Steerlets: A novel approach to rigid-motion covariant multiscale transforms

Robert Azencott, Bernhard G. Bodmann, and Manos Papadakis
Department of Mathematics, University of Houston, 651 Phillip G. Hoffman Hall
Houston, TX 77204-3008, USA

ABSTRACT

We present steerlets, a new class of wavelets which allow us to define wavelet transforms that are covariant with respect to rigid motions in d dimensions. The construction of steerlets is derived from an Isotropic Multiresolution Analysis, a variant of a Multiresolution Analysis whose core subspace is closed under translations by integers and under all rotations. Steerlets admit a wide variety of design characteristics ranging from isotropy, that is the full insensitivity to orientations, to directional and orientational selectivity for local oscillations and singularities. The associated 2D or 3D-steerlet transforms are fast MRA-type of transforms suitable for processing of discrete data. The subband decompositions obtained with 2D or 3D-steerlets behave covariantly under the action of the respective rotation group on an image, so that each rotated steerlet is the linear combination of other steerlets in the same subband.

1 Introduction

The concept of steerlets introduced in this paper presents us with a new class of multidimensional wavelets implementable with fast wavelet transforms. These new wavelets aim at encompassing several properties desirable for data transforms which can be used in texture and shape analysis of 2D and 3D images:

1. The transform should encode typical features in a sufficiently sparse fashion to allow clustering of the feature vectors with sufficiently non-overlapping clusters for most features of interest.

2. The transform must respond predictably under 2D or 3D rigid motions of the entire image or of image patches.

In this paper we want to illustrate the basic ideas and methods of constructing a class of new wavelet transforms that address Property 2. Property 1 is rather more general and it simply says that a transform that does not allow clusters to be formed or if they are formed they overlap in way that leads to misclassification is rather undesirable for the kind of data under consideration. However, Property 2 is not totally unrelated to Property 1. If the transform behaves unpredicatbly when a rigid motion is applied to an image or to an image patch, then the variability of the output data may force clusters to overlap.

The necessity for wavelet transforms to be respond predictably under rigid motions was realized early enough in the wavelet era when the lack of shift-invariance in the lower subband outputs was considered to be a problem in some applications such as noise removal. The fix was easy to discover: shift-invariant wavelet transforms, which are redundant, see e.g., Refs. 10, 20. Although this fix appears to be easy, it required some serious mathematical work to establish that shift-ivantiant fast wavelet transforms are invertible (stability, equivalently boundedness of the inverse transform). Some very nice mathematical studies have been devoted to this subject, e.g., as in Refs. 10, 20. However, rotations still remained a problem. Some partial solutions to this problem, not directly related to wavelets were proposed as early as in mid-1980's as yet there is no general solution of the problem ntiher does exist a mathematical theory of steerable representations. The concept of 'steerable pyramids' was proposed to the image processing community in a series of papers Refs. 14, 21, 27. This approach has since been taken up for several image processing applications, for example image denoising[5, 25] and for texture recognition Refs. 9, 28. Attempts of generalizing and characterizing this approach, particularly in 2D have also been made Refs. 18, 24. One way to address the problem of 'steerability' of multiscale representations is to use representations that are covariant under rigid motions. This is a recipe followed by many authors all, notably in 2D, Refs. 13, 14, 17, 29–33. Most recently the use of the Riesz transform has shown provide steerability rules for 2D and 3D-image representations Refs. 19, 30. Approaching 3D-image analysis from the viewpoint of moments as descriptors of textures and corners, Fickus and coworkers proposed a very different way to obtain steerability rules for moment-based representations of images, Refs. 11, 12. Despite those recent developments, 3D-rigid motion covariant multiscale representations still seem to leave a lot to be developed. This was also true for the mathematical formulation of the concept of steerability described in Ref. 1. The new class of wavelets, called *steerlets*, we present in this paper gives rise to discrete wavelet transforms that are steerable in a particular fashion which takes into account that every image is digitized.

Before stating the main results of this paper we introduce some notation. The *Fourier transform* $\widehat{f} : \mathbb{R}^d \longrightarrow \mathbb{C}$ of $f \in L^1(\mathbb{R}^d)$ is defined by

$$\forall \xi \in \mathbb{R}^d, \; \widehat{f}(\xi) = \int_{\mathbb{R}^d} f(x) e^{-2\pi i x \cdot \xi} dx.$$

Here, $x \cdot \xi$ denotes the standard inner product between two vectors $x, \xi \in \mathbb{R}^d$. The map $f \to \widehat{f}$ restricted to $L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$ extends to a unitary map on $L^2(\mathbb{R}^d)$ denoted by $\mathcal{F}$.

A *frame* for a separable Hilbert space H is a sequence $\{f_i\}_{i \in \mathbb{I}} \subseteq H$, where $\mathbb{I}$ is a countable index set, for which there are constants $A, B > 0$ such that

$$\forall f \in H, \; A\|f\|^2 \leq \sum_{i \in \mathbb{I}} |\langle f, f_i \rangle|^2 \leq B\|f\|^2.$$

A and B are called frame constants. The weaker version of the above inequality where $B > 0$ such that the right-hand side of the above inequality is valid for all $f \in \mathbb{R}^d$ is termed by referring to $\{f_i\}_{i \in \mathbb{I}}$ as a *Bessel* sequence. If $A = B$, we say the frame is *tight*, and a tight frame with $A = B = 1$ is called *Parseval*. A frame is a Riesz basis for H if it is also a minimal family generating H. We use the terms frame (Riesz) sequence when we refer to a countable frame or a Riesz basis of a subspace of a Hilbert space.[26] For $y \in \mathbb{R}^d$, the (unitary) shift operator T_y is defined as $T_y f(x) = f(x - y)$, $f \in L^2(\mathbb{R}^d)$. Set-theoretic complements are indicated by the use of the superscript 'c'. A ball in $\mathbb{R}^d$ centered at x with radius r is denoted by $B(x, r)$. If $r_1 < r_2$ we set $B(x, r_1, r_2) := B(x, r_2) \cap B(x, r_1)^c$. Finally, $e_x(\xi) := e^{-2i\pi(x \cdot \xi)}$, where $x, \xi \in \mathbb{R}^d$.

An image is a function, say f in $L^2(\mathbb{R}^d)$ (here $d = 2, 3$), which belongs to a shift-invariant subspace V. In practical terms this means that we can represent this image by its samples using a more elaborate and general version of the Shannon-Whittaker sampling theorem. With no loss of generality we can assume that the scale has been chosen so that $T_n(V) = V$, where $n \in \mathbb{Z}^d$. We use the τ_a to denote the affine mapping $\tau_a(x) = x - a$, $a, x \in \mathbb{R}^d$. Finally, to avoid unnecessary complications we use only dyadic dilations, but most of the results in this paper hold for more general dilations. We denote by D the dyadic dilation operator, $Df(x) = 2^{d/2} f(2x)$, for $f \in L^2(\mathbb{R}^d)$.

If an image f is rotated, then it is natural to assume that the rotated image belongs to V as well. This results from the requirement that f and every rotation of f should not require different sampling rates when digitizing. This simple requirement naturally leads to a new class of multiresolution analysis, called Isotropic Multiresolution Analysis (IMRA),[26] where the resolution subspaces V_j are not only invariant with respect to the action of the translation group $\{T_{2^{-j}n} : n \in \mathbb{Z}^d\}$, but they are also rotationally invariant meaning $\rho(R)(V_j) = V_j$ for all integers j, where $\rho(R)f(x) = f(R^T x)$ with $f \in L^2(\mathbb{R}^d)$, $x \in \mathbb{R}^d$ and $R \in \mathrm{SO}(d)$; the superscript T of R denotes the transpose of the rotation matrix R. Apparently, all V_j remain invariant under all rotations centered at any point of their respective sampling grid $2^{-j}\mathbb{Z}^d$. All such subspaces have a special structure:

THEOREM 1.1.[26] *Let V be an invariant subspace of $L^2(\mathbb{R}^d)$ under the action of the translation group induced by $\mathbb{Z}^d$, i.e. $T_n(V) = V$, where $n \in \mathbb{Z}^d$. Then V remains invariant under all rotations if and only if there exisits a measurable subset Ω of $\mathbb{R}^d$ satisfying $R(\Omega) = \Omega$ for every rotation R in $SO(d)$ and*

$$V = \{h \in L^2(\mathbb{R}^d) : \hat{h}(\xi) = 0 \textit{ for } \xi \in \Omega^c\}.$$

Such a set Ω is 'radial', or we can more precisely say that it has radial symmetry, equivalently its indicator function χ_Ω is a radial function. An immediate consequence of the previous theorem is that V is translation invariant, that is $T_a(V) = V$ for all $a \in \mathbb{R}^d$.

The most well-known translation invariant subspace is $\{h \in L^2(\mathbb{R}^d) : \hat{h}(\xi) = 0 \text{ for } \xi \notin \left[-\frac{1}{2}, \frac{1}{2}\right)^d\}$ associated with the classical sampling theorem. However, the same sampling theorem would hold for the translation-invariant subspace of $L^2(\mathbb{R}^d)$ given by $\{h \in L^2(\mathbb{R}^d) : \hat{h}(\xi) = 0 \text{ for } \|\xi\| > 1/2\}$. The latter subspace is invariant under translations and all rotations, which we will use in this paper.

2 Steerlets and steerable feature mappings

We begin by introducing a concept of steerability for digital images.

Definition 2.1. *Assume that V is a closed subspace of $L^2(\mathbb{R}^d)$ invariant under the action of all integer translations and all d-dimensional rotations, that is, $T_n f \in V$ for all $f \in V$ and $n \in \mathbb{Z}^d$ as well as $\rho(R)f \in V$ for all $f \in V$ and $R \in SO(d)$. A bounded linear transformation $A : V \to \ell^2(\mathbb{Z}^d, \mathbb{C}^{k_0})$ is a Discrete Steerable Feature Mapping (DSFM) if there exists a shift-invariant representation B of the rotation group on $\ell^2(\mathbb{Z}^d, \mathbb{C}^{k_0})$ and if for all $R \in$SO(d), n and m in $\mathbb{Z}^d$,*

$$(AT_m\rho(R)f)(n) = B(R)(Af)(n-m).$$

In brief, A intertwines ρ and B as well as the integer translations on V and on $\ell^2(\mathbb{Z}^d, \mathbb{C}^{k_0})$. The reader may have noticed a departure from our premise; a DSFM is not explicitly required to be covariant under the action of every translation T_a, $a \in \mathbb{R}^d$. The above definition has been chosen so that it can be verified directly when V is the core subspace of an Isotropic Multiresolution Analysis, defined below, and A associates wavelet coefficients at $n \in \mathbb{Z}^d$ with any function $f \in W_0 := D(V) \cap V^\perp$.

Our starting point for constructing DSFMs is the IMRA.

Definition 2.2. *An Isotropic Multiresolution Analysis (IMRA) of $L^2(\mathbb{R}^d)$ is a sequence $\{V_j\}_{j\in\mathbb{Z}}$ of closed subspaces of $L^2(\mathbb{R}^d)$ satisfying the following conditions:*

1. *$\forall j \in \mathbb{Z}$, $V_j \subset V_{j+1}$ and $D^j(V_0) = V_j$;*
2. *$\cup_{j\in\mathbb{Z}} V_j$ is dense in $L^2(\mathbb{R}^d)$ and $\cap_{j\in\mathbb{Z}} V_j = \{0\}$;*
3. *V_0 is invariant under the action of the groups $\{T_n : n \in \mathbb{Z}^d\}$ and $\{\rho(R) : R \in SO(d)\}$.*

This concept relates to the theory of MRAs because IMRAs are GMRAs of $L^2(\mathbb{R}^d)$ which in addition require the invariance of V_0 under the action of $\{\rho(R) : R \in SO(d)\}$. The term GMRA was coined by Baggett and others[3] in order to signify a more general class of multiresolution analyses than the classical MRAs which require the existence of an orthonormal scaling function generating the core subspace V_0, see also the related work.[6]

Definition 2.3. *A finite family of functions $\{f_i : i = 1, 2 \ldots n\}$ in $L^2(\mathbb{R}^d)$ is called steerable if for every $R \in \mathrm{SO}(d)$ and $i \in \{1, 2 \ldots n\}$, $\rho(R)f_i$ belongs to the linear span of the family.*

The family $\{\psi_i^{(j)} : i = 0, 1, \ldots, n(j), j \in \mathbb{Z}\} \subset L^2(\mathbb{R}^d)$ is called a family of steerlets if $\{D^j T_{m/2} \psi_i^{(j)} : i = 0, 1, \ldots n(j), j \in \mathbb{Z}, m \in \mathbb{Z}^d\}$ is a frame for $L^2(\mathbb{R}^d)$ and if for every $j \in \mathbb{Z}$, $\{\psi_i^{(j)} : i = 0, 1, \ldots, n(j)\}$ is steerable. We will also say that the steerable set $\{\psi_i^{(j)} : i = 0, 1, \ldots, n(j)\}$ is centered at the origin, whereas the corresponding set centered at the point $a \in \mathbb{R}^d$ is $\{T_{-a}\psi_i^{(j)} : i = 0, 1, \ldots, n(j), j \in \mathbb{Z}\} \subset L^2(\mathbb{R}^d)$.

A texture segmentation/identification algorithm can usually be split in two parts, feature extraction and the classification process. Feature extraction utilizes a data representation. Our features must be such that when two patches—small regions of the image—generate 'similar' feature vectors in some statistical sense, then the patches can be considered as coming from the same texture. In this context we are not interested in the global rotations of an image but in rotations of patches. So, we want our steerlets to be well-localized in the space domain. Then, heuristically speaking, if Patch 1 is centered at the origin and is the rotation of Patch 2 centered at the grid point **n** then the steerlet representation of the Patch 2 will be easily predictable from the steerlet representation of Patch 2 if we use steerlets centered at **n**. This can allow us to generate rotation-invariant texture signatures and thus rotation-invariant classifiers. Now the context in which we intent to use steerlets is clear. However, in this paper we will only illustrate the main ideas for their construction.

We start by using an IMRA generated by a very simple function, namely $\varphi = \mathcal{F}^{-1}\chi_{B(0,b_0)}$, where $b_0 < 1/2$. We begin with a simple case because we can prove all claims with standard arguments without having to employ technical results. However, one would only use this IMRA for modelling purposes because the filters associated with it have the poor spatial localization. First, note

$$\mathbb{R}^d = B(0, 2^{j_0} b_0) \cup_{j=j_0}^{\infty} B(0, 2^j b_0, 2^{j+1} b_0) \ , \qquad j_0 \in \mathbb{Z}.$$

This dyadic decomposition is a partition of $\mathbb{R}^d$, so $L^2(\mathbb{R}^d)$ is equal to the orthogonal direct sum of its subspaces X_{j_0} and Y_j, where X_{j_0} contains those functions in $L^2(\mathbb{R}^d)$ vanishing a.e. outside $B_(0, 2^{j_0} b_0)$ and Y_j contain those who vanish a.e. outside $B(0, 2^j b_0, 2^{j+1} b_0)$ respectively. This dyadic decomposition is associated with the IMRA generated by φ because $V_{j_0} = \mathcal{F}^{-1}(X_{j_0})$. Accordingly,

$$W_j := V_{j+1} \cap V_j^{\perp} = \mathcal{F}^{-1}(Y_j) \ .$$

Lemma 2.4 (Benedetto and Li[4]; Han and Larson[16]). *Let $f \in L^2(\mathbb{R}^d)$ and $j \in \mathbb{Z}$. Then, $\{T_n f : n \in \mathbb{Z}^d\}$ is a Parseval frame of its closed linear span if and only if*

$$\sum_{s \in \mathbb{Z}^d} |\hat{f}(\xi + s)|^2 = \chi_C(\xi) \qquad \xi \in \mathbb{T}^d \ ,$$

where C is a measurable subset of $\mathbb{T}^d$.

2.1 A first construction of steerlets

Definition 2.5. *A family of measurable functions $\{p_i^{(j)} : i = 1, \ldots, n(j), j \in \mathbb{Z}\}$ on $\mathbb{R}^d$ is called a family of steerlet modulators with pass band $B(0, b_0, 2b_0)$ if the set satisfies the following two conditions:*

- *Normalized square sum,*

$$\sum_{i=1}^{n(j)} |p_i^{(j)}(\xi)|^2 = 1 \qquad a.e. \ in \quad B(0, b_0, 2b_0); \tag{1}$$

- *For every $i = 1, 2, \ldots, n(j)$ and $R \in SO(3)$ there exist scalars $b_{i,k}^{(j)}(R)$, $1 \le k \le n(j)$ such that*

$$\rho(R) p_i^{(j)}(\xi) = \sum_{k=1}^{n(j)} b_{i,k}^{(j)}(R) p_k^{(j)}(\xi) \qquad a.e. \ in \quad B(0, b_0, 2b_0). \tag{2}$$

If each function in the set is continuous, we call them continuous steerlet modulators.

Next we state and prove the first and simplest steerlet construction.

Proposition 2.6. *Let $\psi_i^{(j)}$ be defined via the Fourier transform*

$$\widehat{\psi_i^{(j)}}(\xi) := 2^{-d/2} p_i^{(j)}(\xi) \chi_{B(0,b_0,2b_0)}(\xi), \qquad \xi \in \mathbb{R}^d, \quad i = 1, 2, \ldots, n(j), \quad j \ge j_0 \ ,$$

and with $\{p_i^{(j)} : i = 1, \ldots, n(j), j \in \mathbb{Z}\}$, the family of steerlet modulators in the preceding definition. The family $\{\psi_i^{(j)} : i = 1, \ldots, n(j), j \in \mathbb{Z}\}$ then enjoys the following properties:

(i) For each $j \ge j_0$, the sets $\{\psi_i^{(j)} : i = 1, 2, \ldots, n(j)\}$ and $\{D^{j_0}\varphi\} \cup \{\psi_i^{(j_0)} : i = 1, 2, \ldots, n(j_0)\}$ are steerable.

(ii) The set $\{D^{j_0} T_n \varphi : n \in \mathbb{Z}^d\} \cup \ \cup_{j=j_0}^{\infty} \{D^j T_{n/2} \psi_i^{(j)} : i = 1, 2, \ldots, n(j), \quad n \in \mathbb{Z}^d\}$ is a Parseval frame of $L^2(\mathbb{R}^d)$.

Proof. (i) Let $j \geq j_0$, $1 \leq i \leq n(j)$ and $R \in \mathrm{SO}(3)$, then,

$$\mathcal{F}\left(\rho(R)\psi_i^{(j)}\right)(\xi) = 2^{-d/2} p_i^{(j)}(R\xi)\chi_{B(0,b_0,2b_0)}(R\xi)$$

which, due to (2) is equal to

$$2^{-d/2}\sum_{k=1}^{n(j)} b_{i,k}^{(j)}(R)p_{i,k}^{(j)}(\xi)\chi_{B(0,b_0,2b_0)}(\xi)\ .$$

This proves the first assertion of (i). The second assertion of (i) follows from the (i) and the fact that φ is radial.

(ii) Since $\widehat{\varphi} = \chi_{B(0,b_0)}$ Lemma 2.4 implies that $\{T_n\varphi : n \in \mathbb{Z}^d\}$ is a Parseval frame of V_0, thus $\{D^{j_0}T_n\varphi : n \in \mathbb{Z}^d\}$ is a Parseval frame of V_{j_0}. Now let $j \geq j_0$. Then,

$$D^j T_{n/2}\psi_i^{(j)} = T_{2^{-(j+1)}n}D^j\psi_i^{(j)}$$

so

$$D^{-(j+1)}T_{2^{-(j+1)}n}D^j\psi_i^{(j)} = T_n D^{-1}\psi_i^{(j)}\ .$$

Thus, $\{D^jT_{n/2}\psi_i^{(j)} : i = 1, 2, \dots, n(j),\ \ n \in \mathbb{Z}^d\}$ is a Parseval frame of W_j if and only if $\{T_nD^{-1}\psi_i^{(j)} : i = 1, 2, \dots, n(j),\ \ n \in \mathbb{Z}^d\}$ is a Parseval frame of $W_{-1} = \mathcal{F}^{-1}(B(0, b_0/2, b_0))$. To establish that this is true we apply Lemma 2.4:

$$\begin{aligned}\sum_{s\in\mathbb{Z}^d} |(\mathcal{F}D^{-1}\psi_i^{(j)})(\xi+s)|^2 &= \sum_{s\in\mathbb{Z}^d} |p_i^{(j)}(2(\xi+s))\chi_{B(0,b_0/2,b_0)}(\xi+2s)|^2 \\ &= \chi_{B(0,b_0/2,b_0)}(\xi)\quad \text{a.e. in } \mathbb{T}^d.\end{aligned}$$

We remark that the last of the previous equalities follows from (1). To complete the proof of (ii) note that $L^2(\mathbb{R}^d)$ is equal to the orthogonal direct sum of V_{j_0} and all of the W_j's with $j \geq j_0$. □

The previous theorem can be generalized to hold for φ and steerlets $\psi_i^{(j)}$ with rapid decay in the spatial domain. These arise when instead of φ and $\psi_i^{(j)}$ proposed above we use smooth versions of the radial windows $\chi_{B(0,b_0)}$ and $B(0, 2^jb_0, 2^{j+1}b_0)$. Here is how this can be done.

2.2 A construction of spatially well-localized steerlets

Definition 2.7. *Let $0 < a_1 < a_0 < 1/4$ and suppose that the low pass filter H_0 is a continuous real-valued $\mathbb{Z}^d$-periodic function such that*

- *$H_0 = 1$ inside the ball of radius a_1;*
- *$H_0 = 0$ on $\mathbb{T}^d \setminus B(0, a_0)$;*
- *$H_0|_{B(0,a_0)}$ is radial.*

Let h be a $\mathbb{Z}^d$-periodic function defined by the identity

$$h(\xi) = \sqrt{1 - H_0^2(\xi)}\,. \tag{3}$$

Let $\{p_i^{(j)}\}$ be a family of continuous, $\mathbb{Z}^d$-periodic steerlet modulators with pass band $B(0, a_1, 2a_0)$. We then define the steerlet filters associated with $\{p_i^{(j)}\}$ and h for each $j \in \mathbb{Z}$ by $H_i^{(j)} = p_i^{(j)}(\xi)h(\xi)$ a.e. in $\mathbb{T}^d$ and by H_0.

Theorem 2.8. *Let $1/8 < a_1 < a_0 < 1/4$. Let $\{p_i^{(j)}\}$ be a family of steerlet modulators with pass band $B(0, a_1, 2a_0)$ and let h be as in the preceding definition, and define the steerlet filters H_0 and $H_i^{(j)}$ associated with these functions accordingly. Let ϕ be defined by $\hat{\phi}(\xi) = H_0(\xi/2)\chi_{\mathbb{T}^d}(\xi)$ a.e. in $\mathbb{R}^d$ and the steerlets $\psi_i^{(j)}$ by*

$$\widehat{\psi_0^{(j)}}(\xi) = 2^{-d/2}h\left(\frac{\xi}{2}\right)\hat{\phi}\left(\frac{\xi}{2}\right) \qquad \widehat{\psi_i^{(j)}}(\xi) = 2^{-d/2}H_i^{(j)}\left(\frac{\xi}{2}\right)\hat{\phi}\left(\frac{\xi}{2}\right) \qquad 1 \leq i \leq n(j),\quad \xi \in \mathbb{R}^d. \tag{4}$$

Then the refinable function ϕ and the steerlets $\psi_i^{(j)}$, where $i = 0, 1, 2, \ldots, n(j)$, defined as above can be chosen to have arbitrarily good spatial localization and the following generalization of Proposition 2.6 is true:

(i) For each scale $j \geq j_0$, the sets $\{\psi_i^{(j)} : i = 0, 1, 2, \ldots, n(j)\}$ and $\{D^{j_0}\phi\} \cup \{D^{j_0}\psi_i^{(j_0)} : i = 0, 1, 2, \ldots, n(j_0)\}$ are steerable.

(ii) The set $\{D^{j_0}T_n\phi : n \in \mathbb{Z}^d\} \cup \cup_{j=j_0}^{\infty}\{D^jT_{n/2}\psi_i^{(j)} : i = 0, 1, 2, \ldots, n(j), \quad n \in \mathbb{Z}^d\}$ is a Parseval frame of $L^2(\mathbb{R}^d)$.

(iii) The set $\cup_{j\in\mathbb{Z}}\{D^jT_{n/2}\psi_i^{(j)} : i = 0, 1, 2, \ldots, n(j), \quad n \in \mathbb{Z}^d\}$ is a Parseval frame of $L^2(\mathbb{R}^d)$.

Apart from item (iii) which directly follows from (ii) because $\lim_{j_0\to-\infty}\mathcal{F}(D^{j_0}\phi)(\xi) = 0$, the proof of this result is non-trivial and, thus is beyond the scope of this publication, so it will appear in a forthcoming paper of our group.

We remark that unlike curvelets[7,8] and shearlets[15,22,23,34] the number of steerlets can vary for each resolution level j in an arbitrary way. This characteristic allows to reduce or increase the number of steerlets used for every resolution level as needed. Steerlets can capture several local characteristics of an image: Local angular oscilation, edges and other singularities alligned in a certain direction.

2.3 Fast Steerlet Transform

Theorem 2.8 allows to use multiple scales. When this is done, we can choose to reduce the redundancy of the transform by compensating the effects of downsampling with a dyadic factor $2^{d/2}$. This allows the transform to preserve energy, as in mathematical terms, the decomposition is an isometry due to (ii) of Theorem 2.8. Due to the construction of steerlets the rotation of steerlet stays within the same scale, while translations of steerlets in general do not share the same property. For this reason if we want the transform to be associated with simple rigid motion covariance properties–which we will discuss in the following section– we must drop the downsampling steps and allow for extra redundancy. In this case we will have to remove the energy adjusting factor $2^{d/2}$. The Fast Steerlet Transform allows for exact reconstruction provided that the input image has a certain bandwidth. For images of higher bandwidth, we have to upsample and low pass the input image before we apply the transform.

With H_0, a_1 and $H_i^{(-1)}$ defined as above. Define, $\mathcal{D}$ and $\mathcal{U}$ the downsampling and upsampling by the factor 2 operators, $X_0\sigma = h_0 * \sigma$ and $X_i^{(-1)}\sigma = h_i^{(-1)} * \sigma$, where $\sigma \in \ell(\mathbb{Z}^d)$ and h_0, $h_i^{(-1)}$ are the $d \times d$ arrays of filter taps of the filters H_0 and $H_i^{(-1)}$ $(i = 0, 1, 2, \ldots, n(-1))$ respectively. One can choose for a single decomposition/reconstruction step any j, not $j = -1$. In other words, any of the steerlet modulators $p_i^{(j)}$ can be chosen for $p_i^{(-1)}$.

Proposition 2.9. Fast Steerlet Transform. *The following are true:*

(i) (Undecimated) *For every digital image $\sigma \in \ell(\mathbb{Z}^d)$ whose Fourier transform $\hat{\sigma}$ vanishes outside $B(0, 2a_1)$ we have*

$$X_0X_0\sigma + \sum_{i=1}^{n(-1)} X_i^{(-1)}X_i^{(-1)}\sigma = \sigma\ .$$

The decomposition in this algorithm is implemented by applying X_0 and $X_i^{(-1)}$ to the input σ.

(ii) (Decimated) *With σ as in (i) we have*

$$(2^{d/2}X_0\mathcal{U})(\mathcal{D}2^{d/2}X_0)\sigma + \sum_{i=1}^{n(-1)} X_i^{(-1)}X_i^{(-1)}\sigma = \sigma\ .$$

The decomposition transform in this algorithm is implemented by applying $2^{d/2}X_0$ and $X_i^{(-1)}$ to the input σ and is an isometry.

Proof. Item (i) follows directly from

$$1 = |H_0(\xi)|^2 + \sum_{i=1}^{n(-1)} \left|H_i^{(-1)}(\xi)\right|^2 \quad \text{for all } |\xi| \leq 2a_1, \tag{5}$$

which in turn follows from the definition of H_0 and (3) and the property of the steerlet modulators.

Recall, that $\{0,1\}^d$ is the set of all arrays with d entries which are all either equal to zero or to one. To prove (ii), first notice that the Fourier transform of $\mathcal{U}\mathcal{D}X_0\sigma$ is equal to $\xi \mapsto \frac{1}{2^d}\sum_{\eta\in\frac{1}{2}\{0,1\}^d} H_0(\xi+\eta)\hat{\sigma}(\xi+\eta)$. Since, H_0 vanishes outside the ball $B(0,1/4)$ we obtain that the Fourier transform of $(2^{d/2}X_0\mathcal{U})(\mathcal{D}2^{d/2}X_0)\sigma$ is equal to $\xi \mapsto |H_0(\xi)|^2\,\hat{\sigma}$ which combined with (5) completes the proof of (ii). The preservation of energy by decimated decomposition transform follows by calculating the energy of the function $\mathcal{U}\mathcal{D}X_0\sigma$ and by observing that $\mathcal{U}$ is an isometry and that each term in the sum $\xi \mapsto \frac{1}{2^{d/2}}\sum_{\eta\in\frac{1}{2}\{0,1\}^d} H_0(\xi+\eta)\hat{\sigma}(\xi+\eta)$ is a translate of the first term, $H_0(\xi)\hat{\sigma}(\xi)$. □

2.4 Discrete Steerability

Throughout this subsection the steerlet modulators $\{p_i^{(j)}\}$ are assumed to have pass band $B(0,a_1,2a_0)$ and the image f is assumed to have Fourier transform vanishing outside $B(0,2a_1)$. This is a convention we adopt since we always assume that an already digitized input image has been sampled at an appropriate sampling rate which we can always set equal to the Nyquist rate. Under this assumption the expansion of f in terms of the Parseval frame of (ii) of Theorem 2.8 and for $j_0 = 0$ is $f = \sum_{k\in\mathbb{Z}^d}\langle f, T_k\phi\rangle T_k\phi$. But $\langle f, T_k\phi\rangle = \int_{\mathbb{R}^d}\hat{f}(\xi)\overline{\hat{\phi}(\xi)}e^{2\pi i(k\cdot\xi)}d\xi = \int_{\mathbb{R}^d}\hat{f}(\xi)e^{2\pi i(k\cdot\xi)}d\xi = f(k)$. This shows that under the proper bandwidth assumptions for the image f we can input its original samples to the Fast Steerlet transform without any initialization.

Next, we consider the high pass output of the transform, Obviously $A_{-1}(T_m f)(n) = A_{-1}f(n-m)$ for every $m \in \mathbb{Z}^d$. Now, let $R \in$SO(3). We say that R_n be a rotation by R centered at n if $R_n(x) = n + R(x-n)$, for all $x \in \mathbb{R}^d$.

Theorem 2.10. *The mapping A_{-1} defined by the undecimated Fast Steerlet Transform $A_{-1}f(n) = (\langle f, T_n\psi_i^{-1}\rangle)_{i=0}^{n(-1)}$, where $n \in \mathbb{Z}^d$, is a Discrete Steerable Feature Mapping, meaning it obeys the desired steerability rules under translations by integers and arbitrary rotations centered at any point where the transform is evaluated.*

Proof. First, let $n = 0$. Then,

$$A_{-1}(\rho(R)f)(0) = (\langle \rho(R)f, \psi_i^{-1}\rangle)_{i=0}^{n(-1)} = (\langle f, \rho(R^T)\psi_i^{-1}\rangle)_{i=0}^{n(-1)} = B^{(-1)}(R^T)(A_{-1}f(0))^T\ ,$$

where for an arbitrary $O \in$SO(3) the $(n(-1)+1)\times(n(-1)+1)$-matrix $B^{(-1)}(O)$ is given by

$$B^{(-1)}(O)_{i,k} = \begin{cases} 1 & \text{if} \quad i = k = 0 \\ 0 & \text{if} \quad i = 0 \;\; k > 0 \;\text{ or }\; k = 0 \;\; i > 0 \\ b_{i,k}^{(-1)}(O) & \text{in all other cases} \end{cases}$$

where $b_{i,k}^{(-1)}(O)$ are as in the continuous steerlet modulators. More generally, when $n \neq 0$ the action of a rotation R_n on f is given by $T_n\rho(R)T_{-n}$. This shows that the steerability rules for rotations and for $n = 0$ apply in this case as well. It is also apparent that translation invariance holds,

$$\begin{aligned} A_{-1}(T_n\rho(R)T_{-n}f)(n) &= (\langle T_n\rho(R)T_{-n}f, T_n\psi_i^{-1}\rangle)_{i=0}^{n(-1)} = (\langle T_{-n}f, \rho(R^T)\psi_i^{-1}\rangle)_{i=0}^{n(-1)} \\ &= B^{(-1)}(R^T)(A_{-1}T_{-n}f(0))^T = B^{(-1)}(R^T)(A_{-1}f(n))^T\ . \end{aligned}$$

If we consider the low pass output of the undecimated transform $L_{-1}f(n) := \langle f, T_n D^*\phi\rangle$, then we see $L_{-1}(T_m f)(n) := L_{-1}f(n-m)$ and for all the actions of all rotations centered at n the output $L_{-1}f(n)$ remains invariant due to the radiality of ϕ. □

3 First Examples

In all examples that follow we include the steerlet modulators $p_i^{(j)}$ only. We also define these modulators on the entire d-dimensional torus $\mathbb{T}^d$. However, these functions are multiplied by the radial window function $hH_0(2\cdot)$. It is the product of the steerlet modulators with the $hH_0(2\cdot)$ that has to be smooth to allow the rapid decay of the resulting steerlets in the spatial domain.

EXAMPLE 3.1. A first example of steerlets in 3 dimensions which are not isotropic is obtained by taking $n(j) = 3$, independent of j, and by defining for $\xi \in \mathbb{T}^3$,

$$\begin{aligned} p_0^{(j)}(\xi) &= \frac{1}{\sqrt{2}} \sin(\beta(\xi)) e^{-i\alpha(\xi)} , \\ p_1^{(j)}(\xi) &= \cos(\beta(\xi)) , \\ p_2^{(j)}(\xi) &= -\frac{1}{\sqrt{2}} \sin(\beta(\xi)) e^{i\alpha(\xi)} , \end{aligned}$$

with spherical coordinates, $\cos(\beta(\xi)) = \xi_3/\|\xi\|$ and $e^{-i\alpha(\xi)} = (\xi_1 + i\xi_2)/\|\xi\|$. This means β measures the angle between the point on the unit sphere $\frac{1}{\|\xi\|}(\xi_1, \xi_2, \xi_3)$ and the 'North pole', namely $(0, 0, 1)$, whereas α is the 'longitude' of the same point measured on the equator of the unit sphere starting from $(1, 0, 0)$. So $-\pi \leq \alpha \leq \pi$ and $0 \leq \beta \leq \pi$. These steerlet modulators are, up to an overall constant factor, spherical harmonics of degree one. It is well known that the span of these spherical harmonics is invariant under rotations. A variation of this set of steerlet modulators is given by

$$\begin{aligned} p_0^{(j)}(\xi) &= 1/\sqrt{2} \\ p_1^{(j)}(\xi) &= 2^{-1/2}\xi_1/\|\xi\| \\ p_2^{(j)}(\xi) &= 2^{-1/2}\xi_2/\|\xi\| \\ p_3^{(j)}(\xi) &= 2^{-1/2}\xi_3/\|\xi\| . \end{aligned}$$

In the latter setting we have one non-directionally selective modulator, $p_0^{(j)}$ which allows to extract non-directional components after applying the high-pass filter $hH_0(2\cdot)$. All other modulators are edge detectors sensitive to each of the main Cartesian axes.

A more general way to produce examples of steerlet modulators in 3D is given by higher order spherical harmonics. These can be used to extract more detailed informations about the singularity structure in an image, and also to improve the angular selectivity in the analysis of features. The details of this will be provided in a forthcoming paper.

4 Acknowledgments

This work was partially supported by a grant from the Institute of Biomedical Imaging Sciences, 'Coronary Arterial Wall Detection Using Clinical CT-Angiography' and NSF-DMS 0807399.

REFERENCES

1. S.K. Alexander, R. Azencott, and M. Papadakis. Isotropic multiresolution analysis for 3D-textures and applications in cardiovascular imaging. In D. Van Der Ville, V. Goyal, and M. Papadakis, editors, *Wavelets XII*, volume 6701, pages 67011S–1 to 67011S–12. SPIE, 2007.
2. J.-P. Antoine, L. Demanet, L. Jacques, and P. Vandergheynst. Wavelets on the sphere: implementation and approximation. *Appl. Comput. Harmon. Anal.*, 13(3):177–200, November 2002.
3. L. Baggett, A. Carey, W. Moran, and P. Ohring. General existence theorems for orthonormal wavelets, an abstract approach. *Publ. Inst. Math. Sci. Kyoto Univ.*, 31:95–111, 1995.
4. J.J. Benedetto and S. Li. The theory of multiresolution analysis frames and applications to filter banks. *Appl. Comput. Harmon. Anal.*, 5:389–427, 1998.

5. A.A. Bharath and J. Ng. A steerable complex wavelet construction and its application to image denoising. *IEEE Trans. Image Processing*, 14(7):948–959, 2005.
6. C. De Boor, R.A. DeVore, and A. Ron. Approximation from shift-invariant subspaces of $l_2(\mathbb{R}^d$. *Trans. Amer. Math. Soc.*, 341(2):787–806, 1994.
7. A. Candes and D. Donoho. Curvelets a surprisingly effective nonadaptive representation for objects with edges. In C. Rabut, A. Cohen, and L. Schumaker, editors, *Curves and Surfaces*, pages 105–120. Vanderbilt University Press, Nashville, TN, 2000.
8. E.J. Candes and L. Demanet. The curvelet representation of wave propagators is optimally sparse. *Comm. Pure and Appl. Math.*, 58(11):1472–1528, 2005.
9. G.Y Chen and P. Bhattacharya. Invariant texture classification using ridgelet packets. In *ICPR*, volume 2, pages 464–467. IEEE, 2006.
10. R. Coifman and D. Donoho. Translation invariant de-noising. In *Wavelets and Statistics*, Lecture Notes in Statistics, pages 125–150. Springer-Verlag, 1995. http://citeseer.ist.psu.edu/coifman95translationinvariant.html.
11. M. Fickus and D.G. Mixon. Isotropic moments over integer lattices. *Appl. Comp. Harmon. Anal.*, 26(1):77–96, 2009.
12. M. Fickus, G.S. Seetharaman, and M.E. Oxley. Multiscale moment transforms over the integer lattice. In D. Van Der Ville, V. Goyal, and M. Papadakis, editors, *Wavelets XII*, volume 6701, pages 67011N–1 to 67011N–8. SPIE, 2007.
13. B. Forster, T. Blu, and M. Unser. A new family of complex rotation-covariant multiresolution bases in 2d. In A. Aldroubi, A. Laine, and M. Unser, editors, *Wavelets X*, volume 5207 of *Proceedings SPIE*, pages 475–479. SPIE, 2003.
14. W.T. Freeman and E.H. Adelson. The design and use of steerable filters. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(9):891–906, 1991.
15. K. Guo, D. Labate, W. Lim, G. Weiss, and E. Wilson. Wavelets with composite dilations and their MRA properties. *Appl. Comput. Harmon. Anal.*, 20:231–249, 2006.
16. D. Han and D.R. Larson. *Frames, Bases and Group Representations*, volume 147 of *Memoirs*. American Mathematical Society, 2000.
17. P. Heinlein, J. Drexl, and W. Schneider. Integrated wavelets for enhancement of microcalcifications in digital mamography. *IEEE Trans. Med. Imag.*, 22(3):402–413, 2003.
18. Yacov Hel-Or and Patrick C. Teo. Canonical decomposition of steerable functions. *J. Math. Imaging Vision*, 9:83–95, 1998.
19. S. Held, M. Storath, B. Forster, and P. Massopust. Steerable wavelet transforms based on the Riesz transform. *IEEE Trans. Image Processing*, 2009. to appear.
20. E. Hernandez, D. Labate, G. Weiss, and E. Wilson. Oversampling, quasi-affine frames and wave packets. *Appl. Comput. Harmon. Anal.*, pages 111–147, 2004.
21. A. Karasaridis and E. Simoncelli. A filter design technique for steerable pyramid image transforms. In *Acoustics Speech and Signal Processing*. ICASP, May7–10 1996. Atlanta, GA.
22. K.G. Kutyniok, D. Labate, and W-Q. Lim. Sparse directional image representations using the discrete shearlet transform. *Trans Amer. Math. Soc*, 2008. to appear.
23. D. Labate, W. Lim, G. Kutyniok, and G. Weiss. Sparse multidimensional representation using shearlets. In M. Papadakis, A. Laine, and M. Unser, editors, *Wavelets XI*, volume 5914 of *SPIE Proceedings*, pages 247–255, January 2005.
24. P. Perona. Deformable kernels for early vision. *IEEE Trans. Pattern Anal. Machine Intell.*, 17(5):488–499, 1995.
25. J. Portilla, V. Strela, M.J. Wainwright, and E.P. Simoncelli. Image denoising using Gaussian scale mixtures in the wavelet domain. *IEEE Trans. Image Processing*, 12(11):1338–1351, 2003. Computer Science techn. rep. nr. TR2002–831, Courant Institute of Mathematical Sciences.
26. J.R. Romero, S. Alexander, S. Baid, S. Jain, and M. Papadakis. The geometry and the analytic properties of Isotropic Multiresolution Analysis. *Advances in Comp. Math.*, 31:283–328, 2009.

27. E.P. Simoncelli and W.T. Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *ICIP 1995*, Washington, DC, October 1995.
28. T.N. Tan. Rotation invariant texture features and their use in automatic script identification. *IEEE Trans. Pattern Anal. Machine Intell.*, 20(7):751–756, 1998.
29. G. Tzagkarakis, B. Beferull-Lozano, and P. Tsakalides. Rotation-invariant texture retrieval with Gaussianized steerable pyramids. *IEEE Trans. Image Processing*, 15(9):2702–2718, 2006.
30. M. Unser, M. Sage, and D. Van De Ville. Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transform. *IEEE Trans. Image Processing*. to appear.
31. D. Van De Ville, T. Blu, and M. Unser. Isotropic polyharmonic B-splines: scaling functions and wavelets. *IEEE Trans. Image Processing*, 14(11):1798–1813, 2005.
32. D. Van De Ville and M. Unser. Complex wavelet bases, steerability, and the Marr-like pyramid. *IEEE Trans. Image Processing*, 17(11):2063–2080, 2008.
33. Y. Wan and R.D. Nowak. Quasi-circular roation invariance in image denoising. In *ICIP*, volume 1, pages 605–609, 1999.
34. S. Yi, D. Labate, G.R. Easley, and H. Krim. A shearlet approach to edge analysis and detection. *IEEE Trans. Image Processing*, 2009. to appear.