2D and 3D Multiscale Geometric Transforms

Jean-Luc Starck

CEA Saclay, service d'Astrophysique

Collaborators: Arnaud Woiselle (SAGEM)
Jalal Fadili (GREYC, Université de Caen)

What is a good sparse representation for data?

A signal *s* (*n* samples) can be represented as sum of weighted elements of a given dictionary

Dictionary (basis, frame)

$$\Phi = \{\phi_1, \ldots, \phi_K\}$$

Atoms

$$s = \sum_{k=1}^{K} \alpha_k \phi_k = \Phi \alpha$$

coefficients

Ex: Haar wavelet

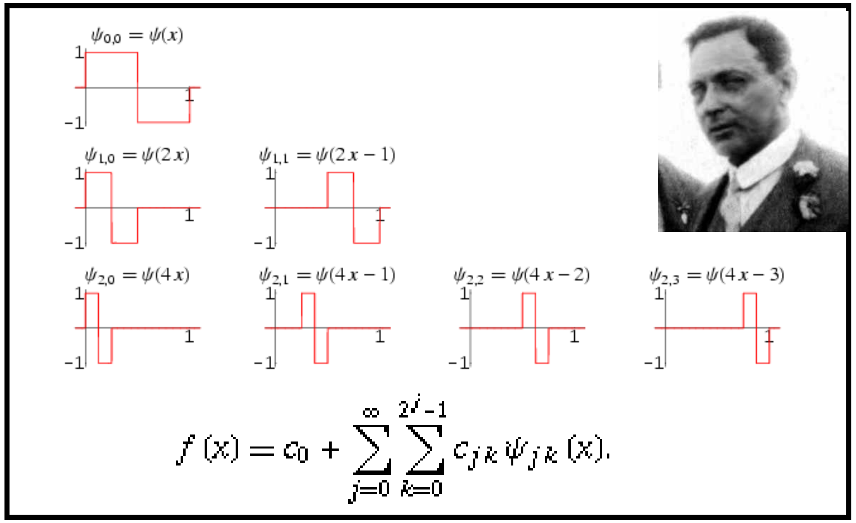

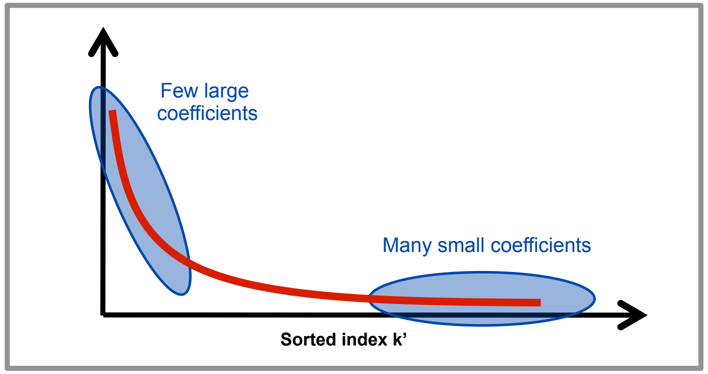

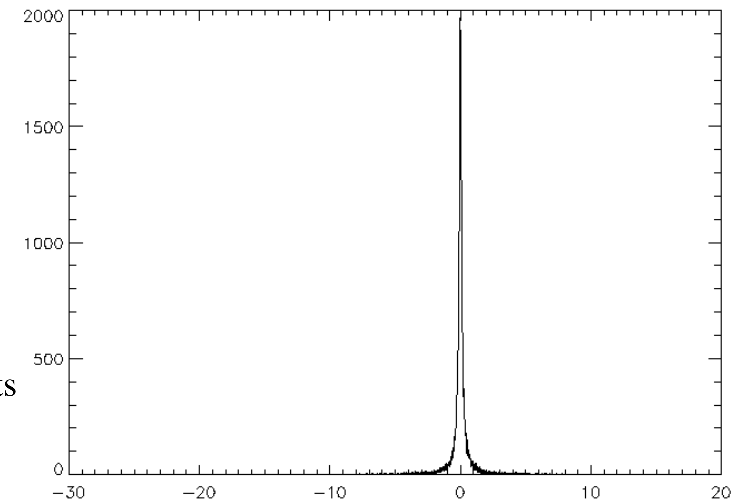

- Fast calculation of the coefficients
- Analyze the signal through the statistical properties of the coefficients
- Approximation theory uses the sparsity of the coefficients

Local DCT

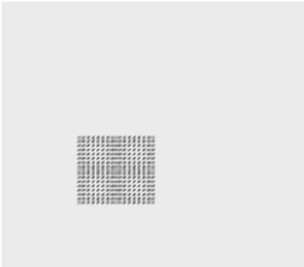

Stationary textures

Locally oscillatory

Wavelet transform

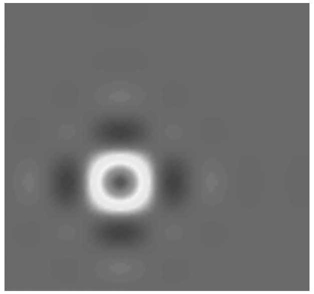

Piecewise smooth

Isotropic structures

Curvelet transform

Piecewise smooth, edge

Sparsity Model 1: Multiscale Transforms

<u>Critical Sampling</u>

(bi-) Orthogonal WT
Lifting scheme construction
Wavelet Packets
Mirror Basis

<u>Redundant Transforms</u>

Pyramidal decomposition (Burt and Adelson)
Undecimated Wavelet Transform
Isotropic Undecimated Wavelet Transform
Complex Wavelet Transform
Steerable Wavelet Transform
Dyadic Wavelet Transform
Nonlinear Pyramidal decomposition (Median)

New Multiscale Construction

Contourlet
Bandelet
Finite Ridgelet Transform
Platelet
(W-)Edgelet
Adaptive Wavelet
Groupelet

Ridgelet
Curvelet (Several implementations)
Wave Atom

Morphological Diversity

- *J.-L. Starck, M. Elad, and D.L. Donoho, Redundant Multiscale Transforms and their Application for Morphological Component Analysis, Advances in Imaging and Electron Physics, 132, 2004.*
- *J.-L. Starck, M. Elad, and D.L. Donoho, Image Decomposition Via the Combination of Sparse Representation and a Variational Approach, IEEE Trans. on Image Proces., 14, 10, pp 1570--1582, 2005.*

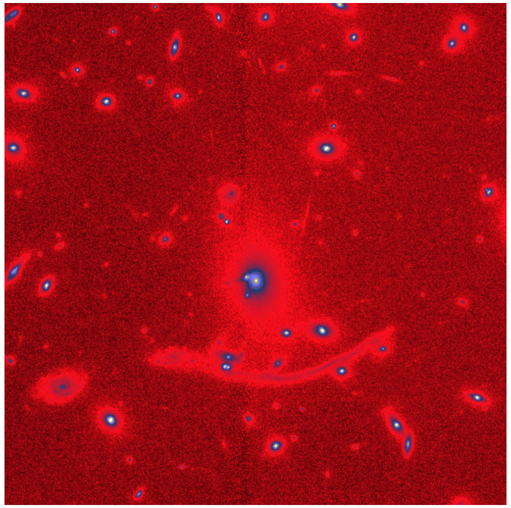

$$\phi = [\phi_1,\ldots,\phi_L], \quad \alpha = \{\alpha_1,\ldots,\alpha_L\}, \quad s = \phi\alpha = \sum_{k=1}^{L} \phi_k \alpha_k$$

Sparsity Model 2: we consider a signal as a sum of K components s_k, $s = \sum_{k=1}^{K} s_k$ each of them being sparse in a given dictionary :

$$s_k = \Phi_k \alpha_k$$

$$s = \sum_{k=1}^{K} s_k = \sum_{k=1}^{K} \Phi_k \alpha_k = \Phi\alpha$$

Sparsity Model 1: we consider a dictionary which has a fast transform/reconstruction operator:

$$\Phi = \{\phi_1, \dots, \phi_K\}$$

$$s = \sum_{k=1}^{K} \alpha_k \phi_k = \Phi\alpha$$

Local DCT

Stationary textures

Locally oscillatory

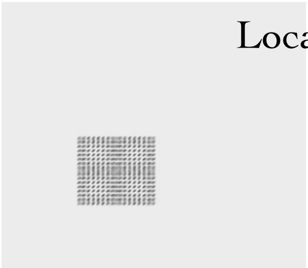

Wavelet transform

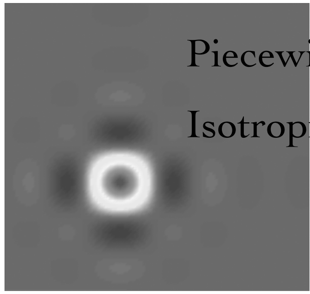

Piecewise smooth

Isotropic structures

Curvelet transform

Piecewise smooth, edge

Sparsity Model 2: Morphological Diversity:

$$\phi = [\phi_1, \dots, \phi_L], \quad \alpha = \{\alpha_1, \dots, \alpha_L\}, \quad s = \phi\alpha = \sum_{k=1}^{L} \phi_k \alpha_k$$

Sparsity Model 3: we adapt/learn the dictionary directly from the data

Model 3 can be also combined with model 2:

G. Peyre, M.J. Fadili and J.L. Starck, , "Learning the Morphological Diversity", **SIAM Journal of Imaging Science**, 3 (3) , pp.646-669, 2010.

Sparsity Model 3: Dictionary Learning

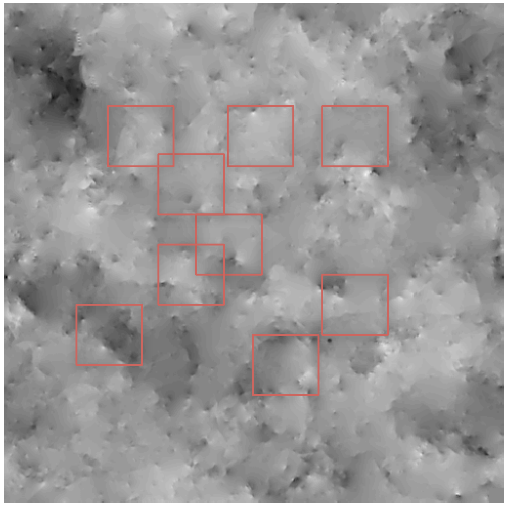

Training basis.

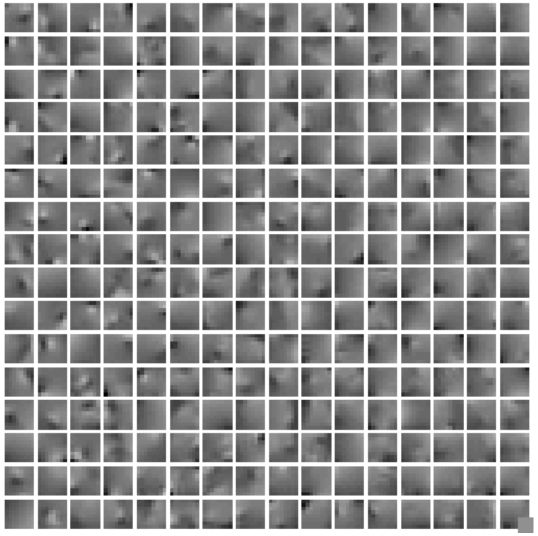

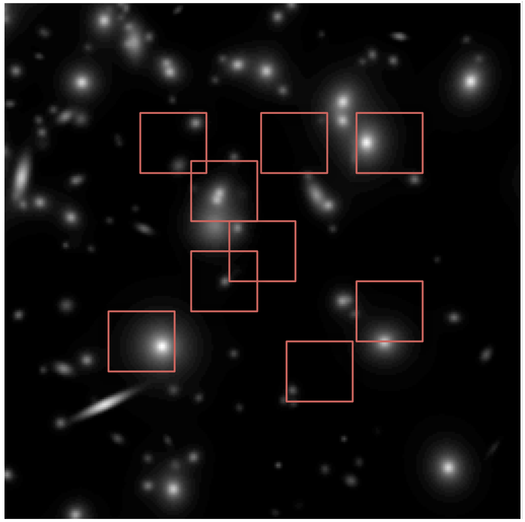

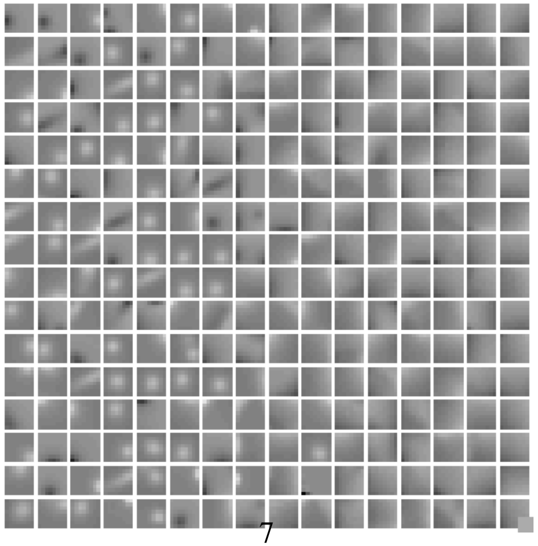

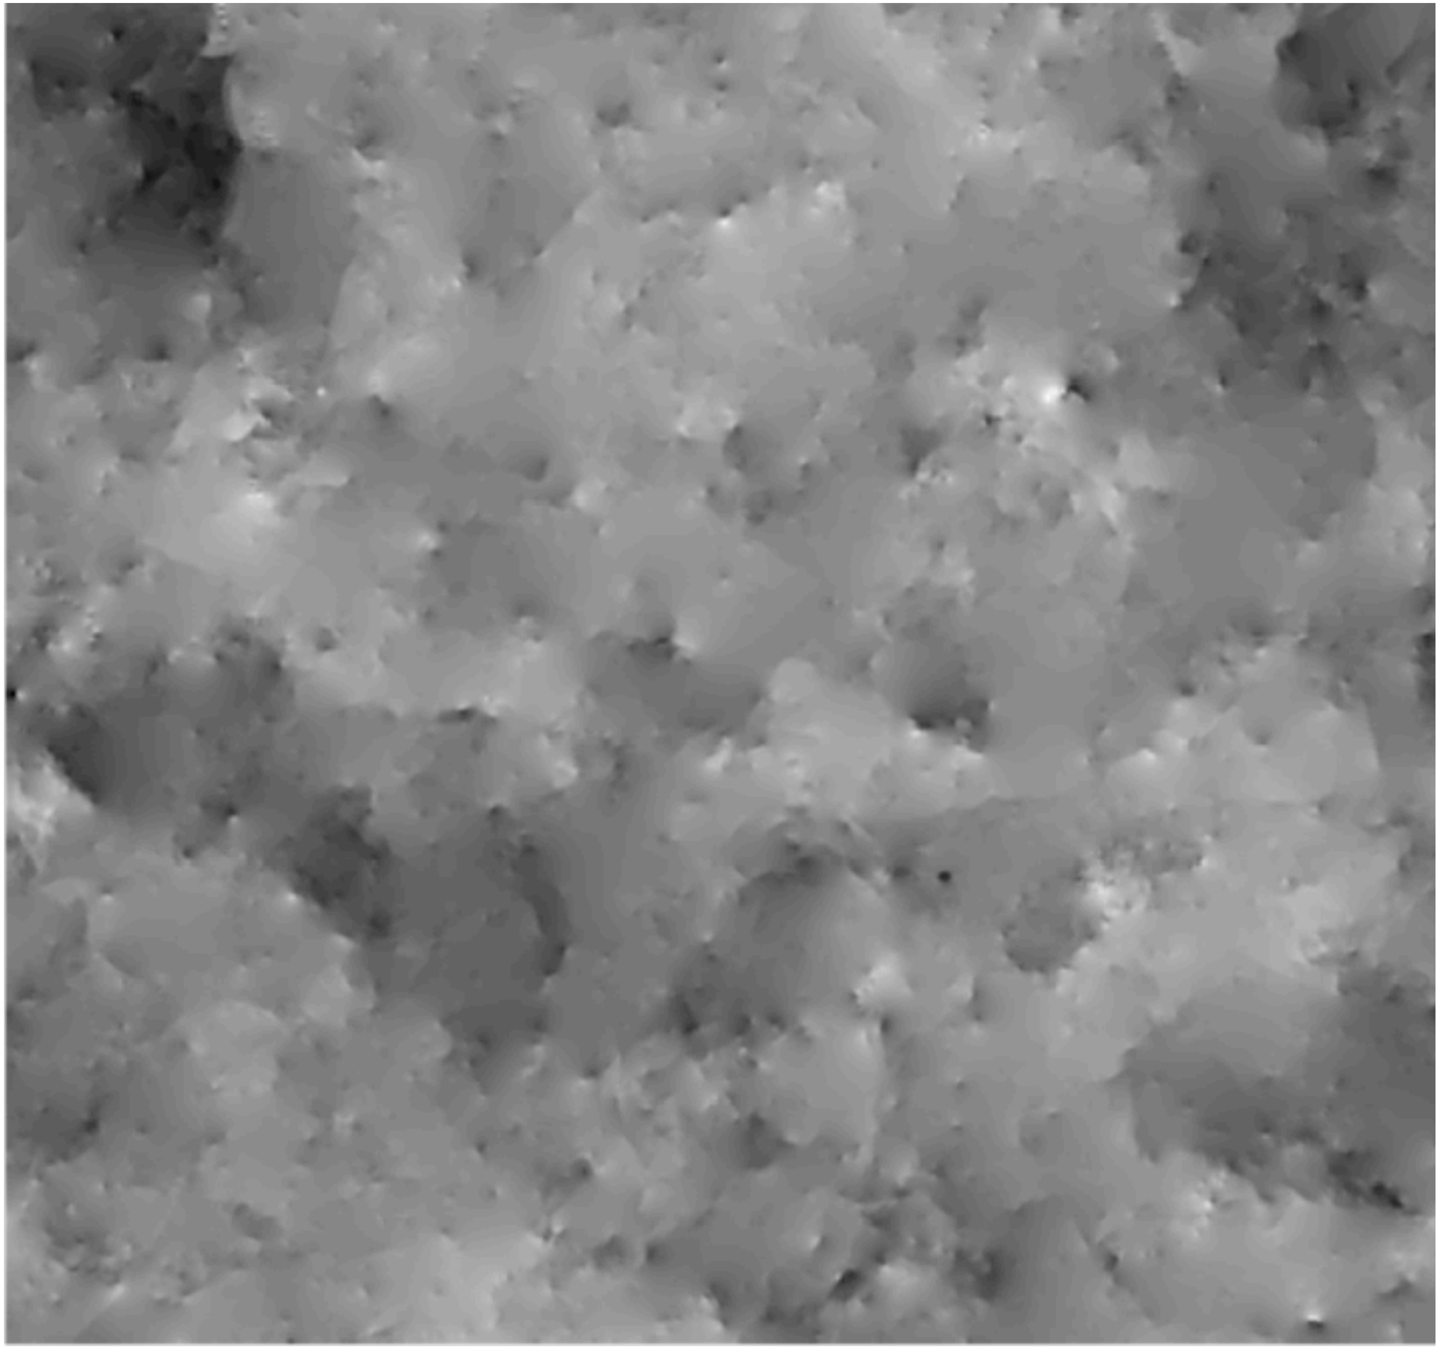

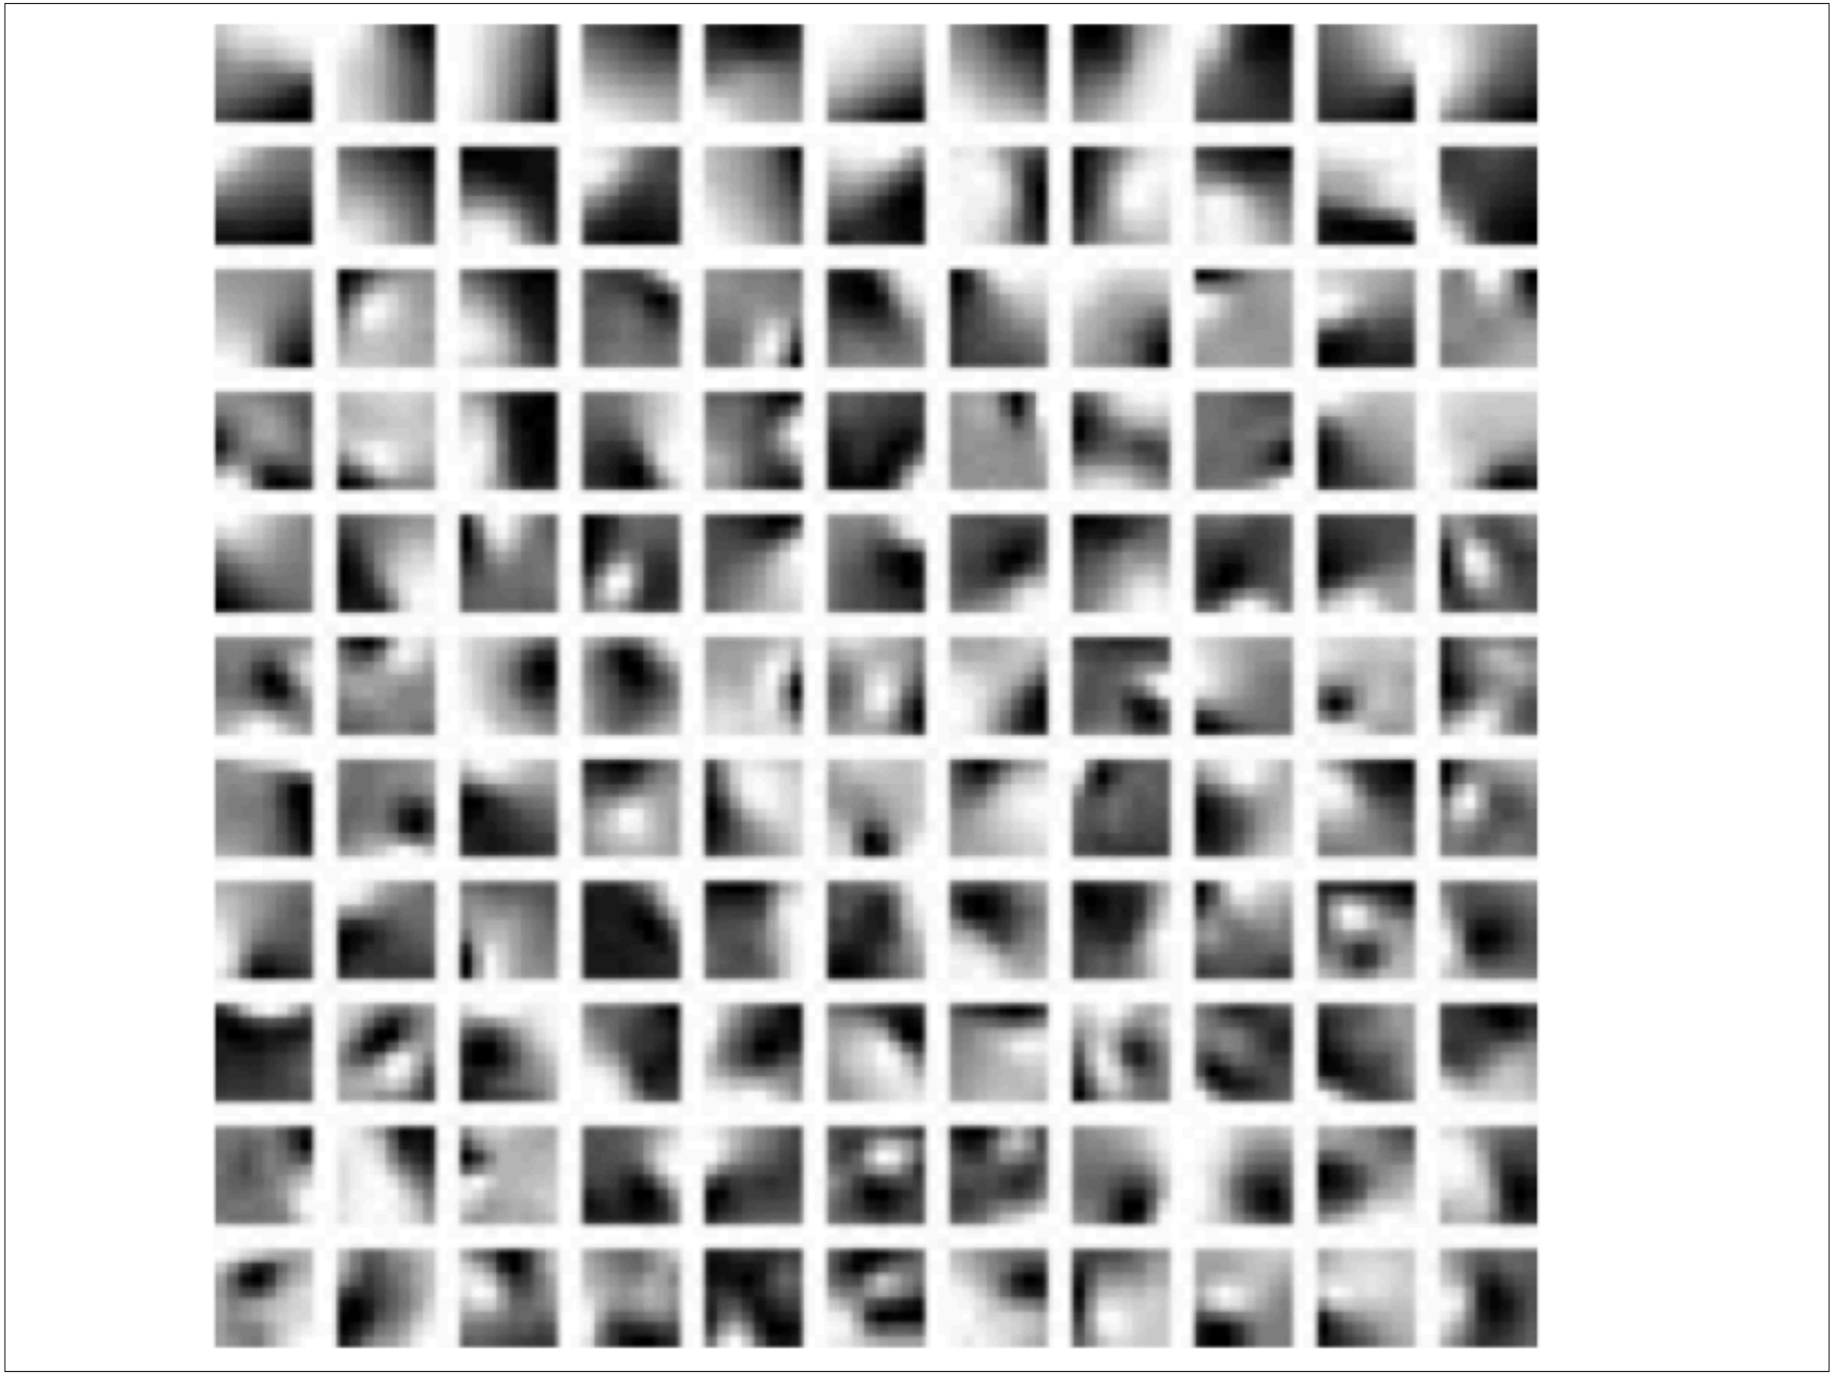

2D and 3D Multiscale Geometric Transforms

- Wavelets 2D
- Ridgelet 2D
- Cuvelet 2D
- BeamCurvelet 3D
- RidCurvelet 3D
- FastCurvelet 3D
- 3D Morphological Diversity

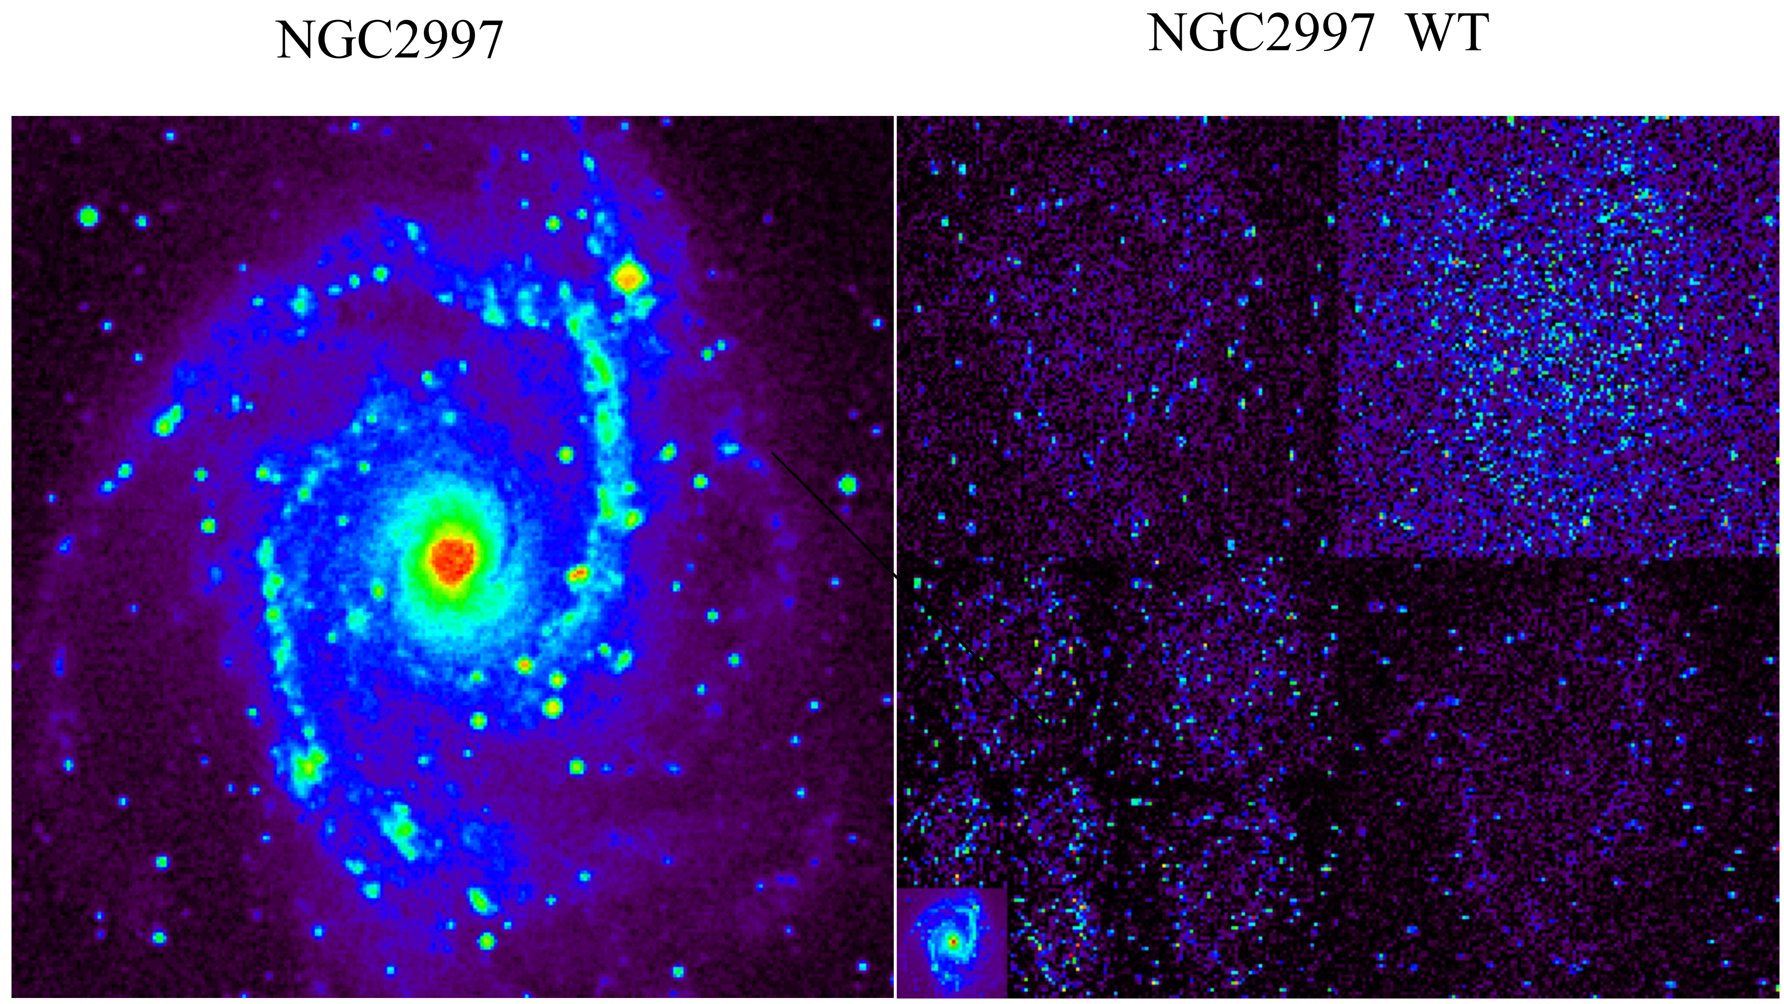
NGC2997
NGC2997 WT

Problems related to the WT

1) Edges representation:
if the WT performs better than the FFT to
represent edges in an image, it is still not optimal.

2) There is only a fixed number of directional elements
independent of scales.

3) Limitation of existing scale concepts:
there is no highly anisotropic elements.

Wavelets and edges

• many wavelet coefficients are needed to account for edges i.e. singularities along lines or curves :

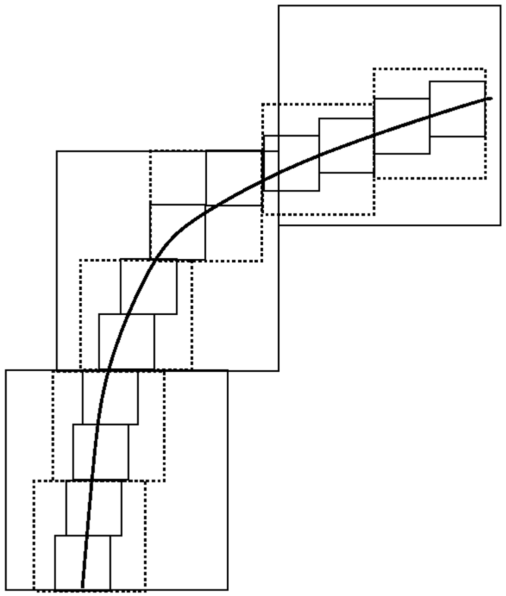

• need dictionaries of strongly anisotropic atoms :

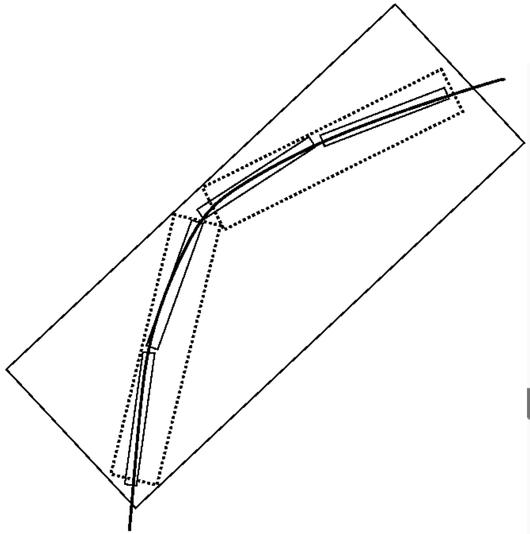

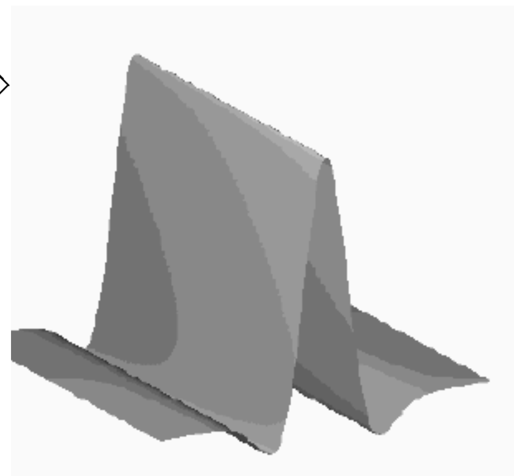

ridgelets, curvelets, contourlets, bandelettes, etc.

SNR = 0.1

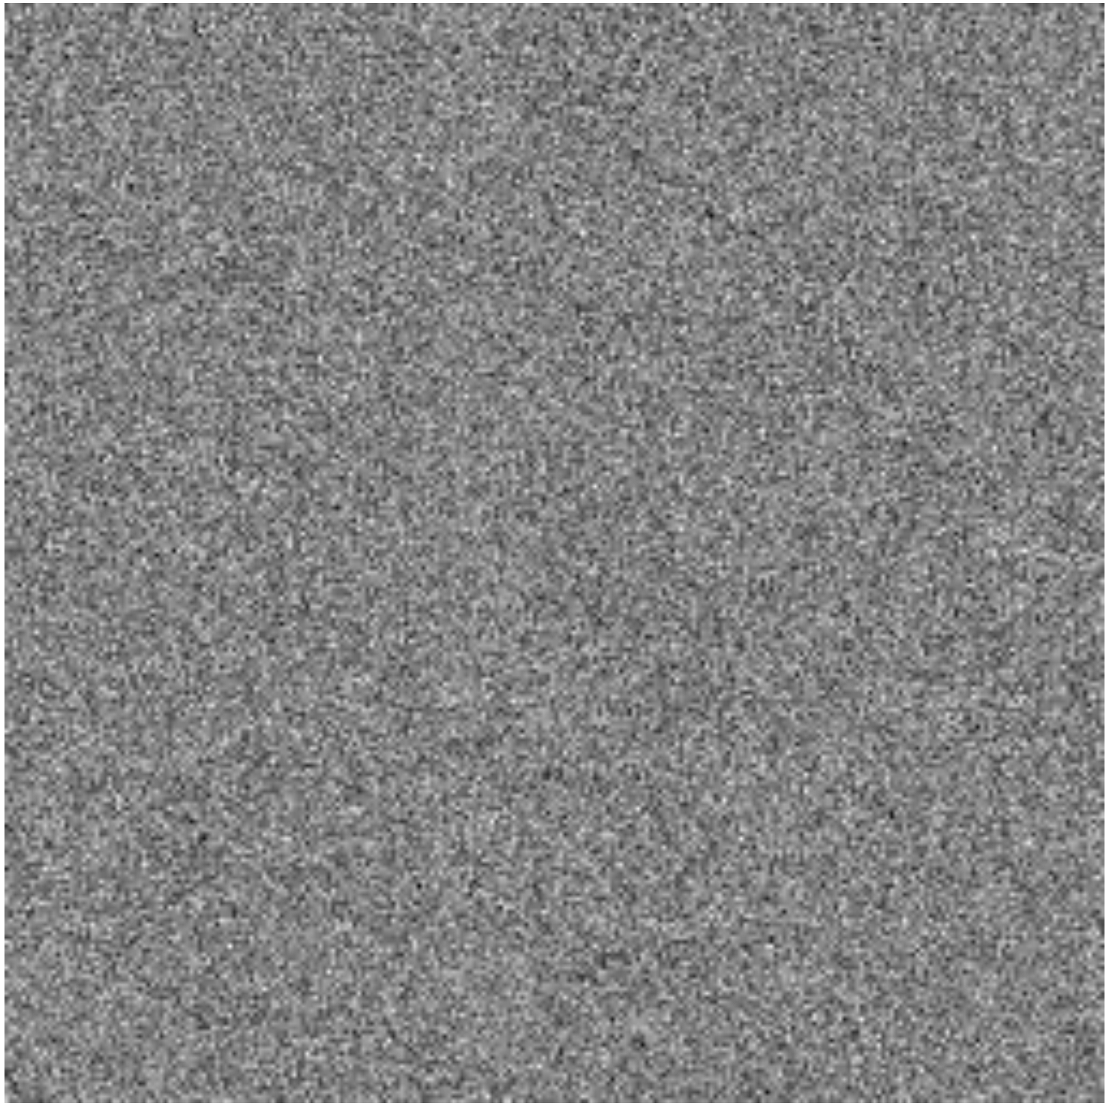

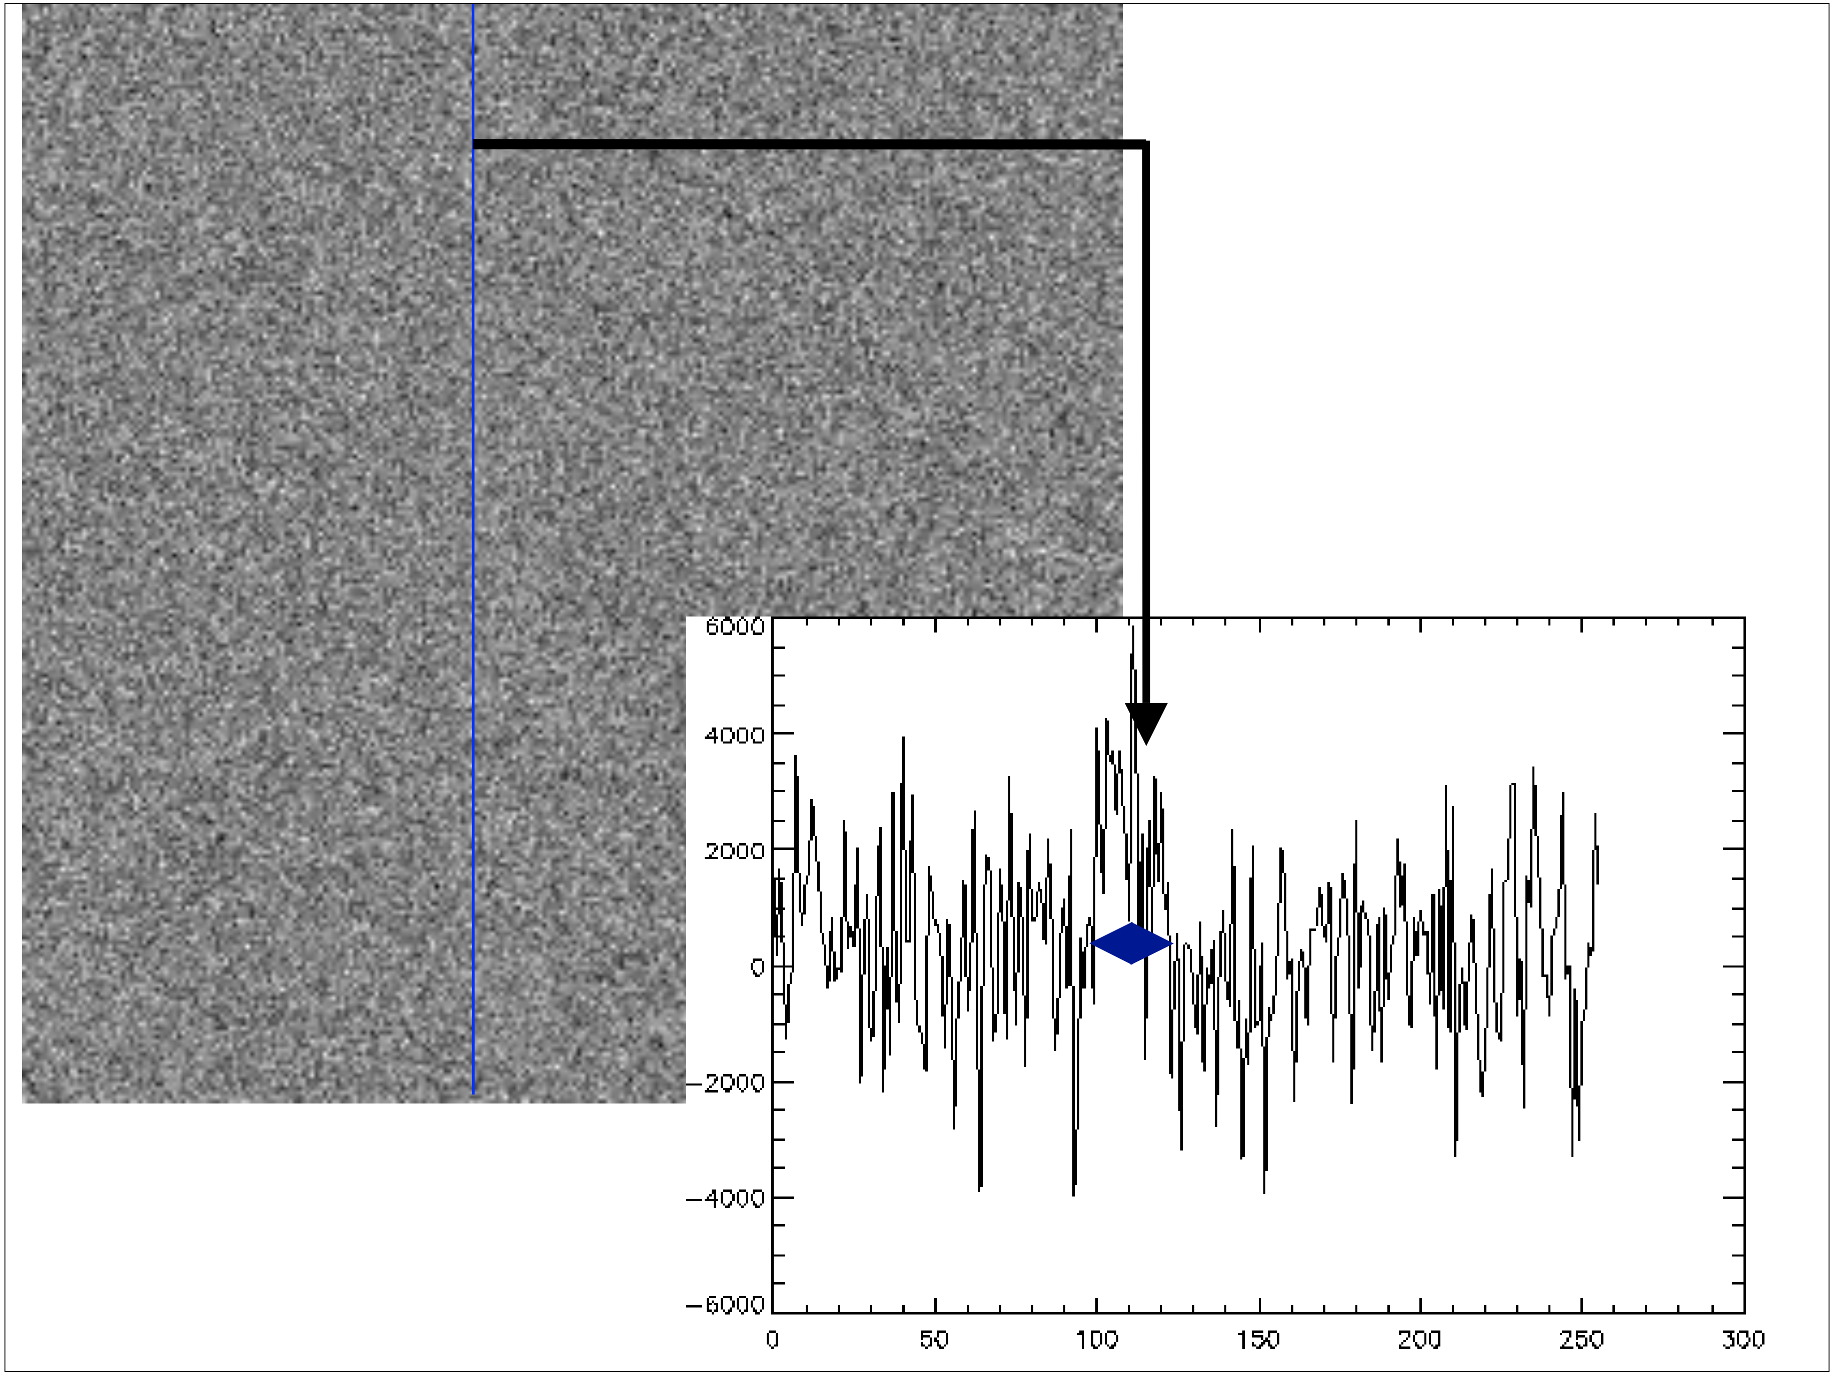
6000
4000
2000
0
−2000
−4000
−6000
0
50
100
150
200
250
300

Undecimated Wavelet Filtering (3 sigma)

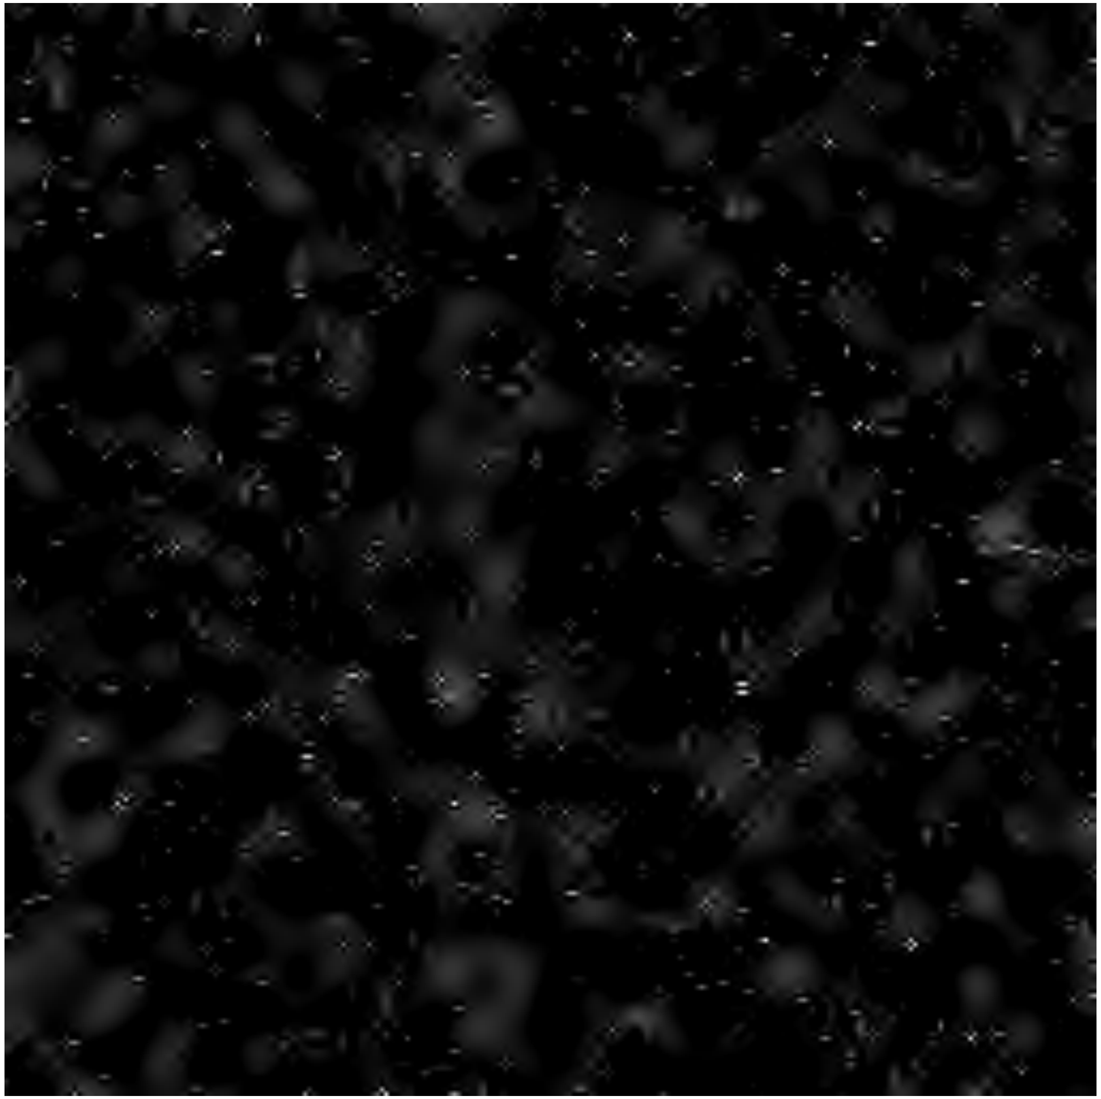

Ridgelet Filtering (5sigma)

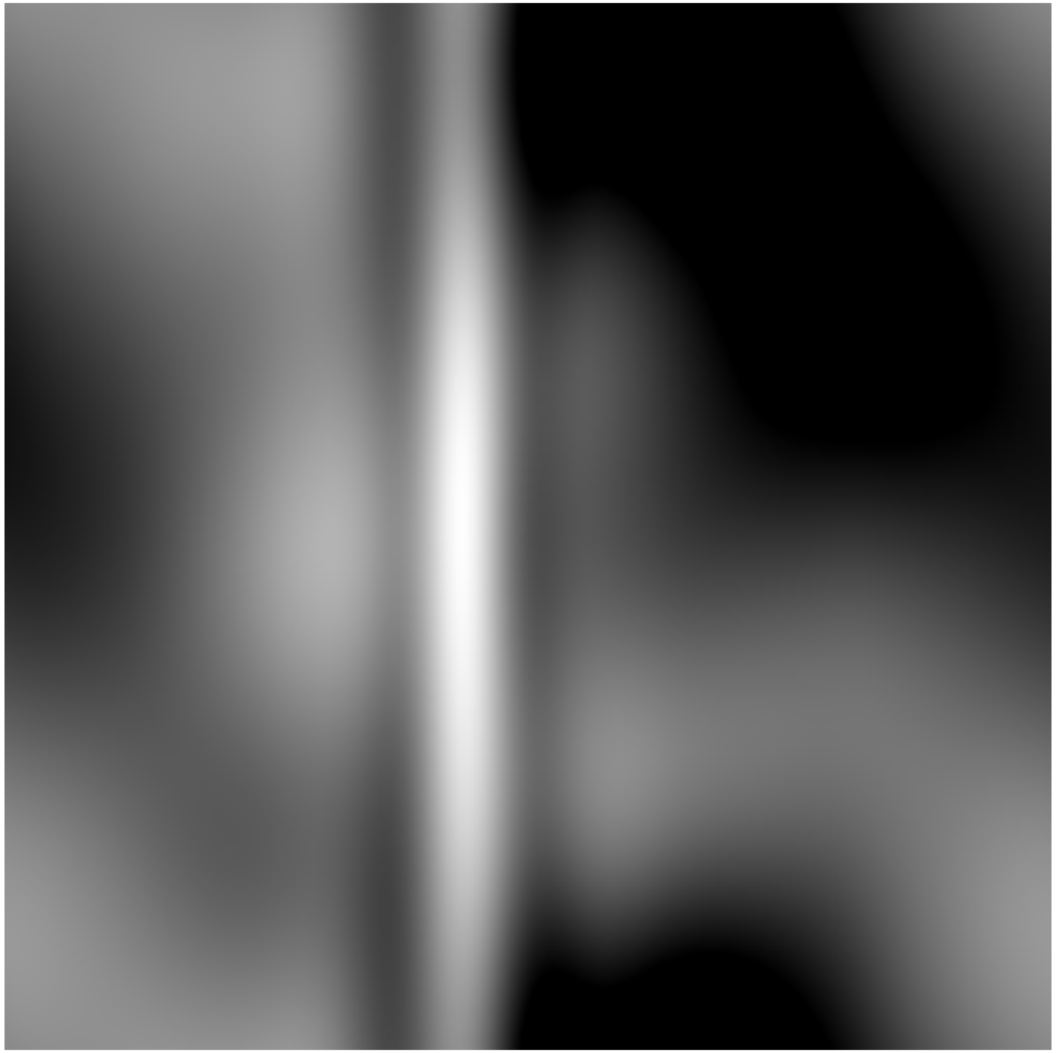

Continuous Ridgelet Transform

Ridgelet Transform (Candes, 1998):

$$R_f(a,b,\theta) = \int \psi_{a,b,\theta}(x) f(x)\,dx$$

Ridgelet function:

$$\psi_{a,b,\theta}(x) = a^{\frac{1}{2}} \psi\left(\frac{x_1 \cos(\theta) + x_2 \sin(\theta) - b}{a} \right)$$

The function is constant along lines. Transverse to these ridges, it is a wavelet.

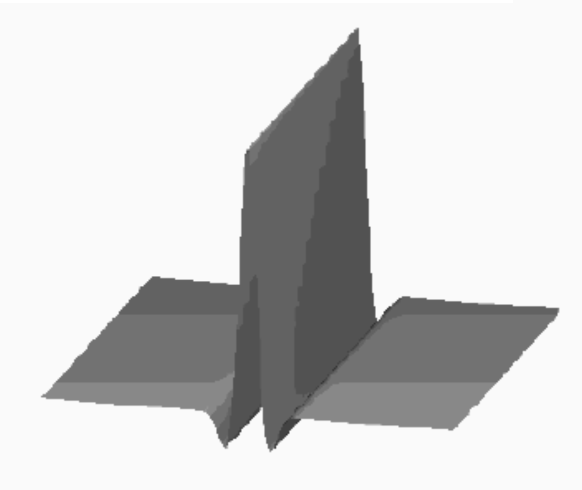

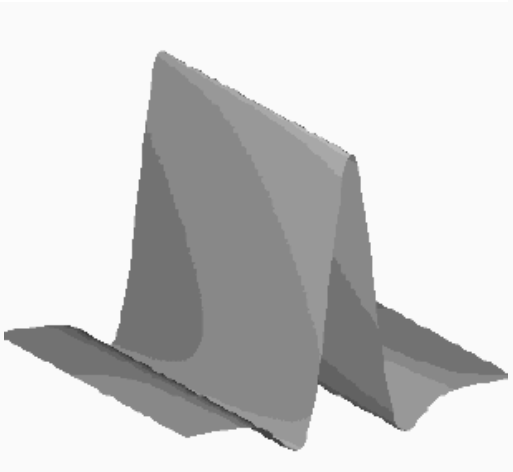

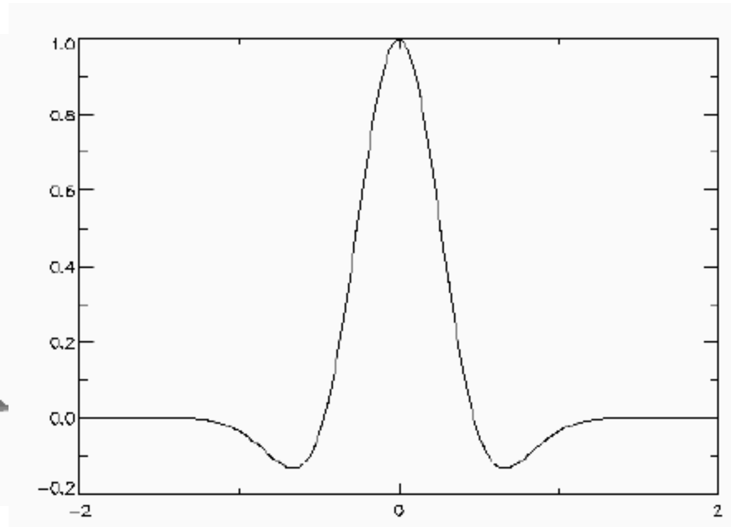

The ridgelet coefficients of an object f are given by analysis

of the Radon transform via: $$R_f(a,b,\theta) = \int Rf(\theta,t)\psi(\frac{t-b}{a})dt$$

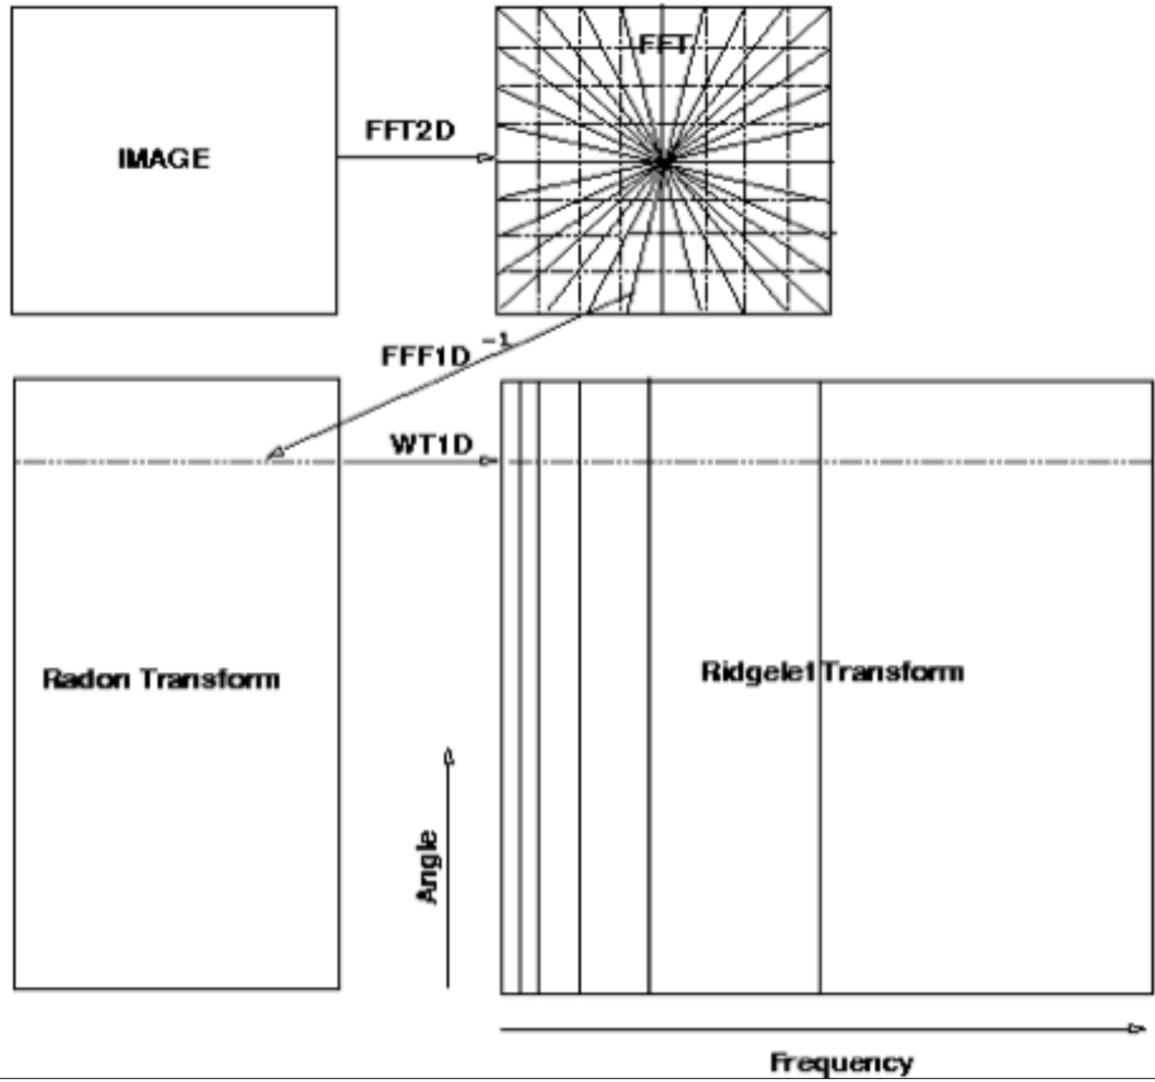

Slant Stack Radon Transform (Averbuch et al, 2001) CUR01-SSR

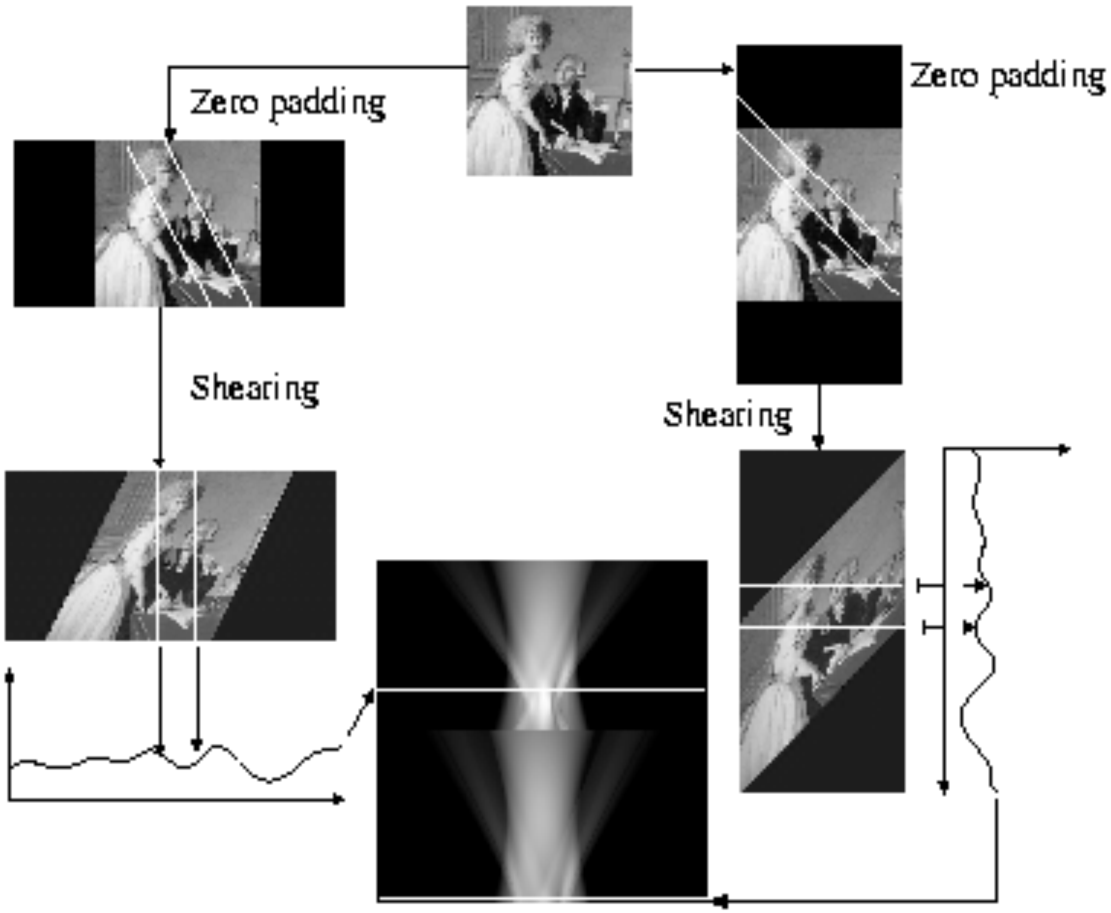

2Nx2N

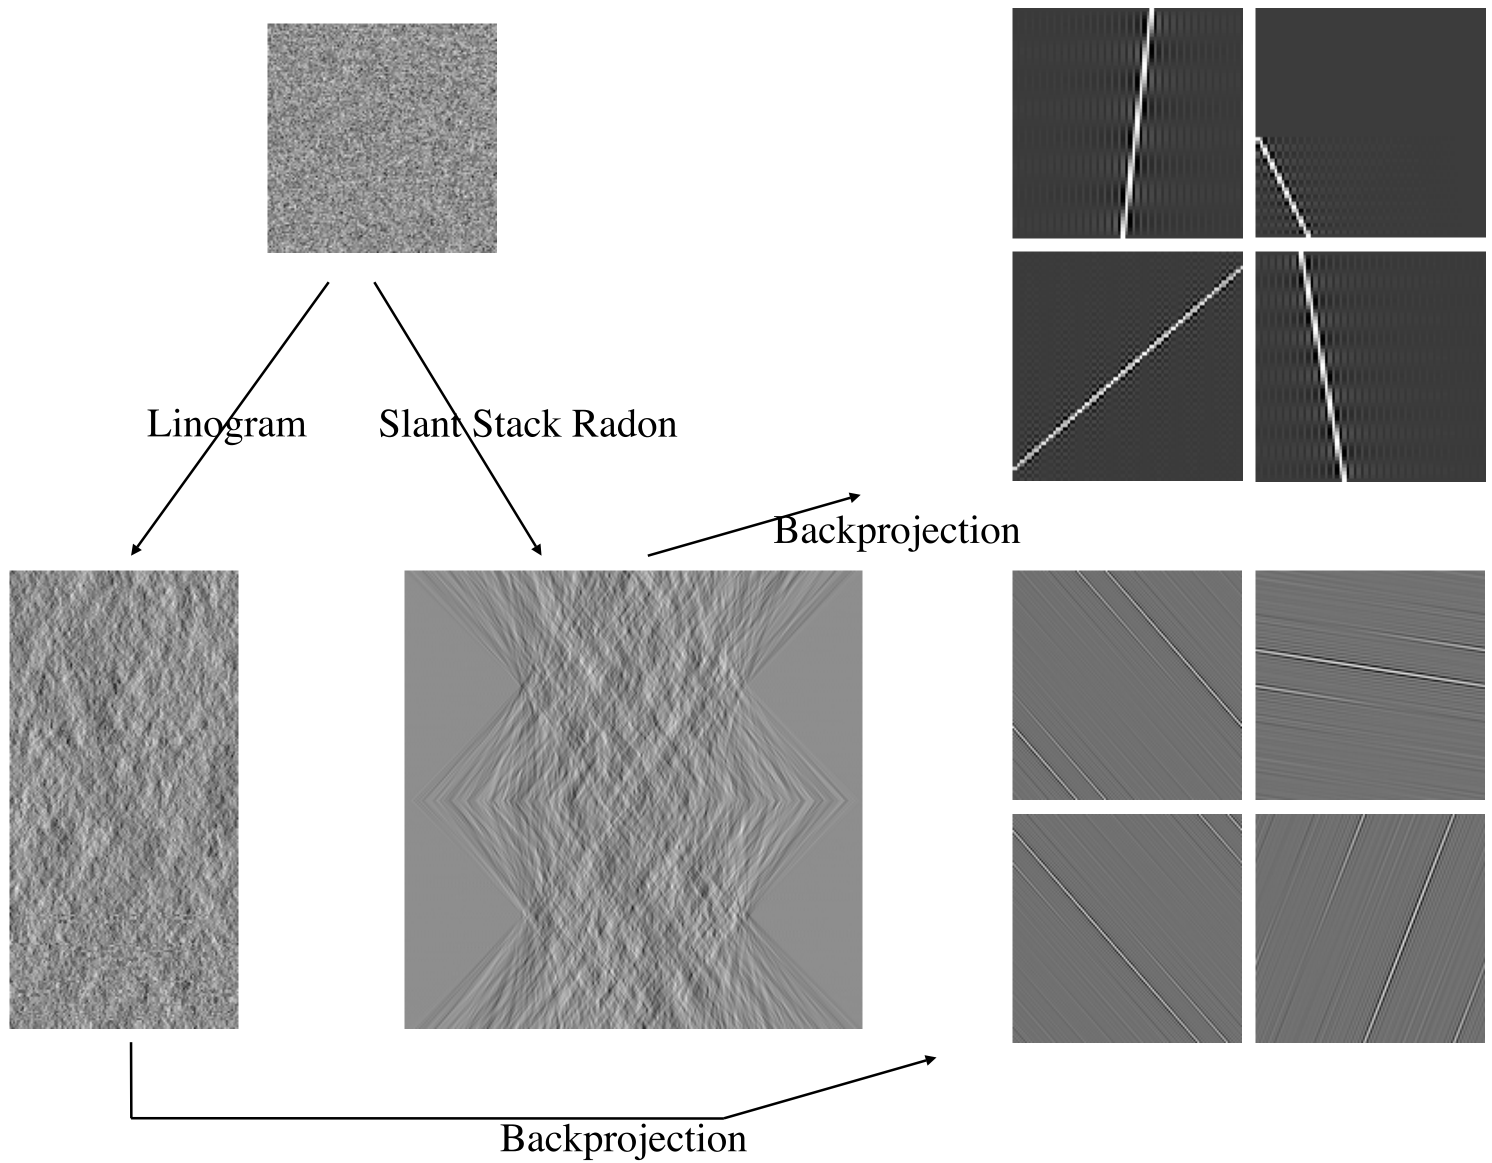
Linogram
Slant Stack Radon
Backprojection
Backprojection

Ridgelet Transform

- Ridgelet transform: Radon + 1D Wavelet

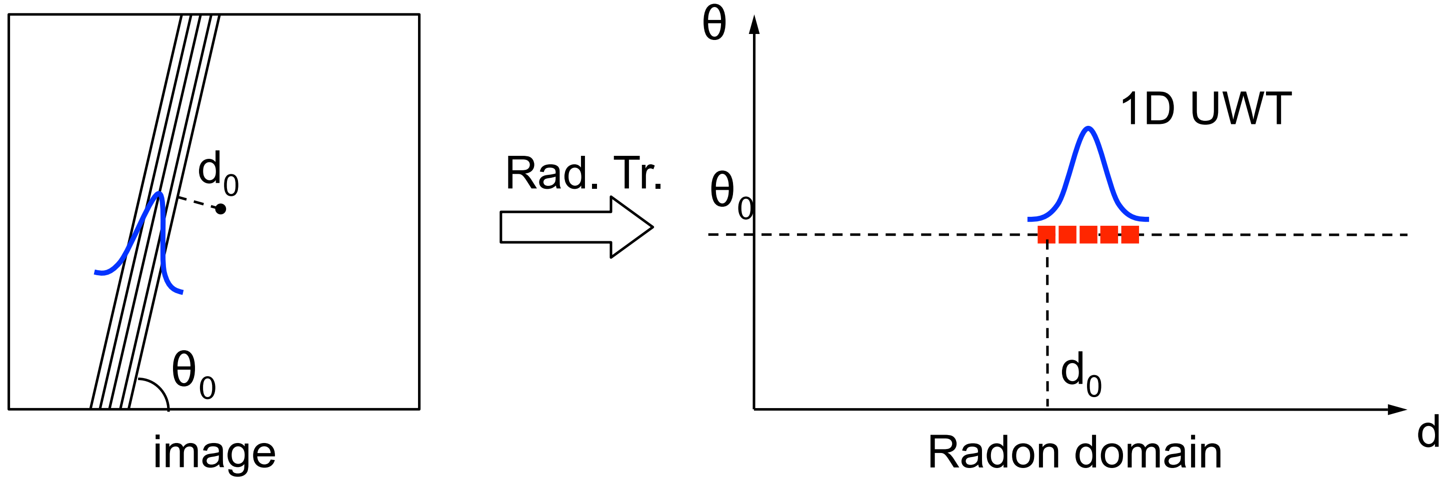

1. Rad. Tr.
2. For each line, apply the same denoising scheme as before

LOCAL RIDGELET TRANSFORM

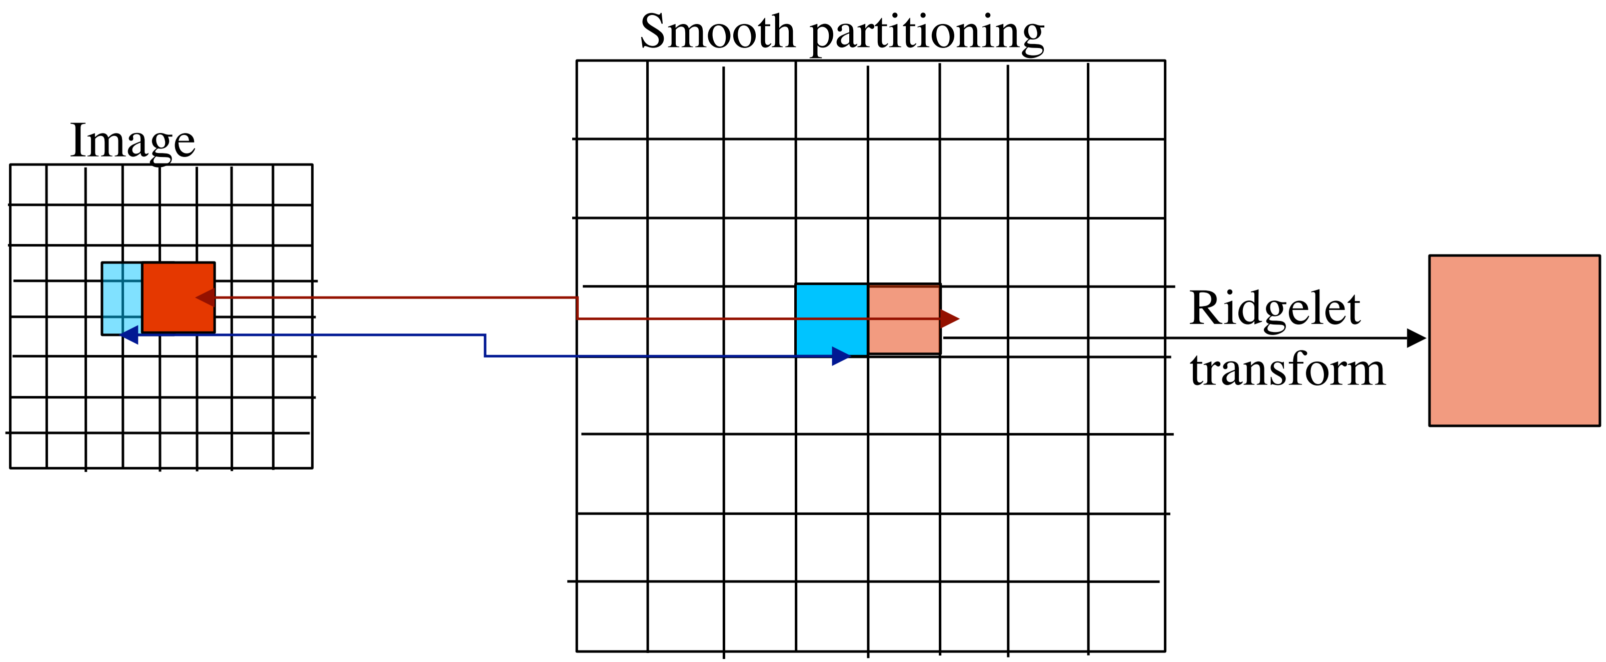

The partitioning introduces a redundancy, as a pixel belongs to 4 neighboring blocks.

Poisson Noise and Line-Like Sources Restoration (MS-VST + Ridgelet)

B. Zhang, M.J. Fadili and J.-L. Starck, "Wavelets, Ridgelets and Curvelets for Poisson Noise Removal" ,ITIP, Vol 17, No 7, pp 1093--1108, 2008.

underlying intensity image

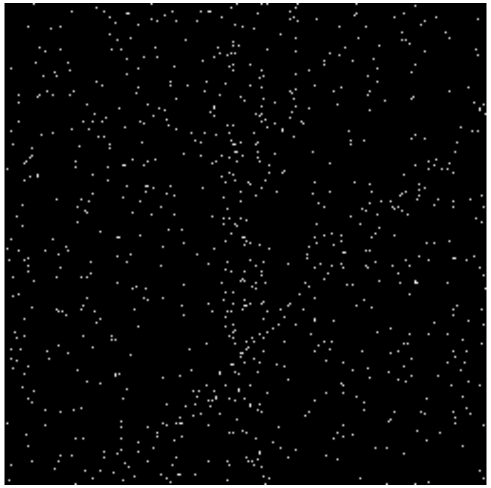

simulated image of counts

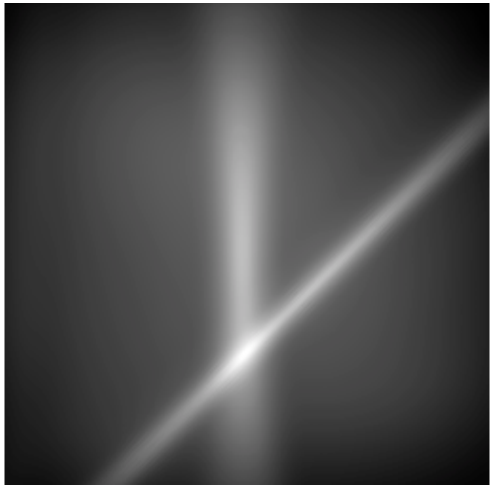

restored image
from the left image of counts

Max Intensity
background = 0.01
vertical bar = 0.03
inclined bar = 0.04

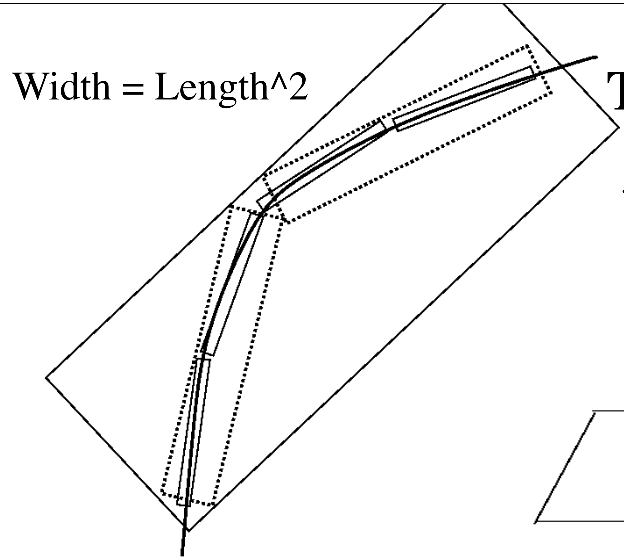

The Curvelet Transform (CUR01)

J.-L. Starck, E. Candes, D.L. Donoho ***The Curvelet Transform for Image Denoising***, *IEEE Transaction on Image Processing, 11, 6*

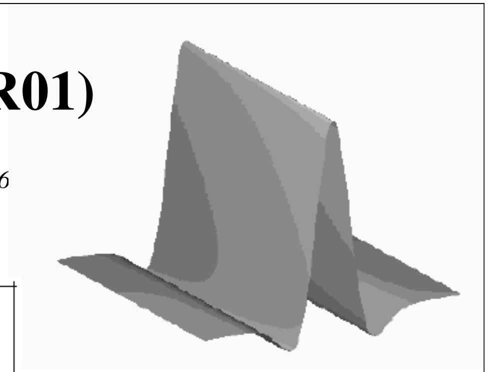

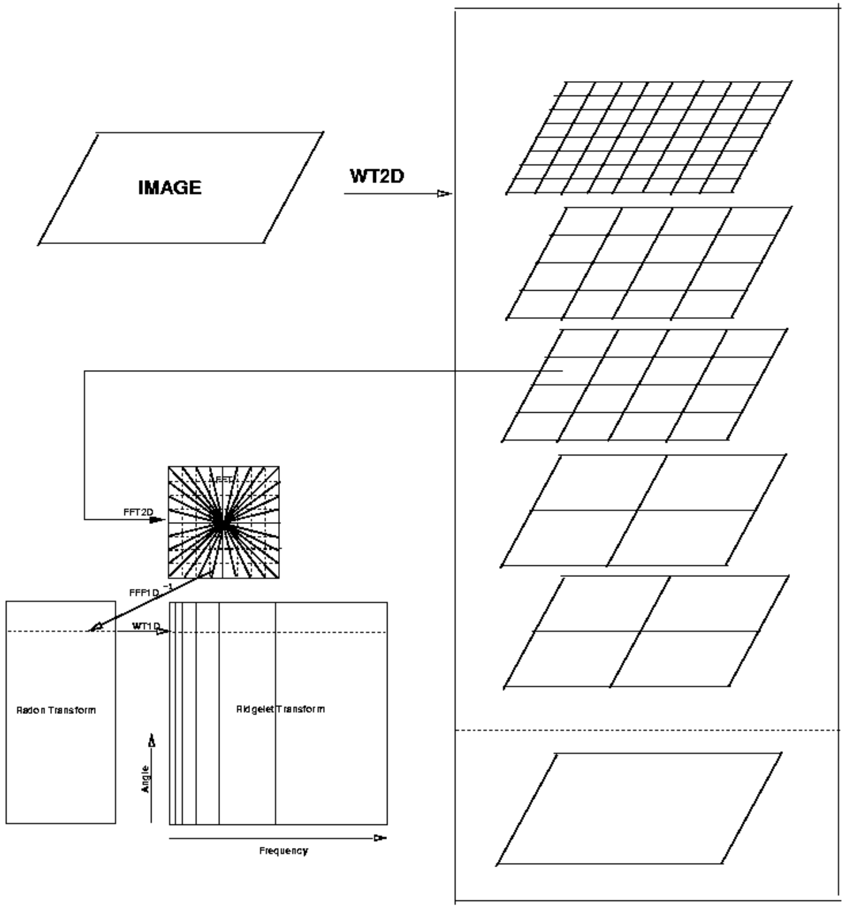

Redundancy $16J + 1$ for J wavelet scales.

Complexity $O(N^2(\log N)^2)$ for $N \times N$ images.

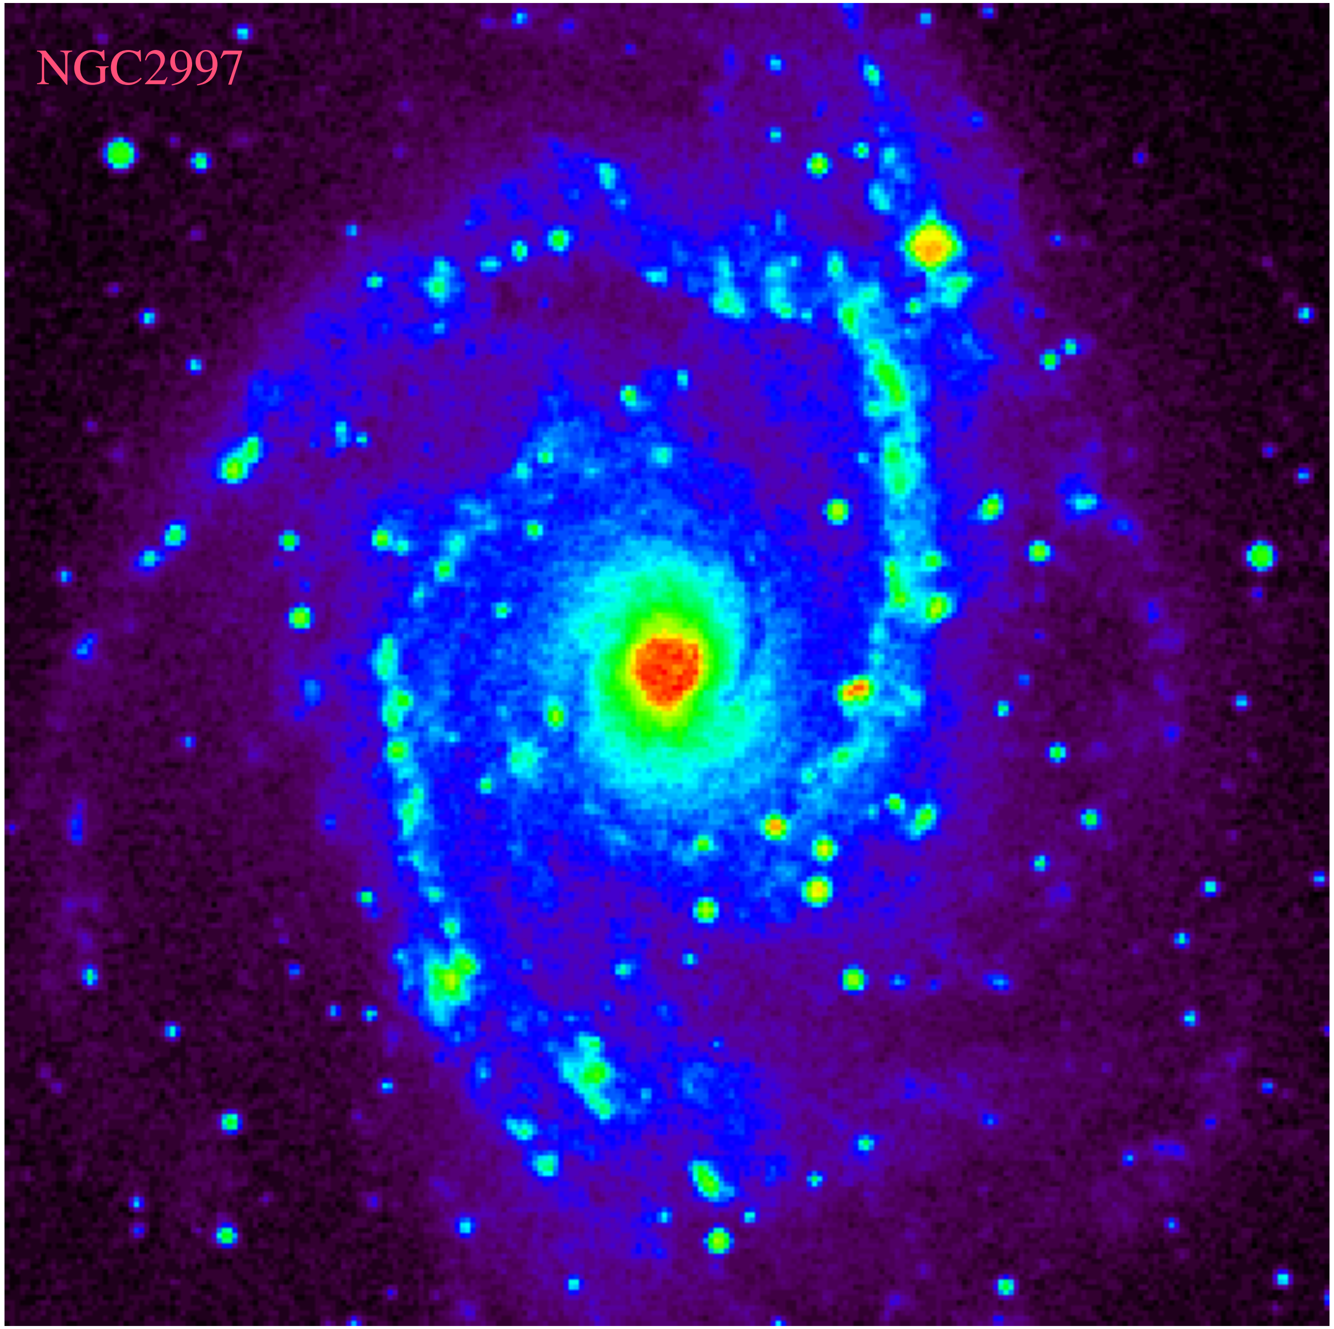
NGC2997

Undecimated Isotropic WT: $$I(k,l) = c_{J,k,l} + \sum_{j=1}^{J} w_{j,k,l}$$

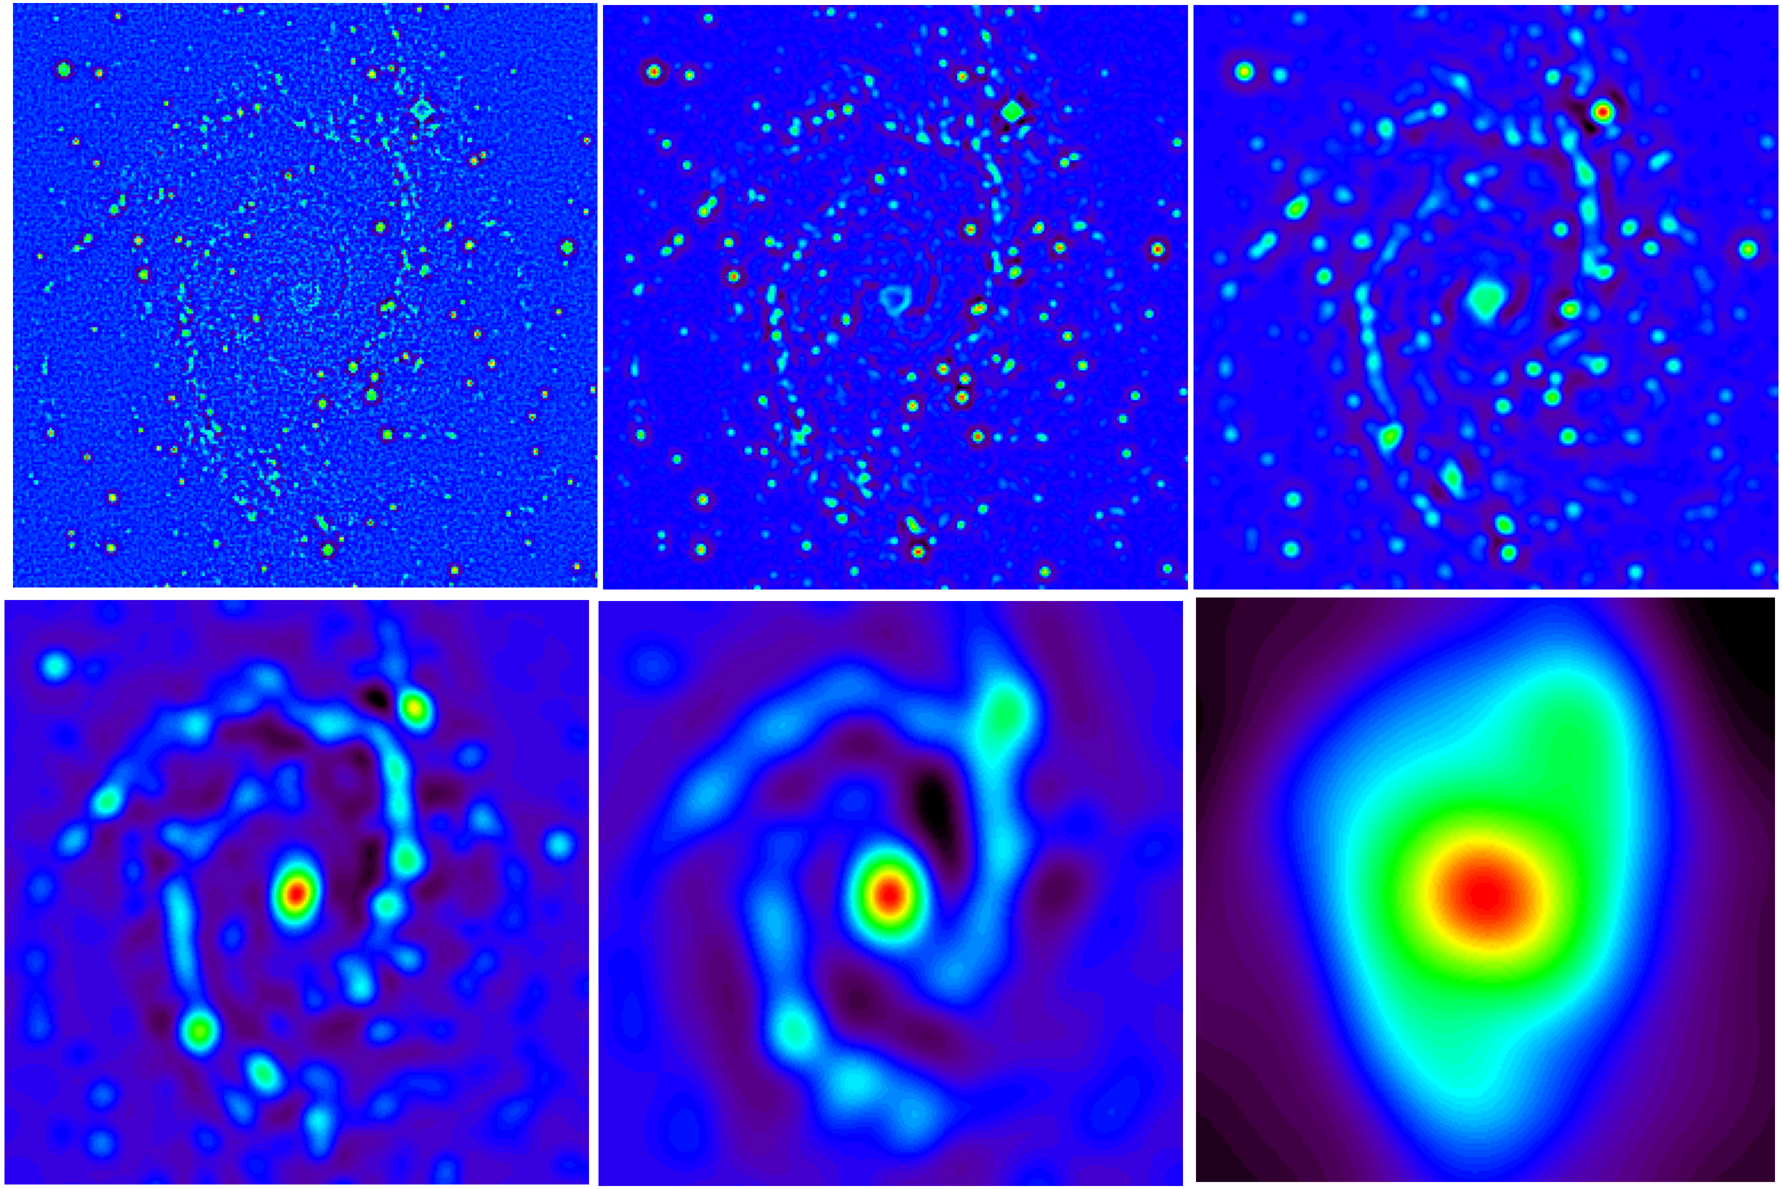

PARTITIONING

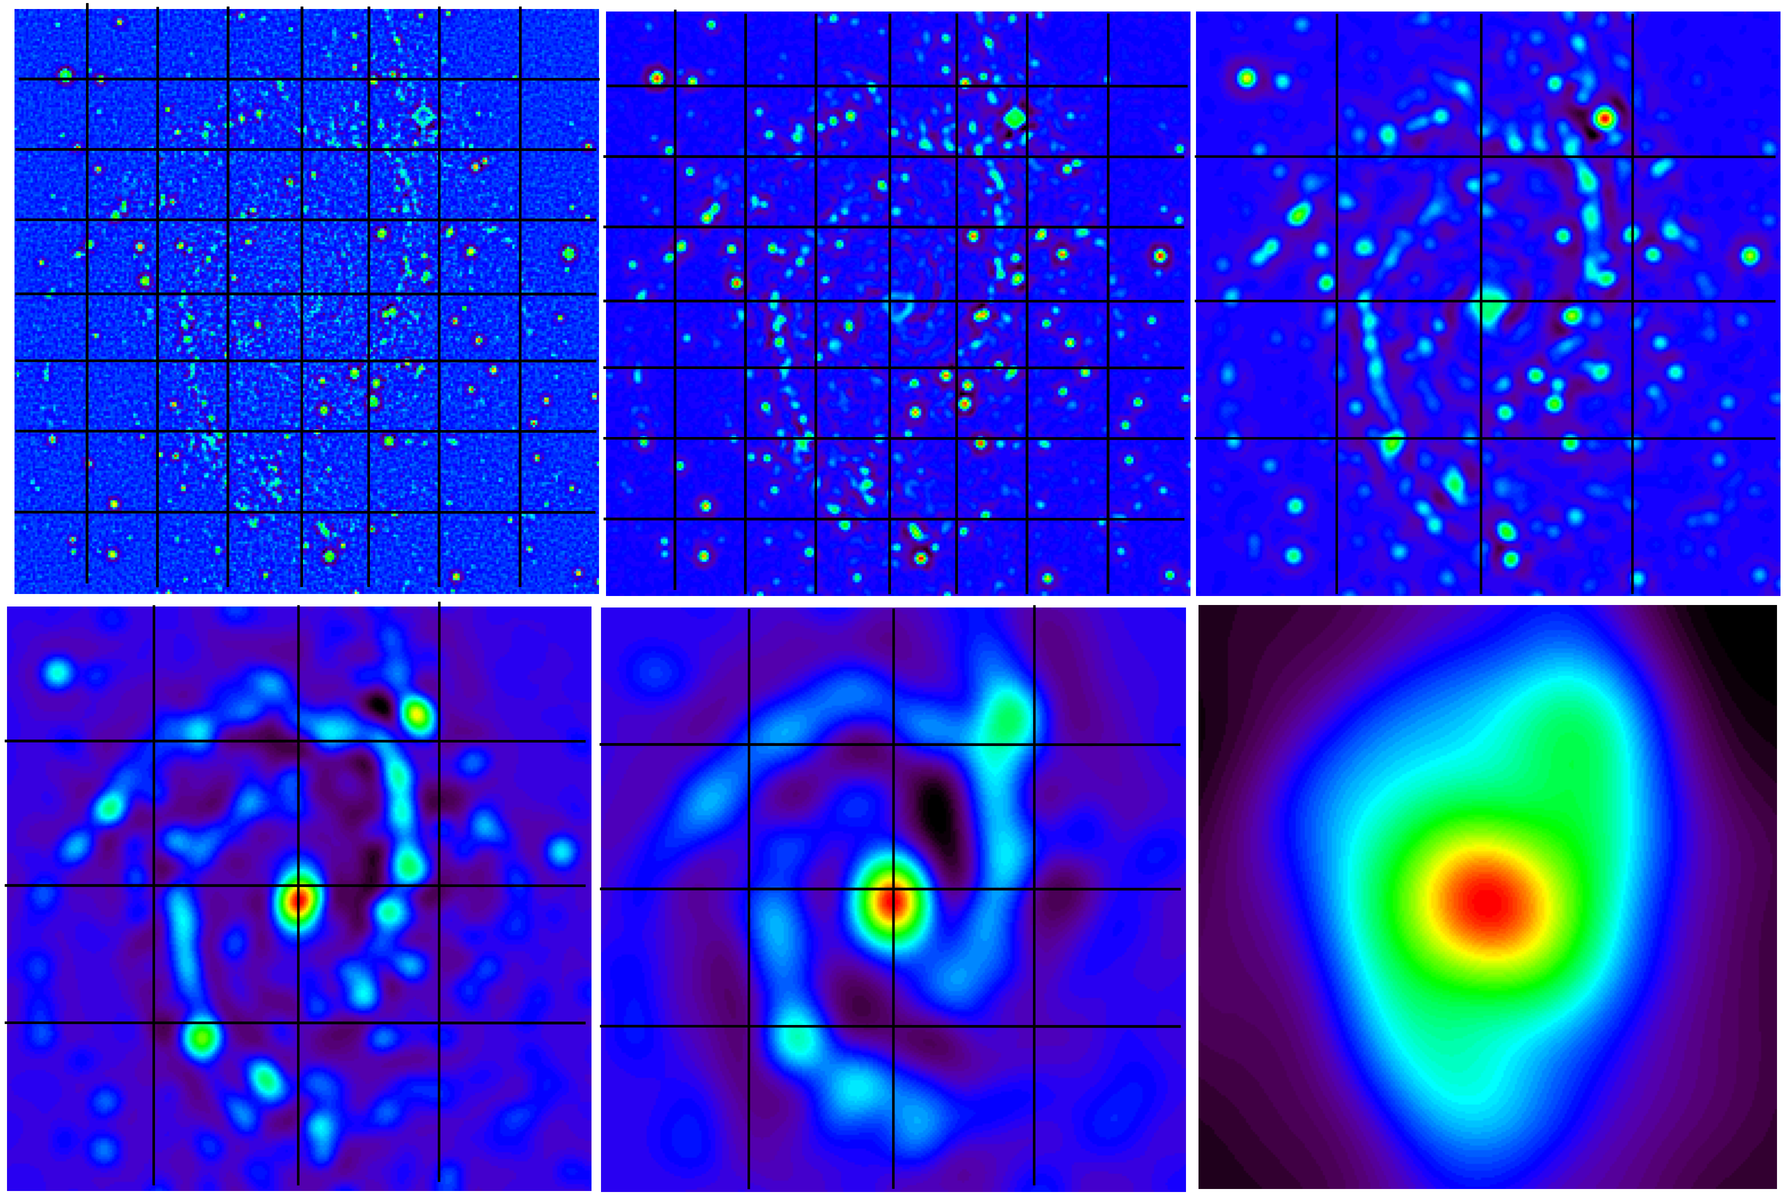

The Fast Curvelet Transform, Candes et al, 2005

CUR03 - Fast Curvelet Transform using the USFFT
CUR04 - Fast Curvelet Transform using the Wrapping and 2DFFT

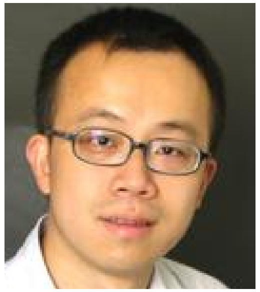

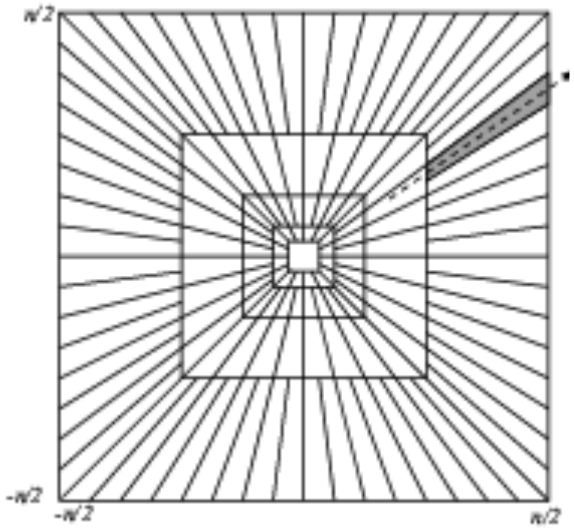

(a)

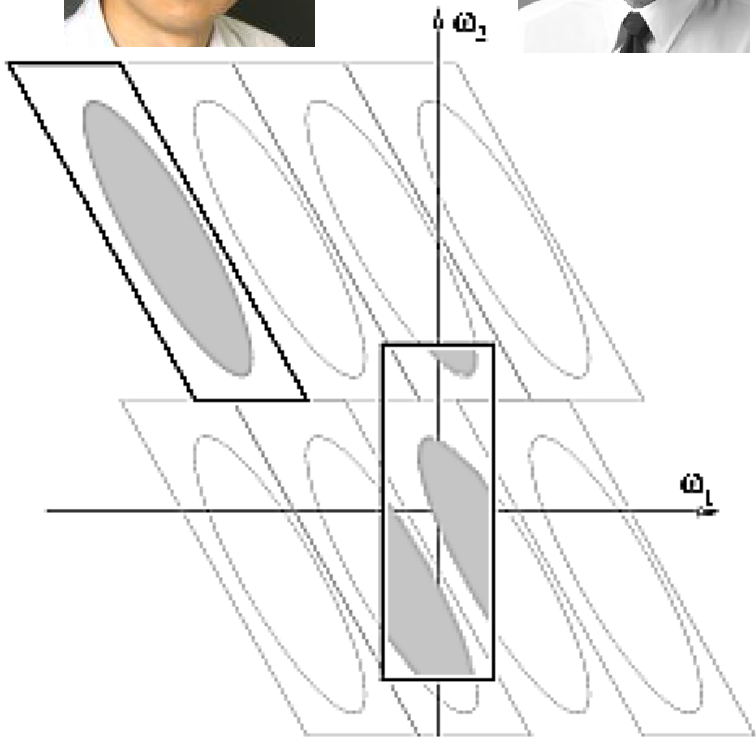

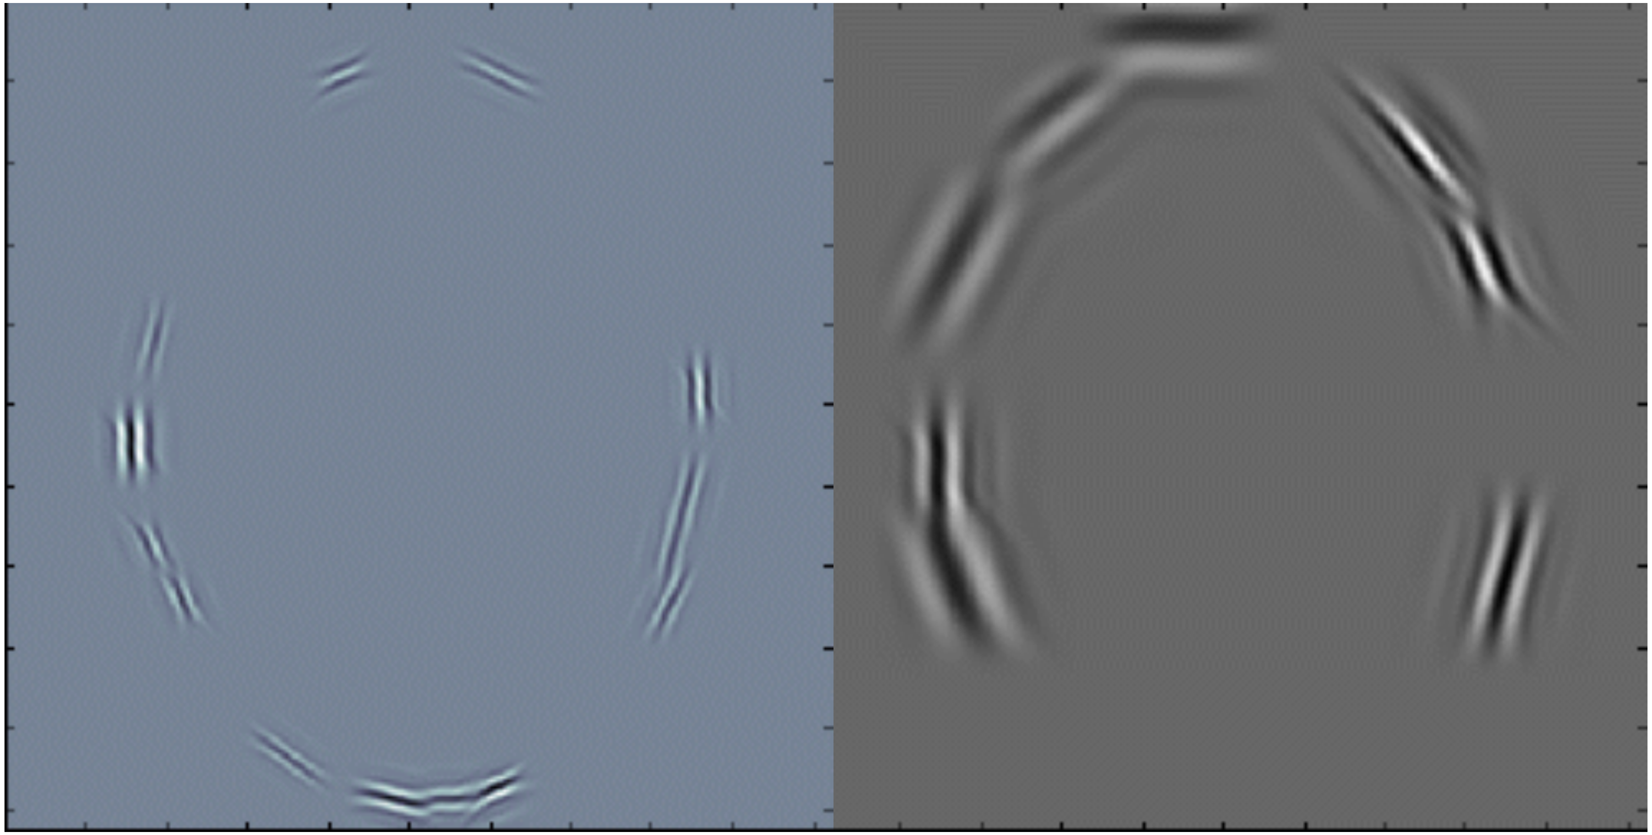

•J.L. Starck, E. Candes, and D.L. Donoho, "**The Curvelet Transform for Image Denoising**", IEEE Transactions on Image Processing , 11, 6, pp 670 -684, 2002.
•J.-L. Starck, M.K. Nguyen and F. Murtagh, "**Wavelets and Curvelets for Image Deconvolution: a Combined Approach**", Signal Processing, 83, 10, pp 2279-2283, 2003.
•J.-L. Starck, E. Candes, and D.L. Donoho, "**Astronomical Image Representation by the Curvelet Tansform**" , Astronomy and Astrophysics, 398, 785--800, 2003.
• J.-L. Starck, F. Murtagh, E. Candes, and D.L. Donoho, "**Gray and Color Image Contrast Enhancement by the Curvelet Transform**", IEEE Transaction on Image Processing, 12, 6, pp 706--717, 2003.

CONTRAST ENHANCEMENT USING THE CURVELET TRANSFORM

J.-L Starck, F. Murtagh, E. Candes and D.L. Donoho, ***"Gray and Color Image Contrast Enhancement by the Curvelet Transform"***,

IEEE Transaction on Image Processing, 12, 6, 2003.

$$\tilde{I} = C_R\left(y_c\left(C_T I\right)\right)$$

$$\begin{cases} y_c(x,\sigma) = 1 & \text{if } x < c\sigma \\ y_c(x,\sigma) = \dfrac{x - c\sigma}{c\sigma}\left(\dfrac{m}{c\sigma}\right)^p + \dfrac{2c\sigma - x}{c\sigma} & \text{if } x < 2c\sigma \\ y_c(x,\sigma) = \left(\dfrac{m}{x}\right)^p & \text{if } 2c\sigma \le x < m \\ y_c(x,\sigma) = \left(\dfrac{m}{x}\right)^s & \text{if } x > m \end{cases}$$

Modified curvelet coefficient

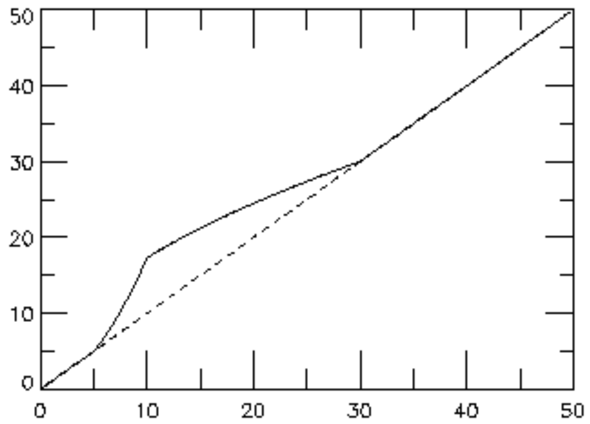

Curvelet coefficient

Contrast Enhancement

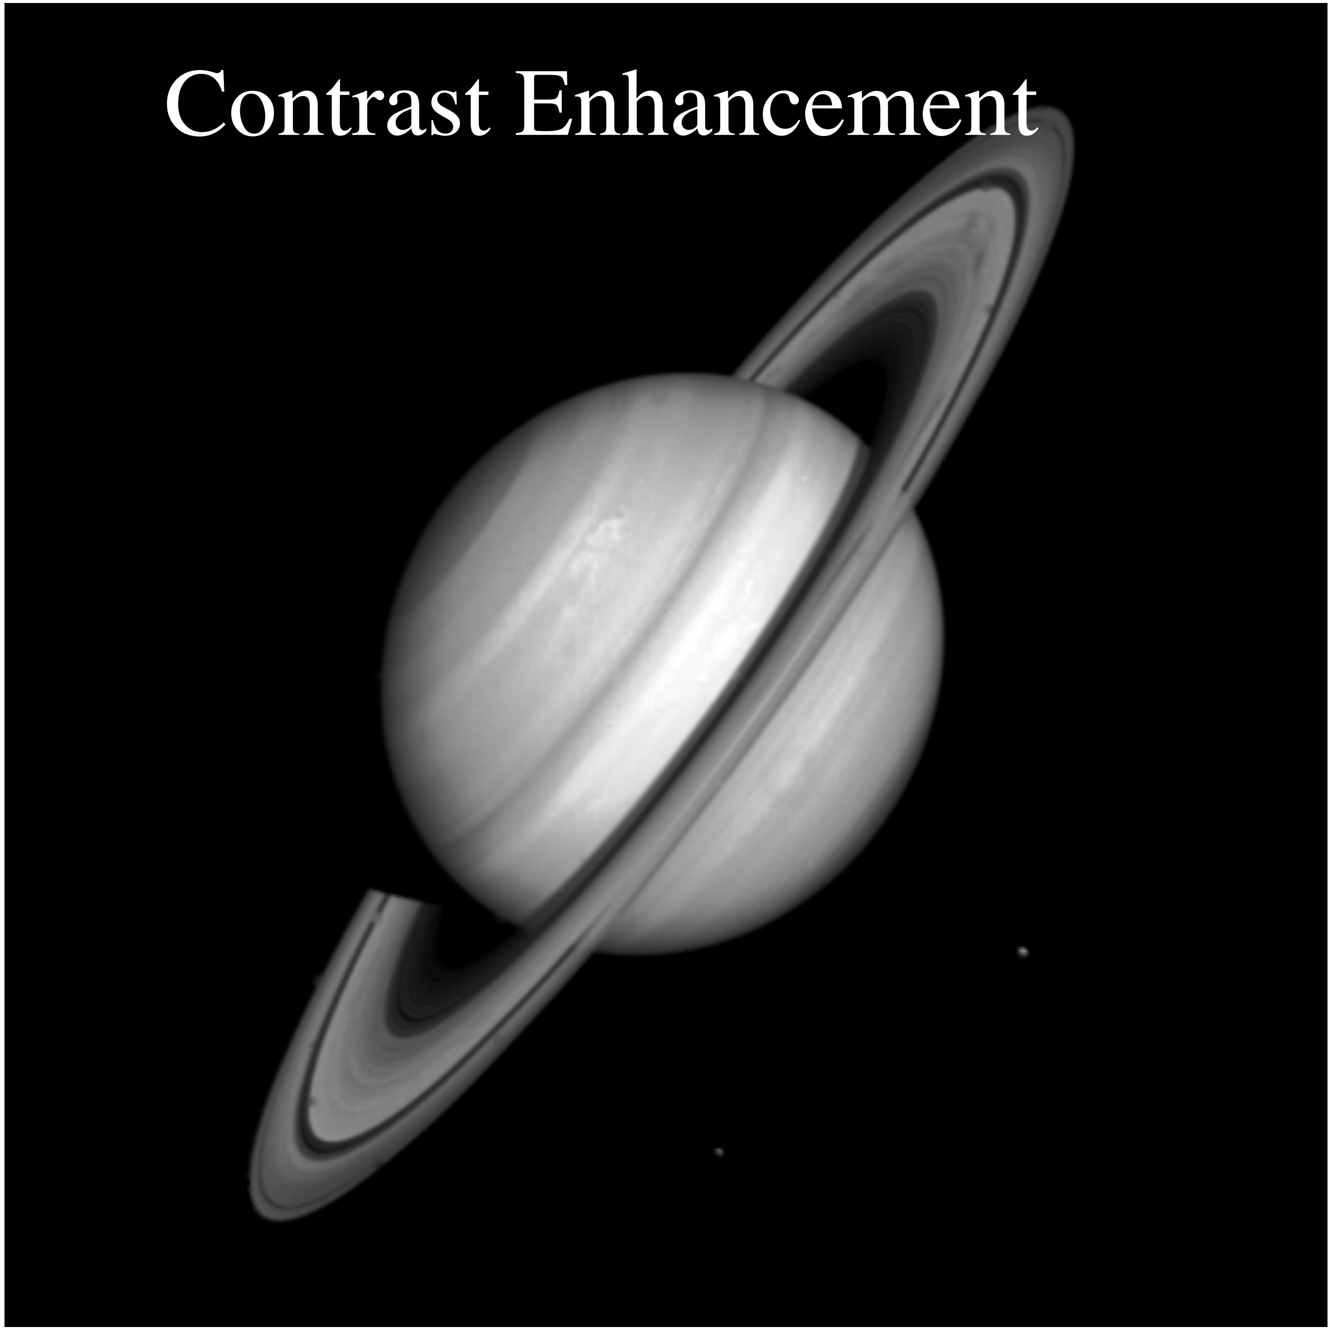

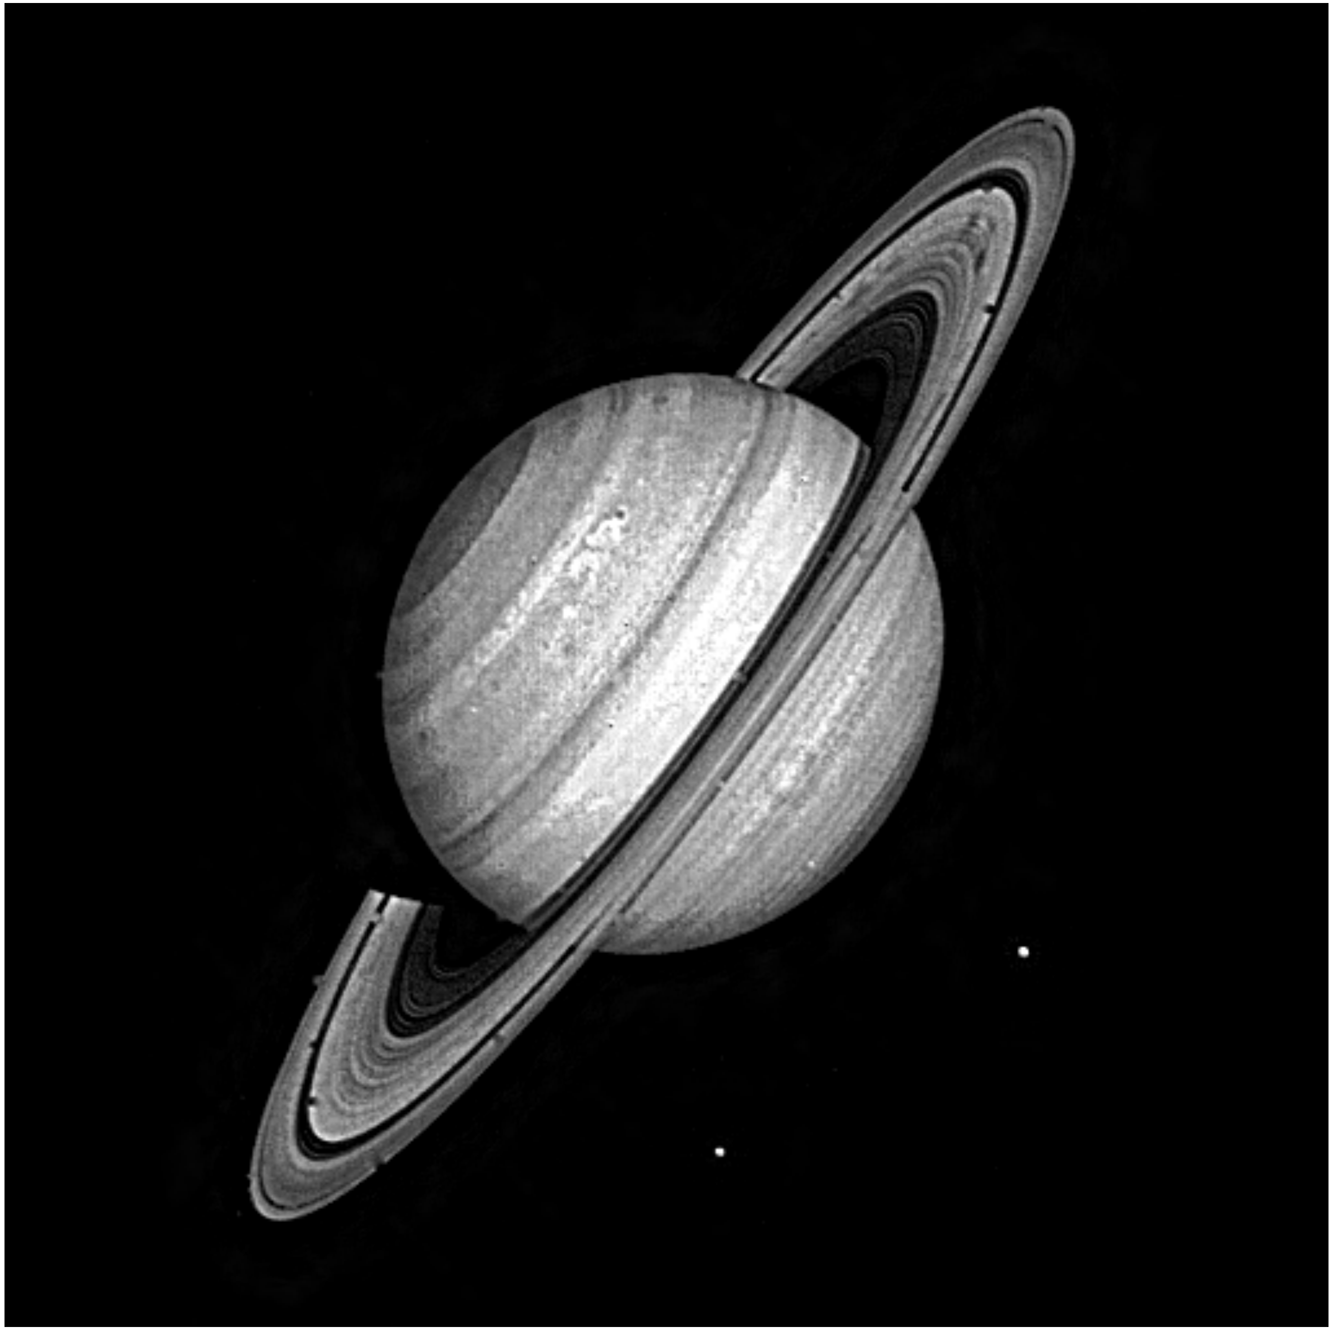

DECONVOLUTION

- E. Pantin, J.-L. Starck, and F. Murtagh, "Deconvolution and Blind Deconvolution in Astronomy", in Blind image deconvolution: theory and applications, pp 277--317, 2007.
- J.-L. Starck, F. Murtagh, and M. Bertero, "The Starlet Transform in Astronomical Data Processing: Application to Source Detection and Image Deconvolution", Springer, Handbook of Mathematical Methods in Imaging, in press, 2011.

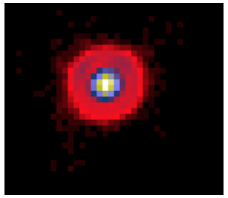

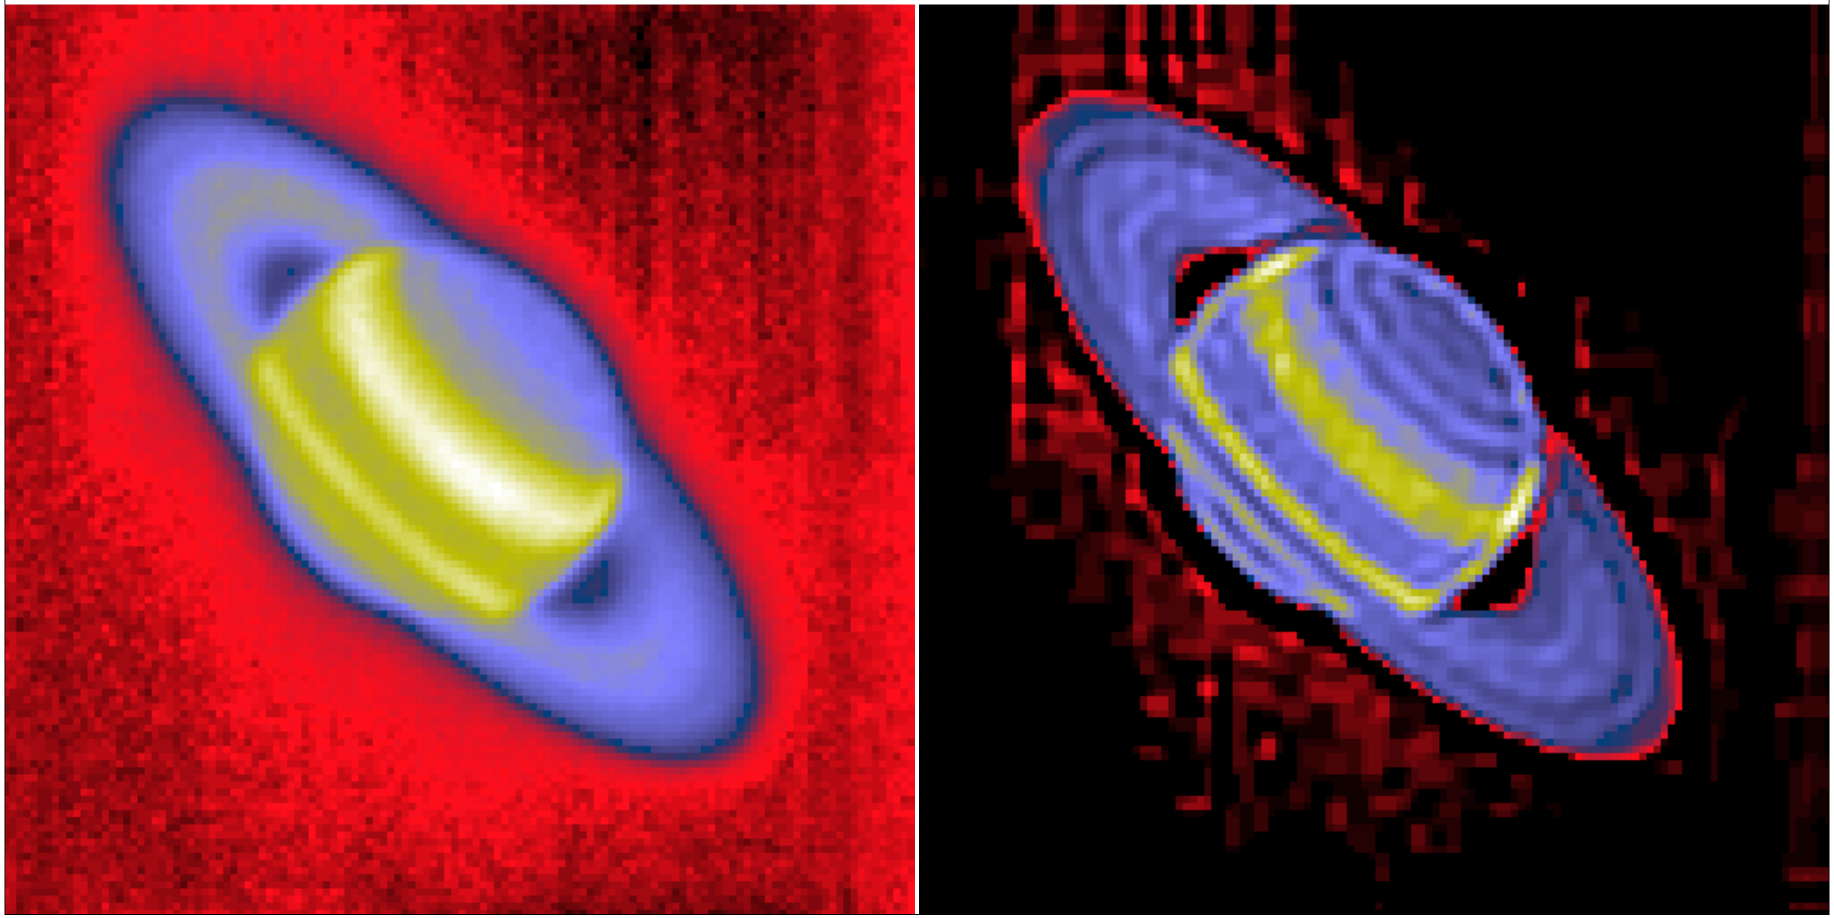

3D Multiscale Geometric Transforms

A. Woiselle, J.L. Starck and M.J. Fadili, "3D curvelet transforms and astronomical data restoration", **Applied and Computational Harmonic Analysis**, Vol. 28, No. 2, pp. 171-188, 2010.

A. Woiselle, J.L. Starck, M.J. Fadili, "3D Data Denoising and Inpainting with the Fast Curvelet transform", **JMIV**, 39, 2, pp 121-139, 2011.

Software: http://jstarck.free.fr/cur3d.html

Curvelet 01 2D ==> 3D
FastCurvelet 3D

3D Wavelets

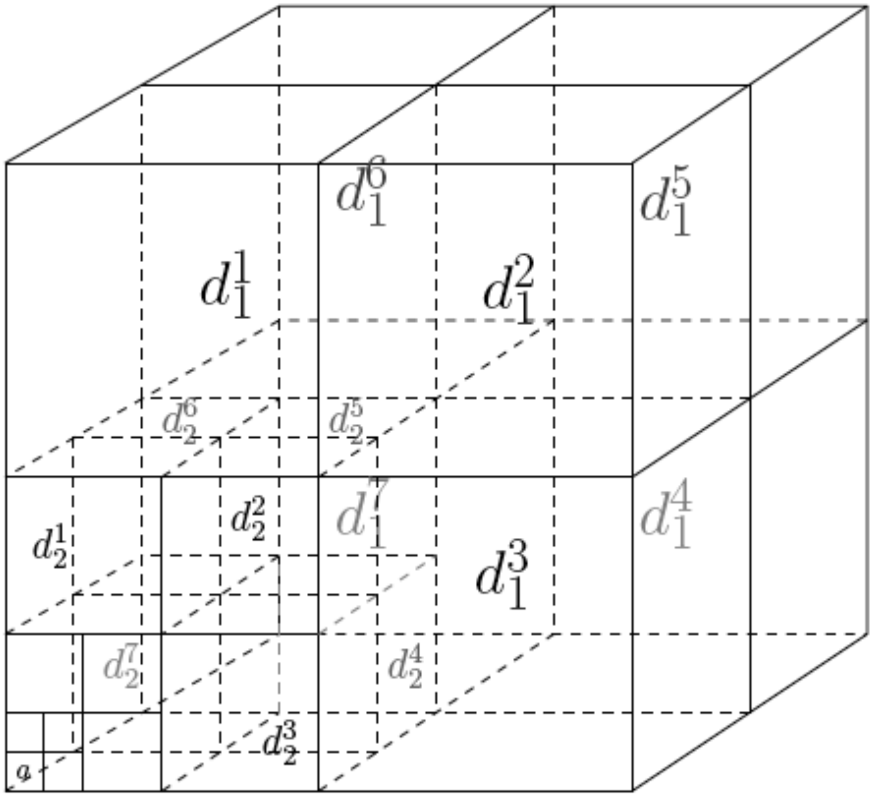

Orthogonal Wavelets

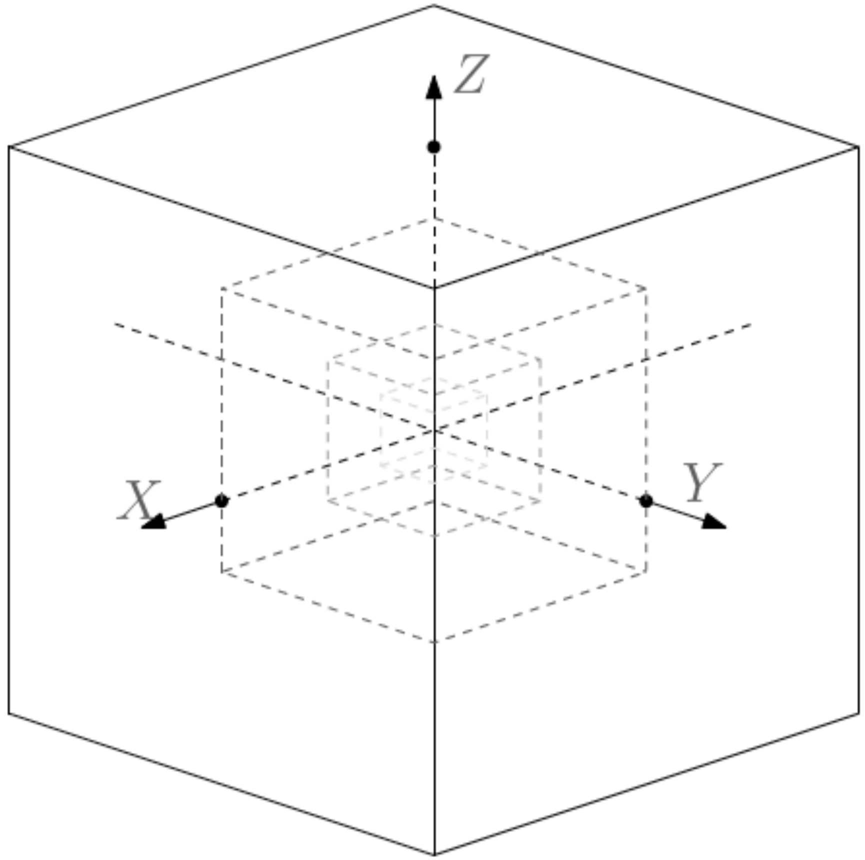

Meyer Wavelets

3D extension of Curvelet

- As in 2D, the 3D first generation curvelet transform we develop is based on the 3D ridgelet transform applied to localized blocks of the output of a 3D wavelet transform.
- The essential ingredient is the projection slice theorem: the m-D FT of the projection of a d-D function onto an m-D linear submanifold is equal to an m-D central slice of the d-D FT parallel to the submanifold.
- Two 3D extensions of the ridgelet transform:
  - Projections along lines (3D partial Radon transform, $d=3$, $m=2$): BeamCurvelets.
  - Projecting along planes (3D Radon transform, $d=3$, $m=1$): RidCurvelets.

3D beamlet transform

- $\gamma \in L_2(\mathbb{R}^2)$ with zero-mean and has sufficient decay (2D wavelet).
- For each scale $s > 0$, position $(b_1, b_2) \in \mathbb{R}^2$ and orientation $(\theta_1, \theta_2) \in [0, 2\pi) \times [0, \pi)$, define the 2D beamlet $\psi^B_{s,b_1,b_2,\theta_1,\theta_2} : \mathbb{R}^3 \to \mathbb{R}$ by

$$\psi^B_{s,b_1,b_2,\theta_1,\theta_2}(\mathbf{x}) = s^{-1/2} \cdot \gamma((-x \sin\theta_1 + y \cos\theta_1 - b_1)/s, \\ (x \cos\theta_1 \sin\theta_2 + y \sin\theta_1 \sin\theta_2 - z \cos\theta_2 - b_2)/s) \; .$$

- The 3D beamlet transform of $f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^3$ is the set of coefficients

$$\left\langle f, \psi^B_{s,b_1,b_2,\theta_1,\theta_2} \right\rangle .$$

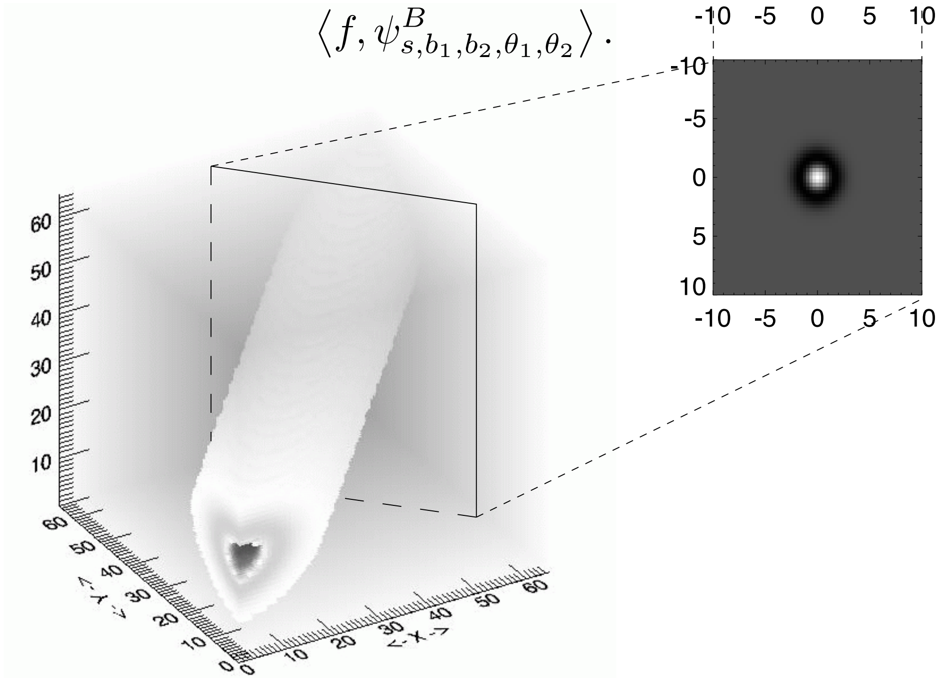

Discrete 3D Beamlet transform

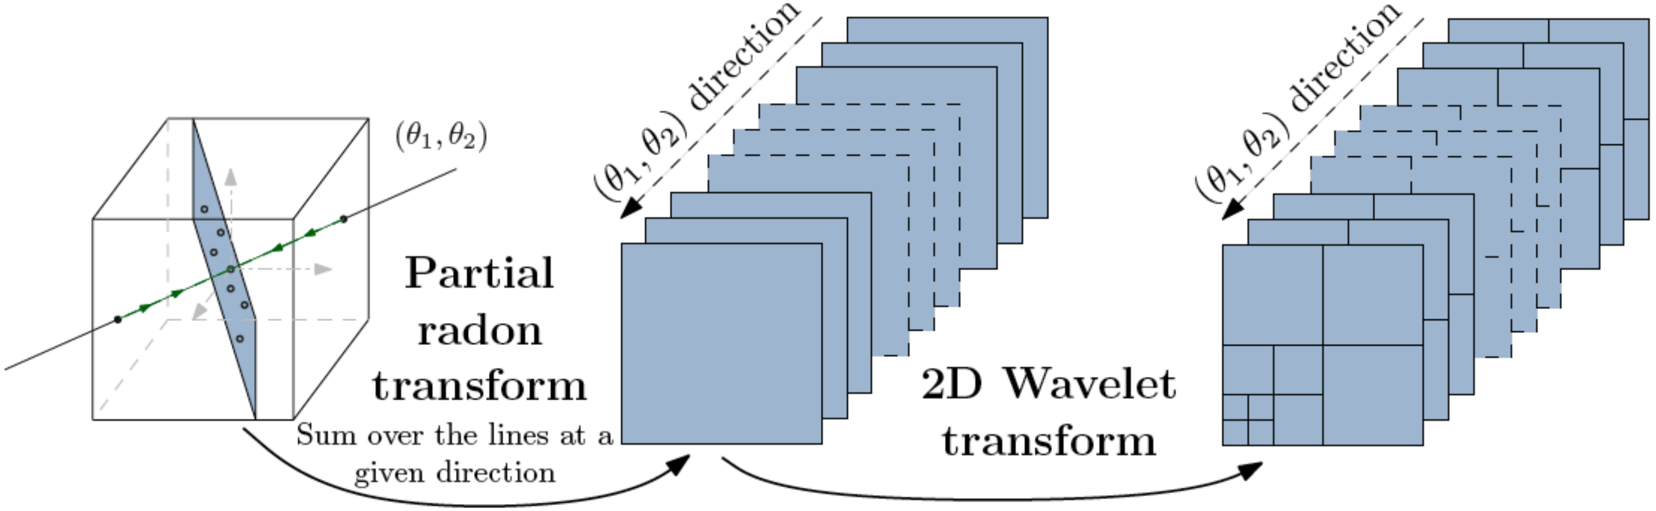

A. Woiselle, J.L. Starck and M.J. Fadili, "3D curvelet transforms and astronomical data restoration", **Applied and Computational Harmonic Analysis**, Vol. 28, No. 2, pp. 171-188, 2010.

Discrete 3D BeamCurvelet transform

Algorithm: Fourier-based implementation.

Data: A data cube and a block size B.

Result: BeamCurvelet transform.

begin

  Apply a 3D isotropic wavelet transform.

  for $j = 1$ *to* J **do**

    Smooth partition of the subband into block cubes of size B.

    for *each block* **do**

      Apply a 3D FFT.

      Extract planes passing through the origin at every angle (θ_1, θ_2).

      for *each plane* (θ_1, θ_2) **do**

        Apply an inverse 2D FFT.

        Apply a 2D wavelet transform to get the BeamCurvelet coefficients.

    if j *is odd* **then** according to the parabolic scaling: B $\leftarrow 2$B.

end

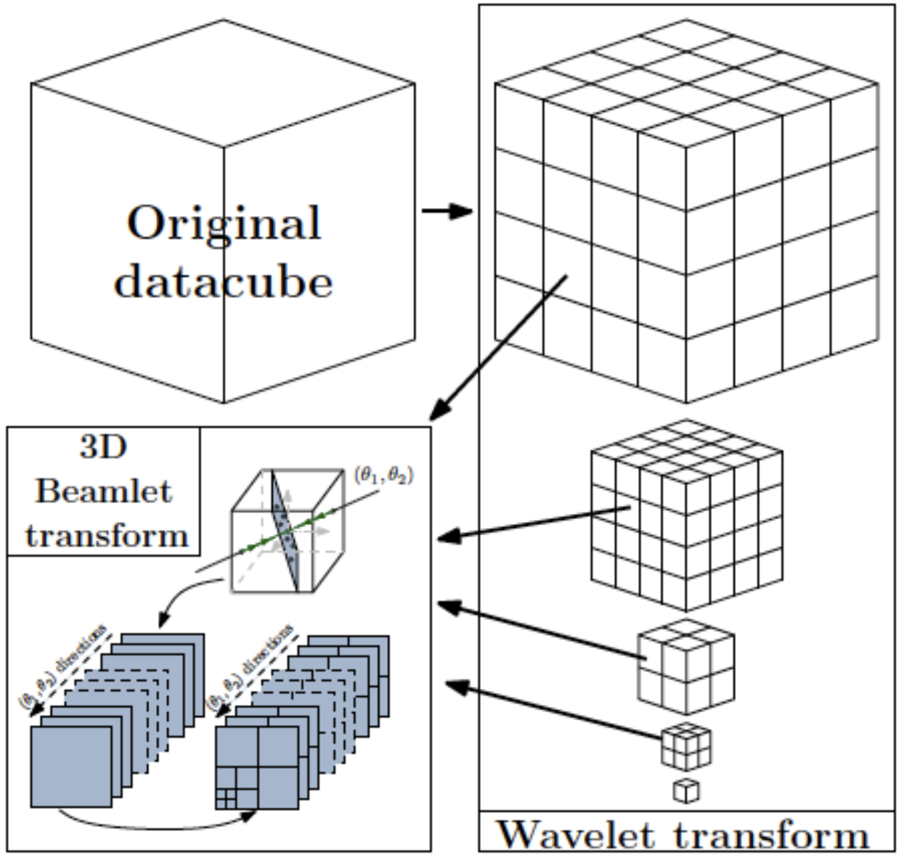

Redundancy $\approx 3\rho^3$B, $\rho \in [1, 2]$.

Complexity $O(N^3 (\log N)^2)$ for $N \times N \times N$ volume.

3D BeamCurvelets

- It is constant along segments of direction (θ_1, θ_2), and a 2D wavelet function transverse to this direction.
- Adapted to filamentary structures in 3D.

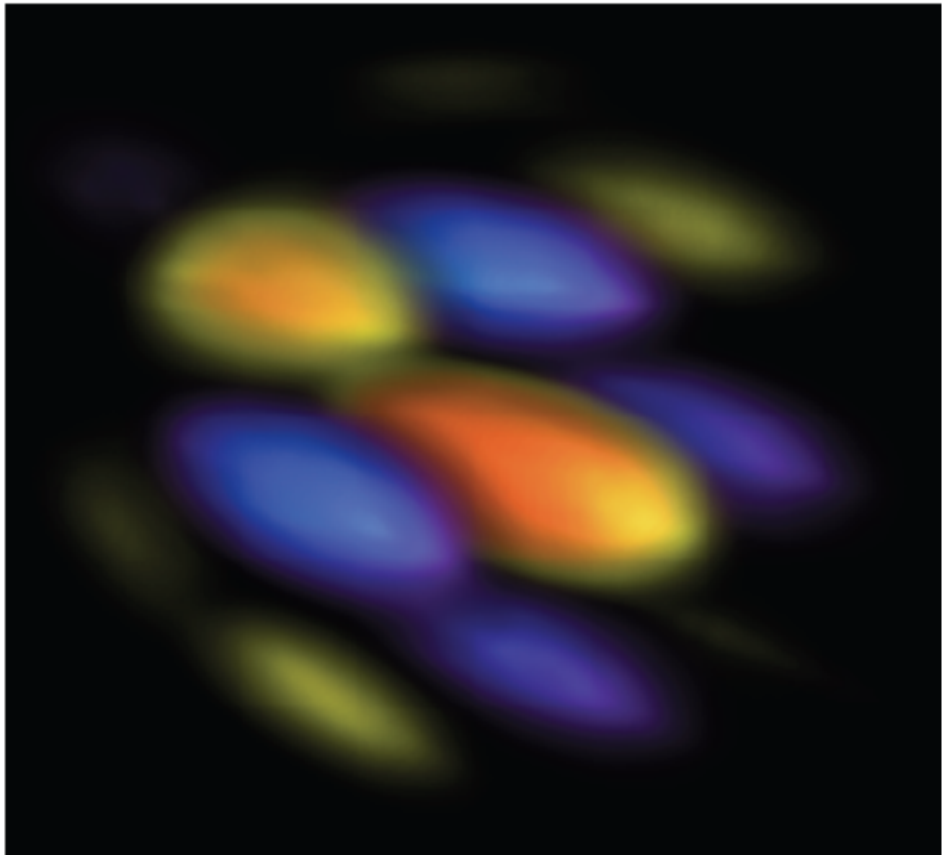

3D BeamCurvelets

Simulation

Noise Data

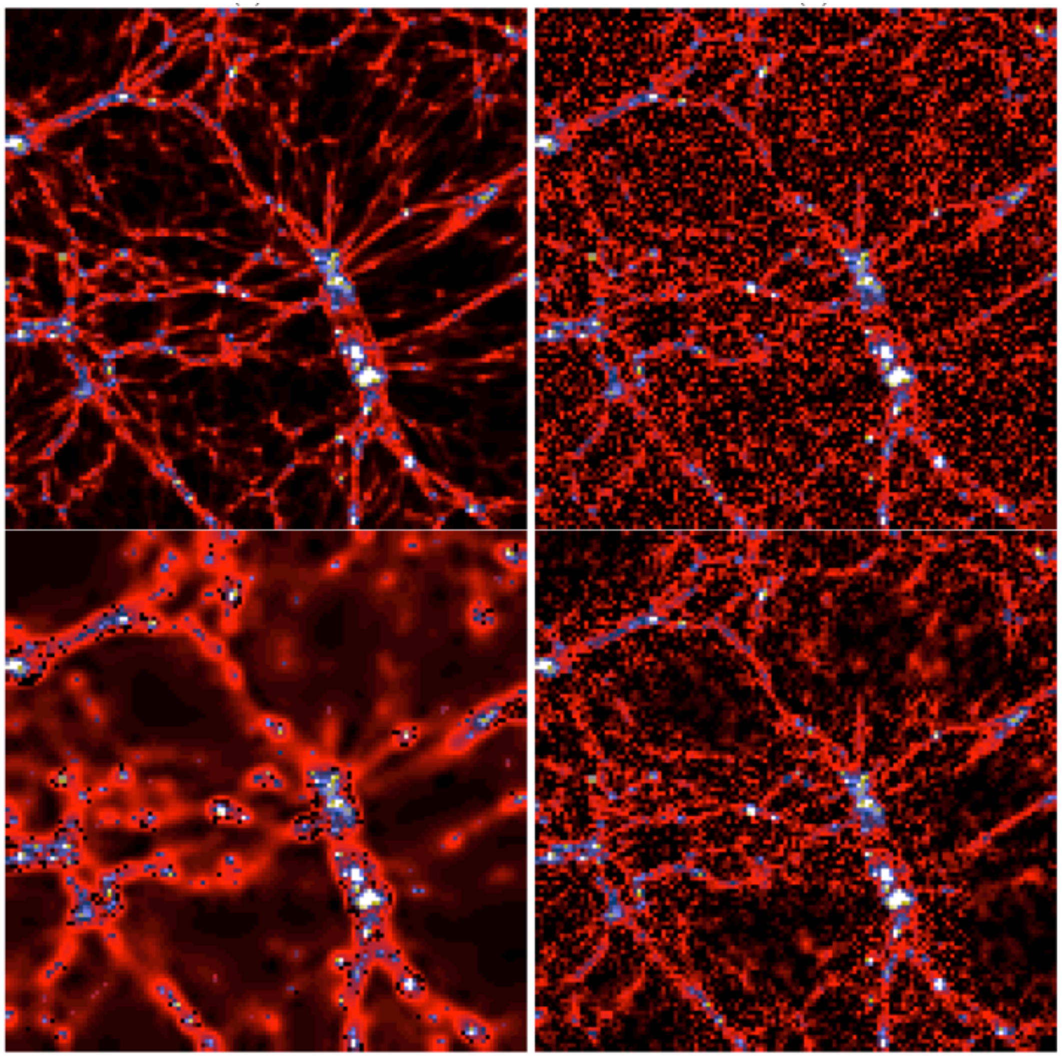

Wavelet Thresholding

BeamCurvelet Thresholding

3D ridgelet transform

- $\psi \in L_2(\mathbb{R})$ with zero-mean and has sufficient decay.
- For each scale $s > 0$, position $b \in \mathbb{R}$ and orientation $(\theta_1, \theta_2) \in [0, 2\pi) \times [0, \pi)$, define the 2D ridgelet $\psi^R_{s,b,\theta_1,\theta_2} : \mathbb{R}^3 \to \mathbb{R}$ by

$$\psi^R_{s,b,\theta_1,\theta_2}(\mathbf{x}) = s^{-1/2} \cdot \psi((x \cos\theta_1 \cos\theta_2 + y \sin\theta_1 \cos\theta_2 + z \sin\theta_2 - b)/s) \; .$$

- The 3D ridgelet transform of $f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^3$ is the set of coefficients

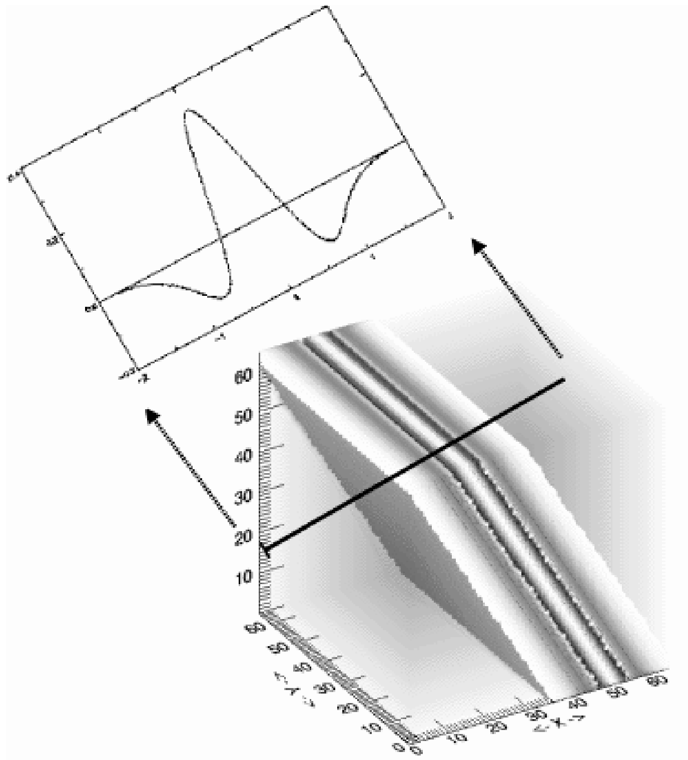

3D Discrete ridgelet transform

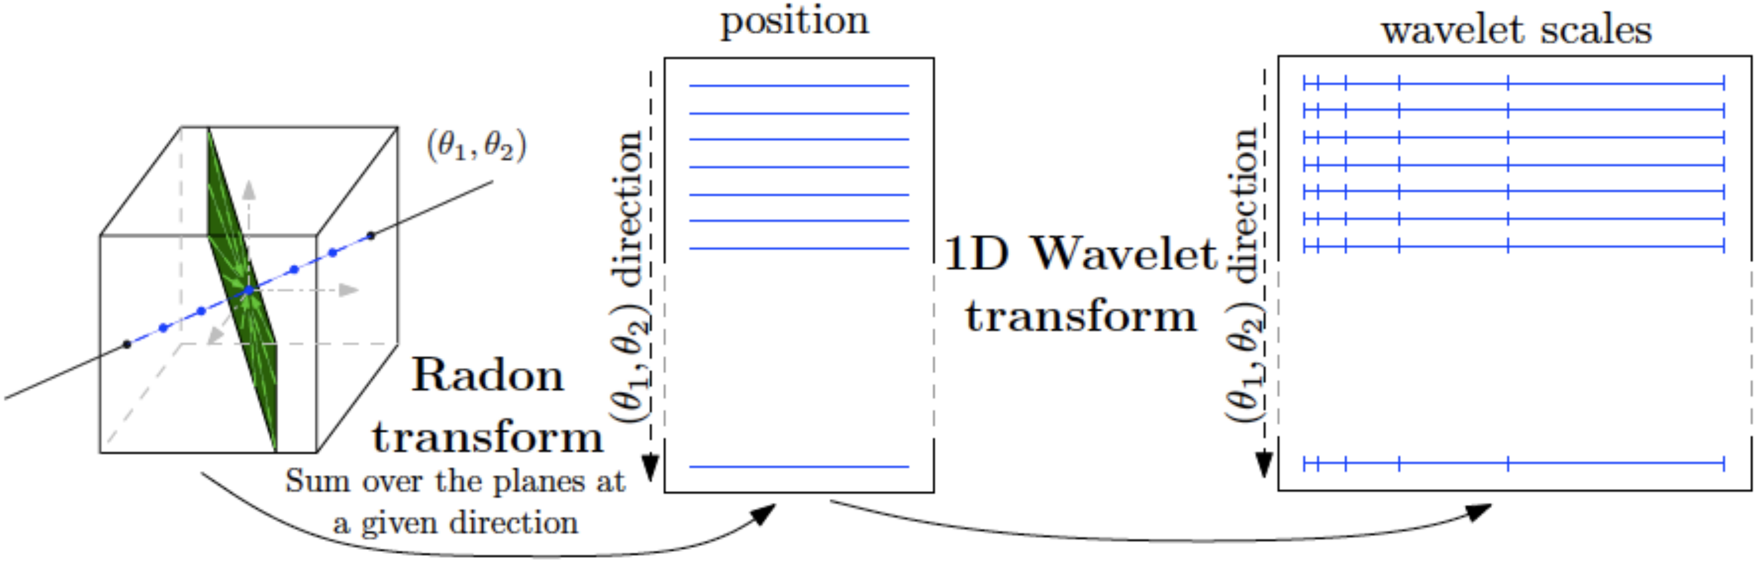

A. Woiselle, J.L. Starck and M.J. Fadili, "3D curvelet transforms and astronomical data restoration", **Applied and Computational Harmonic Analysis**, Vol. 28, No. 2, pp. 171-188, 2010.

Discrete 3D RidCurvelet transform

Algorithm: Fourier-based implementation.

Data: A data cube and a block size B.

Result: RidCurvelet transform.

begin

  Apply a 3D isotropic wavelet transform.

  for $j = 1$ *to* J **do**

    Smooth partition of the subband into block cubes of size B.

    for *each block* **do**

      Apply a 3D FFT.

      Extract lines passing through the origin at every angle (θ_1, θ_2).

      for *each line* (θ_1, θ_2) **do**

        Apply an inverse 1D FFT.

        Apply a 1D wavelet transform to get the RidCurvelet coefficients.

    if j *is odd* **then** according to the parabolic scaling: B $\leftarrow 2$B.

end

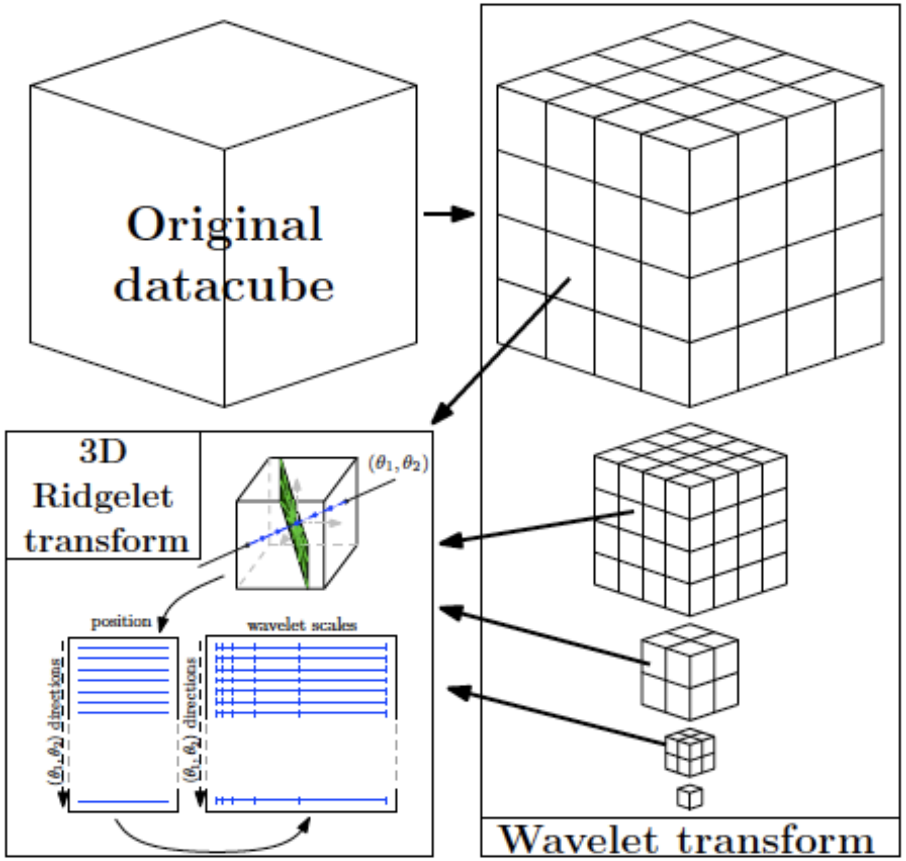

Redundancy $\approx 6\rho^3, \ \rho \in [1, 2]$.

Complexity $O(N^3(\log N)^2)$ for $N \times N \times N$ volume.

3D RidCurvelets

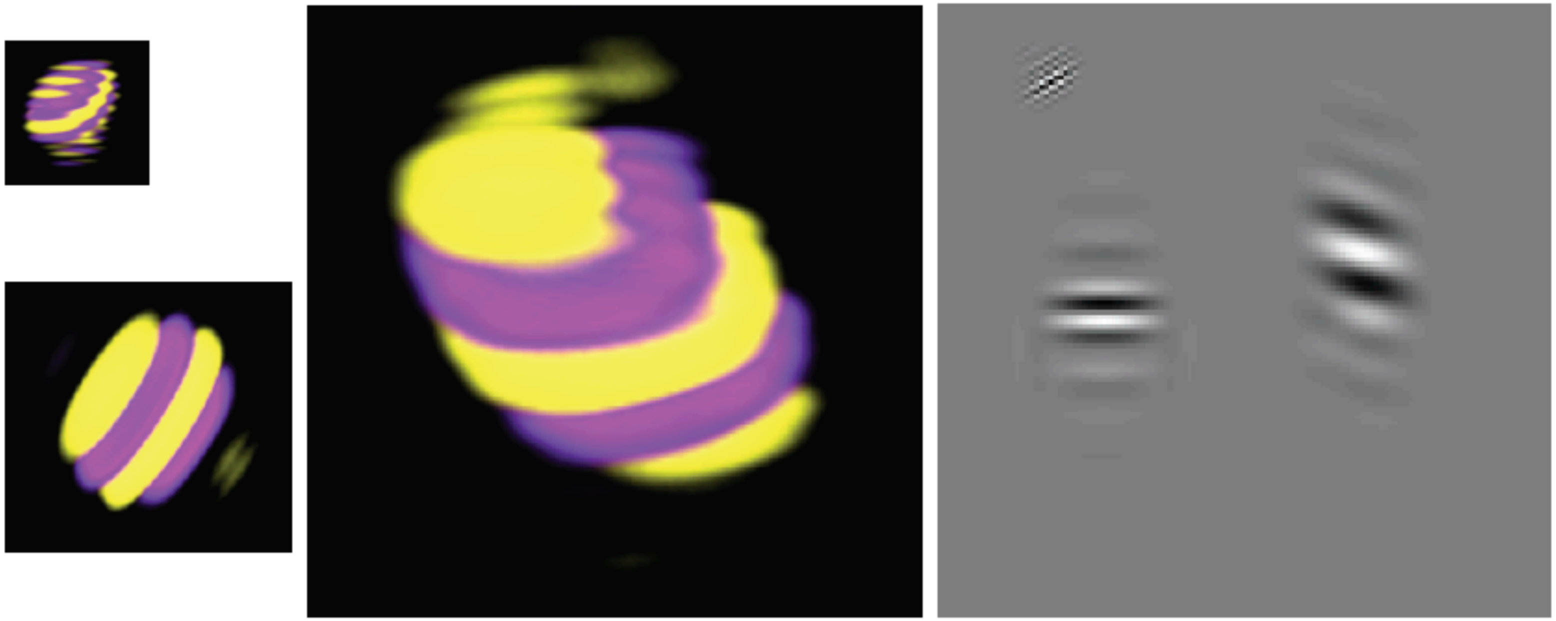

Denoising

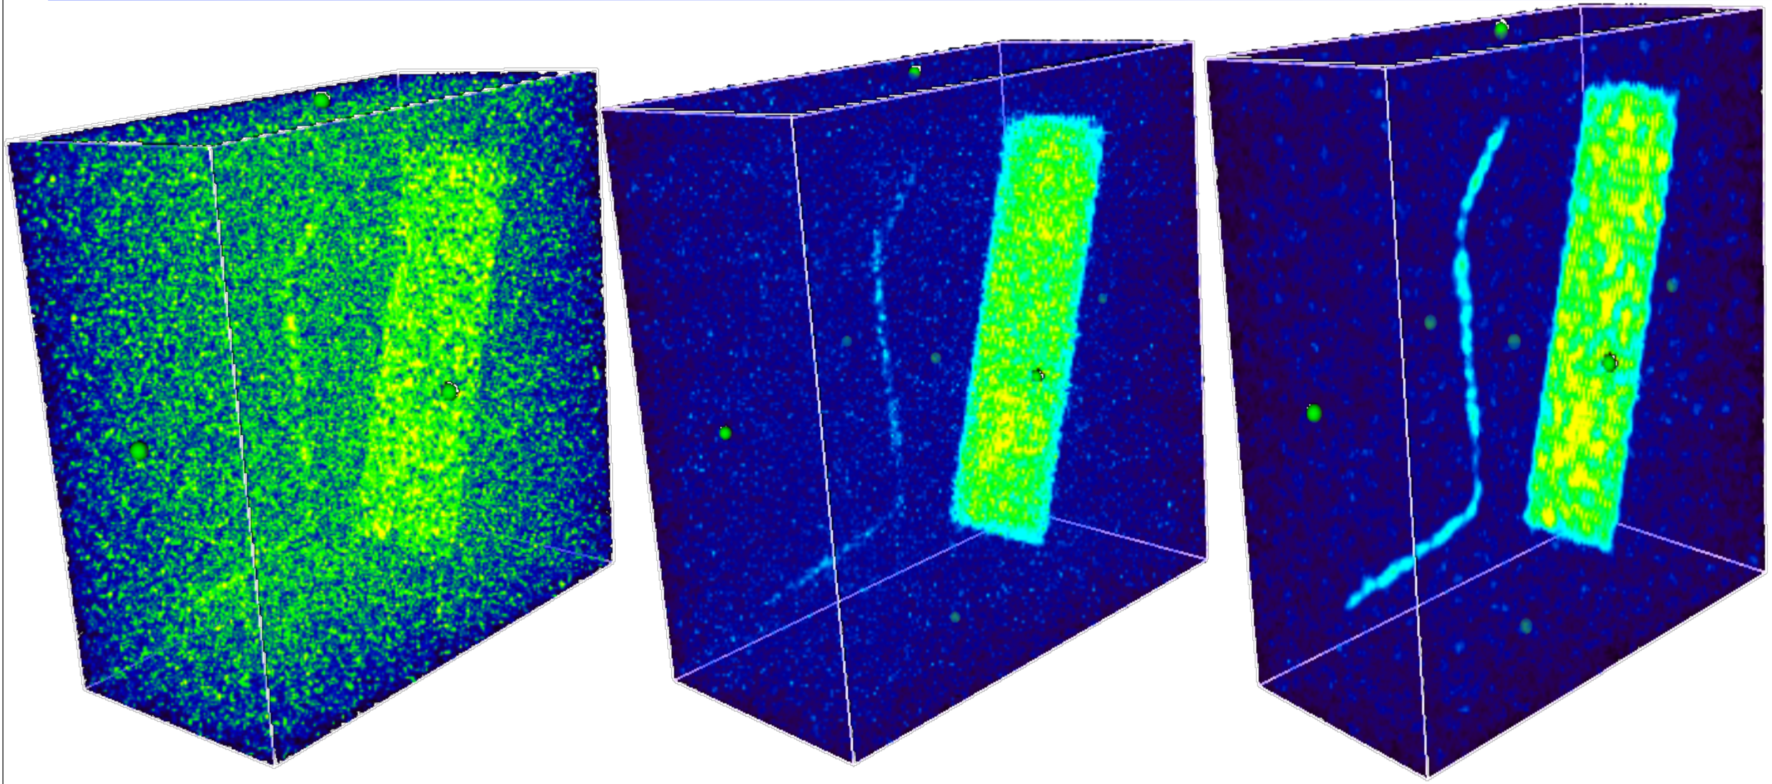

Noisy data **Denoised with RidCurvelets** **Denoised with BeamCurvelets**

3D Plane detection level

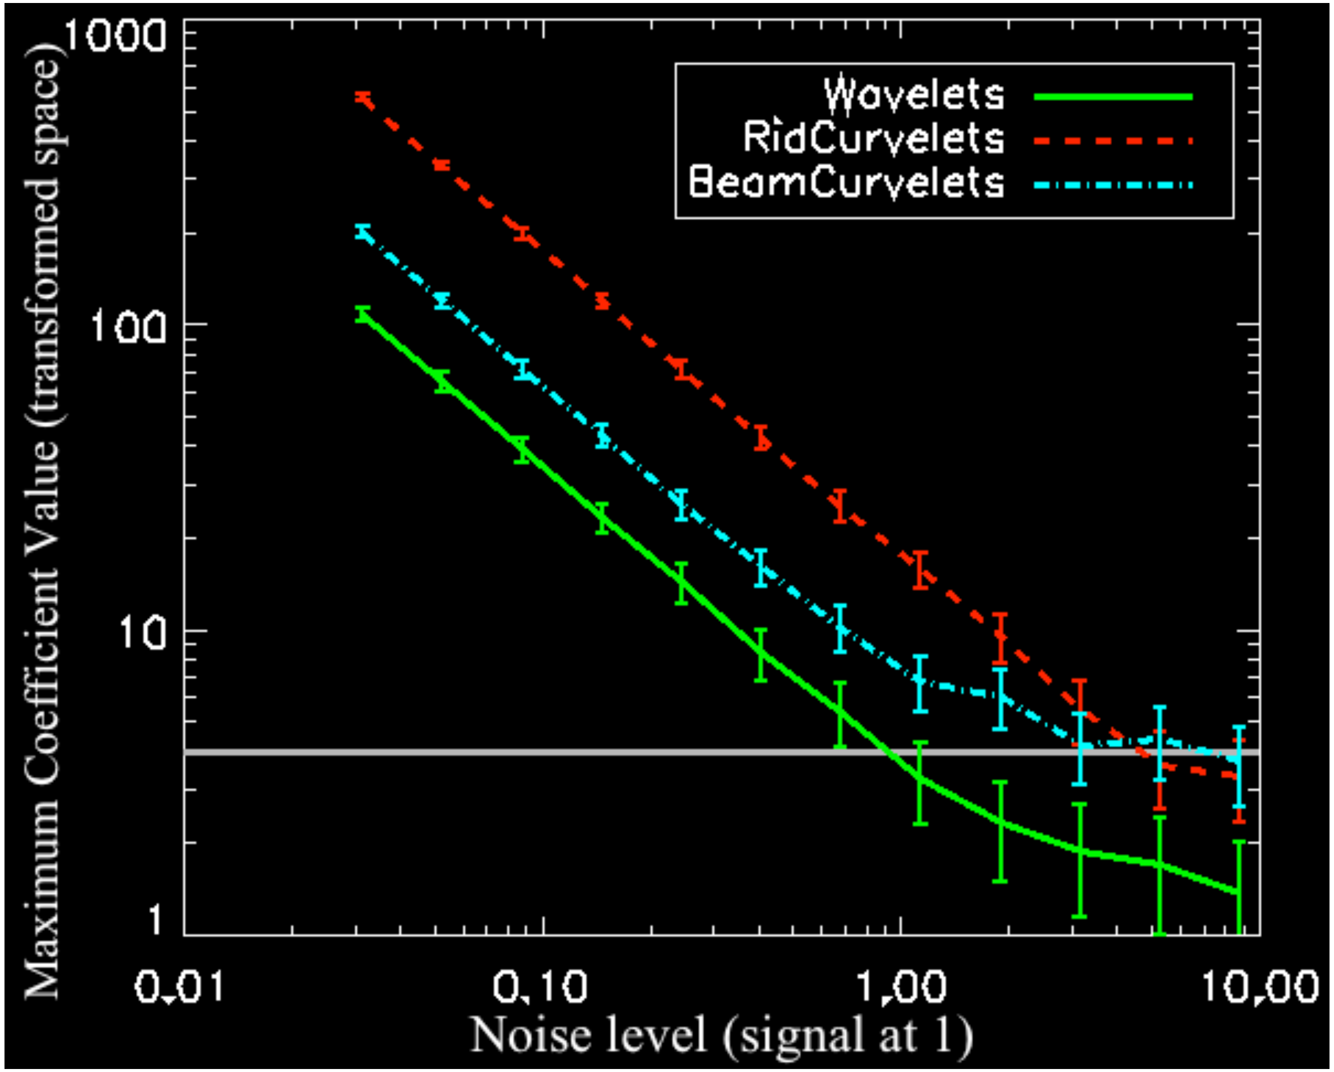

3D Line detection level

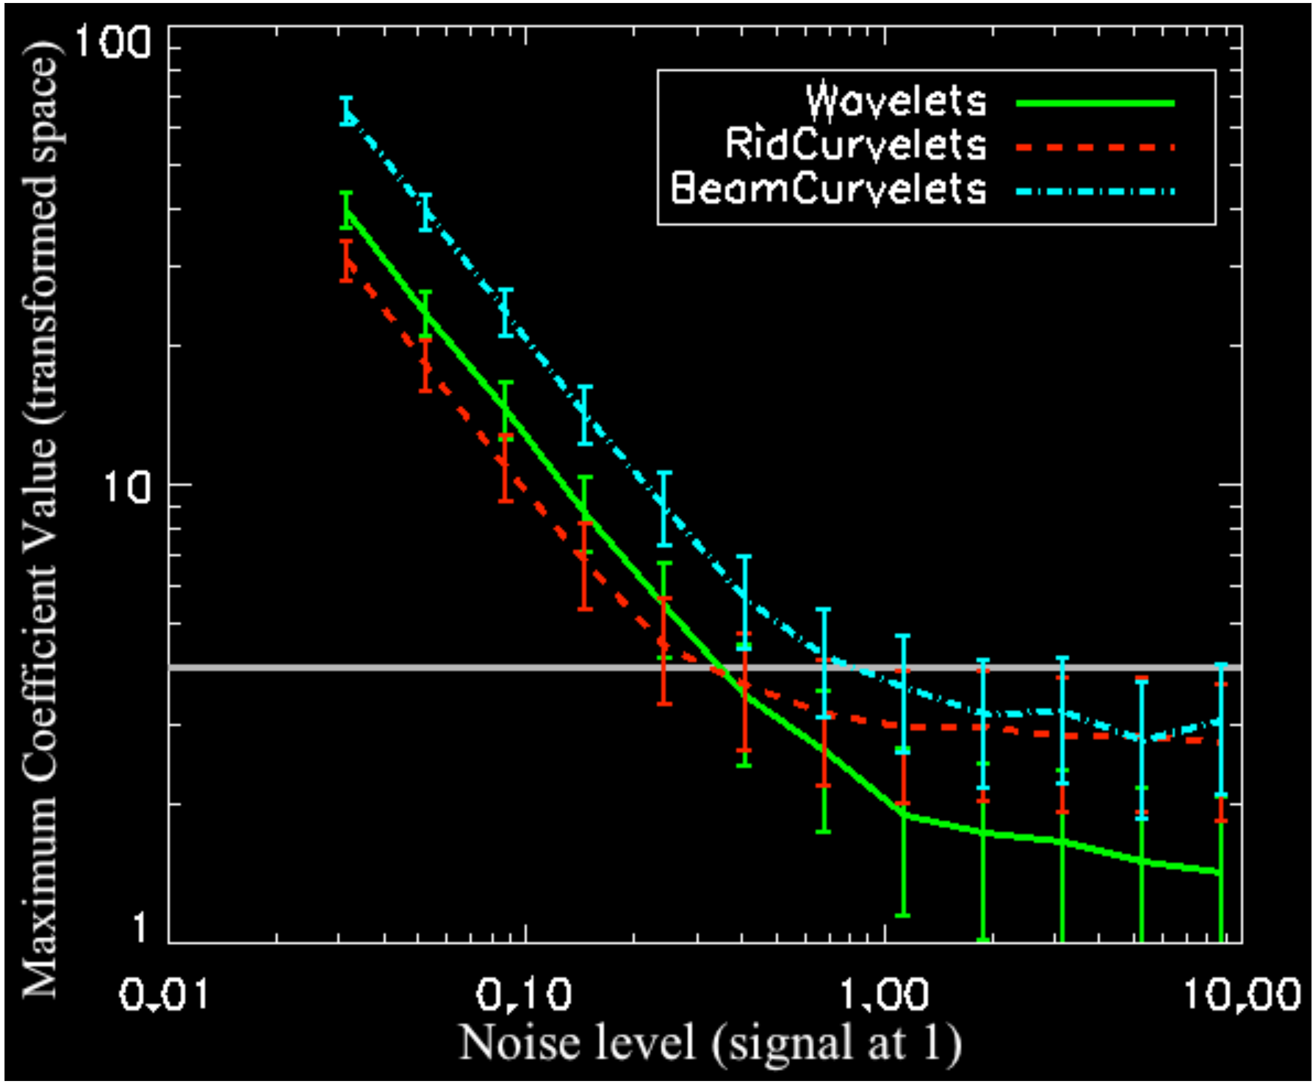

Combined denoising

- Amalgamate several transforms in a single dictionary $\mathbf{\Phi} = [\mathbf{\Phi}_1, \cdots, \mathbf{\Phi}_K]$ to benefit from the best of each transform.
- More flexbility to represent complex geometrical content: the blessing of over-completness.
- We have to solve

$$\min_{\alpha} \ \|\alpha\|_p^p \quad \text{s.t.} \quad \|g - \mathbf{\Phi}\alpha\|_2 \leq \epsilon(\sigma) \ , \ 0 \leq p \leq 1 .$$

- Solutions by e.g.:
  - Convex optimization (monotone operator splitting).
  - Greedy pursuit.

Combined denoising results

Cold Dark Matter simulations: clusters and filamentary structures with density of the filaments 3 orders of magnitude lower than the clusters.

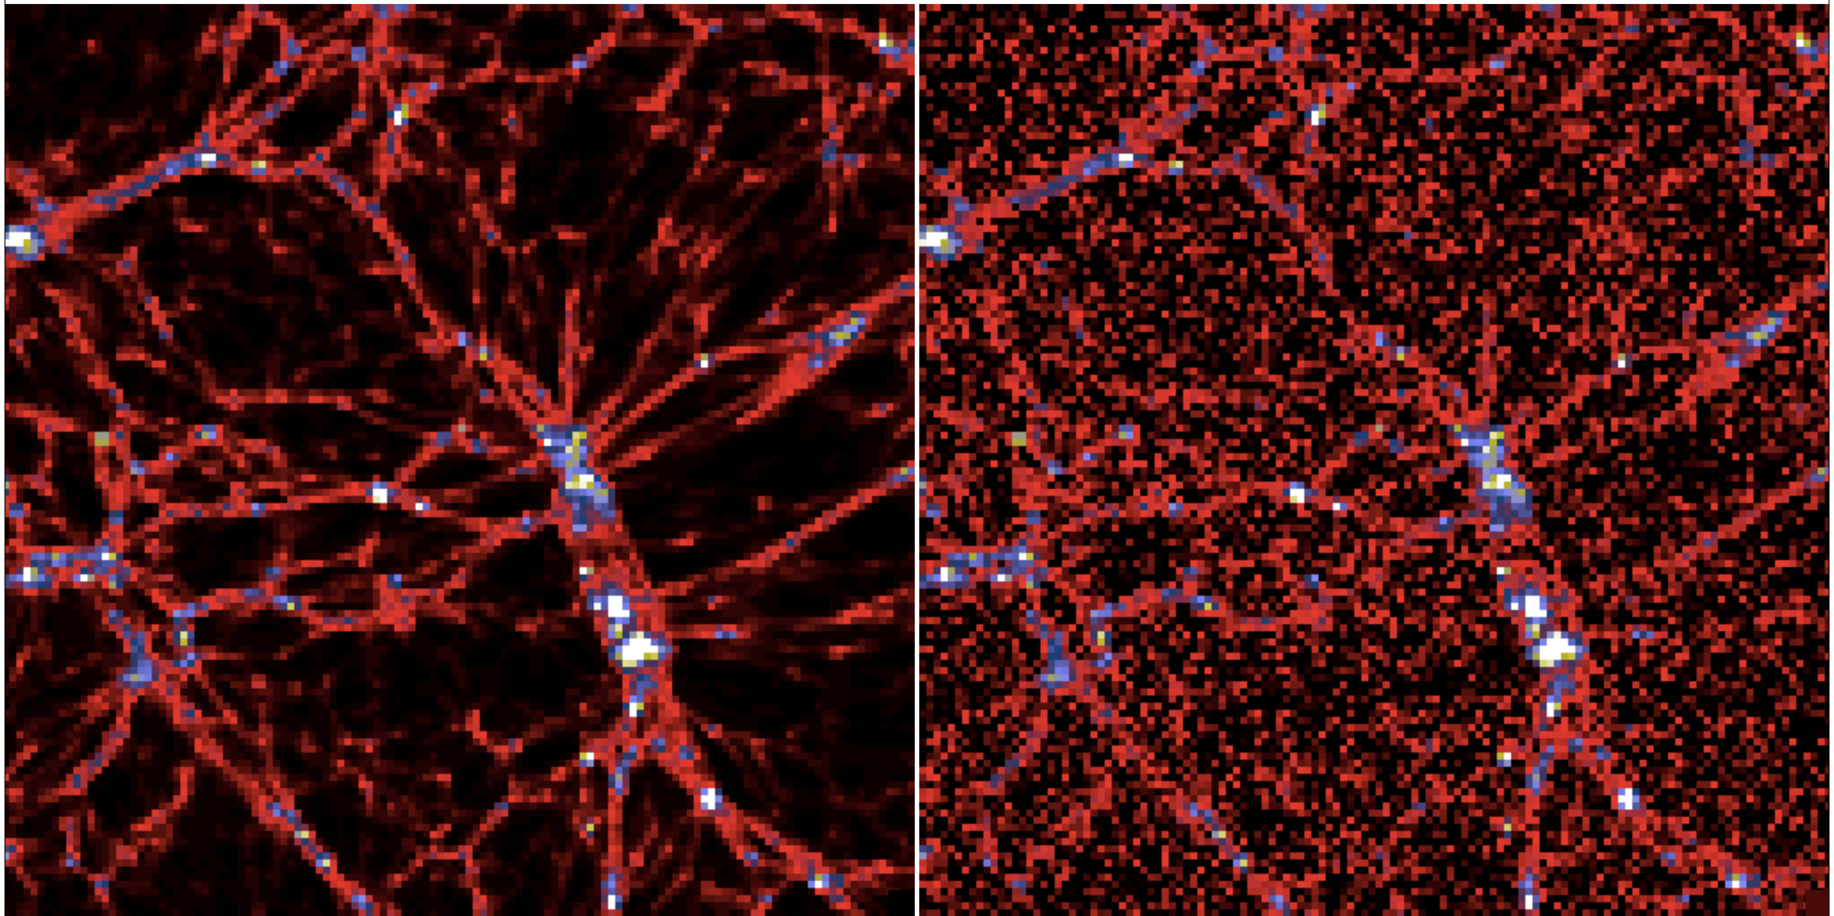

Combined denoising results

3D UDWT

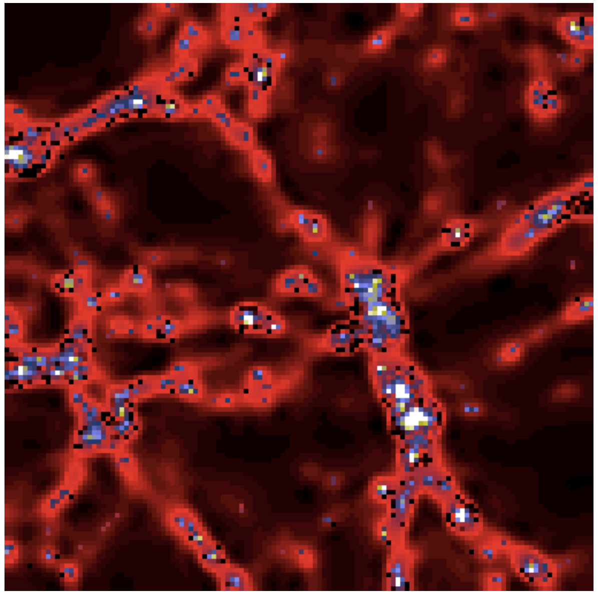

BeamCurvelets

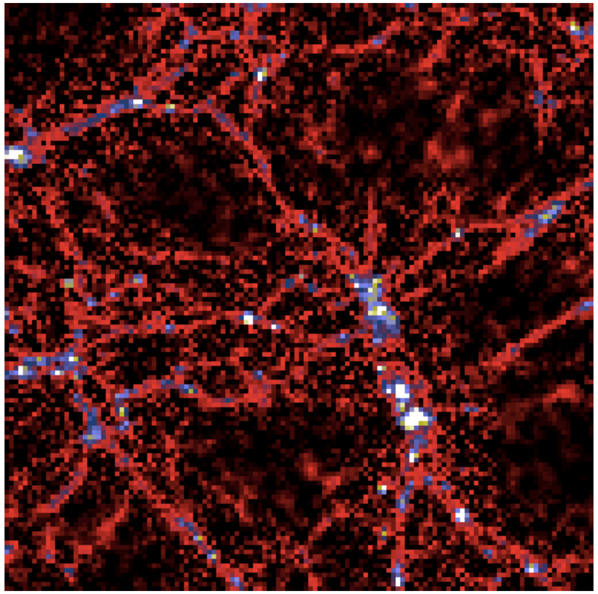

Combined denoising
BeamCurvelets+3D UDWT

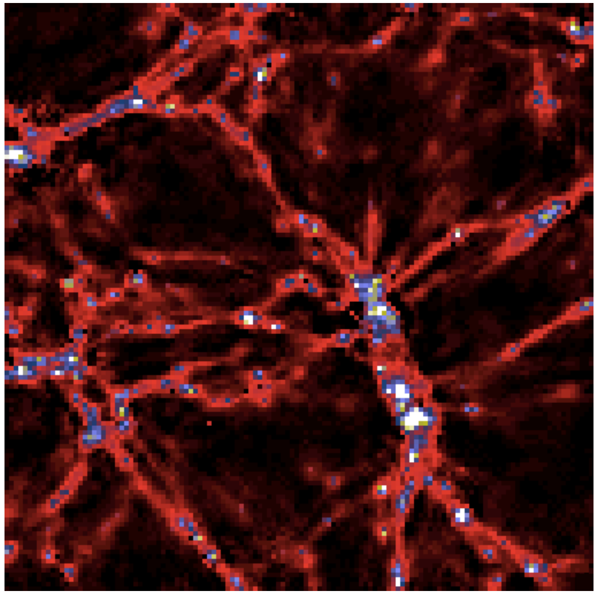

Fast second generation 3D curvelets

- 3D Fast curvelet transform (Candes et al, 2005)
- Curvelab, a C++/Matlab toolbox available at www.curvelet.org..

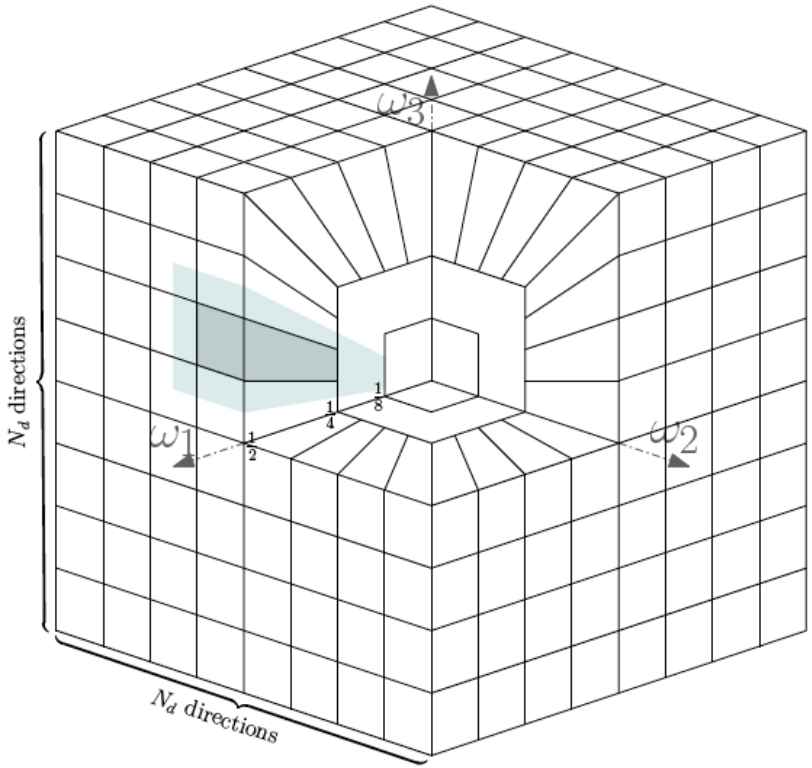

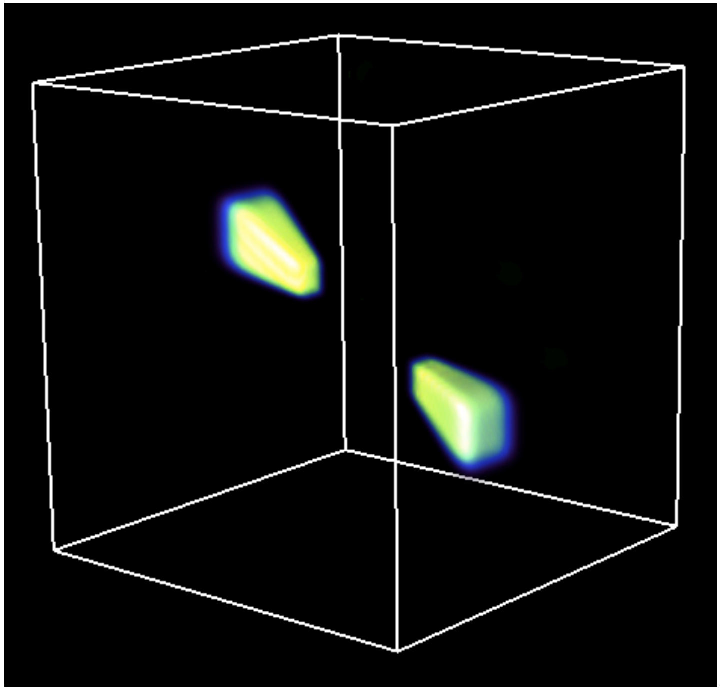

Redundancy: 7.11 in 2D and 24.38 in 3D
with walevets at the finest scale: 3.56 in 2D and 5.42 in 3D

Fast second generation 3D curvelets

- 3D Fast curvelet transform.
- Main differences with Candès et al. CurveLab :
  - Implementation: e.g. wavelet transform, overlapping angular windows.
  - Much less redundant than Candès et al. (2.3-10.3 instead of 5.4-24.4).
  - Faster in practice.

	Original FCT		Proposed FCT	
	C	W	C	W
2-D	7.11	3.56	4.00	2.00
3-D	24.38	5.42	10.29	2.29

. Woiselle, J.L. Starck, M.J. Fadili, "3D Data Denoising and Inpainting with the Fast Curvelet transform", **JMIV**, 39, 2, pp 121-139, 2011.

Fast second generation 3D curvelets

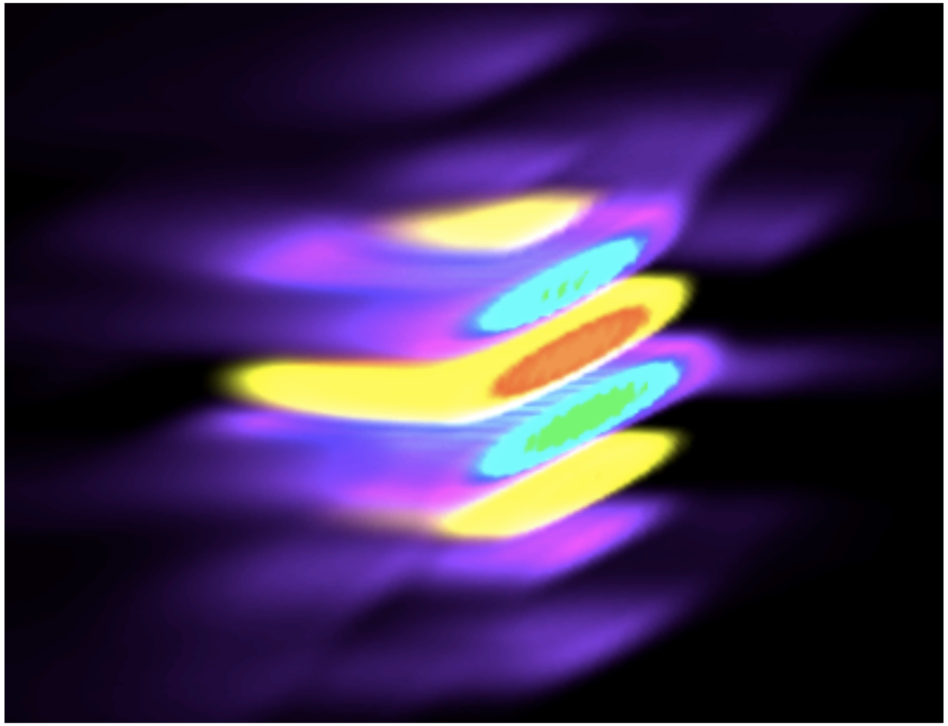

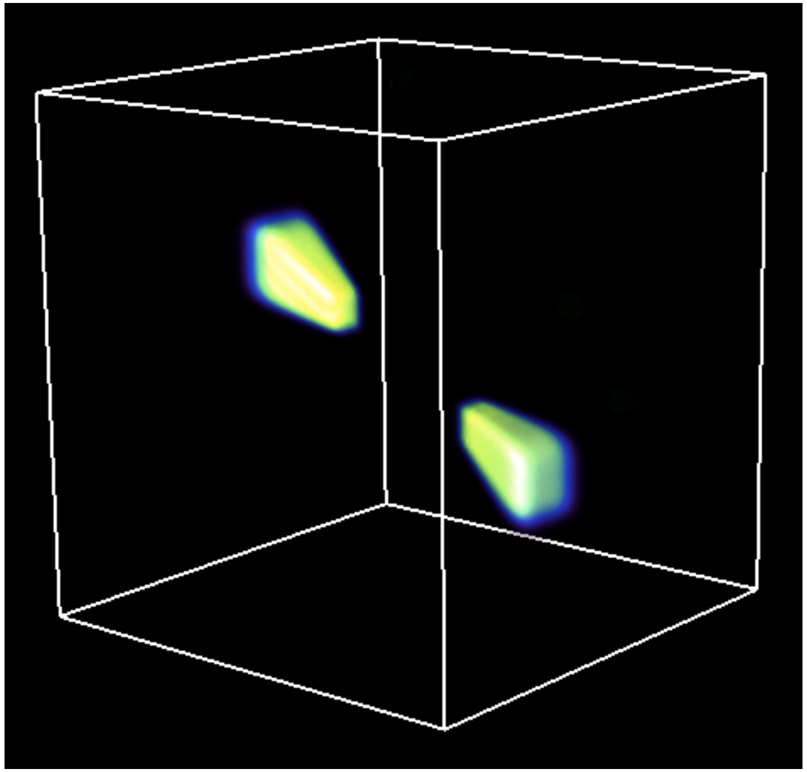

Fast second generation 3D curvelets

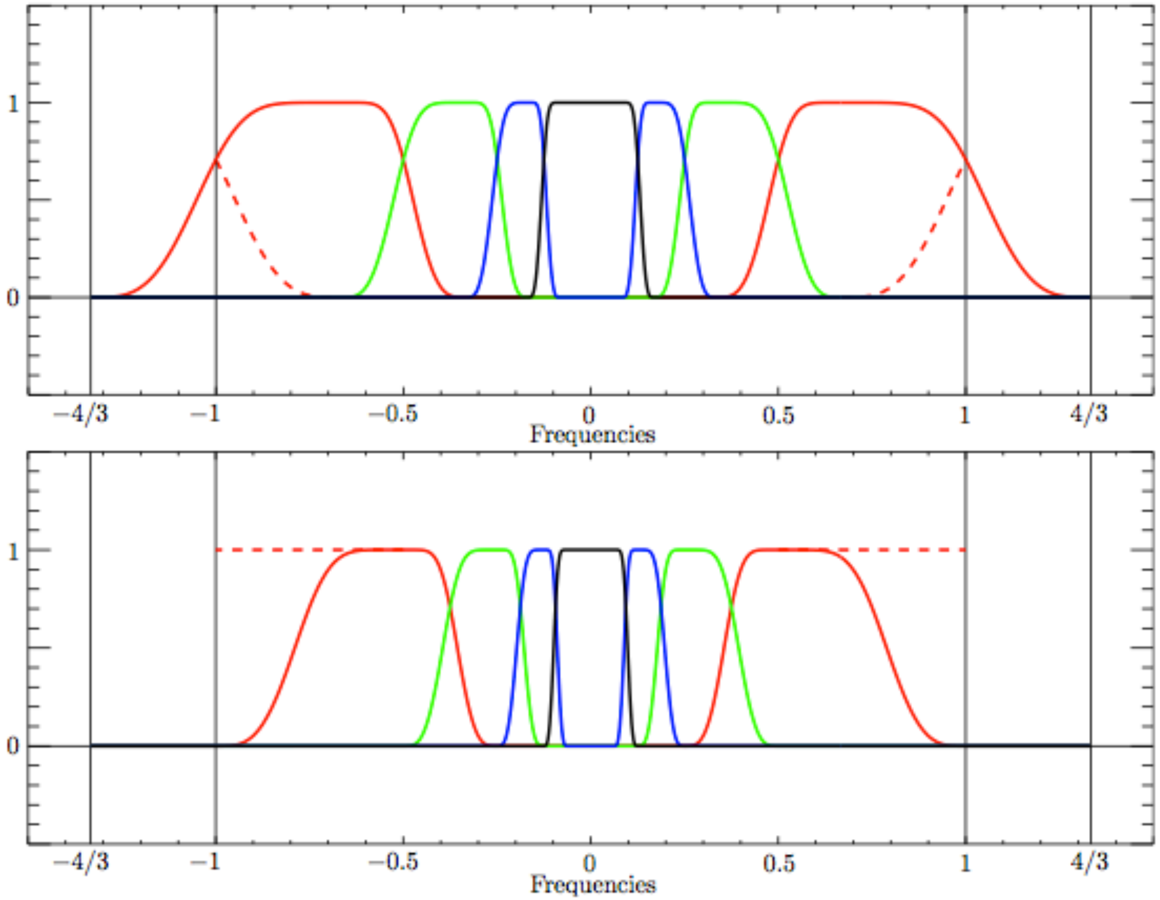

Meyer wavelets functions in Fourier domain. In the discrete case, we only have access to the Fourier samples inside the Shannon band [−1/2, 1/2], while the wavelet corresponding to the finest scale (solid red line) exceeds the Shannon frequency band to 2/3. Top : In the Curvelab implementation, the Meyer wavelet basis is periodized in Fourier, so that the exceeding end of the finest scale wavelet is replaced with the mirrored dashed line on the plot. Bottom : In our implementation, the wavelets are shrunk so that they fit in the [−1/2, 1/2] Shannon band, and the decreasing tail of the finest scale wavelet is replaced by a constant (dashed red line) to ensure a uniform partition of the unity.

Example: denoising

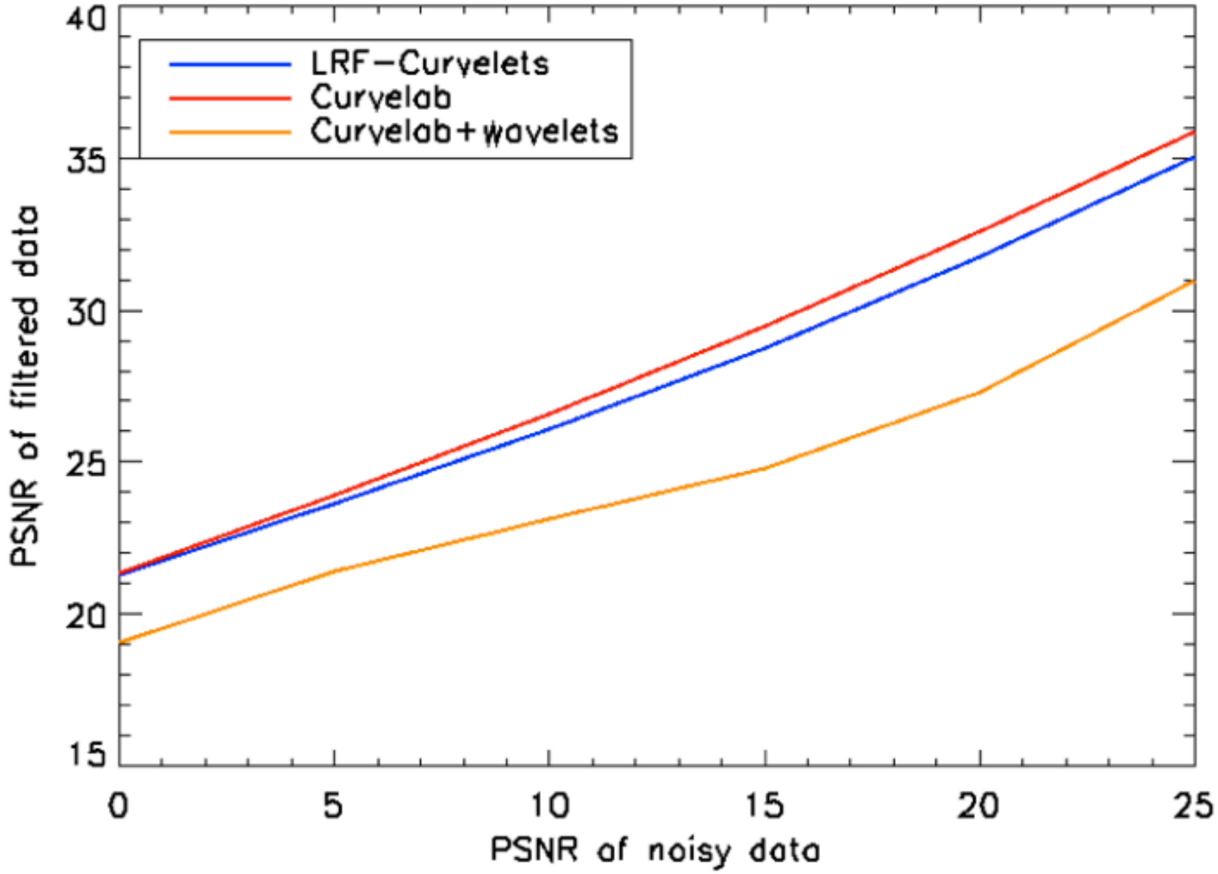

Mean denoising PSNR versus noise level using different FCT implementations. The denoising PSNR was averaged over ten noise realizations and several datasets. The LR-FCT is in blue. Original FCT implementation of Curvelab using curvelets (red) and wavelets (orange) at the finest scale.

Example: video denoising

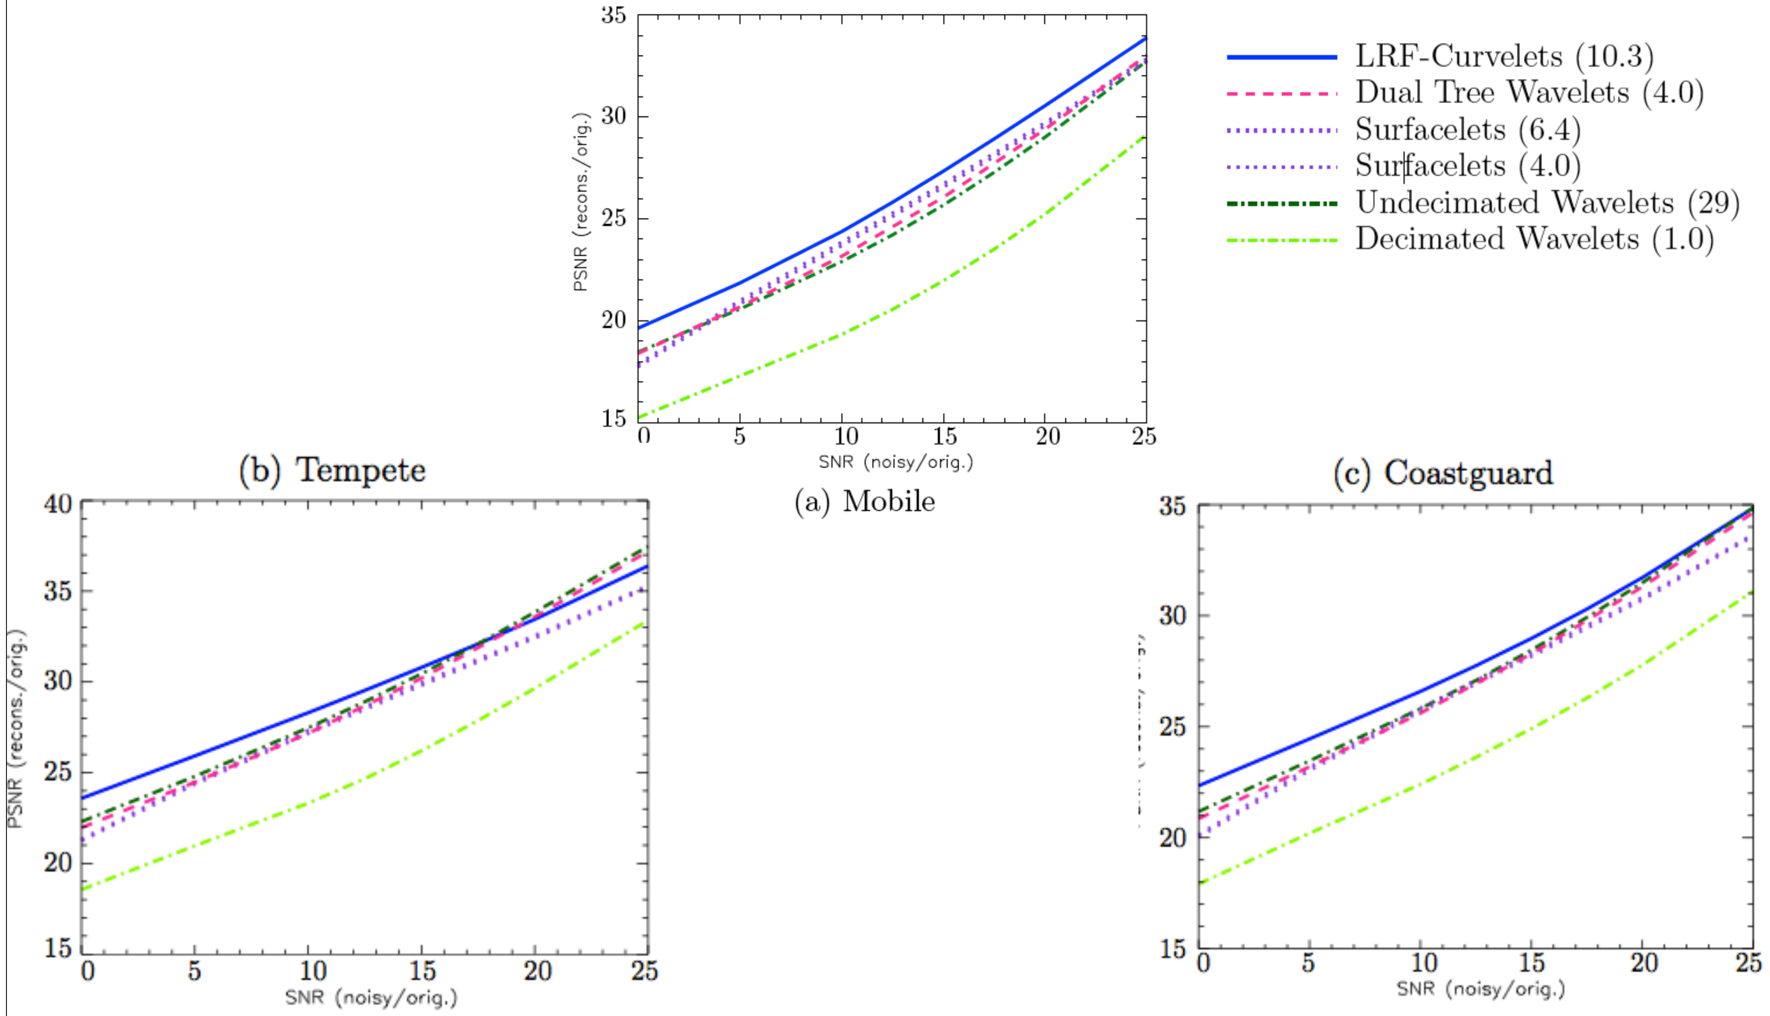

Output PSNR as a function of the input PSNR for three video sequences. (a) mobile, (b) tempete, and (c) coastguard CIF sequence.

Videos available at: www.cipr.rpi.edu

Example: hyperspectral data denoising

DATA: OMEGA spectrometer on Mars Express (www.esa.int/marsexpress) with 128 wavelength from 0.93µm to 2.73µm.

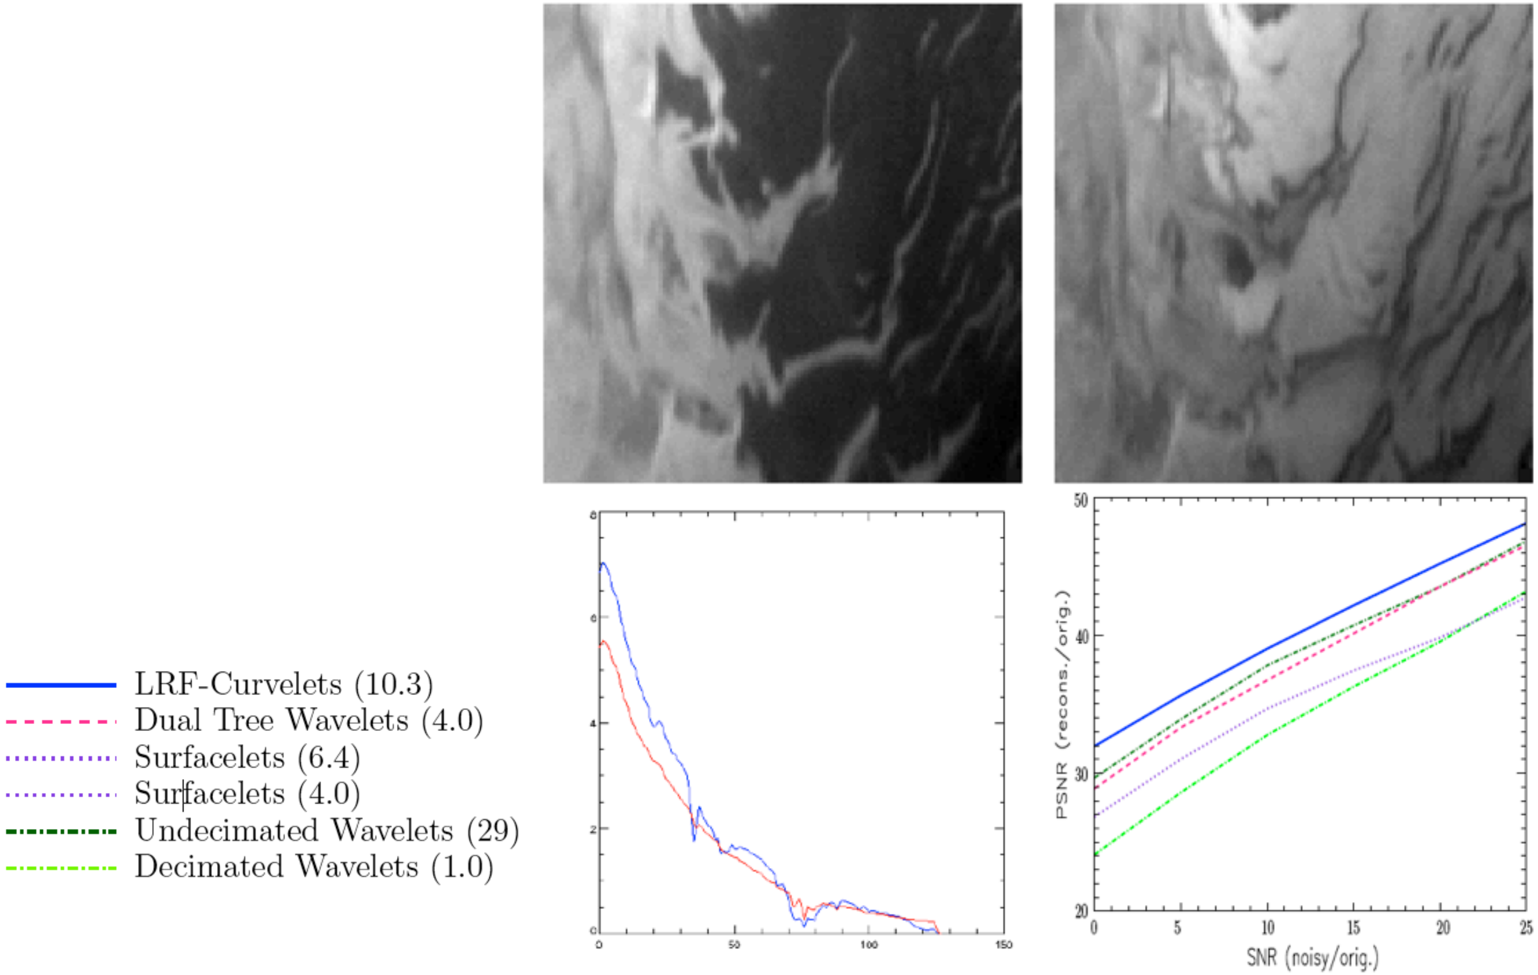

Top row : Mars Express observations at two different wavelengths. Bottom-left : two spectra at two distinct pixels.
Bottom-right : output PSNR as a function of the input PSNR for different transforms

Sparse component separation

- Separate an image into its morphological components from $g = \sum_{k=1}^{K} f_k + \varepsilon = \sum_{k=1}^{K} \mathbf{\Phi}_k \alpha_k + \varepsilon$, each α_k is sparse in $\mathbf{\Phi}_k$ but not (or less) sparse in $\mathbf{\Phi}_{k' \neq k}$.
- A sparse decomposition problem solved by MCA [Starck et al. 2004-2009].

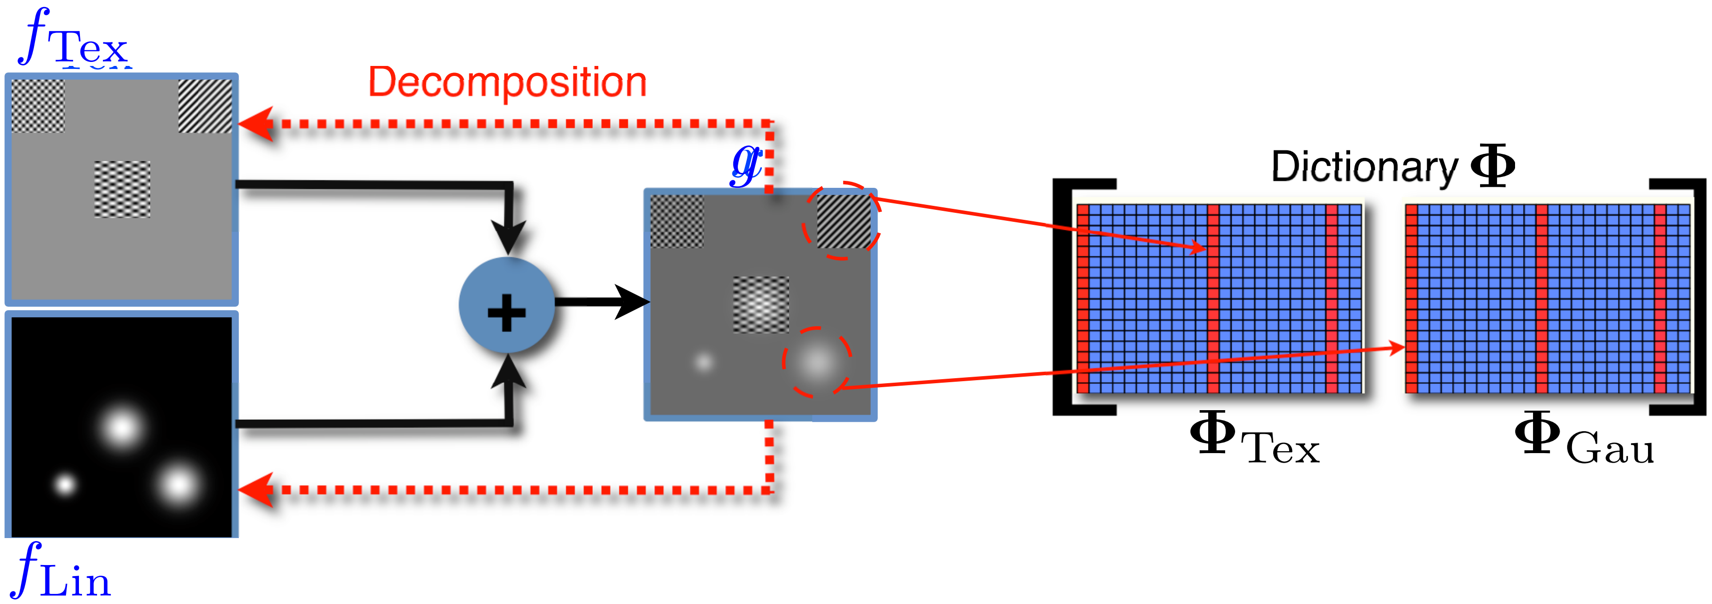

Results

Original (3D shells + Gaussians)

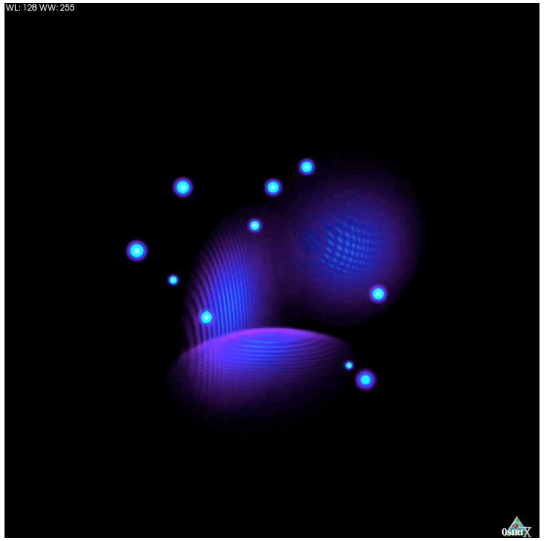

Dictionary

RidCurvelets + 3D UDWT.

Shells

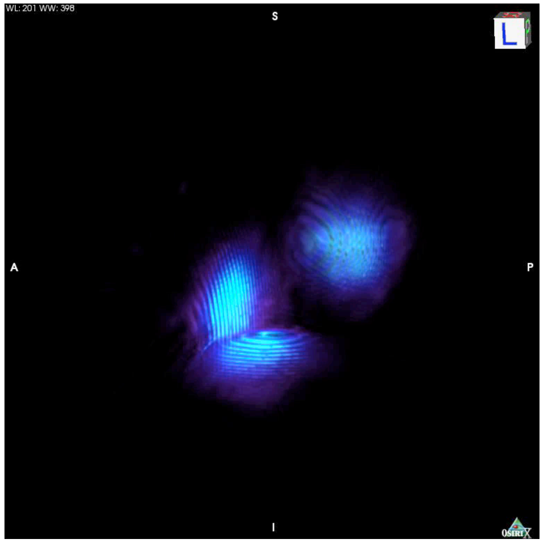

Gaussians

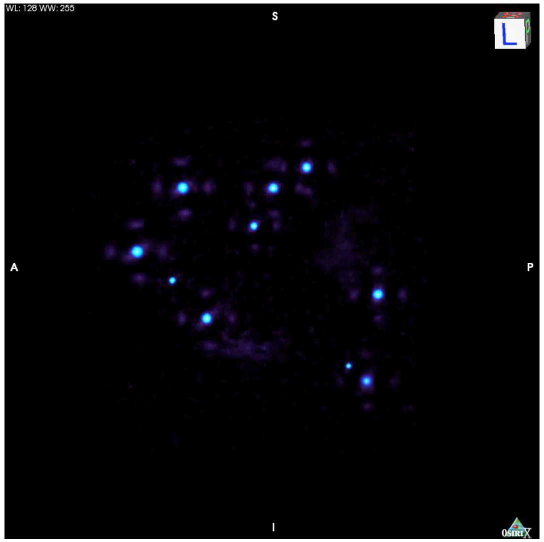

Inpainting

- Restore an image from its degraded version with missing samples:

$$g = \mathbf{M\Phi}\alpha + \boldsymbol{\varepsilon} \ ,$$

- An instance of CS reconstruction.
- Solve

$$\min_{\alpha} \ \|\alpha\|_p^p \quad \text{s.t.} \quad \|g - \mathbf{M\Phi}\alpha\|_2 \leq \epsilon \ .$$

- Can be solved by several algorithms, here we use an adaptation of MCA.

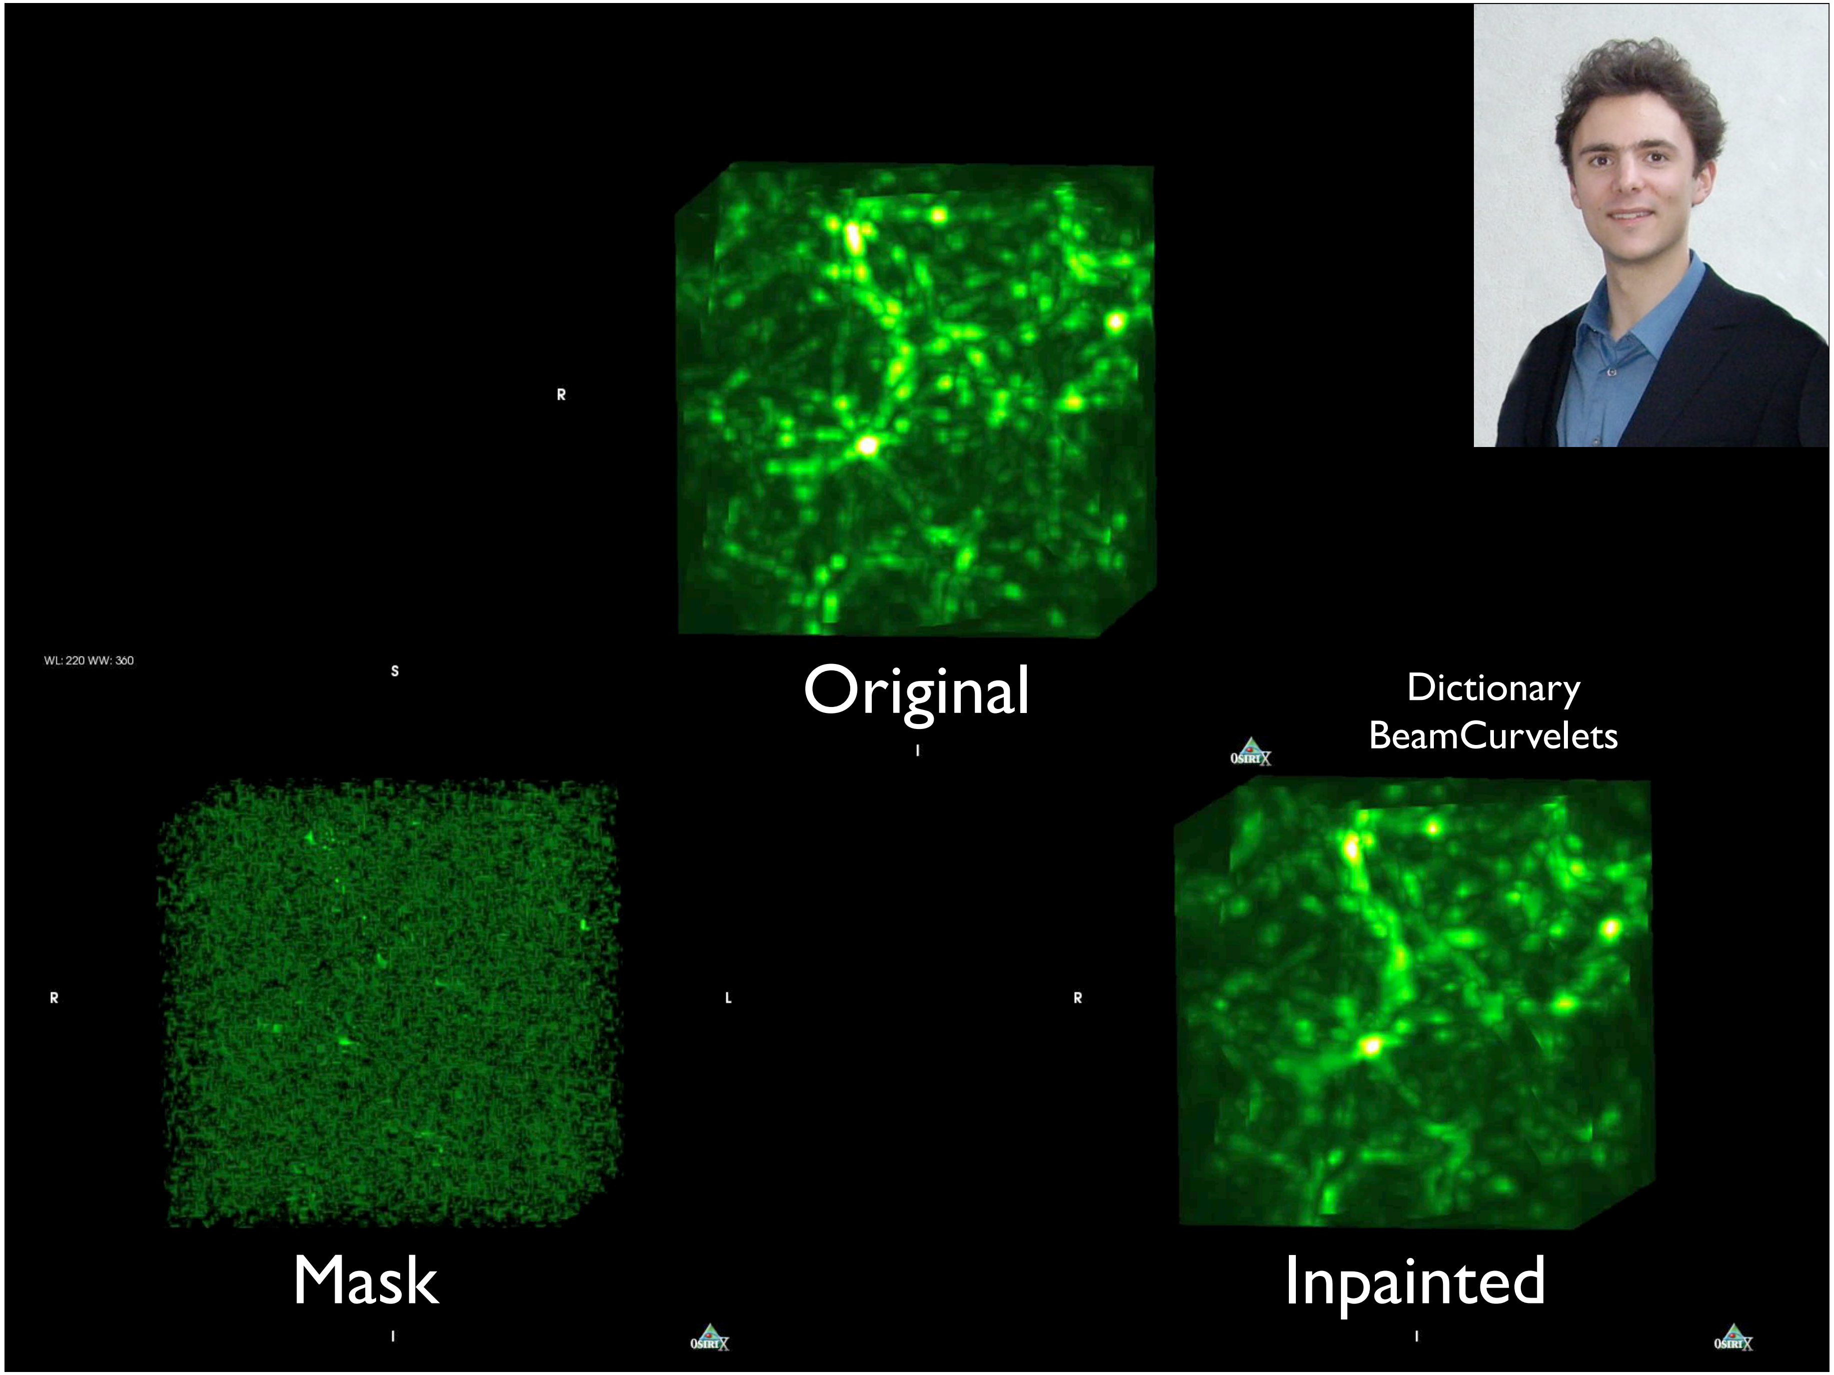
R
WL: 220 WW: 360
S
Original
I
Dictionary
BeamCurvelets
R
L
R
Mask
I
Inpainted
I

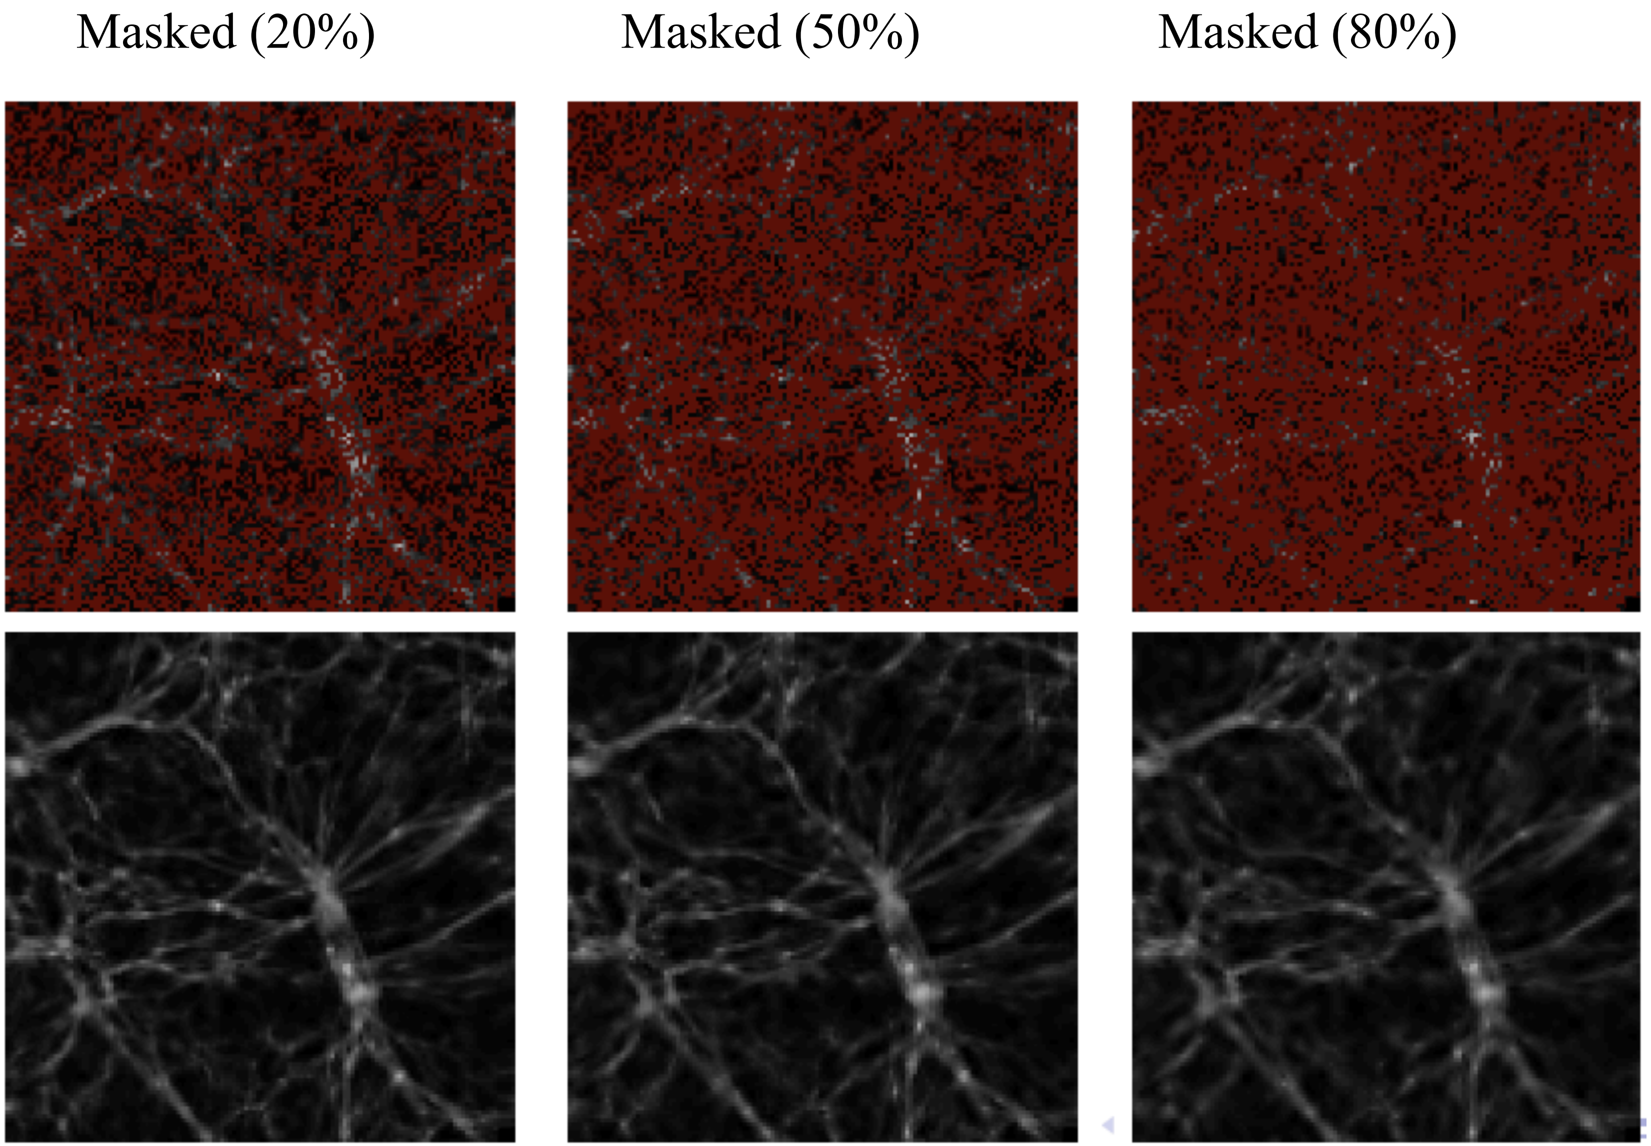

Central slice of the masked CDM data with 20, 50, and 80% missing voxels,
and the inpainted maps. The missing voxels are dark red.

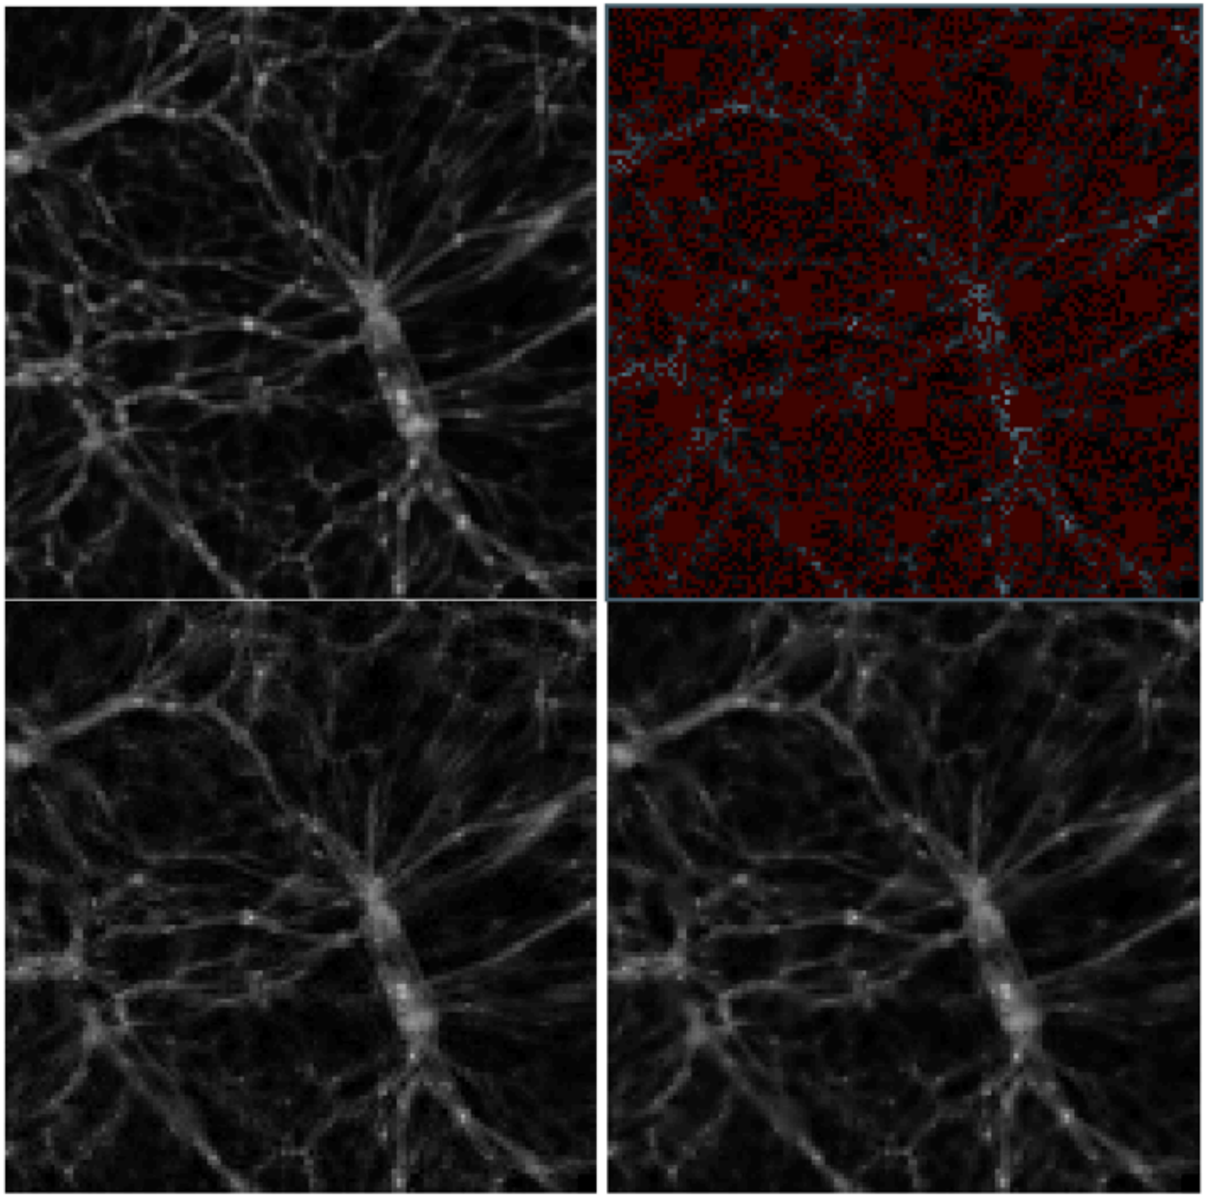

First row : original central frame of the CDM data cube, and degraded version with missing voxels in red.
Bottom row : the filtered results using the RidCurvelets (left) and the BeamCurvelets (right).

A sagittal ((y; z)) slice of the original synthetic MRI volume from BrainWeb.

inpainting results with a FCT +UDWT dictionary.

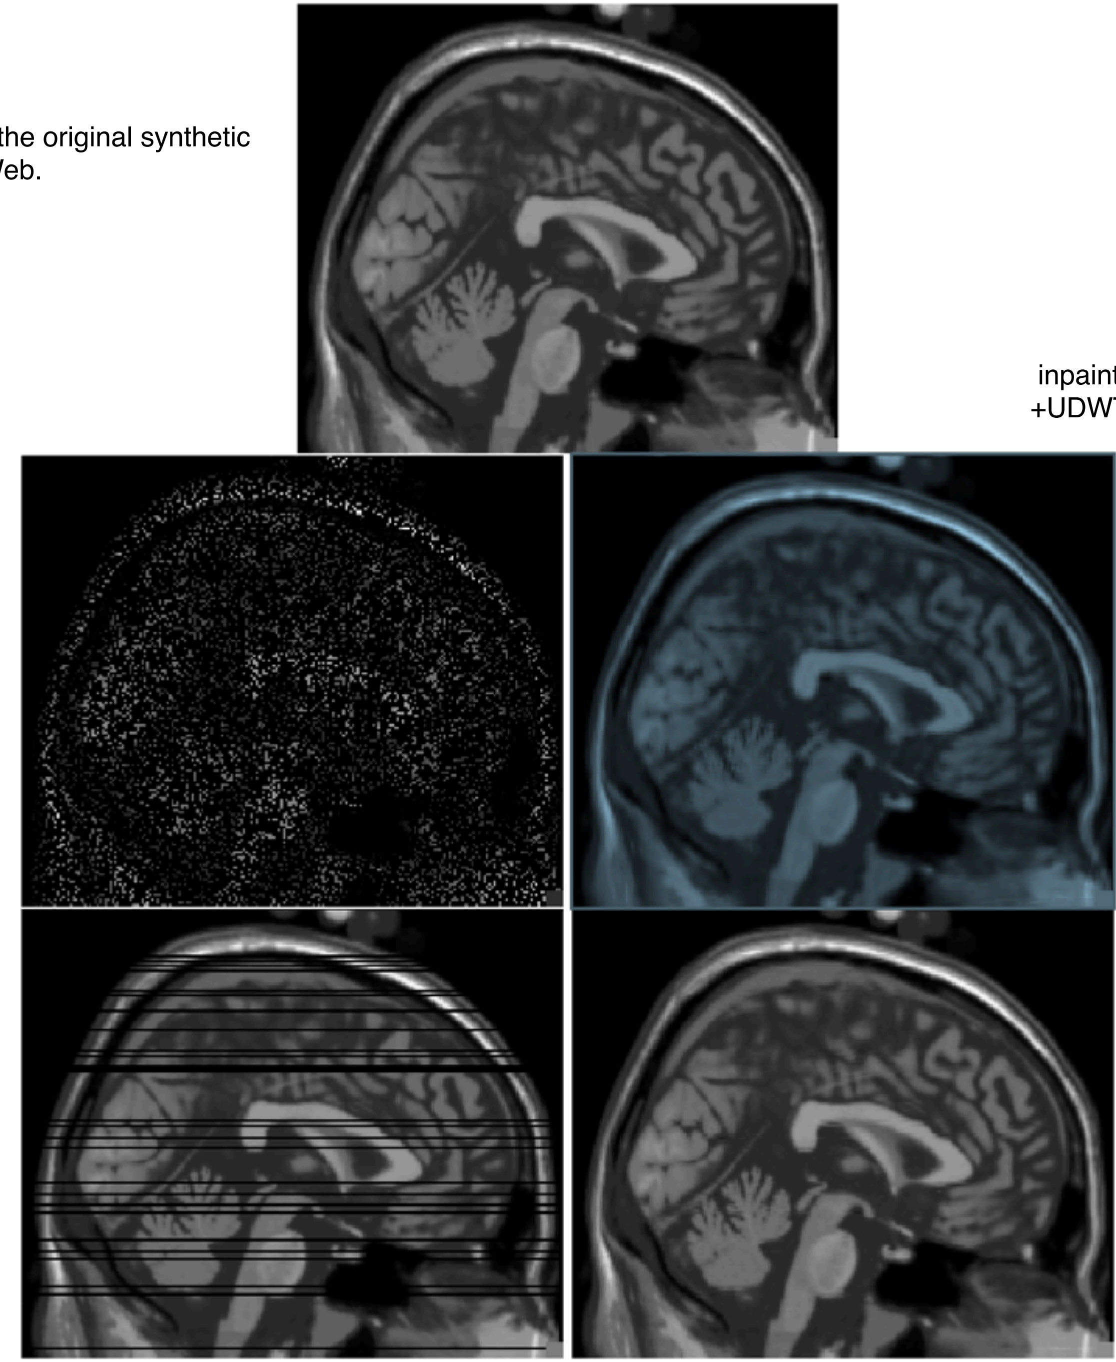

random 80% missing voxels,

10% missing z slices.

Conclusions

- Several 3D multiscale oriented representations.
- Adapted to sparsify several geometrical structures: filamentary and planar segments.
- Fast analysis and synthesis algorithms (FFT-based): parallel implementations.
- A wide variety of applications.

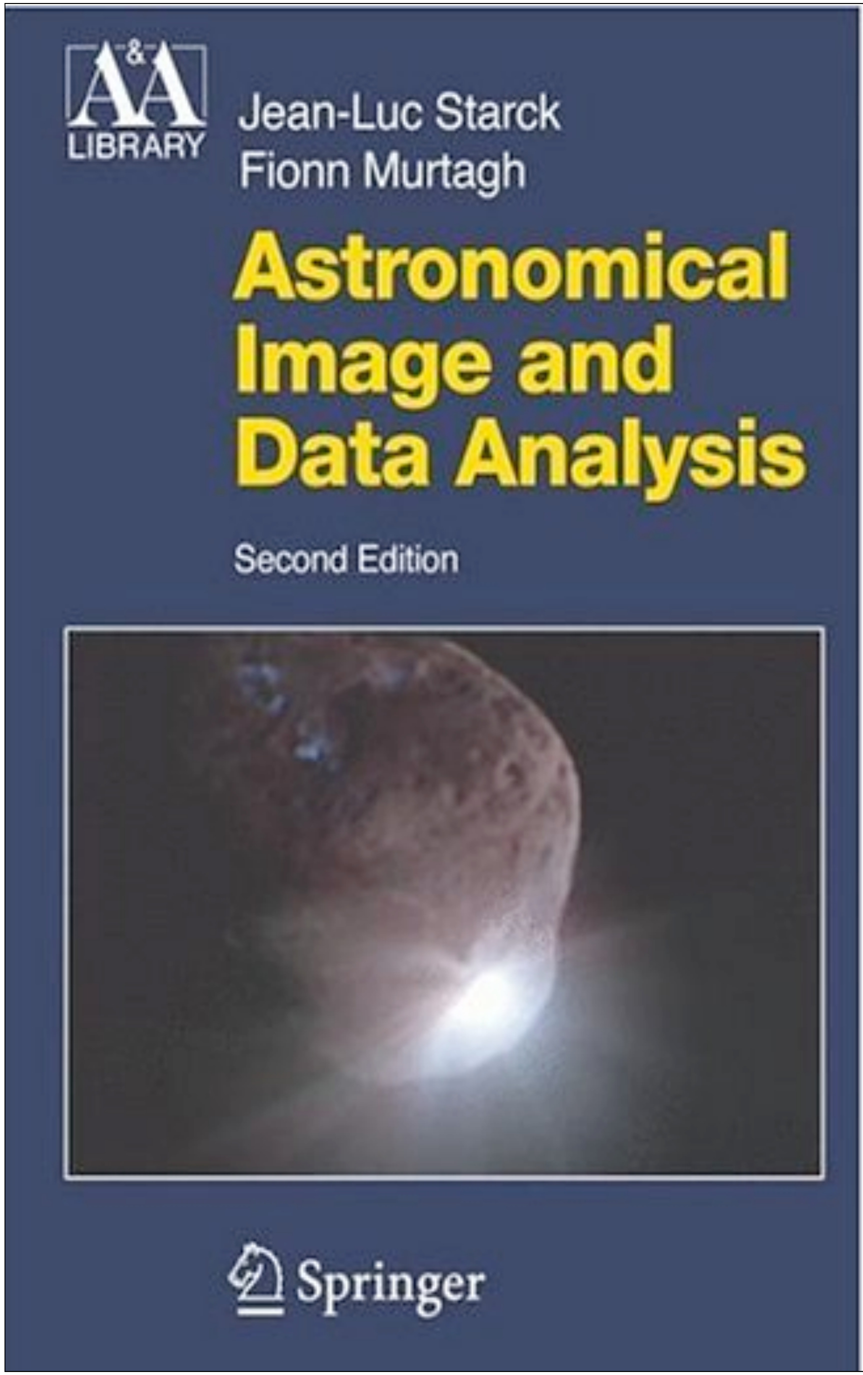
A&A LIBRARY
Jean-Luc Starck
Fionn Murtagh
Astronomical Image and Data Analysis
Second Edition
Springer

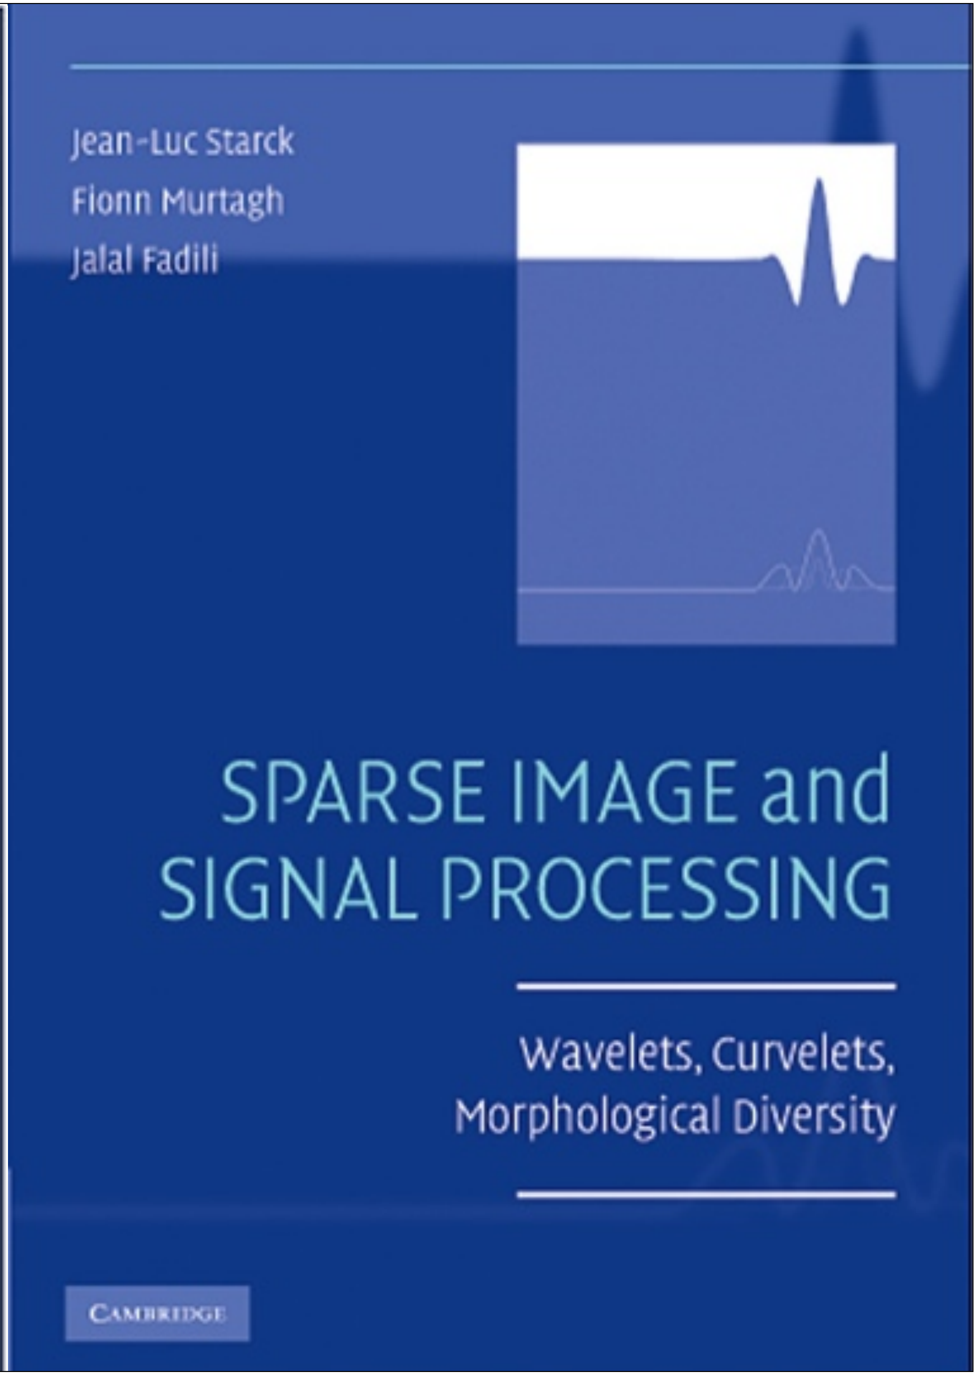
Jean-Luc Starck
Fionn Murtagh
Jalal Fadili
SPARSE IMAGE and SIGNAL PROCESSING
Wavelets, Curvelets, Morphological Diversity
CAMBRIDGE