

BLIND ONE-MICROPHONE SPEECH SEPARATION: A SPECTRAL LEARNING APPROACH

Francis R. Bach

Ecole des Mines de Paris
35, rue Saint-Honoré
77300 Fontainebleau, France
francis.bach@mines.org

Michael I. Jordan

Computer Science and Statistics
University of California
Berkeley, CA 94720, USA
jordan@cs.berkeley.edu

ABSTRACT

We present an algorithm to perform blind, one-microphone speech separation. Our algorithm separates mixtures of speech without modeling individual speakers. Instead, we formulate the problem of speech separation as a problem in segmenting the spectrogram of the signal into two or more disjoint sets. We build feature sets for our segmenter using classical cues from speech psychophysics. We then combine these features into parameterized affinity matrices. We also take advantage of the fact that we can generate training examples for segmentation by artificially superposing separately-recorded signals. Thus the parameters of the affinity matrices can be tuned using recent work on learning spectral clustering [1]. This yields an adaptive, speech-specific segmentation algorithm that can successfully separate one-microphone speech mixtures.

1. INTRODUCTION

The problem of recovering signals from linear mixtures, with only partial knowledge of the mixing process and the signals—a problem often referred to as *blind source separation*—is a central problem in signal processing. It has applications in many fields, including speech processing, network tomography and biomedical imaging [2]. When the problem is over-determined, i.e., when there are no more signals to estimate (the sources) than signals that are observed (the sensors), generic assumptions such as statistical independence of the sources can be used in order to demix successfully [2]. Many interesting applications, however, involve under-determined problems (more sources than sensors), where more specific assumptions must be made in order to demix. In problems involving at least two sensors, progress has been made by appealing to sparsity assumptions [3, 4].

However, the most extreme case, in which there is only one sensor and two or more sources, is a much harder and still-open problem for complex signals such as speech. In this setting, simple generic statistical assumptions do not suffice. One approach to the problem involves a return to the spirit of classical engineering methods such as matched filters, and estimating specific models for specific sources—e.g., specific speakers in the case of speech [5, 6]. While such an approach is reasonable, it departs significantly from the desideratum of “blindness.” In this paper we present an algorithm that is a blind separation algorithm—our algorithm separates speech mixtures from a single microphone without requiring models of specific speakers.

Our approach involves a “discriminative” approach to the problem of speech separation. That is, rather than building a complex model of speech, we instead focus directly on the task of separation and optimize parameters that determine separation performance. We work within a time-frequency representation (a spectrogram), and exploit the sparsity of speech signals in this representation. That is, although two speakers might speak simultaneously, there is relatively little overlap in the time-frequency plane if the speakers are different [5, 4]. We thus formulate speech separation as a problem in segmentation in the time-frequency plane. In principle, we could appeal to classical segmentation methods from vision (see, e.g. [7]) to solve this two-dimensional segmentation problem. Speech segments are, however, very different from visual segments, reflecting very different underlying physics. Thus we must design features for segmenting speech from first principles.

It also proves essential to combine knowledge-based feature design with learning methods. In particular, we exploit the fact that in speech we can generate “training examples” by artificially superposing two separately-recorded signals. Making use of our earlier work on learning methods for spectral clustering [1], we use the training data to optimize the parameters of a spectral clustering algorithm. This yields an adaptive, “discriminative” segmentation algorithm that is optimized to separate speech signals.

We highlight one other aspect of the problem here—the major computational challenge involved in applying spectral methods to speech separation. Indeed, four seconds of speech sampled at 5.5 KHz yields 22,000 samples and thus we need to manipulate affinity matrices of dimension at least $22,000 \times 22,000$. Thus a major part of our effort has involved the design of numerical approximation schemes that exploit the different time scales present in speech signals.

The paper is structured as follows. Section 2 provides a review of basic methodology. In Section 3 we describe our approach to feature design based on known cues for speech separation [8, 9]. Section 4 shows how parameterized affinity matrices based on these cues can be optimized in the spectral clustering setting. We describe our experimental results in Section 5. Most of the results presented in this paper have first been published in [10].

2. SPEECH SEPARATION AS SPECTROGRAM SEGMENTATION

In this section, we first review the relevant properties of speech signals in the time-frequency representation and describe how our




Fig. 1. Spectrogram of speech; (top) single speaker, (bottom) two simultaneous speakers. The gray intensity is proportional to the magnitude of the spectrogram.

training sets are constructed.

2.1. Spectrogram

The spectrogram is a two-dimensional (time and frequency) redundant representation of a one-dimensional signal [11]. Let $f[t]$, $t = 0, \dots, T-1$ be a signal in \mathbb{R}^T . The spectrogram is defined through windowed Fourier transforms and is commonly referred to as a short-time Fourier transform or as Gabor analysis [11]. The value $(Uf)_{mn}$ of the spectrogram at time window n and frequency m is defined as $(Uf)_{mn} = \frac{1}{\sqrt{M}} \sum_{t=0}^{T-1} f[t]w[t-na]e^{i2\pi mt/M}$, where w is a window of length T with small support of length c . We assume that the number of samples T is an integer multiple of a and c . There are then $N = T/a$ different windows of length c . The spectrogram is thus an $N \times M$ image which provides a redundant time-frequency representation of time signals¹ (see Figure 1).

Inversion Our speech separation framework is based on the segmentation of the spectrogram of a signal $f[t]$ in $S \geq 2$ disjoint subsets A_i , $i = 1, \dots, S$ of $[0, N-1] \times [0, M-1]$. This leads to S spectrograms U_i such that $(U_i)_{mn} = U_{mn}$ if $(m, n) \in A_i$ and zero otherwise—note that the phase is kept the same as the one of the original mixed signal. We now need to find S speech signals $f_i[t]$ such that each U_i is the spectrogram of f_i . In general there are no exact solutions (because the representation is redundant), and a classical technique is to find the minimum L_2 norm approximation, i.e., find f_i such that $\|U_i - Uf_i\|^2$ is minimal [11]. The solution of this minimization problem involves the pseudo-inverse of the linear operator U [11] and is equal to $f_i = (U^*U)^{-1}U^*U_i$. By our choice of window (Hanning), U^*U is proportional to the

¹In our simulations, the sampling frequency is $f_0 = 5.5kHz$ and we use a Hanning window of length $c = 216$ (i.e., $43.2ms$). The spacing between window is equal to $a = 54$ (i.e., $10.8ms$). We use a 512-point FFT ($M = 512$). For a speech sample of length 4 sec, we have $T = 22,000$ samples and then $N = 407$, which makes $\approx 2 \times 10^5$ spectrogram pixels.

identity matrix, so that the solution to this problem can simply be obtained by applying the adjoint operator U^* .

Normalization and subsampling There are several ways of normalizing a speech signal. In this paper, we chose to rescale all speech signals as follows: for each time window n , we compute the total energy $e_n = \sum_m |Uf_{mn}|^2$, and its 20-point moving average. The signals are normalized so that the 80% percentile of those values is equal to one.

In order to reduce the number of spectrogram samples to consider, for a given pre-normalized speech signal, we threshold coefficients whose magnitudes are less than a value that was chosen so that the distortion is inaudible.

2.2. Generating training samples

Our approach is based on a learning algorithm that optimizes a segmentation criterion. The training examples that we provide to this algorithm are obtained by mixing separately-normalized speech signals. That is, given two volume-normalized speech signals f_1, f_2 of the same duration, with spectrograms U_1 and U_2 , we build a training sample as $U^{train} = U_1 + U_2$, with a segmentation given by $z = \arg \min\{U_1, U_2\}$. In order to obtain better training partitions (and in particular to be more robust to the choice of normalization), we also search over all $\epsilon \in [0, 1]$ such that the least square reconstruction error of the waveform obtained from segmenting/reconstructing using $z = \arg \min\{\epsilon U_1, (1 - \epsilon)U_2\}$ is minimized. An example of such a partition is shown in Figure 2 (left).

3. FEATURES AND GROUPING CUES FOR SPEECH SEPARATION

In this section we describe our approach to the design of features for the spectral segmentation. We base our design on classical cues suggested from studies of perceptual grouping [12]. Our basic representation is a “feature map,” a two-dimensional representation that has the same layout as the spectrogram. Each of these cues is associated with a specific time scale, which we refer to as “small” (less than 5 frames), “medium” (10 to 20 frames), and “large” (across all frames). (These scales will be of particular relevance to the design of numerical approximation methods in Section 4.3). Any given feature is not sufficient for separating by itself; rather, it is the combination of several features that makes our approach successful.

3.1. Non-harmonic cues

The following non-harmonic cues have counterparts in visual scenes and for these cues we are able to borrow from feature design techniques used in image segmentation [7].

Continuity Two time-frequency points are likely to belong to the same segment if they are close in time or frequency; we thus use time and frequency directly as features. This cue acts at a small time scale.

Common fate cues Elements that exhibit the same time variation are likely to belong to the same source. This takes several particular forms. The first is simply *common offset* and *common onset*. We thus build an offset map and an onset map, with elements that are zero when no variation occurs, and are large when there is a sharp decrease or increase (with respect to time) for that particular time-frequency point. The onset and offset maps are

built using oriented energy filters as used in vision (with one vertical orientation). These are obtained by convolving the spectrogram with derivatives of Gaussian windows [7].

Another form of the common fate cue is *frequency co-modulation*, the situation in which frequency components of a single source tend to move in sync. To capture this cue we simply use oriented filter outputs for a set of orientation angles (8 in our simulations). Those features act mainly at a medium time scale.

3.2. Harmonic cues

This is the major cue for voiced speech [13, 9, 8], and it acts at all time scales (small, medium and large): voiced speech is locally periodic and the local period is usually referred to as the pitch.

Pitch estimation In order to use harmonic information, we need to estimate potentially several pitches. If S pitches are sought, the output that we need to obtain from the pitch extractor is, for each time frame n , the S pitches $\omega_{n1}, \dots, \omega_{nS}$, as well as the strength y_{nms} of the s -th pitch for each frequency m . Several pitch extractors for multiple pitches may be used [14, 15].

Building feature maps from pitch information We build a set of features from the pitch information. Given a time-frequency point (m, n) , let $s(m, n) = \arg \max_s \frac{y_{nms}}{(\sum_{m'} y_{nm's})^{1/2}}$ denote the highest energy pitch, and define the features $\omega_{ns(m,n)}, y_{nms(m,n)}, \frac{y_{nms(m,n)}}{\sum_{m'} y_{nm's(m,n)}}$ and $\frac{y_{nms(m,n)}}{(\sum_{m'} y_{nm's(m,n)})^{1/2}}$. We use a partial normalization with the square root to avoid including very low energy signals, while allowing a significant difference between the local amplitude of the speakers.

Those features all come with some form of energy level and all features involving pitch values ω should take this energy into account when the affinity matrix is built in Section 4. Indeed, the values of the harmonic features have no meaning when no energy in that pitch is present.

4. SPECTRAL CLUSTERING AND AFFINITY MATRICES

Given the features described in the previous section, we now show how to build affinity (i.e., similarity) matrices that can be used to define a spectral segmenter. In particular, our approach builds *parameterized* affinity matrices, and uses a learning algorithm to adjust these parameters.

4.1. Spectral clustering

Given P data points to partition into $S \geq 2$ disjoint groups, spectral clustering methods use an *affinity matrix* W , symmetric of size $P \times P$, that encodes topological knowledge about the problem. Once W is available, it is normalized and its first S (P -dimensional) eigenvectors are computed. Then, forming a $P \times S$ matrix with these eigenvectors as columns, we cluster the P rows of this matrix as points in \mathbb{R}^S using K -means (or a weighted version thereof). These clusters define the final partition [7, 1].

We prefer spectral clustering methods over other clustering algorithms such as K -means or mixtures of Gaussians estimated by the EM algorithm because we do not have any reason to expect the segments of interest in our problem to form convex shapes in the feature representation.

4.2. Parameterized affinity matrices

The success of spectral methods for clustering depends heavily on the construction of the affinity matrix W . In [1], we have shown how learning can play a role in optimizing over affinity matrices. Our algorithm assumes that fully partitioned datasets are available, and uses these datasets as training data for optimizing the parameters of affinity matrices. As we have discussed in Section 2.2, such training data are easily obtained in the speech separation setting. It remains for us to describe how we parameterize the affinity matrices.

From each of the features defined in Section 3, we define a basis affinity matrix $W_j = W_j(\beta_j)$, where β_j is a (vector) parameter. We restrict ourselves to affinity matrices whose elements are between zero and one, and with unit diagonal. We distinguish between harmonic and non-harmonic features. For non-harmonic features, we use a radial basis function to define affinities. Thus, if f_a is the value of the feature for data point a , we use a basis affinity matrix defined as $W_{ab} = \exp(-\|f_a - f_b\|^\beta)$, where $\beta > 1$.

For an harmonic feature, on the other hand, we need to take into account the strength of the feature: if f_a is the value of the feature for data point a , with strength y_a , we use $W_{ab} = \exp(-|g(y_a, y_b) + \beta_3|^\beta \|f_a - f_b\|^\beta)$, where $g(u, v) = (ue^{\beta_5 u} + ve^{\beta_5 v}) / (e^{\beta_5 u} + e^{\beta_5 v})$ ranges from the minimum of u and v for $\beta_5 = -\infty$ to their maximum for $\beta_5 = +\infty$.

Given m basis matrices, we use the following parameterization of W : $W = \sum_{k=1}^K \gamma_k W_1^{\alpha_{k1}} \times \dots \times W_m^{\alpha_{km}}$, where the products are taken pointwise. Intuitively, if we consider the values of affinity as soft boolean variables, taking the product of two affinity matrices is equivalent to considering the conjunction of two matrices, while taking the sum can be seen as their disjunction: our final affinity matrix can thus be seen as a disjunctive normal form. For our application to speech separation, we consider a sum of $K = 3$ matrices, one matrix for each time scale. This has the advantage of allowing different approximation schemes for each of the time scales, an issue we address in the following section.

4.3. Approximations of affinity matrices

The affinity matrices that we consider are huge, of size at least 50,000 by 50,000. Thus a significant part of our effort has involved finding computationally efficient approximations of affinity matrices.

Let us assume that the time-frequency plane is vectorized by stacking one time frame after the other. In this representation, the time scale of a basis affinity matrix W exerts an effect on the degree of “bandedness” of W . The matrix W is said *band-diagonal* with bandwidth B , if for all i, j , $|i - j| \geq B \Rightarrow W_{ij} = 0$. On a small time scale, W has a small bandwidth; for a medium time scale, the band is larger but still small compared to the total size of the matrix, while for large scale effects, the matrix W has no band structure. Note that the bandwidth B can be controlled by the coefficient of the radial basis function involving the time feature n .

For each of these three cases, we have designed a particular way of approximating the matrix, while ensuring that in each case the time and space requirements are *linear* in the number of time frames (see [10] for details).

5. EXPERIMENTS

We have trained our segmenter using data from four different speakers, with speech signals of duration 3 seconds. There were 28 pa-




Fig. 2. (Top) Optimal segmentation for the spectrogram in Figure 1 (bottom), where the two speakers are “black” and “grey;” this segmentation is obtained from the known separated signals. (Right) The blind segmentation obtained with our algorithm.

parameters to estimate using our spectral learning algorithm. For testing, we have used mixes from five speakers that were different from those in the training set.

In Figure 2, for two speakers from the testing set, we show on the left part an example of the segmentation that is obtained when the two speech signals are known in advance (obtained as described in Section 2.2), and on the right side, the segmentation that is output by our algorithm. Although some components of the “black” speaker are missing, the segmentation performance is good enough to obtain audible signals of reasonable quality. The speech samples for this example can be downloaded from `cmm.enscm.fr/~bach/speech/`. On this web site, there are additional examples of speech separation, with various speakers, in French and in English.

An important point is that our method does not require to know the speaker in advance in order to demix successfully; rather, it just requires that the two speakers have distinct and far enough pitches most of the time (another but less crucial condition is that one pitch is not too close to twice the other one).

As mentioned earlier, there was a major computational challenge in applying spectral methods to single microphone speech separation. Using the techniques described in Section 4.3, the separation algorithm has linear running time complexity and memory requirement and, coded in Matlab and C, our implementation is 40 times slower than real time.

6. CONCLUSIONS

We have presented an algorithm to perform blind source separation of speech signals from a single microphone. To do so, we have combined knowledge of physical and psychophysical properties of speech with learning methods. The former provide parameterized affinity matrices for spectral clustering, and the latter make use of our ability to generate segmented training data. The result is an optimized segmenter for spectrograms of speech mixtures. We have successfully demixed speech signals from two speakers using

this approach.

Our work thus far has been limited to the setting of ideal acoustics and equal-strength mixing of two speakers. There are several obvious extensions that warrant investigation. First, the mixing conditions should be weakened and should allow some form of delay or echo. Second, there are multiple applications where speech has to be separated from a non-stationary noise; we believe that our method can be extended to this situation. Third, our framework is based on segmentation of the spectrogram and, as such, distortions are inevitable since this is a “lossy” formulation [6, 4]. We are currently working on post-processing methods that remove some of those distortions. Finally, while running time and memory requirements of our algorithm are linear in the duration of the signal to be separated, the resource requirements remain a concern. We are currently working on further numerical techniques that we believe will bring our method significantly closer to real-time.

7. REFERENCES

- [1] F. R. Bach and M. I. Jordan, “Learning spectral clustering,” in *NIPS 16*, 2004.
- [2] A. Hyvärinen, J. Karhunen, and E. Oja, *Independent Component Analysis*, John Wiley & Sons, 2001.
- [3] M. Zibulevsky, P. Kisilev, Y. Y. Zeevi, and B. A. Pearlmutter, “Blind source separation via multinode sparse representation,” in *NIPS 14*, 2002.
- [4] O. Yilmaz and S. Rickard, “Blind separation of speech mixtures via time-frequency masking,” *IEEE Trans. Sig. Proc.*, vol. 52, no. 7, pp. 1830–1847, 2004.
- [5] S. T. Roweis, “One microphone source separation,” in *NIPS 13*, 2001.
- [6] G.-J. Jang and T.-W. Lee, “A probabilistic approach to single microphone speech separation,” in *Proc. Interspeech 2006*, 2006.