

1 **Clustering of maxima:**

2 **Spatial dependencies among heavy rainfall in France**

3 **ELSA BERNARD, PHILIPPE NAVEAU, \* MATHIEU VRAC**

*Laboratoire des Sciences du Climat et de l'Environnement, CNRS-CEA-UVSQ, Gif-sur-Yvette, France*

4 **OLIVIER MESTRE**

*Direction de la Production, Météo-France, Toulouse, France*

---

\* *Corresponding author address:*

E-mail: naveau@lsce.ipsl.fr

6 One of the main objectives of statistical climatology is to extract relevant information hidden  
7 in complex spatial-temporal climatological datasets. To identify spatial patterns, most well-  
8 known statistical techniques are based on the concept of intra and inter clusters variances  
9 (like the k-means algorithm or EOF's). As analyzing quantitatively extremes like heavy  
10 rainfall has become more and more prevalent for climatologists and hydrologists during  
11 those last decades, finding spatial patterns with methods based on deviations from the mean,  
12 i.e. variances, may not be the most appropriate strategy in this context of studying such  
13 extremes. For practitioners, simple and fast clustering tools tailored for extremes have been  
14 lacking. A possible avenue to bridging this methodological gap resides in taking advantage  
15 of multivariate extreme value theory, a well-developed research field in probability, and to  
16 adapt it to the context of spatial clustering. In this paper, we propose and study a novel  
17 algorithm based on this plan. We compare and discuss our approach with respect to the  
18 classical k-means algorithm throughout the analysis of weekly maxima of hourly precipitation  
19 recorded in France (Fall season, 92 stations, 1993-2011).

# 1. Introduction

Clustering algorithms are routinely run to summarize and visualize important spatial and/or temporal patterns in climate sciences. For example, Stefanon et al. (2012) proposed a method for defining and classifying heatwave events in the Euro-Mediterranean region. Another example corresponds to the use of the k-means algorithm (e.g., Hastie et al. 2009) to provide the different phases of the North Atlantic Oscillation (NAO) (e.g., Cassou et al. 2004). The k-means method is based on the choice of a metric, classically related to a Euclidean (L2) norm, i.e. deviations from the mean behavior like intra and inter variances. In a nutshell, the k-means principle is to find clusters such that the variance within each cluster is minimized. This makes sense for applications that aim at identifying patterns with respect to mean behaviors. In particular, it is ideally suited when the variable of interest follows a mixture of normal distributions because Gaussian random vectors are fully characterized by their mean vectors and their covariances matrix (e.g., von Storch and Zwiers 2002). Coming back to the NAO example, it seems reasonable to implicitly assume that winter monthly sea level pressure means (the k-means inputs in Cassou et al. (2004)) can be represented by a mixture of normal distributions. The Central Limit Theorem (e.g., see page 35 of von Storch and Zwiers 2002) insures the normality of such means within each weather regime. But other atmospheric variables like hourly precipitation amounts may strongly differ from being Gaussian or even a Gaussian mixture. Precipitation intensities take only non-negative values, their probability densities are skewed and their extremes may be heavy tailed (e.g, Katz et al. 2002). In such instances, it is still possible to implement the k-mean algorithm, but one can wonder if the clusters are interpretable when means and variances become ambiguous summaries for skewed and heavy-tailed probability densities. Does this imply that clustering algorithms like the k-means should be discarded? If so, what could be a statistically sound alternative? Answering those types of questions within the context of analyzing maxima is important (e.g., Plaut et al. 2001). Putting into light new spatial or temporal patterns for maxima may help the understanding of climate extremes, provide

useful statistical tools for impact studies and also avoid some erroneous interpretations of extreme events analysis derived from inappropriate clustering techniques.

The statistical analysis of maxima is based on the well-developed Extreme Value Theory (EVT) (e.g., see Resnick 2007; de Haan and Ferreira 2006; Beirlant et al. 2004; Coles 2001). This theory indicates that the Generalized Extreme Value distribution (GEV) represents the ideal candidate for modeling the marginal distribution of block-maxima (as opposed to the peaks-over-threshold approach). This probabilistic framework has been applied in climate studies (e.g., see Kharin et al. 2007). In a spatial context, multivariate EVT also provides a theoretical blueprint to represent dependencies among maxima recorded at different locations. Coles et al. (1999) gives an overview of such dependence measures. For example, it is possible to adapt the variogram, a well-known distance used in geostatistics (e.g., Wackernagel 2003), to EVT. This special variogram called a F-madogram, see Section 2 for details, was proposed by Cooley et al. (2006) and Naveau et al. (2009) who studied a non-parametric approach for estimating pairwise dependence among maxima. It was applied to precipitation maxima measured in Belgium (Vannitsem and Naveau 2007). Those past studies indicate that it is possible to “measure” the distance between two time series of maxima recorded at two different locations and that this measure, the F-madogram, is in compliance with EVT and differs from classical measures of variability like the variance used in the k-means algorithm.

The aim of the present work is to develop a clustering algorithm for maxima based on the F-madogram. A natural strategy could be to simply replace the L2-norm (the variance) in the k-means algorithm by the F-madogram distance. But in the k-means algorithm, new centroids at each time step are obtained by averaging of observations within each cluster. Averages of normally distributed observations remain Gaussian, but averages of GEV distributed maxima do not stay GEV distributed. This poses a problem in terms of interpretability within the EVT framework and leads us to work with the Partitioning Around Medoids (PAM) clustering algorithm proposed by Kaufman and Rousseeuw (1990). Similar

to k-means, PAM is a partitioning algorithm that divides datasets into groups and aims at minimizing an overall distance. Whereas the k-means algorithm represents each cluster center by its mean, the PAM algorithm looks for representative objects (called medoids). This implies that maxima remain maxima and no smoothing (averaging) is performed within PAM.

Our paper is organized as follows. Section 2 recalls some theoretical background about bivariate EVT and makes the necessary links between EVT and the PAM clustering algorithm. Rainfall maxima over the French region are spatially clustered in Section 3. Section 4 leads to a discussion.

## 2. Algorithm description

In terms of notations, the random variable  $M_i$  generically represents weekly maxima of hourly precipitation located at weather station  $i$ . Dividing a region into coherent spatial patterns is a classical endeavor in climatology. To be able to cluster points, we need to assess the strength of the spatial dependence between the maximum  $M_i$  and the maximum  $M_j$ , i.e. how to model their pairwise distribution. Following the mathematical framework of multivariate EVT (e.g., see Resnick 2007; de Haan and Ferreira 2006; Beirlant et al. 2004; Coles 2001; Fougères 2004), it is reasonable to assume that the bivariate vector  $(M_i, M_j)^T$  follows a bivariate EVT distribution

$$\mathbb{P}(M_i \leq u; M_j \leq v) = \exp \left[ -V_{ij} \left( \frac{-1}{\ln F_i(u)}, \frac{-1}{\ln F_j(v)} \right) \right], \quad (1)$$

where  $F_i(u) = \mathbb{P}(M_i \leq u)$  represents the marginal distribution of  $M_i$  and the extremal dependence function  $V_{ij}(\cdot, \cdot)$  is defined as

$$V_{ij}(x, y) = 2 \int_0^1 \max \left( \frac{w}{x}, \frac{1-w}{y} \right) dH_{ij}(w)$$

where  $H_{ij}(\cdot)$  corresponds to any distribution function on  $[0, 1]$  such that its expectation equals to 0.5. This class of distributions arises as the natural non-degenerated limit of rescaled i.i.d.

componentwise maxima of random vectors (de Haan and Ferreira 2006; Resnick 2007). At this stage, such a definition may appear rather obscure and some light can be shed on (1) by looking at the special case where  $u = v$ . Because of the definition of  $V_{ij}$ , we have  $V_{ij}(x, x) = V_{ij}(1, 1)/x$  and it follows from (1) (e.g., Naveau et al. 2009) that

$$\mathbb{P}(M_i \leq u; M_j \leq u) = [\mathbb{P}(M_i \leq u) \mathbb{P}(M_j \leq u)]^{V_{ij}(1,1)/2}. \quad (2)$$

The scalar  $V_{ij}(1, 1)$ , called the “extremal coefficient”, gives partial but paramount information about the degree of dependence between  $M_i$  and  $M_j$  (e.g., see Schlather 2002; Schlather and Tawn 2003). If those two variables are independent, then Equation (2) implies that  $V_{ij}(1, 1) = 2$ . If they are equal, then we have  $V_{ij}(1, 1) = 1$ . Hence, the extremal coefficient can go from one (complete dependence) to two (full independence), and therefore it can capture relevant information about the dependence strength. Another way to interpret the extremal coefficient is to make the connection with a specific variogram of order one. A variogram of order  $p$  is defined as the moment of order  $p$  of the difference between  $M_i$  and  $M_j$ ,  $\mathbb{E}|M_i - M_j|^p$  (e.g., see Wackernagel 2003). Cooley et al. (2006) showed that the “F-madogram” defined as

$$d_{ij} = \frac{1}{2} \mathbb{E}|F_i(M_i) - F_j(M_j)| \quad (3)$$

can be expressed in terms of the extremal coefficient

$$d_{ij} = \frac{1}{2} \frac{V_{ij}(1, 1) - 1}{V_{ij}(1, 1) + 1}. \quad (4)$$

If the two weather stations  $i$  and  $j$  are close to each other and local conditions at both places are basically identical, the precipitation maxima  $M_i$  and  $M_j$  should be similar and  $d_{ij}$  should be close to zero. Equation (4) tells us that the extremal coefficient should be near one. Conversely, if the two locations  $i$  and  $j$  are far away from each other and can be considered as independent, then the extremal coefficient is close to two and Equation (4) implies that the madogram should be equal to 1/6. Besides being an interpretable distance, another advantage of the madogram resides in the fact that its value can be easily inferred

in a non-parametric fashion. The distance  $d_{ij}$  in (3) corresponds to an expectation and can be inferred as a sample mean. Given a sample of maxima  $(M_i^{(t)}, M_j^{(t)})^T$  recorded at two locations  $i$  and  $j$  and at  $T$  different time units, then the definition of the madogram  $d_{ij}$ , provides a natural non-parametric estimator

$$\hat{d}_{ij} = \frac{1}{2T} \sum_{t=1}^T |\hat{F}_i(M_i^{(t)}) - \hat{F}_j(M_j^{(t)})| \quad (5)$$

where  $T$  is the bivariate sample length and  $\hat{F}_i$  is the empirical distribution function

$$\hat{F}_i(u) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}_{\{M_i^{(t)} \leq u\}},$$

where  $\mathbf{1}_{\{M_i^{(t)} \leq u\}}$  represents the indicator function of the event  $\{M_i^{(t)} \leq u\}$ . By plugging  $\hat{d}_{ij}$  in Equation (4), an estimator of the extremal coefficient  $V_{ij}(1, 1)$  is automatically deduced. For the theoretical properties of those estimators, we refer to Cooley et al. (2006) and Naveau et al. (2009).

The definition of the madogram  $d_{ij}$  also emphasizes an essential point concerning the interpretation of our results. Applying to the random variable  $M_i$  its own distribution  $F_i(u) = \mathbb{P}(M_i \leq u)$  in Equation (3) makes the variable  $F_i(M_i)$  uniformly distributed. The same is true for  $F_j(M_j)$ . This implies that the madogram (or equivalently the extremal coefficient) does not depend on the marginal laws and, consequently, it cannot provide information about how much rain can fall at a specific site. It is a dimensionless concept and it only describes the dependence strength. The term copula is often used in the statistical literature to describe this decoupling between margins and the dependence function. This decoupling between the marginals and the dependence strength will be beneficial when we will have to interpret the map of our clustered maxima. To infer the madogram values, we just need to plug in the empirical versions of  $F_j$  and  $F_i$  and compute an average, see Equation (3) and Appendix A. This means that we don't need to fit a GEV at each weather station. This saves computational time and allows weaker modeling assumptions than imposing GEV marginals. Naveau et al. (2009) showed that the dependence  $V(., .)$  can be estimated from the

empirical madogram estimator as the sample size and the block size increase (see Proposition 4 of Naveau et al. (2009)). So, it was not assumed that maxima were GEV distributed, but they only belong to the domain of attraction of max-stable distribution.

Having at our disposal the distance  $d_{ij}$  that is tailored from maxima motivated by (1), we have to choose a clustering algorithm. As already stated in the introduction, the k-means algorithm creates cluster centers by averaging points within a cluster. Such averaging operation destroys the max-stable property encapsulated in (1), since average of more than one maximum is no longer a maximum. As an attractive alternative, the Partitioning Around Medoids (PAM) algorithm proposed by Kaufman and Rousseeuw (1990) is known to preserve the observations at hand, a weekly maximum remains a weekly maximum. The PAM algorithm divides a dataset of  $N$  objects into  $K$  clusters. Three pre-processing steps are needed before implementing PAM. First, the distance matrix  $\{d_{ij}\}$  defined by (3) has to be computed. Second, the number of clusters  $K$  has to be chosen and third, to initialize the PAM algorithm, an initial set of  $K$  medoids has to be randomly selected, i.e. a group of  $K$  randomly chosen stations. Then, the PAM algorithm can be run as follows.

- (A) Form  $K$  clusters by assigning every point to its closest medoid.
- (B) For each cluster, find the new medoid for which the total intra-cluster distance based on  $d_{ij}$  is minimized.
- (C) If at least one medoid has changed, then go back to (A), otherwise end the algorithm.

In summary, PAM proceeds by moving around  $K$  medoids while trying to make the total intra-cluster distance as small as possible. As mentioned previously, the “centers” of the cluster, the so-called medoids, still represent valid weekly precipitation maxima at each step of the algorithm. Consequently, the distance  $d_{ij}$  can always be interpreted via (4) at any stage within the PAM algorithm.

To choose a relevant number  $K$  of clusters and to assess if a weather station is well classified, Rousseeuw (1986) developed the so-called “silhouette coefficient” that compares

167 cluster tightness (small  $d_{ik}$  within the cluster  $k$ ) with cluster dissociation (see  $\delta_{i,-k}$  defined  
168 below). After running the PAM algorithm with a given  $K$ , each location  $i$  is associated with  
169 a medoid  $k$ . The silhouette coefficient for the weather station  $i$  is defined as follows

$$s_i(K) = 1 - (d_{ik} / \delta_{i,-k}),$$

170 where  $d_{ik}$  represents the intra-cluster distance between medoid  $k$  and station  $i$  and  $\delta_{i,-k}$   
171 corresponds to the smallest distance between station  $i$  and all the other medoids but  $k$ .  
172 For the PAM algorithm procedure,  $s_i(K)$  necessarily belongs to the interval  $[-1, 1]$ . If  
173  $s_i(K) \approx 1$ , it means that the intra-cluster distance is much smaller than the inter-cluster  
174 distances. Consequently, the maximum  $M_i$  can be considered as well classified. In contrast,  
175 if  $s_i$  is near zero, the clustering is viewed as non-informative, meaning that  $M_i$  could have  
176 been in an other cluster as well with the same relevancy. To summarize the quality of a  
177 partitionning into  $K$  clusters, one can derive the average silhouette coefficient

$$\bar{s}(K) = \frac{1}{N} \sum_{i=1}^N s_i(K), \quad (6)$$

178 or other statistics from the set  $\{s_1(K), \dots, s_N(K)\}$ . Such summaries will be used in our  
179 application. To implement our approach, a package for the open-source statistical R software  
180 is available the homepage of the second author.

### 181 3. Applications to French precipitation maxima

182 Here we focus on weekly maxima of hourly precipitation at 92 French stations during the  
183 the Fall season (SON) from 1993 to 2011. They were provided by the French meteorological  
184 service, Météo-France. The stations were chosen in function of their quality and to have  
185 a fairly homogeneous coverage of France. To avoid dealing with zero's and in order to be  
186 consistent with EVT, very small values of precipitation (rainfall amounts below 3mm) were  
187 discarded (qqplots and other diagnostics, available upon request, were used to not reject

the hypothesis of GEV distributed marginals). Before applying our PAM approach to those data, we have applied the classical k-means algorithm to those rainfall maxima.

Panel A of Figure 1 displays the outputs into five clusters. The difference between the left and right maps in Panel A is due to the nature of the k-means inputs, raw maxima (left) and their logarithm (right). This discrepancy between the two maps indicates that the choice of the marginal laws has a strong effect on the clustering outputs. For example, rainfall recorded in Brittany along the Atlantic coast is very different (in a distributional sense) from precipitation measured in Corsica, an island in the Mediterranean Sea. This emphasizes that it is unreasonable to “compare apples and oranges”, i.e. to perform clustering on times series with different marginal laws. Quantitatively, this can be assessed by fitting a GEV probability distribution function defined by  $G(x) = \exp\{-[1 + \xi(\frac{x-\mu}{\sigma})]_+^{-1/\xi}\}$  where the real  $\mu$  is the location parameter,  $\sigma$  the positive scale parameter and  $\xi \in \mathbb{R}$  the shape parameter.

Panel B of Figure 1 displays the scale and shape GEV parameters inferred for each location (by probability weighted moments, (e.g., see Dielbolt et al. 2008)), respectively the left and right maps. Panel B indicates well-known climatological results. Fall heavy rainfall intensities are located near the Mediterranean coast while the center and northern part of France have milder extreme precipitation intensities.

Comparing the left of panel A with panels B suggests that the south east region with heavy rainfall, i.e. with large GEV parameters, influences the k-means algorithm. This makes sense because having large scale and shape parameters corresponds to strong variability and the variance is the key clustering criterion for the k-means algorithm. But this also means that this clustering attempts to answer two different questions that may not be linked. The question regarding the intensity of rainfall at a given weather station (a univariate concept based on the marginal distribution) is mixed with the inquiry about the strength is the spatial relationship between two neighboring weather stations (a bivariate distributional concept). This is an undesirable trait that renders the interpretation of those clusters extremely complex.

As previously mentioned, our proposed PAM approach based on the F-madogram is marginal free and implemented via a non-parametric approach. This second point implies that we do not need to fit a GEV distribution at each weather station. This reduces computational time and removes a source of uncertainty (it is always difficult to infer accurately a shape GEV parameter and its associated confidence intervals).

To visualize the differences between the classical k-means approach and our proposed method based on the PAM algorithm, Figure 2 compares the clustering outputs for both methods, maps on the left for our PAM approach and on the right for the k-means algorithm applied on log-precipitation maxima (to reduce the margins problem). Each panel, A, B and C, corresponds to a different number of clusters  $K = 2, 5$  and  $7$ . Each medoid has a diamond shape with a black contour. Each station is linked to its medoid by a grey line if its silhouette coefficient is significant. Otherwise it simply appears as a circle (instead of a diamond).

To determine the 90% confidence level for a fixed  $K$ , our PAM algorithm was rerun after randomly sampling our rainfall data in order to break any spatial dependence. This scheme was repeated 20 times and the 95% quantile from this sample of 20 average silhouette coefficients. At this stage, it is important to emphasize that the k-means and PAM algorithms run without any geographical information, but only rainfall records. So, finding coherent spatial structures from only rainfall measurements was not automatic. From Figure 2, it appears that the PAM and k-means approaches provide strikingly different clusters. This may be one of the most important messages of this work. Choosing a clustering method and a specific metric can have an enormous impact on clustering patterns and lead to potentially different or even conflicting climatological interpretations. For example, PAM with  $K = 2$  (Panel A) divides France into a north-south fashion along the Loire valley line, while the k-mean roughly reproduces the main characteristic of the GEV parameter, see Panel B of Figure 1. This feature is linked to rainfall intensities but not necessarily to spatial precipitation dependencies. For  $K = 5$  (Panel B), PAM isolates the west region above Bordeaux (blue color) from the central region (around Paris), while the k-means emphasizes Corsica

and two Mediterranean cities (blue color), again stressing rainfall intensities. As the number of clusters increases ( $K=7$  in Panel C), sharper regional features appear and are geographically coherent. For  $K = 7$ , k-means starts to break down a little bit by creating clusters without any spatial structure, see the isolated four light orange points in Brittany.

In the south of France, extreme rainfall events in the Fall are usually caused by southern winds, forcing warm and moist air to interact with mountainous areas of Pyrénées, Cévennes and Alps, resulting in severe thunderstorms. A systematic inventory of those situations over 1958-1994 period was studied by Jacq (1994). Those events may be very local in some cases, but often affect one third to one half of the mediterranean coastal area. Large scale extreme events, occurring on both Corsica and Var (around Toulon) or in the Alpes Maritimes (around Nice) regions are very likely to affect the Rhône valley, the Alps and even further west to Montpellier. The "Corsica-Nice-Toulon" cluster does not seem to be very justified climatologically. The Millau, Mende and Carcassonne series should belong to the Mediterranean cluster rather than to the "South West Cluster" (Agen Medoid), which is the case in PAM with  $K=7$ . In the north of France, heavy rainfall is often produced by mid-latitude perturbations. Depending on their tracks, some affect Brittany, while others only influence the north of France and Paris. The very large northern cluster produced by k-means ( $K=2$ ,  $K=5$ ) is not consistent with our understanding of synoptic variability, while PAM clusters can be interpreted easily. Isolating central and eastern clusters (PAM,  $K=7$ ) is coherent with climatic and topographic features.

To complete this example, it is natural to wonder what would be the most appropriate number of clusters. Each boxplot in Figure 3 summarizes the silhouette coefficients distribution for a given  $K$  varying from 2 to 16. Applying Equation (6), the average silhouette coefficient is represented by the solid black line. The dotted line with grey diamonds corresponds to the upper 95% level obtained after randomly reshuffling our precipitation data. This breaks down the spatial structure (figures available upon request) and silhouette coefficients below such thresholds are considered as non-significant, see small circles in Figure 2. Figure

3 does not bring a clear winner here as the largest average silhouette coefficients are very close around 0.12 (for  $K=2$ ) and 0.11 (for  $K=5$ ). In regards with the maps displayed in Figure 2, the spatial patterns for  $K=5$  or even  $K=7$  indicates that the clusters are coherent with geographical features. To keep the maps interpretable and avoid overparametrization, choosing around  $K = 5$  represents a good compromise. Although significant, the silhouette coefficients in this example are not very large and this may be explained by the variable under study. Extreme precipitation events certainly have short range spatial dependences. A finer spatial resolution should give stronger localized structures but such precipitation data at the hourly scale and of high quality are difficult to find at the scale of a country.

## 4. Discussion

By combining two statistical methods, the PAM algorithm with the F-madogram, a simple clustering algorithm for maxima was proposed and studied. Besides being in compliance with EVT, it offers a different perspective for those who are interested in identifying spatial or temporal patterns in statistical climatology. As an illustration, a partitioning of the French region with respect to Fall precipitation maxima was obtained. This clustering strongly differs from a variance based approach like the k-means algorithm. This opens new challenges concerning the analysis of heavy rainfall over France and elsewhere. At the hydrological basin scale, our approach could complement the well-known Regional Frequency Analysis (RFA, see e.g. Gaume et al. (2010)) performed in hydrology to find homogenous regions with respect to extreme events. Despite its name, RFA does not take into account any dependence among maxima. It is a method solely based on marginal probability densities. In contrast, our approach is fully decoupled from the margins and so, it could ideally supplement RFA by making regions based on the dependence strength among maxima.

Taking different block sizes (say a month instead of a week) with different precipitation types (say daily instead of hourly) may provide different clustering patterns. This could

lead to new avenues to explore clustering maps, especially with respect to more traditional approaches.

Another possible direction could be to apply our method within a context of dimension reduction. Currently, very few statistical EVT approaches exist to deal with this issue.

Finally, our approach is computationally fast and could be applied to large datasets like global climate models outputs. For example, it could be used to compare spatial clustering of yearly maxima (or minima) of daily temperatures.

### *Acknowledgments.*

Part of this work has been supported by the EU-FP7 ACQWA Project ([www.acqwa.ch](http://www.acqwa.ch)), by the PEPER-GIS project, by the ANR-MOPERA project, by the ANR-McSim project and by the MIRACCLE-GICC project. The authors would like to thank the contributors of the statistical freely available software R <http://www.r-project.org/>. In particular, we have extensively used the *pam* function from Peter Rousseeuw, Anja Struyf and Mia Hubert, see the *cluster* library (package).

## REFERENCES

- 310 Beirlant, J., Y. Goegebeur, J. Segers, and J. Teugels, 2004: *Statistics of extremes: Theory*  
 311 *and applications*. John Wiley & Sons: New York.
- 312 Cassou, C., L. Terray, J. W. Hurrell, and C. Deser, 2004: North atlantic winter climate  
 313 regimes: Spatial asymmetry, stationarity with time, and oceanic forcing. *JOURNAL OF*  
 314 *CLIMA TE*, **17**, 1055–1068.
- 315 Coles, S., 2001: *An Introduction to Statistical Modeling of Extreme Values*. 208 pp.
- 316 Coles, S., J. Heffernan, and J. Tawn, 1999: Dependence measures for extreme value analysis.  
 317 *Extremes*, **2**, 339–365.
- 318 Cooley, D., P. Naveau, and P. Poncet, 2006: *Dependence in Probability and Statistics*, Lecture  
 319 Notes In Statistics, Vol. 187, chap. Variograms for spatial max-stable random fields, 373–  
 320 90. New York, Springer.
- 321 de Haan, L. and A. Ferreira, 2006: *Extreme Value Theory, An Introduction*. Springer Series  
 322 in Operations Research and Financial Engineering.
- 323 Dielbolt, J., A. Guillou, P. Naveau, and P. Ribereau, 2008: Improving probability-weighted  
 324 moment methods for the generalized extreme value distribution. *RevStat*, **6 (1)**, 33–50.
- 325 Fougères, A.-L., 2004: Multivariate extremes. *Extreme Values in Finance, Telecommunica-*  
 326 *tions, and the Environment*, 373–388.
- 327 Gaume, E., L. Gaál, A. Viglione, J. Szolgay, S. Kohnová, and G. Blöschl, 2010: Bayesian  
 328 mcmc approach to regional flood frequency analyses involving extraordinary flood events  
 329 at ungauged sites. *J. Hydrol.*, doi:10.1016/j.jhydrol.2010.01.008.

- 330 Hastie, T., R. Tibshirani, and J. Friedman, 2009: *The Elements of Statistical Learning: Data*  
331 *Mining, Inference, and Prediction*. 2d ed., Springer Series in Statistics.
- 332 Jacq, V., 1994: Inventaire des situations à précipitations diluviennes, languedoc-  
333 roussillon/paca/corse, période 1958-1994. Publication du Service Central d'Exploitation  
334 de la Météorologie, Météo-France, 190 pp.
- 335 Katz, R., M. Parlange, and P. Naveau, 2002: Extremes in hydrology. *Advances in Water*  
336 *Resources*, **25**, 1287–1304.
- 337 Kaufman, L. and P. Rousseeuw, 1990: *Finding Groups in Data: An Introduction to Cluster*  
338 *Analysis*. Wiley, New-York.
- 339 Kharin, V. V., F. W. Zwiers, X. Zhang, and G. C. Hegerl, 2007: Changes in temperature  
340 and precipitation extremes in the IPCC ensemble of global coupled model simulations. *J.*  
341 *Climate*, **20**, 1419–1444.
- 342 Naveau, P., A. Guillo, D. Cooley, and J. Diebolt, 2009: Modeling pairwise dependence of  
343 maxima in space. *Biometrika*, **96** (1), 1–17, doi:10.1093/biomet/asp001.
- 344 Plaut, G., E. Schuepbach, and M. Doctor, 2001: Heavy precipitation events over a few  
345 alpine sub-regions and the links with large scale circulation, 1971-1995. *Climate Research*,  
346 **17** (285–302).
- 347 Resnick, S., 2007: *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer  
348 Series in Operations Research and Financial Engineering.
- 349 Rousseeuw, P., 1986: Silhouettes: a graphical aid to the interpretation and validation of  
350 cluster analysis. *J. Comp App Math.*, **20** (53-65).
- 351 Schlather, M., 2002: Models for stationary max-stable random fields. *Extremes*, **5**, 33–44.
- 352 Schlather, M. and J. Tawn, 2003: A dependence measure for multivariate and spatial extreme  
353 values: Properties and inference. *Biometrika*, **90** (1), 139–156.

- 354 Stefanon, M., F. D’Andrea, and P. Drobinski, 2012: Heatwave classification over europe and  
355 the mediterranean region. *Environ. Res. Lett.*, **7** (014023), (9pp).
- 356 Vannitsem, S. and P. Naveau, 2007: Spatial dependences among precipitation maxima over  
357 belgium. *Nonlin. Processes Geophys.*, **14**, 621–630.
- 358 von Storch, H. and F. W. Zwiers, 2002: *Statistical analysis in climate research*. 484 pages,  
359 Cambridge University Press.
- 360 Wackernagel, H., 2003: *Multivariate Geostatistics. An Introduction with Applications*. 3d  
361 ed., Springer, Heidelberg.

## List of Figures

- 1 *Weekly maxima of hourly precipitation (Fall season, 92 stations over France, 1993-2011).* PANEL A: The clustering into five classes is obtained with the k-means algorithm applied to the raw maxima (left map) and to their logarithm (right map). This indicates that transforming marginal laws has a strong effect on the clustering. PANEL B: The left and right panels display the estimated scale  $\sigma$  and the shape parameter  $\xi$  after fitting a GEV distribution at each location, respectively. This means that the marginal law behavior varies spatially with heavier extremes in the south of France than in the north. 20
- 2 The left and right maps display the clustering outputs from our PAM algorithm and the Kmeans algorithm, respectively. On the left maps, the medoids are represented by black diamonds and small circles correspond to locations with non-significative silhouette coefficients. The number of clusters  $K$  equals 2, 5 and 7 for panels A, B and C, respectively. 21
- 3 The solid black line represents the average silhouette coefficient defined by (6Algorithm descriptionequation.2.6) in function of the number of clusters. The boxplot summarizes the distribution of silhouette coefficients. The dotted line with grey diamonds corresponds to the upper 95% level obtained after randomly reshuffling our precipitation data (i.e. breaking down the spatial structure). 22

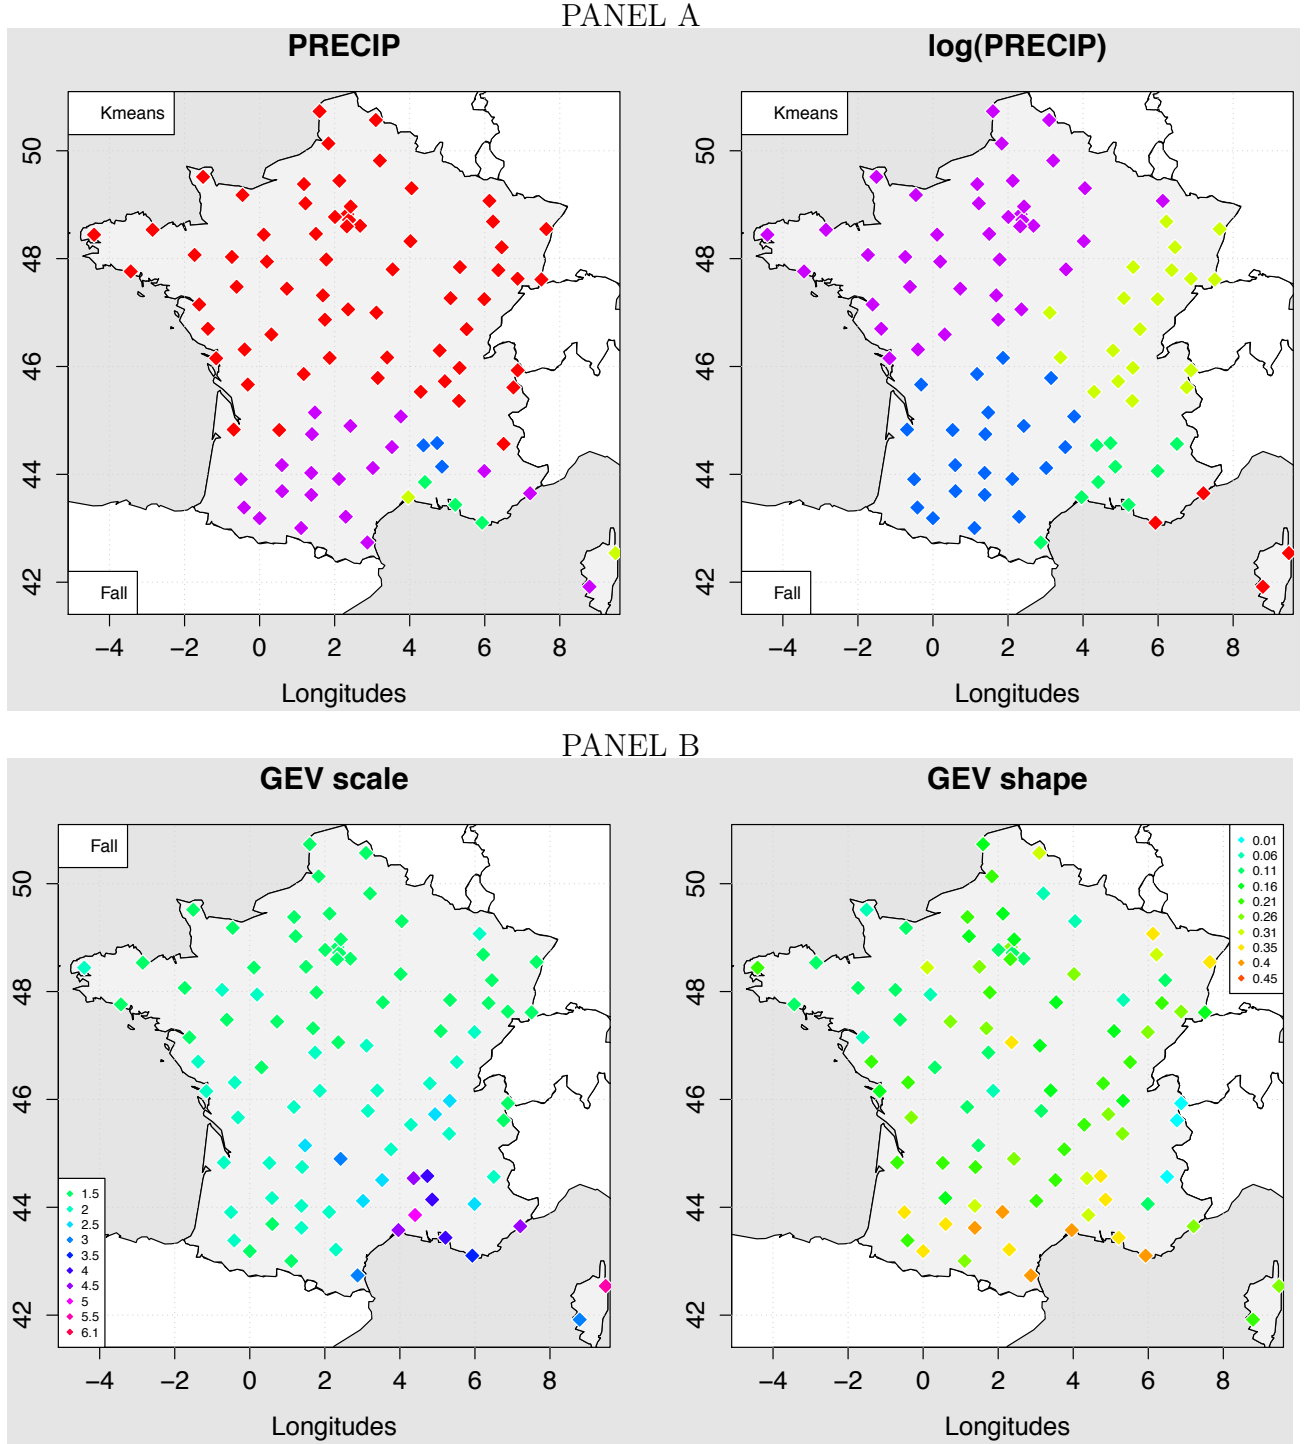


FIG. 1. *Weekly maxima of hourly precipitation (Fall season, 92 stations over France, 1993-2011).* PANEL A: The clustering into five classes is obtained with the k-means algorithm applied to the raw maxima (left map) and to their logarithm (right map). This indicates that transforming marginal laws has a strong effect on the clustering. PANEL B: The left and right panels display the estimated scale  $\sigma$  and the shape parameter  $\xi$  after fitting a GEV distribution at each location, respectively. This means that the marginal law behavior varies spatially with heavier extremes in the south of France than in the north.

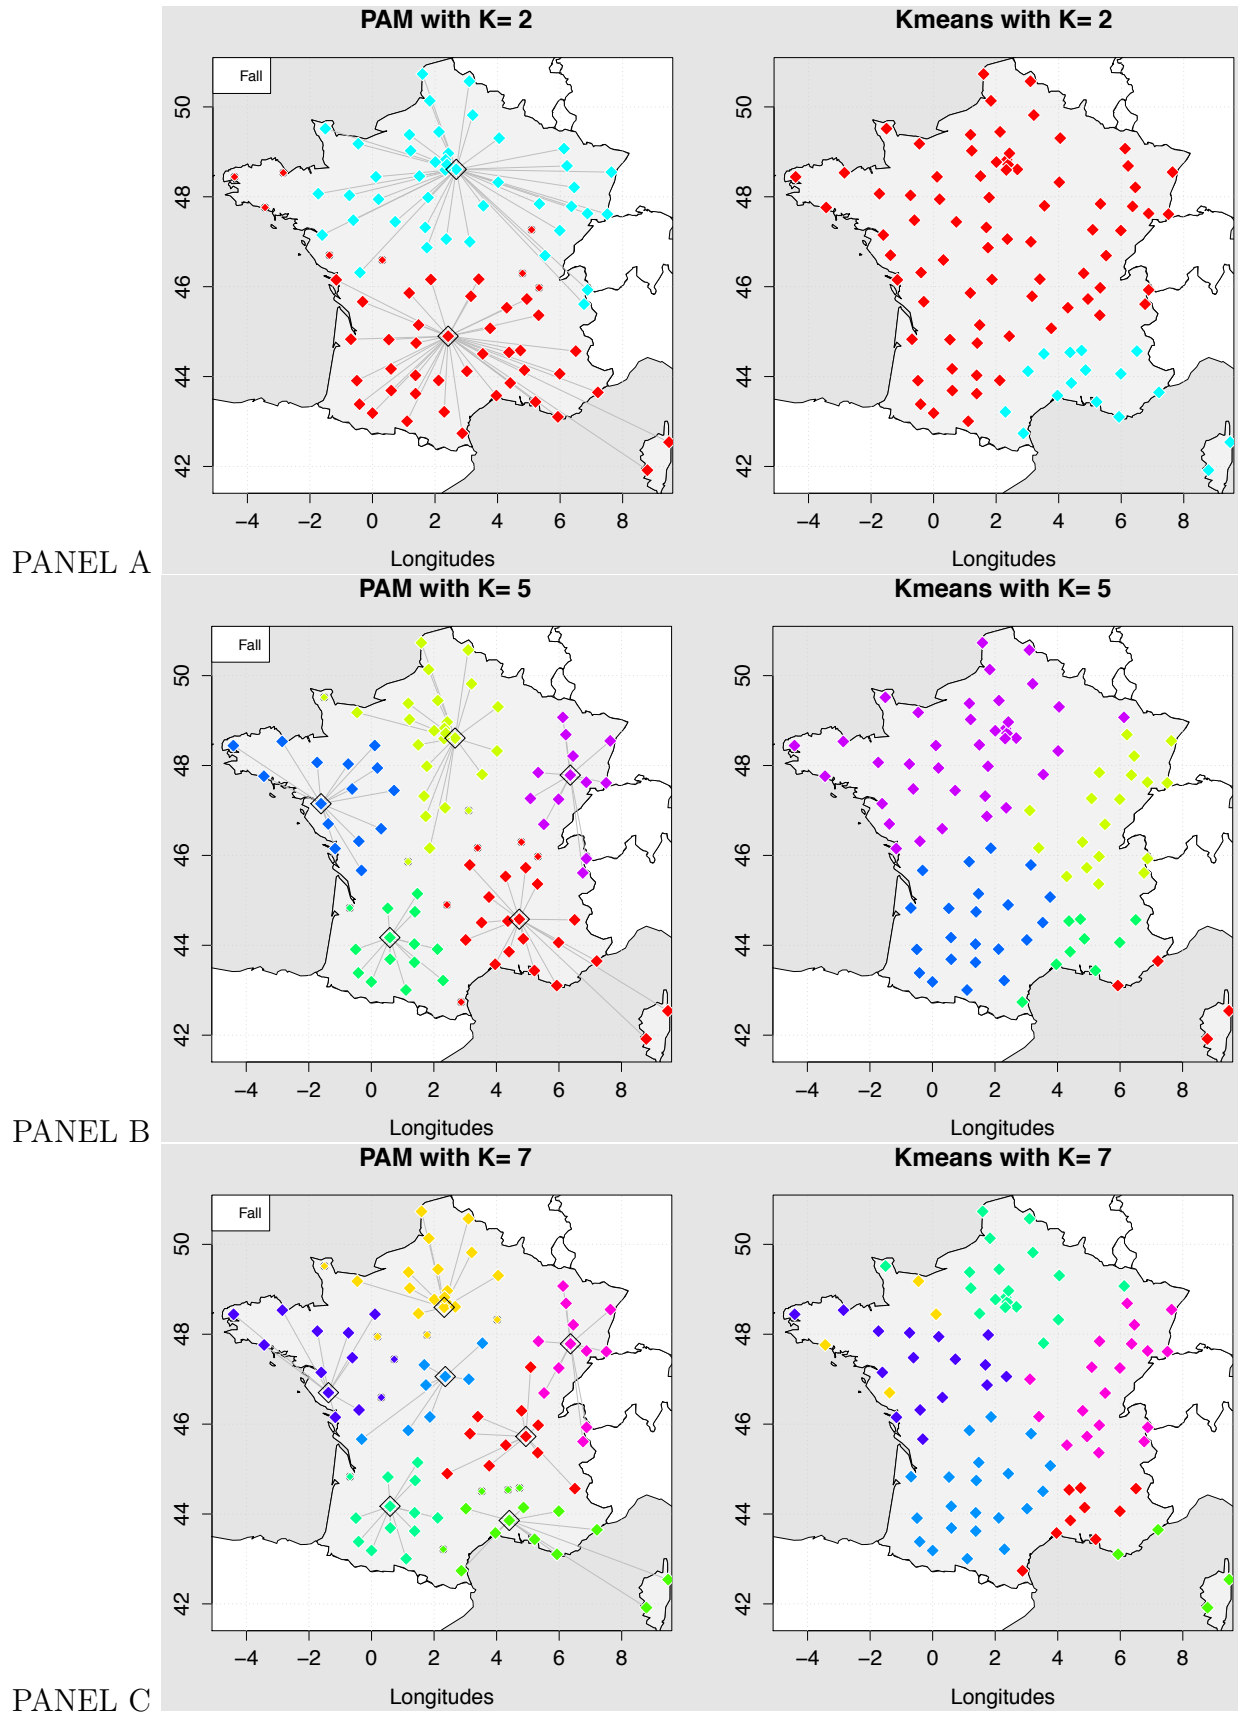


FIG. 2. The left and right maps display the clustering outputs from our PAM algorithm and the Kmeans algorithm, respectively. On the left maps, the medoids are represented by black diamonds and small circles correspond to locations with non-significant silhouette coefficients. The number of clusters  $K$  equals 2, 5 and 7 for panels A, B and C, respectively.

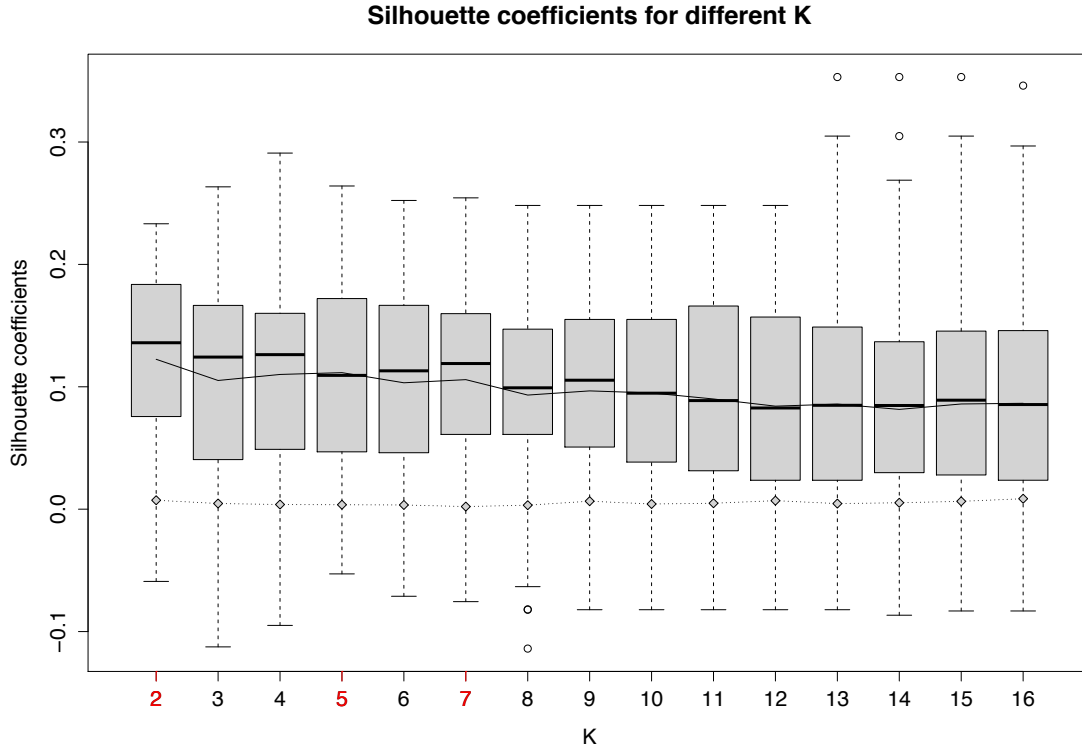


FIG. 3. The solid black line represents the average silhouette coefficient defined by (6) in function of the number of clusters. The boxplot summarizes the distribution of silhouette coefficients. The dotted line with grey diamonds corresponds to the upper 95% level obtained after randomly reshuffling our precipitation data (i.e. breaking down the spatial structure).