



HAL
open science

Bayesian space-time gap filling for inference on extreme hot-spots: an application to Red Sea surface temperatures

Daniela Castro-Camilo, Linda Mhalla, Thomas Opitz

► **To cite this version:**

Daniela Castro-Camilo, Linda Mhalla, Thomas Opitz. Bayesian space-time gap filling for inference on extreme hot-spots: an application to Red Sea surface temperatures. *Extremes*, 2021, 24 (1), pp.105-128. 10.1007/s10687-020-00394-z . hal-03165040

HAL Id: hal-03165040

<https://hal.inrae.fr/hal-03165040v1>

Submitted on 31 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bayesian space-time gap filling for inference on extreme hot-spots: an application to Red Sea surface temperatures

¹Daniela Castro-Camilo, ²Linda Mhalla, ³Thomas Opitz

April 2, 2020

Abstract

We develop a method for probabilistic prediction of extreme value hot-spots in a spatio-temporal framework, tailored to big datasets containing important gaps. In this setting, direct calculation of summaries from data, such as the minimum over a space-time domain, is not possible. To obtain predictive distributions for such cluster summaries, we propose a two-step approach. We first model marginal distributions with a focus on accurate modeling of the right tail and then, after transforming the data to a standard Gaussian scale, we estimate a Gaussian space-time dependence model defined locally in the time domain for the space-time subregions where we want to predict. In the first step, we detrend the mean and standard deviation of the data and fit a spatially resolved generalized Pareto distribution to apply a correction of the upper tail. To ensure spatial smoothness of the estimated trends, we either pool data using nearest-neighbor techniques, or apply generalized additive regression modeling. To cope with high space-time resolution of data, the local Gaussian models use a Markov representation of the Matérn correlation function based on the stochastic partial differential equations (SPDE) approach. In the second step, they are fitted in a Bayesian framework through the integrated nested Laplace approximation implemented in R-INLA. Finally, posterior samples are generated to provide statistical inferences through Monte-Carlo estimation. Motivated by the 2019 Extreme Value Analysis data challenge, we illustrate our approach to predict the distribution of local space-time minima in anomalies of Red Sea surface temperatures, using a gridded dataset (11315 days, 16703 pixels) with artificially generated gaps. In particular, we show the improved performance of our two-step approach over a purely Gaussian model without tail transformations.

Keywords: extreme-value theory, generalised additive modeling, INLA, Red Sea, sea surface temperature, SPDE.

¹School of Mathematics and Statistics, University of Glasgow, UK.

²Department of Decision Sciences, HEC Montréal, Canada

³Biostatistics and Spatial Processes, INRAE, Avignon, France

1 Introduction

Large georeferenced datasets, often based on remote sensing or reanalysis of smaller-sized observation data, have become abundant in the domains of climate, environment, and ecology. However, the spatial and temporal resolution in such datasets may not be fine enough for certain purposes. Moreover, gaps may arise in such datasets when sensors are defective or cannot provide useful data, for instance, due to cloud occlusion or orbital characteristics when using satellite-based instruments. In this setting of gap filling and prediction of partially observed data in the Earth and atmospheric sciences (*e.g.*, Yuan *et al.*, 2011; Henn *et al.*, 2013; Mariethoz *et al.*, 2012; Wang *et al.*, 2012; Henn *et al.*, 2013; Lakshmi, 2016; Xing *et al.*, 2017; Yin *et al.*, 2017; Padhee and Dutta, 2019), we aim to provide probabilistic predictions of summary statistics of the data process for a space-time domain, and in particular to focus on episodes with simultaneous extreme values of a continuous variable such as temperature. By taking into account the spatial and temporal dependence, we seek inference for summary functionals of space-time hot-spots, such as the minimum value of the process over a spatial region during a certain period of time.

Our work was motivated by the 2019 Extreme Value Analysis (EVA) data challenge to which we participated with our team named *BeatTheHeat*. Its goal was to predict the minimum of anomalies of sea surface temperature (SST) of the Red Sea over regions where data have been artificially masked, and predictions for 162,000 space-time cylinders with a spatial radius of 50km and duration of 7 days had to be provided. The Red Sea is home to over 300 species of coral and 700 species of fish, some of them endemic (Spalding *et al.*, 2001; www.fishbase.in). This rich ecosystem is being threatened by global warming. The 2014 International Panel on Climate Change report indicates that the average SST of the Indian Ocean has increased by 0.65°C over the period 1950-2009 (Hoegh-Guldberg *et al.*, 2014). Summertime SSTs for the past decade have remained, on average, 1.46°C above the historical mean (Cantin *et al.*, 2010). Coral reefs are particularly sensitive to modest increases in background seawater temperature; a continued increase of 1°C or more above the historical maximum SST can result in substantial coral damage (Cantin *et al.*, 2010).

There is a need to understand the spatio-temporal dynamics governing SST, and in particular their extremes, in the Red Sea. In this work, we consider temperature anomalies that were constructed by subtracting a spatio-temporal mean from gridded daily SST observations provided by the Operational SST and Sea Ice Analysis (OSTIA; Donlon *et al.* 2012), a satellite-

based data system designed for numerical weather prediction. The data encompasses 31 years (1985–2015) of daily simulations over 16,703 locations in the Red Sea.

Geostatistics provides a wide range of models to describe and predict spatio-temporal phenomena. Historical and current developments are usually centered around Gaussian processes, which have appealing theoretical and computational properties (Cressie and Wikle, 2015). Extreme-value analysis provides more specific tools to model extremes in space and time, although they usually pose a greater computational challenge. In this work, we combine both disciplines to predict high-temperature anomalies over space and time in the Red Sea.

Due to the large size of the dataset, care is needed to develop inference procedures that scale well with large datasets. We propose to model the marginal effects and dependence structure separately following a two-step approach. First, we fit a model for marginal distributions and use it to transform data to a common standard Gaussian scale. Then, stationary Gaussian process models are used locally in time to capture spatio-temporal dependence.

To describe non-stationary marginal anomalies at each site in the first step, we start by removing long-term temporal trends at a yearly scale, assumed to be stationary over space, by correcting data using empirical means and variances. Subsequently, we assume temporal stationarity, and we seek to improve the upper tail representation through a model for data above a stationary threshold close to the global 75%-quantile. We estimate a tail model with space-varying parameters based on a Poisson regression for the rate of threshold exceedances, and on a generalized Pareto (GP) model for the excesses above the threshold. For these marginal models, we compare three approaches that allow generating relatively smooth parameter surfaces. First, we consider the data without marginal transformation as a benchmark to assess the added value of correcting the tail. Second, we rely on the generalized additive modeling (GAM) framework (Wood, 2003, 2006) where we use the latitude, longitude, and distance to the coastline as covariates to describe the spatial pattern in the tails. This framework allows for flexible modeling of covariate effects and has previously been used in univariate and multivariate extreme-value settings to capture non-stationarity in environmental data (Pauli and Coles, 2001; Chavez-Demoulin and Davison, 2005; Jonathan *et al.*, 2014; Mhalla *et al.*, 2019). Third, we implement spatial pooling of data using a nearest-neighbor (NN) approach to infer the tail behavior at each location.

SSTs in the Red Sea are known to change smoothly in space and time (Chaidez *et al.*, 2017). Therefore, in the second step of our approach, we describe the space-time dependence

between Gaussian-transformed SST anomalies by a spatio-temporal stationary Gaussian term characterized by a stationary Matérn covariance structure in space and first-order temporal autoregression. For fast likelihood-based inference and probabilistic prediction with large gridded datasets, we use the stochastic partial differential equations (SPDE) approach. It provides a numerically convenient Gauss–Markov approximation to the Matérn covariance (Lindgren *et al.*, 2011), and we combine it with the computationally-efficient integrated nested Laplace approximation (INLA; Rue *et al.*, 2009, 2017; Opitz, 2017; Opitz *et al.*, 2018; Castro-Camilo *et al.*, 2019) for Bayesian inference.

A numerical comparison of the performances of our three models with different marginal pretransformations (none, GAM, NN) is made using the tail-weighted continuous ranked probability score (twCRPS) (Gneiting and Ranjan, 2011). Our models perform better than all other contributions to the EVA challenge, and the tail correction provides a significant improvement over models without marginal transformation.

The remainder of the paper is organized as follows. Section 2 presents the data through exploratory analyses and graphical illustrations. Section 3 describes our modeling approaches, for which results and prediction performances are discussed in Section 4. Conclusions and an outlook towards further improvements in our model are given in Section 5.

2 Red Sea surface temperature data

We describe the main features of the SST anomalies and provide exploratory graphical support for our modeling choices.

2.1 Description of data and space-time hotspots

The Sea Ice Analysis (OSTIA) system produces daily SST data in the Red Sea at a spatial resolution of $1/20^\circ$. We consider the period between January 1985 and December 2015. We denote by $Y(\mathbf{s}, t)$ the Red Sea surface temperature at location $\mathbf{s} \in \mathcal{S} \subset \mathbb{R}^2$ and time $t \in \mathcal{T} = \{1, \dots, T\}$, where $|\mathcal{S}| = 16703$ and $T = 11315$. Our analysis focuses on anomalies $A(\mathbf{s}, t)$ where

$$A(\mathbf{s}, t) = Y(\mathbf{s}, t) - m(\mathbf{s}, t),$$

and m is a space-time mean surface; see Huser (2020) for details on its estimation based on averages for each pixel-month combination. Anomalies remove recurrent local climatic trends and allow for easier interpretation of short-term variability and large-scale anomalous behavior

with respect to what we observe on average at a given site and time of the year. Large-scale extreme temperatures are likely to negatively affect the environmental and ecological systems in the Red Sea, and we are interested in measuring the impact of extreme anomalies occurring jointly in a specific space-time region.

While the original dataset does not contain any gaps with missing data, such gaps have been inserted artificially in the dataset to compare predictive approaches; see [Huser \(2020\)](#). During each month, the data for a randomly selected subdomain of the Red Sea were removed. Using the remaining observations, our goal is to predict a summary functional of space-time subregions contained in these gaps, *i.e.*, where no observations are available. Specifically, for each pair $(\mathbf{s}, t) \in \mathcal{X} = \mathbb{R}^2 \times \mathcal{T}$, we consider the following cylindrical neighborhood

$$\mathcal{N}(\mathbf{s}, t) = \mathcal{B}(\mathbf{s}, r) \times \{t - 3, t - 2, t - 1, t, t + 1, t + 2, t + 3\} \cap \mathcal{X},$$

where $\mathcal{B}(\mathbf{s}, r)$ is a ball centered at \mathbf{s} with radius $r = 50\text{km}$. We are interested in predicting the distribution of summary functionals describing *extreme space-time hot-spots*, by using the minimum summary:

$$X(\mathbf{s}, t) = \min_{(\mathbf{s}', t') \in \mathcal{N}(\mathbf{s}, t)} A(\mathbf{s}', t').$$

For the EVA competition, the validation dataset of observations to be predicted, $X(\mathbf{s}_i, t_i)$, $i = 1, \dots, 162,000$, consists of 324 distinct time values t_i , and 500 points \mathbf{s}_i for each of these time points. Our aim is to successfully predict these summaries with a focus on extreme quantiles of $X(\mathbf{s}, t)$, which characterize extreme hot-spots of SST anomalies.

The Red Sea spans a relatively large region on the globe, such that it is not straightforward to choose a sound 2D coordinate system for models that take into account spatial distances. While 1 degree of latitude corresponds to 111.2km, we find that 1 degree of longitude corresponds to 96.3km at the northern limit and to 108.8km at the southern limit of the Red Sea. We here switch to a metric system with km unit to avoid artificial anisotropy effects as best as possible, by multiplying latitude degree values with 111.2 and longitude degree values with the average between 96.3 and 108.8.

2.2 Exploratory analysis

Despite the preliminary removal of mean trends according to location and month, strong nonstationarity remains in the data $A(s, t)$, particularly in the upper tail. To illustrate this, [Figure 1](#) shows maps of annual average anomalies in the Red Sea for three different years using

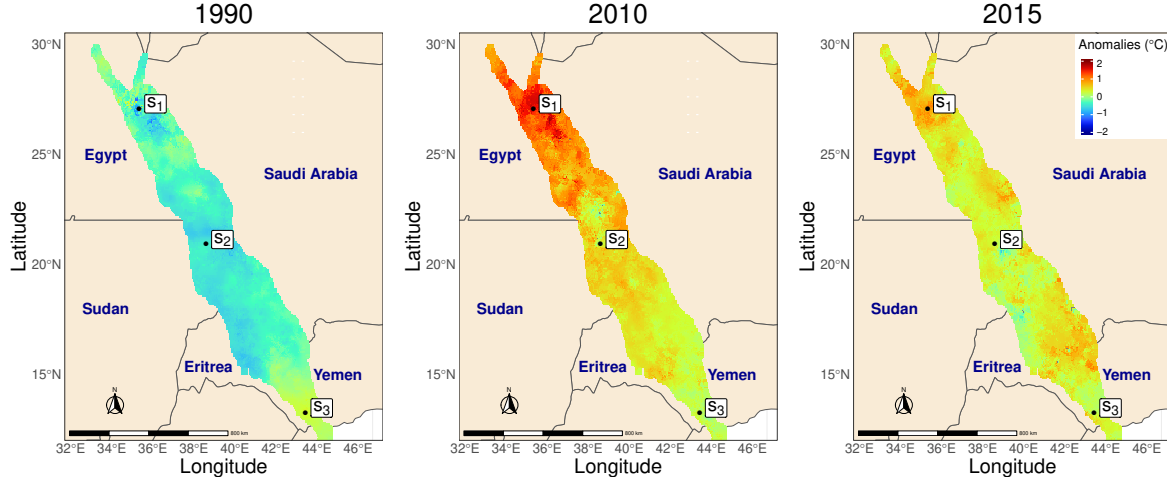


Figure 1: Annual average anomalies for three different years. Data at the three locations indicated on the maps are used in Figure 7.

the same colour scale, highlighting the fact that average anomalies vary strongly between years. We can also see a mild spatial effect in the 2010 map, but overall, the temporal effect seems to drive the nonstationarity observed in the average anomalies.

To further illustrate the behavior of SST anomalies through years, we compute the annual average anomalies along with their standard deviations for every location in each year. These location and scale statistics can be treated as functional data over the years (one function per location), and displayed using the functional boxplots (Sun and Genton, 2011) in Figure 2. To highlight the effect of the years at each location, we also plot the overall statistics by year (*i.e.*, average and standard deviation over all locations for every fixed year). We can see that there is an upward annual trend over the averages, and a mild downward annual trend affecting the location-wise standard deviations of the SST anomalies.

To unveil spatial nonstationarity in the tail of the SST anomalies, we characterize the distribution of threshold exceedances defined as those SST anomalies exceeding 0.49, corresponding to an overall exceedance probability of approximately 20%, and used later for tail modeling. Figure 3 shows the location-wise empirical range of exceedances (*i.e.*, the difference between the maximum and minimum values) as well as the mean exceedance value, $\mathbb{E}\{A(\mathbf{s}, t) \mid A(\mathbf{s}, t) > u\}$. We can see a clear North to South pattern in both maps, consistent with a strong spatial effect in the upper tail of the SST anomalies.

We conclude this section by studying bivariate summaries of extremal dependence. Specifically, we analyze two well-know pairwise measures of tail dependence based on conditional




Figure 2: Annual average anomalies (left) and standard deviations (right). Functional boxplots summarise the behavior of the statistics across all locations: the blue polygons contain the 50% central curves, while the outer coral dashed lines correspond to the minimum and maximum curves. The solid purple lines correspond to the overall statistics by year. The gray dashed lines correspond to the fitted mean and standard deviation as specified by the regression model (1).

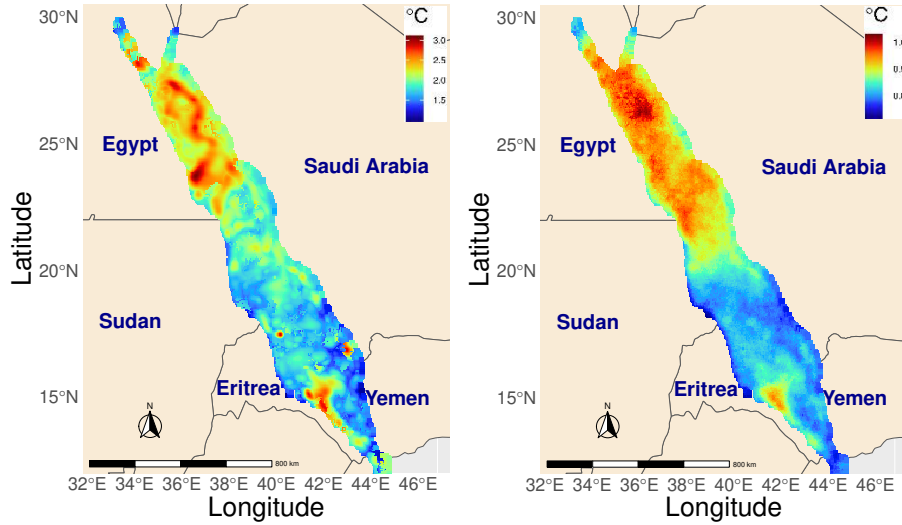


Figure 3: Location-wise empirical range (left) and mean exceedance value (right).

exceedance probabilities, namely

$$\chi_{ij}(u) = \mathbb{P}[F_i\{A(\mathbf{s}_i, \cdot)\} > u \mid F_j\{A(\mathbf{s}_j, \cdot)\} > u],$$

$$\bar{\chi}_{ij}(u) = \frac{2 \log(\mathbb{P}[F_i\{A(\mathbf{s}_i, \cdot)\} > u])}{\log(\mathbb{P}[F_i\{A(\mathbf{s}_i, \cdot)\} > u, F_j\{A(\mathbf{s}_j, \cdot)\} > u])} - 1,$$

where F_i is the distribution function at location i and $A(\mathbf{s}_i, \cdot)$ is the time series at location i . Together, $\chi = \lim_{u \rightarrow 1} \chi_{ij}(u)$ and $\bar{\chi} = \lim_{u \rightarrow 1} \bar{\chi}_{ij}(u)$ summarise the strength of the spatial tail dependence where $\chi > 0$ characterizes *asymptotic dependence*, while $\bar{\chi} < 1$ may arise if in the case of *asymptotic independence* ($\chi = 0$). Figure 4 displays bootstrap estimates and confidence




Figure 4: Empirical estimates of $\chi_{ij}(u)$ (left) and $\bar{\chi}_{ij}(u)$ (right) for the SST anomalies data. The filled regions represent the 95% bootstrap confidence bands.

bands for $\chi_{ij}(u)$ and $\bar{\chi}_{ij}(u)$, where we fix u at three high probabilities $u = 0.99, 0.997, 0.999$ and consider pairs of locations (i, j) separated by different spatial lags. For computational reasons and due to the strong spatial dependence in data, the estimation is conducted on a subset of the available locations, obtained by regular sampling of 1558 locations. The bootstrap procedure relies on 300 bootstrap samples drawn by resampling from the set of pairs of locations. The resulting estimates of $\chi_{ij}(u)$ point towards asymptotic independence as they strongly decrease with larger spatial lags and higher probabilities. The estimates of $\bar{\chi}_{ij}(u)$ remain clearly below 1 for all positive distances, which confirms asymptotic independence in data. They show significantly lower values for higher probabilities, which indicates particularly fast joint tail decay rates. Moreover, this analysis has been done without removing temporal trends in margins, such that even stronger evidence of relatively weak asymptotic independence can be expected after removing such trends.

3 Space-time predictive modeling

In this section, we describe our two-step approach to model trends in marginal distributions and local space-time dependence.

3.1 Marginal modeling

We consider three approaches to pretransforming data before fitting local Gaussian models.

3.1.1 Model 1: no marginal transformations

As a reference model allowing us to assess the utility of marginal pretransformations, we implement an approach where the local Gaussian dependence model is fitted to raw anomalies $A(s, t)$ without preliminary marginal transformations.

3.1.2 Model 2: generalized additive regression

As discussed in Section 2.2, the SST anomalies exhibit a residual temporal trend in the mean and a spatial non-stationary pattern in the tails. We first remove the long-term time trends observed in Figure 2 by fitting a Gaussian location-scale regression model. Recall that we denote by $A(\mathbf{s}, t)$ the SST anomaly at location $\mathbf{s} \in \mathcal{S}$ and time $t \in \mathcal{T}$. Then, we assume the following model structure $A(\mathbf{s}, t) = \mu(t) + e(t)$, where $e(t) \sim \mathcal{N}(0, \sigma^2(t))$ and

$$\begin{aligned} \mu(t) &= \beta_{\mu,i} \mathbb{1}_{\text{yod}_t=i}, \\ \log\{\sigma(t)\} &= \beta_{\sigma,i} \mathbb{1}_{\text{yod}_t=i}, \quad i = 1, \dots, 31 \end{aligned} \tag{1}$$

where yod_t denotes the year of day t and $\beta_{\mu,i}$ and $\beta_{\sigma,i}$ are the average mean and standard deviation of the year $1984 + i$. Owing to the large number of observations in the dataset as well as the spatial pooling of the data, this simple model specification should be able to capture the observed temporal non-stationarity while maintaining the inherent spatio-temporal structure of the large anomalies that we ought to predict.

To reduce the computational cost of fitting the model in (1), we use a sub-sampling strategy by randomly selecting one site out of 50 and keeping all the observations over time. Our approach is sensible since original SST anomalies were obtained using a spatial multi-scale optimal interpolation scheme (Donlon *et al.*, 2012).

As the interest of this work is in the prediction of jointly large anomalies, we propose to improve the tail modeling of the standardized anomalies. More precisely, the SST anomalies $A(\mathbf{s}, t)$ are normalized, in a first step, to a standard Gaussian scale using the fitted model (1), *i.e.*,

$$\tilde{Z}(\mathbf{s}, t) = \frac{A(\mathbf{s}, t) - \hat{\mu}(t)}{\hat{\sigma}(t)}, \quad \mathbf{s} \in \mathcal{S}, t \in \mathcal{T}, \tag{2}$$

where $\tilde{Z}(\mathbf{s}, t)$ are now supposed (approximately) stationary over time. Then, in a second step, we refine the model in the tails by considering a spatial Generalized Pareto (GP) model for the exceedances of a high threshold u . Specifically, we assume a generalized additive model

(GAM) where the scale and shape GP parameters are described by

$$\begin{aligned}\log\{\sigma_{GP}^{\text{GAM}}(\mathbf{s})\} &= f_1(\text{lat}_{\mathbf{s}}, \text{lon}_{\mathbf{s}}) + f_2(\text{dist}_{\mathbf{s}}), \\ \xi^{\text{GAM}}(\mathbf{s}) &= f_3(\text{lat}_{\mathbf{s}}, \text{lon}_{\mathbf{s}}) + f_4(\text{dist}_{\mathbf{s}}),\end{aligned}\tag{3}$$

where $\text{lat}_{\mathbf{s}}, \text{lon}_{\mathbf{s}}$ are the latitude and longitude of the site \mathbf{s} , while $\text{dist}_{\mathbf{s}}$ is its distance to the coast (in meters). Within the GAM framework, the unknown functions f_2 and f_4 are represented using reduced rank smoothing splines, and more specifically, the reduced rank isotropic thin plate splines (Wood, 2017, Chapter 4). The smooth interaction terms f_1 and f_3 are constructed using the tensor product construction where the main effects of lat and lon are excluded (Wood, 2006). Using their respective fixed basis expansions, estimation of the unknown smooth functions boils down to inference on the coefficients of their model matrices. In other words, each smooth term $\mathbf{f}_j = (f_j(x_{j1}), \dots, f_j(x_{jn}))^\top$ can be written as $\mathbf{f}_j = \mathbf{X}_j \boldsymbol{\beta}_j$ where \mathbf{X}_j is the model matrix constructed by evaluating the (known) basis functions at the observed values of the covariates, and $\boldsymbol{\beta}_j$ the vector of its corresponding coefficients to be estimated. To avoid overfitting issues, each smooth function f_j is associated with a smoothing penalty term reflecting its curvature and defined as $\boldsymbol{\beta}^\top \mathbf{S}_j \boldsymbol{\beta}$, where \mathbf{S}_j is a known positive definite matrix that depends on the basis functions. In the special setting of tensor products, the wiggleness of the smooth surface is controlled by two smoothing parameters associated with each marginal direction. Then, given a set of smoothing parameters $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_6)$, the estimated model coefficients are obtained by maximum penalized likelihood

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\text{argmax}} \left\{ \ell(\boldsymbol{\beta}) - \sum_{j=1}^6 \lambda_j \boldsymbol{\beta}^\top \mathbf{S}_j \boldsymbol{\beta} / 2 \right\},\tag{4}$$

where $\ell(\boldsymbol{\beta}) = \log f(\tilde{\mathbf{Z}} \mid f_1, f_2) = \log f(\tilde{\mathbf{Z}} \mid \boldsymbol{\beta})$ is the Generalized Pareto (GP) log-likelihood

$$f\{\tilde{Z}(\mathbf{s}, t) \mid \boldsymbol{\beta}\} = \begin{cases} \frac{1}{\sigma_{GP}^{\text{GAM}}(\mathbf{s})} \left\{ 1 + \xi^{\text{GAM}}(\mathbf{s}) \frac{\tilde{Z}(\mathbf{s}, t) - u}{\sigma_{GP}^{\text{GAM}}(\mathbf{s})} \right\}^{-1/\xi^{\text{GAM}}(\mathbf{s})-1} & \text{if } \xi^{\text{GAM}}(\mathbf{s}) \neq 0, \\ \frac{1}{\sigma_{GP}^{\text{GAM}}(\mathbf{s})} \exp \left\{ -\frac{\tilde{Z}(\mathbf{s}, t) - u}{\sigma_{GP}^{\text{GAM}}(\mathbf{s})} \right\} & \text{if } \xi^{\text{GAM}}(\mathbf{s}) = 0, \end{cases}$$

and $\tilde{\mathbf{Z}} = \{\tilde{Z}(\mathbf{s}, t), \mathbf{s} \in \mathcal{S}, t \in \mathcal{T}\}$. The smoothing parameters are estimated iteratively using the restricted marginal likelihood approach (Wood, 2011; Wood *et al.*, 2016). Precisely, for an intermediate estimate $\hat{\boldsymbol{\beta}}$ obtained through a Newton step based on (4) and for a fixed $\boldsymbol{\lambda}$, Wood *et al.* (2016) propose to undertake a Bayesian view where $\hat{\boldsymbol{\beta}}$ is a posterior mode for $\boldsymbol{\beta}$ and the

estimate of the smoothing parameters $\boldsymbol{\lambda}$ is updated by maximizing a Laplace approximation of the log marginal likelihood of the GP model. Uncertainty measures for the model coefficient estimates include thus the uncertainty associated with the smoothing parameter estimation; see [Wood *et al.* \(2016, Section 4\)](#) for details. This fitting procedure is implemented in the package `mgcv` ([Wood, 2017](#)) available in the R statistical software, and its application in an extreme-value setting is illustrated in [Youngman \(2019\)](#).

Using an *ad-hoc* trial and error procedure, we choose a fixed threshold at the value $u = 0.75$ for all sites and all time points. This value corresponds roughly to the 78.3% empirical quantile of the transformed data in (2). To cope with the computational demands of fitting the GAM model in (3) to all threshold exceedances, we keep all the sites but randomly select one excess out of ten.

Finally, to compensate for the choice of a stationary threshold, we characterize the spatial behavior of the frequency of extreme SST anomalies by modeling the exceedance probability $p^{\text{GAM}}(\mathbf{s}) = \Pr\{Z(\mathbf{s}, t) > 0.75\}$. More specifically, we model the number of exceedances using a Poisson regression model with a spatially varying intensity following the same specification as in (3).

3.1.3 Model 3: Spatial nearest-neighbour modeling

To compare the modeling of the tail behavior with smooth covariate effects described in Section 3.1.2, we propose an alternative local modeling approach. It exploits the large number of observations and captures the spatial pattern in the large anomalies of the normalized observations $\tilde{Z}(\mathbf{s}, t)$ by modeling tail non-stationarity separately for each pixel. Therefore, we pool together the values $\tilde{Z}(\mathbf{s}, t)$ of the 40 nearest neighbor pixels and of the pixel itself. Based on this sample, we then calculate the pixel-specific likelihood estimates of the exceedance probability $p^{\text{NN}}(\mathbf{s})$ and of the generalized Pareto parameters $\sigma_{\text{GP}}^{\text{NN}}(\mathbf{s})$ and $\xi^{\text{NN}}(\mathbf{s})$.

3.2 Local dependence modeling

For predicting nonlinear cluster functionals such as minima or maxima, it is important to provide an appropriate model for all multivariate distributions of the space-time process. In particular, relatively fast and robust geostatistical kriging approaches using only the mean and the variogram ([Cressie and Wikle, 2015](#)) are not suitable. Moreover, the estimation of a dependence model for the full dataset may be impeded by the large size of the dataset,

as in the case of the Red Sea data with more than 150 million observations. Instead, we propose to use a local dependence model for subperiods, including the space-time sets over which we aim to predict cluster functionals. We propose using a stationary Gaussian space-time dependence model for such subperiods, applied either directly to the anomalies $A(s, t)$, or after using a marginal model to transform margins to the standard Gaussian distribution. We then combine a latent variable approach with Gauss-Markov dependence structures in a Bayesian framework to cope with the high space-time resolution and to provide probabilistic predictions for cluster functionals.

As shown by Figure 4, there is strong evidence in favor of asymptotic independence for the Red Sea SST data. Therefore, classical limit models of dependence from extreme-value theory (max-stable processes or generalized Pareto processes, see Davison *et al.*, 2012; Ferreira and De Haan, 2014; Thibaud and Opitz, 2015) would be inappropriate, in addition to being not easily tractable for prediction with large datasets. Gaussian dependence has already proved to be flexible for tails in some cases (Bortot *et al.*, 2000).

A useful choice for modeling spatio-temporal dependence is to link spatial fields through temporal auto-correlation, and in particular first-order autoregression, which leads to sparse spatio-temporal precision matrices (*i.e.*, inverse covariance matrices) in case of sparse spatial precision matrices. We consider the following stationary space-time process defined for discrete time with regular time steps:

$$\begin{aligned} W(\mathbf{s}, 1) &= w_0 + \varepsilon_1(\mathbf{s}), \\ W(\mathbf{s}, t + 1) &= (1 - \rho)w_0 + \rho W(\mathbf{s}, t) + \sqrt{1 - \rho^2} \varepsilon_{t+1}(\mathbf{s}), \quad t = 1, 2, \dots, t_{\max}, \end{aligned} \quad (5)$$

where ε_t , $t = 1, 2, \dots, t_{\max}$ are centered Gaussian random fields with Matérn covariance and $w_0 \in \mathbb{R}$ is the mean parameter. The space-time precision matrix for the Cartesian product of a collection of sites and times corresponds to the Kronecker product of the corresponding purely spatial (Matérn) and purely temporal (AR(1)) precision matrices. In our local models, we use $t_{\max} = 9$, where we have added one day before and after the periods of seven days length, over which we seek to predict cluster functionals such as minima. Using larger temporal buffers around the prediction period would be possible and could further improve the prediction performance in cases where the conditional independence of $W(\mathbf{s}, t + 1)$ and $W(\mathbf{s}, t - 1)$ given $W(\mathbf{s}, t)$, as suggested by first-order autoregression, does not hold. For reasons of computational cost, we do not explore this possibility in our data application.

Using the stochastic partial differential equation (SPDE) approach developed by Lindgren

et al. (2011), we can work with numerically convenient sparse precision matrices that approximate the Matérn covariance function. We shortly recall this framework and refer to Krainski *et al.* (2018) for details. The SPDE is given by

$$(\kappa^2 - \Delta)^{\alpha/2} \tau W(\mathbf{s}) = E(\mathbf{s}), \quad \alpha = \nu + D/2, \quad \mathbf{s} \in \mathbb{R}^d, \quad (6)$$

with the Laplace operator $\Delta y = \sum_{j=1}^D \partial^2 y / \partial^2 x_j$, and a standard Gaussian white noise process $E(\mathbf{s})$. Its unique stationary solution is a zero-mean Gaussian process $W(\mathbf{s})$ with Matérn covariance function with shape $\nu > 0$, fixed to $\nu = 1$ in the following. The marginal variance is $\sigma_W^2 = \Gamma(\nu) / \{\Gamma(\nu + D/2)(4\pi)^{D/2} \kappa^{2\nu} \tau^2\}$, and the *empirical range* parameter is $\sqrt{8\nu} / \kappa^2$, indicating the distance where we observe a correlation of approximately 0.1. With a finite study region, appropriate boundary conditions ensure a unique solution, and we can set the boundary relatively far from the actual study region to render boundary effects negligible within the study area.

The approximate Gauss–Markov solution to (6) is given by constructing a discretization of space through a triangulation mesh, and by using the representation $W(\mathbf{s}) = \sum_{j=1}^{m_W} \omega_j \Psi_j(\mathbf{s})$ with “pyramid-shaped” basis functions $\Psi_j(\mathbf{s})$, one for each mesh node. By solving the SPDE in the subspace spanned by the linear combination $W(\mathbf{s})$, we obtain $\boldsymbol{\omega} = (\omega_1, \dots, \omega_{m_W})^T \sim \mathcal{N}(0, Q_W)$ with explicitly known precision matrix Q_W . The mesh used for the Red Sea in our data application is shown in Figure 5. It contains $m_W = 2746$ nodes, and we choose a fine resolution in the study area (within the inner blue boundary), while a coarser mesh is used in the extension area, which is sufficient to avoid boundary effects but without strongly increasing the number of nodes for the sake of keeping the computational cost as low as possible. With $t_{\max} = 9$ and the intercept w_0 , each local model counts $9 \times 2746 + 1 = 24,715$ latent Gaussian variables.

3.3 Bayesian inference using the integrated nested Laplace approximation

We represent our local dependence model, detailed in Section 3.2, as a latent Gaussian model with Gaussian likelihood as follows. We denote by $\mathbf{Z} = \{Z(\mathbf{s}, t)\}_{t \in \mathcal{T}, \mathbf{s} \in \mathcal{S}}$ the standard Gaussian observations obtained after transforming the raw anomalies using one of the marginal approaches described in Section 3. Precisely, using on the probability integral transform and




Figure 5: Triangulation mesh (km unit) used to discretize the study region for latent variable modeling with Gaussian variables located at the mesh nodes. The inner blue boundary delimits the study area, while the extension to the outer blue boundary ensures approximately stationary behavior of the SPDE solution within the study area.

the fitted tail parameters, we proceed with the following transformation

$$Z(\mathbf{s}, t) = \begin{cases} \Phi^{-1} \left[\tilde{F}_{emp} \left\{ \tilde{Z}(\mathbf{s}, t); \mathbf{s} \right\} \right] & \text{if } \tilde{Z}(\mathbf{s}, t) \leq 0.75, \\ \Phi^{-1} \left[T \left\{ \tilde{Z}(\mathbf{s}, t); \sigma_{GP}(\mathbf{s}), \xi(\mathbf{s}), p(\mathbf{s}) \right\} \right] & \text{if } \tilde{Z}(\mathbf{s}, t) > 0.75, \end{cases}$$

where Φ is the standard Gaussian distribution function and

$$T(z; \sigma_{GP}, \xi, p) = \begin{cases} 1 - p \left(1 + \xi \frac{z - 0.75}{\sigma_{GP}} \right)^{-1/\xi} & \text{if } \xi \neq 0, \\ 1 - p \exp \left(-\frac{z - 0.75}{\sigma_{GP}} \right) & \text{if } \xi = 0, \end{cases}$$

and $\tilde{F}_{emp}(\cdot; \mathbf{s})$ the location-wise empirical distribution, to achieve an improved normalization of the very large anomalies in the dataset. Then, for each $(\mathbf{s}, t) \in \mathcal{S} \times \mathcal{T}$, we define the space-time linear predictor

$$\eta(\mathbf{s}, t) = \eta_0 + W(\mathbf{s}, t),$$

according to (5) with intercept η_0 . We assume a centered Gaussian prior for the intercept, such that the joint distribution of $\boldsymbol{\eta} = \{\eta(\mathbf{s}, t)\}_{t \in \mathcal{T}, \mathbf{s} \in \mathcal{S}}$ is also Gaussian. By assuming that observations $Z(\mathbf{s}, t)$ are conditionally independent given the latent Gaussian random vector $\boldsymbol{\eta}$

and a vector of hyperparameters $\boldsymbol{\theta}$, we use the following hierarchical structure:

$$\begin{aligned} \mathbf{Z} \mid \mathbf{x}, \tau_{\mathbf{z}} &\sim \prod_{t \in \mathcal{T}, \mathbf{s} \in \mathcal{S}} \varphi \{ Z(\mathbf{s}, t) \mid \eta(\mathbf{s}, t), \tau_{\mathbf{z}}^{-1} \}, \\ \boldsymbol{\eta} \mid \boldsymbol{\theta} &\sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{\theta}}, \mathbf{Q}_{\boldsymbol{\theta}}^{-1}), \\ (\tau_{\mathbf{z}}, \boldsymbol{\theta}) &\sim p(\tau_{\mathbf{z}}, \boldsymbol{\theta}), \end{aligned} \tag{7}$$

where $\varphi(\cdot \mid \mu, \sigma^2)$ is the Gaussian density with mean μ and variance σ^2 , $\boldsymbol{\mu}_{\boldsymbol{\theta}}$ is the mean vector and $\mathbf{Q}_{\boldsymbol{\theta}}$ is the precision matrix of $\boldsymbol{\eta}$. The representation in (7) corresponds to a latent Gaussian model (LGM; see, *e.g.*, Rue *et al.*, 2017) where observations follow a Gaussian distribution with mean process $\eta(\mathbf{s}, t)$ and a measurement error of variance $1/\tau_{\mathbf{z}}^{-1}$. The measurement error may not be necessary from a modeling stance when data are very smooth, but its presence in the model is required for implementing LGM-based inference. For smooth data, it is typically estimated to be very small, such that it becomes negligible in practice. Approximate Bayesian inference for LGMs is available through the integrated nested Laplace approximation approach (INLA; Rue *et al.*, 2009) and its implementation, R-INLA (Bivand *et al.*, 2015) in the R statistical software. With INLA, posterior quantities of interest (that usually involve the evaluation of high-dimensional integrals) are numerically approximated using variants of the classical Laplace approximation (Tierney and Kadane, 1986). We use a recent version of INLA calling the PARDISO library for fast numerical matrix computations (De Coninck *et al.*, 2016; Verbosio *et al.*, 2017; Kourounis *et al.*, 2018; van Niekerk *et al.*, 2019). INLA leverages modern numerical techniques for sparse matrices, providing fast and accurate inference for LGMs with many observations and latent variables such as in our case.

Regarding prior distributions, we choose a Gaussian distribution with precision 0.1 for the mean parameter in (5), which provides a moderately informative prior. The hyperparameters $\tau_{\mathbf{z}}$ and $\boldsymbol{\theta} = (\kappa, \sigma_W^2, \rho)$ control the Gaussian likelihood and the latent Gaussian vector $\boldsymbol{\eta}$, respectively. Although non-informative priors are a common choice when little expert knowledge is available, we here select moderately informative prior distributions using the penalized complexity (PC) prior approach (Simpson *et al.*, 2017). This procedure penalizes excessively complex models at a constant rate by putting an exponential prior on a distance to a simpler baseline model (*e.g.*, using the baseline $1/\tau_{\mathbf{z}} = 0$ for the measurement error variance), which helps to stabilize the estimation algorithm by facilitating the convergence of Laplace approximations. Priors then shrink model components toward their base models, thus preventing overfitting. For the precision $\tau_{\mathbf{z}}$ of observations, we set an informative prior, such that the

probability of observing a standard deviation larger than 0.1 is 50%. For the Matérn covariance, we set a prior distribution where the probability that the range is smaller than 500km and that the variance larger than 0.25 is 50% in each case. Finally, an informative PC prior is chosen for the temporal auto-correlation coefficient of the AR(1) process in (5), which reflects the strong temporal dependence: a priori, the probability of observing $|\rho|$ greater than 0.85 is set to 50%.

3.4 Generating predictive samples

In contrast to simulation-based estimation of posterior distribution through techniques such as Markov-chain Monte-Carlo, INLA does not automatically provide a representative sample of latent variables and predictors to calculate posterior quantities of interest. Nevertheless, it is possible to obtain samples of the approximate posterior distribution efficiently: first, we sample a set of hyperparameters according to the approximation to their posterior distribution conducted by INLA; second, conditional on the hyperparameter values, we sample from the Gaussian distributions arising from the Laplace approximation (Krainski *et al.*, 2018).

For each local Gaussian model, we have implemented this simulation procedure to obtain 500 samples of the latent variables of the fitted model. The simulated predictions of $Z(\mathbf{s}, t)$ for unobserved locations and times are part of these samples, and we first backtransform them to the original marginal scale (if marginal transformations were applied preliminarily). Finally, the value of the cluster summary is calculated for each backtransformed posterior sample. This leads to a sample of size 500 for each cluster summary $X(\mathbf{s}, t)$ to be predicted, and we represent predictive cdfs through the empirical cdfs of these samples.

4 Results

We now discuss fitted models and assess the adequacy of our models to predict local space-time hot-spots in the anomalies by comparing goodness-of-fit diagnostics and predictive performance for the cluster summary given by the minimum. Improvements owing to explicit modeling of marginal tails are highlighted.

4.1 Marginal modeling

The first marginal modeling step consists in capturing the temporal non-stationarity in the anomalies, *i.e.*, fitting the parametric regression model in (1). The fitted model parameters

$\hat{\mu}(t)$ and $\hat{\sigma}(t)$ are displayed in Figure 1. The second marginal modeling step is to correct the Gaussian tail using a spatial GP distribution with either the GAM structure in (3) or the nearest-neighbors (NN) approach detailed in Section 3.1.3. Figure 6 shows the fitted GP scale, shape, and exceedance probability across the entire region for both approaches. The fitted scale parameters display a North to South effect that is consistent with the spatial pattern in Figure 3. The shape is consistently negative, which implies a finite upper bound for the distribution of SST.

As a graphical goodness-of-fit test for the GP model, we display in Figure 7 the QQ-plots of the normalized data $Z(\mathbf{s}, t)$ with no tail correction and with a tail correction using both the GAM and the NN approaches. Though the QQ-plots match for non-exceedances, we observe a distinct improvement in the tail region when we resort to extreme-value methods and properly model the spatial pattern in the large normalized anomalies.

4.2 Local Gaussian dependence models

We have fitted 324 local Gaussian models, each spanning over nine days and centered at one of the 324 distinct validation days t_i . We do not report exhaustive results, but to summarize some characteristics of the fits, we show the time series of the following estimated hyperparameters in Figure 8: Matérn range, Matérn standard deviation, and temporal auto-correlation coefficient. Interestingly, some obvious trends arise over time. The three rows of the figure correspond to the three approaches to handling marginal distributions, respectively. Estimated values for the Matérn range and temporal auto-correlation are very similar between the three models. Regarding the Matérn standard deviation, we obtain smaller values for Model 1 without any marginal pretransformation of data to unit variance, and one strongly outlying value arises with Model 2. All three models show the same patterns of temporal trends: ranges tend to increase towards the end of the validation period (years 2007-2015) shown in the plot, which implies that spatial dependence has increased. For the standard deviation, a pattern indicating stronger variability at the beginning and the end of the validation period arises. Finally, the temporal dependence captured in the autoregression coefficient is also increasing towards the end of the validation period, by analogy with the spatial dependence.

Figure 9 illustrates the prediction on the Gaussian scale for a validation day in December 2019. The map on the left shows the available data, pretransformed according to the nearest-neighbor approach for marginal distributions, and the validation locations. The map in the




Figure 6: Generalised Pareto scale, shape, and exceedance probability fitted to exceedances over the value 0.75 in a standard Gaussian scale using the GAM approach (top panels, standard deviations displayed in the bottom pannels) and the NN approach (middle panels).




Figure 7: QQ-plots (Gaussian scale) of the normalized data $Z(\mathbf{s}, t)$ with no tail-correction (top panels) and tail-correction using the GAM approach (middle panels) and the NN approach (bottom panels), at the three locations indicated in Figure 1.

middle shows posterior means of predictions obtained through INLA. Finally, the right-hand map shows the corresponding posterior standard deviations. As expected, prediction uncertainty is lower when locations are close to observed locations but can be relatively high in the case of large gaps, such as the one in the southern part of the Red Sea.

4.3 Score evaluation

We compare Models 1–3 by applying the tail-weighted Continuous Ranked Probability Score (twCRPS) used for ranking predictions of the EVA challenge (Huser, 2020). Given an estimated predictive distribution $\hat{F}_{\mathbf{s},t}$ of the realized cluster summary $x(\mathbf{s}, t)$, here defined as the minimum

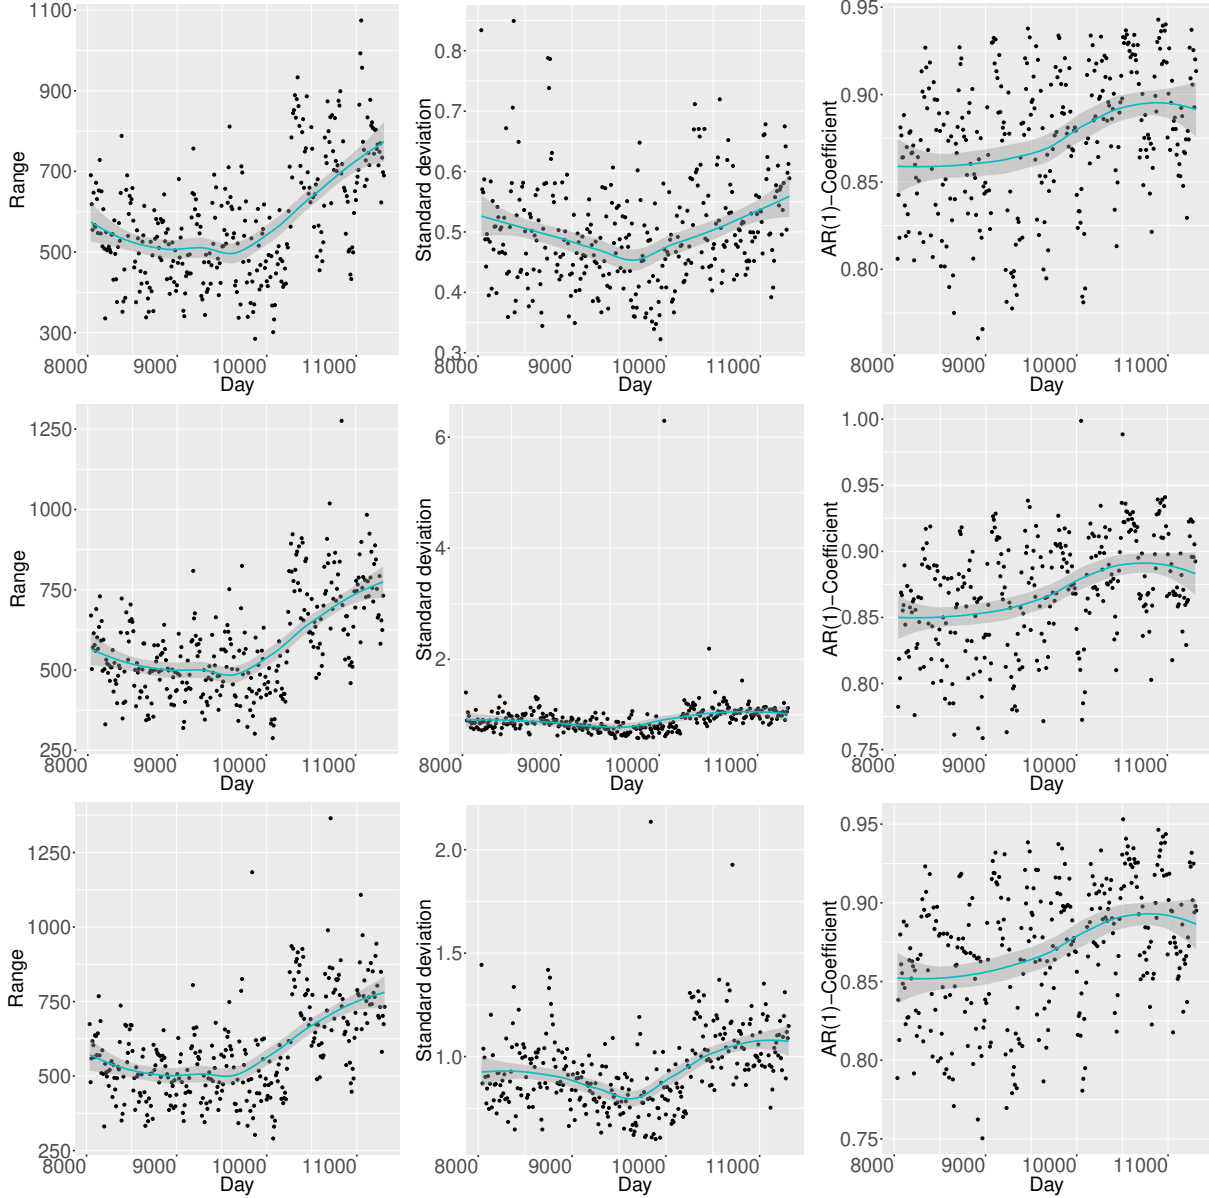


Figure 8: Posterior mean estimates of hyperparameters of the 324 local Gaussian models, and local kernel smoothing. First row: Model without marginal transformation. Second row: marginal GAM. Third row: marginal NN model. Left: Matérn range. Middle: Matérn standard deviation. Right: Temporal autoregression coefficient ρ .

over space-time cylinders, the twCPRS is defined using a weight function $\omega(x)$ as

$$\text{twCPRS}\{\hat{F}_{s,t}, x(s, t)\} = \int_{-\infty}^{\infty} \left[\hat{F}_{s,t}(x) - 1 \{x(s, t) \leq x\} \right]^2 \omega(x) dx, \quad \omega(x) = \Phi\{(x-a)/0.4\}, \quad a \in \mathbb{R}.$$

We consider the weight function of the EVA challenge with $a = 1.5$, corresponding to a Gaussian cdf that puts little weight to values below 1. Moreover, we assess our models by evaluating the twCRPS for weight functions that put stronger weight to more extreme or




Figure 9: Example of prediction through local Gaussian model using the GAM (top) and NN (bottom) approaches. Left: Marginally transformed observations for validation day t_i (in December 2009); black dots indicate validation locations. Middle: Posterior means of the predictions from the local Gaussian model. Right: Posterior standard deviation of the predictions from the local Gaussian model.

less high quantiles by choosing $a = 1.0$, $a = 1.8$, and $a = -\infty$, the latter corresponding to the classical non-weighted CRPS. Table 1 summarizes the scores of Models 1–3 for different tail weight functions. In general, the scores of the three models are relatively close, which confirms that, locally, data are relatively smooth in space and time, such that a Gaussian dependence model performs already very well. The twCRPS scores of the GAM and nearest-neighbor (NN) marginal models are systematically improved as compared to the Gaussian model without marginal transformations, except for the classical CRPS. The performances of GAM and NN are relatively similar, but with better twCRPS for the NN model when putting

	Gaussian	GAM	NN
$a = 1.8$	0.828	0.825	0.770
$a = 1.5$	3.27	3.19	3.07
$a = 1.0$	21.7	21.1	20.9
$a = -\infty$	702	717	717

Table 1: Values of $10^4 \times \text{twCRPS}$ of Models 1–3 using different tail weight functions.

weight on predictions at high quantiles. We conjecture that this is due to some local effects that are more easily captured through local pooling than through covariate modeling, and that the resulting discrepancy seems to be amplified when predicting very high quantiles of the cluster summary. We finally point out that our three models outperform all models developed by other teams for the data of the EVA 2019 challenge (Huser, 2020), with their best score for $a = 1.5$ being 3.67, while a simple benchmark score for $a = 1.5$ was calculated to be 7.89.

5 Discussion

We have developed a probabilistic, likelihood-based approach to fill gaps in space-time datasets with particular attention to extreme values and multivariate distributions. By estimating a generative model in a Bayesian framework, we have simulated from the posterior distribution of unobserved data points to obtain representative samples of the cluster functionals such as the minimum over space-time cylinders. Due to the large size of the dataset, joint estimation of margins and dependence using all data seems out of reach. We therefore consider a two-step approach, where we first remove trends to make data more “Gaussian-like”, especially at high quantiles. The second step consists of a Gaussian model fitted locally in space and time to the transformed data. Although pixel-wise observations seemed to be approximately Gaussian-distributed, the first step in the modeling procedure improved the predictive performance. We believe that much more substantial benefits can be expected from applying the two-step procedure to other data whose features differ more strongly from a stationary Gaussian distribution.

Our approach shows that we do not need an asymptotic model from extreme-value theory for capturing the dependence during extreme episodes. Based on our preliminary analysis showing asymptotic independence, we conjecture that our local Gaussian models are sufficiently flexible to capture spatial and temporal variations during extreme episodes. Moreover, the asymptotic independence in the data suggests that using non-asymptotic models such as

Gaussian random fields might provide a more appropriate framework compared to the asymptotic dependence structure imposed by extreme-value models such as max-stable processes and generalized Pareto processes. For modeling spatio-temporal dependence, Gauss–Markov processes provide good tractability and scalability properties for implementing prediction with large datasets, such that our choice is also of pragmatic nature. The Red Sea SST data of OSTIA are produced using geostatistical interpolation techniques similar to kriging on original data, and they are relatively smooth in space and time due to strong dependence. This is another hint of why the Gaussian dependence models perform very well for prediction.

Due to the artificial generation of gaps in the Red Sea data, we could validate and compare modeling approaches using the true values to be predicted. When this is not the case, external validation could be implemented through cross-validation schemes by predicting cluster functionals for space-time subsets where all observations are available.

The code written for the implementation of the approach presented in this paper is freely available from <https://github.com/dcastrocamilo/EVChallenge2019>.

Acknowledgements

This work started when Daniela Castro-Camilo was a postdoctoral fellow at King Abdullah University of Science and Technology (KAUST). Support from the KAUST Supercomputing Laboratory and access to Shaheen is therefore gratefully acknowledged. Linda Mhalla acknowledges the financial support of the Swiss National Science Foundation.

References

- Bivand, R., Gómez-Rubio, V. and Rue, H. (2015) Spatial data analysis with R-INLA with some extensions. *Journal of Statistical Software* **63**(20), 1–31.
- Bortot, P., Coles, S. and Tawn, J. (2000) The multivariate Gaussian tail model: An application to oceanographic data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **49**(1), 31–049.
- Cantin, N. E., Cohen, A. L., Karnauskas, K. B., Tarrant, A. M. and McCorkle, D. C. (2010) Ocean warming slows coral growth in the central Red Sea. *Science* **329**(5989), 322–325.
- Castro-Camilo, D., Huser, R. and Rue, H. (2019) A spliced gamma-generalized Pareto model for short-term extreme wind speed probabilistic forecasting. *Journal of Agricultural, Biological and Environmental Statistics* **24**(3), 517–534.

- Chaidez, V., Dreano, D., Agusti, S., Duarte, C. M. and Hoteit, I. (2017) Decadal trends in Red Sea maximum surface temperature. *Scientific reports* **7**(1), 1–8.
- Chavez-Demoulin, V. and Davison, A. C. (2005) Generalized additive modelling of sample extremes. *Journal of the Royal Statistical Society, Series C* **54**, 207–222.
- Cressie, N. and Wikle, C. K. (2015) *Statistics for spatio-temporal data*. John Wiley & Sons.
- Davison, A. C., Padoan, S. and Ribatet, M. (2012) Statistical modelling of spatial extremes. *Statistical Science* **27**(2), 161–186.
- De Coninck, A., De Baets, B., Kourounis, D., Verbosio, F., Schenk, O., Maenhout, S. and Fostier, J. (2016) Needles: Toward large-scale genomic prediction with marker-by-environment interaction. *Genetics* **203**(1), 543–555.
- 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.
- Ferreira, A. and De Haan, L. (2014) The generalized Pareto process; with a view towards application and simulation. *Bernoulli* **20**(4), 1717–1737.
- Gneiting, T. and Ranjan, R. (2011) Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics* **29**(3), 411–422.
- Henn, B., Raleigh, M. S., Fisher, A. and Lundquist, J. D. (2013) A comparison of methods for filling gaps in hourly near-surface air temperature data. *Journal of Hydrometeorology* **14**(3), 929–945.
- Hoegh-Guldberg, O., Cai, R., Poloczanska, E. S., Brewer, P., Sundby, S., Hilmi, K., Fabry, V. J. and Jung, S. (2014) The Ocean. In: *Climate Change 2014: Impacts, Adaptation, and Vulnerability. part B: Regional Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. [Barros, V.R., Field, C.B., Dokken, D.J., Mastrandrea, M.D., Mach, K.J., Bilir, T.E., Chatterjee, M., Ebi, K.L., Estrada, Y.O., Genova, R.C., Girma, B., Kissel, E.S., Levy, A.N., MacCracken, S., Mastrandrea, P.R., and White, L.L. (eds.)]. *Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 1655-1731* .
- Huser, R. (2020) EVA 2019 data competition on spatio-temporal prediction of Red Sea surface temperature extremes. *Extremes* pp. 1–14.
- Jonathan, P., Randell, D., Wu, Y. and Ewans, K. (2014) Return level estimation from non-stationary spatial data exhibiting multidimensional covariate effects. *Ocean Engineering* **88**, 520–532.

- Kourounis, D., Fuchs, A. and Schenk, O. (2018) Toward the next generation of multiperiod optimal power flow solvers. *IEEE Transactions on Power Systems* **33**(4), 4005–4014.
- Krainski, E. T., Gómez-Rubio, V., Bakka, H., Lenzi, A., Castro-Camilo, D., Simpson, D., Lindgren, F. and Rue, H. (2018) *Advanced Spatial Modeling with Stochastic Partial Differential Equations Using R and INLA*. Chapman and Hall/CRC.
- Lakshmi, V. (2016) *Remote sensing of hydrological extremes*. Springer.
- Lindgren, F., Rue, H. and Lindström, J. (2011) An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498.
- Mariethoz, G., McCabe, M. F. and Renard, P. (2012) Spatiotemporal reconstruction of gaps in multivariate fields using the direct sampling approach. *Water Resources Research* **48**(10).
- Mhalla, L., de Carvalho, M. and Chavez-Demoulin, V. (2019) Regression-type models for extremal dependence. *Scandinavian Journal of Statistics* **46**(4), 1141–1167.
- van Niekerk, J., Bakka, H., Rue, H. and Schenk, L. (2019) New frontiers in Bayesian modeling using the INLA package in R. *arXiv preprint arXiv:1907.10426* .
- Opitz, T. (2017) Latent Gaussian modeling and INLA: A review with focus on space-time applications. *Journal of the French Statistical Society (Special Issue on Space-Time Statistics)* **158**(3).
- Opitz, T., Huser, R., Bakka, H. and Rue, H. (2018) INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles. *Extremes* **21**(3), 441–462.
- Padhee, S. K. and Dutta, S. (2019) Spatio-temporal reconstruction of MODIS NDVI by regional land surface phenology and harmonic analysis of time-series. *GIScience & Remote Sensing* **56**(8), 1261–1288.
- Pauli, F. and Coles, S. (2001) Penalized likelihood inference in extreme value analyses. *Journal of Applied Statistics* **28**, 547–560.
- Rue, H., Martino, S. and Chopin, N. (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal statistical society: Series B (Statistical Methodology)* **71**(2), 319–392.
- Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P. and Lindgren, F. K. (2017) Bayesian computing with INLA: a review. *Annual Review of Statistics and Its Application* **4**, 395–421.

- Simpson, D., Rue, H., Riebler, A., Martins, T. G., Sørbye, S. H. and Others (2017) Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science* **32**(1), 1–28.
- Spalding, M., Spalding, M. D., Ravilious, C., Green, E. P. *et al.* (2001) *World atlas of coral reefs*. Univ of California Press.
- Sun, Y. and Genton, M. G. (2011) Functional boxplots. *Journal of Computational and Graphical Statistics* **20**(2), 316–334.
- Thibaud, E. and Opitz, T. (2015) Efficient inference and simulation for elliptical Pareto processes. *Biometrika* **102**(4), 855–870.
- Tierney, L. and Kadane, J. B. (1986) Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* **81**(393), 82–86.
- Verbosio, F., Coninck, A. D., Kourounis, D. and Schenk, O. (2017) Enhancing the scalability of selected inversion factorization algorithms in genomic prediction. *Journal of Computational Science* **22**(Supplement C), 99 – 108.
- Wang, G., Garcia, D., Liu, Y., De Jeu, R. and Dolman, A. J. (2012) A three-dimensional gap filling method for large geophysical datasets: Application to global satellite soil moisture observations. *Environmental Modelling & Software* **30**, 139–142.
- Wood, S. N. (2003) Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(1), 95–114.
- Wood, S. N. (2006) Low-rank scale-invariant tensor product smooths for generalized additive mixed models. *Biometrics* **62**(4), 1025–1036.
- Wood, S. N. (2011) Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(1), 3–36.
- Wood, S. N. (2017) *Generalized Additive Models: An Introduction with R*. Second edition. Chapman and Hall/CRC.
- Wood, S. N., Pya, N. and Säfken, B. (2016) Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association* **111**(516), 1548–1563.
- Xing, C., Chen, N., Zhang, X. and Gong, J. (2017) A machine learning based reconstruction method for satellite remote sensing of soil moisture images with in situ observations. *Remote Sensing* **9**(5), 484.

- Yin, G., Mariethoz, G. and McCabe, M. F. (2017) Gap-filling of Landsat 7 imagery using the direct sampling method. *Remote Sensing* **9**(1), 12.
- Youngman, B. D. (2019) Generalized additive models for exceedances of high thresholds with an application to return level estimation for U.S. wind gusts. *Journal of the American Statistical Association* **114**(528), 1865–1879.
- Yuan, H., Dai, Y., Xiao, Z., Ji, D. and Shanguan, W. (2011) Reprocessing the MODIS leaf area index products for land surface and climate modelling. *Remote Sensing of Environment* **115**(5), 1171–1187.