the value of δ . When $\delta \leq 1/2$, R is lighter tailed or equivalent to standard Pareto, which induces asymptotic independence; when $\delta > 1/2$, the converse is true, which induces asymptotic dependence. Specifically, the upper tail dependence parameter $\chi_{X^*} = \frac{2\delta-1}{\delta} E [\min\{g(Z_i), g(Z_k)\}^{(1-\delta)/\delta}]$ when $\delta > 1/2$ and 0 otherwise, while the coefficient of tail dependence is

$$\eta_{X^*} = \begin{cases} 1, & \delta > \frac{1}{2}, \\ \frac{\delta}{1-\delta}, & \frac{\eta_Z}{\eta_Z+1} < \delta \leq \frac{1}{2}, \\ \eta_Z, & \delta \leq \frac{\eta_Z}{\eta_Z+1}, \end{cases}$$

where η_Z is the coefficient of tail dependence for (Z_i, Z_k) (Huser and Wadsworth, 2019).

The model in (5) provides a smooth transition through asymptotically independent and asymptotically dependent submodels. It has many appealing asymptotic properties. However, inference for models of the form (3) is typically made via censored likelihood. This requires computing an integral where the integrand contains the Gaussian distribution function in $|\mathcal{C}|$ dimensions, where $|\mathcal{C}|$ is the number of components below a designated high threshold. Such integrals are computationally prohibitive for even moderately-sized datasets. In Section 3 we introduce a slight alteration to this model to make it tractable while preserving all the desired asymptotic results.

2.2 The Censored Likelihood

In multivariate and spatial extremes, the preferred approach to fitting the dependence structure is using a censored likelihood, which prevents observations from the bulk of the distribution from affecting the estimation of the extremal dependence structure. It provides a reasonable compromise between bias and variance compared to alternative approaches, although different censoring schemes have been adopted (Thibaud and Opitz, 2015; Huser

et al., 2016).

For a process of the form (3) observed at D spatial locations $\mathbf{s}_1, \dots, \mathbf{s}_D \in \mathcal{S}$, we obtain the distribution function by conditioning on R as

$$G(\mathbf{x}^*) = \int_1^{r^*} \Phi_D \left(g^{-1} \left(\frac{\mathbf{x}^*}{r} \right); \Sigma_{\theta_C} \right) f_R(r) dr, \quad (6)$$

where $r^* = \min(x_1^*, \dots, x_D^*)$, and Φ_D denotes the D -variate Gaussian distribution with zero mean and covariance matrix Σ_{θ_C} .

Let $\mathcal{C} \subseteq \{1, \dots, D\}$ be the set of locations with censored observations—that is, the set of locations where the components are below a high threshold; let \mathcal{U} be the set of locations with uncensored observations. For any index set $A, B \subset \{1, \dots, D\}$, denote $\mathbf{x}_A = \{\mathbf{x}_i : i \in A\}$, $\Sigma_{A;B}$ as the matrix Σ restricted to the rows in A and the columns in B , and let $\Sigma_{A|B}$ be the Schur complement of B in $\Sigma_{A;B}$. The likelihood is obtained via taking partial derivatives of (6) with respect to \mathcal{U} :

$$\begin{aligned} \frac{\partial^{|\mathcal{U}|}}{\partial \mathbf{x}_{\mathcal{U}}^*} G(\mathbf{x}^*) &= \int_1^{r^*} \Phi_{|\mathcal{C}|} \left(g^{-1} \left(\frac{\mathbf{x}_{\mathcal{C}}^*}{r} \right) - \Sigma_{\mathcal{C};\mathcal{U}} \Sigma_{\mathcal{U};\mathcal{U}}^{-1} g^{-1} \left(\frac{\mathbf{x}_{\mathcal{U}}^*}{r} \right); \Sigma_{\mathcal{C}|\mathcal{U}} \right) \\ &\quad \times \phi_{|\mathcal{U}|} \left(g^{-1} \left(\frac{\mathbf{x}_{\mathcal{U}}^*}{r} \right), \Sigma_{\mathcal{U}} \right) \prod_{j \in \mathcal{U}} g^{-1'} \left(\frac{x_j^*}{r} \right) r^{-|\mathcal{U}|} f_R(r) dr. \end{aligned} \quad (7)$$

Although a only one-dimensional integral appears in (7), the integrand includes a $|\mathcal{C}|$ -dimensional Gaussian distribution function. When approximating the integral using standard quadrature or Monte Carlo methods, one needs to compute $\Phi_{|\mathcal{C}|}$ for each sample point taken on $(1, r^*)$. This is only feasible when the number of locations D is moderate. Additionally, this calculation will have to be repeated for each time replicate.

To avoid the integrating the process below the threshold, one could instead think of $X^*(\mathbf{s})$ as latent and draw from it using Monte Carlo methods. Consequently there is no need to compute the awkward likelihood (7). However, to update the Markov chain each

time, it is now necessary to draw \mathbf{x}_C^* from a high-dimensional truncated distribution, which might again be computationally intensive.

Therefore, we propose to make a slight adjustment to the model in (3). Our new model is markedly more amenable to higher-dimensional inference, yet it keeps hold of the joint tail decay rates attained in the original model (e.g. [Huser and Wadsworth, 2019](#); [Huser et al., 2017](#)). Equivalently, our new model has non-trivial asymptotically dependent and asymptotically independent submodels with the transition taking place in the interior of the parameter space in the case of our modified version of (5).

3 Model

3.1 Construction

We alter the models in Section 2.1 by adding an independent measurement error term to each component,

$$X(\mathbf{s}_i) := X^*(\mathbf{s}_i) + \epsilon_i = R \cdot g(Z(\mathbf{s}_i)) + \epsilon_i, \quad (8)$$

where $\epsilon_i \stackrel{iid}{\sim} N(0, \tau^2)$, $i = 1, \dots, D$, and distribution of R and the link function remain the same. That is, we add a simple nugget effect to the smooth process $X^*(\mathbf{s}_i)$. When drawing the latent processes below the threshold, we can condition on the smooth process and simply update the noisy one. Because these error terms are independent of each other, there is only a univariate integral involved in the full conditional likelihood. Also, when we update the smooth process $X^*(\mathbf{s})$ given the noisy process $X(\mathbf{s})$, no truncation or censoring is present and it is much easier to sample from the corresponding likelihood. Section 4 contains more details on the Markov Chain Monte Carlo (MCMC) updating scheme, where we show how this small alteration can hugely facilitate inference.

3.2 Dependence Properties

We begin with the model (5) from [Huser and Wadsworth \(2019\)](#), modified as in (8). Recall that $g(Z(\mathbf{s}))$ is a stationary process with standard Pareto margins possessing asymptotic independence; i.e., $P(g(Z(\mathbf{s})) > x) = x^{-1}$ and

$$P(g(Z(\mathbf{s}_i)) > x, g(Z(\mathbf{s}_k)) > x) = \mathcal{L}_Z(x) \cdot x^{-\frac{1}{\eta_Z(h)}}, \quad i \neq k, \quad (9)$$

where $\mathcal{L}_Z(x)$ is slowly varying at infinity, and $\eta_Z(h) = (1 + \rho(h))/2 < 1$ for the Gaussian correlation $\rho(h) < 1$.

Figure 1 illustrates the estimated coefficient of tail dependence η_X as a function of δ for $\eta_Z = 0.1, \dots, 0.9$. For each combination of δ and η_Z , we generate 5,000,000 replicates from model (5) (i.e. $\tau^2 = 0$) and model (8) with $\tau^2 = 1$ respectively. For each replicate, we sample (Z_i, Z_k) from a Gaussian copula with correlation $2\eta_Z - 1$. We then numerically approximate the joint survival probability in (2) for threshold $u = 0.99$ to obtain an estimate of η_X . The left panel of Figure 1 clearly shows that the smooth transition from asymptotic independence to asymptotic dependence takes place around $\delta = 1/2$, confirming the results from [Huser and Wadsworth \(2019\)](#) with a reasonable bias; the right panel shows that adding a measurement error has little effect on the tail dependence because η_X exhibits similar behavior. This result invites investigation of whether the flexible asymptotic properties in [Huser and Wadsworth \(2019\)](#) are preserved in the altered model. In the following, we generalize the problem from the specific model of [Huser and Wadsworth \(2019\)](#) for the process X^* in (8), to any X^* with a wide class of marginal tail behaviors.

The impact of the additive Gaussian nugget effect on the extremal dependence of X depends upon the marginal tail heaviness of X^* : roughly, the heavier the tail of X^* , the less the impact of the noise. Since we employ X as the spatial dependence model, we assume its margins, and those of X^* , are identical over space.

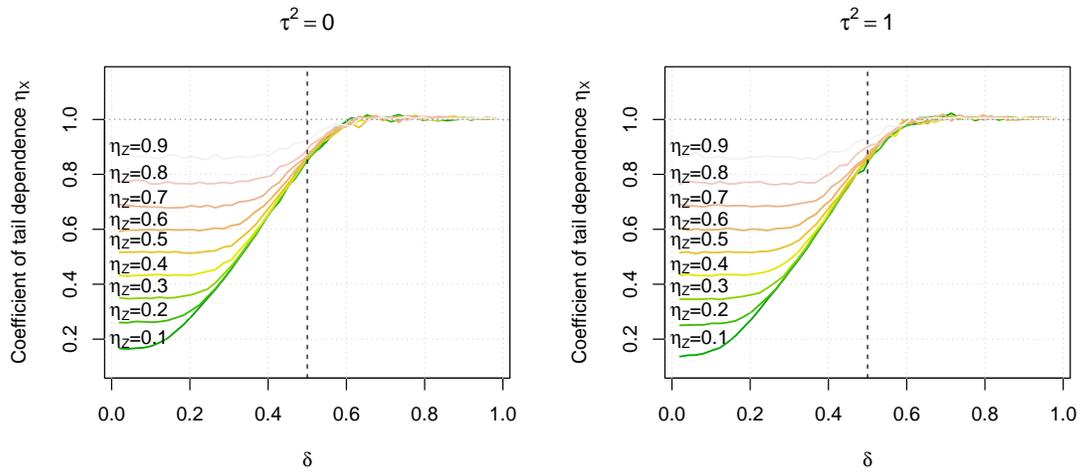


Figure 1: Coefficient of tail dependence approximated for the smooth Gaussian scale mixture processes ($\tau^2 = 0$) and the noisy processes ($\tau^2 = 1$) as a function of $\delta \in (0, 1)$ for $\eta_W = 0.1, \dots, 0.9$. The levels of dependence are similar for two models.

We will focus on the dependence of X under two broad classes of marginal distribution for X^* : regularly varying tails, and Weibull-like tails. A measurable function $u : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is said to be regularly varying at infinity with index κ if for all $x > 0$ $\lim_{t \rightarrow \infty} u(tx)/u(t) = x^\kappa$; we write $u \in \text{RV}_\kappa$. When $\kappa = 0$, the function is slowly varying.

Regularly varying tails are defined through the survival function being regularly varying at infinity, i.e., $P(X^* > x) \in \text{RV}_{-\alpha}$, $\alpha > 0$. Weibull-like tails are defined through the survival function

$$P(X^* > x) \sim u(x) \exp(-\theta x^\alpha), \quad \theta, \alpha > 0, \quad x \rightarrow \infty \quad (10)$$

where $u \in \text{RV}_\kappa$, and α is termed the Weibull index. We also assume that X^* has a density satisfying $f_{X^*}(x) \sim v(x) \exp(-\theta x^\alpha)$, with $v(x) = u(x)(\theta \alpha x^{\alpha-1})$. The main results are now summarized in Proposition 3.1.

Proposition 3.1. *With definitions and notation as above:*

1. *If X^* has a regularly varying tail, or Weibull-like with Weibull index $\alpha < 1$, then $\chi_X = \chi_{X^*}$ and $\eta_X = \eta_{X^*}$.*
2. *If X^* has a Weibull-like tail with Weibull index $\alpha = 1$ then*

$$\chi_X \in [E(e^{\theta \min(\epsilon_1, \epsilon_2)}) / E(e^{\theta \epsilon}), E(e^{\theta \max(\epsilon_1, \epsilon_2)}) / E(e^{\theta \epsilon})] \chi_{X^*}$$

and $\eta_X = \eta_{X^}$. Note if $\chi_{X^*} = 0$ then so is χ_X .*

3. *If X^* has a Weibull-like tail with Weibull index $\alpha > 1$ then:*

(a) If $\alpha \in (1, 2)$, then $\eta_X = \eta_{X^}$*

(b) If $\alpha = 2$, then an interval can be given for η_X (see Expression (27)).

(c) If $\alpha > 2$, then $\eta_X = 1/2$.

The proof of Proposition 3.1 is given in Appendix A.

For the process of Huser and Wadsworth (2019), described in equation (5), X^* always has a regularly varying tail. In this case, Proposition 3.1 part 1 gives $\eta_X(h) = \eta_{X^*}(h)$, and $\chi_X(h) = \chi_{X^*}(h)$, with $\eta_{X^*}(h)$ and $\chi_{X^*}(h)$ given in Section 2.1. This means the flexible asymptotic properties in Huser and Wadsworth (2019) are preserved in the altered model.

Another popular model for spatial data is the t -process (Røislien and Omre, 2006), which is a Gaussian scale mixture for which the mixing variable R^2 follows an inverse gamma distribution. If X^* is a t -process, then it is asymptotically dependent with a regularly varying tail, so $\eta_X(h) = \eta_{X^*}(h) = 1$ and $\chi_X(h) = \chi_{X^*}(h) > 0$ by Proposition 3.1. Similarly, the skew- t process (Padoan, 2011; Morris et al., 2017b) is regularly varying and asymptotically dependent, so the same conclusions apply.

For the Gaussian scale mixture of Huser et al. (2017), X^* either has regularly varying or Weibull-like tails depending on the distribution of the scaling variable R . In particular when $\beta > 0$ in distribution (4), X^* has a Weibull-like tail with Weibull index $\alpha = 2\beta/(\beta+2) < 2$. As such, parts 1, 2 and 3a of Proposition 3.1 are relevant. In all cases $\eta_{X^*}(h) = \eta_X(h) = \{(1 + \rho(h))/2\}^{\beta/(\beta+2)}$, and $\chi_X(h) = \chi_{X^*}(h) = 0$ for $\beta > 0$. When $\beta = 0$ then X^* has a regularly varying tail, and is asymptotically dependent with $\eta_X(h) = \eta_{X^*}(h) = 1$, and

$$\chi_{X^*}(h) = \chi_X(h) = 2\{1 - T_{\gamma+1}((\gamma+1)^{1/2}(1 - \rho(h))/(1 - \rho(h)^2)^{1/2})\} > 0,$$

where T_ν is the cdf of the Student- t distribution with ν degrees of freedom.

We note that the process X^* in equation (3) is constructed only for its dependence properties, and there is no “natural” scale on which to express it. For example, considering the process of Huser and Wadsworth (2019) we could also write

$$X^*(s) = E + V(s), \tag{11}$$

with $E | \delta \sim \text{Exp}(\delta/(1-\delta))$, $V(s) = \log g(Z(s))$, which has the same dependence structure as defined in (3) and (5), since it is obtained through a monotonic marginal transformation. Taking X^* from (11), Proposition 3.1 part 2 gives $\eta_X(h) = \eta_{X^*}(h)$. Asymptotic dependence of X^* implies asymptotic dependence of X , but only bounds on $\chi_X(h)$ are available.

In practice, if choosing a marginal scale for X^* , there may be a trade-off between theoretical desires and computational practicality. Supposing that we wish X to inherit the properties of X^* , a heavy-tailed choice is best.

4 Bayesian Inference

4.1 Hierarchical Model

We define a Bayesian hierarchical model based on the process (8) defined in Section 3.1 and use a MCMC algorithm to fit to the data. For the reasons outlined in Section 2.2, we assume data are censored below a high threshold q . In (8), the mixing parameter $\boldsymbol{\theta}_R$ controls both joint and marginal behavior of the response X , which we would prefer to separate. Therefore, motivated by the theory of univariate extremes, we first assume our observations above the same high threshold q are generalized Pareto distributed, and we include a marginal transformation in the hierarchical model.

Let $\{Y(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$ denote the observed process. We define a marginal transformation $T(y)$ as follows:

$$T(Y(\mathbf{s})) = F_{X|\boldsymbol{\theta}_R, \tau^2}^{-1} \circ F_{Y|q, \sigma, \xi}(Y(\mathbf{s})), \quad (12)$$

where $F_{X|\boldsymbol{\theta}_R, \tau^2}$ is the marginal distribution function for process (8), and

$$F_{Y|q, \sigma, \xi}(y) = \begin{cases} p, & y \leq q, \\ p + (1-p)F_{GPD|q, \sigma, \xi}(y), & y > q, \end{cases} \quad (13)$$

where $p = P(Y(\mathbf{s}_i) \leq q)$, q is a high threshold, and $F_{GPD|q,\sigma,\xi}(y) = 1 - [1 + \xi(y - q)/\sigma]^{-1/\xi}$, with support $\{y \geq q : 1 + \xi(y - q)/\sigma \geq 0\}$. Conditioning on the smooth process $X^*(\mathbf{s}_i) = x^*(\mathbf{s}_i)$, which was not truncated, the censored likelihood for an observation $y(\mathbf{s}_i)$ can be derived as

$$\varphi(y(\mathbf{s}_i) | \mathbf{X}^*, \tau^2, \boldsymbol{\theta}_R, \boldsymbol{\theta}_{GPD}, p) = \begin{cases} \Phi\left(\frac{F_{X|\boldsymbol{\theta}_R, \tau^2}^{-1}(p) - x^*(\mathbf{s}_i)}{\tau}\right) & \text{if } y(\mathbf{s}_i) \leq q, \\ \frac{1}{\tau} \phi\left(\frac{T(y(\mathbf{s}_i)) - x^*(\mathbf{s}_i)}{\tau}\right) \cdot \frac{f_{Y|q,\sigma,\xi}(y(\mathbf{s}_i))}{f_{X|\boldsymbol{\theta}_R, \tau^2}(T(y(\mathbf{s}_i)))} & \text{if } y(\mathbf{s}_i) > q, \end{cases} \quad (14)$$

where $\boldsymbol{\theta}_{GPD} = (\sigma, \xi)$. Note that there are only univariate calculations required in (14), compared to (7) for which we have to estimate the $|\mathcal{C}|$ -dimensional Gaussian distribution functions. In addition, since Y_i and Y_k are independent conditioning on the smooth process ($i \neq k$), the joint likelihood of the whole vector \mathbf{y} is simply $f_{\mathbf{Y}}(\mathbf{y}) = \prod_{i=1}^D \varphi(y(\mathbf{s}_i) | \mathbf{X}^*, \tau^2, \boldsymbol{\theta}_R, \boldsymbol{\theta}_{GPD}, p)$. Likelihoods for independent time replicates are simply multiplied together, and the proportion of censored observations p can be treated as a known parameter or an unknown parameter that enters the hierarchical model. See Appendix A.2 for a complete statement of the hierarchical model. The priors for the model parameters are

$$\tau^2 \sim \text{IG}(a, b), \quad \sigma \sim \text{halfCauchy}(1), \quad \xi \sim \text{Unif}(-0.5, 0.5), \quad (15)$$

where $\text{halfCauchy}(1)$ refers to the positively truncated standard Cauchy distribution (see Gelman, 2006). We implement this methodology for two different Gaussian scale mixture models: that of Huser and Wadsworth (2019), and that of Huser et al. (2017). The priors for $\boldsymbol{\theta}_R$ are $\delta \sim U(0, 1)$ for the former, and $\beta \sim \text{halfCauchy}(1)$ for the latter. The prior for $\boldsymbol{\theta}_C$ depends on the choice of the covariance function. In our implementations, we adopt the Matérn covariance function with $\boldsymbol{\theta}_C = (\rho, \nu)$, where ρ is the range parameter, and ν is the smoothness parameter. The prior is then set to subject to two independent half Cauchy distributions with scale parameter 1.

4.2 Gibbs Sampler

To estimate the posterior distribution of the model parameters, we apply random walk Metropolis (RWM) algorithm using Log-Adaptive Proposals (LAP) as our adaptive tuning strategy (Shaby and Wells, 2010). Since conjugate priors are not available, we use random walk Metropolis-Hastings update steps.

At each MCMC iteration, we first update the smooth process X^* conditioning on the true values for all non-censored sites, current values for X^* and all other model parameters:

$$\varphi(\mathbf{X}^* | \dots) \propto \prod_{i=1}^D \varphi(y_i | \mathbf{X}^*, \tau^2, \boldsymbol{\theta}_R, \boldsymbol{\theta}_{GPD}, p) \cdot \varphi(\mathbf{X}^* | R, \boldsymbol{\theta}_C), \quad (16)$$

where the likelihood function of X^* conditioning on the random scaling factor R is calculated in Appendix A.2. We then update R using its conditional posterior distribution

$$\varphi(R | \dots) \propto \varphi(\mathbf{X}^* | R, \boldsymbol{\theta}_C) \cdot \varphi(R | \boldsymbol{\theta}_R).$$

Since time independence is assumed, we can update \mathbf{X}_t^* and R_t in a parallel fashion across $t = 1, \dots, T$. The other parameters are updated similarly using adaptively-tuned random walk Metropolis-Hastings updates, with the likelihood (14) multiplied by the corresponding priors in (15).

5 Simulation Studies

In this section, we present simulation results and conduct coverage analysis to investigate, firstly, whether the MCMC procedure is able to draw accurate inference on model parameters, and secondly, in the case of the modified version of model (5), to check whether our model captures asymptotic dependence characteristics correctly even when the data-generating model is different from the fitted model.

5.1 Parameter Estimation

To verify the accuracy of inference made by MCMC sampling, we generate data from model (8) in Section 3.1, in both the special case of the Huser et al. (2017) model (4) and the special case of the Huser and Wadsworth (2019) model (5), with the addition of nugget terms. In both cases, we use $D = 200$ sites uniformly drawn from the unit square $[0, 1]^2$, with the latent Gaussian processes $Z(\mathbf{s})$ are generated using a Matérn covariance with smoothness parameter $\nu = 3/2$. The characteristic length scale parameter is set to $\rho = 0.05$ in the case of model (4) and $\rho = 0.1$ in the case of model (5).

For model (4), we use $T = 20$ independent temporal replications, and set $\beta = 0.5$, $\gamma = 1$ (which is fixed during estimation), and nugget variance $\tau^2 = 0.2^2$. This represents a challenging case, with a fairly long tail and nugget that is small relative to the scale of $X^*(\mathbf{s})$. For model (5), we use $T = 40$ independent temporal replications, and set the nugget variance to $\tau^2 = 3^2$. Because the latent $Z(\mathbf{s})$ is transformed to Pareto, $\tau^2 = 3^2$ is still small compared to the scale of smooth process $X^*(\mathbf{s})$; see the Supplementary Material for a more in-depth discussion on the effects of τ^2 . We consider two different scenarios for the dependence parameter δ : $\delta = 0.3$ and $\delta = 0.7$, corresponding to asymptotic independence and asymptotic dependence respectively. Finally, in all cases the processes are marginally transformed to generalized Pareto distribution with $(q, \sigma, \xi) = (11, 1, 0)$, where q and $p = 0.8$, the proportion of censored observations, are treated as known parameters.

The attenuation constants used in the LAP algorithm are $c_0 = 10$, $c_1 = 0.8$. The prior for ρ is `halfCauchy(1)`, and the priors for the other parameters are specified in (15), where $(a, b) = (0.1, 0.1)$ so that the prior for τ^2 is fairly noninformative. We ran each MCMC chain for 400,000 iterations and thinned the results by a factor of 10. The parallelism of updating \mathbf{X}_t^* and R_t is implemented in R via the `foreach` routine with `doParallel` package as a backend (Microsoft Corporation and Weston, 2017).

$$\beta = 0.5$$

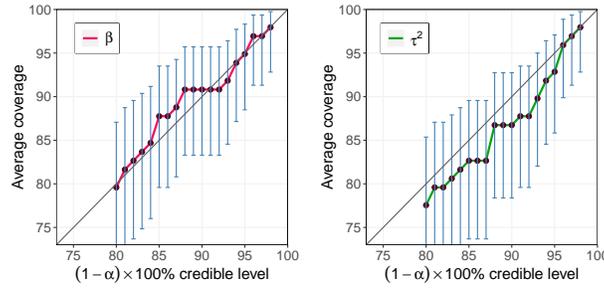


Figure 2: Empirical coverage rates of credible intervals for the [Huser et al. \(2017\)](#) model: $\beta = 0.5$, $\tau^2 = 0.2^2$. The error bars are 95% binomial confidence intervals for the coverage probability.

5.2 Coverage Analysis

We now study the coverage properties of the posterior inference based on the MCMC sampler for the posterior credible intervals with 100 simulated datasets drawn from the [Huser et al. \(2017\)](#) and [Huser and Wadsworth \(2019\)](#) models, under each of the scenarios described in the previous section.

Figures 2 and 3 shows the empirical coverage rates of highest posterior density credible intervals of several sizes, along with standard binomial confidence intervals. In all cases, we can see that the sampler performed well in generating posterior inference that is well calibrated, with close to nominal frequentist coverage. The coverage for larger δ is maybe slightly different than nominal for large α , but overall the results are quite good.

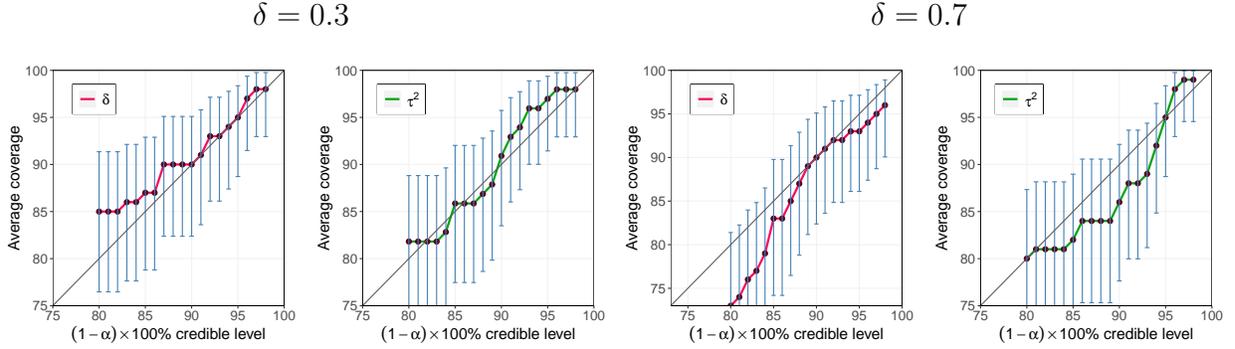


Figure 3: Empirical coverage rates of credible intervals for the [Huser and Wadsworth \(2019\)](#) model under two designs: $\delta = 0.3$ (left) and $\delta = 0.7$ (right). For both designs, $\tau^2 = 3^2$. The error bars are 95% binomial confidence intervals for the coverage probability.

5.3 Simulation with Mis-specified Models

We now fit our model to data generated from other distributions to validate its ability to capture the tail dependence characteristics under mis-specification. We simulate datasets from models referenced in Section 2.1, and use the sampler described in Section 4.2 to fit model (8) with the [Huser and Wadsworth \(2019\)](#) latent process. The data were generated using four different simulation designs:

- Skew- t process from [Morris et al. \(2017a\)](#) with $(a, b) = (6, 16)$ (asymptotically dependent);
- Gaussian scale mixture from [Huser et al. \(2017\)](#) with $\beta = 0$ (asymptotically dependent);
- Gaussian scale mixture from [Huser et al. \(2017\)](#) with $\beta = 1$ or 5 (asymptotically independent).

For each simulation design, we simulate a single dataset using $D = 100$ locations uniformly distributed on $[0, 1]^2$ and $T = 40$ independent time replicates. The Matérn covariance function with $\nu = 3/2$ and $\rho = 1$ is again specified for the latent Gaussian processes. For the skew- t process, $R_t \sim \text{IG}(3, 8)$ to give a t distribution with 6 degrees of freedom, and $\lambda = 3$ to simulate moderate skewness. For the last two designs, the R_t were generated as described in (4), with $\gamma = 1$.

To obtain good starting values for the latent smooth process \mathbf{X}^* for MCMC, we first marginally transform the simulated data to noisy scale mixture variables \mathbf{X} independently at each location using the following procedure. Following the semi-parametric procedure of [Coles and Tawn \(1991\)](#), we estimate each marginal distribution as a blend of the generalized Pareto distribution function above a high marginal threshold, and the empirical distribution function below that threshold. Fixing initial values for (δ, τ^2) , we then transform the margins to noisy scale mixtures via $X_{jt} = F_{X|\delta, \tau^2}^{-1} \circ \hat{F}_{\mathbf{s}_j}(Y_{jt})$. The next step is to run a Metropolis algorithm using the full conditional distribution $\varphi(\mathbf{X}^* | \dots)$ (see (16)) 100 times and save the last random walk states as initial values for \mathbf{X}^* . This procedure is also used for data analysis in Section 6. Finally, with initial values in hand, the datasets from each design were fit using a fully Bayesian approach that simultaneously updates marginal and spatial dependence parameters.

Figure 4 displays the nonparametric and model-based estimates of the upper tail dependence $\chi_u(h)$ defined in (1). To generate nonparametric estimates of $\chi_u(h)$ at distance $h = \|\mathbf{s}_1 - \mathbf{s}_2\|$, we look at all pairs of points whose locations are h apart (within a small tolerance $\varepsilon = 0.005$), and compute the ratio of empirical probabilities $\hat{\chi}_u(h)$. This is similar to an empirical variogram estimator. The nonparametric confidence envelopes are obtained by computing pointwise binomial confidence intervals, pretending that each pair of points is independent from each other pair of points. For parametric estimates, we take samples

from the converged MCMC chain, and use parameters from each iteration to simulate 10^8 pairs of $(X(\mathbf{s}_1), X(\mathbf{s}_2))$ based on our model to generate a smooth $\chi_u(h)$ estimate for each MCMC iteration. Then, combining across MCMC iterations, we compute the pointwise average curves and their credible bands. The results in Figure 4 demonstrate that our model provides a sensible approximation to the extremal dependence structure of the misspecified models. Borrowing strength across locations, the parametric estimators of $\chi_u(h)$ are much more reliable than the nonparametric ones in that they are able to discriminate between the two asymptotic classes via estimating a relatively small number of parameters. Especially for the Huser et al. (2017) models, the dependence strength diminishes gradually as β becomes greater, eventually resembling the Gaussian copula. With the extra tail flexibility, our model is able to accurately capture these features.

6 Data Analysis

We consider daily observations of Fosberg Fire Weather Index (FFWI) from 1974 to 2015 at 93 monitoring stations over parts of the Great Plains, mainly from Central Great Plains to South Texas Plains (Dunn et al., 2012). Figure 9 shows the observation locations as black triangles. The FFWI aims to quantify potential wildfire threat. It is a single number summary calculated from temperature, wind speed, and relative humidity; larger index values signify higher flame lengths and more rapid drying (Fosberg, 1978). Due to human activity and changes in the grassland ecosystem, the Great Plains region is becoming an important high-risk region of large wildfires. According to a in-depth study conducted by Donovan et al. (2017), the average total area burned by fires in Great Plains region between 2005 and 2014 was in the millions of hectares per year. In 2017, 809,380 hectares were lost to wildfires in a single week in Texas, Oklahoma, and Kansas alone (Herskovitz,

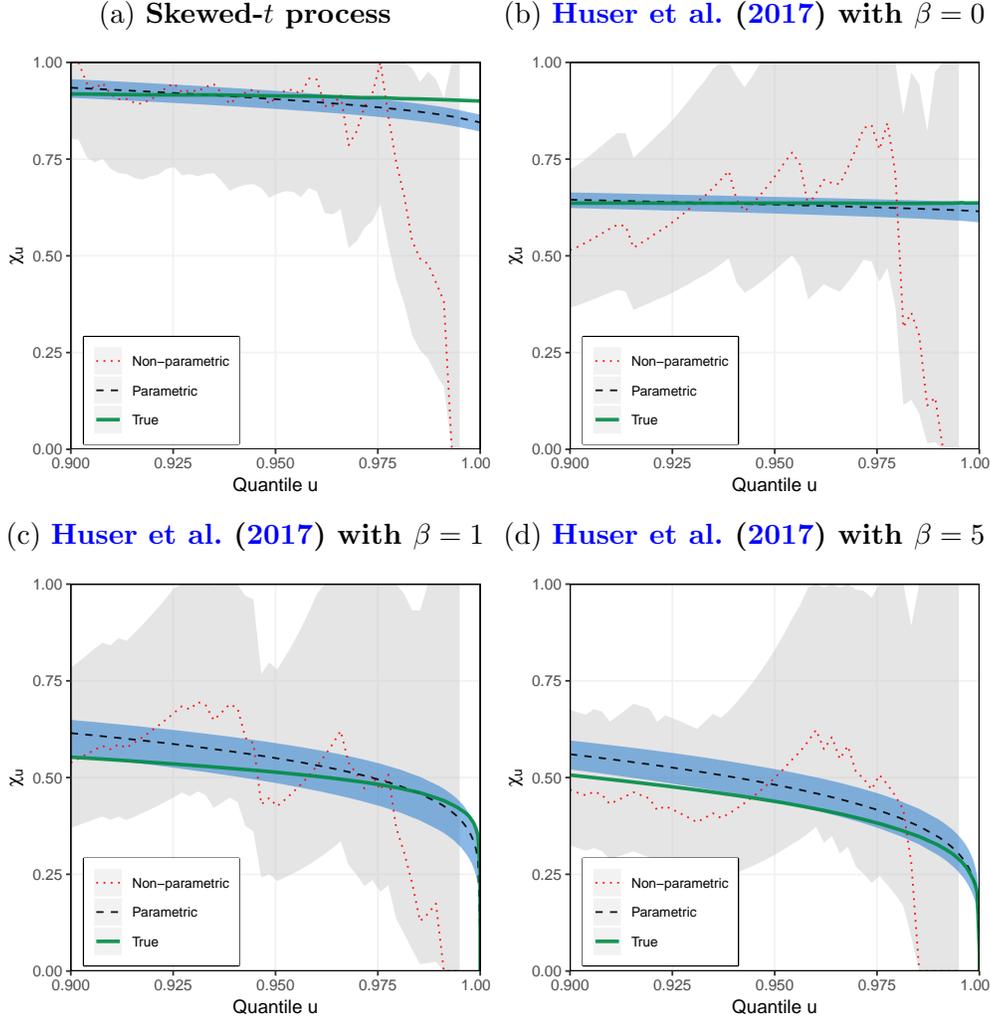


Figure 4: Estimated coefficients $\chi_u(h)$, $u \in [0.9, 1]$, for two points at distance $\|\mathbf{s}_1 - \mathbf{s}_2\| = 1$, using data simulated from the mis-specified models with Matérn correlation function with $\nu = 3/2$ and $\rho = 1$. The two scenarios in the top row are asymptotically dependent, while the scenarios in the bottom row are asymptotically independent. Each simulated data set has 100 uniform locations in $[0, 1]^2$, and 40 time replicates. Solid green lines show true $\chi_u(h)$ function. The black dashed lines show the averaged curve from the posterior samples of fitted model, while the blue shaded areas are 95% credible envelopes. The red dashed lines show the nonparametric estimates from the simulated data sets, while the grey shaded areas are pointwise 95% binomial confidence intervals.

2017). As a consequence of the prevailing hot and dry air, wildfires in this region are particularly more concentrated in the spring, feasting on grasses made dry by long-term drought. Modeling the tail behavior of FFWI and studying the extremes of the process could have major implications for wildfire planning.

To ensure the independence over time and avoid seasonal effects, we take the maxima of the FFWI values over ten-day intervals during the spring season. Figure 5 shows the 50-year return levels estimated using the block maxima with 10-year sliding windows for 12 randomly-selected stations. There is no clear evidence for a systematic increase or decrease in the return levels. Although treating meteorological data as constant over time is often problematic, particularly for temperature data, the behavior in Figure 5 suggests that an assumption of constant marginal parameters in time is appropriate.

To account for the physical features of the terrain in the Great Plains, we describe the scale parameter in $\boldsymbol{\theta}_{GPD}$ by the trend surface:

$$\sigma(\mathbf{s}) = \beta_0 + \beta_1 \text{lon}(\mathbf{s}) + \beta_2 \text{lat}(\mathbf{s}) \quad (17)$$

where $\text{lon}(\mathbf{s})$ and $\text{lat}(\mathbf{s})$ are the longitude and latitude of the stations at which the data are observed. We constrain the joint prior of $(\beta_0, \beta_1, \beta_2)'$ such that the support of $\sigma(\mathbf{s})$ is the positive real line. We model the shape parameter $\xi(\mathbf{s})$ as constant over the spatial domain, as suggested by exploratory analysis (see Appendix B.1).

Similar to the procedure in Section 5.3, to obtain starting values for MCMC, we fit generalized Pareto distributions to model events above the 98% quantile, u_{98} , and empirical distributions to those below u_{98} , of the time series at each station separately, and then use the fitted models to transform the data to have noisy scale mixture distributions. We then ran the MCMC chain for 50,000 iterations thinned by 10 steps and discarded a burn-in period of 25,000 iterations.

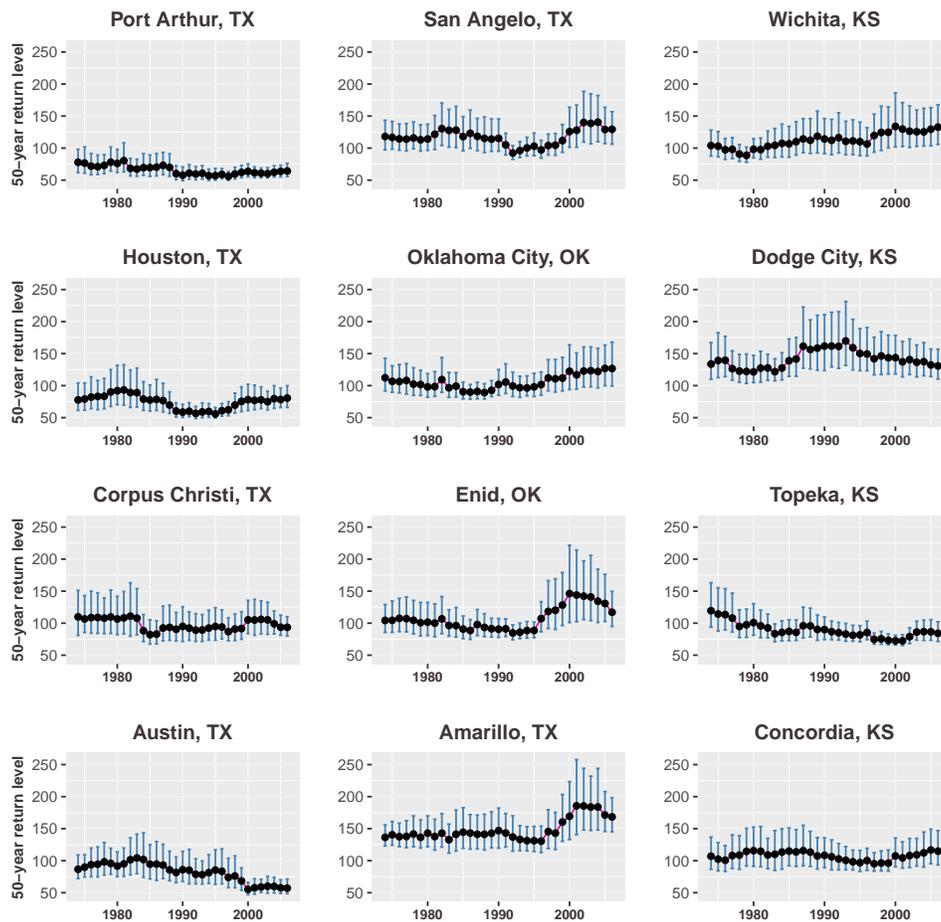


Figure 5: Point estimates and 95% confidence intervals for 50-year return levels of FFWI at 12 randomly selected stations. For each station and for each year, we estimate the return level using annual maxima in a 10-year sliding window. There is no evident systematic trend in the return level, so the assumption of marginal parameters that are constant in time is deemed appropriate.

	ρ	ν	τ	δ
Posterior Mean	0.504	0.344	1.837	0.530
95% Credible Interval	(0.483, 0.524)	(0.324, 0.362)	(0.331, 3.963)	(0.528, 0.531)
	β_0	β_1	β_2	ξ
Posterior Mean	-17.531	-0.284	-0.1873	0.3392
95% Credible Interval	(-17.614, -17.401)	(-0.285, -0.283)	(-0.1874, -0.1871)	(0.3385, 0.3396)

Table 1: Posterior mean and 95% highest posterior density credible interval.

6.1 Model Evaluation

First and foremost, we examine whether the estimate δ falls within $(0, 1/2]$ or $(1/2, 1)$ in accordance to whether the data-generating process is asymptotically independent or dependent. Table 1 reports the posterior means and 95% credible intervals for the model parameters. Trace plots for δ and τ can be found in Appendix B.2. For this dataset, the MCMC results show that the range of the mixing parameter δ is close to the interface between the two dependence class, while demonstrating asymptotic dependence. Nonetheless, the value of δ being close to $1/2$ means that $\chi_u(h)$ will still decrease with u before eventually reaching a positive limit. In addition, the estimated value of the smoothness parameter ν stands out, as it is very small, indicating very rough process realizations.

To better evaluate the model fit, we randomly hold out 5 stations for validation purposes, and exclude them from the MCMC analysis; see the red points in Figure 9. We then compare the empirical distributions of mean, and maxima for each time point at the 5 held-out stations with those simulated with parameters from MCMC samples. Though we modeled our data as censored observations, the model may still work a bit further into the center of the distribution. Results are displayed in Figure 6 where we only show values exceeding 80% threshold for mean, and 90% threshold for maxima. We can see that the fit

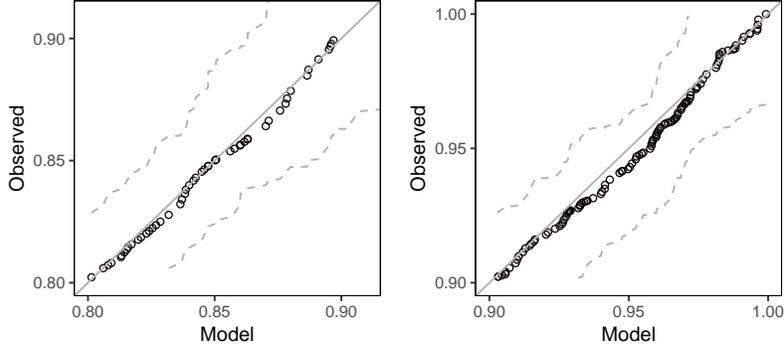


Figure 6: Comparisons of the observed and predicted mean (left) and maxima (right) for 5 locations held out for model validations. Overall 95% confidence envelopes are also shown. Note that the values are transformed marginally to uniform using the empirical distribution functions at each location.

displays a good match against the observed mean and maxima in the upper quantiles.

As a comparison, we change \mathbf{X}^* to be a t process, and re-fit the model using MCMC. We apply proper scoring rules (Gneiting and Raftery, 2007), log scores and continuously-ranked probability scores (CRPS), to compare the quality of probabilistic forecasts. While running MCMC, we interpolate the latent process \mathbf{X}^* at the held-out locations for each iteration using the full conditional likelihood. Plugging the predictive draws at the held-out observations into the equation (14), we obtain the log score (simply the log-likelihood) as

$$V = \log \left(\prod_{t=1}^T \prod_{j=1}^5 \phi(y_t(\mathbf{r}_j) | \mathbf{X}_t^*, \dots) \right) \tag{18}$$

where $\{\mathbf{r}_j | j = 1, \dots, 5\}$ are the validation stations. The left panel of Figure 7 compares the log scores between two models, showing that the transformed Gaussian scale mixture process clearly outperforms the t process. Additionally, we calculate the CPRS (Matheson

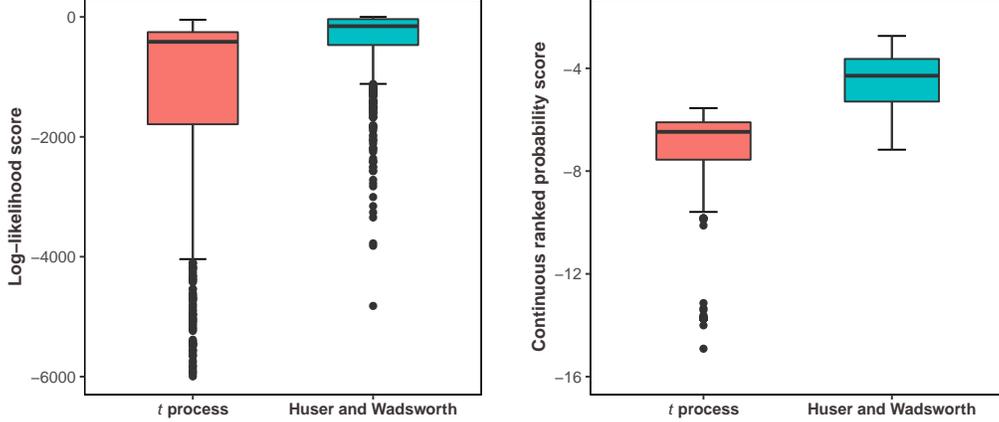


Figure 7: Comparisons of the log-likelihood scores and CRPS between our model and a similar model with t latent process. In both panels, higher values indicate better model fit.

and Winkler, 1976) for both models,

$$\text{CRPS}(F_Y, y) = - \int_{-\infty}^{\infty} (F_Y(z) - \mathbb{1}\{y \leq z\})^2 dz \quad (19)$$

where F_Y is the marginal distribution estimated using parameters using one MCMC iteration, and y is the observed value. The right panel of Figure 7 shows the averaged CRPS for the two models, where our model clearly has better results.

Similar to Figure 4, we show the empirical and model-based values of $\chi_u(h)$ and $\bar{\chi}_u(h)$ for the block maxima of the FFWI in Figure 8, which confirms that our model captures the extremal spatial dependence in the data quite well. The quantity $\bar{\chi}_u(h)$ is an alternative dependence measure useful in the situation $\chi(h) = 0$, and is defined as

$$\bar{\chi}_h(u) = \frac{2 \log(1 - u)}{\log P(F_j(X_j) > u, F_k(X_k) > u)} - 1 \rightarrow 2\eta_X(h) - 1, \quad u \rightarrow 1.$$

Recall that the posterior mean for δ is greater than 0.5, which means $\bar{\chi}_h(u) \rightarrow 1$ as $u \rightarrow 1$. Interestingly, the black dashed curve of the right panel of Figure 8 seems to have a limit less

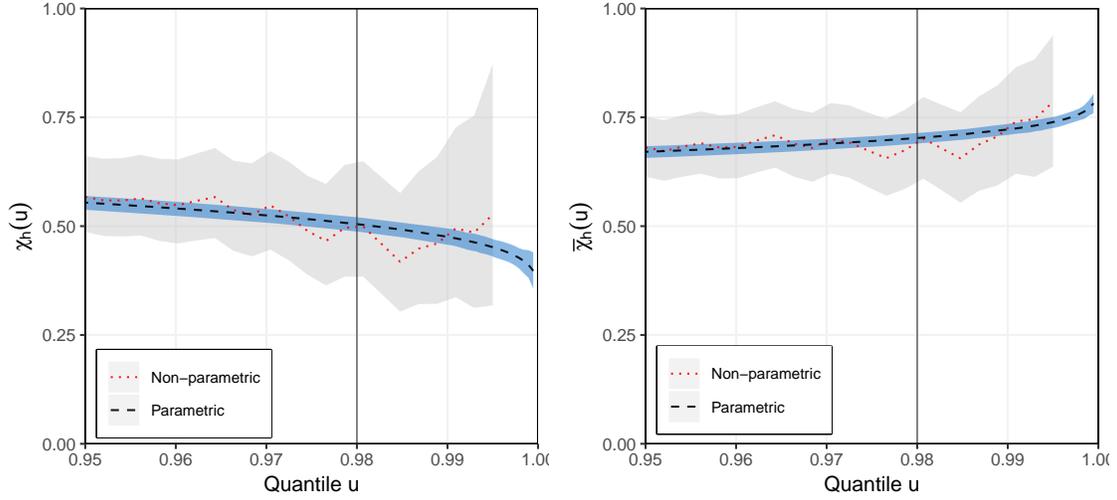


Figure 8: Empirical estimates (dashed red lines and gray envelopes) of $\chi_u(h)$ (left) and $\bar{\chi}_u(h)$ (right) for the FFWI, $u \in [0.95, 0.995]$, for two points at distance $\|\mathbf{s}_1 - \mathbf{s}_2\| = 20km$. The blue shaded areas are 95% credible envelopes obtained from the posterior samples of fitted model, and the black dashed lines show the averaged curve. The vertical line is the threshold used when fitting the dependence model.

than 1. This is because to attain the correct limit in this case, we would need to compute $\bar{\chi}_h(u)$ for values of u that are *very* close to 1, which is very difficult numerically.

6.2 Results

To get an idea of what a realization of the fitted process looks like, the left panel of Figure 9 shows one realization of the latent $X(s)$ scale mixture process using parameters from one MCMC iteration. The extreme values are mainly concentrated in two small regions. The right panel shows the same realization, now marginally transformed to the scale of the FFWI values. Since we modeled our data as partially censored observations, the map here

only displays the areas where threshold exceedances are observed.

A quantity of great interest is areal exceedance probabilities, which represent the amount of territory simultaneously at extreme risk for wildfire. To obtain a Monte Carlo estimate of these joint probabilities, we use parameters from each MCMC sample to simulate 100 processes (on a $15km \times 15km$ grid), and calculate the total area that has FFWDI over a designated threshold. Figure 10 shows the results. These curves represent total area at risk for various FFWDI thresholds. The curves for the higher thresholds decay faster than those for the lower thresholds, which confirms that extreme events simultaneously occurring across large areas becomes less common when the threshold increases. This also shows that the joint tail of the fire threat index exhibits a weakening dependence structure, with more extreme events being more localized. It is not possible to capture this behavior using limiting extreme value models like max-stable or generalized Pareto processes.

7 Discussion

In this paper, we have proposed a new modeling approach, based on the class of transformed Gaussian scale mixture models, which includes those recently proposed by [Huser et al. \(2017\)](#) and [Huser and Wadsworth \(2019\)](#). We added a measurement error to the mixture, hence avoiding the need to calculate the onerous $|\mathcal{C}|$ -dimensional Gaussian distribution function when dealing with the censored likelihood. We also circumvent the need to draw from a high-dimensional truncated distribution by treating the smooth process as latent and updating repeatedly using MCMC. In addition to its computational advantages, the presence of a measurement error term may also make the model more realistic for data collected by real-world instruments. Indeed, nugget effects are ubiquitous in environmental statistics, not just to represent measurements errors, but also as a result of small scale

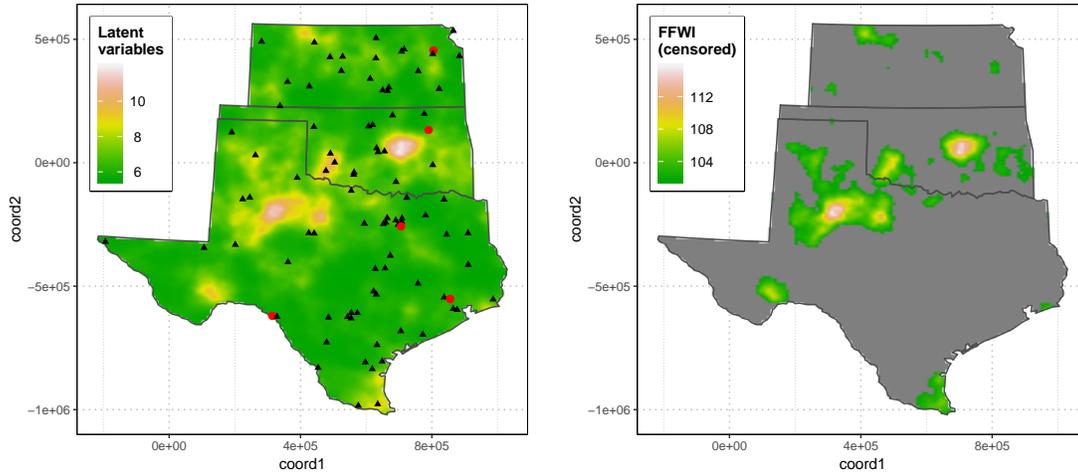


Figure 9: A realization of the process generated from parameters from one MCMC iteration. Left panel shows the latent scale mixture process $X(s)$ in log scale to better visualize spatial patterns, with points showing the stations of the 93 gauges. For model checking, the 88 stations marked by triangles were used to fit the models, and those marked with circles were used to validate the models. Right panel shows the same realization, marginally transformed to the scale of the original FFWI data, and then censored. A projected coordinate reference system, NAD83(HARN), is used here so that the axes are in units of meters.

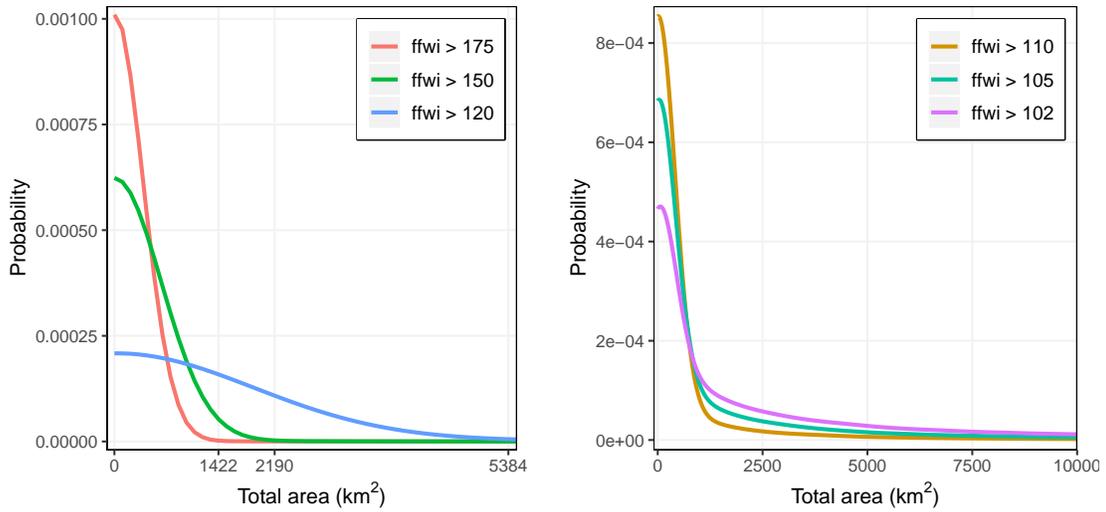


Figure 10: The distribution of the total area that is over a certain threshold. For each MCMC iteration, we simulate 100 processes on a $15km \times 15km$ grid, and count the number of threshold exceedances. The density of the total area is then estimated using Gaussian kernels.

effects that are not included in the large scale model for spatial dependence.

Even with the presence of the measurement error, the model is still able to capture qualitatively different types of sub-asymptotic dependence behavior of spatial processes. In the case of our modification of the [Huser and Wadsworth \(2019\)](#) model, a smooth transition between both extremal dependence paradigms takes place in the interior of the parameter space, which enables inference of the dependence class in a simple manner. We proved that all the appealing asymptotic properties found in the original smooth process are preserved in the modified model, regardless of the size of the measurement error variance.

In the data analysis presented in [Section 6](#), the credible interval for ξ is narrower than we might expect, and it is not clear why this might be the case. In the [Supplementary Material](#), we present some evidence from simulation that the sharpness of the marginal posterior for ξ may indeed accurately reflect the information contained in the data, as viewed through the lens of the model. Nevertheless, we believe this phenomenon requires further investigation.

The model allows inference on spatial extreme-value datasets with relatively large numbers of locations. The computational limitations are similar to those of conventional spatial Gaussian process models. We have defined the model conditionally as a Bayesian hierarchical model, for which standard MCMC techniques can be used to fit the data. Computation is facilitated greatly by parallelizing over time t and migrating some basic linear algebra to C/C++ via [Rcpp](#). Even so, the lack of closed form marginal transformations creates a significant (though embarrassingly parallel) computational challenge that scales with the total number of exceedances, rather than the usual case of scaling with the number of spatial locations.

Despite easing computational limitations associated with the [Huser et al. \(2017\)](#) and [Huser and Wadsworth \(2019\)](#) models, our modified versions inherit the same theoretical

limitations. Neither model is able to account for the possibility of independence between observations as the distance between sites becomes large, nor are they able to transition from asymptotic dependence at short range to asymptotic independence at longer range. [Wadsworth and Tawn \(2019\)](#) presents an alternative approach to modeling spatial extremes that begins to address these issues.

Another interesting possibility to explore would be to include the nugget term inside the link function g . This would result in closed-form marginal distributions in some cases, perhaps making computations easier. However, it would change the dependence structure in ways that would not vanish in the limit, which is not the behavior we were aiming for here, but could be useful nonetheless.

Acknowledgements

We gratefully acknowledge support from NSF grant DMS-1752280 and EPSRC grant EP/P002838/1, along with seed grants from the Institute for CyberScience and the Institute for Energy and the Environment at Pennsylvania State University. Computations for this research were performed on the Pennsylvania State University’s Institute for CyberScience Advanced CyberInfrastructure (ICS-ACI).

References

- Asmussen, S., Hashorva, E., Laub, P. J. and Taimre, T. (2018), ‘Tail asymptotics of light-tailed Weibull-like sums’, Probability and Mathematical Statistics **37**.
- Balkema, A. A., Klüppelberg, C. and Resnick, S. I. (1993), ‘Densities with gaussian tails’, Proceedings of the London Mathematical Society **3**(3), 568–588.

- Breiman, L. (1965), ‘On some limit theorems similar to the arc-sin law’, Theory of Probability & Its Applications **10**(2), 323–331.
- Cline, D. B. and Samorodnitsky, G. (1994), ‘Subexponentiality of the product of independent random variables’, Stochastic Processes and their Applications **49**(1), 75–98.
- Coles, S. G. and Tawn, J. A. (1991), ‘Modelling extreme multivariate events’, Journal of the Royal Statistical Society: Series B (Methodological) **53**(2), 377–392.
- Donovan, V. M., Wonkka, C. L. and Twidwell, D. (2017), ‘Surging wildfire activity in a grassland biome’, Geophysical Research Letters **44**(12), 5986–5993.
- Dunn, R. J., Willett, K. M., Thorne, P. W., Woolley, E. V., Durre, I., Dai, A., Parker, D. E. and Vose, R. E. (2012), ‘Hadisd: a quality-controlled global synoptic report database for selected variables at long-term stations from 1973–2011’, Climate of the Past Discussions **8**(3).
- Engelke, S., Opitz, T. and Wadsworth, J. (2019), ‘Extremal dependence of random scale constructions’, Extremes **22**(4), 623–666.
- Ferreira, A. and de Haan, L. (2014), ‘The generalized pareto process; with a view towards application and simulation’, Bernoulli **20**(4), 1717–1737.
- Flegal, J. M., Jones, G. L. et al. (2010), ‘Batch means and spectral variance estimators in markov chain monte carlo’, The Annals of Statistics **38**(2), 1034–1070.
- Fosberg, M. A. (1978), ‘Weather in wildland fire management: the fire weather index’, US For Serv Reprints of articles by FS employees .

- Gelman, A. (2006), ‘Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper)’, Bayesian Anal. **1**(3), 515–533.
URL: <https://doi.org/10.1214/06-BA117A>
- Gneiting, T. and Raftery, A. E. (2007), ‘Strictly proper scoring rules, prediction, and estimation’, Journal of the American Statistical Association **102**(477), 359–378.
- Herskovitz, J. (2017), ‘U.S. plains wildfires leave thousands of cattle dead’, Reuters . <https://www.reuters.com/article/us-usa-wildfires-idUSKBN16G2XG>.
- Huser, R., Davison, A. C. and Genton, M. G. (2016), ‘Likelihood estimators for multivariate extremes’, Extremes **19**(1), 79–103.
- Huser, R., Opitz, T. and Thibaud, E. (2017), ‘Bridging asymptotic independence and dependence in spatial extremes using gaussian scale mixtures’, Spatial Statistics **21**, 166–186.
- Huser, R. and Wadsworth, J. L. (2019), ‘Modeling spatial processes with unknown extremal dependence class’, Journal of the American Statistical Association **114**, 434–444.
- Joe, H. (1993), ‘Parametric families of multivariate distributions with given margins’, Journal of multivariate analysis **46**(2), 262–282.
- Ledford, A. W. and Tawn, J. A. (1996), ‘Statistics for near independence in multivariate extreme values’, Biometrika **83**(1), 169–187.
- Matheson, J. E. and Winkler, R. L. (1976), ‘Scoring rules for continuous probability distributions’, Management science **22**(10), 1087–1096.

- Microsoft Corporation and Weston, S. (2017), doParallel: Foreach Parallel Adaptor for the 'parallel' Package. R package version 1.0.11.
- Morris, S. A., Reich, B. J., Thibaud, E. and Cooley, D. (2017a), 'A space-time skew-t model for threshold exceedances', Biometrics **73**(3), 749–758.
- Morris, S. A., Reich, B. J., Thibaud, E. and Cooley, D. (2017b), 'A space-time skew-t model for threshold exceedances', Biometrics .
- Opitz, T. (2016), 'Modeling asymptotically independent spatial extremes based on laplace random fields', Spatial Statistics **16**, 1–18.
- Padoan, S. (2011), 'Multivariate extreme models based on underlying skew-t and skew-normal distributions', J. Multivariate Anal. **102**, 977–991.
- Røislien, J. and Omre, H. (2006), 'T-distributed random fields: A parametric model for heavy-tailed well log data', Math. Geol. **38**(7), 821–849.
- Rootzén, H., Segers, J. and Wadsworth, J. L. (2018), 'Multivariate generalized pareto distributions: Parametrizations, representations, and properties', Journal of Multivariate Analysis **165**, 117–131.
- Shaby, B. and Wells, M. T. (2010), 'Exploring an adaptive metropolis algorithm', Currently under review **1**, 1–17.
- Thibaud, E. and Opitz, T. (2015), 'Efficient inference and simulation for elliptical pareto processes', Biometrika **102**(4), 855–870.
- Wadsworth, J. L. and Tawn, J. A. (2012), 'Dependence modelling for spatial extremes', Biometrika **99**(2), 253–272.

Wadsworth, J. L. and Tawn, J. A. (2019), 'Higher-dimensional spatial extremes via single site conditioning'.

A Technical appendix

A.1 Proof of Proposition 3.1

For the proof of Proposition 3.1, we begin by recalling useful results from the literature.

The first is Breiman's lemma, see e.g. [Breiman \(1965\)](#) and [Cline and Samorodnitsky \(1994\)](#), and a corollary for sums of a regularly varying and light-tailed random variables.

Lemma A.1 (Breiman). *Suppose that $Q = ST$ where $P(S > s) \in RV_{-\alpha}$, $\alpha \geq 0$ and $E(T^{\alpha+\delta}) < \infty$ for some $\delta > 0$. Then*

$$P(Q > x) \sim E(T^\alpha)P(S > x), \quad x \rightarrow \infty.$$

Corollary A.1.1. *Suppose that $P(X > x) \in RV_{-\alpha}$, $\alpha \in (0, \infty)$, and ϵ is a random variable such that $E(e^{\delta\epsilon}) < \infty$ for some $\delta > 0$. Then*

$$P(X + \epsilon > x) \sim P(X > x), \quad x \rightarrow \infty.$$

Proof. Since $P(X > x) =: \bar{F}_X(x) \in RV_{-\alpha}$, for positive finite α , $P(e^X > x) = \bar{F}_X(\log(x)) \in RV_0$. By assumption $E(e^{\delta\epsilon}) < \infty$, and so applying Breiman's Lemma to $e^X e^\epsilon$ yields

$$P(e^X e^\epsilon > x) \sim P(e^X > x), \quad x \rightarrow \infty,$$

from which the result follows. □

The second result relates to sums of Weibull-tailed variables, i.e., those with survival functions satisfying (10) and the associated density

$$f_{X^*}(x) \sim v(x) \exp(-\theta x^\alpha), \quad v(x) = u(x)(\theta \alpha x^{\alpha-1}) \in RV_{\kappa+\alpha-1}. \quad (20)$$

The following lemma can be verified directly from Theorem 4.1 of [Asmussen et al. \(2018\)](#), by identifying that the conditions in Section 4 of that paper hold for this subclass when $\alpha > 1$.

Lemma A.2 (Asmussen et al. (2018); Balkema et al. (1993)). Let Y_1, Y_2 be variables with density (20), with regularly varying functions v_j, u_j and parameters $\theta_j > 0, \alpha_j > 1, j = 1, 2$. Then the density and survival function of the convolution satisfy

$$f_{Y_1+Y_2}(x) \sim v_+(x) \exp\{-\psi_+(x)\}, \quad P(Y_1 + Y_2 > x) \sim u_+(x) \exp\{-\psi_+(x)\},$$

where $\psi_+(x) = \theta_1 q_1(x)^{\alpha_1} + \theta_2 q_2(x)^{\alpha_2}$, with $q_1(x), q_2(x)$ determined by solving

$$q_1 + q_2 = x, \quad \theta_1 \alpha_1 q_1^{\alpha_1-1} = \theta_2 \alpha_2 q_2^{\alpha_2-1}, \quad (21)$$

and

$$v_+(x) = \left(\frac{2\pi\psi_+''(x)}{\theta_1 \alpha_1 (\alpha_1 - 1) q_1^{\alpha_1-2} \theta_2 \alpha_2 (\alpha_2 - 1) q_2^{\alpha_2-2}} \right)^{1/2} v_1(q_1) v_2(q_2); \quad u_+(x) = v_+(x) / \psi_+'(x).$$

Finally we note also the following two useful inequalities

$$\begin{aligned} P(X_1^* + \min(\epsilon_1, \epsilon_2) > x, X_2^* + \min(\epsilon_1, \epsilon_2) > x) &\leq P(X_1^* + \epsilon_1 > x, X_2^* + \epsilon_2 > x) \leq \\ &P(X_1^* + \max(\epsilon_1, \epsilon_2) > x, X_2^* + \max(\epsilon_1, \epsilon_2) > x), \end{aligned} \quad (22)$$

$$\begin{aligned} P(\epsilon_1 + \min(X_1^*, X_2^*) > x, \epsilon_2 + \min(X_1^*, X_2^*) > x) &\leq P(X_1^* + \epsilon_1 > x, X_2^* + \epsilon_2 > x) \leq \\ &P(\epsilon_1 + \max(X_1^*, X_2^*) > x, \epsilon_2 + \max(X_1^*, X_2^*) > x), \end{aligned} \quad (23)$$

where the two lower bounds are equal. We can now prove Proposition 3.1.

Proof of Proposition 3.1. 1) When X^* has a regularly varying tail, Corollary A.1.1 provides that $P(X^* + \epsilon > x) \sim P(X^* > x)$. For the existence of η_{X^*} , the variable $\min(X_1^*, X_2^*)$ also has a regularly varying tail. Further,

$$E(e^{\delta \max(\epsilon_1, \epsilon_2)}) = E[\max(e^{\delta \epsilon_1}, e^{\delta \epsilon_2})] \leq E[e^{\delta \epsilon_1} + e^{\delta \epsilon_2}] = 2e^{\delta^2 \sigma^2 / 2} < \infty,$$

so Corollary A.1.1 thus gives $P(X_1^* + \max(\epsilon_1, \epsilon_2) > x, X_2^* + \max(\epsilon_1, \epsilon_2) > x) \sim P(X_1^* > x, X_2^* > x)$, and similarly for the lower bound. Hence $P(X_1^* + \epsilon_1 > x, X_2^* + \epsilon_2 > x) \sim P(X_1^* > x, X_2^* > x)$. Therefore it follows that $\chi_{X^*} = \chi_X$. If $f(x) \sim g(x)$ for $f(x) \rightarrow 0$ then $\log f(x) \sim \log g(x)$, so the result for η follows as well.

When X^* has a Weibull-like tail with Weibull index $\alpha < 1$, then $P(e^{X^*} > x) \sim u(\log x) \exp\{-\theta(\log x)^\alpha\} \in \text{RV}_0$, and the rest of the argument follows as above.

2) When $\alpha = 1$, $P(e^{X^*} > x) \sim u(\log x)x^{-\theta} \in \text{RV}_{-\theta}$. Breiman's lemma now yields

$$P(e^{X^*} e^\epsilon > x) \sim E(e^{\theta\epsilon})P(e^{X^*} > x). \quad (24)$$

As a consequence, inequality (22) does not lead to a precise asymptotic relationship for $P(X_1^* + \epsilon_1 > x, X_2^* + \epsilon_2 > x)$, but rather that it is asymptotically bounded within the range

$$[E[e^{\theta \min(\epsilon_1, \epsilon_2)}], E[e^{\theta \max(\epsilon_1, \epsilon_2)}]] P(\min(X_1^*, X_2^*) > x). \quad (25)$$

Combining (24) and (25), we get the stated bound for χ_{X^*} . It follows also that $\log P(X > x) \sim \log P(X^* > x)$ and $\log P(X_1 > x, X_2 > x) \sim \log P(X_1^* > x, X_2^* > x)$, and hence $\eta_{X^*} = \eta_X$.

3) Here we use Lemma A.2, where different values of q_1, q_2 are found for the three cases $\alpha \in (1, 2)$, $\alpha = 2$ and $\alpha > 2$. To make notation more obvious, we replace q_1, q_2 with q_*, q_ϵ for summation of X^*, ϵ , and $q_*^\wedge, q_\epsilon^\vee$ etc., for summation of $\min(X_1^*, X_2^*)$ and $\max(\epsilon_1, \epsilon_2)$, for example. The three cases are considered for the value of α_* ; we always have $\alpha_\epsilon = 2$.

a) For $\alpha_* \in (1, 2)$,

$$q_*(x) = x - \frac{\alpha_* \theta_*}{2\theta_\epsilon} x^{\alpha_* - 1} [1 + o(1)],$$

$$q_\epsilon(x) = \frac{\alpha_* \theta_*}{2\theta_\epsilon} x^{\alpha_* - 1} [1 + o(1)].$$

Therefore

$$\psi_+(x) = \theta_* q_*(x)^{\alpha_*} + \theta_\epsilon q_\epsilon(x)^2 \sim \theta_* x^{\alpha_*}, \quad (26)$$

which implies $-\log P(X^* + \epsilon > x) \sim -\log P(X^* > x)$.

To understand the joint behavior, we again use (22). Note that

$$P(\max(\epsilon_1, \epsilon_2) > x) = 1 - \Phi(x/\tau)^2 \sim 2\phi(x/\tau)/(x/\tau) = u_\epsilon^\vee(x) e^{-x^2/(2\tau^2)}.$$

$$P(\min(\epsilon_1, \epsilon_2) > x) = [1 - \Phi(x/\tau)]^2 \sim \phi(x/\tau)^2/(x/\tau)^2 = u_\epsilon^\wedge(x) e^{-x^2/\tau^2}.$$

Consequently, $\min(\epsilon_1, \epsilon_2)$ and $\max(\epsilon_1, \epsilon_2)$ both have $\alpha = 2$, with different θ . As θ_ϵ does not affect the leading order behavior of ψ_+ in (26), the tails of $\min(X_1^*, X_2^*) + \min(\epsilon_1, \epsilon_2)$ and $\min(X_1^*, X_2^*) + \max(\epsilon_1, \epsilon_2)$ both have $\psi_+(x) \sim \theta_*^\wedge x^{\alpha_*}$ in the exponent, where

$$P(\min(X_1^*, X_2^*) > x) \sim u_*^\wedge(x) \exp\{-\theta_*^\wedge x^{\alpha_*}\}.$$

Consequently $-\log P(X_1^* + \epsilon_1 > x, X_2^* + \epsilon_2 > x) \sim -\log P(X_1^* > x, X_2^* > x)$ and so $\eta_{X^*} = \eta_X$.

b) When $\alpha_* = \alpha_\epsilon = 2$, solving the equations in (21) provides

$$q_*(x) = \left(\frac{\theta_\epsilon}{\theta_* + \theta_\epsilon} \right) x, \quad q_\epsilon(x) = \left(\frac{\theta_*}{\theta_* + \theta_\epsilon} \right) x,$$

and

$$\psi_+(x) = \left(\frac{\theta_\epsilon \theta_*}{\theta_* + \theta_\epsilon} \right) x^2.$$

For the margins, and maximum $\theta_\epsilon = \theta_\epsilon^\vee = 1/(2\tau^2)$, whilst for the minimum, $\theta_\epsilon^\wedge = 1/\tau^2$.

We can now use both sets of inequalities (22) and (23), to give that the range of η_X is

$$\left[\frac{\theta_*}{2\theta_*^\wedge} \frac{\theta_*^\wedge + 1/\tau^2}{\theta_* + 1/(2\tau^2)}, \min \left\{ \frac{\theta_*}{\theta_*^\wedge} \frac{\theta_*^\wedge + 1/(2\tau^2)}{\theta_* + 1/(2\tau^2)}, \frac{\theta_*}{2\theta_*^\vee} \frac{\theta_*^\vee + 1/\tau^2}{\theta_* + 1/(2\tau^2)} \right\} \right]. \quad (27)$$

Some simplification arises upon noting that $\theta_*^\vee = \theta_*$, since

$$\begin{aligned} \mathrm{P}(\max(X_1^*, X_2^*) > x) &= 2\mathrm{P}(X_1^* > x) - \mathrm{P}(X_1^* > x, X_2^* > x) \\ &= 2\mathrm{P}(X_1^* > x) \{1 - \mathrm{P}(X_2^* > x | X_1^* > x)/2\} \\ &\sim 2\mathrm{P}(X_1^* > x) \{1 - \chi_{X^*}/2\}, \quad x \rightarrow \infty. \end{aligned}$$

As $\tau^2 \rightarrow 0$, i.e., as the nugget effect disappears, the endpoints converge to $\theta_*/\theta_*^\wedge = \eta_{X^*}$. As $\tau^2 \rightarrow \infty$, i.e., as the nugget effect dominates, both endpoints converge to $1/2$.

c) When $\alpha_* > \alpha_\epsilon = 2$, solving the equations in (21) provides

$$\begin{aligned} q_*(x) &= \left(\frac{2\theta_\epsilon}{\theta_*\alpha_*} x \right)^{1/(\alpha_*-1)} [1 + o(1)], \\ q_\epsilon(x) &= x - q_*(x) \sim x, \end{aligned}$$

and $\psi_+(x) \sim \theta_\epsilon x^2$. Using inequality (23),

$$-\log \mathrm{P}(X_1^* + \epsilon_1 > x, X_2^* + \epsilon_2 > x) \sim \theta_\epsilon^\wedge x^2;$$

noting that $\theta_\epsilon = 1/(2\tau^2)$ and $\theta_\epsilon^\wedge = 1/\tau^2$ leads to $\eta_X = 1/2$.

□

A.2 Hierarchical model

Under spatiotemporal setting, we observe $\{Y_t(\mathbf{s}) \mid t = 1, \dots, T, \mathbf{s} \in \mathcal{S}\}$, and the temporal dependence is ignored. Then the hierarchical model can be described as

$$\begin{aligned}
 Y_t(\mathbf{s}) \mid \mathbf{X}_t^*, \tau^2, \boldsymbol{\theta}_R, \boldsymbol{\theta}_{GPD}, p &= T^{-1}(X_t^*(\mathbf{s}) + \epsilon(\mathbf{s})), \\
 \mathbf{X}_t^* \mid R_t, \boldsymbol{\theta}_C &= R_t \cdot g(\mathbf{Z}_t), \\
 R_t \mid \boldsymbol{\theta}_R &\sim F_R, \\
 \tau^2 &\sim IG(a, b), \\
 \boldsymbol{\theta}_{GPD} &= (\sigma, \xi) \sim \text{halfCauchy}(1) \cdot U(\xi; -0.5, 0.5), \\
 \boldsymbol{\theta}_C &= (\rho, \nu) \sim \text{halfCauchy}(1) \cdot \text{halfCauchy}(1),
 \end{aligned} \tag{28}$$

where $\boldsymbol{\theta}_R$ corresponds to $\delta \in [0, 1]$ for the [Huser and Wadsworth \(2019\)](#) model and $\beta \in [0, \infty)$ for the [Huser et al. \(2017\)](#) model, and Z_t is a Gaussian process with Matérn covariance $\boldsymbol{\Sigma}_{\boldsymbol{\theta}_C}$. The priors for $\boldsymbol{\theta}_R$ are $\delta \sim U(0, 1)$ and $\beta \sim \text{halfCauchy}(1)$ respectively.

In Section 4, we have formulated the full conditional likelihood $\varphi(Y_t(\mathbf{s}) \mid \mathbf{X}_t^*, \tau^2, \boldsymbol{\theta}_R, \boldsymbol{\theta}_{GPD}, p)$ for a fixed time and location; see Equation (14). Next we are going to work out the conditional joint likelihood $\varphi(\mathbf{X}_t^* \mid R_t, \boldsymbol{\theta}_C) = \varphi(X_t^*(\mathbf{s}_1), \dots, X_t^*(\mathbf{s}_D) \mid R_t, \boldsymbol{\theta}_C)$ for a fixed time t . Keeping the notations same as the main text, we denote $g(\cdot)$ as the link function that transforms the scale of the latent Gaussian process. For the model of [Huser et al. \(2017\)](#), $g(\cdot)$ is an identity function, and thus \mathbf{X}_t^* conditioning on R_t remains to be a multivariate Gaussian random variable. For the [Huser and Wadsworth \(2019\)](#) model, the conditional density needs more careful scrutiny.

Lemma A.3. *Let $\mathbf{X}^* = (X_1^*, \dots, X_D^*) = (R \cdot g(Z_1), \dots, R \cdot g(Z_D))$, where $(Z_1, \dots, Z_D) \sim N(0, \boldsymbol{\Sigma}_{\boldsymbol{\theta}_C})$. For the [Huser and Wadsworth \(2019\)](#) model, the density conditional on R is*

$$\varphi(x_1^*, \dots, x_D^* \mid R, \boldsymbol{\theta}_C) = \frac{1}{\sqrt{|\boldsymbol{\Sigma}_{\boldsymbol{\theta}_C}|}} \exp \left\{ -\frac{1}{2} g^{-1} \left(\frac{\mathbf{x}^*}{R} \right)^T (\boldsymbol{\Sigma}_{\boldsymbol{\theta}_C}^{-1} - I_D) g^{-1} \left(\frac{\mathbf{x}^*}{R} \right) \right\} \cdot \frac{R^D}{(\prod_{i=1}^D x_i^*)^2},$$

where $g(\cdot) = 1/\{1 - \Phi(\cdot)\}$.

Proof. First note that

$$(Z_1, \dots, Z_D) = (g^{-1}\left(\frac{X_1^*}{R}\right), \dots, g^{-1}\left(\frac{X_D^*}{R}\right)) = (\Phi^{-1}\left(1 - \frac{R}{X_1^*}\right), \dots, \Phi^{-1}\left(1 - \frac{R}{X_D^*}\right)).$$

Fixing R as a constant, we can apply the chain rule to obtain

$$\begin{aligned} \left. \frac{dZ_1}{dX_1^*} \right|_{X_1^*=x} &= \sqrt{2\pi} \exp\left\{\frac{[\Phi^{-1}(1 - R/x)]^2}{2}\right\} \cdot \frac{R}{x^2} \\ &= \sqrt{2\pi} \exp\left\{\frac{g^{-2}(x/R)}{2}\right\} \cdot \frac{R}{x^2}. \end{aligned} \tag{29}$$

Then the Jacobian matrix can be written as

$$\frac{d\mathbf{Z}}{d\mathbf{X}^*} = \begin{bmatrix} \left. \frac{dZ_1}{dX_1^*} \right|_{X_1^*=x} & & \\ & \ddots & \\ & & \left. \frac{dZ_D}{dX_D^*} \right|_{X_D^*=x} \end{bmatrix},$$

and

$$\det\left(\left. \frac{d\mathbf{Z}}{d\mathbf{X}^*} \right|_{\mathbf{X}^*=\mathbf{x}^*}\right) \stackrel{(29)}{=} (\sqrt{2\pi})^D \exp\left\{\frac{1}{2} \sum_{i=1}^D g^{-2}(x_i^*/R)\right\} \cdot \frac{R^D}{(\prod_{i=1}^D x_i^*)^2}$$

Denote f_Z as the density function of $N(0, \Sigma_{\theta_C})$. Then

$$f_Z(g^{-1}(x_1^*/R), \dots, g^{-1}(x_D^*/R)) = \frac{1}{(2\pi)^{D/2} \sqrt{|\Sigma_{\theta_C}|}} \exp\left\{-\frac{1}{2} g^{-1}\left(\frac{\mathbf{x}^*}{R}\right)^T \Sigma_{\theta_C}^{-1} g^{-1}\left(\frac{\mathbf{x}^*}{R}\right)\right\}.$$

Apply the change of variables formula

$$\varphi(\mathbf{x}^* | R, \theta_C) = f_Z(g^{-1}(\mathbf{x}^*/R)) \left| \det\left(\frac{d\mathbf{Z}}{d\mathbf{X}^*}\right) \right|,$$

and we will get the desired result. \square

When updating certain variables for a random walk Metropolis-Hastings step, we calculate the corresponding full conditional likelihood for both current values and proposed

values. Since $\varphi(Y_t(\mathbf{s}) | \mathbf{X}_t^*, \tau^2, \boldsymbol{\theta}_R, \boldsymbol{\theta}_{GPD}, p)$ entails most parameters and latent variables, we will evaluate this density over and over again. Looking back at its analytic form in (14), the most compute-intensive part is to calculate the marginal transformation $T(\cdot)$ as defined in (12), which requires the computation of the marginal quantile function of the noisy process, i.e. $F_{X|\boldsymbol{\theta}_R, \tau^2}^{-1}$. Note that the marginal distribution function $F_{X|\boldsymbol{\theta}_R, \tau^2}$ can be obtained through a convolution:

$$1 - F_{X|\boldsymbol{\theta}_R, \tau^2}(x) = P(X^* + \epsilon > x) = \int_{-\infty}^{\infty} \{1 - F_{X^*|\boldsymbol{\theta}_R}(x - \epsilon)\} \cdot \phi_{\tau}(\epsilon) d\epsilon, \quad (30)$$

where ϕ_{τ} is the density of $N(0, \tau^2)$, and $F_{X^*|\boldsymbol{\theta}_R}$ is the marginal distribution function of the smooth process, whose forms are derived in Huser et al. (2017) and Huser and Wadsworth (2019) for the two models of interest in this paper. Because (30) cannot be further simplified, we compute the improper integral numerically using the QUADPACK algorithms which are implemented within the `gsl_integration` library in C++. The computations in C++ and R are interfaced using the package `Rcpp`.

To compute p th quantile of X , i.e. $F_{X|\boldsymbol{\theta}_R, \tau^2}^{-1}(p)$, which is required for likelihood function evaluations, we first evaluate the distribution function $F_{X|\boldsymbol{\theta}_R, \tau^2}$ at a fine grid of x values. We then perform cubic spline interpolation through the control points to yield a continuous quantile function estimate. Due to the smoothness of the quantile function, numerical experiments showed that this technique suffered no measurable reduction in accuracy relative to the much slower technique of computing quantiles using a numerical root finder.

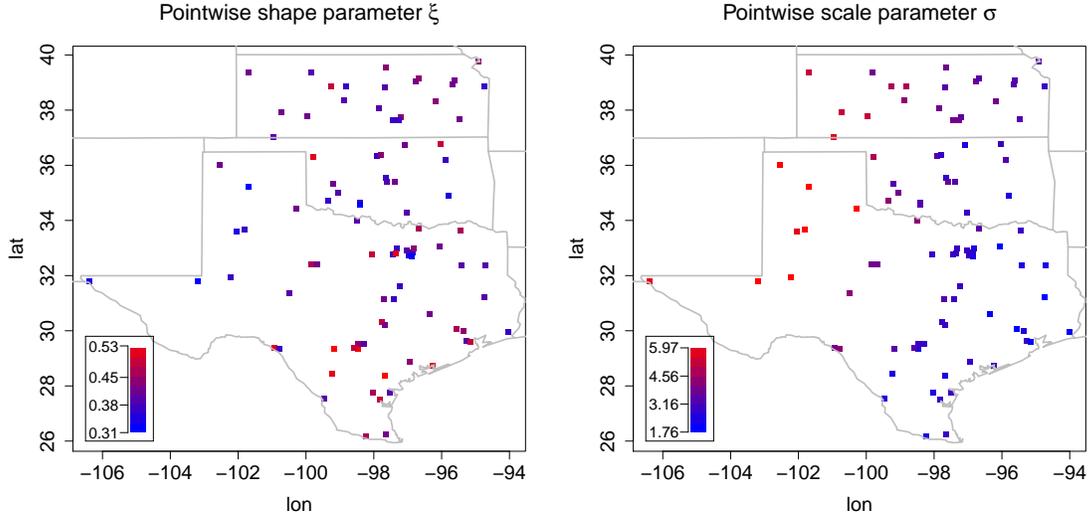


Figure 11: Pointwise generalized Pareto parameters fitted with observations from each station.

B Additional Diagnostics

B.1 Marginal parameters

To examine the trend surfaces of the marginal parameter θ_{GPD} , we fit univariate generalized Pareto distribution to the spring observations at each station over a high threshold $u_{0.8}$. The estimated parameters are plotted in Figure 11. We can see that there is no obvious spatial pattern for the shape parameter, while there is significant longitudinal effect for the scale parameter. This give grounds for applying linear trend surface to scale, and constant surface to shape (see Equation (17)).

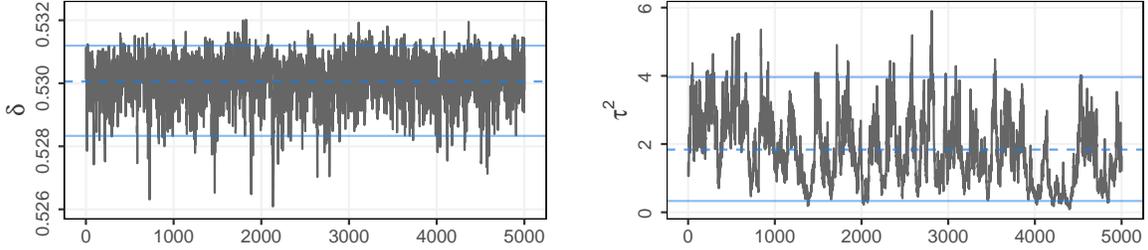


Figure 12: Trace plots for δ and τ^2 .

B.2 MCMC results

Batch means (Flegal et al., 2010) is a convenient way to compute Monte Carlo standard errors for MCMC outputs. If one divides a Markov chain $\{X_n\}$ into k batches of size b , the Monte Carlo estimate of $\mu = E(g(X))$ based on i th batch can be obtained as follows:

$$\hat{\mu}_i = \frac{\sum_{s=(i-1)b+1}^{ib} g(X_s)}{b}, i = 1, \dots, k,$$

and batch means estimate of the Monte Carlo standard error can be defined as

$$\hat{\sigma}^2 = \frac{1}{k(k-1)} \sum_{i=1}^k (\hat{\mu}_i - \hat{\mu})^2,$$

where $\hat{\mu}$ is the overall Monte Carlo estimate. See Figure 13 for batch means standard errors computed periodically for δ and ρ . Other parameters have similar results. We report the stabilized batch means standard errors in Table 2.

For the data analysis, we used 20 cores from the CyberScience Advanced Cyber Infrastructure at Penn State. On average, each iteration takes approximately 1.76 seconds (CPU time). The effective sample size (ESS) per second is also reported in Table 2, in which the marginal parameters has lower values.

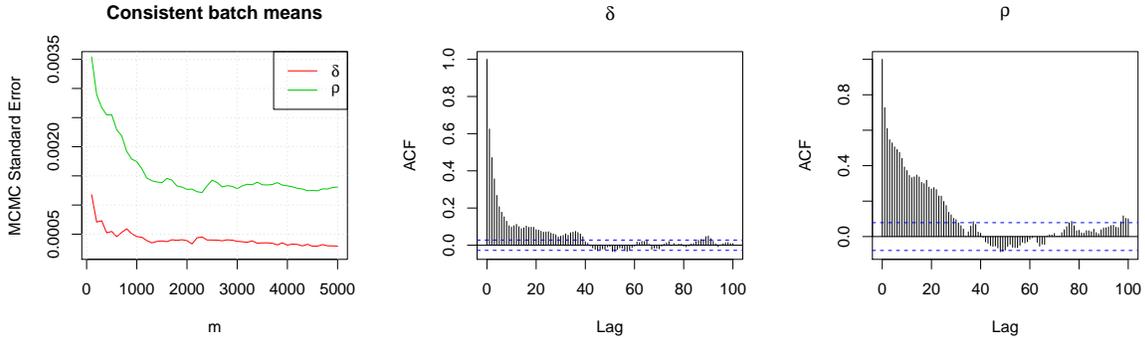


Figure 13: MCMC standard error plot (left) for ρ and δ , which is calculated using consistent batch means for MCMC chains. ACF plots for ρ and δ are displayed on the right.

	ρ	ν	τ	δ	β_0	β_1	β_2	ξ
Monte Carlo SE	0.0013	0.0071	0.3570	0.0003	0.0134	0.0002	0.0001	0.0004
ESS per second	0.223	0.468	0.098	0.190	0.027	0.027	0.031	0.019

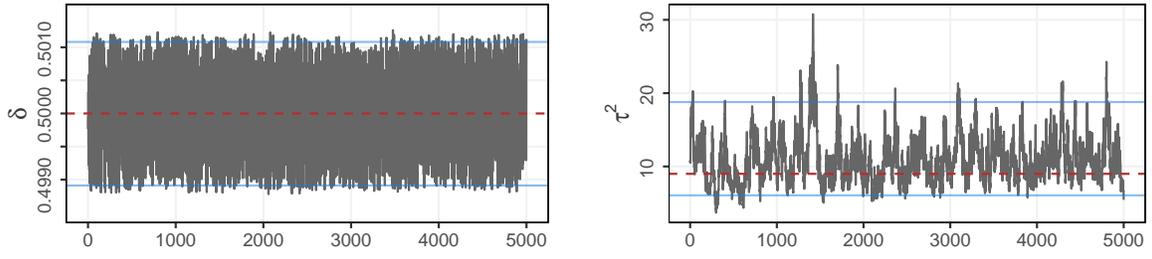
Table 2: Monte Carlo standard error for estimating the mean (computed using stabilized batch means for MCMC chains).

Supplementary Material

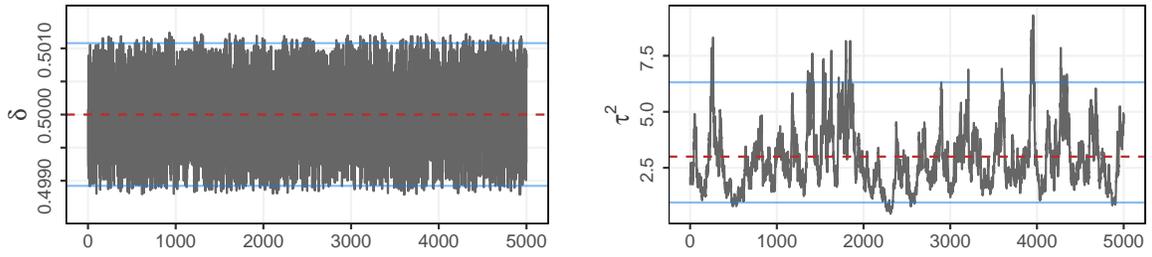
In this document, we examine the performance of the MCMC algorithm when $\tau^2 \rightarrow 0$, i.e., our proposed model (8) converges to the underlying smooth process from [Huser and Wadsworth](#). One may suspect that the Markov chains will mix very poorly when the value of τ^2 is small (although when $\tau = 0$, an alternative MCMC scheme might be possible, wherein components of \mathbf{X}^* could be updated from truncated distributions, albeit more complicated ones). We conduct another simulation study for datasets generated from $\delta = 0.5$ (transition point) and various τ^2 values (9, 3, 0.5, 0.0025). Other parameter settings remain the same as in Section 5.1. The sites and the smooth processes are simulated using the same random number generating seed for each dataset. We want to see the critical point of τ^2 for the algorithm to fail.

Figure 14 displays the results from running the MCMC algorithm, with each row showing one case. The red dashed line signifies the true parameter values, and the blue lines indicates the 95% posterior credible intervals. Each MCMC chain was run for 400,000 iterations. Thinning the results by a factor of 10, we show the last 50,000 iterations. We can see that, while the performance of δ stays stable for different true τ^2 values, the Markov chain for τ^2 converges slower when its true values becomes smaller. In the case where $\tau^2 = 0.0025$, the true value is outside of the 95% credible interval and the chain mixes very poorly. Figure 15 shows trace plots of an X^* at one specific site and time which is censored for all four simulations due to the same collection of sites and underlying smoothing processes. We can see that for larger value of τ^2 , there are more fluctuations for X^* over the threshold, which might be the reason why the MCMC algorithm works better in this case. The dissatisfying performance when τ^2 is very small might also have something to do with the particular sampler we used (random walk M-H with a symmetric proposal

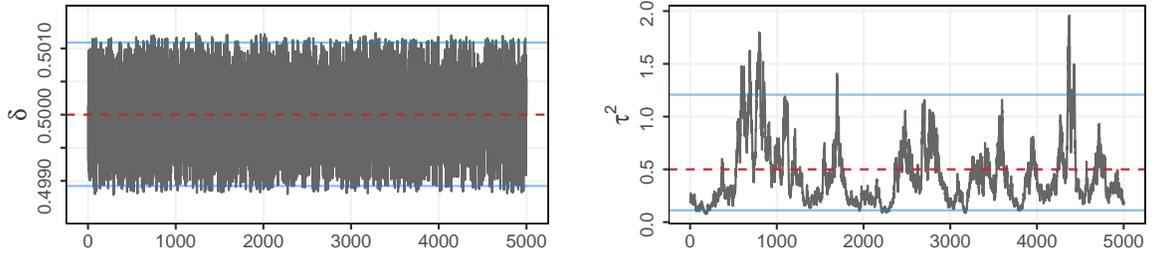
(a) Case $\tau^2 = 9$



(b) Case $\tau^2 = 3$



(c) Case $\tau^2 = 0.5$



(d) Case $\tau^2 = 0.0025$

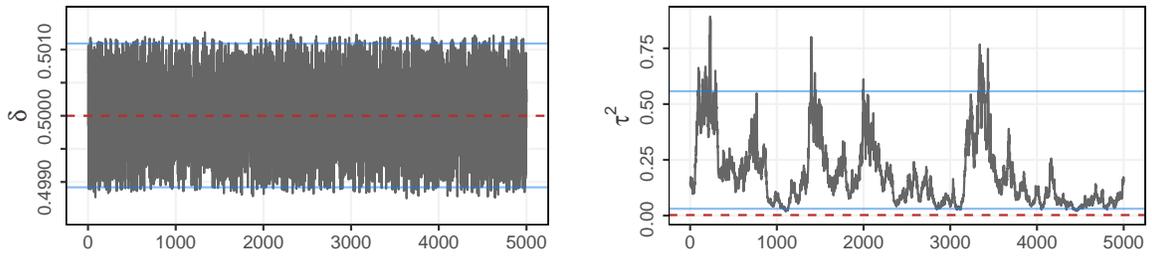


Figure 14: Comparisons of the trace plots of δ and τ^2 . See text for explanations.

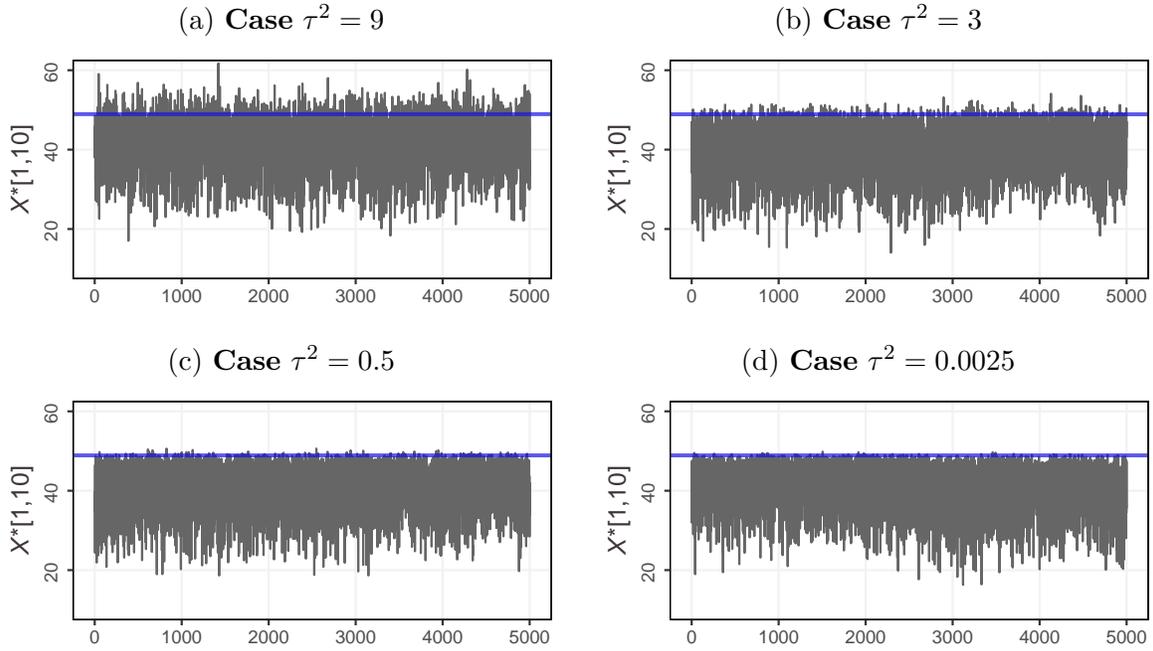


Figure 15: Comparisons of the trace plots of X^* at 1st site and 10th time replicate. The sites and the smooth processes are the same across four simulations, and true $X^*[1, 10]$ is censored.

kernel), given the parameter is so close to the boundary, rather than a problem with the overall modeling scheme.

Next we briefly discuss the narrowness of credible intervals for ξ . The sharpness of the marginal posteriors for ξ may seem unrealistically narrow, but we present some evidence that that they may be accurately reflecting the information contained in the data, as seen through the lens of the model.

For the coverage analysis presented in Section 5.2, the threshold is $u = 0.8$, which resulted in narrow CIs for ξ in the [Huser and Wadsworth \(2019\)](#) model (we did not include the coverage analysis for the marginal parameters in the main text, as this was not the

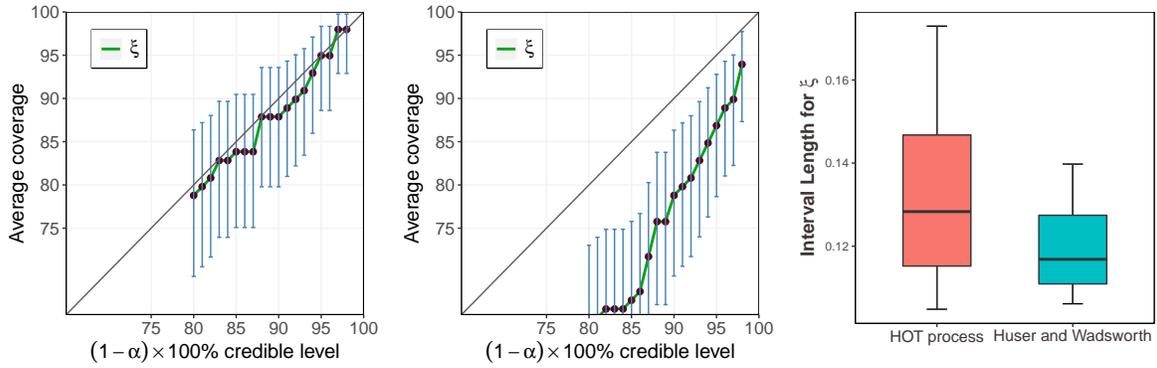


Figure 16: Empirical coverage rates of credible intervals for ξ for [Huser et al. \(2017\)](#) (a.k.a. HOT) model (left) and [Huser and Wadsworth \(2019\)](#) model (middle). The credible interval widths between the two model runs are compared on the right.

focus of the manuscript). While the posterior distributions are not perfectly calibrated, the coverage rates are still acceptable; see Figure 16. The posterior intervals do under-cover, but not by much. For example, 95% credible intervals cover the true shape parameter about 87% of the time. We used the exact same code for the simulation as for the data analysis, so there is reason to think that the posterior intervals for the data analysis are probably okay, even though they are very short. For the [Huser et al. \(2017\)](#) model, in which X^* has a lighter marginal tail, the CIs for ξ appear to be perfectly calibrated. Although the code used for this simulation is not all identical to that of the sampler for the [Huser and Wadsworth \(2019\)](#) model and the data analysis, a large percentage of the code is common across the two samplers.