

Clustering of Maxima: Spatial Dependencies among Heavy Rainfall in France

ELSA BERNARD, PHILIPPE NAVEAU, AND MATHIEU VRAC

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

OLIVIER MESTRE

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

(Manuscript received 14 November 2012, in final form 23 April 2013)

ABSTRACT

One of the main objectives of statistical climatology is to extract relevant information hidden in complex spatial-temporal climatological datasets. To identify spatial patterns, most well-known statistical techniques are based on the concept of intra- and intercluster variances (like the k -means algorithm or EOFs). As analyzing quantitative extremes like heavy rainfall has become more and more prevalent for climatologists and hydrologists during these last decades, finding spatial patterns with methods based on deviations from the mean (i.e., variances) may not be the most appropriate strategy in this context of studying such extremes. For practitioners, simple and fast clustering tools tailored for extremes have been lacking. A possible avenue to bridging this methodological gap resides in taking advantage of multivariate extreme value theory, a well-developed research field in probability, and to adapt it to the context of spatial clustering. In this paper, a novel algorithm based on this plan is proposed and studied. The approach is compared and discussed with respect to the classical k -means algorithm throughout the analysis of weekly maxima of hourly precipitation 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 the 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 intervariances). In a nutshell, the k -means principle finds 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 covariance 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 p. 35 of von Storch and Zwiers 2002) ensures 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 nonnegative 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). Shedding light on new spatial or

Corresponding author address: Philippe Naveau, Laboratoire des Sciences du Climat et de l'Environnement, CNRS-CEA-UVSQ, Orme des Merisiers, 61190 Gif-sur-Yvette, France.
E-mail: naveau@lscce.ipsl.fr

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 event 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) give 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 an F -madogram (see section 2 for details) was proposed by Cooley et al. (2006) and Naveau et al. (2009), who studied a nonparametric 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 the 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 0.5. This class of distributions arises as the natural nondegenerated limit of rescaled independent and identically distributed (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 Eq. (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 Eq. (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 Eq. (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

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

can be expressed in terms of the extremal coefficient

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

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 1. 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 2, and Eq. (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 nonparametric fashion. The distance d_{ij} in Eq. (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 nonparametric estimator

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

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 Eq. (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 Eq. (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 Eq. (3) and appendix A]. This means that we do not 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 $V(\cdot, \cdot)$ 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 that they only belong to the domain of the attraction of max-stable distribution.

Having at our disposal the distance d_{ij} that is tailored from maxima motivated by Eq. (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 an averaging operation destroys the max-stable property encapsulated in Eq. (1), since the average of more than one maximum is no longer a maximum. As an attractive alternative, the 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 preprocessing steps are needed before implementing PAM. First, the distance matrix $\{d_{ij}\}$ defined by Eq. (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:

- (i) Form K clusters by assigning every point to its closest medoid.
- (ii) For each cluster, find the new medoid for which the total intracluster distance based on d_{ij} is minimized.

- (iii) If at least one medoid has changed, then go back to (i), otherwise end the algorithm.

In summary, PAM proceeds by moving around K medoids while trying to make the total intracluster distance as small as possible. As mentioned previously, the centers of the cluster, the so-called medoids, still represent the valid weekly precipitation maxima at each step of the algorithm. Consequently, the distance d_{ij} can always be interpreted via Eq. (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 cluster tightness (small d_{ik} within the cluster k) with cluster dissociation (see $\delta_{i,-k}$ defined below). After running the PAM algorithm with a given K , each location i is associated with a medoid k . The silhouette coefficient for the weather station i is defined as follows

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

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

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

or other statistics from the set $\{s_1(K), \dots, s_N(K)\}$. Such summaries will be used in our application. To implement our approach, a package for the open-source statistical R software is available at the homepage of the second author (<http://www.lsce.ipsl.fr/Phoceo/Pisp/visu.php?id=44&uid=naveau>).

3. Applications to French precipitation maxima

Here we focus on the weekly maxima of hourly precipitation at 92 French stations during the fall season [September–November (SON)] from 1993 to 2011. They were provided by the French meteorological service Météo-France. The stations were chosen in function of

their quality and to have a fairly homogeneous coverage of France. To avoid dealing with zeros and in order to be consistent with EVT, very small values of precipitation (rainfall amounts below 3 mm) were 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.

Figure 1a displays the outputs into five clusters. The difference between the left and right maps in Fig. 1a is a result of 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 time series with different marginal laws). Quantitatively, this can be assessed by fitting a GEV probability distribution function defined by $G(x) = \exp\{-[1 + \xi(x - \mu/\sigma)]_+^{-1/\xi}\}$, where the real μ is the location parameter, σ is the positive-scale parameter, and $\xi \in \mathbb{R}$ is the shape parameter.

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

Comparing the left of Fig. 1a with Fig. 1b suggests that the southeast 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 as 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 nonparametric approach. This second point implies that we do not need to fit a GEV

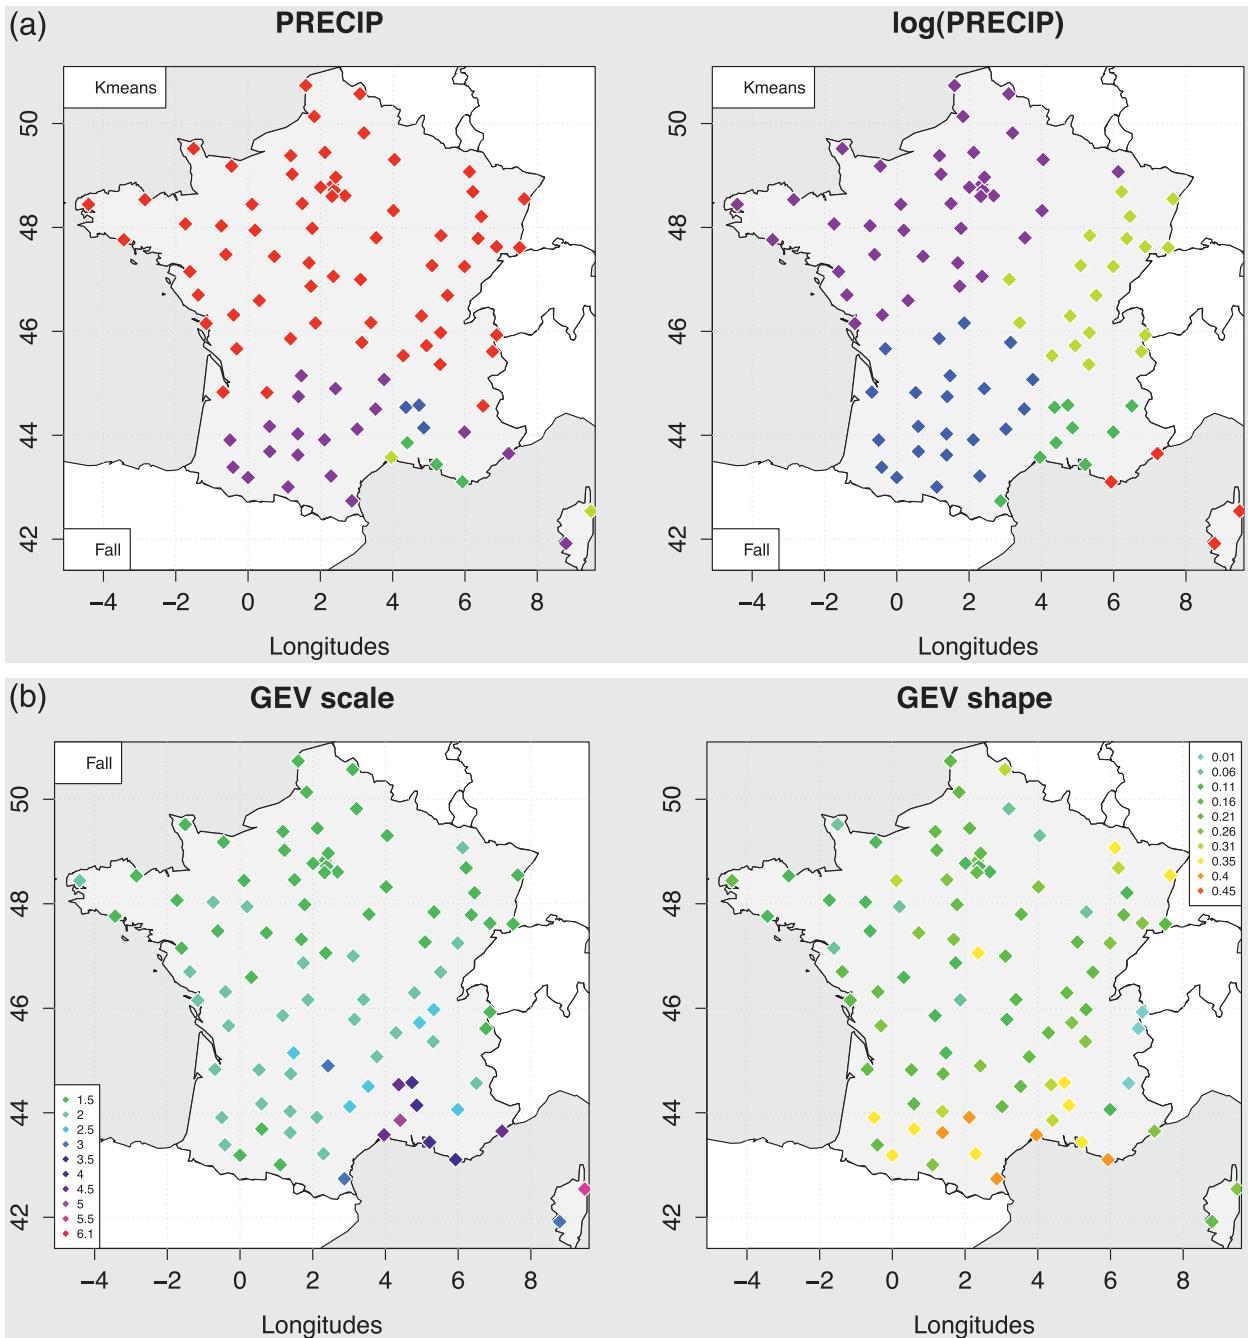


FIG. 1. Weekly maxima of hourly precipitation (fall season, 92 stations over France, 1993–2011). (a) The clustering into five classes is obtained with the k -means algorithm applied (left) to the raw maxima and (right) to their logarithm. This indicates that transforming marginal laws has a strong effect on the clustering. (b) Displayed are (left) the estimated scale σ and (right) the shape parameter ξ after fitting a GEV distribution at each location. This means that the marginal law behavior varies spatially with heavier extremes in the south of France than in the north.

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, Fig. 2 compares the clustering outputs for both methods: maps on the left for our PAM

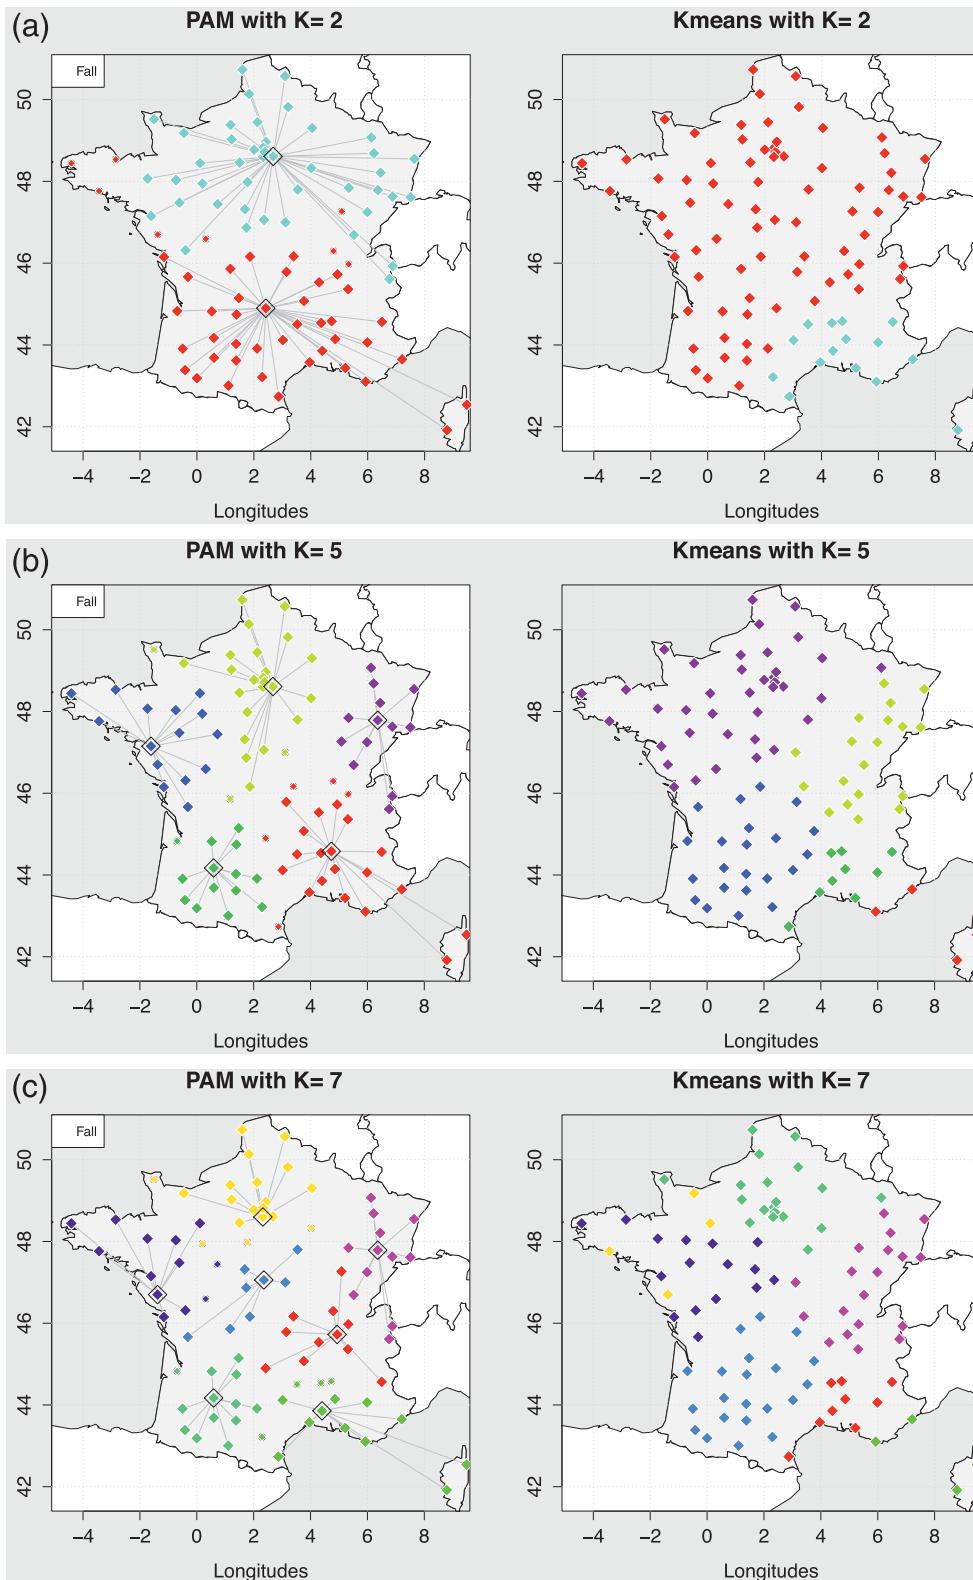


FIG. 2. Displayed are the clustering outputs from our (left) PAM algorithm and the (right) k -means algorithm. (left) The medoids are represented by black diamonds and points not attached to a medoid by a gray line correspond to locations with nonsignificant silhouette coefficients. The number of clusters K equals (a) 2, (b) 5, and (c) 7.

approach and on the right for the k -means algorithm applied on log-precipitation maxima (to reduce the margins problem). Figures 2a–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 gray 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 was taken 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 Fig. 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 $K = 2$ (Fig. 2a) 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 Fig. 1b). This feature is linked to rainfall intensities but not necessarily to spatial precipitation dependencies. For $K = 5$ (Fig. 2b), 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 Fig. 2c), sharper regional features appear and are geographically coherent. For $K = 7$, k means start 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 the Pyrénées, Cévennes, and Alps, resulting in severe thunderstorms. A systematic inventory of those situations over the 1958–94 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 farther 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 “southwest cluster” (Agen medoid), which is the case in PAM with $K = 7$. In the north of France, heavy rainfall is often produced by midlatitude 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 ($K = 2$ and 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 Fig. 3 summarizes the silhouette coefficient distribution for a given K varying from 2 to 16. Applying Eq. (6), the average silhouette coefficient is represented by the solid black line. The dotted line with gray 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 the silhouette coefficients below such thresholds are considered nonsignificant (see small circles in Fig. 2). Figure 3 does not bring a clear winner here as the largest average silhouette coefficients are very close around 0.12 ($K = 2$) and 0.11 ($K = 5$). In regards to the maps displayed in Fig. 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 overparameterization, 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

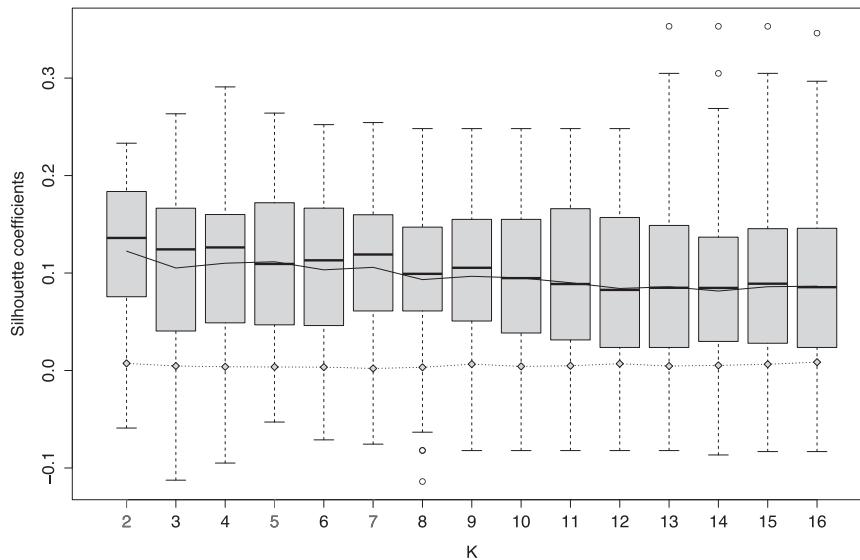


FIG. 3. The solid black line represents the average silhouette coefficient defined by Eq. (6) as a function of the number of clusters. The boxplot summarizes the distribution of silhouette coefficients. The dotted line with gray diamonds corresponds to the upper 95% level obtained after randomly reshuffling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

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.

In addition, 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 such as 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), by the PEPER-GIS Project, by the ANR-MOPERA Project, by the ANR-McSim Project, and by the MIRACCLE-GICC Project. The authors thank the contributors of the

statistical freely available R software (<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

- Beirlant, J., Y. Goegebeur, J. Segers, and J. Teugels, 2004: *Statistics of Extremes: Theory and Applications*. John Wiley and Sons, 514 pp.
- Cassou, C., L. Terray, J. W. Hurrell, and C. Deser, 2004: North Atlantic winter climate regimes: Spatial asymmetry, stationarity with time, and oceanic forcing. *J. Climate*, **17**, 1055–1068.
- Coles, S., 2001: *An Introduction to Statistical Modeling of Extreme Values*. Springer, 208 pp.
- , J. Heffernan, and J. Tawn, 1999: Dependence measures for extreme value analysis. *Extremes*, **2**, 339–365.
- Cooley, D., P. Naveau, and P. Poncet, 2006: Variograms for spatial max-stable random fields. *Dependence in Probability and Statistics*, P. Bertail, P. Doukhan, and P. Soulier, Eds., Lecture Notes in Statistics Series, Vol. 187, Springer, 373–90.
- de Haan, L., and A. Ferreira, 2006: *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering, Vol. 17, Springer, 471 pp.
- Dielbolt, J., A. Guillou, P. Naveau, and P. Ribereau, 2008: Improving probability-weighted moment methods for the generalized extreme value distribution. *Rev. Stat.*, **6**, 33–50.
- Fougères, A.-L., 2004: Multivariate extremes. *Extreme Values in Finance, Telecommunications, and the Environment*, B. Finkenstadt and H. Rootzen, Eds., CRC Press, 373–388.
- Gaume, E., L. Gaál, A. Viglione, J. Szolgay, S. Kohnová, and G. Bloschl, 2010: Bayesian MCMC approach to regional flood frequency analyses involving extraordinary flood events at ungauged sites. *J. Hydrol.*, **394**, 101–117, doi:10.1016/j.jhydrol.2010.01.008.

- Hastie, T., R. Tibshirani, and J. Friedman, 2009: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics, 2nd ed., Springer, 745 pp.
- Jacq, V., 1994: *Inventaire des situations à précipitations diluviennes, Languedoc-Roussillon, PACA, Corse: Période 1958-1994*. Publication du Service Central d'Exploitation de la Météorologie, Vol. 3, Météo-France, 190 pp.
- Katz, R., M. Parlange, and P. Naveau, 2002: Extremes in hydrology. *Adv. Water Resour.*, **25**, 1287–1304.
- Kaufman, L., and P. Rousseeuw, 1990: *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley, 368 pp.
- Kharin, V. V., F. W. Zwiers, X. Zhang, and G. C. Hegerl, 2007: Changes in temperature and precipitation extremes in the IPCC ensemble of global coupled model simulations. *J. Climate*, **20**, 1419–1444.
- Naveau, P., A. Guillou, D. Cooley, and J. Diebolt, 2009: Modelling pairwise dependence of maxima in space. *Biometrika*, **96**, 1–17, doi:10.1093/biomet/asp001.
- Plaut, G., E. Schuepbach, and M. Doctor, 2001: Heavy precipitation events over a few alpine sub-regions and the links with large scale circulation, 1971–1995. *Climate Res.*, **17**, 285–302.
- Resnick, S., 2007: *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer Series in Operations Research and Financial Engineering, Vol. 14, Springer, 404 pp.
- Rousseeuw, P., 1986: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.*, **20**, 53–65.
- Schlather, M., 2002: Models for stationary max-stable random fields. *Extremes*, **5**, 33–44.
- , and J. Tawn, 2003: A dependence measure for multivariate and spatial extreme values: Properties and inference. *Biometrika*, **90**, 139–156.
- Stefanon, M., F. D'Andrea, and P. Drobinski, 2012: Heatwave classification over Europe and the Mediterranean region. *Environ. Res. Lett.*, **7**, 014023, doi:10.1088/1748-9326/7/1/014023.
- Vannitsem, S., and P. Naveau, 2007: Spatial dependences among precipitation maxima over Belgium. *Nonlinear Processes Geophys.*, **14**, 621–630.
- von Storch, H., and F. W. Zwiers, 2002: *Statistical Analysis in Climate Research*. Cambridge University Press, 484 pp.
- Wackernagel, H., 2003: *Multivariate Geostatistics: An Introduction with Applications*. 3rd ed. Springer, 387 pp.