

Sparse Geometric Image Representations With Bandelets

Erwan Le Pennec and Stéphane Mallat, *Fellow, IEEE*

Abstract—This paper introduces a new class of bases, called *bandelet* bases, which decompose the image along multiscale vectors that are elongated in the direction of a geometric flow. This geometric flow indicates directions in which the image gray levels have regular variations. The image decomposition in a bandelet basis is implemented with a fast subband-filtering algorithm. Bandelet bases lead to optimal approximation rates for geometrically regular images. For image compression and noise removal applications, the geometric flow is optimized with fast algorithms so that the resulting bandelet basis produces minimum distortion. Comparisons are made with wavelet image compression and noise-removal algorithms.

Index Terms—Nonlinear filtering and enhancement (2-NFLT), still image coding (1-STIL), wavelets and multiresolution processing (2-WAVP).

I. INTRODUCTION

IMAGE representation in separable orthonormal bases, such as wavelets, local cosine, or Fourier, cannot take advantage of the geometrical regularity of image structures. Sharp image transitions such as edges are expensive to represent, although one could reduce their cost by taking into account the fact that they often have a piecewise regular evolution across the image support. Integrating the geometric regularity in the image representation is, therefore, a key challenge to improve state of the art applications to image compression, denoising, or inverse problems. By reviewing previous approaches, Section II-B explains the difficulties to create stable and efficient geometric representations.

This paper introduces a new class of bases, with elongated multiscale *bandelet* vectors, which are adapted to the image geometry. A bandelet basis is constructed from a *geometric flow* of vectors, which indicate the local directions in which the image gray levels have regular variations. In applications, this geometric flow must be optimized to build bandelet bases that take advantage of the image geometric regularity. To compress an image with a transform code in a bandelet basis, we describe a fast algorithm that computes the geometric flow by minimizing the Lagrangian of the distortion rate. Thresholding estimator in

bandelet bases are also studied for noise removal. A penalized best basis search approach is used to optimize the geometric flow with a fast algorithm.

Bandelet bases are obtained with a *bandeletization* of warped wavelet bases, which takes advantage of the image regularity along the geometric flow. Section III explains how to construct such bases together with their geometric flow and Section IV gives a fast subband filtering algorithm to decompose an image in a discrete bandelet basis. Section V studies applications to image compression and noise removal. In both cases, the geometric flow is optimized with a fast algorithm, that requires $O(N^2(\log_2 N)^2)$ operations for an image of N^2 pixels. Numerical results show that optimized bandelet bases improve significantly image compression and denoising results obtained with wavelet bases. This paper concentrates on algorithms and applications, but mathematical proofs of asymptotic results can be found in [1] and [2].

II. SPARSE IMAGE REPRESENTATIONS

Orthonormal bases are particularly convenient to construct sparse signal approximations for applications such as image compression or noise removal with thresholding estimators. An image f can be approximated in an orthonormal basis $\mathcal{B} = \{g_m\}_m$ by the partial sum

$$f_M = \sum_{m \in I_M} \langle f, g_m \rangle g_m$$

where I_M is the index set of the M largest inner products whose amplitude are above a threshold T_M

$$I_M = \{m \in \mathbb{N} : |\langle f, g_m \rangle| > T_M\}.$$

The resulting approximation error is

$$\|f - f_M\|^2 = \sum_{m \notin I_M} |\langle f, g_m \rangle|^2. \quad (1)$$

For compression applications, the inner products are not just thresholded but quantized and coded. Yet, it has been shown in [3] that for a uniform quantization of step T_M , at high compression rates, the quadratic distortion D is proportional to $\|f - f_M\|^2$ and the total bit budget R is proportional to M . The distortion rate $D(R)$, thus, has an asymptotic decay that is the same as the approximation error $\|f - f_M\|^2$ as a function of M . The efficiency of thresholding estimators that remove additive white noises by representing the signal in the basis \mathcal{B} also depends upon this approximation error [4]. For both applications, given some prior information on the properties of f , we, thus, want to find a basis \mathcal{B} where $\|f - f_M\|^2$ converges quickly to zero when

Manuscript received July 3, 2003; revised March 26, 2004. This work was supported in part by National Science Foundation Grant IIS-0114391. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Giovanni Ramponi.

E. Le Pennec is with the Centre de Mathématiques Appliquées, Ecole Polytechnique, 91128 Palaiseau Cedex, France.

S. Mallat is with the Centre de Mathématiques Appliquées, Ecole Polytechnique, 91128 Palaiseau Cedex, France, and also with the Courant Institute of Mathematical Science, New York University, New York, NY 10012 USA (e-mail: stephane.mallat@polytechnique.fr).

Digital Object Identifier 10.1109/TIP.2005.843753

M increases. This is the case if there exists a small constant C and a large exponent β with

$$\|f - f_M\|^2 \leq C M^{-\beta}. \quad (2)$$

A. Nonlinear Image Approximations With Wavelets

Wavelet bases are particularly efficient to approximate images. A separable wavelet basis is constructed from a one-dimensional (1-D) wavelet $\psi(t)$ and a scaling function $\phi(t)$

$$\psi_{j,m}(t) = \frac{1}{\sqrt{2^j}} \psi\left(\frac{t - 2^j m}{2^j}\right)$$

and $\phi_{j,m}(t) = \frac{1}{\sqrt{2^j}} \phi\left(\frac{t - 2^j m}{2^j}\right)$.

The resulting family of separable wavelets

$$\left\{ \begin{array}{l} \phi_{j,m_1}(x_1) \psi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1) \phi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1) \psi_{j,m_2}(x_2) \end{array} \right\}_{j \in \mathbb{Z}, (m_1, m_2) \in \mathbb{Z}^2} \quad (3)$$

is an orthonormal basis of $\mathbf{L}^2(\mathbb{R}^2)$ [5], [6]. To construct a basis over a subset Ω of \mathbb{R}^2 , one must keep the wavelets whose support are inside Ω and modify appropriately the ones whose supports intersect the boundary of Ω . Several approaches have been developed to do so [7]–[9]. The orthogonality of the basis is obtained with an orthogonalization procedure introduced by Jaffard and Meyer [8]. With a slight abuse of notation, we shall still write $\phi_{j,m_1}(x_1) \psi_{j,m_2}(x_2)$, $\psi_{j,m_1}(x_1) \phi_{j,m_2}(x_2)$, and $\psi_{j,m_1}(x_1) \psi_{j,m_2}(x_2)$ the modified wavelets at the boundary even if they do not have this tensorial structure. The resulting basis of $\mathbf{L}^2(\Omega)$ can be written as follows:

$$\left\{ \begin{array}{l} \phi_{j,m_1}(x_1) \psi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1) \phi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1) \psi_{j,m_2}(x_2) \end{array} \right\}_{(j, m_1, m_2) \in I_\Omega} \quad (4)$$

where I_Ω is an index set that depends upon the geometry of the boundary of Ω .

If the image $f(x_1, x_2)$ is uniformly regular, which is measured by the fact that it is \mathbf{C}^α (α times continuously differentiable) and if the wavelet ψ has $p > \alpha$ vanishing moments, then one can prove [10] that there exists a constant C such that the approximation f_M from M wavelets satisfies

$$\|f - f_M\|^2 \leq C M^{-\alpha}. \quad (5)$$

This decay exponent α is optimal in the sense that no other approximation scheme can increase its value over the set of functions f that are \mathbf{C}^α [10]. However, wavelet bases are not the only bases to achieve the optimal rate (5).

If f is \mathbf{C}^α ($\alpha > 1$) everywhere outside curves of finite length along which it is discontinuous, then the discontinuities create many fine scale wavelet coefficients of large amplitude. As a result, the error decay (5) does not hold. However, one can still prove that there exists a constant C such that

$$\|f - f_M\|^2 \leq C M^{-1}. \quad (6)$$

This result extends to all bounded variation images, which are characterized by the fact that their level set have a finite average length [11]. Moreover, wavelet bases are optimal for bounded variation images in the sense that there exists no basis that leads




Fig. 1. (a) Image which is \mathbf{C}^2 inside and outside a domain Ω . (b) Adapted triangulation that covers the boundary with narrow triangles.

to an approximation error (2) with a decay exponent $\beta > 1$ over all such functions [11].

Yet, this optimality result can be improved by observing that the level sets of many images not only have a finite average length but define regular geometric curves. Exploiting this geometric regularity can improve the representation, as shown by the following example. Let Ω be a subset of $[0, 1]^2$ whose boundary $\partial\Omega$ is a piecewise \mathbf{C}^2 curve, with a finite number of corners, as illustrated in Fig. 1(a). Suppose that $f(x_1, x_2)$ is a \mathbf{C}^2 function inside and outside Ω , which is discontinuous along $\partial\Omega$. Let us construct a triangulation adapted to the image geometry, as illustrated in Fig. 1(b). The boundary $\partial\Omega$ is covered with narrow triangles whose widths are $O(M^{-2})$, and the inside and outside of Ω are covered by large triangles so that the total number of triangles is M . One can define a piecewise linear approximation f_M over these M triangles which satisfies [12]

$$\|f - f_M\|^2 \leq C M^{-2}. \quad (7)$$

The decay rate exponent $\beta = 2$ is better than the exponent $\beta = 1$ obtained in (6) with wavelets and is the same as the optimal exponent (5) obtained for an image f which is \mathbf{C}^2 over its whole support. Hence, the existence of discontinuities does not degrade the asymptotic decay of this approximation.

This simple example shows that exploiting the geometric image regularity can lead to much smaller approximation errors for a fixed number of approximation elements M . However, adaptive triangulations are extremely hard to construct for natural images which generally have a complex geometry [12]. Moreover, one would like to extend this result for regularity indexes $\alpha \geq 2$. If the boundary $\partial\Omega$ is a \mathbf{C}^α curve and if f is \mathbf{C}^α inside and outside Ω then one would like to find a geometric approximation f_M from M elements such that

$$\|f - f_M\|^2 \leq C M^{-\alpha}. \quad (8)$$

We shall see that bandelet bases are able to achieve this optimal decay rate.

B. Geometric Image Representations

The construction of geometric image representations is a very active research area where many beautiful and innovative ideas

have been tested. Summarizing the different approaches will help understand the major difficulties.

In the computer vision community, Carlsson [13] proposed in 1988 an edge-based image representation which measures the image jumps across curves in the images, called edges. An image approximation is then calculated by imposing the same jumps along the edge and by computing values between edges with a diffusion process. Many edge-based image representations have then been elaborated along similar ideas [14], [15], with different edge-detection procedures and image approximations using jump models along these edges. To refine these models, multiscale edge representations using wavelet maxima [16] or an edge-adapted multiresolution [17] have also been studied. Edge-based image representations with noncomplete orthonormal families of foveal wavelets [18] or foot prints [19] have been introduced and studied to reconstruct the main image edge structures. To stabilize the edge detection, global optimization procedures have also been elaborated by Donoho [20], Shukla *et al.* [21] and Wakin *et al.* [22]. The optimal configuration of edges is then calculated with an image segmentation over dyadic squares using fast dynamic programming algorithms over quad trees.

A major difficulty that faces all edge-based approaches is that sharp image transitions often do not correspond to discontinuous jumps along edge curves. On one hand, the optical diffraction produces an averaging effect which blurs the gray-level discontinuities along occlusion boundaries, and on the other hand many sharp transitions are produced by texture variations that are not aggregated along geometric curves. Currently, edge-based algorithms do not seem to outperform separable orthogonal wavelet approximations on complex images such as Lena, over the range of approximation errors where these algorithms are used in applications.

All the approaches previously described are adaptive in the sense that the representation is adapted to a geometry calculated from the image. Surprisingly, a remarkable result of Candès and Donoho [23] shows that one can construct a nonadaptive representation that takes advantage of the image geometric regularity by decomposing it in a fixed basis or frames of curvelets. Curvelet families are composed of multiscale elongated and rotated functions that defines bases or frames of $\mathbf{L}^2(\mathbb{R}^2)$. Candès and Donoho [23] proved that that an approximation f_M with M curvelets of an image f having discontinuities (blurred or not) along \mathbf{C}^2 curves produces an error that satisfies

$$\|f - f_M\|^2 \leq C M^{-2} (\log_2 M)^3. \quad (9)$$

By comparing this to (8), we see that this approximation result is nearly asymptotically optimal up to the $(\log_2 M)^3$ factor. Do and Vetterli [24] used similar ideas to construct contourlets that can be computed with a perfect reconstruction filter bank procedure. However, the beautiful simplicity due to the nonadaptivity of curvelets has a cost: curvelet approximations loose their near optimal properties when the image is composed of edges which are not exactly piecewise \mathbf{C}^2 . If edges are along irregular curves of finite length (bounded variation functions), then curvelets approximations are not as precise as wavelet approximations. If the edges are along curves whose regularity is \mathbf{C}^α with $\alpha > 2$

then the approximation decay rate exponent remains 2 and does not reach the optimal value α .

In image processing applications, we generally do not know in advance the geometric image regularity. It is, therefore, necessary to find approximation schemes that can adapt themselves to varying degrees of regularity. Our goal is, thus, to construct an adaptive image approximation f_M of f , with M parameters, which satisfies an optimal decay rate

$$\|f - f_M\| \leq C M^{-\alpha}.$$

The parameter α is a priori unknown and specifies the geometric image regularity.

III. BANDELETS ALONG GEOMETRIC FLOWS

Instead of describing the image geometry through edges, which are most often ill defined, we characterize the image geometry with a *geometric flow* of vectors. These vectors give the local directions in which the image has regular variations. Orthogonal bandelet bases are constructed by dividing the image support in regions inside which the geometric flow is parallel. Section III-B relates the optimization of the geometric flow to the precision of bandelet image approximations.

A. Block Bandelet Basis

This section describes the construction of a bandelet basis from a wavelet basis that is warped along the geometric flow, to take advantage of the image regularity along this flow. Conditions are imposed on the geometric flow to obtain orthonormal bandelet bases.

In a region Ω , a geometric flow is a vector field $\vec{\tau}(x_1, x_2)$ which gives a direction in which f has regular variations in the neighborhood of each $(x_1, x_2) \in \Omega$. If the image intensity is uniformly regular in the neighborhood of a point, then this direction is not uniquely defined. Some form of global regularity is, therefore, imposed on the flow to specify it uniquely. To construct orthogonal bases with the resulting flow, a first regularity condition imposes that the flow is either parallel vertically, which means that $\vec{\tau}(x_1, x_2) = \vec{\tau}(x_1)$, or parallel horizontally and, hence, $\vec{\tau}(x_1, x_2) = \vec{\tau}(x_2)$. To maintain enough flexibility, this parallel condition is imposed within subregions Ω_i of the image support. The image support \mathcal{S} is, thus, partitioned into regions $\mathcal{S} = \cup_i \Omega_i$, and within each Ω_i the flow is either parallel horizontally or vertically. Fig. 2(a) shows an example of a vertically parallel geometric flow in a region of a real image. If the image intensity f is uniformly regular over a whole region Ω_i then a geometric flow is meaningless and is, therefore, not defined.

Fig. 2(b) gives an example where the image is partitioned into square regions that are small enough so that each region Ω_i includes at most one contour. As a result, the size of the squares becomes smaller in the neighborhood of corners and junctions, up to a minimum size. In a region that does not include any contour, the image intensity is uniformly regular and the flow is, therefore, not defined. In each region including a contour piece, the flow is chosen to be parallel to the tangents of the contour curve, so that the contour corresponds to a line of flow. Bandelets are constructed in these regions by warping separable




Fig. 2. (a) Example of flow in a region. Each arrow is a flow vector $\vec{r}(x_1, x_2)$. (b) Example of an adapted dyadic squares segmentation of an image and its geometric flow.

wavelet bases so that they follow the lines of flow, and by applying a *bandeletization* procedure that takes advantage of the image regularity along the geometric flow. The next section explains how to optimize this image segmentation and compute the flow over each region.

If there is no geometric flow over a region Ω , which indicates that the image restriction to Ω has an isotropic regularity, then this restriction is approximated in the separable wavelet basis (4) of $\mathbf{L}^2(\Omega)$. If a geometric flow is calculated in Ω , this wavelet basis is replaced by a bandelet basis. We first explain how to construct the bandelet basis when the flow is parallel in the vertical direction: $\vec{r}(x_1, x_2) = \vec{r}(x_1)$. We normalize the flow vectors so that it can be written $\vec{r}(x_1) = (1, c'(x_1))$. Let $x_{\min} = \inf_{x_1} \{(x_1, x_2) \in \Omega\}$. A *flow line* is defined as an integral curve of the flow, whose tangents are parallel to $\vec{r}(x_1)$. Since the flow is parallel vertically, a flow line associated to a fixed translation parameter x_2 is a set of point $(x_1, x_2 + c(x_1)) \in \Omega$ for x_1 varying, with

$$c(x) = \int_{x_{\min}}^x c'(u) du.$$

By construction of the flow, the image gray level has regular variations along these flow lines, and hence the warped image

$$Wf(x_1, x_2) = f(x_1, x_2 + c(x_1))$$

is regular along the horizontal lines for x_2 fixed and x_1 varying. As a consequence, if $\Psi(x_1, x_2)$ is a wavelet having several vanishing moments along x_1 for each x_2 fixed, then the inner product

$$\langle Wf, \Psi \rangle = \langle f, W^*\Psi \rangle \quad (10)$$

has a small amplitude. The warping operator W is an orthogonal operator since its adjoint is equal to its inverse

$$W^*f(x_1, x_2) = W^{-1}f(x_1, x_2) = f(x_1, x_2 - c(x_1)).$$

Let us write $W\Omega = \{(x_1, x_2) : (x_1, x_2 + c(x_1)) \in \Omega\}$. Since $W^* = W^{-1}$, (10) suggests decomposing f over a family of *warped wavelets* obtained by applying W^{-1} to each wavelet of an orthonormal basis of $\mathbf{L}^2(W\Omega)$

$$\left\{ \begin{array}{l} \phi_{j,m_1}(x_1) \psi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1) \phi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1) \psi_{j,m_2}(x_2) \end{array} \right\}_{(j,m_1,m_2) \in \mathbf{I}_{W\Omega}}. \quad (11)$$

Since W^{-1} is orthogonal, this yields a warped wavelet orthonormal basis of $\mathbf{L}^2(\Omega)$

$$\left\{ \begin{array}{l} \phi_{j,m_1}(x_1) \psi_{j,m_2}(x_2 - c(x_1)), \\ \psi_{j,m_1}(x_1) \phi_{j,m_2}(x_2 - c(x_1)), \\ \psi_{j,m_1}(x_1) \psi_{j,m_2}(x_2 - c(x_1)) \end{array} \right\}_{(j,m_1,m_2) \in \mathbf{I}_{W\Omega}}. \quad (12)$$

As explained, the regularity of f along the flow lines implies that a warped wavelet coefficient $\langle f, W^{-1}\Psi \rangle$ is small if $\Psi(x_1, x_2)$ has vanishing moments along x_1 for each x_2 . This is valid for $\Psi(x_1, x_2) = \psi_{j,m_1}(x_1)\phi_{j,m_2}(x_2)$ and also for $\Psi(x_1, x_2) = \psi_{j,m_1}(x_1)\psi_{j,m_2}(x_2)$ because the 1-D wavelet $\psi(t)$ has several vanishing moments, but not for $\Psi(x_1, x_2) = \phi_{j,m_1}(x_1)\psi_{j,m_2}(x_2)$ because the scaling function $\phi(t)$ has no vanishing moment. It is, therefore, necessary to replace the family of orthogonal scaling functions $\{\phi_{j,m_1}(x_1)\}_{m_1}$ by an equivalent family of orthonormal functions, that have vanishing moments.

We know that the collection of scaling function $\{\phi_{j,m_1}(x_1)\}_{m_1}$ is an orthonormal basis of a multiresolution space which also admits an orthonormal basis of wavelets $\{\psi_{l,m_1}(x_1)\}_{l>j,m_1}$ [5]. This suggests replacing the orthogonal family $\{\phi_{j,m_1}(x_1)\psi_{j,m_2}(x_2)\}_{j,m_1,m_2}$ by the family $\{\psi_{l,m_1}(x_1)\psi_{j,m_2}(x_2)\}_{j,l>j,m_1,m_2}$ which generates the same space. This is called a *bandeletization*. We shall see that it is implemented with a simple discrete wavelet transform. The functions $\psi_{l,m_1}(x_1)\psi_{j,m_2}(x_2)$ are called bandelets because their support is parallel to the flow lines and is more elongated ($2^l > 2^j$) in the direction of the geometric flow. Inserting these bandelets in the warped wavelet basis (12) yields a bandelet orthonormal basis of $\mathbf{L}^2(\Omega)$

$$\left\{ \begin{array}{l} \psi_{l,m_1}(x_1) \psi_{j,m_2}(x_2 - c(x_1)), \\ \psi_{j,m_1}(x_1) \phi_{j,m_2}(x_2 - c(x_1)), \\ \psi_{j,m_1}(x_1) \psi_{j,m_2}(x_2 - c(x_1)) \end{array} \right\}_{j,l>j,m_1,m_2}. \quad (13)$$

If the geometric flow in Ω is parallel in the horizontal direction, meaning that

$$\vec{r}(x_1, x_2) = \vec{r}(x_2) = (c'(x_2), 1)$$

then the same construction applies by inverting the roles of the variables x_1 and x_2 . Let $x_{\min} = \inf_{x_2} \{(x_1, x_2) \in \Omega\}$ and $c(x) = \int_{x_{\min}}^x c'(u) du$. A warped wavelet basis is constructed from a wavelet basis of

$$W\Omega = \{(x_1, x_2) : (x_1 + c(x_2), x_2) \in \Omega\}$$




Fig. 3. Example of dyadic square image segmentation. Each leaf of the corresponding quad tree corresponds to a square region having the same index number.

and is defined by

$$\left\{ \begin{array}{l} \phi_{j,m_1}(x_1 - c(x_2)) \psi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1 - c(x_2)) \phi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1 - c(x_2)) \psi_{j,m_2}(x_2) \end{array} \right\}_{(j,m_1,m_2) \in \mathbf{I}_{W\Omega}} \quad (14)$$

The bandeletization replaces each family of scaling functions $\{\phi_{j,m_2}(x_2)\}_{m_2}$ by a family of orthonormal wavelets that generates the same space. The resulting bandelet orthonormal basis of $\mathbf{L}^2(\Omega)$ is

$$\left\{ \begin{array}{l} \phi_{j,m_1}(x_1 - c(x_2)) \psi_{j,m_2}(x_2), \\ \psi_{j,m_1}(x_1 - c(x_2)) \psi_{l,m_2}(x_2), \\ \psi_{j,m_1}(x_1 - c(x_2)) \psi_{j,m_2}(x_2) \end{array} \right\}_{j,l > j, m_1, m_2} \quad (15)$$

Given a partition of the image support $\mathcal{S} = \cup_i \Omega_i$ with the corresponding geometric flow, this strategy defines a bandelet or wavelet (if there is no flow) orthonormal basis in each $\mathbf{L}^2(\Omega_i)$. The union of these bases is a block orthonormal basis of $\mathbf{L}^2(\mathcal{S})$ [2]. The orthogonality of the wavelet and bandelet bases can also be relaxed. If the 1-D wavelet ψ and the scaling function ϕ yield a biorthogonal wavelet basis [25] then the same construction defines a biorthogonal bandelet basis of each $\mathbf{L}^2(\Omega_i)$ [1].

B. Optimized Geometry for Approximations

A major difficulty is to compute an appropriate image geometry. For image approximation, the best geometry is the one that leads to an approximation f_M from M parameters that minimizes the approximation error $\|f - f_M\|$. In a bandelet representation, the M parameters include the bandelet coefficients that are used to compute f_M as well as the parameters that specify the image partition and the geometric flow in each region.

To represent the image partition with few parameters, and to be able to compute an optimal partition with a fast algorithm, we restrict ourselves to partitions in squares of varying dyadic sizes. A dyadic square image segmentation is obtained by successive subdivisions of square regions into four squares of twice smaller width. For a square image support of width L , a square region of width $L 2^{-j}$ is represented by a node at the depth j of a quad tree. A square subdivided into four smaller squares corresponds to a node having four children in the quad tree. Fig. 3 gives an example of a dyadic square image segmentation with the corresponding quad tree.

In each region, Ω of the segmentation, one must decide if there should be a geometric flow, if this flow should be parallel in the horizontal or in the vertical direction, and the exact value of this flow. The flow can be written $\vec{\tau}(x_1, x_2) = (1, c'(x_1))$ if

it is parallel in the vertical direction and $\vec{\tau}(x_1, x_2) = (c'(x_2), 1)$ if it is parallel in the horizontal direction. To represent such flows with few parameters, $c'(t)$ is calculated as an expansion over translated B-spline functions dilated by a scale factor 2^l . A B-spline $b(t)$ of degree m is obtained by convolving the indicator function $\mathbf{1}_{[-1/2, 1/2]}$ with itself m times [26]. Over a square Ω of width 2^k , the flow at a scale 2^l is characterized by 2^{k-l} coefficients α_n

$$c'(t) = \sum_{n=1}^{2^{k-l}} \alpha_n b(2^{-l}t - n). \quad (16)$$

The scale parameter 2^l is adjusted through a global optimization of the geometry.

Let M_G be the total number of geometric parameters that specify an image partition $\mathcal{S} = \cup_i \Omega_i$ and a B-spline flow in each region Ω_i . These M_G parameters define a block bandelet basis constructed over this partition. For a threshold T , a bandelet approximation f_M is reconstructed from the M_B bandelet coefficients above T , and the total number of parameters is

$$M = M_G + M_B .$$

Optimizing this approximation means finding a bandelet basis that minimizes $\|f - f_M\|$ for a fixed number of parameters $M = M_B + M_G$. Similarly to the strategy used by Donoho [27], we find a best bandelet basis that minimizes the Lagrangian

$$\mathcal{L}(f, T) = \|f - f_M\|^2 + T^2 M \quad \text{with} \quad M = M_B + M_G \quad (17)$$

The Lagrangian multiplier is T^2 because at the threshold level the squared amplitude T^2 of a bandelet coefficient should increase the Lagrangian by the same amount as an increase by 1 of the number M_B of bandelet coefficients.

Suppose that the image f has contours that are \mathbf{C}^α curves which meet at corners or junctions, and that f is \mathbf{C}^α away from these curves. Although α is an unknown parameter, this procedure leads to a bandelet approximation that has an optimal asymptotic error decay rate

$$\|f - f_M\|^2 \leq C M^{-\alpha} . \quad (18)$$

This result still holds if the contours of f are blurred by an unknown regular smoothing kernel. The proof can be found in [1] and [2] with several variations on the construction of bandelets and on the optimization procedure that minimizes the Lagrangian (17). The best bandelet search adjusts the number of parameters describing the geometric flow in each region, so that

it follows the contours with enough precision. There should typically be a flow line which remains at a distance $O(M^{-\alpha})$ from the contour in each region.

For image compression, the bandelet coefficients and the B-spline coefficients of the flow (16) are quantized and binary encoded, and the image segmentation quadtree is coded. If f has contours which are C^α curves and is C^α away from these curves then [1] proves that for a total of R bits, the image \tilde{f} reconstructed from the quantized bandelet coefficients has a distortion

$$D(R) = \|f - \tilde{f}\|^2 \leq C R^{-\alpha} |\log_2 R|^{\alpha+1}.$$

This distortion-rate decay is close to the Kolmogorov lower bound which is $O(R^{-\alpha})$ for functions that are C^α over their whole support. We now concentrate on discrete fast algorithms. Numerical applications to image compression and noise removal are described in the following sections.

IV. FAST DISCRETE BANDELET TRANSFORM

Bandelets in a region Ω are computed by applying a bandeletization to warped wavelets, which are separable along a fixed direction (horizontal or vertical) and along the flow lines as long as they remain away from the boundary of Ω . A fast discrete bandelet transform can, therefore, be computed from a fast separable wavelet transform along this fixed direction and along the image flow lines. The block bandelet basis of Section III-A is constructed with a warped wavelet basis inside each region. When modifying bandelet coefficients in image processing applications, discontinuities appear along the region boundaries. To avoid these boundary effects, we define a discrete warped wavelet transform which goes across the region boundaries while keeping perfect reconstruction properties and two vanishing moments. No condition is imposed on the shapes of the regions.

The fast discrete bandelet transform associated to an image partition $\cup_i \Omega_i$ includes three steps:

- 1) a resampling, that computes the image sample values along the flow lines in each region Ω_i of the partition;
- 2) a warped wavelet transform with a subband filtering along the flow lines, which goes across the region boundaries;
- 3) a bandeletization that transforms the warped wavelet coefficient to compute bandelet coefficients along the flow lines.

The fast inverse bandelet transform includes the three inverse steps:

- 1) an inverse bandeletization that recovers the warped wavelet coefficient along the flow lines;
- 2) an inverse warped wavelet transform with an inverse subband filtering;
- 3) an inverse resampling which computes the image samples along the original grid from the samples along the flow lines in each region Ω_i

The following three sections describe fast algorithms that implement these three steps, with $O(N^2)$ operations for an image of N^2 pixels.

A. Resampling Along the Geometric Flow

The first step of the discrete bandelet transform computes the image sample values along the flow lines, in each region Ω_i of the partition. We describe its implementation together with the inverse resampling.

In a discrete framework, the geometric flow in a region Ω_i is a vector field $\vec{\tau}_i[n_1, n_2]$ defined over the image sampling grid. If the flow is parallel vertically then

$$\vec{\tau}_i[n_1, n_2] = \vec{\tau}_i[n_1] = (1, c'_i[n_1]) \quad (19)$$

where $c'_i[n_1]$ measures an average relative displacement of the image gray levels in Ω_i along the line n_1 with respect to the line $n_1 - 1$. A discretized flow line in Ω_i is a set of points of coordinates $(k_1, k_2 + c_i[k_1]) \in \Omega_i$ for a fixed integer k_2 and a varying integer k_1 , with

$$c_i[k] = \sum_{p=a_i}^k c'_i[p] \quad (20)$$

and $a_i = \min_{n_1} \{(n_1, n_2) \in \Omega_i\}$. The coordinates of flow lines are stored in a sampling grid array defined for each $(k_1, k_2) \in \mathbb{Z}^2$ by

$$G_i[k_1, k_2] = (k_1, k_2 + c_i[k_1]) \text{ if } (k_1, k_2 + c_i[k_1]) \in \Omega_i$$

and $G_i[k_1, k_2] = \text{nil}$, otherwise.

If the geometric flow is parallel horizontally in Ω_i , then $\vec{\tau}_i[n_1, n_2] = (c'_i[n_2], 1)$. Each flow line is defined by $(k_1 + c_i[k_2], k_2)$ for a fixed k_1 and varying k_2 , where $c_i[k]$ is still defined by (20) with $a_i = \min_{n_2} \{(n_1, n_2) \in \Omega_i\}$. The coordinates of these flow lines are stored in

$$G_i[k_1, k_2] = (k_1 + c_i[k_2], k_2) \text{ if } (k_1 + c_i[k_2], k_2) \in \Omega_i$$

and $G_i[k_1, k_2] = \text{nil}$, otherwise.

Given the original image sample values $f[n_1, n_2]$, at each grid point $G_i[k_1, k_2]$, the resampling computes an interpolated image value that is written $V_i[k_1, k_2]$. For a flow parallel vertically, the grid points $(k_1, k_2 + c_i[k_1]) \in \Omega_i$ are obtained with 1-D translations along x_2 of the integer sampling grid $(n_1, n_2) \in \Omega_i$. If the flow is parallel horizontally, then the 1-D translation is along the x_1 direction.

A 1-D translation by $\tau \in (-1/2, 1/2]$ of a discrete signal $a[n]$ for $1 \leq n \leq P$ is implemented by an operator T_τ which performs an interpolation. This interpolation can be written

$$T_\tau a[n] = \sum_{p=1}^P a[p] \rho_p(n - \tau) \quad (21)$$

where each $\rho_p(t)$ has a support in $[1/2, P+1/2]$ with $\rho_p(p) = 1$ and $\rho_p(n) = 0$ if $n \neq p$ is an integer. In all numerical experiments, this interpolation operator is implemented with cubic splines, using the recursive filtering procedure of Blu *et al.* [26].

For flow parallel vertically, for each k_1 fixed, the grid points $G_i[k_1, k_2] = (k_1, k_2 + c_i[k_1]) \in \Omega_i$ are obtained by translating the points in the integer grid column $(k_1, n_2) \in \Omega_i$ by a sub-pixel shift $\tau[k_1] = n_2 - k_2 - c_i[k_1] \in (-1/2, 1/2]$. The interpolated image values $V_i[k_1, k_2]$ are, thus, obtained by applying the translation operator $T_{\tau[k_1]}$ to each segment of the image column $f[k_1, n_2]$ in Ω_i . If the flow is parallel horizontally, we fix k_2 and the grid points $G_i[k_1, k_2] = (k_1 + c_i[k_2], k_2) \in \Omega_i$ are

obtained by translating the points $(n_1, k_2) \in \Omega_i$ by $\tau[k_2] = n_1 - k_1 - c_i[k_2] \in (-1/2, 1/2]$. The values $V_i[k_1, k_2]$ are then computed by applying $T_{\tau[k_2]}$ to each segment of the image line $f[n_1, k_2]$ in Ω_i .

The inverse discrete bandelet transform computes the image values on the original integer sampling grid (n_1, n_2) from the sample values $V_i[k_1, k_2]$ along the flow lines in each Ω_i . This requires to invert the discrete translation operator T_{τ} . However, unless the interpolation functions $\rho_p(t)$ are periodized sinc functions over the interval, the inverse T_{τ}^{-1} of T_{τ} is an unstable operator which amplifies the highest signal frequencies. We, thus, rather approximate T_{τ}^{-1} by $T_{-\tau}$. Using this stable approximation, the image columns or rows in Ω_i are calculated by applying $T_{-\tau}$ with appropriate values of τ to the rows or columns of $V_i[k_1, k_2]$.

Since T_{τ}^{-1} is approximated by $T_{-\tau}$, the inverse resampling does not recover the original image values. The error depends upon the choice of the interpolation functions $\rho_p(t)$ in (21). Following the analysis of Blu *et al.* [26], cubic splines are chosen because they introduce small errors and produce hardly visible Gibbs-type oscillatory artifacts.

B. Discrete Warped Wavelet and Wavelet Packet Transform

This section explains how to adapt the fast wavelet transform algorithm to compute a warped wavelet transform with a geometric flow computed over a partition of the image support. To avoid creating boundary effects, warped wavelet coefficients are calculated with a subband filtering that goes across the boundaries of the image partition, with an adapted lifting scheme introduced by Bernard [28]. At the boundaries, warped wavelets still have two vanishing moments.

The wavelet coefficients of a discrete image $f[n_1, n_2]$ are computed with a filter bank that convolves the image rows and columns with a pair of perfect reconstruction filters $(h[n], g[n])$ together with a subsampling [29]. These wavelet coefficients are inner products of $f[n_1, n_2]$ with a basis of discrete separable wavelets

$$\left\{ \begin{array}{l} \psi_{j,m_1}[n_1] \phi_{j,m_2}[n_2], \\ \phi_{j,m_1}[n_1] \psi_{j,m_2}[n_2], \\ \psi_{j,m_1}[n_1] \psi_{j,m_2}[n_2] \end{array} \right\}_{j,m_1,m_2}. \quad (22)$$

In the following, we consider the more general case of biorthogonal wavelet bases, where the inverse transform is implemented with a dual pair of filters $(\tilde{h}[n], \tilde{g}[n])$ [29]. All filters are supposed to have a finite impulse response, and we choose the 7–9 CDF filters [25] in all numerical examples.

A warped wavelet transform decomposes the image in a family of warped wavelets in each region Ω_i . Let $c_i[p]$ be an integral curve of the flow in Ω_i , as defined in (20). If the flow is parallel vertically then the warped wavelets in Ω_i can be written

$$\left\{ \begin{array}{l} \psi_{j,m_1}[n_1] \phi_{j,m_2}[n_2 - c_i[n_1]], \\ \phi_{j,m_1}[n_1] \psi_{j,m_2}[n_2 - c_i[n_1]], \\ \psi_{j,m_1}[n_1] \psi_{j,m_2}[n_2 - c_i[n_1]] \end{array} \right\}_{j,m_1,m_2}. \quad (23)$$

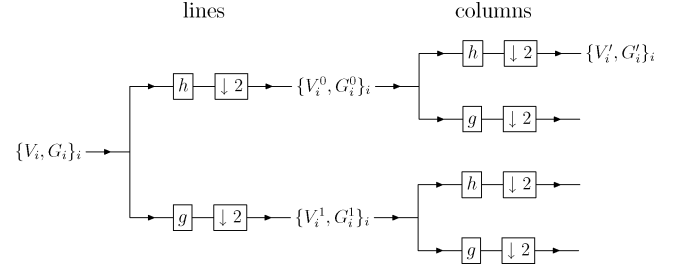


Fig. 4. Warped wavelet transform filters and subsamples by the lines and columns of input values $\{V_i[k_1, k_2]\}_i$, with a lifting scheme that is adapted to the flow sampling grids $\{G_i[k_1, k_2]\}_i$. The same computational block is applied to the output $\{V_i', G_i'\}_i$ of the low-pass filter h .

If the flow is parallel horizontally, then the warped wavelets in Ω_i are

$$\left\{ \begin{array}{l} \psi_{j,m_1}[n_1 - c_i[n_2]] \phi_{j,m_2}[n_2], \\ \phi_{j,m_1}[n_1 - c_i[n_2]] \psi_{j,m_2}[n_2], \\ \psi_{j,m_1}[n_1 - c_i[n_2]] \psi_{j,m_2}[n_2] \end{array} \right\}_{j,m_1,m_2}. \quad (24)$$

Suppose that the flow is parallel horizontally. Since

$$\langle f[n_1, n_2], \Psi[n_1 - c_i[n_2], n_2] \rangle = \langle f[n_1 + c_i[n_2], n_2], \Psi[n_1, n_2] \rangle$$

the image coefficients in the warped wavelet basis (24) are obtained by decomposing the translated image values $V_i[k_1, k_2]$ at the locations $G_i[k_1, k_2] = (k_1 + c_i[k_2], k_2)$ in the separable wavelet basis (22). These wavelet coefficients are, thus, computed by applying the separable wavelet filter bank algorithm along the “lines” and “columns” of the resampled images $V_i[k_1, k_2]$. The same applies to a flow parallel vertically.

In the following, we concentrate on the elementary computational block corresponding to a 1-D warped filtering and subsampling using the filters (h, g) together with its inverse transform. Depending upon the filter bank that cascades these 1-D subband filterings, one can compute a warped wavelet transform or a more general warped wavelet packet transform by using the corresponding wavelet packet filter bank [29], [30]. Fig. 4 illustrates the elementary computational block of a wavelet transform, which performs a subband filtering and subsampling along the lines and columns with (h, g) . A wavelet transform filtering tree applies the same computational block to the output of the filtering by the low-pass filter h along the lines and columns.

Let us concentrate on a filtering and subsampling along the horizontal lines, indexed by varying integers k_1 for k_2 fixed. The same procedure applies to a subband filtering in the vertical direction by exchanging the roles of the horizontal variable k_1 and the vertical variable k_2 . The input of the 1-D subband filtering is a family of sampling grids $\{G_i[k_1, k_2]\}_i$ and their sample values $\{V_i[k_1, k_2]\}_i$, as illustrated in Fig. 4. The output includes two sets of subsampled grids $\{G_i^0[k_1, k_2]\}_i$ and $\{G_i^1[k_1, k_2]\}_i$ that are called respectively even and odd grids, and are defined by

$$G_i^0[k_1, k_2] = G_i[2k_1, k_2] \quad \text{and} \quad G_i^1[k_1, k_2] = G_i[2k_1 + 1, k_2]. \quad (25)$$

We now explain the calculation of the corresponding subband sample values $\{V_i^0[k_1, k_2]\}_i$ and $\{V_i^1[k_1, k_2]\}_i$.

Let us first consider an “inside” point $(2k_1, k_2)$ of a grid G_i such that the supports of h and g centered at this point are entirely included in this same grid. This means that $G_i[2k_1 + l, k_2] \neq \text{nil}$ for all l such that $h[l] \neq 0$ or $g[l] \neq 0$. The horizontal subband filtering of V_i at this location is computed with a standard convolution and subsampling formula along the variable k_1

$$V_i^0[k_1, k_2] = \sum_l h[l - 2k_1] V_i[l, k_2] \quad (26)$$

and

$$V_i^1[k_1, k_2] = \sum_l g[l - 2k_1] V_i[l, k_2]. \quad (27)$$

These inside coefficients are recovered from the subband coefficients with the dual filters

$$V_i[k_1, k_2] = \sum_l \tilde{h}[k_1 - 2l] V_i^0[l, k_2] + \sum_l \tilde{g}[k_1 - 2l] V_i^1[l, k_2]. \quad (28)$$

The main difficulty is to implement a phase-aligned warped subband filtering that computes the inside coefficients with (26) and (27) and which adapts the filtering across the boundaries of sampling grids, while remaining invertible and stable. The bandpass filtering corresponding to (27) should also keep its vanishing moments [29] so that regular signals produce wavelet coefficients of small amplitude. This is achieved with a lifting scheme which is adapted at the boundary of each region.

Let us first consider the case of inside points within each grid. Daubechies and Sweldens [31] have proved that the subband filterings (26) and (27) can be factored into a sequence of lifting steps and a scaling. A lifting is computed with predicting and updating operations that involve the two neighbors of each point. In the horizontal direction, the left and right neighbors of an even grid point $G_i^0[k_1, k_2]$ are odd grid points defined by

$$LG_i^0[k_1, k_2] = G_i^1[k_1 - 1, k_2] \text{ and } RG_i^0[k_1, k_2] = G_i^1[k_1, k_2]. \quad (29)$$

The left and right neighbors of an odd grid point $G_i^1[k_1, k_2]$ are even grid points defined by

$$LG_i^1[k_1, k_2] = G_i^0[k_1, k_2] \text{ and } RG_i^1[k_1, k_2] = G_i^0[k_1 + 1, k_2]. \quad (30)$$

Let us consider sample values $\tilde{V}_i^0[k_1, k_2]$ and $\tilde{V}_i^1[k_1, k_2]$ associated to the subsampled grids $G_i^0[k_1, k_2]$ and $G_i^1[k_1, k_2]$. The left and right neighborhood values are calculated according (29)

$$L\tilde{V}_i^0[k_1, k_2] = \tilde{V}_i^1[k_1 - 1, k_2], \quad R\tilde{V}_i^0[k_1, k_2] = \tilde{V}_i^1[k_1, k_2]$$

and according to (30)

$$L\tilde{V}_i^1[k_1, k_2] = \tilde{V}_i^0[k_1, k_2], \quad R\tilde{V}_i^1[k_1, k_2] = \tilde{V}_i^0[k_1 + 1, k_2].$$

A symmetric predicting operator of parameter α computes $(V_i^0, V_i^1) = P_\alpha(\tilde{V}_i^0, \tilde{V}_i^1)$ defined by

$$\begin{aligned} V_i^0[k_1, k_2] &= \tilde{V}_i^0[k_1, k_2] \\ V_i^1[k_1, k_2] &= \tilde{V}_i^1[k_1, k_2] + \alpha(L\tilde{V}_i^1[k_1, k_2] + R\tilde{V}_i^1[k_1, k_2]). \end{aligned} \quad (31)$$

Its inverse is $P_\alpha^{-1} = P_{-\alpha}$. A symmetric updating operator of parameter β computes $(V_i^0, V_i^1) = U_\beta(\tilde{V}_i^0, \tilde{V}_i^1)$ defined by

$$\begin{aligned} V_i^0[k_1, k_2] &= \tilde{V}_i^0[k_1, k_2] + \beta(L\tilde{V}_i^0[k_1, k_2] + R\tilde{V}_i^0[k_1, k_2]) \\ V_i^1[k_1, k_2] &= \tilde{V}_i^1[k_1, k_2]. \end{aligned} \quad (32)$$

Its inverse is $U_\beta^{-1} = U_{-\beta}$. A scaling operator of parameter ξ computes $(V_i^0, V_i^1) = S_\xi(\tilde{V}_i^0, \tilde{V}_i^1)$ defined by

$$V_i^0[k_1, k_2] = \xi \tilde{V}_i^0[k_1, k_2] \text{ and } V_i^1[k_1, k_2] = \frac{\tilde{V}_i^1[k_1, k_2]}{\xi}.$$

Its inverse is $S_\xi^{-1} = S_{1/\xi}$. The lifting is initiated by a grid splitting $(V_i^0, V_i^1) = \text{Split}(V_i)$ with

$$V_i^0[k_1, k_2] = V_i[2k_1, k_2] \text{ and } V_i^1[k_1, k_2] = V_i[2k_1 + 1, k_2].$$

The inverse is computed by $V_i = \text{Union}(V_i^0, V_i^1)$.

For 7–9 CDF filters, Daubechies and Sweldens have proved [31] that the subband filtering formula (26) and (27) are implemented by the following lifting steps:

$$(V_i^0, V_i^1) = S_\xi U_\delta P_\gamma U_\beta P_\alpha \text{Split}(V_i) \quad (33)$$

with $\alpha \approx -1.5861$, $\beta \approx -0.0530$, $\gamma \approx 0.8829$, $\delta \approx 0.4435$, and $\xi \approx 1.1496$. The inverse of the lifting steps (33) that implements the subband reconstruction (28) is

$$V_i = \text{Union} P_{-\alpha} U_{-\beta} P_{-\gamma} U_{-\delta} S_{1/\xi} (V_i^0, V_i^1). \quad (34)$$

For points near the border of each sampling grid G_i , the subband filtering is calculated with a modified lifting scheme that goes across the boundaries of different grids. This requires to establish a neighborhood relation between sampling points of different grids. To build warped wavelets across regions boundaries that have two vanishing moments, we impose that the left and right neighbors of a point are aligned with this point.

We concentrate on left and right neighbors of even grid points $G_i^0[k_1, k_2]$. The left and right neighbors of odd grid point $G_i^1[k_1, k_2]$ are computed with the same procedure, by exchanging by the roles of even and odd grid points. The left and right neighbors of a point $G_i^0[k_1, k_2]$ inside the same grid of index i are defined by (29). Suppose that $G_i^0[k_1, k_2]$ has a right neighbor in the same grid, which means that $RG_i^0[k_1, k_2] = G_i^1[k_1, k_2] \neq \text{nil}$, but no left neighbor in this grid because $G_i^1[k_1 - 1, k_2] = \text{nil}$. The left neighbor $LG_i^0[k_1, k_2]$ must, therefore, be defined outside this grid, and we shall impose that $G_i^0[k_1, k_2] - LG_i^0[k_1, k_2]$ is collinear and has the same direction as $RG_i^0[k_1, k_2] - G_i^0[k_1, k_2]$. It is calculated as a weighted average of two odd grid points

$$LG_i^0[k_1, k_2] = \epsilon G_{i'}^1[k_1', k_2'] + (1 - \epsilon) G_{i''}^1[k_1'', k_2''] \quad (35)$$

where $G_{i'}^1[k_1', k_2'] \neq \text{nil}$ and $G_{i''}^1[k_1'', k_2''] \neq \text{nil}$ are two points which are on the left and as close as possible to $G_i^0[k_1, k_2]$. The factor $\epsilon \in [0, 1]$ is adjusted so that the triplet $(RG_i^0[k_1, k_2], G_i^0[k_1, k_2], LG_i^0[k_1, k_2])$ is aligned. This is illustrated by Fig. 5. The corresponding left neighbor value is

$$L\tilde{V}_i^0[k_1, k_2] = \epsilon \tilde{V}_{i'}^1[k_1', k_2'] + (1 - \epsilon) \tilde{V}_{i''}^1[k_1'', k_2'']. \quad (36)$$




Fig. 5. Crosses and circles correspond, respectively, to even and odd grid points $G_i^0[k_1, k_2]$ and $G_i^1[k_1, k_2]$. Left and right neighbors are indicated by arrows. Across the boundary of a region, the left neighbor $LG_i^0[k_1, k_2]$ of an even grid point $G_i^0[k_1, k_2]$ is shown as a triangle. The point $LG_i^0[k_1, k_2]$ is aligned with $G_i^0[k_1, k_2]$ and its right neighbor $RG_i^0[k_1, k_2]$, and it is on a line between two odd grid points of another region, shown as circles.

Similarly, suppose that $G_i^0[k_1, k_2]$ has a left neighbor of position $LG_i^0[k_1, k_2] = G_i^1[k_1 - 1, k_2] \neq \text{nil}$ but no right neighbor in the same grid. The right neighbor $RG_i^0[k_1, k_2]$ is calculated so that $G_i^0[k_1, k_2] - RG_i^0[k_1, k_2]$ is collinear and has the same direction as $LG_i^0[k_1, k_2] - G_i^0[k_1, k_2]$. It is obtained as a weighted average of two odd grid points that are on the right of $G_i^0[k_1, k_2]$

$$RG_i^0[k_1, k_2] = \epsilon G_{i'}^1[k_1', k_2'] + (1 - \epsilon) G_{i''}^1[k_1'', k_2''] \quad (37)$$

where $\epsilon \in [0, 1]$ is adjusted so that the triplet $(RG_i^0[k_1, k_2], G_i^0[k_1, k_2], LG_i^0[k_1, k_2])$ is aligned. The corresponding value is

$$R\tilde{V}_i^0[k_1, k_2] = \epsilon \tilde{V}_{i'}^1[k_1', k_2'] + (1 - \epsilon) \tilde{V}_{i''}^1[k_1'', k_2'']. \quad (38)$$

Observe that if the value $\tilde{V}_i^0[k_1, k_2]$ is an affine function of its position $G_i^0[k_1, k_2]$ then $L\tilde{V}_i^0[k_1, k_2]$ and $R\tilde{V}_i^0[k_1, k_2]$ are also affine functions of their positions.

If $G_i^0[k_1, k_2]$ has a no left and no right neighbor in the grid of index i , then these left and right neighbors are computed as weighted averages of odd grid points with (35) and (37), and each factor ϵ is calculated by imposing that $G_i^0[k_1, k_2] - LG_i^0[k_1, k_2]$ and $RG_i^0[k_1, k_2] - G_i^0[k_1, k_2]$ are horizontal vectors (in the case of a horizontal filtering).

With these left and right neighborhood relations, using the predicting operator and updating operators defined by (31) and (32), the lifting (33) implements a subband filtering across the grid boundaries, whose restriction inside each grid is a standard subband filtering. The inverse operator is still given by (34). However, across the boundary of each region, the resulting linear operator that computes the bandpass coefficients has only one vanishing moment, because the predicting and updating parameters do not take into account the fact that the distance may vary between neighbors in different regions. To maintain two vanishing moments across the boundaries of different regions, we use the lifting scheme of Bernard [28], which modifies the predicting and updating parameters α and β , according to the distance between the sampling points. We denote by $\|x - x'\|$ the Euclidean distance between two points in \mathbb{R}^2 . According to [28], the prediction (31) is replaced by

$$V_i^1[k_1, k_2] = \tilde{V}_i^1[k_1, k_2] + 2\alpha \left(\epsilon_1 L\tilde{V}_i^1[k_1, k_2] + (1 - \epsilon_1) R\tilde{V}_i^1[k_1, k_2] \right) \quad (39)$$

with

$$\epsilon_1 = \frac{\|RG_i^1[k_1, k_2] - G_i^1[k_1, k_2]\|}{\|RG_i^1[k_1, k_2] - G_i^1[k_1, k_2]\| + \|LG_i^1[k_1, k_2] - G_i^1[k_1, k_2]\|}.$$

Similarly, the update (32) is replaced by

$$V_i^0[k_1, k_2] = \tilde{V}_i^0[k_1, k_2] + 2\beta \left(\epsilon_0 L\tilde{V}_i^0[k_1, k_2] + (1 - \epsilon_0) R\tilde{V}_i^0[k_1, k_2] \right) \quad (40)$$

with

$$\epsilon_0 = \frac{\|RG_i^0[k_1, k_2] - G_i^0[k_1, k_2]\|}{\|RG_i^0[k_1, k_2] - G_i^0[k_1, k_2]\| + \|LG_i^0[k_1, k_2] - G_i^0[k_1, k_2]\|}.$$

The total number of operations to implement a warped wavelet transform with this modified lifting scheme is at most twice larger than the number of operations to compute a standard separable wavelet transform with a lifting scheme. It, thus, requires $O(N^2)$ operations for an image of N^2 pixels.

We are now going to show that this modified lifting scheme implements a discrete warped wavelet transform which has two vanishing moments, also at the boundary of regions, if the geometric flow has a fixed direction in each region. If the direction of the geometric flow remains constant in each region, but may vary from region to region, our construction of left and right neighbors implies that all triplets of points $(RG_i^1[k_1, k_2], G_i^1[k_1, k_2], LG_i^1[k_1, k_2])$ and $(RG_i^0[k_1, k_2], G_i^0[k_1, k_2], LG_i^0[k_1, k_2])$ are aligned in the plane. To prove that the warped wavelet transform has two vanishing moments, we must verify that a signal whose sample values are an affine functions of their positions (irregularly sampled) produces warped wavelet coefficients that are zeros. If $\tilde{V}_i^1[k_1, k_2]$ and $\tilde{V}_i^0[k_1, k_2]$ are linear functions of their positions $G_i^1[k_1, k_2]$ and $G_i^0[k_1, k_2]$, since the $(RG_i^1[k_1, k_2], G_i^1[k_1, k_2], LG_i^1[k_1, k_2])$ and $(RG_i^0[k_1, k_2], G_i^0[k_1, k_2], LG_i^0[k_1, k_2])$ are aligned and the corresponding values are affine functions of these positions, the prediction and updating operators (39) and (40) compute $V_i^1[k_1, k_2] = (1 + 2\alpha)\tilde{V}_i^1[k_1, k_2]$ and $V_i^0[k_1, k_2] = (1 + 2\beta)\tilde{V}_i^0[k_1, k_2]$. The output values of the prediction and update operators are, thus, independent from the position of the left and right neighbors and is, therefore, the same when all sample values are on a uniform grid. When all samples are on a uniform grid, a lifting implementing a 7–9 wavelet transform produces wavelets coefficients that are zeros when the image is affine, because these wavelets have two vanishing moments. For irregularly sampled values, this result, thus, remains valid since the prediction and update operator outputs the same values when the image is affine.

Inside each region, a warped wavelet transform performs a 1-D wavelet transform along the lines of flow. Since the 7–9 wavelets have four vanishing moments, inside each region, the warped wavelet transform has four vanishing moments with respect to the geometric flow lines. In most cases, the direction of the geometric flow is discontinuous when going from one region to another, and the resulting warped wavelets across such a boundary are not differentiable. Yet, numerical experiments show that it creates hardly visible boundary artifacts when modifying the corresponding warped wavelet coefficients. In particular, we do not see boundary artifacts in compressed images such as the one shown in Fig. 9.

C. Bandedetization

To take advantage of the image regularity along the geometric flow, the bandedetization modifies a warped wavelet basis by transforming 1-D scaling functions into 1-D wavelets. The resulting bandelet coefficients are computed from warped wavelet coefficients with a 1-D discrete wavelet transform along the geometric flow lines.

Let us consider a region Ω_i in which the geometric flow is parallel vertically. The bandedetization is applied only to the warped wavelet coefficients with a scaling function along x_1

$$V_i[k_1, k_2] = \langle f[n_1, n_2], \phi_{j,k_1}[n_1] \psi_{j,k_2}[n_2 - c_i[n_1]] \rangle$$

because the scaling function $\phi_{j,k_1}[n_1]$ cannot take advantage of the geometric image regularity. The bandedetization performs a change of basis with a 1-D discrete wavelet transform along the parameter k_1 , which computes inner products with discrete bandelets at scales $2^l > 2^j$

$$\langle f[n_1, n_2], \psi_{l,p_1}[n_1] \psi_{j,k_2}[n_2 - c_i[n_1]] \rangle.$$

This 1-D wavelet transform is calculated with a wavelet filter bank [5], [29], with the filters (h, g) applied to $V_i[k_1, k_2]$ along the variable k_1 , for each k_2 fixed.

If the geometric flow is parameterized vertically in Ω_i then the bandedetization is applied to the warped wavelet coefficients

$$V_i[k_1, k_2] = \langle f[n_1, n_2], \psi_{j,k_1}[n_1 - c_i[n_2]] \phi_{j,k_2}[n_2] \rangle$$

to compute the bandelet coefficients at scales $2^l > 2^j$

$$\langle f[n_1, n_2], \psi_{j,k_1}[n_1 - c_i[n_2]] \psi_{l,p_2}[n_2] \rangle.$$

These bandelet coefficients are obtained with a one dimensional discrete wavelet transform of $V_i[k_1, k_2]$ along the variable k_2 , for each k_1 .

A bandedetization is computed within each region Ω_i , and not across the boundaries. Indeed, the geometric image regularity is established within each region, not across regions. Computing a bandedetization separately within each region Ω_i does not create boundary effects when processing these coefficients, because this transform is not applied on the image but on warped wavelet coefficients. The bandedetization transforms a biorthogonal warped wavelet basis into a biorthogonal bandelet basis [1].

The bandedetization can also be applied to any warped wavelet packet basis, to take advantage of the regularity of coefficients along the geometric flow. In a region Ω_i whose geometric flow is horizontal, the bandedetization should be applied to coefficients that are inner products with separable wavelet packets including the low-pass scaling signals $\phi_{j,k_1}[n_1]$ along the horizontal direction, to transform these scaling signals into wavelets $\psi_{l,p_1}[n_1]$ for $l > j$. If the geometric flow is vertical, bandedetization is performed on inner products with separable wavelet packets including the low-pass scaling signals $\phi_{j,k_2}[n_2]$ along the vertical direction, to produce wavelets $\psi_{l,p_2}[n_2]$ for $l > j$.

V. FAST GEOMETRIC OPTIMIZATION

A major difficulty of geometric representations is to adapt the geometry to local image structures. For a bandelet transform, the geometry is defined by the image partition in regions Ω_i

and by the geometric flow within each region. This segmented geometric flow is optimized for image compression and noise removal applications.

A. Image Compression

A bandelet transform code is implemented with a scalar quantization and an entropy coding of all coefficients. The geometry is computed by optimizing the resulting distortion rate, with a fast algorithm that requires $O(N^2(\log_2 N)^2)$ operations for an image of N^2 pixels. Numerical comparisons are made with a similar transform code in a wavelet basis.

Let $\mathcal{D} = \{\mathcal{B}^\gamma\}_{\gamma \in \Gamma}$ be the dictionary of all possible biorthogonal bandelet bases, where γ is a parameter that specifies the geometry of the basis. Finding the best geometry for image compression can be interpreted as a search for a best bandelet basis in the dictionary \mathcal{D} . Each bandelet basis is written $\mathcal{B}^\gamma = \{g_m^\gamma\}_{1 \leq m \leq N^2}$ and its biorthogonal basis is written $\tilde{\mathcal{B}}^\gamma = \{\tilde{g}_m^\gamma\}_{1 \leq m \leq N^2}$.

The transform code is implemented with a nearly uniform scalar quantizer $Q(x)$ with bins of size Δ , having a twice larger zero bin: $Q(x) = 0$ if $|x| \leq \Delta$ and $Q(x) = \text{sign}(x)(n+1/2)\Delta$ if $|x| \in [n\Delta, (n+1)\Delta)$ for $n \in \mathbb{N}^*$. The restored image from quantized coefficients is

$$\tilde{f} = \sum_{m=1}^{N^2} Q(\langle f, g_m^\gamma \rangle) \tilde{g}_m^\gamma$$

and the resulting distortion is $D = \|f - \tilde{f}\|^2$. The total number of bits R to code \tilde{f} is equal to the number of bits R_c to code the N^2 quantized coefficients $\{Q(\langle f, g_m^\gamma \rangle)\}_{1 \leq m \leq N^2}$ plus the number of bits to code the geometry of the basis. The distortion D , thus, depends upon R through the value of Δ and through the choice of the geometry.

In a discrete framework, the geometric flow in a region Ω_i is a vector field $\vec{\tau}_i[n_1, n_2]$ defined over the image sampling grid. If the flow is parallel vertically, then $\vec{\tau}_i[n_1, n_2] = \vec{\tau}_i[n_1] = (1, c'_i[n_1])$. As in (16), in a square of size 2^k , this displacement is parameterized by its decomposition coefficients α_n over a family of translated B-splines, which are dilated by a scale factor 2^l

$$c'_i[p] = \sum_{n=1}^{2^{k-l}} \alpha_n b(2^{-l}p - n). \quad (41)$$

In our calculations, we use a linear B-spline: $b(x) = 1 - |x|$ if $|x| < 1$ and $b(x) = 0$ if $|x| \geq 1$. The coefficients α_n are uniformly quantized. The quantization step adjusts the precision of the geometric displacement $c'_i[p]$. It is set to be of the order of 1/8 of a pixel.

To optimize the overall coder, similarly to Section III-B, we use the Lagrangian approach also proposed by Ramchandran and Vetterli [32]. It finds the best basis that minimizes $D(R) + \lambda R$, where λ is a Lagrange multiplier. If $D(R)$ is convex, which is usually the case, by letting λ vary we are guaranteed to minimize $D(R)$ for a fixed R . If $D(R)$ is not convex, then this strategy leads to a $D(R)$ that is at most a factor of 2 larger than the minimum. A new explicit formula is provided to relate λ to the quantization parameter Δ .

For a given image and parameter λ , we want to find the image segmentation $[1, N]^2 = \cup_i \Omega_i$ and the geometric flows in all Ω_i which define a bandelet basis that minimizes $D(R) + \lambda R$. Let us associate each bandelet vector to a single region Ω_i where its support is mostly located. This distortion rate can be decomposed into

$$D + \lambda R = \sum_i (D_i + \lambda R_i) \quad (42)$$

where $D_i = \|f - \tilde{f}\|_{\Omega_i}^2$ is the Euclidean norm restricted to a region Ω_i of the image partition, and R_i is the number of bits needed to code the bandelet coefficients and the geometry associated to Ω_i . It can be decomposed into

$$R_i = R_{s,i} + R_{g,i} + R_{c,i} \quad (43)$$

where $R_{s,i}$ is the number of bits to code the square Ω_i , $R_{g,i}$ to code the geometric flow in Ω_i , and $R_{c,i}$ to code the quantized bandelet coefficients in Ω_i . We now explain how to implement these coding procedures.

The image partition into dyadic squares is represented by a quad tree. Each leaf of the tree corresponds to a region Ω_i of the image partition. The position of each leaf in the quad tree is coded with $R_{s,i}$ bits, using a tree coding algorithm which codes each leaf with a binary word whose length increases with the depth of the leaf.

To code geometric flow in Ω_i with $R_{g,i}$ bits, we first code a variable which indicates if there is a flow and if it is parallel horizontally or vertically. If the geometric flow exists, it is specified by the 2^{k-l} coefficients α_n in (41), where 2^k is the width of the square Ω_i . The adapted scale parameter 2^l in (41) and the quantized coefficients are coded with fixed length codes.

Quantized bandelet coefficients are globally coded over the whole image with an adaptive arithmetic code. In a wavelet basis, as well as in a block cosine basis, it has been shown numerically and theoretically [3] that for most images the total number of bits to code the quantized coefficients is nearly proportional to the number of nonzero quantized coefficients. This remains valid for bandelet coefficients and we, thus, estimate the number of bits $R_{c,i}$ associated to each region Ω_i by

$$R_{c,i} = \gamma_0 M_i \quad (44)$$

where M_i is the number of nonzero quantized coefficients in Ω_i and $\gamma_0 \approx 7$.

The quantization step Δ is related to λ by observing that if $D + \lambda R$ is minimum then

$$\frac{\partial D}{\partial \Delta} = -\lambda \frac{\partial R}{\partial \Delta}. \quad (45)$$

Let $M = \sum_i M_i$ be the total number of nonzero quantized bandelet coefficients of the whole image. Since R depends upon Δ through the $R_{c,i}$, with (44), we verify that

$$\frac{\partial R}{\partial \Delta} = \sum_i \frac{\partial R_{c,i}}{\partial \Delta} = \gamma_0 \frac{\partial M}{\partial \Delta}. \quad (46)$$

When Δ varies, since all quantization bins are uniform outside the zero bin which is twice larger, one can also verify that the variation of D with Δ depends essentially upon the variation of the number M of coefficients which are not quantized to zero. A

coefficient of amplitude Δ is quantized to $\pm 3\Delta/2$, which produces a quadratic error of $\Delta^2/4$. If the quantization bin Δ increases, this same coefficient will be quantized to 0 which increases the quadratic error to Δ^2 , and adds $3\Delta^2/4$ to the distortion D . As a result

$$\frac{\partial D}{\partial \Delta} = -\frac{3\Delta^2}{4} \frac{\partial M}{\partial \Delta}.$$

Inserting this in (45) together with (46) gives

$$\lambda = \frac{3\Delta^2}{4\gamma_0}.$$

This relation specifies the Lagrange multiplier λ as a function of Δ that now remains the only parameter.

To minimize $D + \lambda R = \sum_i (D_i + \lambda R_i)$, we first compute the geometric flow which minimizes $D_i + \lambda R_i$ in each dyadic square Ω_i included in the image support. We shall then use a fast algorithm to find the best segmentation of $[1, N]^2$ which minimizes the overall distortion-rate sum. Let us first consider a square of width 2^k . The distortion rate $D_i + \lambda R_i$ may be minimized by a horizontally or vertically parallel flow or by no flow at all. Each possibility is tested. If the flow exists, it depends upon the scale parameter 2^l in (41) which must also be optimized. Since $1 \leq 2^l \leq 2^k$, the parameter l takes $k+1$ possible values. For a fixed value of l , one can verify that minimizing $D_i + \lambda R_i$ is equivalent to finding a flow that minimizes the distortion D_i for a given number M_i of nonzero quantized bandelet coefficients. This is achieved by finding a flow that follows the directions of geometric image regularity in Ω_i to produce a maximum number of small amplitude bandelet coefficients that are quantized to zero.

If the flow is parallel vertically, then it is defined by $\vec{\tau}_i[n_1, n_2] = (1, c'_i[n_1])$, where $c'_i[n_1]$ is the relative displacement of the image gray levels in Ω_i along the line n_1 with respect to the line $n_1 - 1$. Instead of finding the flow that exactly minimizes $D_i + \lambda R_i$, we compute a flow that is nearly optimal by minimizing the quadratic variation of the image along the flow

$$\mathcal{E}(\vec{\tau}) = \sum_{(n_1, n_2) \in \Omega_i} \left| f \star \theta[n_1, n_2] - f \star \theta[n_1 - 1, n_2 + c'_i[n_1]] \right|^2 \quad (47)$$

with $c'_i[p] = \sum_{n=1}^{2^{k-l}} \alpha_n b(2^{-l}p - n)$. The regularization filter $\theta[n_1, n_2]$ is a separable Gaussian which reduces the effect of the image noise. This minimization is performed with the Lucas and Kanade [33] gradient descent algorithm on the warping parameters α_n . Although the minimization of (47) does not exactly minimize the distortion $D_i + \lambda R_i$, it is proved in [2] that for images having a piecewise regular geometry, the resulting flow remains nearly optimal. If the geometric flow is parallel horizontally in Ω_i then the flow can be written $\vec{\tau}_i[n_1, n_2] = (c'_i[n_2], 1)$ and this flow vector is calculated with a similar procedure by exchanging the roles of the horizontal variable n_1 and the vertical variable n_2 . Whether the flow is parallel horizontally or vertically, with a fixed number of iterations in the Lucas and Kanade [33] gradient descent algorithm, the flow in a square Ω_i of 2^{2k} pixels is computed with $O(2^{2k})$ operations.

It now remains to find the best scale 2^l of the flow in each square Ω_i of size 2^k . The image is segmented into squares of




Fig. 6. (a) Geometric flow segmentation obtained for Barbara and $R = .44$ bits/pixel. (b) The bandelet reconstruction with a PSNR of 31.3 dB.




Fig. 7. Distortion in PSNR of the bandelet coder (full lines) and of the wavelet coder (dashed lines), for the Lena and Barbara images, as a function of the bit rate per pixel R/N^2 . Over all bit rates, the bandelet coder reduces the distortion by approximately .4 dB for Lena and by 1.4 dB for Barbara.

constant size 2^k . For a fixed 2^l , computing the best flow over each of these $N2^{-2k}$ square requires $O(N^2)$ operations. This segmentation and the corresponding flows define a bandelet basis. The fast bandelet transform computes resulting image bandelet coefficients with $O(N^2)$ operations. The quantization and coding procedure yields a value $D_i + \lambda R_i$ for each square Ω_i . By repeating this operation for the k possible values of the scale parameter 2^l , for horizontally and vertically parallel flows, we get the configuration of the geometric flow which minimizes $D_i + \lambda R_i$. This minimum value is finally compared with the value obtained when there is no flow, and the smallest of the two is retained. The distortion rate $D_i + \lambda R_i$ when there is no flow is calculated over each square from the corresponding image coefficients in a separable wavelet basis. Repeating these operations, for all square width $1 \leq 2^k \leq N$, yields an optimal geometric flow that minimizes $D_i + \lambda R_i$ for each dyadic square, with a total of $O(N^2 (\log_2 N)^2)$ operations.

We now find the partition $[1, N]^2 = \cup_i \Omega_i$ which minimizes $\sum_i (D_i + \lambda R_i)$ with a fast bottom up algorithm along the branches of the segmentation quad tree, as in [20]–[22], [32]. For any square Ω_i , a partition into smaller squares $\Omega_i = \cup_l \Omega_l$ gives a better distortion rate if

$$D_i + \lambda R_i \geq \sum_l D_l + \lambda R_l.$$

We begin at a maximum depth of the quad tree corresponding to regions of width $2^K = N2^{-J}$ typically equal to 4. At the

next depth $J - 1$, we compare the distortion rate of each region of size $N2^{-J+1}$ and the sum of the distortion rate of their four subregions, and keep the configuration corresponding to the minimum distortion rate. Again, at the next depth $J + 2$, we compare the distortion rate of each region of size $N2^{-J+2}$ and the sum of the minimum distortion rates for their four subregions, and keep the minimum value together with the optimal configuration. Continuing this aggregation procedure until the top of the tree leads to an optimal partition of the image support $[1, N]^2$ into dyadic regions which minimize the overall distortion rate $D + \lambda R$.

The geometric flow segmentation shown in Fig. 6(a) was obtained when optimizing the compression of the Barbara image for $\Delta = 30$. As expected, the optimization adjusts the dyadic squares so that the parallel geometric flow can follow the geometric directions of the image structures. Fig. 6(b) shows the bandelet compressed image: No blocking artifact can be seen in the reconstruction.

Compression in a bandelet basis is compared with a compression in the 7/9 CDF wavelet basis [25], using the same quantization and adaptive arithmetic coding procedures. We do not incorporate the bit-plane strategy and the contextual coding procedure of JPEG-2000 to compare more easily the performance of the bandelet and wavelet bases themselves. Similar bit plane and contextual coding procedures can also be applied to bandelet coefficients. Fig. 7 compares distortion rate $D(R)$ of the



Fig. 8. Comparison of image compression with (left) bandelet and (right) wavelet bases, with $R = .13$ bits/pixel. The bandelet reconstruction has a PSNR of 30.8 dB and the wavelet reconstruction has a PSNR of 30.3 dB.

bandelet compression algorithm with a wavelet compression for the Lena and Barbara images. For each value of R , the optimization algorithm allocates about 10% of the bits to code the segmented geometric flow and about 90% to code the quantized bandelet coefficients. This proportion may vary from image to image and across bit rates but happens to remain nearly constant over all bit rates for these two images.

The bandelet coder outperforms the wavelet coder by about .5 dB for Lena and 1.5 dB for Barbara. It is important to observe that this remains valid for a bit rate R/N^2 going from .1 bits/pixel to 1 bits/pixel, which covers the whole range of practical applications. From a visual quality point of view, the difference of performance appears clearly in Figs. 8 and 9. Although the bandelet coder introduces errors, the restored images have a regular geometry along the direction of the computed flow, and the resulting error is less visible. On the contrary, wavelets introduce visible ringing effects that are distributed the square grids of the wavelet sampling, which partly destroys the geometric regularity. As a result, the bandelet compressed images have a better visual quality than their wavelet counterparts. For Barbara, the improvements of bandelets over wavelets is larger than for Lena, because of the presence of textures having regular geometric structures. For images having no geometric regularity, the bandelet basis is essentially similar to a wavelet basis and the distortion rate is, therefore, the same.

B. Noise Removal

Thresholding estimators in an orthonormal basis are particularly efficient to remove additive noises if the basis is able to approximate the original signal with few nonzero coefficients [4]. For bandelet bases, this requires to estimate and optimize the geometric flow in presence of additive noise. A penalized estimation finds the “best” bandelet basis which minimizes an empirical risk that is penalized by the complexity of the geometric flow.

We want to estimate a signal $f[n]$ from the noisy data

$$X[n] = f[n] + W[n] \quad (48)$$

where $W[n]$ is a Gaussian white noise of variance σ^2 . A thresholding estimator in a bandelet basis $\mathcal{B} = \{g_m\}_{1 \leq m \leq N^2}$ can be written

$$F = \sum_{m=1}^{N^2} \rho_T(\langle X, g_m \rangle) \check{g}_m \quad (49)$$

where $\rho_T(x)$ is a hard thresholding at T : $\rho_T(x) = x \mathbf{1}_{|x| > T}$. The noise variance σ is estimated with a robust median estimator in a wavelet basis [4], and according to Donoho and Johnstone [34] the threshold is set to $T = \gamma \sigma \sqrt{2 \log_e N^2}$ where γ is a constant that is adjusted.

The expected quadratic risk $\mathbf{E}\{\|F - f\|^2\}$ depends upon f and on the choice of basis. The “best” bandelet basis is the one that minimizes this risk among all possible bandelet bases. In practice, we cannot find this “best” basis because f is unknown, but we can try to estimate a basis which produces a risk that is nearly as small. This requires to optimize the geometric flow of the bandelet basis in presence of noise.

The thresholding estimator (49) can be rewritten as an orthogonal projection $F = P_{\mathcal{M}}(X)$ on a space \mathcal{M} generated by the vectors g_m such that $|\langle X, g_m \rangle| \geq T$. Finding the best thresholding estimator, thus, means finding the best “model” space \mathcal{M} to perform the projection. Model selection procedures have been developed with penalization approaches that introduce a cost that depends upon the “complexity” of the model [35]–[37]. When the noisy data X is obtained by the addition of a Gaussian white noise as in (48), nearly minimax “best” bases are found by minimizing appropriate penalized cost functions [34], [38]. Moulin [39] shows that such thresholding penalized estimators can be obtained by minimizing the Lagrangian of a distortion rate

$$D + \lambda \sigma^2 R \quad \text{with} \quad D = \|X - F\|^2 \quad (50)$$

where R measures the complexity of the model \mathcal{M} as the number of bits needed to code the selected basis \mathcal{B} and the quantized coefficients of X in \mathcal{B} , for a quantization step equal to the threshold T . We are, thus, facing a distortion rate minimization as in the previous section. An analytical calculation of the asymptotic risk obtained for appropriate values of the Lagrange




Fig. 9. Middle column shows different zooms compressed with bandelets using respectively with $R = .22$ bits/pixel for Lena and $R = .40$ bits/pixel for Barbara. Wavelet compression at the same rate are shown on the right column. The left column displays the optimized geometric flow of the bandelet compression.




Fig. 10. Risk in PSNR of (full lines) the bandelet thresholding estimator and of (dashed lines) the wavelet thresholding estimator for the Lena and Barbara images as a function of the PSNR of the original noisy signal. The bandelet estimator reduces the risk by approximately 1 dB for Lena and by 1.8 dB for Barbara.

multiplier λ can be found in [2]. A minimum description length penalization corresponds to $\lambda = 2 \log_e 2$ [39]. In our numerical experiments, we chose a constant that is twice smaller.

The best bandelet basis which minimizes (50) is obtained by optimizing the image segmentation in dyadic squares Ω_i together with their geometric flows. The corresponding bit budget R is calculated in (43). In the context of image compression, given an image segmentation, the flow in each region Ω_i is calculated by minimizing the quadratic image variation along the flow (47). The signal is regularized by a Gaussian filter θ and the displacement parameters $c'_i[p]$ is parameterized in a family of B-splines dilated by 2^l . To estimate the flow in presence of noise, the variance β^2 of the Gaussian filter θ is adjusted to the scale $\beta^2 = 2^{2l}$ in order to filter the noise according to the resolution of the geometric flow. Modulo this modification, the minimization of the distortion rate (50) is achieved by the fast algorithm described in Section V-A for image compression. It re-

quires $O(N^2 (\log_2 N)^2)$ operations to optimize the image segmentation and the geometric flow in each region and compute the corresponding thresholding estimator.

Thresholding estimators are improved by translation invariant procedures which perform a thresholding estimation on each translated version of the image and averages all these estimations after an inverse translation [40]. The following numerical experiments compare the PSNR obtained with a translation invariant thresholding in a 7/9 wavelet basis and a translation invariant thresholding in an optimized bandelet basis, depending upon the value of the noise variance σ . We do not include more sophisticated estimation procedure as in [41], to concentrate on the properties of the bases.

Fig. 10 gives the PSNR of bandelet and wavelet thresholded images for Lena and Barbara, as a function of the PSNR of the original noisy image $X = f + W$. The bandelet estimator outperforms the wavelet estimator by about 1 dB for Lena and 1.8



Fig. 11. Comparison of thresholding estimation with (left) bandelet bases and (right) wavelet bases, from noisy Lena images having a PSNR = 20.2 dB. The PSNR of the bandelet and the wavelet estimations are respectively 30.3 and 29.2 dB.




Fig. 12. Left columns gives zooms of noisy images having a PSNR = 20.19 dB. The middle and left columns are obtained, respectively, with bandelet and wavelet estimators.

dB for Barbara, for nearly all PSNR. As for image compression, the difference of performance between the two estimators appears clearly in Figs. 11 and 12, because the image geometry is better restored.

VI. CONCLUSION

A central idea in the construction of bandelets is to define the geometry as a vector field, as opposed to a set of edge curves. This vector field plays the same role as motion vectors in video image sequences. It indicates the direction of displacement of gray-level values, not in time but in space. Like in video image coding, this geometry is simplified by an image segmentation in squares, whose sizes are adapted to the local image structures.

The geometry of bandelet bases is not calculated a priori but by optimizing the resulting application, whether it is image compression or noise removal, with a fast best basis search algorithm. As a result, bandelet bases clearly improve the image compression and noise removal results obtained with wavelet bases. For video image sequences, a three-dimensional time-space geometric flow should be defined to construct bandelet bases that are adapted to the space-time geometry of the sequence. This is a possible approach to improve the current video compression standard.

REFERENCES

- [1] E. Le Pennec and S. Mallat, "Bandelet image approximation and compression," *SIAM J. Multiscale Modeling Simul.*, to be published.

- [2] C. Dossal, E. Le Pennec, and S. Mallat, "Geometric Image Estimation with Bandelets," Tech. Rep., foCMAP, École Polytech., Palaiseau, France, 2004.
- [3] F. Falzon and S. Mallat, "Analysis of low bit rate image transform coding," *IEEE Trans. Image Process.*, vol. 46, no. 4, pp. 1027–1042, Apr. 1998.
- [4] D. Donoho and I. Johnstone, "Ideal spatial adaptation via wavelet shrinkage," *Biomet.*, vol. 81, pp. 425–455, Dec. 1994.
- [5] S. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *Pami*, vol. 11, no. 7, pp. 674–693, Jul. 1989.
- [6] Y. Meyer, *Wavelets and Operators*. Cambridge, U.K.: Cambridge Univ. Press, 1993.
- [7] A. Cohen, W. Dahmen, and R. DeVore, "Multiscale decompositions on bounded domains," IGP, Tech. Rep. 113, 1995.
- [8] S. Jaffard and Y. Meyer, "Bases d'ondelettes dans des ouverts de \mathbb{R}^n ," *J. Mathématiques Pures et Appliquées*, vol. 68, pp. 95–108, 1989.
- [9] W. Dahmen and R. Schneider, "Wavelets on manifolds I: Construction and domain decomposition," *SIAM J. Math. Anal.*, vol. 31, no. 1, pp. 184–230, 2000.
- [10] R. DeVore, "Nonlinear approximation," *Acta Numer.*, vol. 7, pp. 51–150, 1998.
- [11] A. Cohen, R. DeVore, P. Petrushev, and H. Xi, "Non linear approximation and the space $BV(\mathbb{R}^2)$," *Amer. J. Math.*, no. 121, pp. 587–628, 1999.
- [12] M. Jansen, H. Choi, S. Lavu, and R. Baraniuk, "Multiscale image processing using normal triangulated meshes," presented at the *Int. Conf. Image Processing*, Thessaloniki, Greece, 2001.
- [13] S. Carlsson, "Sketch based coding of gray level images," *IEEE Trans. Image Process.*, vol. 15, no. 1, pp. 57–83, Jan. 1988.
- [14] J. Elder, "Are edges incomplete?," *Int. J. Comput. Vis.*, vol. 34, no. 2/3, pp. 97–122, 1999.
- [15] X. Xue and X. Wu, "Image compression based on multi-scale edge compensation," in *Proc. IEEE Int. Conf. Image Processing*, vol. 3, Oct. 1999, pp. 560–564.
- [16] S. Mallat and S. S. Zhong *et al.*, "Wavelet transform maxima and multiscale edges," in *Wavelets and their Applications*, M. B. Ruskai *et al.*, Eds. Boston, MA: Jones and Bartlett, 1992.
- [17] A. Cohen and B. Matei, "Nonlinear subdivisions schemes: Applications to image processing," in *Tutorial on Multiresolution in Geometric Modeling*, A. Iske, E. Quack, and M. Floater, Eds. New York: Springer, 2002.
- [18] E. Le Pennec and S. Mallat, "Image compression with geometrical wavelets," presented at the *Int. Conf. Image Processing*, Vancouver, BC, Canada, Sep. 2000.
- [19] M. Dragotti and P. L. abd Vetterli, "Wavelet footprints: Theory, algorithms and applications," *IEEE Trans. Signal Process.*, vol. 51, no. 5, pp. 1306–1323, May 2003.
- [20] D. Donoho, "Wedgelets: Nearly-minimax estimation of edges," *Ann. Stat.*, vol. 27, pp. 353–382, 1999.
- [21] R. Shukla, P. L. Dragotti, M. N. Do, and M. Vetterli, "Rate-distortion optimized tree structured compression algorithms," *IEEE Trans. Image Process.*, to be published.
- [22] M. Wakin, J. Romberg, H. Choi, and R. Baraniuk, "Rate-distortion optimized image compression using wedgelets," in *IEEE Int. Conf. Image Process.*, vol. 3, Jun. 2002, pp. III-237–III-240.
- [23] E. Candès and D. Donoho, "Curvelets: A surprisingly effective nonadaptive representation of objects with edges," in *Curves and Surfaces Fitting*, L. L. Schumaker, A. Cohen, and C. Rabut, Eds. Nashville, TN: Vanderbilt Univ. Press, 1999.
- [24] M. N. Do and M. Vetterli, "Contourlets," in *Beyond Wavelets*, J. Stoeckler and G. V. Welland, Eds. New York: Academic, 2002, to be published.
- [25] A. Cohen, I. Daubechies, and J. C. Feauveau, "Biorthogonal bases of compactly supported wavelets," *Comm. Pure Appl. Math.*, vol. 45, pp. 485–560, 1992.
- [26] T. Blu, P. Thevenaz, and M. Unser, "MOMS: Maximal-order interpolation of minimal support," *IEEE Trans. Image Process.*, vol. 10, no. 7, pp. 1069–1080, Jul. 2001.
- [27] D. L. Donoho, "Wedgelets: Nearly-minimax estimation of edges," *Ann. Stat.*, vol. 27, pp. 867–897, 1999.
- [28] C. Bernard and E. Le Pennec, "Adaptation of regular grid filterings to irregular grids," CMAP, École Polytech., Palaiseau, France, Tech. Rep. 500, 2002.
- [29] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. New York: Academic, 1998.
- [30] R. R. Coifman and M. V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 713–718, Mar. 1992.
- [31] I. Daubechies and W. Sweldens, "Factoring wavelet transforms into lifting steps," *J. Fourier Anal. Appl.*, vol. 4, no. 3, pp. 245–267, 1998.
- [32] K. Ramchandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense," *IEEE Trans. Image Process.*, vol. 2, no. 4, pp. 160–175, Apr. 1993.
- [33] B. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. 7th Int. Joint Conf. Artificial Intelligence*, Vancouver, BC, Canada, 1981, pp. 674–679.
- [34] D. Donoho and I. Johnstone, "Ideal denoising in an orthonormal basis chosen from a library of bases," *C. R. Acad. Sci., ser. I*, vol. 319, pp. 1317–1722, 1994, submitted for publication.
- [35] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, pp. 465–471, 1983.
- [36] A. Barron and T. Cover, "Minimum complexity density estimation," *IEEE Trans. Inf. Theory*, vol. 37, pp. 1034–1054, Jul. 1991.
- [37] L. Birgé and P. Massart, "From model selection to adaptive estimation," in *Festschrift for Lucien Le Cam*. New York: Springer, 1996.
- [38] H. Krim and I. Schick, "Minimax description length for signal denoising and optimized representation," *IEEE Trans. Inf. Theory*, vol. 45, no. 3, pp. 898–908, Apr. 1999.
- [39] J. Liu and P. Moulin, "Complexity-regularized image denoising," *IEEE Trans. Image Process.*, vol. 10, no. 6, pp. 841–851, Jun. 2001.
- [40] R. R. Coifman and D. Donoho, "Translation invariant de-noising," Dept. Statistics, Stanford Univ., Stanford, CA, Tech. Rep. 475, 1995.
- [41] J. Portilla, V. Strela, M. Wainwright, and E. P. Simoncelli, "Image denoising using Gaussian scale mixtures in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.



Erwan Le Pennec was born in Courbevoie, France, in 1976. He received the degree from the Ecole Normale Supérieure de Cachan, Cachan, France, the Aggregation of Mathematics degree in 1997, and the Ph.D. degree in applied mathematics from the Ecole Polytechnique, Palaiseau, France, in 2002.

His research interests include wavelets, geometrical representation, nonlinear approximation, data compression, and estimation.



Stéphane Mallat (F'05) was born in Paris, France. He graduated from the Ecole Polytechnique, Paris, France, and the Ecole Nationale Supérieure des Télécommunications, Paris, and received the Ph.D. degree in electrical engineering from the University of Pennsylvania, Philadelphia, in 1984, 1985, and 1988, respectively.

In 1988, he joined the Computer Science Department, Courant Institute of Mathematical Sciences, New York University, New York, where he is now a Research Professor. Since 1995, he has been a

Professor with the Applied Mathematics Department, Ecole Polytechnique and was Chairman of the Department from 1999 to 2002. His research interest includes computer vision, signal processing, and applications of wavelet transforms. He is the author of the book *A Wavelet Tour of Signal Processing* (New York: Academic, 1998).

Dr. Mallat received the 1990 IEEE Signal Processing Society's paper award, the 1993 Alfred Sloan fellowship in Mathematics, the 1997 Outstanding Achievement Award from the SPIE Optical Engineering Society, and the 1997 Blaise Pascal Prize in applied mathematics from the French Academy of Sciences.