

---

# Clustering Noisy Signals with Structured Sparsity Using Time-Frequency Representation

---

**Anonymous Author 1**  
Unknown Institution 1

**Anonymous Author 2**  
Unknown Institution 2

**Anonymous Author 3**  
Unknown Institution 3

## Abstract

We propose a simple and efficient time-series clustering framework particularly suited for low Signal-to-Noise Ratio (SNR), by simultaneous smoothing and dimensionality reduction aimed at preserving clustering information. We extend the sparse K-means algorithm by incorporating structured sparsity, and use it to exploit the multi-scale property of wavelets and group structure in multivariate signals. Finally, we extract features invariant to translation and scaling with the scattering transform, which corresponds to a convolutional network with filters given by a wavelet operator, and use the network’s structure in sparse clustering. By promoting sparsity, this transform can yield a low-dimensional representation of signals that gives improved clustering results on several real datasets.

## 1 Introduction

Clustering of high-dimensional signals, sequences or functional data is a common task that arises in many domains [18, 19]. Such data come up in diverse fields, as in speech analysis, genomics, mass spectrometry, MRI or EEG measurements, and many more.

Clustering seeks to partition data into groups with high overall similarity between members (instances)

of the same group and dissimilarity to members of other groups. For time-series signals, this means partitioning the instances into groups of similarly behaving functions over time, where the measure of similarity is crucial and often application-specific.

In many real-world scenarios, signals are high-dimensional (such as in genomics), noisy (as in low-quality speech recordings), and exhibit non-stationary behavior: for example peaks and other non-smooth local patterns, or changes in frequency over time. In addition, signals are often subject to translations, dilations and deformations. These properties generally make clustering difficult.

Clustering is often based on the pairwise distances between signals - but in the low Signal-to-Noise Ratio (SNR) scenario these distances may be unreliable. An apparent solution is to first smooth each signal, independently of the others - but this may be sub-optimal as potentially important clustering information could be lost when noise is too high (see Figure 1). In addition, typical feature representations and distance measures are not invariant to the kind of transformations that occur to real-world signals, such as translations.

**Our contributions.** We introduce SPARCWave – Sparse Clustering with Wavelets – a framework of methods for sparse clustering of noisy signals, which **(i)** uses “global” cross-signal information for smoothing and dimensionality reduction simultaneously, by sparse clustering using time-frequency representation, **(ii)** is geared towards preserving clustering information (rather than individual signals – see Figure 1), **(iii)** exploits structure in the signal representation, such as multi-scale properties in univariate signals and interdependencies in multivariate signals, and **(iv)**

uses the scattering transform [1, 6, 7], that generates features invariant to translations and dilations and robust to small deformations in both time and frequency. We use the natural structure in the scattering coefficients in our sparse clustering method.

Our methods achieve higher clustering accuracy, compared to methods available in the literature, in simulations on both univariate and multivariate signals, and on real-world datasets from different domains. We implemented our algorithms and simulations in a python software package “SPARCWave” available in the `supp.info.`, and which will be uploaded to `github`.

## 2 Related Work

Much work has been done both recently and in the past on using time-frequency representation of signals to cluster them, primarily with wavelets. Giacomini et al. [14] first perform wavelets-based denoising for each signal, and then use model-based clustering on the union set of non-zero coefficients after shrinkage. However, informative features can be thresholded by single-signal denoising due to the lack of global information highlighting their importance (See Figure 1).

Antoniadis et al. [2] use the fact that the overall energy of a signal  $\mathbf{x} \in \mathbb{R}^T$  can be decomposed into a multi-scale view of the energy, and propose extracting features that involve averaging wavelet coefficients in each scale and measuring each scale’s relative importance. The averaging of *all* wavelets at the same scale in [2] increases SNR and makes the method robust to translations. However, this averaging loses the time information: two different signals with the same power spectrum will have very similar  $\mathbf{v}_{(j)}$  coefficients, which may lead to clustering errors.

Non-rigid transformation are typically present for real world signals, making the computed pairwise distances between signals unreliable if these are not taken into account. A common approach to dealing with deformations is the Dynamic Time Warping approach (DTW) [5], which uses dynamic programming to align pairs of signals. However, usage of DTW for our problem is limited by the fact that SNR is typically too low, such that two individual signals do not contain enough information in order to reliably align them. In

our simulations and real-data applications, we did not obtain satisfactory results with DTW.

### 2.1 Sparse Clustering

Several methods for sparse clustering have recently been proposed. For example, in [3] the authors consider a two-component Gaussian Mixture Model (GMM), and provide an efficient sparse estimation method with theoretical guarantees. While our approach could make use of any sparse clustering method, we chose sparse K-means [26] for its simplicity and ease of implementation. This simplicity allowed us to more readily incorporate our extensions to structured sparsity for both univariate and multivariate signals. We use rich feature representations, and hope for linear cluster separability in this feature space. By using these representations, we are able to achieve good results with the simpler (sparse) K-means, which corresponds to assuming a spherical GMM.

## 3 Methods

We first provide a brief description of wavelets. Next, we introduce our SPARCWave methods - using sparse K-means in the wavelet domain, incorporating group structure for univariate and multivariate signals, and finally applying the scattering transform to obtain transformation-invariant features.

### 3.1 Wavelets Background

Wavelets are smooth and quickly vanishing oscillating functions, often used to represent data such as curves/signals in time. More formally, a wavelet family  $\Psi_{j,k}$  is a set of functions generated by dilations and translations of a unique mother wavelet  $\Psi$ :

$$\Psi_{j,k}(t) = a^{\frac{j}{2}} \Psi(a^j t - k), \quad (1)$$

where  $j, k \in \mathbb{Z}$  and  $a > 0$ . The constant  $a$  represents resolution in frequency,  $j$  represents scale and  $k$  translations. We similarly define a family of functions  $\phi_{0,k}$  derived from a scaling function  $\phi(t)$  by using the dilation and translation formulation given in eq. (1).

Any function  $f \in L^2(\mathbb{R})$  can then be decomposed in




Figure 1: **Simulated cluster data and SPARCWave’s cross-signal smoothing and dimensionality reduction.** Clustering results for six clusters (I - VI): Flat curve, Heavisine, Blocks, Bumps, Doppler, Piecewise Polynomial. We generated each individual signal by applying additive Gaussian noise with  $\sigma = 2.5$ . There are 30 signals for each cluster. We ran SPARCWave as described in the text. (a) True cluster centers. (b) One individual curve from each cluster. SNR is too low to allow individual curve smoothing to reliably estimate cluster centers. (c) Cluster centers returned by SPARCWave resemble the original clusters (d) Wavelet smoothing of returned cluster centers improves estimation of original cluster centers. (e) Wavelet coefficients for each true cluster center from II. – VI. (except for the trivial flat curve) are ordered from the finest (left) to the coarsest level (right), shown in green. The  $w_i$  coefficients fitted in SPARCWave are shown in red. Although a few wavelet coefficients at fine resolutions are large for some of the curves, the informative coefficients for clustering are all at the coarsest levels.

terms of these functions:

$$f(t) = \sum_{k \in \mathbb{Z}} c_k \phi_{0,k}(t) + \sum_{j,k \in \mathbb{Z}} \mathbf{v}_{j,k} \Psi_{j,k}(t), \quad (2)$$

where  $c_k = \langle \phi_{0,k}(t), f(t) \rangle$  are the scaling coefficients, and  $\mathbf{v}_{j,k} = \langle \Psi_{j,k}(t), f(t) \rangle$  are called the wavelet coefficients of  $f(t)$ .

This decomposition has a discrete version known as the Discrete Wavelet Transform (DWT). Given  $x_1, \dots, x_T$  forming a signal  $\mathbf{x} \in \mathbb{R}^T$  sampled at times  $t = 1, \dots, T$ , where  $T = 2^J$ , and taking  $a = 2$ , the DWT of  $\mathbf{x}$  is given by taking  $c_0$  and  $v_{j,k}$  for  $j = 0, \dots, J-1$  and  $k = 0, \dots, 2^j$  in eq. (2), where  $\phi_{0,0}, \Psi_{j,k}$  are evaluated at  $t = 1, \dots, T$  to compute the (discrete) inner products  $\langle \phi_{0,0}, \mathbf{x} \rangle$  and  $\langle \Psi_{j,k}, \mathbf{x} \rangle$ . The DWT can be written conveniently in a matrix form  $\mathbf{v} = \mathbf{W}\mathbf{x}$  where  $\mathbf{W}$  is an orthogonal matrix defined by the chosen  $\Psi, \phi$ , and  $\mathbf{v}$  is the vector of coefficients  $c_0, \Psi_{j,k}$  properly ordered.

Many real-world signals are approximately sparse in the wavelet domain. This property is typically used for signal denoising using a three-step procedure [10]:

$$\hat{\mathbf{x}} \equiv \mathbf{W}^{-1} \eta_\lambda(\mathbf{W}\mathbf{x}) \quad (3)$$

where  $\eta_\lambda$  is a nonlinear shrinkage (smoothing) operator, for example entry-wise soft-threshold operator with parameter  $\lambda$ :  $\eta_\lambda(\mathbf{v}) = \text{sgn}(\mathbf{v})(|\mathbf{v}| - \lambda)_+$ .

### 3.2 Sparse Clustering with Wavelets – Univariate Signals

Consider  $n$  instances (signal vectors)  $\mathbf{x}^{(i)} \in \mathbb{R}^T$ , and let  $\mathbf{v}^{(i)}$  be their DWT transforms. Let  $d_{i_1, i_2, j}$  be the squared (Euclidean) distance between  $\mathbf{v}^{(i_1)}$  and  $\mathbf{v}^{(i_2)}$  over coordinate  $j$ ,  $d_{i_1, i_2, j} \equiv (\mathbf{v}_j^{(i_1)} - \mathbf{v}_j^{(i_2)})^2$ . Let  $C_k$  be the set of indices corresponding to cluster  $k$  with  $|C_k| = n_k$ . In standard K-means clustering, the goal is to maximize the *Between Cluster Sum of Squares* (BCSS), which can be represented as a sum  $1^T \mathcal{D}$  where  $\mathcal{D}$  is a vector representing the contribution of the wavelet coefficient  $j$  to the BCSS, with

$$\mathcal{D}_j \equiv \frac{1}{n} \sum_{i_1, i_2=1}^n d_{i_1, i_2, j} - \sum_{k=1}^K \frac{1}{n_k} \sum_{i_1, i_2 \in C_k} d_{i_1, i_2, j}. \quad (4)$$

Let  $\mathbf{w} \in \mathbb{R}^T$  be a vector of weights, and  $s$  a tuning parameter bounding  $\|\mathbf{w}\|_1$ . In sparse K-means we maximize a *weighted* sum of the contribution from each

coordinate,  $\mathbf{w}^T \mathcal{D}$  [26]. We get the following sparse K-means [26] constrained optimization problem, where sparsity is promoted in the wavelet domain:

$$\begin{aligned} & \operatorname{argmax}_{C_1, C_2, \dots, C_k, \mathbf{w}} \mathbf{w}^T \mathcal{D} \\ & \text{s.t. } \mathbf{w} \in \mathcal{B}_2(0, 1) \cap \mathcal{B}_1(0, s) \cap \mathbb{R}_+^T \end{aligned} \quad (5)$$

where  $\mathcal{B}_p(0, r) \in \mathbb{R}^T$  is the  $L_p$  ball centered at zero with radius  $r$ , and  $\mathbb{R}_+^T$  is the positive orthant,  $\mathbb{R}_+^T \equiv \{\mathbf{w} \in \mathbb{R}^T, w_j \geq 0, \forall j = 1, \dots, T\}$ .

For signals having good sparse approximation in the wavelet domain, we expect that clustering information will also be localized to a few wavelet coefficients, where high  $w_j$  values correspond to informative coefficients (See Figure 1). In [26] the authors propose an iterative algorithm for solving Problem (5), where each iteration consists of updating the clusters  $C_j$  by running K-means, and updating  $\mathbf{w}$  with a closed-form expression. The tuning parameter  $s$  is set by the permutation-based Gap Statistic method [24].

### 3.3 Group Sparse Clustering

In many cases, the features upon which we cluster could have a natural structure. A common example that has seen much interest is group structure, where features come in groups (blocks of pixels in images, neighboring genes, etc.).

Exploiting structured sparsity has been shown to be beneficial in the context of supervised learning (regression, classification), compressed sensing and signal reconstruction [4, 17]. Here, we extend the idea of using group sparsity to clustering, and propose the following **Group-Sparse K-means**.

Let  $\mathcal{G}$  be a partition of  $\{1, \dots, T\}$  and denote by  $\mathbf{v}_{(g)}$  the elements of  $w$  corresponding to group  $g \in \mathcal{G}$ . Let  $|g|$  the number of elements in  $g$ . Define the group-norm with respect to the partition  $\mathcal{G}$  [4],

$$\|\mathbf{v}\|_{\mathcal{G}} \equiv \sum_{g \in \mathcal{G}} \sqrt{|g|} \|\mathbf{v}_{(g)}\|_2 \quad (6)$$

where we multiply the vector norms  $\|\mathbf{v}_{(g)}\|_2$  by  $\sqrt{|g|}$  as in [8] and [11], to give greater penalties to larger groups (other coefficients may also be used).

Using the group norm, we define the group-sparse K-means problem, where the tuning-parameter  $s$  controls group sparsity, with  $\mathcal{B}_{(\mathcal{G})}(0, s)$  the ball of radius  $s$  centered at zero with respect to the group norm  $\|\cdot\|_{\mathcal{G}}$ .

$$\begin{aligned} & \operatorname{argmax}_{C_1, C_2, \dots, C_K, \mathbf{w}} \mathbf{w}^T \mathcal{D} \\ & \text{s.t. } \mathbf{w} \in \mathcal{B}_2(0, 1) \cap \mathcal{B}_{(\mathcal{G})}(0, s) \cap \mathbb{R}_+^T \end{aligned} \quad (7)$$

Problem (7) generalizes the sparse K-means problem in eq. (5), which is obtained here by setting all groups to have size 1.

Problem (7) can be solved iteratively as shown in Algorithm 1. ( $\sqrt{\cdot}$  and  $\odot$  in step (a.) represent element-wise vector operations:  $[\sqrt{\mathbf{w}}]_j = \sqrt{w_j}$  and  $[\mathbf{w} \odot \mathbf{v}]_j = w_j v_j$ ). Optimization *w.r.t*  $\mathbf{w}$  can be done with standard Second Order Cone Programming (SOCP) solvers by introducing auxiliary variables. We used the convex optimization toolbox CVX [15].

---

#### Algorithm 1 Group-Sparse K-means

---

**Input:**  $\mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \dots, \mathbf{v}^{(n)}$ ,  $K$ , *iters*,  $\mathcal{G}$ ,  $s$

**Output:** Clusters  $C_1, C_2, \dots, C_K$

1. **Initialize:** Set  $w_1 = w_2 = \dots = w_T = \frac{1}{\sqrt{T}}$
  2. **for**  $i=1 : \textit{iters}$ 
    - (a) Apply K-means on  $\sqrt{\mathbf{w}} \odot \mathbf{v}^{(1)}, \dots, \sqrt{\mathbf{w}} \odot \mathbf{v}^{(n)}$
    - (b) Hold  $C_1, \dots, C_K$  fixed, solve (7) *w.r.t*  $\mathbf{w}$ , with tuning parameter  $s$  and partition  $\mathcal{G}$
  3. **Return**  $C_1, \dots, C_K$
- 

We next show two applications of group-sparse clustering in the wavelets domain: (i.) using group sparsity to exploit structure of the wavelet coefficients, and (ii.) using group sparsity to exploit correlations between different curves in multivariate signal clustering.

#### 3.3.1 Exploiting Structure of Wavelet Coefficients using Group Sparsity

In our time-series clustering context, each scale  $g$  of the wavelets coefficients can be thought of as a group, together forming a multi-scale view of the signal. Since the elements of each group represent “similar” information about individual signals, it is natural to expect that they also express similar information with respect to the clustering task. Therefore, we define the parti-

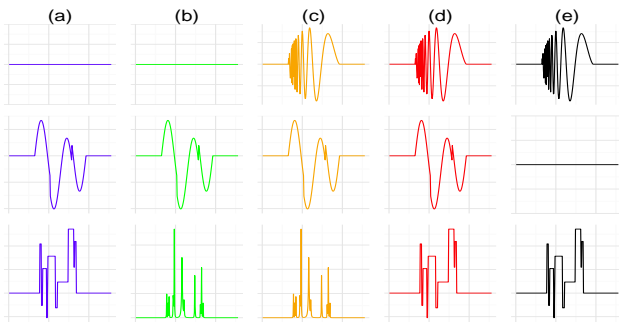


Figure 2: **Multivariate cluster data.** Each column (color) is a multivariate cluster center with  $G$  variables (here  $G = 3$ ). Each instance is  $\mathbf{X}^{(i)} \in \mathbb{R}^{G \times T}$ , with Gaussian additive noise added to each cluster center.

tions  $\mathcal{G} = \{g_j = \{w_{2^j}, \dots, w_{2^{j+1}-1}\}, j = 0, \dots, \log_2(T)\}$  with  $|g_j| = 2^j$ , and solve Problem (7) with the group norm defined by  $\mathcal{G}$ .

Our choice of  $\mathcal{G}$  in this work is not the only one possible exploiting the structure of the wavelets coefficients tree - for example, one may also choose groups corresponding to connected sub-trees (particular nodes and all of their descendent in the tree). Finally, by allowing partitions  $\mathcal{G}$  with group overlaps, it is possible to get a richer structure for wavelet coefficients, for example by using tree-structured sparsity as in [8].

### 3.3.2 Clustering Multivariate Signals with Group Sparsity

In contrast to the univariate signals discussed so far, in many cases signals are multivariate, and are composed of groups of univariate signals. A simple adaptation of group sparsity enables us to make use of this multivariate structure too.

Let each instance  $\mathbf{X}^{(i)}$  be a multivariate signal, represented as a matrix  $\mathbf{X}^{(i)} \in \mathbb{R}^{G \times T}$  where  $G$  is the number of variables and  $T$  as above (See Figure 2). We transform each row of  $\mathbf{X}^{(i)}$  with DWT to get  $\mathbf{V}^{(i)}$ . We then concatenate the rows of each  $\mathbf{V}^{(i)}$  into one vector of length  $G \times T$  and solve Problem (7) with the group norm  $\mathcal{G}$  defined by  $\mathcal{G} = \{g_t = \{t, t + T, \dots, t + GT\}, t = 1, \dots, T\}$ . Each group corresponds to a “vertical time-frequency slice” across the  $G$  univariate signals comprising  $\mathbf{X}^{(i)}$ .

Intuitively, this structure reflects the prior assumption

that if at some point  $t$  on the time-frequency axis we have an “active clustering feature”, then we expect it to be active across all  $G$  univariate signals comprising each  $\mathbf{X}^{(i)}$ , but also enables flexibility by allowing  $\{w_t, w_{t+T}, \dots, w_{t+GT}\}$  to be different.

### 3.4 The Scattering Transform

The wavelet transform is not invariant to translations and scaling. In addition, while the coarse-grain coefficients are stable under small deformations, the fine resolution coefficients are unstable [20].

To overcome the effects of the above transformations we use the scattering transform [20], a recent extension of wavelets which is stable to such deformations. The scattering representation of a signal is built with a non-linear, unitary transform computed in a manner resembling a deep convolutional network where filter coefficients are given by a wavelet operator [1].

More formally, a cascade of three operators - wavelet decomposition, complex modulus, and local averaging - is used as follows. For a function  $f(t)$  the 0-th scattering layer is obtained by a convolution,

$$S^{(0)} \equiv f \star \phi \quad (8)$$

where  $f \star g(t) \equiv \int_{s=-\infty}^{\infty} f(s)g(t-s)ds$  and  $\phi$  is an averaging filter. Applying the filter  $\phi$  provides some local translation-invariance but loses high frequency information - this information is kept by an additional convolution with a wavelet transform described next.

We start by constructing the filters we use in layer 1. Let  $\{\Psi_{j_1}^{(1)}\}_{j_1 \in J_1}$  be a filter-bank containing different dilations of a complex mother wavelet function  $\Psi$  for a set of indices  $J_1$ , with  $\Psi_{j_1}^{(1)}(t) = a_1^{j_1} \Psi(a_1^{j_1} t)$ .

To regain translation invariance lost by applying  $\Psi_{j_1}^{(1)}$ , we perform an additional convolution step with  $\phi$  and obtain the scattering functions of the first layer:

$$S_{j_1}^{(1)} \equiv |f \star \Psi_{j_1}^{(1)}| \star \phi, \quad \forall j_1 \in J_1 \quad (9)$$

where the absolute value  $|\cdot|$  is a contraction operator - this operator reduces pairwise distances between signals, and prevents explosion of the energy propagated to deeper scattering layers [20].

To recover the high frequencies lost by averaging, we again take convolutions of  $|x \star \Psi_{j_1}^{(1)}|$ , with  $\Psi_{j_2}^{(2)}$  for some  $j_2 \in J_2$ . More generally, define  $\Psi_{j_l}^{(l)}(t) = a_l^{j_l} \Psi(a_l^{j_l} t)$  where  $a_l$  determines the dilation frequency resolution at layer  $l$ . For a scattering layer  $l$ , let  $(j_1, \dots, j_l)$  define a *scale path* along the scattering network. This process above continues in the same way, with the scattering functions at the  $l$ -th layer given by

$$S_{j_1, \dots, j_l}^{(l)} \equiv | \dots | f \star \Psi_{j_1}^{(1)} | \star \Psi_{j_2}^{(2)} | \star \Psi_{j_3}^{(3)} | \star \dots \star \Psi_{j_l}^{(l)} | \star \phi,$$

$$\forall (j_1, \dots, j_l) \in J_1 \times J_2 \times \dots \times J_l. \quad (10)$$

The corresponding scattering coefficients  $v_{j_1, \dots, j_l}^{(l)}(t)$  are obtained by evaluating  $S_{j_1, \dots, j_l}^{(l)}$  at time  $t$ .

In practice, we let a resolution parameter  $\Delta_t$  determine the points at which we evaluate  $S_{j_1, \dots, j_l}^{(l)}$ , where for each  $S_{j_1, \dots, j_l}^{(l)}$  we only sample at points  $t = \frac{k\Delta_t}{2}$  with  $k = 0, \dots, \lceil \frac{2T}{\Delta_t} \rceil$ . We stop the process when reaching the final scattering layer  $L$ , obtaining scattering coefficients in layers  $l = 0, 1, \dots, L$ .

We empirically found that in many cases the scattering transform can give a sparse representation of signals, and that this sparsity helps in the clustering task (see Section 5 for real-data results).

### 3.4.1 Exploiting Scattering Group Structure

The scattering coefficients have a natural group structure which we exploit in our clustering. For each function  $S_{j_1, \dots, j_l}^{(l)}$ , we group all coefficients resulting from the evaluation of this function. We thus solve Problem (7) with the following definition of groups:

$$\mathcal{G} = \left\{ g_{j_1, \dots, j_l}, \forall l = 0, \dots, L, \forall (j_1, \dots, j_l) \in J_1 \times \dots \times J_l \right\} \quad (11)$$

where

$$g_{j_1, j_2, \dots, j_l} = \left\{ v_{j_1, \dots, j_l}^{(l)} \left( \frac{k\Delta_t}{2} \right) : k = 0, \dots, \lceil \frac{2T}{\Delta_t} \rceil \right\} \quad (12)$$

## 4 Simulation Results

For univariate signals, we take curves of dimension 256 from [10] (using the *waveband* R package) and pad them with 128 zeros on both ends resulting in  $K = 6$

cluster centers with  $T = 512$  (see Figure 1). Clusters are uniformly sized. We apply independent additive Gaussian noise  $N(0, \sigma^2)$  to generate individual signals. We clustered signals using K-means (picking the best of 100 random starts), sparse K-means on the raw data as in [26] (with the *sparcl* R package, best of 100 random starts, selecting  $s$  using the gap statistic method), and the methods of [14] (using the *curvclust* R package with burn-in period of 100) and [2] (using code provided by the authors). For our group-sparse K-means we used CVXPY [9], with tuning parameter  $s$  selected with the gap statistic method. In all relevant cases, we used the Symmlet-8 wavelet. We used the *Adjusted Rand Index (ARI)* as a measure of clustering accuracy [25]. Results are shown in Figure 3.

Group sparsity improved accuracy substantially when the number of curves  $n$  is particularly low (and thus so is SNR). When  $n$  grows, the more flexible method that does not impose group constraints picks up and the two methods reach about the same accuracy.

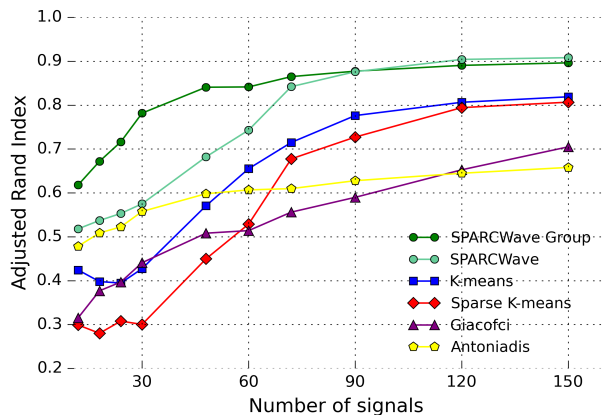


Figure 3: **Univariate simulation results.** Average ARI as a function of number of signals over 1000 simulations, for the simulation described in text with  $\sigma = 2.75$

For multivariate signals, we select curves as depicted in Figure 2. Each univariate signal is of dimension 128 padded with 64 zeros on each end, resulting in  $K = 5$  cluster centers with  $G = 3, T = 256$ . We evaluated Group-Sparse K-means and sparse K-means applied to a concatenation of all signals (termed SPARCWave-Concat). We compared our methods to (i) K-means, (ii) sparse K-means on the concatenated data, (iii) a

Hidden Markov Model approach in [13] (HMM) (code provided by the authors), and (iv) a multivariate PCA-similarity measure [27] to construct a pairwise distance matrix, which we use in spectral clustering [21] (PCA) (self-implementation).

Group sparsity improved clustering accuracy substantially, especially for low  $n$ .

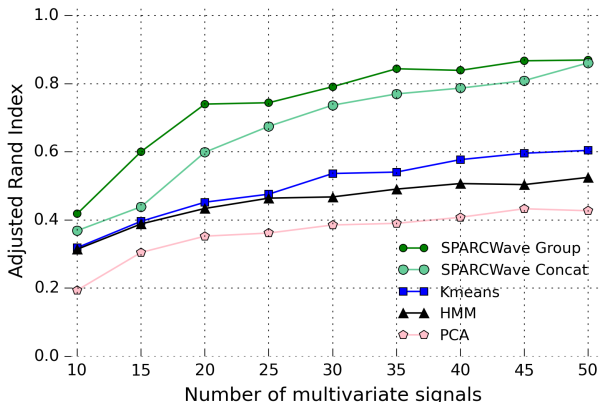


Figure 4: **Multivariate simulation results.** The average  $ARI$  is shown as a function of number of signals over 1000 simulations, for the simulation described in text with  $\sigma = 1.75$ .

## 5 Real-Data Applications

We tested the SPARCWave framework on three real-world datasets. In all cases, using the scattering transform improved clustering accuracy. We used the implementation in the *scatnet* MATLAB software (available at <http://www.di.ens.fr/data/software/>). In all cases, we used complex Morelet wavelets for all scattering layers and a moving average for  $\phi$ . See the Appendix for specific parameters used in our implementations. Results for the different methods on all datasets are shown in Table 1.

### 5.1 Berkeley Growth Study

The Berkeley growth dataset [22] consists of height measurements for 54 girls and 38 boys between the ages of 1 and 18 years giving  $T = 30$  data points for each signal. Of interest to researchers is the “velocity” of the growth curve, i.e. the rate of change. We thus use first-order differences of the curves. However, as a result of time shifts among individuals, simply looking

at the cross-sectional mean fails to capture important common growth [23], which is also not fully captured by the standard wavelet transform. We applied the scattering transform together with sparse clustering, which lead to substantial improvements in accuracy. Applying sparse K-means on the scattering features led to the about the same results as ordinary K-means in this feature space, with  $ARI = 0.87$ ).

The weights  $w_j$  selected in the optimization problems (5) or (7) may be sub-optimal for clustering. The features with high-magnitude coefficients can, however, be used in standard K-means, viewing SPARCWave as a feature-selection step [26]. We can select all non-zero features  $w_j$  (or above some threshold) and apply K-means using only these features. This technique improved accuracy from 0.87 to 0.91.

Group-sparsity clustering with the scattering transform lead to similar results, possibly due to the very short signal length  $T$  (see Table 1 and Figure 5).

### 5.2 Phoneme log-periodograms

The Phoneme dataset contains 4507 noisy log-periodograms corresponding to recordings of speakers of 32 ms duration with  $T = 256$ . There are  $K = 5$  phonemes: “sh” (800 instances), “dcl” (757), “iy” (1163), “aa” (695), “ao” (1022) (See [16] and <http://statweb.stanford.edu/~tibs/ElemStatLearn/datasets/phoneme.info.txt>).

Each instance is a noisy estimate of a spectral-density function. De-noising, as suggested for example in [12], is thus attractive here. By applying the scattering transform on data in the log-frequency domain, we obtain invariance to translation in log-frequency and consequently to scaling in frequency. This is similar to the approach taken in [1], where the scattering transform is applied to log-frequencies. Here too group-sparsity clustering did not change accuracy, possibly due to large sample size (see Table 1 and Figure 7 in the Appendix).

### 5.3 Wheat Spectra

The Wheat dataset consists of near-infrared reflectance spectra of 100 wheat samples with  $T = 701$ , measured in  $2nm$  intervals from 1100 to 2500nm, and an associated response variable, the samples’ moisture

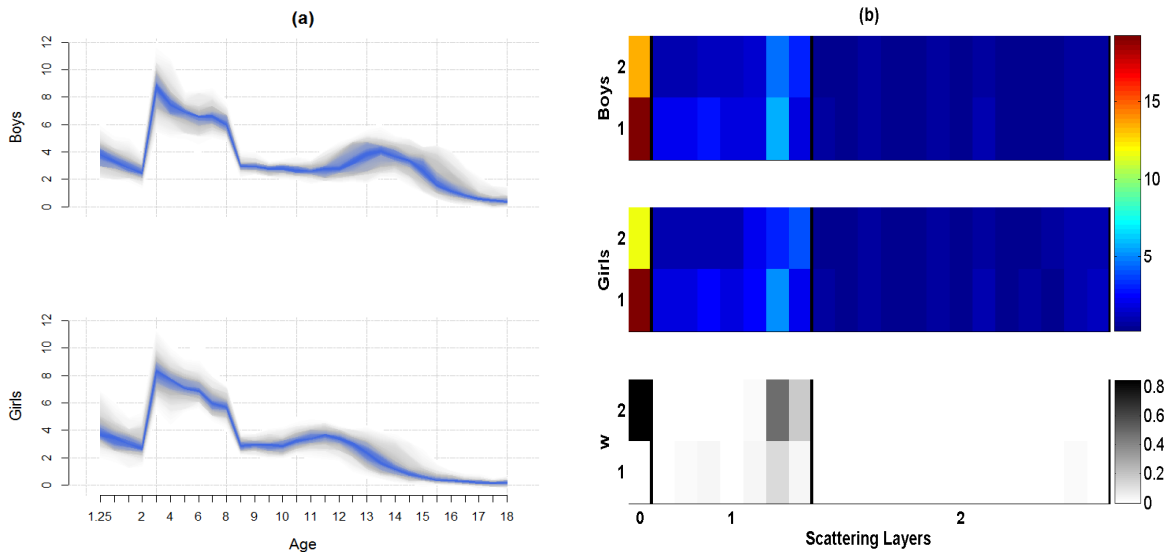


Figure 5: **Clustering results for Berkeley Growth dataset.** (a) Fanplots showing the distribution of first-order differences of the curves in each cluster, representing the growth rate for boys and girls as function of age. Intra-class variation increases at ages around 12 – 16 (indicated by larger faded gray area), motivating the use of the scattering transform which is robust to deformations. (b) In the first two heat maps, scattering coefficients are shown for the growth rate curves in each cluster, with colors representing coefficient magnitude. Each row represents a point  $t$  at which we evaluate the coefficient (see Section 3.4). The third heat map represents the SPARCWave weights  $w$  corresponding to each coefficient.

content. The response variable is clearly divided into two groups: 41 instances with moisture content lower than 15, and 59 with moisture content above 15. We use this grouping to create 2 classes. Applying the scattering transform in combination with sparse clustering improved accuracy significantly. However, using group sparsity reduced accuracy, perhaps because out of 924 features only 2 are found to have non-zero weights (see Table 1 and Figure 8 in the Appendix).

Table 1: Clustering accuracy ( $ARI$ ) for three real datasets. The SPARCScatter method shows superior clustering accuracy over all datasets.

Method	Growth	Phoneme	Wheat
SPARCScatter	<b>0.91</b>	<b>0.73</b>	<b>0.46</b>
SPARCScatter Group	<b>0.91</b>	<b>0.73</b>	0.3
Scattering + K-means	0.87	0.66	0.30
SPARCWave	0.62	0.30	0.31
Antoniadis [2]	0.62	0.34	0.35
Giacofci [14]	0.58	0.69	0.30
K-means	0.58	0.68	0.31

## 6 Conclusion and Further Work

We proposed a method for time-series clustering that uses a “built-in” shrinkage of wavelet coefficients based

on their contribution to the clustering information. We extended the sparse K-means framework by incorporating group structure and used it to exploit wavelet multi-resolution properties in univariate signals, and multivariate structure. An interesting future direction is to adapt this approach to other sparse clustering methods. Another direction is to incorporate richer structures, such as tree-sparsity in the wavelet and scattering coefficients, long-range dependencies and interdependencies in multivariate signals. In this work we applied sparse clustering to one-dimensional signals. The wavelets transform is widely used to represent two-dimensional signals such as images. In addition, a two-dimensional scattering transform achieved excellent results in supervised image classification tasks due to its invariance to translation, dilation and deformation [6, 7]. It is thus natural to apply our approach to sparse clustering of images and other multi-dimensional datasets. Finally, there are few known theoretical guarantees for sparse-clustering methods, and it would be interesting to develop such guarantees to the framework in [26] and its group extension we have proposed.



## References

- [1] J. Andén and S. Mallat. Deep scattering spectrum. *IEEE Transactions on Signal Processing*, 62(16):4114–4128, 2014.
- [2] A. Antoniadis, X. Brossat, J. Cugliari, and J.-M. Poggi. Clustering functional data using wavelets. *International Journal of Wavelets, Multiresolution and Information Processing*, 11(01):1350003, 2013.
- [3] M. Azizyan, A. Singh, and L. Wasserman. Efficient sparse clustering of high-dimensional non-spherical gaussian mixtures. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.
- [4] F. Bach, R. Jenatton, J. Mairal, G. Obozinski, et al. Convex optimization with sparsity-inducing norms. *Optimization for Machine Learning*, pages 19–53, 2011.
- [5] D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pages 359–370. Seattle, WA, 1994.
- [6] J. Bruna and S. Mallat. Classification with scattering operators. *arXiv preprint arXiv:1011.3023*, 2010.
- [7] J. Bruna and S. Mallat. Invariant scattering convolution networks. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1872–1886, 2013.
- [8] C. Chen and J. Huang. Compressive sensing MRI with wavelet tree sparsity. In *Advances in neural information processing systems*, pages 1124–1132, 2012.
- [9] S. Diamond, E. Chu, and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization, version 0.2. <http://cvxpy.org/>, May 2014.
- [10] D. L. Donoho. De-noising by soft-thresholding. *IEEE Transactions on Information Theory*, 41(3):613–627, 1995.
- [11] J. Friedman, T. Hastie, and R. Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.
- [12] H.-Y. Gao. Choice of thresholds for wavelet estimation of the log spectrum. *Stanford Technical Report*, 1993.
- [13] S. Ghassempour, F. Girosi, and A. Maeder. Clustering multivariate time series using hidden markov models. *International journal of environmental research and public health*, 11(3):2741–2763, 2014.
- [14] M. Giacomci, S. Lambert-Lacroix, G. Marot, and F. Picard. Wavelet-based clustering for mixed-effects functional models in high dimension. *Biometrics*, 69(1):31–40, 2013.
- [15] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- [16] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: Data mining, inference and prediction, 2nd ed.* Springer, New York, 2009.
- [17] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. *The Journal of Machine Learning Research*, 12:3371–3412, 2011.
- [18] J. Jacques and C. Preda. Functional data clustering: a survey. *Advances in Data Analysis and Classification*, pages 1–25, 2014.
- [19] W. T. Liao. Clustering of time series data: a survey. *Pattern recognition*, 38(11):1857–1874, 2005.
- [20] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- [21] A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 2:849–856, 2002.
- [22] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis, 2nd ed.* Springer, New York, 2006.
- [23] R. Tang and H.-G. Müller. Pairwise curve synchronization for functional data. *Biometrika*, 95(4):875–889, 2008.
- [24] R. Tibshirani, G. Walther, and T. Hastie. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):411–423, 2001.
- [25] S. Wagner and D. Wagner. *Comparing clusterings: an overview.* Universität Karlsruhe, Fakultät für Informatik Karlsruhe, 2007.
- [26] D. M. Witten and R. Tibshirani. A framework for feature selection in clustering. *Journal of the American Statistical Association*, 105(490):713–726, 2010.
- [27] K. Yang and C. Shahabi. A PCA-based similarity measure for multivariate time series. In *Proceedings of the 2nd ACM international workshop on Multimedia databases*, pages 65–74. ACM, 2004.