



Unsupervised Vehicle Recognition Using Incremental Reseeding of Acoustic Signatures

Justin Sunu^{1(✉)}, Allon G. Percus^{1(✉)}, and Blake Hunter^{2(✉)}

¹ Claremont Graduate University, Claremont, CA 91711, USA
justinsunu@gmail.com, allon.percus@cgu.edu

² Microsoft, Redmond, WA 98052, USA
blake.hunter@microsoft.com

Abstract. Vehicle recognition and classification have broad applications, ranging from traffic flow management to military target identification. We demonstrate an unsupervised method for automated identification of moving vehicles from roadside audio sensors. Using a short-time Fourier transform to decompose audio signals, we treat the frequency signature in each time window as an individual data point. We then use a spectral embedding for dimensionality reduction. Based on the leading eigenvectors, we relate the performance of an incremental reseeded algorithm to that of spectral clustering. We find that incremental reseeded accurately identifies individual vehicles using their acoustic signatures.

Keywords: Spectral clustering · Machine learning · Vehicle audio

1 Introduction

Recognizing and distinguishing moving vehicles based on their audio signals are problems of broad interest. Applications range from traffic analysis and urban planning to military vehicle recognition. Audio data sets are small compared to video data, and multiple audio sensors can be placed easily and inexpensively. However, challenges arise due to equipment as well as to the underlying physics. Microphone sensitivity can result in disruption from wind and ambient noise. The Doppler shift can make a vehicle's acoustic signature differ according to its position.

In order to interpret acoustic signatures, one must extract information contained within the raw audio data. A natural feature extraction method is the short-time Fourier transform (STFT), using time windows large enough to carry sufficient frequency information but small enough to localize vehicle events. STFT has been used previously for classifying cars vs. motorcycles with principle component analysis [13], for characterizing ϵ -neighborhoods in vehicle frequency signatures [14], for vehicle classification based on power spectral density [1], and for vehicle engine identification [8, 12]. Other related feature extraction

approaches have included the wavelet transform [1,2] and the one-third-octave filter bands [10].

Our study is motivated by previous work that uses a spectral embedding approach to identify different individual vehicles [11]. Representing each time window as an individual data point, we define a similarity measure between two points based on the cosine distance between their sets of Fourier coefficients, and then cluster according to the symmetric normalized graph Laplacian [9]. In this paper, we relate the eigenvectors of the Laplacian to a recently proposed clustering method, incremental reseeding (INCREs) [3], that iteratively propagates cluster labels across a graph. We compare the performance of INCREs with spectral clustering on the vehicle audio data. We find that both are promising unsupervised methods for vehicle identification, with INCREs correctly clustering 91.7% of the data points in a sequence of passages of three different vehicles.

2 Algorithms

Our clustering algorithms are based on the use of spectral embedding for dimensionality reduction [11]. Consider a signal of length n , with feature vector $\mathbf{x}_i \in \mathbb{R}^m$ associated with data point $i \in \{1, \dots, n\}$. A spectral embedding represents data as vertices on a weighted graph, with edge weights S_{ij} expressing a similarity measure between data points i and j . The graph is encoded using the symmetric normalized graph Laplacian matrix [9]

$$\mathbf{L}_s = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{S} \mathbf{D}^{-1/2} \quad (1)$$

where \mathbf{D} is a diagonal matrix with $D_{ii} = \sum_j S_{ij}$.

We use this embedding for vehicle identification with two related clustering methods: spectral clustering, and a recently developed incremental reseeding approach [3].

2.1 Spectral Clustering

The eigenvectors of \mathbf{L}_s associated with the leading nontrivial eigenvalues $\lambda_2, \dots, \lambda_k$ form a $(k-1)$ -dimensional approximation to \mathbf{x}_i . The approximation is justified when the spectral gap $|\lambda_{k+1} - \lambda_k|$ is large, which occurs when the data naturally form k clusters [9]. Spectral clustering uses k -means to cluster this \mathbb{R}^{k-1} projection of the data.

2.2 Incremental Reseeding (INCREs) Algorithm

The INCREs algorithm [3] is a diffusive method that propagates cluster labels across the graph specified by \mathbf{L}_s . The approach (Algorithm 1) is incremental: it plants cluster seeds among nodes, grows clusters from these seeds, and then reseeds among the growing clusters.

Algorithm 1. INGRES

```

1: Input Similarity matrix  $S$ , number of clusters  $k$ , number of iterations  $s$ 
2: Initialize Randomly partition data points into  $k$  clusters
3: for  $i = 1$  to  $s$  do
4:   PLANT: Randomly choose certain points in each cluster as seeds
5:   GROW: Apply Laplacian  $\mathbf{L}_s$  to matrix of seed labels
6:   HARVEST: Assign data points to cluster with largest label value
7: end for

```

To elaborate on the algorithm, an indicator matrix of seed labels is constructed in the PLANT step: $U_{ij} = 1$ if point i is picked as a seed for cluster j , and 0 otherwise. By applying the Laplacian in the GROW step, $\mathbf{L}_s \mathbf{U}$ becomes a real-valued matrix representing the seeds propagating to neighbors according to similarity. Finally, the HARVEST step updates the assignment of data points to clusters, by finding for each point i the cluster j maximizing the propagated label matrix element $(\mathbf{L}_s \mathbf{U})_{ij}$. The process then repeats, with new seeds drawn from the updated clusters.

Since the process of seed propagation is governed by the graph Laplacian, INGRES is closely connected with spectral clustering. Eigenvectors of \mathbf{L}_s are organized hierarchically: the second eigenvector separates data into two clusters at the coarsest resolution, the third eigenvector identifies a third cluster at a finer resolution, and so on. An application of INGRES with parameter k propagates seeds with k different labels through the graph, resulting in k clusters governed by the spectral properties of \mathbf{L}_s .

To illustrate the relation between the two algorithms, consider a similarity matrix S_{ij} given in block matrix form, with added “salt and pepper” noise. This is shown in Fig. 1, with lighter colors representing greater similarities. Figure 2 shows the second and third eigenvectors of \mathbf{L}_s , along with the results of INGRES using $k = 2$ and $k = 3$. The binary clustering results of both methods split the data into the same two larger classes, using only the second eigenvector for spectral clustering and using $k = 2$ for INGRES. When the third eigenvector is also used for spectral clustering and when $k = 3$ is used for INGRES, the two methods find the same subdivision of one of these two larger classes.

In less straightforward clustering examples, the reseeding process can allow INGRES to learn partitions that are not apparent to spectral clustering. Furthermore, the formulation of INGRES allows it to be applied even in cases of larger datasets where eigenpairs cannot be easily computed.

3 Data and Feature Extraction

Our audio data consists of recordings, provided by the US Naval Air Systems Command, of different vehicles moving multiple times through a parking lot at approximately 15mph. The original dataset consists of MP4 videos taken from a roadside camera; we extract the dual channel audio signal, and average the

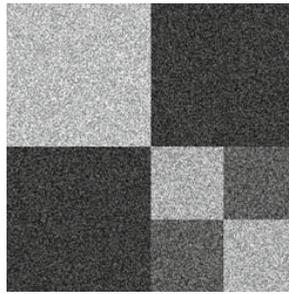


Fig. 1. Synthetic similarity matrix with salt and pepper noise. White represents a similarity value of 1; black represents a similarity value of 0.

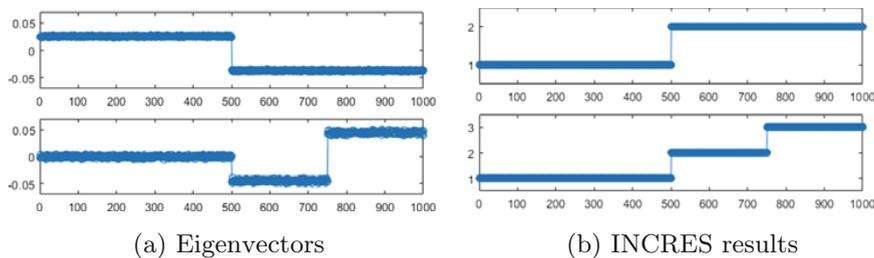


Fig. 2. (a) 2nd and 3rd eigenvectors of \mathbf{L}_s , and (b) INCREs results with $k = 2$ and $k = 3$, using the synthetic similarity matrix in Fig. 1. Note similar separation into two and three clusters.

channels together into a single channel. The audio signal has a sampling rate of 48,000 samples per second. Video information is used only to ascertain ground truth (vehicle identification) for training data.

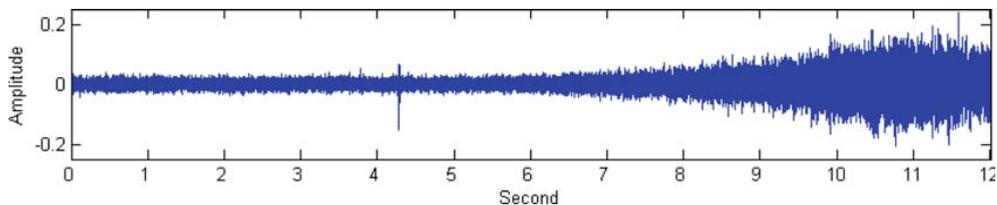


Fig. 3. Raw audio signal for a vehicle passage.

Each extracted audio signal is a sequence of a vehicle approaching from a distance, becoming audible after 5 or 6 s, passing the microphone after 10 s, and then leaving. An example of the raw audio signal is shown in Fig. 3. We form a composite sequence, shown in Fig. 4, from multiple passages of three different vehicles (a white truck, black truck, and jeep), selecting the two seconds where the vehicle is closest to the camera. The goal is to test the clustering algorithm's ability to differentiate the vehicles.

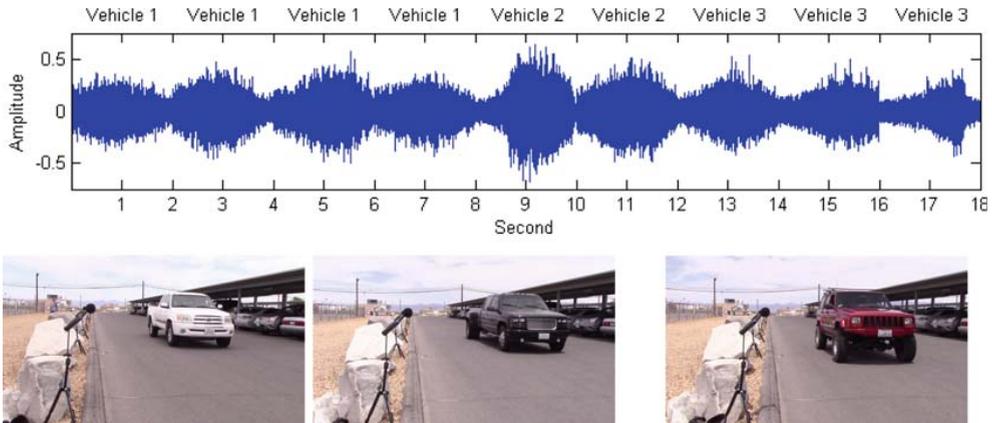


Fig. 4. Raw audio signal for composite data. Images show the three different vehicles, as seen in accompanying video (not used for analysis).

We preprocess the data by grouping audio samples into larger windows. With windows of $1/8$ of a second, or 6000 samples per window, we find both a sufficient number of windows and sufficient information per window. While there is no clear standard in the literature, this window size is comparable to those used in other studies [13]. Discontinuities between successive windows can in some cases be reduced by applying a Hamming window as a filter, or by allowing overlap between windows [13]. However, in our study we found no conclusive benefit from either of these, and therefore used standard box windows with no overlap.

Relevant features are extracted from the raw audio signal using the short-time Fourier transform (STFT). The Fourier decomposition contains 6000 symmetric coefficients, leaving 3000 usable coefficients. Figure 5 shows the first 1000 Fourier coefficients, using a moving mean of 5 samples, for a time window representing a truck passing and a time window representing a sedan passing, both in similar positions. Note that a clear frequency signature is apparent for each vehicle, with much of the signal concentrated within the first 250 coefficients or 2000 Hz.

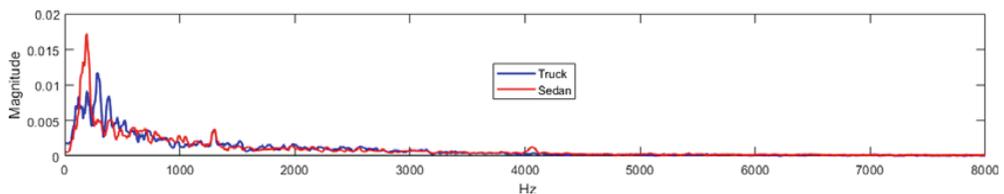


Fig. 5. First 1000 Fourier coefficients for a truck and a car, after applying a moving mean of size 5.

Each time window in the audio signal is taken as an independent data point to be clustered: we define the feature vector $\mathbf{x}_i \in \mathbb{R}^m$ as the set of m Fourier

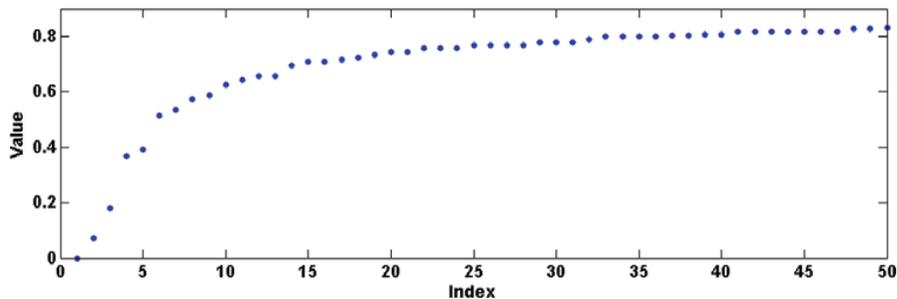


Fig. 6. Spectrum of \mathbf{L}_s for vehicle data. Largest gap is after third eigenvalue.

coefficients associated with that window. Since many of these coefficients are relatively insignificant, we consider the cosine distance measure between data points

$$d_{ij} = 1 - \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|}. \quad (2)$$

We then construct an M -nearest neighbor graph, where the edge $\{i, j\}$ is present if j is among the M closest neighbors of i or vice-versa, for a fixed value of M . Following standard methods [9], the similarity S_{ij} is taken to be a Gaussian function of distance,

$$S_{ij} = e^{-d_{ij}^2/\sigma_i^2}, \quad (3)$$

where σ_i is defined adaptively [15] as the distance to vertex i 's M th neighbor.

4 Results

Our composite vehicle dataset contained the 18s of raw audio shown in Fig. 4, resulting in $n = 144$ data points each representing 1/8-second time windows. We used only the first $m = 1500$ Fourier coefficients, reducing the risk of artifacts from window discontinuities. We set $M = 15$ for the M -nearest neighbor graph, so that neighborhoods contain the 16 data points used in the 2-second clips of a single vehicle passage.

Figure 6 shows the eigenvalues of the Laplacian for the vehicle data. The largest gap follows the third eigenvalue, consistent with three clusters representing the three vehicles actually present in the data. We therefore set $k = 3$ for both spectral clustering and INCRES.

Figure 7 shows the second and third eigenvectors of \mathbf{L}_s and Fig. 8 shows typical results of INCRES for $k = 2$ and $k = 3$ (INCRES is stochastic, but results vary little from run to run). As in our earlier synthetic example, the second eigenvector and $k = 2$ INCRES result provide comparable binary separations of the data. Thresholding the eigenvector just above zero would place all of the vehicle 1 data in one cluster, and most of the vehicle 2 and 3 data in the other cluster (the exceptions are primarily data points at the beginning and end of a vehicle passage, where the signal is weakest). The third eigenvector

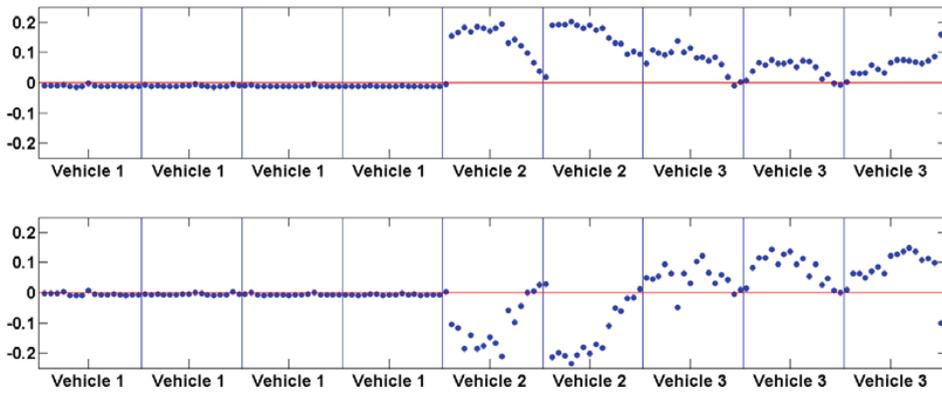


Fig. 7. 2nd and 3rd eigenvectors of L_S for vehicle data.

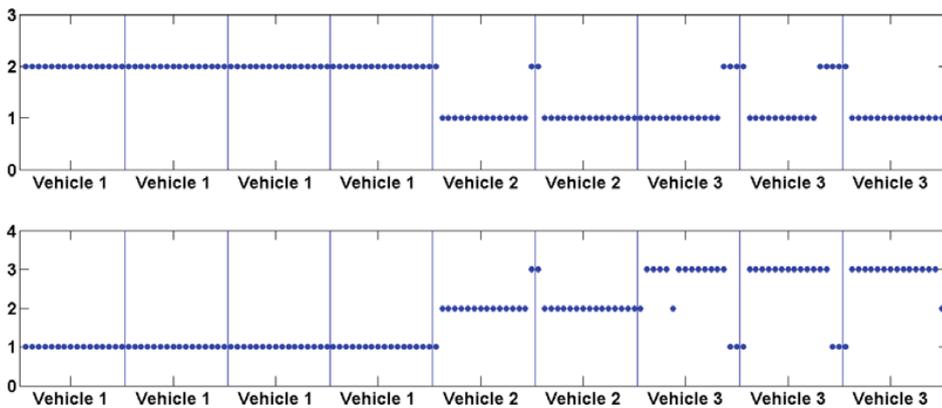


Fig. 8. INGRES results for vehicle data.

mostly distinguishes vehicle 2 (negative values) and vehicle 3 (positive values). The $k = 3$ INGRES result recognizes the three vehicles very accurately, and is discussed below.

Note that unlike in the straightforward synthetic data problem, the third eigenvector is not by itself sufficient to separate the three clusters. Figure 9 shows the results of k -means clustering, with $k = 3$, on the third eigenvector alone. While all vehicle 1 data points are clustered together, a significant fraction of vehicle 2 and 3 data points are incorrectly placed in that cluster as well.

Figure 10 shows results of the more conventional spectral clustering method, using k -means on the \mathbb{R}^2 projection of the data given by the 2nd and 3rd eigenvectors. The inclusion of the 2nd eigenvector is sufficient to cluster the vast majority of vehicle 2 and 3 data points correctly.

Tables 1 and 2 interpret the spectral clustering results of Fig. 10 and the INGRES $k = 3$ results of Fig. 8 as classifications. Both methods classify all of vehicle 1 correctly, but INGRES performs noticeably better than spectral clustering on vehicle 2, and they perform comparably on vehicle 3. Overall purity

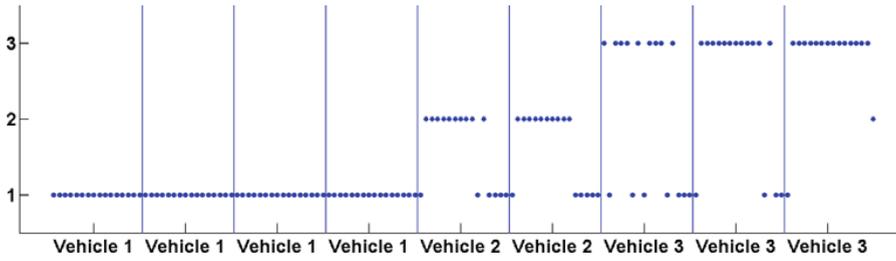


Fig. 9. k -means on third eigenvector of L_s for vehicle data.

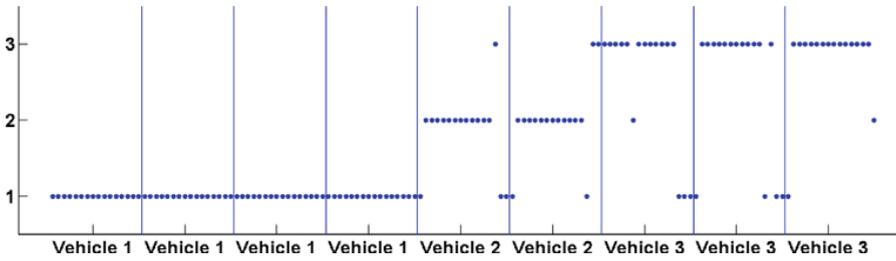


Fig. 10. k -means on second and third eigenvectors of L_s (standard spectral clustering) for vehicle data.

Table 1. Vehicle clustering results using spectral clustering.

| True | Obtained cluster | | |
|-------------------------|-------------------------|-------------------------|------------------|
| | Vehicle 1 (white truck) | Vehicle 2 (black truck) | Vehicle 3 (jeep) |
| Vehicle 1 (white truck) | 64 | 0 | 0 |
| Vehicle 2 (black truck) | 5 | 24 | 3 |
| Vehicle 3 (jeep) | 8 | 2 | 38 |

Table 2. Vehicle clustering results using INCRES with $k = 3$.

| True | Obtained cluster | | |
|-------------------------|-------------------------|-------------------------|------------------|
| | Vehicle 1 (white truck) | Vehicle 2 (black truck) | Vehicle 3 (jeep) |
| Vehicle 1 (white truck) | 64 | 0 | 0 |
| Vehicle 2 (black truck) | 1 | 29 | 2 |
| Vehicle 3 (jeep) | 6 | 3 | 39 |

scores are 87.5% for spectral clustering, and 91.7% for INCRES, with misclassifications again occurring primarily at the beginning or end of a vehicle passage.

The runtimes for both algorithms in Matlab were both under one second, insignificant compared to the minutes needed for processing the initial data.

5 Conclusions

We have presented a method to identify moving vehicles from audio recordings, by clustering their frequency signatures with an incremental reseeded method (INCRESES) [3]. Motivated by the approach in [11], we decompose the audio signal with a short-time Fourier transform (STFT), and treat each 1/8-second time window as an individual data point. We then apply a spectral embedding and consider the symmetric normalized graph Laplacian. We find that spectral clustering, which uses the leading eigenvectors of the Laplacian, correctly clusters 87.5% of the data points. INCRESES, which directly uses the Laplacian to construct a random walk on the graph, correctly clusters 91.7% of the data points. Almost all incorrectly clustered points lie at the very beginning or very end of a vehicle passage, when the vehicle is furthest from the recording device. The vast majority of data points result in correct vehicle recognition.

We observe that there is a close relation between the k th eigenvector and the INCRESES output for k clusters. This suggests that clustering results might be improved by simultaneously taking the INCRESES output for 2 through k clusters, and then using k -means on this \mathbb{R}^{k-1} projection of the data just as spectral clustering does on the 2nd through k th eigenvectors. While doing so does not noticeably change our INCRESES $k = 3$ results, the difference could be significant for larger values of k . This could be tested, using a dataset with a larger number of vehicles.

Finally, we note that, since time windows are treated as independent data points, our approach ignores most temporal information. Explicitly taking advantage of the time-series nature of our data in the clustering algorithm could improve results, by clustering data points according not only to their own frequency signatures but also to those of preceding or subsequent time windows. Furthermore, while the STFT is a standard method for processing audio signals, it suffers from two drawbacks: the use of time windows imposes a specific time scale for resolving the signal that may not always be the appropriate one, and vehicle sounds may contain too many distinct frequencies for the Fourier decomposition to yield easily learned signatures. These difficulties may best be addressed by using multiscale techniques such as wavelet decompositions that have been proposed for vehicle detection and classification [1, 2], as well as more recently developed sparse decomposition methods that learn a set of basis functions from the data [4–7].

Acknowledgments. The authors are grateful to Dr. Arjuna Flenner at the US Navy’s Naval Air Systems Command (NAVAIR) for having supplied the vehicle data. We also wish to thank the anonymous reviewers for their comments and suggestions, which helped improve the clarity of the paper.

References

1. Aljaafreh, A., Dong, L.: An evaluation of feature extraction methods for vehicle classification based on acoustic signals. In: IEEE International Conference on Networking, Sensing and Control (2010)
2. Averbuch, A., Zheludev, V.A., Rabin, N., Schclar, A.: Wavelet-based acoustic detection of moving vehicles. *Multidimension. Syst. Signal Process.* **20**, 55–80 (2009)
3. Bresson, X., Hu, H., Laurent, T., Szlam, A., von Brecht, J.: An incremental reseed-ing strategy for clustering. *Mathematics and Visualization* (2018, to appear)
4. Chui, C.K., Mhaskar, H.: Signal decomposition and analysis via extraction of fre-quencies. *Appl. Comput. Harmonic Anal.* **40**(1), 97–136 (2016)
5. Daubechies, I., Lu, J., Wu, H.T.: Synchrosqueezed wavelet transforms: an empiri-cal mode decomposition-like tool. *Appl. Comput. Harmonic Anal.* **30**(2), 243–261 (2011)
6. Gilles, J.: Empirical wavelet transform. *IEEE Trans. Sign. Process.* **61**(16), 3999–4010 (2013)
7. Hou, T.Y., Shi, Z.: Adaptive data analysis via sparse time-frequency representa-tion. *Adv. Adapt. Data Anal.* **03**(01n02), 1–28 (2011)
8. Kozhisseri, S., Bikdash, M.: Spectral features for the classification of civilian vehi-cles using acoustic sensors. In: IEEE Workshop on Computational Intelligence in Vehicles and Vehicular Systems (2009)
9. Luxburg, U.V.: A tutorial on spectral clustering. *Stat. Comput.* **17**(4), 395–416 (2007)
10. Rahim, N.A., Paulraj, P.M., Adom, A.H., Kumar, S.S.: Moving vehicle noise clas-sification using multiple classifiers. In: IEEE Student Conference on Research and Development (2011)
11. Sunu, J., Percus, A.G.: Dimensionality reduction for acoustic vehicle classifica-tion with spectral embedding. In: IEEE International Conference on Networking, Sensing and Control, pp. 129–133 (2018)
12. Wierzchowska, A., Kubera, E., Słowik, T., Skrzypiec, K.: Spectral features for audio based vehicle and engine classification. *J. Intell. Inf. Syst.* **50**(2), 265–290 (2018)
13. Wu, H., Siegel, M., Khosla, P.: Vehicle sound signature recognition by frequency vector principle component analysis. *IEEE Trans. Instrum. Meas.* **48**, 1005–1009 (1999)
14. Yang, S.S., Kim, Y.G., Choi, H.: Vehicle identification using wireless sensor net-works. In: IEEE SoutheastCon (2007)
15. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. *Adv. Neural Inf. Proc-ess. Syst.* **17**, 1601–1608 (2004)