Spatiotemporal modeling of hydrological return levels: A quantile regression approach

Spatio-temporal modelling of hydrological return levels. A quantile regression approach

Maria Franco-Villoria[1]*, Marian Scott[2] and Trevor Hoey[3]

[1]Department of Economics and Statistics "Cognetti de Martiis", University of Torino, Italy

[2]School of Mathematics and Statistics, University of Glasgow, UK

[3]School of Geographical and Earth Sciences, University of Glasgow, UK

Abstract

Extreme river flows can lead to inundation of floodplains, with consequent impacts for society, the environment and the economy. Extreme flows are inherently difficult to model, being infrequent, irregularly spaced and affected by non-stationary climatic controls. To identify patterns in extreme flows a quantile regression approach can be used. This paper introduces a new framework for spatio-temporal quantile regression modelling, where the regression model is built as an additive model that includes smooth functions of time and space, as well as space-time interaction effects. The model exploits the flexibility that P-splines offer and can be easily extended to incorporate potential covariates. We propose to estimate model parameters using a penalized least squares regression approach as an alternative to linear programming methods, classically used in quantile parameter estimation. The model is illustrated on a data set of flows in 98 rivers across Scotland.

Keywords: P-splines; PIRLS; hydrometric time series; extreme values

*Corresponding author. Email: maria.francovilloria@unito.it

1 Introduction

The occurrence of extreme events in the environment, i.e. those that deviate considerably from expected average levels, for example extreme temperatures, river flows, wave heights, or pollutant concentrations, has received increasing interest over the last decade. The last report by the Intergovernmental Panel on Climate Change (IPCC) recognizes the effect that climate change may have on the "frequency, intensity, spatial extent, duration, and timing of extreme events" (IPCC, 2014). Understanding the spatial and temporal structure of extremes is essential for planning purposes. An example of the severe consequences of extreme events, including loss of life, are the floods that hit the UK, especially Northern England and Scotland, in December 2015. These events prompted the National Flood Resilience Review (HM Government, 2016), a "review of how we assess flood risk, reduce the likelihood of flooding, and make the country as resilient as possible to flooding" and the latest in a series of reviews following damaging events in the UK and internationally (Evans et al., 2004; Pitt, 2008; Georgi et al., 2016).

It is now recognized that statistical methods specifically developed to analyze extreme values (over time and/or space) are needed. One common approach for dealing with spatial data is the use of geostatistical models (Diggle and Ribeiro Jr., 2007) that assume the data to be a realization of an underlying spatial Gaussian random field; these kinds of models have also been extended to the spatio-temporal case (see, e.g. Cressie and Wikle (2011)). However, the Gaussian assumption is not realistic when modelling extreme values, whose distribution is known to be skewed (Coles, 2004). The increasing interest in extreme values, especially in environmental applications, has led to the development of new statistical models specifically designed for spatial (and spatio-temporal) extremes. These include latent variables (Davison et al., 2012; Cooley et al., 2007), copula models (Fuentes et al., 2012) and the more recent max-stable processes (Davison et al., 2012; Davison and Gholamrezaee, 2012). A max-stable process can be thought of as the limiting distribution of the pointwise maxima of independent copies of a process. Max-stable processes are best characterized via their spectral representation, for which several models have been

proposed. In particular, the Smith (Smith, 1990), Schlather (Schlather, 2002) and Brown-Resnick (Brown and Resnick, 1977) models are widely used, where a composite marginal likelihood, built using the pairwise marginal distributions, is used for parameter estimation (Davison et al., 2012; Davison and Gholamrezaee, 2012; Wadsworth and Tawn, 2012). Good reviews of these models can be found in Davison et al. (2012) and Davison and Gholamrezaee (2012). More recently, several authors have considered the extension of max-stable processes to model extremes over space and time (Davis et al., 2013; Embrechts et al., 2016). Max-stable processes are based on asymptotic extreme value theory and hence can be useful when the aim is to estimate very extreme events. When interest lies in less extreme values, a further approach is that of quantile regression (Koenker, 2005). An extreme event in this case is characterized as a value falling in the upper (or lower, e.g. to model drought conditions) tail of the distribution. Quantile regression allows estimation of the relationship between response and explanatory variables at any percentile of the distribution of the response (conditioned on the explanatory variables). As a result, rates of change in the response variable can be estimated for the whole distribution and not only at the mean. The regression coefficients are estimated minimizing an objective function that is defined in terms of the sum of weighted absolute residuals. Quantile regression has been mostly developed in the case of independent observations, but a number of approaches for spatial and spatio-temporal quantile regression can be found in the literature. The paper by He et al. (1998) is, to our knowledge, the first one to extend the 1-dimensional quantile regression problem into a 2-dimensional context; the objective function can be re-written in terms of bivariate smoothing splines (Koenker, 2005; He et al., 1998) and then minimized using linear programming methods. Penalization and fitting of splines coefficients becomes considerably more complex. This initial paper on bivariate quantile regression was followed by the work of Hallin et al. (2009), who propose using a local linear regression approach in which the regression coefficients are allowed to vary spatially. Quantile regression has also been considered from a Bayesian perspective; Lee and Neocleous (2010) and Neelon et al. (2015) develop a quantile regression model for spatial and spatio-temporal count data respectively, incorporating a spatial

autoregressive term in the predictor to deal with spatial correlation. Reich et al. (2011) and Reich (2012) propose, respectively, spatial and spatio-temporal quantile regression models, in which the spatial structure is introduced via the covariance function of the Gaussian spatial processes associated with the parameters of the model; regarding the temporal structure, the quantile function at each spatial location is defined as a linear function of time within a hierarchical model. Any residual spatial correlation is then modelled via a spatial copula. More recently, Sun et al. (2016) introduce temporal and spatial dependence using a fused adaptive Lasso penalty.

In this paper, we propose a spatio-temporal quantile regression model for river flow data; by re-expressing the quantile model as a weighted linear regression, parameter estimates can be obtained using penalized iterative re-weighted least squares (PIRLS). The paper introduces a new framework for spatio-temporal quantile regression modelling, where the regression model is built as an additive model that includes smooth functions of time and space, as well as space-time interaction effects. While inclusion of bivariate smooth functions in a quantile regression setting has been considered before (He et al., 1998), the case of higher dimensional smooth functions, needed, for example, for the space-time interaction term, has not been addressed in the literature. The model exploits the flexibility that P-splines offer and can be easily extended to incorporate potential covariates. We propose to estimate model parameters using a penalized least squares regression approach as an alternative to linear programming methods classically used in quantile parameter estimation. The fitting procedure is simple and computationally efficient and allows modelling strategies already available for mean regression (e.g. varying coefficient models) to be adapted to the case of quantile regression. This presents a clear advantage over linear programming methods given the increasing complexity and availability of data and the interest in extreme events. By considering a fully spatio-temporal model rather than modelling one river a time, information is borrowed across rivers; this means a more efficient use of the available data, fundamental when dealing with short records and/or aiming to estimate very high (low) quantiles.

In particular, quantile regression can be a useful modelling strategy for extreme river flow

values, given the direct link between quantile estimates and return levels, which, in turn, are used in risk assessment. In the extreme value theory (Coles, 2004) the quantile z_p of a generalized extreme value distribution G is defined as the return level with return period $1/(1-p)$; the latter, also known as recurrence interval, is defined as the amount of time (on average) until the value z_p is likely to be equalled or exceeded; i.e. the long term average of the time intervals between successive exceedances of a peak of magnitude z_p (Coles, 2004). The flexible spatio-temporal regression model introduced in Section 3 can be seen as a tool for modelling return levels under non-stationary conditions. Such a tool is valuable as environmental variables often exhibit non-stationarity; for example, ongoing climate change has been shown to affect both the means and extreme values of climatic and river flow time series (Arnell and Gosling, 2013, 2016). By fixing the quantile to be estimated, we are fixing the probability of the estimated value being exceeded and hence the return period, to then estimate the associated return level, that is allowed to vary both in time and space. The proposed model is illustrated on a large set of Scottish rivers.

In Scotland, the Scottish Flood Risk Management Act (The Scottish Government, 2010) was introduced in 2010 following the requirements of the European Union Directive on Flood Risk Management (2007/60/EC). This new piece of legislation was partly motivated by changes in Scotland's river flow and rainfall regimes. Evidence from a number of published studies (Black, 1996; Black and Burns, 2002; Werritty, 2002) reports increased variability and statistically significant changes in annual peak-over-threshold magnitude and frequency of events and in annual maxima trends in Scotland during 1956-1995. In particular, dry (1960s-1970s) and wet periods (late 1980s-early 1990s) have been identified. However, no general trend seems to hold across the country as observed changes are not homogeneous, neither in frequency, with estimates of changes in return periods changing with location, nor in time, with seasonal changes in extreme rainfall being more pronounced than annual ones, especially in autumn (Fowler and Wilby, 2010). Climate model predictions (Fowler and Kilsby, 2003; Fowler and Wilby, 2010; Fowler and Ekström, 2009) suggest that these differences are likely to continue and/or increase in the near future, with fairly reliable estimates over the winter months but greater uncertainty over

the summer. Despite these concerns, there seems to be a research gap in the recent literature on Scottish hydrology; the most complete paper on Scottish rivers dates back to 1997 (Black and Werritty, 1997), while most UK based studies are limited to England and Wales (Hannaford and Buys, 2012) and are focused on individual catchments. The work done by Prosdocimi et al. (2013) includes a relatively large number of Scottish gauging stations, but analysis is limited to one gauging station at a time and autumn and spring are disregarded, even though significant rainfall increases in western Scotland have been identified in those seasons (Jenkins et al., 2009; Werritty, 2002). The study by Hannaford and Buys (2010) includes twenty Scottish gauging stations, but the analysis is performed on the UK as a whole. Results from some of the previously mentioned studies suggest regional variation in hydrological trends across the UK as a whole and also within Scotland. Although the focus of this paper is development of a spatio-temporal modelling methodology, the results presented here also update the literature and extend understanding of extreme river flows in Scotland. Rather than working on one gauging station at a time, we follow a spatio-temporal modelling approach that takes into account possible dependencies among stations. Further, important changes in river flow may not be detected at the annual scale (Hannaford and Buys, 2012); we use daily data, without restricting the series to annual maxima or peak-over-threshold data, the usual practice in extreme value analysis. The dataset considered is very rich, with nearly 100 gauging stations, and the results show how the spatio-temporal trend in extreme flows has changed in recent years (1996-2013), helping to fill in the current research gap. Identifying these changes is important to understand the effect of climate change (Prosdocimi et al., 2013) and to investigate the validity of model projections or for historical model run validation (Hannaford and Buys, 2012).

The paper is organized as follows. The dataset is introduced in Section 2. Section 3 describes the proposed methodology for fitting a spatio-temporal quantile regression model. The performance of the model is illustrated in Section 4 on a set of Scottish rivers, for which the 95% quantile of river flow is estimated. A simulation study is presented in Section 5. Finally, the main results and discussion are summarized in Section 6.

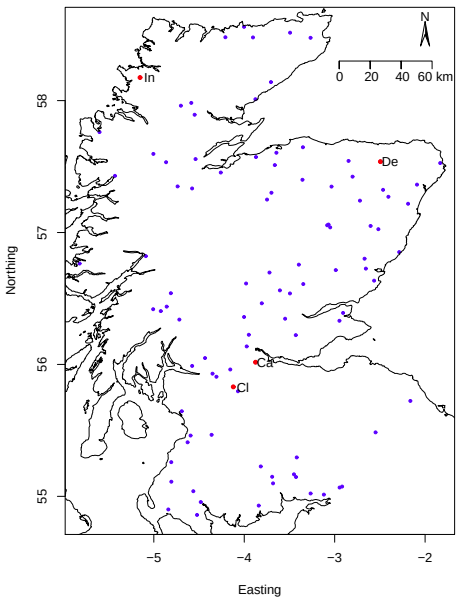

Figure 1: Location of the 98 selected gauging stations. In red, location of rivers Inver (In), Clyde (Cl), Carron (Ca) and Deveron (De).

2 Data

Data, in the form of daily river flow (m^3/s), were provided by the Scottish Environment Protection Agency (SEPA) and the National River Flow Archive (NRFA). Most of the records are available online on the NRFA webpage. Our data set consists of 98 gauging stations, selected on the basis of geographical location, quality and length of the records, covering the period 1^{st} January 1996- 31^{st} December 2013. The spatial locations can be seen in Figure 1. Forty-three series contained missing values. However, the missing proportions ($< 0.1\%$) were small enough not to be a concern, and missing values were imputed using linear interpolation. The interpolation was done separately for each month to better reproduce the variability of the series; i.e. missing values in January were imputed using only recorded values in January, and so on. Since the distribution of river flow is very skewed, a log transformation was used.

In particular, four rivers have been chosen for illustrative purposes and are discussed in more detail in Section 4: the river Inver (North West), Clyde (South West), Carron (South East) and Deveron (North East). The gauging stations are marked in red in Figure 1 while data are

shown in Figure 2. The main characteristics of these rivers (catchment area, maximum elevation, mean flow, 95th (Q95) and 99th (Q99) quantiles of river flow and mean flow/catchment area) are summarized in Table 1. Rivers Clyde and Deveron have higher flow values on average, as they are the largest rivers out of the four rivers considered, with a catchment area of 1903.1 km^2 and 954.9 km^2 respectively. Rivers Inver and Carron, on the other hand, have a catchment area of 137.5 km^2 and 122.3 km^2 respectively. All four rivers exhibit a clear strong seasonal pattern, with greater variability in rivers Clyde and Carron, which are located on the southern part of the country. While there is no apparent long term trend, Figure 2 shows extreme flows in all four rivers, but these are not all coincident. This suggests that a simple statistical model with only a time trend and a spatial trend will probably not be enough to capture the complexity of the data, but inclusion of a time-space interaction may be required. Spatial differences can be partly explained by the predominant rainfall pattern in Scotland, wetter in the West and dryer in the East, as seen in Figure 3.

Prior to model fitting, the mean flow was removed at each individual location as a way of standardizing the data. This was done to account for differences in flow values due to catchment size.

3 Methodology

We introduce a new approach that builds upon the idea of approximating the absolute residuals with the squared residuals, as suggested in Reiss and Huang (2012). This approximation, motivated by the fact that the check function in Equation (1) is not differentiable at zero, ensures differentiability everywhere. This way, instead of using linear programming methods to estimate the model parameters, a weighted least squares approach is preferred, exploiting the fact that the objective function in Equation (1) is a weighted sum of absolute residuals. In their paper, Reiss and Huang (2012) consider a very simple model where the response depends on a single covariate, while we introduce a flexible spatio-temporal model in a generalized additive model

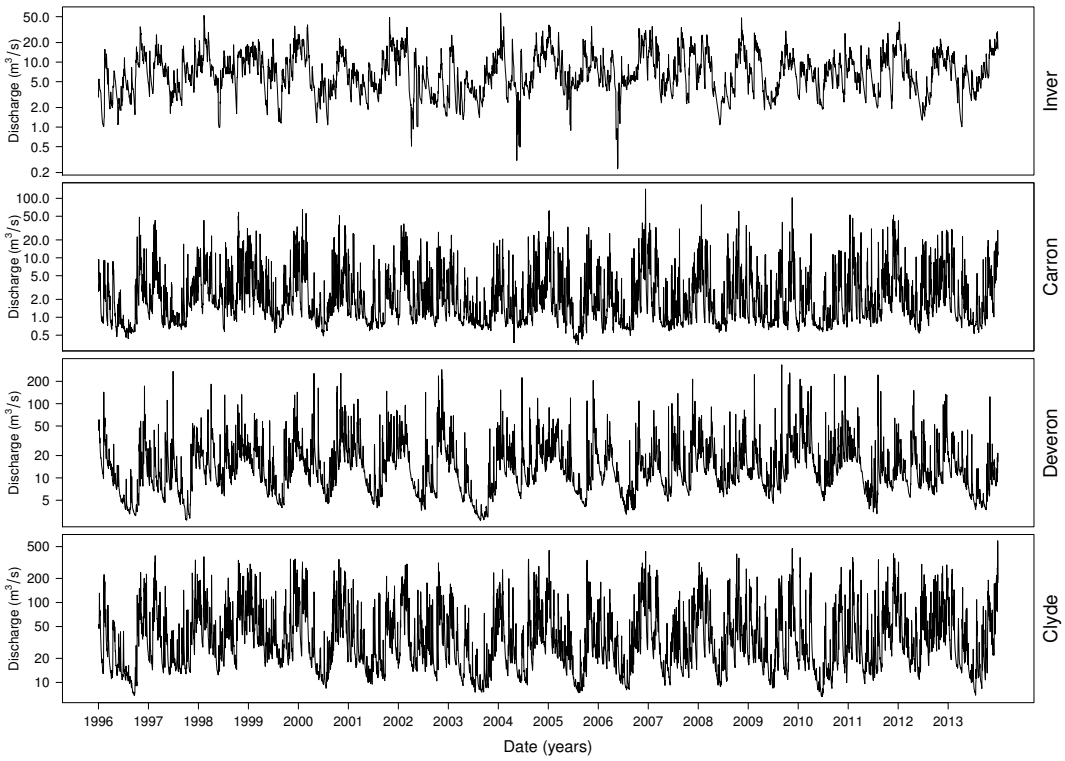

Figure 2: Time series of flow (m^3/s) ordered by flow magnitude for the rivers Inver, Carron, Deveron and Clyde.

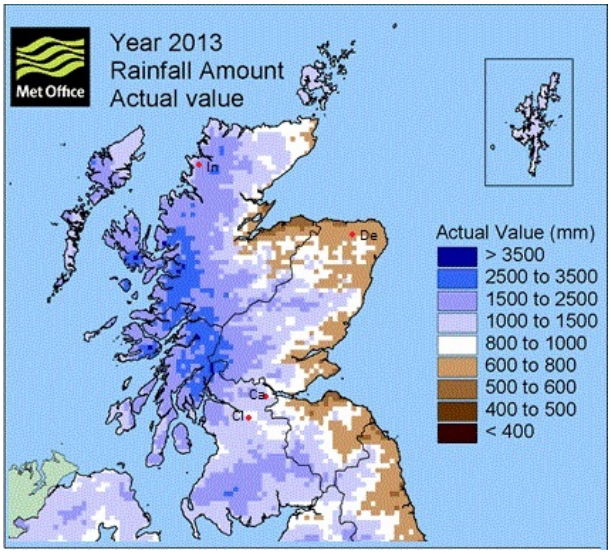

Figure 3: Average annual rainfall in 2013 (*Source:* modified from *https://www.metoffice.gov.uk/climate/uk/summaries*). In red, location of rivers Inver (In), Clyde (Cl), Carron (Ca) and Deveron (De).

framework.

The classic simple linear regression model of y over x aims to estimate $E[Y|X = x] = \alpha + \beta x$. Instead, the aim now is to estimate a quantile of the distribution of $Y|X = x$ rather than the mean. The conditional quantile of a random variable Y with cumulative distribution function F_Y can be expressed as $Q_Y(\tau|\boldsymbol{X} = \boldsymbol{x}) = F_Y^{-1}(\tau|\boldsymbol{X} = \boldsymbol{x})$, where $\tau \in (0,1)$, and $\boldsymbol{X} = (X_1, \ldots, X_p)$ is a vector of explanatory variables (Koenker, 2005; Cade and Noon, 2003). As in classic linear regression, the errors are assumed to be independent, but here no assumptions are made regarding their distribution. The general objective function in a quantile linear regression model is defined as:

$$R(\boldsymbol{\beta}) = \sum_{i=1}^{n} \rho_\tau \left(y_i - (x_{1,i}\beta_1 + \ldots + x_{p,i}\beta_p) \right), \tag{1}$$

where $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)^T$, $\rho_\tau(u) = u(\tau - I(u < 0))$ is the check function and I is an indicator function (Koenker, 2005; Koenker and Hallock, 2001; Koenker and Bassett, 1978). Estimating the parameters $\boldsymbol{\beta}$ in Equation (1) involves minimizing a sum of weighted absolute deviations, where the weights are asymmetric functions of τ. The function $R(\boldsymbol{\beta})$ is piecewise linear and continuous, being differentiable at every point except at those whose residuals are zero. Until now, linear programming methods, like the simplex method or the interior point method (Koenker, 2005; Koenker and Hallock, 2001) are used to estimate $\boldsymbol{\beta}$.

3.1 A spatio-temporal quantile regression model

A simple spatio-temporal additive model (main effects model) for river flow can be expressed as:

$$Q_{\log(flow)_i}(\tau|t_i, d_i, z_i) = s_1(t_i) + s_2(d_i) + s_3(z_i), \quad i = 1, \ldots, n \tag{2}$$

where $Q_{log(flow)}(\tau|t, d, z)$ is the τ^{th} quantile of the (conditional) distribution of log(flow), $s_1(t)$, $s_2(d)$ are smooth functions of time and day of the year and $s_3(z)$ is a bivariate smooth function of easting and northing coordinates. These three terms represent the temporal, seasonal and spatial trends in river flow respectively. In our particular application, the trend ($s_1(t)$), seasonal

($s_2(d)$) and spatial ($s_3(z)$) terms were built as smooth functions of t=time (1996 to 2013), d=day within the year (1 to 365) and z=(easting, northing), respectively. The model assumes the seasonal component $s_2(d)$ to be constant over the years and the temporal trend to be the same at all spatial locations (or similarly, the spatial trend is assumed to be constant over time). A preliminary exploratory analysis of the data revealed a seasonal cycle that changes from year to year, suggesting that this might be too simple a model for the Scottish river flow dataset. We introduce an interaction term $s_4(t, d)$ to adjust for yearly changes in the seasonal pattern:

$$Q_{\log(flow)_i}(\tau|t_i, d_i, z_i) = s_1(t_i) + s_2(d_i) + s_3(z_i) + s_4(t_i, d_i). \tag{3}$$

Similarly, Model (3) can be further extended to include space-time $s_5(t, z)$ and space-season $s_6(d, z)$ interactions so that the full model becomes:

$$Q_{\log(flow)_i}(\tau|t_i, d_i, z_i) = s_1(t_i) + s_2(d_i) + s_3(z_i) + s_4(t_i, d_i) + s_5(t_i, z_i) + s_6(d_i, z_i). \tag{4}$$

Each of the univariate smooth functions can be rewritten as a linear combination of k cubic B-spline basis functions $B_1(t), \ldots, B_k(t)$ (Eilers and Marx, 2009, 2010), so that $s_1(t) = \boldsymbol{B_1\theta_1}$ and $s_2(d) = \boldsymbol{B_2\theta_2}$, where $\boldsymbol{\theta_1}$, $\boldsymbol{\theta_2}$ are $k_1 \times 1$, $k_2 \times 1$ vectors of coefficients respectively, $\boldsymbol{B_1}$, $\boldsymbol{B_2}$ are the matrices of basis functions ($\dim(\boldsymbol{B_1}) = n \times k_1$, $\dim(\boldsymbol{B_2}) = n \times k_2$), k_1, k_2 are the number of basis functions used in each case and n is the total number of observations.

The bivariate smooth function $s_3(z) = s_3(easting, northing)$ can be expressed in terms of the tensor product of the marginal B-splines basis on the individual variables *easting* and *northing* (Wood, 2006; Eilers and Marx, 2003, 2010; He et al., 1998) so that $s_3(z) = \boldsymbol{B_3\theta_3}$ where $\boldsymbol{\theta_3}$ is a $(k_{east} \times k_{north}) \times 1$ vector of coefficients and $\boldsymbol{B_3}$ is a $n \times (k_{east} \times k_{north})$ matrix. The interaction terms $s_4(t, d)$, $s_5(t, z)$ and $s_6(d, z)$ can be built in a similar way in terms of the tensor product of the corresponding marginal basis matrices as detailed above.

The full model can be expressed in matrix form as:

$$\boldsymbol{y} = \boldsymbol{B\theta} + \boldsymbol{\epsilon}$$

where the design matrix $\boldsymbol{B} = [\boldsymbol{B_1} \quad \boldsymbol{B_2} \quad \boldsymbol{B_3} \quad \boldsymbol{B_4} \quad \boldsymbol{B_5} \quad \boldsymbol{B_6}]$ is the matrix that results from combining the individual matrices column-wise and $\boldsymbol{\theta} = [\boldsymbol{\theta_1} \quad \boldsymbol{\theta_2} \quad \boldsymbol{\theta_3} \quad \boldsymbol{\theta_4} \quad \boldsymbol{\theta_5} \quad \boldsymbol{\theta_6}]^T$ is the vector of coefficients. By expressing the model as a linear model, the parameters can be estimated easily using efficient matrix-vector operations. A penalty term can be added to control for the amount of smoothness, which can be tuned individually for each term in the model. This means that a vector of smoothing parameters $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \boldsymbol{\lambda_3}, \boldsymbol{\lambda_4}, \boldsymbol{\lambda_5}, \boldsymbol{\lambda_6})$ needs to be specified. In this case, a second order different penalty on the spline coefficients is used following Eilers and Marx (2009); further, a periodicity constraint can be imposed on the seasonal term $s_2(d)$ to ensure continuity of the seasonal cycle over the years. This can be implemented, for cubic B-splines, by forcing the first and last three spline coefficients to be the same. The penalty term can be expressed as $\boldsymbol{\theta^T P \theta}$ where $\boldsymbol{P}$ is a block diagonal penalty matrix, each block corresponding to the penalization of the individual smooth terms included in the model. For the univariate terms $s_1(t)$, $s_2(d)$ the corresponding penalty matrix can be easily built as $\boldsymbol{D_{do}^T D_{do}}$ following Eilers and Marx (2009), where $\boldsymbol{D_{do}}$ is a difference matrix of order do ($do = 2$ for a second order penalty) of dimension $(k-2) \times k$ with k the number of basis functions.

For the spatial term $s_3(east, north)$, the penalty is constructed by penalizing individually the rows and columns of matrix $\boldsymbol{B}_3$ (Wood, 2006; Eilers and Marx, 2003, 2010; He et al., 1998). Penalty terms for $s_4(t, d)$, $s_5(t, z)$ and $s_6(d, z)$ can be built similarly.

The vector of model parameters $\boldsymbol{\theta}$ is estimated using the penalized iterative weighted regression approach described below. Assuming the vector of smoothing parameters $\boldsymbol{\lambda}$, to be fixed,

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \epsilon R^k}{\operatorname{argmin}} \left[\sum_{i=1}^{n} \rho_\tau(y_i - \boldsymbol{B_i \theta}) + \boldsymbol{\lambda \theta^T P \theta} \right] \tag{5}$$

where k is the total number of coefficients and $\boldsymbol{B}_i$ represents the i^{th} row of matrix $\boldsymbol{B}$. We

propose translating the minimization problem in Equation (5) into a penalized least squares problem that can be solved using penalized iterative reweighted least squares (PIRLS, Wood (2006)). The (approximate) objective function expressed in matrix form becomes:

$$||\boldsymbol{W}(\boldsymbol{y} - \boldsymbol{B\theta})||^2 + \boldsymbol{\lambda\theta^T P\theta} \tag{6}$$

where $\boldsymbol{W}$ is a diagonal matrix of weights calculated iteratively following Equation (7):

$$w_i^{(j)} = \frac{\tau - I\left[(y_i - \boldsymbol{B_i}\hat{\boldsymbol{\theta}}^{(\boldsymbol{j})}) < 0\right]}{2(y_i - \boldsymbol{B_i}\hat{\boldsymbol{\theta}}^{(\boldsymbol{j})})} \tag{7}$$

for $i = 1, \ldots,$n. A large upper bound is set for the weights to avoid residuals close to zero, which would result in the check function not being differentiable. Truncating all weight values above the given upper bound does not have a large effect on the fitting process as we are just forcing very small residuals, whose contribution to the objective function is negligible, to be even smaller. The estimated vector of parameters $\hat{\boldsymbol{\theta}}$ at iteration (j) can be computed as:

$$\hat{\boldsymbol{\theta}}^{(\boldsymbol{j})} = (\boldsymbol{B^T W}^{(\boldsymbol{j-1})}\boldsymbol{B} + \boldsymbol{\lambda P})^{-\mathbf{1}}\boldsymbol{B^T W}^{(\boldsymbol{j-1})}\boldsymbol{y} \tag{8}$$

Convergence of the algorithm is defined based on the objective function $R(\boldsymbol{\theta})$ defined in Equation (1); the algorithm stops when the difference between $R(\boldsymbol{\theta}^{(j-1)})$ and $R(\boldsymbol{\theta}^{(j)})$ is smaller than some predefined small tolerance. Results from a simulation study suggest no differences in the fitted model for tolerance values of 10^{-2} and above. Identifiability of the single components of the model is ensured by including a ridge penalty (Eilers and Marx, 2002).

The smoothing parameters are chosen to minimize a modified version of the Schwarz information criterion (SIC, Koenker et al. (1994)):

$$SIC(\boldsymbol{\lambda}) = \log\left[\frac{1}{n}\sum_{i=1}^{n}\rho_\tau(y_i - \hat{y}_i)\right] + \frac{1}{2n}df_\lambda \log n \tag{9}$$

where the approximated degrees of freedom df_λ can be calculated as the trace of the smoothing matrix $\boldsymbol{S} = \boldsymbol{B}(\boldsymbol{B^T W B} + \boldsymbol{\lambda P})^{-1}\boldsymbol{B^T W}$ (Hastie and Tibshirani, 1990). As previously stated, Model (4) is fitted assuming independent observations. In this case, standard errors for the fitted values $\hat{y}$ can be obtained as:

$$se(\hat{\boldsymbol{y}}) = \sqrt{diag(\boldsymbol{S S^T})\hat{\sigma}^2}$$

where $\boldsymbol{S}$ is the smoothing matrix at the last iteration and $\hat{\sigma}^2$ is an estimate of the residual variance that can be obtained as $\hat{\sigma}^2 = \frac{RSS}{df_{error}}$, with $RSS = \boldsymbol{Y^T(I-S)W(I-S)Y}$ and $df_{error} = n - df_\lambda$. However, if some spatio-temporal dependence structure is left in the residual term, standard error calculation needs to be adjusted accordingly. Let $\boldsymbol{V}$ be the correlation matrix. Adjusted standard errors can be estimated as:

$$se(\hat{\boldsymbol{y}}) = \sqrt{diag(\boldsymbol{S V S^T})\hat{\sigma}^2}. \tag{10}$$

4 Results

Model (4) was fitted to the data set described in Section 2, comprising 98 rivers with 6570 daily observations per river spanning 18 years, from 1^{st} January 1996 to 31^{st} December 2013 ($n = 643860$ observations in total). The 29th of February was removed from the dataset after ensuring that no relevant information (i.e. no extreme values) was lost. Since we are mostly interested in extreme values, a value of τ=0.95 was chosen to fit a model for the 95^{th} quantile of logged river flow. The trend ($s_1(t)$), seasonal ($s_2(d)$) and spatial ($s_3(z)$) terms in Model (4) were built as smooth functions of t=time (1996 to 2013), d=day within the year (1 to 365) and z=(easting, northing), respectively. When using penalized splines, the usual practice is to choose an arbitrary large number of basis functions and then control the amount of smoothness by penalizing the spline coefficients. The definition of *large*, however, depends on the application at hand. We fitted the model using an increasing number of basis functions until the percentage change in SIC was smaller than 1%. In this case, the number of basis functions was chosen to

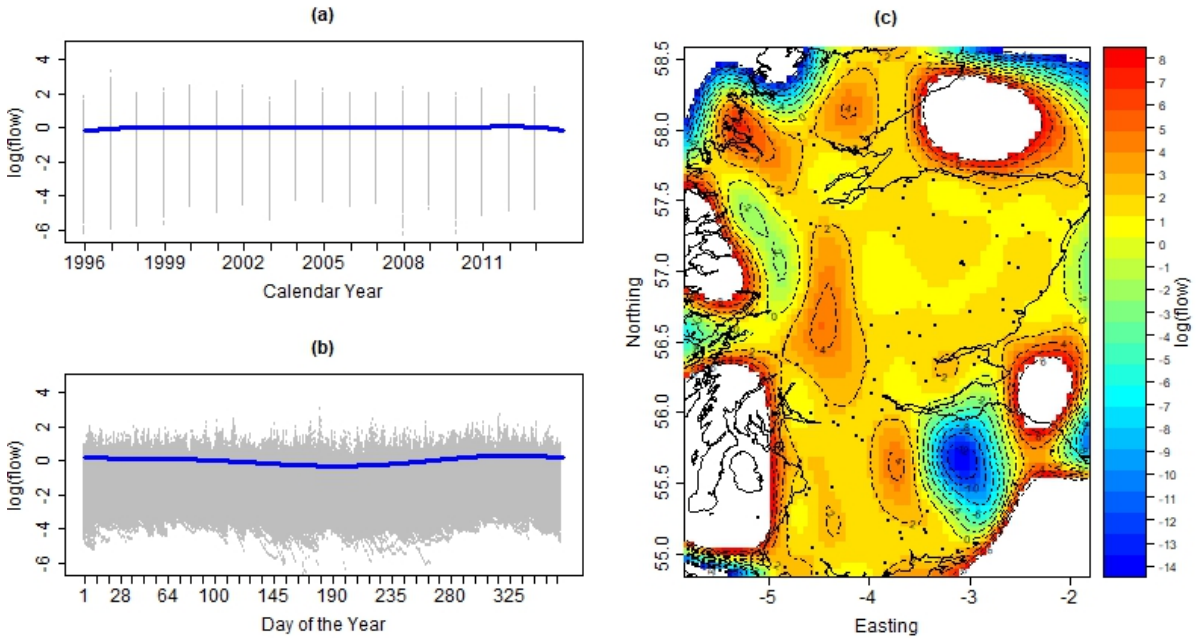

Figure 4: Estimated trend (a), seasonal (b) and spatial (c) components along with partial residuals.

be $k_1 = 12$ for the trend component, $k_2 = 6$ for the seasonal component and $k_3 = 12^2$ for the space component, while $k_4 = 12 \times 6$, $k_5 = 12^3$ and $k_6 = 6 \times 12^2$ were used for the interaction terms $s_4(t, d)$, $s_5(t, z)$ and $s_6(d, z)$ respectively. A second order penalty was imposed on the spline coefficients. This is a commonly used smoothness penalty that corresponds to penalizing the roughness of a curve, measured as the integral square of the second derivative of the curve (Eilers and Marx, 1996). Smoothing parameters were chosen based on the Schwarz Information Criteria (Equation (9)), using a restricted grid of values for $\boldsymbol{\lambda}$. The trend, seasonal and spatial main effects are shown in Figure 4.

Overall, the estimated trend (Figure 4 (a)) appears to be fairly flat. There is a seasonal effect (Figure 4 (b)), as expected, with lower values during the summer (reaching a minimum at the beginning of July) and higher values during the winter months. The estimated spatial pattern (Figure 4 (c)) suggests a slight East-West gradient, with greater values on the Western side. Figure 5(a) shows the seasonal adjustment that needs to be made to the overall seasonal pattern ($s_4(t, d)$) in four different years (1998, 2000, 2003 and 2012); it can be seen that there is

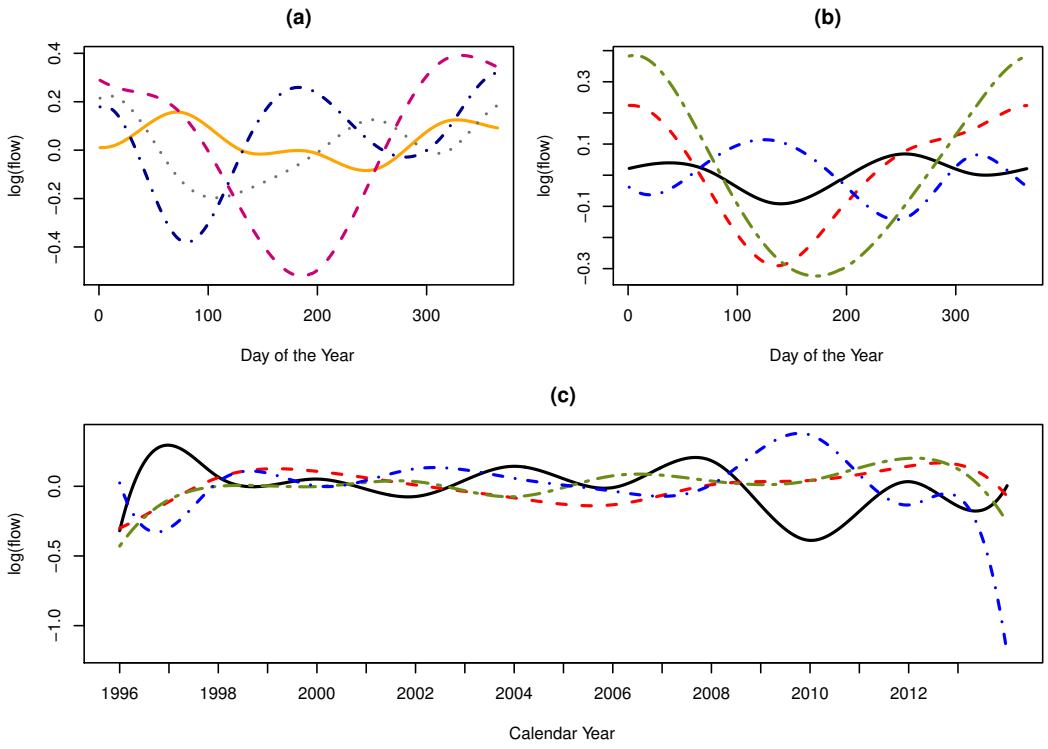

Figure 5: (a) Season-year interaction for 4 different years (1998 (orange), 2000 (pink), 2003 (grey) and 2012 (dark blue), (b) season-space interaction and (c) trend-space interaction for rivers Inver (black), Clyde (red), Carron (blue) and Deveron (green).

variation from year to year and, in particular, that the seasonal pattern in 2000 (pink curve) is very different from the rest.

Following Model (4), we can identify not only long term trends over the whole period/spatial region but also assess how the spatial distribution of extreme flows has changed over time by visual inspection of the interaction term $s_5(t, z)$, that can also be interpreted as the adjustment that needs to be made to the temporal trend $s_1(t)$ shown in Figure 4(a) at each location; these adjustments can be seen in Figure 5(c) for the four rivers described in Section 2, while the corresponding seasonal adjustment (interaction term $s_6(d, z)$) can be seen in Figure 5(b). While the trend appears to be very homogeneous over space, with little adjustment needed with respect to the overall trend shown in Figure 4(a), seasonality varies considerably among the different rivers (Figure 5(b)), supporting the idea that no unique seasonal pattern is valid over the whole country.

One can graphically show the fitted values from a spatial or a temporal point of view. The estimated 95^{th} quantile of river flow over Scotland is shown in Figure 6 for four different time points, namely 1^{st} of January, April, July and October, over years 1996 to 2000. These dates were chosen to illustrate possible differences between seasons. We can see how the contrast between East and West is more pronounced in some periods than others, and that there is a clear difference in the summer (1^{st} of July) and winter (1^{st} of January) values between 1996 and the remaining years. A model without a time-space interaction term would not have been able to identify a spatial pattern changing over time.

Even though the data were log transformed initially, it is possible to show the fitted model at each gauging station in the original scale as quantile regression is invariant to monotonic transformations. An example for the fitted model (with and without interaction) can be seen in Figure 7 for four different rivers located in various parts of the country (see Figure 1 for locations). While for River Carron the fitted models with and without interaction are very close, for the River Inver there are clear differences between the two models, especially at the beginning and end of the record.

Model (4) was estimated assuming independence; residual correlation was investigated by means of empirical variograms and autocorrelation plots. Once the spatial trend is taken into account, there is no spatial structure left in the residuals, as suggested by the flat empirical variograms (not included here). This is expected given the spatial flexibility that we have allowed for in the model by incorporating the interaction terms space-time and space-season. On the other hand, the residual autocorrelation plots indicated the presence of temporal correlation. Hence, an AR(1) process was assumed on the residuals, a common choice for environmental processes. The standard errors of the fitted model were adjusted for residual correlation following Equation (10) and are illustrated in Figure 7. As expected, the pointwise confidence bands become wider once correlation is taken into account.

The full model includes 2826 parameters. With these specifications, running the model took $\approx$ 10 hours on a CPU with 2.0Ghz (128Gb RAM). The PIRLS algorithm converged after 43

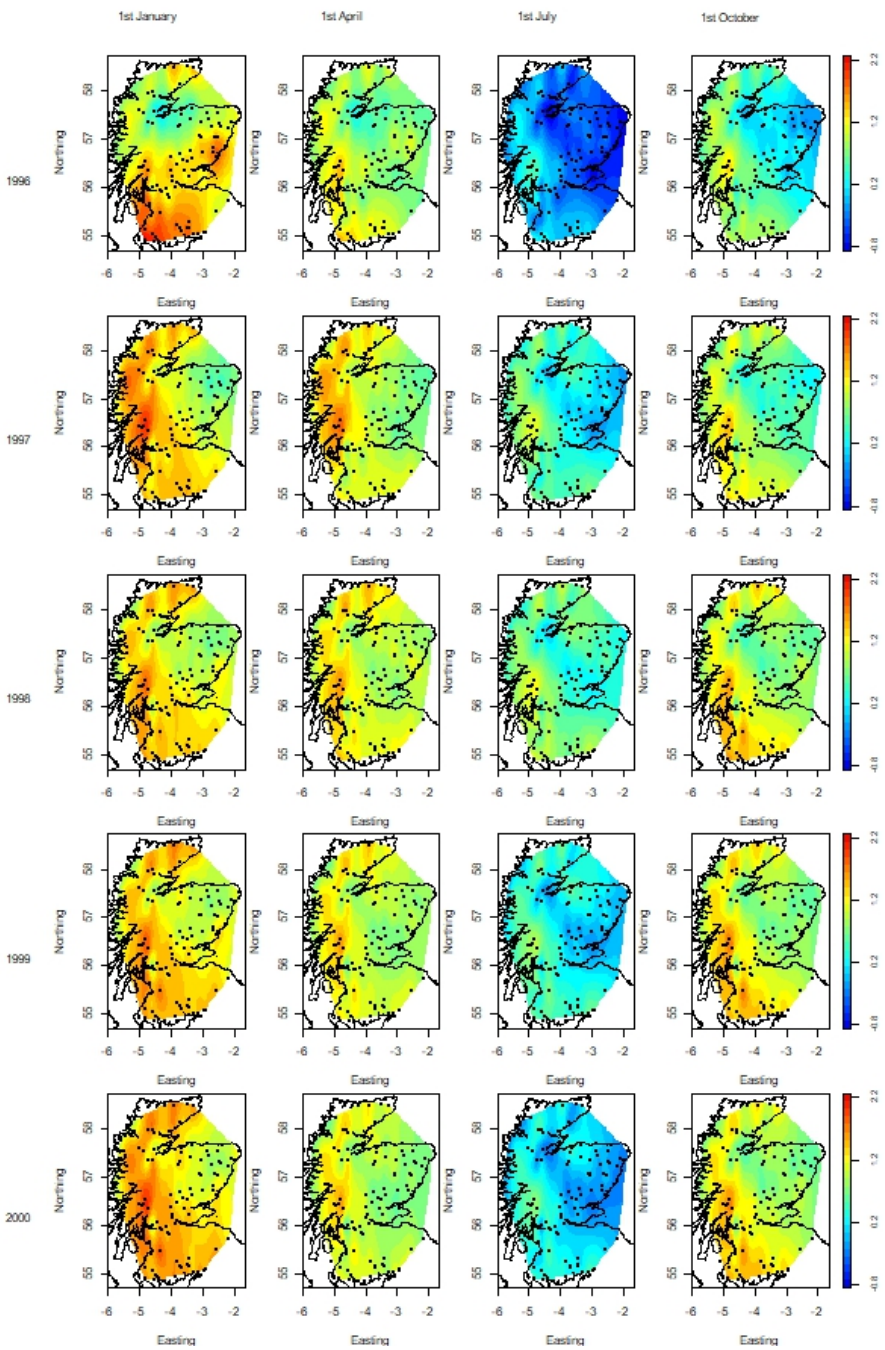

Figure 6: Estimated 95^{th} quantile of (logged) river flow at four different dates (columns) and five different years (rows)

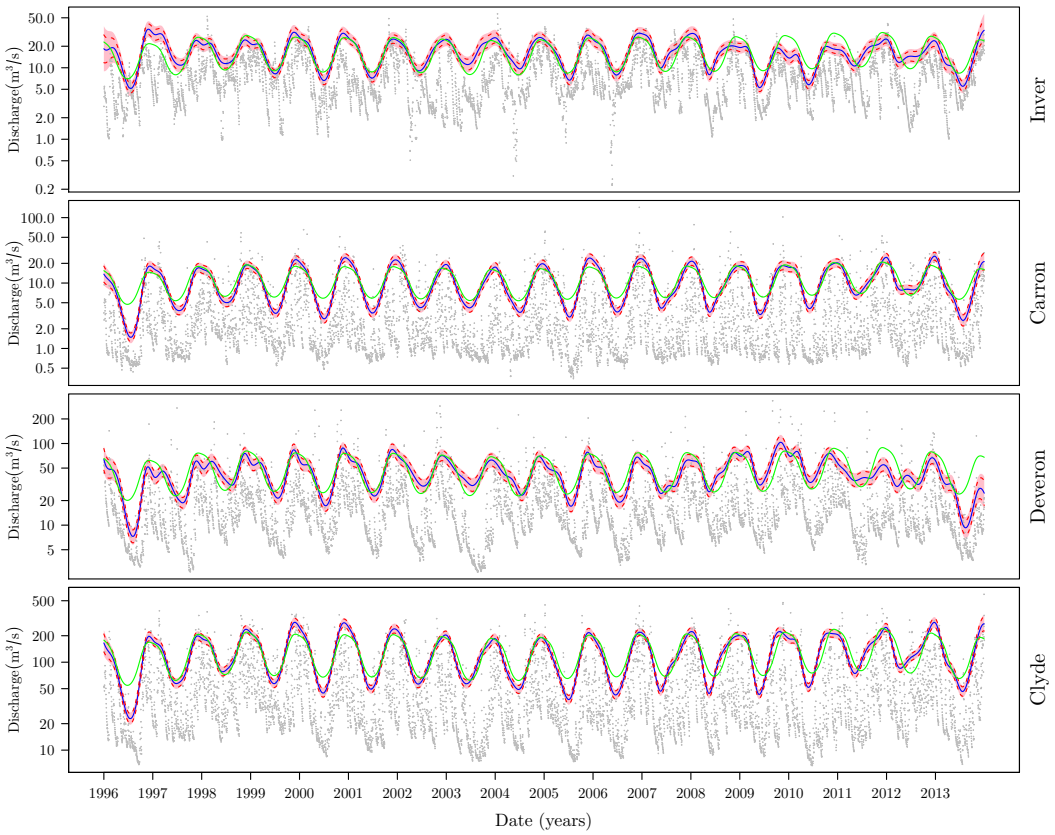

Figure 7: Fitted model (blue), 95% approximated confidence interval under independence (red) and dependence (pink) and fitted model without the interaction terms (green) for rivers Inver (North West), Carron (South East), Deveron (North East) and Clyde (South West).

iterations, with a tolerance value set to 10^{-2}. The approximation of the check-function as a weighted sum of squared residuals was also investigated. At the last iteration, the weighted residuals (i.e. using the approximation) were similar to those obtained using the check function. As a way of informally assessing whether the model was appropriate or not, the proportion of residuals above and below the fitted surface was calculated; these were equal to 5.2% and 94.8% respectively, not far from the theoretical expected values of 5% and 95%.

5 Simulation Study

A simulation study was run in order to evaluate the performance of the model proposed in Section 3.1 and compare it to that of the `R` package `quantreg` (Koenker, 2018) for the 95^{th} quantile. We simulated daily data under 4 different scenarios with varying sample size in time and space. For the time component, we consider 5 and 10 years of data, while for the space component we consider 20 and 50 locations on a spatial irregular grid on $[0, 1] \times [0, 1]$. Data were generated as the sum of a smooth time trend $s(t)$, a seasonal cycle $s(d)$ and a spatial component $s(z)$ as follows:

$$
\begin{aligned}
s(t) &= 1 + 0.3t + 0.7t^2, \\
s(d) &= \sin(2\pi d/365 + \pi/2), \\
s(z) &= s(z_1, z_2) = 0.2z_1 + 0.25z_2 + 3z_1z_2.
\end{aligned}
$$

Random noise was simulated from an asymmetric Laplace distribution $ALD(\mu, \sigma, q)$, a commonly used distribution in quantile regression, with $\mu = 0$ and $q = 0.95$. Regarding σ, two different levels of noise were considered in each scenario, corresponding to a signal to noise ratio (SNR) of two and four. Table 7 summarizes the different scenarios considered while and example of simulated data is illustrated in Figure 8.

Model (2) was fitted for the 95^{th} quantile to 100 simulated data sets under each scenario using the method described in Section 3.1 with $k = 5$ for all terms. The `rqss` function from the `quantreg R` package was also used to estimate the same model. Smoothing parameters

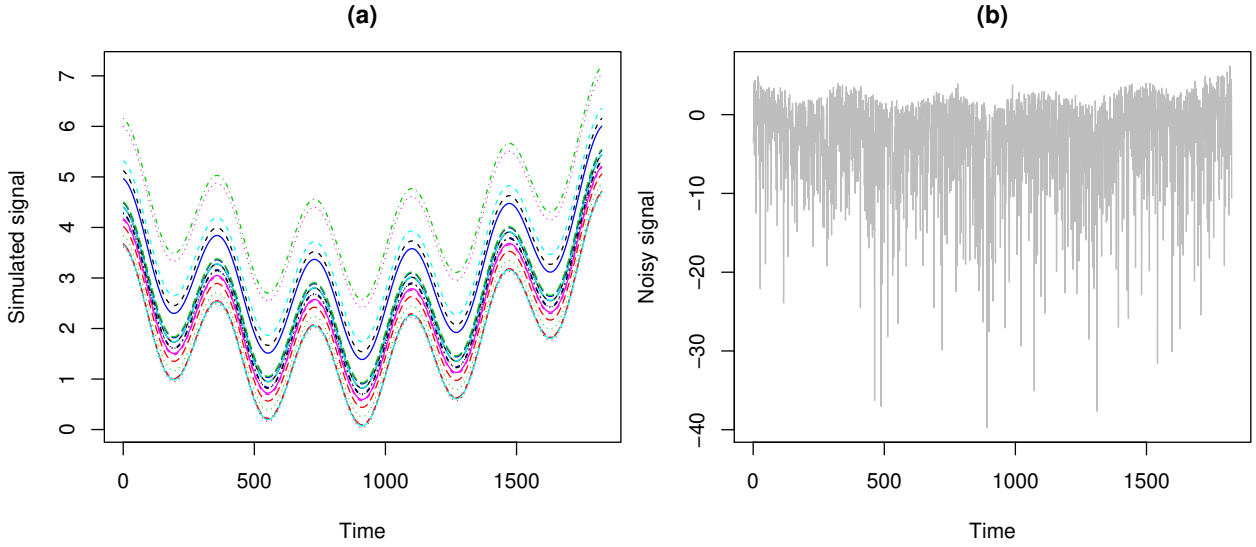

Figure 8: (a) Simulated signal under scenario A, (b) noisy signal simulated at one spatial location under scenario A4.

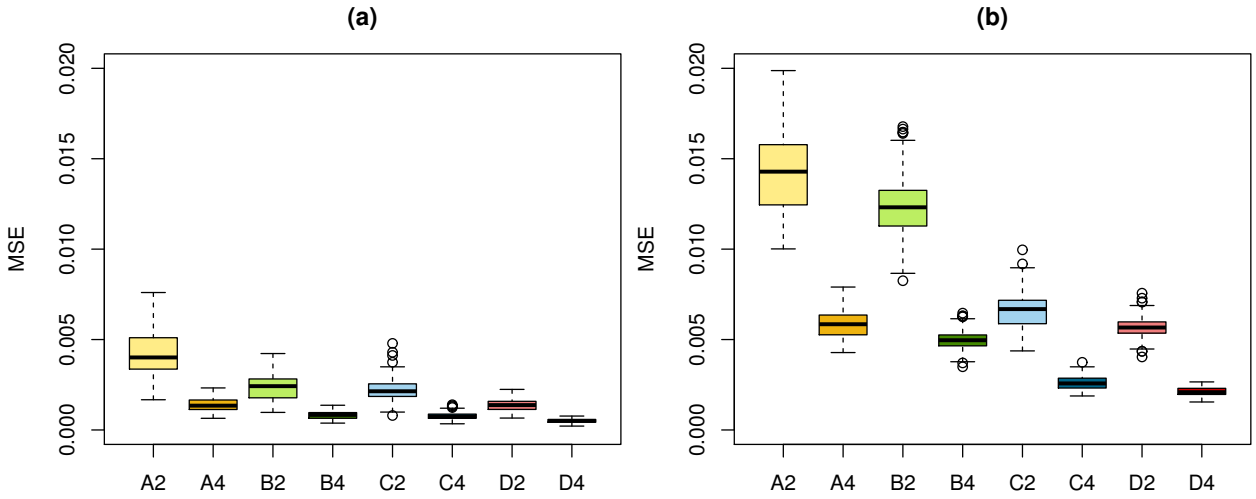

Figure 9: Boxplots of MSE under the different simulation scenarios when the 95^{th} quantile model is estimated using the method proposed in this paper (a) and the `quantreg` package (b).

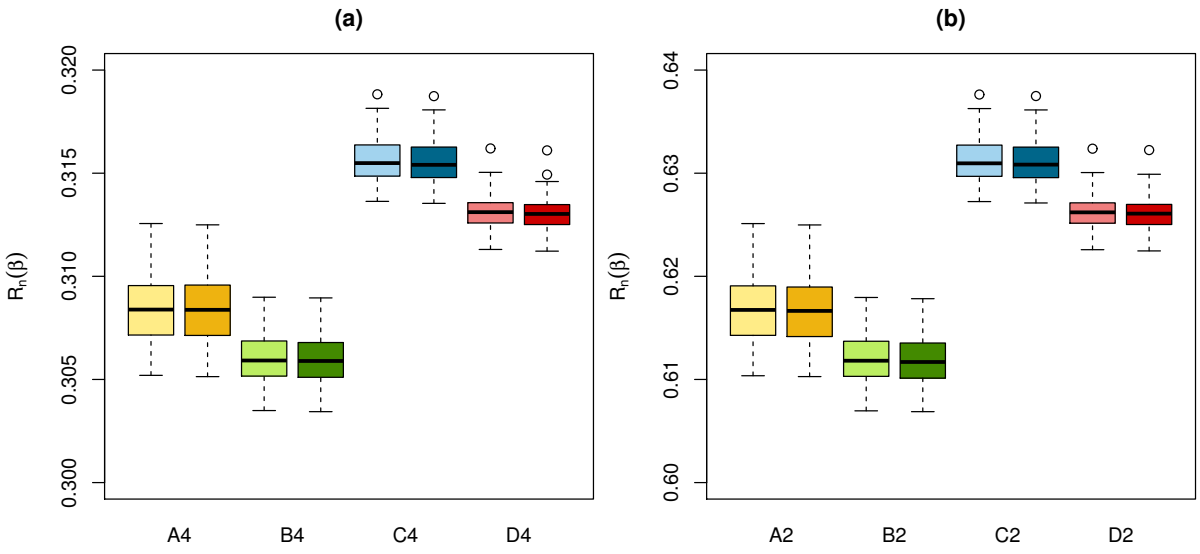

Figure 10: Boxplots of $R_n(\beta)$ under the different simulation scenarios when the 95^{th} quantile model is estimated using the method proposed in this paper (lighter colours) and the `quantreg` package (darker colours) with SNR=4 (a) and SNR= 2 (b).

in both cases were chosen based on SIC. Boxplots of the mean square error (MSE) between the simulated and estimated signal are shown in Figure 9, where it can be clearly seen that the MSE is considerably lower for the proposed method (panel (a)) in all simulation scenarios. In both panels (a) and (b) the MSE is smaller when the signal to noise ratio is greater and decreases as the sample size increases (both in time and space), as one would expect. From this simulation study, it can be concluded that for nonparametric spatio-temporal modelling our method performs better than classical quantile regression estimation. The poorer performance obtained using `quantreg` might be partly explained by the fact that is not possible to constrain the seasonal component to be cyclic.

To assess how good the approximation that the proposed method makes is (Equation (6)), we compared the following quantity for both approaches:

$$R_n(\beta) = \frac{1}{n}\sum_{i=1}^{n} \rho_\tau(y_i - \hat{y}_i)$$

where ρ_τ is the check function in Equation (1), y_i are the (noisy) simulated values, $\hat{y}_i$ are the

fitted values and n is the total number of observations. As can be seen in Figure 10, there are no differences in $R_n(\beta)$ between the two approaches, suggesting very good agreement between the original check function and its approximated version.

For the more complex model including the interaction terms, it is not possible to compare the performance of the proposed method with `quantreg` as the latter does not allow inclusion of smooth functions of dimension greater than two.

6 Discussion

Environmental processes are highly variable in both space and time and, as such, investigation of trends in their means may not be sufficiently informative, especially when the interest is in extreme values. On the other hand, river flow records tend to be relatively short, usually spanning about 20-30 years of data; modelling extremes, such as the 1 in 100 year design event, from such short records using the classical extreme value theory approach can be problematic, as the number of observations is dramatically reduced when considering the annual maxima series or only the peaks above a certain threshold. This paper proposes a spatio-temporal quantile regression approach for modelling extreme river flow values. Although the work was motivated by a case study on river flow in Scotland, the proposed model is applicable to other areas where quantiles might be of interest, e.g. air pollution. The model is built in a generalized additive model framework that allows inclusion of three-variate smooth functions to account for space-time interaction effects. The number of observations is crucial for reliable estimation in a quantile regression setting, especially when building models for very high (or very low) quantiles. Rather than modelling one river at a time, we use a single model for the whole dataset; this way we borrow information across locations, particularly important when dealing with short records.

Despite quantile regression being increasingly popular in environmental applications (see e.g. Weerts et al. (2011); Reich (2012); Planque and Buffaz (2008)), applications on spatio-temporal data are still very limited. We propose a flexible spatio-temporal model that takes into

account potential space-time interactions with the novelty that model parameters are estimated via penalized iterative re-weighted least squares. We believe this has two main advantages with respect to using linear programming methods; first, it is a more intuitive way of thinking about parameter estimation for those familiar with classic regression and least squares estimation. Second, this regression-like framework can be easily extended to incorporate covariates or to build more complex models, e.g. varying coefficient models. Further, estimating such a complex model is not possible using the classical quantile regression approach as currently implemented in the `R` package `quantreg`, where neither smooth functions depending on more than two covariates nor periodicity constraints can be included in the model. For a simpler spatio-temporal model with no interaction terms, results from a simulation study showed that our proposed method outperforms the standard quantile regression approach.

The focus of this work is on extreme values and hence we have only considered modelling single quantiles high up in the tail of the distribution (e.g. 90^{th}, 95^{th} quantile) separately. Even though the modelling approach used here does not allow fitting various quantiles simultaneously, it is easy to implement and therefore can be fitted for a range of quantiles independently to investigate, for instance, how the spatial distribution of extreme values changes with respect to the mean. Nevertheless, the issue of quantile crossing was not investigated and is left for further research. Schnabel and Eilers (2013) proposed the so-called "quantile sheets" to estimate a range of quantiles simultaneously by means of a bivariate smooth function of a covariate and the probability τ using P-splines; non-crossing of quantile curves is ensured by forcing monotonicity in the τ direction in the penalty. Finally, assessing goodness of fit for quantile models remains an area of open research; usual indices of performance such as the root mean square error or the correlation coefficient between observed and predicted values are of no use here. In this paper, goodness of fit was informally assessed based on the expected and observed proportion of positive and negative residuals. For linear quantile regression models, an equivalent to the coefficient of determination R^2 was proposed by Koenker and Machado (1999).

The proposed model was applied to a set of Scottish rivers spanning 18 years of data. In

Scotland, studies assessing changes in extreme river flows have been carried out mostly in the time domain. Spatially, differences in frequency and magnitude of river flow extreme events have been found between the East and the West (Black and Burns, 2002; Black, 1996), as well as in trends in annual maxima series (Black, 1996). In particular, two 'micro-climates' have been identified over 1980-2000, wetter in the North-West, with significant spring and autumn increases in the West and winter increases in the North, and drier in the South-East, especially in the summer months (Jenkins et al., 2009; Werritty, 2002). Even though we only have data for years 1996-2000 within the time frame 1980-2000, our results (partially shown in Figure 6) also support a clear East/West gradient in the 95^{th} quantile. Downscaled projections from global circulation models for the 2050s predict that observed trends are to continue in the near future (Fowler and Wilby, 2010; Werritty, 2002); however, while estimates over the winter months are fairly reliable there is great uncertainty over the summer. Regional climate models predict increases in winter extreme rainfall (Fowler and Kilsby, 2003) and an increase of annual runoff of 5-15% across the country for the 2050s but that could locally exceed 25% (Werritty, 2002). Given the uncertainty of the predictions and the heterogeneity of the observed changes over time and space, it is important to gain a better understanding of the spatio-temporal pattern of extreme river flows and we believe that the model proposed in this paper could prove useful in doing so. Given the observed and projected seasonal differences in river flow, it is common practice to divide the year into seasons to then model each season independently, see e.g. Hannaford and Buys (2012), where seasonal quantiles were calculated to investigate the presence of trends in the quantiles. Our approach avoids this division of the years into seasons, for which there is no general agreement (and hence a different division of the year may yield different results) and which may be changing due to climate change, and directly models the quantile of interest across time so that the temporal evolution of the spatial trend in extremes can be assessed.

Since estimates of flood risk derived from river flow data are usually based on relatively short records, the current procedure recommends adding a safety margin of 20% of the expected flow level to ensure that design infrastructure (e.g. flood barriers) can cope with an unexpected

extreme flow (HM Government, 2016). On the other side, after the 2007 UK floods (Pitt, 2008), action was taken and the plan was to lower the threshold at which alarms were issued. The results from the Scottish data presented in this paper suggest that neither of these adjustments may be sufficient to provide accurate flood warning as trends for high quantiles have not changed in a homogeneous way spatially. These spatial differences have potentially important implications for decision making, where optimizing the balance between expenditure and risk reduction is a critical part of the decision process.

Additional information and supporting material for this article is available online at the journal's website.

7 Acknowledgements

The authors would like to thank Alistair Cargill from SEPA, the NRFA for providing the data and the Kelvin-Smith postgraduate studentship scheme from the University of Glasgow for partially funding this research.

References

Arnell, N. and S. Gosling (2013). The impacts of climate change on river flow regimes at the global scale. *Journal of Hydrology 486*, 351–364.

Arnell, N. and S. Gosling (2016). The impacts of climate change on river flood risk at the global scale. *Climate Change 134*, 387–401.

Black, A. (1996). Major flooding and increased flood frequency in Scotland since 1988. *Physics and Chemistry of the Earth 20*(5-6), 463–468.

Black, A. and C. Burns (2002). Re-assessing the flood risk in Scotland. *The Science of the Total Environment 294*, 169–184.

Black, A. and A. Werritty (1997). Seasonality of flooding: a case study of North Britain. *Journal of Hydrology 195*, 1–25.

Brown, B. and S. Resnick (1977). Extreme values of independent stochastic processes. *Journal of Applied Probability 14*, 732–739.

Cade, B. and B. Noon (2003). A gentle introduction to quantile regression for ecologists. *Frontiers in Ecology and the Environment 1*(8), 412–420.

Coles, E. (2004). *An Introduction to Statistical Modelling of Extreme Values*. Springer Series in Statistics.

Cooley, D., D. Nychka, and P. Naveau (2007). Bayesian spatial modeling of extreme precipitation return levels. *Journal of the American Statistical Association 102*(479), 824–840.

Cressie, N. and C. Wikle (2011). *Statistics for Spatio-temporal Data*. New York: Wiley.

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

Davison, A. and M. Gholamrezaee (2012). Geostatistics of extremes. *Proceedings of the Royal Society A 468*, 581–608.

Davison, A., S. Padoan, and M. Ribatet (2012). Statistical modelling of spatial extremes. *Statistical Science*.

Diggle, P. and P. Ribeiro Jr. (2007). *Model-based Geostatistics*. Springer Series in Statistics.

Eilers, P. and B. Marx (1996). Flexible smoothing with B-splines and penalties. *Statistical Science 11*, 89–102.

Eilers, P. and B. Marx (2002). Generalized linear additive smooth structures. *Journal of Computational and Graphical Statistics 11*, 758–783.

503 Eilers, P. and B. Marx (2003). Multivariate calibration with temperature interaction using two-
504 dimensional penalized signal regression. *Chemometrics and Intelligent Laboratory Systems 66*,
505 159–174.

506 Eilers, P. and B. Marx (2009). The craft of smoothing. Course hand-out (Dublin, October 2009).

507 Eilers, P. and B. Marx (2010). Splines, knots, and penalties. *Wiley Interdisciplinary Reviews:*
508 *Computational Statistics 2*(6), 637–653.

509 Embrechts, P., E. Koch, and C. Robert (2016). Space-time max-stable models with spectral
510 separability. *Advances of Applied Probability 48*(A), 77–97.

511 Evans, E., R. Ashley, J. Hall, E. Penning-Rowsell, P. Sayers, C. Thorne, and A. Watkinson
512 (2004). Foresight future flooding. Technical report, Office of Science and Technology, London.

513 Fowler, H. and M. Ekström (2009). Multi-model ensemble estimates of climate change impacts
514 on UK seasonal precipitation extremes. *International Journal of Climatology 29*, 385–416.

515 Fowler, H. and C. Kilsby (2003). Implications of changes in seasonal and annual extreme rainfall.
516 *Geophysical Research Letters 30*(13), 1720. DOI:10.1029/2003GL017327.

517 Fowler, H. and R. Wilby (2010). Detecting changes in seasonal precipitation extremes using re-
518 gional climate model projections: Implications for managing fluvial flood risk. *Water Resources*
519 *Research 46*, W03525.

520 Fuentes, M., J. Henry, and B. Reich (2012). Nonparametric spatial models for extremes: appli-
521 cation to extreme temperature data. *Extremes*, 1–27. 10.1007/s10687-012-0154-1.

522 Georgi, B., S. Isoard, M. Asquith, C. Garzillo, R. Swart, and J. Timmerman (2016). Urban
523 adaptation to climate change in Europe 2016 - transforming cities in a changing climate.
524 Technical report, European Environmental Agency.

525 Hallin, M., Z. Lu, and K. Yu (2009). Local linear spatial quantile regression. *Bernoulli 15*,
526 658–686.

Hannaford, J. and G. Buys (2010). Modification of climate-river flow associations by basin properties. *Journal of Hydrology 389*, 186–204.

Hannaford, J. and G. Buys (2012). Trends in seasonal river flow regimes in the UK. *Journal of Hydrology 475*, 158–174.

Hastie, T. and R. Tibshirani (1990). *Generalized Additive Models*. Chapman and Hall.

He, X., P. Ng, and S. Portnoy (1998). Bivariate quantile smoothing splines. *Journal of the Royal Statistical Society B 60*(3), 537–550.

HM Government (2016). National flood resilience review. Technical report, Cabinet Office, Department for Environment, Food & Rural Affairs, The Rt Hon Ben Gummer MP, and The Rt Hon Andrea Leadsom MP.

IPCC (2014). Climate change 2014: Synthesis report. Contribution of working groups i, ii and iii to the fifth assessment report of the Intergovernmental Panel on Climate Change. Technical report, Intergovernmental Panel on Climate Change [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)], Geneva, Switzerland, 151 pp.

Jenkins, G. J., J. M. Murphy, D. M. H. Sexton, J. A. Lowe, P. Jones, and C. G. Kilsby (2009). UK Climate Projections: Briefing report. Technical report, Met Office Hadley Centre, Exeter, UK.

Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs.

Koenker, R. (2018). quantreg: Quantile Regression. R package version 5.35 https://CRAN.R-project.org/package=quantreg

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Koenker, R. and K. Hallock (2001). Quantile regression. *The Journal of Economic Perspectives 15*(4), 43–156.

Koenker, R. and J.A.F. Machado (2001). Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association 94*(448), 1296–1310.

Koenker, R., P. Ng, and S. Protnoy (1994). Quantile smoothing splines. *Biometrika 81*(4), 673–680.

Lee, D. and T. Neocleous (2010). Bayesian quantile regression for count data with application to environmental epidemiology. *Journal of the Royal Statistical Society. Series C 59*(5), 905–920.

Neelon, B., F. Li, L. Burgette, and S. Neelon (2015). A spatiotemporal quantile regression model for emergency department expenditures. *Statistics in Medicine 34*, 2559–2575.

Pitt, M. (2008). The Pitt review: Learning lessons from the 2007 floods. Technical report, London: Cabinet Office.

Planque, B. and L. Buffaz (2008). Quantile regression models for fish recruitment environment relationships: four case studies. *Marine Ecology Progress Series 357*, 213–223.

Prosdocimi, I., T. Kjeldsen, and C. Svensson (2013). Non-stationarity in annual and seasonal series of peak flow and precipitation in the UK. *Natural Hazards and Earth System Sciences (Discussions) 1*, 5499–5544.

Reich, B. (2012). Spatiotemporal quantile regression for detecting distributional changes in environmental processes. *Journal of the Royal Statistical Society. Series C 61*(4), 535–553.

Reich, B., M. Fuentes, and D. Dunson (2011). Bayesian spatial quantile regression. *Journal of the American Statistical Association 106*(493), 6–20.

Reiss, P. and L. Huang (2012). Smoothness selection for penalized quantile regression splines. *The International Journal of Biostatistics*. DOI:10.1002/env.1106.

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

572 Schnabel, S. and P. Eilers (2013). Simultaneous estimation of quantile curves using quantile
573 sheets. *Advances in Statistical Analysis 97*, 77–87.

574 Smith, R. (1990). Max-stable processes and spatial extremes. Unpublished manuscript.

575 Sun, Y., H. Wang, and M. Fuentes (2016). Fused adaptive Lasso for spatial and temporal quantile
576 function estimation. *Technometrics 58*(1), 127–137.

577 The Scottish Government (2010). Electronic Resource [Accessed 13/09/2010].

578 Wadsworth, J. and J. Tawn (2012). Dependence modelling for spatial extremes.
579 *Biometrika 99*(2), 253–272.

580 Weerts, A., H. Winsemius, and J. Verkade (2011). Estimation of predictive hydrological uncer-
581 tainty using quantile regression: examples from the national flood forecasting system (England
582 and Wales). *Hydrology and Earth System Sciences 15*, 255–265.

583 Werritty, A. (2002). Living with uncertainty: climate change, river flows and water resource
584 management in Scotland. *The Science of the Total Environment 294*, 29–40.

585 Wood, S. (2006). *Generalized Additive Models - An Introduction with R*. Champman and
586 Hall/CRC.

River	Catchment area (km^2)	Max elevation (mAOD)	Mean flow (m^3/s)	Q95[a] (m^3/s)	Q99[b] (m^3/s)	mean flow/ catchment area
Inver	137.5	1108.6	8.182	20.500	28.852	0.060
Clyde	1903.1	745.2	49.362	159.710	263.334	0.026
Carron	122.3	561.9	3.664	13.821	28.569	0.030
Deveron	954.9	754.2	18.605	50.921	107.931	0.019

[a] 95th quantile of river flow, [b] 99th quantile of river flow

Table 1: Main characteristics of rivers Inver, Clyde, Carron and Deveron.

Scenario	Time	Space	n	SNR	
				2	4
A	5 years	20 locations	36500	A2	A4
B	10 years	20 locations	73000	B2	B4
C	5 years	50 locations	91250	C2	C4
D	10 years	50 locations	182500	D2	D4

Table 2: Scenarios considered in the simulation study.