

ORIGINAL ARTICLE OPEN ACCESS

Spatial Extremal Modelling: A Case Study on the Interplay Between Margins and Dependence

Lydia Kakampakou¹  | Emma S. Simpson²  | Jennifer L. Wadsworth¹ 

¹School of Mathematical Sciences, Lancaster University, Lancaster, UK | ²Department of Statistical Science, University College London, London, UK

Correspondence: Lydia Kakampakou (l.kakampakou1@lancaster.ac.uk)

Received: 6 June 2024 | **Revised:** 14 September 2024 | **Accepted:** 29 September 2024

Keywords: environmental extremes | extremal dependence | marginal modelling | non-stationarity | spatial extremes

ABSTRACT

It is no secret that statistical modelling often involves making simplifying assumptions when attempting to study complex stochastic phenomena. Spatial modelling of extreme values is no exception, with one of the most common such assumptions being stationarity in the marginal and/or dependence features. If non-stationarity has been detected in the marginal distributions, it is tempting to try to model this while assuming stationarity in the dependence, without necessarily putting this latter assumption through thorough testing. However, margins and dependence are often intricately connected and the detection of non-stationarity in one feature might affect the detection of non-stationarity in the other. This work is an in-depth case study of this interrelationship, with a particular focus on a spatio-temporal environmental application exhibiting well-documented marginal non-stationarity. Specifically, we compare and contrast four different marginal detrending approaches in terms of our post-detrending ability to detect temporal non-stationarity in the spatial extremal dependence structure of a sea surface temperature dataset from the Red Sea.

1 | Introduction

In environmental applications, it is not uncommon to observe non-stationary behaviour in the extreme values of a random variable. However, extreme value modelling often relies on a simplifying assumption of stationarity. In practice, non-stationarity may be present in different aspects of the data generating process; the more aspects it affects, the greater the degree of modelling complexity. Marginal non-stationarity refers to the situation where realisations of a univariate random variable are non-identically distributed. In environmental data, this varying behaviour typically depends on time, although other covariates, that themselves evolve with time, are normally the drivers of change. Various methods for modelling marginal non-stationarity of extremes have been well explored in the literature. One of the main approaches is to incorporate covariates directly into the parameters of univariate models for extreme value analysis. Typically, this will be the generalised Pareto model for excesses of a high threshold: See, for example, Davison

and Smith (1990), Chavez-Demoulin and Davison (2005) and Youngman (2019). Section 2.1 contains further details on these and other methods for modelling non-stationarity.

Non-stationarity may also be encountered in the extremal dependence structure of a random vector. That is, the extremal dependence between two or more components of a random vector may change with covariates. In such instances, it is common in analyses of multivariate or spatial extremes to handle marginal and dependence modelling as a two-step procedure, implementing marginal detrending methods first and then applying non-stationary methods for dependence modelling to the marginally detrended data. Although alternative methods that model marginal and copula behaviour jointly in a single step exist, they are not as numerous, and their application in extreme value modelling is scarce in practice. One reason for this is that univariate extreme value distributions are suited to marginal upper tails, while multivariate extremal dependence models rely on potentially different upper tail definitions. A

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Stat* published by John Wiley & Sons Ltd.

mismatch between marginal and copula support impedes joint modelling considerations. When the multivariate definition of extremity follows as a direct generalisation of the univariate one, as in the case of modelling maxima, Padoan, Ribatet, and Sisson (2010) and Ribatet, Cooley, and Davison (2012) develop such joint modelling procedures for max-stable processes, in a frequentist and Bayesian realm respectively. In a similar manner, multivariate generalised Pareto distributions (Rootzén and Tajvidi 2006; Rootzén et al. 2018) allow for joint marginal-dependence modelling, which could be extended to cover non-stationarity, while similar ideas could potentially be applied to spatial generalised Pareto processes (Buishand et al. 2008; Ferreira and de Haan 2014). However, both max-stable and generalised Pareto processes yield very restrictive dependence structures and are therefore unlikely to capture the complex dependence features often present in environmental applications (Huser and Wadsworth 2022). Even putting these issues aside, pursuing joint marginal-copula modelling brings about another important risk, even in the stationary case, that misspecification in one form — the margins or dependence — can affect the other. Such spillover is likely worsened when non-stationarity is present. In this paper, we adopt the more established two-stage approach, thus making the task of disentangling the relationship between margins and dependence, which is the focus of this work, as clear and straightforward as possible.

Available methods for non-stationary dependence modelling include the following. In a bivariate setting, de Carvalho and Davison (2014) develop a semi-parametric model that incorporates covariates into the modelling of the so-called *spectral density*, thus capturing the effect of explanatory variables on joint extremes. In a multivariate setting, Castro-Camilo, de Carvalho, and Wadsworth (2018) and Mhalla, de Carvalho, and Chavez-Demoulin (2019) also focus on spectral density functions with the former developing non-parametric regression-based methods and the latter building a vector generalised additive model to capture covariate-varying extremal dependence. Lee et al. (2024) model covariate(time)-varying extremal dependence in various summary statistics. Murphy-Barltrop and Wadsworth (2024) model temporally varying extremal dependence in a bivariate setting via non-stationary extensions of the so-called *angular dependence function* and *return curves*, while Mhalla, Opitz, and Chavez-Demoulin (2019) also make use of non-stationary angular dependence functions to account for covariate influence.

In a spatial setting, modelling changes in extremal dependence using space-related covariates, which may be expected when the spatial domain is large or morphologically diverse, has received the most attention; see, for example, Huser and Genton (2016), Blanchet and Davison (2011) and Richards and Wadsworth (2021). Much less attention has been dedicated, however, to time-related covariates, which, given the ongoing climate crisis, may have considerable effects on the spatial extremal dependence structure over time, especially in applications spanning a large temporal domain. A noteworthy exception of research in this direction is that of Healy et al. (2024) who model temporal non-stationarity in both marginal and dependence features. They combine information on Irish temperatures from both observational and climate-model sources and consider

time-varying dependence via the framework of *r*-Pareto processes (Dombry and Ribatet 2015).

Of the aforementioned examples, Healy et al. (2024) is the only one dealing with spatial aspects of time-varying extremal dependence; the remaining do so in a multivariate setting. To the best of our knowledge, there is no other contribution where this problem is tackled in a fully spatial context available in the literature. What is more, Healy et al. (2024) focus on a class of models that can only accommodate a particular type of dependence, namely asymptotic dependence (see Section 2.2 for details), which, although appropriate for their application, is often unrealistic for similar spatial environmental datasets.

It is this lack of available literature that provided the initial motivation for the present work. Our study is focused on a dataset with well documented marginal non-stationarities (Huser and Genton 2016), which was also believed to showcase non-stationary behaviour in its spatial extremal dependence structure. The data comprise daily sea surface temperature values from the Gulf of Suez spanning across 31 years (1985 to 2015 inclusive). As is common in analyses of extremes, we focused our attention on the summer data (July–August–September) from that period, to reduce the degree of seasonality to be modelled, while considering the period of highest temperatures. The original data are available for the whole of the Red Sea, on a daily basis and a very fine resolution grid of $0.05^\circ \times 0.05^\circ$ (Donlon et al. 2012). Initial analysis of spatial dependence changes over time for this dataset suggested the need for adapting spatial dependence models to account for time-varying dependence (see top left plot of Figure 3).

However, following our initial exploratory analysis, we tested various approaches to modelling marginal non-stationarity, finding that subsequent results concerning changes in spatial extremal dependence were highly sensitive to the marginal model used. We thus came to the conclusion that the deciding factor behind whether or not temporal non-stationarity is detectable in the spatial extremal dependence structure of the data critically relies on the marginal detrending approach pursued and that any resultant analysis is, in our case, extremely sensitive to marginal modelling choices. As such, instead of focusing on adapting existing models to accommodate changes in spatial dependence over time, the aim of this paper is to serve as a cautionary tale of the misleading effects the marginal and dependence structure interrelationship can have on analyses of extremes. Our main objective is to investigate the effects of four different marginal detrending approaches on spatial extremal dependence. In particular, we use graphical tools to assess whether the apparent spatial extremal dependence changes over time under each marginal detrending procedure. We note here that the dataset in question displays strong autocorrelation; such a feature will remain after modelling trends in the data, but can affect visual assessment of stationarity. Graphical tools should therefore be interpreted with some degree of caution. We further assess the impact of each marginal model on the spatial extremal dependence under a (potentially false) assumption of temporal stationarity by applying the conditional spatial extremes model (Wadsworth and Tawn 2022) to all four marginally treated versions of the Gulf of Suez dataset and comparing model-based dependence features. This particular spatial model

was selected because of its computational and modelling flexibility, which allow for increased scalability (in terms of the number of spatial locations that can be handled) compared to most alternative models available, as well as simpler characterisation/modelling of different types of extremal dependence.

The remainder of the paper is outlined as follows. Section 2 provides background information for our statistical approach. Section 3 describes the four marginal detrending methods employed. Section 4 is dedicated to the analysis of the Gulf of Suez data. Finally, we conclude in Section 5 with a short discussion.

2 | Statistical Background

2.1 | Marginal Modelling of Extremes

Marginal modelling in extremes concerns the study of univariate extreme values, which appear either in isolation, or as components of a higher dimensional random vector. When studying a random variable X , asymptotic theory dictates that the distribution of threshold exceedances $X - u | X > u$ of a sufficiently high threshold u can usually be approximated by the limiting form G in (1), known as the generalised Pareto distribution (GPD):

$$G(x) = \begin{cases} 1 - (1 + \xi x / \sigma_u)_+^{-1/\xi}, & \text{for } \xi \neq 0, \sigma_u > 0, x > 0, \\ 1 - \exp(-x / \sigma_u), & \text{for } \xi \rightarrow 0, \sigma_u > 0, x > 0, \end{cases} \quad (1)$$

for $a_+ = \max(a, 0)$; we write $X - u | X > u \sim \text{GPD}(\sigma_u, \xi)$. The GPD model is suitable for capturing the tail behaviour of stationary processes. When marginal non-stationarity is present, Davison and Smith (1990) proposed incorporating covariates into the parameters of the GPD. That is, given a marginally non-stationary process $\{X_t\}$ with associated covariates \mathbf{Z}_t , one can model $(X_t - u | X_t > u, \mathbf{Z}_t = \mathbf{z}_t) \sim \text{GPD}(\sigma_u(\mathbf{z}_t), \xi(\mathbf{z}_t))$ for a high enough threshold u . Note that covariate modelling here is performed according to the principles of generalised linear models, where typically we set $\log \sigma_u(\mathbf{z}_t) = \boldsymbol{\beta}_\sigma^T \mathbf{z}_t$, $\xi(\mathbf{z}_t) = \boldsymbol{\beta}_\xi^T \mathbf{z}_t$. Subsequent work has considered more flexible covariate formulations, such as those provided by generalised additive models (GAMs); see, for example, Chavez-Demoulin and Davison (2005) and Youngman (2019).

One problem with such techniques is that their definition of extremity often depends on a constant threshold, but under non-stationarity, the extremes are suitably defined through the corresponding covariate levels as well. To address this issue, Northrop and Jonathan (2011) and Jonathan, Ewans, and Randell (2013) define covariate-varying thresholds. Another concern is that these techniques focus on modelling non-stationary behaviour only in the tail of a process. However, it is often of interest to model non-stationarity in the body of the data as well, since extremes across components of a random vector do not necessarily occur simultaneously. That is, we may observe extremity in one component but moderate levels of another. If marginal non-stationarity in the body of the distribution is neglected, then we cannot model this dependence appropriately.

To overcome these issues, Eastoe and Tawn (2009) borrow techniques from time-series modelling. They introduce a two-step

procedure comprising a pre-processing step, aiming to remove most of the non-stationarity present in the body of the data, and a non-stationary GPD step, to remove any leftover trends in the tails. The pre-processing step consists of applying a Box-Cox location-scale transformation of the form

$$\frac{X_t^{\lambda(\mathbf{z}_t)} - 1}{\lambda(\mathbf{z}_t)} = \mu(\mathbf{z}_t) + \sigma(\mathbf{z}_t) R_t \quad (2)$$

to the marginally non-stationary process $\{X_t\}$, where $\{R_t\}$ is approximately stationary and λ , μ , and $\log(\sigma)$ are linear functions of covariates \mathbf{z}_t . The process $\{R_t\}$ is assigned a Gaussian distribution, so that specification (2) has a likelihood that is maximised to provide parameter estimates. This is then followed by a non-stationary GPD fit to the residual process $\{R_t\}$, using the methodology of Davison and Smith (1990). If any trends are identified in the scale and shape parameters, $\sigma_u(\mathbf{z}_t)$ and $\xi(\mathbf{z}_t)$ respectively, they are removed via the probability integral transformation, using a semi-parametric approach to estimation of the distribution of R_t (Coles and Tawn 1991). This consists of a rank transform for data below a high marginal threshold, u , which are now assumed stationary, and the GPD cumulative distribution function for values above u . To make things more explicit, let r_1, \dots, r_n be observations from the process $\{R_t\}$ at times 1, \dots , n , and define

$$\hat{F}_{R_t}(r | \mathbf{Z}_t = \mathbf{z}_t) = \begin{cases} \sum_{i=1}^n \mathbb{1}(r_i \leq r) / (n+1), & r \leq u \\ 1 - (1-u) \left[1 + \hat{\xi}(\mathbf{z}_t)(r-u) / \hat{\sigma}_u(\mathbf{z}_t) \right]_+^{-1/\hat{\xi}(\mathbf{z}_t)}, & r > u, \end{cases}$$

where $\hat{\sigma}_u(\mathbf{z}_t)$ and $\hat{\xi}(\mathbf{z}_t)$ are the maximum likelihood estimates of the scale and shape parameters respectively, which are estimated by the non-stationary GPD step. Applying \hat{F}_{R_t} to the marginal process $\{R_t\}$ achieves standardisation to a standard uniform scale. Finally, we note that some more recent work replaces Eastoe and Tawn (2009)'s pre-processing step with a more elaborate one which, instead of assuming linear parametric forms for the covariate functions, makes use of the GAM framework to allow parameters $\mu(\mathbf{z}_t)$, $\sigma(\mathbf{z}_t)$ of (2) and σ_u, ξ of the non-stationary version of (1) to vary smoothly with covariates (Murphy-Bartrop and Wadsworth 2024). This extra flexibility is desirable when the observed non-stationarities are too complicated to be adequately captured by parametric forms.

2.2 | Measures of Extremal Dependence

A key endeavour of many multivariate analyses is to characterise the extremal dependence between $D \geq 2$ random variables. Apart from the obvious exploratory benefits of understanding the data better by uncovering its dependence structure before attempting any modelling, another reason why such a task is of importance is that many of the available models—and their underpinning assumptions—can only accommodate data that fall under a single category of dependence classes. Hence, D -variate measures of extremal dependence are often examined for dependence-class assessment and categorisation in multivariate analyses. In the spatial case however, one can often assume that bivariate distributions ($D = 2$) of the spatial process at

location pairs are very informative and, thus, looking at bivariate versions of extremal dependence measures for different pairs is all one needs to characterise the extremal dependence of the process at given distances.

A widely used dependence measure in practice is the so-called *tail dependence coefficient*, typically denoted by χ . Let $\{Y(\mathbf{s}): \mathbf{s} \in S \subset \mathbb{R}^2\}$ be a stationary and isotropic spatial random field and $(Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_D))$ the D -dimensional random vector comprised of the realisations of $\{Y(\mathbf{s})\}$ at the observed spatial locations $\mathbf{s}_i, i = 1, \dots, D$. Letting $Y(\mathbf{s}_i) \sim F$ denote the common marginal distribution, the coefficient of tail dependence $\chi(h_{jk})$ is given by $\chi(h_{jk}) = \lim_{u \rightarrow 1} \chi_u(h_{jk})$, with

$$\chi_u(h_{jk}) = \Pr\{F(Y(\mathbf{s}_j)) > u | F(Y(\mathbf{s}_k)) > u\}, \quad u \in [0,1], \quad (3)$$

where $h_{jk} = \|\mathbf{s}_j - \mathbf{s}_k\|$ is the distance between sites j and $k (j \neq k, j, k \in \{1, \dots, D\})$; we will take this to be the Euclidean distance for the remainder of the paper. The quantity $\chi(h_{jk})$ is the limiting probability of the process being extreme at location \mathbf{s}_j given that it is also extreme at location \mathbf{s}_k and characterises the extremal dependence between the pair (j, k) at distance h_{jk} apart. A process is said to be *asymptotically dependent* (AD) in the spatial domain S if $\chi(h_{jk}) > 0$ for all $h_{jk}, j, k \in \{1, \dots, D\}$. In practical terms, asymptotic dependence (AD) means that the variables, or realisations of the spatial process at different locations, can take their most extreme values simultaneously. Conversely, the case where $\chi(h_{jk}) = 0$ for all $h_{jk}, j, k \in \{1, \dots, D\}$, defines an *asymptotically independent* (AI) process in the spatial domain S , which means that the spatial extent of the extremes becomes more and more localised as the level of extremity increases. In such cases, the measure $\chi(h_{jk})$ does not give us any additional information on the joint tail dependence. A complementary statistic to $\chi(h_{jk})$ that is more informative under asymptotic independence (AI) is the *residual tail dependence coefficient*, $\eta(h_{jk}) \in (0,1]$, of Ledford and Tawn (1996), which may be defined through the relation

$$\Pr\{F(Y(\mathbf{s}_j)) > u, F(Y(\mathbf{s}_k)) > u\} = \mathcal{L}(1-u)(1-u)^{1/\eta(h_{jk})}, \quad u \rightarrow 1, \quad (4)$$

where \mathcal{L} is a *slowly varying function* such that $\mathcal{L}(cr)/\mathcal{L}(r) \rightarrow 1$ as $r \rightarrow 0$ for all positive constants c . Note that relation (4) implies $\eta(h_{jk}) = \lim_{u \rightarrow 1} \eta_u(h_{jk})$, with

$$\eta_u(h_{jk}) = \frac{\log(1-u)}{\log(\Pr\{F(Y(\mathbf{s}_j)) > u, F(Y(\mathbf{s}_k)) > u\})}, \quad u \in [0,1], \quad (5)$$

$j \neq k, j, k \in \{1, \dots, D\}$. The boundary case $\eta(h_{jk}) = 1$ corresponds to $\chi(h_{jk}) = \lim_{u \rightarrow 1} \mathcal{L}(1-u)$. Hence, when $\eta(h_{jk}) = 1$ for all $h_{jk}, j, k \in \{1, \dots, D\}$, and $\mathcal{L}(1-u) \rightarrow 0$ as $u \rightarrow 1$ we get $\chi(h_{jk}) > 0$ for all $h_{jk}, j, k \in \{1, \dots, D\}$; that is, the process is AD in the spatial domain S . The more interesting case where either $\eta(h_{jk}) < 1$ for all $h_{jk}, j, k \in \{1, \dots, D\}$, or $\eta(h_{jk}) = 1$ for all $h_{jk}, j, k \in \{1, \dots, D\}$, and $\mathcal{L}(1-u) \rightarrow 0$ as $u \rightarrow 1$, leads to $\chi(h_{jk}) = 0$ for all $h_{jk}, j, k \in \{1, \dots, D\}$, and signifies that the process is AI in S , which can be further categorised into *negative extremal association* when $\eta(h_{jk}) \in (0,1/2)$, *near independence* when $\eta(h_{jk}) = 1/2$ and *positive extremal association* when

$\eta(h_{jk}) \in (1/2,1]$. The latter case is of particular interest to environmental applications which usually exhibit positive dependence over space. Because we cannot typically estimate the limiting quantities $(\chi(h_{jk}), \eta(h_{jk}))$, the behaviour of $(\chi_u(h_{jk}), \eta_u(h_{jk}))$ as $u \rightarrow 1$ is used instead to help inform us about the extremal dependence structure of the process.

As already mentioned, distinguishing between extremal dependence classes is important for modelling. Doing so in practice, however, is a non-trivial task. With spatial data, for example, we might have $\chi_u(h_{jk}) > 0$ at all observable high quantiles, u , and pairs of sites, but $\chi_u \searrow 0$ as $u \nearrow 1$. It is frequently encountered in practice that the dependence of environmental processes weakens as events become more extreme with the very severe extreme events becoming more spatially localised (Huser and Wadsworth 2022). One might also observe that the dependence class of the process is not the same for all distances, but changes from AD to AI after a distance Δ , or with respect to other covariates. This creates additional modelling difficulties and highlights the importance of thorough exploratory investigation of the dependence structure as a first step to overcoming said difficulties.

2.3 | Conditional Spatial Extremes Model

So-called *conditional extremes methods* for extremal dependence modelling were first developed in a multivariate setting by Heffernan and Tawn (2004) and further studied by Heffernan and Resnick (2007). In contrast to previously existing methods that relied on the limiting behaviour of a D -dimensional random vector when all of its components become simultaneously extreme at the same rate, the authors developed novel methodology based on a limiting assumption on the joint tail of a vector given that one of its components is extreme. This alternate limiting assumption of the Heffernan and Tawn (2004) model is much less restrictive, allowing for increased flexibility both in dependence modelling and in fitting scalability. With these benefits in mind, Wadsworth and Tawn (2022) extended the work of Heffernan and Tawn (2004) and Heffernan and Resnick (2007) to a spatial setting, termed the conditional spatial extremes (CSE) model.

Let $\{Y(\mathbf{s}): \mathbf{s} \in S \subset \mathbb{R}^2\}$ be a stationary and isotropic spatial random field with exponential-tailed margins, that is, $\Pr(Y(\mathbf{s}_j) > x) \sim ce^{-x}, c > 0$, as $x \rightarrow \infty$; a typical marginal choice is standard Laplace, which we also adopt in our subsequent analyses. For a high marginal threshold t , Wadsworth and Tawn (2022) show that for many dependence structures, a reasonable assumption is

$$\{Y(\mathbf{s}) | Y(\mathbf{s}_0) > t: \mathbf{s} \in S\} \stackrel{d}{\approx} \{a_{\mathbf{s}-\mathbf{s}_0}(Y(\mathbf{s}_0)) + b_{\mathbf{s}-\mathbf{s}_0}(Y(\mathbf{s}_0))Z^0(\mathbf{s}): \mathbf{s} \in S\}, \quad (6)$$

which can be used to model the behaviour of the process $\{Y(\mathbf{s})\}$ given an extreme value has occurred at some location \mathbf{s}_0 . Specifying functional forms for $a_{\mathbf{s}-\mathbf{s}_0}$ and $b_{\mathbf{s}-\mathbf{s}_0}$, where $a_{\mathbf{s}-\mathbf{s}_0}: \mathbb{R} \rightarrow \mathbb{R}, a_0(y) = y$, and $b_{\mathbf{s}-\mathbf{s}_0}: \mathbb{R} \rightarrow (0, \infty), \mathbf{s} \in S$, and a distributional form for the so-called *residual process* $\{Z^0\}$ — satisfying the condition $Z^0(\mathbf{s}_0) = 0$ — completes the model. The dependence structure of the residual process $\{Z^0\}$ is assumed

to be Gaussian in nature, while its margins are modelled via the so-called *delta-Laplace* distribution. Inference for the model in (6) is performed based on a composite likelihood scheme, combining likelihood contributions stemming from different conditioning locations (all or a subset of available locations can be selected), thus overcoming the problem of choice of conditioning location which is often not obvious in applications. Additional details on the CSE model are given in Appendix A.

3 | Marginal Detrending Methods

In Sections 3.1–3.4, we outline four approaches considered for capturing marginal non-stationarities in our dataset. In each case, our objective is to transform the data to follow a stationary standard Laplace distribution, that is, with distribution function

$$F_L(x) = \begin{cases} 0.5\exp\{x\}, & x \leq 0, \\ 1 - 0.5\exp\{-x\}, & x > 0, \end{cases}$$

enabling subsequent application of the CSE model. We do this in two steps. We first transform our data to the standard uniform scale by implementing each detrending procedure outlined below. Then, assuming all trends have been successfully removed and that, therefore, our data are now stationary, we apply F_L^{-1} , the Laplace inverse cumulative distribution function, to the uniform data to bring them to the required marginal scale. Figure 2 displays the data on Laplace margins following each of the four processing techniques. As noted in Section 1, strong temporal dependence between observations within a year can hinder visual assessment of stationarity from such plots. All marginal detrending procedures are implemented site-wise.

3.1 | Margins A

The first marginal detrending approach considered comes from the simple but crude assumption that data are stationary within a year—facilitated by considering data only from a single season—but that their distribution may change over the years. It consists of a simple rank transform applied in yearly blocks, whereby data corresponding to the same year are grouped, rank-transformed together and collated post-transformation. Implementing this transformation separately for each year allows for long term and inter-annual trends—which, according to Figure 1, both seem to be present for the Gulf of Suez dataset—without the need to explicitly specify the form of such trends.

The main merit of such an approach lies precisely in its simplicity, which in turn leads to very fast and straightforward implementation. On the other hand, performing the rank operation yearly could lead to cruder empirical approximations of the underlying marginal distributions—given only 92 data points per year are available in our case—and introduces the repeated values that can be seen in the top-left panel of Figure 2 across the full, transformed time series.




FIGURE 1 | Time series plot at location s_{35} on the original marginal scale.

3.2 | Margins B

The Eastoe and Tawn (2009) approach has been widely used in applications to achieve data detrending at a marginal level (e.g., Winter, Tawn, and Brown 2016; Winter et al. 2017). We employ this, in particular, its more flexible GAM extension, as our second detrending technique. As evidenced in Figure 1, the Gulf of Suez dataset is characterised by strong temporal trends of both cyclical (short-term) and linear (long-term) nature. We endeavour to remove such trends site-wise, via the GAM modelling framework of Wood (2017), as implemented in Murphy-Bartrop and Wadsworth (2024).

To be precise, let us denote the temperature process by $\{Y_t(\mathbf{s})\}$, $t = 1, \dots, T$. We are interested in applying detrending methods to all marginal time series $Y_{i,t} = Y_t(\mathbf{s}_i)$, for all $i \in \{1, \dots, D\}$. For every $Y_{i,t}$, we allow the location, μ_p , and scale, σ_p , parameters of (2) to vary smoothly with covariates $\mathbf{z}_t = \{1, \text{gmt}_t, d_t\}$, where gmt_t denotes global mean temperature anomalies⁴ at time t and $d_t = \{1, \dots, 92\}$ denotes the day index of the process at time t . The reason gmt_t is chosen instead of t is that we found it explained more of the observed trends than simply t . Moreover, its use has been advocated by others (e.g., D'Arcy, Tawn, and Sifnioti 2022; Healy et al. 2024) for better capturing long-term trends in environmental applications. Note that we do not include a Box-Cox parameter in this procedure for simplicity, as there is no obvious need for a changing shape parameter. Aiming to capture the long-term linear trend, a thin plate regression spline is fitted for the covariate gmt_t , while a cubic regression spline of dimension 92—found to result in the most stationary-looking data amongst a number of different choices examined—is used to capture the seasonal trends through the covariate d_t . The μ_i and σ_i parameters are estimated by means of the R package `mgcv` (Wood 2003; 2011) and subsequently removed, resulting in residual time series $R_{i,t} = (Y_{i,t} - \mu_i(\mathbf{z}_t))/\sigma_i(\mathbf{z}_t)$ for all i . We then proceed by applying a non-stationary GPD to $R_{i,t}$ using a constant 90% marginal threshold, thus aiming to remove any remaining trends in the tails. A likelihood ratio test is used to determine whether no, linear, seasonal or both trends should be added to the scale parameter of the GPD for each site \mathbf{s}_i . The shape parameter in all margins is assumed to be invariant to temporal

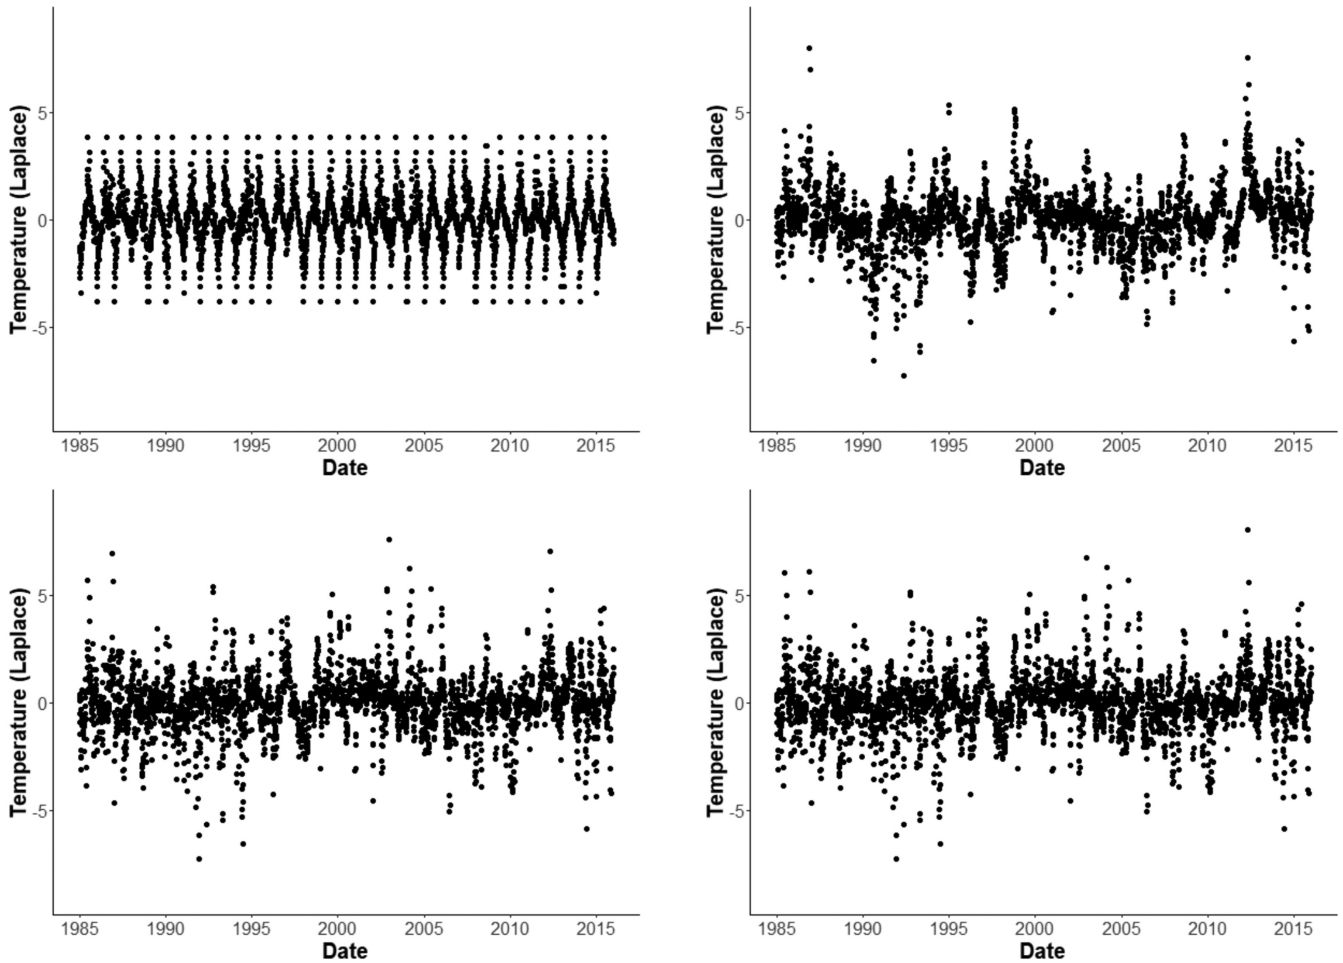


FIGURE 2 | Time series for location s_{35} in standard Laplace margins. [Left to right and top to bottom] A–B–C–D margins. Detrending is performed based on a 90% constant marginal threshold for margins A–C, while for margins D, an 83% threshold was automatically selected for this location.

trends. The detrended data are visualised on the standard Laplace scale in the top right plot of Figure 2.

3.3 | Margins C

The third marginal detrending method examined is motivated by the unsatisfactory degree of detrending achieved in margins B. In particular, there seems to be an inconsistency in the number of exceedances of a fixed high threshold observed per year, with some years experiencing multiple exceedances and some very few, if any. It appears therefore that the detrending method described in Section 3.2 fails to capture the full structural complexity of the margins, resulting in leftover marginal ‘irregularities’ manifesting as inter-year variability. Eastoe (2019) and Clarkson, Eastoe, and Leeson (2023) have reported similar findings in their analyses of a river flow and a US temperature dataset, respectively. To account for leftover inter-year variability they develop a random-effects-based methodology. Inspired by their method, we introduce an additional intermediate step to the aforementioned margins B procedure, acting as a crude approximation of the Eastoe (2019) approach.

Let \mathcal{Y}_k denote the set of time points t corresponding to year k , where $k = \{1, \dots, 31\}$. Since each year is comprised of 92 summer-month

observations, $\mathcal{Y}_1 = \{1, \dots, 92\}$, $\mathcal{Y}_2 = \{93, \dots, 184\}$, and so forth. The residual process $\{R_{i,t}\}$ is then adapted as follows:

$$\frac{\{R_{i,t}\}_{t \in \mathcal{Y}_k} - m_k}{s_k}, \tag{7}$$

where m_k is the yearly mean and s_k is the yearly standard deviation. Both quantities are estimated over all locations $i \in \{1, \dots, D\}$ and using all data points corresponding to time $t \in \mathcal{Y}_k$. The transformation in (7) is applied to each $i \in \{1, \dots, D\}$. This intermediate step is then followed by the final GPD step, for which we introduce a high but constant 90% marginal threshold, similar to our approach for margins B.

3.4 | Margins D

Our final detrending scheme builds upon the methodology for margins C, replacing the constant high marginal threshold feature with the automated threshold selection routine of Murphy, Tawn, and Varty (2024), thus allowing for a different marginal threshold for each spatial location s_i , $i \in \{1, \dots, D\}$.

The Murphy, Tawn, and Varty (2024) method is aimed at addressing the bias-variance trade-off that is inherent to the

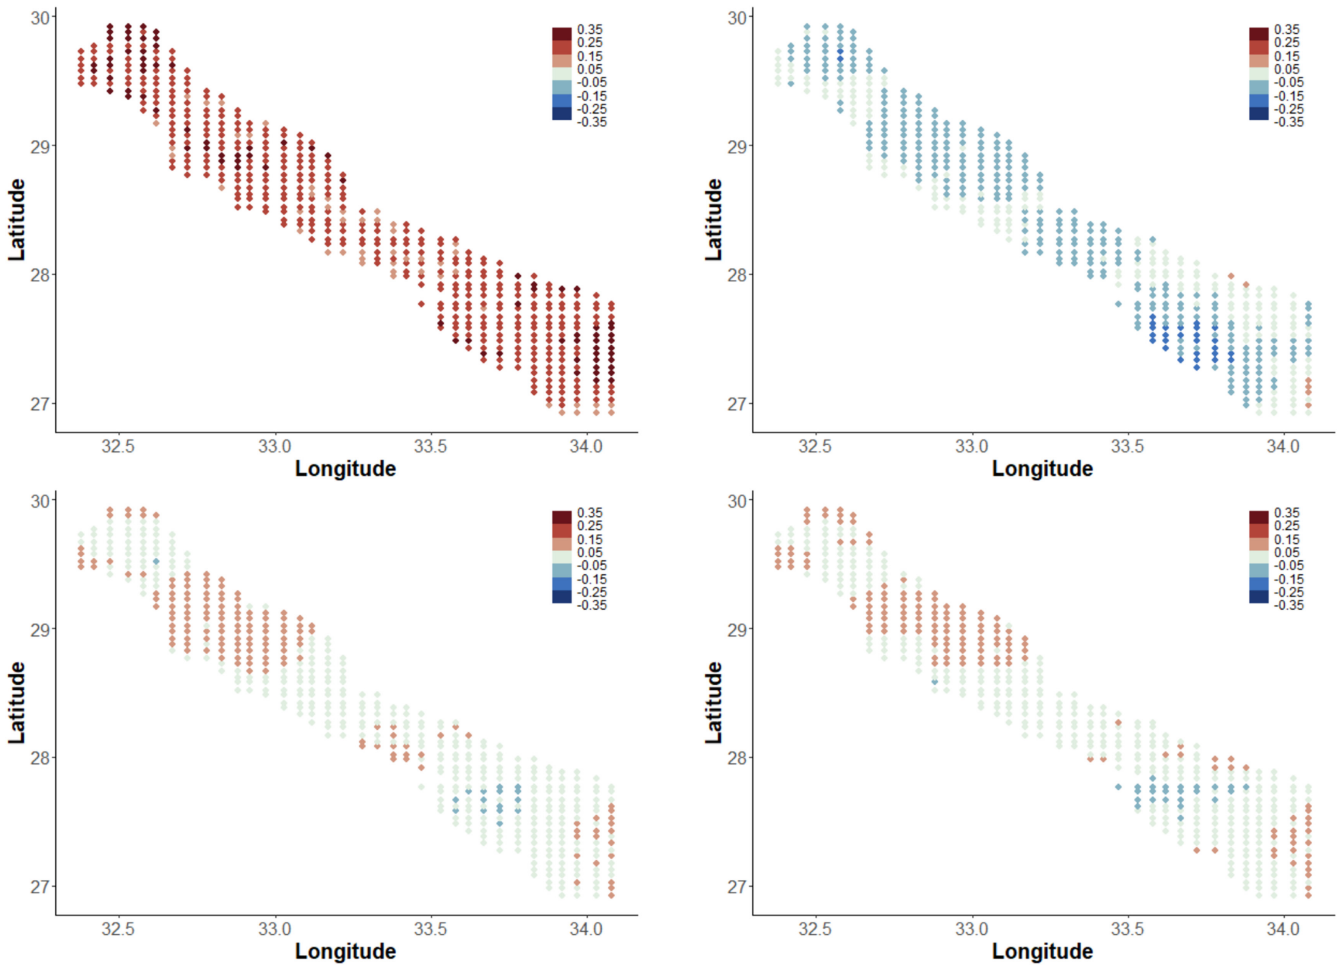


FIGURE 3 | Differences in $\tilde{\chi}_{0.95}(\mathbf{s}_k)$ between periods (1985–1989) and (2011–2015) for all spatial locations $\mathbf{s}_k, k \in \{1, \dots, D\}$. [Left to right and top to bottom] A–B–C–D margins.

problem of threshold choice for GPD model fitting; namely that too low a threshold is likely to violate the assumptions upon which asymptotic justification for the model is based—thus introducing bias in the GPD model fit—while too high a threshold can lead to fewer excesses being used to fit the GPD model—hence contributing to higher parameter uncertainty. The method focuses on selecting a constant threshold for independent and identically distributed univariate time series. Threshold selection is performed based on an algorithm that minimises the so-called *expected quantile discrepancy* between the sample quantiles and the fitted GPD model quantiles. For more details, we refer the interested reader to the original paper.

We use code provided as supplementary material to the Murphy, Tawn, and Varty (2024) paper to apply their method to the pre-whitened and yearly adjusted series in (7) used in margins C. A histogram of the resulting thresholds across all D locations is given in Appendix B. These thresholds are subsequently used for the implementation of the final GPD step, in accordance with all the aforementioned marginal procedures. We note that neither the independence nor the identical distribution assumptions of the Murphy, Tawn, and Varty (2024) method are likely to hold in our case. The former is violated by the presence of temporal dependence mentioned in Section 1, while in the case of the latter,

we allow for its possible violation, should the likelihood ratio test—performed as part of the GPD fitting step—deem a non-stationary fit appropriate. However, given the dimensionality of our data and the fact that few, if any, alternative automated approaches for appropriate threshold selection exist, we treat margins D and their subsequent analysis as a sensitivity experiment of the effect threshold choice may have on marginal features. Visualisation of the resulting time series in standard Laplace margins (bottom right plot of Figure 2) suggests potentially little effect of threshold sensitivity given the similarity of the series to the respective one for margins C (bottom left plot of Figure 2).

4 | Spatial Analysis of the Gulf of Suez Dataset

The Red Sea temperature dataset introduced in Section 1 has been analysed in numerous other studies including Huser and Genton (2016), Simpson and Wadsworth (2021) and Simpson, Opitz, and Wadsworth (2023). In agreement with those studies, we found in exploratory analysis that the whole dataset showcases non-trivial non-stationarity in its spatial extremal dependence across space, with the Gulf of Suez having a slightly different behaviour to the north part of the Red Sea (latitude $> 21.3^\circ$), which in turn is significantly different to the south part of the Red Sea (latitude $< 21.3^\circ$). Hence, the decision to focus on the Gulf of Suez

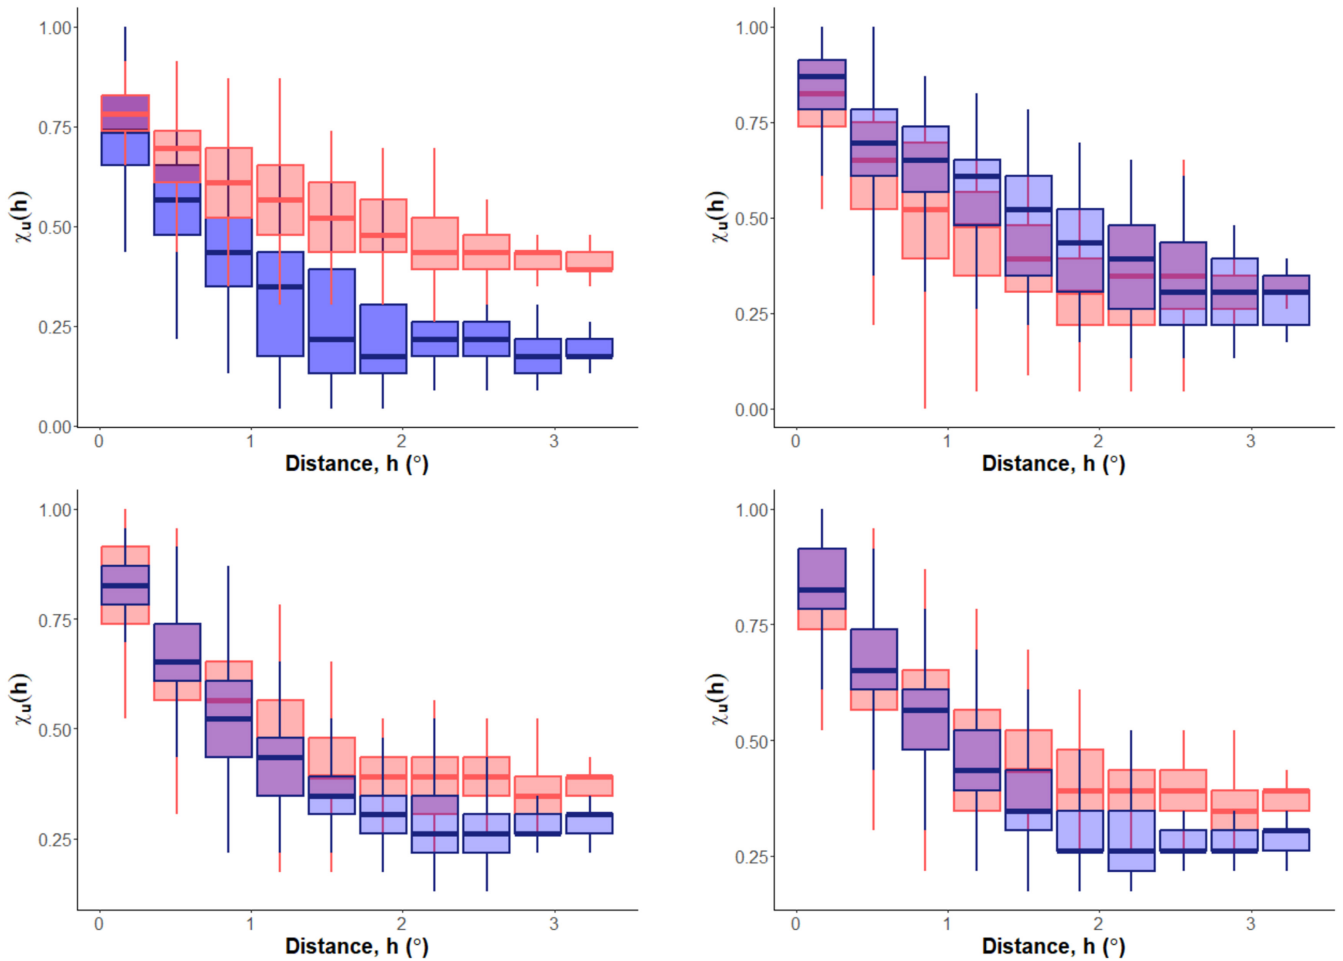


FIGURE 4 | Boxplots of $\chi_u(h_{jk})$ estimates grouped in 10 equidistant distance blocks. Boxplots in red are based on $\chi_u(h_{jk})$ values for the period 1985–1989, while boxplots in blue are based on $\chi_u(h_{jk})$ values for years 2011–2015. All estimates are calculated using $u = 0.95$. [Left to right and top to bottom] A–B–C–D margins.

where the assumption of stationarity in spatial extremal dependence over space seems reasonable. This smaller dataset now comprises a total of $D = 510$ spatial locations, for which marginal time series consisting of $T = 2852$ data points (92 per year) are available.

Having obtained marginally detrended versions of the dataset, as described in Sections 3.1–3.4, we perform exploratory investigations of the spatial extremal dependence structure. We focus on two key measures to assess changing spatial extremal dependence over time. The first measure examined is an average version of the χ_u and η_u measures defined in (3) and (5), respectively. To be precise, we define $\tilde{\chi}_u(\mathbf{s}_k) = \frac{1}{D-1} \sum_{j \neq k} \chi_u(h_{jk})$ and $\tilde{\eta}_u(\mathbf{s}_k) = \frac{1}{D-1} \sum_{j \neq k} \eta_u(h_{jk})$. Note that \mathbf{s}_k is the conditioning site from the definition of $\chi_u(h_{jk})$ in (3) and is therefore included as an argument in the definition of $\tilde{\chi}_u(\mathbf{s}_k)$ and $\tilde{\eta}_u(\mathbf{s}_k)$ to highlight that averaging is performed for each conditioning location separately. These quantities allow us to assess average changes in the extremal dependence of the entire spatial domain over time when computed over and contrasted against different time periods. For example, Figure 3 shows differences in $\tilde{\chi}_u(\mathbf{s}_k)$ between periods 1985–1989 and 2011–2015, that is,

$\tilde{\chi}_{u,(1985-1989)}(\mathbf{s}_k) - \tilde{\chi}_{u,(2011-2015)}(\mathbf{s}_k)$ where $\tilde{\chi}_{u,(A-B)}(\mathbf{s}_k)$ represents the measure $\tilde{\chi}_u(\mathbf{s}_k)$ calculated using data in the time period $A - B$. The quantity $\tilde{\chi}_{u,(1985-1989)}(\mathbf{s}_k) - \tilde{\chi}_{u,(2011-2015)}(\mathbf{s}_k)$ is calculated empirically for all spatial locations $\mathbf{s}_k, k = \{1, \dots, D\}$, using $u = 0.95$. A similar plot is obtained, respectively, for the quantity $\tilde{\eta}_{u,(1985-1989)}(\mathbf{s}_k) - \tilde{\eta}_{u,(2011-2015)}(\mathbf{s}_k)$ and is available in Figure C2 of Appendix C. The second measure we look at is pairwise $\chi_u(h_{jk})$ estimates from (3). We compute those empirically for the same two periods and threshold level ($u = 0.95$) and plot them against distance, thus assessing how the extremal dependence of our dataset changes with respect to distance and over time. These estimates are grouped in ten equidistant distance blocks and visualised via boxplots in Figure 4 for better clarity of representation.

It is important to note the considerably different messages Figures 3 and 4 convey concerning the change in dependence over time: When $\tilde{\chi}_{u,(1985-1989)} - \tilde{\chi}_{u,(2011-2015)}$ is computed using margins A (top left plot of Figure 3), a weakening in spatial extremal dependence is suggested, with $\tilde{\chi}_u$'s being smaller for the later period, 2011–2015, compared to the earlier period, 1985–1989, across all available locations. The opposite however is suggested by the plot corresponding to margins B (top right plot of Figure 3), with a sign reversal in $\tilde{\chi}_u$ suggesting

TABLE 1 | Table of CSE model coefficients.

Margins	$\hat{\kappa}$	$\hat{\lambda}$	$\hat{\beta}$	$\hat{\phi}_Z$	$\hat{\mu}_Z$	$\hat{\delta}$	$\hat{\nu}_Z$	$\hat{\sigma}_Z$
A	2.00	1.02	0.68	76.1	-8.51	1.03	1.34	5.21
B	0.901	3.23	1.000	85.2	-163	1.04	1.64	7.28
C	0.975	1.63	1.000	15.4	-7.81	0.829	1.64	1.84
D	0.968	2.26	1.000	15.9	-11.4	0.982	1.57	1.80

Note: Parameters κ and λ relate to the function $a_{s-s_0}(\cdot)$, parameter β relates to the function $b_{s-s_0}(\cdot)$ and parameters ϕ_Z , μ_Z , δ , ν_Z and σ_Z relate to the structure of the process Z^0 ; see Appendix A as well as Wadsworth and Tawn (2022) for more details.

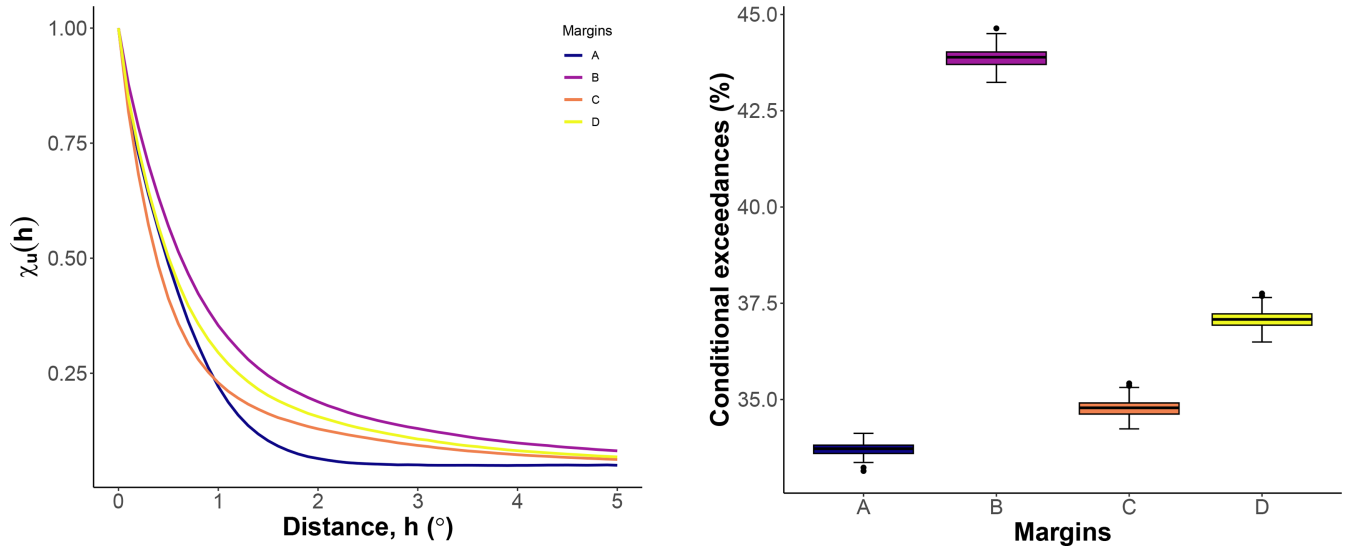


FIGURE 5 | [Left] Model $\chi_u(h)$ estimates for $u = 0.95$ plotted over distance. Lines correspond to mean $\chi_u(h)$ estimates. [Right] Boxplots of model simulated conditional exceedances. All estimates are obtained based on 200 simulations of 10^4 random fields each.

a strengthening of spatial extremal dependence over time. Finally, plots based on margins C and D (bottom left and right plots of Figure 3, respectively) suggest little to no change over time. The same conclusion holds when examining $\tilde{\eta}_{u,(1985-1989)} - \tilde{\eta}_{u,(2011-2015)}$ values (see Figure C2) or plots of χ_u over distance shown in Figure 4.

Given the non-negligible discrepancies between detrending procedures, we aim to quantify what repercussions these might have on a more practical level by fitting the same spatial model to the four datasets presented in Section 3. In doing so, we are not attempting to model potential non-stationarity in the spatial dependence, but rather to highlight the different conclusions obtained concerning spatial extremal dependence from the different marginal models. In practice, this feature is not usually considered in extremal analyses of spatial processes. We fit the CSE model of Wadsworth and Tawn (2022), introduced in Section 2.3 and detailed in Appendix A. Parameter estimation is achieved via a composite maximum likelihood scheme—information from all locations acting as the conditioning site was allowed to contribute to the inference procedure—using code provided as supplementary material from Wadsworth and Tawn (2022). The resulting estimates are presented in Table 1. A constant high threshold corresponding to the 95% quantile

of the standard Laplace distribution was used to obtain these results.

Combined parameter estimates from the CSE model are not immediately interpretable in terms of the strength of spatial dependence. We therefore assess this by simulating from the fitted models to obtain model-based estimates of $\chi_u(h)$, which we compare graphically in the left plot of Figure 5. Additionally, the right hand plot of Figure 5 provides percentages of conditional threshold exceedances obtained from our model simulations. That is, conditional on a single location s_0 (selected to lie approximately in the centre of the spatial domain) being extreme, the percentage of exceedances of the 95% quantile observed in the remaining locations was calculated. Both plots once again succeed in conveying the disagreement between the same spatial analysis of differently detrended datasets, with the weakest overall dependence for margins A and the strongest for margins B, and highlight the important influence marginal considerations can inflict on dependence features and vice versa. Note that these summaries do not include the effects of parameter uncertainty in the model fits, so may be interpreted with some caution. Uncertainty quantification could potentially be obtained by block bootstrap but is both computationally intensive and difficult to do appropriately given the likely non-stationary nature

of some of the margins. Therefore this is not performed for this investigative analysis.

5 | Discussion

The focus of this paper has primarily been on unpicking the margins-dependence interrelationship in the analysis of spatial extremes. By examining four different marginal detrending techniques, each with different merits and drawbacks, and comparing their effects on the dependence characteristics of the correspondingly detrended datasets we have established two key findings.

Our first message is that margins and dependence are intricately connected. Their interaction affects both apparent non-stationarity in spatial extremal dependence over time, as well as spatial extremal dependence estimates in model fits assuming temporal stationarity. This is highlighted both by visual exploratory tools such as the ones presented in Figures 3, C2 and 4 and model fits like the ones implemented in Section 4 and summarised in Table 1 and Figure 5. The high sensitivity of conclusions to the marginal model is something we should be aware of when performing similar analyses.

Our second finding is that the task of marginal detrending itself can be a very difficult undertaking in practice. One reason is that there is no single approach to marginal modelling. We strive for the most stationary data possible, but it is difficult to guarantee that this has been achieved to an acceptable level, especially with a large number of sites (in this case, > 500) to consider. Another reason is that, when attempting to analyse a dataset in practice, there is usually no way of knowing a priori the level of its underlying complexity and therefore making an informed decision on the best detrending approach to adopt. One has to rely on exploratory—principally graphical—tools that might not straightforwardly reveal the existence and extent of non-stationarity in the dataset and can be plagued by other complicating factors unrelated to non-stationarity, such as temporal dependence discussed in Section 1.

Author Contributions

Lydia Kakampakou: writing – original draft, formal analysis, methodology, software, visualization, validation, writing – review and editing. **Emma S. Simpson:** conceptualization, methodology, supervision, writing – review and editing. **Jennifer L. Wadsworth:** conceptualization, methodology, supervision, writing – review and editing.

Acknowledgements

We are grateful to the two anonymous reviewers and editors for constructive comments and suggestions that have improved this article. This paper is based on work completed while Lydia Kakampakou was part of EPSRC funded project EP/W524438/1. The Red Sea data were provided by GHRSSST, Met Office and CMEMS.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Sea surface temperature data can be downloaded at <https://marine.copernicus.eu>.

Endnotes

1Data form part of the HadCRUT5 dataset (Morice et al. 2021).

References

- Blanchet, J., and A. C. Davison. 2011. “Spatial Modeling of Extreme Snow Depth.” *The Annals of Applied Statistics* 5, no. 3: 1699–1725.
- Buishand, T. A., L. de Haan, and C. Zhou. 2008. “On Spatial Extremes: With Application to a Rainfall Problem.” *The Annals of Applied Statistics* 2, no. 2: 624–642.
- Castro-Camilo, D., M. de Carvalho, and J. L. Wadsworth. 2018. “Time-Varying Extreme Value Dependence With Application to Leading European Stock Markets.” *The Annals of Applied Statistics* 12, no. 1: 283–309.
- Chavez-Demoulin, V., and A. C. Davison. 2005. “Generalized Additive Modelling of Sample Extremes.” *Journal of the Royal Statistical Society. Series C: Applied Statistics* 54, no. 1: 207–222.
- Clarkson, D., E. F. Eastoe, and A. Leeson. 2023. “The Importance of Context in Extreme Value Analysis With Application to Extreme Temperatures in the U.S. and Greenland (With Discussion).” *Journal of the Royal Statistical Society. Series C: Applied Statistics* 72, no. 4: 829–843.
- Coles, S. G., and J. A. Tawn. 1991. “Modelling Extreme Multivariate Events.” *Journal of the Royal Statistical Society. Series B: Methodological* 53, no. 2: 377–392.
- D’Arcy, E., J. A. Tawn, and D. E. Sifnioti. 2022. “Accounting for Climate Change in Extreme Sea Level Estimation.” *Water* 14, no. 19: 2956.
- Davison, A. C., and R. L. Smith. 1990. “Models for Exceedances Over High Thresholds (With Discussion).” *Journal of the Royal Statistical Society. Series B: Methodological* 52, no. 3: 393–442.
- de Carvalho, M., and A. C. Davison. 2014. “Spectral Density Ratio Models for Multivariate Extremes.” *Journal of the American Statistical Association* 109, no. 506: 764–776.
- Dombry, C., and M. Ribatet. 2015. “Functional Regular Variations, Pareto Processes and Peaks Over Threshold.” *Statistics and Its Interface* 8, no. 1: 9–17.
- Donlon, C. J., M. Martin, J. D. Stark, J. Roberts-Jones, E. Fiedler, and W. Wimmer. 2012. “The Operational Sea Surface Temperature and Sea Ice Analysis (OSTIA) System.” *Remote Sensing of Environment* 116: 140–158.
- Eastoe, E. F. 2019. “Nonstationarity in Peaks-Over-Threshold River Flows: A Regional Random Effects Model.” *Environmetrics* 30, no. 5: e2560.
- Eastoe, E. F., and J. A. Tawn. 2009. “Modelling Non-Stationary Extremes With Application to Surface Level Ozone.” *Journal of the Royal Statistical Society. Series C: Applied Statistics* 58, no. 1: 25–45.
- Ferreira, A., and L. de Haan. 2014. “The Generalized Pareto Process; With a View Towards Application and Simulation.” *Bernoulli* 20, no. 4: 1717–1737.
- Healy, D., J. A. Tawn, P. Thorne, and A. Parnell. 2024. “Inference for Extreme Spatial Temperature Events in a Changing Climate With Application to Ireland.” *Journal of the Royal Statistical Society. Series C: Applied Statistics (to appear)* 2024: qlae047.
- Heffernan, J. E., and S. I. Resnick. 2007. “Limit Laws for Random Vectors With an Extreme Component.” *The Annals of Applied Probability* 17, no. 2: 537–571.

- Heffernan, J. E., and J. A. Tawn. 2004. "A Conditional Approach for Multivariate Extreme Values (With Discussion)." *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 66, no. 3: 497–546.
- Huser, R., and M. G. Genton. 2016. "Non-Stationary Dependence Structures for Spatial Extremes." *Journal of Agricultural, Biological, and Environmental Statistics* 21, no. 3: 470–491.
- Huser, R., and J. L. Wadsworth. 2022. "Advances in Statistical Modeling of Spatial Extremes." *Wiley Interdisciplinary Reviews: Computational Statistics* 14, no. 1: e1537.
- Jonathan, P., K. Ewans, and D. Randell. 2013. "Joint Modelling of Extreme Ocean Environments Incorporating Covariate Effects." *Coastal Engineering* 79: 22–31.
- Ledford, A. W., and J. A. Tawn. 1996. "Statistics for Near Independence in Multivariate Extreme Values." *Biometrika* 83, no. 1: 169–187.
- Lee, J., M. de Carvalho, A. Rua, and J. Avila. 2024. "Bayesian Smoothing for Time-Varying Extremal Dependence." *Journal of the Royal Statistical Society. Series C: Applied Statistics* 73: 581–597.
- Mhalla, L., M. de Carvalho, and V. Chavez-Demoulin. 2019. "Regression-Type Models for Extremal Dependence." *Scandinavian Journal of Statistics* 46, no. 4: 1141–1167.
- Mhalla, L., T. Opitz, and V. Chavez-Demoulin. 2019. "Exceedance-Based Nonlinear Regression of Tail Dependence." *Extremes* 22, no. 3: 523–552.
- Morice, C. P., J. J. Kennedy, N. A. Rayner, et al. 2021. "An Updated Assessment of Near-Surface Temperature Change From 1850: The HadCRUT5 Data Set." *Journal of Geophysical Research: Atmospheres* 126, no. 3: e2019JD032361.
- Murphy, C., J. A. Tawn, and Z. Varty. 2024. "Automated Threshold Selection and Associated Inference Uncertainty for Univariate Extremes." arXiv 2310.17999.
- Murphy-Barltrop, C. J. R., and J. L. Wadsworth. 2024. "Modelling Non-Stationarity in Asymptotically Independent Extremes." *Computational Statistics & Data Analysis* 199: 108025.
- Northrop, P. J., and P. Jonathan. 2011. "Threshold Modelling of Spatially Dependent Non-Stationary Extremes With Application to Hurricane-Induced Wave Heights (With Discussion)." *Environmetrics* 22, no. 7: 799–816.
- Padoan, S. A., M. Ribatet, and S. A. Sisson. 2010. "Likelihood-Based Inference for Max-Stable Processes." *Journal of the American Statistical Association* 105, no. 489: 263–277.
- Ribatet, M., D. Cooley, and A. C. Davison. 2012. "Bayesian Inference From Composite Likelihoods, With an Application to Spatial Extremes." *Statistica Sinica* 22, no. 2: 813–845.
- Richards, J., and J. L. Wadsworth. 2021. "Spatial Deformation for Nonstationary Extremal Dependence." *Environmetrics* 32, no. 5: e2671.
- Rootzén, H., J. Segers, and J. L. Wadsworth. 2018. "Multivariate Generalized Pareto Distributions: Parametrizations, Representations, and Properties." *Journal of Multivariate Analysis* 165: 117–131.
- Rootzén, H., and N. Tajvidi. 2006. "Multivariate Generalized Pareto Distributions." *Bernoulli* 12, no. 5: 917–930.
- Simpson, E. S., T. Opitz, and J. L. Wadsworth. 2023. "High-Dimensional Modeling of Spatial and Spatio-Temporal Conditional Extremes Using INLA and Gaussian Markov Random Fields." *Extremes* 26: 669–713.
- Simpson, E. S., and J. L. Wadsworth. 2021. "Conditional Modelling of Spatio-Temporal Extremes for Red Sea Surface Temperatures." *Spatial Statistics* 41: 100482.
- Wadsworth, J. L., and J. A. Tawn. 2022. "Higher-Dimensional Spatial Extremes via Single-Site Conditioning." *Spatial Statistics* 51: 100667.
- Winter, H. C., S. J. Brown, and J. A. Tawn. 2017. "Characterising the Changing Behaviour of Heatwaves With Climate Change." *Dynamics and Statistics of the Climate System* 1, no. 1: dzw006.
- Winter, H. C., J. A. Tawn, and S. J. Brown. 2016. "Modelling the Effect of the El Niño-Southern Oscillation on Extreme Spatial Temperature Events Over Australia." *The Annals of Applied Statistics* 10, no. 4: 2075–2101.
- Wood, S. N. 2003. "Thin-Plate Regression Splines." *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 65, no. 1: 95–114.
- 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, no. 1: 3–36.
- Wood, S. N. 2017. *Generalized Additive Models: An Introduction With R*. Chapman and Hall/CRC.
- 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, no. 528: 1865–1879.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.