Development of a Digital Shearlet Transform Based on Pseudo-Polar FFT

Gitta Kutyniok[a], Morteza Shahram[b], and David L. Donoho[b]

[a]Institute of Mathematics, University of Osnabrück,
49069 Osnabrück, Germany;

[b]Department of Statistics, Stanford University,
Stanford, CA 94305, USA

ABSTRACT

`Shearlab` is a Matlab toolbox for digital shearlet transformation of two-D (image) data we developed following a rational design process. The Pseudo-Polar FFT fits very naturally with the continuum theory of the Shearlet transform and allows us to translate Shearlet ideas naturally into a digital framework. However, there are still windows and weights which must be chosen. We developed more than a dozen performance measures quantifying precision of the reconstruction, tightness of the frame, directional and spatial localization and other properties. Such quantitative performance metrics allow us to: (a) tune parameters and objectively improve our implementation; and (b) compare different directional transform implementations. We present and interpret the most important performance measures for our current implementation.

Keywords: Density Compensation, Digital Shearlet Theory, Fast Pseudo-Polar Transform, Monte Carlo Estimates, Parabolic Scaling, Pseudo-Polar Grid, Reproducible Research, Test Measures, Shearlet Transform, Testing Environment

1. INTRODUCTION

Applied harmonic analysts introduced in recent years several approaches for directional representations of image data, each one with the intent of efficiently representing highly anisotropic image features. Examples include curvelets,[3–5] contourlets,[6] and shearlets.[12,13,15,16] These proposals are inspired by elegant results in theoretical harmonic analysis, which study functions defined on the continuum plane (i.e. not digital images) and address problems of efficiently representing certain types of functions and operators. One set of inspiring results concerns the possibility of highly compressed representation of 'cartoon' images, i.e. functions which are piecewise smooth with singularities along smooth curves. Another set of results concerns the possibility of a highly compressed representation of wave propagation operators. In 'continuum theory', anisotropic directional transforms can significantly outperform wavelets in important ways.

Accordingly, one hopes that a digital implementation of such ideas would also deliver performance benefits over wavelet algorithms in real-world settings. Anticipated applications include[14] where missing sensors cause incomplete measurements, and the problem of texture/geometry separation in image processing - for example in astronomy when images of galaxies require separated analyses of stars, filaments, and sheets.[7,18]

In many cases, however, there are no publicly available implementations of such ideas, or the available implementations are only sketchily tested or the available implementations are only vaguely related to the continuum transforms they are reputed to represent. Accordingly, we have not yet seen a serious exploration of the potential benefit of such transforms, carefully comparing the expected benefits with those delivered by specific implementations.

In this paper we aim at providing both:

Further author information: (Send correspondence to G. Kutyniok)
D.L.D.: E-mail: donoho@stanford.edu, Telephone: 1 650 723 3350
G.K.: E-mail: kutyniok@uni-osnabrueck.edu, Telephone: 49 541 969 3516
M.S.: E-mail: mshahram@stanford.edu, Telephone: 1 650 725 2236

(1) A rationally designed shearlet transform implementation.

(2) An comprehensive framework for quantifying performance of directional representations in general.

Our ultimate goal is public release of a reproducibly-tested rationally-designed implementation.

For (1), we developed an implementation of the digital shearlet transform based on a digital shearlet theory which is a very natural digitalization of the existing shearlet theory for continuous data. Other parabolic-scaling transforms - we think here of curvelets - are inherently based on operations (rotation) which translate awkwardly into the digital realm. In contrast, when we consider shearlets, rotations are replaced by shearing, which has a natural digital realization.

The framework in (2) has three benefits. First, it provides quantitative performance measures which we use to tune the parameters of our implementation. This allows us to specify 'recommended choices' for the parameters of our implementation. Second, the same 'measure and tune' approach may be useful to other implementers of directional transforms. Third, we show a way to improve the level of intellectual seriousness in applied mathematics which pretends to work in image processing. We believe that widespread adoption of this measure and tune framework can be very valuable, since many supposedly scientific presentations are now little more than vague, numbing 'advertising' or 'marketing' pitches. They could instead offer quantitative comparisons between algorithms, and thereby be far more informative. In fact the combination of quantitative evaluation with reproducible research[11] would be particularly effective at producing both intellectual seriousness and rapid progress.

1.1. Desiderata

Our rational design process considers a list of desiderata that a digital implementation of the shearlet transform may satisfy.

[D1] *Algebraic Exactness.* The transform should give exact reconstruction, up to numerical accuracy.

[D2] *Isometry.* The transform should create a (near-) isometry between the coefficient domain and the image domain.

[D3] *Time-Frequency-Localization.* The frame elements should be spatially anisotropic at fine scales, and well-localized. They should also be localized to the expected anisotropic regions in the frequency domain.

[D4] *Geometric Faithfulness.* The transform should analyze geometric objects in a way familiar from continuum theory: a shearing of the input image should be mirrored in a simple shift of the transform coefficients; and edges in the spatial domain should be mapped to ridges in the coefficient domain.

[D5] *Sparsification.* The transform should analyze objects with the same quantitative properties that are familiar from continuum theory; for example, coefficients of smooth functions should be sparse and coefficients of functions which are smooth away from smooth singularities should be sparse.

[D6] *Robustness.* The reconstruction should degrade gracefully under perturbations such as (hard) thresholding and quantization of coefficients.

[D7] *Speed.* The implementation should run in $O(N^2 log(N))$ flops, where N^2 is the number of digital points of the input image, and the actual performance at specific N such as 512 and 1024 should reflect reasonable leading coefficients.

Although this project is interested in shearlet systems, the desiderata are relevant for other directional transforms, and other transforms could (and should) be considered from this viewpoint as well.

1.2. Shearlets: A Directional Multiscale Representation System based on Parabolic Scaling

The continuum shearlet transform[13,15] (CShT) for functions in $L^2(\mathbb{R}^2)$ uses a two-parameter dilation group, where one parameter indexes scale, and the second parameter indexes orientation. For each $a > 0$ and $s \in \mathbb{R}$, let D_a denote the *parabolic scaling matrix* and S_s denote the *shear matrix* of the form

$$D_a = \begin{pmatrix} a & 0 \\ 0 & \sqrt{a} \end{pmatrix} \qquad \text{and} \qquad S_s = \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix},$$

respectively. To provide an equal treatment of the x- and y-axis, we split the frequency plane into the following four cones $\mathcal{C}_1 - \mathcal{C}_4$:

$$\mathcal{C}_\iota = \begin{cases} \{(\xi_1,\xi_2) \in \mathbb{R}^2 : \xi_1 \geq 1,\, |\xi_1/\xi_2| \geq 1\} & : \quad \iota = 1, \\ \{(\xi_1,\xi_2) \in \mathbb{R}^2 : \xi_2 \geq 1,\, |\xi_1/\xi_2| \leq 1\} & : \quad \iota = 2, \\ \{(\xi_1,\xi_2) \in \mathbb{R}^2 : \xi_1 \leq -1,\, |\xi_1/\xi_2| \geq 1\} & : \quad \iota = 3, \\ \{(\xi_1,\xi_2) \in \mathbb{R}^2 : \xi_2 \leq -1,\, |\xi_1/\xi_2| \leq 1\} & : \quad \iota = 4. \end{cases}$$

We now present the definition for a scaling of 4^j. A shearlet system can be defined similarly for a scaling of 2^j. However, in this case the odd scales have to be handled particularly carefully.

Let now $\psi_1 \in L^2(\mathbb{R})$ be a wavelet with $\hat{\psi}_1 \in C^\infty(\mathbb{R})$ and supp $\hat{\psi}_1 \subseteq [-4, -\frac{1}{4}] \cup [\frac{1}{4}, 4]$, and let $\psi_2 \in L^2(\mathbb{R})$ be a 'bump' function satisfying $\hat{\psi}_2 \in C^\infty(\mathbb{R})$ and supp $\hat{\psi}_2 \subseteq [-1, 1]$. We define $\psi \in L^2(\mathbb{R}^2)$ by

$$\hat{\psi}(\xi) = \hat{\psi}(\xi_1, \xi_2) = \hat{\psi}_1(\xi_1)\, \hat{\psi}_2(\tfrac{\xi_2}{\xi_1}). \tag{1}$$

For cone $\mathcal{C}_1$, at scale $j \geq 0$, orientation $k = -2^j, \ldots, 2^j$, and spatial position $m \in \mathbb{Z}^2$, the associated *shearlets* are then defined by their Fourier transforms

$$\begin{aligned} \hat{\sigma}_\eta(\xi) &= 2^{-j\frac{3}{4}} \hat{\psi}(S_k' D_{4^{-j}} \xi) \chi_{\mathcal{C}_1}(\xi) \exp\{-i(D_{4^{-j}} S_k m)' \xi\} \\ &= 2^{-j\frac{3}{4}} \hat{\psi}_1(\xi_1/4^j) \hat{\psi}_2(k + 2^j \xi_2/\xi_1) \chi_{\mathcal{C}_1}(\xi) \exp\{-i(D_{4^{-j}} S_k m)' \xi\}, \end{aligned}$$

where $\eta = (j, k, m, \iota)$ index scale, orientation, position, and cone. The shearlets for $\mathcal{C}_2 - \mathcal{C}_4$ are defined likewise by symmetry, as illustrated in Figure 1, and we denote the resulting *shearlet system* by

$$\{\sigma_\eta : \eta \in \mathbb{N}_0 \times \{-2^j, \ldots, 2^j\} \times \mathbb{Z}^2 \times \{1, \ldots, 4\}\}. \tag{2}$$

The definition shows that shearlets live on anisotropic regions of width 2^{-2j} and length 2^{-j} at various orienta-

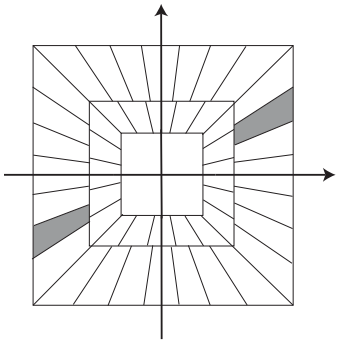

Figure 1. The tiling of the frequency domain induced by shearlets.

tions, which are parametrized by slope rather than angle as for second generation curvelets.

Setting

$$\mathcal{C} = \bigcup_{\iota=1}^{4} \mathcal{C}_\iota,$$

let $L^2(\mathcal{C}^\vee) = \{f \in L^2(\mathbb{R}^2) : \text{supp } \hat{f} \subset \mathcal{C}\}$.

THEOREM 1.1. *[13, Thm. 3] The system* (2) *is a Parseval frame for* $L^2(\mathcal{C}^\vee)$.

The low frequency band can be appropriately filled in to obtain a Parseval frame for $L^2(\mathbb{R}^2)$. Finally, the generating window allows in fact more freedom than (1), but in this paper we restrict ourselves to this (customary) choice.

1.3. Ingredients of the Digital Shearlet Transform

The shearlet transform for continuous data (see Figure 1) implicitly induces a trapezoidal tiling of frequency space which is evidently not cartesian. Introducing a special set of coordinates on the continuum 2D frequency space the CShT can be represented as a cascade of five operations:

- classical Fourier transformation,
- change of variables to pseudo-polar coordinates,
- weighting by a radial 'density compensation' factor,
- decomposition into rectangular tiles,
- inverse Fourier transform of each tile.

The weights referred to in the third step are simply the square root of the Jacobian of the change of variables mapping and amount to a form of *density compensation*.

Surprisingly, this process admits a natural translation into the digital domain. The key observation is that the pseudo-polar coordinates are naturally compatible with digital image processing (compare Figure 2). Since there is a fast pseudo-polar fourier transform,[2] this suggests we can easily and naturally get a faithful DShT using the PPFT.

Using the PPFT requires care. The PPFT as recently presented[2] is not an isometry. Fortunately, the same researchers earlier[1] explored the possibility of isometry and in experiments conducted during that earlier work, it was demonstrated that near-isometry and even isometry were possible by combining oversampling with appropriate weighting. One should work in a *weighted* pseudo-polar transform space, using weights which correspond roughly to the density compensation weights underlying the continuous change of variables. Loosely speaking, the sampling points of the PPFT are highly nonuniform and it is valuable to downweight points in regions of very high density. In fact, our own experiments found that if we use the PPFT with radial oversampling by a factor of at least 8, we found weights giving exact isometry; however, the weights which do this in the discrete world are not derivable from simple density compensation arguments.

Summarizing, the DShT of an $N \times N$ image cascades the following steps:

1) PPFT with oversampling factor of 8 in the radial direction.
2) Multiplication by 'density-compensation-style' weights.
3) Decomposing the pseudo-polar-indexed array into rectangular subbands.
4) Applying the 2D inverse fast Fourier transform to each subband.

This is an exact analogy of the CShT, in which the steps of Fourier transformation and pseudo-polar coordinate change are collapsed into one step. With a careful choice of the weights and the windows, this transform is an isometry. Hence the inverse transform can be computed by merely taking the adjoint in each step.

1.4. Performance Measurement

The above sketch does not uniquely specify an implementation; there is freedom in choice of weights and windows. How can we decide if one choice is better than another one? It seems that currently researchers often use overall system performance on isolated tasks, such as denoising and compression of specific standard images like 'Lena', 'Barbara', etc. However, overall system performance for a system made up of a cascade of steps seems very opaque and at the same time very particular. It seems far better from an intellectual viewpoint to carefully decompose the performance according to a more insightful array of tests, each one motivated by a particular well-understood property we are trying to obtain.

We have developed quantitative performance measures inspired by desiderata we presented in Subsection 1.1. Each performance measure produces a real value or by a real-valued curve, thus providing a standardized framework for evaluation and, especially, comparison.

1.5. Related Work

Several research teams have previously designed algorithms for anisotropic directional transforms: we mention the curvelet implementation 'CurveLab',[3] the contourlet implementation,[6] as well as two shearlet implementations.[12,17] In our opinion, these pioneer efforts demonstrate real progress in directional representation, but further progress is needed and careful attention to design, tuning and testing is essential.

1.6. Contribution of this Paper

The contributions of this paper two-fold. Firstly, we introduce a digital shearlet transform which is *rationally designed* based on a *natural* digitization of shearlet theory. Secondly, we provide a variety of *quantitative performance measures* for directional representations, which allow tuning and comparsion of implementations. Our digital shearlet implementation was tuned utilizing this framework, so we can provide the user community with an optimized representation.

All presented algorithms and tests as well as codes for the displayed figures and tables are provided at URL `ShearLab.org`[9] in the spirit of reproducible research.[11]

2. AN ISOMETRIC PSEUDO-POLAR TRANSFORM

2.1. General Condition on Density Compensation

Given an $N \times N$ image I, it is well known that the Fourier transform $\hat{I}$ of I evaluated on a rectangular $N \times N$ grid is an isometry:

$$\sum_{u,v=-N/2}^{N/2-1} |I(u,v)|^2 = \frac{1}{N^2} \sum_{\omega_x,\omega_y=-N/2}^{N/2-1} |\hat{I}(\omega_x,\omega_y)|^2. \tag{3}$$

We aim to obtain a similar formula for the Fourier transform of I evaluated on the pseudo-polar grid. For this, we first extend the definition of the pseudo-polar grid slightly by introducing an oversampling parameter $R > 0$ in radial direction. This new grid, which we will denote in the sequel by Ω_R, is defined by

$$\Omega_R = \Omega_R^1 \cup \Omega_R^2, \tag{4}$$

where

$$\begin{aligned} \Omega_R^1 &= \{(-4\ell k/(RN), 2k/R) : -N/2 \le \ell \le N/2,\ -RN/2 \le k \le RN/2\}, \\ \Omega_R^2 &= \{(2k/R, -4\ell k/(RN)) : -N/2 \le \ell \le N/2,\ -RN/2 \le k \le RN/2\}. \end{aligned}$$

This grid is illustrated in Figure 2. Notice that the 'original' pseudo-polar grid (see[2]) is a special case of this definition for $R = 2$.

Choosing the weights carefully, a 'Plancherel theorem' similar to (3) can be proved (see[10]) for the pseudo-polar grid $\Omega_R = \Omega_R^1 \cup \Omega_R^2$ defined in (4).

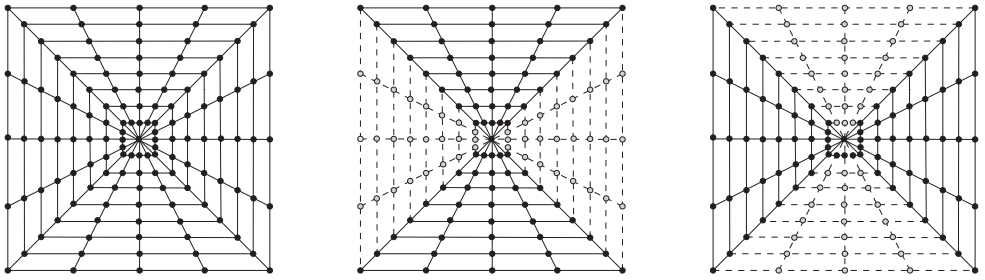

Figure 2. The pseudo-polar grid for $N = 4$ and $R = 4$.

THEOREM 2.1. *Let N be even, and let $w : \Omega_R \to \mathbb{R}^+$ be a weight function. Then*

$$\sum_{u,v=-N/2}^{N/2-1} |I(u,v)|^2 = \sum_{(\omega_x,\omega_y)\in\Omega_R} w(\omega_x,\omega_y)\cdot|\hat{I}(\omega_x,\omega_y)|^2$$

holds if and only if, for all $-N/2 \le u, v \le N/2-1$, the weights w satisfy

$$\begin{aligned}\delta(u,v) \;=\; & (2N+1)\cdot w(0,0)\\ & +4\cdot\sum_{\ell=0,N/2}\sum_{k=1}^{RN/2} w(2k/R,-4\ell k/(RN))\cdot\cos(2ku/(R(2N+1)))\cdot\cos(4\ell kv/(RN(2N+1)))\\ & +8\cdot\sum_{\ell=1}^{N/2-1}\sum_{k=1}^{RN/2} w(2k/R,-4\ell k/(RN))\cdot\cos(2ku/(R(2N+1)))\cdot\cos(4\ell kv/(RN(2N+1)))\end{aligned}$$

and, for all $(\omega_x,\omega_y)\in\Omega_R$, the weights w satisfy the symmetry conditions

$$w(\omega_x,\omega_y) = w(\omega_y,\omega_x),\ w(\omega_x,\omega_y) = w(-\omega_y,\omega_x),\ w(\omega_x,\omega_y) = w(-\omega_x,\omega_y) \text{ and } w(\omega_x,\omega_y) = w(\omega_x,-\omega_y).$$

Moreover, in general, R needs to be at least 8 for such weights to exist.

In fact, the stated conditions the weights have to satisfy lead to a linear system of equations with $RN^2/8+1/16$ variables and N^2 equations. In general, we need the oversampling factor $R \ge 8$ or else we can't have a uniquely solvable system.

2.2. Recommended Choice of Weights

By extensive computational experiments we found weights which provide exact isometry; however, the calculation of such weights is a heavy overhead and the weights themselves are not smooth. Non-smooth weights are problematic; they act as windows on the Fourier side and, seemingly, non-smooth weights could cause synthesis elements to have poor spatial decay. To avoid high complexity and sporadic weighting patterns, we relax the requirement for exact isometric weighting, and instead represent the weights not as the solution of a large system of equations, but instead in terms of an undercomplete basis for functions on the pseudo-polar grid. We compute the coefficients in this expansion once for a given problem size; then hardwire them in the algorithm. For the present algorithm, we expand using 7 functions $w_1,\dots,w_7$ on the pseudo-polar grid such that $\sum_{j=1}^{7} w_j(\omega_x,\omega_y) \ne 0$ for each $(\omega_x,\omega_y)\in\Omega_R$. The center

$$w_1 = 1_{(0,0)} \quad \text{and} \quad w_2 = 1_{\{(\omega_x,\omega_y):|k|=1\}},$$

the boundary

$$w_3 = 1_{\{(\omega_x,\omega_y):|k|=NR/2 \text{ and } \omega_x=\omega_y\}} \quad \text{and} \quad w_4 = 1_{\{(\omega_x,\omega_y):|k|=NR/2 \text{ and } \omega_x\ne\omega_y\}},$$

and the seam lines

$$w_5(\omega_x,\omega_y) = |k| \cdot 1_{\{(\omega_x,\omega_y):1<|k|<NR/2 \text{ and } \omega_x=\omega_y\}} \quad \text{and} \quad w_6 = 1_{\{(\omega_x,\omega_y):|k|=NR/2-3 \text{ and } \omega_x=\omega_y\}}$$

have to be treated with much care compared to the 'interior'

$$w_7(\omega_x,\omega_y) = |k| \cdot 1_{\{(\omega_x,\omega_y):1<|k|<NR/2 \text{ and } \omega_x\neq\omega_y\}}.$$

Notice that here we use (ω_x,ω_y) and (k,ℓ) interchangeably.

The weighting generated by our recommended choice of coefficients for $w_1,\ldots,w_7$ is displayed in Figure 3. The different regions, in particular, the seam lines can clearly be seen.

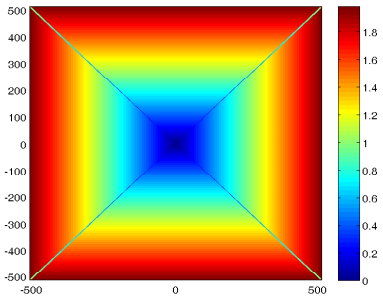

Figure 3. Recommended weighting of the pseudo-polar grid.

Certainly, the isometry suffers from this restricted choice. But tuning the algorithm, we found that with respect to the trade-off between the closeness to isometry and complexity of the algorithm, this was the best choice.

3. QUANTITATIVE TEST MEASURES

To study our implementation and the extent to which it satisfies our desiderata [D1]-[D7], we have defined performance measures covering each desideratum; in Section 4 we provide numerical results illustrating how our current implementation performs as of July 6, 2009. These performance measures provide

- a means to quantify the performance of our algorithm,
- a framework for tuning and improving our algorithm,
- a basis for comparison of parabolic scaling algorithms in general.

In the following, P shall denote the operator defined by the code which computes the pseudo-polar tranform, w shall denote the weighting applied to the values on the pseudo-polar grid, W shall be the windowing with additional iFFT, and S shall denote the operator defined by the code which computes the shearlet transform as well as its adjoint will be denotde by $S^\star$.

[D1] *Algebraic Exactness.*
We generate a sequence of 5 random images $I_1,\ldots,I_5$ on a pseudo-polar grid for $N = 256$ and $R = 8$ with standard normally distributed entries. Our Monte Carlo estimate for the operator norm $\|W^\star W - Id\|_{op}$ is

$$M_{[D1]} = \max_{i=1,\ldots,5} \frac{\|W^\star W I_i - I_i\|_2}{\|I_i\|_2}.$$

This measure controls the extent to which one can exactly reconstruct an image by standard iterative algorithms; if it is less than 1, then standard relaxation converges. Controlling $M_{[D1]}$ also controls the speed of conjugate-gradient-type solvers.

[D2] *Isometry.*

- *Closeness to tightness.* Generate a sequence of 5 random images $I_1, \ldots, I_5$ of size 256×256 with standard normally distributed entries. Our quality measure is a Monte Carlo estimate for the operator norm $\|P^* wP - Id\|_{op}$ given by

$$M_{[D2],1} = \max_{i=1,\ldots,5} \frac{\|P^* wPI_i - I_i\|_2}{\|I_i\|_2}.$$

- *Quality of preconditioning.* We measure the eigenvalue spread of the Gram operator $P^\star wP$ given by

$$M_{[D2],2} = \frac{\lambda_{\max}(P^\star wP)}{\lambda_{\min}(P^\star wP)}.$$

- *Tight Frame Property.* Generate a sequence of 5 random images $I_1, \ldots, I_5$ of size 256×256 with standard normally distributed entries. We compute a Monte Carlo estimate for the operator norm $\|S^* S - Id\|_{op}$ given by

$$M_{[D2],3} = \max_{i=1,\ldots,5} \frac{\|S^* SI_i - I_i\|_2}{\|I_i\|_2}.$$

[D3] *Time-Frequency-Localization.*
Let I be a shearlet in a 256×256 image centered at the origin $(129, 129)$ with slope 0 of scale 5. We make four measurements:

- *Decay in Spatial Domain.* We compute the decay rates $d_1, \ldots, d_{256}$ along lines parallel to the y-axis starting from the line $[129, :]$ and the decay rates $d_{257}, \ldots, d_{512}$ with x and y interchanged. By decay rate, for instance, for the line $[129 : 256, 1]$, we first compute the smallest monotone majorant $M(x, 1)$, $x = 129, \ldots, 256$ – note that we could also choose an average amplitude here or a different 'envelope' – for the curve $|I(x, 1)|$, $x = 129, \ldots, 256$. Then the decay rate is defined to be the average slope of the line, which is a least square fit to the curve $\log(M(x, 1))$, $x = 129, \ldots, 256$. Based on these decay rates, we choose our measure to be

$$M_{[D3],1} = \max_{i=1,\ldots,512} d_i.$$

- *Decay in Frequency Domain.* Here we intend to check whether the Fourier transform of I is compactly supported and also the decay. For this, let $\hat{I}$ be the 2D-FFT of I and compute the decay rates d_i, $i = 1, \ldots, 512$ as before. Then we define the measure by

$$M_{[D3],2} = \max_{i=1,\ldots,512} d_i.$$

- *Smoothness in Spatial Domain.* We will measure smoothness by Hölder regularity and minimize $|I(u) - I(v)|/|u - v|^{\tilde{\alpha}}$ over all $u \neq v$ with respect to $\tilde{\alpha}$. The value of α, for which the minimum is reached, will define our measure

$$M_{[D3],3} = \alpha.$$

- *Smoothness in Frequency Domain.* We compute the smoothness now for $\hat{I}$, the 2D-FFT of I to obtain the new α, which will be our measure

$$M_{[D3],4} = \alpha.$$

[D4] *Geometric Fathfulness.* We measure performance in three ways. The first two consider images $I_1, \ldots, I_8$ be 256×256 of a planar heaviside function with discontinuity passing through the origin $(129, 129)$ and having one of the five slopes $[-1, -0.5, 0, 0.5, 1]$, as well as three other images produced by turning such images on their sides. We let $c_{i,j}$ be the associated shearlet coefficients for image I_i and scale j. We consider:

- *Decay of significant coefficients.* Consider the curve $j \mapsto \frac{1}{8}\sum_{i=1}^{8} \max|c_{i,j}|$, where the coefficients $c_{i,j}$ are associated with shearlets aligned with the line. Let d be the average slope of the line, which is a least square fit to log of this curve, and define

$$M_{[D4],1} = d.$$

- *Decay of insignificant coefficients.* Consider the curve $j \mapsto \frac{1}{8}\sum_{i=1}^{8} \max|c_{i,j}|$, where the coefficients $c_{i,j}$ are associated with shearlets except the ones aligned with the line. Let d be the average slope of the line, which is a least square fit to log of this curve, and define

$$M_{[D4],2} = d.$$

- *Shear Invariance.* For this measure, now let I be an 256×256 image with a cross of the form $[129, :] \cup [:, 129]$. For each scale j we choose the set $S_j = \{s : 2^{-j}s \in \mathbb{Z} \text{ and } \hat{\psi}(S_{2^{-j}s}^T A_4^{-j} \cdot) \text{ does not touch the seam line}\}$, and generate a sequence of images $I_{s,j} = I(S_s \cdot)$, $s \in S_j$. Our quality measure will then be the curve

$$j \mapsto \max_{s \in S_j} \frac{\|SI_{s,j} - C_{j,s}SI\|_2}{\|I\|_2},$$

where $C_{j,s}$ shifts the shear parameter of the coefficients by $2^j s$.

[D5] *Robustness.* We consider two performance measures:

- *Thresholding.* Let I be a 256×256 digital array sampling a Gaussian bump with mean 0 and variance 256 on a regular grid $\{-128, 127\}^2$. Our performance measure for $k = 1, 2$ is the curve

$$p_k \mapsto \frac{\|S^\star \operatorname{thres}_{p_k} SI - I\|_2}{\|I\|_2},$$

where

  - $\operatorname{thres}_{p_1}$ discards $100 \cdot (1 - 2^{-p_1})$ percent of the coefficients ($p_1 = [2 : 2 : 20]$),
  - $\operatorname{thres}_{p_2}$ sets all those coefficients to zero with absolute values below the threshold $m/2^{p_2}$ with m being the maximal absolute value of all coefficients. ($p_2 = [0.5 : 0.5 : 5]$)

- *Quantization.* Let I contain the pixel values of a Gaussian bump with mean 0 and variance 256 sampled on a 256×256 grid $\{-128, 127\}^2$. Our performance measure is the curve

$$[5 : -0.5 : 0.5] \ni q \mapsto \frac{\|S^\star \operatorname{quant}_q SI - I\|_2}{\|I\|_2},$$

where $\operatorname{quant}_q(c) = \operatorname{round}(c/(m/2^q)) \cdot (m/2^q)$ and m being the maximal absolute value of all coefficients.

[D6] *Speed.* Generate a sequence of 5 random images I_i, $i = 5, \ldots, 9$ of size $2^i \times 2^i$ with standard normally distributed entries. Let s_i be the measured execution time in CPU seconds for the Shearlet Transform S applied to I_i. Our hypothesis is that the speed behaves like $s_i = c \cdot (2^{2i})^d$; 2^{2i} being the size of the input. Let now $\tilde{d}_a$ be the average slope of the line, which is a least square fit to the curve $i \mapsto \log(s_i)$. Let also f_i be the measured execution time in CPU seconds for the 2D-FFT applied to I_i, $i = 5, \ldots, 9$. Our three performance measures quantify the exponential order, and the leading constant, in the computational complexity, and also a comparison with the running time of the 2D-FFT:

$$M_{[D6],1} = \frac{\tilde{d}_a}{2\log 2}, \quad M_{[D6],2} = \frac{1}{5}\sum_{i=5}^{9} \frac{s_i}{(2^{2i})^{M_{[D6],1}}} \quad \text{and} \quad M_{[D6],3} = \frac{1}{5}\sum_{i=5}^{9} \frac{s_i}{f_i}.$$

4. TEST RESULTS

Our development effort has reached a mature state for the most important criteria, namely [D1]-[D2] and [D6]. We are improving performance on the other criteria daily, and the reader should consult the values reported in our talk at the SPIE conference, or else at the website, `Shearlet.org`. In this article, we present current values for the performance measures [D1]-[D2] and [D6]; they offer the most lasting information we can present at this time.

4.1. Results for Tests [D1]-[D2]

Table I presents measurements concerning algebraic exactness $M_{[D1]}$, the isometry of the pseudo-polar transform $M_{[D2],1}$ and $M_{[D2],2}$, and the isometry of the shearlet transform $M_{[D2],3}$, which our recommended version achieves.

TABLE I
RESULTS FOR [D1]-[D2]

$M_{[D1]}$	$M_{[D2],1}$	$M_{[D2],1}$	$M_{[D2],3}$
0.0023147827	0.016212074	1.6584499	0.017172738

The tightness deficiency of our recommended transform is only 0.017. The successive values in the table indicate how key steps in the implementation (weighting or windowing) perform, hence where efforts to further improve this error should be directed. Clearly, density-compensation weighting produces an error of about 0.016; with the recommended weighting, the quotient of the eigenvalues of the Gram matrix is about 1.658, significantly far from 1. The error of 0.016 should be compared to the error 0.0023 arising from the windowing, which has a smaller order of magnitude. Evidently, progress on the choice of weights could have a significant impact on the tightness error. Keep in mind that there is a trade-off between the sophistication of the weights, the running time of the algorithm, and the smoothness of the shearlet basis elements; compare Subsection 2.2.

Figure 4 shows the error matrix for the three tests.

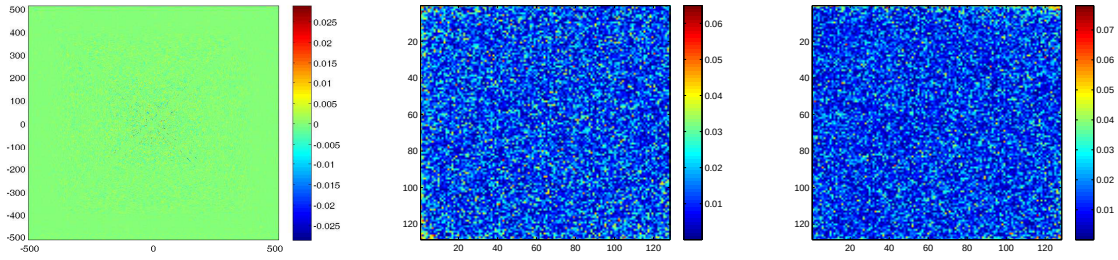

Figure 4. Error image for tests [D1], and [D2], with $N = 128$, $R = 8$, and $\beta = 2$.

4.2. Timing Tests – Desideratum [D6]

The timing performance measures - the complexity $M_{[D6],1}$, the constant $M_{[D6],2}$, and the comparison with 2D-FFT $M_{[D6],3}$ - are presented in Table III.

TABLE III
RESULTS FOR [D6]

$M_{[D6],1}$	$M_{[D6],2}$	$M_{[D6],3}$
0.42944697	1.0258062	255.77436

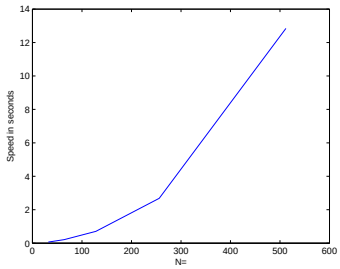

Figure 5. Timing of the forward Shearlet Transform for $R = 8$ and $\beta = 2$. Horizontal axis: N. Vertical axis: Execution Time in Seconds

The timing data are depicted in Figure 5.

Superficially, these measures would seem to say that our implementation has *sublinear* running time characteristics, scaling approximately as N, for an $N \times N$ image. However, we believe this is simply a start-up effect caused by the significant overhead that the implementation imposes on very small problems, which becomes less important for large problems. For the same reason, we don't believe that the factor of 250 comparison to running time of the FFT adequately describes the comparison for large problem sizes N; the comparison gets much more favorable to the DShT as N grows large.

In particular, we emphasize that the *correct* way to use shearlets and other directional ideas is not as doctrinaire as we have been testing here. It is well known that, at the pixel scale, 'real' images are filled with so-called 'jaggy' effects that aren't related to underlying properties of the scene; they instead reflect the digitization of the scene. Accordingly, a mixed transform is appropriate for 'real' image content, in which the finest-scale subband uses wavelet subbands rather than shearlet subbands, and coarser-scale subbands use shearlets. As it turns out, most of the expense in our implementation occurs at the finest scale, which can be skipped entirely in such a mixed transform.

5. DISCUSSION

Our work makes several contributions.

5.1. Rational Design of Digital Parabolic Scaling Algorithm

We implemented the shearlet transform following a rational design process:

- A digitization framework was provided, which offers a natural digitization of the continuum transform.
- Our implementation follows this digital theory, with various choices of windowing guided by quantitative performance.

This process makes it quantitatively clear that we are obtaining relatively fast and relatively accurate transforms.

5.2. Standardized Framework for Evaluation and Comparison

A similar process can be followed for other mathematically defined parabolic scaling systems such as, for instance, curvelets or contourlets. It provides several benefits.

Firstly, researchers can objectively optimize the performance of a given algorithm. A researcher can tune the parameters of his algorithm to improve performance in a rigorous, quantitative sense. Vague claims which we often hear in conference - such as 'Now the algorithm seems to perform better' can then be replaced by standardized quantitative statements.

Secondly, researchers can objectively compare several different representations and/or implementations within a standardized framework. This adds a new chapter to the philosophy of reproducible research,[11] since once

implemented, everybody from the community has not only access to newly designed algorithms, but also the means for a fair comparison.

Our general hope is to systematize the design of directional transform algorithms. Certainly, other testing measures can be imagined and added to the list. However, we believe that this present list contains and formalizes the core problems digital parabolic scaling algorithms can encounter. Our testing environment is publicly provided at URL `ShearLab.org`.[9]

By sharing the complete code for the tests, we would like to encourage researchers, who believe that important aspects of a parabolic scaling algorithm are left out, to define additional tests and add those to our testing framework. Also we share the hope that researchers will support the idea of having a standardized elaborate test dictionary, which greatly extends the common noise suppression and compression tests using 'Lena', etc. The reader shall be reminded that our tests are specifically designed for parabolic scaling algorithms.

5.3. Results for Present Version

We found the performance metric 'measure and tune' framework very helpful in designing our current implementation. Several components of our implementation were optimized - in particular, the choice of weights, the tiling of the pseudo-polar grid, and choice of the window functions. We presented tables giving values of performance metrics for our currently recommended code. Our code with the recommended choice of parameters is publicly available at URL `ShearLab.org`.[9]

We don't claim the results for our algorithm are the best ones which can ever be achieved! We promote further tuning of this implementation as well as bringing in totally new ideas to further optimize the performance measures.

6. CONCLUSIONS

We have developed a digital shearlet transform based on a carefully defined digital shearlet theory, which is the natural digitization of the continuum shearlet transform. Our digital shearlet transform was rationally designed, by studying the effect of various choices – windowing and weighting – on the accuracy and speed of the implementation. The main steps of our approach are a weighted pseudo-polar Fourier transform to achieve an almost-isometric mapping into the pseudo-polar domain, thereby using a particular oversampling strategy, as well as a careful selection of tiling windows. We defined performance measures which could be used to study not only our shearlet implementation, but indeed any digital implementation of a general parabolic scaling algorithm. Our digital shearlet transform was tuned to optimize its performance by those measures. We presented several of the most important performance measures, documenting a successful implementation strategy.

ACKNOWLEDGEMENTS

GK would like to thank Demetrio Labate and Wang-Q Lim for inspiring discussions on shearlet theory and algorithmic aspects. GK would also like to thank the Department of Statistics at Stanford University and the Department of Mathematics at Yale University for their hospitality and support during her visits. This work was partially supported by Deutsche Forschungsgemeinschaft (DFG) Heisenberg Fellowship KU 1446/8-1.

REFERENCES

1. A. Averbuch, R. R. Coifman, D. L. Donoho, M. Israeli, J. Walden, and Y. Shkolnisky, *Fast slant stack: A notion of radon transform for data in a cartesian grid which is rapidly computible, algebraically exact, geometrically faithful and invertible*, preprint.
2. A. Averbuch, R. R. Coifman, D. L. Donoho, M. Israeli, and Y. Shkolnisky, *A framework for discrete integral transformations I – the pseudo-polar Fourier transform*, SIAM J. Sci. Comput. **30** (2008), 764–784.
3. E. J. Candès, L. Demanet, D. L. Donoho and L. Ying, *Fast discrete curvelet transforms,* Multiscale Model. Simul. **5** (2006), 861–899.
4. E. J. Candès and D. L. Donoho, *Continuous curvelet transform: I. Resolution of the wavefront set*, Appl. Comput. Harmon. Anal. **19** (2005), 162–197.

5. E. J. Candès and D. L. Donoho, *Continuous curvelet transform: II. Discretization of frames*, Appl. Comput. Harmon. Anal. **19** (2005), 198–222.
6. M. N. Do and M. Vetterli, *The contourlet transform: an efficient directional multiresolution image representation*, IEEE Trans. Image Process. **14** (2005), 2091–2106.
7. D. L. Donoho and G. Kutyniok, *Microlocal Analysis of the Geometric Separation Problem*, preprint.
8. D. L. Donoho and G. Kutyniok, *Sparsity Equivalence of Anisotropic Decompositions*, preprint.
9. D. L. Donoho, G. Kutyniok, and M. Shahram, `www.ShearLab.org`.
10. D. L. Donoho, G. Kutyniok, and M. Shahram, *ShearLab: A Rational Design of a Digital Parabolic Scaling Algorithm*, preprint.
11. D. L. Donoho, A. Maleki, M. Shahram, V. Stodden, and I. Ur-Rahman, *Fifteen years of Reproducible Research in Computational Harmonic Analysis,* Computing in science and engineering **11** (2009), 8–18.
12. G. Easley, D. Labate, and W. Lim, *Sparse Directional Image Representations using the Discrete Shearlet Transform*, Appl. Comput. Harmon. Anal. **25** (2008), 25–46.
13. K. Guo, G. Kutyniok, and D. Labate, *Sparse Multidimensional Representations using Anisotropic Dilation and Shear Operators*, Wavelets and Splines (Athens, GA, 2005), Nashboro Press, Nashville, TN (2006), 189-201.
14. F. J. Herrmann and G. Hennenfent, *Non-parametric seismic data recovery with curvelet frames*, Geophys. J. Int. **173** (2008), 233-248.
15. G. Kutyniok and D. Labate, *Resolution of the Wavefront Set using Continuous Shearlets*, Trans. Amer. Math. Soc. **361** (2009), 2719–2754.
16. G. Kutyniok, D. Labate, and W. Lim, `www.Shearlet.org`.
17. W. Lim, *Discrete Shearlet Transform: New Multiscale Directional Image Representation* Proceedings of SAMPTA'09, Marseille 2009.
18. J.-L. Starck, E. Candès, and D. L. Donoho, *Astronomical image representation by the curvelet tansform*, Astronomy and Astrophysics **398** (2003), 785–800.