

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



***National Institute for Applied Statistics Research
Australia***

University of Wollongong, Australia

Working Paper

08-23

**Likelihood-Free Parameter Estimation with
Neural Bayes Estimators**

Matthew Sainsbury-Dale, Andrew Zammit-Mangion, and Raphaël Huser

*Copyright © 2023 by the National Institute for Applied Statistics Research Australia, UOW.
Work in progress, no part of this paper may be reproduced without permission from the Institute.*

National Institute for Applied Statistics Research Australia, University of Wollongong,
Wollongong NSW 2522, Australia T: +61 2 42215076. E: karink@uow.edu.au

Likelihood-Free Parameter Estimation with Neural Bayes Estimators

Matthew Sainsbury-Dale¹, Andrew Zammit-Mangion¹, and Raphaël Huser²

¹School of Mathematics and Applied Statistics, University of Wollongong, Australia

²Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Saudi Arabia

Abstract

Neural Bayes estimators are neural networks that approximate Bayes estimators. They are fast, likelihood-free, and amenable to rapid bootstrap-based uncertainty quantification. In this paper, we aim to increase the awareness of statisticians to this relatively new inferential tool, and to facilitate its adoption by providing user-friendly open-source software. We also give attention to the ubiquitous problem of estimating parameters from replicated data, which we address in the neural network setting using permutation-invariant neural networks. Through extensive simulation studies we demonstrate that neural Bayes estimators can be used to quickly estimate parameters in weakly-identified and highly-parameterised models with relative ease. We illustrate their applicability through an analysis of extreme sea-surface temperature in the Red Sea where, after training, we obtain parameter estimates and bootstrap-based confidence intervals from hundreds of spatial fields in a fraction of a second.

Keywords: amortised inference, deep learning, exchangeable data, extreme-value model, permutation invariant, point estimation, spatial statistics

1 Introduction

The most popular methods for estimating parameters in parametric statistical models are those based on the likelihood function. However, it is not always possible to formulate the likelihood function (Diggle and Gratton, 1984; Lintusaari et al., 2017) and, even when it is available, it may be computationally intractable. For example, popular models for spatial extremes are max-stable processes, for which the number of terms involved in the likelihood function grows more than exponentially fast with the number of observations (Padoan et al., 2010; Huser et al., 2019). One common workaround is to replace the full likelihood with a composite likelihood (e.g., Cox and Reid, 2004; Varin and Vidoni, 2005; Varin et al., 2011), but this usually results in a loss of statistical efficiency (Huser and Davison, 2013; Castruccio et al., 2016), and it is not always clear how one should construct the composite likelihood.

The related Vecchia approximation (Vecchia, 1988) has been applied successfully both in Gaussian (e.g., Stein et al., 2004) and max-stable (Huser et al., 2022) settings, but this approximation still trades statistical efficiency for computational efficiency.

To bypass these challenges, several model-fitting methods have been developed that preclude the need to evaluate the likelihood function. The most popular of these so-called likelihood-free methods is approximate Bayesian computation (ABC; Beaumont et al., 2002; Sisson et al., 2018). In its simplest form, ABC involves sampling parameters from the prior, simulating from the model, and retaining parameters as an approximate sample from the posterior if the simulated data are “similar” to the observed data, with similarity typically assessed by comparing low-dimensional summary statistics. ABC methods are sensitive to the choice of summary statistics, and they are notoriously difficult to calibrate: for example, the number of summary statistics and the tolerance used when comparing statistics affect both the computational efficiency and the statistical efficiency of ABC.

Recently, neural networks have emerged as a promising approach to likelihood-free inference. In this work, we focus on neural networks that map data to parameter point estimates; we refer to such neural networks as *neural point estimators*. Neural point estimators date back to at least Chon and Cohen (1997), but they have only been adopted widely in recent years, for example in applications involving models for stock returns (Creel, 2017); population-genetics (Flagel et al., 2018); time series (Rudi et al., 2021); spatial fields (Gerber and Nychka, 2021; Banesh et al., 2021; Lenzi et al., 2023); spatio-temporal fields (Zammit-Mangion and Wikle, 2020); and agent-based models (Gaskin et al., 2023). The computational bottleneck in neural point estimation is the training procedure, whereby the neural network ‘learns’ a useful mapping between the sample space and the parameter space. Importantly though, this training cost is amortised: for a given statistical model and prior distribution over the parameters, a trained neural point estimator with sufficiently flexible architecture can be re-used for new data sets at almost no computational cost, provided that each data replicate has the same format as those used to train the point estimator (e.g., images of a pre-specified width and height). Uncertainty quantification of the estimates proceeds naturally through the bootstrap distribution, which is essentially available “for free” with neural point estimators since they are so fast to evaluate. As we shall show, neural point estimators can be trained to approximate Bayes estimators and, in this case, we refer to them as *neural Bayes estimators*.

Parameter estimation from replicated data is commonly required in statistical applications, and we therefore give this topic particular attention. Neural point estimation with replicated data is not straightforward; for example, Gerber and Nychka (2021) considered two approaches to handling replicated data, both with drawbacks. In their first approach, they trained a neural point estimator for a single spatial field, applied it to each field independently, and averaged the resulting estimates; we call this the ‘one-at-a-time’ approach. This approach does not reduce the bias commonly seen in small-sample estimators as the sample size increases, and it is futile when the parameters are unidentifiable from a single replicate. In their second approach, they adapted the size of the neural estimator’s input layer to match the number of independent replicates. The resulting estimator is not invariant to the ordering of the replicates; it results in an explosion of neural-network parameters with increasing number of replicates; and it requires a different architecture for every possible sample size, which reduces its applicability and generality.

In this paper, we first clarify the connection between neural point estimators and classical Bayes estimators, which is sometimes misconstrued or ignored in the literature. Second, we propose a novel way to perform neural Bayes estimation from independent replicates by leveraging permutation-invariant neural networks, constructed using the DeepSets framework (Zaheer et al., 2017). To the best of our knowledge, this paper is the first to explore its use and the related practical considerations in a point estimation context. We show that these architectures lead to a substantial improvement in estimation accuracy when compared to those that do not account for replication appropriately. Third, we discuss important practicalities for designing and training neural Bayes estimators and, in particular, we describe a way to construct an estimator that is approximately Bayes for any sample size. For illustration, we estimate parameters in Gaussian and max-stable processes, as well as the highly-parameterised spatial conditional extremes model (Wadsworth and Tawn, 2022). We use the latter model in the analysis of sea-surface temperature extremes in the Red Sea where, using a neural Bayes estimator, we obtain estimates and bootstrap confidence intervals from hundreds of spatial fields in a fraction of a second. A primary motivation of this work is to facilitate the adoption of neural point estimation by statisticians and, to this end, we accompany the paper with user-friendly open-source software in the Julia (Bezanson et al., 2017) and R (R Core Team, 2023) programming languages: our software package, **NeuralEstimators**, is available at <https://github.com/msainsburydale/NeuralEstimators>, and can be used in a wide range of applied settings.

The remainder of this paper is organised as follows. In Section 2, we outline the theory underlying neural Bayes estimators and discuss their implementation. In Section 3, we conduct extensive simulation studies that clearly demonstrate the utility of neural Bayes estimators. In Section 4, we use the spatial conditional extremes model to analyse sea-surface temperature in the Red Sea. In Section 5, we conclude and outline avenues for future research. A supplement is also available that contains more details and figures.

2 Methodology

In Section 2.1, we introduce neural Bayes estimators. In Section 2.2, we discuss the use of permutation-invariant neural networks for neural Bayes estimation from replicated data. In Section 2.3, we describe the general workflow for implementing neural Bayes estimators, and discuss some important practical considerations.

2.1 Neural Bayes estimators

A parametric statistical model is a set of probability distributions on a sample space \mathcal{S} , where the probability distributions are parameterised via some p -dimensional parameter vector $\boldsymbol{\theta}$ on a parameter space Θ (McCullagh, 2002). Suppose that we have data, which we denote as \mathbf{Z} , from one such distribution. Then, the goal of parameter point estimation is to come up with an estimate of the unknown $\boldsymbol{\theta}$ from \mathbf{Z} using an estimator,

$$\hat{\boldsymbol{\theta}} : \mathcal{S} \rightarrow \Theta,$$

which is a mapping from the sample space to the parameter space.

Estimators can be constructed within a decision-theoretic framework. Assume that the sample space is $\mathcal{S} = \mathbb{R}^n$, and consider a non-negative loss function, $L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{Z}))$, which assesses an estimator $\hat{\boldsymbol{\theta}}(\cdot)$ for a given $\boldsymbol{\theta}$ and data set $\mathbf{Z} \sim f(\mathbf{z} | \boldsymbol{\theta})$, where $f(\mathbf{z} | \boldsymbol{\theta})$ is the probability density function of the data conditional on $\boldsymbol{\theta}$. An estimator’s risk function is its loss averaged over all possible data realisations,

$$R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\cdot)) \equiv \int_{\mathcal{S}} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{z})) f(\mathbf{z} | \boldsymbol{\theta}) d\mathbf{z}. \quad (1)$$

So-called Bayes estimators minimise a (weighted) average of (1) known as the Bayes risk,

$$r_{\Omega}(\hat{\boldsymbol{\theta}}(\cdot)) \equiv \int_{\Theta} R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\cdot)) d\Omega(\boldsymbol{\theta}), \quad (2)$$

where $\Omega(\cdot)$ is a prior measure for $\boldsymbol{\theta}$. Note that in (1) and (2) we deviate slightly from classical notation and notate the estimator as $\hat{\boldsymbol{\theta}}(\cdot)$, rather than just $\hat{\boldsymbol{\theta}}$, to stress that it is constructed as a function of the data.

Bayes estimators are theoretically attractive: for example, unique Bayes estimators are admissible and, under suitable regularity conditions and the squared-error loss, consistent and asymptotically efficient (Lehmann and Casella, 1998, Ch. 5, Thm. 2.4; Ch. 6, Thm. 8.3). Further, for a large class of prior distributions, every set of conditions that imply consistency of the maximum likelihood estimator also imply consistency of Bayes estimators (Strasser, 1981). Unfortunately, however, Bayes estimators are typically unavailable in closed form for the complex models often encountered in practice. A way forward is to assume a flexible parametric model for $\hat{\boldsymbol{\theta}}(\cdot)$, and to optimise the parameters within that model in order to approximate the Bayes estimator. Neural networks are ideal candidates, since they are universal function approximators (e.g., Hornik et al., 1989; Zhou, 2018), and because they are fast to evaluate, usually involving only simple matrix-vector operations.

Let $\hat{\boldsymbol{\theta}}(\mathbf{Z}; \boldsymbol{\gamma})$ denote a *neural point estimator*, that is, a neural network that returns a point estimate from data \mathbf{Z} , where $\boldsymbol{\gamma}$ contains the neural-network parameters (i.e., the so-called “weights” and “biases”). Bayes estimators may be approximated with $\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma}^*)$ by solving the optimisation problem,

$$\boldsymbol{\gamma}^* \equiv \arg \min_{\boldsymbol{\gamma}} r_{\Omega}(\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma})). \quad (3)$$

Typically, $r_{\Omega}(\cdot)$ defined in (2) cannot be directly evaluated, but it can be approximated using Monte Carlo methods. Specifically, given a set ϑ of K parameter vectors sampled from the prior $\Omega(\cdot)$ and, for each $\boldsymbol{\theta} \in \vartheta$, J samples from $f(\mathbf{z} | \boldsymbol{\theta})$ collected in the set $\mathcal{Z}_{\boldsymbol{\theta}}$,

$$r_{\Omega}(\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma})) \approx \frac{1}{K} \sum_{\boldsymbol{\theta} \in \vartheta} \frac{1}{J} \sum_{\mathbf{Z} \in \mathcal{Z}_{\boldsymbol{\theta}}} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{Z}; \boldsymbol{\gamma})). \quad (4)$$

Note that (4) does not involve evaluation, or knowledge, of the likelihood function.

The surrogate objective function (4) can be straightforwardly minimised with respect to $\boldsymbol{\gamma}$ using back-propagation and stochastic gradient descent with deep-learning software packages such as **Flux** (Innes, 2018). For sufficiently expressive architectures, the point estimator targets a Bayes estimator with respect to $L(\cdot, \cdot)$ and $\Omega(\cdot)$. We therefore call the fitted

neural point estimator a *neural Bayes estimator*. Note that neural Bayes estimators, like all other neural networks, are function approximators that, at least in theory, can approximate continuous functions arbitrarily well. However, the discrepancy between $\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma}^*)$ and the true Bayes estimator will depend on a number of factors, such as the neural-network architecture and the specific optimisation procedure used to minimise (4). We provide more discussion on these practical considerations in Section 2.3.

Like Bayes estimators, neural Bayes estimators target a specific point summary of the posterior distribution. For instance, the 0–1, absolute-error, and squared-error loss functions lead to neural Bayes estimators that approximate the posterior mode, median, and mean, respectively. Further, posterior quantiles, which can be used to construct credible intervals, may be estimated by using the quantile loss function (e.g., Cressie, 2023, eqn. 7). This important link between neural point estimators and Bayes estimators is often overlooked in the literature.

2.2 Neural Bayes estimators for replicated data

While neural Bayes estimators have been employed, sometimes inadvertently, in various applications, much less attention has been given to neural Bayes estimators for replicated data. In what follows, we denote a collection of m independent and identically distributed replicates $\mathbf{Z}_1, \dots, \mathbf{Z}_m$ as $\mathbf{Z}^{(m)} \equiv (\mathbf{Z}'_1, \dots, \mathbf{Z}'_m)'$. In Section 2.2.1 we show how the so-called DeepSets framework is ideally placed to construct Bayes estimators for replicated data; in Section 2.2.2 we discuss approaches one can adopt to design neural point estimators that are approximately Bayes for any number of replicates; and in Section 2.2.3 we illustrate the neural Bayes estimator for replicated data on a relatively simple problem where the Bayes estimator is known in closed form.

2.2.1 The DeepSets Framework

Parameter estimation from replicated data is commonly required in statistical applications and, as discussed in Section 1, naïve approaches to constructing neural Bayes estimators for replicated data can lead to bias, an explosion in the number of neural-network parameters, and an inability to generalise to different sample sizes. We overcome these issues by leveraging an important property of Bayes estimators for replicated data, namely that they are invariant to permutations of the replicates. In Section S1 of the Supplementary Material, we give a formal proof of this property: in particular, we prove that if a Bayes estimator for replicated data is unique, then it will also be permutation invariant.

The permutation-invariance property of Bayes estimators for replicated data has largely been ignored to date in the context of neural point estimation. We propose enforcing our neural Bayes estimators to be permutation invariant by couching them in the DeepSets framework (Zaheer et al., 2017), a computationally convenient special case of Janossy pooling (Murphy et al., 2019). In this framework, the Bayes estimator is represented by

$$\begin{aligned} \hat{\boldsymbol{\theta}}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}) &= \boldsymbol{\phi}(\mathbf{T}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}_\psi); \boldsymbol{\gamma}_\phi), \\ \mathbf{T}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}_\psi) &= \mathbf{a}(\{\boldsymbol{\psi}(\mathbf{Z}_i; \boldsymbol{\gamma}_\psi) : i = 1, \dots, m\}), \end{aligned} \tag{5}$$

where $\boldsymbol{\psi} : \mathbb{R}^n \rightarrow \mathbb{R}^q$ and $\boldsymbol{\phi} : \mathbb{R}^q \rightarrow \mathbb{R}^p$ are (*deep*) neural networks parametrised by $\boldsymbol{\gamma}_\psi$ and $\boldsymbol{\gamma}_\phi$,




Figure 1: Schematic of the DeepSets representation. Each independent replicate Z_i is transformed independently using the function $\psi(\cdot)$. The set of transformed inputs are then aggregated elementwise using a permutation-invariant function, $\mathbf{a}(\cdot)$, yielding the summary statistic \mathbf{T} . Finally, the summary statistic is mapped to parameter estimates $\hat{\theta}$ by the function $\phi(\cdot)$. Many classical estimators take this form; in this work, we use neural networks to model $\psi(\cdot)$ and $\phi(\cdot)$.

respectively, $\gamma \equiv (\gamma'_\phi, \gamma'_\psi)'$, and $\mathbf{a} : (\mathbb{R}^q)^m \rightarrow \mathbb{R}^q$ is a permutation-invariant (*set*) function. Figure 1 shows a schematic of this representation. Each element of $\mathbf{a}(\cdot)$, $a_j(\cdot)$, $j = 1, \dots, q$, can be chosen to be a simple aggregation function, such as elementwise addition, average, or maximum, but may also be parameterised (Soelch et al., 2019); in this paper, we use the elementwise average. Beyond its attractive parsimonious form, (5) has appealing theoretical properties. For example, Han et al. (2022) show that functions of the form (5) can approximate any continuously differentiable permutation-invariant real-valued function for sufficiently large q , the dimension of $\mathbf{T}(\cdot)$; see also Wagstaff et al. (2019, 2022) for relevant discussion. The permutation-invariance property of Bayes estimators (Section S1) coupled with the representational capacity of (5) make DeepSets a principled and theoretically motivated framework for constructing neural Bayes estimators for replicated data.

Furthermore, the representation (5) is similar in form to many classical estimators when viewed as a nonlinear mapping $\phi(\cdot)$ of summary statistics $\mathbf{T}(\cdot)$. For example, best (i.e., minimum variance) unbiased estimators for exponential family models are of the form (5) where $\mathbf{T}(\cdot)$ is sufficient for θ (Casella and Berger, 2001, Ch. 7). This connection provides an additional rationale for adopting the structure in (5), and provides interpretability: the functions $\psi(\cdot)$ and $\mathbf{a}(\cdot)$ together extract summary statistics from the data, while $\phi(\cdot)$ maps these learned summary statistics to parameter estimates.

Some summary statistics are available in closed form, simple to compute, and highly informative (e.g., sample quantiles). Denote these statistics as $\mathbf{S}(\cdot)$. One may choose to explicitly incorporate these statistics by making $\phi(\cdot)$ in (5) a function of both $\mathbf{T}(\cdot)$ (learned) and $\mathbf{S}(\cdot)$ (user-defined). The estimator remains permutation invariant provided that $\mathbf{S}(\cdot)$ is permutation invariant. Since $\mathbf{T}(\cdot)$ can theoretically approximate well any summary statistic as a continuous function of the data, the choice to include $\mathbf{S}(\cdot)$ in (5) is mainly a practical one that could be useful in certain applications.

2.2.2 Variable sample size

Estimators of the form (5) can be applied to data sets of arbitrary size. However, the Bayes estimator for replicated data is a function of the number of replicates (see, e.g., (8) in the illustrative example of Section 2.2.3). Therefore, a neural Bayes estimator of the form (5) trained on data sets with m replicates will generally not be Bayes for data sets containing $\tilde{m} \neq m$ replicates.

There are at least two approaches that could be adopted if one wishes to use a neural Bayes estimator of the form (5) with data sets containing an arbitrary number of replicates, m . First, one could train l neural Bayes estimators for different sample sizes, or groups thereof (e.g., a small-sample estimator and a large-sample estimator). Specifically, for sample-size change-points $m_1 < m_2 < \dots < m_{l-1}$, one could construct a piecewise neural Bayes estimator,

$$\hat{\boldsymbol{\theta}}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}^*) = \begin{cases} \hat{\boldsymbol{\theta}}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}_{\tilde{m}_1}^*) & m \leq m_1, \\ \hat{\boldsymbol{\theta}}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}_{\tilde{m}_2}^*) & m_1 < m \leq m_2, \\ \vdots & \\ \hat{\boldsymbol{\theta}}(\mathbf{Z}^{(m)}; \boldsymbol{\gamma}_{\tilde{m}_l}^*) & m > m_{l-1}, \end{cases} \quad (6)$$

where, here, $\boldsymbol{\gamma}^* \equiv (\boldsymbol{\gamma}_{\tilde{m}_1}^*, \dots, \boldsymbol{\gamma}_{\tilde{m}_l}^*)$, and where $\boldsymbol{\gamma}_{\tilde{m}}^*$ are the neural-network parameters optimised for sample size \tilde{m} ; $\hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma}_{\tilde{m}}^*)$ is then assumed to be approximately Bayes over the range of sample sizes for which it is then applied in (6). Typically, $\tilde{m}_1 \leq m_1, m_1 < \tilde{m}_2 \leq m_2$, and so on. We find that this approach works well in practice, and that it is less computationally burdensome than it first appears when used in conjunction with pre-training (see Section 2.3). Note that the relative influence of the prior distribution diminishes as the sample size increases, so that only a single “large sample” estimator is needed. Alternatively, one could treat the sample size as a random variable, M , with support over a set of positive integers, \mathcal{M} , in which case (1) becomes

$$R(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\cdot; \boldsymbol{\gamma})) \equiv \sum_{m \in \mathcal{M}} \Pr(M = m) \left(\int_{S^m} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{z}^{(m)}; \boldsymbol{\gamma})) f(\mathbf{z}^{(m)} | \boldsymbol{\theta}) d\mathbf{z}^{(m)} \right). \quad (7)$$

This does not materially alter the workflow, except that one must sample the number of replicates before simulating data for (4). These two approaches can also be combined, so that each sub-estimator in (6) is trained using (7). We illustrate the importance in accounting for the dependence of the Bayes estimator on m in Section S2 of the Supplementary Material.

2.2.3 Illustrative example

We now present a relatively simple example that demonstrates that the DeepSets representation can approximate well Bayes estimators for replicated data. Consider the problem of estimating θ from data Z_1, \dots, Z_m that are independent and identically distributed (i.i.d.) according to a uniform distribution on $[0, \theta]$. Suppose that we choose a conjugate Pareto(α, β) prior for θ with shape $\alpha = 4$ and scale $\beta = 1$; that is, $\Pr(\theta \leq x) = 1 - (x/\beta)^{-\alpha}$, $x \geq \beta$. Under the absolute-error loss, the unique Bayes estimator is the posterior median which,

for this model, is available in closed form: for a Pareto(α, β) prior distribution, the Bayes estimator given m independent replicates $\mathbf{Z}^{(m)} \equiv (Z_1, \dots, Z_m)'$ is

$$\hat{\theta}_{\text{Bayes}}(\mathbf{Z}^{(m)}) = 2^{\frac{1}{\alpha+m}} \max(Z_1, \dots, Z_m, \beta), \quad (8)$$

which is clearly invariant to permutations of the observations, and can be expressed as a scalar rendition of (5), where $\psi(z) \equiv z$, $a(\{\cdot\}) \equiv \max(\{\cdot\})$, and $\phi(t) \equiv 2^{1/(\alpha+m)} \max(t, \beta)$. Note that when $a(\{\cdot\})$ is instead chosen to be the set average, alternative nonlinear functions $\psi(z)$ and $\phi(t)$ exist such that the estimator (5) approximates (8) arbitrarily well by approximating the maximum as a log-sum-exp.

For this model, the ‘one-at-a-time’ estimator, which applies the single-replicate Bayes estimator to each replicate independently and averages the resulting estimates, is given by

$$\hat{\theta}_0(\mathbf{Z}^{(m)}) = \frac{1}{m} \sum_{i=1}^m 2^{\frac{1}{\alpha+1}} \max(Z_i, \beta). \quad (9)$$

Although (9) is permutation invariant, it is clearly very different to (8). In Section S3 of the Supplementary Material we show that the one-at-a-time estimator for this model is not consistent, while the Bayes estimator is consistent. We compare our proposed estimator to the one-at-a-time estimator in this example and in the remainder of the paper, to stress the importance of properly accounting for replication when constructing neural Bayes estimators.

We now proceed with the design of the neural Bayes estimator as though we had no knowledge of its closed-form expression. We construct our neural Bayes estimator to have the form (5), where we use three-layer neural networks to model $\phi(\cdot; \gamma)$ and $\psi(\cdot; \gamma)$, and we let $J = 1$ and $K = 10^6$ in (4) when training the network. For simplicity, we consider a single value for m and fix $m = 10$. We compare the distributions of the estimators at $\theta = 4/3$ and $m = 10$ by simulating 30,000 data sets of size $m = 10$ from a uniform distribution on $[0, 4/3]$ and then applying the estimators to each of these simulated data sets. Figure 2 shows the kernel-smoothed distribution of the true Bayes estimator (8), the neural Bayes estimator, the one-at-a-time estimator (9), and the maximum likelihood estimator which, for this model, is given by $\max(Z_1, \dots, Z_m)$. The distribution of our neural Bayes estimator is clearly very similar to that of the true Bayes estimator, while the one-at-a-time estimator is clearly biased. The kernel-smoothed distribution of the maximum likelihood estimator emphasises that (neural) Bayes estimators are influenced by the prior distribution; in this case, the prior distribution leads to a Bayes estimator with a bimodal distribution due to a point mass in the lower endpoint of the support. This example also shows clearly that neural point estimators trained using (4) target the Bayes estimator, and not necessarily the maximum likelihood estimator; this notion is sometimes misconstrued in the literature.

2.3 Implementation

Neural Bayes estimators are conceptually simple and can be used in a wide range of problems where other approaches are computationally infeasible. They also have marked practical appeal, as the general workflow for their construction is only loosely connected to the model being considered. The workflow for constructing a neural Bayes estimator is as follows:




Figure 2: Kernel density approximations to the distribution of the Bayes estimator (green line), our neural Bayes estimator (red line), the maximum likelihood estimator (purple line), and the one-at-a-time estimator (orange line), for θ from $\text{Unif}(0, \theta)$ data, where $\theta = 4/3$ (grey dashed line) and where the sample size is $m = 10$.

- (a) Define the prior, $\Omega(\cdot)$.
- (b) Choose a loss function, $L(\cdot, \cdot)$, typically the absolute-error or squared-error loss.
- (c) Design a suitable neural-network architecture for the neural point estimator $\hat{\theta}(\cdot; \gamma)$.
- (d) Sample parameters from $\Omega(\cdot)$ to form training/validation/test parameter sets.
- (e) Given the above parameter sets, simulate data from the model, to form training/validation/test data sets.
- (f) Train the neural network (i.e., estimate γ) by minimising the loss function averaged over the training sets. That is, perform the optimisation task (3) where the Bayes risk is approximated using (4) with the training sets. During training, monitor performance and convergence using the validation sets.
- (g) Assess the fitted neural Bayes estimator, $\hat{\theta}(\cdot; \gamma^*)$, using the test set.

We elaborate on the steps of this workflow below. A crucial factor in the viability of any statistical method is the availability of software; to facilitate the adoption of neural Bayes estimators by statisticians, we accompany the paper with user-friendly open-source software in the Julia (Bezanson et al., 2017) and R (R Core Team, 2023) programming languages (available at <https://github.com/msainsburydale/NeuralEstimators>) that can be used in a wide range of settings.

Defining the prior. Prior distributions are typically determined by the applied problem being tackled. However, the choice of prior has practical implications on the training phase of the neural network. For example, an informative prior with compact and narrow support reduces the volume of the parameter space that must be sampled from when evaluating (4). In this case, a good approximation of the Bayes estimator can typically be obtained with smaller values of K and J in (4) than those required under a diffuse prior. This consideration

is particularly important when the number of parameters, p , is large, since the volume of the parameter space increases exponentially with p . On the other hand, if the neural Bayes estimator needs to be re-used for several applications, it might be preferable to employ prior distributions that are reasonably uninformative. [Lenzi et al. \(2023\)](#) suggest using likelihood-based estimates to elicit an informative prior from the data: however, this requires likelihood estimation to be feasible in the first place, which is often not the case in applications for which neural Bayes estimators are attractive.

Designing the neural-network architecture. The main consideration when designing the neural-network architecture is the structure of the data. For example, if the data are gridded, a convolutional neural network (CNN) may be used, while a dense neural network (DNN; also known as a multi-layer perceptron, MLP) is more appropriate for unstructured data (for further examples, see [Goodfellow et al., 2016](#)). When estimating parameters from replicated data using the representation (5), the architecture of $\psi(\cdot)$ is dictated by the structure of the data (since it acts on the data directly), while $\phi(\cdot)$ is typically a DNN (since $\mathbf{T}(\cdot)$ is a fixed-dimensional vector). Note that the training cost of the neural Bayes estimator is only amortised if the structure of new data sets conforms with the chosen architecture (otherwise the neural Bayes estimator will need to be re-trained with a new architecture). Once the general neural-network class is identified, one must specify architectural hyperparameters, such as the number of layers and number of neurons in each layer. We found that our results were not overly-sensitive to the specific choice of hyperparameters. In practice, the network must be sufficiently large so that the universal approximation theorem applies, but not so large as to make training prohibitively expensive. We give details on the architectures used in our experiments in Section 3.

Simulating parameters/data and training the network. In standard applications of neural networks, the amount of training data is fixed. One of the biggest risks one faces when fitting a large neural network is that the amount of training data is too “small”, which can result in the neural network overfitting the training set ([Goodfellow et al., 2016](#), Section 5.2). Such a neural network is said to have a high generalisation error. However, when constructing a neural Bayes estimator, we are able to simulate as much training data as needed. That is, we are able to set K and J in (4) as large as needed to avoid overfitting. The amount of training data needed would depend on the model, the number of parameters, the neural-network architecture, and the optimisation algorithm that is used. Providing general guidance is therefore difficult; however our experience is that K needs to be at least 10^4 – 10^5 for the neural network to have low generalisation error. Note that J can be kept small (on the order of 10^0 – 10^1) since data are simulated for every sampled parameter vector. One also has the option to simulate training data “on-the-fly”, in the sense that new training data are simulated continuously during training ([Chan et al., 2018](#)). This approach facilitates the use of large networks with a high representational capacity, since then the data used in stochastic-gradient-descent updates are always different from those in previous updates (see Section S4 of the Supplementary Material for more details and for an illustration). On-the-fly simulation also allows the data to be simulated “just-in-time”, in the sense that the data can be simulated from a small batch of parameters, used to train the estimator, and then removed

from memory; this can reduce memory requirements when a large amount of training data are required to avoid overfitting. Chan et al. (2018) also continuously refresh the parameter vectors in the training set, which has similar benefits. Keeping these parameters fixed, however, allows computationally expensive terms, such as Cholesky factors, to be reused throughout training, which can substantially reduce the training time with some models (Gerber and Nychka, 2021).

The parameters of a neural network trained for one task can be used as initial values for the parameters of another neural network intended for a slightly different task. This is known as pre-training (Goodfellow et al., 2016, Ch. 8). Pre-training is ideal when developing a piecewise neural Bayes estimator (6), whereby one may randomly initialise and train $\hat{\theta}(\cdot; \gamma_{\tilde{m}_1})$, use the optimised parameters $\gamma_{\tilde{m}_1}^*$ as initial values when training $\hat{\theta}(\cdot; \gamma_{\tilde{m}_2})$, where $\tilde{m}_2 > \tilde{m}_1$, and so on. Since each estimator need only account for a larger sample size than that used by its predecessor, scant computational resources are needed to train subsequent estimators. This approach can be useful even when only a single large-sample estimator is needed: doing most of the learning with small, computationally cheap sample sizes can substantially reduce training time when compared to training a single large-sample estimator from scratch. We demonstrate the benefits of this strategy through an illustration in Section S5 of the Supplementary Material.

Assessing the estimator. Once a neural Bayes estimator is trained, its performance needs to be assessed on unseen test data. Simulation-based (empirical) methods are ideal, since simulation is already required for constructing the estimator. Neural Bayes estimators are very fast to evaluate, and can therefore be applied to thousands of simulated data sets at almost no computational cost. One can therefore quickly and accurately assess the estimator with respect to any property of its sampling distribution (e.g., bias, variance, etc.).

3 Simulation studies

We now conduct several simulation studies that clearly demonstrate the utility of neural Bayes estimators in increasingly complex settings. In Section 3.1, we outline the general setting. In Section 3.2 we estimate the parameters of a Gaussian process model with three unknown parameters. Since the likelihood function is available for this model, here we compare the efficiency of our neural Bayes estimator to that of a gold-standard likelihood-based estimator. In Section 3.3 we consider a spatial extremes setting and estimate the two parameters of Schlather’s max-stable model (Schlather, 2002); the likelihood function is computationally intractable for this model, and we observe substantial improvements over the classical composite-likelihood approach. In Section 3.4 we consider the highly-parameterised conditional extremes model (Wadsworth and Tawn, 2022). We provide simulation details and density functions for each model in Section S6 of the Supplementary Material. For ease of notation, we omit dependence on the neural-network parameters γ in future references to the neural Bayes estimators, so that $\hat{\theta}(\cdot; \gamma)$ is written simply as $\hat{\theta}(\cdot)$.

3.1 General setting

Across the simulation studies we assume, for ease of exposition, that our processes are spatial. Our spatial domain of interest, \mathcal{D} , is $[0, 16] \times [0, 16]$, and we simulate data on a regular grid with unit-square cells, yielding $16^2 = 256$ observations per spatial field. CNNs are a natural choice for regularly-spaced gridded data, and we therefore use a CNN architecture, summarised in Table S3 of the Supplementary Material. To implement the neural point estimators, we use the accompanying package **NeuralEstimators** that is written in **Julia** and which leverages the package **Flux** (Innes, 2018). We conduct our studies using a workstation with an AMD EPYC 7402 3.00GHz CPU with 52 cores and 128 GB of CPU RAM, and an NVIDIA Quadro RTX 6000 GPU with 24 GB of GPU RAM. All results presented in the remainder of this paper can be generated using reproducible code at <https://github.com/msainsburydale/NeuralBayesEstimators>.

We consider two neural point estimators which are both based on the architecture given in Table S3 of the Supplementary Material. The first estimator, $\hat{\theta}_0(\cdot)$, is the permutation-invariant one-at-a-time neural point estimator considered by Gerber and Nychka (2021). The second estimator, $\hat{\theta}_{\text{DS}}(\cdot)$, is the piecewise neural point estimator (6), where each sub-estimator employs the DeepSets representation (5) with $\psi(\cdot)$ and $\phi(\cdot)$ constructed using the first four and last two rows of Table S3, respectively. Five sub-estimators are used, with training sample sizes $\tilde{m}_1 = 1$, $\tilde{m}_2 = 10$, $\tilde{m}_3 = 35$, $\tilde{m}_4 = 75$, and $\tilde{m}_5 = 150$, and with sample-size changepoints $m_1 = 1$, $m_2 = 20$, $m_3 = 50$, and $m_4 = 100$.

We assume that the parameters are *a priori* independent and uniformly distributed on an interval that is parameter dependent. We train the neural point estimators under the absolute-error loss. We set K in (4) to 10,000 and 2,000 for the training and validation *parameter sets*, respectively, and we keep these sets fixed during training. We construct the training and validation *data sets* by simulating $J = 10$ sets of m model realisations for each parameter vector in the training and validation parameter sets. During training, we fix the validation data, but simulate the training data on-the-fly; hence, in this paper, we define an epoch as a pass through the training sets when doing stochastic gradient descent, after which the training data (i.e., the $J = 10$ data sets at each of the 10,000 parameter samples) are refreshed. We cease training when the risk evaluated using the validation set has not decreased in 5 consecutive epochs.

We compare the neural Bayes estimators to the maximum *a posteriori* (MAP) estimator,

$$\hat{\theta}_{\text{MAP}}(\mathbf{Z}^{(m)}) \equiv \arg \max_{\theta \in \Theta} \sum_{i=1}^m \ell(\theta; \mathbf{Z}_i) + \log p(\theta), \quad (10)$$

where $\ell(\theta; \cdot)$ is the log-likelihood function for a single replicate and $p(\theta)$ is the prior density. We solve (10) using the true parameters as initial values for the Nelder–Mead algorithm. We advantage the competitor MAP estimator, which minimises (2) under the 0–1 loss, by assessing all estimators with respect to the 0–1 loss (we assign zero loss if an estimate is within 10% of the true value), with $K = 500$ test parameter vectors.

3.2 Gaussian process model

Spatial statistics is concerned with modelling data that are collected across space; reference texts include [Cressie \(1993\)](#), [Banerjee et al. \(2004\)](#), and [Diggle and Ribeiro \(2007\)](#). Here, we consider a classical spatial model, the linear Gaussian-Gaussian model,

$$Z_{ij} = Y_i(\mathbf{s}_j) + \epsilon_{ij}, \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad (11)$$

where $\mathbf{Z}_i \equiv (Z_{i1}, \dots, Z_{in})'$ are data observed at locations $\{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ on a spatial domain \mathcal{D} , $\{Y_i(\cdot)\}$ are i.i.d. spatially-correlated mean-zero Gaussian processes, and $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$ is Gaussian white noise. The covariance function, $C(\mathbf{s}, \mathbf{u}) \equiv \text{cov}(Y_i(\mathbf{s}), Y_i(\mathbf{u}))$, for $\mathbf{s}, \mathbf{u} \in \mathcal{D}$ and $i = 1, \dots, m$, is the primary mechanism for capturing spatial dependence. Note that, since $\{Y_i(\cdot)\}$ are i.i.d., we have that $\text{cov}(Y_i(\mathbf{s}), Y_{i'}(\mathbf{u})) = 0$ for all $i \neq i'$ and $\mathbf{s}, \mathbf{u} \in \mathcal{D}$. Here, we use the popular isotropic Matérn covariance function,

$$C(\mathbf{s}, \mathbf{u}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\|\mathbf{s} - \mathbf{u}\|}{\rho} \right)^\nu K_\nu \left(\frac{\|\mathbf{s} - \mathbf{u}\|}{\rho} \right), \quad (12)$$

where σ^2 is the marginal variance, $\Gamma(\cdot)$ is the gamma function, $K_\nu(\cdot)$ is the Bessel function of the second kind of order ν , and $\rho > 0$ and $\nu > 0$ are the range and smoothness parameters, respectively. We follow [Gerber and Nychka \(2021\)](#) and fix $\sigma^2 = 1$. This leaves three unknown parameters that need to be estimated: $\boldsymbol{\theta} \equiv (\sigma_\epsilon, \rho, \nu)'$.

Fixing ν simplifies the estimation task, since the remaining parameters are well-identified from a single replicate; we consider this model, which was also considered by [Gerber and Nychka \(2021\)](#), in Section S7 of the Supplementary Material. Here, we consider the case where all three parameters are unknown. Estimation in this case is more challenging since ρ and ν are only weakly identifiable from a single replicate ([Zhang, 2004](#)), but inference on ν is important since it controls the differentiability of the process ([Stein, 1999](#), Ch. 2).

We use the priors $\sigma_\epsilon \sim \text{Unif}(0.1, 1)$, $\rho \sim \text{Unif}(2, 10)$, and $\nu \sim \text{Unif}(0.5, 3)$. The total training time for $\hat{\boldsymbol{\theta}}_0(\cdot)$ is 15 minutes. The training time for $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$, despite it consisting of several sub-estimators and being trained with large sample sizes, is only slightly more, at 33 minutes, due to the pre-training strategy discussed in Section 2.3. This modest increase in training time is a small cost when accounting for the improvement in statistical efficiency, as shown in Figure 3. Clearly, $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ substantially improves over $\hat{\boldsymbol{\theta}}_0(\cdot)$, and performs similarly well to the MAP estimator. The neural point estimators only take 0.1 and 2 seconds to yield estimates from all the test data when $m = 1$ and $m = 150$, respectively, while the MAP estimator takes between 600 and 800 seconds.

Next, we compare the empirical joint distributions of the estimators for a single parameter vector. Figure 4 shows the true parameters (red cross) and corresponding estimates obtained by applying each estimator to 100 sets of $m = 150$ independent model realisations. The MAP estimator and $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ are both approximately unbiased, and both are able to capture the correlations expected when using the parameterisation (12); both are clearly much better than $\hat{\boldsymbol{\theta}}_0(\cdot)$. Empirical joint distributions for additional parameter vectors are shown in Figure S11 of the Supplementary Material, and lead to similar conclusions to those drawn here. Overall, these results show that $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ is a substantial improvement over the prior art, $\hat{\boldsymbol{\theta}}_0(\cdot)$, and that $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ is competitive with the likelihood-based estimator for this model.




Figure 3: The risk with respect to the 0–1 loss against the number of replicates, m , evaluated using the parameter vectors in the test parameter set, for the estimators considered in Section 3.2. The estimators $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{\text{DS}}(\cdot)$, and the MAP estimator are shown in green, orange, and purple, respectively.




Figure 4: The empirical joint distribution of the estimators considered in Section 3.2 for a single parameter vector. The true parameters are shown in red, while estimates from $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{\text{DS}}(\cdot)$, and the MAP estimator are shown in green, orange, and purple, respectively. Each estimate was obtained from a simulated data set of size $m = 150$.

3.3 Schlather’s max-stable model

We now consider models used for spatial extremes, useful reviews for which are given by Davison and Huser (2015), Davison et al. (2019) and Huser and Wadsworth (2022). Max-stable processes are the cornerstone of spatial extreme-value analysis, being the only possible non-degenerate limits of properly renormalised pointwise maxima of i.i.d. random fields. However, their practical use has been severely hampered due to the computational bottleneck in evaluating their likelihood function. They are thus natural models to consider in our experiments, and we here consider a fairly simple max-stable model with only two parameters. We consider Schlather’s max-stable model (Schlather, 2002), given by

$$Z_{ij} = \bigvee_{k \in \mathbb{N}} \zeta_{ik} \max\{0, Y_{ik}(\mathbf{s}_j)\}, \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad (13)$$

where \bigvee denotes the maximum over the indexed terms, $\mathbf{Z}_i \equiv (Z_{i1}, \dots, Z_{in})'$ are observed at locations $\{\mathbf{s}_1, \dots, \mathbf{s}_n\} \subset \mathcal{D}$, $\{\zeta_{ik} : k \in \mathbb{N}\}$ for $i = 1, \dots, m$ are i.i.d. Poisson point processes on $(0, \infty)$ with intensity measure $d\Lambda(\zeta) = \zeta^{-2}d\zeta$, and $\{Y_{ik}(\cdot) : i = 1, \dots, m, k \in \mathbb{N}\}$ are i.i.d. mean-zero Gaussian processes scaled so that $\mathbb{E}[\max\{0, Y_{ik}(\cdot)\}] = 1$. Here, we model each $Y_{ik}(\cdot)$ using the Matérn covariance function (12), with $\sigma^2 = 1$. Hence, $\boldsymbol{\theta} \equiv (\rho, \nu)'$.

We use the same uniform priors for ρ and ν as in Section 3.2. Realisations from the present model, here expressed on unit Fréchet margins, tend to have highly varying magnitudes, and we reduce this variability by log-transforming our data to the unit Gumbel scale. The total training time for $\hat{\boldsymbol{\theta}}_0(\cdot)$ and $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ is 14 and 66 minutes, respectively.

As in Section 3.2, we assess the neural point estimators by comparing them to a likelihood-based estimator. For Schlather’s model (and other max-stable models in general), the full likelihood function is computationally intractable, since it involves a summation over the set of all possible partitions of the spatial locations (see, e.g., Padoan et al., 2010; Huser et al., 2019, and the references therein). A popular substitute is the pairwise likelihood (PL) function, a composite likelihood formed by considering only pairs of observations; specifically, the pairwise log-likelihood function for the i th replicate is

$$\ell_{\text{PL}}(\boldsymbol{\theta}; \mathbf{z}_i) \equiv \sum_{j=1}^{n-1} \sum_{j'=j+1}^n \log f(z_{ij}, z_{ij'} \mid \boldsymbol{\theta}), \quad (14)$$

where $f(\cdot, \cdot \mid \boldsymbol{\theta})$ denotes the bivariate probability density function for pairs in \mathbf{z}_i . Hence, in this subsection, we compare the neural Bayes estimators to the pairwise MAP (PMAP) estimator, that is, (10) with the full log-likelihood function $\ell(\boldsymbol{\theta}; \cdot)$ replaced by $\ell_{\text{PL}}(\boldsymbol{\theta}; \cdot)$. Often, both computational and statistical efficiency can be drastically improved by using only a subset of pairs that are within a fixed cut-off distance, d (see, e.g., Bevilacqua et al., 2012; Sang and Genton, 2014). A line-search for d (see Figure S12 of the Supplementary Material for details) shows that, here, $d = 3$ units (used hereafter) provides good results.

The left and centre panels of Figure 5 show the estimators’ risk against the number of independent replicates. For small samples, both neural point estimators improve over the PMAP estimator. For moderate-to-large samples, $\hat{\boldsymbol{\theta}}_0(\cdot)$ hits a performance plateau, while $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ continues to substantially outperform the PMAP estimator. The run time for the neural point estimators to estimate all test parameters scales linearly between 0.1 and 1.5 seconds for $m = 1$ and $m = 150$, respectively, while the PMAP estimator takes between 750 and 1900 seconds. The right panel of Figure 5 shows the empirical joint distribution of the estimators for a single parameter vector, where each estimate was obtained from $m = 150$ replicates. Again, $\hat{\boldsymbol{\theta}}_0(\cdot)$ is strongly biased, while the PMAP estimator is unbiased but is less efficient than $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$. Empirical joint distributions for additional parameter vectors are shown in Figure S13 of the Supplementary Material, and lead to similar conclusions to those drawn here. Overall, the proposed estimator, $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$, is statistically and computationally superior to the likelihood-based technique for Schlather’s max-stable model.

3.4 Spatial conditional extremes model

While max-stable processes are asymptotically justified for modelling spatial extremes defined as block maxima, they have strong limitations in practice (Huser and Wadsworth, 2022).




Figure 5: Diagnostic plots for the simulation study of Section 3.3. (Left and centre) The risk with respect to the 0–1 loss against the number of replicates, m , evaluated using the parameter vectors in the test parameter set. (Right) True parameters (red) and corresponding estimates, each of which was obtained using a simulated data set of size $m = 150$. In all panels, $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{\text{DS}}(\cdot)$, and the pairwise MAP estimator are shown in green, orange, and purple, respectively.

Beyond the intractability of their likelihood function in high dimensions, max-stable models have an overly rigid dependence structure, and in particular cannot capture a property known as asymptotic tail independence. Therefore, different model constructions, justified by alternative asymptotic conditions, have been recently proposed to circumvent the restrictions imposed by max-stability. In particular, the spatial conditional extremes model, first introduced by [Wadsworth and Tawn \(2022\)](#) as a spatial extension of the multivariate [Heffernan and Tawn \(2004\)](#) model, and subsequently studied and used in applications by, for example, [Richards et al. \(2022\)](#) and [Simpson et al. \(2023\)](#), is especially appealing. This model has a flexible dependence structure capturing both asymptotic tail dependence and independence ([Wadsworth and Tawn, 2022](#)) and leads to likelihood-based inference amenable to higher dimensions. Nevertheless, parameter estimation remains challenging because this model is complex and is typically highly parameterised. Here, we consider a version of the spatial conditional extremes model that involves eight dependence parameters in total.

Our formulation of the model is similar to that originally proposed by [Wadsworth and Tawn \(2022\)](#). Specifically, we model the process, expressed on Laplace margins, conditional on it exceeding a threshold, u , at a conditioning site, $\mathbf{s}_0 \in \mathcal{D}$, as

$$Z_{ij} \mid Z_{i0} > u \stackrel{\text{d}}{=} a(\mathbf{h}_j, Z_{i0}) + b(\mathbf{h}_j, Z_{i0})Y_i(\mathbf{s}_j), \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad (15)$$

where ‘ $\stackrel{\text{d}}{=}$ ’ denotes equality in distribution, Z_{i0} is the datum at \mathbf{s}_0 , $\mathbf{h}_j \equiv \mathbf{s}_j - \mathbf{s}_0$, and $Z_{i0} - u \mid Z_{i0} > u$ is a unit exponential random variable that is independent of the residual process, $Y_i(\cdot)$, which we describe below. We model $a(\cdot, \cdot)$ and $b(\cdot, \cdot)$ using parametric forms proposed by [Wadsworth and Tawn \(2022\)](#), namely

$$\begin{aligned} a(\mathbf{h}, z) &= z \exp\{-\|\mathbf{h}\|/\lambda\}^\kappa, & \lambda > 0, \quad \kappa > 0, \\ b(\mathbf{h}, z) &= 1 + a(\mathbf{h}, z)^\beta, & \beta > 0, \end{aligned}$$

where $\mathbf{h} \equiv \mathbf{s} - \mathbf{s}_0$ for $\mathbf{s} \in \mathcal{D}$ and $z \in \mathbb{R}$. We construct the residual process, $Y_i(\cdot)$, by first defining $\tilde{Y}_i^{(0)}(\cdot) \equiv \tilde{Y}_i(\cdot) - \tilde{Y}_i(\mathbf{s}_0)$, with $\tilde{Y}_i(\cdot)$ a mean-zero Gaussian process with Matérn

covariance function and unit marginal variance, and then marginally transforming it to the scale of a delta-Laplace (generalised Gaussian) distribution (Subbotin, 1923) which, for $y \in \mathbb{R}$ and parameters $\mu \in \mathbb{R}, \tau > 0, \delta > 0$, has density function

$$f_S(y \mid \mu, \tau, \delta) = \frac{\delta}{2\tau\Gamma(1/\delta)} \exp\left(-\left|\frac{y - \mu}{\tau}\right|^\delta\right).$$

We model δ as decaying from 2 to 1 as the distance to \mathbf{s}_0 increases; specifically,

$$\delta(\mathbf{h}) = 1 + \exp\left\{-\left(\|\mathbf{h}\|/\delta_1\right)^2\right\}, \quad \delta_1 > 0. \quad (16)$$

We defer to Wadsworth and Tawn (2022) for model justification and interpretation of model parameters. The threshold u is a modelling decision made to ensure that the data are sufficiently extreme and that, in a real-data setting, we have sufficiently many fields available for estimation: here, we set u to the 0.975 quantile of the unit-Laplace distribution. For simplicity, we consider \mathbf{s}_0 fixed and in the centre of \mathcal{D} .

We again use uniform priors, with $\kappa \sim \text{Unif}(1, 2)$, $\lambda \sim \text{Unif}(2, 5)$, $\beta \sim \text{Unif}(0.05, 1)$, $\mu \sim \text{Unif}(-0.5, 0.5)$, $\tau \sim \text{Unif}(0.3, 0.9)$, $\delta_1 \sim \text{Unif}(1.3, 3)$, and with the same priors for ρ and ν as used in the preceding sections. We use the cube-root function as a variance-stabilising transformation. The training time for $\hat{\boldsymbol{\theta}}_0(\cdot)$ and $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ is 22 and 43 minutes, respectively.

Figure 6 shows the risk against the sample size. Figure 7 shows the empirical joint distribution of the estimators for a single parameter vector (panels in the lower triangle) and simulations from the corresponding model (panels in the upper triangle). The estimator $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ is approximately unbiased for all parameters, and captures the expected negative correlation between ρ and ν . Overall, the estimator $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ is clearly appropriate for this highly parameterised model (unlike $\hat{\boldsymbol{\theta}}_0(\cdot)$), which is an important result for this framework.

4 Application to Red Sea surface temperature

We now apply our methodology to the analysis of sea-surface temperature data in the Red Sea, which have also been analysed by Hazra and Huser (2021), Simpson and Wadsworth (2021), and Simpson et al. (2023), among others, and have been the subject of a competition in the prediction of extreme events (Huser, 2021). The data set we analyse comprises daily observations from the years 1985 to 2015, for 16703 regularly-spaced locations across the Red Sea; see Donlon et al. (2012) for further details. Following Simpson et al. (2023), we focus on a southern portion of the Red Sea, consider only the summer months to approximately eliminate the effects of seasonality, and retain only every third longitude and latitude value; this yields a data set with 678 unique spatial locations that are regularly-spaced but contained within an irregularly-shaped spatial domain, \mathcal{D} . To account for \mathcal{D} lying away from the equator, we follow Simpson et al. (2023) in scaling the longitude and latitude so that each unit distance in the spatial domain corresponds to approximately 100 km.

To model these data, we use the spatial conditional extremes model described in Section 3.4. We transform our data to the Laplace scale and set the threshold, u in (15), to the 95th percentile of the transformed data. This yields 141 spatial fields for which the transformed datum at the conditioning site, \mathbf{s}_0 , here chosen to lie in the centre of \mathcal{D} , is




Figure 6: The risk with respect to the 0–1 loss against the number of replicates, m , evaluated using the parameter vectors in the test parameter set, for the estimators considered in Section 3.4. The estimators $\hat{\theta}_0(\cdot)$ and $\hat{\theta}_{DS}(\cdot)$ are shown in green and orange, respectively.

greater than u ; that is, we estimate parameters based on $m = 141$ replicates of the spatial process. A randomly-selected sample of these extreme fields are shown in the upper panels of Figure S14 of the Supplementary Material.

The irregular shape of \mathcal{D} means that CNNs are not directly applicable. We use the standard technique of padding empty regions with zeros, so that each field is a 29×37 rectangular array consisting of a data region and a padded region, as shown in Figure S15 of the Supplementary Material. We use the same prior distributions as in Section 3.4, but with those associated with range parameters scaled appropriately. Our architecture is the same as that given in Table S3, but with an additional convolutional layer that transforms the input array to dimension 16×16 . We validate our neural Bayes estimator using the approach taken in Section 3 (figures omitted for brevity).

Once training is complete, we can compute parameter estimates for the observed data. Table 1 gives estimates and 95% confidence intervals for the estimates. These confidence intervals are obtained using the non-parametric bootstrap procedure described in Section S8 of the Supplementary Material, which accounts for temporal dependence between the spatial fields. Estimation from a single set of 141 fields takes only 0.008 seconds, meaning that bootstrap confidence intervals can be obtained very quickly. This is clearly an advantage of using neural point estimators, as opposed to likelihood-based techniques for which uncertainty assessment in complex models is usually a computational burden.

Next, following Simpson et al. (2023), we separate \mathcal{D} into 17 non-overlapping regions, which are shown in the left panel of Figure 8. Given that the transformed datum at \mathbf{s}_0 exceeds u , we estimate (and quantify the uncertainty in) the proportion of locations in each region that also exceed u , using both model-based and empirical methods. These are

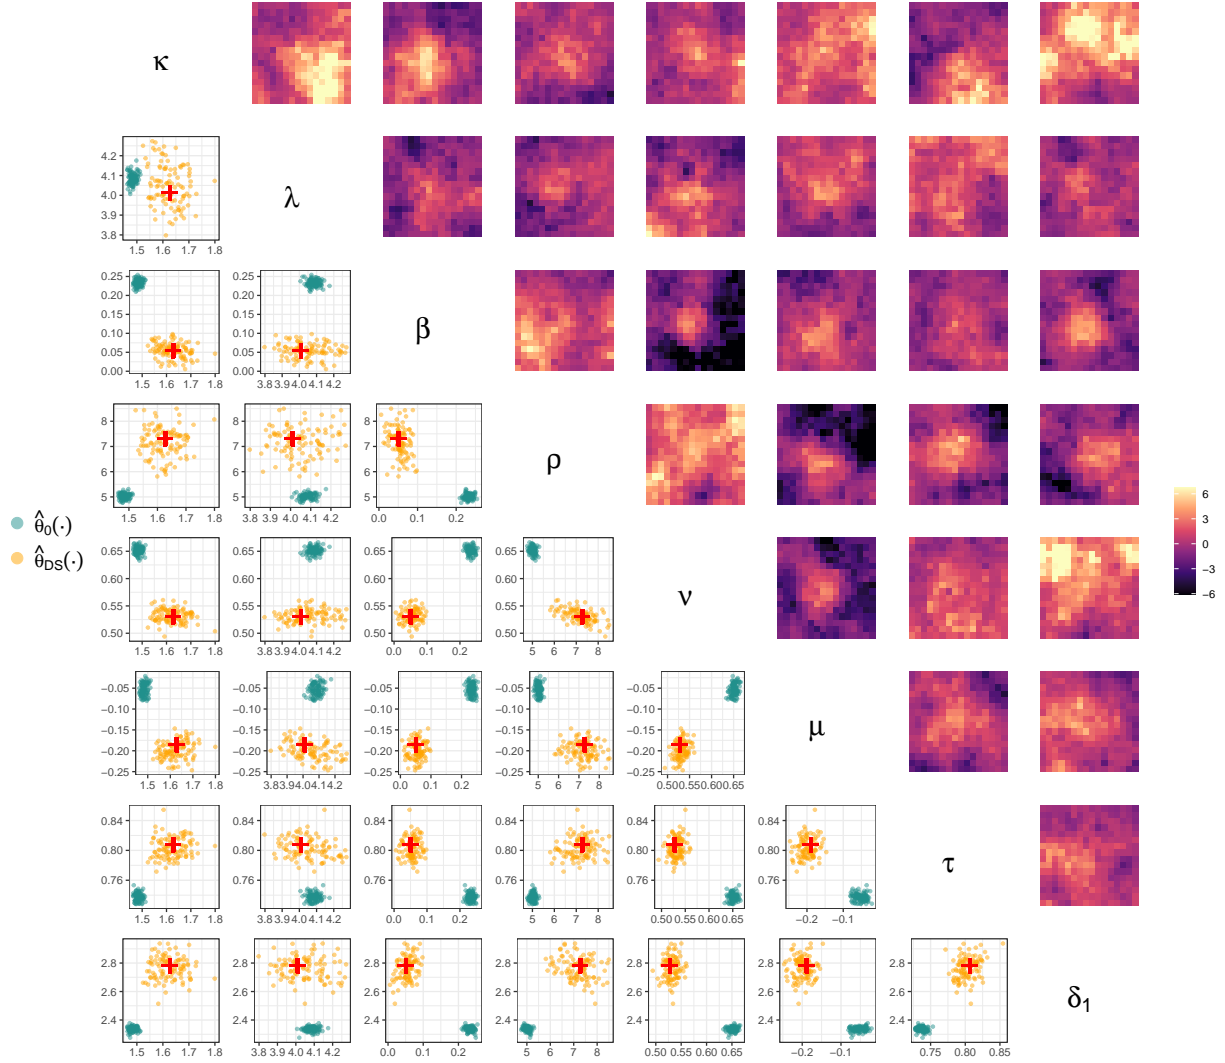


Figure 7: (Lower triangle) The empirical joint distribution of the estimators considered in Section 3.4 for a single parameter vector. The true parameters are shown in red, while estimates from $\hat{\theta}_0(\cdot)$ and $\hat{\theta}_{DS}(\cdot)$ are shown in green and orange, respectively. Each estimate was obtained from a simulated data set of size $m = 150$. (Upper triangle) Simulations from the model.

Table 1: Parameter estimates and 95% bootstrap confidence intervals (provided via the 2.5 and 97.5 percentiles of the bootstrap distribution) for the Red Sea data set of Section 4.

	κ	λ	β	ρ	ν	μ	τ	δ_1
Estimate	1.00	2.74	0.24	1.34	0.78	0.09	0.56	1.22
2.5%	0.88	2.13	0.14	0.95	0.70	0.04	0.49	0.91
97.5%	1.19	3.79	0.40	1.85	0.87	0.14	0.65	1.63

shown in the right panel of Figure 8, and indicate overall agreement between the model and the observed data, with some minor lack of fit at very short distances. The empirical estimates are not monotonically decreasing as a function of distance; this could be due to complex spatial dynamics or non-stationarity in the data, or it could simply be an artefact




Figure 8: (Left) The spatial domain of interest for the Red Sea study of Section 4, with the conditioning site, \mathbf{s}_0 , shown in red, and the remaining locations separated into 17 regions; the region labels begin at 1 in the centre of the domain and increase with distance from the centre. (Right) The estimated proportion of locations for which the process exceeds u given that it is exceeded at \mathbf{s}_0 (points) and corresponding 95% confidence intervals (vertical segments) using model-simulated data sets (blue) and bootstrap samples of the observed data set (red).

of sampling variability. In either case, the fit is reasonable, which suggests that our neural Bayes estimator has provided reasonable parameter estimates for this data set.

5 Conclusion

Neural Bayes estimators are a class of likelihood-free estimators that approximate Bayes estimators using neural networks. Their connection to classical estimation theory is often under-appreciated, and this article serves to increase the awareness and adoption of this powerful estimation tool by statisticians. This paper also proposes a principled way to construct neural Bayes estimators for replicated data via the DeepSets architecture, which has the same structure as that of well-known conventional estimators, such as best unbiased estimators with exponential family models. Using these estimators that are able to automatically learn suitable summary statistics from the data, we jointly estimate the range and smoothness parameters in a Gaussian process model and in Schlather’s max-stable model, and estimate parameters in the highly-parameterised spatial conditional extremes model. These estimators are implemented with just a few lines of code, thanks to the package **NeuralEstimators**, which is released with this paper and is available at <https://github.com/msainsburydale/NeuralEstimators>.

As with all estimation methods, neural Bayes estimators come with advantages and disadvantages that will either make them highly applicable, or impractical, depending on the application. At the time of writing we see five main drawbacks of neural Bayes estimators. First, it is unclear how one could verify that a trained neural Bayes estimator is ‘close’ to the true Bayes estimator, and more theory is needed to provide guarantees and guidelines on how to choose the number of training samples in practice. This limitation can be somewhat mitigated by running empirical checks (as we do in our examples); these checks are not com-

putationally costly given the amortised nature of the neural Bayes estimator. Second, the training cost is amortised only if the estimator is used repeatedly for the same estimation problem; in applications where estimation needs to be done just once, and the likelihood function is available and computationally tractable, classical likelihood-based methods supported by theoretical guarantees and extensive software availability might be more attractive. Third, one needs to be able to simulate relatively easily from the data generating process; although this sampling phase is parallelisable, this step could make the use of neural Bayes estimators with some models impractical. Fourth, the fast construction of neural Bayes estimators requires graphics processing units (GPUs); however, GPUs suitable for deep learning are now commonplace in high-end workstations, and this hardware is only needed for training, and not for evaluating the estimator once trained. Fifth, the volume of the parameter space increases exponentially fast with the number of parameters; we therefore expect training the neural Bayes estimator to be more difficult with highly-parameterised models, especially those with non-orthogonal parametrisations. On the other hand, neural point estimators have substantial advantages that will, in our opinion, make them the preferred and the *de facto* option in several applications in the near future. First, they are particularly useful when the likelihood function is unavailable or computationally intractable. Second, due to their amortised nature, they are ideal for settings in which the same statistical model must be fit repeatedly (e.g., when solving inverse problems with remote sensing data; see [Cressie, 2018](#)), or in scenarios where accurate bootstrap-based uncertainty estimates are needed.

There are many avenues for future research. In this work, we have illustrated neural Bayes estimation with spatial data observed over a regular grid. Parameter estimation from irregular spatial data is an important problem, and one that we consider briefly with DNNs in Section S9 of the Supplementary Material. An attractive way forward in this regard is the use of graph neural networks (GNNs; [Wu et al., 2021](#)), which generalise the convolution operation to irregular data; their use in the context of parameter estimation is the subject of ongoing work. Ways to incorporate covariate information also need to be explored. A possible criticism of neural Bayes estimators is that they are not robust to model misspecification since neural networks are, generally, poor extrapolators. However, we did not find this to be the case in our work. In Section S10 of the Supplementary Material, we provide some empirical evidence that neural Bayes estimators can be used on data that are very different to those used during training, but further research on this topic is needed. Neural Bayes estimators of the form (5) are ideally placed for online learning problems; investigating their potential in this context is the subject of future work. Finally, while this paper focuses on neural point estimation, there is a growing literature on neural approaches that approximate the full posterior distribution ([Radev et al., 2022](#); [Pacchiardi and Dutta, 2022](#)): this is also a promising avenue for future research.

Acknowledgements

Matthew Sainsbury-Dale’s and Andrew Zammit-Mangion’s research was supported by an 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. Raphaël Huser was partially supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No. OSR-CRG2020-4394. The authors would like to acknowledge Johan Barthelemy, NVIDIA, SMART Infrastructure Facility of the University of Wollongong, as well as a 2021 Major Equipment Grant from the University of Wollongong’s University Research Committee, that together have provided access to extensive GPU computing resources. The authors would like to thank Yi Cao for technical support. Thanks also go to Emma Simpson, Jennifer Wadsworth, and Thomas Opitz for providing access to the preprocessed Red Sea data set. The authors are also grateful to Noel Cressie and Jonathan Rougier for providing comments on the manuscript. We are also grateful to the reviewers and the editors for their helpful comments and suggestions that improved the quality of the manuscript.

References

- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC Press, Boca Raton, FL.
- Banesh, D., Panda, N., Biswas, A., Roekel, L. V., Oyen, D., Urban, N., Grosskopf, M., Wolfe, J., and Lawrence, E. (2021). Fast Gaussian Process estimation for large-scale in situ inference using convolutional neural networks. In *IEEE International Conference on Big Data*, pages 3731–3739.
- Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics*, 162:2025–2035.
- Bevilacqua, M., Gaetan, C., Mateu, J., and Porcu, E. (2012). Estimating space and space-time covariance functions for large data sets: A weighted composite likelihood approach. *Journal of the American Statistical Association*, 107:268–280.
- Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59:65–98.
- Casella, G. and Berger, R. (2001). *Statistical Inference*. Duxbury, Belmont, CA, second edition.
- Castruccio, S., Huser, R., and Genton, M. G. (2016). High-order composite likelihood inference for max-stable distributions and processes. *Journal of Computational and Graphical Statistics*, 25:1212–1229.
- Chan, J., Perrone, V., Spence, J., Jenkins, P., Mathieson, S., and Song, Y. (2018). A likelihood-free inference framework for population genetic data using exchangeable neural networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Chon, K. and Cohen, R. (1997). Linear and nonlinear ARMA model parameter estimation using an artificial neural network. *IEEE Transactions on Biomedical Engineering*, 44:168–174.
- Cox, D. and Reid, N. (2004). A note on pseudolikelihood constructed from marginal densities. *Biometrika*, 3:729–737.
- Creel, M. (2017). Neural nets for indirect inference. *Econometrics and Statistics*, 2:36–49.
- Cressie, N. (1993). *Statistics for Spatial Data*. Wiley, Hoboken, NJ, revised edition.
- Cressie, N. (2018). Mission control: A statistical scientist’s role in remote sensing of atmospheric carbon dioxide. *Journal of the American Statistical Association*, 113(521):152–168.
- Cressie, N. (2023). Decisions, decisions, decisions in an uncertain environment. *Environmetrics*, 34:e2767.
- Davison, A. C. and Huser, R. (2015). Statistics of extremes. *Annual Review of Statistics and its Application*, 2:203–235.

- Davison, A. C., Huser, R., and Thibaud, E. (2019). Spatial extremes. In Gelfand, A. E., Fuentes, M., Hoeting, J. A., and Smith, R. L., editors, *Handbook of Environmental and Ecological Statistics*, pages 711–744. Chapman & Hall/CRC Press, Boca Raton, FL.
- Diggle, P. J. and Gratton, R. J. (1984). Monte Carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society B*, 46:193–227.
- Diggle, P. J. and Ribeiro, P. J. (2007). *Model-Based Geostatistics*. Springer, New York, NY.
- Donlon, C. J., Martin, M., Stark, J., Roberts-Jones, J., Fiedler, E., and Wimmer, W. (2012). The operational sea surface temperature and sea ice analysis (OSTIA) system. *Remote Sensing of Environment*, 116:140–158.
- Flagel, L., Brandvain, Y., and Schrider, D. R. (2018). The unreasonable effectiveness of convolutional neural networks in population genetic inference. *Molecular Biology and Evolution*, 36:220–238.
- Gaskin, T., Pavliotis, G. A., and Girolami, M. (2023). Neural parameter calibration for large-scale multiagent models. *Proceedings of the National Academy of Sciences*, 120(7):e2216415120.
- Gerber, F. and Nychka, D. W. (2021). Fast covariance parameter estimation of spatial Gaussian process models using neural networks. *Stat*, 10:e382.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge, MA.
- Han, J., Li, Y., Lin, L., Lu, J., Zhang, J., and Zhang, L. (2022). Universal approximation of symmetric and anti-symmetric functions. *Communications in Mathematical Sciences*, 20:1397–1408.
- Hazra, A. and Huser, R. (2021). Estimating high-resolution Red Sea surface temperature hotspots, using a low-rank semiparametric spatial model. *Annals of Applied Statistics*, 15:572–596.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values. *Journal of the Royal Statistical Society B*, 66:497–546.
- Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2:359–366.
- Huser, R. (2021). Editorial: EVA 2019 data competition on spatio-temporal prediction of Red Sea surface temperature extremes. *Extremes*, 24:91–104.
- Huser, R. and Davison, A. C. (2013). Composite likelihood estimation for the Brown–Resnick process. *Biometrika*, 100:511–518.
- Huser, R., Dombry, C., Ribatet, M., and Genton, M. G. (2019). Full likelihood inference for max-stable data. *Stat*, 8:e218.
- Huser, R., Stein, M. L., and Zhong, P. (2022). Vecchia likelihood approximation for accurate and fast inference in intractable spatial extremes models. *arXiv:2203.05626v1*.
- Huser, R. and Wadsworth, J. (2022). Advances in statistical modeling of spatial extremes. *Wiley Interdisciplinary Reviews: Computational Statistics*, 14:e1537.
- Innes, M. (2018). Flux: Elegant machine learning with Julia. *Journal of Open Source Software*, 3:602.
- Lehmann, E. L. and Casella, G. (1998). *Theory of Point Estimation*. Springer, New York, NY, 2nd edition.
- Lenzi, A., Bessac, J., Rudi, J., and Stein, M. L. (2023). Neural networks for parameter estimation in intractable models. *Computational Statistics & Data Analysis*, 185:107762.
- Lintusaari, J., Gutmann, M., Dutta, R., Kaski, S., and Corander, J. (2017). Fundamentals and recent developments in approximate Bayesian computation. *Systematic Biology*, 66:66–82.
- McCullagh, P. (2002). What is a statistical model. *The Annals of Statistics*, 30:1225–1310.

- Murphy, R. L., Srinivasan, B., Rao, V. A., and Ribeiro, B. (2019). Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs. In *7th International Conference on Learning Representations, ICLR*.
- Pacchiardi, L. and Dutta, R. (2022). Likelihood-free inference with generative neural networks via scoring rule minimization. *arXiv:2205.15784*.
- Padoan, S. A., Ribatet, M., and Sisson, S. A. (2010). Likelihood-based inference for max-stable processes. *Journal of the American Statistical Association*, 105:263–277.
- Radev, S. T., Mertens, U. K., Voss, A., Ardizzone, L., and Köthe, U. (2022). BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33:1452–1466.
- R Core Team (2023). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Richards, J., Tawn, J. A., and Brown, S. (2022). Modelling extremes of spatial aggregates of precipitation using conditional methods. *Annals of Applied Statistics*, 16:2693–2713.
- Rudi, J., Julie, B., and Lenzi, A. (2021). Parameter estimation with dense and convolutional neural networks applied to the FitzHugh-Nagumo ODE. In Bruna, J., Hesthaven, J., and Zdeborova, L., editors, *Proceedings of the 2nd Annual Conference on Mathematical and Scientific Machine Learning*, volume 145 of *Proceedings of Machine Learning Research*, pages 1–28. PMLR.
- Sang, H. and Genton, M. G. (2014). Tapered composite likelihood for spatial max-stable models. *Spatial Statistics*, 8:86–103.
- Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes*, 5:33–44.
- Simpson, E. S., Opitz, T., and Wadsworth, J. L. (2023). High-dimensional modeling of spatial and spatio-temporal conditional extremes using INLA and Gaussian Markov random fields. *Extremes*, to appear.
- Simpson, E. S. and Wadsworth, J. L. (2021). Conditional modelling of spatio-temporal extremes for Red Sea surface temperatures. *Spatial Statistics*, 41:100482.
- Sisson, S. A., Fan, Y., and Beaumont, M. (2018). *Handbook of Approximate Bayesian Computation*. Chapman & Hall/CRC Press, Boca Raton, FL.
- Soelch, M., Akhundov, A., van der Smagt, P., and Bayer, J. (2019). On Deep Set learning and the choice of aggregations. In Tetko, I. V., Kurková, V., Karpov, P., and Theis, F. J., editors, *Proceedings of the 28th International Conference on Artificial Neural Networks, ICANN*, Lecture Notes in Computer Science, pages 444–457. Springer.
- Stein, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, New York, NY.
- Stein, M. L., Chi, Z., and Welty, L. J. (2004). Approximating likelihoods for large spatial data sets. *Journal of the Royal Statistical Society B*, 66:275–296.
- Strasser, H. (1981). Consistency of maximum likelihood and Bayes estimates. *The Annals of Statistics*, 9:1107–1113.
- Subbotin, M. T. (1923). On the law of frequency of errors. *Matematicheskii Sbornik*, 31:296–301.
- Varin, C., Reid, N., and Firth, D. (2011). An overview of composite likelihood methods. *Statistica Sinica*, 21:5–42.
- Varin, C. and Vidoni, P. (2005). A note on composite likelihood inference and model selection. *Biometrika*, 92:519–528.
- Vecchia, A. V. (1988). Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society B*, 50:297–312.
- Wadsworth, J. L. and Tawn, J. A. (2022). Higher-dimensional spatial extremes via single-site conditioning. *Spatial Statistics*, 51:100677.

- Wagstaff, E., Fuchs, F. B., Engelcke, M., Osborne, M., and Posner, I. (2022). Universal approximation of functions on sets. *Journal of Machine Learning Research*, 23:1–56.
- Wagstaff, E., Fuchs, F. B., Engelcke, M., Posner, I., and Osborne, M. (2019). On the limitations of representing functions on sets. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 6487–6494. PMLR.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32:4–24.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. (2017). Deep sets. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Zammit-Mangion, A. and Wikle, C. K. (2020). Deep integro-difference equation models for spatio-temporal forecasting. *Spatial Statistics*, 37:100408.
- Zhang, H. (2004). Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association*, 99:250–261.
- Zhou, D. (2018). Universality of deep convolutional neural networks. *Applied and Computational Harmonic Analysis*, 48:787–794.

Supplementary Material for “Likelihood-Free Parameter Estimation with Neural Bayes Estimators”

In Section S1, we prove that Bayes estimators are invariant to permutations of replicated data. In Section S2, we illustrate how neural Bayes estimators depend on the sample size. In Section S3, we show that the Bayes estimator is consistent for θ in the example of Section 2.2.3 of the main text, but that the one-at-a-time estimator is not. In Sections S4 and S5, we conduct experiments to illustrate the benefits of “on-the-fly” simulation and pre-training, respectively. In Section S6, we provide implementation details for the models considered in Section 3 of the main text. In Section S7, we describe our analysis for the Gaussian process model with known smoothness parameter. In Section S8, we detail the bootstrap techniques used in Section 4 of the main text. In Section S9, we analyse sea-surface temperature data in the Red Sea, where the data are measured irregularly in space. In Section S10, we provide empirical evidence that our neural Bayes estimators can be used on data that are different to those used during training. Finally, in Section S11, we provide additional exploratory and diagnostic plots for the experiments given in the main text.

S1 Permutation-invariance of Bayes estimators

Theorem 1. *Let $\mathcal{P} \equiv \{P_{\theta} : \theta \in \Theta\}$ denote a class of distributions parameterised by θ , and let $\Omega(\cdot)$ denote a prior measure for θ . For a strictly convex loss function $L(\cdot, \cdot)$, the Bayes estimator $\hat{\theta}^*(\cdot)$ for replicated data (i.e., data that are conditionally independent given θ) is unique and permutation invariant with probability 1 provided that*

- (i) *the Bayes risk of $\hat{\theta}^*(\cdot)$ is finite, and*
- (ii) *the distribution $P_{\theta}(\cdot)$ is absolutely continuous with respect to the marginal distribution $\int P_{\theta}(\cdot) d\Omega(\theta)$.*

Proof. The theorem is a straightforward extension of Lehmann and Casella (1998, Ch. 4, Cor. 1.4) for the special case where the posterior distribution is invariant to the ordering of the data, as is the case with conditionally-independent replicates. Here, for ease of exposition, we give the proof for the case where the prior distribution admits a density $p(\theta)$ and where the distribution of the data under $P_{\theta}(\cdot)$ admits a density $f(\cdot | \theta)$ with respect to Lebesgue measure. This proof is based on the following two properties:

- (a) (Permutation invariance) The posterior density for θ is given by

$$p(\theta | \mathbf{z}^{(m)}) = \frac{p(\theta) \prod_{i=1}^m f(\mathbf{z}_i | \theta)}{\int_{\Theta} p(\theta) \prod_{i=1}^m f(\mathbf{z}_i | \theta) d\theta} = p(\theta | \mathbf{z}_{\pi}^{(m)}), \quad (\text{S1})$$

where $\mathbf{z}_{\pi}^{(m)} \equiv \pi(\mathbf{z}^{(m)}) = (\mathbf{z}'_{\pi(1)}, \dots, \mathbf{z}'_{\pi(m)})'$ is a permutation under $\pi(\cdot)$ of the conditionally independent replicates in $\mathbf{z}^{(m)} \equiv (\mathbf{z}'_1, \dots, \mathbf{z}'_m)'$.

- (b) (Minimisation of posterior expected loss) For a given $\mathbf{z}^{(m)}$, a Bayes estimator $\hat{\boldsymbol{\theta}}^*(\mathbf{z}^{(m)})$ is a minimiser over $\hat{\boldsymbol{\theta}}(\mathbf{z}^{(m)})$ of

$$\int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}(\mathbf{z}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta}. \quad (\text{S2})$$

Now, by optimality of the Bayes estimator, we have that

$$\begin{aligned} \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta} &\leq \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^* \circ \pi(\mathbf{z}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta} \\ &= \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}_{\pi}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta}, \end{aligned}$$

since $\hat{\boldsymbol{\theta}}^* \circ \pi(\cdot)$ is a different (potentially non-Bayes) estimator formed by the composition of the Bayes estimator and a permutation function. Similarly, for observations $\mathbf{z}_{\pi}^{(m)}$, we have

$$\begin{aligned} \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}_{\pi}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}_{\pi}^{(m)})d\boldsymbol{\theta} &= \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}_{\pi}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta} \\ &\leq \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^* \circ \pi^{-1}(\mathbf{z}_{\pi}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta} \\ &= \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta}, \end{aligned}$$

where the first equality is due to (S1). Therefore,

$$\begin{aligned} \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta} &= \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^*(\mathbf{z}_{\pi}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta} \\ &= \int_{\Theta} L(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^* \circ \pi(\mathbf{z}^{(m)}))p(\boldsymbol{\theta} | \mathbf{z}^{(m)})d\boldsymbol{\theta}. \end{aligned}$$

Since $L(\cdot, \cdot)$ is strictly convex, if conditions (i) and (ii) hold, the Bayes estimator is unique with probability 1 (Lehmann and Casella, 1998, Ch. 4, Cor. 1.4), and

$$\hat{\boldsymbol{\theta}}^*(\mathbf{z}^{(m)}) = \hat{\boldsymbol{\theta}}^* \circ \pi(\mathbf{z}^{(m)}), \quad (\text{S3})$$

with probability 1 for any $\mathbf{z}^{(m)}$ and permutation function $\pi(\cdot)$. \square

S2 Variable sample sizes

Recall that our neural Bayes estimators based on the DeepSets representation can be applied to data sets of arbitrary size. However, the neural Bayes estimator for replicated data is only (approximately) Bayes for the number of replicates used during training.

To demonstrate this dependence on m , we trained neural Bayes estimators with a range of sample sizes, where the inferential target is θ from replicated data generated from a $N(0, \theta)$ distribution. We set the prior distribution over the variance θ to be an InverseGamma(2, 2)




Figure S1: The Bayes risk $r_{\Omega}(\hat{\theta}(\cdot; \cdot))$ plotted against the sample size m , for the theoretical Bayes estimator (green) and several neural Bayes estimators (orange, red, purple) for θ from $N(0, \theta)$ data.

distribution. The estimators $\hat{\theta}(\cdot; \gamma_5^*)$ and $\hat{\theta}(\cdot; \gamma_{150}^*)$ were trained with training data sets consisting of exactly 5 or 150 independent replicates, respectively, while the estimator $\hat{\theta}(\cdot; \gamma_{1:150}^*)$ was trained with training data sets containing a variable number of replicates and with M in (7) a discrete uniform random variable with support between 1 and 150 inclusive. Figure S1 shows that estimators trained with fixed m are (approximately) Bayes only for that choice of m . On the other hand, the estimator $\hat{\theta}(\cdot; \gamma_{1:150}^*)$ is not (approximately) Bayes for any m , but performs well for all $m \in \{1, \dots, 150\}$. Combining the approaches of (6) and (7) is therefore an effective strategy for constructing a neural Bayes estimator that is quasi-Bayes for all sample sizes.

S3 Bayes estimator vs. one-at-a-time estimator

The model of Section 2.2.3 of the main text is given by,

$$\begin{aligned} Z_1, \dots, Z_m \mid \theta &\stackrel{\text{iid}}{\sim} \text{Unif}(0, \theta), \\ \theta \mid \alpha, \beta &\sim \text{Pareto}(\alpha, \beta). \end{aligned}$$

Here, we analyse the consistency of two estimators for the above model. Specifically, we consider the Bayes estimator and the ‘one-at-a-time’ estimator, which applies the single-replicate Bayes estimator to each replicate independently and averages the resulting estimates. For this model, under the absolute-error loss, these estimators are respectively given by,

$$\begin{aligned} \hat{\theta}_{\text{Bayes}}(\mathbf{Z}^{(m)}) &= \text{median}(\theta \mid \mathbf{Z}) = 2^{\frac{1}{\alpha+m}} \max(Z_1, \dots, Z_m, \beta), \\ \hat{\theta}_0(\mathbf{Z}^{(m)}) &= \frac{1}{m} \sum_{i=1}^m 2^{\frac{1}{\alpha+1}} \max(Z_i, \beta) = \frac{2^{\frac{1}{\alpha+1}}}{m} \sum_{i=1}^m \max(Z_i, \beta), \end{aligned}$$

where $\mathbf{Z}^{(m)} \equiv (Z_1, \dots, Z_m)'$ is the observed data containing m independent replicates.

Consistency of $\hat{\theta}_{\text{Bayes}}(\cdot)$. Consider first the distribution of $\max(Z_1, \dots, Z_m)$:

$$\begin{aligned} \Pr(\max(Z_1, \dots, Z_m) \leq x) &= \Pr(Z_1 \leq x, \dots, Z_m \leq x) \\ &= F_{Z_i}(x)^m \\ &= \begin{cases} 0, & x < 0 \\ \left(\frac{x}{\theta}\right)^m, & x \in [0, \theta] \\ 1, & x > \theta \end{cases} \\ &\xrightarrow{m \rightarrow \infty} \begin{cases} 0, & x < \theta \\ 1, & x \geq \theta. \end{cases} \end{aligned}$$

Therefore, $\max(Z_1, \dots, Z_m) \xrightarrow{D} \theta$ (convergence in distribution) which is a fixed value, and therefore implies $\max(Z_1, \dots, Z_m) \xrightarrow{P} \theta$ (convergence in probability), which in turn implies that $\max(Z_1, \dots, Z_m, \beta) \xrightarrow{P} \theta$ for $\beta < \theta$. Thus, the estimator,

$$\hat{\theta}_{\text{Bayes}}(\mathbf{Z}^{(m)}) = 2^{\frac{1}{\alpha+m}} \max(Z_1, \dots, Z_m, \beta) \xrightarrow{P} \theta,$$

and is therefore consistent for θ .

Inconsistency of $\hat{\theta}_0(\cdot)$. From basic properties of the uniform distribution, we have that

$$\begin{aligned} \mathbb{E}(\max(Z_i, \beta)) &= \mathbb{E}(Z_i) \Pr(Z_i > \beta) + \beta \Pr(Z_i < \beta) \\ &= \frac{\theta}{2} \left(1 - \frac{\beta}{\theta}\right) + \beta \frac{\beta}{\theta} \\ &= \frac{\theta}{2} - \frac{\beta}{2} + \frac{\beta^2}{\theta}, \end{aligned}$$

and therefore

$$\begin{aligned} \mathbb{E}(\hat{\theta}_0(\mathbf{Z}^{(m)})) &= \mathbb{E} \left(2^{\frac{1}{\alpha+1}} \frac{1}{m} \sum_{i=1}^m \max(Z_i, \beta) \right) \\ &= 2^{\frac{1}{\alpha+1}} \left(\frac{\theta}{2} - \frac{\beta}{2} + \frac{\beta^2}{\theta} \right) \neq \theta, \end{aligned}$$

for all $\alpha, \beta > 0$ and for all $m > 0$. Therefore, $\hat{\theta}_0(\cdot)$ is not consistent for θ .

It can be shown that $\hat{\theta}_{\text{Bayes}}(\cdot)$ has an asymptotic generalised extreme value (GEV) distribution, while $\hat{\theta}_0(\cdot)$ is asymptotically normal; therefore these two estimators also have very different distributional properties.

S4 Further details on “simulation-on-the-fly”

For many models of interest, data can be simulated periodically during training in a technique sometimes referred to as “simulation-on-the-fly” (Chan et al., 2018). Here, we conduct an

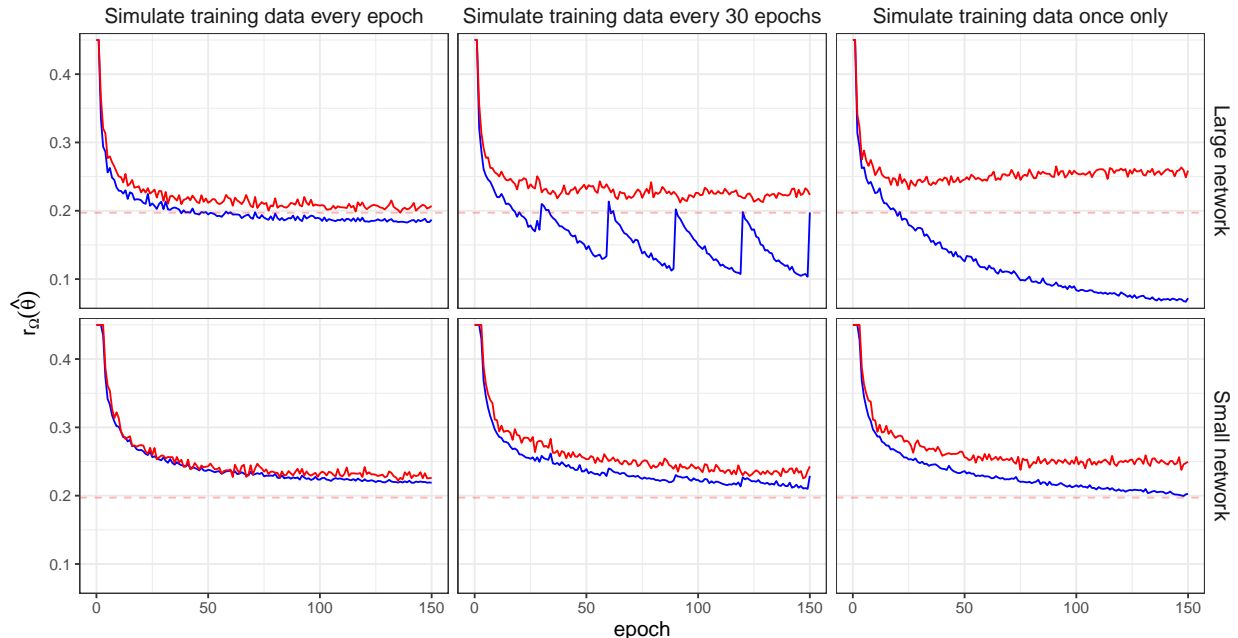


Figure S2: The risk with respect to the absolute error loss approximated using the training (blue) and validation (red) data sets for two neural network architectures (rows) and for three simulation-on-the-fly training routines (columns). The horizontal red dashed line shows the minimum validation risk achieved in the study, and all axes are fixed to facilitate comparison.

experiment to illustrate the benefits of this strategy. We train several neural Bayes estimators for the parameters of the Gaussian process model with known smoothness, as detailed in Section S7. We consider three simulation-on-the-fly training routines; refreshing the training data every epoch (using our definition of “epoch” as given in Section 3.1 of the main text), refreshing the training data every 30 epochs, and keeping the training data fixed. We also consider two neural network architectures that differ only in the width of each layer; the first contains roughly 600,000 trainable parameters, while the second contains an order-of-magnitude fewer trainable parameters. All other components of the experiment are held fixed.

Figure S2 shows the risk evaluated during training for each combination of simulation-on-the-fly training routine and network architecture. Comparing the columns of the figure reveals that regularly refreshing the training data reduces overfitting, as can be seen from the reduced discrepancy between the training and validation risks. This property allows one to use larger neural networks with higher representational capacity that are prone to overfitting when the training data are fixed. This leads to an overall lower out-of-sample error with the larger network (compare the two panels in the first column of Figure S2).

S5 Motivation for pre-training

Recall that the parameters of a neural network trained for one task can be used as initial values for the parameters of another neural network intended for a slightly different task.




Figure S3: The risk function with respect to the absolute error loss evaluated on the validation set for two neural Bayes estimators, each trained with $m = 30$ independent replicates available for each parameter configuration. One estimator (solid line) was pre-trained with a neural Bayes estimator trained with $m = 1$ replicate, while the other (dashed-dotted line) was randomly initialised. The horizontal dashed line shows the minimum risk achieved by the pre-trained estimator.

This is known as pre-training (Goodfellow et al., 2016, Ch. 8). To illustrate the benefits of pre-training in the context of neural Bayes estimation, Figure S3 shows the validation risk during training for two neural Bayes estimators for the Gaussian process model of Section 3.2 of the main text. Both networks were trained with $m = 30$ independent replicates for each parameter configuration, but one of these was pre-trained with a neural Bayes estimator trained with $m = 1$ replicate. Although both estimators eventually converge to a similar validation risk, the pre-trained estimator converges in drastically fewer epochs.

S6 Model simulation and density functions

We now describe how we simulate from the models considered in Section 3 of the main text, and we provide the density functions required for the likelihood-based estimators.

S6.1 Gaussian process model

Recall our formulation of the Gaussian process model from Section 3.2 of the main text. We consider a classical spatial model, the linear Gaussian-Gaussian model,

$$Z_{ij} = Y_i(\mathbf{s}_j) + \epsilon_{ij}, \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad (11 \text{ revisited})$$

where $\mathbf{Z}_i \equiv (Z_{i1}, \dots, Z_{in})'$ are data observed at locations $\{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ on a spatial domain \mathcal{D} , $\{Y_i(\cdot)\}$ are i.i.d. spatially-correlated mean-zero Gaussian processes, and $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$ is Gaussian white noise. An important component of the model is the covariance function, $C(\mathbf{s}, \mathbf{u}) \equiv \text{cov}(Y_i(\mathbf{s}), Y_i(\mathbf{u}))$, for $\mathbf{s}, \mathbf{u} \in \mathcal{D}$ and $i = 1, \dots, m$. Here, we use the popular

Algorithm S1 Simulation from the mean-zero Gaussian process model

Require: Parameters ρ , ν , and σ_ϵ ; sample size m .

- 1: Compute \mathbf{C} , the covariance matrix with entries $C(\mathbf{s}_j, \mathbf{s}_{j'})$ for $j, j' = 1, \dots, n$.
 - 2: Compute \mathbf{L} , the lower Cholesky factor of \mathbf{C} .
 - 3: **for** $i = 1, \dots, m$ **do**
 - 4: Simulate $\mathbf{w}_i \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$.
 - 5: Simulate $\boldsymbol{\epsilon}_i \sim \mathbf{N}(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I})$.
 - 6: Set $\mathbf{y}_i = \mathbf{L}\mathbf{w}_i$.
 - 7: Set $\mathbf{z}_i = \mathbf{y}_i + \boldsymbol{\epsilon}_i$.
 - 8: **end for**
 - 9: Return $\mathbf{z} \equiv (\mathbf{z}'_1, \dots, \mathbf{z}'_m)'$.
-

isotropic Matérn covariance function which, for $\mathbf{h} \equiv \mathbf{s} - \mathbf{u}$, is

$$C(\mathbf{s}, \mathbf{u}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\|\mathbf{h}\|}{\rho} \right)^\nu K_\nu \left(\frac{\|\mathbf{h}\|}{\rho} \right), \quad (12 \text{ revisited})$$

with σ^2 the marginal variance, $\Gamma(\cdot)$ is the gamma function, $K_\nu(\cdot)$ the Bessel function of the second kind of order ν , and $\rho > 0$ and $\nu > 0$ the range and smoothness parameters, respectively. Following Gerber and Nychka (2021), we fix $\sigma^2 = 1$. This leaves three (possibly) unknown parameters that need to be estimated: $\boldsymbol{\theta} \equiv (\sigma_\epsilon, \rho, \nu)'$.

Let \mathbf{L} denote the lower Cholesky factor of the covariance matrix

$$\mathbf{C} \equiv (C(\mathbf{s}_j, \mathbf{s}_{j'}) : j, j' = 1, \dots, n).$$

Then, simulation from (11) proceeds with Algorithm S1 and, for a single replicate, the log-likelihood function for (11) is

$$\begin{aligned} \ell(\boldsymbol{\theta}; \mathbf{z}_i) &= -\frac{n}{2} \log 2\pi - \frac{1}{2} \log |\boldsymbol{\Sigma}_i| - \frac{1}{2} \mathbf{z}'_i \boldsymbol{\Sigma}_i^{-1} \mathbf{z}_i \\ &= -\frac{n}{2} \log 2\pi - \sum_{j=1}^n \log L_{jj} - \frac{1}{2} \mathbf{u}'_i \mathbf{u}_i, \end{aligned}$$

where $\boldsymbol{\Sigma} \equiv \text{cov}(\mathbf{z}_i, \mathbf{z}_i) = \mathbf{C} + \sigma_\epsilon^2 \mathbf{I}$ with \mathbf{I} the identity matrix, L_{jj} is the j th diagonal element of \mathbf{L} , and \mathbf{u}_i is the solution to $\mathbf{L}\mathbf{u}_i = \mathbf{z}_i$ which, since \mathbf{L} is lower-triangular, can be computed using an efficient forward solve.

S6.2 Schlather's max-stable model

Recall Schlather's max-stable model from Section 3.3 of the main text, given by

$$Z_{ij} = \bigvee_{k \in \mathbb{N}} \zeta_{ik} \max\{0, Y_{ik}(\mathbf{s}_j)\}, \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad (13 \text{ revisited})$$

where \bigvee denotes the maximum over the indexed terms, $\mathbf{Z}_i \equiv (Z_{i1}, \dots, Z_{in})'$ are data observed at locations $\{\mathbf{s}_1, \dots, \mathbf{s}_n\} \subset \mathcal{D}$, $\{\zeta_{ik} : k \in \mathbb{N}\}$ for $i = 1, \dots, m$ are i.i.d. Poisson point

Algorithm S2 Simulation from Schlather’s max-stable model

Require: Positive R ; parameters ρ and ν ; sample size m .

- 1: Compute \mathbf{C} , the covariance matrix with entries $C(\mathbf{s}_j, \mathbf{s}_{j'})$ for $j, j' = 1, \dots, n$.
 - 2: Compute \mathbf{L} , the lower Cholesky factor of \mathbf{C} .
 - 3: **for** $i = 1, \dots, m$ **do**
 - 4: Set $\mathbf{z}_i = \mathbf{0}$.
 - 5: Simulate $\zeta_i^{-1} \sim \text{Exp}(1)$.
 - 6: **while** $R\zeta_i > \min\{\mathbf{z}_i\}$ **do**
 - 7: Denote the current iteration by k .
 - 8: Set $\mathbf{y}_{ik} = \mathbf{L}\mathbf{w}_k$ where $\mathbf{w}_k \sim \text{N}(\mathbf{0}, \mathbf{I})$.
 - 9: Set $z_{ij} = \max(z_{ij}, \zeta_i y_{ikj})$ for $j = 1, \dots, n$.
 - 10: Update ζ_i^{-1} by $\zeta_i^{-1} + e_k$, where $e_k \sim \text{Exp}(1)$.
 - 11: **end while**
 - 12: Set $\mathbf{z}_i = (z_{i1}, \dots, z_{in})'$.
 - 13: **end for**
 - 14: Return $\mathbf{z} \equiv (\mathbf{z}'_1, \dots, \mathbf{z}'_m)'$.
-

processes on $(0, \infty)$ with intensity measure $d\Lambda(\zeta) = \zeta^{-2}d\zeta$, and $\{Y_{ik}(\cdot) : i = 1, \dots, m, k \in \mathbb{N}\}$ are i.i.d. mean-zero Gaussian processes scaled so that $\mathbb{E}[\max\{0, Y_{ik}(\cdot)\}] = 1$. We model each $Y_{ik}(\cdot)$ using the Matérn covariance function, (12), with $\sigma^2 = 1$. Hence, $\boldsymbol{\theta} \equiv (\rho, \nu)'$.

We use Algorithm S2 for approximate simulation from (13), which was given by Schlather (2002) (see also Dey and Yan, 2016, alg. 1.2.2). The tuning parameter R involves a trade-off between computational efficiency (favouring small R) and accuracy (favouring large R). Schlather (2002) recommends the use of $R = 3$; conservatively, we set $R = 3.5$.

Recall that we compare our neural Bayes estimator to the pairwise maximum *a posteriori* (MAP) estimator, which is simply the MAP estimator with the log-likelihood function replaced with the pairwise log-likelihood function. Further recall that, in general, the pairwise log-likelihood function for the i th replicate is

$$\ell_{\text{PL}}(\boldsymbol{\theta}; \mathbf{z}_i) \equiv \sum_{j=1}^{n-1} \sum_{j'=j+1}^n \log f(z_{ij}, z_{ij'} \mid \boldsymbol{\theta}), \quad (14 \text{ revisited})$$

where $f(\cdot, \cdot \mid \boldsymbol{\theta})$ denotes the bivariate probability density function for pairs in \mathbf{z}_i . The bivariate cumulative distribution function for max-stable models with unit Fréchet margins (e.g., Schlather’s model) is of the form (Huser, 2013, pg. 231–232),

$$F(z_1, z_2 \mid \boldsymbol{\theta}) = \exp\{-V(z_1, z_2)\},$$

where $V(\cdot, \cdot)$ is the so-called exponent function. Then, by the chain rule,

$$f(z_1, z_2 \mid \boldsymbol{\theta}) = \frac{\partial}{\partial z_1 \partial z_2} F(z_1, z_2) = \{V_1(z_1, z_2)V_2(z_1, z_2) - V_{12}(z_1, z_2)\} \exp\{-V(z_1, z_2)\},$$

where V_1 , V_2 , and V_{12} denote the partial derivatives of V with respect to z_1 , z_2 , and both z_1 and z_2 , respectively. Now, for $(z_1, z_2)'$ drawn from Schlather’s max-stable model (13) and

observed at locations $\mathbf{s}_1, \mathbf{s}_2 \in \mathcal{D}$, respectively, we have

$$\begin{aligned} V(z_1, z_2) &= \left(\frac{1}{z_1} + \frac{1}{z_2} \right) \left[1 - \frac{1}{2} \left(1 - \frac{1}{z_1 + z_2} \{ z_1^2 - 2z_1z_2\psi + z_2^2 \}^{1/2} \right) \right], \\ V_1(z_1, z_2) &= -\frac{1}{2z_1^2} + \frac{1}{2} \left(\frac{\psi}{z_1} - \frac{z_2}{z_1^2} \right) (z_1^2 - 2z_1z_2\psi + z_2^2)^{-1/2}, \\ V_2(z_1, z_2) &= V_1(z_2, z_1), \\ V_{12}(z_1, z_2) &= -\frac{1}{2} (1 - \psi^2) (z_1^2 - 2z_1z_2\psi + z_2^2)^{-3/2}, \end{aligned}$$

where $\psi \equiv \text{corr}(Y(\mathbf{s}_1), Y(\mathbf{s}_2))$ depends on $\boldsymbol{\theta}$.

S6.3 Spatial conditional extremes model

Recall our formulation of the spatial conditional extremes model from Section 3.4, which describes the behaviour of a process conditional on it being greater than some threshold, u , at a conditioning site, $\mathbf{s}_0 \in \mathcal{D}$. Specifically,

$$Z_{ij} \mid Z_{i0} > u \stackrel{\text{d}}{=} a(\mathbf{h}_j, Z_{i0}) + b(\mathbf{h}_j, Z_{i0}) Y_i(\mathbf{s}_j), \quad i = 1, \dots, m, \quad j = 1, \dots, n, \quad (15 \text{ revisited})$$

where ‘ $\stackrel{\text{d}}{=}$ ’ denotes equality in distribution, $Z_{i0} \equiv Z_i(\mathbf{s}_0)$ is the value of the data process at the conditioning site, $\mathbf{h}_j \equiv \mathbf{s}_j - \mathbf{s}_0$, and $Z_{i0} - u \mid Z_{i0} > u$ is a unit exponential random variable that is independent of the residual process, $Y_i(\cdot)$. We model $a(\cdot, \cdot)$ and $b(\cdot, \cdot)$ using parametric forms proposed by [Wadsworth and Tawn \(2022\)](#),

$$\begin{aligned} a(\mathbf{h}, z) &= z \exp\{ -(\|\mathbf{h}\|/\lambda)^\kappa \}, \quad \lambda > 0, \quad \kappa > 0, \\ b(\mathbf{h}, z) &= 1 + a(\mathbf{h}, z)^\beta, \quad \beta > 0, \end{aligned}$$

where $\mathbf{h} \equiv \mathbf{s} - \mathbf{s}_0$ for $\mathbf{s} \in \mathcal{D}$ and $z \in \mathbb{R}$. We construct the residual process, $Y_i(\cdot)$, by first constructing the process $\tilde{Y}_i^{(0)}(\cdot) \equiv \tilde{Y}_i(\cdot) - \tilde{Y}_i(\mathbf{s}_0)$, where $\tilde{Y}_i(\cdot)$ is a mean-zero Gaussian process with Matérn covariance function and with unit marginal variance, and then marginally transforming it to the scale of a Subbotin distribution. We model δ as decaying from 2 to 1 as the distance to \mathbf{s}_0 increases; specifically,

$$\delta(\mathbf{h}) = 1 + \exp\{ -(\|\mathbf{h}\|/\delta_1)^2 \}, \quad \delta_1 > 0. \quad (16 \text{ revisited})$$

To describe simulation from (15), it is necessary to make several definitions. First, let $\tilde{C}(\cdot, \cdot)$ denote the covariance function of the i.i.d. processes $\{\tilde{Y}_i(\cdot) : i = 1, \dots, m\}$. Second, define $\tilde{Y}_i^{(01)}(\cdot) \equiv \tilde{Y}_i^{(0)}(\cdot)/\tilde{\sigma}^{(0)}(\cdot)$, where $\tilde{\sigma}^{(0)}(\cdot)$ is the standard-deviation process of $\tilde{Y}^{(0)}(\cdot)$. Finally, since $\tilde{Y}_i^{(01)}(\cdot)$ is a mean-zero Gaussian process with unit marginal variance, we obtain the Subbotin-scale residual process as $Y_i(\cdot) \equiv t(\tilde{Y}_i^{(01)}(\cdot))$, where $t(\cdot) \equiv Q_S(\Phi(\cdot))$, $\Phi(\cdot)$ is the distribution function of the standard univariate Gaussian distribution, and $Q_S(\cdot)$ is the quantile function of the delta-Laplace distribution (see below). Then, simulation from (15) proceeds with Algorithm S3.

Algorithm S3 Simulation from the spatial conditional extremes model

Require: Parameters $\rho, \nu, \kappa, \lambda, \beta, \mu, \tau$, and δ_1 ; extremal threshold u ; sample size m .

- 1: Compute $\boldsymbol{\delta} = (\delta(\mathbf{h}_1), \dots, \delta(\mathbf{h}_n))'$.
 - 2: Compute $\tilde{\boldsymbol{\sigma}}^{(0)} = (\sigma^{(0)}(\mathbf{h}_1), \dots, \sigma^{(0)}(\mathbf{h}_n))'$.
 - 3: Compute $\tilde{\mathbf{C}} \equiv (\tilde{C}(\mathbf{s}_j, \mathbf{s}_{j'}) : j, j' = 1, \dots, n)$.
 - 4: Compute $\tilde{\mathbf{L}}$, the lower Cholesky factor of $\tilde{\mathbf{C}}$.
 - 5: **for** $i = 1, \dots, m$ **do**
 - 6: Simulate $z_{i0} = u + e_i$, where $e_i \sim \text{Exp}(1)$.
 - 7: Simulate $\tilde{\mathbf{y}}_i = \tilde{\mathbf{L}}\mathbf{w}_i$, where $\mathbf{w}_i \sim \text{N}(\mathbf{0}, \mathbf{I})$.
 - 8: Set $\tilde{\mathbf{y}}_i^{(0)} = \tilde{\mathbf{y}}_i - \tilde{\mathbf{y}}_{i0}$, where subtraction is elementwise.
 - 9: Set $\tilde{\mathbf{y}}_i^{(01)} = \tilde{\mathbf{y}}_i^{(0)} / \tilde{\boldsymbol{\sigma}}^{(0)}$, where division is elementwise. Note that the element of $\tilde{\boldsymbol{\sigma}}^{(0)}$ associated with \mathbf{s}_0 is equal to zero; since $Z_{i0} \equiv Z_i(\mathbf{s}_0)$ is simulated from independently, this is not a problem in practice.
 - 10: Set $\mathbf{y}_i = t(\tilde{\mathbf{y}}_i^{(01)})$ where $t(\cdot)$ is applied elementwise.
 - 11: Set $z_{ij} = a(\mathbf{h}_j, z_{i0}) + b(\mathbf{h}_j, z_{i0})y_{ij}$, for $j = 1, \dots, n$.
 - 12: Set $\mathbf{z}_i = (z_{i0}, z_{i1}, \dots, z_{in})'$.
 - 13: **end for**
 - 14: Return $\mathbf{z} \equiv (\mathbf{z}'_1, \dots, \mathbf{z}'_m)'$.
-

S6.3.1 The delta-Laplace distribution

A delta-Laplace (generalised Gaussian) distribution (Subbotin, 1923) is parameterised by a location parameter $\mu \in \mathbb{R}$, a scale parameter $\tau > 0$, and a shape parameter $\delta > 0$. For $y \in \mathbb{R}$ and $p \in [0, 1]$, its density, distribution, and quantile functions are

$$\begin{aligned} f_S(y \mid \mu, \tau, \delta) &= \frac{\delta}{2\tau\Gamma(1/\delta)} \exp\left(-\left|\frac{y-\mu}{\tau}\right|^\delta\right), \\ F_S(y \mid \mu, \tau, \delta) &= \frac{1}{2} + \text{sign}(y-\mu) \frac{1}{2\Gamma(1/\delta)} \gamma\left(1/\delta, \left|\frac{y-\mu}{\tau}\right|^\delta\right), \\ Q_S(p \mid \mu, \tau, \delta) &= \text{sign}(p-0.5) Q_G\left(2|p-0.5|; \frac{1}{\delta}, \frac{1}{(k\tau)^\delta}\right)^{1/\delta} + \mu, \end{aligned}$$

where $\gamma(\cdot)$ is the unnormalised incomplete lower gamma function and $Q_G(\cdot)$ is the quantile function of the Gamma distribution.

S7 Gaussian process model with known smoothness

We now consider the Gaussian process model of Section 3.2 of the main text with the smoothness parameter, ν , fixed to 1. This leaves two estimable parameters, σ_ϵ and ρ , which are usually well-identified from just a single field. Note that this model was also considered by Gerber and Nychka (2021).




Figure S4: The training/validation/test parameter sets used in Section S7, shown in purple, green, and yellow, respectively. The small subset of test parameters chosen to represent a range of model behaviours is highlighted in red.

Rather than randomly sampling from a prior measure $\Omega(\cdot)$, Gerber and Nychka (2021) used a stratified, deterministic design (we defer to Gerber and Nychka, 2021, for details). The use of a deterministic design still implies a prior measure but, unless the design is relatively simple (e.g., a regularly spaced grid), this prior is difficult to formulate analytically. Although one should sample from $\Omega(\cdot)$, to facilitate a comparison with the work of Gerber and Nychka (2021), here, we use the same training and test parameter sets that were used in that work, which contain 40,200 and 900 parameter configurations, respectively. These parameter sets are shown in Figure S4, with a small subset of test parameters chosen to represent a range of model behaviours highlighted in red.

In our experiments, we compare the neural Bayes estimators to a likelihood-based estimator, namely, the maximum *a posteriori* (MAP) estimator. The MAP estimator requires the prior measure, $\Omega(\cdot)$, to be known; since this is not known here, and because we found that prior information is needed to avoid unreasonable maximum likelihood estimates in certain cases, here we let $\sigma_\epsilon \sim \text{Unif}(0.001, 1.5)$ and $\rho \sim \text{Unif}(1, 30)$.

Figure S5 shows the test risk with respect to the squared-error loss and the prior implied by the deterministic parameter-space design, against the number of independent replicates, m . We also compare the empirical joint distributions of the estimators for several test parameters chosen to represent a range of model behaviours: Figure S6 shows the true parameters (red cross) and estimates obtained by applying each estimator to 100 sets of $m = 150$ independent model realisations. Overall, these results are in agreement with those given for the more complicated Gaussian process model considered in the main text, namely, that $\hat{\theta}_{\text{DS}}(\cdot)$ is a significant improvement over the prior art, $\hat{\theta}_0(\cdot)$.




Figure S5: The risk with respect to the 0–1 loss against the number of replicates, m , evaluated using the parameter vectors in the test parameter set, for the estimators considered in Section S7. The estimators $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{\text{DS}}(\cdot)$, and the MAP estimator are in green, orange, and purple, respectively.




Figure S6: The empirical joint distribution of the estimators considered in Section S7, where each panel corresponds to one test parameter vector (see Figure S4) intentionally chosen to capture a range of model behaviours (red). In accordance with Figure S4, the panels are ordered such that ρ increases moving from bottom to top, and σ_ϵ increases moving from left to right. The estimators $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{\text{DS}}(\cdot)$, and the MAP estimator are shown in green, orange, and purple, respectively. Each estimate was obtained from a simulated data set of size $m = 150$.

S8 Bootstrapping in the Red Sea study

Here, we detail the bootstrap techniques used in Section 4 of the main text. Denote the observed data set as \mathbf{z}^{obs} (note that \mathbf{z}^{obs} contains m replicates). To quantify the uncertainty associated with the estimate $\hat{\boldsymbol{\theta}}(\mathbf{z}^{\text{obs}})$ and the model-based estimates of the proportion of threshold exceedances in each region, we take the following approach.

- (i) Generate B pseudo-replicate data sets from the original data set (we set $B = 400$) and denote these as $\mathbf{z}_b^{\text{obs}}$, $b = 1, \dots, B$. To account for temporal dependence, evidence of which is displayed in Figure S16, we use a non-parametric-block bootstrap (Huser and Davison, 2014). Specifically, based on the assumption that fields between seasons are independent, we sample entire seasons with replacement.
- (ii) Using the pseudo-replicate data, compute bootstrap estimates $\hat{\boldsymbol{\theta}}(\mathbf{z}_1^{\text{obs}}), \dots, \hat{\boldsymbol{\theta}}(\mathbf{z}_B^{\text{obs}})$; the 2.5 and 97.5 percentiles yield 95% bootstrap confidence intervals for $\boldsymbol{\theta}$.
- (iii) For each bootstrap estimate, simulate some large number, 1000 say, fields from $f(\mathbf{z}^{(m)} | \hat{\boldsymbol{\theta}}(\mathbf{z}_b^{\text{obs}}))$, generating B simulated data sets each with 1000 field replicates.
- (iv) Using these simulated data sets, compute B estimates of the proportion of threshold exceedances in each region, from which bootstrap confidence intervals are obtained through the appropriate quantiles.

Note that the finite sample size of \mathbf{z}^{obs} makes the empirical estimates of the proportion of threshold exceedances in each region susceptible to noise; the pseudo replicates of \mathbf{z}^{obs} also account for this uncertainty, since they are each approximately of the same size as \mathbf{z}^{obs} (approximately since the blocks vary in size).

S9 Red Sea data: Irregularly-spaced locations

Here we consider a second data set in the Red Sea, obtained by randomly selecting 678 irregularly-spaced locations from the full data set. To cater for irregular spatial data, which cannot be processed with CNNs directly, Gerber and Nychka (2021) suggest passing the empirical variogram as input to a CNN. The variogram, however, is based on the second moment of the data only, and it is unlikely to contain sufficient information for reliably estimating parameters in complex models. Instead, we construct our neural Bayes estimator using dense neural networks (DNNs). Specifically, we use a piecewise neural Bayes estimator in the same form as $\hat{\boldsymbol{\theta}}_{\text{DS}}(\cdot)$ of the main text, with $\boldsymbol{\psi}(\cdot)$ and $\boldsymbol{\phi}(\cdot)$ constructed using the first two and last two rows of Table S1, respectively.

DNN-based estimators ignore the dependence structure present in spatial data, treating each datum as an independent input, and this typically results in a loss of statistical efficiency (see, e.g., Rudi et al., 2021). It is difficult to directly compare the CNN-based estimator of Section 4 of the main text with a DNN-based estimator for the present data, since these estimators are trained and assessed on different data sets. Hence, for the purpose of comparison, we also train a DNN-based estimator using the regular Red Sea data set considered in Section 4 of the main text. We denote the CNN-based and DNN-based estimators trained on the regular data set by $\hat{\boldsymbol{\theta}}_{\text{CNN:reg}}(\cdot)$ and $\hat{\boldsymbol{\theta}}_{\text{DNN:reg}}(\cdot)$, respectively, and we denote the DNN-based estimator trained for the irregular data set by $\hat{\boldsymbol{\theta}}_{\text{DNN}}(\cdot)$. Figure S7 shows the risk with

Table S1: Summary of the neural network architecture used Section S9. We report the dimension of the input array and output array of each layer. Here, the dimension of the input array to the first layer, 678, is the number of spatial locations, n , in the data set considered in Section S9, and the dimension of the array output from the final layer is the number of parameters, p , in the spatial conditional extremes model.

layer type	activation function	input shape	output shape	weights
dense	ReLU	[678]	[300]	203,700
dense	ReLU	[300]	[100]	30,100
dense	ReLU	[100]	[50]	5,050
dense	identity	[50]	[8]	408
total trainable parameters:				239,258

Table S2: Parameter estimates and 95% bootstrap confidence intervals (provided via the 2.5 and 97.5 percentiles of the bootstrap distribution) for the Red Sea data set considered in Section S9.

	κ	λ	β	ρ	ν	μ	τ	δ_1
Estimate	1.04	3.29	0.29	2.76	0.55	-0.02	0.63	2.03
2.5%	0.95	2.38	0.19	2.11	0.49	-0.10	0.53	1.40
97.5%	1.19	4.17	0.52	3.61	0.63	0.06	0.75	2.67

respect to the absolute-error loss against the number of independent replicates, m , available for the estimation of each parameter in the test set. The DNN-based estimators perform similarly irrespective of whether the data are regular or irregular and, in agreement with the observation of Rudi et al. (2021), they are less efficient than the CNN-based estimator.

Table S2 gives estimates from the irregularly-spaced Red Sea data set, which are similar to those obtained from the regularly-spaced data set (Table 1). One disagreement lies in the estimates for ν and ρ , but these parameters are negatively correlated and, hence, the two sets of estimates are likely to lead to similar models, as reflected by the diagnostic plots given in Figures S8 and S9. Overall, the DNN-based estimator seems to be adequate for doing parameter estimation from these irregular spatial data. However, it is clearly sub-optimal, and future work will investigate the use of graph neural networks (GNNs; see, e.g., Wu et al., 2021), which generalise the convolution operation to irregular spatial data.




Figure S7: The risk with respect to the absolute-error loss against the number of replicates, m , available for each parameter vector in the test set, for the estimators considered in Section S9. The estimators $\hat{\theta}_{\text{DNN}}(\cdot)$, $\hat{\theta}_{\text{DNN:reg}}(\cdot)$, and $\hat{\theta}_{\text{CNN:reg}}(\cdot)$ are shown in green, purple, and orange, respectively.

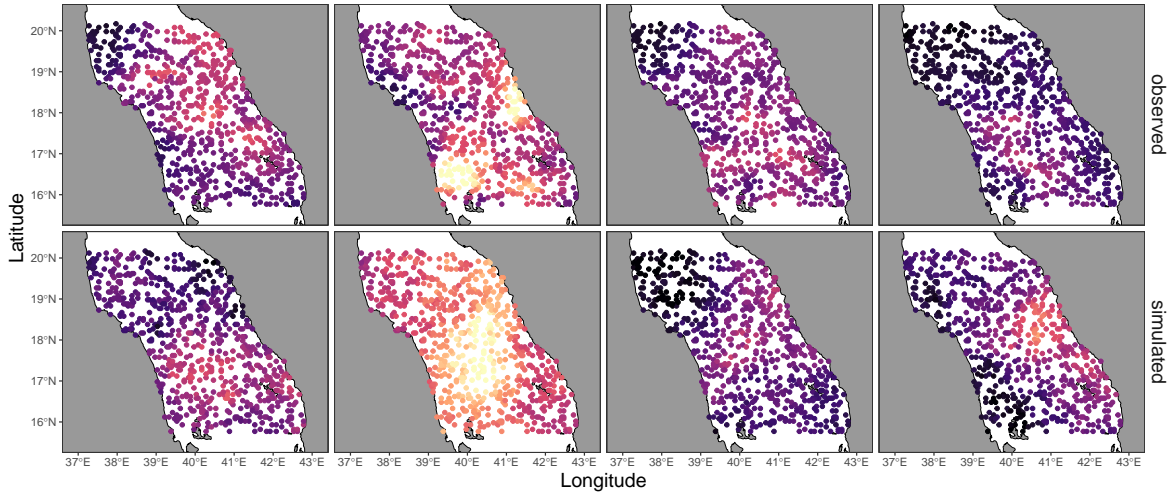


Figure S8: (Top row) Randomly-selected fields from the Red Sea data set considered in Section S9, where $Z(\mathbf{s}_0) > u$. (Bottom row) Simulations from the fitted model. The data are overlaid on a map where the grey areas correspond to land (Sudan and Eritrea to the West, Saudi Arabia to the East) and the white areas correspond to sea.

S10 Robustness to model misspecification

When the model is misspecified, the maximum likelihood (ML) estimator, under mild regularity conditions, asymptotically follows a Normal distribution centred on the value of θ that

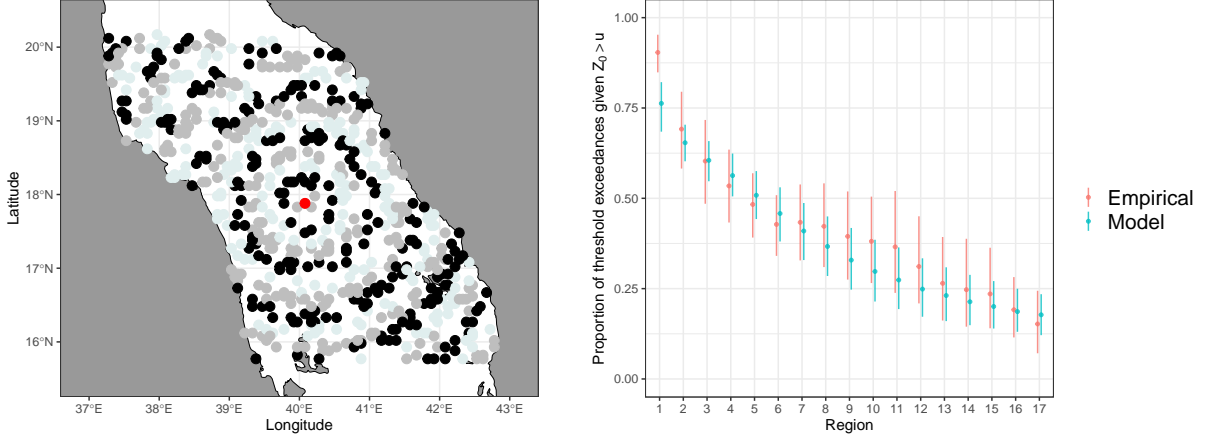


Figure S9: (Left) The spatial domain of interest for the Red Sea study of Section S9, with the conditioning site, \mathbf{s}_0 , shown in red, and the remaining locations separated into 17 regions; the region labels begin at 1 in the centre of the domain and increase with distance from the centre. (Right) The estimated proportion of locations (using $\hat{\theta}_{\text{DNN}}(\cdot)$) on the irregularly-spaced data) for which the process exceeds u given that it is exceeded at \mathbf{s}_0 (points) and corresponding 95% confidence intervals (vertical segments) using model-simulated and empirical bootstrap data sets.

minimises the Kullback-Leibler divergence between the assumed model and the true model (Davison, 2003, pg. 147). Due to their asymptotic relation to the ML estimator, we can therefore expect that a large class of Bayes estimators exhibit similar asymptotic behaviour under model misspecification. During training, however, neural Bayes estimators are only exposed to simulations from the assumed model and, hence, there is a question as to whether they are suitable when the true model from which the observed data are generated is different from the model that they are trained on. Here, we provide empirical evidence that neural Bayes estimators can be used on data that are very different to those used during training.

We trained a neural Bayes estimator under a linear model, and applied it to several non-linear data sets. The assumed model is,

$$\mathbf{Z} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where $\mathbf{Z} \equiv (Z_1, \dots, Z_n)'$ is the data vector, $\boldsymbol{\beta} \equiv (\beta_0, \beta_1)'$ is a vector of regression parameters, \mathbf{X} is a design matrix containing a column of ones and a covariate with values between -1 and 1 , and $\boldsymbol{\epsilon} \sim \text{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Here, we set $n = 100$ and $\sigma = 0.05$. We used a conjugate Gaussian prior, $\boldsymbol{\beta} \sim \text{N}(\mathbf{0}, \mathbf{I})$.

Figure S10 shows the fitted lines using estimates from the neural Bayes estimator (red) and the posterior median of $\boldsymbol{\beta}$ computed analytically (blue) given data simulated from several non-linear models (black points). In each case, the neural Bayes estimator provides a similar fit to that based on the analytic posterior median. Hence, at least in this simple setting, neural Bayes estimators appear to have some extrapolative power beyond the sample space of the assumed model. The good results we obtained in the real experimental study are further evidence that neural Bayes estimators are robust to deviations from the true model.




Figure S10: Fitted lines using the neural Bayes estimator (red) and the posterior median computed analytically (blue) given data simulated from several non-linear models (black points).

S11 Additional figures and tables

Table S3: Summary of the neural network architecture used in Section 3 of the main text. We report the dimension of the input array and output array of each layer. A padding of 0 and a stride of 1 are used in each 2D convolution layer. For all layers except the final layer we use a rectified linear unit (ReLU) activation function. For the final layer, we use an identity activation function. Recall that p denotes the number of parameters in the given statistical model.

Layer type	Input dimension	Output dimension	Kernel size	Parameters
2D convolution	[16, 16, 1]	[7, 7, 64]	10×10	6,464
2D convolution	[7, 7, 64]	[3, 3, 128]	5×5	204,928
2D convolution	[3, 3, 128]	[1, 1, 256]	3×3	295,168
vec(\cdot)	[1, 1, 256]	[256]		0
dense	[256]	[500]		128,500
dense	[500]	[p]		$501p$
Total trainable parameters:				$635,060 + 501p$

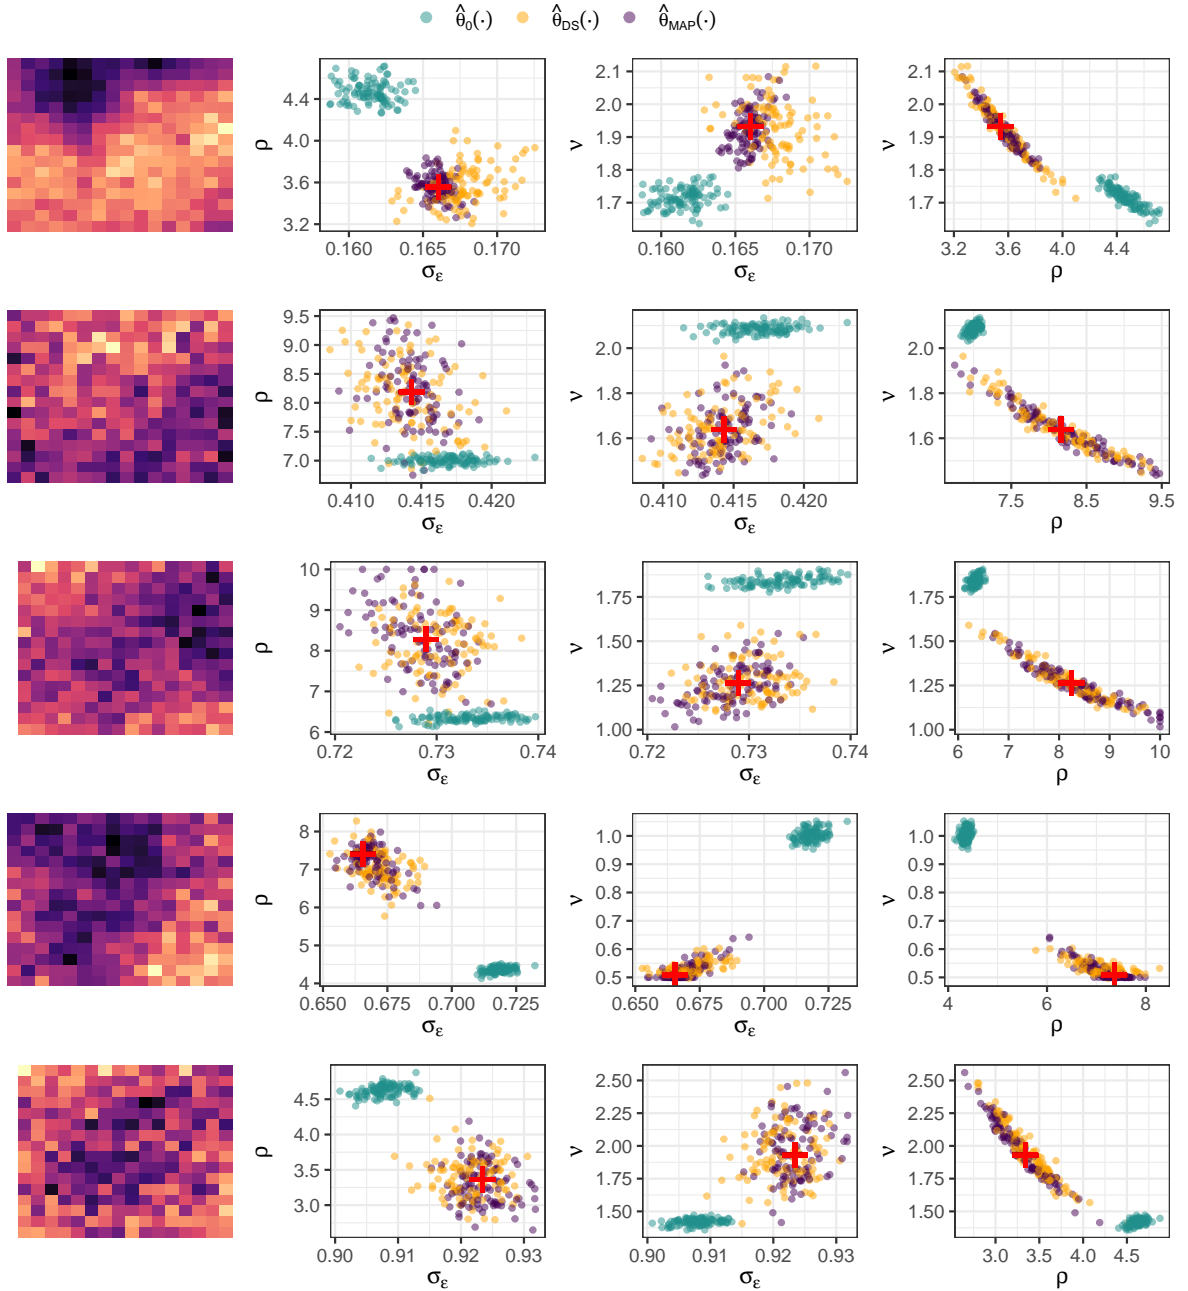


Figure S11: The empirical joint distribution of the estimators considered for the Gaussian process model of Section 3.2 of the main text. Each row corresponds to a randomly chosen parameter vector. The left panel of each row shows a single realisation from the model under the true parameters of the given row. Within each of the remaining panels, the true parameters are shown as red crosses, while estimates from $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{DS}(\cdot)$, and the MAP estimator are shown in green, orange, and purple, respectively. Each estimate was obtained from a simulated data set of size $m = 150$.




Figure S12: The risk with respect to the 0–1 loss against the number of replicates, m , evaluated using the parameter vectors in the test parameter set, for several pairwise MAP (PMAP) estimators considered for Schlather’s max-stable model of Section 3.3 of the main text. We denote the PMAP estimator based on the pairwise likelihood (PL) function that uses only a subset of pairs within a fixed cut-off distance, d , by PL_d . In this case, taking $d = 3$ provides the best results.




Figure S13: The empirical joint distribution of the estimators for Schlather’s max-stable model of Section 3.3 of the main text. Each column corresponds to a randomly chosen parameter vector. Within each panel, the true parameters are shown in red, while estimates from $\hat{\theta}_0(\cdot)$, $\hat{\theta}_{DS}(\cdot)$, and the PMAP estimator are shown in green, orange, and purple, respectively. Each estimate was obtained from a simulated data set of size $m = 150$.




Figure S14: (Top row) Randomly-selected extreme fields from the Red Sea data set considered in Section 4, where $Z(\mathbf{s}_0) > u$. (Bottom row) Simulations from the fitted model. The data are overlaid on a map where the grey areas correspond to land (Sudan and Eritrea to the West, Saudi Arabia to the East) and the white areas correspond to sea.




Figure S15: Data and padded regions for the Red Sea application of Section 4 of the main text. Padding is necessary to use a convolutional neural network, which requires a rectangular array.




Figure S16: Extreme fields from the year 2015 for the Red Sea data set considered in Section 4 of the main text; here, we show the cube-root of the Laplace-scale data to prevent the colour scale being dominated by outliers. We account for the temporal dependence present in these fields using the block bootstrap approach described in Section S8.

References

- Chan, J., Perrone, V., Spence, J., Jenkins, P., Mathieson, S., and Song, Y. (2018). A likelihood-free inference framework for population genetic data using exchangeable neural networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Davison, A. C. (2003). *Statistical Models*. Cambridge University Press, Cambridge, England, first edition.
- Dey, D. K. and Yan, J. (2016). *Extreme Value Modeling and Risk Analysis: Methods and Applications*. Chapman & Hall/CRC Press, Boca Raton, FL.
- Gerber, F. and Nychka, D. W. (2021). Fast covariance parameter estimation of spatial Gaussian process models using neural networks. *Stat*, 10:e382.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge, MA.
- Huser, R. (2013). *Statistical Modeling and Inference for Spatio-Temporal Extremes*. PhD thesis, Swiss Federal Institute of Technology, Lausanne, Switzerland.
- Huser, R. and Davison, A. C. (2014). Space-time modelling of extreme events. *Journal of the Royal Statistical Society B*, 76:439–461.
- Lehmann, E. L. and Casella, G. (1998). *Theory of Point Estimation*. Springer, New York, NY, 2nd edition.
- Rudi, J., Julie, B., and Lenzi, A. (2021). Parameter estimation with dense and convolutional neural networks applied to the FitzHugh-Nagumo ODE. In Bruna, J., Hesthaven, J., and Zdeborova, L., editors, *Proceedings of the 2nd Annual Conference on Mathematical and Scientific Machine Learning*, volume 145 of *Proceedings of Machine Learning Research*, pages 1–28. PMLR.
- Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes*, 5:33–44.
- Subbotin, M. T. (1923). On the law of frequency of errors. *Mathematicheskii Sbornik*, 31:296–301.
- Wadsworth, J. L. and Tawn, J. A. (2022). Higher-dimensional spatial extremes via single-site conditioning. *Spatial Statistics*, 51:100677.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32:4–24.