

Neural Bayes estimators for censored inference with peaks-over-threshold models

Jordan Richards

JORDAN.RICHARDS@ED.AC.UK

*School of Mathematics and Maxwell Institute for Mathematical Sciences,
University of Edinburgh,
Edinburgh, United Kingdom.*

*Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division,
King Abdullah University of Science and Technology (KAUST),
Thuwal, 23955-6900, Saudi Arabia.*

Matthew Sainsbury-Dale

MATTHEW.SAINSBURYDALE@KAUST.EDU.SA

*Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division,
King Abdullah University of Science and Technology (KAUST),
Thuwal, 23955-6900, Saudi Arabia.*

*School of Mathematics and Applied Statistics,
University of Wollongong,
Wollongong, Australia.*

Andrew Zammit-Mangion

AZM@UOW.EDU.AU

*School of Mathematics and Applied Statistics,
University of Wollongong,
Wollongong, Australia.*

Raphaël Huser

RAPHAEL.HUSER@KAUST.EDU.SA

*Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division,
King Abdullah University of Science and Technology (KAUST),
Thuwal, 23955-6900, Saudi Arabia.*

Editor: Pradeep Ravikumar

Abstract

Making inference with spatial extremal dependence models can be computationally burdensome since they involve intractable and/or censored likelihoods. Building on recent advances in likelihood-free inference with neural Bayes estimators, that is, neural networks that approximate Bayes estimators, we develop highly efficient estimators for censored peaks-over-threshold models that use augmented data to encode censoring information in the neural network input. Our new method provides a paradigm shift that challenges traditional censored likelihood-based inference methods for spatial extremal dependence models. Our simulation studies highlight significant gains in both computational and statistical efficiency, relative to competing likelihood-based approaches, when applying our novel estimators to make inference with popular extremal dependence models, such as max-stable, r -Pareto, and random scale mixture process models. We also illustrate that it is possible to train a single neural Bayes estimator for a general censoring level, precluding the need to retrain the network when the censoring level is changed. We illustrate the efficacy of our estimators by making fast inference on hundreds-of-thousands of high-dimensional spatial

extremal dependence models to assess extreme particulate matter 2.5 microns or less in diameter (PM_{2.5}) concentration over the whole of Saudi Arabia.

Keywords: censored data; convolutional neural network; likelihood-free inference; neural estimation; spatial extremes; threshold exceedance

1. Introduction

Due to the flexibility and computational scalability of neural networks, recent years have seen a dramatic uptake in their application in the study of spatial extreme values, with their uses ranging from extreme quantile regression (Cannon, 2010; Vasiliades et al., 2015; Galib et al., 2022; Wilson et al., 2022; Pasche and Engelke, 2024; Richards et al., 2023a; Cisneros et al., 2024; Richards and Huser, 2022) to extremal dependence characterisation (Ahmed et al., 2022), extreme event simulation (Boulaguiem et al., 2022; Zhang et al., 2023), and parameter estimation (McDonald et al., 2022; Lenzi et al., 2023; Majumder and Reich, 2023; Walchessen et al., 2024; Majumder et al., 2024; Rai et al., 2024; Sainsbury-Dale et al., 2024a; Zammit-Mangion et al., 2025). A promising new approach to likelihood-free inference for spatial (and space-time) dependence models is neural point estimation. In neural point estimation, one trains a neural network to take in data as input and provide parameter point estimates as output; see, for example, Cremanns and Roos (2017), Zammit-Mangion and Wikle (2020), Gerber and Nychka (2021), and also Lenzi et al. (2023) and Sainsbury-Dale et al. (2024a) who consider inference with spatial extremal processes. Neural point estimators are accurate, very fast to evaluate (once trained) and they can be trained to be approximately Bayes; we refer to such estimators as neural Bayes estimators (NBEs, Sainsbury-Dale et al., 2024a). They are also “amortised” in the sense that a trained estimator may be re-used repeatedly at almost no computational cost; once trained, these estimators generate estimates via a single function evaluation and are very fast. Moreover, they provide estimates without requiring evaluation of a likelihood function. This is particularly helpful for models where the likelihood function is computationally intractable or expensive to evaluate, as is often the case with extreme value models (e.g., Padoan et al., 2010; Castruccio et al., 2016; Huser and Wadsworth, 2022).

Models for spatial extremal dependence can be broadly classified into two categories: models for maxima and models for peaks-over-threshold. Popular models for the former are max-stable processes (MSPs), which arise as the only possible non-degenerate limits of suitably normalised pointwise maxima of independent and identically distributed (i.i.d.) random fields. Likelihood-based inference for MSPs can be computationally difficult, as the number of likelihood terms for a d -dimensional observation vector grows super-exponentially with d (Padoan et al., 2010). Since their popularisation by de Haan (1984), Schlather (2002), and Padoan et al. (2010), many efforts have been made to improve the computational speed of making inference with MSPs, and to increase the feasible dimensionality for which they can be applied; see, for example, Genton et al. (2011), Huser and Davison (2013), and recent work by Hector and Reich (2024) and Huser et al. (2024). Despite these efforts, traditional techniques for making inference with MSPs remain limited to relatively low-dimensional settings or make use of approximations that lead to a substantial loss in statistical efficiency.

Popular models for spatial peaks-over-threshold data include r -Pareto processes (Ferreira and de Haan, 2014; de Fondeville and Davison, 2018), inverted MSPs (Wadsworth

and Tawn, 2012), random scale/location mixtures (Huser et al., 2017; Krupskii et al., 2018; Huser and Wadsworth, 2019), and the spatial conditional extremes model (Wadsworth and Tawn, 2022). It has been shown that likelihood estimators for these models can be highly biased if spatial extreme events include marginally non-extreme observations (Huser et al., 2016), and this sub-asymptotic misspecification problem is often solved by treating non-extreme observations as censored when making inference. Although the spatial conditional extremes model usually suffers less from these bias issues, inference with this model may still require censored likelihood estimation in certain applications (e.g., for rainfall extremes, see Richards et al., 2022). Moreover, whilst MSPs are not intended to be used with peaks-over-threshold data (being limits for maxima, not threshold exceedances), their joint tail structure is in one-to-one correspondence with r -Pareto processes. It is therefore possible to fit models designed for maxima to high peaks-over-threshold data by censoring low values as, for example, done by Huser and Davison (2014). In these peaks-over-threshold inference approaches, censoring is key to avoid non-extreme values from adversely affecting the estimation of the extremal dependence structure, but can lead to computationally burdensome inference. While computations can be simplified through the use of a nugget effect (see, e.g., Zhang et al., 2022), such approaches are still computationally demanding in high dimensions. Censoring is therefore an important technique when making inference with models for spatial extremes, but comes at a high (often prohibitive) computational cost.

A way forward to circumvent the computational issues that arise when censoring data is to use likelihood-free inference methods, where the likelihood function need not be evaluated. NBEs are an attractive choice since, as discussed above, they are likelihood-free and optimal in a Bayes sense. Lenzi et al. (2023) and Sainsbury-Dale et al. (2024a) illustrate the efficacy of their use when estimating the parameters for several spatial extremal processes, including the Schlather (Schlather, 2002) and Brown–Resnick (Brown and Resnick, 1977) classes of max-stable processes, and the spatial conditional extremes model (Wadsworth and Tawn, 2022). However, the treatment of censored data in this framework has not yet been addressed. In this paper we develop a class of NBEs that allow for, and efficiently extract information from, censored data. The resulting NBEs can be used to make inference from censored data with a variety of popular peaks-over-threshold models. We observe substantial gains in statistical and computational efficiency (when evaluating a trained NBE) over competing censored likelihood-based approaches. Our proposed framework allows us to make inference with spatial extremal dependence models of a dimension that is much higher than was previously possible when used in an “operational” manner, that is, when used to repeat an estimation task thousands of times within a short time frame (e.g., a few minutes). In our application, we fit upwards of 130 million separate extremal dependence models using five different trained NBEs; the scale of such a study is unprecedented in literature on models for spatial extreme values. Our methodology has been incorporated into the user-friendly open-source software package `NeuralEstimators` (Sainsbury-Dale, 2024), which is available in `Julia` and `R`.

The remainder of this paper is organised as follows. In Section 2, we review NBEs and outline our new methodology for handling censored inputs. In Section 3, we summarise some popular spatial extremal dependence models and discuss estimation strategies based on censored likelihood maximisation. In Section 4, we conduct simulation studies to illustrate the efficacy of our inference framework, and in Section 5, we apply our estimator to analyse

extremal dependence of high-resolution Saudi Arabian particulate matter data. The paper concludes with a discussion in Section 6. Reproducible code is available from <https://github.com/Jbrich95/CensoredNeuralEstimators>.

2. Methodology

In Section 2.1, we review NBEs. In Section 2.2, we detail our methodology for adapting NBEs to allow them to handle censored data as input.

2.1 Neural Bayes estimators

A parametric statistical model is defined as a set of probability distributions \mathcal{P} on a sample space Ω , parameterised by a parameter vector $\boldsymbol{\theta} \in \mathbb{R}^p$ such that $\mathcal{P} := \{F(\cdot | \boldsymbol{\theta}) : \boldsymbol{\theta} \in \Theta\}$, where Θ is the parameter space (see McCullagh, 2002). For ease of exposition, we take $\Omega \subseteq \mathbb{R}^d$ and assume that $F(\cdot | \boldsymbol{\theta})$ admits a density $f(\cdot | \boldsymbol{\theta})$ with respect to Lebesgue measure. Consider m mutually independent realisations from $F(\cdot | \boldsymbol{\theta}) \in \mathcal{P}$; a point estimator $\hat{\boldsymbol{\theta}}(\cdot)$ is any mapping from Ω^m to Θ . Denote the realisations as $\mathbf{Z} := (\mathbf{Z}'_1, \dots, \mathbf{Z}'_m)'$. The performance of an estimator $\hat{\boldsymbol{\theta}}(\cdot)$, for a given $\boldsymbol{\theta}$ and \mathbf{Z} , can be evaluated using a loss function $L : \Theta \times \Theta \rightarrow [0, \infty)$. The risk of the estimator at $\boldsymbol{\theta}$ is then defined as

$$R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\cdot)) := \int_{\Omega^m} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{z}))f(\mathbf{z} | \boldsymbol{\theta})d\mathbf{z}, \quad (1)$$

which represents the expected loss over all possible data realisations. The Bayes risk is a weighted average of (1) over $\boldsymbol{\theta} \in \Theta$, with respect to some prior measure $\Pi(\cdot)$, defined as

$$r_{\Pi}(\hat{\boldsymbol{\theta}}(\cdot)) := \int_{\Theta} R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\cdot))d\Pi(\boldsymbol{\theta}). \quad (2)$$

An estimator $\hat{\boldsymbol{\theta}}(\cdot)$ that minimises (2) is a *Bayes estimator* with respect to $L(\cdot, \cdot)$ and $\Pi(\cdot)$. Bayes estimators have attractive theoretical properties under mild regularity conditions, including asymptotic normality, efficiency, and consistency under the squared-error loss (see, e.g., Lehmann and Casella, 1998).

An NBE is a neural network that approximately minimises (2). Bayes estimators for independent replicates satisfy permutation invariance; that is, $\hat{\boldsymbol{\theta}}(\mathbf{Z}) = \hat{\boldsymbol{\theta}}(\mathbf{Z}^*)$ where \mathbf{Z}^* denotes any permutation of the independent replicates in \mathbf{Z} (Sainsbury-Dale et al., 2024a). To construct parsimonious NBEs that have this property, Sainsbury-Dale et al. (2024a) propose using a neural network architecture known as DeepSets (Zaheer et al., 2017). Consider the functions $\boldsymbol{\psi} : \mathbb{R}^d \mapsto \mathbb{R}^q$ and $\boldsymbol{\phi} : \mathbb{R}^q \mapsto \mathbb{R}^p$, and a permutation-invariant set function $\boldsymbol{\alpha} : (\mathbb{R}^q)^m \mapsto \mathbb{R}^q$. Here, we let the j -th component of $\boldsymbol{\alpha}$, $\boldsymbol{\alpha}_j(\cdot)$, return the element-wise average over its input set (of cardinality m) for $j = 1, \dots, q$. We represent $\boldsymbol{\phi}(\cdot)$ and $\boldsymbol{\psi}(\cdot)$ as neural networks, and collect in $\boldsymbol{\gamma} := (\boldsymbol{\gamma}'_{\boldsymbol{\phi}}, \boldsymbol{\gamma}'_{\boldsymbol{\psi}})'$ their trainable parameters (the so-called “weights” and “biases”). Then, a permutation invariant NBE is of the form

$$\hat{\boldsymbol{\theta}}(\mathbf{Z}; \boldsymbol{\gamma}) = \boldsymbol{\phi}(\mathbf{T}(\mathbf{Z}; \boldsymbol{\gamma}_{\boldsymbol{\psi}}); \boldsymbol{\gamma}_{\boldsymbol{\phi}}), \quad \text{with} \quad \mathbf{T}(\mathbf{Z}; \boldsymbol{\gamma}_{\boldsymbol{\psi}}) = \boldsymbol{\alpha}(\{\boldsymbol{\psi}(\mathbf{Z}_t; \boldsymbol{\gamma}_{\boldsymbol{\psi}}) : t = 1, \dots, m\}). \quad (3)$$

Permutation invariance allows the NBE to handle independent replicates parsimoniously, which is key for inference with replicated data and spatial extremal dependence models.

As our interest is in estimating dependence only, we hereon assume that $\boldsymbol{\theta}$ determines only the dependence characteristics of \mathbf{Z} , not its marginal behaviour. Following Sainsbury-Dale et al. (2024a), we use a densely-connected neural network for $\phi(\cdot)$; the architecture for $\psi(\cdot)$ is application-dependent (see the example in Section 4.1).

An NBE is trained by finding the neural-network parameters $\boldsymbol{\gamma}^*$ that minimise (2) in the estimator space spanned by $\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma})$. As (2) cannot be directly evaluated, it is approximated using Monte Carlo methods. Given a set of K parameter vectors, $\{\boldsymbol{\theta}^{(k)} : k = 1, \dots, K\}$, drawn from the prior $\Pi(\cdot)$, we simulate, for each k , m mutually independent random samples $\mathbf{Z}^{(k)}$ from $F(\cdot | \boldsymbol{\theta}^{(k)})$. Then we approximate (2) using the estimator parameterised by $\boldsymbol{\gamma}$ as

$$\hat{r}_{\Pi}(\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma})) = \frac{1}{K} \sum_{k=1}^K L(\boldsymbol{\theta}^{(k)}, \hat{\boldsymbol{\theta}}(\mathbf{Z}^{(k)}; \boldsymbol{\gamma})) \approx r_{\Pi}(\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma})), \quad (4)$$

and we obtain estimates $\boldsymbol{\gamma}^* = \operatorname{argmin}_{\boldsymbol{\gamma}} \hat{r}_{\Pi}(\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma}))$ using back propagation and stochastic gradient descent.

When training an NBE for a large number of replicates m , computational speed-ups can be achieved with the use of pre-training (Goodfellow et al., 2016, Ch. 8). Let $\boldsymbol{\gamma}_m^*$ denote the optimal NBE weights for datasets containing m replicates. We employ pre-training by first finding $\boldsymbol{\gamma}_{\tilde{m}}^*$ for some smaller sample size $\tilde{m} < m$, and then using these as initial estimates when finding $\boldsymbol{\gamma}_m^*$. In practice, we run a recursive pre-training scheme for a sequence of increasing sample sizes, $\tilde{\mathbf{m}} := (\tilde{m}_1, \tilde{m}_2, \dots, m)'$; for details, see Section 4.1. Throughout we assume that m equals the number of observations from which inference needs to be made.

2.2 Neural Bayes estimators for censored data

In Section 2.2.1, we review censored likelihood-based inference and detail a way of encoding relevant censoring information into the input of an NBE. In Section 2.2.2, we describe how to build NBEs that can be used with a range of censoring levels.

2.2.1 ENCODING CENSORED VALUES

Suppose that we have (time-)replicated observations $\mathbf{Z}_t := (Z_{t,1}, \dots, Z_{t,d})'$ for $t = 1, \dots, m$, sampled independently from the joint distribution $F(\cdot | \boldsymbol{\theta})$, and that we want to estimate the unknown parameter vector $\boldsymbol{\theta}$ by prioritising calibration in the joint upper tail. In order to prevent low values from affecting the estimation of extremal dependence, the standard approach in statistics of extremes is to treat all observations falling below some high marginal quantile as *censored*. That is, while we know which observations are treated as censored or not, the exact values of the censored observations are assumed to be unknown. Mathematically, we consider a vector of censoring thresholds $\mathbf{c} := (c_1, \dots, c_d)'$, where here $c_j = F_j^{-1}(\tau | \boldsymbol{\theta})$ is set to be the marginal τ -quantile of $Z_{t,j}$, $j = 1, \dots, d$, for some probability $\tau \in [0, 1)$ close to one. Let $\mathbf{I}_t := (I_{t,j} : j = 1, \dots, d)'$, with $I_{t,j} = \mathbb{1}\{Z_{t,j} \leq c_j\}$, denote the vector of indicator variables with j -th component equal to one if the corresponding element of \mathbf{Z}_t is censored and zero otherwise, and let $d_{t-} := \sum_{j=1}^d I_{t,j} \leq d$ denote the number of censored components. We rearrange the observation vector as $\mathbf{Z}_t = (\mathbf{Z}'_{t-}, \mathbf{Z}'_{t+})'$, where $\mathbf{Z}_{t-} := (Z_{t,j} : Z_{t,j} \leq c_j, j = 1, \dots, d)'$ is the subvector of censored observations, and $\mathbf{Z}_{t+} := (Z_{t,j} : Z_{t,j} > c_j, j = 1, \dots, d)'$ is the subvector of uncensored observations. We do the same for the vector of censoring thresholds $\mathbf{c}_t := (\mathbf{c}'_{t-}, \mathbf{c}'_{t+})'$, where

$\mathbf{c}_{t-} := (c_j : Z_{t,j} \leq c_j, j = 1, \dots, d)'$ and $\mathbf{c}_{t+} := (c_j : Z_{t,j} > c_j, j = 1, \dots, d)'$. Note that \mathbf{c}_t is a permutation of \mathbf{c} , with the ordering time-dependent. Traditional censoring in statistics of extremes assumes that $\mathbf{Z}_t^{\text{cens}} := (\mathbf{I}'_t, \mathbf{Z}'_{t+})'$ is the only information available for inference at time t . The corresponding likelihood contribution is given in the following proposition.

Proposition 1 *Let $\mathbf{Z}_t = (Z_{t,1}, \dots, Z_{t,d})'$, $d \in \mathbb{N}, t = 1, \dots, m$, be random vectors defined by a statistical model parameterised by $\boldsymbol{\theta}$ with joint distribution $F(\cdot \mid \boldsymbol{\theta})$ and density $f(\cdot \mid \boldsymbol{\theta})$. Let $\mathbf{Z}_{t+} := (Z_{t,j} : Z_{t,j} > c_j, j = 1, \dots, d)'$ denote the subvector of uncensored values of \mathbf{Z}_t , where $\mathbf{c} := (c_1, \dots, c_d)'$ is a vector of censoring levels with time-dependent subvector $\mathbf{c}_{t-} = (c_j : Z_{t,j} \leq c_j, j = 1, \dots, d)'$. Given indicators $\mathbf{I}_t := (\mathbb{1}\{Z_{t,j} \leq c_j\} : j = 1, \dots, d)'$, the likelihood contribution of $\mathbf{Z}_t^{\text{cens}} := (\mathbf{I}'_t, \mathbf{Z}'_{t+})'$ is*

$$\int_{-\infty}^{\mathbf{c}_{t-}} f(\mathbf{z}_t \mid \boldsymbol{\theta}) d\mathbf{z}_{t-} = \frac{\partial^{|\mathcal{J}_t|}}{\prod_{j \in \mathcal{J}_t} \partial z_j} F(\mathbf{z} \mid \boldsymbol{\theta}) \Big|_{\mathbf{z}=\mathbf{b}_t}, \quad t = 1, \dots, m, \quad (5)$$

where $\mathbf{z}_t = (z_{t,1}, \dots, z_{t,d})'$ is the realisation of \mathbf{Z}_t , $\mathcal{J}_t := \{j : z_{t,j} > c_j\}$ denotes the indices of the uncensored components of \mathbf{z}_t with subvector $\mathbf{z}_{t-} := (z_{t,j} : z_{t,j} \leq c_j, j = 1, \dots, d)'$, and $\mathbf{b}_t := (\max\{z_{t,j}, c_j\} : j = 1, \dots, d)'$.

For completeness, we provide a formal proof of (5) in Appendix A. Note that (5) is a d_{t-} -variate integral of the joint density; if d_{t-} is large, (5) quickly becomes computationally intractable and this often prevents censored likelihood inference in moderate-to-large dimensions d . While a different censoring framework for extremal processes has been used by Wadsworth and Tawn (2012), Opitz (2016), and Wadsworth et al. (2017), the formulation in (5) is usually preferred since it is more robust against model misspecification in the bulk, and thus less subject to sub-asymptotic estimation bias (Huser et al., 2016). In this paper, we seek to develop novel NBE methodology that is able to handle censored inputs in a way similar to the censored likelihood contribution in (5).

A possible approach could be to discard all censored observations, \mathbf{Z}_{t-} , and train our NBE based on the uncensored data only, namely \mathbf{Z}_{t+} . However, this approach inevitably leads to a loss of information, particularly for large τ close to one, given that the indicator vector \mathbf{I}_t is then ignored: it is impossible to recover the censored likelihood contribution (5) from the marginal density of \mathbf{Z}_{t+} alone. We note that the information carried by \mathbf{I}_t may not be negligible (especially for large d); in the spatial setting, for example, larger (respectively smaller) clusters of zeros or ones are indicative of stronger (respectively weaker) spatial dependence. Moreover, this ‘‘data removal’’ approach requires a specialised neural network architecture that can efficiently handle partially-observed vectors with varying patterns of missingness, given that the censored observations may change with every time replicate. Such an architecture can be difficult to design and implement.

Here, we propose a different approach that circumvents the aforementioned issues. We first standardise the data to an appropriate scale by marginally transforming the data $\mathbf{Z}_t = (\mathbf{Z}'_{t-}, \mathbf{Z}'_{t+})'$ to $\mathbf{Z}_t^* := ((\mathbf{Z}_{t-}^*)', (\mathbf{Z}_{t+}^*)')'$ with marginal distribution $H(\cdot)$. We then replace the censored component \mathbf{Z}_{t-}^* with a constant $\zeta_\tau < H^{-1}(\tau)$ (which may vary with the censoring probability τ). In other words, we construct the vector $\tilde{\mathbf{Z}}_t^* := (\boldsymbol{\zeta}'_\tau, (\mathbf{Z}_{t+}^*)')'$, where $\boldsymbol{\zeta}_\tau := (\zeta_\tau, \dots, \zeta_\tau)'$ denotes a vector of appropriate length, so that $\tilde{\mathbf{Z}}_t^*$ is of dimension d and of the same format as the original vector \mathbf{Z}_t . Motivated by Proposition 1, we then propose to

train the NBE using an augmented dataset $\mathbf{A}_t := (\mathbf{I}'_t, (\tilde{\mathbf{Z}}_t^*)')'$, $t = 1, \dots, m$. The new input is passed to the NBE by replacing \mathbf{Z}_t in (3) with \mathbf{A}_t arranged in a multi-dimensional array with two outer dimensions corresponding to \mathbf{I}_t and $\tilde{\mathbf{Z}}_t^*$, respectively. While the choice of marginal distribution $H(\cdot)$ and constant ζ_τ are, in principle, arbitrary, a judicious choice can facilitate the training of the NBE. In particular, we find that the estimator is more easily trained when $H(\cdot)$ has a finite lower bound $z_H < \infty$ and when $\zeta_\tau \leq z_H$. For example, one could choose $H(\cdot)$ to be the unit exponential distribution function and $\zeta_\tau \leq 0$; in this way, ζ_τ is outside the possible set of uncensored values, which helps the neural network distinguish between the two data types, and thereby learn to treat censored data differently to fully-observed data. We further note that while the augmented data set \mathbf{A}_t includes additional (artificial) data, when compared to $\mathbf{Z}_t^{\text{cens}}$ used to define the censored likelihood contribution (5), these artificial values are fixed and independent of $\boldsymbol{\theta} \in \Theta$. During training, the NBE therefore learns to ignore these censored values; in our simulation study in Section 4, we verify that our censored NBE has indeed a similar performance to censored likelihood estimators in cases where the latter can be computed. A similar point estimation approach (that uses both data and an indicator variable) was proposed by Pagendam et al. (2023) and Wang et al. (2024) in the context of missing (not censored) data. There is a subtle but fundamental difference between censoring and missingness; while the censoring mechanism is completely determined by the data and the censoring level τ (and is thus itself dependent on $\boldsymbol{\theta}$), data missingness mechanisms are often independent of the data generating process.

2.2.2 VARYING CENSORING LEVEL

To accommodate varying censoring levels, we feed τ as an input to the outer neural network of the NBE; that is, with a slight abuse of notation, we replace (3) with

$$\hat{\boldsymbol{\theta}}(\mathbf{A}, \tau; \boldsymbol{\gamma}) = \boldsymbol{\phi}(\mathbf{T}(\mathbf{A}; \boldsymbol{\gamma}_\psi), \tau; \boldsymbol{\gamma}_\phi), \quad (6)$$

where $\mathbf{T}(\mathbf{A}; \boldsymbol{\gamma}_\psi)$ is as defined in (3) but with \mathbf{Z} replaced by the augmented data set $\mathbf{A} := (\mathbf{A}'_1, \dots, \mathbf{A}'_m)'$ and with \mathbf{Z}_t replaced with \mathbf{A}_t . The NBE $\hat{\boldsymbol{\theta}}(\mathbf{A}, \tau; \boldsymbol{\gamma})$ is approximately Bayes for any $\mathbf{A} \in \Omega_{\mathbf{A}, \tau}$ and $\tau \in \mathcal{T} \subseteq [0, 1)$ if it minimises, with respect to $\boldsymbol{\gamma}$, the Bayes risk,

$$r_{\Pi}^{\text{cens}}(\hat{\boldsymbol{\theta}}(\cdot, \cdot; \boldsymbol{\gamma})) := \int_{\Theta} \int_{\mathcal{T}} \int_{\Omega_{\mathbf{A}, \tau}} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{a}, \tau; \boldsymbol{\gamma})) f(\mathbf{a} \mid \tau, \boldsymbol{\theta}) p(\tau) d\mathbf{a} d\tau d\Pi(\boldsymbol{\theta}), \quad (7)$$

where $p(\tau)$ is a prior density for τ , and $\Omega_{\mathbf{A}, \tau}$ is the sample space of the augmented dataset \mathbf{A} , which is τ -dependent. This adaptivity to τ may be useful where, for example, data-collection devices only store values that exceed a certain threshold, but where the threshold is variable. In Appendix B, we show that this augmented Bayes estimator is invariant to the choice of $p(\tau)$, provided that \mathcal{T} is a subset of its support and that τ is independent of $\boldsymbol{\theta}$ (as is always the case in applications of the peaks-over-threshold censoring framework).

To train the NBE, we use a Monte Carlo approximation to (7) similar to (4), obtained by sampling $\boldsymbol{\theta}$ and τ from their prior distributions, and simulating \mathbf{A} by conditioning on these samples.

3. Spatial extremal processes

Here we concisely describe the spatial extremal processes that we consider in our simulations and application; we provide details of the processes and their inference using censored likelihoods. Section 3.1 details max-stable, r -Pareto, and inverted max-stable processes. In Appendix C, we give further details on r -Pareto and Matérn Gaussian processes in Sections C.3 and C.2, respectively, with the latter used to build the random scale mixture processes described in Section 3.2. For comprehensive reviews of the considered processes, see Davison et al. (2012), Davison et al. (2019) and Huser and Wadsworth (2022). For the software used for inference and simulation, see Section C.4 of Appendix C.

3.1 Max-stable, r -Pareto, and inverted max-stable processes

Suppose that $\{X_t(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$ are i.i.d. stochastic processes with continuous sample paths on $\mathcal{S} \subset \mathbb{R}^2$ and there exists sequences of functions $a_m(\mathbf{s}) > 0$ and $b_m(\mathbf{s})$, such that $\{a_m(\mathbf{s})\}^{-1} \left[\max_{t=1, \dots, m} \{X_t(\mathbf{s})\} - b_m(\mathbf{s}) \right]$, converges weakly, as $m \rightarrow \infty$, to a process $\{Z(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$ with non-degenerate margins. Then $Z(\mathbf{s})$ is a max-stable process (MSP) with generalised extreme value margins (de Haan, 1984); see, for example, Davison and Huser (2015) and Section C.1 of Appendix C. Following marginal standardisation to unit Fréchet, denoted by $\tilde{Z}(\mathbf{s})$, and under mild regularity conditions (de Haan, 1984; de Haan and Ferreira, 2006), a MSP can be constructed through the spectral representation

$$\tilde{Z}(\mathbf{s}) = \sup_{k \geq 1} R_k W_k(\mathbf{s}), \tag{8}$$

where $\{R_k\}_{k \in \mathbb{N}}$ are points of a Poisson process on $(0, \infty)$ with intensity $r^{-2} dr$ and $\{W_k(\mathbf{s})\}_{k \in \mathbb{N}}$ are i.i.d. copies of a non-negative stochastic process $W(\cdot)$ satisfying $\mathbb{E}[W(\mathbf{s})] = 1$ for all $\mathbf{s} \in \mathcal{S}$.

Specification of the process $W(\mathbf{s})$ in (8) leads to a limited selection of parametric models for MSPs (see e.g., Smith, 1990; Schlather, 2002; Kabluchko et al., 2009; Opitz, 2013). Whilst Sainsbury-Dale et al. (2024a) consider inference for the Schlather process (Schlather, 2002), we instead focus on the more flexible Brown–Resnick model (Brown and Resnick, 1977; Kabluchko et al., 2009; Engelke et al., 2011); Lenzi et al. (2023) have previously considered neural estimators for this process, albeit in the uncensored setting. The Brown–Resnick model is defined by specifying $W(\mathbf{s}) = \exp\{\varepsilon(\mathbf{s}) - \sigma^2(\mathbf{s})\}$ for $\varepsilon(\mathbf{s})$ a zero-mean Gaussian process (GP) with variance $\sigma^2(\mathbf{s}) > 0$ for all $\mathbf{s} \in \mathcal{S}$, and which satisfies $\varepsilon(\mathbf{0}) = 0$ almost surely. Extremal dependence in \tilde{Z} is controlled by the semivariogram of the underlying process $\varepsilon(\cdot)$, denoted $\gamma(\mathbf{s}_i, \mathbf{s}_j)$. A typical model for γ is the stationary and isotropic semivariogram $\gamma(\mathbf{s}_i, \mathbf{s}_j) = (\|\mathbf{s}_i - \mathbf{s}_j\|/\lambda)^\kappa$, with range and smoothness parameters $\lambda > 0$ and $\kappa \in (0, 2]$, respectively; in this case, $\sigma^2(\mathbf{s}) = 2\gamma(\mathbf{s}, \mathbf{0}) = 2(\|\mathbf{s}\|/\lambda)^\kappa$. We adopt this model and hence inference for this process requires estimation of the unknown parameter vector $\boldsymbol{\theta} = (\lambda, \kappa)'$. Note that whilst we focus on a single class of MSP, it is trivial to extend our censored NBE framework to other parametric max-stable processes, provided that simulation is feasible.

As for r -Pareto processes, they arise as the only possible limits of properly renormalised threshold exceedances defined in terms of a homogeneous risk functional $r(\cdot)$ that determines

the ‘magnitude’ of the process. Their characterization is intrinsically linked to max-stable processes; in particular, r -Pareto processes can be sampled by generating the random functions $R_k W_k(\cdot)$ in (8) and retaining only those such that $r\{R_k W_k(\cdot)\} > 1$. For more details on risk functionals and r -Pareto processes, see Section C.3 of Appendix C.

We quantify the strength of asymptotic dependence through the measure $\chi(\cdot, \cdot)$, the upper-tail index (Joe, 1997). Assuming $Z(\mathbf{s}_i) \sim F_i$, $Z(\mathbf{s}_j) \sim F_j$ with continuous F_i, F_j , then $\chi(\mathbf{s}_i, \mathbf{s}_j) = \lim_{q \rightarrow 1} \Pr[F_i\{Z(\mathbf{s}_i)\} > q \mid F_j\{Z(\mathbf{s}_j)\} > q]$, where the process $Z(\cdot)$ exhibits asymptotic independence between sites \mathbf{s}_i and \mathbf{s}_j if $\chi(\mathbf{s}_i, \mathbf{s}_j) = 0$, and asymptotic dependence otherwise. Max-stable and r -Pareto processes are either inherently asymptotically dependent (i.e., $\chi(\mathbf{s}_i, \mathbf{s}_j) > 0$ for all $\mathbf{s}_i, \mathbf{s}_j \in \mathcal{S}$) or exhibit exact independence everywhere; this limitation may not be appropriate for some environmental applications where extreme events become increasingly localised with severity (Huser and Wadsworth, 2019). Wadsworth and Tawn (2012) propose an alternative collection of flexible sub-asymptotic models for spatial extremal processes, which exhibit asymptotic independence (but not exact independence). One such model is the inverted max-stable process (IMSP), which is obtained by applying a monotonically-decreasing marginal transformation to a MSP. An IMSP $Y(\mathbf{s})$ with standard exponential margins can be constructed from a MSP $\tilde{Z}(\cdot)$ by setting $Y(\mathbf{s}) = 1/\tilde{Z}(\mathbf{s})$ for all $\mathbf{s} \in \mathcal{S}$: if $\tilde{Z}(\cdot)$ is a Brown–Resnick process, then we term $Y(\cdot)$ an “inverted Brown–Resnick process”. Inference for an inverted Brown–Resnick process requires estimation of the same parameter vector, $\boldsymbol{\theta} = (\lambda, \kappa)'$, as its max-stable counterpart.

Inference for MSPs is computationally problematic as the full d -dimensional density is generally intractable even for moderate d (i.e., for $d > 12$, see Castruccio et al., 2016); note that, by construction, IMSPs inherit the intractable likelihood and computational issues of MSPs (Huser and Wadsworth, 2022). Computational speed-ups can be exploited in certain cases, see, for example, Stephenson and Tawn (2005), Davison and Gholamrezaee (2012), Wadsworth and Tawn (2014) and Dombry et al. (2017), but their use can lead to biased parameter estimation (Wadsworth, 2015; Huser et al., 2016) and still incur high computational demand in relatively low dimension, that is, $d < 20$ (Huser et al., 2019). Intractability of the full likelihood has motivated the use of composite pairwise likelihood methods, developed by Padoan et al. (2010), for inference with MSPs; these approaches falsely assume mutual independence between all observation pairs and construct the (surrogate) pairwise likelihood as a (potentially weighted) product over all bivariate densities. Under mild regularity conditions, pairwise likelihood inference leads to valid inference, whereby the corresponding estimators exhibit consistency and asymptotic normality (Padoan et al., 2010; Varin et al., 2011), but they suffer from a reduction in efficiency relative to methods using the full likelihood (Huser and Davison, 2013; Castruccio et al., 2016). Computational expense for the pairwise likelihood grows with rate $\mathcal{O}(d^2)$, so a common approach is to use only observation pairs with sampling locations with pairwise distance subceeding a fixed cut-off distance, h_{max} say; this has been shown to improve both statistical, and computational, efficiency of the pairwise likelihood estimator (Bevilacqua et al., 2012; Sang and Genton, 2014). Due to their popularity, we adopt censored pairwise likelihood estimators as the competing likelihood approach for MSPs, IMSPs, and GPs, in the simulation study detailed in Section 4.

The aforementioned issues with likelihood-based inference with MSPs and IMSPs are compound with the requirement of left-censoring for unbiased inference (see Section 2.2.1);

for expressions of the censored pairwise likelihood for the Brown–Resnick process and associated IMSP, see Huser and Davison (2013) and Thibaud et al. (2013), respectively.

3.2 Random scale mixtures

Max-stable and r -Pareto processes can be constructed from scale mixture processes of the form $Z(\mathbf{s}) = RW(\mathbf{s})$, for an asymptotically independent stochastic process $W(\mathbf{s})$ with light or bounded upper-tails and an independent Pareto-tailed random variable R ; such mixture-based processes exhibit asymptotic dependence due to the heavy-tailed random variable R (Huser and Wadsworth, 2022). Modification of R and/or $W(\cdot)$ can lead to processes with more flexible extremal dependence than the inherent asymptotic dependence of MSPs and r -Pareto processes. Popular examples include Gaussian scale (Huser et al., 2017) and location (Krupskii et al., 2018; Castro-Camilo and Huser, 2020) mixtures, but we constrain our focus to the random scale mixture process proposed by Huser and Wadsworth (2019), which provides a model that parsimoniously bridges the gap between asymptotic dependence and asymptotic independence in the interior of the parameter space.

Following Huser and Wadsworth (2019), we define a process $\{Z(\mathbf{s})\} = R^\delta \{W(\mathbf{s})^{1-\delta}\}$ for $\delta \in [0, 1]$, and where $R \geq 1$ is a unit Pareto random variable and $\{W(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$ is a stochastic process with unit Pareto margins that exhibits asymptotic independence, for example, a marginally transformed GP or IMSP. The parameter δ controls the extremal dependence class of the process Z ; if $\delta < 1/2$, then Z inherits the extremal dependence class of W and, hence, exhibits asymptotic independence; if $\delta \geq 1/2$, R^δ is heavier-tailed than $W(\mathbf{s})^{1-\delta}$, and so Z is asymptotically dependent. As this model interpolates between an asymptotic dependence, and asymptotic independence, regime, it provides a flexible extremal dependence model that does not require choosing the extremal dependence class of data a priori. We follow Huser and Wadsworth (2019) and take $W(\cdot)$ to be a (marginally transformed) GP; we adopt the Matérn class with range $\lambda > 0$ and smoothness $\kappa > 0$, described in Section C.2 of Appendix C. Hence, inference for Z follows by estimating $\boldsymbol{\theta} = (\lambda, \kappa, \delta)'$. Note that whilst R and $W(\cdot)$ have unit Pareto margins, this does not hold for $Z(\cdot)$; hereafter we assume that the margins of $Z(\cdot)$ have been standardised to unit Pareto.

Huser and Wadsworth (2019, 2022) use the full d -dimensional likelihood for inference, rather than the pairwise alternative described in Section 3.1; this is approximated via a quasi-Monte Carlo algorithm (described in Section C.3 of Appendix C). We note that Zhang et al. (2022) adapt the model of Huser and Wadsworth (2019) to a hierarchical Bayesian setting, where the full d -dimensional censored likelihood becomes more easily amenable to high dimensions by adding a tiny noise term; however, as we are considering the Huser and Wadsworth (2019) model in the original frequentist setting, we use the quasi-Monte Carlo estimator as the competing likelihood approach for this model (see Section 4). Note that the full likelihood does not scale well with d , and so previous applications of this model have been limited to moderate dimensions (mostly $d \leq 30$, and $d \leq 200$ for Zhang et al., 2022); in Section 5, we illustrate that much higher dimensional inference ($d > 1000$) is possible using our fast NBE-based approach.

4. Simulation studies

We now conduct several simulation studies to highlight the efficacy of our novel NBE against competing censored (pseudo) likelihood approaches. In Section 4.1, we outline the general setting for the simulation studies. In Section 4.2, we consider studies where the censoring threshold τ is fixed, while in Section 4.3 we consider studies where τ varies during training.

4.1 General setting

We illustrate our methodology by application to a variety of spatial extremal processes $\{Z(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$, described in Section 3, indexed by the spatial location $\mathbf{s} \in \mathcal{S} \subset \mathbb{R}^2$. Inference on $Z(\cdot)$ is made using replicated observations $\mathbf{Z}_t = (Z_{t,1}, \dots, Z_{t,d})'$ for $t = 1, \dots, m$, where each \mathbf{Z}_t is an independent realisation of $Z(\cdot)$ at gridded sampling locations $\{\mathbf{s}_1, \dots, \mathbf{s}_d\} \subset \mathcal{S}$ at time t . The domain of interest, \mathcal{S} , is taken to be a 16×16 regular grid on $[0, 16] \times [0, 16]$, which gives $|\mathcal{S}| = 16^2 = 256$ sampling locations.

Our NBEs use the architecture given in Table 1, with $\psi(\cdot)$ and $\phi(\cdot)$ in (3) constructed using the first four and last two rows, respectively, of Table 1. Following Sainsbury-Dale et al. (2024a), we use a convolutional neural network (CNN) to model $\psi(\cdot)$; for details, see Gu et al. (2018). Note that CNNs are only applicable if the inputs are completely observed over a regular grid; hereon we assume this to be the case, but discuss alternative approaches for non-gridded data in Section 6. We use $m = 200$ independent replicates throughout, and the pre-training scheme described in Section 2.1, with $\tilde{\mathbf{m}} = (10, 50, 100, 200)'$; $L(\cdot, \cdot)$ in (2) is taken to be the absolute-error loss. NBEs are trained and validated with parameter sets of length K and $K/5$, respectively, where K differs between studies. For each sampled parameter vector $\boldsymbol{\theta}^{(k)}$, $k = 1, \dots, K$, we generate a random censoring level τ together with m mutually independent random samples of the augmented (censored) data, i.e., pairs $\{\mathbf{A}^{(k)}, \tau^{(k)}\}$; see Section 2.2.2. For the general case where the NBE is dependent on the censoring level (as in (6)), the empirical marginal risk for the parameter corresponding to the i -th component $\theta_i^{(k)}$ of $\boldsymbol{\theta}^{(k)}$ is $K^{-1} \sum_{k=1}^K L(\theta_i^{(k)}, [\hat{\boldsymbol{\theta}}(\mathbf{A}^{(k)}, \tau^{(k)}; \boldsymbol{\gamma})]_i)$; we also consider the fixed- τ case, which is a special case where the prior density $p(\tau)$ is degenerate at a fixed value of τ , i.e., all sampled values of τ are the same. All NBEs and likelihood-based estimators are tested by evaluating this empirical marginal risk (under both the absolute and squared-error loss) on a test set of 1000 parameter vectors; as the absolute and squared-error losses are additively separable, summing the empirical marginal risks for $i = 1, \dots, p$ gives a Monte Carlo approximation of the risk in (7). We assume that the parameters are a priori independent with marginal prior distributions $\lambda \sim \text{Unif}(2, 10)$, $\kappa \sim \text{Unif}(0.5, 2)$ and $\delta \sim \text{Unif}(0, 1)$ across all models.

We conduct our experiments using the package `NeuralEstimators` (Sainsbury-Dale, 2024), which is available in Julia and R. NBEs are trained on GPUs randomly selected from within KAUST’s IbeX cluster, whilst likelihood-based estimators are evaluated on CPUs from KAUST’s Shaheen II supercomputer, see <https://docs.hpc.kaust.edu.sa/systems/ibex> and <https://www.hpc.kaust.edu.sa/content/shaheen-ii-0> for details (last accessed 19/12/2024).

Table 1: Neural network architecture used in Section 4. Each convolution filter uses zero padding and unit stride. The function $\text{vec}(\cdot)$ refers to a flattening of an array to a vector. All layers, except the final layer, used rectified linear unit (ReLU) activation functions; the final layer uses the identity function. Note that p is the number of model parameters.

layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[16, 16, 2]	[7, 7, 64]	10×10	$6400c + 64$
2D convolution	[7, 7, 64]	[3, 3, 128]	5×5	204,928
2D convolution	[3, 3, 128]	[1, 1, 256]	3×3	295,168
$\text{vec}(\cdot)$	[1, 1, 256]	[256]		0
dense	[256]	[500]		128,500
dense	[500]	[p]		$501p$
total trainable parameters:				$641,460 + 501p$

4.2 Fixed censoring threshold

We here consider a fixed censoring threshold of $\tau = 0.9$ that does not change during training. We compare our NBEs to estimators based on both the pairwise and full likelihood.

4.2.1 COMPARISON TO PAIRWISE LIKELIHOOD ESTIMATORS

We compare NBEs to censored pairwise likelihood estimators for fitting Matérn GPs (see Section C.2 of Appendix C), MSPs and IMSPs; note that inference with each requires estimation of a range and smoothness parameter, λ and κ , respectively. For the censored pairwise likelihoods (CPL), we consider two cut-off distances (see Section 3.1): $h_{max} = 3$ units and $h_{max} = \infty$ units (i.e., where all pairs are used).

Each NBE takes under three hours to train, with $K = 100,000$ training samples, and estimation (given a single parameter set) is incredibly fast, taking on average 0.0016 seconds. In comparison, estimation with the pairwise likelihood is much slower; for $h_{max} = \infty$, the GP and MSP/IMSP required approximately 10 and 20 minutes on average, respectively, with these averages lowering to 2.1 and 1.8 minutes when $h_{max} = 3$; thus, the NBE provides a speed-up of the order of 67,500–750,000 (post-training) in this case.

Results for the three considered processes are given in Tables 2 and A1, which present the marginal test risk with respect to the absolute and squared-error loss, respectively. We observe lower marginal test risk (with respect to both loss functions) for both parameters when using an NBE compared to the censored pairwise likelihood, suggesting that the former is more statistically efficient. Overall, our proposed NBE for censored data exhibits substantial increases in both statistical and computational efficiency relative to the competing likelihood-based approach.

Figure 1 presents the joint distribution of parameter estimates for a single test parameter vector with 1000 sets of replicates. The NBE captures the dependence in estimates of $(\lambda, \kappa)'$ that is also present when utilising the competing pairwise likelihood estimator ($h_{max} = 3$), but our NBE clearly has lower sampling variability than the pairwise likelihood estimator.

Table 2: Results for the simulation study detailed in Section 4.2.1. Marginal test risks (s.d.) under the absolute-error loss are provided separately for λ ($\times 10^{-1}$) and κ ($\times 10^{-2}$). Test risks under the squared-error loss are presented in Table A1 of Appendix E. The lowest value in each column is given in bold.

	GP		MSP		IMSP	
	λ	κ	λ	κ	λ	κ
NBE	3.1 (0.1)	4.5 (0.1)	2.4 (0.1)	1.8 (0.1)	2.6 (0.1)	2.2 (0.1)
CPL ($h_{max} = 3$)	4.4 (0.1)	5.7 (0.2)	3.5 (0.1)	2.2 (0.1)	4.6 (0.2)	3.2 (0.1)
CPL ($h_{max} = \infty$)	9.3 (0.2)	16.0 (0.5)	4.3 (0.1)	6.4 (0.2)	5.4 (0.2)	6.8 (0.2)




Figure 1: Empirical joint distribution from 1000 parameter estimates for (left) GP, (centre) MSP and (right) IMSP, given replicates from a single test parameter set (green cross). Black and blue points denote estimates from the CPL with cut-off distances $h_{max} = \infty$ and $h_{max} = 3$, respectively, whilst the brown points correspond to estimates from the NBE.

4.2.2 COMPARISON TO FULL LIKELIHOOD ESTIMATORS

We now consider inference for the HW and r -Pareto processes detailed in Sections 3.1–3.2. Typical likelihood-based inference for both processes uses the full d -dimensional likelihood, approximated via quasi-Monte Carlo methods; we compare the efficacy of such an approach against our proposed NBE. Inference with the r -Pareto process requires estimation of the Brown–Resnick variogram parameters described in Section 3.1.

For the HW model, previous attempts at inference have been limited to moderate dimension $d \leq 30$ due to the computational expense of evaluating the likelihood (with the exception of Zhang et al. (2022), who pushed it to $d \leq 200$ using specialised Bayesian methods); we find that likelihood inference for replicates on the previously defined 16×16 grid ($d = 256$) is not computationally feasible. Hence, we perform inference for the model simulated on a smaller spatial domain \mathcal{S}^* , which is taken to be a regular 6×6 grid on $[0, 16] \times [0, 16]$ (with $d = 6^2 = 36$). Whilst we compare NBE- and likelihood-based inference on this smaller grid, we further illustrate that inference for the process observed on the larger 16×16 grid is feasible with an NBE; for the r -Pareto process, however, we only

consider the larger grid. We slightly adjust the architecture in Table 1 to account for the smaller size of \mathcal{S}^* relative to \mathcal{S} ; this architecture is given in Table A6.

We use $K = 100,000$ ($K = 500,000$) for the r -Pareto process and HW process on \mathcal{S} (and \mathcal{S}^*). As relatively less information is contained within a single replicate, we found that larger K was required to train the NBEs with the smaller domain \mathcal{S}^* ; see also, Section 5. NBE estimation time using either architectures (Table 1 or Table A6 of Appendix E) is comparable with that described in Section 4.2.1. Estimation time with the full likelihood is unpredictable as it is dependent on the number of censored locations within a field, which varies between replicates. For the r -Pareto process, the estimation time varies between approximately 2.5 and 12 hours, and for the HW process (on \mathcal{S}^*), it varies between approximately 1 and 10 hours, but with some cases requiring upwards of 24 hours. These times clearly show that likelihood-based inference for a HW process observed on \mathcal{S} is not computationally feasible, even when using astute quasi-Monte Carlo approximations.

Results for the r -Pareto process and HW processes are presented in Tables A2–A5 of Appendix E. For the r -Pareto process, our NBE significantly outperforms the likelihood-based estimator; the marginal test risk (with respect to both considered losses) is lower for both parameters, suggesting a gain in statistical efficiency when using the NBE. For the HW process (with domain \mathcal{S}^*), the NBE is more computationally efficient than the likelihood-based approach, and it exhibits gains in statistical efficiency, especially for the parameter δ considered the most difficult parameter to infer (Huser and Wadsworth, 2019). Computational limitations mean that we are prohibited from using the full likelihood-based approach for inference of the process on the larger domain \mathcal{S} , but the NBE does not have the same limitation.

Figure 2 presents the joint distribution of parameter estimates, for both the r -Pareto, and HW, processes, given a single test parameter vector with 1000 sets of replicates. For the r -Pareto process, we observe that the NBE provides unbiased estimation of $\boldsymbol{\theta} = (\lambda, \kappa)'$, with a reduction in the variance of the estimator relative to the likelihood-based approach. For the HW process, we observe that, when considering \mathcal{S}^* , the joint distribution of parameter estimates is similar for NBE- and likelihood-based inference; when the size of \mathcal{S} grows, a substantial reduction in estimator variance is observed.

4.3 Varying censoring threshold

We now investigate the amortised nature of NBEs, in particular their efficiency in handling variable censoring levels using a single estimator (trained only once); we constrain our focus to censoring levels $\tau \in [0.85, 0.95]$, which is the typical range for such values in applications of spatial peaks-over-threshold models, and set $K = 125,000$ training samples throughout. We consider inference for the IMSP and HW processes, defined over the 16×16 grid \mathcal{S} , as considered in Section 4.2.2 and with the same priors on their respective parameter sets. An estimator, accommodating variable τ , is constructed using the scheme described in Section 2.2.2: the “ τ -random” NBE is trained by sampling τ from a $\text{Unif}(0.85, 0.95)$ prior. For testing, we construct 1000 parameter sets with a fixed τ drawn randomly from $\text{Unif}(0.85, 0.95)$ and for comparison, we also train an NBE for this fixed τ (as in Section 4.2), denoted the “ τ -fixed” NBE. The marginal test risk is compared across the two different estimators, with the test value of τ being 0.893 and 0.919 for the IMSP and HW process,




Figure 2: Empirical joint distribution with 1000 parameter estimates for (left) r -Pareto and (centre/right) HW processes, given replicates from a single test parameter set (green cross); left/centre and right panels plot estimates of $(\lambda, \kappa)'$ and $(\kappa, \delta)'$, respectively. Black points denote estimates from the likelihood inference scheme. Brown points are NBE-based estimates for processes observed on the larger domain \mathcal{S} ; yellow points are NBE-based estimates for the HW process observed on the small domain \mathcal{S}^* .

Table 3: Results for the variable τ simulation study detailed in Section 4.3. Marginal test risks (s.d.) under the absolute-error loss are provided separately for λ ($\times 10^{-1}$), κ ($\times 10^{-2}$) and δ ($\times 10^{-2}$). Test risks under the squared-error loss are presented in Table A7 of Appendix E. The lowest value in each column is given in bold.

τ	IMSP		HW		
	λ	κ	λ	κ	δ
random	2.97 (0.07)	2.41 (0.06)	2.62 (0.07)	2.13 (0.05)	2.98 (0.09)
fixed	3.03 (0.07)	2.11 (0.05)	2.75 (0.06)	2.41 (0.06)	3.25 (0.10)

respectively. Note that the training parameter sets are the same across both types of NBE. The NBE architecture for τ -fixed is given in Table 1; the NBE architecture for τ -random differs slightly to Table 1, with an input dimension of 257 in the first dense layer. Table 3 presents the marginal test risk estimates for all processes and estimators; these are evaluated for τ at the τ -fixed level (see above). For the IMSP, the marginal test risk for λ and κ are minimised by the τ -random and τ -fixed NBE, respectively; for the HW process, the marginal test risks for all parameters are minimised by the NBE trained with τ varying during training (see Section 2.2.2). Overall, the results in Table 3 suggest that the τ -random NBE often performs slightly better than the τ -fixed NBE, although the differences in the marginal test risk are quite small and not always statistically significant. Therefore, in practice, both NBEs perform almost equally well. This has substantial implications for the generalisability of our estimator and highlights its amortised nature; a single estimator can be trained for a general τ , rather than an estimator for every choice of τ , which greatly reduces the computational expense of training.




Figure 3: Empirical joint distribution with 1000 parameter estimates given $m = 200$ replicates from a single test parameter set (green cross); top and centre rows give estimates of $(\lambda, \kappa)'$ for the IMSP and HW processes, respectively, and the bottom row illustrates $(\kappa, \delta)'$ for the HW process. Black and orange points denote estimates using a τ -fixed and τ -random NBE, respectively. Test parameter sets and censoring level τ differ across panels; note that τ -fixed was trained for the fixed τ reported in panels in the first column.

To further evidence the efficacy of the new estimator, we illustrate in Figure 3 the joint distribution of parameter estimates from the two NBEs, for three different test parameter sets and true censoring levels; results for three additional test sets are provided in Figure A1 of Appendix D. We observe good performance for τ -random, the NBE trained for general τ , across all test parameter sets and censoring levels; the estimates appear generally unbiased

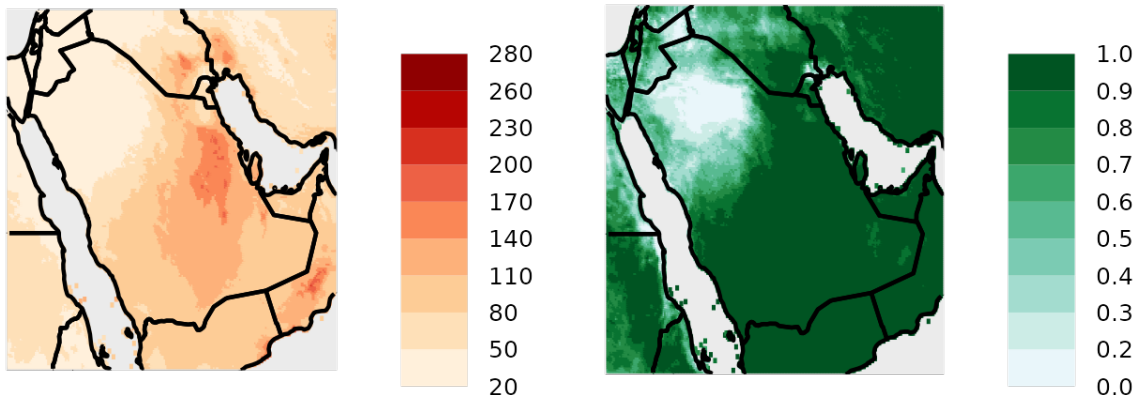


Figure 4: Observation of (left) surface average PM_{2.5} concentration ($\mu\text{g}/\text{m}^3$) for July 2012 and (right) its corresponding standardised field.

as the test parameter estimates are centred on the true parameter set, except in the cases where the true parameter set is close to the boundaries of the prior support (e.g., κ close to 0.5); unsurprisingly, the efficacy of the τ -fixed NBE is only comparable with the τ -random NBE when its training τ is equal (or very close) to the true censoring level (the first column in Figure 3). The τ -fixed NBE cannot extrapolate to previously unseen censoring levels, and so it should be applied only with the specific τ used during training.

5. Application

Inhalation of ambient fine micro-scale particulate matter (PM_{2.5}) can lead to fatal respiratory and cardiovascular diseases (Polichetti et al., 2009; Xing et al., 2016); globally, poor air quality is responsible for millions of deaths per annum (Lelieveld et al., 2019). A country that is particularly affected by poor air quality is Saudi Arabia, due in part to its heavy oil combustion rates, industrial and traffic emissions (Khodeir et al., 2012), and its arid climate, which is conducive to frequent dust storms (Alharbi et al., 2013). In order to mitigate the impacts of extreme air pollution in Saudi Arabia, an understanding of the spatial characteristics of extreme air pollution across the entire region is required; whilst studies of the spatial characteristics of PM_{2.5} across the Arabian Peninsula have been conducted by, for example, Munir et al. (2017) and Tariq et al. (2022), they often focus on quantifying marginal trends or investigating the spatial profile of events through empirical methods, rather than considering extremal dependence. We provide the first analysis of the spatial extremal dependence of PM_{2.5} concentrations across Saudi Arabia using asymptotically-justified spatial extremal dependence models.

Data are taken from version 5 (V5.GL.02) of the global ground level product developed by Van Donkelaar et al. (2021), which consists of monthly mean surface PM_{2.5} concentrations ($\mu\text{g}/\text{m}^3$) from 1998 to 2020 ($m = 276$); the product is constructed by combining satellite data from a number of sources with ground level observations, through a geographically-weighted regression, and samples over land only. Our observations are arranged on a regular 242×182 grid encompassing Saudi Arabia; each observation corresponds to the average over

a 0.1° by 0.1° spatial grid-cell. As our interest is in modelling extremal dependence only, we use a simple rank transform to standardise the data to $\text{Unif}(0, 1)$ margins; to account for spatial non-stationarity in the margins, we perform the transformation site-wise. The observed field for July 2012, which exhibits the largest pointwise value over the observation period, and its standardised values are presented in Figure 4.

We assess extremal dependence in the data through fits of an anisotropic extension of the Huser and Wadsworth (2019) process. Dependence in our model is represented as a function of the great-circle distance between transformed sites,

$$\mathbf{s}^* = \begin{pmatrix} 1 & 0 \\ 0 & 1/\alpha \end{pmatrix} \begin{pmatrix} \cos \omega & -\sin \omega \\ \sin \omega & \cos \omega \end{pmatrix} \mathbf{s}, \quad (9)$$

where $\omega \in [-\pi/2, 0]$ and $\alpha > 0$ control the rotation and coordinate stretching effect, respectively; hence inference requires estimation of $\boldsymbol{\theta} = (\lambda, \kappa, \delta, \alpha, \omega)'$, for $(\lambda, \kappa, \delta)'$ in Section 3.2. To account for spatial non-stationarity in the data, we perform local extremal dependence model fits. We train a single estimator for a $G \times G$ regular grid of sites, with $G \in \{4, 8, 16, 24, 32\}$, and apply it on all subsets of similarly arranged data; the value of G determines the level of smoothing in maps of the parameter estimates. The estimator is only applied to subsets of data that contain no values sampled over water, and so we require between 16,393 ($G = 32$) and 35,589 ($G = 4$) model fits for each G ; although only five estimators are trained for the analysis, they facilitate estimation, and parameter uncertainty assessment, for over 130,000 local spatial models. Recall from Section 3.2 that applications of the isotropic Huser and Wadsworth (2019) model using the full likelihood have been limited to low dimension $d \leq 30$ (or $d \leq 200$, in the special case of Zhang et al., 2022); here we consider much higher d up to $32^2 = 1024$.

Each NBE is trained on a reference domain \mathcal{S} taken to be a $G \times G$ regular grid of locations, which is centred within the spatial domain of interest, see Figure 4. We fix $\tau = 0.9$ throughout and consequently adopt the τ -fixed estimator, as discussed in Section 4.3; we choose an NBE optimised for fixed rather than variable τ as the latter is more computationally expensive to train and the gains in its use are comparably slight (relative to the gains in using an NBE over the competing likelihood approach, see Table 2). We train each NBE using $\tilde{\mathbf{m}} = (46, 138, 276)$. To improve model training and reduce computational memory requirements, we adopt “simulation-on-the-fly” (for an overview, see Chan et al., 2018; Gerber and Nychka, 2021; Sainsbury-Dale et al., 2024a). For training, we begin with an initial $|\vartheta_{\text{train}}| = K$ and $|\vartheta_{\text{val}}| = K/5$ training and validation parameter sets, respectively. Before every 30th epoch, the parameters are refreshed and new values are drawn from the prior and, at the end of every fifth epoch, new training data are refreshed using the current parameter sets. We use the maximum, computationally-feasible value of K , which, due to increasing computational expense, changes with the dimension G ; we take K equal to 750,000, 330,000, 100,000 and 38,000 for dimension 4 or 8, 16, 24 and 32, respectively. The architecture of each NBE is also dependent on G ; these are given in Appendix E. Parameter sets are a priori independent with priors $\lambda \sim \text{Unif}(20, 1250)$, $\kappa \sim \text{Unif}(0.1, 4)$, $\delta \sim \text{Unif}(0, 1)$, $\alpha \sim \text{Unif}(0.5, 3.5)$ and $\omega \sim \text{Unif}(-\pi/2, 0)$. Although the various NBEs take up to 72 hours to train, their evaluation on observed datasets is extremely fast: a single evaluation takes approximately 1×10^{-3} and 4.1×10^{-3} seconds for $G = 4$ and $G = 32$, respectively.

For brevity, we constrain our focus to the case $G = 16$; results for all other values of G are provided in Appendix D. As the final activation function in each architecture is the identity (see Tables A8–A11 of Appendix E), the parameter estimates are not guaranteed to lie in their respective prior support; we simply truncate such values. Although we could choose an architecture that enforces parameter estimates to lie within their prior support, we found that training of the NBE was more numerically stable with our approach. Figures A2–A6 in Appendix D provide the joint distribution of parameter estimates for four different test parameter sets and across all values of G . We observe good performance for the NBEs across all test parameter sets, except in cases where the test parameter set is on the boundary of the prior parameter space.

We illustrate estimates of $(\lambda, \kappa, \delta)'$ and $(\alpha, \omega)'$ parameters for $G = 16$ in Figure 5. As expected, we observe a clear increase in spatial smoothness of parameter estimates as G increases. Estimates of λ are high ($> 600\text{km}$) across most of the Arabian Peninsula, suggesting strong spatial persistence in $\text{PM}_{2.5}$ extremes; smaller estimates of λ are observed in parts of Iran, Jordan and Oman, and close to the coastline of the Arabian Gulf, with particularly large estimates in the Empty Quarter to the south of Saudi Arabia and north-east of Yemen. Estimates of κ illustrate the smoothness of the extremal process, with larger values in the Empty Quarter suggesting that the process here is smoother. Recall that the parameter δ determines if the process is asymptotically dependent ($\delta \geq 1/2$) or asymptotically independent ($\delta < 1/2$); Figure 5 illustrates that Arabian $\text{PM}_{2.5}$ is composed of a mixture of locally asymptotically dependent, and asymptotically independent, processes. In central Saudi Arabia and Iran, the estimated process is asymptotically dependent whilst it appears to be increasingly asymptotically independent as we move towards the coast. For other countries such as Iraq, Jordan and Sudan, the estimated process exhibits asymptotic independence. For the anisotropy parameters, $(\alpha, \omega)'$, we observe in Figure 5 larger estimates of the stretch parameter α along the coastline, and large parts of Iran. A directional effect on the extremal dependence structure corresponds to values of $\alpha \neq 1$; we observe strong directionality in the south of Yemen and Iraq, as well as the north of Saudi Arabia.

Parameter uncertainty assessment is performed using a non-parametric bootstrap. Observed fields are sampled with replacement and used to obtain 1000 bootstrap parameter estimates. These are obtained relatively quickly as the estimators do not require re-training; all model fits can indeed be performed using the same five trained estimators, each taking only a few millisecond to evaluate post-training. For all G and each model parameter, marginal 2.5% and 97.5% bootstrap quantile estimates are provided in Figures A9–A11 of Appendix D. Note that upwards of 130 million individual extremal dependence models are estimated to perform our analysis (using only the five trained estimators); to the best of our knowledge, this is a far greater number than any other study utilising the models described in Section 3, and shows that NBEs enable large-scale bootstrap studies with complex models that were impossible to run previously. While here we use NBEs to fit a large ensemble of models, we still expect to see significant computational gains (compared to a likelihood-based approach) when using NBEs to fit a single model. For $G \geq 8$, fitting the Huser and Wadsworth (2019) process model using the full likelihood-based approach is expected to take substantially longer than the total training and evaluation time for a single NBE (> 72 hours; see comparative computation times for $G = 6$ in Section 4.2.2).

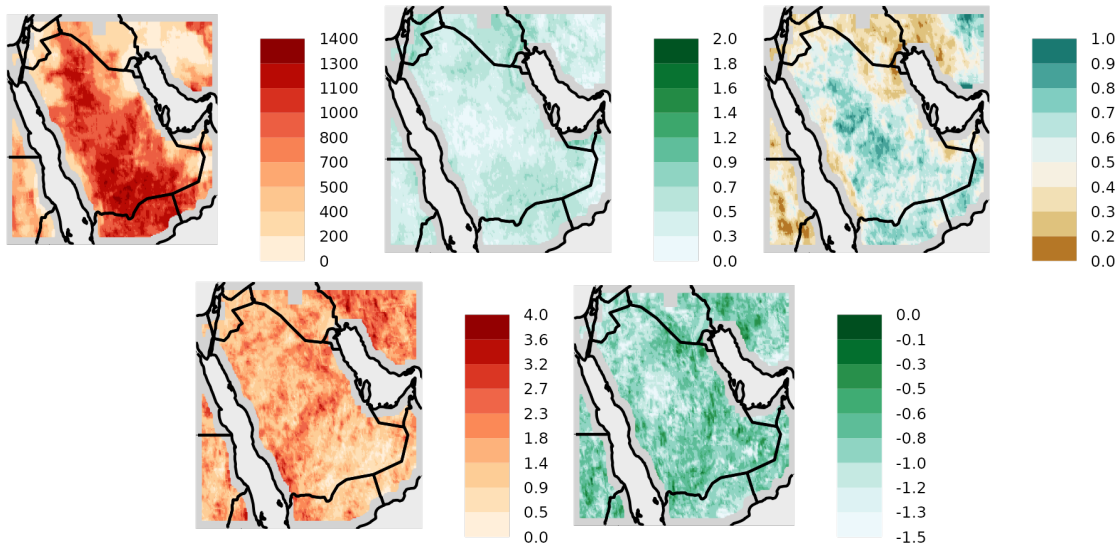


Figure 5: Parameter estimates for the HW process model with $G = 16$. The top and bottom rows give estimates of dependence and anisotropy parameters, $(\lambda, \kappa, \delta)'$ and $(\alpha, \omega)'$, respectively. Note that every pixel corresponds to a parameter estimate for a model fitted locally over the surrounding $G \times G$ grid of pixels.

6. Discussion

We develop neural Bayes estimators (NBEs) for performing fast likelihood-free inference with popular peaks-over-threshold models, providing a paradigm shift challenging traditional statistical inference with censored data. Our novel approach to handling censored inputs for neural networks facilitates the construction of efficient estimators for a wide range of spatial extremal dependence models, including max-stable, r -Pareto, and random scale mixture processes, and we illustrate substantial gains in both the computational, and statistical, efficiency of NBEs relative to competing likelihood-based approaches. We have illustrated our proposed methodology with an application to Saudi Arabian $\text{PM}_{2.5}$ concentrations, where we show that high-dimensional models can be fitted in a matter of milliseconds (post-training). The ability to fit hundreds-of-thousands of models accurately in a reasonable time frame allows us to study our dataset on a scale that is much higher than was previously possible with standard likelihood-based inference techniques, and we reveal new insights into the extremal dependence of air pollution across the Arabian Peninsula.

The efficacy and speed of NBEs opens the door to fitting complex extreme value process models to data of unprecedented size in an operational and online manner (thanks to their amortised nature). Although the training phase may require substantial time, it needs to be performed only once, after which an NBE can be reapplied, in perpetuity, for a negligible computational cost. Simulation of data during training can be costly, particularly when the size K of the training data or number of sampling locations d are large. Approximate methods for data simulation may alleviate this challenge.

In the simulation study in Section 4.3, we demonstrate that it is possible to build an NBE for an arbitrary, user-specified censoring level. However, as we take the censoring threshold to be a function of a single probability τ , our proposed framework does not cover site-specific or random censoring that may arise, for example, when modelling rainfall (Richards et al., 2022, 2023b). Estimation under general, user-specified censoring schemes is a future consideration. Our framework also has the constraint that \mathcal{S} is a regular grid with fixed dimension. To increase the generality of our framework, this constraint should be removed as it implies that for new data observed over a different domain \mathcal{S} , we must retrain a brand new estimator from scratch. For gridded domains, we may be able to build a CNN-based estimator that handles arbitrary grid dimensions by exploiting the scalable CNNs developed by He et al. (2015), but this would still constrain \mathcal{S} to be a regular grid. Sainsbury-Dale et al. (2024b) instead leverage graph neural networks (see the review by Wu et al., 2020) to build NBEs which can be applied to irregular data and which are not wedded to a specific number d or a specific configuration of spatial locations. Combining these estimators with our censoring framework is treated as future work. We also note that, whilst here we focus on inference for spatial processes with NBEs, our methodology can be easily adapted to other application types which utilise left-, right-, or interval-censored data (as, e.g., in survival data analysis, Klein and Moeschberger, 2003), or to other types of neural inference methods (see the review by Zammit-Mangion et al., 2025), by adapting the architecture and inputs.

We did not provide a comparative study to assess the sensitivity of our NBEs to the choice of neural network architecture; such a study, and theoretical results pertaining to this question, are left as future work. Our choice of architecture was based on that proposed by Gerber and Nychka (2021) for data observed on a 16×16 grid; we made minor adjustments for different grid dimensions. With this architecture we were able to construct NBEs that significantly outperform (in terms of both computational speed and estimation accuracy) the gold-standard likelihood-based inference approach. This suggests that this architecture might be suitable for many applications of a similar nature.

We have only considered stationary models. Non-stationary spatial processes, such as those proposed by Huser and Genton (2016), Zammit-Mangion and Rougier (2020), Richards and Wadsworth (2021), and Shao et al. (2024), provide a more realistic assessment of environmental risk and local impacts for data observed over complex domains. However, as these models are often characterised by a high-dimensional parameter set, the training of NBEs in this context may require more complex neural-network architectures and larger training datasets; we here limit the parameter-set dimension to five and, whilst this is representative for many stationary spatial extremal dependence models, some models, like the spatial conditional extremes model of Wadsworth and Tawn (2022), can have dozens of parameters (see, e.g., Richards et al., 2022, 2023b). Whilst Sainsbury-Dale et al. (2024a) illustrate inference of models with up to eight parameters, the extension of NBEs to higher dimensions remains the subject of future work.

Throughout, we have used uniform priors for model parameters; these are the most commonly used in practice and, by equipping them with a reasonably large support, they allow us to construct NBEs that can be widely applied in various applications. Similarly to traditional Bayes estimators, NBEs are optimal with respect to the chosen prior measure, and are thus affected by the choice of prior. Finally, we note that we have here only

considered settings where the process model is well-specified. Investigating the efficacy of our trained estimators under model misspecification (see, e.g., Radev et al., 2023; Schmitt et al., 2023) remains future work.

Acknowledgments

This publication is based upon work supported by King Abdullah University of Science and Technology (KAUST) Research Funding under Awards No. ORA-2022-5336 and RFS-OFP2023-5550. Additionally, Jordan Richards and Raphaël Huser were supported by funding from KAUST under Award No. OSR-CRG2020-4394, while Matthew Sainsbury-Dale and Andrew Zammit-Mangion were supported by Australian Research Council (ARC) Discovery Early Career Research Award DE180100203. Matthew Sainsbury-Dale’s research was also supported by an Australian Government Research Training Program Scholarship and a 2022 Statistical Society of Australia (SSA) top-up scholarship. Support from the KAUST Supercomputing Laboratory is gratefully acknowledged.

Appendix

Appendix A provides proof of Proposition 1. Appendix B provides a discussion and proof of the Bayes estimators’ invariance to the distribution of the marginal censoring level. Appendix C provides further details for the spatial extremal dependence models that we consider in our simulation studies. Appendices D and E provide supplementary figures and tables, respectively.

Appendix A. Proof of Proposition 1 for the likelihood function of censored data

Here we provide a proof of Proposition 1, which derives the likelihood function of censored data. Recall Proposition 1, as stated in the main paper:

Proposition 1 *Let $\mathbf{Z}_t = (Z_{t,1}, \dots, Z_{t,d})'$, $d \in \mathbb{N}, t = 1, \dots, m$, be random vectors defined by a statistical model parameterised by $\boldsymbol{\theta}$ with joint distribution $F(\cdot \mid \boldsymbol{\theta})$ and density $f(\cdot \mid \boldsymbol{\theta})$. Let $\mathbf{Z}_{t+} := (Z_{t,j} : Z_{t,j} > c_j, j = 1, \dots, d)'$ denote the subvector of uncensored values of \mathbf{Z}_t , where $\mathbf{c} := (c_1, \dots, c_d)'$ is a vector of censoring levels with time-dependent subvector $\mathbf{c}_{t-} = (c_j : Z_{t,j} \leq c_j, j = 1, \dots, d)'$. Given indicators $\mathbf{I}_t := (\mathbb{1}\{Z_{t,j} \leq c_j\} : j = 1, \dots, d)'$, the likelihood contribution of $\mathbf{Z}_t^{\text{cens}} := (\mathbf{I}'_t, \mathbf{Z}'_{t+})'$ is*

$$\int_{-\infty}^{\mathbf{c}_{t-}} f(\mathbf{z}_t \mid \boldsymbol{\theta}) d\mathbf{z}_{t-} = \frac{\partial^{|\mathcal{J}_t|}}{\prod_{j \in \mathcal{J}_t} \partial z_j} F(\mathbf{z} \mid \boldsymbol{\theta}) \Big|_{\mathbf{z}=\mathbf{b}_t}, \quad t = 1, \dots, m, \quad (5)$$

where $\mathbf{z}_t = (z_{t,1}, \dots, z_{t,d})'$ is the realisation of \mathbf{Z}_t , $\mathcal{J}_t := \{j : z_{t,j} > c_j\}$ denotes the indices of the uncensored components of \mathbf{z}_t with subvector $\mathbf{z}_{t-} := (z_{t,j} : z_{t,j} \leq c_j, j = 1, \dots, d)'$, and $\mathbf{b}_t := (\max\{z_{t,j}, c_j\} : j = 1, \dots, d)'$.

Proof As in the main paper, we split the observation vector as $\mathbf{Z}_t = (\mathbf{Z}'_{t-}, \mathbf{Z}'_{t+})'$, its realised value as $\mathbf{z}_t = (\mathbf{z}'_{t-}, \mathbf{z}'_{t+})'$, and the censoring threshold vector as $\mathbf{c}_t = (\mathbf{c}'_{t-}, \mathbf{c}'_{t+})'$. Let $\mathbf{i}_t = (i_{t,1}, \dots, i_{t,d})$ be the realised value of \mathbf{I}_t ; that is, for \mathbf{Z}_t arranged as $\mathbf{Z}_t = (\mathbf{Z}'_{t-}, \mathbf{Z}'_{t+})'$,

we have $\mathbf{i}_t = (\mathbf{1}'_{t-}, \mathbf{0}'_{t+})'$, where $\mathbf{1}_{t-}$ and $\mathbf{0}_{t+}$ are vectors of ones and zeros of dimensions $d_{t-} := \sum_{j=1}^d i_{t,j}$ and $d_{t+} := d - d_{t-}$, respectively, which denote the numbers of censored and uncensored values, respectively. Then, since $\mathbf{z}_{t+} > \mathbf{c}_{t+}$ by definition, we note that

$$\begin{aligned} \Pr(\mathbf{I}_t = \mathbf{i}_t, \mathbf{Z}_{t+} \leq \mathbf{z}_{t+}) &= \Pr(\mathbf{Z}_{t-} \leq \mathbf{c}_{t-}, \mathbf{Z}_{t+} > \mathbf{c}_{t+}, \mathbf{Z}_{t+} \leq \mathbf{z}_{t+}) \\ &= \Pr(\mathbf{Z}_{t-} \leq \mathbf{c}_{t-}, \mathbf{c}_{t+} < \mathbf{Z}_{t+} \leq \mathbf{z}_{t+}) \\ &= \Pr(\mathbf{Z}_t \leq \mathbf{b}_t) - \Pr(\mathbf{Z}_t \leq \mathbf{b}_t^*), \end{aligned} \quad (\text{A.1})$$

where $\mathbf{b}_t^* = (b_{t,1}^*, \dots, b_{t,d}^*)'$ and $\mathbf{b}_t = (b_{t,1}, \dots, b_{t,d})'$ are vectors with components $b_{t,j}^* = c_j$ if $z_{t,j} > c_j$ and $b_{t,j}^* = -\infty$, otherwise, and $b_{t,j} = \max\{z_{t,j}, c_j\}$ for $j = 1, \dots, d$. Therefore, because the rightmost term in (A.1) does not depend on \mathbf{z}_{t+} , the censored density/likelihood contribution (obtained by differentiating (A.1) with respect to \mathbf{z}_{t+}) is

$$\frac{\partial^{d_{t+}}}{\partial \mathbf{z}_{t+}} \{\Pr(\mathbf{Z}_t \leq \mathbf{b}_t) - \Pr(\mathbf{Z}_t \leq \mathbf{b}_t^*)\} = \frac{\partial^{|\mathcal{J}_t|}}{\prod_{j \in \mathcal{J}_t} \partial z_j} F(\mathbf{z} \mid \boldsymbol{\theta}) \Big|_{\mathbf{z}=\mathbf{b}_t},$$

which corresponds to the rightmost expression in the proposition. Leveraging the fact that the joint distribution function can be obtained as an integral of the joint density function, this censored likelihood contribution can be rewritten differently as

$$\frac{\partial^{|\mathcal{J}_t|}}{\prod_{j \in \mathcal{J}_t} \partial z_j} F(\mathbf{z} \mid \boldsymbol{\theta}) \Big|_{\mathbf{z}=\mathbf{b}_t} = \int_{-\infty}^{\mathbf{c}_{t-}} f(\mathbf{z}_t \mid \boldsymbol{\theta}) d\mathbf{z}_{t-},$$

which concludes the proof. ■

Appendix B. Invariance of Bayes estimator to the distribution of the censoring level

Here we give the conditions for which the Bayes estimator is invariant to the distribution of the censoring level, $p(\tau)$. For ease of exposition, we assume that hereon all distributions admit a density with respect to Lebesgue measure. Notation used in the proof is defined in Section 2.

Theorem 1 *Let $L : \Theta \times \Theta \rightarrow [0, \infty)$ denote a strictly convex nonnegative loss function and let $p(\boldsymbol{\theta} \mid \mathbf{a}, \tau)$ denote the posterior density of $\boldsymbol{\theta}$ corresponding to the augmented data $\mathbf{A} = \mathbf{a}$ and censoring level τ . Assume that the Bayes estimate $\hat{\boldsymbol{\theta}}^*$ has finite posterior expected loss $\int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*) p(\boldsymbol{\theta} \mid \mathbf{a}, \tau) d\boldsymbol{\theta}$ for all fixed $\mathbf{a} \in \Omega_{\mathbf{A}, \tau}$ and $\tau \in \mathcal{T} \subseteq [0, 1)$. Then the Bayes estimator $\hat{\boldsymbol{\theta}}^*(\mathbf{a}, \tau)$, $\mathbf{a} \in \Omega_{\mathbf{A}, \tau}$, $\tau \in \mathcal{T}$, is invariant to the distribution of τ , $p(\tau)$, provided that $p(\tau) > 0$ for $\tau \in \mathcal{T}$ and that τ and $\boldsymbol{\theta}$ are independent.*

Proof The proof follows a similar argument to that of Sainsbury-Dale et al. (2024b, Theorem 1). For a given $\tau \in \mathcal{T}$ and augmented data $\mathbf{a} \in \Omega_{\mathbf{A}, \tau}$, the Bayes estimate $\hat{\boldsymbol{\theta}}^*$ minimises the posterior expected loss, that is,

$$\hat{\boldsymbol{\theta}}^* = \operatorname{argmin}_{\boldsymbol{\theta}} \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}) p(\boldsymbol{\theta} \mid \mathbf{a}, \tau) d\boldsymbol{\theta}, \quad (\text{A.2})$$

and, by assumption, $\int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*) p(\boldsymbol{\theta} \mid \mathbf{a}, \tau) d\boldsymbol{\theta} < \infty$. Consider now the Bayes estimator $\hat{\boldsymbol{\theta}}^*(\mathbf{a}, \tau)$ that returns the Bayes estimate for each $\mathbf{a} \in \Omega_{\mathbf{A}, \tau}$ and censoring level $\tau \in \mathcal{T}$. Since the posterior expected loss is nonnegative, if $p(\tau) > 0$ for $\tau \in \mathcal{T}$ then it is also the case that

$$\hat{\boldsymbol{\theta}}^*(\cdot) = \operatorname{argmin}_{\hat{\boldsymbol{\theta}}(\cdot)} \int_{\mathcal{T}} \int_{\Omega_{\mathbf{A}, \tau}} \left[\int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{a}, \tau)) p(\boldsymbol{\theta} \mid \mathbf{a}, \tau) d\boldsymbol{\theta} \right] p(\mathbf{a} \mid \tau) p(\tau) d\mathbf{a} d\tau;$$

see Zammit-Mangion et al. (2025) for details. Applying Bayes rule and assuming τ and $\boldsymbol{\theta}$ are independent yields Equation (7) of the main text, thus completing the proof. \blacksquare

Appendix C. Spatial extremal dependence model details

C.1 Univariate extreme value distributions

We give here the forms of the generalised extreme value (GEV) and generalised Pareto (GPD) distributions used for modelling univariate extreme values; see Coles (2001) for details.

The generalised extreme value GEV(μ, σ, ξ) distribution function is

$$F_{\text{GEV}}(z \mid \mu, \sigma, \xi) = \begin{cases} \exp \left[- \left\{ 1 + \xi \left(\frac{z - \mu}{\sigma} \right) \right\}_+^{-1/\xi} \right], & \xi \neq 0, \\ \exp \left\{ - \exp \left(- \frac{z - \mu}{\sigma} \right) \right\}, & \xi = 0, \end{cases} \quad (\text{A.3})$$

with $\{y\}_+ = \max\{0, y\}$, location, scale, and shape, parameters $\mu \in \mathbb{R}, \sigma > 0$, and $\xi \in \mathbb{R}$, respectively, and support $\{z \in \mathbb{R} : 1 + \xi(z - \mu)/\sigma > 0\}$.

A GPD(σ, ξ) random variable has distribution function $F_{\text{GPD}}(z) = (1 + \xi z/\sigma)^{-1/\xi}$ and support $z \geq 0$ for $\xi \geq 0$ and $0 \leq z \leq -\sigma/\xi$ for $\xi < 0$.

C.2 Matérn Gaussian processes

Gaussian processes (GPs) have extensive use in classical spatial statistics due to their convenient statistical properties (see Gelfand and Schliep, 2016), but are not necessarily appropriate for modelling spatial extremes due to their inability to capture asymptotic dependence (Davison et al., 2013), that is, a non-zero limit probability of joint threshold exceedances at any two sites. However, we still consider their inference as they have been previously applied to model spatial extremal dependence and, more importantly, they provide the foundation for many other spatial extremal dependence models (Sang and Gelfand, 2010), for example, MSPs and IMSPs (see Section 3.1 of the main text), random scale/location mixtures (see Section 3.2 of the main text) and the spatial conditional extremes model (Wadsworth and Tawn, 2022).

Unlike MSPs and IMSPs, the full d -dimensional density for a GP is tractable with computationally inexpensive evaluation, and many methods have been developed to facilitate their fast evaluation for large d (see, e.g., Aune et al., 2014). However, under the censoring regime described in Section 2.2.1 of the main text, evaluation of the full likelihood becomes computationally expensive. Hence, for modelling of spatial extremes, inference for GPs usually proceeds with the same censored weighted pairwise likelihood (with cut-off distance h_{max})

as described in Section 3.1 of the main text (see, e.g., Davison et al., 2013; Thibaud et al., 2013; Richards et al., 2022, 2023b). Due to their popularity, we adopt censored pairwise likelihood estimators as the competing likelihood approach for MSPs, IMSPs, and GPs, in the simulation study detailed in Section 4 of the main text.

We constrain our focus to isotropic and stationary Matérn GPs. Specifically, we consider Gaussian copula models where, after arbitrary marginal transformation, the distribution of the process observed at a finite collection of sites $\{\mathbf{s}_1, \dots, \mathbf{s}_d\}$ is zero mean, unit variance, d -variate Gaussian with the correlation between the process observed at sites $\mathbf{s}_i, \mathbf{s}_j \in \mathcal{S}$ determined by the isotropic Matérn correlation function

$$C(\mathbf{s}_i, \mathbf{s}_j) = 2^{1-\kappa} \{\Gamma(\kappa)\}^{-1} (\|\mathbf{s}_i - \mathbf{s}_j\|/\lambda) K_\kappa(\|\mathbf{s}_i - \mathbf{s}_j\|/\lambda),$$

where $\Gamma(\cdot)$ is the gamma function, $K_\kappa(\cdot)$ is the Bessel function of the second kind of order κ , and $\lambda > 0$ and $\kappa > 0$ are range and smoothness parameters, respectively; with the marginal transformations assumed known, inference for these processes involves estimation of $\boldsymbol{\theta} = (\lambda, \kappa)'$. Similar processes were inferred using (uncensored) NBEs by Gerber and Nychka (2021) and Sainsbury-Dale et al. (2024a).

C.3 r -Pareto processes

Whilst max-stable processes are limits for renormalised maxima, the generalised Pareto process describes the limit of renormalised exceedances, defined as $X \mid \{\sup_{\mathbf{s} \in \mathcal{S}} X(\mathbf{s}) > u\}$ for large u , where the process $X(\cdot)$ has unit Pareto margins; these can be seen as analogues of the univariate generalised extreme value (GEV), and generalised Pareto (GPD), distributions respectively. For details on the univariate extreme value distributions, see Section C.1. Following Buishand et al. (2008) and Ferreira and de Haan (2014), consider the process

$$\left[1 + \xi(\mathbf{s}) \left\{ \frac{X(\mathbf{s}) - b_m(\mathbf{s})}{a_m(\mathbf{s})} \right\}_+ \right]^{1/\xi(\mathbf{s})} \mid \sup_{\mathbf{s} \in \mathcal{S}} \frac{X(\mathbf{s}) - b_m(\mathbf{s})}{a_m(\mathbf{s})} > 1, \quad (\text{A.4})$$

where a_m and b_m are as defined in Section 3.1, $\xi(\mathbf{s}) \in \mathbb{R}$ is the associated GEV or GPD shape parameter (see Section C.1) at site $\mathbf{s} \in \mathcal{S}$, and $z_+ = \max\{0, z\}$; then if this process converges weakly, as $m \rightarrow \infty$, to a process with non-degenerate margins, this process is a generalised Pareto process. As with max-stable processes, we typically consider the process on standardised margins; a transformation of $X(\mathbf{s})$ to unit Pareto margins, denoted $\bar{X}(\mathbf{s})$, leads to a standard Pareto process with (A.4) becoming $\bar{X}(\mathbf{s})/m \mid \sup_{\mathbf{s} \in \mathcal{S}} \bar{X}(\mathbf{s}) > m$; this same standardisation is used to create a max-stable process with unit Fréchet margins.

Dombry and Ribatet (2015) extend formulation (A.4) by replacing the conditioning event with an exceedance of some functional of the entire process. They term these processes “ ℓ -Pareto processes”, whilst later developments by de Fondeville and Davison (2018, 2022) refer to them as “ r -Pareto processes”; we adopt the latter terminology. Consider a risk functional $r(\cdot)$, which is homogeneous of order 1 (i.e., $r(\cdot)$ satisfies, for any $c > 0$, that $r(cX) = cr(X)$), and takes in as input the entire process \bar{X} and returns some scalar value; examples include $r(\bar{X}) = \int_{\mathcal{S}} \bar{X}(\mathbf{s}) d\mathbf{s}$ or $r(\bar{X}) = \sup_{\mathbf{s} \in \mathcal{S}} \bar{X}(\mathbf{s})$, or their finite-dimensional equivalents. An r -Pareto process, denoted $Z(\mathbf{s})$ and with standard unit Pareto marginals, is defined by $Z(\mathbf{s}) = \lim_{r_0 \rightarrow \infty} \bar{X}(\mathbf{s})/r_0 \mid r(\bar{X}) > r_0$. Similarly to (8), we can construct an r -Pareto process through the scale mixture $Z(\mathbf{s}) = RW(\mathbf{s})$, but here R is a unit Pareto

random variable which is independent of the non-negative stochastic process $W(\mathbf{s})$, which satisfies $r(W) = 1$ almost surely; hence we can relate each r -Pareto process to a max-stable counterpart. We constrain our focus to inference for r -Pareto processes that arise from log-Gaussian random functions, that is, with Brown–Resnick max-stable counterparts, as these are a popular choice (see, e.g, Healy et al., 2024; Béwentaoré and Barro, 2022); however, we note that extremal- t counterparts are an alternative (Thibaud and Opitz, 2015).

Likelihood-based inference for r -Pareto processes follows with specification of a risk functional and use of a censored likelihood. In the case where $r(\bar{X}) = \max_{i=1,\dots,d} \bar{X}(\mathbf{s}_i)$, the full d -dimensional censored likelihood can be represented as a censored Poisson process likelihood (Wadsworth and Tawn, 2014; Asadi et al., 2015); hereon we use this risk functional for inference. Evaluation of the full censored likelihood is computationally troublesome as it requires evaluation of d_- -variate Gaussian distributions for $1 \leq d_- \leq d - 1$. To circumvent this issue, de Fondeville and Davison (2018) detail a quasi-Monte Carlo algorithm to approximate these integrals and hasten computations, see, also, Genz and Bretz (2009); we take their approach to inference as the competing likelihood-based approach in Section 4 of the main text. Inference requires estimation of the Brown–Resnick variogram parameters described in Section 3.1 of the main text.

C.4 Computational details

For full details on the software used for likelihood-based inference and simulation of the considered spatial extremal processes, see Belzile et al. (2023).

To simulate from the Brown–Resnick process, we use the `SpatialExtremes` package in R (Ribatet, 2022); realisations from the associated inverted Brown–Resnick process are computed by taking the marginal reciprocal. Simulation of the Brown–Resnick r -Pareto processes and their inference with the full d -dimensional censored likelihood is conducted using the `mvPOT` package (de Fondeville and Belzile, 2022). For the HW process, both simulation and inference are conducted using the R package `SpatialADAI` (Wadsworth, 2018), with inference using the quasi-Monte Carlo approximation requiring the `mvPOT` package (de Fondeville and Belzile, 2022).

Appendix D. Supplementary figures

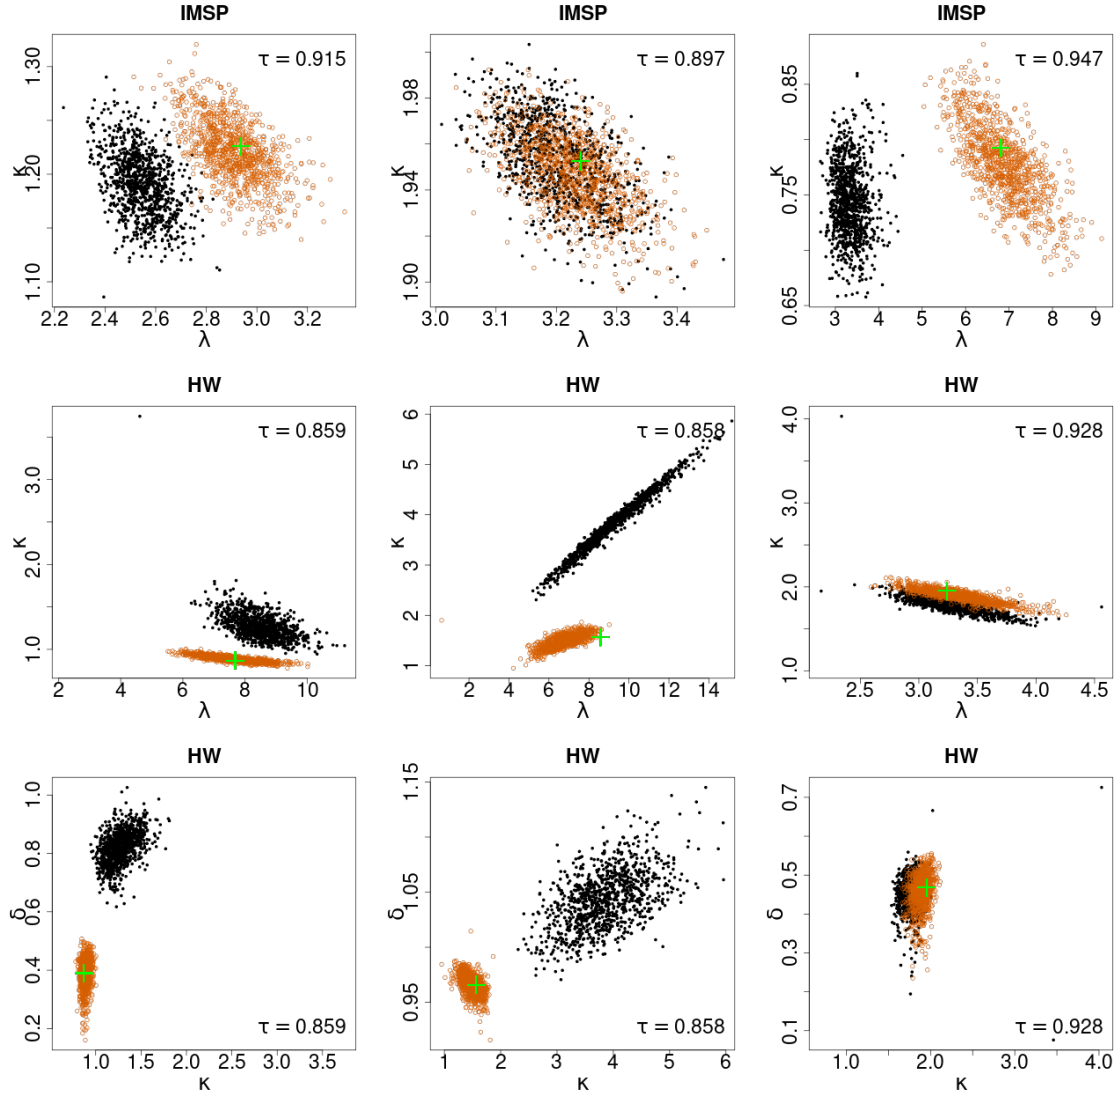


Figure A1: Joint distribution of 1000 parameter estimates given $m = 200$ replicates from a single test parameter set (green cross); top and centre rows give estimates of $(\lambda, \kappa)'$ for the IMSP and HW processes, respectively, and the bottom row illustrates $(\kappa, \delta)'$ for the HW process. Black, orange and blue points denote estimates using a τ -fixed, τ -random and τ -sequence NBE, respectively (see Section 4.3 of the main text). The test parameter sets and censoring level τ differ across panels.

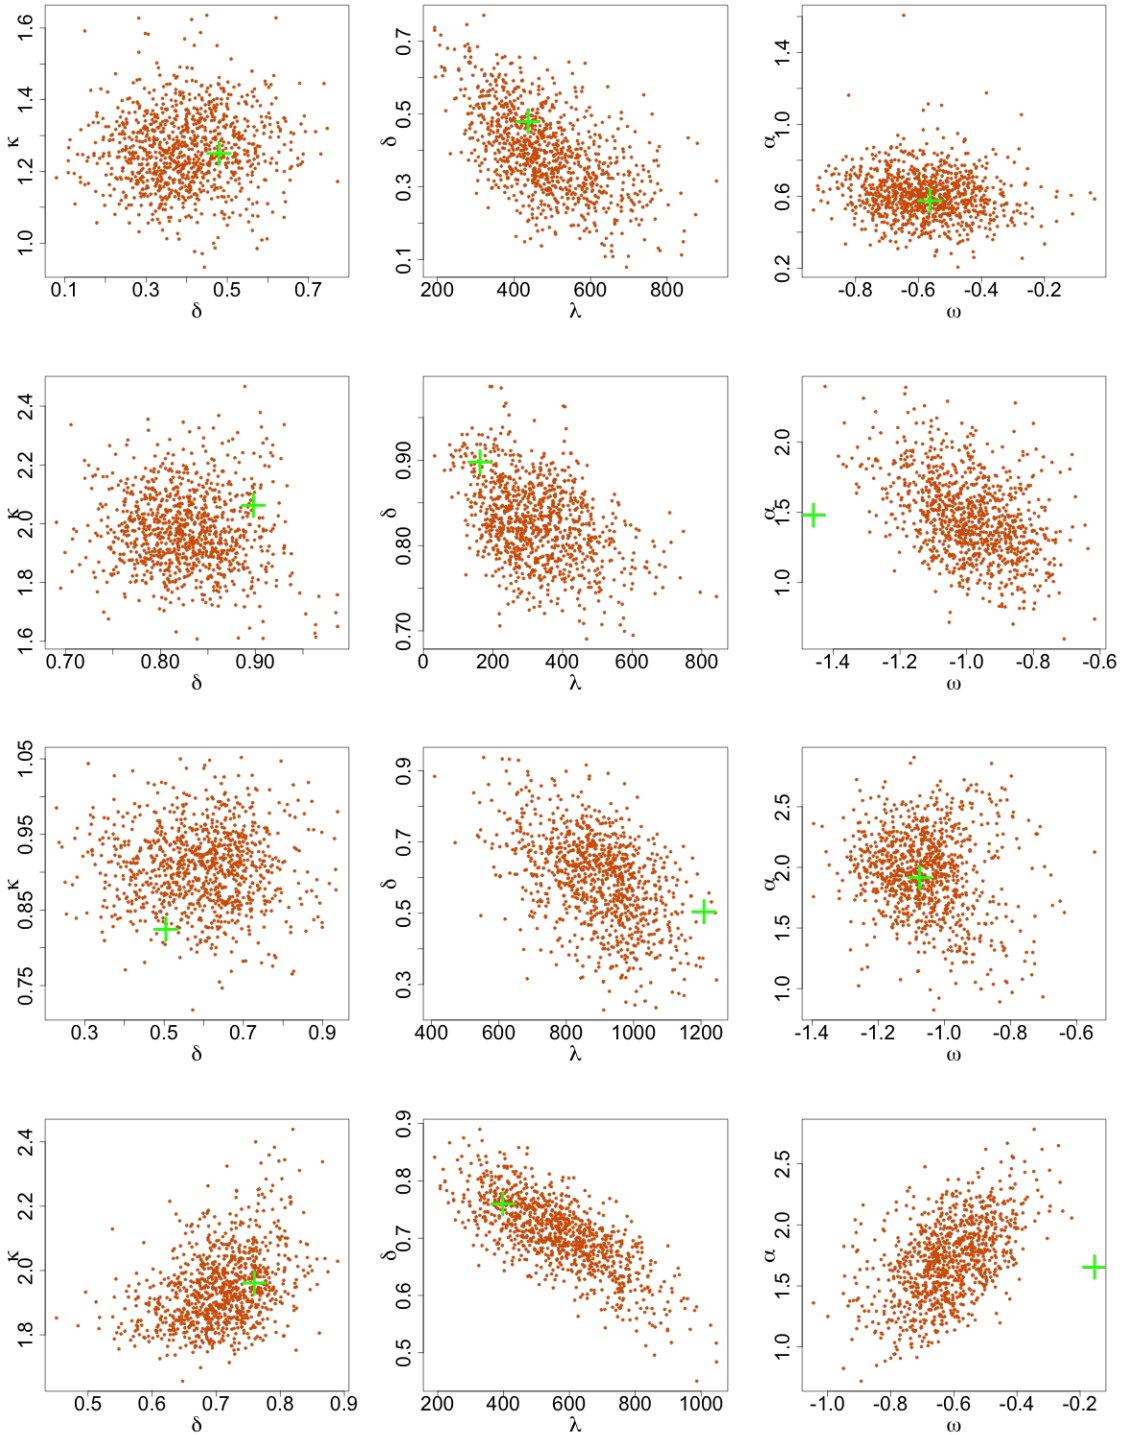


Figure A2: Joint distribution of 1000 parameter estimates (orange) when $G = 4$ in Section 5 of the main text. Four test parameter sets (green) are drawn randomly from the prior and differ across rows. The left, centre and right columns give estimates of $(\delta, \kappa)'$, $(\lambda, \delta)'$ and $(\omega, \alpha)'$, respectively.

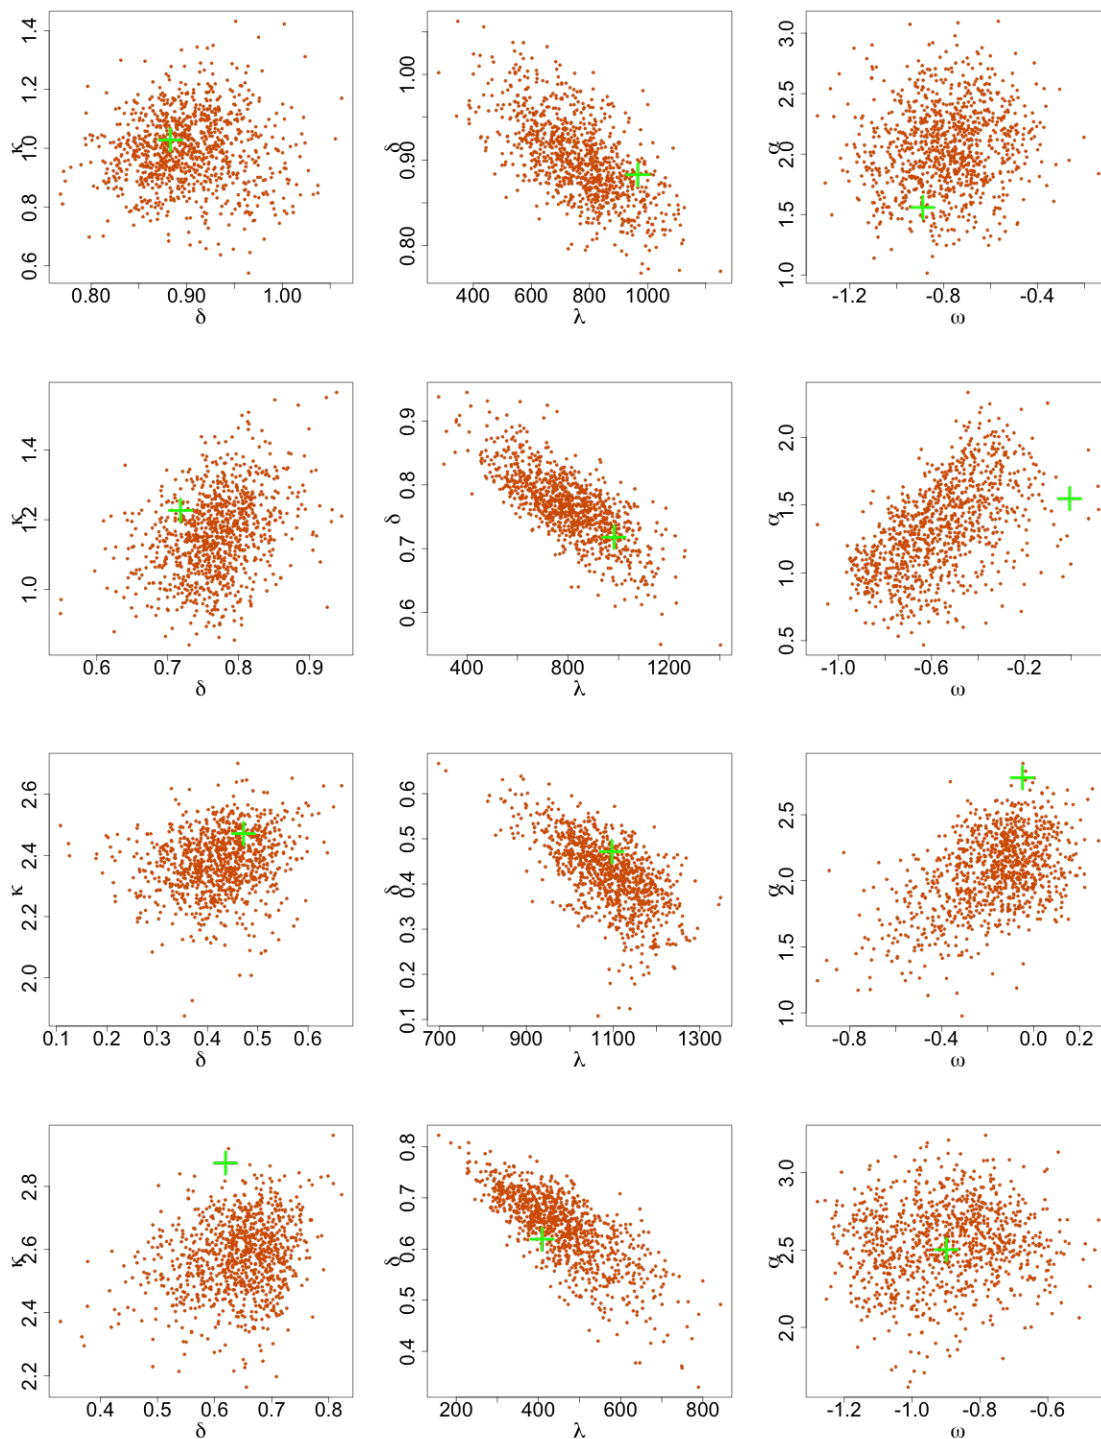


Figure A3: Joint distribution of 1000 parameter estimates (orange) when $G = 8$ in Section 5 of the main text. Four test parameter sets (green) are drawn randomly from the prior and differ across rows. The left, centre and right columns give estimates of $(\delta, \kappa)'$, $(\lambda, \delta)'$ and $(\omega, \alpha)'$, respectively.

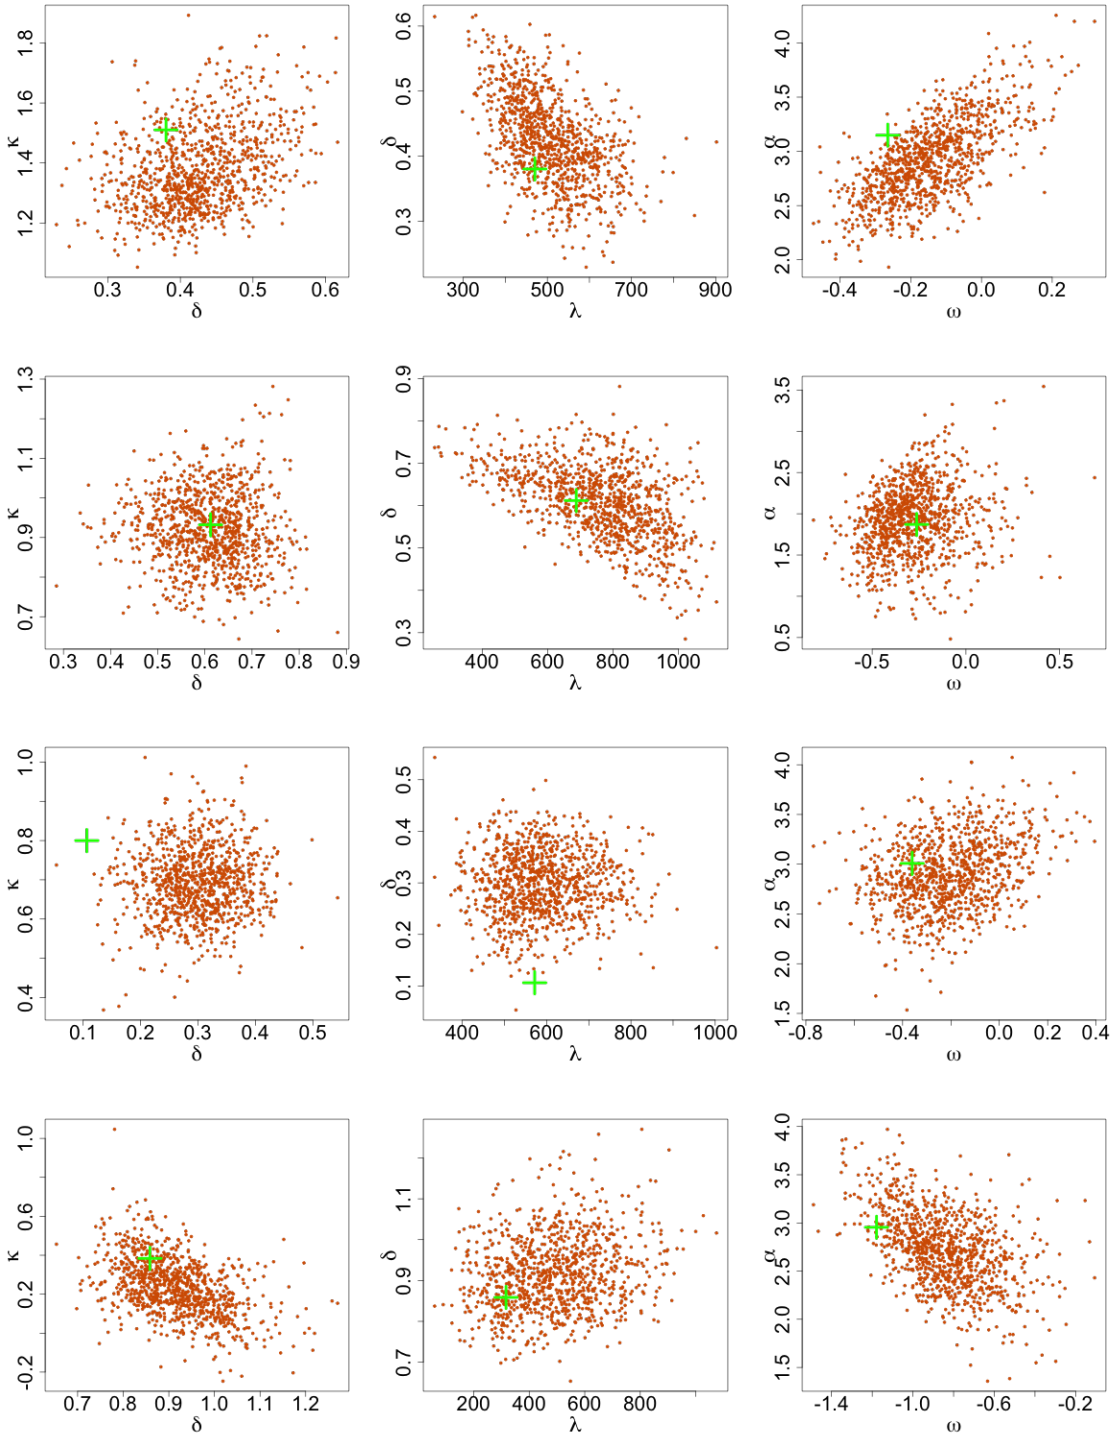


Figure A4: Joint distribution of 1000 parameter estimates (orange) when $G = 16$ in Section 5 of the main text. Four test parameter sets (green) are drawn randomly from the prior and differ across rows. The left, centre and right columns give estimates of $(\delta, \kappa)'$, $(\lambda, \delta)'$ and $(\omega, \alpha)'$, respectively.

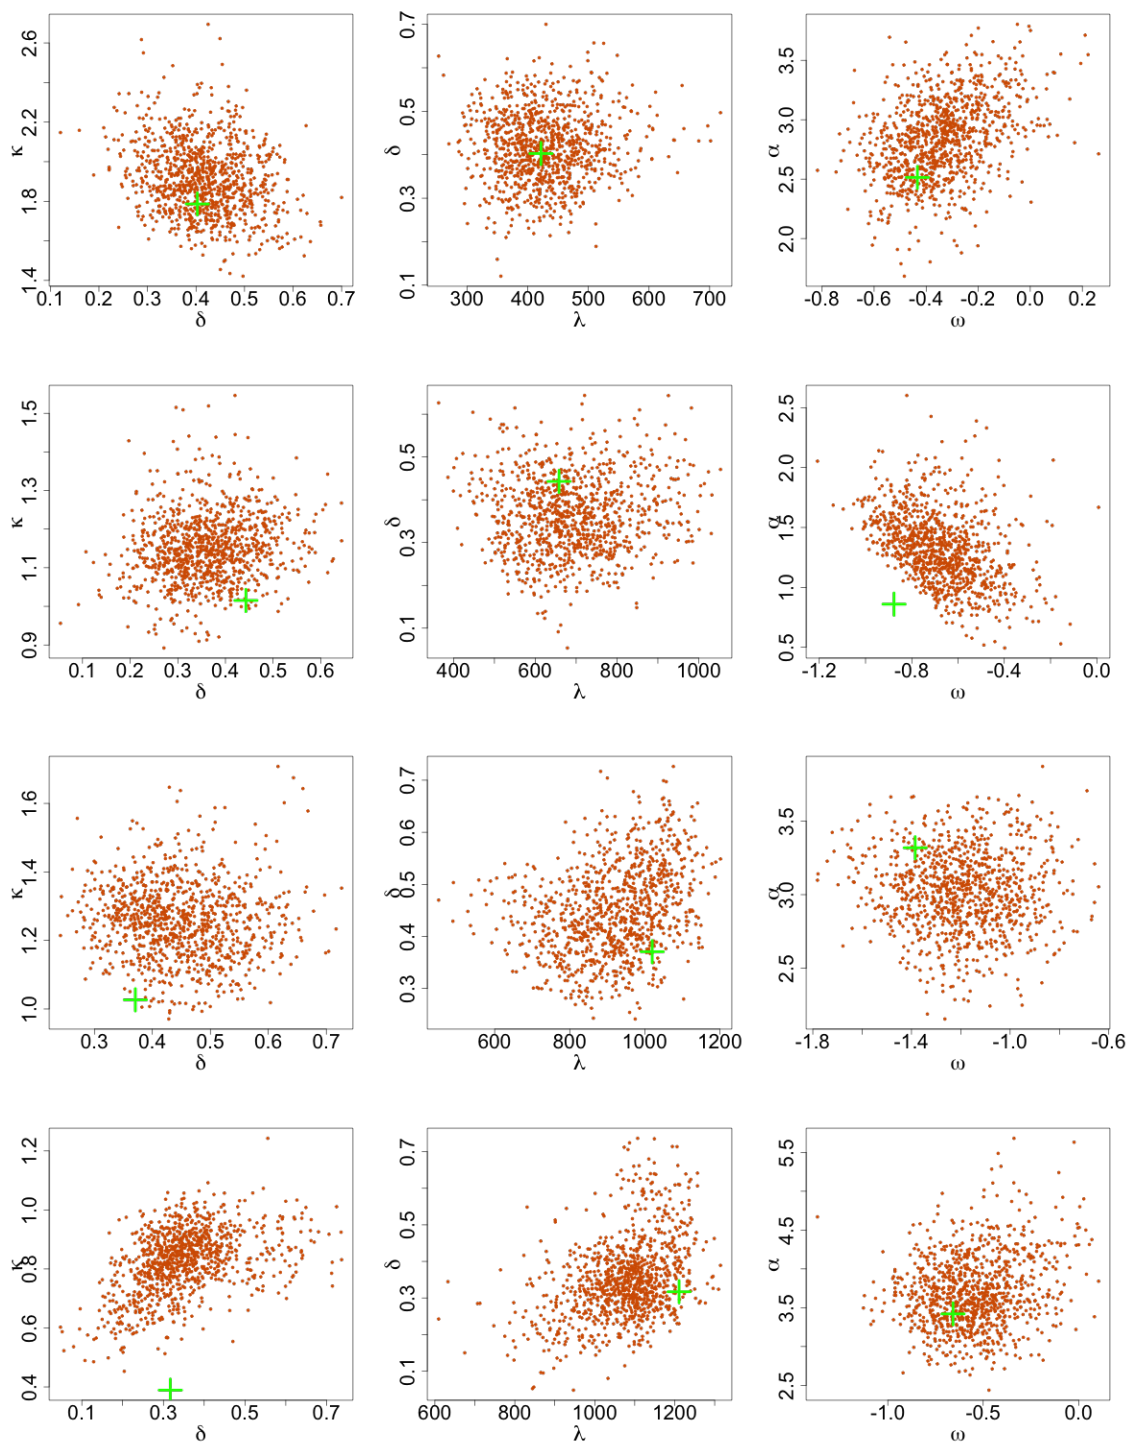


Figure A5: Joint distribution of 1000 parameter estimates (orange) when $G = 24$ in Section 5 of the main text. Four test parameter sets (green) are drawn randomly from the prior and differ across rows. The left, centre and right columns give estimates of $(\delta, \kappa)'$, $(\lambda, \delta)'$ and $(\omega, \alpha)'$, respectively.

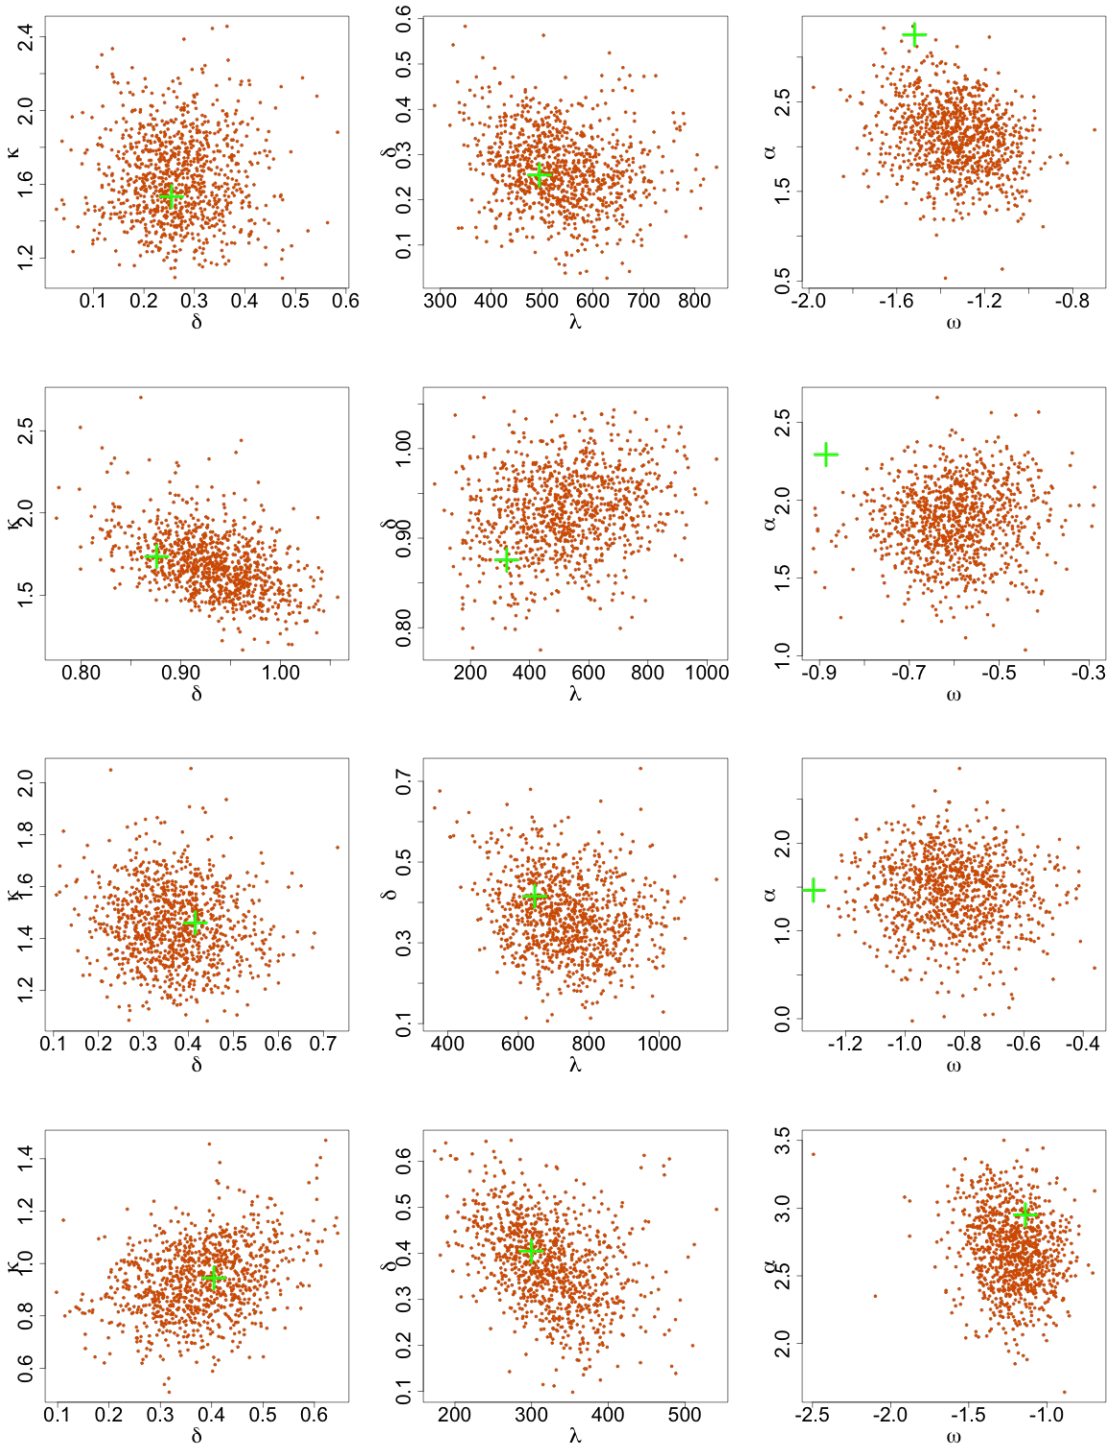


Figure A6: Joint distribution of 1000 parameter estimates (orange) when $G = 32$ in Section 5 of the main text. Four test parameter sets (green) are drawn randomly from the prior and differ across rows. The left, centre and right columns give estimates of $(\delta, \kappa)'$, $(\lambda, \delta)'$ and $(\omega, \alpha)'$, respectively.

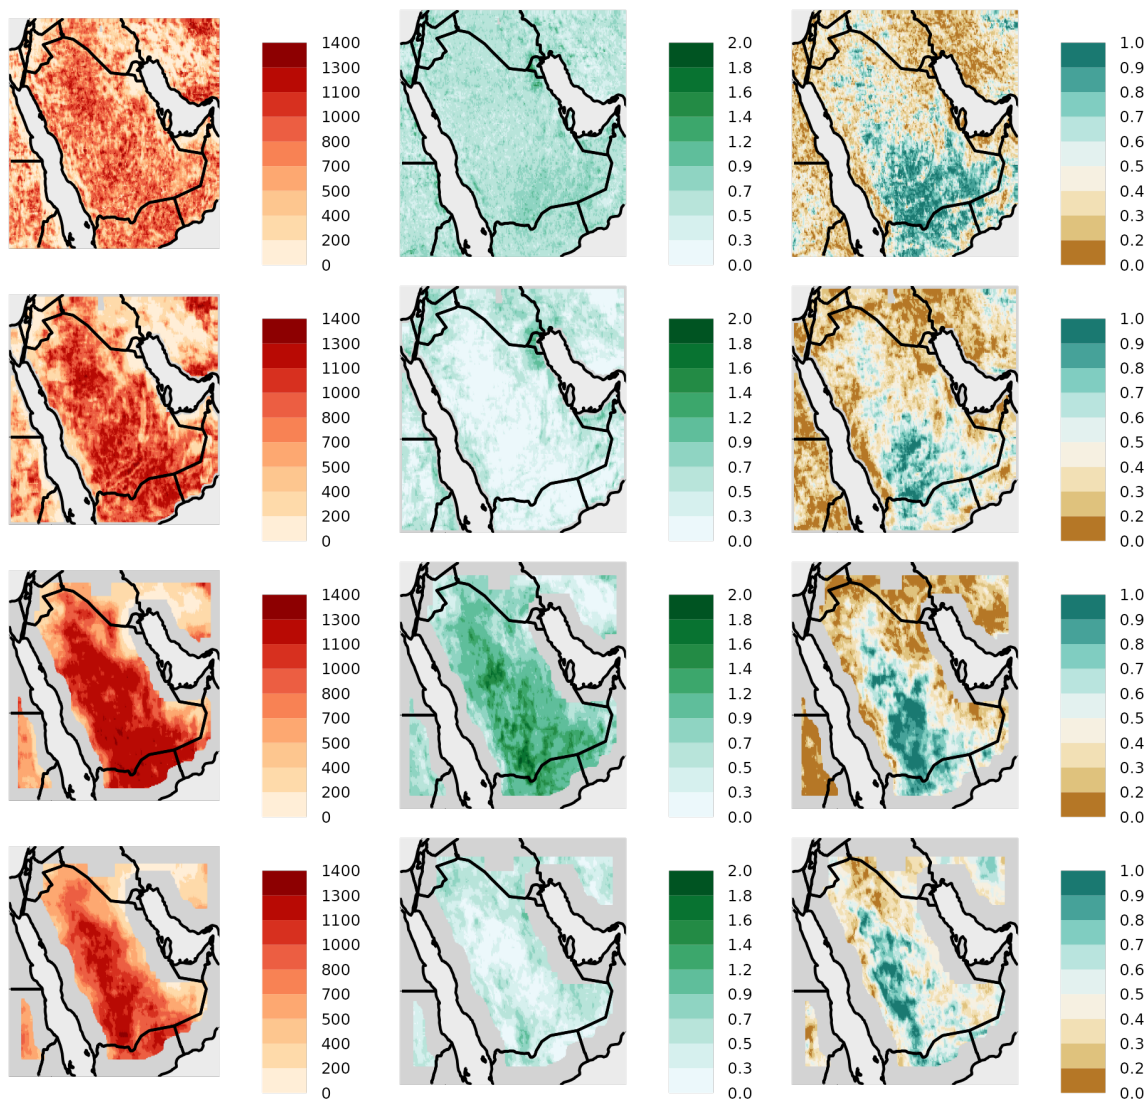


Figure A7: Extremal dependence parameter estimates for the HW model used in the application, Section 5 of the main text. The smoothing level G changes with the row (1st, 4; 2nd, 8; 3rd, 24; 4th, 32). The left, centre and right columns illustrate estimates of λ , κ and δ , respectively.

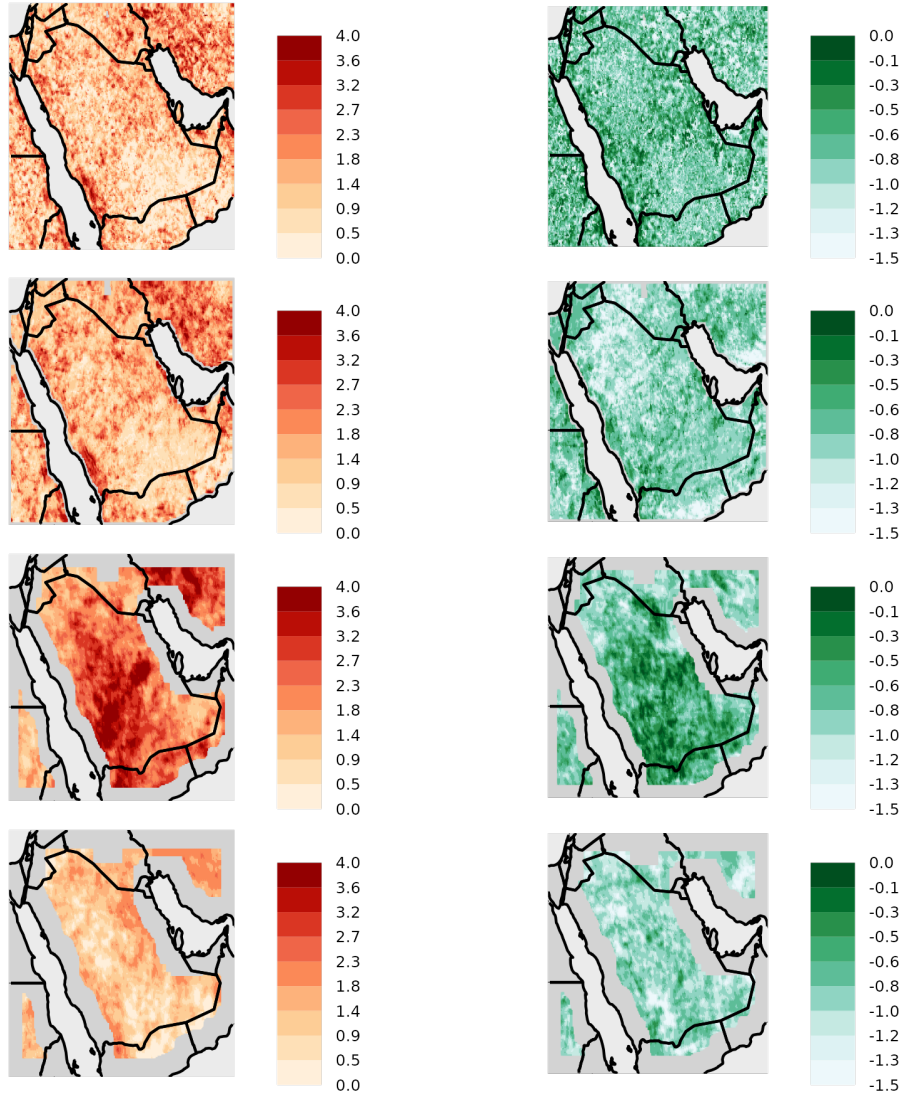


Figure A8: Anisotropy parameter estimates for the HW model used in the application, Section 5. The smoothing level G changes with the row (1st, 4; 2nd, 8; 3rd, 24; 4th, 32). The left and right columns illustrate estimates of α and ω , respectively.

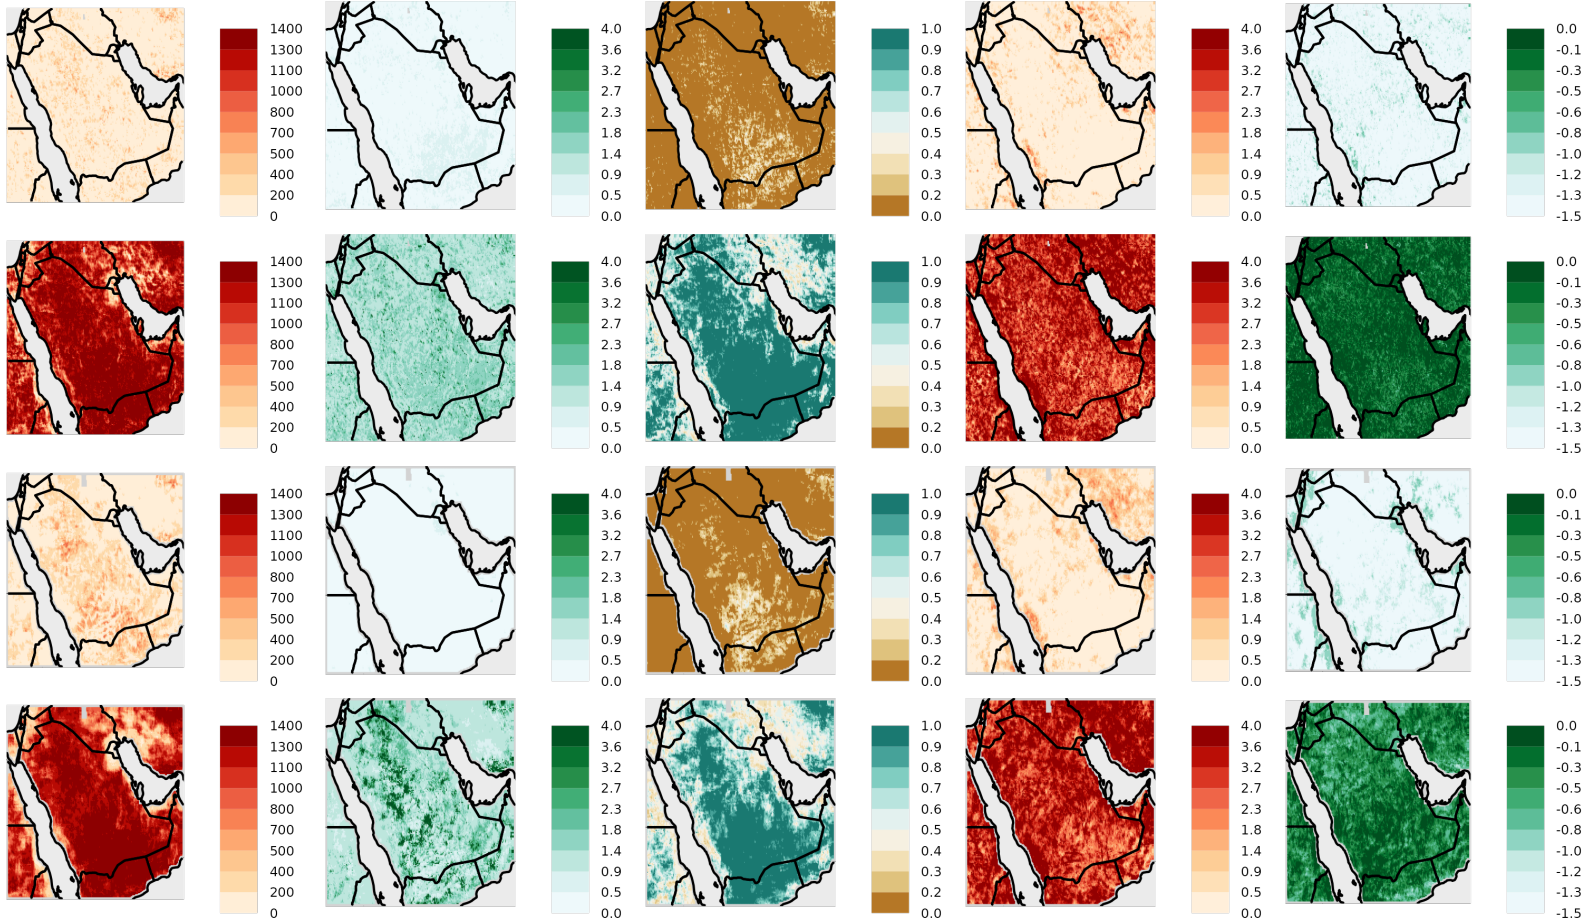


Figure A9: Parameter uncertainty assessment for $G = 4$ (top two rows) and $G = 8$ (last two rows) in the application, Section 5. First and second (third and fourth) rows give estimates of the 2.5% and 97.5% bootstrap quantiles, respectively, of (left to right) λ , κ , δ , α , and ω .

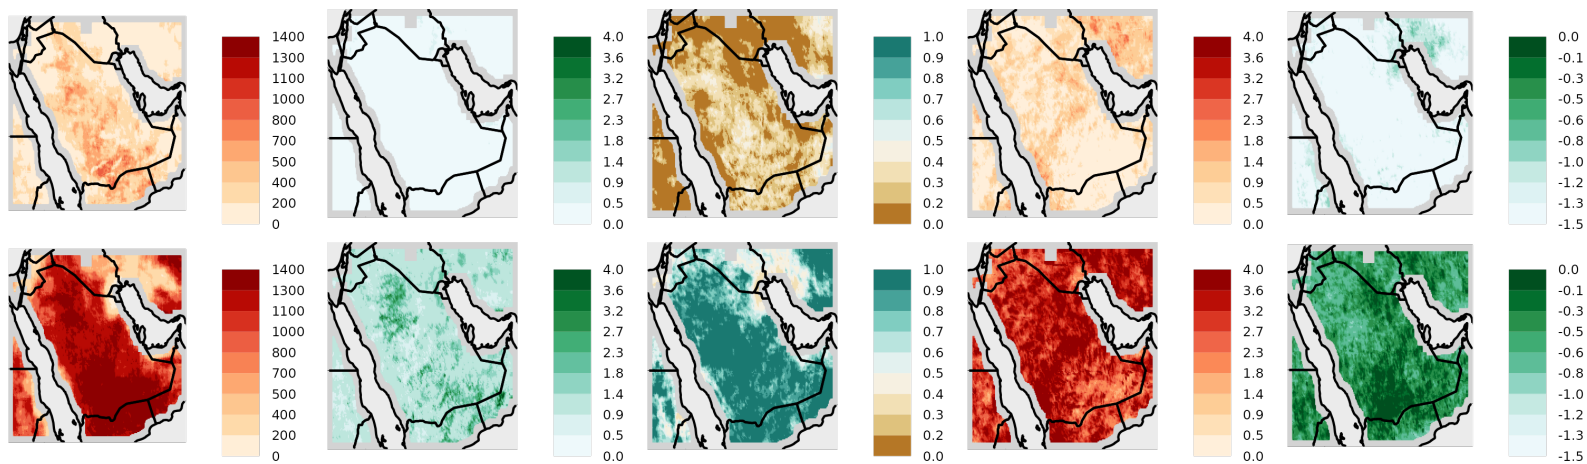


Figure A10: Parameter uncertainty assessment for $G = 16$ in the application, Section 5. Top and bottom rows give estimates of the 2.5% and 97.5% bootstrap quantiles, respectively, of (left to right) λ , κ , δ , α , and ω .

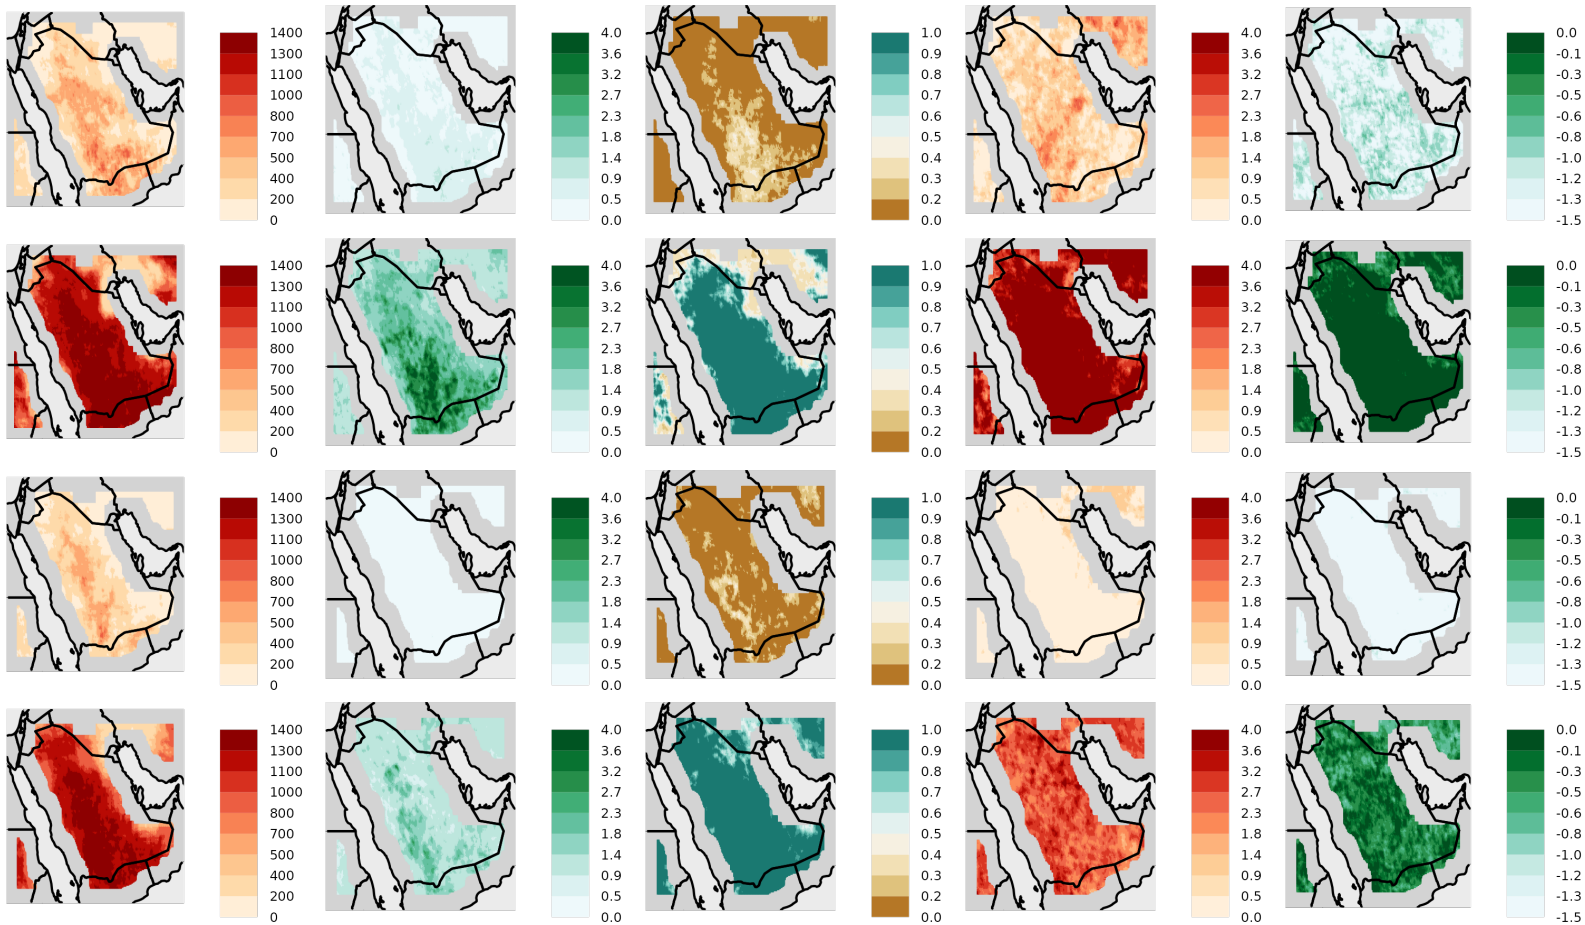


Figure A11: Parameter uncertainty assessment for $G = 24$ (first two rows) and $G = 32$ (last two rows) in the application, Section 5. First and second (third and fourth) rows give estimates of the 2.5% and 97.5% bootstrap quantiles, respectively, of (left to right) λ , κ , δ , α , and ω .

Appendix E. Supplementary tables

Table A1: Results for the simulation study detailed in Section 4.2.1 of the main text. Marginal test risks (s.d.) under the squared-error loss are provided separately for λ ($\times 10^{-1}$) and κ ($\times 10^{-2}$). The lowest value in each column is given in bold.

	λ	κ
GP		
NBE	1.74 (0.09)	0.37 (0.02)
CPL ($h_{max} = 3$)	3.74 (0.24)	0.71 (0.05)
CPL ($h_{max} = \infty$)	14.62 (0.69)	4.84 (0.27)
IMSP		
NBE	1.38 (0.10)	0.078 (0.004)
CPL ($h_{max} = 3$)	4.82 (0.39)	0.160 (0.008)
CPL ($h_{max} = \infty$)	6.13 (0.44)	0.758 (0.036)
MSP		
NBE	1.18 (0.10)	0.055 (0.003)
CPL ($h_{max} = 3$)	2.51 (0.18)	0.080 (0.003)
CPL ($h_{max} = \infty$)	3.75 (0.25)	0.686 (0.035)

Table A2: Results for the r -Pareto process simulation study detailed in Section 4.2.2 of the main text. Marginal test risks (s.d.) under the absolute-error loss are provided separately for λ ($\times 10^{-1}$) and κ ($\times 10^{-2}$). The lowest value in each column is given in bold.

	λ	κ
NBE	2.29 (0.07)	1.87 (0.05)
Quasi-Monte Carlo censored likelihood	4.92 (0.20)	4.31 (0.23)

Table A3: Results for the r -Pareto process simulation study detailed in Section 4.2.2 of the main text. Marginal test risks (s.d.) under the squared-error loss are provided separately for λ ($\times 10^{-1}$) and κ ($\times 10^{-2}$). The lowest value in each column is given in bold.

	λ	κ
NBE	0.99 (0.06)	0.060 (0.003)
Quasi-Monte Carlo censored likelihood	6.30 (0.64)	0.731 (0.101)

Table A4: Results for the HW process simulation study detailed in Section 4.2.2 of the main text; the process is observed over either \mathcal{S} , a 16×16 regular grid, or \mathcal{S}^* , a 6×6 regular grid. Marginal test risks (s.d.) under the *absolute-error* loss are provided separately for λ ($\times 10^{-1}$), κ ($\times 10^{-2}$) and δ ($\times 10^{-2}$). The lowest value in each column is given in bold.

	λ	κ	δ
NBE (\mathcal{S})	5.77 (0.19)	7.42 (0.31)	3.94 (0.17)
NBE (\mathcal{S}^*)	6.30 (0.19)	9.49 (0.26)	3.24 (0.11)
Quasi-Monte Carlo censored likelihood	5.82 (0.18)	9.53 (0.36)	4.80 (0.30)

Table A5: Results for the HW process simulation study detailed in Section 4.2.2 of the main text; the process is observed over either \mathcal{S} , a 16×16 regular grid, or \mathcal{S}^* , a 6×6 regular grid. Marginal test risks (s.d.) under the *squared-error* loss are provided separately for λ ($\times 10^{-1}$), κ ($\times 10^{-2}$) and δ ($\times 10^{-2}$). The lowest value in each column is given in bold.

	λ	κ	δ
NBE (\mathcal{S})	6.87 (0.47)	1.54 (0.25)	0.46 (0.14)
NBE (\mathcal{S}^*)	7.42 (0.50)	1.57 (0.09)	0.23 (0.02)
Quasi-Monte Carlo censored likelihood	6.61 (0.48)	2.19 (0.28)	1.13 (0.24)

Table A6: Neural network architecture for the HW process observed on the smaller domain \mathcal{S}^* , see Section 4.2.2 of the main text. Each convolution filter uses zero padding and unit stride. The function $\text{vec}(\cdot)$ refers to a flattening of an array to a vector. All layers, except the final layer, used rectified linear unit (ReLU) activation functions; the final layer uses the identity function.

layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[6, 6, 2]	[4, 4, 128]	3×3	2,432
2D convolution	[4, 4, 128]	[2, 2, 256]	3×3	147,584
2D convolution	[2, 2, 256]	[1, 1, 256]	2×2	262,400
$\text{vec}(\cdot)$	[1, 1, 256]	[256]		0
dense	[256]	[500]		128,500
dense	[500]	[3]		1503
total trainable parameters:				542,419

Table A7: Results for the variable τ simulation study detailed in Section 4.3. Marginal test risks (s.d.) under the squared-error loss are provided separately for λ ($\times 10^{-1}$), κ ($\times 10^{-2}$) and δ ($\times 10^{-2}$). The lowest value in each column is given in bold.

τ	IMSP		HW		
	λ	κ	λ	κ	δ
random	1.40 (0.07)	0.089 (0.004)	1.14 (0.06)	0.070 (0.004)	0.16 (0.01)
fixed	1.43 (0.07)	0.068 (0.003)	1.18 (0.05)	0.092 (0.004)	0.21 (0.01)

Table A8: Neural network architectures used for the application, see Section 5 of the main text, when $G = 4$. Each convolution filter uses zero padding and unit stride. The function $\text{vec}(\cdot)$ refers to a flattening of an array to a vector. All layers, except the final layer, used rectified linear unit (ReLU) activation functions; the final layer uses the identity function.

$G = 4$				
layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[4, 4, 2]	[3, 3, 64]	2×2	576
2D convolution	[3, 3, 64]	[2, 2, 128]	2×2	32,896
2D convolution	[2, 2, 128]	[1, 1, 256]	2×2	131,328
$\text{vec}(\cdot)$	[1, 1, 256]	[256]		0
dense	[256]	[500]		128,500
dense	[500]	[5]		2505
total trainable parameters:				295,805

Table A9: Neural network architectures used for the application, see Section 5 of the main text, when $G = 8$. Each convolution filter uses zero padding and unit stride. The function $\text{vec}(\cdot)$ refers to a flattening of an array to a vector. All layers, except the final layer, used rectified linear unit (ReLU) activation functions; the final layer uses the identity function.

$G = 8$				
layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[8, 8, 2]	[4, 4, 64]	5×5	3,264
2D convolution	[4, 4, 64]	[2, 2, 128]	3×3	73,856
2D convolution	[2, 2, 128]	[1, 1, 256]	2×2	131,328
$\text{vec}(\cdot)$	[1, 1, 256]	[256]		0
dense	[256]	[500]		128,500
dense	[500]	[5]		2505
total trainable parameters:				339,453

Table A10: Neural network architectures used for the application, see Section 5 of the main text, when $G = 16$. Each convolution filter uses zero padding and unit stride. The function $\text{vec}(\cdot)$ refers to a flattening of an array to a vector. All layers, except the final layer, used rectified linear unit (ReLU) activation functions; the final layer uses the identity function. Note that the layer denoted by the asterisk (*) features in Table A11.

$G = 16$				
layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[16, 16, 2]	[7, 7, 64]	10×10	12,864
2D convolution*	[7, 7, 64]	[3, 3, 128]	5×5	204,928
2D convolution	[3, 3, 128]	[1, 1, 256]	3×3	295,168
$\text{vec}(\cdot)$	[1, 1, 256]	[256]		0
dense	[256]	[500]		128,500
dense	[500]	[5]		2505
total trainable parameters:				643,965

Table A11: Neural network architectures used for the application, see Section 5 of the main text, when $G = 24$ and $G = 32$. Each convolution filter uses zero padding and unit stride. The function $\text{vec}(\cdot)$ refers to a flattening of an array to a vector. All layers, except the final layer, used rectified linear unit (ReLU) activation functions; the final layer uses the identity function.

$G = 24$				
layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[24, 24, 2]	[16, 16, 32]	9×9	5,216
2D convolution	[16, 16, 32]	[7, 7, 64]	10×10	204,864
As in * and subsequent layers of Table A10.				
total trainable parameters:				841,181
$G = 32$				
layer type	input dimension	output dimension	filter dimension	parameters
2D convolution	[32, 32, 2]	[16, 16, 32]	17×17	18,528
2D convolution	[16, 16, 32]	[7, 7, 64]	10×10	204,864
As in * and subsequent layers of Table A10.				
total trainable parameters:				854,493

References

- M. Ahmed, V. Maume-Deschamps, and P. Ribereau. Recognizing a spatial extreme dependence structure: A deep learning approach. *Environmetrics*, 33(4):e2714, 2022.
- B. Alharbi, A. Maghrabi, and N. Tapper. The March 2009 dust event in Saudi Arabia: Precursor and supportive environment. *Bulletin of the American Meteorological Society*, 94(4):515–528, 2013.
- P. Asadi, A. C. Davison, and S. Engelke. Extremes on river networks. *The Annals of Applied Statistics*, 9(4):2023 – 2050, 2015.
- E. Aune, D. P. Simpson, and J. Eidsvik. Parameter estimation in high dimensional Gaussian distributions. *Statistics and Computing*, 24:247–263, 2014.
- L. R. Belzile, C. Dutang, P. J. Northrop, and T. Opitz. A modeler’s guide to extreme value software. *Extremes*, 26(4):595–638, 2023.
- M. Bevilacqua, C. Gaetan, J. Mateu, and E. Porcu. Estimating space and space-time covariance functions for large data sets: a weighted composite likelihood approach. *Journal of the American Statistical Association*, 107(497):268–280, 2012.
- S. Béwentaoré and D. Barro. Space-time trend detection and dependence modeling in extreme event approaches by functional peaks-over-thresholds: application to precipitation in Burkina Faso. *International Journal of Mathematics and Mathematical Sciences*, 2022, 2022.
- Y. Boulaguiem, J. Zscheischler, E. Vignotto, K. van der Wiel, and S. Engelke. Modeling and simulating spatial extremes by combining extreme value theory with generative adversarial networks. *Environmental Data Science*, 1:e5, 2022.
- B. M. Brown and S. I. Resnick. Extreme values of independent stochastic processes. *Journal of Applied Probability*, 14(4):732–739, 1977.
- T. A. Buishand, L. de Haan, and C. Zhou. On spatial extremes: With application to a rainfall problem. *The Annals of Applied Statistics*, 2(2):624 – 642, 2008.
- A. J. Cannon. A flexible nonlinear modelling framework for nonstationary generalized extreme value analysis in hydroclimatology. *Hydrological Processes: An International Journal*, 24(6):673–685, 2010.
- D. Castro-Camilo and R. Huser. Local likelihood estimation of complex tail dependence structures, applied to U.S. precipitation extremes. *Journal of the American Statistical Association*, 115(531):1037–1054, 2020.
- S. Castruccio, R. Huser, and M. G. Genton. High-order composite likelihood inference for max-stable distributions and processes. *Journal of Computational and Graphical Statistics*, 25(4):1212–1229, 2016.

- J. Chan, V. Perrone, J. Spence, P. Jenkins, S. Mathieson, and Y. Song. A likelihood-free inference framework for population genetic data using exchangeable neural networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- D. Cisneros, J. Richards, A. Dahal, L. Lombardo, and R. Huser. Deep graphical regression for jointly moderate and extreme Australian wildfires. *Spatial Statistics*, page 100811, 2024.
- S. Coles. *An Introduction to Statistical Modeling of Extreme Values*, volume 208. Springer, 2001.
- K. Cremanns and D. Roos. Deep Gaussian covariance network. *arXiv preprint arXiv:1710.06202*, 2017.
- A. C. Davison and M. M. Gholamrezaee. Geostatistics of extremes. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 468(2138):581–608, 2012.
- A. C. Davison and R. Huser. Statistics of extremes. *Annual Review of Statistics and its Application*, 2:203–235, 2015.
- A. C. Davison, S. A. Padoan, and M. Ribatet. Statistical modeling of spatial extremes. *Statistical Science*, 27(2):161–186, 2012.
- A. C. Davison, R. Huser, and E. Thibaud. Geostatistics of dependent and asymptotically independent extremes. *Mathematical Geosciences*, 45:511–529, 2013.
- A. C. Davison, R. Huser, and E. Thibaud. Spatial extremes. In A. E. Gelfand, M. Fuentes, J. A. Hoeting, and R. L. Smith, editors, *Handbook of Environmental and Ecological Statistics*, pages 711–744. CRC Press, 2019.
- R. de Fondeville and L. Belzile. *mvPot: Multivariate Peaks-over-Threshold Modelling for Spatial Extreme Events*, 2022. URL <https://CRAN.R-project.org/package=mvPot>. R package version 0.1-5.
- R. de Fondeville and A. C. Davison. High-dimensional peaks-over-threshold inference. *Biometrika*, 105(3):575–592, 2018.
- R. de Fondeville and A. C. Davison. Functional peaks-over-threshold analysis. *Journal of the Royal Statistical Society: Series B*, 4(84):1392–1422, 2022.
- L. de Haan. A spectral representation for max-stable processes. *The Annals of Probability*, 12(4):1194–1204, 1984.
- L. de Haan and A. Ferreira. *Extreme Value Theory: An Introduction*, volume 3. Springer, 2006.
- C. Dombry and M. Ribatet. Functional regular variations, Pareto processes and peaks over threshold. *Statistics and Its Interface*, 8(1):9–17, 2015.
- C. Dombry, S. Engelke, and M. Oesting. Bayesian inference for multivariate extreme value distributions. *Electronic Journal of Statistics*, 11(2):4813–4844, 2017.

- S. Engelke, Z. Kabluchko, and M. Schlather. An equivalent representation of the Brown–Resnick process. *Statistics & probability letters*, 81(8):1150–1154, 2011.
- A. Ferreira and L. de Haan. The generalized Pareto process; with a view towards application and simulation. *Bernoulli*, 20(4):1717–1737, 2014.
- A. H. Galib, A. McDonald, T. Wilson, L. Luo, and P.-N. Tan. DeepExtrema: A deep learning approach for forecasting block maxima in time series data. *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)*, pages 2980–2986, 2022.
- A. E. Gelfand and E. M. Schliep. Spatial statistics and Gaussian processes: A beautiful marriage. *Spatial Statistics*, 18:86–104, 2016.
- M. G. Genton, Y. Ma, and H. Sang. On the likelihood function of Gaussian max-stable processes. *Biometrika*, pages 481–488, 2011.
- A. Genz and F. Bretz. *Computation of Multivariate Normal and t Probabilities*, volume 195. Springer Science & Business Media, 2009.
- F. Gerber and D. Nychka. Fast covariance parameter estimation of spatial Gaussian process models using neural networks. *Stat*, 10(1):e382, 2021.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT press, 2016.
- J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, et al. Recent advances in convolutional neural networks. *Pattern Recognition*, 77: 354–377, 2018.
- K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, 2015.
- D. Healy, A. Parnell, P. Thorne, and J. Tawn. Inference for extreme spatial temperature events in a changing climate with application to Ireland. *Journal of the Royal Statistical Society: Series C*, 2024. To appear.
- E. C. Hector and B. J. Reich. Distributed inference for spatial extremes modeling in high dimensions. *Journal of the American Statistical Association*, 119(546):1297–1308, 2024.
- R. Huser and A. C. Davison. Composite likelihood estimation for the Brown–Resnick process. *Biometrika*, 100(2):511–518, 2013.
- R. Huser and A. C. Davison. Space–time modelling of extreme events. *Journal of the Royal Statistical Society: Series B*, 76(2):439–461, 2014.
- R. Huser and M. G. Genton. Non-stationary dependence structures for spatial extremes. *Journal of Agricultural, Biological, and Environmental Statistics*, 21:470–491, 2016.
- R. Huser and J. L. Wadsworth. Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 114(525):434–444, 2019.

- R. Huser and J. L. Wadsworth. Advances in statistical modeling of spatial extremes. *Wiley Interdisciplinary Reviews: Computational Statistics*, 14(1):e1537, 2022.
- R. Huser, A. C. Davison, and M. G. Genton. Likelihood estimators for multivariate extremes. *Extremes*, 19:79–103, 2016.
- R. Huser, T. Opitz, and E. Thibaud. Bridging asymptotic independence and dependence in spatial extremes using Gaussian scale mixtures. *Spatial Statistics*, 21:166–186, 2017.
- R. Huser, C. Dombry, M. Ribatet, and M. G. Genton. Full likelihood inference for max-stable data. *Stat*, 8(1):e218, 2019.
- R. Huser, M. L. Stein, and P. Zhong. Vecchia likelihood approximation for accurate and fast inference with intractable spatial max-stable models. *Journal of Computational and Graphical Statistics*, 33:978–990, 2024.
- H. Joe. *Multivariate Models and Multivariate Dependence Concepts*. CRC press, 1997.
- Z. Kabluchko, M. Schlather, and L. de Haan. Stationary max-stable fields associated to negative definite functions. *The Annals of Probability*, 37(5):2042–2065, 2009.
- M. Khodeir, M. Shamy, M. Alghamdi, M. Zhong, H. Sun, M. Costa, L.-C. Chen, and P. Maciejczyk. Source apportionment and elemental composition of PM2.5 and PM10 in Jeddah City, Saudi Arabia. *Atmospheric Pollution Research*, 3(3):331–340, 2012.
- J. P. Klein and M. L. Moeschberger. *Survival Analysis: Techniques for Censored and Truncated Data*, volume 1230. Springer, 2003.
- P. Krupskii, R. Huser, and M. G. Genton. Factor copula models for replicated spatial data. *Journal of the American Statistical Association*, 113(521):467–479, 2018.
- E. L. Lehmann and G. Casella. *Theory of Point Estimation*. Springer, 2nd edition, 1998.
- J. Lelieveld, K. Klingmüller, A. Pozzer, U. Pöschl, M. Fnais, A. Daiber, and T. Münzel. Cardiovascular disease burden from ambient air pollution in Europe reassessed using novel hazard ratio functions. *European Heart Journal*, 40(20):1590–1596, 2019.
- A. Lenzi, J. Bessac, J. Rudi, and M. L. Stein. Neural networks for parameter estimation in intractable models. *Computational Statistics & Data Analysis*, 185:107762, 2023.
- R. Majumder and B. J. Reich. A deep learning synthetic likelihood approximation of a non-stationary spatial model for extreme streamflow forecasting. *Spatial Statistics*, 55:100755, 2023.
- R. Majumder, B. J. Reich, and B. A. Shaby. Modeling extremal streamflow using deep learning approximations and a flexible spatial process. *The Annals of Applied Statistics*, 18(2):1519–1542, 2024.
- P. McCullagh. What is a statistical model? *The Annals of Statistics*, 30(5):1225–1310, 2002.

- A. McDonald, P.-N. Tan, and L. Luo. COMET flows: Towards generative modeling of multivariate extremes and tail dependence. *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)*, pages 3328–3334, 2022.
- S. Munir, T. M. Habeebullah, A. M. Mohammed, E. A. Morsy, M. Rehan, K. Ali, et al. Analysing PM2.5 and its association with PM10 and meteorology in the arid climate of Makkah, Saudi Arabia. *Aerosol and Air Quality Research*, 17(2):453–464, 2017.
- T. Opitz. Extremal t processes: Elliptical domain of attraction and a spectral representation. *Journal of Multivariate Analysis*, 122:409–413, 2013.
- T. Opitz. Modeling asymptotically independent spatial extremes based on Laplace random fields. *Spatial Statistics*, 16:1–18, 2016.
- S. A. Padoan, M. Ribatet, and S. A. Sisson. Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, 105(489):263–277, 2010.
- D. Pagendam, S. Janardhanan, J. Dabrowski, and D. MacKinlay. A log-additive neural model for spatio-temporal prediction of groundwater levels. *Spatial Statistics*, 55:100740, 2023.
- O. C. Pasche and S. Engelke. Neural networks for extreme quantile regression with an application to forecasting of flood risk. *The Annals of Applied Statistics*, 18(4):2818–2839, 2024.
- G. Polichetti, S. Cocco, A. Spinali, V. Trimarco, and A. Nunziata. Effects of particulate matter (PM10, PM2.5 and PM1) on the cardiovascular system. *Toxicology*, 261(1-2):1–8, 2009.
- S. T. Radev, M. Schmitt, V. Pratz, U. Picchini, U. Köthe, and P.-C. Bürkner. JANA: Jointly amortized neural approximation of complex Bayesian models. In *Uncertainty in Artificial Intelligence*, pages 1695–1706. PMLR, 2023.
- S. Rai, A. Hoffman, S. Lahiri, D. W. Nychka, S. R. Sain, and S. Bandyopadhyay. Fast parameter estimation of generalized extreme value distribution using neural networks. *Environmetrics*, 35(3):e2845, 2024.
- M. Ribatet. *SpatialExtremes: modelling spatial extremes*, 2022. URL <https://cran.r-project.org/package=SpatialExtremes>. R package version 2.1-0.
- J. Richards and R. Huser. Regression modelling of spatiotemporal extreme U.S. wildfires via partially-interpretable neural networks. *arXiv preprint arXiv:2208.07581*, 2022.
- J. Richards and J. L. Wadsworth. Spatial deformation for nonstationary extremal dependence. *Environmetrics*, 32(5):e2671, 2021.
- J. Richards, J. A. Tawn, and S. Brown. Modelling extremes of spatial aggregates of precipitation using conditional methods. *The Annals of Applied Statistics*, 16(4):2693–2713, 2022.

- J. Richards, R. Huser, E. Bevacqua, and J. Zscheischler. Insights into the drivers and spatiotemporal trends of extreme mediterranean wildfires with statistical deep learning. *Artificial Intelligence for the Earth Systems*, 2(4):e220095, 2023a.
- J. Richards, J. A. Tawn, and S. Brown. Joint estimation of extreme spatially aggregated precipitation at different scales through mixture modelling. *Spatial Statistics*, 53:100725, 2023b.
- M. Sainsbury-Dale. *NeuralEstimators: Likelihood-Free Parameter Estimation using Neural Networks*, 2024. URL <https://CRAN.R-project.org/package=NeuralEstimators>. R package version 0.1-2.
- M. Sainsbury-Dale, A. Zammit-Mangion, and R. Huser. Likelihood-free parameter estimation with neural Bayes estimators. *The American Statistician*, 78(1):1–14, 2024a.
- M. Sainsbury-Dale, A. Zammit-Mangion, J. Richards, and R. Huser. Neural Bayes estimators for irregular spatial data using graph neural networks. *Journal of Computational and Graphical Statistics*, 2024b. To appear.
- H. Sang and A. E. Gelfand. Continuous spatial process models for spatial extreme values. *Journal of Agricultural, Biological, and Environmental Statistics*, 15:49–65, 2010.
- H. Sang and M. G. Genton. Tapered composite likelihood for spatial max-stable models. *Spatial Statistics*, 8:86–103, 2014.
- M. Schlather. Models for stationary max-stable random fields. *Extremes*, 5:33–44, 2002.
- M. Schmitt, P.-C. Bürkner, U. Köthe, and S. T. Radev. Detecting model misspecification in amortized Bayesian inference with neural networks. In *DAGM German Conference on Pattern Recognition*, pages 541–557. Springer, 2023.
- X. Shao, A. Hazra, J. Richards, and R. Huser. Flexible modeling of nonstationary extremal dependence using spatially fused LASSO and Ridge penalties. *Technometrics*, 2024. To appear.
- R. L. Smith. Max-stable processes and spatial extremes. *Unpublished manuscript*, 205:1–32, 1990.
- A. Stephenson and J. Tawn. Exploiting occurrence times in likelihood inference for componentwise maxima. *Biometrika*, 92(1):213–227, 2005.
- S. Tariq, A. Mariam, U. Mehmood, et al. Spatial and temporal variations in PM2.5 and associated health risk assessment in Saudi Arabia using remote sensing. *Chemosphere*, 308:136296, 2022.
- E. Thibaud and T. Opitz. Efficient inference and simulation for elliptical Pareto processes. *Biometrika*, 102(4):855–870, 2015.
- E. Thibaud, R. Mutzner, and A. C. Davison. Threshold modeling of extreme spatial rainfall. *Water Resources Research*, 49(8):4633–4644, 2013.

- A. Van Donkelaar, M. S. Hammer, L. Bindle, M. Brauer, J. R. Brook, M. J. Garay, N. C. Hsu, O. V. Kalashnikova, R. A. Kahn, C. Lee, et al. Monthly global estimates of fine particulate matter and their uncertainty. *Environmental Science & Technology*, 55(22):15287–15300, 2021.
- C. Varin, N. Reid, and D. Firth. An overview of composite likelihood methods. *Statistica Sinica*, 21(1):5–42, 2011.
- L. Vasiliades, P. Galiatsatou, and A. Loukas. Nonstationary frequency analysis of annual maximum rainfall using climate covariates. *Water Resources Management*, 29:339–358, 2015.
- J. L. Wadsworth. On the occurrence times of componentwise maxima and bias in likelihood inference for multivariate max-stable distributions. *Biometrika*, 102(3):705–711, 2015.
- J. L. Wadsworth. SpatialADAI: R package. <https://www.lancaster.ac.uk/~wadswoj1/SpatialADAI>, 2018. Accessed: 2024-18-12.
- J. L. Wadsworth and J. A. Tawn. Dependence modelling for spatial extremes. *Biometrika*, 99(2):253–272, 2012.
- J. L. Wadsworth and J. A. Tawn. Efficient inference for spatial extreme value processes associated to log-Gaussian random functions. *Biometrika*, 101(1):1–15, 2014.
- J. L. Wadsworth and J. A. Tawn. Higher-dimensional spatial extremes via single-site conditioning. *Spatial Statistics*, 51:100677, 2022.
- J. L. Wadsworth, J. A. Tawn, A. C. Davison, and D. M. Elton. Modelling across extremal dependence classes. *Journal of the Royal Statistical Society: Series B*, 79(1):149–175, 2017.
- J. Walchessen, A. Lenzi, and M. Kuusela. Neural likelihood surfaces for spatial processes with computationally intensive or intractable likelihoods. *Spatial Statistics*, 62:100848, 2024.
- Z. Wang, J. Hasenauer, and Y. Schälte. Missing data in amortized simulation-based neural posterior estimation. *PLOS Computational Biology*, 20(6):e1012184, 2024.
- T. Wilson, P.-N. Tan, and L. Luo. DeepGPD: A deep learning approach for modeling geospatio-temporal extreme events. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(4):4245–4253, 2022.
- Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2020.
- Y.-F. Xing, Y.-H. Xu, M.-H. Shi, and Y.-X. Lian. The impact of PM_{2.5} on the human respiratory system. *Journal of Thoracic Disease*, 8(1):E69, 2016.
- M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. R. Salakhutdinov, and A. J. Smola. Deep sets. *Advances in Neural Information Processing Systems*, 30, 2017.

- A. Zammit-Mangion and J. Rougier. Multi-scale process modelling and distributed computation for spatial data. *Statistics and Computing*, 30(6):1609–1627, 2020.
- A. Zammit-Mangion and C. K. Wikle. Deep integro-difference equation models for spatio-temporal forecasting. *Spatial Statistics*, 37:100408, 2020.
- A. Zammit-Mangion, M. Sainsbury-Dale, and R. Huser. Neural methods for amortized inference. *Annual Review of Statistics of Its Application*, 12, 2025. To appear.
- L. Zhang, B. A. Shaby, and J. L. Wadsworth. Hierarchical transformed scale mixtures for flexible modeling of spatial extremes on datasets with many locations. *Journal of the American Statistical Association*, 117(539):1357–1369, 2022.
- L. Zhang, X. Ma, C. K. Wikle, and R. Huser. Flexible and efficient spatial extremes emulation via variational autoencoders. arXiv preprint arXiv:2307.11390, 2023.