

Morphological Component Analysis and Inpainting on the Sphere: Application in Physics and Astrophysics

P. Abrial¹, Y. Moudden¹, J.L. Starck¹, B. Afeyan³
J. Bobin¹, J. Fadili² and M.K. Nguyen⁴

¹DAPNIA/SEDI-SAP, Service d'Astrophysique, CEA-Saclay, 91191 Gif-sur-Yvette, France

²GREYC - CNRS UMR 6072, 10 Bd Maréchal Juin, 14050 Caen, France

³Polymath Research Inc., Pleasanton, CA, 94566, USA

⁴ETIS - CNRS UMR 8051/ENSEA/Universit de Cergy-Pontoise, France

September 4, 2007

Abstract

Morphological Component Analysis (MCA) is a new method which takes advantage of the sparse representation of structured data in large overcomplete dictionaries to separate features in the data based on the diversity of their morphology. It is an efficient technique in such problems as separating an image into texture and piecewise smooth parts or for inpainting applications. The MCA algorithm consists of an iterative alternating projection and thresholding scheme, using a successively decreasing threshold towards zero with each iteration. In this paper, the MCA algorithm is extended to the analysis of spherical data maps as may occur in a number of areas such as geophysics, astrophysics or medical imaging. Practically, this extension is made possible thanks to the variety of recently developed transforms on the sphere including several multiscale transforms such as the undecimated isotropic wavelet transform on the sphere, the ridgelet and curvelet transforms on the sphere. An MCA-inpainting method is then directly extended to the case of spherical maps allowing us to treat problems where parts of the data are missing or corrupted. We demonstrate the usefulness of these new tools of spherical data analysis by focusing on a selection of challenging applications in physics and astrophysics.

Keywords and Phrases. : Spherical Data, Morphological Component Analysis, Sparse representation, Inpainting.

1 Introduction

A usual task in processing signals, images as well as spherical data maps, is to decompose the data into its elementary building blocks. This can be formulated as an inverse problem where the data is assumed to have been generated according to the following model :

$$y = \sum_i \alpha_i \phi_i + \eta \quad (1)$$

that is a linear combination of relevant waveforms $\phi_i \in \mathbb{R}^n$ with weights α_i . Here η represents possible contamination by additive, typically Gaussian white noise. Given data $y \in \mathbb{R}^n$, one then wants to recover the underlying structures that is to say estimate a set of waveforms ϕ_i that build the data and

their corresponding weights $\tilde{\alpha}_i$. The solution to this estimation problem will depend heavily on the available prior information. Of interest here is the case where one is given *a priori* a set of waveforms from which to select a good subset. This set may be a basis, a frame or several bases or frames grouped into a large redundant dictionary.

Possible dictionaries in 1D and 2D include Fourier and related bases, wavelet bases, as well as other more recent multiscale systems such as the ridgelet [17] and curvelet frames [24, 52], etc. Depending on the morphology of the data, each of these dictionaries will have different performance characteristics in a non-linear approximation scheme. For instance, sparse approximations of piecewise smooth signals or images with point singularities are easily obtained using wavelets. However these are no longer optimal in the case of piecewise smooth images with singularities along smooth curves or edges. Such images are more efficiently approximated using curvelets which are highly anisotropic and thus exhibit high directional selectivity. Digital implementations of both ridgelet and curvelet transforms and their application to image denoising are described in [52].

Available transforms in the spherical topology include the spherical harmonics and several wavelet transforms. Software packages such as Healpix¹ [36] or Glesp [29] provide approximate digital spherical harmonic transform routines based on their specific pixelization schemes. Schröder and Sweldens [49] have developed an orthogonal wavelet transform on the sphere based on the Haar wavelet function which then suffers from the poor frequency domain localization properties of the primitive Haar function and from the problems inherent in orthogonal decomposition (*e.g.* lack of translation invariance). A few papers describe continuous wavelet transforms on the sphere [2, 18, 39, 59, 9] which have been extended to directional wavelet transforms [60, 45]. Although useful for data analysis, these continuous transforms lack an inverse transform and hence are clearly not suitable for restoration or synthesis purposes.

In their pioneering work, Freedman and Maier [33, 34] gave a wavelet transform/reconstruction scheme on the sphere which is based on the spherical harmonic transform. Following this idea, Starck *et al.* [56] have proposed a new invertible isotropic undecimated wavelet transform (UWT) on the sphere which preserves the same desirable properties as the standard isotropic UWT for flat 2D maps [57]: the reconstruction is simple and immediate since it is just the addition of all the wavelet bands with the coarsest scale. Based on this new decomposition, other multiscale transforms such as the pyramidal wavelet transform, the ridgelet transform and the curvelet transform have been successfully constructed on the sphere [56]. Each of these decompositions on the sphere will sparsely represent parts of the image based on their morphological properties. Wavelets will easily detect more or less isotropic localized structures, while curvelets are better suited for efficiently detecting highly anisotropic objects.

A data set y has an exact representation over any complete basis of the data space, or several such exact representations in the case of redundant overcomplete dictionaries. However, these representations are not equally interesting in terms of data modeling or feature detection. In fact, a strong *a priori* is to favor representations of y that use only a small number of waveforms leading to a more concise and possibly more interpretable representation of the data. In fact, building sparse representations or approximations is the (he)art of structured data processing: the design of good detection, denoising, restoration and compression algorithms relies on the availability of good dictionaries and good selection algorithms. Indeed, selecting the smallest subset of waveforms from a large dictionary, that will linearly combine to reproduce the salient features of a given signal or image, is a hard combinatorial problem. Several *pursuit* algorithms have been proposed that can help build very sparse decompositions such as the greedy Matching Pursuit (MP) [42] algorithm which refines the signal approximation by picking at each iteration the one waveform which best correlates with the current approximation error. Basis Pursuit

¹<http://www.eso.org/science/healpix>

(BP) [20] is a global procedure which seeks an approximation \tilde{y} to y by solving the linear programming problem:

$$\min_{\alpha} \|\alpha\|_{\ell_1} \text{ subject to } y = \Phi\alpha. \quad (2)$$

where the ℓ_1 norm measures sparsity in place of the ℓ_0 counting norm. In the presence of noise, a noise-aware variant of BP, known as BPDN (for BP denoising), can be stated as a convex quadratic programming problem and solved using the Interior Point method [20]. The BPDN problem can also be written in the augmented Lagrangian form:

$$\min_{\alpha} \|y - \Phi\alpha\|_{\ell_2}^2 + \lambda \cdot \|\alpha\|_{\ell_1} \quad (3)$$

Among all possible solutions, the chosen one has the minimum ℓ_1 norm. This choice of ℓ_1 norm is very important. An ℓ_2 norm, as used in the method of frames [21], does not favor sparsity [20]. A number of recent results prove that these algorithms will recover the unique maximally sparse decomposition provided this solution is sparse enough and the dictionary is sufficiently incoherent [25, 31, 37, 26, 35]. Nevertheless, in problems involving large data sets (e.g. images, spherical maps), BP or MP synthesis algorithms are computationally prohibitive.

Morphological Component Analysis (MCA) is a recent faster alternative described in [54] that constructs a sparse representation of a signal or an image assuming that it is a combination of morphologically distinct features which are sparsely represented in different dictionaries associated with fast transform algorithms. For instance, images commonly combine contours and textures: the former are well accounted for using curvelets, while the latter may be well represented using local cosine functions. In searching for a sparse decomposition of a signal or image y , it is assumed that y is a sum of K components $(s_k)_{1,\dots,K}$, where each can be described as $s_k = \Phi_k \alpha_k$ with a possibly over-complete dictionary Φ_k and a sparse vector of coefficients α_k . It is further assumed that for any given component the sparsest decomposition over the proper dictionary yields a highly sparse description, while its decomposition over the other dictionaries, $\Phi_{k' \neq k}$, is non sparse. Thus, the different Φ_k can be seen as discriminating between the different components of the initial signal. MCA achieves its sparse decomposition relying on an iterative thresholding algorithm with a successively decreasing threshold [8] thus refining the current approximation by including finer structures alternately in the different morphological components. Based on MCA, it has also been shown that we can derive a very efficient inpainting method [32].

This paper: Motivated by the success of MCA in signal and image processing, the purpose of this contribution is to take advantage of the variety of transforms on the sphere recently made available [56] to extend the applicability of MCA to the analysis of spherical maps which are commonly recorded in a number of areas such as geophysics, astrophysics or medical imaging. As in the case of Euclidean 2D images, we further extend the MCA algorithm on the sphere in order to perform inpainting tasks on the sphere. The proposed numerical tools are shown to be valuable in several selected applications in physics and astrophysics. The construction of the undecimated isotropic wavelet and curvelet transforms on the sphere is reviewed in the next section. Sections 3 and 4 describe the extension to the sphere of the MCA algorithm and of its modification for inpainting purposes.

2 Transforms on the Sphere

2.1 Isotropic Undecimated Wavelet Transform on the Sphere (UWTS)

The undecimated isotropic transform on the sphere described in [56] is similar in many respects to the regular *à trous* isotropic wavelet transform. It is obtained using a zonal scaling function $\phi_{l_c}(\vartheta, \varphi)$ which depends only on colatitude ϑ and is invariant with respect to a change in longitude φ . It follows that the spherical harmonic coefficients $\hat{\phi}_{l_c}(l, m)$ of ϕ_{l_c} vanish when $m \neq 0$ which makes it simple to compute the spherical harmonic coefficients $\hat{c}_0(l, m)$ of $c_0 = \phi_{l_c} * f$ where $*$ stands for convolution :

$$\hat{c}_0(l, m) = \widehat{\phi_{l_c} * f}(l, m) = \sqrt{\frac{4\pi}{2l+1}} \hat{\phi}_{l_c}(l, 0) \hat{f}(l, m) \quad (4)$$

A possible scaling function [57], defined in the spherical harmonics representation, is $\phi_{l_c}(l, m) = \frac{2}{3} B_3(\frac{2l}{l_c})$ where B_3 is the cubic B-spline compactly supported over $[-2, 2]$. Denoting $\phi_{2^{-j}l_c}$ a rescaled version of ϕ_{l_c} with cut-off frequency $2^{-j}l_c$, a multi-resolution decomposition of f on a dyadic scale is obtained recursively :

$$\begin{aligned} c_0 &= \phi_{l_c} * f \\ c_j &= \phi_{2^{-j}l_c} * f = c_{j-1} * h_{j-1} \end{aligned} \quad (5)$$

where the zonal low pass filters h_j are defined by

$$\hat{H}_j(l, m) = \sqrt{\frac{4\pi}{2l+1}} \hat{h}_j(l, m) = \begin{cases} \frac{\hat{\phi}_{\frac{l_c}{2^{j+1}}}(l, m)}{\hat{\phi}_{\frac{l_c}{2^j}}(l, m)} & \text{if } l < \frac{l_c}{2^{j+1}} \quad \text{and} \quad m = 0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

The cut-off frequency is reduced by a factor of 2 at each step so that in applications where this is useful such as compression, the number of samples could be reduced by the same factor. Using a pixelization scheme such as Healpix [36], this can easily be done by dividing by 2 the Healpix *nside* parameter when computing the inverse spherical harmonics transform. Of course, this is only an approximate *Sampling Theorem* but it proved sufficient for numerical purposes. However, in the present isotropic undecimated transform, no downsampling is performed and the maps have the same number of pixels on each scale. Hence the orthogonality requirement is relaxed, which provides us with a higher degree of freedom in the choice and design of the wavelet function ψ_{l_c} to be used with the scaling function ϕ_{l_c} . As in the *à trous* algorithm, the wavelet coefficients can be defined as the difference between two consecutive resolutions, $w_{j+1}(\vartheta, \varphi) = c_j(\vartheta, \varphi) - c_{j+1}(\vartheta, \varphi)$ which defines a zonal wavelet function ψ_{l_c} as

$$\hat{\psi}_{\frac{l_c}{2^j}}(l, m) = \hat{\phi}_{\frac{l_c}{2^{j-1}}}(l, m) - \hat{\phi}_{\frac{l_c}{2^j}}(l, m) \quad (7)$$

This particular decomposition is readily inverted by summing the coefficient maps on all wavelet scales

$$f(\vartheta, \varphi) = c_J(\vartheta, \varphi) + \sum_{j=1}^J w_j(\vartheta, \varphi) \quad (8)$$

where we have made the simplifying assumption that f is equal to c_0 . Obviously, other wavelet functions could be used just as well. Also, because of the redundancy of the described decomposition, the inverse transform is not unique and in fact this can profitably be used to impose additional constraints on the synthesis functions (*e.g.* smoothness, positivity) used in the reconstruction [55].

2.2 Curvelets on the Sphere

The 2D ridgelet transform [17] was developed in an attempt to overcome some limitation inherent in former multiscale methods *e.g.* the 2D wavelet, when handling smooth images with edges *i.e.* singularities along smooth curves. Ridgelets are translation invariant *ridge* functions with a wavelet profile in the normal direction. Although ridgelets provide sparse representations of smooth images with straight edges, they fail to efficiently handle edges along curved lines. This is the framework for curvelets which were given a first mathematical description in [16]. Basically, the curvelet dictionary is a multiscale pyramid of localized directional functions with anisotropic support obeying a specific parabolic scaling such that at scale 2^{-j} , its length is $2^{-j/2}$ and its width is 2^{-j} . This is motivated by the parabolic scaling property of smooth curves. Other properties of the curvelet transform and decisive optimality results in approximation theory are reported in [16, 14]. Notably, curvelets provide optimally sparse representations of manifolds which are smooth away from edge singularities along smooth curves. Several digital curvelet transforms [24, 52, 22, 13] have been proposed which attempt to preserve the essential properties of the continuous curvelet transform and many papers [51, 54, 38, 58, 53] report on their successful application in image processing experiments. The contourlet transform in [22] is a filter bank implementation of the curvelet transform. Candès *et al.* implement fast discrete curvelet transforms in [13]. Their design is very close to the mathematical transformation as they resort to a paving of the Fourier plane into pseudo-polar wedges delimited by dyadic concentric squares. The so-called first generation discrete curvelet described in [24, 52] is a very different construction : it consists in applying the ridgelet transform to sub-images of a wavelet decomposition of the original image. By construction, the sub-images are well localized in space and frequency and the subsequent ridgelet transform provides the necessary directional sensitivity. This latter implementation in combination with the good geometric properties of the Healpix pixelisation scheme, inspired the digital curvelet transform on the sphere [56].

Partitioning using the Healpix representation. The Healpix representation [36] is a curvilinear partition of the sphere into quadrilateral pixels of exactly equal area but with varying shape. The base resolution divides the sphere into 12 quadrilateral faces of equal area placed on three rings around the poles and equator. Each face is subsequently divided into n_{side}^2 pixels following a hierarchical quadrilateral tree structure. The geometry of the Healpix sampling grid makes it easy to partition a spherical map into blocks of a specified size. We first extract the twelve base-resolution faces, and each face is then decomposed into smoothly (*i.e.* an apodizing window is used) overlapping blocks of side-length B_j pixels in such a way that the overlap between two vertically adjacent blocks is a rectangular array of size $B_j \times B_j/2$, as in the 2D digital curvelet transform [24, 52, 53]. With this scheme however, there is no overlapping between blocks belonging to different base-resolution faces. This may result in blocking artifacts in different data processing tasks such as image denoising *via* non linear filtering. A simple way around this difficulty, which however results in greater redundancy, is to work with various rotations of the data with respect to the sampling grid.

Ridgelet transform The continuous ridgelet transform is the application of a 1D-wavelet transform to the angular slices of the Radon transform [17]. A digital implementation described in [53] consists in i) first computing the 2D Fourier transform of the image, ii) next extracting lines going through the origin in the frequency plane, iii) then computing the 1D inverse Fourier transform of each of these lines, iv) and finally computing the 1D wavelet transform of the extracted lines. The first three steps implement the discrete Radon transform method called the *linogram*. Other implementations of the Radon transform, such as the *Slant Stack Radon Transform* [23], can be used as well, provided they offer

an exact reconstruction. Applying the ridgelet transform on sub-images as suggested in [15] provides the means to localize and limit the extent of the ridgelet functions. More details on the implementation of the digital ridgelet transform can be found in [52].

Curvelet transform Finally, the proposed digital curvelet transform algorithm on the sphere is as follows:

1. Apply the undecimated isotropic wavelet transform on the sphere with J scales,
2. Initialize the block size $B_1 = B_{min}$,
3. For $j = 1, \dots, J$ do,
 - partition the subband w_j with a block size B_j and apply the digital ridgelet transform to each block,
 - if j modulo 2 = 1 then $B_{j+1} = 2B_j$,
 - else $B_{j+1} = B_j$.

where smaller j means smaller length scales and higher frequency bands. Doubling the sidelength of the localizing window *at every other* dyadic sub-band (*i.e.* when the spatial frequency has been divided by four) approximately preserves the specific parabolic scaling property of the continuous curvelets. Nevertheless, the proportionality factor between *width* and *length*² has to be decided on in practice. The proposed implementation uses the default value $B_{min} = 16$ pixels. Figure 1 shows the backprojection of a few curvelet coefficients at different scales and orientations on the sphere. The digital curvelet transform on the sphere is clearly invertible in the sense that each step of the overall transform is itself invertible. The curvelet transform on the sphere has a redundancy factor of $16J + 1$ when J scales are used, which may be a problem for handling huge data sets such as from the future Planck-Surveyor experiment. This can be reduced by substituting the pyramidal wavelet transform to the undecimated wavelet transform in the above algorithm. More details on the wavelet, ridgelet, curvelet algorithms on the sphere can be found in [56] and software related to these new transforms is available from the web page <http://jstarck.free.fr>.

Local DCT on the Sphere The Healpix partitioning of the sphere into warped square blocks allows for even more numerical local transformations on the sphere. For instance, a local DCT on the sphere is readily obtained by applying the standard DCT to a complete set of possibly overlapping blocks of maybe different sizes. This heuristic revealed useful in a particular application described in section 3.3. Although it is difficult to fully characterize the geometry of such digital transformations on the sphere, they profitably enrich the set of tools available for spherical data analysis.

3 MCA on the Sphere

3.1 Principle and algorithm

For a given spherical map y modeled as a linear combination of K spherical maps $(s_k)_{k=1, \dots, K}$, $y = \sum_{k=1}^K s_k$, having different morphologies, MCA assumes that a dictionary of bases $\{\Phi_1, \dots, \Phi_K\}$ exists such that, for each k , s_k is sparse in Φ_k while its representation in the other $\Phi_{k'}$ ($k' \neq k$) is not sparse : $\forall k' \neq k, \|\Phi_k^T s_k\|_0 < \|\Phi_{k'}^T s_k\|_0$, where $\|x\|_0$ denotes the ℓ_0 pseudo-norm of the vector x . The problem




Figure 1: Curvelets on the Sphere - The above figures show the backprojection of various curvelet coefficients at different scales and orientations on the sphere. Each map is obtained by setting all but one of the curvelet coefficients to zero, and applying an inverse curvelet transform. Depending on the scale and the position of the non zero curvelet coefficient, the reconstructed image presents a feature with a given width, length and orientation.

is to separate the mixture y into its constitutive morphological components $(s_k)_{k=1, \dots, K}$ relying on the discriminating power of the different dictionaries Φ_k . Ideally, the α_k are the solutions to :

$$\min_{\alpha_1, \dots, \alpha_K} \sum_{k=1}^K \|\alpha_k\|_0 \quad \text{subject to} \quad y = \sum_{k=1}^K \Phi_k \alpha_k \quad (9)$$

While sensible from the point of view of the desired solution, the problem formulated in Equation (9) is non-convex and combinatorial by nature. Its complexity grows exponentially with the number of columns in the overall dictionary (NP-hard problem). Motivated by recent equivalence results *e.g.* in [25], the MCA algorithm seeks a solution to the following minimization problem:

$$\min_{s_1, \dots, s_K} \lambda \sum_{k=1}^K \|\alpha_k\|_1 + \left\| y - \sum_{k=1}^K s_k \right\|_2^2 \quad \text{with} \quad s_k = \Phi_k \alpha_k \quad (10)$$

where an ℓ_1 sparsity measure is substituted to the ℓ_0 counting norm following a prescription of the Basis Pursuit algorithm [20]. In the above, the equality constraint was relaxed and again $s_k = \Phi_k \alpha_k$. In the case where each Φ_k is an orthonormal basis, a block-coordinate solution to the above problem is given by the following set of coupled equations:

$$\forall k, s_k = r_k - \frac{\lambda_k}{2} \Phi_k \text{sign}(\Phi_k^T s_k) \quad \text{with} \quad r_k = y - \sum_{k' \neq k} s_{k'} \quad (11)$$

This can be solved efficiently using the iterative *Block-Coordinate Relaxation Method* [11] in conjunction with, at a given k , a soft-thresholding of the decomposition of r_k over Φ_k . However, when non-unitary

or redundant transforms are used, the above is no longer strictly valid. Nevertheless, simple shrinkage still gives satisfactory results as explained in [30]. Finally, denoting by \mathbf{T}_k and \mathbf{R}_k the forward and inverse transforms associated with the redundant dictionary Φ_k , MCA seeks a solution to problem (10) with the following algorithm:

1. Set the number of iterations I_{\max} and the initial thresholds $(\lambda_k^{(0)})_k$
2. While $\lambda_k^{(t)}$ is greater than a given lower bound λ_{\min} (e.g. can depend on the noise standard deviation),
 - Proceed with the following iteration to estimate components $(s_k)_{k=1,\dots,K}$ at iteration t :
 - For $k = 1, \dots, K$
 - Compute the residual term $r_k^{(t)}$ assuming the current estimates $\tilde{s}_{k' \neq k}^{(t-1)}$ of $s_{k' \neq k}$, are fixed:

$$r_k^{(t)} = y - \sum_{k' \neq k} \tilde{s}_{k'}^{(t-1)}$$
 - Estimate the current coefficients of $\tilde{s}_k^{(t)}$ by thresholding with threshold $\lambda_k^{(t)}$:

$$\tilde{\alpha}_k^{(t)} = \delta_{\lambda_k^{(t)}}(\mathbf{T}_k r_k^{(t)})$$
 - Get the new estimate of s_k by reconstructing from the selected coefficients $\tilde{\alpha}_k^{(t)}$:

$$\tilde{s}_k^{(t)} = \mathbf{R}_k \tilde{\alpha}_k^{(t)}$$
 - Decrease the thresholds λ_k following a given strategy.

3.2 Thresholding strategy

The operator δ in the above algorithm is a soft thresholding operator as a result of the use of an ℓ_1 sparsity measure in approximation to the ideal ℓ_0 norm. In practice, hard thresholding leads generally to better results [54]. The final threshold should vanish in the noiseless case or it may be set to a multiple of the noise standard deviation in the presence of noise as in common detection or denoising methods. The way the threshold is decreased along the iterations of the proposed iterative thresholding scheme is paramount in terms of performance of the MCA separation mechanism. The original algorithm [54] used a linear strategy :

$$\lambda^{(t)} = \lambda^{(0)} - (t - 1) \frac{\lambda^{(0)} - \lambda_{\min}}{I_{\max} - 1} \quad (12)$$

where $\lambda^{(0)}$ is the initial threshold, and I_{\max} is the number of iterations. The first threshold can be set automatically to a large enough value such as the maximum of all coefficients $\lambda^{(0)} = \max_k \|\mathbf{T}_k y\|_{\infty}$. But there is no way to estimate the minimum number of iterations yielding a successful separation. Too small a number of iterations leads to bad separation while too large a number is computationally costly. Further, experiments have clearly shown that the optimal number of iterations depends on the data. We recently focused on devising some new data adaptive thresholding strategies to speed up the MCA decomposition preserving the quality of the component separation. Hereafter we describe two promising strategies, namely MAD and MOM, in the case where $K = 2$; generalizing to $K \geq 2$ is straightforward.

MAD Consider a map y such that $y = s_1 + s_2 = \Phi_1 \alpha_1 + \Phi_2 \alpha_2$ where s_1 and s_2 have similar ℓ_2 norm and $\alpha_{k=1,2} = \Phi_{k=1,2}^T s_{k=1,2}$ are sparse. When both $\Phi_{k=1,2}$ are orthonormal bases, decomposing y in Φ_1 leads to $y \Phi_1^T = \alpha_1 + \Phi_1^T \Phi_2 \alpha_2$. Provided the mutual coherence [12, 37, 25] of Φ_1 and Φ_2 is low, y_2 has no particular structure in Φ_1 and hence it is tempting to model $\Phi_1^T s_2$ as a Gaussian noise. Its standard deviation can be estimated using a robust estimator such as the Median Absolute Deviation (MAD) [27]. It follows that estimating the significant entries $\tilde{\alpha}_1$ in α_1 is a denoising problem readily solved by thresholding $\Phi_1^T y$ with a threshold $k\sigma$ (typically k is in the range 3 to 4). The next step is to project the residual $y - \tilde{s}_1 = y - \Phi_1 \tilde{\alpha}_1$ on Φ_2 and so on. Clearly, the variance of the residual decreases

along iterations and so this provides a simple strategy to adaptively control the threshold in the MCA algorithm. In practice, this strategy remains fruitful in the case of redundant dictionaries. Donoho *et al.* in [28] have recently focused on an iterative thresholding scheme applied to solving under-determined linear sparse problems in which they use a similar rule to manage their decreasing threshold.

MOM Let $\tilde{s}_1^{(t)}$ and $\tilde{s}_2^{(t)}$ denote the current estimates of components s_1 and s_2 at the t^{th} iteration of the MCA decomposition of y . The current residual is $r^{(t)} = y - \tilde{s}_1^{(t)} - \tilde{s}_2^{(t)}$. In the strategy coined MOM as in "Mean of Max", the value of the threshold at iteration t is given by :

$$\lambda^{(t)} = \frac{1}{2} \left[\|\Phi_1^T (y - \tilde{s}_1^{(t-1)} - \tilde{s}_2^{(t-1)})\|_\infty + \|\Phi_2^T (y - \tilde{s}_1^{(t-1)} - \tilde{s}_2^{(t-1)})\|_\infty \right] \quad (13)$$

which is easily computed at each step of the iterative process. When one considers more than two dictionaries, one should take the mean of the two largest decomposition coefficients of the full residual over two distinct dictionaries. The intuition underlying this strategy is that the next significant coefficients to be selected should be attached to the dictionary in which the projection of the full residual has coefficients of largest amplitudes. Assuming the coefficients selected at iteration t are in Φ_1 , it can be shown, under some conditions on the sparsity of the components and the mutual coherence of the dictionary [8], that the proposed strategy fixes the threshold so that :

$$\|\Phi_1^T \Phi_2 \bar{\alpha}_2^{(t-1)}\|_\infty < \lambda_1^{(t)} < \|\bar{\alpha}_1^{(t-1)}\|_\infty, \bar{\alpha}_{k=1,2}^{(t-1)} = \alpha_{k=1,2} - \tilde{\alpha}_{k=1,2}^{(t-1)} \quad (14)$$

hence avoiding false detections (upper bound) and ensuring that at least one coefficient is selected (lower bound). This thresholding strategy can easily be made more or less conservative depending on the desired decomposition speed. With these new thresholding strategies, MCA is a fast and robust algorithm to achieve sparse decompositions in redundant dictionaries and a practical alternative to other well-known sparse decomposition algorithms [8].

Example

The spherical maps shown on figure 2 illustrate a simple numerical experiment. We applied the proposed Morphological Component Analysis on the Sphere to synthetic data resulting from the linear mixture of components respectively sparse in the spherical harmonics and the isotropic wavelet representations. The method was able to separate the data back into its original constituents. A more involved application is described in the next section.

3.3 Application in Physics

In Inertial Confinement Fusion (ICF) a spherical shell is irradiated by laser energy directly or after the laser energy has been converted to soft X-rays [3]. Either way, the aim is to implode the capsule which contains a shell of nuclear fusion fuel (deuterium and tritium) ready to ignite if, after it has been imploded, its density is high enough and a hot spot in its center becomes hot enough to cause a propagating nuclear burn wave to travel through the rest of the fuel. This ultimate energy source will not work if during the implosion hydrodynamic instabilities develop which can break apart the shell before it assembles at the center and a hot spot forms [41]. Hydrodynamic instabilities such as Rayleigh-Taylor occur due to nonuniformities in the laser spatial profile or imperfections in the composition of multiple surfaces which make up the layers of thin material that surround the nuclear fuel. Very small amplitude imperfections




Figure 2: Simple toy experiment with MCA on the sphere - The top map shows a linear combination of a spherical harmonic function and a localized Gaussian-like function on the sphere. The bottom maps show the resulting separated components that were obtained using the proposed Morphological Component Analysis on the sphere.




Figure 3: **left** : Surface structures of ICF spherical shells measured on the nanometer scale are a superposition of global scale variations, isolated bumps and scratches as well as artifacts which look like interference patterns on intermediate scales. **right** : Coarsest scale of the undecimated isotropic wavelet transform of the surface measurements of an ICF target.

initially can result in the ultimate failure of the target due to the large compression ratios involved in ICF. It is therefore extremely important to characterize the inner and outer surfaces of ICF shell targets so as to know whether they are worthy of consideration for ICF implosions. One day in a reactor setting tens of thousands of targets will have to be imploded daily so that checking each one is totally out of the question. Instead, very good target fabrication quality control processes have to be adopted so that confidence levels in proper performance will be high. A major step along this path to fusion energy then is to understand why imperfections occur and to correct the systematic elements and control the harm done by random sources.

Fine structures on the surfaces of spherical shells can be measured on the nanometer scale, among others, by atomic force microscopy or phase shifting spherical diffractive optical interferometry. An example of such measurements is shown on figure 3. As can be seen from the figure, there appears to be a superposition of global scale variations, isolated bumps and scratches as well as artifacts which look like interference patterns on intermediate scales of localization. The latter must be isolated and eliminated from consideration when deciding the readiness of the target for implosion. We have achieved the morphological feature separation by first doing an isotropic wavelet transform on the spherical data and subtracting the coarsest scale information. MCA on the sphere was used on the rest of the image using the undecimated wavelet and the local cosine transforms on the sphere. The isolated bumps were thus identified and the measurement technique caused artifacts were removed easily. The resulting bumps added to the coarsest scale, is the clean data with the interference patterns and artifacts removed as shown in figure 4. The spherical harmonic decomposition of the cleaned image gives rise to coefficients of various ℓ modes which will be amplified by the implosion process which can now be assessed correctly using numerical hydrodynamics simulation generated growth factors. If the bumps are clustered and not randomly distributed, then systematic errors in the manufacturing process can be tracked down. A code called MODEM has been put together to study such target surface data and extract the localized bump statistics including their correlations in height, size and relative location. For more details see [1].

4 Inpainting on the Sphere

4.1 Algorithm

Named after the expert recovery process used for the restoration of deteriorated masterpieces, inpainting refers to a set of techniques used to alter images in a way that is undetectable to people who are unaware of the original images. There are numerous applications among which removing scratches or objects in digitized photographs, removing overlaid text or graphics, filling-in missing blocks in unreliably transmitted images, predicting values in images for better compression or image upsampling. Inpainting algorithms strive to interpolate through the gaps in the image relying on the available pixels, the continuation of edges, the periodicity of textures, etc. The preservation of edges and texture, in other words discontinuities, across gaps has attracted much interest, and many contributions have been proposed to solve this interpolation task. Non-texture image inpainting has received considerable interest and excitement since the pioneering paper by Masnou and Morel [43, 44] who proposed variational principles for image disocclusion. A recent wave of interest in inpainting has started from the recent contributions of Sapiro *et al.* [4, 5, 6], followed by Chan and Shen [19]. In these works, authors point to the importance of geometry and design anisotropic diffusion PDEs to fill in gaps by smooth continuation of isophotes. PDE methods have been shown to perform well on piecewise smooth functions.

A very different approach is the inpainting algorithm based on MCA described in [32] which has proved capable of filling in holes in either texture or cartoon content in 2D images. To make the link




Figure 4: **top** : Spherical map obtained by subtracting the coarse scale map on the right of figure 3 from the initial map on the left of figure 3. **bottom** : Component maps separated by the MCA method on the sphere : interference patterns and measurement artifacts were grabbed by the local cosine functions on the sphere (left) while the isolated bumps were caught using the undecimated wavelet on the sphere (right). Adding back the coarse scale on the right of figure 3 to the latter map results in a clean map of the surface structures of an ICF spherical shell with the interference patterns and artifacts removed.

between building sparse representations and inpainting, consider the effect of a rectangular gap on the set of Fourier coefficients of a monochromatic sinewave : because of the non-locality of the Fourier basis functions it takes a large number of coefficients to account for the gap, which is known as the Gibbs effect. Seeking a sparse representation of the incomplete sine-wave outside the gap, that is without fitting the gap, enables the recovery of the complete monochromatic sinewave.

Following [32], an inpainting algorithm on the sphere is readily built from the Morphological Component Analysis on the sphere described in the previous section. Consider a discrete spherical data map y and a binary map M such that ones in M indicate that the corresponding pixels in y are valid data while zeros indicate invalid data. The objective function of MCA (eq. 10) can be modified as follows :

$$\min_{s_1, \dots, s_n} \lambda \sum_{k=1}^K \|\alpha_k\|_1 + \left\| M \odot \left(y - \sum_{k=1}^K s_k \right) \right\|_2^2 \quad \text{with } s_k = \Phi_k \alpha_k. \quad (15)$$

where \odot stands for entry-wise multiplication. Thus we are preventing the sparse model under construction from attempting to fit the invalid data. Other constraints can be easily imposed on the interpolated sparse components. For instance, in [32], a total variation penalty is shown to enhance the recovery of piece-wise smooth components. Asking for the regularity across the gaps of some localized statistics (*e.g.* enforcing that the empirical variance of a given inpainted sparse component be *nearly equal* outside and inside the masked areas) are other possible constraints. In practice, because of the lack of accuracy of some digital transformations we used in the spherical topology, additional constraints, which may be relaxed close to convergence, were also found useful in some cases to stabilize the described iterative algorithms.

It is proposed that a solution to the above minimization problem can be reached using the same iterative thresholding process as in the MCA algorithm detailed in the previous section, with the only required modification consisting in *masking* the full residual using M after each residual estimation. The MCA-inpainting algorithm is as follows :

1. Set the number of iterations I_{\max} and the initial thresholds $\lambda^{(0)}$
2. While $\lambda_k^{(t)}$ is greater than a given lower bound λ_{\min} (e.g. can depend on the noise standard deviation),
 - Proceed with the following iteration to estimate components $(s_k)_{k=1, \dots, K}$ at iteration t :
 - For $k = 1, \dots, K$
 - Compute the residual term $r^{(t)}$:

$$r^{(t)} = y - \sum_k \tilde{s}_k^{(t-1)}$$
 - Estimate the current coefficients of $\tilde{s}_k^{(t)}$ by thresholding with threshold $\lambda_k^{(t)}$:

$$\tilde{\alpha}_k^{(t)} = \delta_{\lambda_k^{(t)}} \left(\mathbf{T}_k \left(M \odot r^{(t)} + \tilde{s}_k^{(t-1)} \right) \right)$$
 - Get the new estimate of s_k by reconstructing from the selected coefficients $\tilde{\alpha}_k^{(t)}$:

$$\tilde{s}_k^{(t)} = \mathbf{R}_k \tilde{\alpha}_k^{(t)}$$
 - Decrease the thresholds λ_k following a given strategy.

The different thresholding strategies described in the previous section can be used in the proposed MCA inpainting iterative thresholding algorithm.

Example

A simple numerical experiment is shown on figure 5. Starting with a full satellite view of the Earth², an incomplete spherical map was obtained by randomly masking some of the pixels. In fact, as much as

²availble from : http://www.nasa.gov/vision/earth/features/bmng_gallery_4.html




Figure 5: Application of the proposed MCA-inpainting algorithm on the sphere. **top left** : original satellite view of the Earth ($mean = 76.9, \sigma = 47.7$). **top right** : incomplete map retaining 40 percent of the original pixels. **bottom left** : inpainted map. **bottom right** : map of reconstruction errors ($mean = 0.0, \sigma = 2.86$ empirically estimated from the reconstructed pixels only).

sixty percent of the pixels were masked. Using both the spherical harmonics transform and the curvelet transform on the sphere within the proposed MCA inpainting algorithm, it is possible to fill in the missing pixels in a visually undetectable way. The residual map is shown at the bottom right of figure 5.

4.2 Application in Astrophysics

A major issue in modern cosmology is the measurement and the statistical characterization (spatial power spectrum, Gaussianity) of the slight fluctuations in the Cosmic Microwave Background radiation field. These are indeed strongly related to the cosmological scenarios describing the properties and evolution of our Universe. Some 370 000 years after the 'Big Bang', when the temperature of the Universe was around 3000 K, thermal energy was no longer sufficient to keep electrons and positively charged particles apart so they combined. Photons were then set free in a nearly transparent Universe. Since the Universe further expanded, these photons are now in the microwave range but they should still be distributed according to a Black Body emission law. Indeed, before recombination, the Universe was a highly homogeneous opaque plasma in near thermal equilibrium in which photons and charged particles were highly interacting. Hence the slight fluctuations in matter density from which such large scale structures as galaxies or clusters of galaxies have evolved, are also imprinted on the distribution of photons.

The Cosmic Microwave Background (CMB) was first observed in 1965 by Penzias and Wilson confirming a prediction made by Gamow in the late 1940's. But it was not until the early 1990's that evidence for small fluctuations in the CMB sky could finally be found thanks to the observations made by COBE [50]. This was confirmed by several subsequent observations and recently by NASA's Wilkin-




Figure 6: **left** : CMB data map provided by the WMAP team. Areas of significant foreground contamination in the galactic region and at the locations of strong radio point sources have been masked out. **right** : Map obtained by applying the proposed MCA-inpainting algorithm on the sphere to the former incomplete WMAP CMB data map.

son Microwave Anisotropy Probe³. Full-sky multi-spectral observations with unprecedented sensitivity and angular resolution are expected from the ESA’s PLANCK⁴ mission, which is to be launched in 2008. The statistical analysis of this data set will help set tighter bounds on major cosmological parameters.

There are nonetheless a few practical issues and notably that several other astrophysical sources also emit radiation in the frequency range used for CMB observations [10]. Separating back the observed mixtures into maps of the different astrophysical contributions in order to isolate the CMB properly is a difficult inverse problem for which methods and algorithms are being actively designed (see *e.g.* [47, 7, 46, 48] and references therein). The estimated spherical CMB maps will inevitably be contaminated by some level of residual contributions, most significantly in the galactic region and at the locations of strong radio point sources. Therefore, it is common practice to mask out that part of the data (*e.g.* using the mask shown on figure 7 upper left, provided by the WMAP team) in order to reliably assess the non-gaussianity of the CMB field through estimated higher order statistics (*e.g.* skewness, kurtosis) in various representations (*e.g.* *wavelet*, *curvelet*, *etc.*) [51, 40]. But the gaps in the data thus created need to be handled properly as the detection of non-gaussianity in CMB would have a major scientific impact.

The proposed MCA-inpainting on the sphere was used here successfully to fill in the masked regions in order to restore the stationarity of the observed CMB field and lower the impact of the incompleteness of the data set on the estimated measures of non-gaussianity or any other non-local statistical test. The experiment was conducted on several simulations of full-sky Gaussian CMB maps. A typical CMB map (the CMB data map disclosed by the WMAP consortium) is shown on figure 6 along with the map obtained as a result of the inpainting process allowing for a first visual assessment of the quality of the proposed method. Figure 7 shows the wavelet decomposition of the inpainted map. We can see that the mask is not visible at all in the different scales. Here we have applied the MCA-Inpainting algorithm with 200 iterations and a single transform which was the Spherical Harmonic Decomposition. A more quantitative evaluation of the proposed inpainting algorithm is reported on figure 8 where plots of the estimated measures of non-Gaussianity on both the original map and the inpainted map are given. These reveal no significant discrepancy: we believe that the proposed method will help discriminate between truly non-Gaussian CMB and non-Gaussianity related to the non-stationarity of incomplete maps. This will be further investigated in the future.

³The WMAP data and mask we used here are available online at <http://map.gsfc.nasa.gov/>

⁴<http://astro.estec.esa.nl/Planck>




Figure 7: **left** : Mask provided by the WMAP team. The dark blue pixels indicate areas of high level foreground contamination in the WMAP CMB data map. **From top to bottom and left to right** : Maps of the wavelet decomposition on seven scales of the inpainted WMAP CMB map shown on the right of figure 6. From the visual point of view, the masked area cannot be distinguished anymore in the wavelet scales of the inpainted map.




Figure 8: Horizontally is the scale number increasing for lower frequencies. **Left** : skewness of the wavelet coefficients in a given scale of the original complete simulated spherical CMB map (\times) and of the inpainted map (\diamond). **Right** : kurtosis of the wavelet coefficients in a given scale of the original complete simulated spherical CMB map (\times) and of the inpainted map (\diamond). Error bars were estimated on a small set of fifteen simulated complete CMB maps.

5 Conclusion

This paper presented an extension of the Morphological Component Analysis method to the case of spherical maps. This was made possible thanks to the wealth of multiscale analysis tools and discrete transforms newly made available for the representation, analysis and synthesis of data on the sphere. An inpainting algorithm on the sphere was also described building on the ideas of MCA. The difficulties in porting MCA and the related inpainting algorithm to the spherical topology are less theoretical than practical. Indeed, the proposed algorithms are iterative and some of the digital transforms we used on the sphere were insufficiently accurate. As quickly mentioned, additional constraints were sometimes necessary to stabilize the convergence of the algorithms. However, lingering over such practicalities is not the purpose of this paper especially since they did not deeply affect the algorithms. We reported on a few applications of the proposed methods to challenging data analysis problems in physics and astrophysics. We expect these tools to be valuable in many other applications such as in areas where very smooth and precise motion is required (*e.g.* moving chips on a conveyor belt in an ultra clean room in the semiconductor industry) : hundred micron scale balls or ball bearings are used which must not have any imperfections on their surfaces above the nano-scale. Characterizing such tiny spherical objects with that kind of precision requires optical or X ray techniques which inherently produce artifacts. But since the morphology of the artifacts (interference pattern like repeated ring or arrays of light and dark regions) is far different than the bumps and scratches one is trying to avoid, they can be isolated using MCA techniques on the sphere and thus help assess the manufacturing process and make the changes required to meet specifications. Similar considerations are at play when treating ICF targets for laser fusion. The results reported in this paper allow us to expect that the described extensions of MCA and inpainting to the sphere will bear much fruit in the study of CMB, non-Gaussianity, and related matters such as astrophysical component separation [7] in full-sky multichannel observations of the celestial sphere in the microwave range.

Acknowledgments

Some of the results in this paper have been derived using the HEALPix (Górski, Hivon, and Wandelt 1999).

References

- [1] B. Afeyan, K. Won, J. L. Starck, R. Stevens, E. Mapoles, M. Johnson, and S. Haan. Modem: Morphological diversity extraction method to identify, classify and characterize surface defects of spherical icf targets via afm and phase shifting diffraction interferometry. In *Proceedings of the 17th Target Fabrication Meeting, San Diego, CA*, pages 1–5, oct 2006.
- [2] J.-P. Antoine and P. Vandergheynst. Wavelets on the 2-sphere: a group-theoretical approach. *Appl. Comput. Harmon. Anal.*, 7(3):262–291, 1999.
- [3] S. Atzeni and J. Meyer ter Vehn. *The Physics of Inertial Fusion*. Oxford University Press, 2004.
- [4] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera. Filling-in by joint interpolation of vector fields and grey levels. *IEEE Trans. Image Processing*, 10:1200–1211, August 2001.
- [5] M. Bertalmio, A.L. Bertozzi, and G. Sapiro. Navier-stokes, fluid dynamics, and image and video inpainting. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, 2001.
- [6] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. *Comput. Graph. (SIGGRAPH 2000)*, pages 417–424, July 2000.
- [7] J. Bobin, Y. Moudden, J.-L. Starck, and M. Elad. Morphological diversity and source separation. *IEEE Signal Processing Letters*, 13(7):409–412, july 2006.
- [8] J. Bobin, J.-L. Starck, J. Fadili, Y. Moudden, and D.L. Donoho. Morphological component analysis: An adaptive thresholding strategy. *IEEE Transactions on Image Processing*, 2007. to appear.
- [9] I. Bogdanova, P. Vandergheynst, J.-P. Antoine, L. Jacques, and M. Mrovidone. Stereographic wavelet frames on the sphere. *Applied and Computational Harmonic Analysis*, 19(2):223–252, 2005.
- [10] F. R. Bouchet and R. Gispert. Foregrounds and CMB experiments I. Semi-analytical estimates of contamination. *New Astronomy*, 4:443–479, September 1999.
- [11] A.G. Bruce, S. Sardy, and P. Tseng. Block coordinate relaxation methods for nonparametric signal de-noising. *Proceedings of the SPIE - The International Society for Optical Engineering*, 3391:75–86, 1998.
- [12] A.M. Bruckstein and M. Elad. A generalized uncertainty principle and sparse representation in pairs of \mathbf{r}^n bases. *IEEE Transactions on Information Theory*, 48:2558–2567, 2002.
- [13] E. Candes, L. Demanet, D. Donoho, and L. Ying. Fast discrete curvelet transforms. *SIAM Multiscale Modeling and Simulation*, 2006. to appear.
- [14] E. Candès and D. Donoho. Curvelets and curvilinear integrals. *J. Approx. Theory*, 113(1):59–90, 2001.

- [15] E. J. Candès. Harmonic analysis of neural networks. *Applied and Computational Harmonic Analysis*, 6:197–218, 1999.
- [16] E. J. Candès and D. L. Donoho. Curvelets – a surprisingly effective nonadaptive representation for objects with edges. In A. Cohen, C. Rabut, and L.L. Schumaker, editors, *Curve and Surface Fitting: Saint-Malo 1999*, Nashville, TN, 1999. Vanderbilt University Press.
- [17] E.J. Candès and D. Donoho. Ridgelets: the key to high dimensional intermittency? *Philosophical Transactions of the Royal Society of London A*, 357:2495–2509, 1999.
- [18] L. Cayón, J. L. Sanz, E. Martínez-González, A. J. Banday, F. Argüeso, J. E. Gallegos, K. M. Górski, and G. Hinshaw. Spherical Mexican hat wavelet: an application to detect non-Gaussianity in the COBE-DMR maps. *Monthly Notices of the Royal Astronomical Society*, 326:1243–1248, October 2001.
- [19] T. Chan and J. Shen. Local inpainting models and tv inpainting. *SIAM J. Appl. Math.*, 62:1019–1043, July 2001.
- [20] S.S. Chen, D.L. Donoho, and M.A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20:33–61, 1998.
- [21] I. Daubechies. Time-frequency localization operators: A geometric phase space approach. *IEEE Transactions on Information Theory*, 34:605–612, 1988.
- [22] M.N. Do and M. Vetterli. The contourlet transform: an efficient directional multiresolution image representation. *Processing, IEEE Transactions on Image Processing*, 14(12):2091 – 2106, 2005.
- [23] D. Donoho and A.G. Flesia. Digital Ridgelet Transform Based on True Ridge Functions. In J. Schmeidler and G.V. Welland, editors, *Beyond Wavelets*. Academic Press, 2002.
- [24] D.L. Donoho and M.R. Duncan. Digital curvelet transform: strategy, implementation and experiments. In H.H. Szu, M. Vetterli, W. Campbell, and J.R. Buss, editors, *Proc. Aerosense 2000, Wavelet Applications VII*, volume 4056, pages 12–29. SPIE, 2000.
- [25] D.L. Donoho and M. Elad. Optimally sparse representation in general (non-orthogonal) dictionaries via ℓ^1 minimization. *Proc. Nat. Aca. Sci.*, 100:2197–2202, 2003.
- [26] D.L. Donoho, M. Elad, and V. Temlyakov. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Trans. On Information Theory*, 52:6–18, january 2006.
- [27] D.L. Donoho and I.M. Johnstone. Ideal spatial adaptation via wavelet shrinkage. *Biometrika*, 81:425–455, 1994.
- [28] D.L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck. Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit. *IEEE Trans. on Information Theory*, 2006. submitted.
- [29] A. G. Doroshkevich, P. D. Naselsky, O. V. Verkhodanov, D. I. Novikov, V. I. Turchaninov, I. D. Novikov, P. R. Christensen, and L.-Y. Chiang. Gauss-legendre sky pixelization (glesp) scheme for cmb maps. *International Journal of Modern Physics D*, 14(2):275–290, 2005. also available at <http://arXiv.org/abs/astro-ph/0305537>.

- [30] M. Elad. Why simple shrinkage is still relevant for redundant representations? *IEEE Transactions on Information Theory*, 52(12):5559–5569, 2006.
- [31] M. Elad and A.M. Bruckstein. A generalized uncertainty principle and sparse representation in pairs of bases. *IEEE Transactions on Information Theory*, 48:2558–2567, 2002.
- [32] M. Elad, J.-L. Starck, P. Querre, and D. Donoho. Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA). *Applied and Computational Harmonic Analysis*, 19(3):340–358, 2005. to appear.
- [33] W. Freeden and T. Maier. On multiscale denoising of spherical functions: Basic theory and numerical aspects. *Electronic Transactions on Numerical Analysis (ETNA)*, 14:40–62, 2002.
- [34] W. Freeden, T. Maier, and S. Zimmermann. A survey on wavelet methods for (geo)applications. *Revista Mathematica Complutense*, 16(1):277–310, 2003.
- [35] J.J. Fuchs. Recovery of exact sparse representations in the presence of bounded noise. *IEEE Trans. On Information Theory*, 51:3601–3608, october 2005.
- [36] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *Astrophysical Journal*, 622:759–771, April 2005.
- [37] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *IEEE Transactions on Information Theory*, 49(12):3320–3325, 2003.
- [38] F. Herrmann, P.P. Moghaddam, and C. Stolk. The contourlet transform: an efficient directional multiresolution image representation. *Applied and Computational Harmonic Analysis*, 2005. submitted.
- [39] M. Holschneider. Wavelet analysis on the sphere. *J. Math. Phys.*, 37(8):4156–4165, 1996.
- [40] J. Jin, J.-L. Starck, D.L. Donoho, N. Aghanim, and O. Forni. Cosmological non-gaussian signatures detection: Comparing performance of different statistical tests. *EURASIP Journal on Applied Signal Processing*, 2005(15):2470–2485, August 2005.
- [41] J.D. Lindl. *Inertial Confinement Fusion: The Quest for Ignition and Energy Gain Using Indirect Drive*. AIP Press, 1997.
- [42] S. Mallat and Z. Zhang. Matching pursuit with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, december 1993.
- [43] S. Masnou and J.M. Morel. Level lines based disocclusion. In *IEEE International Conference on Image Processing*, volume III, pages 259–263, 1998.
- [44] S. Masnou and J.M. Morel. Disocclusion: A variational approach using level lines. *IEEE Trans. Image Process.*, 11(2):68–76, 2002.
- [45] J. McEwen, M. Hobson, A. Lasenby, and D. Mortlock. A fast directional continuous spherical wavelet transform. *MNRAS*, 2004. submitted.

- [46] Y. Moudden, J.-F. Cardoso, J.-L. Starck, and J. Delabrouille. Blind component separation in wavelet space: Application to CMB analysis. *EURASIP Journal on Applied Signal Processing*, 2005(15):2437–2454, August 2005.
- [47] G. Patanchon, J.-F. Cardoso, J. Delabrouille, and P. Vielva. Cosmic microwave background and foregrounds in Wilkinson Microwave Anisotropy Probe first-year data. *MNRAS*, 364:1185–1194, December 2005.
- [48] S. Pires, J. B. Juin, D. Yvon, Y. Moudden, S. Anthoine, and E. Pierpaoli. Sunyaev-Zel’dovich cluster reconstruction in multiband bolometer camera surveys. *AA*, 455:741–755, August 2006.
- [49] P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics Proceedings (SIGGRAPH 95)*, pages 161–172, 1995.
- [50] G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson. Structure in the COBE differential microwave radiometer first-year maps. *Astrophysical Journal Letters*, 396:L1–L5, September 1992.
- [51] J.-L. Starck, N. Aghanim, and O. Forni. Detecting cosmological non-gaussian signatures by multi-scale methods. *Astronomy and Astrophysics*, 416:9–17, 2004.
- [52] J.-L. Starck, E. Candès, and D.L. Donoho. The curvelet transform for image denoising. *IEEE Transactions on Image Processing*, 11(6):131–141, 2002.
- [53] J.-L. Starck, E. Candès, and D.L. Donoho. Astronomical image representation by the curvelet transform. *Astronomy and Astrophysics*, 398:785–800, 2003.
- [54] J.-L. Starck, M. Elad, and D.L. Donoho. Redundant multiscale transforms and their application for morphological component analysis. *Advances in Imaging and Electron Physics*, 132, 2004.
- [55] J.-L. Starck, J. Fadili, and F. Murtagh. The undecimated wavelet decomposition and its reconstruction. *IEEE Transactions on Image Processing*, 16(2):297–309, 2007.
- [56] J.-L. Starck, Y. Moudden, P. Abrial, and M. Nguyen. Wavelets, ridgelets and curvelets on the sphere. *Astronomy and Astrophysics*, 446(3):1191–1204, February 2006.
- [57] J.-L. Starck, F. Murtagh, and A. Bijaoui. *Image Processing and Data Analysis: The Multiscale Approach*. Cambridge University Press, 1998.
- [58] J.-L. Starck, F. Murtagh, E. Candes, and D.L. Donoho. Gray and color image contrast enhancement by the curvelet transform. *IEEE Transactions on Image Processing*, 12(6):706–717, 2003.
- [59] Y. Wiaux, L. Jacques, and P. Vandergheynst. Correspondence Principle between Spherical and Euclidean Wavelets. *Astrophysical Journal*, 632:15–28, October 2005.
- [60] Y. Wiaux, L. Jacques, and P. Vandergheynst. Fast directional correlation on the sphere with steerable filters. *Astrophys. J.*, 652:820, 2006.