ESTIMATION OF SPATIO-TEMPORAL EXTREMES VIA GENERATIVE NEURAL NETWORKS

Christopher Bülte[*]
Chair of Statistical Methods and Econometrics
Karlsruhe Institute of Technology (KIT)
Karlsruhe, Germany
`buelte@math.lmu.de`

Lisa Leimenstoll
Chair of Statistical Methods and Econometrics
Karlsruhe Institute of Technology (KIT)
Karlsruhe, Germany
`lisa.leimenstoll@kit.edu`

Melanie Schienle
Chair of Statistical Methods and Econometrics
Karlsruhe Institute of Technology (KIT)
Karlsruhe, Germany
`melanie.schienle@kit.edu`

July 12, 2024

ABSTRACT

Recent methods in modeling spatial extreme events have focused on utilizing parametric max-stable processes and their underlying dependence structure. In this work, we provide a unified approach for analyzing spatial extremes with little available data by estimating the distribution of model parameters or the spatial dependence directly. By employing recent developments in generative neural networks we predict a full sample-based distribution, allowing for direct assessment of uncertainty regarding model parameters or other parameter dependent functionals. We validate our method by fitting several simulated max-stable processes, showing a high accuracy of the approach, regarding parameter estimation, as well as uncertainty quantification. Additional robustness checks highlight the generalization and extrapolation capabilities of the model, while an application to precipitation extremes across Western Germany demonstrates the usability of our approach in real-world scenarios.

1 Introduction

As the frequency of extreme weather events rises, it becomes increasingly crucial to understand and detect them at the earliest opportunity. Statistical models provide a way to enhance their interpretability and offer insights into the connections between extreme events. Since geophysical data is often coupled across both space and time this poses challenges for modeling, often leading to highly complex statistical models. For spatial data, such as precipitation, a common way to describe and analyze extremes are max-stable processes, which arise as the unique limit of pointwise maxima of random fields. These processes are an essential tool in analyzing spatial extremes (Davison et al., 2012), as they allow for flexible modeling of the underlying dependence structure. However, when it comes to modeling these extremes, usually only a few observations are available, even less so as the underlying process is usually changing across time. For that reason traditional statistical methods often fail to identify parameters correctly, particularly as these models are high dimensional and complex. Furthermore, estimating parameters becomes especially challenging when dealing with extreme values. Therefore, specifying a distribution rather than relying on point estimators can be beneficial for quantifying uncertainty. While there was a recent focus on new methods for parameter estimation, mainly

---

[*]Current affiliation: Ludwig-Maximilians-Universität, Munich

employing neural networks, these typically fall short of providing adequate uncertainty estimates, as they are usually based on bootstrap sampling.

In this work, we propose an estimation approach, based on generative neural networks, that allows for estimation and uncertainty quantification not only for the model parameters, but also for the spatial dependence across extremes. This is achieved by training a neural network on simulated max-stable processes and predicting a full (sample-based) parameter distribution of the corresponding max-stable model. This has the advantage that any functional of interest, e.g. mean or confidence intervals can be obtained easily from the predicted samples. Our work combines recent efforts in neural networks for point estimation (Lenzi et al., 2023; Sainsbury-Dale et al., 2022) with advancements in training generative neural networks with proper scoring rules (Pacchiardi and Dutta, 2022; Chen et al., 2022). Using these methods, we offer a complete approach to estimate the full parameter distribution, allowing for deterministic and probabilistic evaluation of parameters or corresponding functionals of the underlying model simultaneously. Furthermore, we extend this approach to nonparametrically estimate the distribution of the pairwise extremal coefficient function, which is a commonly applied tool to characterize the extremal dependence between two locations of a max-stable process. This unified approach offers a way to identify parameters in various max-stable models and to assess uncertainty in spatial dependence estimations, even if the underlying model is unknown, which is often the case with real data. Comparison to established methods, such as the ABC method, highlights the advantages of our approaches, especially considering uncertainty quantification. Furthermore, we evaluate the model predictions against several scenarios of misspecification, which naturally occur if working with real data.

From a statistical viewpoint it is natural to provide a maximum likelihood estimation for the parameters of a statistical model. However, in the case of max-stable processes even in moderate dimensions, the likelihood contains too many terms to be evaluated in reasonable time. Improvements can be made by considering occurrence times of maxima (Stephenson and Tawn, 2005; Huser et al., 2019), leading to a computable likelihood function but a highly biased estimator (Huser et al., 2016). A more common approach proposed by Padoan et al. (2010) is to replace the full likelihood by the approximate pairwise likelihood, which however comes with a loss of statistical efficiency (Huser and Davison, 2013; Castruccio et al., 2016). Other methods include using the related Vecchia approximation (Huser et al., 2022) or an expectation-maximization algorithm (Huser et al., 2019), but parameter estimation in max-stable processes remains an ongoing field of research in extreme value theory.

Different approaches have been proposed that avoid evaluating the likelihood function altogether, often referred to as likelihood-free methods. These are most often based on the assumption that it is possible to simulate from the given model and use these simulations to produce an estimate of the true parameters. The most popular of these methods is the approximate Bayesian computation (ABC) framework (Beaumont et al., 2002; Franks, 2020). By sampling parameters from some prior distribution, simulating from the model and minimizing some suitable summary statistics between the simulations and the observations, the method retrieves a posterior parameter distribution. However, while summary statistics have been developed for special purposes, such as max-stable processes (Erhardt and Smith, 2012; Fearnhead and Prangle, 2012), the choice is not obvious and requires calibration. In addition, the approach requires a large amount of simulations in order to produce a good estimate, making it infeasible.

Quite recently, neural networks and deep learning approaches in general have gained increasing popularity in likelihood-free inference. Creel (2017) train a neural network, similar to the ABC approach, on a highly informative summary statistics and apply their method to two econometric models. In a very similar approach, Rai et al. (2023) use a summary statistic, based on extreme quantiles to estimate parameters of an extreme value distribution. Their results suggest similar accuracy, as compared to classical ML-estimation, but an increase in computational speed. Considering spatial data, Gerber and Nychka (2021) use a neural network to perform local covariance estimation for spatial Gaussian processes, while Lenzi et al. (2023) estimate the parameters of a max-stable process directly from the observations by utilizing a convolutional neural network. Sainsbury-Dale et al. (2022) propose a so-called neural Bayes estimator that minimizes the Bayes risk of an estimator to estimate parameters of different spatial models, including max-stable processes. The main advantage of such approaches is that they are likelihood-free, work with very small sample sizes and can be much faster than classical methods.

In comparison to most of the previous methods, our approach has the advantage that it gives access to the full parameter distribution, thus requiring no additional steps for assessing the parameter uncertainty. Furthermore, we extend the estimation procedure to directly estimating the extremal coefficient function, which is an important tool to analyze the dependence structure across the spatial extremes. The remainder of this article is organized as follows. Section 2 outlines the theory regarding max-stable processes, as well as theoretical background on our approach. Section 3 entails the specific implementation of the approach, evaluation metrics and extensive results of simulation studies, as well as additional robustness checks. Section 4 illustrates how the approach can be applied to a real data scenario, by estimating the dependence structure of precipitation extremes across Western Germany, while Section 5 concludes and future research is outlined.

2 Methodology

As mentioned in the introduction, we focus on modelling spatial extremes using max-stable processes, which arise as the unique non-degenerate limit of renormalized pointwise block maxima of i.i.d. random fields and are therefore a widely used tool for analyzing spatial extremes (Davison et al., 2012; Davison and Huser, 2015). Section 2.1 contains the necessary theoretical background on max-stable processes its characteristics and measures of spatial dependence, while Section 2.2 introduces our proposed modeling and estimation framework of spatial extremes.

2.1 Setup

In this article, we consider the following definition of a max-stable process, due to Schlather (2002). Let $\mathcal{X} \subseteq \mathbb{R}^d$, $\{\xi_i, \ i \in \mathbb{N}\}$ denote the points of a nonnegative Poisson process on $(0, \infty)$ with intensity measure $d\Lambda(\xi) = \xi^{-2} d\xi$ and let $\boldsymbol{Y}(\cdot)$ denote a nonnegative stochastic process defined on $\mathbb{R}^d$ such that $\mathbb{E}[\boldsymbol{Y}(\boldsymbol{x})] = 1, \ \forall \boldsymbol{x} \in \mathbb{R}^d$. Then

$$\boldsymbol{Z}(\boldsymbol{x}) = \max_{i \geq 1} \xi_i \boldsymbol{Y}_i(\boldsymbol{x}), \quad \boldsymbol{x} \in \mathcal{X} \tag{1}$$

defines a max-stable process with unit Fréchet margins,[2] where $\boldsymbol{Y}_i(\cdot)$ are i.i.d. copies of $\boldsymbol{Y}(\cdot)$. Different suitable choices of $\boldsymbol{Y}(\cdot)$ lead to different max-stable processes. We will focus on two different models: The Brown-Resnick (Kabluchko et al., 2009) and the Schlather (Schlather, 2002) model.

- The Brown-Resnick model arises if $Y_i(\boldsymbol{x}) = \exp\{\epsilon_i(\boldsymbol{x}) - \gamma(\boldsymbol{h})\}$ is chosen in (1), where ϵ_i are independent copies of a centered Gaussian process with (semi-) variogram $\gamma(\boldsymbol{h})$ and spatial separation $\boldsymbol{h}$. A typical choice is $\gamma(\boldsymbol{h}) = (\|\boldsymbol{h}\|/\lambda)^\nu$, with range parameter $\lambda > 0$ and smoothness parameter $\nu \in (0, 2]$. Due to their flexibility, Brown-Resnick models are often applied in practice (compare Thibaud et al., 2016; Oesting et al., 2017).

- The Schlather model also comes from representation (1), with $Y_i(\boldsymbol{x}) = \sqrt{2\pi} \max\{0, \epsilon_i(\boldsymbol{x})\}$, where $\epsilon_i(x)$ are i.i.d. copies of a standard Gaussian process with correlation function $\rho(\boldsymbol{h})$. The correlation function is usually chosen from a choice of valid parametric families. Common examples are the powered exponential correlation function

$$\rho(\boldsymbol{h}) = \exp\left(-(\|\boldsymbol{h}\|/\lambda)^\nu\right), \quad \lambda > 0, \ \nu \in (0, 2],$$

or the Whittle-Matérn correlation function

$$\rho(\boldsymbol{h}) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\boldsymbol{h}}{\lambda}\right)^\nu K_\nu \left(\frac{\boldsymbol{h}}{\lambda}\right), \quad \lambda > 0, \ \nu > 0,$$

where Γ is the gamma function and K_ν is the modified Bessel function of the third kind with order ν. In both cases, λ and ν again denote the range and smoothness parameter respectively. The max-stable processes of the Schlather model are isotropic and stationary. The Schlather model has been applied for example to precipitation maxima (Davison and Gholamrezaee, 2012) or to temperature minima (Erhardt and Smith, 2012).

From Equation 1 one can derive the joint cumulative distribution of $\boldsymbol{Z}(\boldsymbol{x})$ at a finite collection of spatial sites $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k\} \subset \mathcal{X}$ and the corresponding probability density function as

$$f(z_1, \ldots, z_k; \boldsymbol{\gamma}) = \exp\left(-V(z_1, \ldots, z_k)\right) \sum_{\pi \in \mathcal{P}_k} (-1)^{|\pi|} \prod_{j=1}^{|\pi|} V_{\pi_j}(z_1, \ldots, z_k), \tag{2}$$

where $\boldsymbol{\gamma} = (\lambda, \nu)^T$ is the parameter vector, $\mathcal{P}_k$ denotes the set of all partitions $\{\pi_1, \ldots, \pi_k\}$ of the set $\{x_1, \ldots, x_k\}$ and $|\pi| = l$ is the size of the partition π, while V denotes the so called exponent measure (de Haan and Ferreira, 2006) and $V_{\pi_j} = \frac{\partial^{|\pi_j|}}{\partial z_{\pi_j}} V(z_1, \ldots, z_k)$ its partial derivative of with respect to the variables indexed by the set π_j. For reasons of notation, the dependence of the functions V and w on the unknown parameter $\boldsymbol{\gamma}$ is omitted. The number of terms involved in Equation 2 quickly explodes, as it is summed over the set of all possible partitions. Even if V is analytically available, the number of elements $\pi \in \mathcal{P}_k$, called the Bell number, is too large to make the expression computationally tractable. Castruccio et al. (2016) conclude that even for processes where closed form expressions are available, calculating the full likelihood is not possible for $k > 12$.

Even though the full form might not be available, one can still take advantage of the concept of maximum-likelihood estimation and its properties. Typically, one considers the pairwise likelihood (Padoan et al., 2010; Davis et al., 2013), which is defined as

$$\ell_p(\boldsymbol{\gamma}; \boldsymbol{z}) = \sum_{i=1}^{k-1} \sum_{j=i+1}^{k} w_{i,j} \log f(z_i, z_j; \boldsymbol{\gamma}), \tag{3}$$

---

[2]This means that $\mathbb{P}(\boldsymbol{Z}(\boldsymbol{x}) \leq z) = \exp(-1/z), \ z > 0$.

3

where $\boldsymbol{z} = (z_1, \ldots, z_k)$ is a single observation and $f(\cdot, \cdot; \boldsymbol{\theta})$ is the bivariate pdf, obtained from Equation 2. In order to reduce computational complexity and improve statistical efficiency of the estimator, the weights $w_{i,j}$ can be chosen accordingly, for example based on a cutoff distance (Padoan et al., 2010). Since this estimator is unbiased and converges to a normal distribution, confidence intervals and standard errors can be retrieved easily. However, if the sample size is small, the pairwise likelihood can be highly biased.

As described earlier, measuring and analyzing the dependence structure across spatial extremes, is of great importance. A useful quantity for that matter is the so-called pairwise extremal coefficient function, defined by

$$\theta(h) = -z \log \mathbb{P}(\boldsymbol{Z}(\boldsymbol{x}_1) \leq z, \boldsymbol{Z}(\boldsymbol{x}_2) \leq z) = \mathbb{E}[\max\{\boldsymbol{Y}(\boldsymbol{x}_1), \boldsymbol{Y}(\boldsymbol{x}_2)\}], \tag{4}$$

where $Y(\cdot)$ comes from the representation in (1) and $h = \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2$. As can be seen, $\theta(h)$ is directly related to the probability that two spatial sites do not exceed a common threshold z and provides a measure of the dependence between two spatial locations. $\theta(h)$ lies in the range $[1, 2]$, with the lower bound corresponding to complete dependence and the upper bound to independence of the two spatial locations and is analytically available for a wide range of models.

2.2 Estimation framework

This section first outlines our approach for modeling spatial extremes by estimating the parameter distribution of max-stable models using generative neural networks. Afterwards it demonstrates how this approach can be extended to directly estimate the pairwise extremal coefficient function and its corresponding distribution. Let $\boldsymbol{Z}(\cdot)$ denote a

1. Draw parameter from prior distribution

2. Simulate max-stable process $\boldsymbol{Z} \sim \mathbb{P}(\cdot \mid \boldsymbol{\gamma})$

3. Feed through neural network

4. Estimate of posterior parameter distribution
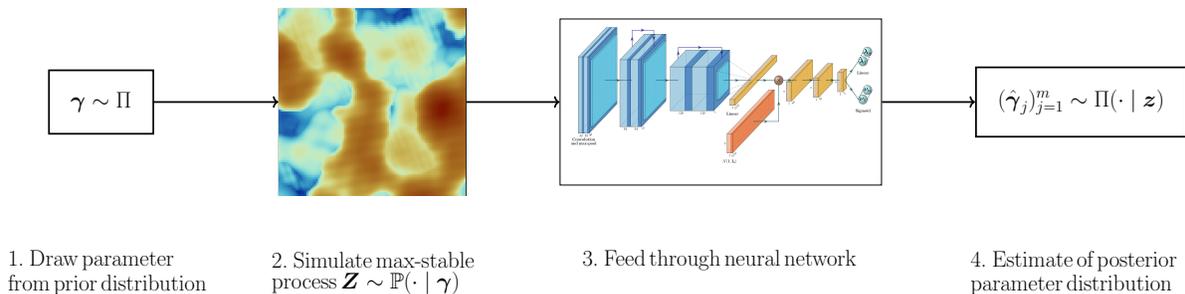
Figure 1: The figure shows the parameter estimation setup, for a two-dimensional max-stable process. In the training phase, all four steps are run multiple times, and the network is trained on the loss between the drawn parameters and the estimated parameters. For inference, only steps 3 and 4 are required.

max-stable process on $\mathcal{X} \subseteq \mathbb{R}^2$ and $\Gamma \subseteq \mathbb{R}_+^2$ denote the corresponding parameter space. The distribution and likelihood function of the model are denoted by $\mathbb{P}(\cdot \mid \boldsymbol{\gamma})$ and $p(\cdot \mid \boldsymbol{\gamma})$, where $\boldsymbol{\gamma} = (\lambda, \nu)^\top \in \Gamma$ is the parameter vector of the max-stable process. Similarly, Π and π denote the prior distribution and density on the parameter space Γ, while $\Pi(\cdot \mid \boldsymbol{z})$ and $\pi(\cdot \mid \boldsymbol{z})$ denote the respective posterior distribution and density given the observation $\boldsymbol{z}$. Classical parameter point estimation aims at estimating $\boldsymbol{\gamma}$ from the data $\boldsymbol{z}$ using a mapping from the sample space to the parameter space $\hat{\gamma} : \mathcal{X} \to \Gamma$. However, as we are interested in information on the full posterior distribution of the parameter $\boldsymbol{\gamma}$, we require a mapping from the sample space to the full posterior distribution, e.g.

$$\mathbb{Q}(\boldsymbol{\gamma} \mid \cdot) : \mathcal{X} \to \mathcal{P}, \ \boldsymbol{z} \mapsto \Pi(\boldsymbol{\gamma} \mid \boldsymbol{z}),$$

where $\mathcal{P}$ denotes a suitable class of probability measures. It is important to keep in mind that so far we are only working with simulated processes $\boldsymbol{z}$ and therefore have full access to the data generating process, as well as the true parameters and corresponding functionals, as depicted in Figure 1. This is especially relevant, as we are working with statistics of extremes, which in application settings are usually not elicitible, i.e. they cannot be the unique minimizer of some loss function (Brehmer and Strokorb, 2019).

In order to estimate the parameter distribution, we follow Pacchiardi and Dutta (2022) and utilize a generative neural network to learn an approximate posterior distribution $\mathbb{Q}_\phi(\cdot \mid \boldsymbol{z})$ given a (simulated) observation $\boldsymbol{z}$, where ϕ denotes the networks parameters. The network can be defined via the mapping $F_\phi : \mathcal{S} \times \mathcal{X} \to \Gamma$ that transforms samples from some probability distribution $\mathbb{P}_{\boldsymbol{s}}$ over the space $\mathcal{S}$, usually a Gaussian distribution, conditioned on the observed max-stable process $\boldsymbol{z} \in \mathcal{X}$. Samples of the approximate posterior distribution $\mathbb{Q}_\phi(\cdot \mid \boldsymbol{z})$ are then obtained by a forward pass through the neural network. In a general setting, we are working with finite observations of parameter-simulation

pairs $(\boldsymbol{\gamma}_i, \boldsymbol{z}_i)_{i=1}^n$, where $\boldsymbol{\gamma}_i \sim \Pi$ is generated from the prior parameter distribution and $\boldsymbol{z}_i \sim \mathbb{P}(\cdot | \boldsymbol{\gamma}_i)$ is simulated from the underlying model. Using these samples, the aim is to train the network parameters ϕ, such that it matches the true distribution, e.g. $\mathbb{Q}_\phi(\cdot \mid \boldsymbol{z}) \approx \Pi(\cdot \mid \boldsymbol{z})$ for all $\boldsymbol{z} \in \mathcal{X}$. Different options are available for training a generative neural network in order to produce an approximate posterior, such as adversarial training (Ramesh et al., 2022) or invertible neural networks (Radev et al., 2022). However, we focus on a recently developed method that trains a generative network by minimizing a proper scoring rule. The resulting network has the advantage that it does not suffer from mode collapse and can be trained in a fairly simple way. This method was introduced by Pacchiardi et al. (2024) and has been successfully applied for example to multivariate forecasting (Chen et al., 2022).

In order to obtain the underlying distribution, a suitable scoring function has to be chosen as the loss function of the network. Scoring rules are a class of functions that assign a numerical score to the distance between a probability distribution and a realized observation. More precisely, a scoring rule $S(\mathbb{P}, \boldsymbol{y})$ (Gneiting and Raftery, 2007) measures the discrepancy between a probability distribution $\mathbb{P}$ and a corresponding observation $\boldsymbol{y}$ of a random variable $\boldsymbol{Y}$. If $\boldsymbol{Y} \sim \mathbb{P}'$, the expected scoring rule is given as $S(\mathbb{P}, \mathbb{P}') := \mathbb{E}_{\boldsymbol{Y} \sim \mathbb{P}'} S(\mathbb{P}, \boldsymbol{Y})$. If the expected scoring rule is minimized by the true distribution $\mathbb{P}'$, the scoring rule S is said to be *proper*. If the attained minimum is unique, it is called *strictly proper*. For our setting of max-stable processes this means that a proper scoring rule attains its minimum, when the true parameter originates from the predicted distribution, e.g. $\boldsymbol{\gamma} \sim \mathbb{Q}_\phi(\cdot \mid \boldsymbol{z})$. Utilizing the notion of proper scoring rules, our generative network can be trained by solving

$$\operatorname*{argmin}_{\phi} \mathbb{E}_{\boldsymbol{\gamma} \sim \Pi} \mathbb{E}_{\boldsymbol{Z} \sim \mathbb{P}(\cdot | \boldsymbol{\gamma})} S(\mathbb{Q}_\phi(\cdot \mid \boldsymbol{Z}), \boldsymbol{\gamma}),$$

where $\mathbb{Q}_\phi(\cdot \mid \boldsymbol{z})$ denotes the approximate posterior of the network (Pacchiardi and Dutta, 2022). The solution of the above expression leads to $\mathbb{Q}_\phi(\cdot \mid \boldsymbol{z}) = \Pi(\cdot \mid \boldsymbol{z})$ almost everywhere. The most widely used scoring rule to evaluate probabilistic forecasts is the Continuous Ranked Probability Score (CRPS, Gneiting and Raftery, 2007). When dealing with observations $\boldsymbol{z} \in \mathbb{R}^m$ instead of $z \in \mathbb{R}$ the *energy score* generalizes the CRPS and can be applied to distributional forecasts of vector-valued quantity instead of single values (Gneiting and Raftery, 2007). While in principle many choices of Scoring Rules are available (compare Pacchiardi et al., 2024), we focus on using the energy score, as it is strictly proper under mild regularity conditions and therefore admits a unique minimum. In the same setting as above, the energy score is given by $\operatorname{ES}(\mathbb{P}, \boldsymbol{y}) = \mathbb{E}\big[\|\boldsymbol{Y} - \boldsymbol{y}\|^\beta\big] - \frac{1}{2}\mathbb{E}\big[\|\boldsymbol{Y} - \boldsymbol{Y}'\|^\beta\big]$, where $\boldsymbol{Y}, \boldsymbol{Y}' \overset{i.i.d}{\sim} \mathbb{P}$ and $\beta \in (0, 2)$. However, a closed form solution of the scoring rule is usually not admissible, and it needs to be replaced by an unbiased estimator that is evaluated using m samples generated by the neural network.[3]

Transferring this to our setting of max-stable processes, consider a true parameter $\boldsymbol{\gamma} \sim \Pi(\cdot \mid \boldsymbol{z})$ and samples from the approximate posterior of the neural network $(\hat{\boldsymbol{\gamma}}_j)_{j=1}^m \sim \mathbb{Q}_\phi(\cdot \mid \boldsymbol{z})$. Using $\beta = 1$ we utilize the following unbiased estimator (Pacchiardi and Dutta, 2022):

$$\operatorname{ES}(\mathbb{Q}_\phi(\cdot \mid \boldsymbol{z}), \boldsymbol{\gamma}) = \frac{1}{m} \sum_{j=1}^m \|\hat{\boldsymbol{\gamma}}_j - \boldsymbol{\gamma}\|_2 - \frac{1}{2m(m-1)} \sum_{\substack{j,k=1 \\ k \neq j}}^m \|\hat{\boldsymbol{\gamma}}_j - \hat{\boldsymbol{\gamma}}_k\|_2. \tag{5}$$

The main advantage of such an approach is that by using the estimate of the posterior distribution, any functional of interest, for example confidence intervals can be derived. This is especially important, since we are trying to replace classical likelihood-estimation, which supplies these information in general. A visualization of the complete estimation workflow can be found in Figure 1.

With a simple extension, this approach can also be used to estimate the pairwise extremal coefficient function directly, only requiring little changes in the setup. For a given max-stable process, consider $\theta : (0, \infty) \longrightarrow [1, 2]$, $h \mapsto \theta(h)$ to be the corresponding true pairwise extremal coefficient function. We now want to sample from the (functional) posterior distribution using the same approach as before. To achieve that goal, we discretize the function $\theta(h)$ with a finite number of points $h_1, h_2, \ldots, h_k$, $k \in \mathbb{N}$ and function values $\theta(h_1), \ldots, \theta(h_k)$. In addition we can always set $h_0 = 0$ and $\theta(h_0) = 1$. We now use our network to estimate the vector $\boldsymbol{\theta}(h_{disc}) := (\theta(h_1), \ldots, \theta(h_k))^T$ and we restrict the values to lie between $(1, 2)$. As in the parameter estimation setup we use the energy score as a loss function to train our model but calculate the score with the estimated samples $\hat{\boldsymbol{\theta}}_j(h_{disc}) := \left(\hat{\theta}_j(h_1), \ldots, \hat{\theta}_j(h_k)\right)^T$, $j = 1, \ldots, m$ of the function values $\boldsymbol{\theta}(h_{disc})$. The adjusted estimator for the functional energy score is given as

$$\operatorname{ES}_\theta := \frac{1}{m} \sum_{j=1}^m \left\|\hat{\boldsymbol{\theta}}_j(h_{disc}) - \boldsymbol{\theta}(h_{disc})\right\|_2 - \frac{1}{2m(m-1)} \sum_{\substack{j,l=1 \\ l \neq j}}^m \left\|\hat{\boldsymbol{\theta}}_j(h_{disc}) - \hat{\boldsymbol{\theta}}_l(h_{disc})\right\|_2. \tag{6}$$

---

[3]Contrary to for example GANs, this approach therefore requires a large enough amount of output samples in order to converge.

With the introduced framework, we can now estimate posterior parameter distributions of max-stable processes and functionals of interest with or without an underlying model assumption.

3 Simulation studies

In this section we conduct several simulation studies that highlight the usability of the proposed approach in the setting of max-stable processes. We apply our method to the previously introduced Brown-Resnick and Schlather model and provide a detailed analysis of the methods robustness. An additional focus is put on proper evaluation of the models, incorporating measures of spatial dependence and evaluation of uncertainty. Finally, we provide insights in how the models can be trained more data efficiently.

3.1 Neural network model
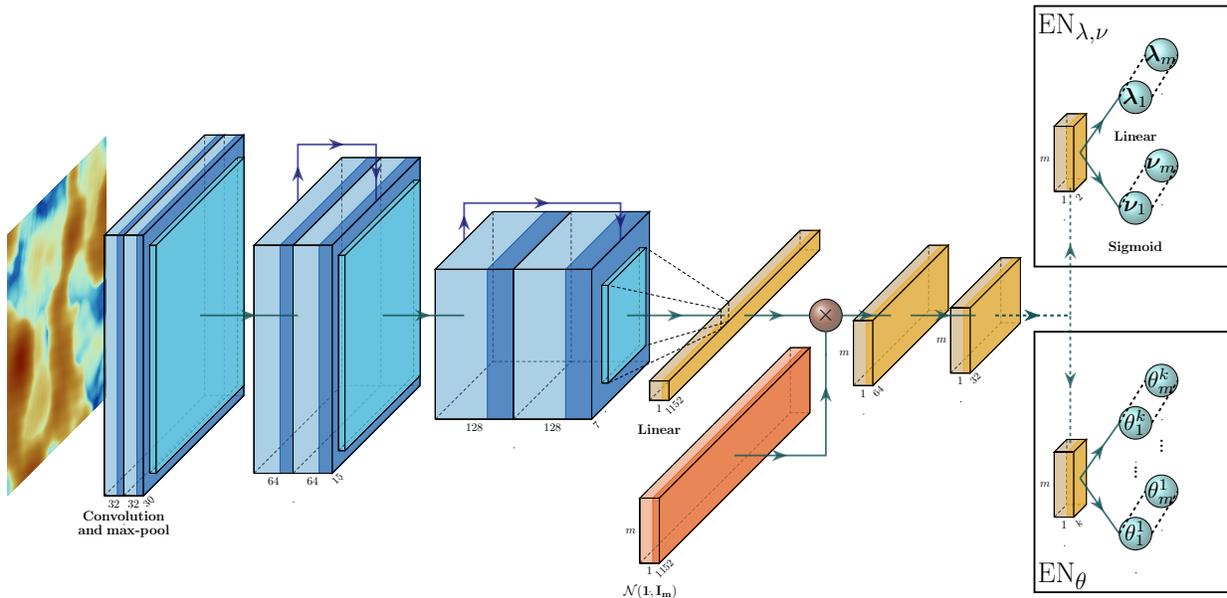
Figure 2: The figure shows the proposed model architecture. The spatial field is fed through three blocks convolutional and max-pooling layers. Across the blocks, the output size decreases, while the channel size increases. In the second and third block, residual connections are added, marked by the arrows on top. After the convolutional layers the network is flattened and fed through some final linear layers, where Gaussian noise is multiplied on top to finally create m output samples. For parameter prediction, samples of λ, ν are created, while for the direct estimation of the extremal coefficient function, sample points of the function are predicted as $\theta_j^i := \hat{\theta}_j(h_i)$.

In order to assess the performance of our proposed approach we evaluate different simulation scenarios, based on max-stable models. We use the previously introduced Brown-Resnick and Schlather model with $k = 900$ spatial locations uniformly distributed on the domain $\mathcal{D} = [0, 30] \times [0, 30]$. As we are assuming that the underlying process changes across time, our approach aims at providing estimates from a single process observation.[4] For the simulations, similar to Erhardt and Smith (2012), we aim at using minimally informative priors by utilizing $\Pi \sim \mathcal{U}(\boldsymbol{a}, \boldsymbol{b})$, where the parameters differ across the simulation scenarios. The realizations of the processes are obtained using the R-package *SpatialExtremes* (Ribatet, 2022).

As we are working with two-dimensional spatial data on a regular grid, convolutional neural networks (CNNs) are a natural choice of architecture and have been successfully applied to max-stable processes (Lenzi et al., 2023; Sainsbury-Dale et al., 2022). Figure 2 shows a visualization of the network. All layers are equipped with ReLu activation functions, except for the final layers. Since $\nu \in (0, 2]$, we transform it to the unit interval and use a sigmoid activation function.

---

[4]If this assumption is relaxed, the network can easily be adjusted to processing multiple samples, by using three dimensional convolutions.

For λ we employ a log-transform and a linear activation function, similar to Lenzi et al. (2023), while for the values of $\theta(h_1), \ldots, \theta(h_n)$ no transformation is required and we use a sigmoid activation function scaled to the range $(1, 2)$. We specify the network to create $m = 500$ samples from the posterior distribution by sampling from a latent space $\mathcal{N}(\mathbf{1}, \boldsymbol{I}_m)$ and multiplying it to a linear layer (compare Figure 2).

For training the network, the main bottleneck lies in the simulation of the processes used as training data. While different methods have been proposed and evaluated, such as using an informative prior (Lenzi et al., 2023) or simulating new data during training (Sainsbury-Dale et al., 2022), we want to focus on utilizing techniques from image augmentation. Image augmentation refers to the process of rotating, flipping or distorting images and using them as additional training data (Perez and Wang, 2017). While not all of these methods can be used, as we need to preserve the structure of the max-stable process, we can still take advantage of the general idea. As the Brown-Resnick and Schlather model are stationary and isotropic, meaning that the properties of the processes across the spatial domain only depend on their distance, we can utilize image augmentation techniques that do not distort the distance between spatial locations. For that purpose we use an image rotation of $180°$ and vertical and horizontal image flips with probability 0.3 and 0.2 respectively.

The network is trained minimizing the Energy Score with the RMSProp optimizer implemented in PyTorch and a learning rate of $7e^{-4}$ using a learning rate scheduler that stops training if the metrics do not improve. In each epoch the weights are updated across a data batch of size 100. The parameters were chosen based on some minor experiments, extensive hyperparameter tuning is left as future work. Studies were conducted using a workstation with an Intel XEON E5-2680 2.50 GHz CPU with 40 cores and an NVIDIA GeForce RTX 2080 with 8GB of GPU RAM. Reproducible code can be found at `http://www.github.com`.

As a comparison, we implement several benchmark methods. First, we use the previously described pairwise likelihood method, where, following Lenzi and Rue (2023), the optimizer is run from 20 starting values, from which the 5 best estimates are again used as starting values, leading to the final estimate. The cutoff for weights is chosen as 5 and the estimator is fitted using the function *fitmaxstab* of the R-package *SpatialExtremes* (Ribatet, 2022). Furthermore we employ the ABC method, which is based on the idea of generating many simulations of a process and comparing those to the observed data, using some summary statistic S. Using a pre-specified cutoff distance, the algorithm outputs m samples with the best value of S. Following Erhardt and Smith (2012), we use the tripletwise extremal coefficient as a summary statistic. As the method is not feasible for a large number of spatial locations, the max-stable processes are downsampled to a grid of size 5×5, using bilinear interpolation. We generate 50000 independent simulations with 25 processes each to compare against the observed data, where the cutoff is chosen such that the algorithm results in $m = 500$ samples. Finally, we also compare our methods to a "regular" implementation of the CNN with the same hyperparameters but minimizing the mean squared error, which is the setting in Lenzi et al. (2023). For the rest of this article, we will refer to the networks as $\text{EN}_{\lambda,\nu}$ and EN_θ for the parameter estimation and direct estimation, respectively.[5]

3.2 Evaluation

In order to completely assess the adequacy of the estimations, a detailed evaluation procedure for the max-stable processes is required, which is introduced in this section. While the network outputs a full (sample-based) distribution, one is generally more interested in specific functionals or characteristics of it. For that reason, we restrict the evaluation to important and commonly used quantities. However, since our method is not specifically tailored to the energy score, any functional can be evaluated, as long as it admits some sort of empirical estimator. Specifically, we provide evaluation for the mean and interval predictions for γ and $\theta(h)$, as well as an additional measure determining the fit of the complete distribution. By providing these different performance metrics, we can assess the predictive power of the estimator, while simultaneously analyzing the uncertainty in the predictions.

To evaluate the parameter estimation, we employ the typical mean squared error (MSE) as a metric, which is given by $\text{MSE}_\lambda(\hat{\boldsymbol{\gamma}}, \boldsymbol{\gamma}) := \mathbb{E}[\|\hat{\lambda} - \lambda\|^2]$ for the parameter λ and for ν similarly. In order to assess the uncertainty of the prediction, recall that the model outputs m samples from the posterior parameter distribution. In order to obtain uncertainty intervals for each parameter separately, the corresponding empirical quantiles can be estimated from the model output. A typical way to assess the quality of a prediction interval is to use the so-called interval score (IS Gneiting and Raftery, 2007). Let λ be the true range parameter and l_α, u_α denote the lower and upper endpoints of the predictive interval, e.g. the quantiles at $\left(\frac{\alpha}{2}, 1 - \frac{\alpha}{2}\right)$. Then the interval score for that parameter is given by

$$IS_{\alpha,\lambda}(l_\alpha, u_\alpha) := (u_\alpha - l_\alpha) + \frac{2}{\alpha}(l_\alpha - x)\mathbb{1}\{\lambda < l_\alpha\} + \frac{2}{\alpha}(\lambda - u_\alpha)\mathbb{1}\{\lambda > u_\alpha\}. \tag{7}$$

The interval score is a proper scoring rule with respect to the corresponding quantiles and a lower score corresponds to a better prediction. For the parameter ν the score is calculated in the same way. While the interval score is

---

[5]EN as an abbreviation for energy network.

helpful to quantify the uncertainty estimate, it can only be used for each parameter separately. In order to assess the full predicted distribution $\mathbb{Q}_\phi(\cdot|\mathbf{z})$ of the model we can use the already established energy score (5). As described previously, the energy score measures the discrepancy between the true parameter $\boldsymbol{\gamma}$ and the predicted posterior samples $(\hat{\boldsymbol{\gamma}})_{j=1}^m \sim \mathbb{Q}_\phi(\cdot \mid \mathbf{z})$.

So far, the focus was set on evaluating the fit of the parameters λ, ν that specify the max-stable process. However, as the relation between the model and its parameters is highly nonlinear, a seemingly bad parameter estimate might still lead to a good dependence estimate. In addition, focusing only on the parameter fit does not allow for comparison between different max-stable models. For that purpose, we extend the MSE to measure the error in the estimated spatial dependence. Denote by $\theta(h; \boldsymbol{\gamma})$ the pairwise extremal coefficient function (4) of the max-stable model specified by $\boldsymbol{\gamma}$. Then we can use the following integrated metric based on the L^2 norm:

$$\mathrm{MSE}_\theta(\hat{\boldsymbol{\gamma}}, \boldsymbol{\gamma}) = \mathbb{E}\left[\int_0^\infty \|\theta(h; \hat{\boldsymbol{\gamma}}) - \theta(h; \boldsymbol{\gamma})\|^2 \, dh \right]. \tag{8}$$

This expression can be interpreted as taking the pointwise mean squared error of the extremal coefficient at h and average it over all possible distances $h \in \mathbb{R}_+$. This allows for evaluating how the estimated parameters translate to the spatial dependence and to compare different max-stable models. As the parameters could differ or might not have the same interpretation, evaluating the MSE on the parameters alone is not enough for a full assessment of the model prediction.

Similar as before, we also want to assess the uncertainty in the prediction of the spatial dependence. For that purpose, we can use the previously established interval score (7) to build pointwise confidence intervals for $\theta(h)$. The resulting interval specifies bounds in which the true function lies with probability α. For a fixed distance h, the interval score can be calculated by taking the empirical α-quantile of the functions $\theta(h; \hat{\boldsymbol{\gamma}}_i)$, $i = 1, .., m$. The interval score over the whole function, can again be achieved via integration as

$$IIS_\alpha := \int_{h>0} IS_\alpha(\hat{l}_\alpha(h), \hat{u}_\alpha(h); \theta(h; \boldsymbol{\gamma})) \, dh, \tag{9}$$

where $\hat{l}_\alpha(h), \hat{u}_\alpha(h)$ are the empirical quantiles evaluated at h and $\theta(h; \boldsymbol{\gamma})$ is the true extremal coefficient function. Utilizing Equation 6 we can again use the energy score to assess the fit of the whole predictive distribution.

As the pairwise likelihood and the regular CNN do not incorporate uncertainty estimations, the corresponding metrics do not apply. Furthermore, since the probabilistic methods generate multiple samples as the output, these need to be aggregated for evaluation. For the MSE of the parameters, the estimation is simply given by the sample mean, while for the extremal coefficient function, a pointwise mean function is used. As the direct estimator is discretized at $h_1, \ldots, h_k$ points, the integral metrics are approximated via a sum over those support points in order to make all methods comparable.

3.3 Results

As a first comparison, for both max-stable models a test set of parameters and simulated processes is generated by

$$(\lambda_i^{\text{test}}, \nu_i^{\text{test}}) : \ \lambda_i^{\text{test}} \sim \mathcal{U}(0.5, 5), \ \nu_i^{\text{test}} \sim \mathcal{U}(0.3, 1.8), \quad i = 1, \ldots, 250, \tag{10}$$

which covers a wide range of parameter combinations and different spatial dependencies. The training set of size $n = 5000$ is generated from the same parameter space, with 20% of the data used as validation data. The direct estimation of $\boldsymbol{\theta}(h_{disc})$ requires an upper bound of h, since the network can only be trained for a finite number of points. We decide to consider simply $h \in \left(0, \sqrt{30+30}\right]$, which is the maximal spatial separation across the max-stable processes on the 30×30 grid. Taking an interval distance $dh = 0.1$ leads to the support points $h_1 = 0.1, h_2 = 0.2, \ldots, h_{425} = 42.5$.

One problem of directly estimating values of the pairwise extremal coefficient function is that there is no guarantee for the resulting function to still be monotone, which is a main characteristic of $\theta(h)$. However, this issue can be solved by permutation of the estimated values. As our approach not only predicts a single function $\theta(h_i)$, $i = 1, \ldots, 425$ but a full sample distribution, the order of permutation and functional calculation needs to be considered. If one is interested for example in the mean prediction of $\theta(h)$, first the mean is calculated $\overline{\hat{\theta}(h_i)}, i = 1, \ldots, 425$ and afterwards the resulting functional is permuted, such that $\overline{\hat{\theta}(h_1)} \leq \overline{\hat{\theta}(h_2)} \leq \ldots \leq \overline{\hat{\theta}(h_{425})}$ (Appendix B includes further details).

The full results for both max-stable models are given in Table 1. As can be seen, the energy network has the best performance across most metrics and both models. Especially for the error regarding the extremal coefficient function, the method is much better than the benchmark methods, with significant additional improvements by the direct estimation. However, for the parameter estimation, the performance of the regular CNN is usually quite similar and in
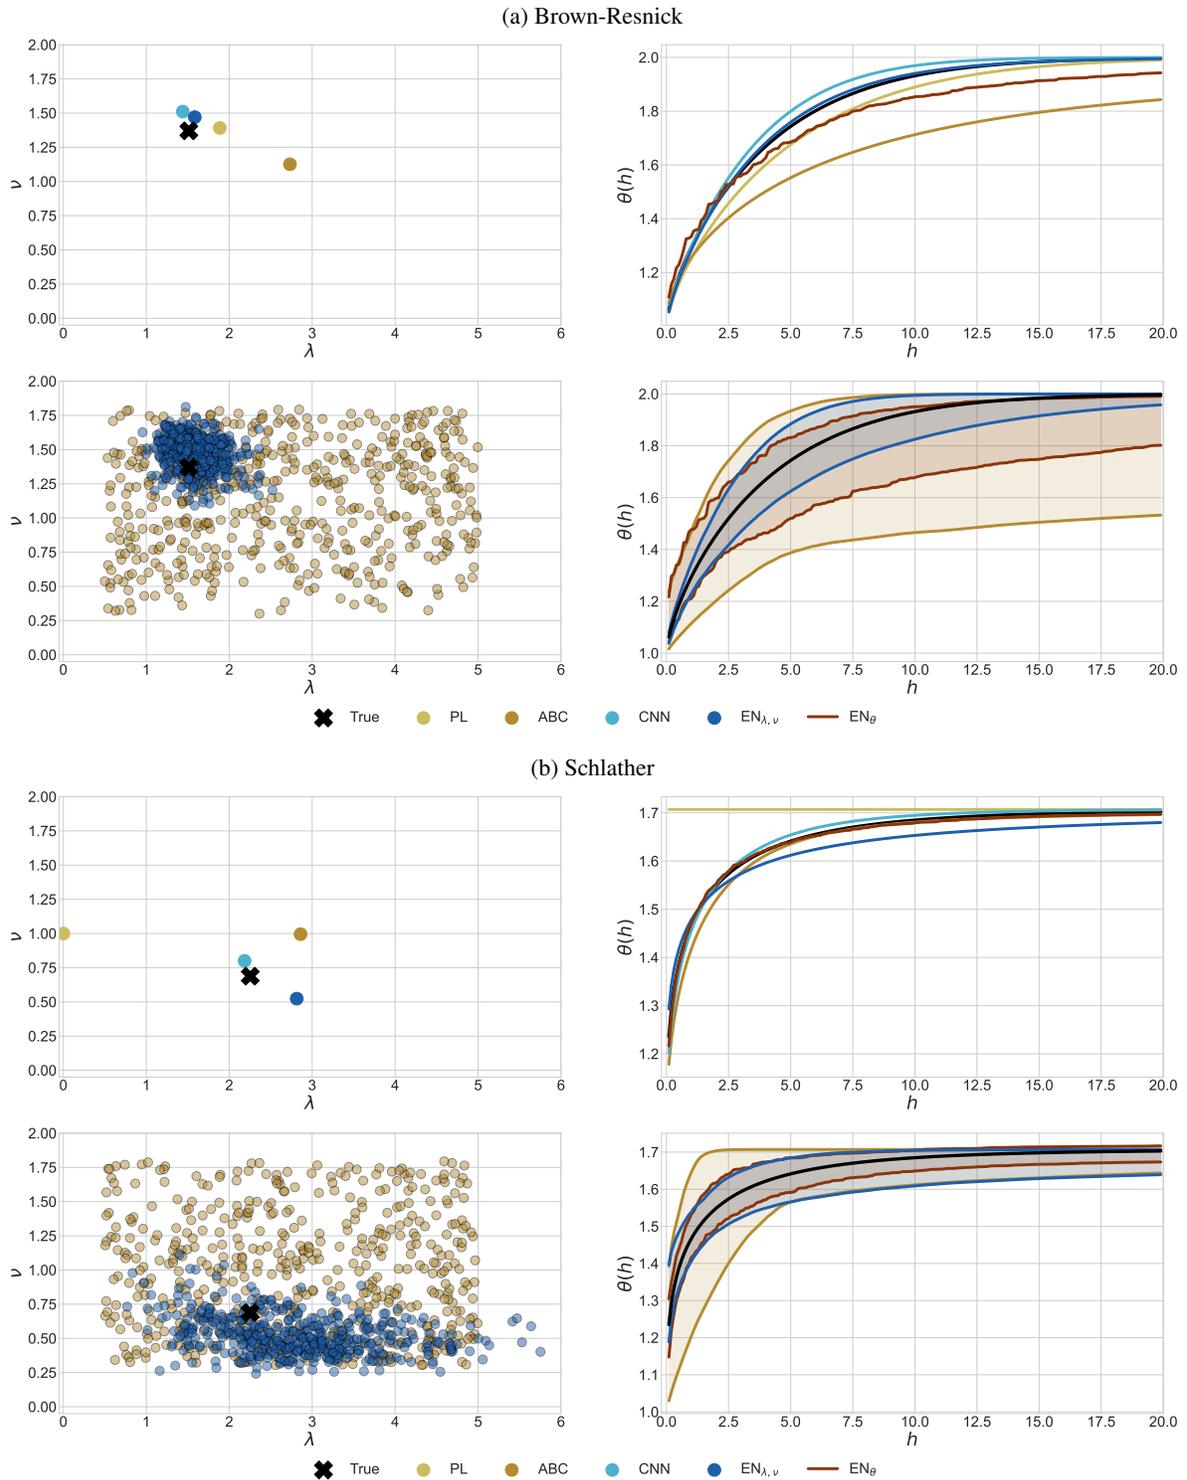
Figure 3: The figure visualizes the different estimation methods for the max-stable models using a selected test sample $((\lambda, \nu) = (1.51, 1.37)$ for the Brown-Resnick and $(\lambda, \nu) = (2.25, 0.69)$ for the Schlather model). In each figure the upper left panel shows the different location estimates, while the upper right panel shows the estimated extremal coefficient functions. The lower left panel shows the sample-based distribution estimates of the ABC and $EN_{\lambda,\nu}$ method and the lower right panel shows the estimated pointwise confidence intervals ($\alpha = 0.05$) for the extremal coefficient function.

		PL	**CNN**	**ABC**	**EN$_{\lambda,\nu}$**	**EN$_\theta$**
Brown-Resnick	**MSE$_\lambda$**	3.19 (20.44)	0.43 (0.89)	1.61 (1.79)	**0.38 (0.85)**	-
	MSE$_\nu$	0.13 (0.22)	0.02 (0.04)	0.26 (0.32)	**0.01 (0.02)**	-
	MSE$_\theta$	0.71 (1.20)	0.19 (0.34)	1.33 (1.67)	**0.15 (0.25)**	0.16 (0.27)
	IS$_{0.05,\lambda}$	-	-	4.59 (3.03)	**3.03 (6.43)**	-
	IS$_{0.05,\nu}$	-	-	3.43 (6.53)	**0.55 (0.49)**	-
	IIS$_{0.05}$	-	-	53.83 (93.44)	**10.66 (14.19)**	13.71 (23.00)
	ES$_{\lambda,\nu}$	-	-	0.85 (0.43)	**0.34 (0.32)**	-
	ES$_\theta$	-	-	1.92 (0.68)	1.89 (0.78)	**0.74 (0.59)**
Powexp	**MSE$_\lambda$**	8.35 (6.86)	**0.47 (0.96)**	1.60 (1.49)	0.66 (1.24)	-
	MSE$_\nu$	0.78 (0.85)	**0.03 (0.05)**	0.19 (0.17)	0.05 (0.08)	-
	MSE$_\theta$	0.29 (0.17)	**0.01 (0.02)**	0.05 (0.04)	0.02 (0.03)	**0.01 (0.02)**
	IS$_{0.05,\lambda}$	-	-	4.30 (0.33)	**3.45 (4.05)**	-
	IS$_{0.05,\nu}$	-	-	1.47 (0.26)	**1.35 (2.21)**	-
	IIS$_{0.05}$	-	-	4.22 (3.28)	**2.98 (3.67)**	3.58 (5.09)
	ES$_{\lambda,\nu}$	-	-	0.82 (0.32)	**0.46 (0.36)**	-
	ES$_\theta$	-	-	0.45 (0.16)	0.47 (0.14)	**0.23 (0.14)**

Table 1: The table shows the different metrics for the different estimation methods and max-stable models. All metrics are negatively oriented, with the best model highlighted in bold and standard deviation given in brackets.

some cases even better. The pairwise likelihood approach leads to very large errors in comparison, especially for the powered exponential model. While the metrics of the ABC method are lower than those of the PL approach they still do not compare to the neural network methods.

A selected visualization of the different estimation methods is shown in Figure 3. For the Brown-Resnick model, it is clear that the energy network leads to the best estimation. While the regular CNN and the PL estimation are also quite close to the original parameters, the corresponding extremal coefficient function is much better represented by the estimates of EN$_{\lambda,\nu}$. Furthermore, if looking at the uncertainty, the estimated samples have a much lower spread than those of the ABC method. This also corresponds to the prediction intervals for $\theta(h)$. The same goes for the EN$_\theta$, although due to the nature of the approach the estimated function is not smooth. While for all three methods, the pointwise prediction intervals of $\theta(h)$ seem to be adequate, the intervals of EN$_{\lambda,\nu}$ are much sharper and therefore preferable. The direct estimate leads to even sharper prediction intervals but the true function is not always entailed, as represented in the integrated interval score. For the Schlather model, the pairwise likelihood approach does not lead to usable estimations, as it always seems to predict the range as $\lambda = 0$, meaning that all spatial locations are independent, which also leads to an incorrect extremal coefficient function. Again, the predictions provided by EN$_{\lambda,\nu}$ seem to fit the best, as they scatter closely around the true parameter and lead to a good representation of $\theta(h)$.

3.4 Robustness

While, the previous analysis was performed under optimal conditions, where one assumes the correct model and parameter space a priori, in practice the true model is not only unknown, it might even be unobservable or non-existent. Therefore it is important to check the robustness of the approach against model misspecification. In order to do so, we consider three different scenarios:

1. Misspecified parameter range.
2. Misspecified correlation function.
3. Unspecified model.

An additional scenario of an overspecified model can be found in Appendix D. Evaluating these scenarios will give insights on how the different estimators are able to extrapolate across the parameters and across different models. A robust estimator should still work reliably in a setting, where the underlying model and corresponding parameters are unknown. The numerical results for the first and second scenario are shown in Table 2.

Scenario #1 In the first scenario, we analyze how the different estimators are able to deal with a misspecified parameter space in the setting of a Brown-Resnick model. For that purpose, the set of parameters in the training data is chosen as a subset from the test data, forcing the models to predict parameters for previously unseen processes. The training range is chosen as $\lambda^{\text{train}} \in [0.5, 5]$, $\nu^{\text{train}} \in [0.3, 1.8]$, while the test set covers $\lambda^{\text{test}} \in (0, 0.5) \cup (5, 10]$, $\nu^{\text{test}} \in (0, 0.3) \cup (1.8, 2]$.
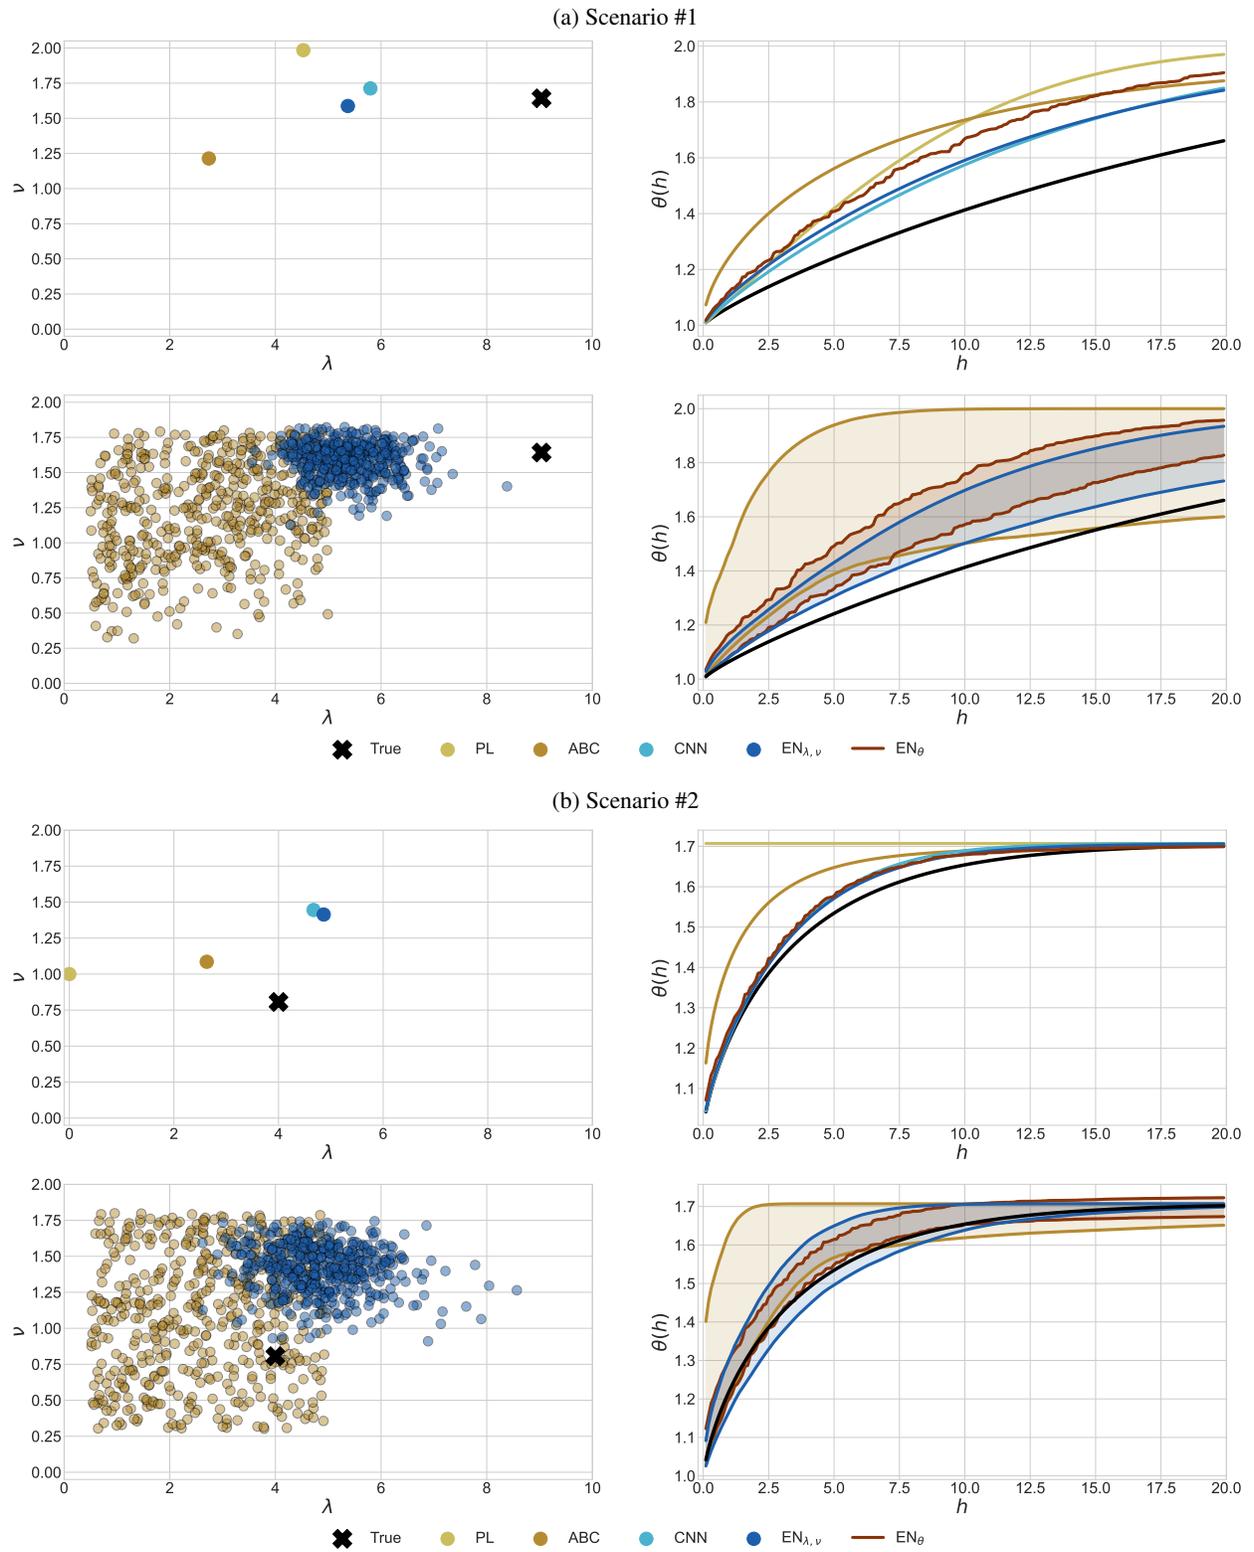
Figure 4: The figure visualizes the different estimation methods for the robustness scenarios using a selected test sample $((\lambda, \nu) = (9.04, 1.64)$ for scenario #1 and $(\lambda, \nu) = (4.00, 0.81)$ for scenario #2). The plot division is the same as in Figure 3.

		PL	CNN	ABC	$EN_{\lambda,\nu}$	EN_θ
#1	$\mathbf{MSE}_\lambda$	13.07 (29.13)	9.68 (10.65)	18.01 (15.15)	**9.47 (10.15)**	-
	$\mathbf{MSE}_\nu$	0.23 (0.55)	0.06 (0.10)	0.52 (0.58)	**0.03 (0.06)**	-
	$\mathbf{MSE}_\theta$	0.91 (1.46)	0.78 (0.83)	2.65 (2.48)	**0.55 (0.52)**	0.76 (0.65)
	$\mathbf{IS}_{0.05,\lambda}$	-	-	87.49 (68.32)	**59.76 (56.30)**	-
	$\mathbf{IS}_{0.05,\nu}$	-	-	9.12 (12.24)	**1.22 (1.00)**	-
	$\mathbf{IIS}_{0.05}$	-	-	111.21 (100.04)	**54.54 (54.08)**	92.36 (78.12)
	$\mathbf{ES}$	-	-	3.25 (1.75)	**2.24 (1.63)**	-
	$\mathbf{ES}_\theta$	-	-	2.88 (0.99)	2.75 (1.17)	**2.03 (1.16)**
#2	$\mathbf{MSE}_\lambda$	9.61 (7.28)	2.37 (2.82)	**1.70 (1.47)**	1.76 (1.87)	-
	$\mathbf{MSE}_\nu$	0.94 (0.89)	0.17 (0.17)	0.17 (0.16)	**0.16 (0.16)**	-
	$\mathbf{MSE}_\theta$	0.67 (0.42)	**0.03 (0.04)**	0.17 (0.19)	0.04 (0.06)	0.05 (0.09)
	$\mathbf{IS}_{0.05,\lambda}$	-	-	**4.31 (0.41)**	12.38 (14.42)	-
	$\mathbf{IS}_{0.05,\nu}$	-	-	**1.46 (0.29)**	4.683 (5.57)	-
	$\mathbf{IIS}_{0.05}$	-	-	13.80 (16.79)	**7.21 (12.46)**	13.12 (18.78)
	$\mathbf{ES}$	-	-	**0.83 (0.32)**	0.92 (0.54)	-
	$\mathbf{ES}_\theta$	-	-	0.84 (0.58)	0.66 (0.42)	**0.45 (0.46)**

Table 2: The table shows the evaluation metrics across the robustness scenarios #1 and #2. All metrics are negatively oriented, with the best model highlighted in bold and standard deviation given in brackets.

The results in Table 2 show that the $EN_{\lambda,\nu}$ obtains the lowest errors for almost all metrics, while the EN_θ obtains similar errors for the MSE and a lower error for the functional energy score. In comparison to the previous results in Table 1 the errors are generally quite high, especially for the IS and the IIS, which indicates that the predictive intervals might not be very adequate. However, a direct comparison is not possible, as in this scenario the magnitude of the parameters is larger, which leads to larger errors. Figure 4a shows a visualization of a selected test sample. It is clear that all point estimators, as well as the functional estimations of $\theta(h)$ do not fit to the real value very well. However, the estimations of $EN_{\lambda,\nu}$ do shift to the true parameter, extrapolating to values outside of the training range of the model. For this example, no method provides adequate confidence intervals for $\theta(h)$, as the true function is covered almost nowhere by the intervals. Still, the quantile predictions of $EN_{\lambda,\nu}$ are better in terms of the interval score and distance to the true function, as is also reflected in the numerical results.

Scenario #2 In the second scenario, we analyze how the estimators are able to deal with a misspecified model, by using a Schlather process with a wrong correlation function. For that purpose, a test set is generated with the same parameters as before for a Schlather model with a Whittle-Matérn kernel, while the training set is generated via a powered exponential kernel. Table 2 shows that the PL approach has a significantly higher MSE than the other models, while the CNN and the ABC method seem to perform quite well. While at a first look, the $EN_{\lambda,\nu}$ seems to have quite high errors in estimating the parameters, it results in really low errors for metrics regarding the extremal coefficient function. Generally, the approach corresponds to high errors in all parameter related metrics (including the energy score) and a very low error in all metrics regarding θ. The EN_θ does not lead to much improvements, except in the functional energy score. Figure 4b shows a selected visualization of the model estimations. As can be seen, the $EN_{\lambda,\nu}$ extrapolates to some previously unknown parameter range leading to a bad parameter estimation but a very good representation of the spatial dependence. The normal CNN displays a similar behavior, with a good fit of the spatial dependence. The $EN_{\lambda,\nu}$ also produces valid confidence intervals, while those of the ABC method do not lead to good coverage in this example. This is also the case for the direct estimation, which leads to quite narrow confidence intervals, but does not cover the true function adequately.

Scenario #3 In the case where the true model is unknown, it can be advantageous to train the network on different underlying models, so that it can react flexibly to given test data. For that purpose, training data is generated from the Brown-Resnick and from the Schlather model with powered exponential and Whittle-Matérn kernel. Combining these models covers many different data scenarios and thus should lead to the neural network being able to generalize its predictions across different models. We consider two different training sets, containing 1666 datapoints and 5000 datapoints of each model, respectively. The first case corresponds to the same amount of data points as in the previous scenarios, while the second case also captures the effect of increasing the amount of data. The trained models are tested on data from the Brown-Resnick model and from the Schlather model with powered exponential kernel. The results are shown in Table 3.

Comparing these metrics with those in Table 1, shows that the scores for the Brown-Resnick model for both the $EN_{\lambda,\nu}$ and the EN_θ are higher when training on different models, unless the dataset is enlarged at the same time. For the

			CNN	$\text{EN}_{\lambda,\nu}$	EN_{θ}
1666 dat-apoints	**Brown-Resnick**	$\textbf{MSE}_{\theta}$	0.33 (0.6)	0.37 (0.63)	**0.29 (0.45)**
		$\textbf{IIS}_{0.05}$	-	**15.94 (19.62)**	20.96 (33.60)
		$\textbf{ES}_{\theta}$	-	1.91 (0.96)	**1.00 (0.79)**
	Powexp	$\textbf{MSE}_{\theta}$	**0.02 (0.02)**	**0.02 (0.03)**	0.03 (0.04)
		$\textbf{IIS}_{0.05}$	-	**3.70 (8.13)**	5.80 (8.16)
		$\textbf{ES}_{\theta}$	-	0.47 (0.20)	**0.35 (0.22)**
5000 dat-apoints	**Brown-Resnick**	$\textbf{MSE}_{\theta}$	0.20 (0.35)	0.23 (0.40)	**0.17 (0.30)**
		$\textbf{IIS}_{0.05}$	-	**13.94 (19.45)**	15.39 (32.37)
		$\textbf{ES}_{\theta}$	-	1.89 (0.91)	**0.74 (0.66)**
	Powexp	$\textbf{MSE}_{\theta}$	0.02 (0.02)	0.02 (0.02)	**0.01 (0.02)**
		$\textbf{IIS}_{0.05}$	-	**2.57 (2.51)**	2.76 (3.63)
		$\textbf{ES}_{\theta}$	-	0.46 (0.18)	**0.21 (0.16)**

Table 3: The table shows the different metrics for the different estimation methods and max-stable models for training on 1666 and 5000 datapoints each. All metrics are negatively oriented, with the best model highlighted in bold and standard deviation given in brackets.

Schlather model with powered exponential kernel the scores in Table 3 are similar to those in Table 1, except for the interval score. Increasing the number of data points improves the estimation, although only for the Schlather model the obtained scores are lower than in Table 1. However, a higher amount of data also improves the estimation for the Brown-Resnick model, so that the scores in Table 3 lie in the same range as the ones in Table 1. This is a advantageous when dealing with real data, where the true model is not known and cannot be used to determine the extremal coefficient function, as required for $\text{EN}_{\lambda,\nu}$. Altogether, these results suggest that the EN_{θ} might benefit from training on different max-stable models as it can then predict the extremal coefficient function of a given dataset regardless of its true underlying model.

The previous results indicate that our approach is able to learn the characteristics of the underlying misspecified model and can extrapolate to the correct model. By providing several scenarios of misspecification, we cover different cases that are of relevance in an application scenario and show that the networks still produce robust and reliable results. A closer look on the energy scores for single observations gives further insights into cases where the estimation seems to fail. For the Schlather model with identical test and training parameters, Figure 5a shows that for EN_{θ} high energy scores don't arise for certain values of λ and ν while for $\text{EN}_{\lambda,\nu}$ the scores are higher on the boundary of the test parameter range. This is most likely due to the fact that the estimations of $\text{EN}_{\lambda,\nu}$ highly depend on the parameters, as opposed to EN_{θ}. Analyzing robustness scenario #2, Figure 5b shows that for the EN_{θ} the energy score is high whenever λ and ν take on large values. Inspecting the pairwise extremal coefficient function, it is clear that in this area in particular, the values for the powered exponential kernel and for the Whittle-Matérn kernel differ significantly. The same effect can be seen for the $\text{EN}_{\lambda,\nu}$, but additionally the scores are also higher in the lower left corner. A visualization of the extremal coefficient function can be found in Figure 17 in Appendix C. The analysis of the energy scores based on the parameters λ and ν shows that depending on the given data one approach might be preferable over the other. While the energy scores for the EN_{θ} in the normal cases do not depend on parameter values, the $\text{EN}_{\lambda,\nu}$ performs worse on the margin of the trained parameter range. Also for the different robustness scenarios a certain model can perform better depending on the true parameter. Further examples and insides are given in Appendix C.

(a) Schlather - powered exponential

(b) Scenario #2
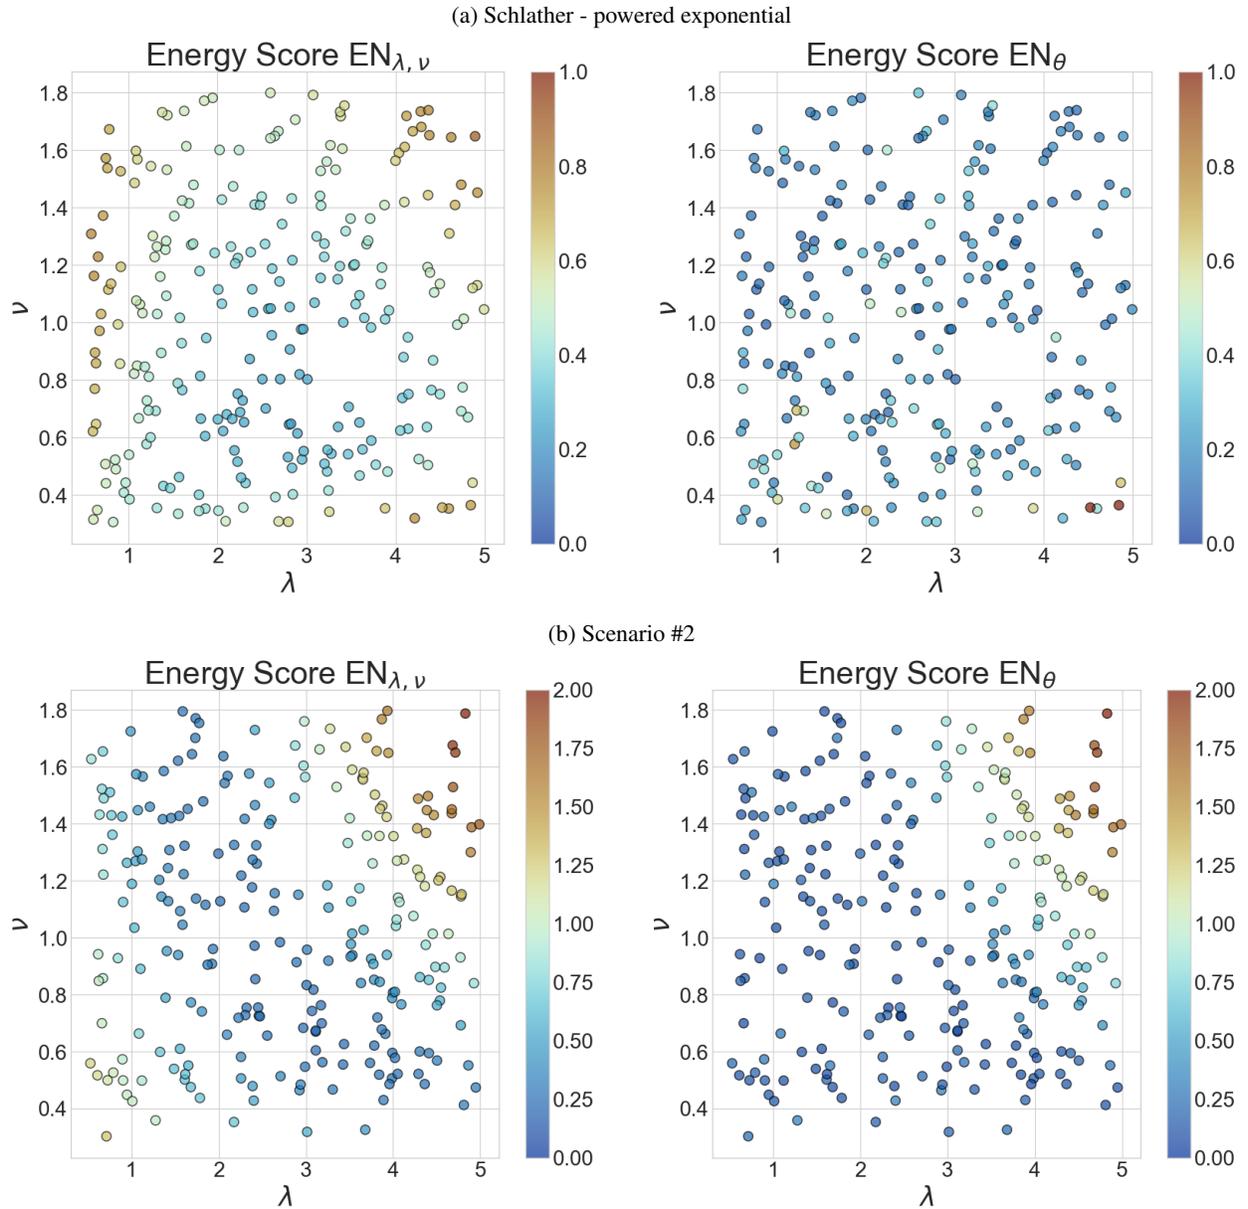
Figure 5: The figure visualizes the energy score across the parameters (λ, ν) from the test data. The upper panels display results from the optimal case with test data from the Schlather model with powered exponential kernel, while the lower panels show results for robustness scenario #2 with test data from the Schlather model with Whittle-Matérn kernel.

14

4 Application to German precipitation extremes

In this section, we apply the developed methodology to precipitation maxima across Western Germany. The dataset used is the HYRAS dataset,[6] which stems from the German National Meteorological service (DWD) and covers a reanalysis of mean daily precipitation from 1931 until 2023 across the whole of Germany on a 1km × 1km resolution. Our analysis is restricted to precipitation maxima over the summer months June, July and August, in order to reduce seasonality effects, similar to the setting described by Forster and Oesting (2022). Data from 2021-2023 is used for assessing the prediction quality, while data from 1932-2020 is used for fitting the GEV parameters, which are required to transform to unit Fréchet margins. After the processing, the resulting data is on a 30 × 30 grid of around 122km × 78km, covering the Rhineland with the cities Bonn and Cologne, as well as the Ahr valley, as depicted in Figure 6.
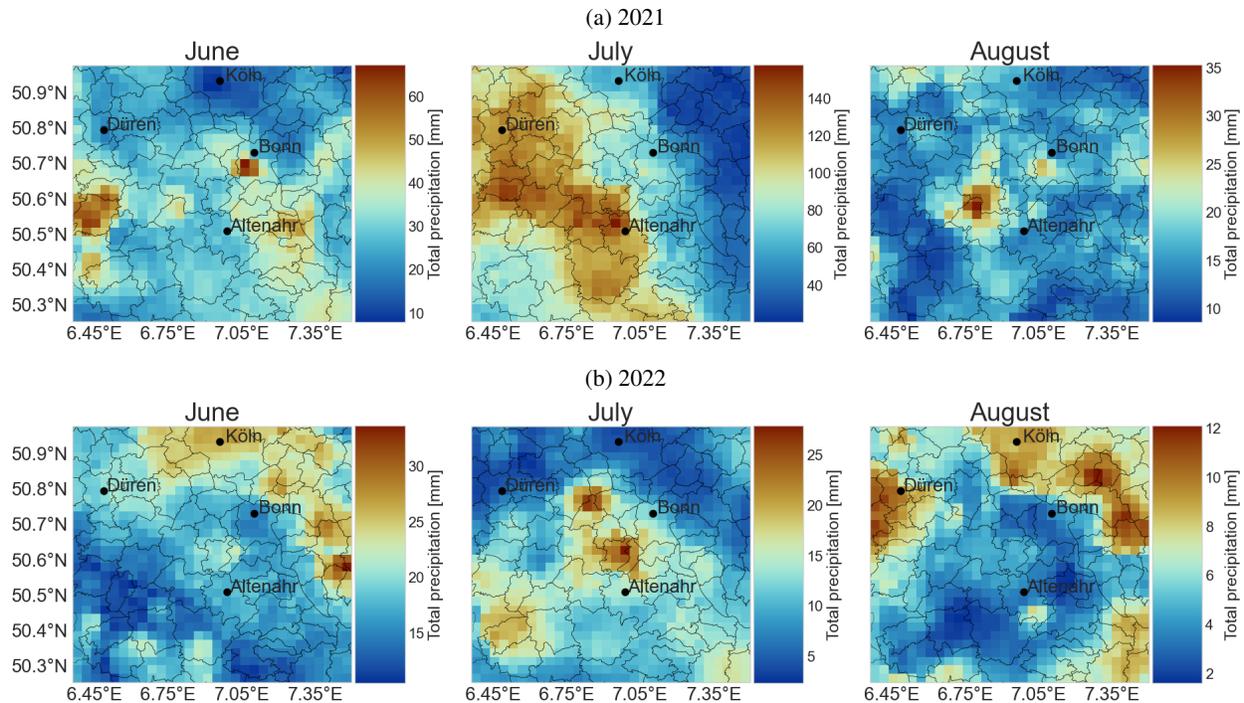
Figure 6: The figure shows the maximum precipitation in mm, aggregated over the three summer months of the years 2021 and 2022.

The specified area covers the region where in July 2021 an extremely heavy precipitation event took place, with over 150 mm precipitation on an extensive area in around 15h to 18h. The resulting floods, mainly concerning the Ahr valley, led to at least 180 fatalities, 40.000 people affected and an estimated damage of around EUR 32 billion[7]. For an overview and a description of the event see for example Bosseler et al. (2021) or Mohr et al. (2023). However, it is worth mentioning that the monthly precipitation maxima in July 2021 was the highest recorded maxima, across all available data. While the magnitude of this event is quite extreme, the spatial dependence might be very similar to other observations.

In order to transform the data into unit Fréchet margins, we fit a response surface (Ribatet, 2013) that includes additional covariates. Extensive model selection is omitted and we follow an approach similar to Davison and Gholamrezaee (2012); Sang and Gelfand (2010). More details can be found in Appendix E. Contrary to the simulation setting, with the actual data we do not know the underlying model or parameters. Although we have thoroughly assessed the performance of the approach previously, for this setting different means of evaluation are required. Furthermore, recall that so far the training parameters were usually chosen from some pre-specified parameter range. In order to keep the setting similar to the simulation studies, the data grid was transformed to units of $3.4km$, which allows to use the same

---

[6]Deutscher Wetterdienst, HYRAS - Hydrometeorologische Rasterdaten, version 5.0 (https://www.dwd.de/DE/leistungen/hyras/hyras.html)

[7]Munich Re, Hurricanes, cold waves, tornadoes: Weather disasters in USA dominate natural disaster losses in 2021, *Press report* 10.01.2022 (link, accessed on 07.11.2023)
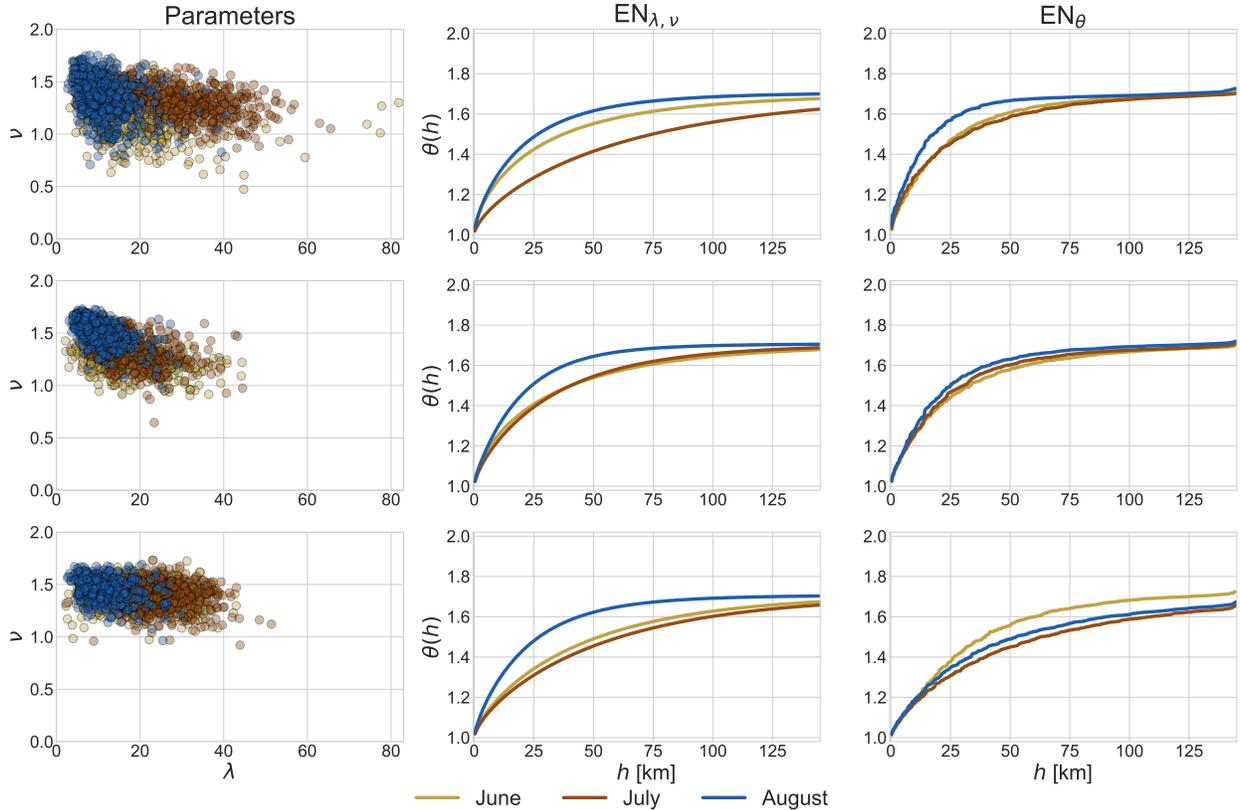
Figure 7: The figure shows the parameter estimates for a powered exponential model for all three months and years using the proposed approaches. The years are from 2021-2023 from top to bottom.

support points for estimating EN_θ.[8] In order to cover a large range of dependencies, we then choose to simulate data from $\lambda \in (0, 50]$, $\nu \in (0, 2]$.

As our data stems from a fairly limited area, we do not necessary expect any spatial maxima to be independent. For that reason, a Schlather model seems appropriate, as it describes processes that never reach full independence. Therefore, we mainly visualize results for the Schlather model in this section, while additional figures can be found in Appendix A. Figure 7 shows the estimations from both approaches for the case of a powered exponential model. It is visible that both estimations are quite consistent, in estimating very similar extremal coefficient functions with only minor fluctuations throughout the months. In addition, the results also seem to be similar across the different years, with small variations. The same goes for the parameter estimations of the $EN_{\lambda,\nu}$ approach.

In order to quantify the model fit, we propose to use the so-called *logarithmic score*, which is a strictly proper scoring rule[9] and can be interpreted as the Kullback-Leibler divergence for a sample observation. For the purpose of this application, one can calculate the logarithmic score by calculating the parametrized bivariate density of the realized observation:

$$\text{LogS}(F, \boldsymbol{z_1}, \boldsymbol{z_2}) = -\log f(\boldsymbol{z_1}, \boldsymbol{z_2}; \boldsymbol{\gamma}),$$

where F is a predictive CDF of a max-stable process with bivariate density f (dependent on the parameter vector $\boldsymbol{\gamma}$) and $\boldsymbol{z_1}, \boldsymbol{z_2}$ are observed processes. This calculation results in a logarithmic score for each summer month and year, aggregated in Table 4. The lowest score is obtained by the pairwise likelihood approach and the Brown-Resnick model, although the differences between all scores are quite low. It was anticipated that the pairwise likelihood approach would perform very well here, as it is based on the bivariate density. However, as this method depends on assuming an underlying model and might be biased due to the small sample size, additional evaluation tools should be considered.

For that matter, we employ the so-called (binned) F-madogram (Cooley et al., 2006), which is typically used in practice for estimating the spatial dependence and is directly related to the pairwise extremal coefficient function. Since the

---

[8]This corresponds to a maximum spatial separation between two locations of $\|\boldsymbol{h}_{max}\| \approx 42.5$.

[9]For a more precise characterization see Gneiting and Raftery (2007).

	PL	**CNN**	**EN$_{\lambda,\nu}$**
Brown-Resnick	**5.890**	5.948	6.321
Powered exponential	6.008	5.987	6.020
Whittle-Matérn	-	6.013	6.030

Table 4: The table shows the mean logarithmic score for three different methods and two models across the observed test data. For the EN$_{\lambda,\nu}$ the mean prediction is used. The lowest score is highlighted in bold. The pairwise likelihood score is not always computable, as the methods predicts values that lead to a ill-defined density.
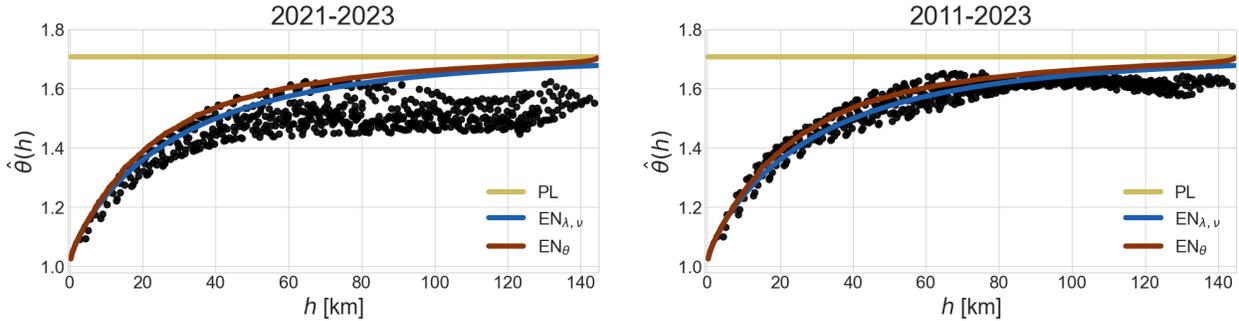
Figure 8: The figure shows the different estimates for the extremal coefficient function. The black dots are the binned F-madogram estimates and the lines correspond to the pointwise mean of the estimated extremal coefficient functions. The left panels shows F-madogram estimate with data from 2021-2023 and the right panel with data from 2011-2023.

evaluation period (2021-2023) is quite short, the estimator will likely be highly biased due to small sample size. Therefore, we provide an additional estimate that is based on the years 2011-2023. For each model the pointwise mean of $\theta(h)$ is taken across all three summer months, to compare against the F-madogram. The different estimates for $\theta(h)$ are visualized in Figure 8.

The results suggest that our model is able to correctly approximate the spatial dependence in the data by fitting a parametric max-stable model or a model-free extremal coefficient function, depending on the approach. As the pairwise likelihood estimator does not lead to sensible results due to optimization issues in small sample sizes, it does not provide a usable model fit. For the powered exponential model, visualized in Figure 8, both approaches seem to be fairly consistent and fit well to the empirical madogram estimation. In addition, both methods are able to produce prediction intervals for $\theta(h)$. Figure 9 shows a selected prediction for the spatial dependence of the extreme precipitation event in July 2021 with corresponding confidence bands. The function estimated by EN$_\theta$ shows a steeper ascent as compared to the estimation by EN$_{\lambda,\nu}$, indicating a higher spatial independence for the precipitation across the distance h. In addition it seams to stagnate for large distances, while the function estimated by EN$_{\lambda,\nu}$ is still growing, indicating that full spatial independence of precipitation maxima is only reached for locations with spatial separation far extending the data domain. Additional visualizations of other max-stable models can be found in Appendix A, while Appendix F includes a comparison of simulated processes and actual observations.
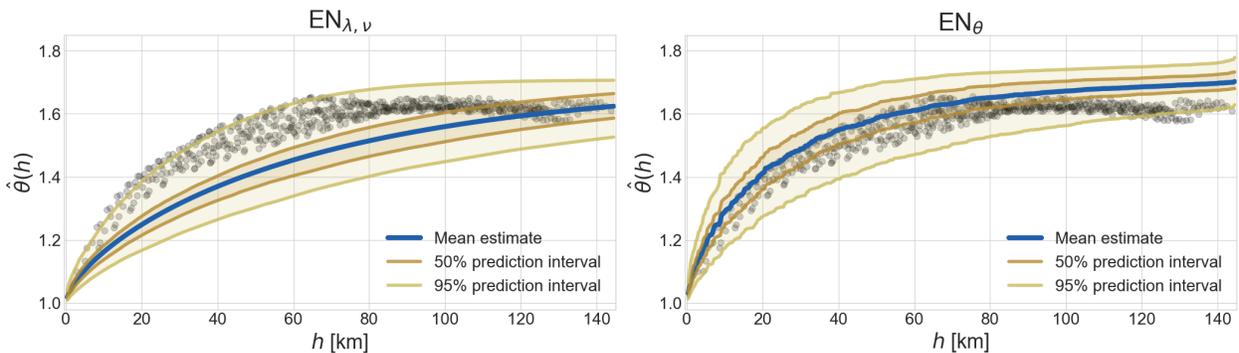

Figure 9: The figure shows estimations of the extremal coefficient function for July 2021. The black dots display the madogram estimate over the last ten years as a reference.

17

5 Discussion

In this work, we introduce a generative neural network-based approach for analyzing spatial extremes by estimating the full parameter distribution of max-stable models, where the likelihood function is not available. With the proposed approach, one can perform parameter estimation, where classical theory tends to fail, due to an infeasible likelihood, under very small sample sizes. Our method is not only able to predict parameters of the underlying process, but predicts a whole parameter distribution. This allows for uncertainty quantification with no additional cost, which is an important step in replacing classical statistical methods that would provide standard errors or other measures of uncertainty. Finally, we extend our approach to directly predict the distribution of the pairwise extremal coefficient function, which is an important tool in assessing the spatial dependence of the underlying process. With this extension no underlying parametric model needs to be imposed and the functional and corresponding uncertainty can be estimated directly.

We demonstrate the effectiveness of the two neural networks in several simulation studies of different max-stable processes and provide several evaluation procedures, covering uncertainty quantification, as well as evaluation of the predicted spatial dependence. In comparison with different benchmarks our models show preferable metrics across different scenarios. To validate the robustness of the approach, several analyses are performed under the assumption of model misspecification. The corresponding results and metrics suggest that in several scenarios the networks are able to project and extrapolate onto the true model, even if it was unseen during the training phase. By comparing and analyzing the two different methods, we provide a general framework that can be adjusted to the specific scenario and needs, for example if one is directly interested in parameter dependent functionals of the model. While both approaches outperform the implemented benchmarks, the $\text{EN}_{\lambda,\nu}$ provides better interval estimations in general, while the EN_θ leads to a significantly smaller Energy Score. To validate our approach, we provide an empirical study of monthly summer precipitation maxima across Germany. Although in that setting the true model is unknown, we present several ways to verify the predictions of the networks, specifically regarding the estimation of the spatial dependence. While this application is mainly a demonstration of the models capabilities and might lack further evaluation compared to other literature regarding extreme event analysis, it highlights that the developed method is suitable for estimating spatial environmental processes and seems to have multiple benefits compared to classical methods.

While this work focused on specific types of max-stable processes, the general approach can easily be modified for other relevant intractable models, such as three dimensional max-stable processes or epidemiological models (Lawson, 2018), as well as other parameter-dependent functionals of interest. Changing the underlying process would usually just require a change in the network architecture, whereas the general approach stays the same. This indicates a promising area of research, as neural networks might be increasingly useful for parameter estimation in complex settings. Especially with more sophisticated simulations methods, such as Markov Chain Monte Carlo, one could extend the approach to highly complex physical processes. Extending the approach to a more general application independent framework is therefore a promising direction of research. Another step in that direction is to discover ways to make the approach more automatic. For example, one could drop the need to specify a prior parameter range by implementing some iterative approach that converges automatically to the best estimation. An extension for the proposed network would be quite natural, as the predicted parameter samples can be used to simulate new processes in an iterative manner until some stopping criterion is used. This could make the simulation process more efficient (compare Sainsbury-Dale et al., 2022), but an unlimited amount of available samples might be unrealistic from an application point of view. Finally, more research is required in analyzing the extrapolation capacities of the models. While the previous results are highly promising, neural networks can struggle with extrapolation and one should be interested in cases, where and why the estimations fail. While these are all promising directions of further research, the developed framework shows to be very suitable for parameter and functional estimation in complex intractable models and has the possibility to be applied to new scenarios in the near future.

References

Beaumont, M. A., Zhang, W., and Balding, D. J. Approximate Bayesian Computation in Population Genetics. *Genetics*, 162(4):2025–2035, 12 2002.

Bosseler, B., Salomon, M., Schlüter, M., and Rubinato, M. Living with urban flooding: A continuous learning process for local municipalities and lessons learnt from the 2021 events in germany. *Water*, 13(19):2769, 2021.

Brehmer, J. R. and Strokorb, K. Why scoring functions cannot assess tail properties. *Electronic Journal of Statistics*, 13 (2):4015 – 4034, 2019.

Castruccio, S., Huser, R., and Genton, M. G. High-order composite likelihood inference for max-stable distributions and processes. *Journal of Computational and Graphical Statistics*, 25(4):1212–1229, 2016.

Chen, J., Janke, T., Steinke, F., and Lerch, S. Generative machine learning methods for multivariate ensemble post-processing. *working paper*, 2022. URL `https://arxiv.org/pdf/2211.01345.pdf`.

Cooley, D., Naveau, P., and Poncet, P. Variograms for spatial max-stable random fields. In *Dependence in Probability and Statistics*, volume 187, pages 373–390. 2006.

Creel, M. Neural nets for indirect inference. *Econometrics and Statistics*, 2:36–49, 2017.

Davis, R. A., Klüppelberg, C., and Steinkohl, C. Statistical inference for max-stable processes in space and time. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 75(5):791–819, 2013.

Davison, A. C. and Gholamrezaee, M. M. Geostatistics of extremes. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 468(2138):581–608, 2012.

Davison, A. C. and Huser, R. Statistics of extremes. *Annual Review of Statistics and Its Application*, 2(1):203–235, 2015.

Davison, A. C., Padoan, S. A., and Ribatet, M. Statistical modeling of spatial extremes. *Statistical Science*, 27(2), 2012.

de Haan, L. and Ferreira, A. *Extreme Value Theory: An Introduction*. SpringerLink Bücher. Springer New York, New York, NY, 2006.

Erhardt, R. J. and Smith, R. L. Approximate bayesian computing for spatial extremes. *Computational Statistics & Data Analysis*, 56(6):1468–1481, 2012.

Fearnhead, P. and Prangle, D. Constructing summary statistics for approximate bayesian computation: semi-automatic approximate bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74 (3):419–474, 2012.

Forster, C. and Oesting, M. Non-stationary max-stable models with an application to heavy rainfall data. *working paper*, 2022. URL `https://arxiv.org/pdf/2212.11598.pdf`.

Franks, J. J. Handbook of approximate bayesian computation. *Journal of the American Statistical Association*, 115 (532):2100–2101, 2020.

Gerber, F. and Nychka, D. Fast covariance parameter estimation of spatial gaussian process models using neural networks. *Stat*, 10(1), 2021.

Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

Huser, R. and Davison, A. C. Composite likelihood estimation for the brown-resnick process. *Biometrika*, 100(2): 511–518, 2013.

Huser, R., Davison, A. C., and Genton, M. G. Likelihood estimators for multivariate extremes. *Extremes*, 19(1):79–103, 2016.

Huser, R., Dombry, C., Ribatet, M., and Genton, M. G. Full likelihood inference for max–stable data. *Stat*, 8(1), 2019.

Huser, R., Stein, M. L., and Zhong, P. Vecchia likelihood approximation for accurate and fast inference in intractable spatial extremes models. *working paper*, 2022. URL `https://arxiv.org/pdf/2203.05626.pdf`.

Kabluchko, Z., Schlather, M., and de Haan, L. Stationary max-stable fields associated to negative definite functions. *The Annals of Probability*, 37(5), 2009.

Lawson, A. B. *Bayesian disease mapping: hierarchical modeling in spatial epidemiology*. CRC press, 2018.

Lenzi, A. and Rue, H. Towards black-box parameter estimation. *working paper*, 2023. URL `https://arxiv.org/pdf/2303.15041`.

Lenzi, A., Bessac, J., Rudi, J., and Stein, M. L. Neural networks for parameter estimation in intractable models. *Computational Statistics & Data Analysis*, 185:107762, 2023.

Mohr, S., Ehret, U., Kunz, M., Ludwig, P., Caldas-Alvarez, A., Daniell, J. E., Ehmele, F., Feldmann, H., Franca, M. J., Gattke, C., Hundhausen, M., Knippertz, P., Küpfer, K., Mühr, B., Pinto, J. G., Quinting, J., Schäfer, A. M., Scheibel, M., Seidel, F., and Wisotzky, C. A multi-disciplinary analysis of the exceptional flood event of July 2021 in central Europe – Part 1: Event description and analysis. *Natural Hazards and Earth System Sciences*, 23(2):525–551, 2023.

Oesting, M., Schlather, M., and Friederichs, P. Statistical post-processing of forecasts for extremes using bivariate brown-resnick processes with an application to wind gusts. *Extremes*, 20(2):309–332, 2017.

Pacchiardi, L. and Dutta, R. Likelihood-free inference with generative neural networks via scoring rule minimization. *working paper*, 2022. URL `https://arxiv.org/abs/2205.15784`.

Pacchiardi, L., Adewoyin, R. A., Dueben, P., and Dutta, R. Probabilistic forecasting with generative networks via scoring rule minimization. *Journal of Machine Learning Research*, 25(45):1–64, 2024.

Padoan, S. A., Ribatet, M., and Sisson, S. A. Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, 105(489):263–277, 2010.

Perez, L. and Wang, J. The effectiveness of data augmentation in image classification using deep learning. *working paper*, 2017. URL `https://arxiv.org/pdf/1712.04621.pdf`.

Radev, S. T., Mertens, U. K., Voss, A., Ardizzone, L., and Köthe, U. Bayesflow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(4):1452–1466, 2022.

Rai, S., Hoffman, A., Lahiri, S., Nychka, D. W., Sain, S. R., and Bandyopadhyay, S. Fast parameter estimation of generalized extreme value distribution using neural networks. *working paper*, 2023. URL `https://arxiv.org/pdf/2305.04341.pdf`.

Ramesh, P., Lueckmann, J.-M., Boelts, J., Tejero-Cantero, Á., Greenberg, D. S., Goncalves, P. J., and Macke, J. H. GATSBI: Generative adversarial training for simulation-based inference. In *International Conference on Learning Representations*, 2022.

Ribatet, M. Spatial extremes: Max-stable processes at work. *Journal de la Société Française de Statistique,*, 154(2): 156–177, 2013.

Ribatet, M. Spatialextremes: Modelling spatial extremes, 2022. URL `https://CRAN.R-project.org/package=SpatialExtremes`.

Sainsbury-Dale, M., Zammit-Mangion, A., and Huser, R. Likelihood-free parameter estimation with neural bayes estimators. *working paper*, 2022. URL `https://arxiv.org/pdf/2208.12942`.

Sang, H. and Gelfand, A. E. Continuous spatial process models for spatial extreme values. *Journal of Agricultural, Biological, and Environmental Statistics*, 15(1):49–65, 2010.

Schlather, M. Models for stationary max-stable random fields. *Extremes*, 5(1):33–44, 2002.

Stephenson, A. and Tawn, J. Exploiting occurrence times in likelihood inference for componentwise maxima. *Biometrika*, 92(1):213–227, 2005.

Thibaud, E., Aalto, J., Cooley, D. S., Davison, A. C., and Heikkinen, J. Bayesian inference for the brown–resnick process, with an application to extreme low temperatures. *The Annals of Applied Statistics*, 10(4), 2016.

A   Additional Visualizations
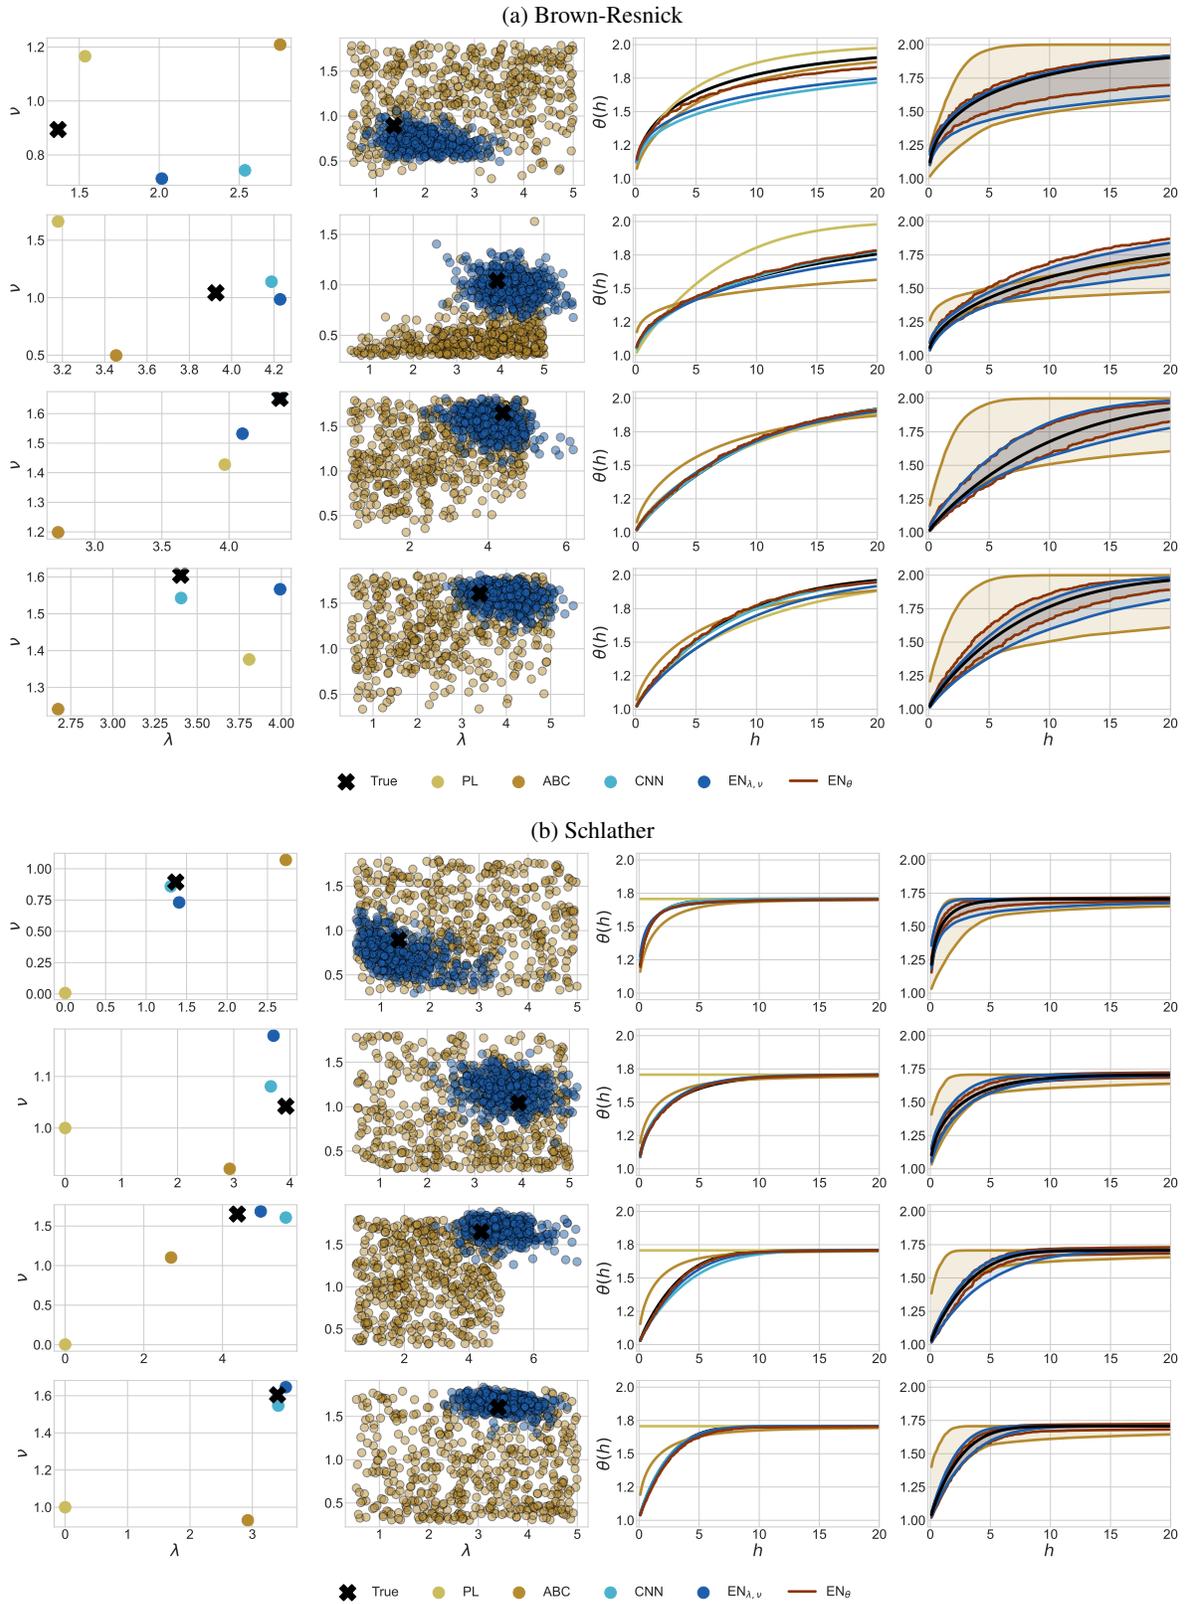
Figure 10: The figure visualizes the different estimation methods for the max-stable models using four randomly drawn test samples. The plot division is the same as in Figure 3.
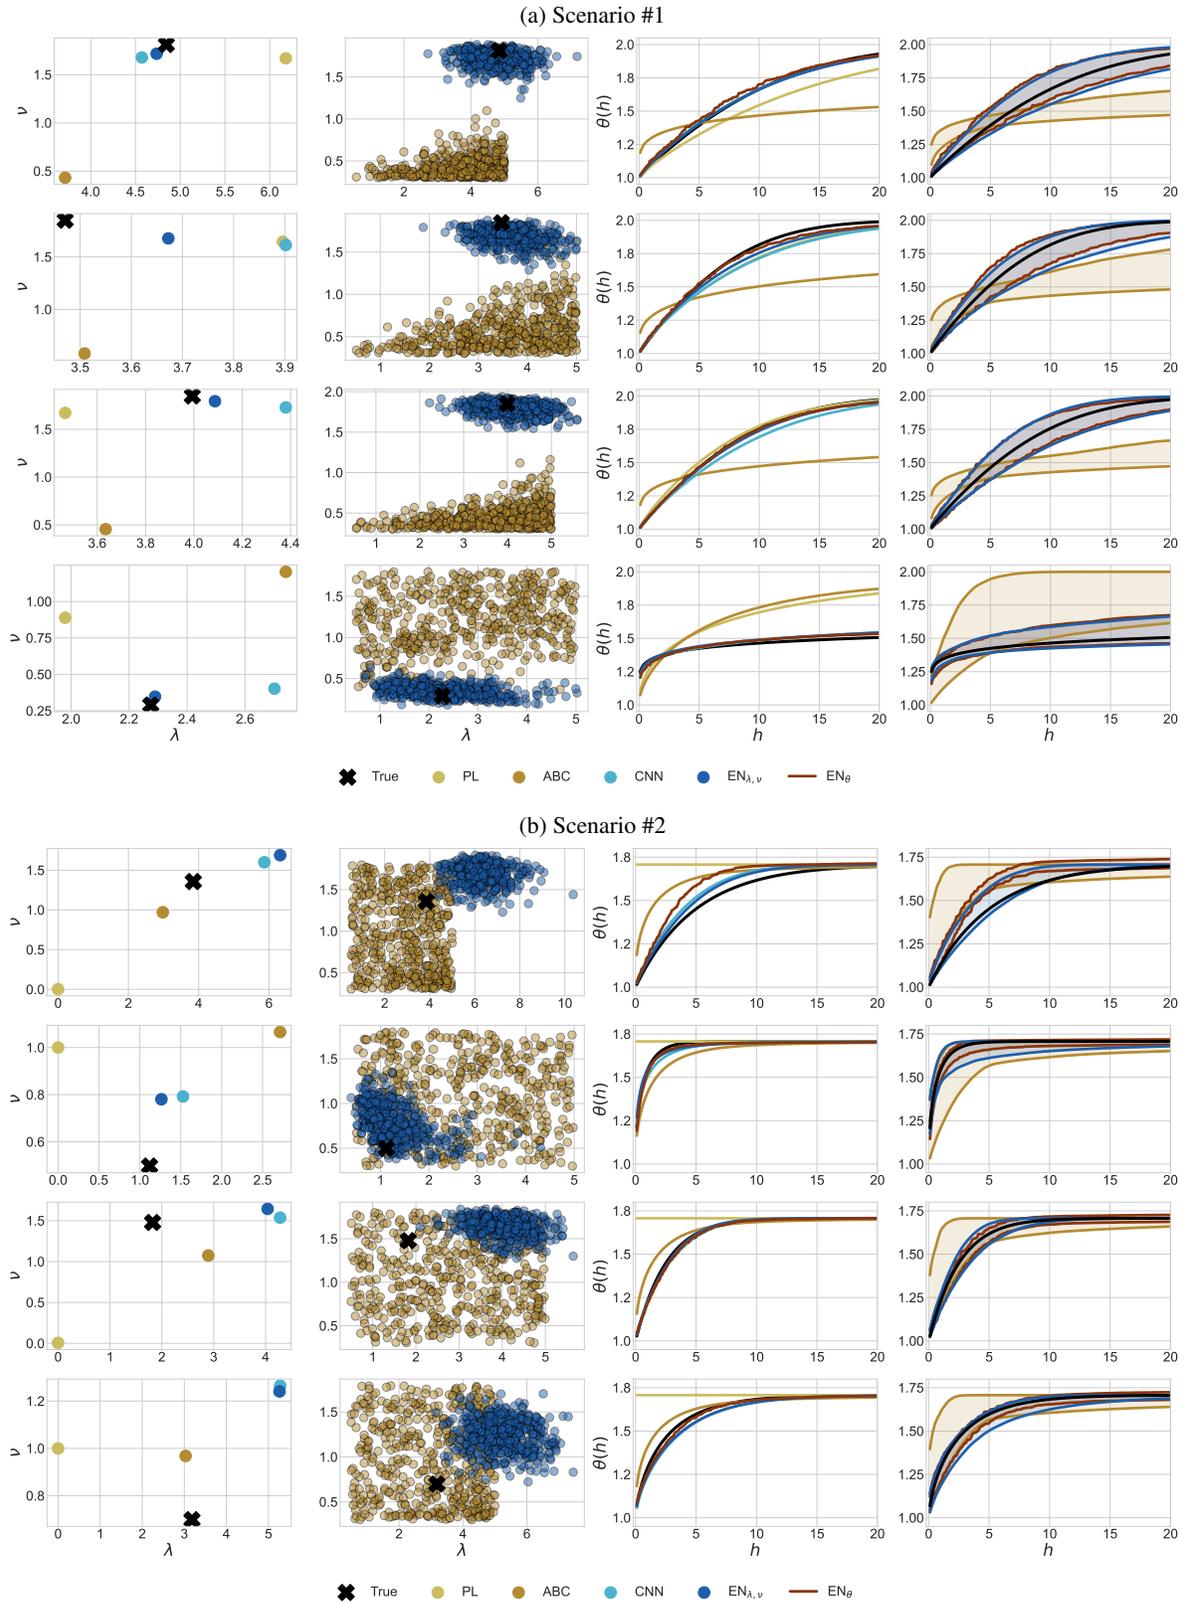
Figure 11: The figure visualizes the different estimation methods for the two robustness scenarios using four randomly drawn test samples. The plot division is the same as in Figure 3.
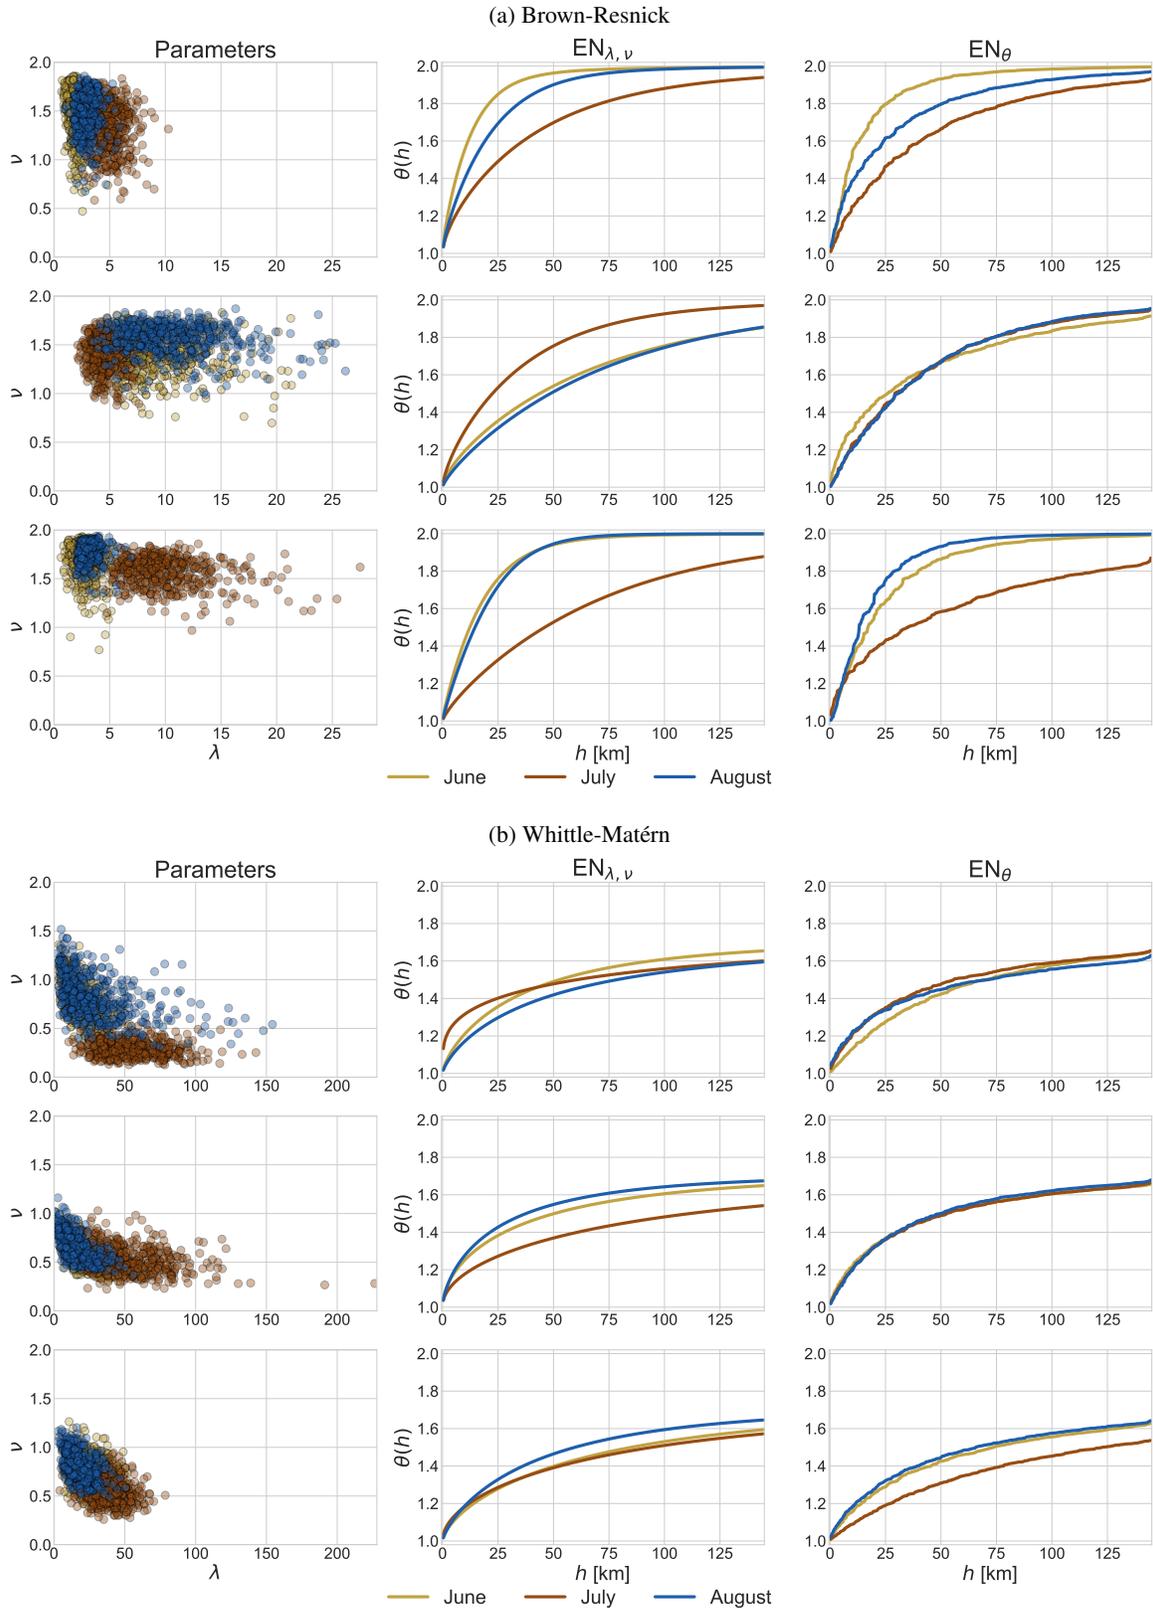
Figure 12: The figure shows the parameter estimates for different models for all three months and years using the proposed approaches. The years are from 2021-2023 from top to bottom.
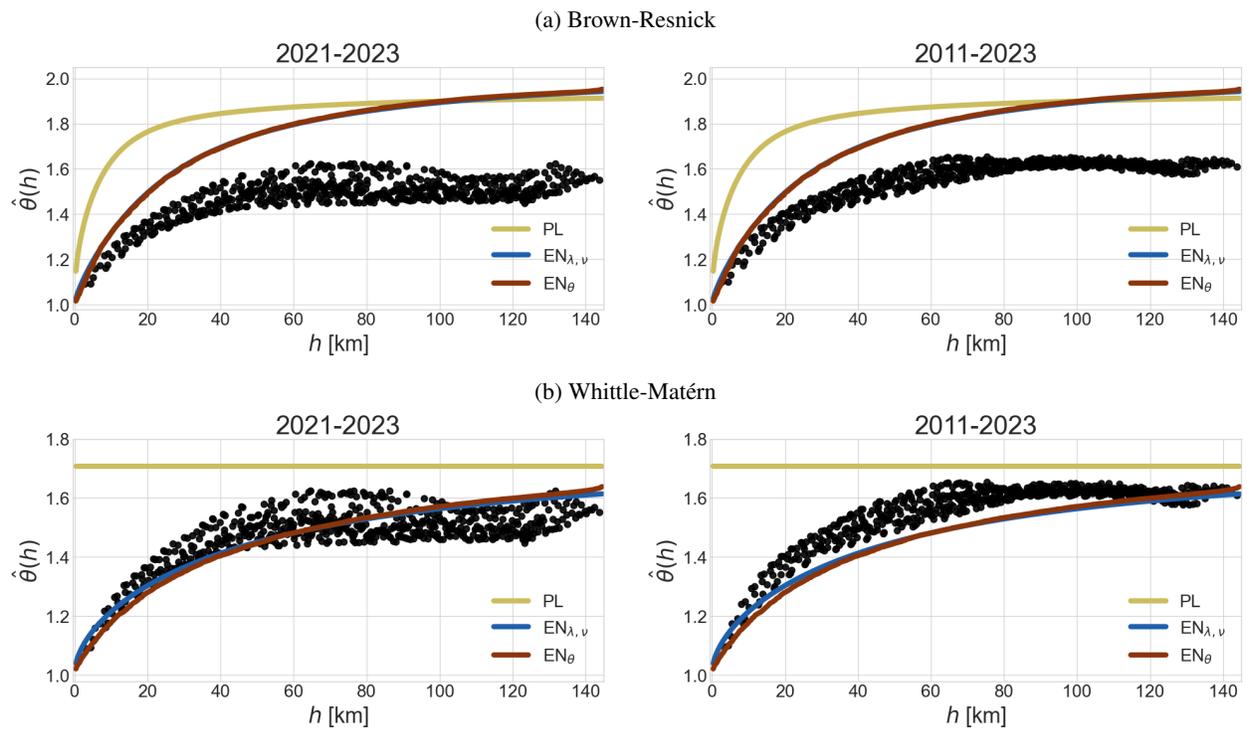
Figure 13: The figure shows the different estimates for the extremal coefficient function. The black dots are the binned F-madogram estimates and the lines correspond to the pointwise mean of the estimated extremal coefficient functions. The left panels shows F-madogram estimate with data from 2021-2023 and the right panel with data from 2011-2023.

B Retaining monotonicity

For the direct estimation of $\boldsymbol{\theta}(h_{disc})$ with EN_θ it can not longer be guaranteed that $\hat{\theta}_r(h_i) \leq \hat{\theta}_r(h_j)$ for $i \leq j$, $\forall r = 1, \ldots, k$. The problem can be solved by sorting the estimated values in ascending order. Since the estimated values can fluctuate greatly, sorting for all $r = 1, \ldots, k$ leads to an increase in the function at the upper bound of the discretization (here 42.5). Figure 14 visualizes this problem and also shows a possible solution. Instead of sorting first all of the values, the functional of interest, e.g. mean or a certain quantile, is computed and afterwords the sorting of for example $\overline{\hat{\theta}(h_1)}, \ldots, \overline{\hat{\theta}(h_k)}$ is performed. This way the increase of the function near the margin is flattened and the estimated functional of the pairwise extremal coefficient function is still monotone.
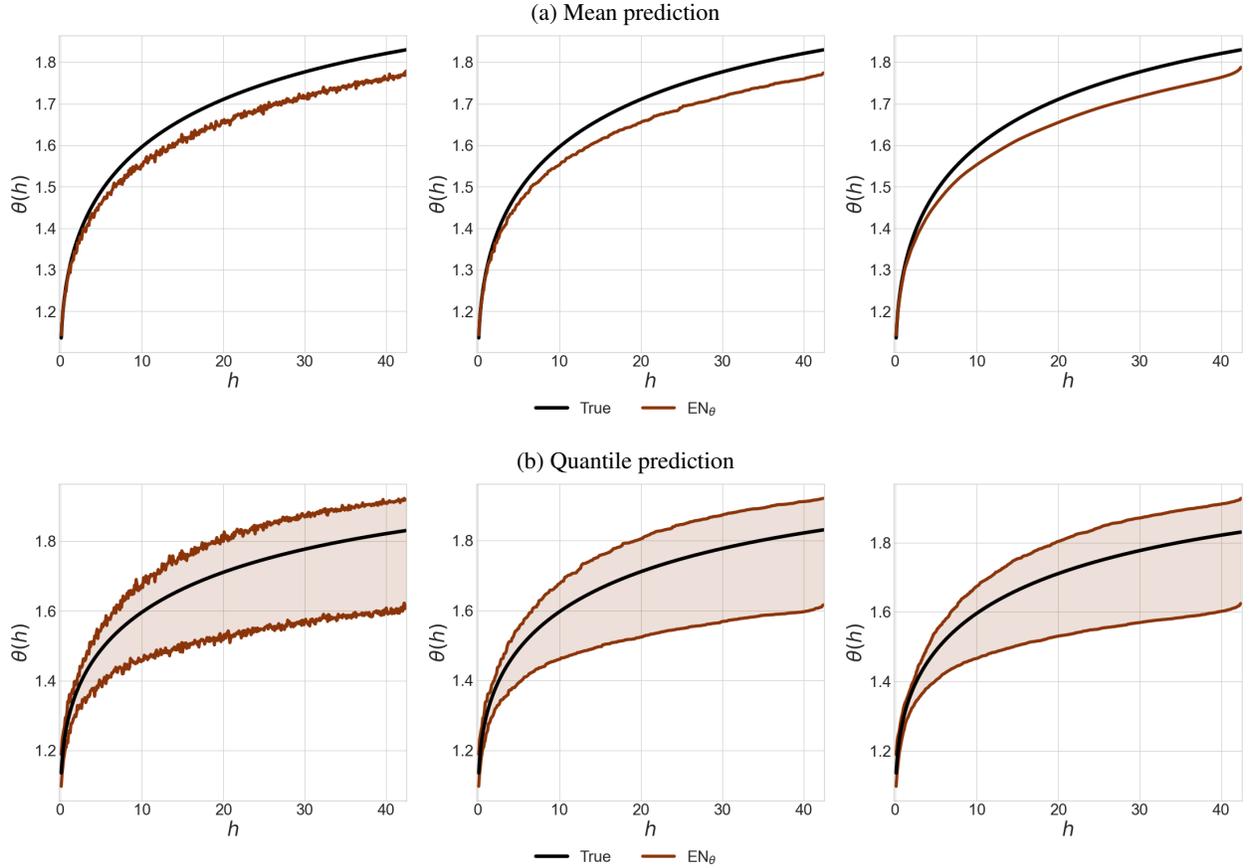

Figure 14: Estimates of pairwise extremal coefficient function of EN_θ with different ways of sorting. On the left, the values $\hat{\theta}(h_i)$, $i = 1, \ldots, k$ are not sorted, only the functionals are calculated of the given sample. In the middle the functional is calculated and afterwards the values are sorted in ascending order, while on the right the values are first sorted and then the functional is calculated.

C   Energy Score analysis of test data

Similarly to Figure 5 in section 3 a visualization of the energy scores for single observations for the Brown-Resnick model is given in Figure 15 and for robustness scenario #1 in Figure 16. As in the case of the Schlather model with powered exponential kernel for EN_θ high energy scores don't arise for certain values of λ and ν while for $EN_{\lambda,\nu}$ the scores are higher on the margins of the test parameter range, especially in the upper left and the lower right corner, as visualized in Figure 15.
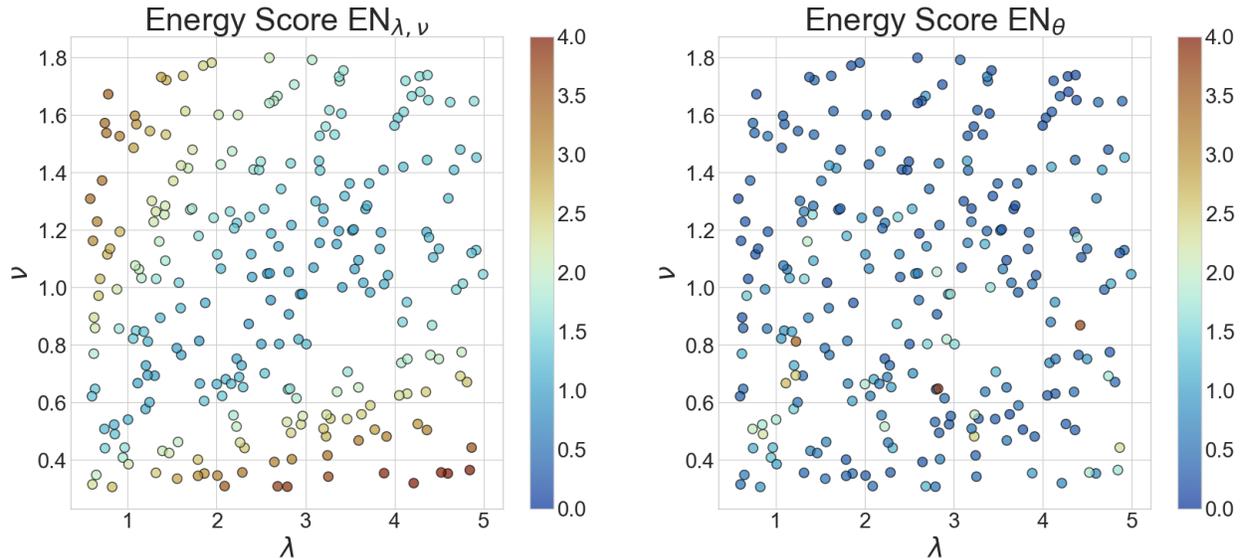
Figure 15: Energy score for parameters (λ, ν) and regular test data from the Brown-Resnick model.

In the case of robustness scenario #1, where the range of trained and tested parameters differ, the Energy score is larger for test parameters further away from the training parameters. For the EN_θ parameters especially in the upper right corner correspond to higher energy scores. A possible explanation can be obtained by considering the values of the pairwise extremal coefficient function in Figure 17. The first column shows that there exists nearly a full dependency, which is not the case for lower values of λ and ν. For the $EN_{\lambda,\nu}$ the energy scores look similar to those in Figure 15 although the values are slightly higher.
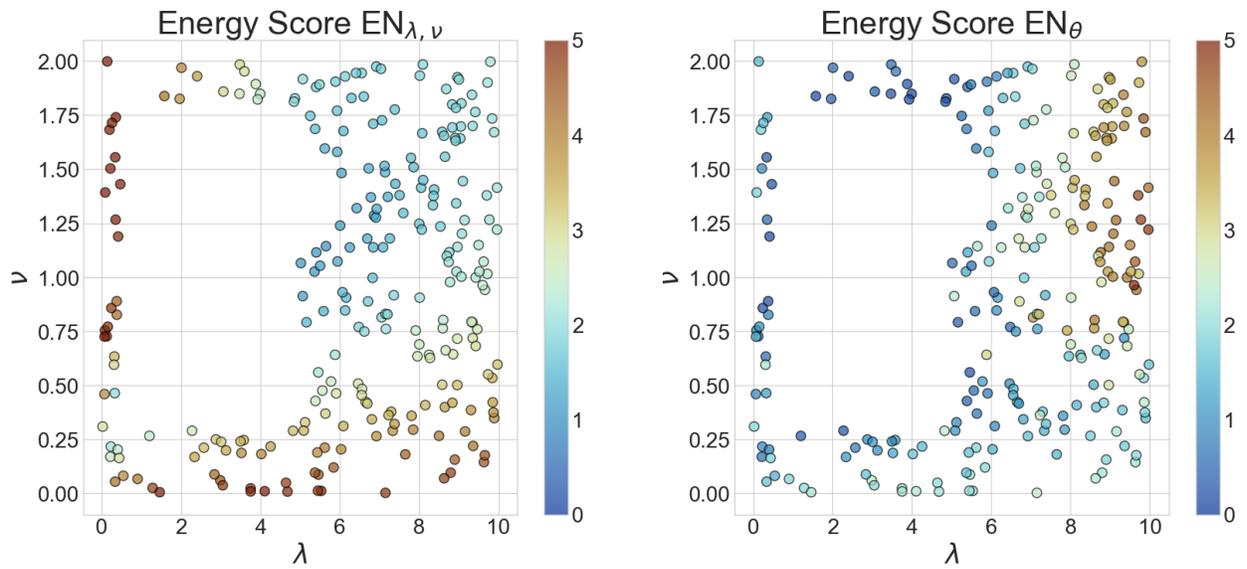
Figure 16: Energy scores for parameters (λ, ν) and robustness scenario #1 with test data from the Brown-Resnick model.
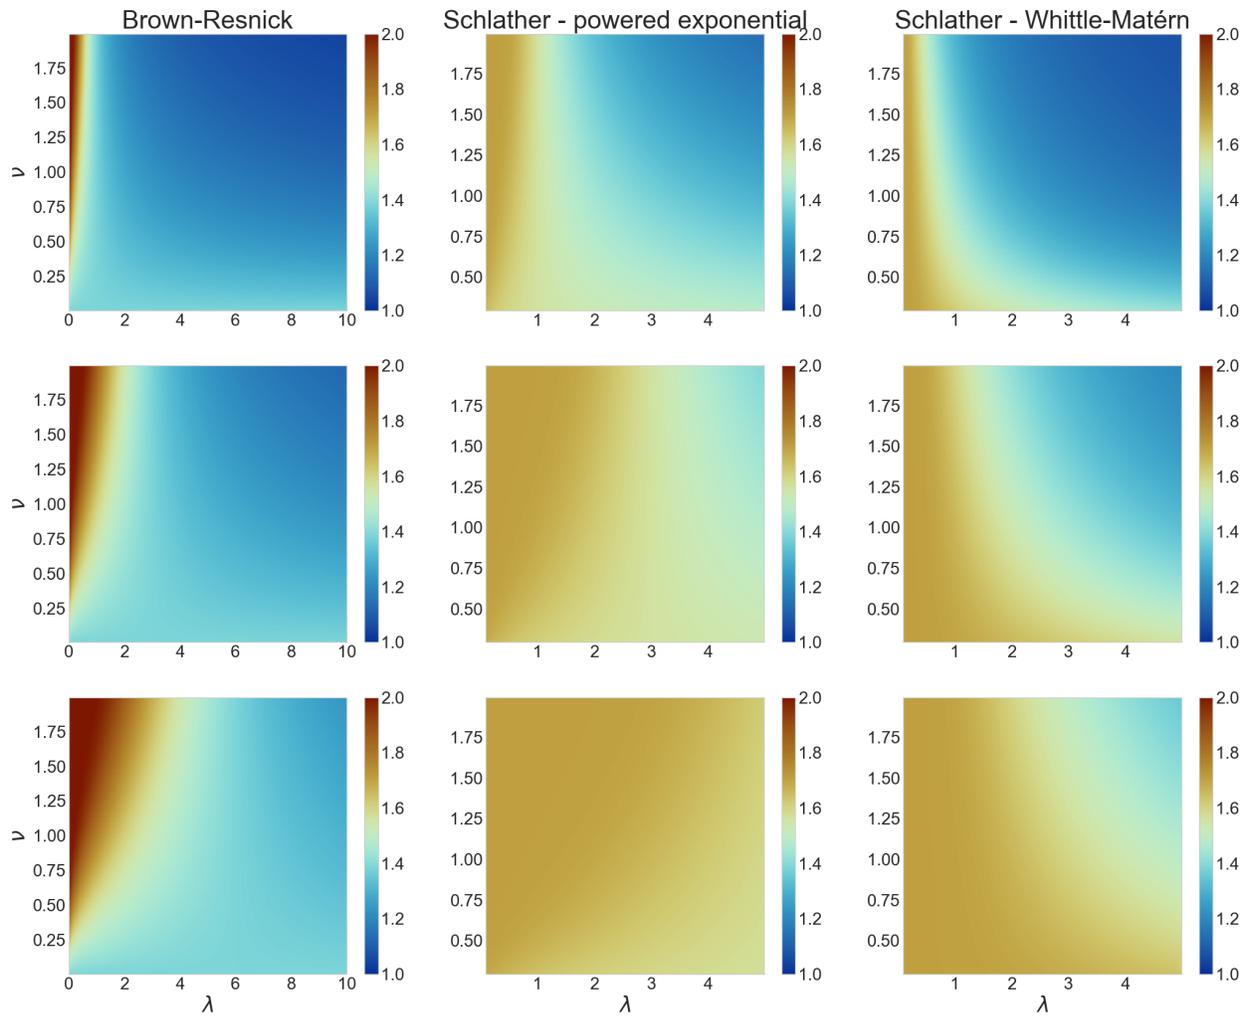
Figure 17: Visualization of the pairwise extremal coefficient function in dependence of λ and ν for different models and distances h. In top row h is set to 1, in the middle $h = 3$ and in the bottom row $h = 6$.

D   Additional robustness scenario

In this additional robustness scenario, we investigate whether the methods are able to correctly estimate an over-parametrized model. For that purpose, the methods are trained on a Brown-Resnick model, while the true processes stem from a Smith model, which can be seen as a special case of the prior. The bivariate CDF of a Brown-Resnick process is given by

$$\mathbb{P}(\boldsymbol{Z}(\boldsymbol{x}_1) \leq z_1, \boldsymbol{Z}(\boldsymbol{x}_2) \leq z_2) = \exp\left(-\frac{1}{z_1}\Phi\left(\frac{a}{2} + \frac{1}{a}\log\left(\frac{z_2}{z_1}\right)\right) - \frac{1}{z_2}\Phi\left(\frac{a}{2} + \frac{1}{a}\log\left(\frac{z_1}{z_2}\right)\right)\right),$$

with $a^2 = 2\gamma(\boldsymbol{h}) = 2(\|\boldsymbol{h}\|/\lambda)^\nu$. In the Smith model it holds that $a^2 = \boldsymbol{h}^\top\Sigma^{-1}\boldsymbol{h}$ with covariance matrix $\Sigma \in \mathbb{R}^{d\times d}$. Now consider a Smith model with diagonal covariance matrix $\Sigma = \text{diag}(\sigma)$. Then $a^2 = \|\boldsymbol{h}\|^2/\sigma$, which corresponds to a Brown-Resnick process with $\nu = 2$ and $\lambda = \sqrt{2\sigma}$. A good estimator should be able to cope with the additional degree of freedom and therefore always predict the smoothness parameter as $\nu = 2$. Again, a test set of size $n = 250$ is simulated based on the Smith model with $\sigma \sim \mathcal{U}(0.5, 5)$, while the different methods are optimized under the assumption of a Brown-Resnick model with $\lambda \sim \mathcal{U}(0.5, 5)$, $\nu \sim \mathcal{U}(0, 2)$.

The results in Table 5 show that the $\text{EN}_{\lambda,\nu}$ has the lowest error for all parameter-related metrics. Especially the error of ν is significantly lower as compared to the other methods, indicating that the $\text{EN}_{\lambda,\nu}$ is able to correctly identify the fixed parameter $\nu = 2$. The CNN and PL approach also display fairly low metrics, while the ABC method does not seem to provide adequate predictions. The EN_θ leads to significantly improvements regarding the metrics involving the extremal coefficient function. Figure 18 shows a visualization of the estimates. All methods except ABC provide fairly accurate parameter point estimates. Furthermore, the predictive distribution of the $\text{EN}_{\lambda,\nu}$ seems to be very concentrated, indicating that the model is confident in the true parameter lying in that range. This is also reflected in the narrow predictive intervals for the extremal coefficient function. These results indicate that the $\text{EN}_{\lambda,\nu}$ is able to provide valid predictions even if operating with an additional degree of freedom. The same goes for the EN_θ, although the predictive intervals are marginally wider than of the $\text{EN}_{\lambda,\nu}$.

		PL	CNN	ABC	$\text{EN}_{\lambda,\nu}$	EN_θ
	MSE_λ	0.23 (1.12)	0.15 (0.20)	0.89 (1.11)	**0.11 (0.21)**	-
	MSE_ν	0.16 (0.39)	0.05 (0.05)	0.81 (0.68)	**0.02 (0.02)**	-
	$\text{MSE}_{\theta(h)}$	0.25 (0.84)	0.08 (0.07)	1.83 (2.31)	0.04 (0.04)	**0.03 (0.03)**
Scenario	$\text{IS}_{0.05,\lambda}$	-	-	4.22 (0.89)	**1.43 (1.34)**	-
#3	$\text{IS}_{0.05,\nu}$	-	-	9.28 (14.58)	**2.18 (1.20)**	-
	$\text{IIS}_{0.05}$	-	-	79.17 (146.55)	**3.36 (3.56)**	3.87 (2.98)
	ES	-	-	0.89 (0.40)	**0.20 (0.13)**	-
	ES_θ	-	-	2.11 (0.40)	0.62 (0.34)	**0.35 (0.14)**

Table 5: The table shows the evaluation metrics across the different scenarios for robustness checks. All metrics are negatively oriented, with the best model highlighted in bold and standard deviation given in brackets.
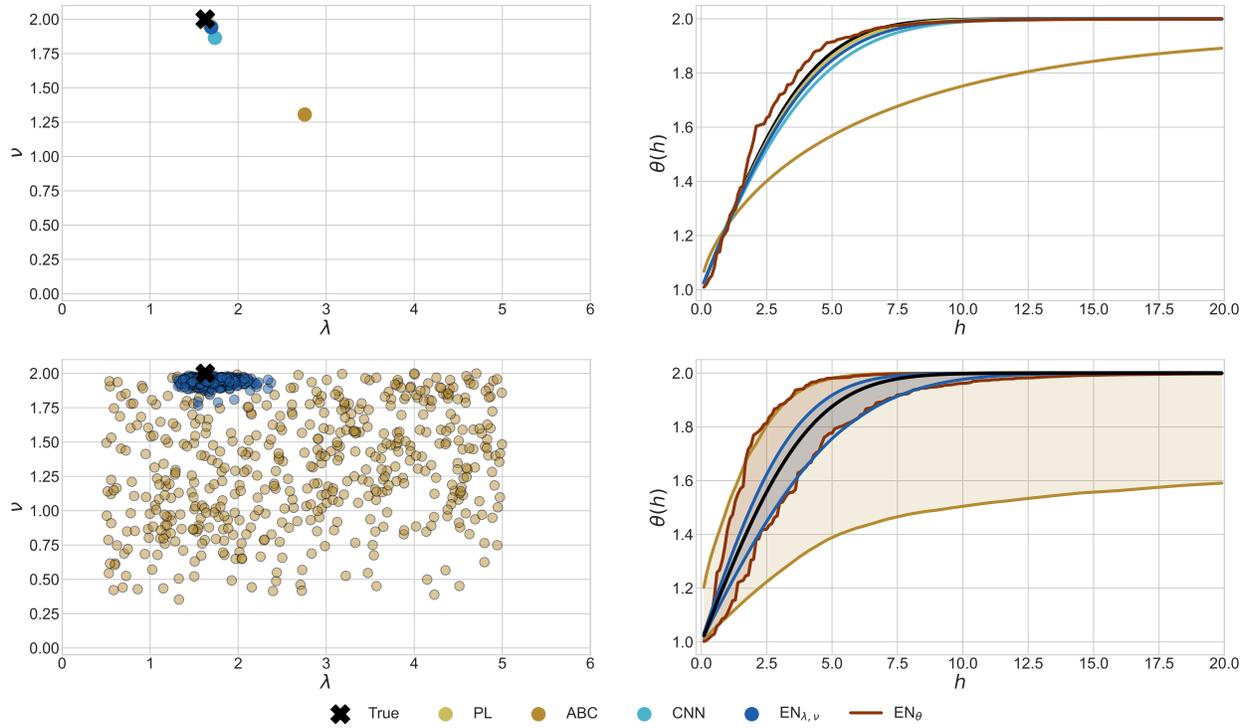
Figure 18: The figure visualizes the different estimation methods for the additional robustness scenario using a selected test sample $((\lambda, \nu) = (1.63, 2.00))$. The plot division is the same as in Figure 3.
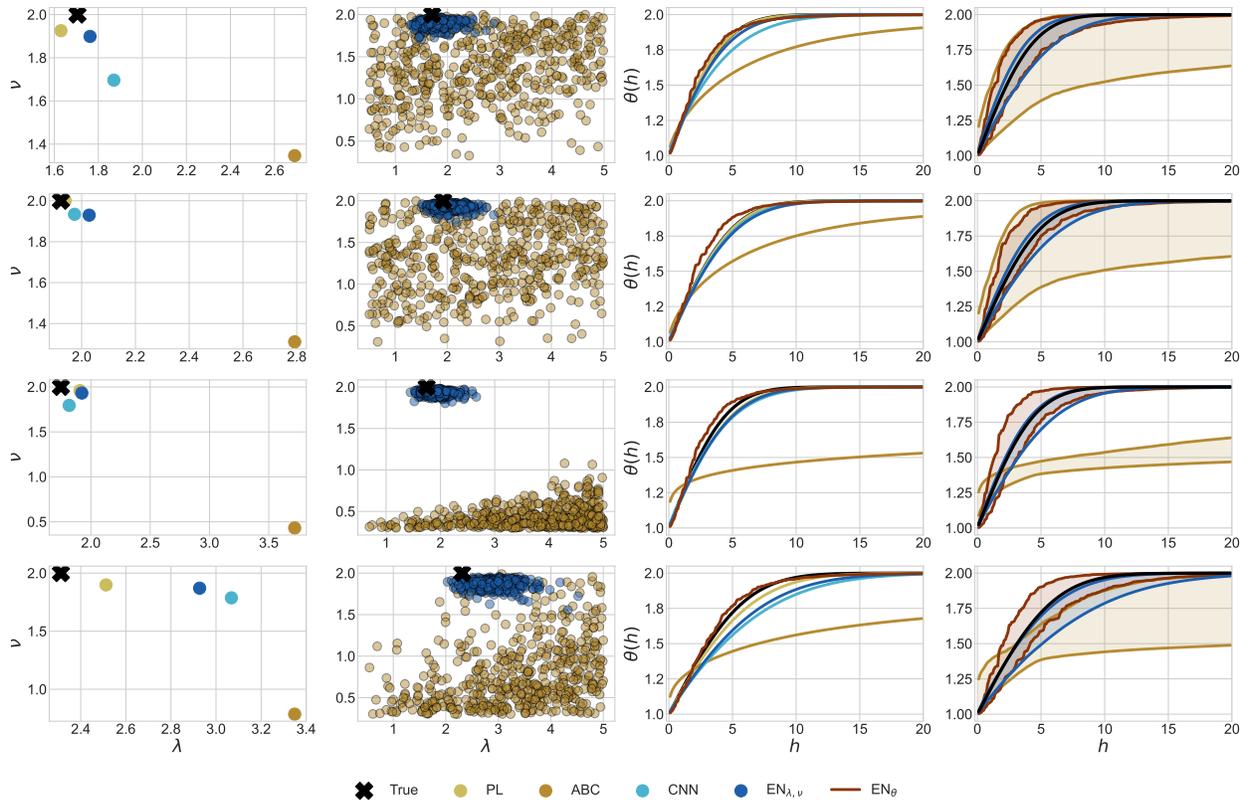

Figure 19: The figure visualizes the different estimation methods for the additional robustness scenario using four randomly drawn test samples. The plot division is the same as in Figure 3.

31

E   GEV fit

The GEV parameters are modeled by the following equations

$$\mu(i,t) = \beta_{0,\mu} + \beta_{1,\mu}\text{lat}(i) + \beta_{2,\mu} + \text{lon}(i) + \beta_{3,\mu}t$$
$$\sigma = \beta_{0,\sigma}$$
$$\gamma = \beta_{0,\gamma},$$

where $i = 1, \ldots, 900$ is the index of the corresponding location, t is the year of the observation and lat, lon describe the latitude and longitude, respectively. While this model can only describe linear relationships across the covariates, this usually suffices in practice, although in principle more sophisticated approaches are possible. The model is fitted using the function *fitspatgev* of the *SpatialExtremes* package (Ribatet, 2022).

The estimated parameters and the corresponding standard errors are shown in Table 6. First of, note that the shape

	$\beta_{0,\mu}$	$\beta_{1,\mu}$	$\beta_{2,\mu}$	$\beta_{3,\mu}$	$\beta_{0,\sigma}$	$\beta_{0,\gamma}$
Estimation	64.7996	-0.9997	0.0149	0.0011	7.0045	0.1052
Standard error	27.3724	0.5433	0.3192	0.0104	0.2126	0.0224

Table 6: The estimated GEV parameters and corresponding standard errors of the model described above. The parameters were fit on the three summer months over the years of 1931-2020.

parameter is estimated as $\gamma = 0.1052 > 0$, which indicates that the data can best be described using a Fréchet distribution. This makes sense, since the Fréchet distribution has a left endpoint, which is reasonable since precipitation can only take nonnegative values. Using the estimates from Table 6, the observed precipitation fields are transformed to unit Fréchet margins.

F   Simulated precipitation over Germany
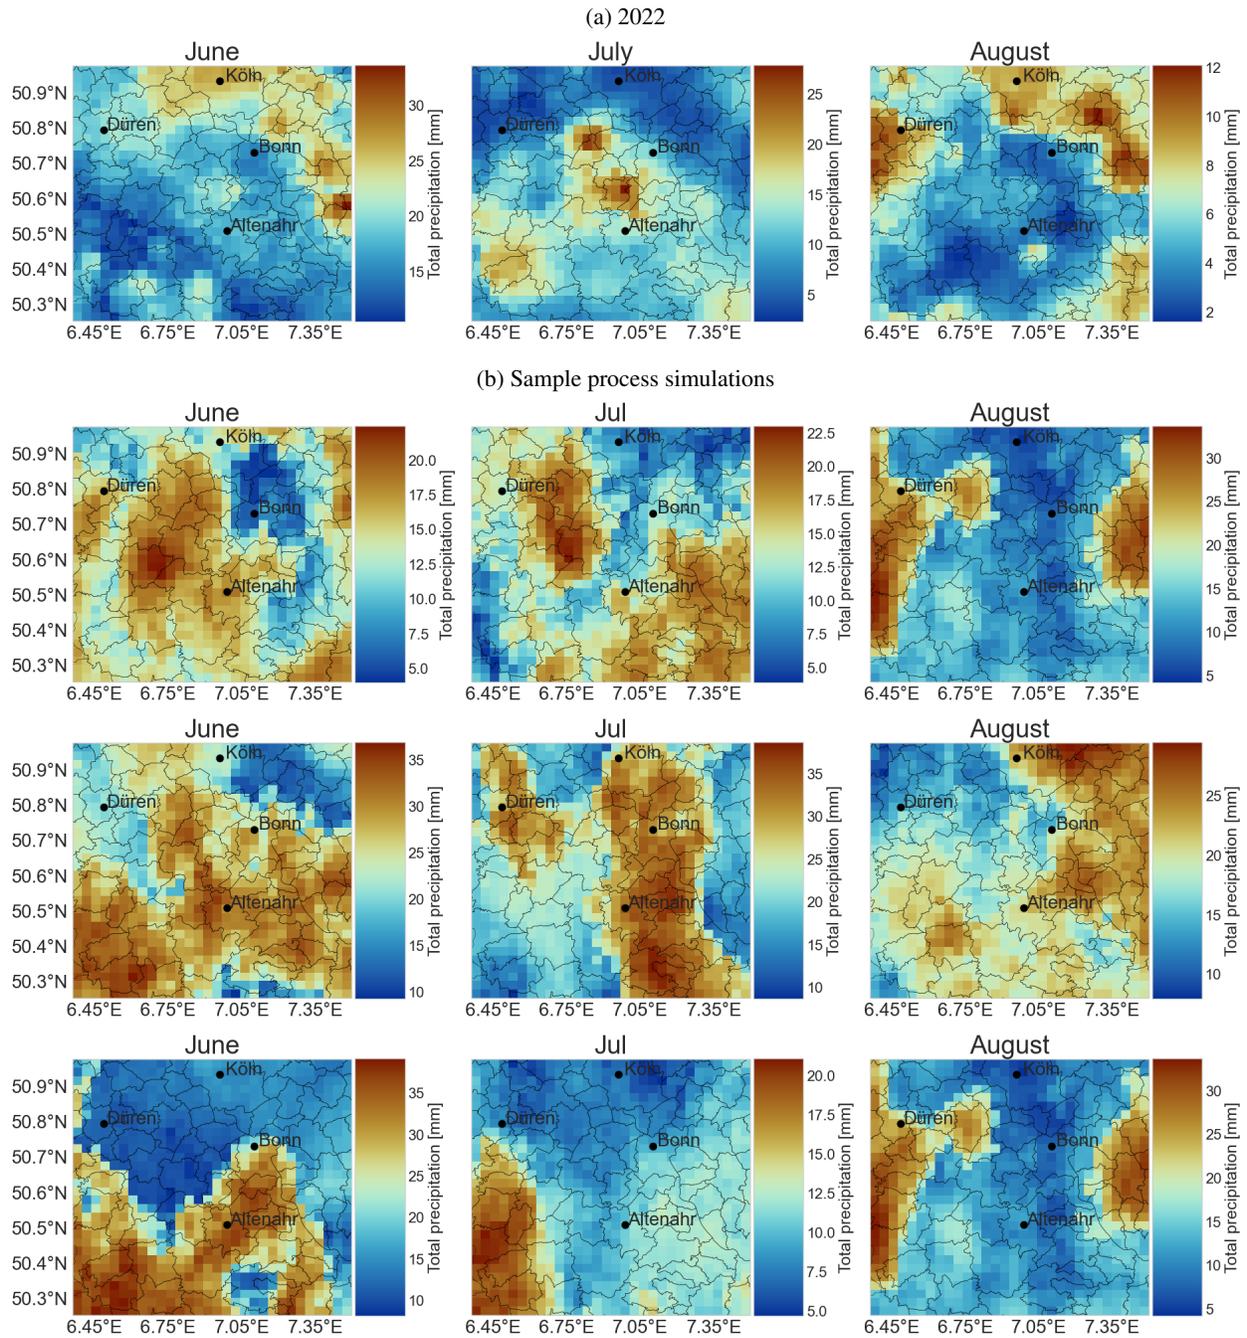
Figure 20: The figure shows the observed precipitation maxima in 2022 (top row) and corresponding simulations from an estimated Schlather model with powered exponential correlation function (bottom rows). The simulations have been transformed back to the original GEV surface.