Evaluation of classical spatial-analysis schemes of extreme rainfall

Davide Ceresetti, Eugen Ursu, Julie Carreau, Sandrine Anquetin,
Jean-Dominique Creutin, Laurent Gardes, Stéphane Girard, Gilles Molinie

Evaluation of classical spatial-analysis schemes of extreme rainfall

D. Ceresetti[1,*], E. Ursu[2,*], J. Carreau[3,*], S. Anquetin[1], J. D. Creutin[1], L. Gardes[2], S. Girard[2], and G. Molinié[1]

[1]LTHE, Laboratoire d'étude des transferts en hydrologie et environnement, UMR5564, Université de Grenoble/CNRS, France
[2]Team Mistis, INRIA Rhône-Alpes and LJK, 655, avenue de l'Europe, Montbonnot, 38334 Saint-Ismier Cedex, France
[3]HSM Montpellier, Université de Grenoble/CNRS, France
[*]These authors contributed equally to this work.

Correspondence to: G. Molinié (gilles.molinie@ujf-grenoble.fr)

Abstract. Extreme rainfall is classically estimated using raingauge data at raingauge locations. An important related issue is to assess return levels of extreme rainfall at ungauged sites. Classical methods consist in interpolating extreme-value models. In this paper, such methods are referred to as regionalization schemes. Our goal is to evaluate three classical regionalization schemes. Each scheme consists of an extreme-value model (block maxima, peaks over threshold) taken from extreme-value theory plus a method to interpolate the parameters of the statistical model throughout the Cévennes-Vivarais region. From the interpolated parameters, the 100-yr quantile level can be estimated over this whole region. A reference regionalization scheme is made of the couple block maxima/kriging, where kriging is an optimal interpolation method. The two other schemes differ from the reference by replacing either the extreme-value model block maxima by peaks over threshold or kriging by a neural network interpolation procedure. Hyper-parameters are selected by cross-validation and the three regionalization schemes are compared by double cross-validation. Our evaluation criteria are based on the ability to interpolate the 100-yr return level both in terms of precision and spatial distribution. It turns out that the best results are obtained by the regionalization scheme combining the peaks-over-threshold method with kriging.

1 Introduction

The statistical frequency analysis of extreme rainfall can be seen as a normalization procedure allowing site-to-site comparison. Therefore, several applications rely on rainfall frequencies (or equivalently return period), such as the design of civil equipments (bridge, dike, dam,...) or climate studies and specifically variable trends due to climate change (see Zwiers and Kharin, 1998; Kharin and Zwiers, 2005, for examples).

Estimating the frequency of a particular event implies the knowledge of its background distribution. Concerning rainfall the most reliable and longest data series are those recorded by raingauges. As one very long series is more efficient than several shorter ones, numerous methods have been elaborated to gather in the same set rainfall series with common statistical properties (Carreau and Girard, 2011; Cunnane, 1988; Daouia et al., 2011; Hosking and Wallis, 1997; Gardes and Girard, 2008; Gardes and Girard, 2010; Neppel et al., 2011). Hosking and Wallis (1997) coined the term "regional frequency analyses" for such methods. Regional frequency analysis is by extension able to provide spatial representation of extreme rainfall, which is of prime interest for example to compare with model outputs not collocated with reference raingauges.

Another point is extreme areal integration of the rainfall intensities. An elegant procedure to estimate extreme values of areal intensities is to fit a spatial extreme-value model (max-stable process, de Haan 1984; de Haan and Pickands, 1986) to extreme rainfall collected asynchronously at different raingauge stations. Such models have to take into account the spatial and temporal dependence between extreme point rainfall (Aryal et al., 2009; Frei and Schar, 2001; Katz, 2010). This technique is relatively recent and not widely used in practical applications. In addition, the heterogeneity of extreme rainfall in the study region (Ceresetti et al., 2010) makes the application of max-stable processes very challenging. As a consequence, the current study focuses on more classical methods.

The MedUp project was dedicated to the assessment of uncertainties in climate projections of extreme-weather impacts. In such a context, it is required to implement a statistical tool to analyse extreme rainfall in the study region. In the MedUp framework, extreme-value statistics were used to compute severity diagrams. Ramos et al. (2005) have shown the importance of considering extremes of areal rainfall to assess the storm severity and propose to display the return period as a function of the temporal and spatial scales of integration: the so-called severity diagrams. To compute severity diagrams, Ceresetti et al. (2011) use the dynamic scaling framework to convert maximum areal rainfall intensities into point ones and then extreme-value theory to calculate return periods. In this project, severity diagrams have further been used to assess extreme rainfall from the large amount of outputs yielded by ensemble simulations (Norbiato et al., 2007; Ceresetti et al., 2011; Vié et al., 2012).

Hosking and Wallis (1997) describe step by step the implementation of the regional frequency analysis. After screening the data, they propose to delineate homogeneous regions for extreme rainfall, then to choose the appropriate probability function in order to fit heavy rainfall frequencies and to interpolate the probability function parameters allowing to compute regional return levels or periods. The latter two steps, the models of extreme point rainfall and the interpolation schemes, have been subject to discussions in the literature.

There are two families of probability models asymptotically derived from point-rainfall frequency analysis. If the point-rainfall set is made of maxima drawn from fixed-length periods of time, a suitable model is the so-called generalized extreme-value model (GEV). If the set is made of excesses regarding a given threshold, the theoretical model is the generalized Pareto distribution (GPD) (see Coles, 2001, for an example).

Rainfall interpolation has been a recurrent topic in literature for decades (see Ripley, 1981, or Myers, 1994, for an overview). Rainfall interpolation methods have been classified as deterministic (inverse distance, multiquadratic, spline) and optimal (kriging, Chiles, 2001). Optimal is used in the sense that interpolation weights are selected so as to optimize some criterion of best fit at the data points. Creutin and Obled (1982) and Borga and Vizzaccaro (1997) noted that these methods could be equivalent in certain conditions. Creutin and Obled (1982) noted that one advantage of optimal interpolation is the possibility to compute the estimation variance. Borga and Vizzaccaro (1997), who compared multiquadratic (deterministic) and kriging (optimal) interpolations, found them equivalent in terms of performance statistics except in the case of low-density raingauge networks. In such a case, kriging performs better than multiquadratic. Kriging allows accounting for the spatial dependence between observations, because the weights of the linear interpolation coefficients depend on a spatial structure function, the so-called variogram. To overcome the computation time

necessary to fit a variogram model (referred to above as implementation cost), Bastin et al. (1984) and Lebel and Bastin (1985) propose using a climatological variogram determined only once for any random variable, while Goovaerts (2000) implemented an automatic fitting procedure for variograms, using a weighted sum-of-squares minimization where the weights are functions of the inter-distance between data and of the number of couples for each inter-distance. In addition, kriging allows taking into account covariates which are easier to sample than rainfall (like altitude) (Goovaerts, 2000; Prudhomme and Reed, 1999). Thanks to all these functionalities, kriging is an attractive interpolator of rainfall data and it is widely used in this research field.

Another optimal but automatic (no hand tuning) technique is neural networks. They are flexible regression estimators useful for spatial interpolation (Bishop, 2006). Neural networks are non-linear, non-parametric estimators which can in theory approximate any continuous function (Hornik et al., 1989). However, neural networks, like any other standard regression method, do not take into account the spatial dependence structure in the data. The performances of neural networks for spatial interpolation are evaluated in Snell et al. (2000), showing that, usually, the predictive accuracy of neural networks is superior to that of the traditional methods (nearest neighbors, spatial average, inverse-distance weighted average).

To interpolate rainfall data, two state-of-the-art methods are kriging and neural networks. Despite its benefits, the two major steps in kriging are sources of uncertainties. Any application of kriging requires the choice of a spatial dependence model (variogram model) and of its estimation (often handmade). Secondly, kriging assumes a linear relationship between rainfall parameters and the altitude covariate, which is valid only for some time steps and raingauge environments (Ceresetti et al., 2010; Molinié et al., 2012).

Therefore, the purpose of this study is to assess the uncertainty in regional frequency analysis of rainfall intensities due to two major steps: (i) the choice of the extreme-value density function (EV) and (ii) the interpolation scheme. We consider the most widespread EV and interpolation schemes, i.e. GEV and kriging, as the reference regional-frequency analysis scheme RS_{ref}. We compare RS_{ref} with two other schemes: RS2, which differs from RS_{ref} in its EV model, and RS3, which differs in its interpolation scheme. We have implemented the comparison of the three RS models with a double cross-validation procedure. This procedure allows the selection of the complexity level of the models, such as the variogram model, and the fair comparison of the models' performance on out-of-sample data (Bishop, 2006).

This paper is organized as follows. Section 2 reviews the two extreme-value models and the kriging and neural-networks interpolation methods. Then, Sect. 3 details the practical implementation of the three regionalization schemes. In Sect. 4, we discuss hyper-parameter selection and model comparison. The application to the

Cévennes-Vivarais rainfall data is detailed in Sect. 5 and some conclusions are drawn in Sect. 6.

2 Methodology

2.1 Statistical modelling of extreme rainfall

Given a sample set $z_1, ..., z_n$ of independent observations with common cumulative distribution function F, our goal is to model the tail distribution of F in order to estimate extreme quantiles q_{α_n}, defined by $1 - F(q_{\alpha_n}) = \alpha_n \in (0, 1)$. A quantile is called "extreme" if $\alpha_n < 1/n$. In this situation, the true quantile q_{α_n} is, in most cases, larger than the maximum observation. In this paper, the observations will be the rainfall intensity records measured at a given station.

We often speak of "T-yr return level" instead of quantile. The T-yr return level is the amount of rain exceeded on average every T years, and it corresponds to the quantile of order $\alpha_n = \tilde{T}/(nT)$, where $\tilde{T}$ is the number of years of measurements. Two methods of quantile estimation are considered: block maxima and peaks over threshold (POT). These approaches are based on the generalized extreme-value (GEV) distribution (Coles (2001), Sect. 3.1.3), which is given by

$$G_\xi(t) = \begin{cases} \exp\left\{-(1+\xi t)^{-1/\xi}\right\} & \text{if } \xi \neq 0, \quad (1a) \\ \exp\{-\exp(-t)\} & \text{if } \xi = 0, \quad (1b) \end{cases}$$

where the parameter ξ is referred to as the extreme-value index. The two methods differ first by the way heavy rainfall is selected in the data set and second by the expression of the density-function model.

2.1.1 1-yr maxima approach

In extreme-value theory, the block-maxima approach consists in fitting a GEV distribution, $G_\xi((.-\mu)/\sigma)$, to a set of maxima each drawn from a data block. In this study, data blocks span one year (from January to December), including thus an average of 365.25 records of daily rainfall amounts $(z_1, ..., z_{365})$. Thus, the block-maxima approach will be called in the following the 1-yr maxima approach. For each year $i = 1, ..., \tilde{T}$ (with $\tilde{T}$ as the number of years of observations), the maximum $m_i = \max(z_{(i-1)s+1}, ..., z_{i365})$ is extracted. The GEV distribution (Eq. 1) provides a model for the distribution of the 1-yr maxima. We propose estimating the three parameters ξ, μ and σ by $\hat{\xi}_1$, $\hat{\mu}_1$ and $\hat{\sigma}_1$, which maximize the likelihood

$$l_1(\xi, \mu, \sigma; \{m_i\}_{i=1,...,r}) = \prod_{i=1}^{r} \frac{1}{\sigma} g_\xi\left(\frac{m_i - \mu}{\sigma}\right), \quad (2)$$

where g_ξ is the first derivative of G_ξ. When the observations are the rainfall intensities collected at a single station, the easiest way to constitute the data blocks is to group the observations year by year. In this case, the T-yr return level is estimated by
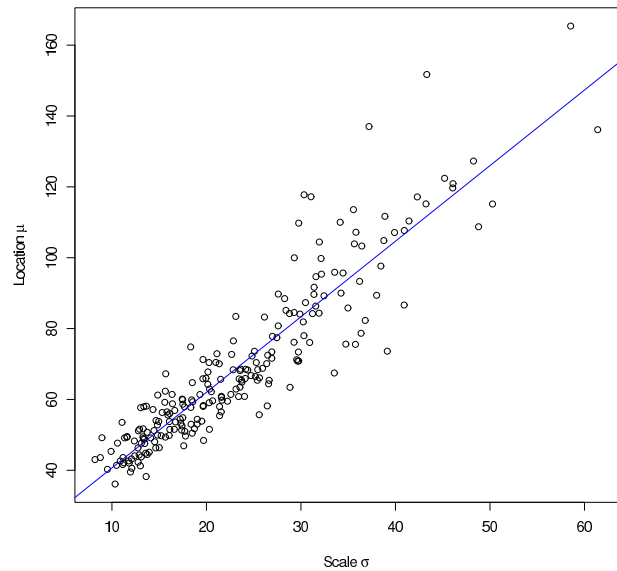
Fig. 1. Scatter plot of GEV location parameters (μ) against scale parameters (σ) estimated for every station. The line added corresponds to linear regression.

$$\hat{q}_T^{(1)} = \hat{\mu}_1 - \frac{\hat{\sigma}_1}{\hat{\xi}_1}\left[1 - \left\{-\log\left(1 - \frac{1}{T}\right)\right\}^{-\hat{\xi}_1}\right]. \quad (3)$$

At each station, the GEV distribution is fitted to these annual maxima by maximizing the likelihood function l_1. For each station $s \in \{1, ..., S\}$ in the training set, three GEV parameters $(\mu^{(s)}, \sigma^{(s)}, \xi^{(s)})$ have to be estimated. In order to get some insight into our data, a preliminary analysis is performed using the complete data set. In Fig. 1, the GEV location and scale parameters $\mu^{(s)}$ and $\sigma^{(s)}$ are estimated at each station $s \in \{1, ..., S\}$ and are plotted one against the other together with the regression line. Since the coefficient of determination R^2 is equal to 0.838, it seems reasonable to assume a linear relationship of the form $\mu = a + b\sigma$ between these two parameters. This preliminary analysis leads us to consider the following constrained GEV model for the annual maxima at the stations:

$$\mu^{(s)} = a + b\sigma^{(s)}, \quad \xi^{(s)} = \xi,$$

where $s \in \{1, ..., S\}$. The parameters of this model, a, b, ξ, $\sigma = \{\sigma^{(s)}\}_{s=1}^{S}$, are estimated by maximizing the function (see Eq. 2)

$$\prod_{s=1}^{S} l_1(\xi, a + b\sigma^{(s)}, \sigma^{(s)}; \{m_{s,j}\}_{j=1,...,n_s}), \quad (4)$$

where $m_{s,j}$ is the maximum observation of the j-th year for the s-th site, S is the number of sites and n_s represents the number of years of measure at the s-th site. From

the estimates $\hat{a}$, $\hat{b}$, $\hat{\xi}$ and $\hat{\sigma}^{(s)}$, $s = 1, \ldots, S$, a prediction at an ungauged site $\boldsymbol{x}$ can be computed by interpolating the scale parameter $\sigma(\boldsymbol{x})$. Then one gets $\hat{\mu}(\boldsymbol{x}) = \hat{a} + \hat{b}\hat{\sigma}(\boldsymbol{x})$ and $\hat{\xi}(\boldsymbol{x}) = \hat{\xi}$.

2.1.2 Peaks-over-threshold model

In the "peaks-over-threshold" approach (POT), one models the statistical behaviour of observations exceeding a given threshold ($\{z_i^* = z_i - u | z_i > u, i \in \{1, \ldots, n\}\}$, u being the threshold). Denoting by F_u the distribution function of the exceedances above the threshold u, it is well known (see Coles, 2001, Theorem 4.1) that for u large enough, there exists $\beta(u) > 0$ such that

$$1 - F_u(t) \approx -\log G_\xi \left(\frac{t - u}{\beta(u)} \right). \tag{5}$$

The associated likelihood is then given by

$$l_2(\xi, \beta(u); \{z_i^*\}_{i=1,\ldots,d})$$
$$= \prod_{i=1}^{d} \frac{1}{\beta(u)} g_\xi \left(\frac{z_i^* - u}{\beta(u)} \right) \bigg/ G_\xi \left(\frac{z_i^* - u}{\beta(u)} \right),$$

where the threshold u is fixed a priori and d is the number of exceedances. The corresponding maximum likelihood estimators are denoted by $\hat{\xi}_2$ and $\hat{\beta}_2$. Using the relationship in Eq. (5), the tail of the distribution function F can be approximated by

$$1 - F(t) = (1 - F(u))(1 - F_u(t - u))$$
$$\approx -(1 - F(u)) \log G_\xi((t - u)/\beta(u)),$$

and estimating $1 - F(u)$ by d/n, the observed proportion of exceedances, the T-yr return level (i.e. the quantile of order $\tilde{T}/(nT)$) is estimated by

$$\hat{q}_T^{(2)} = u - \frac{\hat{\beta}_2}{\hat{\xi}_2} \left[1 - \left(\frac{\tilde{T}}{dT} \right)^{-\hat{\xi}_2} \right]. \tag{6}$$

The practical implementation of POT is used to estimate the three parameters u, β and ξ at each gauged site. The first step consists in declustering the data by filtering dependent data to obtain a set of independent exceedances (Coles, 2001, Sect. 5.3.2). To determine clusters of extremes, one needs to fix a first very low threshold. Starting from the maximum, the first cluster is defined as the largest set of consecutive observations such that at most r values fall below the threshold. Removing this first cluster from the initial data, the next cluster is determined in the same way. This procedure is iterated until all the remaining observations fall below the threshold. Finally, we only keep the maximum observation in each cluster in our estimation procedure. The second step is the choice of a sufficiently high threshold u for each station. The threshold is fixed to the $1 - p$ empirical quantile of daily rainfall amounts, for a small p. The remaining parameters β and ξ can then be estimated by the method of moments or by Maximum Likelihood Estimation.

2.1.3 Comparison between the two approaches

Let us note that, when T is large, one has

$$\hat{q}_T^{(1)} \simeq \hat{\mu}_1 - \frac{\hat{\sigma}_1}{\hat{\xi}_1} \left[1 - \left(\frac{1}{T} \right)^{-\hat{\xi}_1} \right],$$

and thus the analytical expression of the estimated quantile using the 1-yr maxima approach is similar to Eq. (6) derived from the peaks-over-threshold model. Parameters u and $\hat{\mu}_1$ can be interpreted as location parameters, while $\hat{\beta}_3$ and $\hat{\sigma}_1$ can be interpreted as scale parameters. In both cases, $\hat{\xi}_3$ and $\hat{\xi}_1$ correspond to the extreme-value index. As a consequence, both approaches use three parameters with analogous interpretations. However, as their fittings do not rely on the same data set, they may therefore yield different results.

2.2 Interpolation methods

The two approaches described above yield estimates of the extreme-value density parameters u, ξ and σ or β at each gauged station. From Eqs. (3) or (6), return levels at each gauge station can then be derived. These so-called local estimates have to be interpolated to perform a regional analysis. As explained in the introduction (Sect. 1), kriging with an external drift depending on the altitude and neural networks are implemented because of their popularity and ability to perform interpolation with covariates. Kriging and neural network interpolation perform better when applied to Gaussian random variables. Since rainfall intensity is known to be non-Gaussian, the interpolation with both methods is applied to the extreme-value density parameters which are approximately normal. Return-period levels are then calculated.

2.2.1 Kriging with an external drift

Kriging is a linear, unbiased and optimal interpolation method (see e.g. Chiles, 2001 for more details). Kriging with an external drift (Goovaerts, 1997; Wackernagel, 1998) is a non-stationary method which generalizes the kriging interpolation. It consists in estimating $\hat{z}$, a realization of the random variable Z at an ungauged location x_0 as a function of: (i) observed realizations of Z (for instance, the return level known at the sites $\boldsymbol{x}_s$, $s = 1, \ldots, S$) and (ii) a covariate Y better documented than Z. Goovaerts (2000) showed that kriging rainfall fields with external drift yields better results than co-kriging when the covariate includes the altitude. Ceresetti et al. (2010) and Molinié et al. (2012) show a close correlation between daily extreme rainfall and terrain elevation. Therefore, the altitude of a digital terrain model is used in this study as the covariate Y to parametrize the mean drift of Z in space.

Kriging with an external drift requires to break Z down into the sum of its mean $\bar{Z}$ and residual Z_r. The mean $\bar{Z}$ is derived from the covariate Y. The simple kriging interpolation

Nat. Hazards Earth Syst. Sci., 12, 3229–3240, 2012

www.nat-hazards-earth-syst-sci.net/12/3229/2012/

of the residuals Z_r yields an estimated residual $\hat{Z}_r$ at location $\boldsymbol{x}_0$,

$$\hat{z}_r(\boldsymbol{x}_0) = \sum_{s=1}^{S} \lambda_s(\boldsymbol{x}_s) z_r(\boldsymbol{x}_s), \tag{7}$$

where λ_s are weights applied to the sample z_s at location $\boldsymbol{x}_s$, $s = 1, \dots, S$. They are determined by solving the kriging linear system

$$\begin{cases} \sum_{\beta=1}^{S} \lambda_\beta \gamma(\boldsymbol{x}_s - \boldsymbol{x}_\beta) + \mu_0 + \mu_1 Y(\boldsymbol{x}_s) = \gamma(\boldsymbol{x}_s - \boldsymbol{x}_0) \\ \sum_{\beta=1}^{S} \lambda_\beta = 1 \\ \sum_{\beta=1}^{S} \lambda_\beta Y(\boldsymbol{x}_\beta) = Y(\boldsymbol{x}_0) \end{cases} \tag{8}$$

for $s = 1, \dots, S$, where μ_0 and μ_1 are local coefficients required to express the local mean of the field Z as a linear combination of the auxiliary variable Y. Note that γ is the variogram of the residues. In practice it is an analytical expression of the variogram model chosen between four classical models called "spherical", "Gaussian", "exponential" and "power law" models (Chiles, 2001, p. 83). The variogram model parameters are considered as hyper-parameters of the RS schemes. Four sets of parameters are estimated. A set is computed by fitting a variogram model to the sample variogram (Eq. 9) expressing the spatial correlation of the field at discrete spatial lags δ,

$$\Gamma(\delta) = \frac{1}{2N(\delta)} \sum_{k=1}^{N(\delta)} [z(\boldsymbol{x}_k) - z(\boldsymbol{x}_k + \delta)]^2, \tag{9}$$

where $N(\delta)$ is the number of sample pairs separated by the spatial lag δ. The most appropriate model is determined via the cross-validation procedure detailed in Sect. 4.1.

In our context, the kriging interpolation scheme implies (Chiles, 2001, p. 154–171)

- the intrinsic hypothesis: the covariance function depends only on the lag δ; and

- the co-location of the random variable Z and of the covariate Y. The latter must also be known at the estimation points.

Moreover, if the random field Z is normally distributed, it is well known that (i) kriging performs better (unbiased estimation ensured), and (ii) the estimations of the two parameters of the normal distribution converge.

In practical applications, the Gaussianity of the field is rarely verified on real data. An appropriate anamorphosis is often a log transformation of Z. In addition, a continuous variogram model has to be fitted to the empirical variogram computed for a finite number of distance ranges.

2.2.2 Neural networks

Feed-forward one-hidden-layer neural networks (Bishop, 2006) are a flexible class of algorithms which allows for non-linear non-parametric interpolation. By non-parametric we mean an algorithm which does not assume a specific parametric shape, such as a polynomial, for instance. Given a set of covariates $y_1, \dots, y_d$ measured at a site $\boldsymbol{x}$ (for instance, we might consider the longitude, latitude and altitude at the site), the neural network will output $\hat{z}_1, \dots, \hat{z}_k$ (which could be the three GEV parameters predicted for this site). We implemented the feed-forward neural network in a standard way as follows. Let H be the number of hidden units. Each hidden unit a_h with $h = 1, \dots, H$ computes a linear combination of the covariates y_i, which is then non-linearly transformed by means of the hyperbolic tangent tanh,

$$a_h = \tanh\left(\sum_{i=1}^{d} v_{h,i} y_i + v_{h,0}\right), \tag{10}$$

where $v_{h,i}$ are the neural network weights linking the covariates to the hidden units. When the number of hidden units is selected according to the data at hand, the hyperbolic tangent confers the universal approximation property to the neural network (Hornik et al., 1989). Similarly to a_h, each output $\hat{z}_j$ is given by a linear combination of the hidden units a_h, which are then transformed by a function $g(\cdot)$ in order to ensure range constraint (e.g. the exponential could be used to impose positivity of the GEV scale parameter),

$$\hat{z}_j = g\left(\sum_{h=1}^{H} w_{j,h} a_h + w_{j,0}\right), \tag{11}$$

where $j = 1, \dots, k$, k being the number of outputs computed by the neural network. The weights $v_{h,i}$ in Eq. (10) and $w_{j,h}$ in Eq. (11), which we collect in ω, are determined by minimizing a cost function $\text{Err}(\omega)$. We chose to minimize the sum-of-squares error, so that for every site $\boldsymbol{x}_n$, $n = 1, \dots, N$, the predicted GEV parameters $\hat{z}_i(\boldsymbol{x}_n)$ matched as closely as possible the GEV parameters estimated at site $\boldsymbol{x}_n$:

$$\text{Err}(\omega) = \frac{1}{2} \sum_{n=1}^{N} \sum_{i=1}^{k} (z_i(\boldsymbol{x}_n) - \hat{z}_i(\boldsymbol{x}_n))^2. \tag{12}$$

The minimization of $\text{Err}(\omega)$ requires gradient-descent-type optimization algorithms, such as the conjugate gradient algorithm. Efficient computation of the gradient of $\text{Err}(\omega)$ is done via the back-propagation algorithm, see Rumelhart et al. (1986). In order to avoid local minima, the optimization is re-started 10 times with different initial weight values and the weights which give the smallest error are kept. To obtain error bounds from neural networks, the error function should be modified in order to account for our distributional assumptions on the error process, see Cawley et al. (2007).

3 Regionalization Schemes (RS)

In this section, we compare three regionalization schemes, each of them being a combination of one model to estimate the distribution of extreme observations (Sect. 2.1) and another to interpolate the parameters of this distribution (Sect. 2.2). As explained in the introduction, one reference scheme, RS_{ref}, is compared to two others, RS2 and RS3, each of them differing from RS_{ref} either in the extreme-value model or in the interpolation method.

In RS_{ref}, the 1-yr maxima model is coupled to kriging with an external drift. In RS2, kriging is replaced by a neural network interpolation. In RS3, the 1-yr maxima model of RS_{ref} is replaced by the POT method, whereas the interpolation model is the same as in RS_{ref} (kriging).

The 1-yr maxima model (RS_{ref}, RS2) leads to the estimation of three parameters (μ, σ, ξ) for the distribution of annual maxima by maximizing the log-likelihood function in Eq. (2). The POT method involves a distribution for exceedances with parameters u, β and ξ, where u is fixed a priori; only β and ξ are determined by log-likelihood estimation (Eq. 6) and then interpolated with kriging. The interpolation of each parameter is made by choosing the variogram model that best fits the sample variogram.

In scheme RS2, neural networks, implemented as described in Sect. 2.2.2, serve to predict the parameters of a constrained 1-yr maxima model, see Sect. 2.1.1. Given the covariates (longitude, latitude, altitude) at a site s, the neural network provides an estimate $\hat{\sigma}^{(s)}$ of the scale parameter. The location parameter is then estimated by $\hat{\mu}^{(s)} = a + b\hat{\sigma}^{(s)}$ and the tail index by a constant ξ. Then, the neural network is used to calculate a, b and ξ by minimizing the sum-of-squares error function between the parameters μ, σ and ξ estimated for each raingauge using the maximum likelihood estimation described in Sect. 1 and the neural network estimates with the constrained model.

Note that instead of performing this two-stage optimization procedure (maximizing the log-likelihood at each site to obtain local estimates of μ, σ and ξ and then minimizing the sum-of-squares error to determine the neural network weights), we could optimize the neural network weights ω by including them in the log-likelihood function of Eq. (4). In practice, this one-stage optimization approach did not improve our results much (the error decreased by less than 1%) and is much more computationally demanding, so we did not pursue it further.

4 Hyper-parameter selection and model comparison

For each of the regionalization schemes, we are faced with two modelling questions: (a) how to select the optimal complexity level of the models (such as the number of hidden units in the neural network) and (b) how to perform a fair comparison of the RS schemes. Cross- and double validation

are procedures by which we can answer questions (a) and (b), respectively.

4.1 Cross-validation

Typically, determining the optimal complexity level of models involves a bias/variance trade-off. The complexity level depends on a hyper-parameter, which is often proportional to the number of parameters in the model and controls its ability to approximate complex functions. For a neural network, the hyper-parameter is the number of hidden units. Too low a complexity level and the model will systematically deviate from the underlying process generating the data. This phenomenon is called bias. On the other hand, if the complexity level is too high, this might induce over-fitting, i.e. the model is fitting too precisely the data points. This means that the model has a high variance in the sense that it depends closely on the data set. Low bias, good approximation of the underlying process together with low variance, i.e. stability across different data sets, are desirable. As just described, these are conflicting properties. One way to get around this is by choosing the complexity level so as to minimize the expected error (the average error on new data from the same underlying process). The expected error can be decomposed into a sum of bias, variance plus noise (the noise is the residual part of the error that cannot be fit). It follows that the complexity level which minimizes the expected error minimizes simultaneously bias and variance. One way to estimate the expected error is via the cross-validation method. For example, we describe the hyper-parameter selection with m-fold cross-validation for the number of hidden units in a neural network. Let $\mathcal{H}$ contain all likely values for the number of hidden units, typically chosen as powers of 2. We proceed as follows:

1. Split the data set into m subsets, called folds: $L_1, \ldots, L_m$.

2. Leave one of the folds, L_ℓ, aside and denote by $L_{-\ell}$ the $m - 1$ remaining folds.

3. For each $H \in \mathcal{H}$, find the weights ω_H which minimize $E_{rr}(\omega)$ in Eq. (12) on $L_{-\ell}$.

4. Use again Eq. (12) to evaluate the performance $E_\ell(\omega_H)$ on the left-aside fold L_ℓ.

5. Return to step 2 until all m folds have been left aside in turn.

6. Choose the number of hidden units:

$$H^* = \arg\min_H \sum_{\ell=1}^{m} E_\ell(\omega_H)$$

This procedure is also applied to kriging. For the latter, $\mathcal{H}$ is the set of variogram models. The chosen model is the one minimizing the error between the variogram model estimated using the training set and applied on the independent samples of the test set. The variogram parameters are computed using a dichotomy procedure.

4.2 Double cross-validation

To perform a fair comparison among the schemes (question b above), we would need a test set, i.e. data that were used neither for training nor for hyper-parameter selection. To this end, a second loop in the cross-validation method is introduced. This double-loop cross-validation is called "double cross-validation". The second or outer loop goes as follows. The data are split into five folds. Keep four of these folds for hyper-parameter selection and model calibration and one last fold for testing. Alternate until all folds have been used for testing. The cross-validation described in steps 1 to 6 performs hyper-parameter selection on each of the possible combinations of the four folds, i.e. for each evaluation of the outer loop. Different models or schemes can be compared by looking at their performance on the test set (the combination of each of the five folds left aside in turn in the outer loop of the double cross-validation method). The threshold level selection in RS3 also raises a bias/variance trade-off. A higher threshold means lower bias with respect to the extremal behaviour, but a higher threshold also implies fewer observations available for the estimation and this results in larger estimation variance. The double cross-validation procedure described in this section can serve proper threshold selection as well. Double cross-validation can be computationally intensive, especially if there are many hyper-parameters to choose. Also, this method could result in different hyper-parameter values selected in each iteration of the outer loop of the double cross-validation. This could serve to check whether the model is stable across data sets. However, since the cross-validation error is a noisy estimate of the expected error, it is normal to observe some variability in the selection of hyper-parameters.

5 Application in the Cévennes-Vivarais region

In this section, the three previously defined schemes RS_{ref}, RS2 and RS3 are used to estimate the return levels for $T = 100\,\text{yr}$ in the Cévennes-Vivarais region. Section 5.1 presents the data. In Sect. 5.2 we compare these models using the double cross-validation method described in Sect. 4. We then plot return level maps according to each scheme in Sect. 5.3.

5.1 Data

This study relies on daily rainfall records from 1958 to 2008. The data have been collected by the French meteorological service Météo-France. In order to get longer rainfall series, rainfall records of neighbouring raingauges are concatenated if they do not overlap in time. Two raingauges are considered neighbours if their spatial distance is smaller than 1 km and their altitude difference is smaller than 100 m. This results in 225 rainfall series, whose locations are displayed in Fig. 5. The data points are distributed quite homogeneously in a $160 \times 200\,\text{km}^2$ area. The average of the minimum distance between raingauges is about 5 km. It has been noticed in Molinié et al. (2012) that the raingauge density increases with the altitude from 1 per $100\,\text{km}^{-2}$ in the range 0–500 m above sea level (a.s.l.) to 3 per $100\,\text{km}^{-2}$ above 900 m a.s.l.

5.2 Regionalization scheme comparison

The double cross-validation method described in Sect. 4.2 is used to perform hyper-parameter selection for each scheme and for evaluation and comparison of test data. In the outer loop, the initial 225 raingauges are split into five subsets, each of which contains data from 45 gauges. We perform m-fold cross-validation (with $m = 20$) at each step of the outer loop. Thus, hyper-parameter selection is performed on four of the folds, which makes for a data set of 180 gauges. In RS_{ref}, the hyper-parameter is the variogram model; it is the number of hidden units in RS2, while in RS3, a two-dimensional grid search determines as the first hyper-parameter the probability p to exceed the threshold and as the second hyper-parameter the variogram model.

Some additional details of the hyper-parameters choice for the three regionalization schemes are reported below:

- For RS_{ref} and RS3, four variogram models are considered (Chiles, 2001, p. 80–85): (i) spherical, of easy parametrization even if not physically based; (ii) exponential, related to Markov processes; (iii) Gaussian, expressing the covariance of infinite differentiable stationary random functions; (iv) power law, related to scale invariance.

- For RS2, the number of hidden units of the single-hidden-layer neural network is chosen in the set $\mathcal{H} = \{2, 4, 8, 16, 32\}$.

- For RS3, $p \in \{0.005, 0.010\ldots, 0.070\}$. These values cover an average number of events per year ranging from 1.8 to 25.

In order to compare the three regionalization schemes, a distance between local, i.e. at-site, and RSi (i standing for ref, 2 and 3) estimates of the 100-yr return level is introduced:
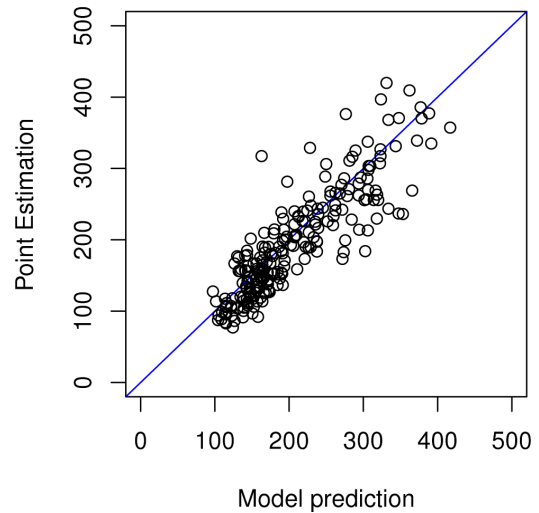
$$\Delta(\text{RS}i) = \sqrt{\sum_{\ell=1}^{5} \sum_{n \in L_\ell} \left(\hat{q}_{100,n}^{(\ell)} - \tilde{q}_{100,n} \right)^2},$$

where $\hat{q}_{100,n}^{(\ell)}$ is the 100-yr return level estimator provided by RSi on L_ℓ for the n-th station, and $\tilde{q}_{100,n}$ is the local estimator given by Eqs. (3) or (6). For the three regionalization
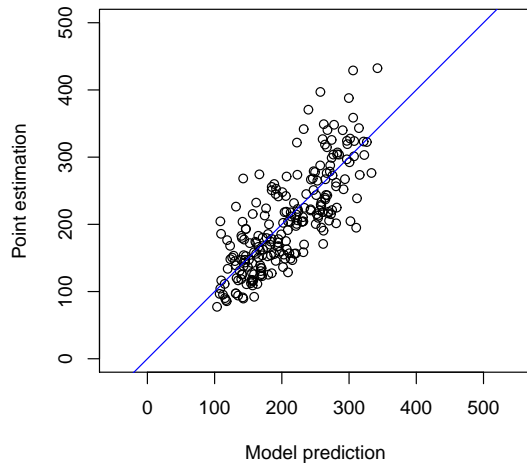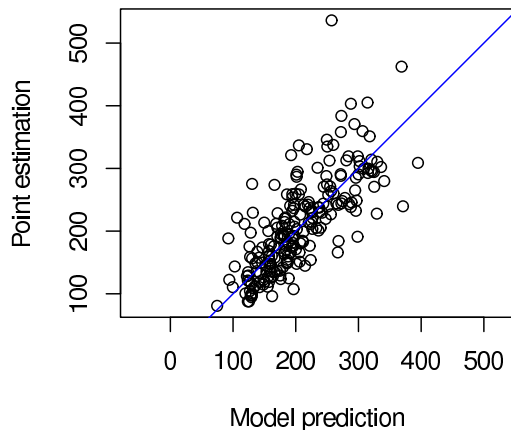
Fig. 2. Scatter plot of 100-yr return levels estimated locally versus the return levels predicted by RS_{ref}. Pearson's correlation is 0.73.

Fig. 4. Scatter plot of 100-yr return levels estimated locally versus the return levels predicted by RS3. Pearson's correlation is 0.87.

Fig. 3. Scatter plot of 100-yr return levels estimated locally versus the return levels predicted by RS2. Pearson's correlation is 0.79.

schemes, the following distances are obtained: $\Delta(RS_{ref}) = 58$, $\Delta(RS2) = 50$ and $\Delta(RS3) = 39$. According to the measure Δ, the best regionalization scheme is provided by RS3 for this data set. A visual assessment of Δ is displayed in Figs. 2, 3 and 4, where the return levels estimated with models RS_{ref}, RS2 and RS3, respectively, are also plotted versus the local estimates. It appears that, in all cases, the points are approximately aligned on the line $x = y$ and thus all the three schemes provide good regionalizations of the return levels. These results have been quantified using Pearson's correlation coefficient; RS3 gets the best result.

5.3 Return level maps

To plot the return level maps, we calibrate each RSi on the complete database, which contains 225 stations.

The choice of the hyper-parameters is done by performing cross-validation on the full database with the same number of folds and the same grid search as in the cross-validation procedure of Sect. 5.2. The selected hyper-parameters are displayed in Table 1. The return level maps for daily rainfall intensities characterized by a return period of $T = 100\,\mathrm{yr}$ are reported in Figs. 6, 7 and 8 for the RS_{ref}, RS2 and RS3 schemes, respectively. The three maps show similar patterns for the 100-yr return level of daily rainfall. South of the southwest–northeast diagonal, the return level is well correlated with the mountain slope until it reaches the Rhône River valley and the Mediterranean shore. In these regions, despite low altitude variations, the return level gradient is kept constant in direction and collinear to the Cévennes main-slope gradient. This is certainly due to the forcing mechanisms of heavy precipitation, such as the cold pool, which enhances air lifting well ahead of the mountain slopes (Nuissier et al., 2008; Ducrocq et al., 2008). North of this diagonal, the 100-yr return level is no more correlated with the altitude. This is the result of the well-known mountain-shading effect. It takes place at the north of the maps because air fluxes producing heavy rainfall at the daily time scale are from a sector from the south to southeast (Nuissier et al., 2011). A close look at the return level maps shows that over the plain and even over the foothills, all the scheme (RS_{ref}, RS2 and RS3) estimates are very similar.

The main differences between return levels yielded by the three schemes are located over the Cévennes mountain ridge. Each RSi estimates the maximum 100-yr return levels over the three higher peaks: Mont Aigoual, Mont Lozère and Serre de la croix de Beauzon. Molinié et al. (2012) discuss the mechanisms involving the heavy daily rainfall over these peaks. It turns out that Mont Lozère is the highest peak of the mountain range, whereas Mont Aigoual is the most
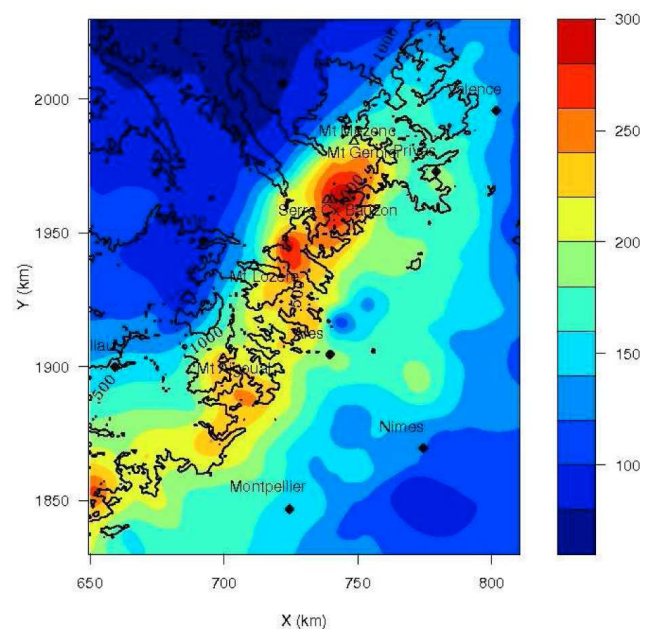
Table 1. Hyper-parameters selected by the inner cross-validation over the 225 stations.

	Number of hidden units H^*	Probability of exceedance p	Variogram models
RS$_{ref}$ (1-yr max. + kriging)	–	–	Spherical
RS2 (1-yr max. + neural networks)	4	–	–
RS3 (P.O.T. + kriging)	–	0.04	Power law
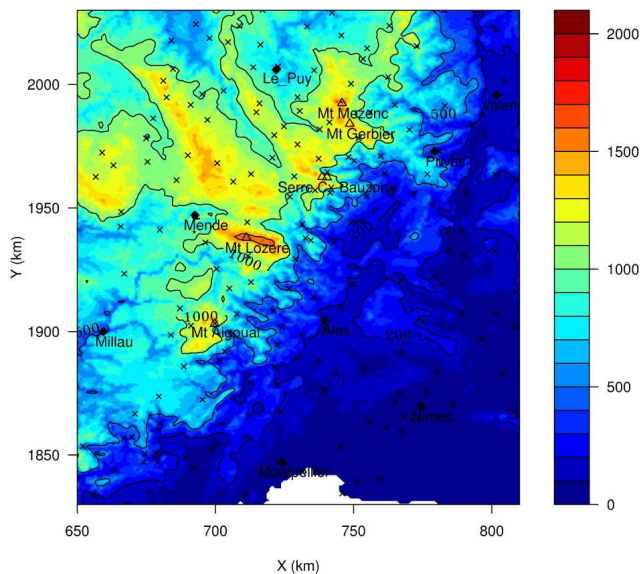
Fig. 5. Altitude map of the study region. The levels 200, 500 and 1000 m are also represented as iso-contour lines. The crosses represent the raingauge network composed by 225 daily gauges. Black diamonds represent cities, triangles the main mountain peaks.

Fig. 6. 100-yr return levels estimated by RS$_{ref}$ on the Cévennes-Vivarais data (mm).

exposed to southerly unstable air fluxes (see Fig. 5). Serre de la croix de Beauzon, even though less exposed to moist and warm air masses because of its northern location, features very deep (more than 600 m) valleys oriented west–east. Godart et al. (2011) show that stationary rain bands triggered on the shoulders of mountain ridges yielding to these deep valleys (Miniscloux et al., 2001; Anquetin et al., 2003) could conceivably enhance the rainfall regime.

Each of the schemes RS$_{ref}$, RS2 and RS3 predicts the maximum return levels over Serre de la Croix Beauzon and Mont Aigoual, even though of different magnitude. Another difference between the three schemes is for Mont Lozère, where RS3 does not predict a return level as high as over Mont Aigoual or Serre de la Croix de Beauzon, comparatively to RS$_{ref}$ and RS2. Molinié et al. (2012) display a map of medians of maximum daily intensities on which the Mont Lozère is at a lower level than Mont Aigoual and the Serre de la Croix de Beauzon, similarly to RS3. Moreover, the highest of the medians is 450 mm day^{-1} and the overall maximum of daily rainfall amounts ever measured over the region
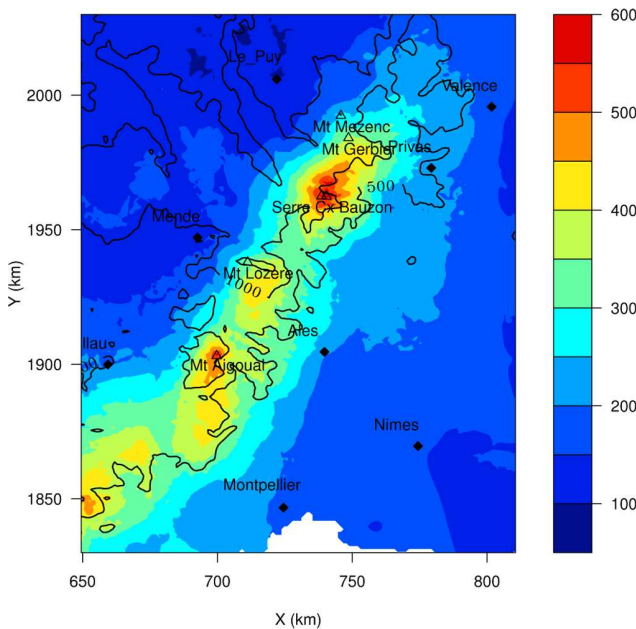
in the 50-yr period from 1958 to 2008 is 519 mm day^{-1}. Therefore, as the predicted maximum return levels are about 390 mm day^{-1} in RS$_{ref}$, about 450 mm day^{-1} in RS2 and 600 mm day^{-1} in RS3, this RS3 prediction is again closer to observations.

Ceresetti et al. (2010) identified a difference in extreme rainfall distributions between the plain and mountainous part of the study region. The return level distributions over the plain have heavier tails than the ones over the mountain and are well modelled by power law. All regionalization schemes are able to fit the extreme rainfall distribution of the plain region, while RS3 provides the best fit over the mountain range.

As stated in Sect. 2.2.1, the standard deviation due to kriging can be assessed thanks to the variogram structure function. The kriging coefficient of variation (standard deviation normalized by the 100-yr return level) is displayed in Fig. 9. It is shown that the uncertainty is relatively large, as it lies in the range of 25 to 40 % of the return level in the region of interest. This is a result of the inhomogeneity of the return
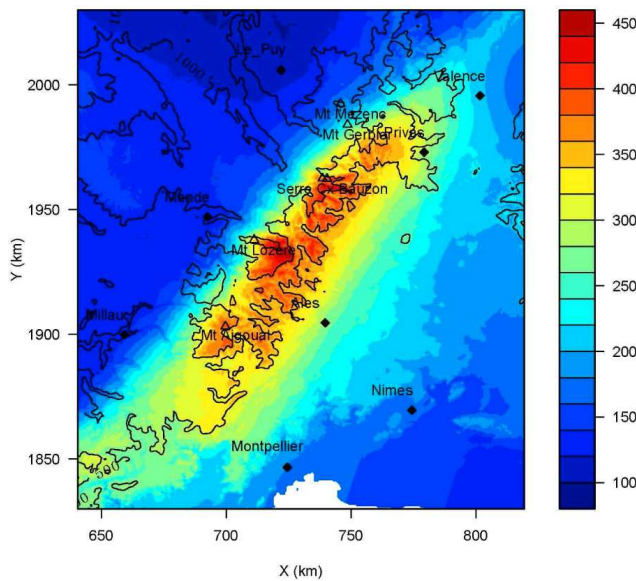
Fig. 7. 100-yr return levels estimated by RS2 on the Cévennes-Vivarais data (mm).

Fig. 9. Kriging error for RS3 evaluated as the kriging standard-deviation normalized by the 100-yr return levels.
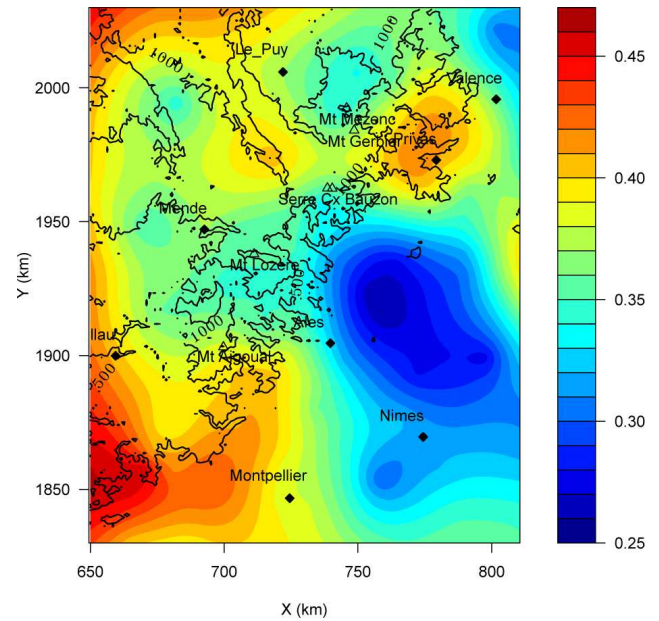

Fig. 8. 100-yr return levels estimated by RS3 on the Cévennes-Vivarais data (mm).

level field at the raingauge network resolution, of the order of 5 km. Therefore, the uncertainty on return levels is at least as large as the kriging uncertainty.

6   Discussion and conclusion

This paper's goal is to evaluate a widely used scheme for regional analysis of extreme rainfall intensities, the so-called

reference scheme RS$_{ref}$. The reference scheme couples a 1-yr model of heavy rainfall to the kriging interpolator. The comparison between the reference RS$_{ref}$ and the second regionalization scheme RS2 intends to assess the interpolation method so that RS2 is made of the 1-yr maxima point-rainfall analysis and of a neural-network interpolator. Comparing the third scheme RS3 and RS$_{ref}$ is meant to evaluate the extreme point-rainfall analysis method so that RS3 couples a POT selection of heavy rainfall to kriging.

The use of a double cross-validation algorithm allows assessing the point prediction of the regionalization scheme. In terms of the distance Δ and Pearson's statistic, RS3 performs better.

When comparing the maps of the 100-yr return level of daily rainfall yielded by the three schemes, it turns out that RS3 is again the closest to the current knowledge on extreme rainfall in the region. Since RS$_{ref}$ and RS2 provide similar results, the major issue in regional analysis of extreme rainfall seems to be the statistical model for point-extreme rainfall. It appears that the 1-yr maxima approach does not provide return levels as high as the POT method over the localized three mountain peaks of the region.

The two extreme-value models have different features: The 1-yr maxima approach is easy to use since it does not require any hyper-parameter selection. Conversely, POT is more flexible, since it sets a threshold to choose the number of extremal events. The two interpolation methods present different skills. Kriging is the best linear interpolation and uses the spatial correlation to infer the spatial distribution of the data; its implementation is limited to data respecting a number of constraints. Neural networks are a very flexible

method that provides similar results. It does not use spatial correlation and does not return the interpolation error, but the data are not submitted to the constraints of kriging. The two methods gave comparable results in this application; kriging seems to work better in densely gauged networks and when the field is well correlated.

Edited by: D. Lambert
Reviewed by: M. J. Pereira and one anonymous referee

References

Anquetin, S., Miniscloux, F., and Creutin, J. D.: Numerical simulation of orographic rainbands, J. Geophys. Res., 108, D8 doi: 10.1029/2002JD001593, 2003.

Bastin, G., Lorent, B., and Gevers, M.: Optimal estimation of the average rainfall and optimal selection of rain gauge locations, Water Resour. Res., 20, 463–470, 1984.

Bishop, C. M.: Pattern Recognition and Machine Learning, Information Science and Statistics, Springer, New-York, 2006.

Borga, M. and Vizzaccaro, A.: On the interpolation of hydrologic variables: formal equivalence of multiquadratic surface fitting and kriging, J. Hydrol., 195, 160–171, doi:10.1016/S0022-1694(96)03250-7, 1997.

Carreau J. and Girard, S.: Spatial extreme quantile estimation using a weighted log-likelihood approach, Journal de la Société Française de Statistique, 152, 66–83 , 2011.

Cawley, G. C., Janacek, G. J., Haylock, M. R., and Dorling, S. R.: Predictive uncertainty in environmental modelling, Neural Networks, 20, 537–549, 2007

Ceresetti, D., Molinié, G., and Creutin, J.: Scaling properties of heavy rainfall at short duration: a regional analysis, Water Resour. Res., 46, (W09531), doi:10.1029/2009WR008603, 2010.

Ceresetti, D., Anquetin, S., Molinié, G., Leblois, E., and Creutin, J.: Severity diagrams: a new approach for the multi-scale evaluation of extreme rainfall events, Weather Forecast., 27, 174–188, doi:10.1175/WAF-D-11-00003.1, 2011.

Chiles, J. P. and Delfiner, P.: Geostatistics – Model Spatial Uncertainty, Wiley Series in Probability and Statistics, John Wiley & Sons, Inc., New York, 1999.

Coles, S.: An Introduction to Statistical Modeling of Extreme Values, Springer Series in Statistics, Springer-Verlag, London, 2001.

Creutin, J. D. and Obled, C.: Objective analyses and mapping techniques for rainfall fields: an objective comparison, Water Resour. Res., 18, 413–431, 1982.

Cunnane, C.: Methods and merits of regional frequency analysis, J. Hydrol., 100, 269–290, 1988.

Daouia, A., Gardes, L., Girard, S., and Lekina, A.: Kernel estimators of extreme level curves, Test, 20, 311–333, 2011.

de Haan, L.: A spectral representation for max-stable processes, Ann. Probab., 12, 1194–1204, 1984.

de Haan, L. and Pickands, J.: Stationary min-stable stochastic processes, Probab. Theory Relat. Fields, 72 477–492, 1986.

Ducrocq, V., Nuissier, O., Ricard, D., Lebeaupin, C., and Thouvenin, T.: A numerical study of three catastrophic precipitating events over southern France, ii: Mesoscale triggering and stationarity factors, Q. J. Roy. Meteorol. Soc., 134, 131–145, 2008.

Gardes, L. and Girard, S.: A moving window approach for nonparametric estimation of the conditional tail index, J. Multivariate Analys., 152, 2368–2388, 2008.

Gardes, L. and Girard, S.: Conditional extremes from heavy-tailed distributions, an application to the estimation of extreme rainfall return levels, Extremes, 13, 177–204, 2010.

Godart, A., Anquetin, S., Leblois, E., and Creutin, J.: Contribution of rainfall associated with shallow banded orographic convection to the rainfall climatology of a mediterranean region, J. Appl. Meteorol. Climatol., 50, 2235–2246, doi:10.1175/JAMC-D-10-05016.1, 2011.

Goovaerts, P.: Geostatistics for Natural Resources Evaluation, Oxford University press, New York, 1997.

Goovaerts, P.: Geostatistical approaches for incorporating elelevation into the spatial interpolation of rainfall, J. Hydrol., 228, 113–129, 2000.

Hosking, J. and Wallis, J.: Regional Frequency Analysis, Cambridge University Press, 224 pp., 1997.

Hornik, K., Stinchcombe, M., and White, H.: Multilayer feedforward networks are universal aproximators, Neural Networks, 2, 359–366, 1989.

Frei, C. and Schär, C.: A precipitation climatology of the Alps from high-resolution rain-gauge observations, Int. J. Climatol., 18, 873–900, 1998.

Kharin, V. V. and Zwies, W. F.: Estimating extremes in transient climate change simulations, J. Climate, 18, 1156–1173, 2005.

Lebel, T. and Bastin, G.: Variogram identification by the mean-squared interpolation error method with application to hydrologic fields, J. Hydrol., 77, 31–56, 1985.

Miniscloux, F., Creutin, J. D, and Anquetin, S.: Geostatistical analysis of orographic rainbands, J. Appl. Meteorol., 40, 1835–1854, 2001.

Molinié, G., Ceresetti, D., Anquetin, S., Creutin, J. D., and Boudevillain, B.: Rainfall regimes in a mediterranean mountainous region. J. Appl. Meteorol. Climate, 51, 3 doi:10.1175/2011JAMC2691.1, 2012.

Myers, D. E.: Spatial interpolation: an overview, Geoderma, 62, 17–28, 1994.

Neppel, L., Pujol, N., and Sabatier, R.: A multivariate regional test for detection of trends in extreme rainfall: the case of extreme daily rainfall in the French Mediterranean area, Adv. Geosci., 26, 145–148, doi:10.5194/adgeo-26-145-2011, 2011.

Norbiato, D., Borga, M., Sangati, M., and Zanon, F.: Regional frequency analysis of extreme precipitation in the eastern italian

alps and the august 29, 2003 flash flood, J. Hydrol., 345, 149–166, doi:10.1016/j.jhydrol.2007.07.009, 2007.

Nuissier, O., Ducrocq, V., Ricard, D., Lebeaupin, C., and Anquetin, S.: A numerical study of three catastrophic precipitating events over southern France, i: Numerical framework and synoptic ingredients, Q. J. Roy. Meteorol. Soc., 134, 111–130, 2008.

Nuissier, O., Joly, B., Ducrocq, V., and Arbogast, P.: A statistical downscaling to identify the large scale circulation patterns associated with heavy precipitation events over southern France, Q. J. Roy. Meteorol. Soc., 137, 1812–1827, doi:10.1002/qj.866, 2011.

Prudhomme, C. and Reed, D.: Mapping extreme rainfall in a mountainous region using geostatistical techniques: a case study in scotland. Int. J. Climatol., 19, 1337–1356, 1999.

Ramos, M. H., Creutin, J. D., and Leblois, E.: Visualization of storm severity, J. Hydrol., 315, 295–307, 2005.

Ripley, B. D.: Spatial Statistics, John Wiley & Sons Inc., New-York, 1981.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J.: Learning internal representations by error propagation, Parallel Distributed Processing: Explorations in the Macrostructure of Cognition, Bradford Books, 1986

Snell, E. S., Gopal, S., and Kaufmann, R. K.: Spatial interpolation of surface air temperatures using artificial neural networks: evaluating their use for downscaling GCMs, J. Climate, 13, 886–895, 2000.

Vié, B., Molinié, G., Nuissier, O., Vincendon, B., Ducrocq, V., Bouttier, F., and Richard, E.: Hydro-meteorological evaluation of a convection-permitting ensemble prediction system for Mediterranean heavy precipitating events, Nat. Hazards Earth Syst. Sci., 12, 2631–2645, doi:10.5194/nhess-12-2631-2012, 2012.

Wackernagel, H.: Multivariate Geostatistics, 2nd Edn., Springer, Berlin, 1998.

Zwiers, F. W. and Kharin. V. V.: Changes in the extremes of the climate simulated by CCC CGM2 under CO_2 doubling, J. Climate, 11, 2200–2223, 1998.