

Blind Source Separation by Sparse Decomposition

University of New Mexico Technical Report No. CS99-1

Michael Zibulevsky
michael@cs.unm.edu

Barak A. Pearlmutter
bap@cs.unm.edu

Computer Science Dept, FEC 313
University of New Mexico
Albuquerque, NM 87131 USA

July 19, 1999

Abstract

The blind source separation problem is to extract the underlying source signals from a set of their linear mixtures, where the mixing matrix is unknown. This situation is common, *eg* in acoustics, radio, and medical signal processing. We exploit the property of the sources to have a sparse representation in a corresponding (possibly overcomplete) signal dictionary. Such a dictionary may consist of wavelets, wavelet packets, *etc.*, or be obtained by learning from a given family of signals. Starting from the *maximum posteriori* framework, which is applicable to the case of more sources than mixtures, we derive a few other categories of objective functions, which provide faster and more robust computations, when there are an equal number of sources and mixtures. Our experiments with artificial signals and with musical sounds demonstrate significantly better separation than other known techniques.

1 Introduction

We consider the problem of blind separation of source signals from a set of their linear mixtures, including the case when the number of sources is larger than the number of mixtures. This work can be considered a natural generalization of the Bell-Sejnowski Infomax [2] and maximum posteriori [13, 15] methods of blind source separation. We assume that the source signals have a sparse representation in a corresponding (possibly overcomplete) signal dictionary. In this way *independence* and *sparsity* are not required from the signals themselves, but rather from their decomposition coefficients, which is more natural in many practical cases. On the other hand our approach may be considered an extension of basis pursuit [7] to the case of signal mixtures.

This paper is organized as follows. Section 2 gives the problem formulation and assumptions. In Section 3 we present the *maximum posteriori* approach, which is applicable to the case of more sources than mixtures. In Section 4 we derive another objective function, which provides more robust computations when there are an equal number of sources and mixtures. Section 5 presents sequential source extraction using quadratic programming with non-convex quadratic constraints. Finally, in Section 6 we derive a faster method for non-overcomplete dictionaries and demonstrate high-quality separation of synthetically mixed musical sounds.

2 Problem Formulation and Assumptions

Let $x(t)$ be an N -dimensional vector of sensor signals which is an instantaneous linear mixture of M unknown “independent” sources $s(t)$, possibly corrupted by additive noise $\xi(t)$:

$$x(t) = As(t) + \xi(t) \tag{1}$$

We will estimate the unknown mixing matrix of real numbers A (up to row permutation and scaling) and the source signal $s(t)$ (up to component permutation and scaling.)

We take advantage of the fact that many natural sources of signal have

sparse representation in the proper signal dictionary φ :

$$s_i(t) = \sum_{k=1}^K C_{ik} \varphi_k(t) \quad (2)$$

The functions $\varphi_k(t)$ are called *atoms* or *elements* of the dictionary. These elements do not have to be linearly independent, and instead may form an overcomplete dictionary. Important examples are wavelet-related dictionaries (wavelet packets, *etc.*) [7, 16]. Sparsity means that only a small number of the coefficients C_{ik} differ significantly from zero.

In the discrete time case $t = 1, 2, \dots, T$ we will use matrix notation. For example, X will be a $N \times T$ matrix, with the signal $x_i(t)$ in row i . S will be an $M \times T$ matrix with underlying source $s_j(t)$ in row j , and Φ a $K \times T$ matrix with atom $\varphi_k(t)$ in row k , so that

$$X = AS + \xi \quad (3)$$

$$S = C\Phi \quad (4)$$

We suppose that the coefficients C_{ik} are statistically independent random variables with a probability density function (pdf) of an exponential type

$$p_i(C_{ik}) \propto e^{-\beta_i h(C_{ik})} \quad (5)$$

This kind of distribution is widely used for modeling sparsity [15, 18]. A reasonable choice of $h(c)$ may be

$$h(c) = |c|^{1/\gamma} \quad \gamma \geq 1 \quad (6)$$

or its smooth approximations. In our computations we use a family of convex smooth approximations to the absolute value

$$h_1(c) = |c| - \log(1 + |c|) \quad (7)$$

$$h_\lambda(c) = \lambda h_1(c/\lambda) \quad (8)$$

where λ is a proximity parameter: $h_\lambda(c) \rightarrow |c|$ when $\lambda \rightarrow 0$.

We also suppose *a priori* that the matrix A is uniformly distributed over the range of interest, and that the noise $\xi(t)$ is a spatially and temporally uncorrelated Gaussian process¹ with zero mean and variance σ^2 .

¹The assumption of the noise whiteness is for simplicity of exposition and can be easily removed.

3 Maximum Posteriori Approach

We wish to maximize the posterior distribution $P(A, C | X)$

$$\max_{A, C} P(A, C | X) \propto \max_{A, C} P(X | A, C) P(A) P(C) \quad (9)$$

where $P(X | A, C)$ is the conditional probability of observing X given A and C . Taking into account (3), (4), and the white Gaussian noise, we get

$$P(X | A, C) \propto \prod_{i,t} \exp - \frac{(X_{it} - (AC\Phi)_{it})^2}{2\sigma^2} \quad (10)$$

By the statistical independence of the coefficients C_{jk} and (5), the prior pdf of C is

$$P(C) \propto \prod_{j,k} \exp(-\beta_j h(C_{jk})) \quad (11)$$

If the prior pdf $P(A)$ is uniform, it can be dropped² from (9). In this way we are left with the problem

$$\max_{A, C} P(X | A, C) P(C). \quad (12)$$

By substituting (10) and (11) into (12), taking the logarithm, and inverting the sign, we obtain the following optimization problem

$$\min_{A, C} \frac{1}{2\sigma^2} \|AC\Phi - X\|_F^2 + \sum_{j,k} \beta_j h(C_{jk}) \quad (13)$$

where $\|A\|_F = \sqrt{\sum_{i,j} A_{ij}^2}$ is the Frobenius matrix norm.

One can consider this objective as a generalization of [18, 17] by incorporating the matrix Φ , or as a generalization of [7] by including the matrix A . One problem with such a formulation is that it can lead to the degenerate solution $C = 0$ and $A = \infty$. We can overcome this difficulty in various ways. The first approach is to force each row A_i of the mixing matrix A to be bounded in norm,

$$\|A_i\| \leq 1 \quad i = 1, \dots, N. \quad (14)$$

²Otherwise, if $P(A)$ is some other known function, we should use (9) directly.

The second way is to restrict the norm of the rows C_j from below

$$\|C_j\| \geq 1 \quad j = 1, \dots, M. \quad (15)$$

A third way is to reestimate the parameters β_j based on the current values of C_j . For example, this can be done using sampling variance as follows: for a given function $h(\cdot)$ in the distribution (5), express the variance of C_{jk} as a function $f_h(\beta)$. An estimate of β can be obtained by applying the corresponding inverse function to the sampling variance,

$$\hat{\beta}_j = f_h^{-1}(K^{-1} \sum_k C_{jk}^2) \quad (16)$$

In particular, when $h(c) = |c|$, $\text{var}(c) = 2\beta^{-2}$ and

$$\hat{\beta}_j = \frac{2}{\sqrt{K^{-1} \sum_k C_{jk}^2}} \quad (17)$$

Substituting $h(\cdot)$ and $\hat{\beta}$ into (13), we obtain

$$\min_{A,C} \frac{1}{2\sigma^2} \|AC\Phi - X\|_F^2 + \sum_j \frac{2 \sum_k |C_{jk}|}{\sqrt{K^{-1} \sum_k C_{jk}^2}} \quad (18)$$

A remarkable property of this objective function is its invariance to rescaling of the rows of C when a corresponding inverse rescaling is applied to the columns of A .

Computational experiments with more sources than mixtures In our experiments we used the standard wavelet packet dictionary with the basic wavelet *symmlet-8*. When the signal length is 64 samples, this dictionary consists of 448 atoms *i.e.* it is overcomplete by a factor of seven. Examples of atoms and their images in the time-frequency phase plane [9, 16] are shown in Figure 1. We used the ATOMIZER [8] and WAVELAB [4] MATLAB packages for fast multiplication by Φ and Φ^T .

We created three sparse sources (Figure 2, top), each composed of two or three atoms. The first two sources have significant cross-correlation, equal to 0.34, which makes separation difficult for conventional methods. Two synthetic sensor signals (Figure 2, center) were obtained as a linear mixture of

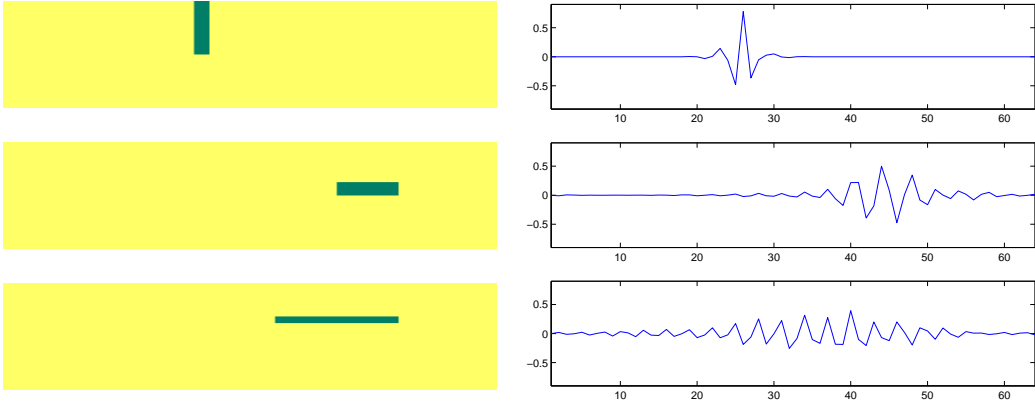


Figure 1: Examples of atoms: time-frequency phase plane (left) and time plot (right.)

the sources. In order to measure the accuracy of the separation, we normalized the original sources with $\|S_j\|_2 = 1$, and the estimated sources with $\|\tilde{S}_j\|_2 = 1$. The error was then computed as

$$\text{Error} = \frac{\|\tilde{S}_j - S_j\|_2}{\|S_j\|_2} \cdot 100\% \quad (19)$$

We tested two methods with this data. The first method used the objective function (13) and the constraints (15), while the second method used the objective function (18). As a tool for constrained optimization we used the PBM method [3]. Unconstrained optimization was produced by the method of conjugate gradients using the TOMLAB package [10]. The same tool was used for internal unconstrained optimization in PBM.

In all the cases we used $h_\lambda(\cdot)$ defined by (7) and (8), with the parameter $\lambda = 0.01$. Another parameter $\sigma^2 = 0.0001$. The resulting errors of the source estimates were 0.09% and 0.02% by the first and the second method respectively. The estimated sources are shown in the bottom three traces of Figure 2. They are visually indistinguishable from the original sources, shown in top three traces of Figure 2.

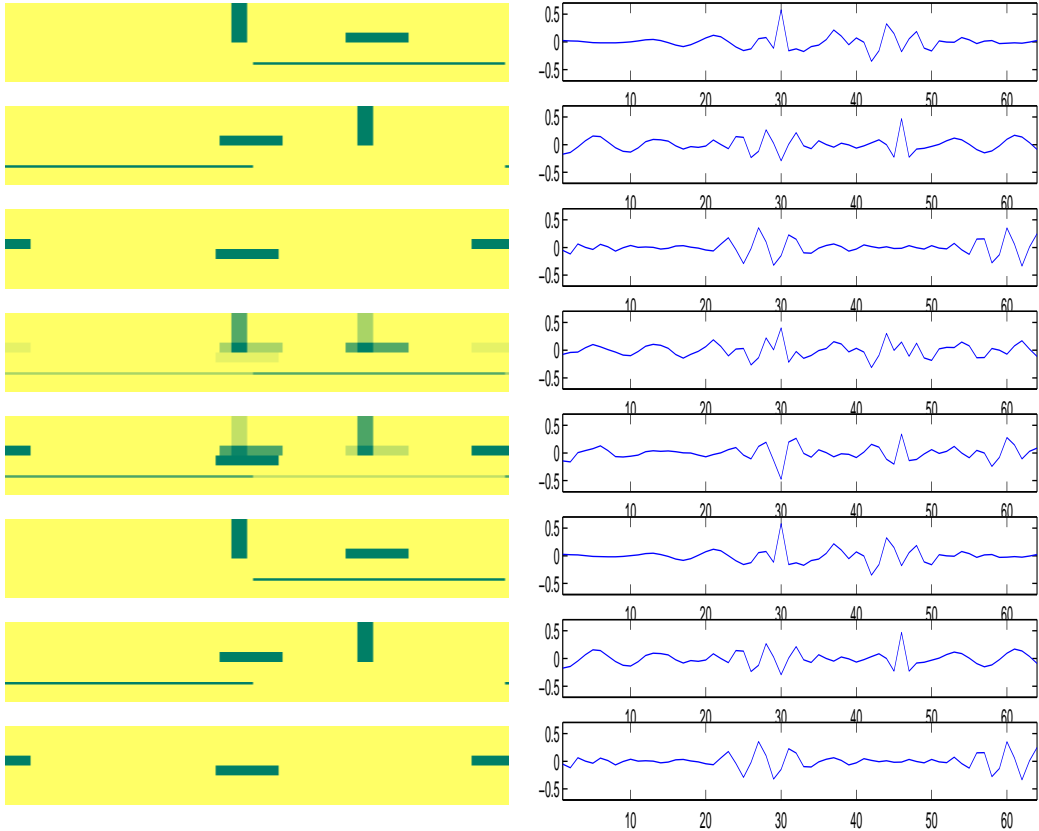


Figure 2: Sources (top three), mixtures (center two), reconstructed sources (bottom three), in both time-frequency phase plane (left) and time domain (right).

4 Equal number of sources and sensors: more robust formulations

The main difficulty in a maximization problem like (13) is the bilinear term $AC\Phi$, which destroys the convexity of the objective function and makes convergence unstable when optimization starts far from the solution. In this section we consider more robust formulations for the case when the number of sensors is equal to the number of sources, $N = M$, and the mixing matrix is invertible $W = A^{-1}$.

In the case when the noise is small and the matrix A is far from singular,

WX gives a reasonable estimate of the source signals S . Taking into account (4), we obtain a least square term $\|C\Phi - WX\|_F^2$, so the separation objective may be written

$$\min_{W,C} \frac{1}{2} \|C\Phi - WX\|_F^2 + \mu \sum_{j,k} \beta_j h(C_{jk}) \quad (20)$$

We also need to add a constraint, which enforces the non-singularity of W . For example, we can restrict from below its minimal singular value $r_{min}(W)$:

$$r_{min}(W) \geq 1 \quad (21)$$

It can be shown, that in the noiseless case, $\sigma \approx 0$, the problem (20)–(21) is equivalent to the maximum posteriori formulation (13) with the constraint $\|A\|_2 \leq 1$. Another possibility for ensuring the non-singularity of W is to subtract $K \log |\det W|$ from the objective

$$\min_{W,C} -K \log |\det W| + \frac{1}{2} \|C\Phi - WX\|_F^2 + \mu \sum_{j,k} \beta_j h(C_{jk}) \quad (22)$$

When the noise is zero and Φ is the identity matrix, we can substitute $C = WX$ and obtain the Bell-Sejnowski Infomax objective [2]

$$\min_W -K \log |\det W| + \sum_{j,k} \beta_j h((WX)_{jk}) \quad (23)$$

Computational experiments with equal number of sources and sensors We created two sparse sources (Figure 3, top) with strong cross-correlation of 0.52. Separation, produced by minimization of the objective function (22), gave an error of 0.23%. For comparison we tested the JADE [6, 5], FastICA [12, 11] and Bell-Sejnowski Infomax [2, 1] algorithms on the same signals. The Resulting relative errors (Figure 4) confirm the significant superiority of the sparse decomposition approach.

5 Sequential Extraction of Sources via Quadratic Programming

Let us ask what is the most “sparse” signal one can obtain by a linear combination of the sensor signals $s = w^T X$. By sparsity, as before, we mean the

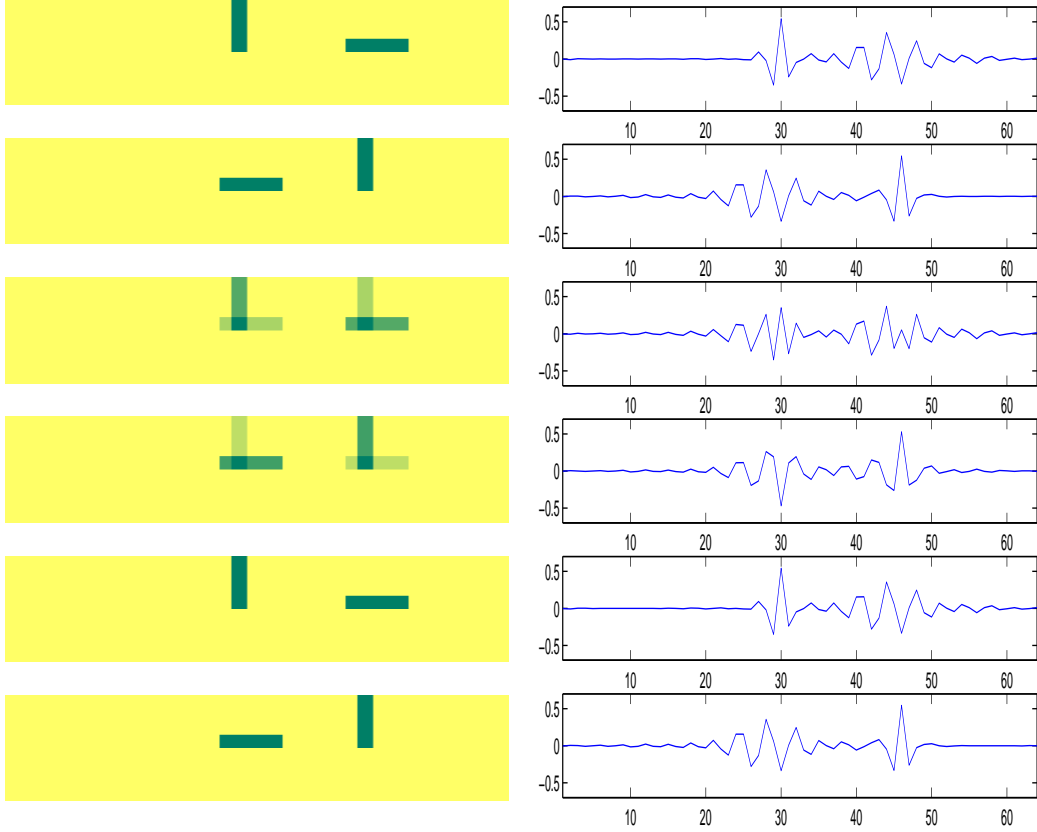


Figure 3: Sources (top two), mixtures (center two), reconstructed sources (bottom two), in both time-frequency phase plane (left) and time domain (right).

ability of the signal to be approximated by a linear combination of a small number of dictionary elements ϕ_k

$$s \approx c^T \Phi, \quad c \text{ sparse}$$

This will lead us to the following objective:

$$\min_{w,c} \frac{1}{2} \|c^T \Phi - w^T X\|_2^2 + \mu \sum_k h(c_k), \quad (24)$$

where the term $\sum_k h(c_k)$ may be considered as a penalty on non-sparsity. In order to avoid the trivial solution of $w = 0$ and $c = 0$, we need to add a

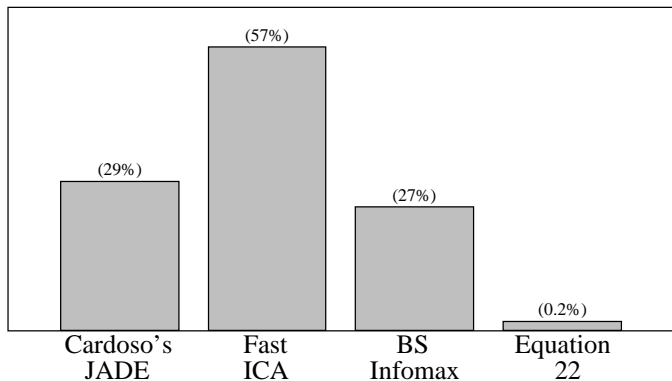


Figure 4: Percent relative error of separation of the artificial sparse sources recovered by (1) JADE, (2) Fast ICA, (3) Bell-Sejnowski Infomax, (4) Equation 22.

constraint that separates w from zero. It could be, for example,

$$\|w\|_2^2 \geq 1, \quad (25)$$

A similar constraint can be used as a tool to extract all the sources sequentially: the new separation vector w^j should have a component of unit norm in the subspace orthogonal to the previously extracted vectors w^1, \dots, w^{j-1}

$$\|(I - P^{j-1})w^j\|_2^2 \geq 1, \quad (26)$$

where P^{j-1} is an orthogonal projector onto $\text{Span}\{w^1, \dots, w^{j-1}\}$.

When $h(c_k) = |c_k|$, we can use the standard substitution

$$\begin{aligned} c &= c^+ - c^-, \quad c^+ \geq 0, \quad c^- \geq 0 \\ \hat{c} &= \begin{pmatrix} c^+ \\ c^- \end{pmatrix} \quad \text{and} \quad \hat{\Phi} = \begin{pmatrix} \Phi \\ -\Phi \end{pmatrix}, \end{aligned}$$

that transforms (24) and (26) into the following quadratic program

$$\begin{aligned} \min_{w, \hat{c}} \quad & \frac{1}{2} \|\hat{c}^T \hat{\Phi} - w^T X\|_2^2 + \mu e^T \hat{c} \\ \text{subject to:} \quad & \|w\|_2^2 \geq 1, \\ & \hat{c} \geq 0, \end{aligned}$$

where e is a vector of ones.

6 Fast Solution in Non-overcomplete Dictionaries

In important applications, the sensor signals may have hundreds of channels and hundreds of thousands of samples. This may make separation computationally difficult. Here we present an approach which compromises between statistical and computational efficiency. In our experience this approach provides high quality of separation in reasonable time. Suppose the dictionary is “complete”, *i.e.* the it forms a basis in the space of discrete signals. This means that the matrix Φ is square and non-singular. As examples of such a dictionary one can think of the Fourier basis, Gabor basis, various wavelet-related bases, *etc.* We can also obtain an “optimal” dictionary by learning from given family of signals [15, 14, 18, 17].

Let us denote the dual basis

$$\Psi = \Phi^{-1} \quad (27)$$

and suppose that coefficients of decomposition of the sources

$$C = S\Psi \quad (28)$$

are sparse and statistically independent. This assumption is reasonable for properly chosen dictionaries, although of course we would lose the advantages of overcompleteness [15].

Let Y be the decomposition of the sensor signals

$$Y = X\Psi \quad (29)$$

Multiplying both sides of (3) by Ψ from the right and taking into account (28) and (29), we obtain

$$Y = AC + \zeta, \quad (30)$$

where ζ is decomposition of the noise

$$\zeta = \xi\Psi. \quad (31)$$

In this paper we consider an “easy” situation, when ζ is a white noise, that requires orthogonality of Ψ . We can see that all the objective functions from the sections 3–5 remain valid if we remove from them Φ (substituting

instead the identity matrix) and replace the sensor signals X by their decomposition Y . For example, maximum posteriori objectives (13) and (18) are transformed into

$$\min_{A,C} \frac{1}{2\sigma^2} \|AC - Y\|_F^2 + \sum_{j,k} \beta_j h(C_{jk}) \quad (32)$$

and

$$\min_{A,C} \frac{1}{2\sigma^2} \|AC - Y\|_F^2 + \sum_j \frac{2 \sum_k |C_{jk}|}{\sqrt{K^{-1} \sum_k C_{jk}^2}} \quad (33)$$

The objective (22) becomes

$$\min_{W,C} -K \log |\det W| + \frac{1}{2} \|C - WY\|_F^2 + \mu \sum_{j,k} \beta_j h(C_{jk}) \quad (34)$$

In this case when the noise is zero, we can substitute $C = WY$ and obtain the Bell-Sejnowski Infomax objective [2]

$$\min_W -K \log |\det W| + \sum_{j,k} \beta_j h((WY)_{jk}) \quad (35)$$

Also other known methods (for example, [13, 15]), which normally assume sparsity of source signals, may be directly applied to the decomposition Y of the sensor signals. This may be more efficient than the traditional approach, and the reason is obvious: typically, a properly chosen decomposition gives significantly higher sparsity than the raw signals had originally. Also, statistical independence of the coefficients is a more reasonable assumption than statistical independence of the raw signal samples.

Computational experiments with musical sound sources In our experiments we artificially mixed seven 5-second fragments of musical sound recordings taken from commercial digital audio CDs. Each of them included 40k samples after down-sampling by a factor of 5. (Figure 5).

The easiest way to perform sparse decomposition of such sources is to compute a *spectrogram*, the coefficients of a time-windowed discrete Fourier transform. (We used the function SPECGRAM from the MATLAB signal processing toolbox with a time window of 1024 samples.) The sparsity of the spectrogram coefficients (the histogram in Figure 6, right) is much higher than the sparsity of the original signal (Figure 6, left)

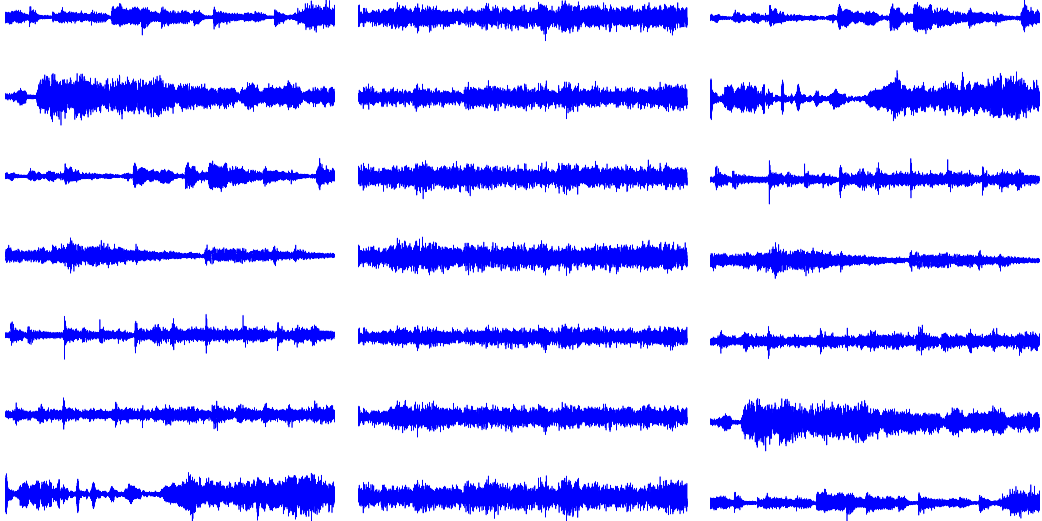


Figure 5: Separation of musical recordings taken from commercial digital audio CDs (five second fragments.) Original sources (left); mixtures (center); separated sources (right).

In this case Y (29) is a real matrix, with separate entries for the real and imaginary components of each spectrogram coefficient of the sensor signals X . We used the objective function (35) with $\beta_j = 1$ and $h_\lambda(\cdot)$ defined by (7),(8) with the parameter $\lambda = 10^{-4}$. Unconstrained minimization was performed by a BFGS Quasi-Newton algorithm (MATLAB function FMINU.)

This algorithm separated the sources with a relative error of 0.67% for the least well separated source (error computed according to (19).) We also applied the Bell-Sejnowski Infomax algorithm [2] implemented in [1] to the spectrogram coefficients Y of the sensor signals. Separation errors were slightly larger: 0.9%.

For comparison we tested JADE [6, 5], FastIca [12, 11] and Bell-Sejnowski Infomax algorithms on the same signals. Resulting relative errors (Figure 7) confirm the significant (by a factor of more than 10) superiority of the sparse decomposition approach.

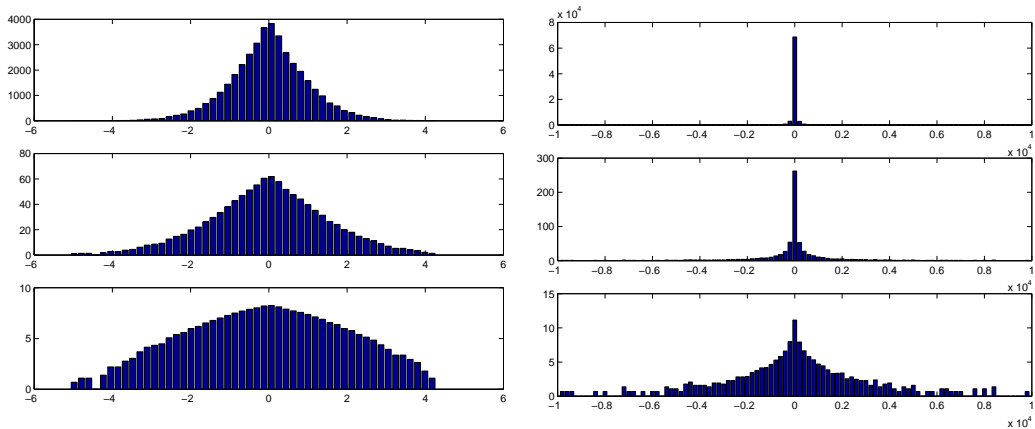


Figure 6: Histogram of sound source values (left) and spectrogram coefficients (right), shown with linear y-scale (top), square root y-scale (center) and logarithmic y-scale (bottom).

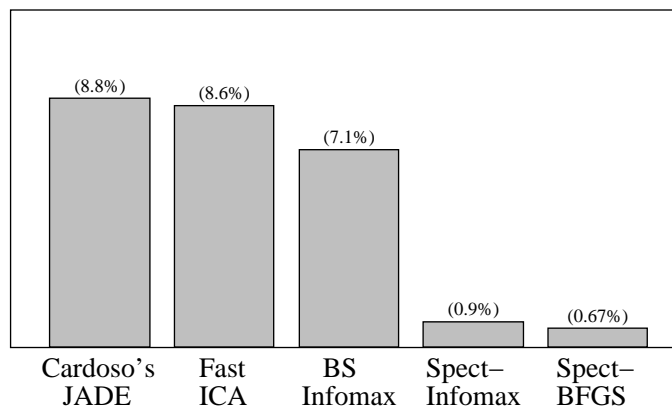


Figure 7: Percent relative error of separation of seven musical sources recovered by (1) JADE, (2) Fast ICA, (3) Bell-Sejnowski Infomax, (4) Infomax, applied to the spectrogram coefficients, (5) BFGS minimization of the objective (35) with the spectrogram coefficients.

7 Future research

We should mention an alternative to the maximum posteriori approach (12). Considering the mixing matrix A as a parameter, we can estimate it by maximizing the probability of the observed signal X

$$\max_A \left[P(X | A) = \int P(X | A, C)P(C)dC \right]$$

The integral over all possible coefficients C may be approximated, for example, by Monte-Carlo sampling or by a matching Gaussian, in the spirit of [15, 14]. It would be interesting to compare this possibility to the other methods presented in this paper.

Another important direction give us the problem of blind separation-deconvolution of convolutive mixtures of signals (see for example [2].) In this case the matrices A and W will have linear filters as an elements, and multiplication by the element will mean convolution. Even in this matrix-of-filters context most of the formulae in this paper remain valid.

8 Conclusions

In this paper we showed that the use of sparse decomposition in a corresponding signal dictionary provides high-quality blind source separation. The maximum posteriori framework gives the most general approach, including the situation of more sources than sensors. Faster and computationally robust solutions are possible in the case of an equal number of sources and sensors. We can also extract the sources sequentially using quadratic programming with non-convex quadratic constraints. Finally, the fastest solution may be obtained using non-overcomplete dictionaries. Our experiments with artificial signals and digitally mixed musical sounds demonstrate a high quality of source separation, compared to other known techniques.

References

- [1] ICA/EEG toolbox. Computational Neurobiology Laboratory, the Salk Institute. http://www.cnl.salk.edu/~tewon/ica_cnl.html.

- [2] Anthony J. Bell and Terrence J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, 1995.
- [3] A. Ben-Tal and M. Zibulevsky. Penalty/barrier multiplier methods for convex programming problems. *SIAM Journal on Optimization*, 7(2):347–366, 1997.
- [4] J. Buckheit, S.S. Chen, D.L. Donoho, I. Johnstone, and J. Scargle. About wavelab. Technical report, Department of Statistics, Stanford University, 1995. <http://www-stat.stanford.edu/~donoho/Reports/>.
- [5] Jean-François Cardoso. JADE for real-valued data. <http://sig.enst.fr:80/~cardoso/guidesepsou.html>.
- [6] Jean-François Cardoso. High-order contrasts for independent component analysis. *Neural Computation*, 11(1):157–192, 1999.
- [7] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. 1996. <http://www-stat.stanford.edu/~donoho/Reports/>.
- [8] S. S. Chen, D. L. Donoho, M. A. Saunders, I. Johnstone, and J. Scargle. About atomizer. Technical report, Department of Statistics, Stanford University, 1995. <http://www-stat.stanford.edu/~donoho/Reports/>.
- [9] R.R. Coifman and M.V. Wickerhauser. Entropy-based algorithms for best-basis selection. *IEEE Transactions on Information Theory*, 38:713–718, 1992.
- [10] K. Holmstrom and M. Bjorkman. The TOMLAB NLPLIB. *Advanced Modeling and Optimization*, 1:70–86, 1999. <http://www.ima.mdh.se/tom/>.
- [11] A. Hyvärinen. The Fast-ICA MATLAB package. <http://www.cis.hut.fi/~aapo/>.
- [12] A. Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, 1999.

- [13] T.W. Lee, M.S. Lewicki, M. Girolami, and T.J. Sejnowski. Blind source separation of more sources than mixtures using overcomplete representations. *IEEE Sig. Proc. Lett.*, 1998. to appear.
- [14] M.S. Lewicki and B.A. Olshausen. A probabilistic framework for the adaptation and comparison of image codes. *Journal of the Optical Society of America*, 1999. in press.
- [15] M.S. Lewicki and T.J. Sejnowski. Learning overcomplete representations. *Neural Computation*, 1998. to appear.
- [16] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1998.
- [17] B.A. Olshausen and D.J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, 1996.
- [18] B.A. Olshausen and D.J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37:3311–3325, 1997.