



HAL
open science

A Review of Bandlet Methods for Geometrical Image Representation

Stéphane Mallat, Gabriel Peyré

► **To cite this version:**

Stéphane Mallat, Gabriel Peyré. A Review of Bandlet Methods for Geometrical Image Representation. Numerical Algorithms, 2007, 44 (3), pp.205-234. hal-00359744

HAL Id: hal-00359744

<https://hal.science/hal-00359744v1>

Submitted on 9 Feb 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Review of Bandlet Methods for Geometrical Image Representation

Stéphane Mallat
CMAP, École Polytechnique
91128 Palaiseau Cedex, France

Gabriel Peyré
CEREMADE, Université Paris Dauphine.
75775 Paris Cedex, France

February 26, 2008

Abstract

This article reviews bandlet approaches to geometric image representations. Orthogonal bandlets using an adaptive segmentation and a local geometric flow are well suited to capture the anisotropic regularity of edge structures. They are constructed with a “bandletization” which is a local orthogonal transformation applied to wavelet coefficients. The approximation in these bandlet bases exhibits an asymptotically optimal decay for images that are regular outside a set of regular edges. These bandlets can be used to perform image compression and noise removal. More flexible orthogonal bandlets with less vanishing moments are constructed with orthogonal grouplets that group wavelet coefficients along a multiscale association field. Applying a translation invariant grouplet transform over a translation invariant wavelet frame leads to state of the art results for image denoising and super-resolution.

1 Geometry of Images and Textures

Taking advantage of geometrical structures in natural images is crucial to improve the state of the art in image processing. But geometry is also the bottleneck in other scientific areas and similar ideas emerge in various fields such as turbulence in fluid dynamics or visual coding in the cortex. At a first glance, geometry might seem restricted to a well defined set of curves along which the image is singular. Figure 1 (a) shows an example of such a simple geometric image where the relevant information is only carried along a set of edges. Natural images are however much more complex than cartoon images such as the one depicted on figure 1 (b). They carry a textural content that is neither pure noise nor regular edge curves, see figure 1 (c). Natural phenomena such as seismic, wood growth or fluids dynamic is often responsible for the emergence of this textural content. They lead to turbulent dynamics that creates most of the complexity of these geometric textures, see figure 1 (d-f).

Geometric structures exist in a lot of signal modalities and carry most of the perceptual information. The motion of objects in a movie is described using an optical flow that follows the 3D geometry of the signal, see figure 2 (a-b). Natural sounds also exhibit geometric patterns in the time-frequency plane where evolving harmonics follows geometric paths, see for instance a bird sing spectrogram, figure 2 (c). All these geometric cues are essential for human perception and should be exploited by modern signal processing methods.

From a mathematical point of view, classical tools from differential geometry can characterize the geometry of contours when the edge curves are well defined. However, for




Figure 1: *Examples of images with varying geometric complexity. (a): simple geometric image. (b): cartoon image. (c): natural image. (d): seismic image. (e): wood texture. (f): vorticity field of a fluid.*




Figure 2: *(a): Sample image of a movie. (b) Corresponding optical flow. (c): Spectrogram of a bird's sing.*

natural images with a varying blurring and turbulent textures, the local description of geometric regularity is ill-posed and cannot lead to robust and efficient algorithms. Thanks to the wavelet transform, harmonic analysis brings a first answer to the representation of the regular parts of images and texture patterns. This is the reason why orthogonal wavelets bases are at the heart of JPEG2000, the latest image compression standard. Wavelets are however sub-optimal to compress the geometrically regular part of images as explained in section 2.

Bridging the gap between geometric representations and harmonic analysis is a major issue in image processing. A compact representation of geometric structures would have applications for traditional image processing tasks such as inverse problems or compression, but would also ease learning algorithms in computer vision. Compressing with minimum loss the geometry of images is at the heart of industrial problems. Satellite imaging requires compression of urban geometric patterns with increasing resolution and medical imaging requires a fine rendering of vessels and other tubular structures. High-definition numeric video requires the upsampling of movies where the 3D geometry is crucial to solve the aliasing problem.

Section 2 studies the wavelet representation and explains its inefficiencies on geometric images. Both the finite elements and the curvelets schemes, exposed in sections 3.1 and 3.2 enjoy a better approximation rate than orthogonal wavelets on geometrical images. The orthogonal bandlet approximation scheme is explained in section 4. Bandlet bases are obtained through a hierarchical cascade of orthogonal elementary operators. The geometry parameterizes these operators to adapt the representation to the local anisotropy of geometric images. This hierarchical cascade leads to fast algorithms that compute the decomposition of an image in an adapted bandlet basis. The optimality of the adapted bandlet representation is proved for the approximation, compression, and estimation of geometrically regular images.

An orthogonal bandlet basis has a constrained geometry that is a source of inefficiency to capture the turbulent geometry of natural textures. This issue is solved with a multiscale association field that drives an adapted grouplet transform, presented in section 5. This transform can be implemented in a translation-invariant manner and over coefficients of a multiscale wavelet transform. This leads to a tight frame of adapted grouping-bandlets. This tight frame gives state of the art results in denoising and super-resolution applications and can be used in computer graphics applications such as texture synthesis and image inpainting.

2 Image Representation in a Wavelet Basis

Decomposing a function in an orthogonal basis allows to define a sparse representation using a simple wavelet thresholding. In particular, orthogonal wavelet bases define optimal approximations for classes of piecewise regular functions. In this section we review the main properties and limitations of wavelets to approximate geometrical singularities.

The best approximation f_M of a function f with M coefficients in an orthogonal basis $B = \{g_\mu\}_\mu$ is computed using the largest M coefficients above some threshold T :

$$f_M \stackrel{\text{def.}}{=} \sum_{|\langle f, g_\mu \rangle| > T} \langle f, g_\mu \rangle g_\mu \quad \text{with} \quad M \stackrel{\text{def.}}{=} \text{Card} \{ \mu \mid |\langle f, g_\mu \rangle| > T \}, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ is the inner product. The approximation error is then:

$$\|f - f_M\|^2 = \sum_{|\langle f, g_\mu \rangle| \leq T} |\langle f, g_\mu \rangle|^2.$$

A signal model defines a set Θ such that $f \in \Theta$. Optimizing the representation is then equivalent to maximizing the decay of the error $\|f - f_M\|^2$ when M increases, for all $f \in \Theta$. Asymptotically, one looks for bases B such that $\|f - f_M\|^2 = O(M^{-\beta})$ for the largest possible β .

If some basis B reaches an optimal error decay on Θ , one can prove that a compression algorithm (resp. a denoising algorithm) that quantizes (resp. thresholds) the coefficients in this basis is optimal on Θ . The approximation problem is thus at the heart of both compression and restoration problems.

2.1 1D Wavelets Bases

A wavelet basis B of $L^2([0, 1])$ is obtained by dilating and translating a function ψ [?, ?, ?, ?]

$$B \stackrel{\text{def.}}{=} \{ \psi_{j,n} \mid j \leq 0, n = 0 \dots 2^{-j} - 1 \} \quad \text{with} \quad \psi_{j,n}(x) \stackrel{\text{def.}}{=} 2^{-j/2} \psi(2^{-j}x - n),$$

with slight modifications for functions $\psi_{j,n}$ whose support intersect the boundary of $[0, 1]$. Wavelet are oscilating function with vanishing moments. A wavelet has p vanishing moments if it is orthogonal to polynomials up to degree $p - 1$:

$$\forall k \leq p - 1, \quad \int_0^1 \psi(x) x^k dx = 0.$$

Daubechies [?] shown that one can build such a wavelet that has a compact support and that generates an orthogonal basis The support of $\psi_{j,n}$ is thus proportional to 2^j and is localized around $2^j n \in [0, 1]$.

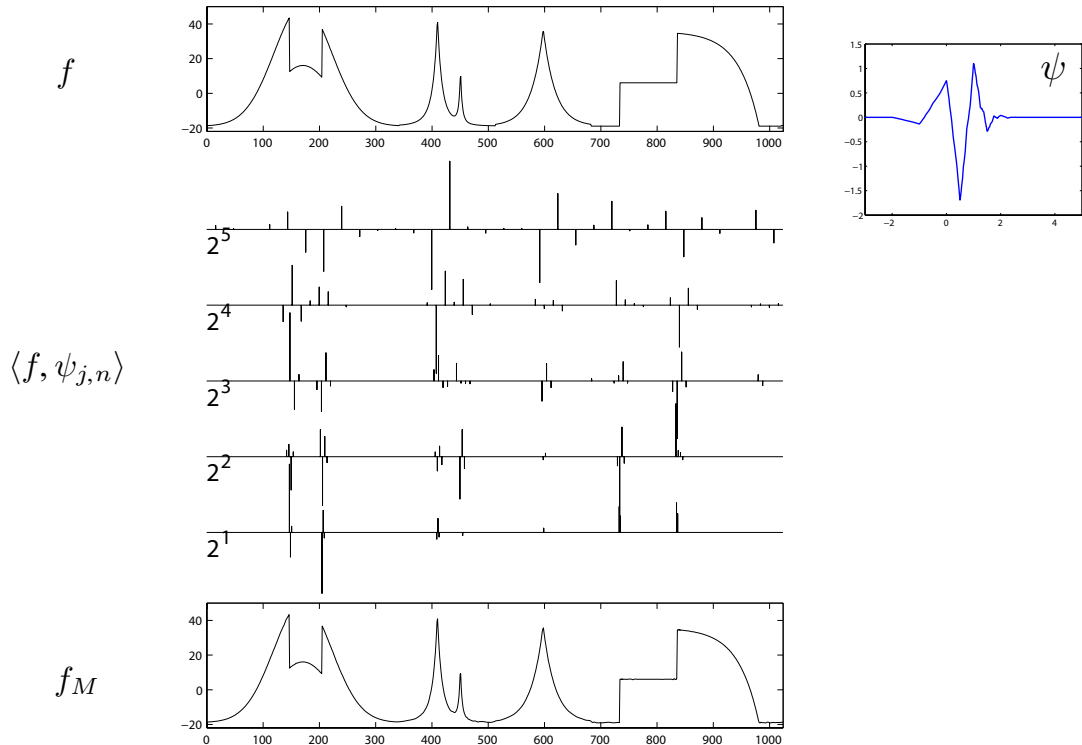


Figure 3: *Function f , wavelet transform and approximation f_M obtained by keeping the 10% largest wavelet coefficients.*

The construction of multiresolution spaces shows the simplicity of the wavelet transform [?, ?]. The existence of a fast algorithm comes from the fact that this transform can be factored in a product of elementary orthogonal operators. These operators are computed numerically with discrete convolutions with quadrature mirror filters that are dilated by inserting zeros. The cascade of these orthogonal filtering steps implements the fast wavelet transform that requires $O(N)$ operations for a signal of length N [?].

Figure 3 shows a piecewise regular function together with its wavelet coefficients $\langle f, \psi_{j,n} \rangle$. One can see that there are few large coefficients localized in the neighborhood of singularities. Indeed, if f is C^α in an interval that contains the support of a function $\psi_{j,n}$ then the wavelet coefficient is small for small scale 2^j : $|\langle f, \psi_{j,n} \rangle| = O(2^{j(\alpha+1/2)})$. If f is piecewise C^α and has a finite number of singularities, one can show that the approximation f_M in (1) obtained with the largest M wavelet coefficients satisfies

$$\|f - f_M\|^2 = O(M^{-2\alpha}). \quad (2)$$

This asymptotic decay is optimal and is equal with the one obtained if f has no singularity. The existence of a finite number of singularities thus does not affect the asymptotic

precision of a wavelet approximation. Figure 3 shows f_M computed with the 10% largest wavelet coefficients.

2.2 2D Wavelet Bases

Wavelet bases of $L^2([0, 1]^2)$ are obtained by translating and dilating 3 elementary wavelets $\{\psi^H, \psi^V, \psi^D\}$ which oscillate in the horizontal, vertical and diagonal directions. These wavelets are separable products of mono-dimensional wavelet functions. Figure 4 shows an example of 2D wavelets of compact support.




Figure 4: *Example of a 3-tuple of wavelets in 2D.*

A wavelet orthogonal basis of $L^2([0, 1]^2)$ can be written as

$$B = \left\{ \psi_{jn}^k(x) = 2^{-j} \psi^k(2^{-j}x - n) = 2^{-j} \psi^k(2^{-j}x_1 - n_1, 2^{-j}x_2 - n_2) \right\}_{\substack{k=H,V,D \\ j < 0, 2^j n \in [0, 1]^2}}.$$

Figure 5 (b,c) shows the wavelet coefficients along the three directions. These coefficients have been thresholded in order to keep only the 10% and 2% largest coefficients in (b) and (c). One can see on the zoom on f_M that with only 10% of the coefficients, one gets an accurate reconstruction and that the quality gets lower when the number of coefficients diminishes. The JPEG2000 image compression standard decomposes an image in a wavelet basis and performs a quantization and an entropic coding of the coefficients in order to optimize the binary code.

Figure 6 shows an application of wavelet bases to the denoising of images. Image 6 (b) is corrupted with a gaussian white noise W of variance σ . Image 6 (c) is a linear estimate obtained using a convolution with an optimized filter. Such a linear estimation suppresses a part of the noise but also smooths the image singularities which creates a blurry image. Image 6 (d) shows the wavelet coefficients of the noisy image. These coefficients are thresholded at a level $T = 3\sigma$ in order to keep only the largest coefficients. The restored image 6 (f) is obtained using the inverse wavelet transform of thresholded coefficients. As one can see, the noise has disappeared in homogeneous regions and edges are better reconstructed because their wavelet coefficients are kept by the thresholding.

The asymptotic accuracy of estimation and compression algorithms in an orthogonal wavelet basis depends upon the approximation power of this basis. If $f \in L^2([0, 1]^2)$ is a C^α image, then its approximation f_M in (1) with M wavelet coefficient satisfies:

$$\|f - f_M\|^2 = O(M^{-\alpha}). \quad (3)$$

This result is however no more valid if f is discontinuous along an edge. If f is piecewise regular meaning that it is C^α ($\alpha > 1$) outside a set of curves with finite length (contours), then the error decay satisfies only:

$$\|f - f_M\|^2 = O(M^{-1}). \quad (4)$$




Figure 5: *Approximation of an image in a wavelet basis with a varying number of coefficients. (a): Original image and zoom below. (b): 10% largest coefficients and zoom of the reconstruction. (c): 2% largest coefficients and zoom.*

Unlike the mono-dimensional case, the existence of singularities controls the asymptotic decay of the error, which becomes much slower. The result (4) is a special case of a general result for functions with bounded variations [?]. Figure 7 shows the wavelet coefficients of a piecewise regular image. Large coefficients are localized along the contours (black and white coefficients), so the number of these coefficients is proportional to the length of contours. These coefficients are responsible for the slower decay of the approximation.

The goal of a geometric representation is to take advantage of the geometric regularity of the image “singularities” to enhance the approximation result (4). In particular, one would like to obtain an approximation that satisfies $\|f - f_M\|^2 = O(M^{-\alpha})$ as if there was no singularity in the image. This is indeed the result (2) obtained for piecewise regular one dimensional functions.

3 Geometric Representations of Images

A simple model of geometrically regular images is defined as a function f that is uniformly C^α outside a set of curves which are themselves C^α with $\alpha > 1$. These curves correspond to the contours of objects that create occlusions. To model diffraction phenomena, the singularities of f can be blurred by an unknown convolution kernel. The triangles images of figure 9 are examples of geometrically regular images.

3.1 Finite Elements

A thresholding in a wavelet basis is equivalent to an approximation with finite element having a square support, such that the size of the elements is refined near the singularities, as shown on figure 9 (c). To enhance the performance of this kind of approximation, it is




Figure 6: *Denoising with a thresholding in a wavelet basis.*

necessary to adapt the geometry of the finite elements, using for instance an anisotropic adaptive triangulation.

Given a triangulation of $[0, 1]^2$ with M triangles, one can define an approximation \tilde{f}_M of f which is the best piecewise linear approximation on this triangulation. The goal of an adaptive triangulation is to optimize the shape of the triangles in order to minimize the approximation error $\|f - \tilde{f}_M\|$. Near a discontinuity, the triangle should be thin and stretched along the singularity curve, as displayed on figure 10. The lengths of the triangles should be of order M^{-1} and their widths should be of order M^{-2} . If f is C^2 outside a set of C^2 contours, then one has for such an adapted triangulation

$$\|f - \tilde{f}_M\|^2 = O(M^{-2}). \quad (5)$$

This construction can be generalized by replacing triangles by higher order geometric primitives whose boundaries are polynomial curves of degree α , as shown on the right side of figure 10. The adapted approximation using polynomials defined on M such higher order primitives leads to an approximation error $\|f - \tilde{f}_M\|^2 = O(M^{-\alpha})$ for a function f that is C^α outside a set of C^α contours.

Adaptive triangulations have proven very useful in numerical analysis where shocks or boundary layers require anisotropic refinement, see for instance the work of Aguilar and Goodman [?]. However, it exists currently no algorithm that can guaranty such an approximation result as (5) for functions as complex as images [?]. Indeed, the connectivity and shape of the triangulation should adapt itself to the local regularity of the image. When an image is smoothed by an unknown kernel of width s , the triangulation should depend on s in order to get the result of equation (5), as shown on figure 11. To reach an error decay of $O(M^{-2})$, in the neighborhood of a contour smoothed by a kernel of width s , the triangle should have a length of order $s^{1/4}M^{-1/2}$ and a width of order $s^{3/4}M^{-1/2}$. The scale s is most of the time unknown and one thus need an automatic algorithm to devise




Figure 7: A geometrically regular image together with its wavelet coefficients. For wavelet coefficients, gray corresponds to a coefficient near zero, white to a large positive value and black to a large negative value.




Figure 8: (a) A geometrically regular image. (b) Approximation with a triangulation. (c) Approximation with a wavelet basis (only the support of the basis functions used for the approximation are displayed).

the size of the triangles.

This analysis shows that it is possible to reach approximation error bounds that decay faster than wavelets approximation by adapting the representation to the geometry of the image. However the finite element approach does not yet come with algorithms that can handle complex images.

3.2 Curvelets

The curvelets basis of Candès and Donoho [?] brings a mathematical and algorithmic solution to the problem of approximating geometric images whose contours are C^2 . Unlike wavelets, curvelets are functions whose support are elongated like the anisotropic triangles of figure 9 (b). A curvelet is a function $\psi_{\theta,j,u}(x)$ whose support is centered around u , with length proportional to 2^j , a width proportional to 2^{2j} and an orientation θ . Figure 12 shows some examples of curvelets.

Candès et Donoho have build a Riesz basis of $L^2([0,1]^2)$ using curvelets. If f is C^2 with C^2 contours, they have shown that a thresholding of the curvelets coefficients defines an approximation f_M that satisfies

$$\|f - f_M\|^2 = O(M^{-2}(\log M)^3).$$

Up to a $\log(M)$ factor, one recovers the result (5) obtained using an optimal triangulation, but this time with an algorithmic approach. The beauty of this result comes from its simplicity. Unlike an optimal triangulation that has to adapt the aspect of its elements,




Figure 9: Approximation with finite elements on a triangulation for a function without and with additional blurring.




Figure 10: Finite elements for the approximation around a singularity curve.




Figure 11: Aspect ratio of triangles for the approximation of a blurred contour.

the curvelets basis is a priori fixed and the thresholding of the curvelets coefficients is enough to adapt the approximation to the geometry of the image. This simplicity however has a downside. The curvelets approximation is only optimal for piecewise C^α functions with $\alpha = 2$, but it is no more optimal for $\alpha > 2$ or for less regular functions such as bounded variation functions. For now, it does not exist orthogonal basis of curvelets, which makes them less efficient than wavelet to compress natural images.

3.3 Adaptive Representations

Many adaptive geometric representations have been proposed recently with good results in image processing. Instead of decomposing an image in a fixed a priori basis, an adaptive algorithm modifies the representation using a geometry computed from the image. The wedgelets of Donoho [?] segments the support of the image in dyadic adapted squares. On each square, the image is approximated with a constant value on each side of a straight edge. The direction of this estimated edge is optimized using the local content of the image. This approach is generalized by Shukla et al. [?] that replace constant values by polynomials and the straight edges by polynomial curves. This kind of approach is efficient as long as the geometry of the image is not too complex and edges are not blurred.

To enhance wavelets representations, Wakin et al. [?] and Dragotti et Vetterli [?] have proposed to approximate the wavelet coefficients using adaptive vector quantization techniques. Following the work of Matei and Cohen [?] on adaptive lifting schemes, new lifting algorithms have also been proposed to predict wavelet coefficients from their neighbors.



Figure 12: *Examples of curvelets.*

These works are mostly algorithmic and do not provide mathematical bounds. They use the fact that wavelet coefficients inherit some regularity from the image geometric regularity. Similar ideas are at the core of the bandlets construction.

4 Orthogonal Bandlets

A sparse representation takes advantage of some kind of regularity of the function to approximate. Wavelet bases exploit the isotropic regularity on square domains of varying sizes. Geometric regularity along edges in images is an anisotropic regularity. Although the image may be discontinuous across a contour, the image can be differentiable in a direction parallel to the tangent of the edge curve. The bandlet transform exploits such an anisotropic regularity by constructing orthogonal vectors that are elongated in the direction where the function has a maximum of regularity.

The first bandlet bases constructed by Erwan Le Pennec [?, ?] have brought optimal approximation results for geometrically regular functions. Later works have enriched this construction thanks to the use of a multiscale geometry defined over the coefficients of a wavelet basis [?, ?]. These multiscale bandlet bases are described in section 4.3.

4.1 Regularity of Wavelet Coefficients

The wavelet transform can be factored in a product of elementary orthogonal operators, obtained by dilating “quadrature mirror filters”. Orthogonal bandlet bases are obtained from a wavelet basis by using an additional cascade of orthogonal operators parameterized by the local geometry of the image.

The wavelet representation is both sparse and structured. For a geometrically regular image, figure 13 shows that for each scale the large coefficients are localized near the singularity curves. If K is the size of the support of the wavelets functions ψ^k , the large coefficients are localized in tubes of width K , as shown on figure 13. Those coefficients are compressed using an orthogonal “bandletization” operator that exploits the underlying geometric regularity.

Wavelet coefficients of f can be written as uniform samples from the function f regularized with a wavelet kernel ψ_j^k whose support has a width of 2^j :

$$\langle f, \psi_{jn}^k \rangle = f_j(2^j n) \quad \text{where} \quad f_j(x) \stackrel{\text{def.}}{=} f * \psi_j^k(x) \quad \text{and} \quad \psi_j^k(x) = \frac{1}{2^j} \psi^k\left(\frac{-x}{2^j}\right).$$

The convolution guarantees that f_j is at least as regular as ψ_j^k . The function f_j also inherits the regularity of f in the direction parallel to the edge. Figure 13 shows an example of a




Figure 13: Wavelet coefficients at a given scale 2^j are uniformly sampled from a regularized function $f * \psi_j^k(x)$ shown on the right.

set of coefficients near an edge. In order to derive an adapted approximation of the wavelet coefficients, we now study the regularity of f_j by bounding its derivatives.

In the following we study the regularity of f_j on some small square $S \subset [0, 1]^2$ of width λ . If f is a C^α function outside a C^α edge curve parameterized horizontally by $x_2 = \gamma(x_1)$, then one can control the derivatives of f_j along a vector field $(1, \tilde{\gamma}'(x_1))$ close to $(1, \gamma'(x_1))$. The same construction applies for vertically parameterized geometries. Figure 14 (a) shows an example of such an approximate flow $\tilde{\gamma}'$. This flow defines a warping operator w

$$(\tilde{x}_1, \tilde{x}_2) = w(x_1, x_2) \stackrel{\text{def.}}{=} (x_1, x_2 - \tilde{\gamma}(x_1)), \quad \text{where} \quad \tilde{\gamma}(x_1) = \int_0^{x_1} \tilde{\gamma}'(t) dt. \quad (6)$$

As shown on figure 14 (b,c), this warping modifies the sampling location and aligns the geometrical flow with the horizontal direction. The regularity of f_j along the flow can be formulated using the derivatives of the warped function $f_{jw}(\tilde{x}) \stackrel{\text{def.}}{=} f_j(w^{-1}(\tilde{x}))$ in the horizontal direction. Indeed, if the approximated flow $\tilde{\gamma}'$ is close from the original flow γ'

$$\forall (x_1, x_2) \in S, \quad |\gamma'(x_1) - \tilde{\gamma}'(x_1)| \leq (1 + \|\gamma\|_{C^\alpha}) \lambda^{\alpha-1}, \quad (7)$$

then the resulting warped wavelet coefficients $f_{jw}(\tilde{x}) \stackrel{\text{def.}}{=} f_j(w^{-1}(\tilde{x}))$ satisfy

$$\forall i_1, i_2 \leq \alpha, \forall \tilde{x}, \quad \left| \frac{\partial^{i_1+i_2} f_{jw}}{\partial \tilde{x}_1^{i_1} \partial \tilde{x}_2^{i_2}}(\tilde{x}) \right| \leq C 2^j (1 + \|\gamma\|_{C^\alpha}^\alpha) 2^{-j(i_1/\alpha - i_2)}. \quad (8)$$

where C is a constant that depends only on f . The bound of equation (8) states the anisotropic regularity of a set of wavelet coefficients.

4.2 Polynomial Approximation of Wavelets Coefficients

In order to capture the regularity stated by equation (8), one can use a piecewise polynomial approximation \tilde{f}_{M_j} of f_j defined on M_j small bands of length λ and width μ that follow locally the approximated flow $\tilde{\gamma}'$. Figure 15 (c) shows an example of layout of bands that follows the geometry of the image. Note that at a distance farther than K from any edge curve, one does not need to define bands since the wavelet coefficients are close to zero and for small scale 2^j one can approximate these coefficients $\langle f, \psi_{jn}^k \rangle$ by $\tilde{f}_{M_j}(2^j n) = 0$.

A local Taylor approximation of f_j in each band of length λ and width μ proves that the error of a polynomial approximation in each band can be made as low as

$$|f_j(x) - \tilde{f}_{M_j}(x)| \leq C \left\| \frac{\partial^\alpha f_{jw}}{\partial \tilde{x}_1^\alpha} \right\|_\infty \lambda^\alpha + C \left\| \frac{\partial^\alpha f_{jw}}{\partial \tilde{x}_2^\alpha} \right\|_\infty \mu^\alpha, \quad (9)$$

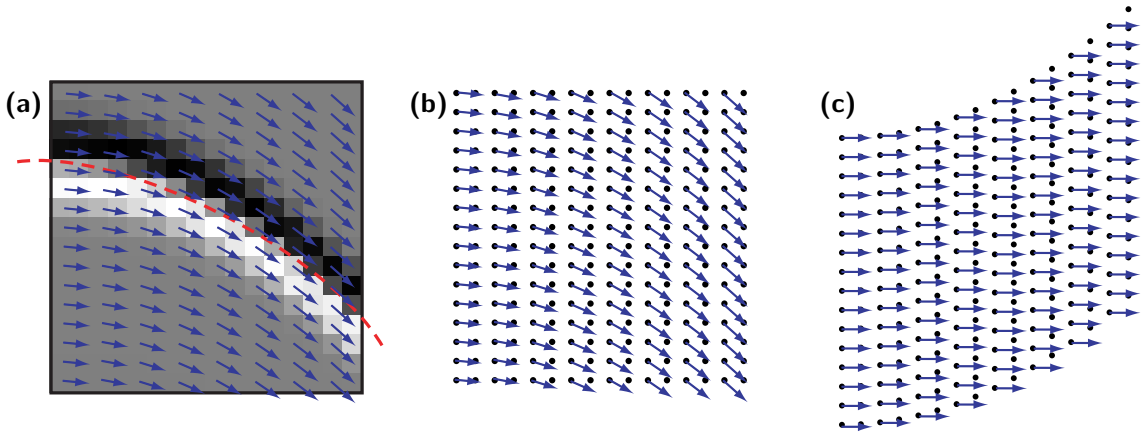


Figure 14: (a) Wavelets coefficients and geometric flow $\tilde{\gamma}'$. (b) Sampling position and geometric flow. (c) Warped sampling position and warped constant flow (horizontal).

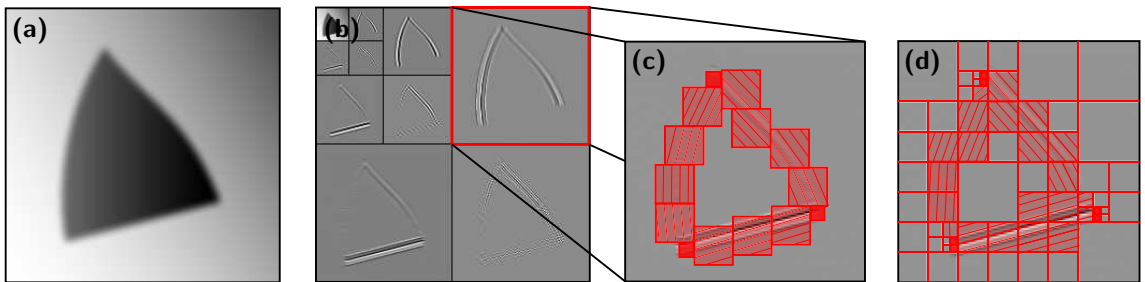


Figure 15: (a) Geometric images. (b) Wavelets coefficients. (c) Layout of bands of length λ and width μ . (d) Layout of bands inside a dyadic subdivision (bandlet approximation).

where C is a constant that depends only on f .

In order to optimally capture the anisotropy of f_j , one needs to scale the length λ and the width μ of these bands according to

$$\left\| \frac{\partial^\alpha f_{jw}}{\partial \tilde{x}_1^\alpha} \right\|_\infty \lambda^\alpha = \left\| \frac{\partial^\alpha f_{jw}}{\partial \tilde{x}_2^\alpha} \right\|_\infty \mu^\alpha \implies 2^{-j} \lambda^\alpha = 2^{-j\alpha} \mu^\alpha. \quad (10)$$

This optimal aspect ratio of the band is exactly the one derived in section 3.1 for the finite element approximation of contours, but this time for a geometric image smoothed by a kernel ψ_j of width 2^j . As detailed in [?], the approximation error of such a polynomial approximation satisfies $\|f_j - f_{M_j}\| = O(M_j^{-\alpha})$. To approximate the original image f , its wavelets coefficients $\langle f, \psi_{jn}^k \rangle$ are approximated using f_{M_j} for each scale 2^j . The resulting approximated function \tilde{f}_M can be shown to satisfies $\|f - \tilde{f}_M\| = O(\log^\alpha(M)M^{-\alpha})$, where M is the total number of bands used to define the polynomial approximations for each relevant scale 2^j .

4.3 Orthogonal Bandlets Approximation

The polynomial approximation scheme presented in the previous section is able to recover, up to a $\log(M)$ factor, the optimal approximation rate one would have for an uniformly regular image. It is very similar to the finite elements scheme presented in section 3.1, except that the approximation is now defined over the wavelet domain. This scheme does not provide an algorithm to compute the approximation, which is needed for applications.

The bandlet approximation scheme [?] solves this issue by computing the polynomial approximation by a thresholding in an orthogonal Alpert basis [?]. The Alpert transform can be thought as a polynomial wavelet transform adapted to an irregular sampling such as the one depicted on figure 14 (c). It is obtained by orthogonalizing multiresolution space of polynomials defined on the irregular sampling grid. The resulting vectors are not samples of a regular function but have vanishing moments on the sampling grid, which is the relevant property to approximate the warped wavelet coefficients. A vector corresponding to a sampling of a function with an anisotropic regularity is well approximated with a few vectors from the Alpert basis. This bandletization of wavelet coefficients using an Alpert transform defines a set of bandlet coefficients. These coefficients can be written as inner products $\langle f, b_{j,\ell,n}^k \rangle$ of the original image f with bandlet functions that are linear combinations of wavelet functions

$$b_{j,\ell,n}^k(x) = \sum_p a_{\ell,n}[p] \psi_{j,p}^k(x) .$$

The $a_{\ell,n}[p]$ are the coefficient of the Alpert transform, which depends on the local geometric flow $\tilde{\gamma}'$ since this flow defines the warped sampling locations depicted on figure 14 (c). The bandlet function is defined at some location n and scale 2^j . The Alpert transform introduces a new scale factor $2^\ell > 2^j$ which defines the elongation of the bandlet function. The bandlet $b_{j,\ell,n}(x)$ inherits the regularity of the wavelets $\psi_{j,p}^k(x)$.

Approximated Segmented Flow. The family of orthogonal bandlets depends on the local adapted flow $\tilde{\gamma}'$ defined over small squares $S \subset [0, 1]^2$ for each scale 2^j and orientation k . This parallel flow is characterized by an integral curve $\tilde{\gamma}$, already defined in equation (6), that one can see in dashed plot on the left of figure 14. As stated by equation (7), this integral curve does not need to be strictly parallel to the contour. This is due to the bidimensional regularization introduced by the smoothing of $f_j = f * \psi_j^k$ with the wavelet ψ_j^k . Locally, it is thus enough to use a polynomial approximation $\tilde{\gamma}'$ of γ' that will ensure that condition (7) is satisfied.




Figure 16: (a): Wavelet coefficients of the image. (b): Example of segmentation of the wavelet transform in dyadic squares of varying sizes. An adapted flow $\tilde{\gamma}'$ is computed over each square.

In order to approximate the geometry by a polynomial flow, one needs to segment the set of wavelet coefficients in squares S . For each scale 2^j and orientation k of the wavelet transform, this segmentation is obtained using a recursive subdivision in dyadic squares of various sizes, as shown on figure 16. Such a subdivision defines a quadtree that specifies if a square S should be further subdivided in four sub-squares of size twice smaller. If there

is no specific direction of regularity inside a square, which is the case either in uniformly regular regions or at the vicinity of edge junctions, then there is no geometric directional regularity to exploit. It is thus not necessary to modify the wavelet basis. In this case, no flow is defined, and it corresponds to the “regular” and “corner” squares of figure 16. One only needs to compute the adapted flow $\tilde{\gamma}'$ in “edge” squares in order to obtain a bandlet basis that exploits the anisotropic regularity of the image. In the following we denote by Γ_j^k the segmentation together with the adapted flows $\tilde{\gamma}'$ chosen in each square S of the segmentation.

A different geometry Γ_j^k can thus be chosen for each scale 2^j and orientation k in order to adapt to the evolution of the geometric structures through scales. The set of all geometries $\Gamma = \{\Gamma_j^k\}_j^k$ consists of all the quadtree segmentation and the adapted flow inside all the squares of the segmentation. Each potential geometry Γ corresponds to a bandlet basis $\mathcal{B}(\Gamma)$ and the set of bandlet bases $\mathcal{D} = \{\mathcal{B}(\Gamma)\}_\Gamma$ defines the bandlet dictionary. In order to approximate a function up to some predefined precision T on the bandlet coefficients, this dictionary can be kept of size polynomial in T^{-1} . We now explain how to compute a basis $\mathcal{B}(\Gamma^*)$ adapted to some function f with a fast algorithm.

Best Bandlet Approximation. A bandlet basis $\mathcal{B}(\Gamma)$ depends on the geometry Γ of the local flows defined by a dyadic segmentation of the wavelet coefficients and the choice of a polynomial flow inside each square (or no flow in regular and corner squares). The goal being to optimize the approximation of f , the best geometry Γ^* is the one that produces the approximation f_M of f with the lowest error for a fixed number M of parameters needed to describe f_M .

Let M_g be the number of parameters needed to specify the geometry Γ of the flow. This includes the parameters of the quadtree for each scale 2^j and orientation k and the parameters of the polynomial flow $\tilde{\gamma}'$ for each square. This geometry defines a bandlet basis $\mathcal{B}(\Gamma) = \{b_\nu\}_\nu$ of $L^2([0, 1]^2)$, where $\nu = (k, j, \ell, n)$ indexes each bandlet function. Let M_b be the number of bandlet satisfying $|\langle f, b_{\ell,j,n}^k \rangle| > T$ for some threshold T . The approximation

$$f_M = \sum_{|\langle f, b_\nu \rangle| > T} \langle f, b_\nu \rangle b_\nu$$

is defined by $M = M_b + M_g$ parameters. Among all geometries and thus all bandlet bases, one needs to find a bandlet basis that minimizes the error $\|f - f_M\|^2$ for a fixed number $M = M_b + M_g$ of parameters. This problem of “best orthogonal basis search” can be solved by minimizing the Lagrangian:

$$\mathcal{L}(\mathcal{B}(\Gamma), f, T) = \|f - f_M\|^2 + M T^2 = \sum_{|\langle f, b_\nu \rangle| \leq T} |\langle f, b_\nu \rangle|^2 + M T^2. \quad (11)$$

An approximation theorem [?, ?] shows that if f is uniformly C^α outside a set of curves that are themselves C^α then the best bandlet basis $\mathcal{B}(\Gamma^*)$ that minimizes the Lagrangian (11) defines an approximation f_M that satisfies

$$\|f - f_M\|^2 = O(M^{-\alpha}). \quad (12)$$

This result is still valid if f is regularized by a smoothing kernel that models the effects of diffraction during image acquisition. One can notice that the bandlet approximation does not require to know the value of α as long as $\alpha < p$, where p is the number of vanishing moments of the wavelet and Alpert bases. This adaptivity is the key to the efficiency of bandlet for natural images.

A best basis search algorithm allows to compute the best bandlet basis $\mathcal{B}(\Gamma^*)$ adapted to an image f in $O(NT^{(p-1)^2})$ operations where N is the number of pixels in the image [?, ?]. This algorithm relies on the embedded structure of the dyadic segmentation and on the fast Alpert transform algorithm.

4.4 Applications of Orthogonal Bandlets

Compression of images and surfaces. Image compression in a bandlet basis $\mathcal{B}(\Gamma) = \{b_\nu\}_\nu$ is a straightforward application of bandlet approximation. Similarly to the compression in wavelet bases, it requires to quantize uniformly the bandlet coefficients with a quantization step T

$$f_R \stackrel{\text{def.}}{=} \sum_\nu Q_T(\langle f, b_\nu \rangle) b_\nu, \quad (13)$$

where R is the number of bits needed to specify f_R and Q_T is a uniform quantizer defined by

$$Q_T(x) = qT, \quad \text{if} \quad (q - 1/2)T \leq x \leq (q + 1/2)T. \quad (14)$$

The distortion of this coding scheme is $D_b(R) = \|f - f_R\|^2$ and for a given bit budget R one thus needs to find the best bandlet basis $\mathcal{B}(\Gamma^*)$ that give the lowest distortion. Similarly to the bandlet optimization for function approximation, this can be achieved by minimizing the lagrangian \mathcal{L} of equation (11). Using the approximation result (12) one can show [?, ?] that if f is C^α outside a set of C^α curves, then the distortion in the adapted bandlet basis $\mathcal{B}(\Gamma^*)$ satisfies

$$D_b(R) = \|f - f_R\|^2 = O(R^{-\alpha} |\log(R)|^\alpha).$$

In a wavelet basis, the approximation (4) leads to a coding error that decays asymptotically much slower: $D_w(R) = O(R^{-1} |\log(R)|)$, see [?].

Figure 17 compares the compressed image f_R obtained with an average of $R = 0.2$ bit per pixel for a coding in a wavelet and a bandlet basis [?]. The distortion is lower in a bandlet basis than in a wavelet basis, which can be seen on the better reconstruction of the geometrical structures of the image. For the compression of 3D surfaces used in computer graphics, surfaces are locally parameterized on a 2D plane [?], and classical schemes from image processing can be used to compress them. Bandlet bases [?] adapt themselves to the geometry of the surface to compress, which typically exhibits both sharp features and smoothed edges. Figure 18 compares the reconstruction of surfaces in wavelets and in bandlets and shows the corresponding enhancement of the PSNR as a function of the number of bits R . Note that the error for surfaces is measured with the Hausdorff distance, which is the relevant distortion for computer graphics applications.

Image denoising. The enhancement of the approximation performances using bandlet bases has also applications in image denoising [?]. In denoising applications, one seeks to recover f from an observation $Y = f + W$ where W is a gaussian white noise of variance σ . A thresholding algorithm in a bandlet basis $\mathcal{B}(\Gamma) = \{b_\nu\}_\nu$ defines an estimator of f

$$\tilde{f} \stackrel{\text{def.}}{=} \sum_\nu S_T(\langle Y, b_\nu \rangle) b_\nu$$

where the threshold is set to $T = \lambda\sigma$ and where the thresholding is defined as

$$S_T(x) = \begin{cases} x & \text{if } |x| < T, \\ 0 & \text{otherwise.} \end{cases}$$




Figure 17: Comparison of the compression using wavelets and bandlets at 0.2 bits/pixel.

Donoho and Johnstone [?] proved that $\lambda = \sqrt{2 \log_e(N)}$, where N is the number of pixels, is asymptotically optimal when N increases.

This estimator can be computed in a best bandlet basis $\mathcal{B}(\Gamma^*)$ computed using the minimization of a Lagrangian similar to (11) but with a multiplier T that is tuned to reach the optimal risk decay. If the function f is C^α outside a set of C^α edges, then one can prove [?] that this estimator in the best basis $\mathcal{B}(\Gamma^*)$ has an average risk that satisfies

$$E(\|f - \tilde{f}\|^2) = O(|\log(\sigma)|^{\frac{1}{2\alpha+1}} \sigma^{\frac{2\alpha}{2\alpha+1}}). \quad (15)$$

This decay of the risk is asymptotically optimal up to a $\log(\sigma)$ factor for the class of geometrically regular functions. This best bandlet basis estimator corresponds to a model selection process where the model is defined by an adapted geometry.

Image restoration. Inverse problem are others applications where geometry plays an important role. The inversion of the tomography operator R is an important problem in medical imaging. The measuring process can be modeled as

$$Y = Rf + W$$

where R is the Radon transform and W is an additional gaussian white noise of variance σ^2 . The radon transform is defined as

$$(Rf)(t, \theta) = \int f(x) \delta(x_1 \cos \theta + x_2 \sin \theta - t) dx.$$

so that the value $(Rf)(t, \theta)$ sums the contributions of the original function f along lines parameterized by its slope θ and abscissa t . The inverse operator R^{-1} is unbounded and makes the inverse problem of recovering f ill-posed. A direct inversion of the Radon operator $R^{-1}Y = f + R^{-1}W$ considerably amplifies the noise $R^{-1}W$. A thresholding algorithm in a best bandlet basis $\mathcal{B}(\Gamma^*) = \{b_\nu\}_\nu$ defines an estimator of f

$$\tilde{f} \stackrel{\text{def.}}{=} \sum_\nu S_{T_\nu}(\langle R^{-1}Y, b_\nu \rangle) b_\nu.$$




Figure 18: Comparison on two examples of the Hausdorff distortion for surfaces compression using wavelets and bandlets.

For the tomography inversion, the thresholds T_ν depends upon the scales 2^j in order to match the amplification of noise $R^{-1}W$. For an index $\nu = (j, k, n, \ell)$, the threshold is set to $T_\nu \stackrel{\text{def.}}{=} \sigma 2^{-j/2}$. If f is C^α outside a set of C^α curves, one can prove that the average risk $E(\|f - \tilde{f}\|^2)$ of this estimator has the asymptotic decay (15) which is optimal up to a $|\log(\sigma)|$ factor [?].

5 Grouping Bandlets

The geometry Γ of orthogonal bandlets is described by a locally parallel flow $\tilde{\gamma}'$ over each square of a dyadic segmentation. Such a geometry is suitable for the approximation of geometrically regular images, but lacks flexibility to represent the complex geometry of turbulent textures. Junctions are not explicitly modeled and require a fine recursive segmentation to be isolated from the remaining contours. Furthermore, the segmentation in small square areas forbid to take advantage of the long range regularity of fine elongated structures such as the hair texture or the wood patterns in figure 19.




Figure 19: Example of textures with complex geometric pattern with long range regularity.

Grouplets [?] are constructed using a geometry inspired from the Gestalt theory [?]. This theory states a set of grouping laws that are supposed to be applied recursively during the human perception of a natural scene. Similarly, a grouplet transform uses a multiscale association field in order to group together coefficients in the direction specified by the flow. These recursive groupings allow to take into account junctions and long range regularities.

Similarly to the orthogonal bandlets introduced in section 4.3, this grouplet trans-

form can be used to define grouping bandlets. The grouplet scheme is applied over a set of wavelet coefficients and performs a bandletization operation similar to the directional Alpert transform, but this time along an association field and not a locally parallel geometric flow. This process defines grouping bandlet functions that are combinations of the wavelet functions located along the association field using a multiscale transformation similar to the Haar transform. This scheme can be orthogonal if the critical sub-sampling is performed during the wavelet and the grouplet transforms or it can be translation-invariant if the transforms are not sub-sampled.

5.1 Orthogonal Grouplets

Haar transform. At the first scale $2^\ell = 1$, the Haar transform of a signal $a[n]$ groups each odd coefficient $a[2n + 1]$ with the even neighboring coefficient $a[2n]$ and associates to this pair a mean and a difference:

$$M = \frac{a[2n] + a[2n + 1]}{2}, \quad D = \frac{a[2n] - a[2n + 1]}{\sqrt{2}}.$$

An “in place” transform stores the mean by replacing the even coefficients $s[2n] = M$ and the difference by replacing the odd coefficients $s[2n + 1] = D$. This orthogonal elementary operator is applied in a hierarchical manner on the mean coefficients by doubling the scale 2^ℓ at each iteration. At a scale 2^ℓ , the mean already computed in positions $a[2^\ell 2n]$ and $a[2^\ell(2n + 1)]$ are grouped together in order to compute new means and details:

$$M = \frac{a[2^\ell 2n] + a[2^\ell(2n + 1)]}{2}, \quad D = (a[2^\ell 2n] - a[2^\ell(2n + 1)])\sqrt{2^{(\ell-1)}},$$

and these values are stored in place: $s[2^\ell 2n] = M$ and $s[2^\ell(2n + 1)] = D$. At this stage, M is equal to the mean of the signal values over the interval $[2^{\ell+1}n, 2^{\ell+1}n + 2^{\ell+1}[$ and D is proportional to the difference of the means on $[2^{\ell+1}n, 2^{\ell+1}n + 2^\ell[$ and on $[2^{\ell+1}n + 2^\ell, 2^{\ell+1}n + 2^{\ell+1}[$.

Bandletization with a multiscale grouping. A bandletization by grouping applies this Haar transform over pairs of points that are neighbors according to some association field. Although this field could link arbitrary computed means, this field should group together points that have similar neighborhoods in order to exploit the geometry of the signal.

For a multidimensional signal (image or video), the sampling grid \mathcal{G}_0 is divided in two pre-defined sub-grids that we call the “even sub-grid” $\mathcal{G}_{1,e}$ and the “odd sub-grid” $\mathcal{G}_{1,o}$, in analogy to the one dimensional case. A weight $s[n]$ is initially set to 1 for points $n \in \mathcal{G}_0$. This weight represents the number of coefficients aggregated by the means computed during the computation. Each point $m_o \in \mathcal{G}_{1,o}$ of the odd grid is associated to a point $m_e \in \mathcal{G}_{1,e}$ of the even grid. This association is optimized so that the value of $a[n]$ for n in the vicinity of m_e are as close as possible to the value $s[p]$ for p in the vicinity of m_o . The vector $A_1[m_o] = m_e - m_o$ is the association field that dictate the grouping between each point of the odd grid and some point in the even grid. In practice, the associated point m_e can be computed with a best fit of radius P :

$$m_e = \operatorname{argmin}_{m \in \mathcal{G}_{1,e}} \sum_{|n| < P} |a[m - n] - a[m_o - n]|^2. \quad (16)$$

This kind of “block matching” association is used in video processing to compute the movement of structures in movies. This is the so-called optical flow, see figure 2 (a,b), but one could use other schemes to optimize the association field.

Like in a Haar transform, a weighted mean and a weighted difference are computed between the values of the signal that are grouped together:

$$M = \frac{s[m_e] a[m_e] + s[m_o] a[m_o]}{s[m_e] + s[m_o]} \quad (17)$$

$$D = (a[m_e] - a[m_o]) \frac{\sqrt{s[m_e]s[m_o]}}{\sqrt{s[m_e] + s[m_o]}} \quad (18)$$

The “in place” transform stores the differences on the odd grid and the means on the even grid, with a weight that sums the weight of the to associated points:

$$a[m_o] = D, \quad a[m_e] = M \quad \text{with} \quad s[m_e] = s[m_e] + s[m_o]. \quad (19)$$

One can check [?] that the matrix that transforms $(a[m_e], a[m_o])$ is orthogonal and the difference coefficient $a[m_o] = D$ is the inner product of the original signal with a unit-norm vector having one vanishing moment and a support equal to $s[m_o]$. As for the Haar transform, this process is repeated iteratively by doubling the scale at each step.

At some grouplet scale 2^ℓ , a mean signal $a[m]$ has been computed during the previous iterations on an even grid $\mathcal{G}_{\ell,e}$. This grid is sub-divided in an “even sub-grid” $\mathcal{G}_{\ell+1,e}$ and an “odd sub-grid” $\mathcal{G}_{\ell+1,o}$. Each point $m_o \in \mathcal{G}_{\ell+1,o}$ of the odd grid is associated to a point $m_e \in \mathcal{G}_{\ell+1,e}$ of the even grid, which is optimized so that the values $a[n]$ near $n = m_e$ are close to the values near $n = m_o$. This association field is denoted as $A_\ell[m_o] = m_e - m_o$. One can use for instance a block matching similar to equation (16) to compute this association between m_o and m_e . One then computes new means and differences using equations (17) and (18). Those values are respectively stored in the even sub-grid $\mathcal{G}_{\ell+1,e}$ and the odd sub-grid $\mathcal{G}_{\ell+1,o}$, together with an update of the weights using equation (19). This cascade of orthogonal operators decomposes the original signal in an orthogonal basis called grouplet basis that is adapted to the signal geometry. Figure 22 (c) shows examples of grouplet vectors which have 1 vanishing moments.




Figure 20: Column embedded grids subsample the columns of $\mathcal{G}_{\ell-1}$ to define the grid \mathcal{G}_ℓ (empty circles) and the complementary grid $\tilde{\mathcal{G}}_\ell$ (filled circles). The association field groupings are illustrated by arrows.

The splitting $\mathcal{G}_{\ell,e} = \mathcal{G}_{\ell+1,e} \cup \mathcal{G}_{\ell+1,o}$ can be performed freely and one can devise a scheme for any targeted application. Figure 20 shows an example of horizontal associations where the “even sub-grid” $\mathcal{G}_{1,e}$ is the set of even columns (black dots) and the “odd sub-grid” $\mathcal{G}_{1,o}$ is the set of odd columns (white circles). This kind of splitting scheme is relevant for 2D

signals where the geometric structures propagate in the horizontal direction. This is indeed the case for seismic imaging, as one can see on the association fields computer in figure 21. For other applications, one can design a more isotropic splitting scheme that does not favor any orientation.




Figure 21: *Grouping of an association field at scales 2^1 , 2^2 and 2^3 computed by block matching for a seismic image shown in transparency.*

On figure 22, one can see the grouplet coefficients obtained by transforming the image according to the associations fields displayed on figure 21. The detail coefficients require the same amount of storage as the original image, but for a better understanding of their structure, figure 22 (b) rearranges them from the coarse scale $2^\ell = 2$ to fine scales from left to right. One can see that most of these transformed coefficients are gray (near zero) which was not the case of the original image. It means that the transform has been able to exploit the anisotropic regularity of the seismic image in order to reach a sparser representation than the pixel values. It is however unclear how to adjust the complexity of the association fields so that it can be coded with few bits to reach good compression results with grouplets.




Figure 22: *(a): Original seismic image. (b): Orthogonal grouplet coefficients over 6 scales, displayed with the same dynamic range as in (a). (c) Example of grouplet vectors.*

A denoising of the image can be performed by using the thresholding technic defined in (13). Grouplet coefficients below the noise level are thresholded to zero. The association fields $\{A_\ell\}_\ell$, which define the geometry of the transform, are computed over the noisy image. Thanks to the block matching procedure of equation (16) the estimation can be made robust by choosing a radius P large enough. This thresholding denoising is equivalent to performing automatically an adaptive averaging of the signal along the directions of regularity estimated by the association fields. Figure 23 gives an example from [?] for denoising a seismic image.




Figure 23: (a): Original image. (b): Noisy seismic image (PSNR=26dB). (c): Orthogonal grouplet coefficients. (d): Thresholded grouplet coefficients. (e): Image reconstructed from thresholded orthogonal coefficients (PSNR=27.3dB). (f): Image reconstructed from thresholded tight frame coefficients (PSNR=29.5dB).

5.2 Translation Invariant Grouplet Tight Frame

For denoising and computer graphics applications, the strict orthogonality of the grouplet transform is a source of inefficiencies since it forbids a translation invariant processing of the image. In order to solve this issue, one should remove the sub-sampling of the Haar transform in order to have a stable redundant transform. The same ideas carry over the grouplet setting by replacing the splitting of the grids by a partial ordering $m \prec m'$ of the points $m, m' \in \mathcal{G}_0$ on the original grid. This ordering can be thought as a 1D traversal of the grid points that ensures that a point m satisfying $m \prec m'$ is processed before m' during the grouplet computations.

For each scale 2^ℓ , the current mean values $a[m]$ is defined on the whole grid \mathcal{G}_0 . Each grid point $m \in \mathcal{G}_0$ is associated to a point $\tilde{m} \in \mathcal{G}_0$ located before: $\tilde{m} \prec m$. The association field is defined $A_\ell[m] = m - \tilde{m}$ and the optimization of \tilde{m} is carried using a block matching similar to (16) under the restriction that $\tilde{m} \prec m$. We further require that $|m - \tilde{m}| \approx 2^\ell$ in order to force the averaging of the grouping process to be performed over an increasing distance. This grouplet transform is not computed “in place” since it increases the number of coefficients. The grouped points (m, \tilde{m}) are processed in order to update “in place” the weight s and mean a values and to extract a new detail coefficient $d_\ell[m]$ that are stored in a different array.

$$d_\ell[m] = (a[m] - a[\tilde{m}]) \frac{\sqrt{s[m]s[\tilde{m}]}}{\sqrt{s[m] + s[\tilde{m}]}}$$

$$a[\tilde{m}] = \frac{s[m]a[m] + s[\tilde{m}]a[\tilde{m}]}{s[m] + s[\tilde{m}]}$$

$$s[\tilde{m}] = s[m] + s[\tilde{m}].$$

Once the process has been performed over L grouplet scales, the recursion is stopped and the remaining coarse scale layer is kept with a renormalization $a_L[m] = a[m]\sqrt{s[m]s[m]}$.

The translation invariant grouplet transform maps the original signal $a[m]$ to the set of coefficients $\{d_\ell, a_L\}_{\ell \leq L}^{m \in \mathcal{G}_0}$. The overall process is stable and one can prove [?] that it satisfies an energy conservation:

$$\|a\|^2 = \sum_{m \in \mathcal{G}_0, \ell \leq L} \frac{1}{2^\ell} |d_\ell[m]|^2 + \sum_{m \in \mathcal{G}_0} \frac{1}{2^L} |a_L[m]|^2. \quad (20)$$

This means that $\{d_\ell[m], a_L[m]\}_{m \in \mathcal{G}_0}^{1 \leq \ell \leq L}$ can be interpreted as the signal coefficients in a grouplet tight frame. A thresholding can be applied over these tight frame coefficients in order to perform denoising. Figure 23 shows that the translation invariance brings a significant improvement with respect to the original orthogonal grouplet denoising and improves the PSNR by 2.2dB.

5.3 From Grouplets to Bandlets

A grouping bandlet transform is obtained by applying the grouplet bandletization process to the set of coefficient of a multiscale representation. One applies the orthogonal grouplet transform over each scale 2^j and orientation k of an orthogonal wavelet transform in order to get the decomposition of the image on an orthogonal basis $\mathcal{B}(\Gamma) = \{b_\nu\}_\nu$ of grouping bandlets. The index $\nu = (j, k, \ell, m)$ refers to the wavelet scale 2^j , wavelet orientation k , grouplet scale 2^ℓ and grouplet position m . Similarly to the original orthogonal bandlet exposed in section 4.3, the cascade of the orthogonal wavelet transform and the orthogonal grouplet transform defines an orthogonal grouping bandlet transform. The geometry $\Gamma = \{A_\ell^{j,k}\}_\ell^{j,k}$ is now composed of the association fields $A_\ell^{j,k}$ computing during the grouplet transforms for each wavelet scale 2^j , orientation k and grouplet scale 2^ℓ .

Another option is to use a translation invariant wavelet tight frame and to apply the translation invariant grouplet tight frame over each scale and orientation. Similarly to the orthogonal bandlet bases, the set of association fields is denoted as Γ . Those geometries parameterizes the grouping bandlet transform whose coefficients are the inner product with bandlet function $\{b_\nu\}_\nu$. The cascade of the wavelet and grouplet tight frames defines a grouping bandlet tight frame of $L^2([0, 1]^2)$. Figure 24 shows an example of such a decomposition, where the association fields are displayed for various wavelets orientation and scales. For such application of the bandlet translation invariant transform, the partial ordering \prec is set in order to match the direction k of the wavelet (either horizontal, vertical or diagonal), see [?].

The grouping bandlet bases are more flexible than the orthogonal bandlet bases exposed in section 4.3. If one impose multiscale association fields that follow the integral lines of a locally parallel vector field $\tilde{\gamma}'$, then the grouplet transform is equivalent to the original Alpert bandletization with 1 vanishing moment. However, the bandletization with grouping is more general since the association fields can deviate from the integral lines in order to converge to singularity points such as junction or crossings. The following applications to denoising and synthesis show that this flexibility is indeed crucial when one is not concerned with image compression.

5.4 Applications of Grouplets

Image denoising. The flexibility of the association process of grouping bandlets makes it efficient for the denoising of textures with a complex geometry. Figure 25 compares the denoising obtained on the textured hat of the Lena image using a thresholding in wavelet and grouping bandlet tight frames. The bandlet transform is able to recover the

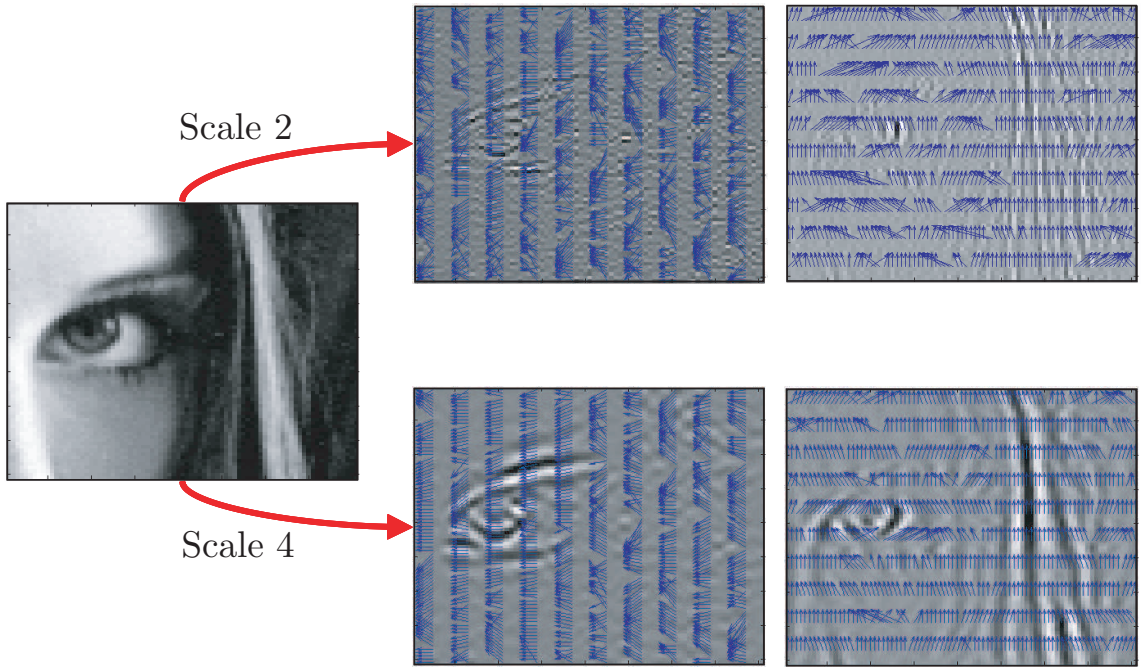


Figure 24: *Multiscale association fields for various scale and orientation of a translation invariant wavelet transform. Left: original image. Center: horizontal details. Right: vertical details.*

fine structures of the texture, which is not possible with wavelets that perform an isotropic regularization not suited to the directional oscillations of the texture.

Texture synthesis. To perform texture synthesis one can exploit the sparsity of the representation of a given input texture f in a grouping bandlet frame $\mathcal{B}(\Gamma) = \{b_\nu^\Gamma\}_\nu$, where the association fields Γ are computed during the grouping bandlet transform of f . The texture synthesis is performed by modifying the original geometry Γ into $\tilde{\Gamma}$ using some user interaction. Figure 26, left, shows some examples of vector fields defined by the user that can be used to construct associations fields along integral line of the flow. This new geometry defines a new bandlet frame $\mathcal{B}(\tilde{\Gamma})$ adapted to the texture to synthesize. Figure 26, right, shows some example synthesis \tilde{f} where the coefficients $\langle \tilde{f}, b_\nu^{\tilde{\Gamma}} \rangle$ are realizations of a random variable whose histogram matches the one of the original texture coefficients $\{\langle f, b_\nu^\Gamma \rangle\}_\nu$, using an algorithm introduced in [?].

Image inpainting. In inpainting applications, one needs to fill some hole $\Omega \subset [0, 1]^2$ where pixels of an image f are missing. One thus looks for an image \tilde{f} such that $\tilde{f}(x) = f(x)$ for $x \notin \Omega$ and the overall function \tilde{f} should have the same geometrical regularity as f outside Ω . This is enforced by imposing that the bandlet image representation is sparse which is obtained by minimizing the ℓ^1 norm of bandlet coefficients. This solution is thus calculated with the following minimization

$$\tilde{f} = \operatorname{argmin}_{g, \Gamma} \sum_{\nu} |\langle g, b_\nu^\Gamma \rangle| \quad \text{subject to} \quad \forall x \notin \Omega, \tilde{f}(x) = f(x), \quad (21)$$

with some additional constrained on Γ that can be enforced by using a large enough radius P during the block matching computation (16) of the association fields Γ . The minimization




Figure 25: *Comparison of the denoising with a translation invariant wavelet transform and a translation invariant grouping bandlet transform.*

of equation (21) is hard to solve and in practice, one can use an iterative thresholding procedure similar to the morphological component analysis of Starck et al. [?]. The resulting algorithm fills the hole with some random noise and perform iterative denoising using a decreasing threshold T . Between each iteration, the known values $\tilde{f}(x) = f(x)$ of the pixels $x \notin \Omega$ are enforced.




Figure 26: *Left: original texture, whose adapted bandlet frame is $\mathcal{B}(\Gamma)$. Center: geometrical flow used to define the multiscale association field $\tilde{\Gamma}$ for the synthesis. Right: texture synthesized as a realization of a random field over the coefficients in the bandlet frame $\mathcal{B}(\tilde{\Gamma})$.*




Figure 27: *Iteration of the inpainting process that modifies the image to obtain a sparse representation in an adapted grouping bandlet basis.*