

## The Hierarchical Spectral Merger Algorithm: A New Time Series Clustering Procedure

Carolina Euán

Centro de Investigación en Matemáticas, México  
King Abdullah University of Science and Technology, Saudi Arabia  
UC Irvine Space-Time Modeling Group, University of California, USA

Hernando Ombao

Department of Statistics, University of California, Irvine, USA  
King Abdullah University of Science and Technology, Saudi Arabia  
UC Irvine Space-Time Modeling Group

Joaquín Ortega

Centro de Investigación en Matemáticas, México

**Abstract:** We present a new method for time series clustering which we call the Hierarchical Spectral Merger (HSM) method. This procedure is based on the spectral theory of time series and identifies series that share similar oscillations or waveforms. The extent of similarity between a pair of time series is measured using the total variation distance between their estimated spectral densities. At each step of the algorithm, every time two clusters merge, a new spectral density is estimated using the whole information present in both clusters, which is representative of all the series in the new cluster. The method is implemented in an R package *HSMClust*. We present two applications of the HSM method, one to data coming from wave-height measurements in oceanography and the other to electroencefalogram (EEG) data.

**Keywords:** Hierarchical spectral merger clustering; Time series clustering; Hierarchical clustering; Total variation distance; Time series; Spectral analysis.

---

The authors would like to thank the reviewers for their comments which led to improvements in this work.

Corresponding Author's Address: C. Euán, King Abdullah University of Science and Technology, Saudi Arabia, Thuwal 23955-6900, e-mail: [carolina.euancampos@kaust.edu.sa](mailto:carolina.euancampos@kaust.edu.sa).

Published online: 12 April 2018

## 1. Introduction

The subject of time series clustering is an active research area with applications in many fields. Finding similarity between time series frequently plays a central role in many applications. In fact, time series clustering problems arise in a natural way in a wide variety of fields, including economics, finance, medicine, ecology, environmental studies, engineering, and many others. A recent work, developed by Krafty (2016), uses a clustering model to develop a discriminant analysis of stochastic cepstra. Time series clustering is, in general, not an easy task due to the potential complexity of the data and the difficulty of defining an adequate notion of similarity between time series.

In Liao (2005), and subsequently in Caiado, Maharaj, and D'Urso (2015), there are three approaches to time series clustering: methods based on the comparison of raw data, methods based on parameters from models adjusted to the data, and feature-based methods where the similarity between time series is measured through features extracted from the data. The first clustering approach compares raw data and may not be computationally scalable for long time series. The second approach, based on parameters, is one of the most frequently used. However, it has the limitation of requiring a parametric model and this might suffer from model misspecification. Montero and Vilar (2014) present an R package (TSclust) for time series clustering with a wide variety of alternative procedures.

Here, we consider the problem of clustering stationary time series and our proposed method uses the spectral density as the central feature for classification purposes. To build a clustering method one needs to measure the similarity between the spectral densities. Our method uses the total variation (TV) distance as a measure of dissimilarity, which was proposed in Alvarez-Esteban, Euán, and Ortega (2016). This distance requires the use of normalized spectral densities, which is equivalent to standardizing the time series so that their variances are equal to one. Thus, it is the distribution of the power across different frequencies of the time series that is considered the fundamental feature for clustering purposes rather than the magnitude of the oscillations. The use of periodogram and normalized periodogram (which are estimators of the spectrum and the spectral density, respectively) for clustering was first proposed in Caiado, Crato, and Peña (2006, 2009). In fact, to our knowledge Caiado et al. (2006) were the first to propose the use of frequency features to find similarities between time series.

In Alvarez-Esteban et al. (2016), the TV distance was used to build a dissimilarity matrix consisting of the distances between all the spectral densities, which was then fed into a classical hierarchical agglomerative algorithm with the complete and average linkage functions. The method intro-

duced in this work which we call the Hierarchical Spectral Merger (HSM) method, is a new clustering algorithm that takes advantage of the spectral theory of time series. The key difference with classical hierarchical agglomerative clustering is that every time a cluster is updated, a new representative (new estimated spectral density) is computed. Each time two clusters are joined in the HSM procedure, the information available in all the series belonging to both clusters is merged to produce a new estimate of the common spectral density for the new cluster. The proposed approach is appealing because it is intuitive and the updated spectral estimates are smoother, less noisy and hence give better estimates of the TV distance. Thus, every time two clusters merge, the dissimilarity matrix reduces its size in one unit. In contrast, for the classical hierarchical agglomerative algorithm the dissimilarity matrix is the same throughout the procedure, and the distances between clusters at each step are calculated using linear combinations of the individual distances; the linear combination used depends on the linkage function that is chosen (single, complete, average, Ward, etc.). These methods are based on geometrical ideas which, in some cases, may not be meaningful for clustering time series, since a linear combination of the individual distances may not have a clear meaning in terms of the spectral densities.

We will present two applications of the HSM method: one to data coming from wave-height measurements in oceanography and the other to electroencephalogram (EEG) data. Some of the numerical experiments described in Section 3 are related to these applications.

The rest of the paper is organized as follows: Section 2 describes the Hierarchical Spectral Merger procedure. Section 3 presents some numerical experiments which compare the HSM method with some of the existing algorithms and considers the problem of deciding how many clusters there are in a given sample. Finally, Section 4 gives some examples of the use of the HSM algorithm. The paper ends with some discussion of the results and some ideas about future work.

## 2. HSM: A Method for Time Series Clustering

Our goal is to develop a method that finds groups or clusters that represent spectrally synchronized time series. The algorithm we introduce, known as the Hierarchical Spectral Merger (HSM) method, uses the TV distance as a dissimilarity measure and proposes a new clustering procedure. This algorithm is a modification of the usual agglomerative hierarchical procedure, taking advantage of the spectral point of view for the analysis of time series.

The first question when building a clustering algorithm is how to measure the dissimilarity between the objects one is considering. In our case,




Figure 1. The total variation distance measures the similarity between two densities. The shaded area represents the common area below the two densities, which we denote as  $\delta$ .

this amounts to measuring the dissimilarity between the spectral densities estimated from the original time series, for which we use the TV distance between the normalized spectra. In what follows, Section 2.1 presents the TV distance while Section 2.2 describes in detail the HSM method.

## 2.1 Total Variation (TV) Distance

In general, the TV distance can be defined for any pair of probability measures that live on the same  $\sigma$ -algebra of sets. We will be focusing here on the case where these probability measures are defined on the real line and have density functions with respect to the Lebesgue measure. The TV distance between two probability densities,  $f$  and  $g$ , is defined as

$$d_{TV}(f, g) = 1 - \int \min\{f(\omega), g(\omega)\} d\omega. \quad (2..1)$$

Equation (2.1) suggests a graphical interpretation of the TV distance. If  $f$  and  $g$  are probability density functions and  $d_{TV}(f, g) = 1 - \delta$  then  $\int \min\{f(\omega), g(\omega)\} d\omega = \delta$  and this means that the common area below the two densities is equal to  $\delta$ , which corresponds to the shaded area in Figure 1. When the area shared by the two densities increases then the TV distance decreases.

Since spectral densities are not probability densities, they have to be normalized by dividing the estimated spectral density by the sample variance  $\hat{\gamma}(0)$ , since  $\hat{\gamma}(0) = \int \tilde{f}(\omega) d\omega$ . Thus, we denote  $\hat{f}(\omega) = \tilde{f}(\omega)/\hat{\gamma}(0)$  to be the normalized estimated spectral density.

In comparison with other similarity measures, the TV distance has some desirable properties. (1) The TV distance is a pseudo metric. It satisfies the symmetry condition and the triangle inequality, which are two reasonable properties expected from a similarity measure. In this sense, the TV distance may be a better choice than the Kullback-Leibler (KL) divergence

(although one can symmetrize the KL divergence). (2) The TV distance is bounded,  $0 \leq d_{TV}(f, g) \leq 1$  and can be interpreted in terms of the common area between two densities. Having a bounded range ( $[0, 1]$ ) is a desirable feature, since this gives a very intuitive sense to the values attained by the similarity measure. A value near 0 means that the densities are similar while a value near 1 indicates they are highly dissimilar. In contrast, both the  $L^2$  distance and the Kullback-Leibler divergence are not bounded above and thus lack this intuitive interpretation.

## 2.2 The Hierarchical Spectral Merger (HSM) Algorithm

There are two general families of clustering algorithms: partitioning and hierarchical. Among partitioning algorithms, the K-means and K-medoids are the most commonly used. Maharaj and D'Urso (2012) proposed a K-means fuzzy clustering method based on wavelets coefficients and made a comparison with other K-means algorithms. For the hierarchical clustering algorithms, the typical examples are agglomerative with single-linkage or complete-linkage (Xu and Wunsch, 2005). Storage and computational properties of the hierarchical clustering methods are reviewed in Contreras and Murtagh (2015).

The hierarchical spectral merger algorithm has two versions: the first, known as *single version*, updates the spectral estimate of the cluster from a concatenation of the time series, and the second, known as *average version*, updates the spectral estimate of the cluster from a weighted average of the spectral estimate obtained from each signal in the cluster.

**Hierarchical Spectral Merger Algorithm.** Let  $\{X_i = (X_i(1), \dots, X_i(T)), i = 1, \dots, N\}$  be a set of time series. The procedure starts with  $N$  clusters, each cluster being a single time series.

**Step 1.** Suppose there are  $k$  clusters. For each cluster, estimate the spectral density (using some smoothing or averaging method) and represent each cluster by a common normalized spectral density  $\hat{f}_j(\omega)$ ,  $j = 1, \dots, k$ .

**Step 2.** Compute the TV distance between these  $k$  spectral densities.

**Step 3.** Find the two clusters that have smallest TV distance.

**Step 4.** Merge the time series in the two clusters with the smallest TV distance and replace the two clusters with the newly combined single cluster.

**Step 5.** Repeat Steps 1-4 until there is only one cluster left.

In Step 1, at the first iteration the spectral density for each time series is initially estimated using the smoothed periodogram. In our case, we used a lag window estimator with a Parzen window. In further iterations, however, the way the common spectral density is obtained depends on the version

Table 1. Hierarchical Spectral Merger Algorithm proposed using the Total Variation distance and the estimated spectra.

---



---

**Algorithm:**

---



---

1. Initial clusters:  $\mathbf{C} = \{C_i\}$ ,  $C_i = X_i$ ,  $i = 1, \dots, N$ .  
Dissimilarity matrix entry between clusters  $i$  and  $j$ ,
 
$$D_{ij} = d(C_i, C_j) := d_{TV}(\hat{f}_i, \hat{f}_j),$$
 TV distance between the corresponding estimated normalized spectral densities  $\hat{f}_i$  using the signals in each cluster.
2. **for**  $k$  *in*  $1 : N - 1$
3.  $(i_k, j_k) = \underset{ij}{\operatorname{argmin}} D_{ij}$ ;  $\min_k = \min_{ij} D_{ij}$       #Find the closest clusters
4.  $C_{new} = C_{i_k} \cup C_{j_k}$   
#Join the closest clusters
5.  $D^{new} = D \setminus \{D_{i_k}, \cup D_{j_k}, \cup D_{i_k}, \cup D_{j_k}\}$       #Delete rows and columns  $i_k, j_k$
6. **for**  $j$  *in*  $1 : N - k - 1$
7.  $D_{(N-k)j}^{new} = D_{j(N-k)}^{new} = d_{TV}(C_{new}, C_j)$       #Compute new distances
8. **end**
9.  $D = D^{new}$ ;  $\mathbf{C} = (\mathbf{C} \setminus \{C_{i_k}, C_{j_k}\}) \cup C_{new}$       #New matrix  $D$  and new clusters
10. **end**

---

of the algorithm. When two clusters merge, there are two options, either (1) for the single version, we concatenate the signals in both clusters and compute the smoothed periodogram with the concatenated signal; or (2) for the average version, we take the weighted average over all the estimated spectra for each signal in the new cluster as the new estimated spectra.

**Remarks.** (1) The value in Step 3 represents the “cost” of joining two clusters, i.e., having  $k - 1$  clusters versus  $k$  clusters. This cost may be useful in deciding whether or not to join the two clusters. (2) Both versions of the algorithm compute the TV distance between the new and the old clusters based on updated estimated spectral densities, which is the main difference with classical hierarchical algorithms. While a hierarchical algorithm has a dissimilarity matrix of size  $N \times N$  during the whole procedure, the proposed method reduces this size to  $(N - k) \times (N - k)$  at the  $k$ -th iteration. Table 1 gives a summary of the hierarchical merger algorithm.




Figure 2. Estimated spectra. (a.) Different colors correspond to different time series, (b.) Red spectra (continuous curve) are from realizations of the AR(2) model with activity at alpha (8-12 Hz) and beta (12-30 Hz) bands and black spectra (dashed curve) are from realizations of the AR(2) model with activity at alpha and gamma (30-50 Hz) bands. (See online version for color.)

**Illustration.** Consider two different AR(2) models with their spectra concentrated at 10 Hertz, however, one also contains power at 21 Hertz while the other has power at 40 Hertz. We simulated three time series for each process, 10 seconds of each one with a sampling frequency of 100 Hertz ( $t = 1, \dots, 1000$ ). Figure 2(a) shows the estimated spectra for each series and Figure 2(b) shows by different lines (continuous and dashed) which one belongs to the first or second process. If we only look at the spectra, it is hard to recognize the number of clusters and their memberships. The method might have difficulty in clustering some cases, like the red and purple spectra.

The step-by-step result of the HSM method is shown in Figure 3. We start with six clusters; at the first iteration we find the closest spectra, represented in Figure 3(a) with the same color (red). After the first iteration we merge these time series and get 5 estimated spectra, one per cluster, Figure 3(b) shows the estimated spectra where the new cluster is represented by the dashed red curve. We can follow the procedure in Figures 3(c), 3(d), 3(e) and 3(f). In the end, the proposed clustering algorithm reaches the correct solution, Figures 3(g) and 2(b) coincide. Also, the estimated spectra for the two clusters, Figure 3(h), is better than any of the initial spectra and we can identify the dominant frequency bands for each cluster.

### 3. Numerical Experiments

We now investigate the finite sample performance of the HSM clustering algorithm. First, we explain the simulation methods based on the spectrum that we will use in some of the experiments. Then, we present the results of the experiments, assuming that the number of clusters is known.

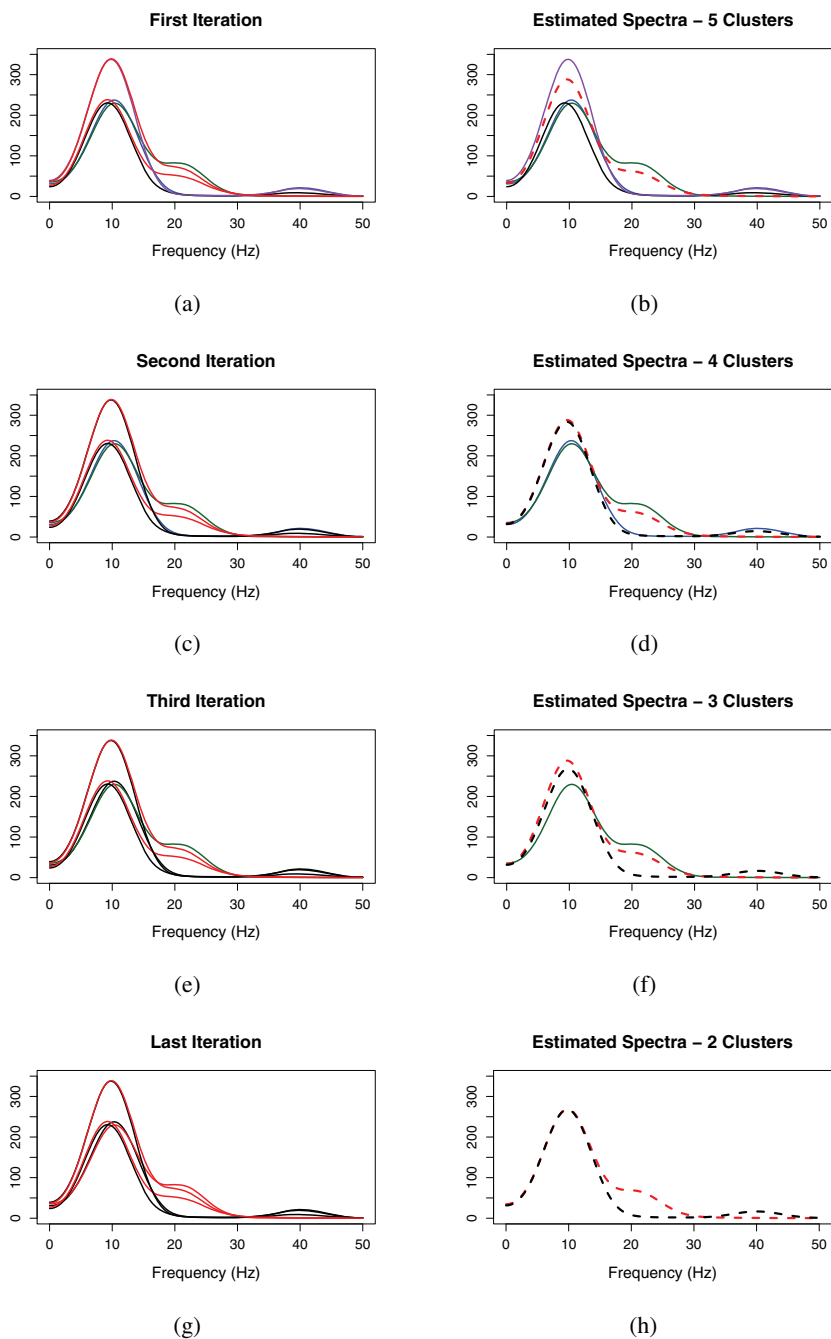


Figure 3. Dynamics of the hierarchical spectral merger algorithm. (a), (c), (e) and (g) show the clustering process for the spectra. (b), (d), (f) and (h) show the evolution of the estimated spectra, which improves as more series are merged in the same cluster. (See online version for color.)






Figure 4. JONSWAP spectrum with different parameters. (See online version for color.)

Finally, we explore the case of unknown number of clusters and possible criteria to choose it.

*Simulation based on a parametric family of spectral densities.*

There exist several parametric families of spectral densities of frequent use in oceanography, which have an interpretation in terms of the behavior of sea waves (Ochi, 1998). Motivated by the applications, we will simulate time series (Gaussian process) using spectra from one of these families. This methodology is already implemented by Brodtkorb, Johannesson, Lindgren, Rychlik, Rydén, and Sjö (2010) in the WAFO toolbox for MATLAB.<sup>1</sup>

An example of a group of parametric densities is the JONSWAP (Joint North-Sea Wave Project) spectral family, which is given by the formula

$$S(\omega) = \frac{g^2}{\omega^5} \exp(-5\omega_p^4/4\omega^4)\gamma^{\exp(-(\omega-\omega_p)^2/2\omega_p^2s^2)}, \quad \omega \in (-\pi, \pi),$$

where  $g$  is the acceleration of gravity,  $s = 0.07$  if  $\omega \leq \omega_p$  and  $s = 0.09$  otherwise;  $\omega_p = \pi/T_p$  and  $\gamma = \exp(3.484(1 - 0.1975(0.036 - 0.0056T_p/\sqrt{H_s})T_p^4/(H_s^2)))$ . The parameters for the model are the significant wave

---

1. The software WAFO developed by the Wafo group at Lund University of Technology, Sweden, available at <http://www.maths.lth.se/matstat/wafo> was used for the calculation of all spectral densities and associated spectral characteristics. The data for station 160 were furnished by the Coastal Data Information Program (CDIP), Integrative Oceanographic Division, operated by the Scripps Institution of Oceanography, under the sponsorship of the U.S. Army Corps of Engineers and the California Department of Boating and Waterways (<http://cdip.ucsd.edu/>).

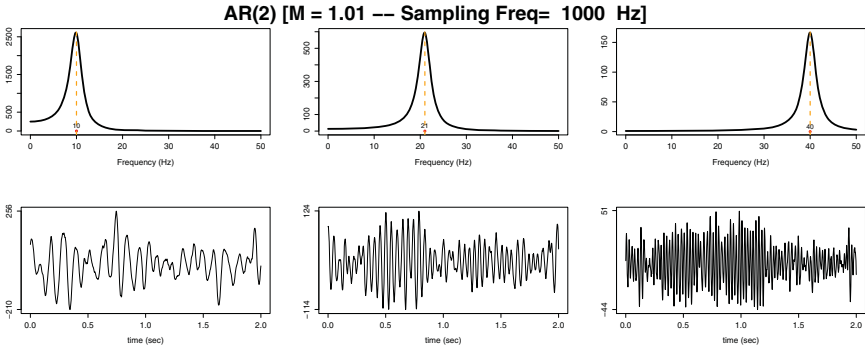


Figure 5. Top: Spectra for the AR(2) process with different peak frequency;  $\eta = 10, 21, 40$ . Bottom: Realizations from the corresponding AR(2) process.

height  $H_s$ , which is defined as 4 times the standard deviation of the time series, and the spectral peak period  $T_p$ , which is the period corresponding to the modal frequency of the spectrum. Figure 4 shows some examples of this family with different values for the parameters  $T_p$  and  $H_s$ .

*Simulation based on AR(2) processes.* Consider the stationary second order auto-regressive AR(2) model which is defined as

$$Z_t = \phi_1 Z_{t-1} + \phi_2 Z_{t-2} + \epsilon_t, \quad (3.1)$$

where  $\epsilon_t$  is a white noise process. The characteristic polynomial for this model is  $\phi(z) = 1 - \phi_1 z - \phi_2 z^2$ . The roots of the polynomial equation (assumed with magnitude bigger than 1 for causality) indicate the properties of the oscillations. If the roots, denoted  $z_0^1$  and  $z_0^2$  are complex-valued then they must be complex-conjugates, i.e.,  $z_0^1 = \overline{z_0^2}$ . These roots have a polar representation whose magnitudes and phases are, respectively,

$$|z_0^1| = |z_0^2| = M, \quad \arg(z_0) = \frac{2\pi\eta}{F_s}, \quad (3.2)$$

where  $F_s$  denotes the sampling frequency (in Hertz);  $M$  is the amplitude or magnitude of the root ( $M > 1$  for causality); and  $\eta$  is the frequency index ( $\eta \in (0, F_s/2)$ ). The spectrum of the AR(2) process with polynomial roots as above will have peak frequency at  $\eta$ . The peak becomes broader as  $M \rightarrow \infty$ , and it becomes narrower (or more concentrated around  $\eta$ ) as  $M \rightarrow 1^+$ . Then, given  $(\eta, M, F_s)$ , we take

$$\phi_1 = \frac{2 \cos(w_0)}{M} \quad \text{and} \quad \phi_2 = - \left( \frac{1}{M^2} \right), \quad (3.3)$$

where  $w_0 = \frac{2\pi\eta}{F_s}$ . If one computes the roots of the characteristic polynomial with the coefficients in (3.3), they satisfy (3.2).

To illustrate the type of oscillatory patterns that can be observed in time series from processes with corresponding spectra, we plot in Figure 5 the spectral densities (top) for different values of  $\eta$ ,  $M = 1.01$  and  $F_s = 1000$  Hertz; and examples of generated time series (bottom). Larger values of  $\eta$  gives rise to faster oscillations (higher frequencies) of the signal.

### 3.1 Experimental Design

We consider two different experiments. The first one is motivated by applications in oceanography, where the differences between spectra could be produced by a small change in the modal frequency. The second experiment was designed to test if the algorithms are able to distinguish between unimodal and bimodal spectra. For all the experiments, the lengths of the time series were  $T = 500, 1000,$  and  $2000$ .

- **Experiment 1** is based on two different JONSWAP (Joint North-Sea Wave Project) spectra (i.e. two clusters). The spectral densities considered have significant wave height  $H_s$  equal to 3, the first has a peak period  $T_p$  of  $3.6\sqrt{H_s}$  while for the second  $T_p = 4.1\sqrt{H_s}$ . Figure 6(a) exhibits the JONSWAP spectra, showing that the curves are close to each other. Five series from each spectrum were simulated and  $N = 500$  replicates of this experiment were made. In this case the sampling frequency was set to 1.28 Hertz, which is a common value for wave data recorded using sea buoys. This experiment was carried out in Alvarez-Esteban et al. (2016) to compare several clustering procedures.
- **Experiment 2** is based on the AR(2) process. Let  $Z_t^j, j = 1, 2, 3,$  be AR(2) processes with  $M_j = 1.1$  for all  $j$  and peak frequency  $\eta_j = .1, .13, .16$  for  $j = 1, 2, 3,$  respectively.  $Z_t^j$  represents a latent signal oscillating at a pre-defined band. Define the observed time series to be a mixture of these latent AR(2) processes.

$$\begin{pmatrix} X_t^1 \\ X_t^2 \\ \vdots \\ X_t^K \end{pmatrix}_{K \times 1} = \begin{pmatrix} e_1^T \\ e_2^T \\ \vdots \\ e_K^T \end{pmatrix}_{K \times 3} \begin{pmatrix} Z_t^1 \\ Z_t^2 \\ Z_t^3 \end{pmatrix}_{3 \times 1} + \begin{pmatrix} \varepsilon_t^1 \\ \varepsilon_t^2 \\ \vdots \\ \varepsilon_t^K \end{pmatrix}_{K \times 1} \quad (3.4)$$

where  $\varepsilon_t^j$  is Gaussian white noise,  $X_t^j$  is a signal with oscillatory behavior generated by the linear combination  $e_i^T \mathbf{Z}_t$  and  $K$  is the number of clusters. In this experiment, we set  $K = 3$ , with  $e_1^T = c(1, 0, 0)$ ,  $e_2^T$




Figure 6. Spectra used in the simulation study to compare the HSM method with other similarity measures. Each spectrum, with different color and line type, corresponds to a cluster. (See online version for color.)

$= c(0, 1, 0)$  and  $e_3^T = c(0, 1, 1)$ . We simulate five replicates of each signal  $X_t^i$ , so, we have three clusters with five members each. Figure 6(b) plots the three different spectra. For this experiment  $N = 1000$  replicates were made, and the sampling frequency was set to 1 Hertz.

### 3.2 Comparative Study

To compare clustering results, we must take into account the “quality of the clustering” produced which depends on both the similarity measure and the clustering algorithm used. The HSM method has two main features: The use of the TV distance as a similarity measure and the hierarchical spectral merger algorithm. The HSM method will be compared with the usual hierarchical agglomerative clustering algorithm using the complete and average linkage functions, which are among the standard clustering procedures used in the literature. Pérttega Díaz and Vilar (2010) proposed two simulation tests to compare the performance of several dissimilarity measures for time series clustering. We compare the HSM method with competitors that were based on the spectral density and performed well in Pérttega Díaz and Vilar’s experiments. In addition, we also considered the distance based on the cepstral coefficients, which was used in Maharaj and D’Urso (2011), and the symmetric version of the Kulback-Leibler divergence, that was used in a hierarchical clustering algorithm in Shumway and Stoffer (2011).

Let  $I_X(\omega_k) = T^{-1} \left| \sum_{t=1}^T X_t e^{-i2\pi\omega_k t} \right|^2$  be the periodogram for time series  $X$ , at frequencies  $\omega_k = k/T$ ,  $k = 1, \dots, n$  with  $n = [(T-1)/2]$ , and  $NI_X$  be the normalized periodogram, i.e.  $NI_X(\omega_k) = I_X(\omega_k)/\hat{\gamma}_0^X$ , with  $\hat{\gamma}_0^X$  the sample variance of time series  $X$  (notice that  $\hat{\gamma}_0^X =$

$\sum_{k=-(n-1)}^n I_X(\omega_k)$ ). The estimator of the spectral density  $\hat{f}_X$  is the smoothed periodogram using a Parzen window with bandwidth equal to  $100/T$ , normalized by dividing by  $\hat{\gamma}_0^X$ .

The dissimilarity criteria in the frequency domain considered were:

- The Euclidean distance between the normalized estimated spectra:

$$d_{NIP}(X, Y) = \frac{1}{n} \left( \sum_k (\hat{f}_X(\omega_k) - \hat{f}_Y(\omega_k))^2 \right)^{1/2}.$$

- The Euclidean distance between the logarithm of the normalized estimated spectra:

$$d_{LNIP}(X, Y) = \frac{1}{n} \left( \sum_k (\log \hat{f}_X(\omega_k) - \log \hat{f}_Y(\omega_k))^2 \right)^{1/2}.$$

- The square of the Euclidean distance between the cepstral coefficients

$$d_{CEP}(X, Y) = \sum_k^p (\theta_k^X - \theta_k^Y)^2 \text{ where, } \theta_0 = \int_0^1 \log I(\lambda) d\lambda \text{ and } \theta_k = \int_0^1 \log I(\lambda) \cos(2\pi k\lambda) d\lambda.$$

- The symmetric Kullback-Leibler distance between the normalized estimated spectra:

$$d_{KL}(X, Y) = \int \hat{f}_X(\omega) \log \left( \frac{\hat{f}_X(\omega)}{\hat{f}_Y(\omega)} \right) d\omega,$$

$$d_{SKL}(X, Y) = d_{KL}(X, Y) + d_{KL}(Y, X).$$

We also added the clustering method proposed by Alvarez-Esteban et al. (2016), that uses the TV distance in a hierarchical clustering algorithm. All these dissimilarity measures were compared with the HSM method using normalized estimated spectra.

To evaluate the rate of success, we considered the following index which has been already used for comparing different clustering procedures (Pértega Díaz and Vilar, 2010; Gavrilov et al., 2000). Let  $\{C_1, \dots, C_g\}$  and  $\{G_1, \dots, G_k\}$ , be the set of the  $g$  true groups and a  $k$ -cluster solution, respectively.

Then,  $\text{Sim}(G, C) = \frac{1}{g} \sum_{i=1}^g \max_{1 \leq j \leq k} \text{Sim}(G_j, C_i)$ , where

$$\text{Sim}(G_j, C_i) = \frac{2|G_j \cup C_i|}{|G_j| + |C_i|}.$$

Note that this similarity measure will return 0 if the two clusterings are completely dissimilar and 1 if they are the same.

In the comparative study, each simulation setting was replicated  $N$  times, and the rate of success for each one was computed. The mean values for this index are shown in Tables 2 and 3, and boxplots of the values obtained are shown in Figures 7 and 8. The simulation settings were the same in all cases and the clustering algorithm is hierarchical with the complete link (C) or average link function (A). For the HSM method, we used the notation HSM1 when we used the single version, and HSM2 when we used the average version. The clustering method of Alvarez-Esteban et al. (2016) is denoted by TV.

Table 2. Mean values of the similarity index obtained using different distances and the two proposed methods in Experiment 1. The number of replicates is  $N = 500$ .

Experiment 1												
$T$	NP		LNP		CEP		TV		SKL		HSM1	HSM2
	C	A	C	A	C	A	C	A	C	A		
500	0.979	0.979	0.772	0.787	0.597	0.588	0.988	0.992	0.994	0.995	0.989	0.988
1000	0.998	0.998	0.851	0.916	0.825	0.867	0.999	0.999	0.999	1	0.999	0.999
2000	1	1	0.932	0.976	0.908	0.958	1	1	1	1	1	1

Table 3. Mean values of the similarity index obtained using different distances and the two proposed methods in Experiment 2. The number of replicates is  $N = 1000$ .


Experiment 2												
$T$	NP		LNP		CEP		TV		SKL		HSM1	HSM2
	C	A	C	A	C	A	C	A	C	A		
500	0.864	0.802	0.949	0.914	0.895	0.849	0.930	0.905	0.952	0.941	0.836	0.838
1000	0.961	0.947	0.996	0.997	0.974	0.976	0.990	0.990	0.994	0.996	0.983	0.983
2000	0.995	0.997	1	1	0.999	0.999	0.999	0.999	0.999	1	0.999	0.999

We can see from the boxplots corresponding to **Experiment 1** that the CEP distance has many values smaller than 0.9 even in the case of  $T = 2000$  with the complete link. In **Experiment 2**, the HSM method did not have a good performance for small-sized time series. It is necessary to have  $T = 1000$ , for the HSM method to identify the clusters more precisely. The NP distance has the worst performance overall, using either the complete or average link functions.


In general, the rates of success for the HSM methods are very close to one. In some cases the HSM method has the best results, and when it is not the case, the rates obtained by the HSM method are close to the best. The methods based on logarithms, such as the LNP and CEP, have in some cases a good performance but in others their performance is very poor. Compared to the symmetric Kullback-Leibler distance, there is no clear winner but the method proposed here still has the advantage of being easily interpretable because the KL (or symmetric KL) cannot indicate if the dissimilarity value is large since it belongs to the range  $[0, \infty)$ .

### 3.3 When the Number of Clusters is Unknown


For real data the number of clusters is usually unknown. Thus, an objective criterion is needed to determine the optimal number of clusters.



(c)  $T = 500$ .




(d)  $T = 1000$ .




(e)  $T = 2000$ .


Figure 7. Boxplots of the rate of success for the replicates under the simulation setting of Experiment 1, by using different distances.



(f)  $T = 500$ .



(g)  $T = 1000$ .



(h)  $T = 2000$ .

Figure 8. Boxplots of the rate of success for the replicates under the simulation setting of Experiment 2, by using different distances.

As mentioned in Step 3 of our algorithm, the TV distance computed before joining two clusters can be used as a criterion.

We present two options for selecting the number of clusters. The first is an empirical criterion while the second one is based on a bootstrap procedure.

### 3.3.1 Empirical Criterion

Consider the following experimental design, similar to that of **Experiment 2**. Let  $Z_t^j$  be an AR(2) process with  $M_j = 1.01$  for all  $j$  and peak frequencies  $\eta_j = 2, 6, 10, 21$  and  $40$  for  $j = 1, \dots, 5$ , respectively. Define the observed time series to be a mixture of these latent AR(2) processes.

$$\begin{pmatrix} X_t^1 \\ X_t^2 \\ \vdots \\ X_t^K \end{pmatrix}_{K \times 1} = \begin{pmatrix} e_1^T \\ e_2^T \\ \vdots \\ e_K^T \end{pmatrix}_{K \times 5} \begin{pmatrix} Z_t^1 \\ Z_t^2 \\ \vdots \\ Z_t^5 \end{pmatrix}_{5 \times 1} + \begin{pmatrix} \varepsilon_t^1 \\ \varepsilon_t^2 \\ \vdots \\ \varepsilon_t^K \end{pmatrix}_{K \times 1} \quad (3..5)$$

where  $\varepsilon_t^j$  is Gaussian white noise,  $X_t^j$  is a signal with oscillatory behavior generated by the linear combination  $e_i^T \mathbf{Z}_t$ ,  $K$  is the number of spectrally synchronized groups or clusters and  $n$  denotes the number of replicates of each signal  $X_t^i$ .

Figure 9 displays the graphs corresponding to the minimum values of the TV distance between clusters as a function of the number of clusters; Figure 9(a) corresponds to the experimental design just described; Figure 9(b) corresponds to the experimental design with the same coefficients as the first one but  $K = 5$ ,  $n = 3$  and 500 draws; Figure 9(c) corresponds to a design with  $K = 5$ ,  $n = 3$  and 500 draws, with coefficients  $e_1^T = (1/2 \ 1 \ 0 \ 0 \ 0)$ ,  $e_2^T = (0 \ 1 \ 1/2 \ 0 \ 0)$ ,  $e_3^T = (0 \ 0 \ 1/2 \ 1 \ 0)$ ,  $e_4^T = (0 \ 0 \ 0 \ 1 \ 1/2)$ ,  $e_5^T = (0 \ 1 \ 0 \ 1 \ 0)$ . Finally, in Figure 9(d) we set  $K = 10$ ,  $n = 15$  and 500 draws but the coefficients  $e_i$  are drawn from a Dirichlet distribution with parameters  $(.2, .2, .2, .2, .2)$ . All these curves are decreasing and the speed of decrease slows down after the true number of clusters, even when the signals involved in each experiment are different. This “elbow” that seems to appear with the true number of clusters can be used as an empirical criteria to decide the number of clusters. Analogous results were obtained with several different simulation schemes. Similar criteria are frequently used in cross validation methods or as an empirical tool to select a number of factors in a factor model (see Cattell, 1966). Related to clustering methods, the scree plot based on the sum of squared errors (SSE), has been considered when a K-means algorithm is used (Tibshirani, Walther, and Hastie, 2001). Thorndike (1953) and Goutte et al. (1999) discussed in-






Figure 9. Trajectory of the value of the TV distance that is achieved by the algorithm.

tuitive reasons for choosing the number of clusters based on a within-group criteria that should increase when the number of clusters decreases. In this sense we propose this empirical criteria to get the number of clusters in the analysis of real data.

### 3.3.2 Bootstrapping

The second option is based on a bootstrap procedure to approximate the distribution of the TV distance between estimated spectral densities, that was proposed in Euán (2016) and will be reported in detail in a different manuscript. Here, we will use this methodology to approximate the distribution of the total variation distance between two clusters. Then, we use this approximated distribution to choose a threshold for the TV distance between estimated spectra to decide whether or not the clusters should be merged.

The algorithm proposed in Euán (2016) to obtain a bootstrap sample of the TV distance is the following.

1. From  $X_1(t)$  and  $X_2(t)$ , estimate  $\hat{f}^{X_1}(\omega)$  and  $\hat{f}^{X_2}(\omega)$  and take  $\hat{f}(\omega) = \frac{\hat{f}_N^{X_1}(\omega) + \hat{f}_N^{X_2}(\omega)}{2}$ .
2. Draw  $Z(1), \dots, Z(T) \sim N(0, 1)$  i.i.d random variables, then estimate  $\hat{f}_N^Z(\omega)$  using also the smoothed periodogram.
3. The bootstrap spectral density will be  $\hat{f}_N^B(\omega) = \hat{f}_N(\omega) \hat{f}_N^Z(\omega)$ .

4. Repeat 2 and 3 and estimate  $\hat{d}_{TV}$  using the bootstrap spectral densities, i.e.,

$$d_{TV}^B = d_{TV}(\hat{f}_N^{B_1}, \hat{f}_N^{B_2}),$$

where  $\hat{f}_N^{B_i}$ ,  $i = 1, 2$ , are two bootstrap spectral densities using different replicates of the process  $Z(\cdot)$ .

The bootstrap spectral density presented in 2 and 3 is motivated by the method presented in Kreiss and Paparoditis (2015). Euán (2016) shows that this procedure produces a consistent estimation of the distribution of the TV distance. To extend this procedure to choose the number of clusters, we first note that due to the hierarchical structure of the algorithms used in all the methods proposed, the following tests are equivalent:

$$\begin{array}{ll} H_0 : k - 1 \text{ Clusters} & \text{vs} \quad H_A : k \text{ Clusters,} \\ H_0 : 1 \text{ Cluster} & \text{vs} \quad H_A : 2 \text{ Clusters,} \end{array}$$

since the  $(k - 1)$  clusters are built by joining two of the  $k$  clusters.

In addition, we will use this option in the method presented by Alvarez-Esteban et al. (2016) to select the number of clusters and compare with the HSM method. The distribution of the total variation distance between two clusters depends on the clustering procedure. When using the HSM method we aim to approximate the distribution of the distance between the mean spectra in each cluster while for the hierarchical clustering with the TV distance, we need to produce samples from each cluster to approximate the distribution of the distance calculated through the link function.

The procedure of this test is as follows:

- Run the clustering procedure, either the HSM method or hierarchical clustering with average or complete linkage.
- Identify the two clusters that are joined to get the  $(k - 1)$  clusters.
- Under the null hypothesis where the two clusters should be merged, denote the common representative spectral estimate to be  $\hat{f}$ .
- Simulate the spectra with the bootstrap procedure to compute the TV distance. We consider two cases:

*Case 1.* When using the HSM method simulate two spectral densities from the common spectra  $f$  and compute the TV distance between them. Repeat this procedure  $M$  times.

*Case 2.* When using hierarchical clustering with the TV distance simulate two sets of spectral densities of size  $g_1$  and  $g_2$  from the common spectra  $f$ , where  $g_i$  are the number of members in cluster  $i = 1, 2$  (clusters to be joined). Compute the link function (complete or average) between these two sets of spectra using the TV distance.

- Run the test with the bootstrap sample.

**Remark.** Notice that this test assumes that there exists a common spectra  $f$  within each cluster. In practice, it is possible that the spectra in one cluster could vary slightly and in that setting we could cast this under mixed effects model (Krafty, Hall, and Guo, 2011).

To investigate the performance of this procedure, we used **Experiments 1** and **2**. We considered the TV distance to feed a hierarchical algorithm with two different link functions, namely, average and complete. Also, we considered the HSM method with the average version. In this case, we use  $N = 500$  replicates for each experiment.

Tables 4 and 5 present the proportion of times that the null hypothesis is rejected. To reject we used the bootstrap quantile of probability  $\alpha$ . We did not expect to have a proportion of rejection equal to  $\alpha$ , since in the case of using the complete or average link, these values are not a direct observation of the TV distance. However, we expect to have a good performance. In general, it could be possible to overestimate the number of clusters.

In **Experiment 1** the true number of clusters is 2. From Table 4, we observe that all methods rejected the hypothesis of one cluster, at all the significance levels. This means that the procedures did not underestimate the number of clusters. To test 2 versus 3 clusters, the proportion of rejection is high when we used the average link function, except in the case  $\alpha = 0.01$ . If we used the complete link, the results are better. However, the best results are given by the HSM method.

In **Experiment 2** the true number of clusters is 3. This is a more difficult case, since the spectra are very close. From Table 5 when testing 2 versus 3 clusters, we observe that the complete and average link functions did not underestimate the number of clusters. However, the HSM method did not distinguish 3 clusters at a level  $\alpha = 0.01$ , but the results are better at higher levels. For testing 3 versus 4 clusters, HSM performed the best, followed by the complete link. Again it was necessary to have a small value of  $\alpha$  in order that the average link gave a reasonable performance.

Figure 10 shows the p-values obtained comparing the value from each simulation with the bootstrap distribution. We confirm the fact that the underestimation of the number of clusters has low probability, almost zero in some cases, for the three methods. When the number of clusters to test is the correct one, 2 in **Experiment 1** and 3 in **Experiment 2**, the p-values are widely distributed in the case of the complete link and HSM method. With the average link, the p-values are smaller compared to the other methods. In general, this test has a good performance when one uses the complete link or the HSM method.

Table 4. Proportion of times that the null hypothesis is rejected. Complete corresponds to the TV distance in hierarchical algorithm with the complete link function, Average with the average link, and HSM is the hierarchical spectral merger with the average version.

Experiment 1				
Test	$\alpha$	Complete	Average	HSM
1 cluster vs 2 clusters	0.01	1	1	1
	0.05	1	1	1
	0.1	1	1	1
2 clusters vs 3 clusters	0.01	0.052	0.154	0.008
	0.05	0.206	0.492	0.058
	0.1	0.382	0.670	0.164

Table 5. Proportion of times that the null hypothesis is rejected, in Experiment 2.

Experiment 2				
Test	$\alpha$	Complete	Average	HSM
2 clusters vs 3 clusters	0.01	0.968	1	0.25
	0.05	1	1	0.924
	0.1	1	1	0.998
3 clusters vs 4 clusters	0.01	0.072	0.18	0.002
	0.05	0.228	0.924	0.050
	0.1	0.376	0.998	0.106

#### 4. Analysis of the Ocean Waves and EEG Data

We developed the *HSMClust* package written in R that implements our proposed clustering method. The package can be downloaded from <http://ucispacetime.wix.com/spacetime#project-a/cxl2>.

**Example 1: EEG Data.** Our first data example is the resting-state EEG data from a single subject. The goal here is to cluster resting-state EEG signals from different channels that are spectrally synchronized, i.e., that show similar spectral profiles. This subject is from the healthy population and the EEG clustering here will serve as a “standard” to which the clustering of stroke patients (with severe motor impairment) will be compared. Data were collected at 1000 Hz and pre-processed in the following way. The continuous EEG signal was low-pass filtered at 100 Hz, segmented into non-overlapping 1 second epochs, and detrended. The original number of channels 256 had to be reduced to 194 because of the presence of artifacts that could not be corrected (e.g., loose leads).

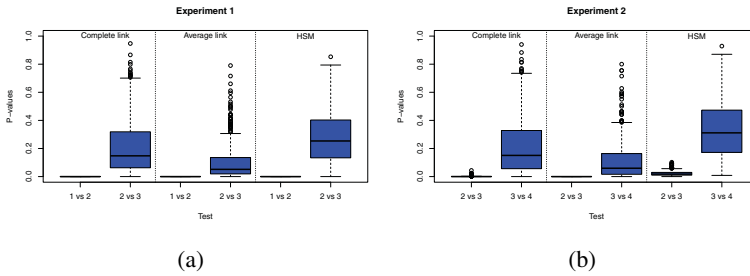



Figure 10. P-values obtained in the test of number of clusters using bootstrap samples.

The EEG channels were grouped into 19 pre-defined regions in the brain as specified in Wu et al. (2014): prefrontal (left-right), dorsolateral prefrontal (left-right), pre-motor (left-right), supplementary motor area (SMA), anterior SMA, posterior SMA, primary motor region (left-right), parietal (left-right), lateral parietal (left-right), media parietal (left-right) and anterior parietal (left-right). Figure 11 shows the locations of these regions on the cortical surface.


We present the results for subject BLAK at epoch 25 of the resting state. The interpretations will employ the usual division in frequency ranges for the analysis of spectral densities of EEG data: Delta (0.1-3 Hertz), Theta (4-7 Hertz), Alpha (8-12 Hertz), Beta (13-30 Hertz) and Gamma (31-50 Hertz).

Figure 12(a) shows the minimum value of the TV distance, in this case the “elbow” appears around 6 clusters. We analyze the results for 7 clusters as well, in this case the channels in two of the clusters were grouped into one since there were no significant evidence to reject 6 clusters over 7. Figures 12(b) and 12(c) show the shape of the mean normalized spectra by cluster and the location at the cortical surface. Many of the channels at the occipital and left premotor regions belong to cluster 6 (pink), which is dominated by frequencies at the theta and alpha bands (4-12 Hz). Cluster 2 (red) is conformed by the channels at the right premotor region and this cluster is only influenced by the delta band (0-4 Hz). Clusters 1 (black) and 5 (sky blue) are the only ones with influence of the beta band and they are located at the left motor and left anterior parietal regions.

The HSM method captures the behavior of the EEG during the resting state through the cluster membership of the EEG channels. In addition, the HSM method also identifies the frequency bands that primarily drive the formation of the clusters. The clusters produced are consistent for the most part with the anatomical-based parcellation of the cortical surface and thus cluster formation based on the spectra of the EEGs could be used to recover the spatial structure of the underlying brain process.



(a) Brain Regions



(b) EEGs

Figure 11. (a) Brain regions: Left/Right Prefrontal (L\_Pf, R\_Pr), Left/Right Dorsolateral Prefrontal (L\_dPr, R\_dPr), Left/Right Pre-motor (L\_PMd, R\_PMd), Supplementary Motor Area (SMA), anterior SMA (aSMA), posterior SMA (pSMA), Left/Right Primary Motor Region (L\_M1, R\_M1), Left/Right Parietal (L\_Pr, R\_Pr), Left/Right Lateral Parietal (L\_latPr, R\_latPr), Left/Right Media Parietal (L\_medPr, R\_medPr), Left/Right Anterior Parietal (L\_antPr, R\_antPr). Gray squared channels do not belong to any of these regions; Light blue region corresponds to right and left occipital and light green region corresponds to central occipital. (a) A recorded EEG signal in different regions. (See online version for color.)

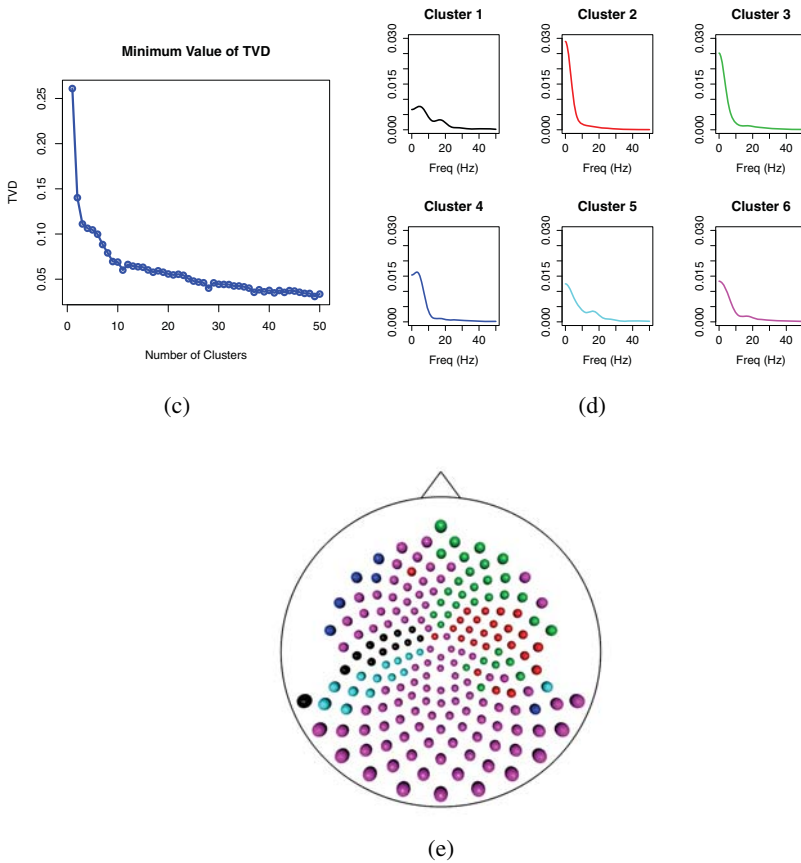


Figure 12. Results using the *HSM* function. (a) The trajectory of the minimum value of the TV distance, (b) Spectrum using all signals at each cluster and (c) Location of the clusters on the scalp. (See online version for color.)

**Example 2. Sea Waves Data.** As a second example we consider wave-height data from the Coastal Data Information Program (CDIP) buoy 160 (44098 for the National Data Buoy Center) situated off the coast of New Hampshire, USA, at a water depth of 76.5 m., when a storm took place in December 2010. The data correspond to a period of 69 hours, divided into 138 intervals of 30 minutes and recorded with a frequency of 1.28 Hz, between the 26th and the 29th of December.

The use of stochastic processes for the analysis and modeling of sea waves has a long history, starting with the work of Pierson (1955) and Longuet-Higgins (1957). A model commonly used is that stationary Gaussian random process. For this particular sea waves data, the presence of a storm implies that the sea surface will be very unstable and the hypothesis of stationarity will only be valid for short periods of time.




Figure 13. Significant wave height (black continuous line) and peak frequency (blue dashed line) for the 69 hour period (left). Trajectory for the minimum TV distance (right). (See online version for color.)

Sea-height data from buoys are usually divided into short intervals, between 20 and 30 minutes of duration, which are considered to be long enough to get a reasonable estimation of the spectral density, and short enough for stationarity to hold. For each interval, several numerical parameters are calculated from the estimated spectral density, which give a summary of the sea conditions. Two very important ones are the significant wave height and the peak frequency, or equivalently, the peak period. The former is a standard measure of sea severity and is defined as four times the standard deviation of the time series. The latter is just the modal spectral frequency, or the corresponding period. Figure 13 (left) presents the evolution of both parameters for the time interval being considered.

In Alvarez-Esteban et al. (2016) a segmentation method for long time series based on a clustering procedure was proposed and applied to a similar problem, but in a different location and different sea conditions. We will apply the same method with the clustering procedure introduced in this work. The main idea is that the spectral density should be able to identify the time intervals coming from the same stationary process. Thus, since the data are already divided into 30-minute intervals, a clustering procedure on the corresponding spectral densities should identify intervals having similar oscillatory behavior. If these intervals are contiguous in time, then they are considered part of a longer stationary interval. More details on the procedure can be found in Alvarez-Esteban et al. (2016).

As in the previous example, the trajectory for the minimum TV distance (Figure 13, right) is used to get a first estimate for the number of clusters, which in this case is 15. A bootstrap test supports this choice ( $p$ -value = 0.01). The segmentation resulting from this clustering is shown in Figure 14, where the different clusters correspond to different colors in the vertical lines.

It is interesting to compare these results with a similar analysis of a different data set in Alvarez-Esteban et al. (2016), even though the cluster-






Figure 14. Segmentation for the 96 hour period. Different colors correspond to different clusters. (See online version for color.)

Table 6. Frequency table for the duration of consecutive intervals within a common cluster.

Time (h)	0.5	1	1.5	2	2.5	3	5	6.5	7.5	8
Number of intervals	9	6	3	3	3	2	1	1	2	1

ing procedures differ. In that case, for a time series of 96 hours recorded in Waimea Bay, Hawaii in January, 2003, only 5 or 6 clusters were detected, depending on the linkage function (average or complete) used for the hierarchical agglomerative procedure. During the period studied for the Hawaii data there is also a storm, but of a much smaller magnitude, with a significant wave height always below 5 m. In contrast, for the data considered in this paper, significant wave height surpasses 7 m. at the peak of the storm. The intensity of the storm is a possible factor in the presence of more clusters of shorter duration. Table 6 gives the duration of consecutive intervals within a common cluster, with a mean value of 2.23 h and a median of 1.5 h.

Figure 15 presents the normalized spectral densities for the 15 clusters. The graphics show that bimodal spectral densities characterize some clusters, even though in some cases the secondary mode is much smaller than the main one. Bimodal spectra correspond to the presence of a second train of waves with different dominating frequency, and the presence of this secondary train of waves, even if it is small, is captured by the method as an important feature of the spectra. In other cases there are differences on the dispersion of the spectral density or in the location of the modal frequency. The procedure employed yields a spectral characterization of the different clusters, which can be linked to the statistical characteristics of their duration, something which is useful for the design of marine structures and energy converters.

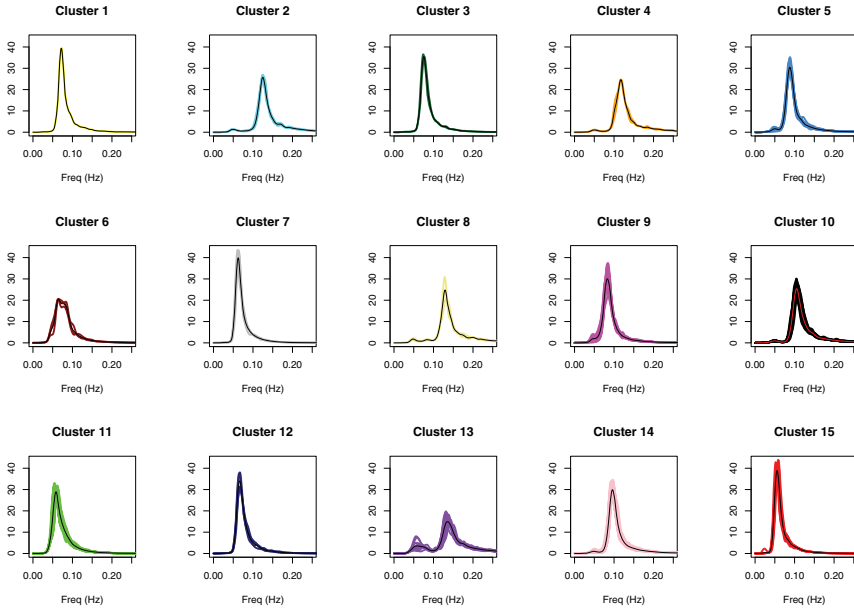


Figure 15. Normalized spectral densities for the 15 clusters, the mean spectral density is represented in black. (See online version for color.)

## 5. Discussion and Future Work

A new clustering procedure for time series based on the spectral densities and the total variation (TV) distance, the Hierarchical Spectral Merger (HSM) algorithm, was introduced. Using numerical experiments, the proposed HSM method was compared with other available procedures, and the results show that its performance is comparable to the best, in addition to having an intuitive interpretation. Applications to two data from different scientific fields were also presented, showing that the method has wide applicability in many different areas.

The HSM method was developed with the goal of building clusters with similar oscillatory behavior. Thus, two time series in the same cluster need not be correlated. Experiment 2 in the simulation study is an example where the simulated time series were independent but had the same spectral density. On the other hand, correlated time series may belong to different clusters, as occurs with the clusters built using coefficients  $e_2$  and  $e_3$ . As regards the analysis of real data, the main interest in the EEG case was to identify regions that show similar spectral behavior during resting state, i.e. those having the same active frequency bands. In the case of ocean waves, we were interested in identifying time intervals during which waves have

similar oscillatory behavior, with a view to detecting changes in the state of the sea. In this case, the presence of swell combined with local (in time) winds, may produce consecutive time intervals during which wave height time series may be correlated but have different oscillatory behavior.

However, the method is not free from limitations and further developments are needed. The fact that each cluster has a characteristic spectral density, with respect to which all distances are measured, may provide a way of correcting classification mistakes that occur due to the evolution of the clusters. This methodology would give a more robust clustering method.

The HSM method is based on estimating the spectral density where stationarity of the time series is assumed. However a modification of the method can be developed to address non-stationary cases. This could be done using evolutionary spectral features as in Maharaj (2002). Alternatively, we can consider the evolutionary wavelet spectra (instead of the Fourier spectra) for clustering (Maharaj and Alonso, 2007; Maharaj, D'Urso, and Galagedera, 2010; Maharaj and Alonso, 2014). The combination of the evolutionary wavelet spectra and the hierarchical spectral merger algorithm will give a more robust clustering results on non-stationarity cases. These extensions are left for future work.

### Acknowledgements

This work was partially supported by 1) CONACYT, México, scholarship as visiting research student, 2) CONACYT, México, Proyectos 169175 Análisis Estadístico de Olas Marinas, Fase II and 234057 Análisis Espectral, Datos Funcionales y Aplicaciones, and 3) Centro de Investigación en Matemáticas (CIMAT), A.C. Euán and Ortega wish to thank Prof Pedro C. Alvarez Esteban for several fruitful conversations on the methodology proposed on this paper. C. Euán wishes to thank to UC Irvine Space Time Modeling Group for the invitation to collaborate as a visiting scholar in their research group.

The research conducted at the UC Irvine Space-Time Modeling group (PI: Ombao) is supported in part by the National Science Foundation Division of Mathematical Sciences and the Division of Social and Economic Sciences. The authors thank Dr. Steven C. Cramer of the UC Irvine Department of Neurology for sharing the EEG data that was used in this paper. This work was done while J.O. was visiting, on sabbatical leave from CIMAT and with support from CONACYT, México, the Departamento de Estadística e I.O., Universidad de Valladolid. Their hospitality and support is gratefully acknowledged.

## References

- ALVAREZ-ESTEBAN, P.C., EUÁN, C., and ORTEGA, J. (2016), “Time Series Clustering Using the Total Variation Distance with Applications in Oceanography”, *Environmetrics*, 27, 355–369.
- BRODTKORB, P.A., JOHANNESSEN, P., LINDGREN, G., RYCHLIK, I., RYDÉN, J., and SJÖ, E. (2010), “WFO - A Matlab Toolbox for Analysis of Random Waves and Loads”, in *Proceedings of the 10th International Offshore and Polar Engineering Conference*, Vol. 3, Seattle, USA, pp. 343–350.
- CAIADO, J., CRATO, N., and PEÑA, D. (2006), “A Periodogram-Based Metric for Time Series Classification”, *Computational Statistics and Data Analysis*, 50, 2668–2684.
- CAIADO, J., CRATO, N., and PEÑA, D. (2009), “Comparison of Time Series with Unequal Length in the Frequency Domain”, *Communications in Statistics - Simulation and Computation*, 38, 527–540.
- CAIADO, J., MAHARAJ, E.A., and D’URSO, P. (2015), “Time Series Clustering”, in *Handbook of Cluster Analysis*, eds. C. Hennig, M. Meila, F. Murtagh, and R. Rocci, Handbooks of Modern Statistical Methods, Chap. 12, Chapman and Hall/CRC, pp. 241–263.
- CATTELL, R.B. (1966), “The Scree Test For The Number Of Factors”, *Multivariate Behavioral Research*, 1, 245–276.
- CONTRERAS, P., and MURTAGH, F. (2015), “Hierarchical Clustering”, in *Handbook of Cluster Analysis*, eds. C. Hennig, M. Meila, F. Murtagh, and R. Rocci, Handbooks of Modern Statistical Methods, Chap. 12, Chapman and Hall/CRC, pp. 103–123.
- EUÁN, C. (2016), “Detection of Changes in Time Series: A Frequency Domain Approach”, PhD dissertation, CIMAT.
- GAVRILOV, M., ANGUELOV, D., INDYK, P., and MOTWANI, R. (2000), “Mining the Stock Market: Which Measure is Best”, in *Proceedings of the 6th ACM International Conference on Knowledge Discovery and Data Mining*, pp. 487–496.
- GOUTTE, C., TOFT, P., ROSTRUP, E., NIELSEN, F., and HANSEN, L.K. (1999), “On Clustering fMRI Time Series”, *NeuroImage*, 9, 298–310.
- KRAFTY, R.T. (2016), “Discriminant Analysis of Time Series in the Presence of Within-Group Spectral Variability”, *Journal of Time Series Analysis*, 37, 435–450.
- KRAFTY, R.T., HALL, M., and GUO, W. (2011), “Functional Mixed Effects Spectral Analysis”, *Biometrika*, 98, 583–598.
- KREISS, J.-P., and PAPANICOLAOU, E. (2015), “Bootstrapping Locally Stationary Processes”, *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 77, 267–290.
- LIAO, T.W. (2005), “Clustering of Time Series Data – A Survey”, *Pattern Recognition*, 38, 1857–1874.
- LONGUETT-HIGGINS, M. (1957), “The Statistical Analysis of a Random Moving Surface”, *Philosophical Transactions of the Royal Society of London, Series A*, 249, 321–387.
- MAHARAJ, E., D’URSO, P., and GALAGEDERA, D. (2010), “Wavelet-Based Fuzzy Clustering of Time Series”, *Journal of Classification*, 27, 231–275.
- MAHARAJ, E.A. (2002), “Comparison of Non-Stationary Time Series in the Frequency Domain”, *Computational Statistics and Data Analysis*, 40, 131–141.
- MAHARAJ, E.A., and ALONSO, A.M. (2007), “Discrimination of Locally Stationary Time Series Using Wavelets”, *Computational Statistics and Data Analysis*, 52, 879–895.

- MAHARAJ, E.A., and ALONSO, A.M (2014), “Discriminant Analysis of Multivariate Time Series: Application to Diagnosis Based on ECG Signals”, *Computational Statistics and Data Analysis*, 70, 67–87.
- MAHARAJ, E.A., and D’URSO, P. (2011), “Fuzzy Clustering of Time Series in the Frequency Domain”, *Information Sciences*, 181, 1187–1211.
- MAHARAJ, E.A., and D’URSO, P. (2012), “Wavelets-Based Clustering of Multivariate Time Series”, *Fuzzy Sets and Systems*, 193, 33–61.
- MONTERO, P., and VILAR, J. (2014), “TsClust: An R package for Time Series Clustering”, *Journal of Statistical Software*, 62(1), 1–43
- OCHI, M.K. (1998), *Ocean Waves: The Stochastic Approach*, Cambridge, U.K: Cambridge University Press.
- PÉRTEGA DÍAZ, S., and VILAR, J.A. (2010), “Comparing Several Parametric and Non-parametric Approaches to Time Series Clustering: A Simulation Study”, *Journal of Classification*, 27, 333–362.
- PIERSON, W.J. (1955), “Wind-Generated Gravity Waves”, *Advances in Geophysics*, 2, 93–178.
- R CORE TEAM (2014), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- SHUMWAY, R.H., and STOFFER, D.S. (2011), *Time Series Analysis and Its Applications. With R Examples* (3rd ed.), New York: Springer.
- THORNDIKE, R.L. (1953), “Who Belongs in the Family”, *Psychometrika*, 18(4), 267–276.
- TIBSHIRANI, R., WALTHER, G., and HASTIE, T. (2001), “Estimating the Number of Clusters in a Data Set via the Gap Statistic”, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63, 411–423.
- WU, J., SRINIVASAN, R., KAUR, A., and CRAMER, S.C. (2014), “Resting-State Cortical Connectivity Predicts Motor Skill Acquisition”, *NeuroImage*, 91, 84–90.
- XU, R., and WUNSCH, D. (2005), “Survey of Clustering Algorithms”, *IEEE Transactions on Neural Networks*, 16, 645–678.