

Image Decomposition and Separation Using Sparse Representations: An Overview

This overview paper points out that signal and image processing, as well as many other important areas of engineering, can benefit from the techniques it discusses.

By M. JALAL FADILI, JEAN-LUC STARCK, JÉRÔME BOBIN, AND YASSIR MOUDDEN

ABSTRACT | This paper gives essential insights into the use of sparsity and morphological diversity in image decomposition and source separation by reviewing our recent work in this field. The idea to morphologically decompose a signal into its building blocks is an important problem in signal processing and has far-reaching applications in science and technology. Starck *et al.* [1], [2] proposed a novel decomposition method—morphological component analysis (MCA)—based on sparse representation of signals. MCA assumes that each (monochannel) signal is the linear mixture of several layers, the so-called morphological components, that are morphologically distinct, e.g., sines and bumps. The success of this method relies on two tenets: sparsity and morphological diversity. That is, each morphological component is sparsely represented in a specific transform domain, and the latter is highly inefficient in representing the other content in the mixture. Once such transforms are identified, MCA

is an iterative thresholding algorithm that is capable of decoupling the signal content. Sparsity and morphological diversity have also been used as a novel and effective source of diversity for blind source separation (BSS), hence extending the MCA to multichannel data. Building on these ingredients, we will provide an overview the generalized MCA introduced by the authors in [3] and [4] as a fast and efficient BSS method. We will illustrate the application of these algorithms on several real examples. We conclude our tour by briefly describing our software toolboxes made available for download on the Internet for sparse signal and image decomposition and separation.

KEYWORDS | Blind source separation; image decomposition; morphological component analysis; sparse representations

I. INTRODUCTION

Although mathematics has its million-dollar problems, in the form of Clay Math Prizes, there are several billion-dollar problems in signal and image processing. Famous ones include the cocktail party problem (separate a speaker voice from a mixture of other recorded voices and background sounds at a cocktail party). These signal-processing problems seem to be intractable according to orthodox arguments based on rigorous mathematics, and yet they keep cropping up in problem after problem.

One such fundamental problem involves decomposing a signal or image into superposed contributions from different sources; think of symphonic music, which may involve superpositions of acoustic signals generated by many different instruments—and imagine the problem of separating these contributions. More abstractly, we can see many forms of media content that are superpositions of contributions from different “content types,” and we can imagine wanting to separate out the contributions from

Manuscript received March 10, 2009; revised May 29, 2009; accepted June 1, 2009. Date of publication September 29, 2009; date of current version May 19, 2010. This work was supported by NatImages ANR under Grant ANR-08-EMER-009.

M. J. Fadili is with GREYC CNRS, ENSICAEN, Image Processing Group, University of Caen, 14050 Caen Cedex, France (e-mail: Jalal.Fadili@greyc.ensicaen.fr).

J.-L. Starck is with AIM, CEA/DSM, CNRS, University Paris Diderot, F-91191 Gif-sur-Yvette Cedex, France (e-mail: jstarck@cea.fr).

J. Bobin is with the Department of Applied and Computational Mathematics, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: bobin@acm.caltech.edu).

Y. Moudden is with DSM/IRFU/SEDI, CEA/Saclay, F-91191 Gif-sur-Yvette, France (e-mail: yassir.moudden@cea.fr).

Digital Object Identifier: 10.1109/JPROC.2009.2024776

each. We can easily see a fundamental problem; for example, an N -pixel image created by superposing K different types offers us N data (the pixel values), but there may be as many as $N \cdot K$ unknowns (the contribution of each content type to each pixel). Traditional mathematical reasoning—in fact, the fundamental theorem of linear algebra—tells us not to attempt this: there are more unknowns than equations. On the other hand, if we have prior information about the underlying object, there are some rigorous results showing that such separation might be possible—using a sparsity prior.

The idea to morphologically decompose a signal into its building blocks is an important problem in signal and image processing. Successful methods for signal or image separation can be applied in a broad range of areas in science and technology including biomedical engineering, medical imaging, speech processing, astronomical imaging, remote sensing, communication systems, etc. An interesting and complicated image content separation problem is the one targeting decomposition of an image to texture and piece-wise-smooth (cartoon) parts. A functional-space characterization of oscillating textures was proposed in [5] and was used for variational cartoon + texture image decomposition [6]. Since then, we have witnessed a flurry of research activity in this application field.

In [1] and [2], the authors proposed a novel decomposition method—morphological component analysis (MCA)—based on sparse representation of signals. MCA assumes that each signal is the linear mixture of several layers, the so-called morphological components, that are morphologically distinct, e.g., sines and bumps. The success of this method relies on the assumption that for every component behavior to be separated, there exists a dictionary of atoms that enables its construction using a sparse representation. It is then assumed that each morphological component is sparsely represented in a specific transform domain. And when all transforms (each one attached to a morphological component) are amalgamated in one dictionary, each one must lead to sparse representation over the part of the signal it is serving while being highly inefficient in representing the other content in the mixture. If such dictionaries are identified, the use of a pursuit algorithm searching for the sparsest representation leads to the desired separation. MCA is capable of creating atomic sparse representations containing as a by-product a decoupling of the signal content.

Over the last few years, the development of multichannel sensors motivated interest in methods for the coherent processing of multivariate data. Consider a situation where there is a collection of signals emitted by some physical sources. These could be, for example, different brain areas emitting electric signals; people speaking in the same room, thus emitting speech signals; or radiation sources emitting their electromagnetic waves.

Assume further that there are several sensors or receivers. These sensors are in different positions, so that each records a mixture of the original source signals with different weights. The so-called blind source separation (BSS) problem is to find the original sources or signals from their observed mixtures, without prior knowledge of the mixing weights, and by knowing very little about the original sources. Some specific issues of BSS have already been addressed, as testified by the wide literature in this field. In this context, as clearly emphasized by previous work, it is fundamental that the sources to be retrieved present some quantitatively measurable diversity or contrast (e.g., decorrelation, independence, morphological diversity, etc.). The seminal work of [7] and [8] paved the way for the use of sparsity in BSS. Recently, sparsity and morphological diversity have emerged as a novel and effective source of diversity for BSS for both underdetermined and overdetermined BSS; see the comprehensive review in [4]. Building on the sparsity and morphological diversity ingredients, the authors proposed the generalized MCA (GMCA) as a fast and efficient multichannel sparse data-decomposition and BSS method [3], [4], [9].

A. Organization of the Paper

Our intent in this paper is to provide an overview of the recent work in monochannel image decomposition and multichannel source separation based on the concepts of sparsity and morphological diversity. The first part of this paper is devoted to monochannel sparse image decomposition, and the second part to blind sparse source separation. In this review, our goal is to highlight the essential concepts and issues, and to describe the main algorithms. Several applications to real data are given in each part to illustrate the capabilities of the proposed algorithms. We conclude our tour by providing pointers to our software toolboxes that implement our algorithms and reproduce the experiments on sparse signal and image decomposition and source separation.

B. Notations

The ℓ_p -norm of a (column or row) vector x is $\|x\|_p := (\sum_i |x[i]|^p)^{1/p}$ with the usual adaptation when $p = \infty$ and $\|x\|_0 := \lim_{p \rightarrow 0} \|x\|_p$ is the ℓ_0 pseudonorm, i.e., the number of nonzero components. Bold symbols represent matrices and \mathbf{X}^T is the transpose of \mathbf{X} . The Frobenius norm of \mathbf{X} is $\|\mathbf{X}\|_F = \text{Trace}(\mathbf{X}^T \mathbf{X})^{1/2}$. The k th entry of y_i (respectively, y^j) is $y_i[k]$ (respectively, $y^j[k]$), where y_i is the i th row and y^j is the j th column of \mathbf{Y} . An atom is an elementary signal-representing template. Examples might include sinusoids, monomials, wavelets, and Gaussians. A dictionary $\Phi = [\phi^1, \dots, \phi^L]$ defines a $N \times L$ matrix whose columns are unit ℓ_2 -norm atoms ϕ^i . When the dictionary has more columns than rows, it is called overcomplete or redundant. We are mainly interested here in overcomplete dictionaries.

II. MONOCHANNEL SPARSE IMAGE DECOMPOSITION

A. Morphological Component Analysis

Suppose that the N -sample signal or image x is the linear superposition of K morphological components, possibly contaminated with noise

$$y = \sum_{k=1}^K x_k + \varepsilon, \quad \sigma_\varepsilon^2 = \text{Var}[\varepsilon] < +\infty.$$

The MCA framework aims at solving the inverse problem that consists in recovering the components $(x_k)_{k=1,\dots,K}$ from their observed linear mixture, as illustrated in the top of Fig. 1. MCA assumes that each component x_k can be sparsely represented in an associated basis Φ_k , i.e.,

$$x_k = \Phi_k \alpha_k, \quad k = 1, \dots, K$$

where α_k is a sparse coefficient vector (sparse means that only a few coefficients are large and most are negligible).

Thus, a dictionary can be built by amalgamating several transforms (Φ_1, \dots, Φ_K) such that, for each k , the representation of x_k in Φ_k is sparse and not, or at least not as sparse, in other Φ_l , $l \neq k$. In other words, the subdictionaries (Φ_1, \dots, Φ_K) must be mutually incoherent. Thus, the dictionary Φ_k plays a role of a discriminant between content types, preferring the component x_k over the other parts. This is a key observation for the success of the separation algorithm. Owing to recent advances in computational harmonic analysis, many novel representations, including the wavelet transform, curvelet, contourlet, steerable, or complex wavelet pyramids, were shown to be very effective in sparsely representing certain kinds of signals and images. Thus, for decomposition purposes, the dictionary will be built by taking the union of one or several (sufficiently incoherent) transforms, generally each corresponding to an orthogonal basis or a tight frame.

However, the augmented dictionary $\Phi = [\Phi_1 \dots \Phi_K]$ will provide an overcomplete representation of x . Because there are more unknowns than equations, the system $x = \Phi \alpha$ is underdetermined. Sparsity can be used to find a unique solution, in some idealized cases; there is an extensive literature on the subject, and the

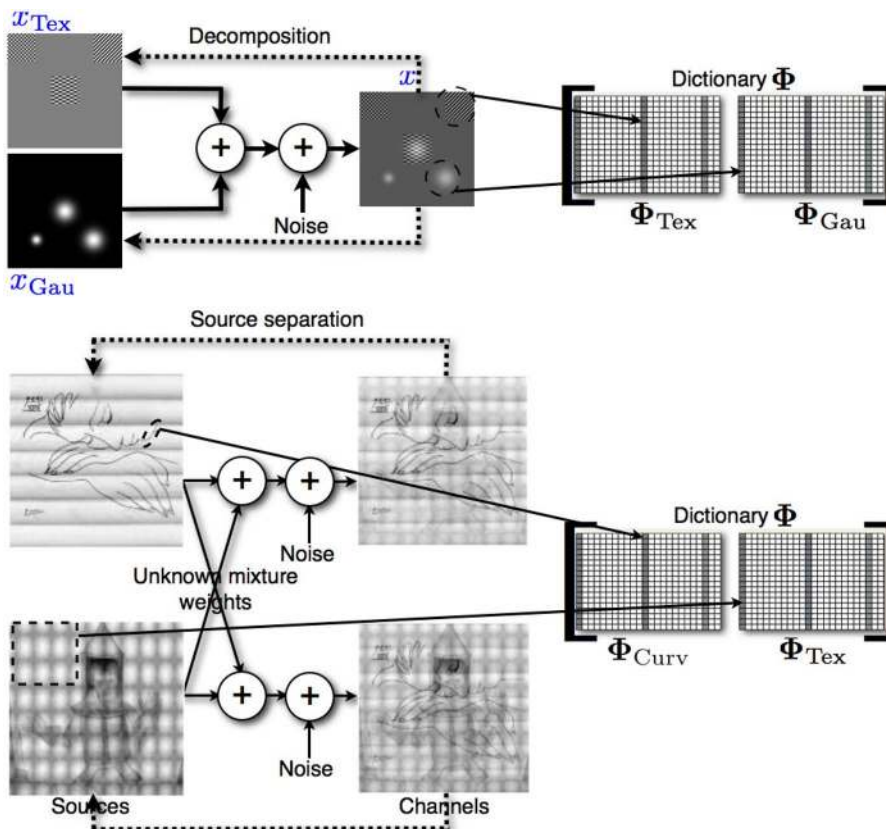


Fig. 1. Illustration of the (top) image decomposition and (bottom) BSS problems using sparsity and morphological diversity. For the bottom part, each source is itself a mixture of morphological components [see (4)] to be isolated.

interested reader may refer to the comprehensive review paper [10].

In [1] and [2], it was proposed to solve this under-determined system of equations and recover the morphological components $(x_k)_{k=1,\dots,K}$ by solving the following constrained optimization problem:

$$\min_{\alpha_1, \dots, \alpha_K} \sum_{k=1}^K \|\alpha_k\|_p^p \text{ such that } \left\| y - \sum_{k=1}^K \Phi_k \alpha_k \right\|_2 \leq \tau \quad (1)$$

where $\|\alpha\|_p^p$ is the penalty quantifying sparsity (the most interesting regime is for $0 \leq p \leq 1$) and τ is typically chosen as a constant time $\sqrt{N}\sigma_\epsilon$. The constraint in this optimization problem accounts for the presence of noise and model imperfection. If there is no noise and the linear superposition model is exact ($\tau = 0$), an equality constraint is substituted for the inequality constraint. This formulation is flexible enough to incorporate external forces that direct the morphological components to better suit their expected content; these forces will fine-tune the separation process to achieve its task. As an example for such successful external force, [1] proposed to add a total variation penalty [11] to the cartoon part in order to direct this component to fit the piecewise-smooth model.

B. MCA Algorithm

Equation (1) is not easy to solve in general, especially for $p < 1$ (for $p = 0$, it is even NP-hard). Nonetheless, if all components $x_l = \Phi_l \alpha_l$ but the k th are fixed, then it can be proved that the solution α_k is given by hard thresholding (for $p = 0$) or soft thresholding (for $p = 1$) the marginal residuals $r_k = y - \sum_{l \neq k} \Phi_l \alpha_l$. These marginal residuals r_k

are relieved from the other components and are likely to contain mainly the salient information of x_k . This intuition dictates a coordinate relaxation algorithm that cycles through the components at each iteration and applies a thresholding to the marginal residuals. This is what justifies the steps of the MCA algorithm summarized in Algorithm 1, where $\text{TH}_\lambda(\alpha)$ denotes component-wise thresholding with threshold λ : hard thresholding $\text{HT}_\lambda(u) = u$ if $|u| > \lambda$ and zero otherwise, or soft-thresholding $\text{ST}_\lambda(u) = u \max(1 - \lambda/|u|, 0)$.

Beside coordinate relaxation, another important ingredient of MCA is iterative thresholding with *varying threshold*. Thus, MCA can be viewed as a stagewise hybridization of matching pursuit (MP) [12] with block-coordinate relaxation [13] to (approximately) solve (1). The adjective *stagewise* is because MCA exploits the fact that the dictionary is structured (union of transforms), and the atoms enter the solution by groups rather than individually unlike MP. As such, MCA is a *salient-to-fine* process where, at each iteration, the most salient content of each morphological component is iteratively computed. These estimates are then progressively refined as the threshold λ decreases towards λ_{\min} .

In the noiseless case, a careful analysis of the recovery properties (uniqueness and support recovery) of the MCA algorithm and its convergence behavior when all Φ_k are orthobases can be found in [9] and [14].

C. Dictionary Choice

From a practical point of view, given a signal x , we will need to compute its forward (or analysis) transform by multiplying it by Φ^T . We also need to reconstruct any signal from its coefficients α . In fact, the matrix Φ and its adjoint Φ^T corresponding to each transform are never explicitly constructed in memory.

Algorithm 1: MCA decomposition algorithm.

Task: Signal/image decomposition.

Parameters: The signal/image x , the dictionary $\Phi = [\Phi_1 \dots \Phi_K]$, number of iterations N_{iter} , stopping threshold λ_{\min} , threshold update schedule.

Initialization:

- Initial solution $x_k^{(0)} = 0, \forall k$.
- Initial residual $r^{(0)} = y$.
- Initial threshold: let $k^* = \max_k \|\Phi_k^T y\|_\infty$, set $\lambda^{(0)} = \max_{k \neq k^*} \|\Phi_k^T y\|_\infty$.

Main iteration:

for $t = 1$ **to** N_{iter} **do**

for $k = 1$ **to** K **do**

 Compute marginal residuals $r_k^{(t)} = r^{(t-1)} + x_k$.

 Update k -th component coefficients by thresholding $\alpha_k^{(t)} = \text{TH}_{\lambda^{(t-1)}}(\Phi_k^T r_k^{(t)})$.

 Update k -th component $x_k^{(t)} = \Phi_k \alpha_k^{(t)}$.

 Update the residuals $r^{(t+1)} = y - \sum_{k=1}^K x_k^{(t)}$.

 Update the threshold $\lambda^{(t)}$ according to the given schedule.

if $\lambda^{(t)} \leq \lambda_{\min}$ **then stop**.

Output: Morphological components $(x_k^{(N_{\text{iter}})})_{k=1,\dots,K}$.

Rather, they are implemented as fast implicit analysis and synthesis operators taking a signal vector x and returning $\Phi^T x = \mathbf{T}x$ (analysis side), or taking a coefficient vector α and returning $\Phi\alpha$ (synthesis side). In the case of a simple orthogonal basis, the inverse of the analysis transform is trivially $\mathbf{T}^{-1} = \Phi$; whereas, assuming that Φ is a tight frame implies that the frame operator satisfies $\Phi\Phi^T = c\mathbf{I}$, where $c > 0$ is the tight frame constant. Hence, $\mathbf{T}^T = \Phi$ is the Moore–Penrose pseudoinverse transform (corresponding to the minimal dual synthesis frame), up to the constant c . In other words, the pseudoinverse reconstruction operator \mathbf{T}^+ corresponds to $c^{-1}\Phi$. It turns out that $\mathbf{T}^+\alpha$ is the reconstruction operation implemented by most implicit synthesis algorithms.

Choosing an appropriate dictionary is a key step towards a good sparse decomposition. Thus, to represent efficiently isotropic structures, a qualifying choice is the wavelet transform [15], [16]. The curvelet system [17] is a very good candidate for representing piecewise smooth (C^2) images away from C^2 contours. The ridgelet transform [18] has been shown to be very effective for sparsely representing global lines in an image. For locally oscillating textures, one can think of the local discrete cosine transform (DCT) [15], waveatoms [19] or brushlets [20]. These transforms are also computationally tractable particularly in large-scale applications and, as stated above, never explicitly implement Φ and \mathbf{T} . The associated implicit fast analysis and synthesis operators have typical complexities of $O(N)$, with N the number of samples or pixels (e.g., orthogonal or biorthogonal wavelet transform) or $O(N \log N)$ (e.g., ridgelets, curvelets, and local DCT, waveatoms).

What happens if none of the known fixed transforms can efficiently sparsify a morphological component; e.g., a complex natural texture. In [21], the authors have extended the MCA algorithm to handle the case where the dictionary attached to each morphological component is not necessarily fixed *a priori* as above but learned from a set of exemplars in order to capture complex textural patterns.

D. Thresholding Strategy

In practice, hard thresholding leads to better results. Furthermore, in [14], we empirically showed that the use of hard thresholding is likely to provide the ℓ_0 -sparsest solution. As far as the thresholding decreasing strategy is concerned, there are several alternatives. For example, in [1] and [2], linear and exponential decrease were advocated. In [14], a more elaborated strategy coined MOM (for mean-of-max) was proposed.

E. Handling Bounded Noise

MCA handles in a natural way data perturbed by additive noise ε with bounded variance σ_ε^2 . Indeed, as MCA is a coarse-to-fine iterative procedure, bounded noise can be handled just by stopping iterating when the residual

is at the noise level. Assuming that the noise variance σ_ε^2 is known, the algorithm may be stopped at iteration t when the ℓ_2 -norm of the residual satisfies $\|r^{(t)}\|_2 \leq \sqrt{N}\sigma_\varepsilon$. Alternatively, one may use a strategy reminiscent of denoising methods by taking $\lambda_{\min} = \tau\sigma_\varepsilon$, where τ is a constant, typically between three and four.

F. Applications

Fig. 2 shows examples of application of the MCA sparse decomposition algorithm to three real images: (a)–(c) Barbara, (d)–(f) X-ray riser image, and (g)–(j) an astronomical image of the galaxy SBS 0335-052. The riser in Fig. 2(d) is made of a composite material layer, a layer of steel-made fibers having opposite lay angles, and lead-made markers used as a reference to calibrate the X-ray camera. The observed image is corrupted by noise. The structures of interest are the curvilinear fibers. The astronomical image of Fig. 2(g) is contaminated by noise and a stripping artifact; the galaxy of interest is hardly visible in the original data.

The dictionaries used for the three images are, respectively: local DCT + curvelets for Barbara to decompose it into cartoon + texture parts; translation invariant wavelets + curvelets for the riser image; and ridgelets + curvelets + translation invariant wavelets for the astronomical image. The details of the experimental setup including the parameters of the dictionaries for each image are found in [22]. From Fig. 2(e) and (f), one can clearly see how MCA managed to get rid of the lead-made markers while reconstructing the curvilinear fibers structure. In Fig. 2(j), the galaxy has been well detected in the wavelet space while the stripping artifact was remarkably captured and removed owing to ridgelets and curvelets.

III. MULTICHANNEL SPARSE SOURCE SEPARATION

A. The Blind Source Separation Problem

In the BSS setting, the instantaneous linear mixture model assumes that we are given m observations (channels) $\{y_1, \dots, y_m\}$, where each y_j is a row-vector of size N ; each channel is the linear mixture of n sources

$$\forall j \in \{1, \dots, m\}, \quad y_j = \sum_{i=1}^n a_j[i]s_i + \varepsilon_j \quad (2)$$

or, equivalently, in matrix form

$$\mathbf{Y} = \mathbf{AS} + \mathbf{E} \quad (3)$$

where $\mathbf{Y} = [y_1^T, \dots, y_m^T]^T$ is the $m \times N$ measurement matrix whose rows are the channels y_j , $\mathbf{S} = [s_1^T, \dots, s_n^T]^T$ is the

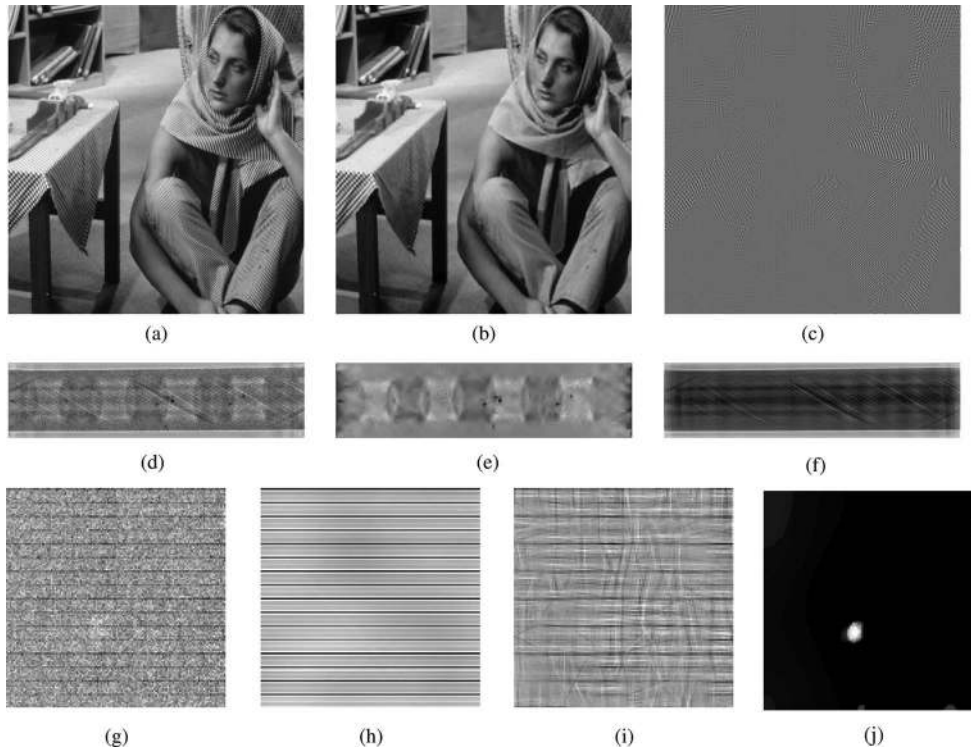


Fig. 2. MCA of three real two-dimensional images. *Barbara*: (a) original, (b) cartoon component (curvelets), (c) texture (local DCT). *X-ray riser image*: (d) observed, (e) isotropic structures and background (wavelets), (f) curvilinear fibers (curvelets). *Galaxy SBS 0335-052*: (g) observed, (h) ridgelet component, (i) curvelet component, (j) detected galaxy (wavelets).

A source matrix whose rows are the sources s_i , and \mathbf{A} is the $m \times n$ mixing matrix. \mathbf{A} defines the contribution of each source to each channel. An $m \times N$ matrix \mathbf{E} is added to account for instrumental noise or model imperfections. See Fig. 1.

Source separation techniques aim at recovering the original sources \mathbf{S} by taking advantage of some information contained in the way the signals are mixed. In the blind approach, both the mixing matrix \mathbf{A} and the sources \mathbf{S} are unknown. Source separation is overwhelmingly a question of contrast and diversity. Indeed, source separation boils down to devising quantitative measures of diversity or contrast to extricate the sources. Typical measures of contrast are statistical criteria such as independence (i.e., independent component analysis (ICA) [23]) or sparsity and morphological diversity; see [3], [4], [7], [24], and [25] and references therein.

B. Generalized Morphological Component Analysis

The GMCA framework assumes that the observed data \mathbf{Y} is a linear instantaneous mixture of unknown sources \mathbf{S} with an unknown mixing matrix \mathbf{A} , as in (3). For notational convenience, the dictionaries in the multi-channel case will be transposed versions of those considered in the single-channel case in Section II; each dictionary Φ_k is now a matrix whose rows are unit-norm

atoms. Thus, we let Φ be the concatenation of K transforms $\Phi = [\Phi_1^T, \dots, \Phi_K^T]^T$.

The GMCA framework assumes *a priori* that the sources $(s_i)_{i=1, \dots, n}$ are sparse in the dictionary Φ , $\forall i$ $s_i = \alpha_i \Phi$, where α_i is sparse (or compressible). More precisely, in the GMCA setting, each source is modeled as the linear combination of K morphological components, where each component is sparse in a specific basis

$$\forall i \in \{1, \dots, n\}; \quad s_i = \sum_{k=1}^K x_{ik} = \sum_{k=1}^K \alpha_{ik} \Phi_k. \quad (4)$$

GMCA seeks an unmixing scheme, through the estimation of \mathbf{A} , which leads to the sparsest sources \mathbf{S} in the dictionary Φ . This is expressed by the following optimization problem:

$$\min_{\mathbf{A}, \alpha} \sum_{i=1}^n \sum_{k=1}^K \|\alpha_{ik}\|_p^p \text{ such that } \|\mathbf{Y} - \mathbf{A}\alpha\Phi\|_F \leq \sigma$$

$$\text{and } \|a^i\|_2 = 1 \quad \forall i = 1, \dots, n \quad (5)$$

where typically $p = 0$ or a relaxed version with $p = 1$. But other sparsity regularization terms can be used in (5), e.g.,

mixed-norms [26]. The unit ℓ_2 -norm constraint on the columns of \mathbf{A} avoids the classical scale indeterminacy of the product $\mathbf{A}\mathbf{S}$ in (3). The reader may have noticed that the MCA problem (1) is a special case of the GMCA problem (5) when there is only one source $n = 1$ and one channel $m = 1$ (no mixing). Thus GMCA is indeed a multichannel generalization of MCA.

Equation (5) is a difficult nonconvex optimization problem even for convex penalties $p \geq 1$. More conveniently, the product $\mathbf{A}\mathbf{S}$ can be split into $n \cdot K$ multichannel morphological components: $\mathbf{A}\mathbf{S} = \sum_{i,k} a^i x_{ik} = \mathbf{A}\alpha\Phi = \sum_{i,k} a^i \alpha_{ik} \Phi_k$. Based on this decomposition, GMCA yields an alternating minimization algorithm to estimate iteratively one term at a time. It has been shown in [3] that estimating the morphological component $x_{ik} = \alpha_{ik} \Phi_k$ assuming \mathbf{A} and $x_{\{pq\} \neq \{ik\}}$ can be obtained through iterative thresholding for $p = 0$ and $p = 1$.

Now, considering fixed $\{a^p\}_{p \neq i}$ and \mathbf{S} , updating the column a^i is then just a least squares estimate

$$a^i = \frac{1}{\|s_i\|_2^2} \left(\mathbf{Y} - \sum_{p \neq i} a^p s_p \right) s_i^T. \quad (6)$$

This estimate is then projected onto the unit sphere to meet the unit ℓ_2 -norm constraint in (5).

C. GMCA Algorithm

The GMCA algorithm is summarized in Algorithm 2. In the same vein as MCA, GMCA also relies on a salient-to-fine strategy. More precisely, GMCA is an iterative thresholding algorithm such that, at each iteration, it first computes coarse versions of the morphological components for a fixed source s_i . These raw sources are estimated from their most significant coefficients in Φ . Then, the corresponding

column a^i is estimated from the most significant features of s_i . Each source and its corresponding column of \mathbf{A} is then alternately and progressively refined as the threshold λ decreases towards λ_{\min} . This particular iterative thresholding scheme provides robustness to noise, model imperfections, and initialization by working first on the most significant features in the data and then progressively incorporating smaller details to finely tune the model parameters. As a multichannel extension of MCA, GMCA is also robust to noise and can be used with either linear or exponential decrease of the threshold [3], [4]. Moreover, hard thresholding leads to its best practical performance.

D. Unknown Number of Sources

In BSS, the number of sources n is assumed to be a fixed known parameter of the problem. This is rather an exception than a rule, and estimating n from the data is a crucial and strenuous problem. Only a few works have attacked this issue. One can think of using model selection criteria such as the minimum description length used in [27]. In [4], a sparsity-based method to estimate n within the GMCA framework was proposed. Roughly speaking, this selection procedure uses GMCA to solve a sequence of problems (5) for each constraint radius $\sigma(q)$ with increasing q , $1 \leq q \leq m$. In [4], it was argued to set $\sigma(q)$ to the Frobenius-norm of the error when approximating the data matrix \mathbf{Y} with its largest q singular vectors; see [4] for further details.

E. Hyperspectral Data

In standard BSS, \mathbf{A} is often seen as a mixing matrix of small size $m \times n$. On the other hand, there are applications in which one deals with data from instruments with a very large number of channels m , which are well organized according to some physically meaningful index. A typical example is *hyperspectral* data, where images are collected in

Algorithm 2: GMCA BSS algorithm.

Task: Blind source separation.

Parameters: Observed channels \mathbf{Y} , the dictionary $\Phi = [\Phi_1^T \cdots \Phi_K^T]^T$, number of sources n , number of iterations N_{iter} , initial and stopping thresholds $\lambda^{(0)}$ and λ_{\min} , threshold update schedule.

Main iteration:

for $t = 1$ **to** N_{iter} **do**

for $i = 1$ **to** n **do**

for $k = 1$ **to** K **do**

 Compute marginal residuals $\mathbf{R}_{ik}^{(t)} = \mathbf{Y} - \sum_{\{p,q\} \neq \{i,k\}} a^{p(t-1)} x_{pq}^{(t-1)}$.

 Update (i, k) -th coefficients by thresholding $\alpha_{ik}^{(t)} = \text{TH}_{\lambda^{(t-1)}} \left(a^{i(t-1)T} \mathbf{R}_{ik}^{(t)} \Phi_k^T \right)$.

 Update (i, k) -th component by reconstructing from the selected coefficients $x_{ik}^{(t)} = \alpha_{ik}^{(t)} \Phi_k$.

 Update i -th source $s_i^{(t)} = \sum_{k=1}^K x_{ik}^{(t)}$.

 Update i -th mixing matrix column $a^{i(t)} = \frac{1}{\|s_i^{(t)}\|_2} \left(\mathbf{Y} - \sum_{p \neq i} a^{p(t-1)} s_p^{(t)} \right) s_i^{(t)T}$ and normalize to a unit norm.

 Update the threshold $\lambda^{(t)}$ according to the given schedule.

if $\lambda^{(t)} \leq \lambda_{\min}$ **then stop.**

Output: Sources $\left(s_i^{(N_{\text{iter}})} \right)_{i=1, \dots, n}$ and mixing matrix $\mathbf{A}^{(N_{\text{iter}})}$.

a large number (i.e., hundreds) of contiguous regions of the electromagnetic spectrum. Regardless of other definitions appearing in other scientific communities, the term “hyperspectral” is used here for multichannel data model (3) where the number of channels m is large and these channels achieve a uniform sampling of some meaningful physical index (e.g., wavelength, space, time), which we refer to as the *spectral* dimension. For such data, it then makes sense to consider the regularity of the spectral signatures $(a^i)_{i=1,\dots,n}$. For instance, these spectral signatures may be known a priori to have a sparse representation in some specified possibly redundant dictionary Ψ of *spectral* waveforms.

In [4] and [28], the authors propose a modified GMCA algorithm capable of handling hyperspectral data. This is achieved by assuming that each rank-one matrix $\mathbf{X}_i = a^i s_i$ has a sparse representation in the multichannel dictionary $\Phi \otimes \Psi$ [9], [29] (\otimes is the Kronecker product); i.e., $\forall i$ $s_i = \alpha_i \Phi$ and $a^i = \Psi \beta^i$, where α_i and β_i are both sparse. The separation optimization problem for hyperspectral data is then

$$\min_{\alpha, \beta} \frac{1}{2} \left\| \mathbf{Y} - \sum_{i=1}^n \Psi \beta^i \alpha_i \Phi \right\|_F^2 + \sum_{i=1}^n \lambda_i \|\beta^i \alpha_i\|_1 \quad (7)$$

where a Laplacian prior is imposed on each α_i conditionally on β_i , and vice versa. Remarkably, this joint prior preserves the scale invariance of (3). Equation (7) is again a nonconvex optimization problem for which no closed-form solution exists. In the line of the GMCA algorithm, thanks to the form of the ℓ_1 -penalty in (7), a block-relaxation iterative thresholding algorithm was proposed in [4] and [28] that alternately minimizes (7) with respect to α and β . It was shown by these authors that the update equations on the coefficient matrices are

$$\begin{aligned} \alpha^{(t+1)} &= \text{ST}_{\eta^{(t)}} \left(\left(\beta^{(t)T} \beta^{(t)} \right)^{-1} \beta^{(t)T} \Psi^T \mathbf{Y} \Phi^T \right) \\ \beta^{(t+1)} &= \text{ST}_{\zeta^{(t)}} \left(\Psi^T \mathbf{Y} \Phi^T \alpha^{(t)T} \left(\alpha^{(t)} \alpha^{(t)T} \right)^{-1} \right). \end{aligned} \quad (8)$$

$\eta^{(t)}$ is a vector of length n and entries $\eta^{(t)}[i] = \lambda^{(t)} \|\beta^{(t)}\|_1 / \|\beta^{(t)}\|_2^2$; $\zeta^{(t)}$ has length m and entries $\zeta^{(t)}[j] = \lambda^{(t)} \|\alpha_j^{(t)}\|_1 / \|\alpha_j^{(t)}\|_2^2$; and $\lambda^{(t)}$ is a decreasing threshold. The multichannel soft-thresholding operator ST_{η} acts on each row i with threshold $\eta[i]$ and ST_{ζ} acts on each column j with threshold $\zeta[j]$.

F. Applications

BSS: We first report a simple BSS application. Fig. 3 shows (a) two original sources and (b) the two noisy

mixtures [signal-to-noise ratio (SNR) = 10 dB]. GMCA was applied with a dictionary containing curvelets and local DCT. As a quantitative performance measure of BSS methods, we use the mixing matrix criterion $\Delta_{\mathbf{A}} = \|\mathbf{I}_n - \hat{\mathbf{P}} \mathbf{A}^+ \mathbf{A}\|_1$, where $\hat{\mathbf{A}}^+$ is the pseudoinverse of the estimate of the mixing matrix \mathbf{A} and $\hat{\mathbf{P}}$ is a matrix that reduces the scale/permutation indeterminacy of the mixing model. Fig. 3(c) compares GMCA to popular BSS techniques in terms of $\Delta_{\mathbf{A}}$ as the SNR increases. We compare GMCA to ICA JADE [30], relative Newton algorithm (RNA) [31] that accounts for sparsity, and EFICA [32]. Both RNA and EFICA were applied after “sparsifying” the data via an orthonormal wavelet transform. It can be seen that JADE performs rather badly, while RNA and EFICA behave quite similarly. GMCA seems to provide much better results, especially at high noise level.

Color Image Denoising: GMCA can be applied to color image denoising. This is illustrated in Fig. 4, where the original RGB image is shown in (a). Fig. 4(b) shows the RGB image obtained using a classical undecimated wavelet-domain hard thresholding on each color plane independently. GMCA is applied to the RGB color channels using the curvelet dictionary. In the notation of Section III-A, we have $m = 3$ channels (color planes); each color channel is y_j , $j = 1, 2, 3$, $n = 3$ sources; and \mathbf{A} plays the role of the color space conversion matrix. Unlike classical color spaces (e.g., YUV, YCC), where the conversion matrix from RGB is fixed, the color space conversion matrix is here estimated by GMCA from the data. As such, GMCA is able to find adaptively the appropriate color space corresponding to the color image at hand. Once \mathbf{A} is estimated by GMCA, we applied the same undecimated wavelet-based denoising to the estimated sources. The denoised data are obtained by coming back to the RGB space via the estimated mixing matrix. Fig. 4(c) shows the GMCA-based denoised image. Clearly, denoising in the “GMCA color space” is substantially better than in the RGB space (or other color spaces such as YUV or YCC; see [3]).

Hyperspectral Data Processing: In this experiment, we consider $m = 128$ mixtures of $n = 5$ source images. The sources are drawn at random from a set of structured images shown in Fig. 5(a). For the spectra (i.e., columns of \mathbf{A}), we randomly generated sparse coefficient vectors $(i = 1, \dots, n)$ ($i = 1, \dots, n$) with independent Gaussian-distributed nonzero entries and then applied the inverse orthogonal wavelet transform to these sparse vectors to get the spectra. Φ was chosen as the curvelet dictionary. Fig. 5(b) gives four typical noisy observed channels with SNR = 20 dB. The sources recovered using GMCA (Algorithm 2) and its hyperspectral extension (iteration of Section III-E) are shown, respectively, in Fig. 5(c) and (d). Visual inspection shows that GMCA is outperformed

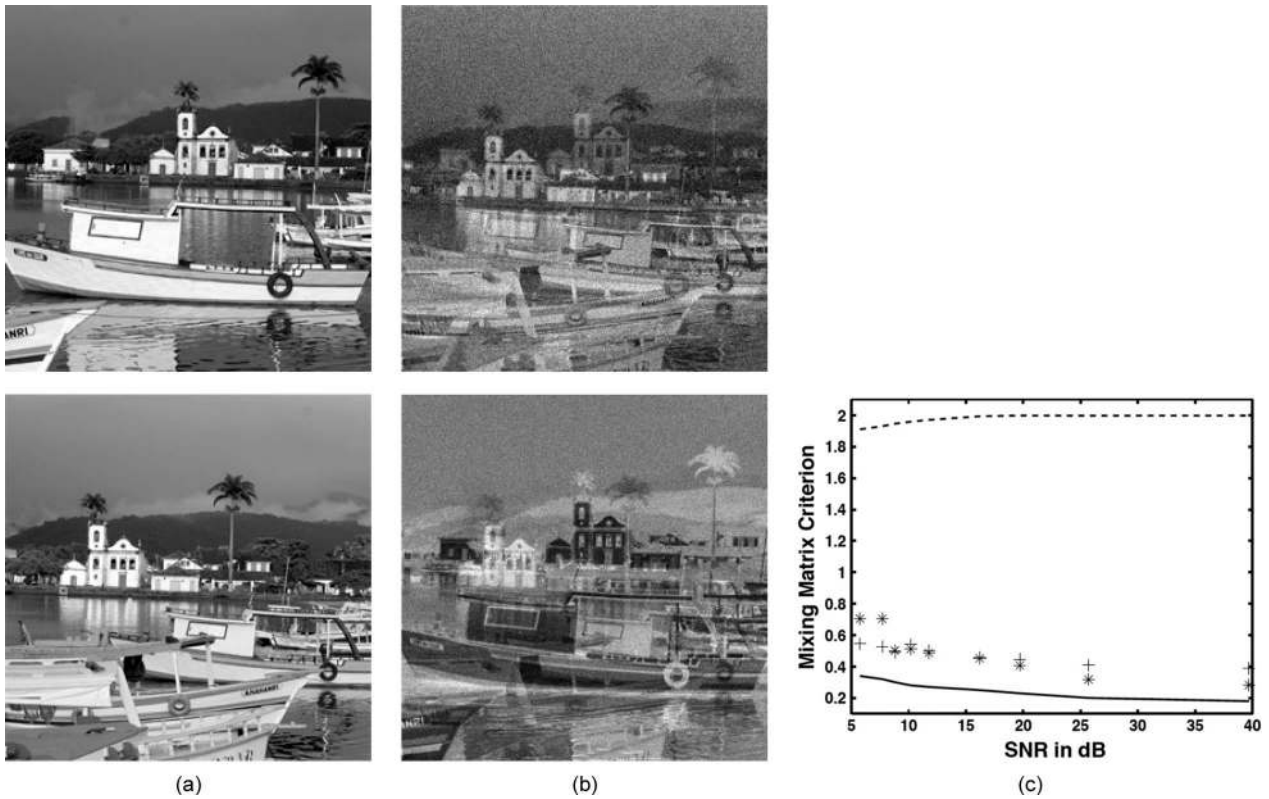


Fig. 3. Example of BSS with (a) two sources and (b) two noisy mixtures. (c) depicts the evolution of the mixing matrix criterion Δ_A with input SNR [solid line: GMCA; dashed line: JADE; (+): RNA; (*): EFICA].



Fig. 4. (a) Original RGB image with additive Gaussian noise SNR = 15 dB. (b) Wavelet-based denoising in the RGB space. (c) Wavelet-based denoising in the “GMCA color space.”

by hyperspectral GMCA, which better accounts for both spatial and spectral sparsity.

IV. REPRODUCIBLE RESEARCH SOFTWARE

Following the philosophy of reproducible research, two toolboxes, MCALab and GMCALab [22], are made

available freely for download.¹ MCALab and GMCALab have been developed to demonstrate key concepts of MCA and GMCA and make them available to interested researchers and technologists. These toolboxes are libraries of MATLAB routines that implement the decomposition

¹<http://www.morphologicaldiversity.org>.

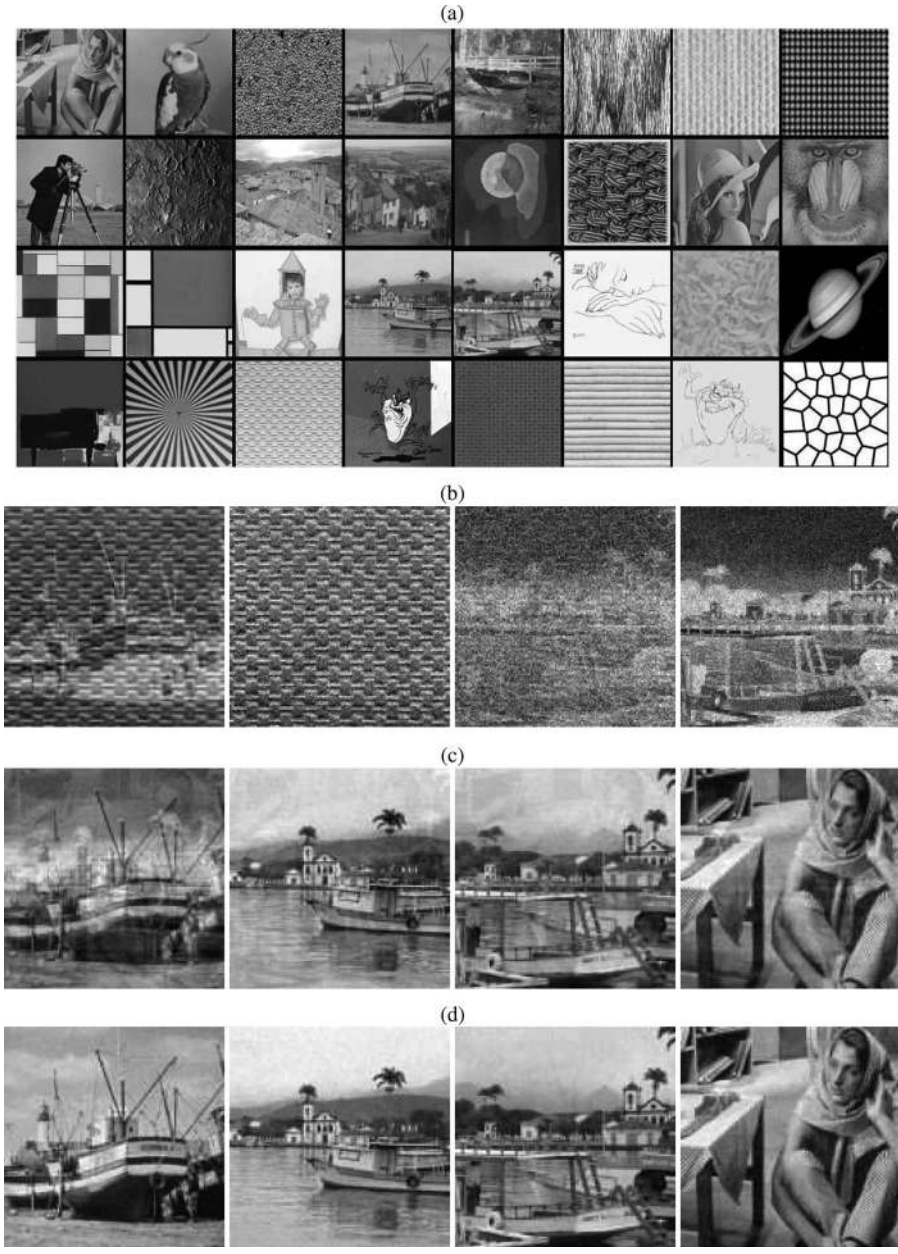


Fig. 5. (a) Images used as sources. (b) Four noisy mixtures out of $m = 128$ (SNR = 20 dB). (c) Recovered sources using GMCA. (d) Recovered sources using hyperspectral GMCA.

and source separation algorithms overviewed in this paper. They contain a variety of scripts to reproduce the figures in our own papers, as well as other exploratory examples not included in the papers.

V. CONCLUSION

In this paper, we gave an overview of how sparsity and morphological diversity can be used advantageously to

regularize image decomposition and blind source separation problems. We also reported several numerical experiments to illustrate the wide applicability of the algorithms described. We believe that this is an exciting field where many interesting problems are still open. Among them, we may cite, for instance, the theoretical guarantees of the sparsity-regularized BSS problem and sharper theoretical guarantees for the decomposition problem by exploiting geometry. ■

REFERENCES

- [1] J.-L. Starck, M. Elad, and D. Donoho, "Image decomposition via the combination of sparse representations and variational approach," *IEEE Trans. Image Process.*, vol. 14, no. 10, pp. 1570–1582, 2005.
- [2] J.-L. Starck, M. Elad, and D. Donoho, "Redundant multiscale transforms and their application for morphological component analysis," in *Advances in Imaging and Electron Physics*, vol. 132, P. Hawkes, Ed. New York: Academic/Elsevier, 2004, vol. 132.
- [3] J. Bobin, J.-L. Starck, M. J. Fadili, and Y. Moudden, "Sparsity and morphological diversity in blind source separation," *IEEE Trans. Image Process.*, vol. 16, pp. 2662–2674, Nov. 2007.
- [4] J. Bobin, J.-L. Starck, Y. Moudden, and M. J. Fadili, "Blind source separation: The sparsity revolution," in *Advances in Imaging and Electron Physics*, P. Hawkes, Ed. New York: Academic/Elsevier, 2008, pp. 221–298.
- [5] Y. Meyer, "Oscillating patterns in image processing and in some nonlinear evolution equations," in *15th Dean Jacqueline B. Lewis Memorial Lectures*, 2001.
- [6] L. A. Vese and S. J. Osher, "Modeling textures with total variation minimization and oscillating patterns in image processing," *J. Sci. Comput.*, vol. 19, no. 1–3, pp. 553–572, 2003.
- [7] M. Zibulevsky and B. Pearlmutter, "Blind source separation by sparse decomposition," *Neural Comput.*, vol. 13/4, 2001.
- [8] P. Bofill and M. Zibulevsky, "Underdetermined blind source separation using sparse representations," *Signal Process.*, vol. 81, no. 11, pp. 2353–2362, 2001.
- [9] J. Bobin, Y. Moudden, M. J. Fadili, and J.-L. Starck, "Morphological diversity and sparsity for multichannel data restoration," *J. Math. Imag. Vision*, vol. 10, no. 2, pp. 149–168, 2009.
- [10] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Rev.*, vol. 51, no. 1, pp. 34–81, 2009.
- [11] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation noise removal algorithm," *Phys. D*, vol. 60, pp. 259–268, 1992.
- [12] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [13] S. Sardy, A. Bruce, and P. Tseng, "Block coordinate relaxation methods for nonparametric wavelet denoising," *J. Comput. Graph. Statist.*, vol. 9, no. 2, pp. 361–379, 2000.
- [14] J. Bobin, J.-L. Starck, M. J. Fadili, Y. Moudden, and D. L. Donoho, "Morphological component analysis: An adaptive thresholding strategy," *IEEE Trans. Image Process.*, vol. 16, pp. 2675–2681, Nov. 2007.
- [15] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. New York: Academic, 1998.
- [16] J.-L. Starck, F. Murtagh, and A. Bijaoui, *Image Processing and Data Analysis: The Multiscale Approach*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [17] E. Candès, L. Demanet, D. Donoho, and L. Ying, "Fast discrete curvelet transforms," *SIAM Multiscale Model. Simul.*, vol. 5, pp. 861–899, 2005.
- [18] E. Candès and D. Donoho, "Ridgelets: The key to high dimensional intermittency?" *Phil. Trans. Royal Soc. London A*, vol. 357, pp. 2495–2509, 1999.
- [19] L. Demanet and L. Ying, "Wave atoms and sparsity of oscillatory patterns," *Appl. Comput. Harmon. Anal.*, vol. 23, no. 3, pp. 368–387, 2007.
- [20] F. G. Meyer and R. R. Coifman, "Brushlets: A tool for directional image analysis and image compression," *Appl. Comput. Harmon. Anal.*, vol. 5, pp. 147–187, 1997.
- [21] G. Peyré, M. J. Fadili, and J.-L. Starck, "Learning adapted dictionaries for geometry and texture separation," in *Proc. SPIE Conf. Wavelets XII*, San Diego, CA, Aug. 2007.
- [22] M. J. Fadili, J.-L. Starck, M. Elad, and D. L. Donoho, "MCALab: Reproducible research in signal and image decomposition and inpainting," *IEEE Comput. Sci. Eng.*, to be published.
- [23] A. Hyvärinen, J. Karhunen, and E. Oja, *Independent Component Analysis*. New York: Wiley, 2001.
- [24] P. G. Georgiev, F. Theis, and A. Cichocki, "Sparse component analysis and blind source separation of underdetermined mixtures," *IEEE Trans. Neural Netw.*, vol. 16, no. 4, pp. 992–996, 2005.
- [25] J. Bobin, Y. Moudden, J.-L. Starck, and M. Elad, "Morphological diversity and source separation," *IEEE Signal Process. Lett.*, vol. 13, no. 7, pp. 409–412, 2006.
- [26] M. Kowalski, E. Vincent, and R. Gribonval, "Under-determined source separation via mixed-norm regularized minimization," in *Proc. EUSIPCO'08*, Lausanne, Switzerland, Aug. 2008.
- [27] R. Balan, "Estimator for number of sources using minimum description length criterion for blind sparse source mixtures," in *Independent Component Analysis and Signal Separation*, vol. 4666, M. E. Davies, C. J. James, S. A. Abdallah, and M. D. Plumbley, Eds. New York: Springer, 2007, pp. 333–340.
- [28] Y. Moudden, J. Bobin, J.-L. Starck, and M. J. Fadili, "Dictionary learning with spatio-spectral sparsity constraints," in *Proc. SPARS'09*, St. Malo, France, Apr. 2009.
- [29] R. Gribonval and M. Nielsen, "Beyond sparsity: Recovering structured representations by l_1 -minimization and greedy algorithms: Application to the analysis of sparse underdetermined ica," *Adv. Comput. Math.*, vol. 28, no. 1, pp. 23–41, 2008.
- [30] J.-F. Cardoso, "Blind signal separation: Statistical principles," *Proc. IEEE (Special Issue on Blind Identification and Estimation)*, vol. 9, no. 10, pp. 2009–2025, Oct. 1998.
- [31] M. Zibulevsky, "Blind source separation with relative newton method," in *Proc. ICA2003*, 2003, pp. 897–902.
- [32] Z. Koldovsky, P. Tichavsky, and E. Oja, "Efficient variant of algorithm fastica for independent component analysis attaining the cramer-rao lower bound," *IEEE Trans. Neural Netw.*, vol. 17, pp. 1265–1277, 2006.

ABOUT THE AUTHORS

M. Jalal Fadili graduated from the Ecole Nationale Supérieure d'Ingénieurs (ENSI) de Caen, Caen, France. He received the M.Sc. and Ph.D. degrees in signal and image processing from the University of Caen.

He was a Research Associate with the University of Cambridge, Cambridge, U.K., from 1999 to 2000, where he was a MacDonnel-Pew Fellow. He has been an Associate Professor of signal and image processing since September 2001 at ENSI. He was a Visitor at several universities (QUT-Australia, Stanford University, California Institute of Technology, EPFL). His research interests include statistical approaches in signal and image processing, inverse problems, computational harmonic analysis, optimization theory, and sparse representations. His areas of application include medical and astronomical imaging.



Jean-Luc Starck received the Ph.D. degree from the University Nice-Sophia Antipolis, France, and the Habilitation degree from the University Paris XI, Orsay, France.

He was a Visitor with the European Southern Observatory (ESO) in 1993, the University of California, Los Angeles, in 2004, and the Statistics Department, Stanford University, in 2000 and 2005. He has been a Researcher with CEA, France, since 1994. His research interests include image processing, statistical methods in astrophysics, and cosmology. He is an expert in multiscale methods such as wavelets and curvelets. He is a Leader of the Multiresolution project at CEA and is a core Team Member of the PLANCK ESA project. He has published more than 100 papers in different areas in scientific journals. He is also author of *Image Processing and Data Analysis: The Multiscale Approach* (Cambridge, U.K.: Cambridge University Press, 1998) and *Astronomical Image and Data Analysis*, 2nd ed. (New York: Springer, 2006).



Jérôme Bobin graduated from the Ecole Normale Supérieure (ENS) de Cachan, France, in 2005. He received the M.Sc. degree in signal and image processing from ENS de Cachan and University Paris XI, Orsay, France, and the Ph.D. degree from Paris XI in 2008. He received the Agrégation de Physique degree in 2004.

He is now a Postdoctoral Researcher with ACM, California Institute of Technology, Pasadena, CA. His research interests include statistics, information theory, multiscale methods, and sparse representations in signal and image processing.



Yassir Moudden graduated in electrical engineering from SUPELEC, Gif-sur-Yvette, France. He received the M.S. degree in physics from University Paris VII, France, in 1997 and the Ph.D. degree in signal processing from the University Paris XI, Orsay, France.

He was a Visitor at the University of California, Los Angeles, in 2004 and is currently with CEA, France, working on applications of signal processing to astronomy. His research interests include signal and image processing, data analysis, statistics, and information theory.

