A panorama on 2D directional wavelets: contourlets, curvelets, shearlets, *lets, etc.

Laurent Jacques, Laurent Duval†, Caroline Chaux, Gabriel Peyré

UCL, IFPEN, AMU, Dauphine

05/11/2013

Séminaire de l'Institut Langevin (ondes et images)

Wavelets for the eye


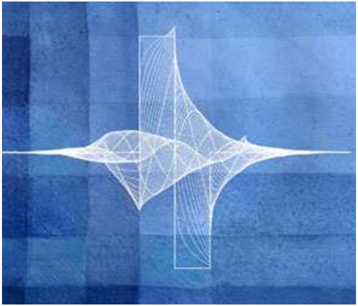

Artlets: painting wavelets

Wavelets for signals

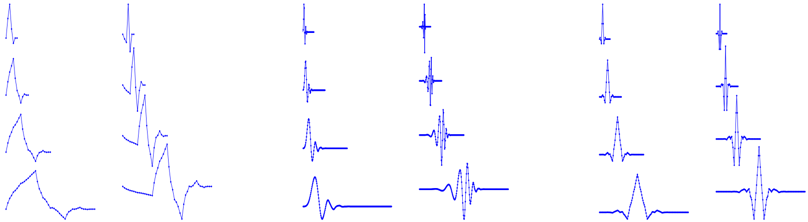

1D scaling functions and wavelets

Wavelets for images

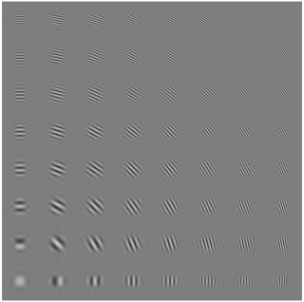

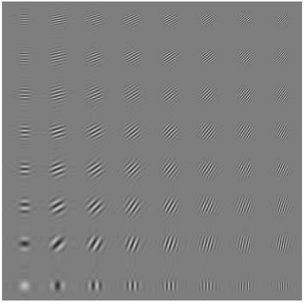

2D scaling functions and wavelets

1D signals

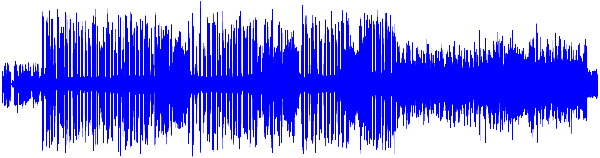

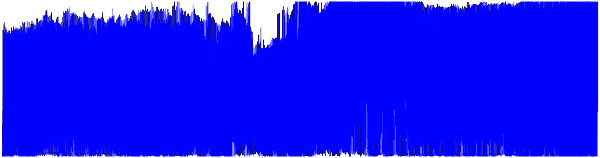

2D images

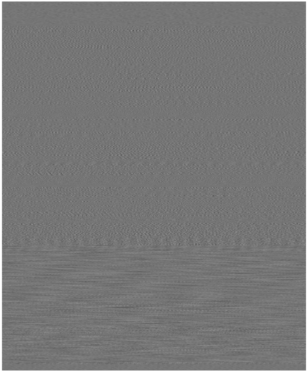

1.5D signals: motivations for 2D directional "wavelets"

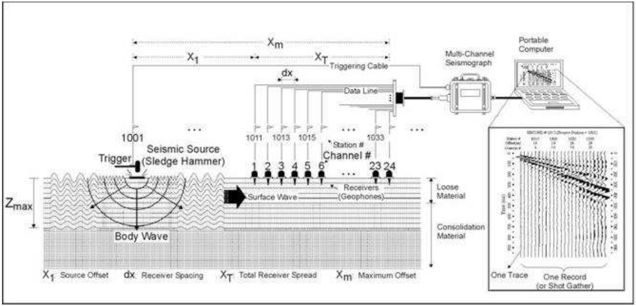

Figure : Geophysics: seismic data recording (surface and body waves)

1.5D signals: motivations for 2D directional "wavelets"

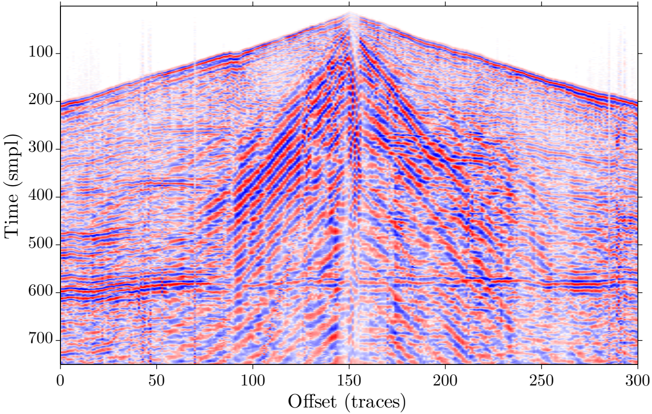

Figure : Geophysics: surface wave removal (before)

1.5D signals: motivations for 2D directional "wavelets"

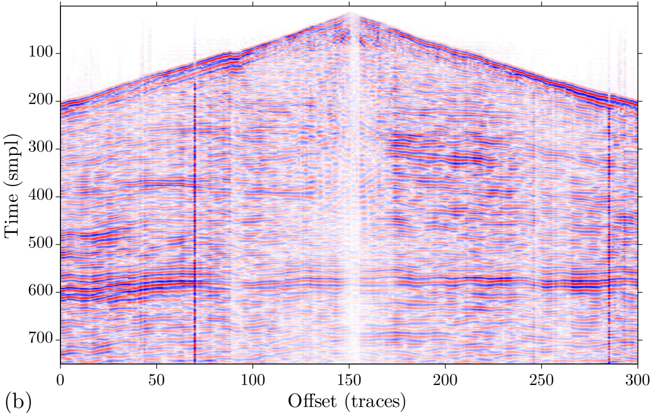

Figure : Geophysics: surface wave removal (after)

1.5D signals: motivations for 2D directional "wavelets"

Issues here:

- different types of waves on seismic "images"
  - appear hyperbolic [layers], linear [noise] (and parabolic)
- not the standard "mid-amplitude random noise problem"
- kind of halfway between signals and images (1.5D)
- yet local, directional, frequency-limited, scale-dependent structures to separate

Agenda

- To survey 15 years of improvements in 2D wavelets
  - with spatial, directional, frequency selectivity increased
  - yielding sparser representations of contours and textures
  - from fixed to adaptive, from low to high redundancy
  - generally fast, compact (if not sparse), informative, practical
  - requiring lots of hybridization in construction methods
- Outline
  - introduction + early days ($\leqslant$ 1998)
  - fixed: oriented & geometrical (selected):
    - $\pm$ separable (Hilbert/dual-tree wavelet)
    - isotropic non-separable (Morlet-Gabor)
    - anisotropic scaling (ridgelet, curvelet, contourlet, shearlet)
  - (hidden bonuses):
    - adaptive, lifting, meshes, spheres, manifolds, graphs
  - conclusions

In just one slide

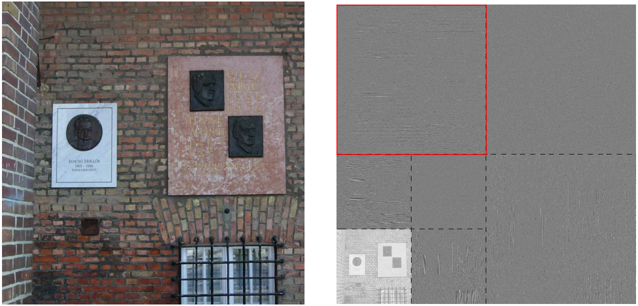

Figure : A standard, “dyadic”, separable wavelet decomposition

Where do we go from here? 15 years, 300+ refs in 30 minutes

Images are pixels (but...):

$$\tilde{X} = \begin{pmatrix} 67 & 93 & 129 & 155 \\ 52 & 97 & 161 & 207 \\ 33 & 78 & 143 & 188 \\ 22 & 48 & 84 & 110 \end{pmatrix}$$

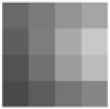

Figure : Image as a (canonic) linear combination of pixels

- suffices for (simple) data (simple) manipulation
  - counting, enhancement, filtering
- very limited in higher level understanding tasks
  - looking for other (meaningful) linear combinations
  - what about
    $67 + 93 + 52 + 97$, $67 + 93 - 52 - 97$
    $67 - 93 + 52 - 97$, $67 - 93 - 52 + 97$?

Images are pixels (but...):

A review in an active research field:

- (partly) inspired by:
  - early vision observations [Marr *et al.*]
  - sparse coding: wavelet-like oriented filters and receptive fields of simple cells (visual cortex) [Olshausen *et al.*]
  - a widespread belief in sparsity
- motivated by image handling (esp. compression)
- continued from the first successes of wavelets (JPEG 2000)
- aimed either at pragmatic or heuristic purposes
  - known formation model *or* unknown information
- developed through a quantity of *-lets and relatives

Images are pixels, wavelets are legion

Room(let) for improvement:

Activelet, AMlet, Armlet, Bandlet, Barlet, Bathlet, Beamlet, Binlet, Bumplet, Brushlet, Caplet, Camplet, Chirplet, Chordlet, Circlet, Coiflet, Contourlet, Cooklet, Craplet, Cubelet, CURElet, Curvelet, Daublet, Directionlet, Dreamlet, Edgelet, FAMlet, FLaglet, Flatlet, Fourierlet, Framelet, Fresnelet, Gaborlet, GAMlet, Gausslet, Graphlet, Grouplet, Haarlet, Haardlet, Heatlet, Hutlet, Hyperbolet, Icalet (Icalette), Interpolet, Loglet, Marrlet, MIMOlet, Monowavelet, Morelet, Morphlet, Multiselectivelet, Multiwavelet, Needlet, Noiselet, Ondelette, Ondulette, Prewavelet, Phaselet, Planelet, Platelet, Purelet, QVlet, Radonlet, RAMlet, Randlet, Ranklet, Ridgelet, Riezlet, Ripplet (original, type-I and II), Scalet, S2let, Seamlet, Seislet, Shadelet, Shapelet, Shearlet, Sinclet, Singlet, Slantlet, Smoothlet, Snakelet, SOHOlet, Sparselet, Spikelet, Splinelet, Starlet, Steerlet, Stockeslet, SURE-let (SURElet), Surfacelet, Surflet, Symmlet, S2let, Tetrolet, Treelet, Vaguelette, Wavelet-Vaguelette, Wavelet, Warblet, Warplet, Wedgelet, Xlet, not mentioning all those not on -let!

Now, some reasons behind this quantity

Images are pixels, but altogether different

Figure : Different kinds of images

Images are pixels, but altogether different

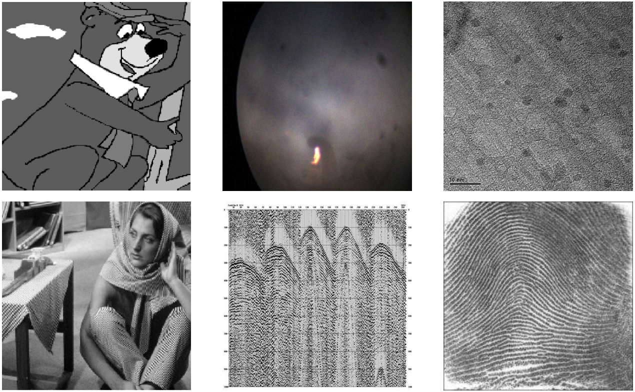

Figure : Different kinds of images

Images are pixels, but might be described by models

To name a few:

- edge cartoon + texture:

[Meyer-2001]

$$\inf_u E(u) = \int_\Omega |\nabla u| + \lambda \|v\|_*, f = u + v$$

- edge cartoon + texture + noise:

[Aujol-Chambolle-2005]

$$\inf_{u,v,w} F(u,v,w) = J(u) + J^* \left(\frac{v}{\mu}\right) + B^* \left(\frac{w}{\lambda}\right) + \frac{1}{2\alpha}\|f - u - v - w\|_{L^2}$$

- heuristically: piecewise-smooth + contours + geometrical textures + noise (or unmodeled)

Images are pixels, but resolution/scale helps with models

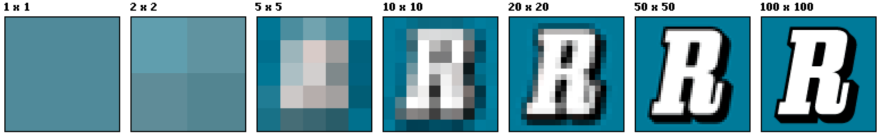

Figure : Notion of sufficient resolution [Chabat *et al.*, 2004]

- coarse-to-fine and fine-to-coarse relationships
- discrete 80's wavelets were not bad for: piecewise-smooth (moments) + contours (gradient-behavior) + geometrical textures (oscillations) + noise
- not enough for complicated images (poor sparsity decay)

Images are pixels, but might be described by models

Standard model: smooth w/ M-term approximation (2D $\neq$ 1D)

- $C^\alpha \to O(M^{-\alpha})$ (standard wavelets)
- piecewise $C^\alpha \to O(M^{-1})$ (standard wavelets)
- piecewise $C^2 \to O(M^{-2})$ (greedy triangles)

Question:

- piecewise $C^2 \to O(M^{-2} f(M))$ w/ fixed directional wavelet?

Images are pixels, but sometimes deceiving

Figure : Real world image and illusions

Images are pixels, but sometimes deceiving

Figure : Real world image and illusions

Images are pixels, but sometimes deceiving

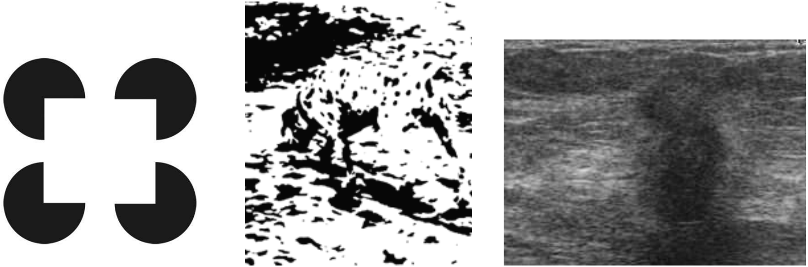

Figure : Real world image and illusions

Images are pixels, but resolution/scale helps

To catch important "objects" in their context

- use scales or multiresolution schemes,
- combine w/ various of description/detection/modeling methods:
  - smooth curve or polynomial fit, oriented regularized derivatives (Sobel, structure tensor), discrete (lines) geometry, parametric curve detectors (e.g. Hough transform), mathematical morphology, empirical mode decomposition, local *frequency estimators*, Hilbert and Riesz (analytic and monogenic), quaternions, Clifford algebras, optical flow approaches, smoothed random models, generalized Gaussian mixtures, warping operators, etc.

Images are pixels, and need efficient descriptions

Depend on application, with sparsity priors:

- compression, denoising, enhancement, inpainting, restoration, contour detection, texture analysis, fusion, super-resolution, registration, segmentation, reconstruction, source separation, image decomposition, MDC, learning, etc.

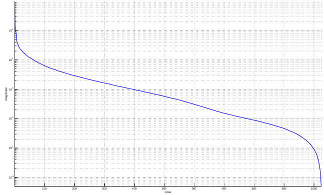

Figure : Image (contours/textures) and decaying singular values

Images are pixels: a guiding thread (GT)

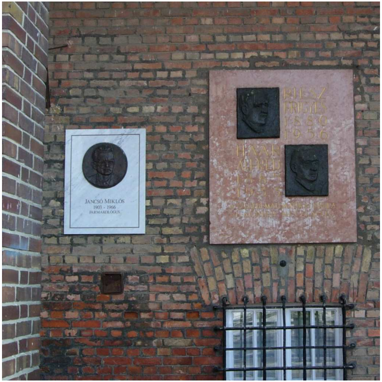

Figure : Memorial plaque in honor of A. Haar and F. Riesz: *A szegedi matematikai iskola világhírű megalapítói*, courtesy Prof. K. Szatmáry

Guiding thread (GT): early days

Fourier approach: critical, orthogonal

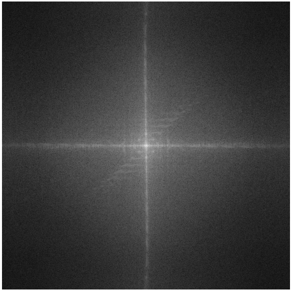

Figure : GT luminance component amplitude spectrum (log-scale)

Fast, compact, practical but not quite informative (not local)

Guiding thread (GT): early days

Scale-space approach: (highly)-redundant, more local

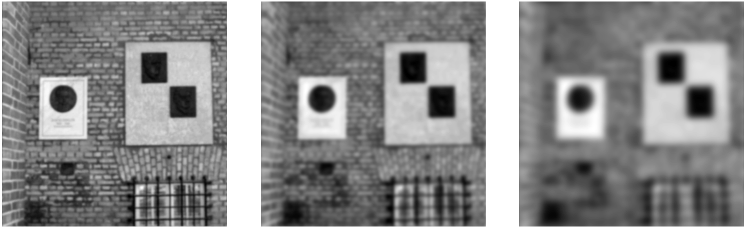

Figure : GT with Gaussian scale-space decomposition

Gaussian filters and heat diffusion interpretation
Varying persistence of features across scales $\Rightarrow$ redundancy

Guiding thread (GT): early days

Pyramid-like approach: (less)-redundant, more local

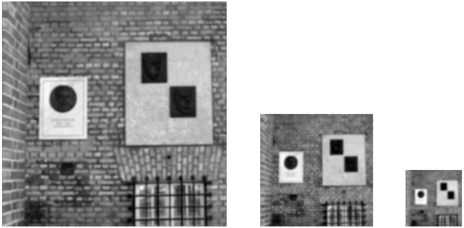

Figure : GT with Gaussian scale-space decomposition

Gaussian pyramid

Varying persistence of features across scales + subsampling

Guiding thread (GT): early days

Differences in scale-space with subsampling

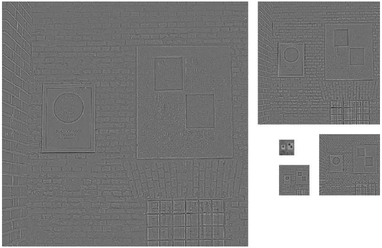

Figure : GT with Laplacian pyramid decomposition

Laplacian pyramid: complete, reduced redundancy, enhances image singularities, low-activity regions/small coefficients, algorithmic

Guiding thread (GT): early days

Isotropic wavelets (more axiomatic)

Consider

Wavelet $\psi \in \mathbb{L}^2(\mathbb{R}^2)$ such that $\psi(\boldsymbol{x}) = \psi_{\rm rad}(\|\boldsymbol{x}\|)$, with $\boldsymbol{x} = (x_1, x_2)$, for some radial function $\psi_{\rm rad} : \mathbb{R}_+ \to \mathbb{R}$ (with adm. conditions).

Decomposition and reconstruction

For $\psi_{(\boldsymbol{b},a)}(\boldsymbol{x}) = \frac{1}{a}\psi(\frac{\boldsymbol{x}-\boldsymbol{b}}{a})$, $W_f(\boldsymbol{b}, a) = \langle \psi_{(\boldsymbol{b},a)}, f \rangle$ with reconstruction:

$$f(\boldsymbol{x}) = \frac{2\pi}{c_\psi} \int_0^{+\infty}\int_{\mathbb{R}^2} W_f(\boldsymbol{b}, a)\ \psi_{(\boldsymbol{b},a)}(\boldsymbol{x})\ \mathrm{d}^2\boldsymbol{b}\,\frac{\mathrm{d}a}{a^3} \qquad (1)$$

if $c_\psi = (2\pi)^2 \int_{\mathbb{R}^2} |\hat{\psi}(\boldsymbol{k})|^2/\|\boldsymbol{k}\|^2\ \mathrm{d}^2\boldsymbol{k} < \infty$.

Guiding thread (GT): early days

Wavelets as multiscale edge detectors: many more potential wavelet shapes (difference of Gaussians, Cauchy, etc.)

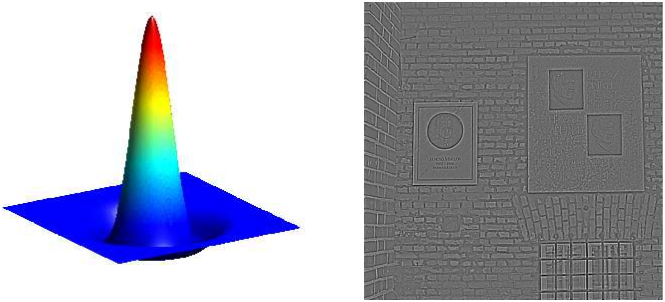

Figure : Example: Marr wavelet as a singularity detector

Guiding thread (GT): early days

Definition

The family $\mathcal{B}$ is a frame if there exist two constants $0 < \mu_1 \leqslant \mu_2 < \infty$ such that for all f

$$\mu_1 \|f\|^2 \leqslant \sum_{\boldsymbol{m}} |\langle \psi_{\boldsymbol{m}}, f \rangle|^2 \leqslant \mu_2 \|f\|^2$$

Possibility of discrete orthogonal bases with $O(N)$ speed. In 2D:

Definition

Separable orthogonal wavelets: dyadic scalings and translations $\psi_{\boldsymbol{m}}(\boldsymbol{x}) = 2^{-j}\psi^k(2^{-j}\boldsymbol{x} - \boldsymbol{n})$ of three tensor-product 2-D wavelets

$$\psi^V(\boldsymbol{x}) = \psi(x_1)\varphi(x_2),\ \psi^H(\boldsymbol{x}) = \varphi(x_1)\psi(x_2),\ \psi^D(\boldsymbol{x}) = \psi(x_1)\psi(x_2)$$

Guiding thread (GT): early days

So, back to orthogonality with the discrete wavelet transform: fast, compact and informative, but... is it sufficient (singularities, noise, shifts, rotations)?

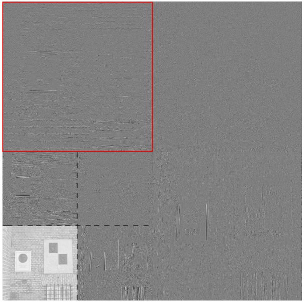

Figure : Discrete wavelet transform of GT

Oriented, $\pm$ separable

To tackle orthogonal DWT limitations

- 1D, orthogonality, realness, symmetry, finite support (Haar)

Approaches used for simple designs (& more involved as well)

- relaxing properties: IIR, biorthogonal, complex
- M-adic MRAs with M integer > 2 or $M = p/q$
- hyperbolic, alternative tilings, less isotropic decompositions
- with pyramidal-scheme: steerable Marr-like pyramids
- relaxing critical sampling with oversampled filter banks
- complexity: (fractional/directional) Hilbert, Riesz, phaselets, monogenic, hypercomplex, quaternions, Clifford algebras

Oriented, $\pm$ separable

Illustration of a combination of Hilbert pairs and M-band MRA

$$\widehat{\mathcal{H}\{f\}}(\omega) = -\imath\,\mathrm{sign}(\omega)\widehat{f}(\omega)$$

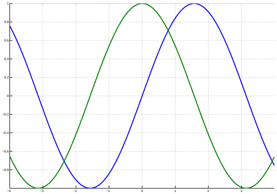

Figure : Hilbert pair 1

Oriented, $\pm$ separable

Illustration of a combination of Hilbert pairs and M-band MRA

$$\widehat{\mathcal{H}\{f\}}(\omega) = -\imath\,\mathrm{sign}(\omega)\widehat{f}(\omega)$$

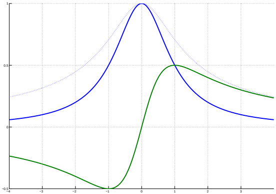

Figure : Hilbert pair 2

Oriented, $\pm$ separable

Illustration of a combination of Hilbert pairs and M-band MRA

$$\widehat{\mathcal{H}\{f\}}(\omega) = -\imath\, \mathrm{sign}(\omega) \widehat{f}(\omega)$$

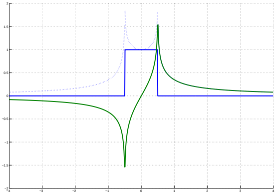

Figure : Hilbert pair 3

Oriented, ± separable

Illustration of a combination of Hilbert pairs and M-band MRA

$$\widehat{\mathcal{H}\{f\}}(\omega) = -\imath\,\mathrm{sign}(\omega)\widehat{f}(\omega)$$

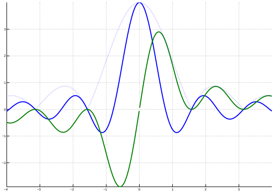

Figure : Hilbert pair 4

Oriented, $\pm$ separable

Illustration of a combination of Hilbert pairs and M-band MRA

$$\widehat{\mathcal{H}\{f\}}(\omega) = -\imath\,\mathrm{sign}(\omega)\widehat{f}(\omega)$$

Compute two wavelet trees in parallel, wavelets forming Hilbert pairs, and combine, either with standard 2-band or 4-band

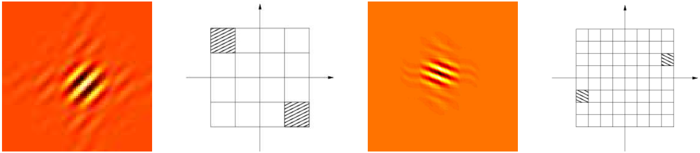

Figure : Dual-tree wavelet atoms and frequency partinioning

Oriented, ± separable

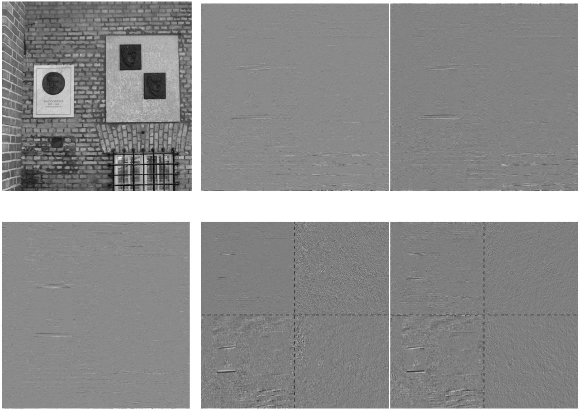

Figure : GT for horizontal subband(s): dyadic, 2-band and 4-band DTT

Oriented, $\pm$ separable

Figure : GT (reminder)

Oriented, ± separable

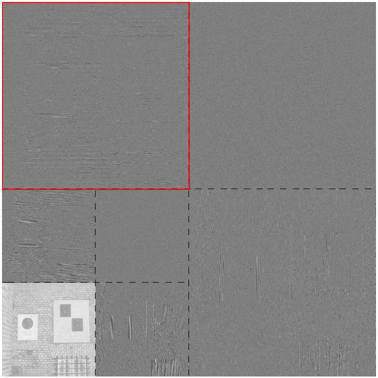

Figure : GT for horizontal subband(s) (reminder)

Oriented, ± separable

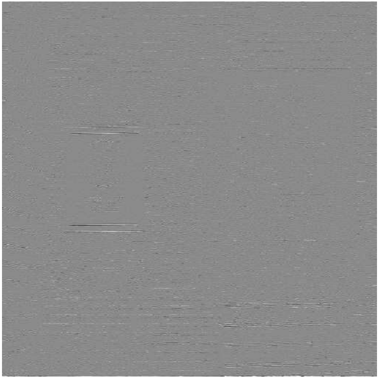

Figure : GT for horizontal subband(s): 2-band, real-valued wavelet

Oriented, ± separable

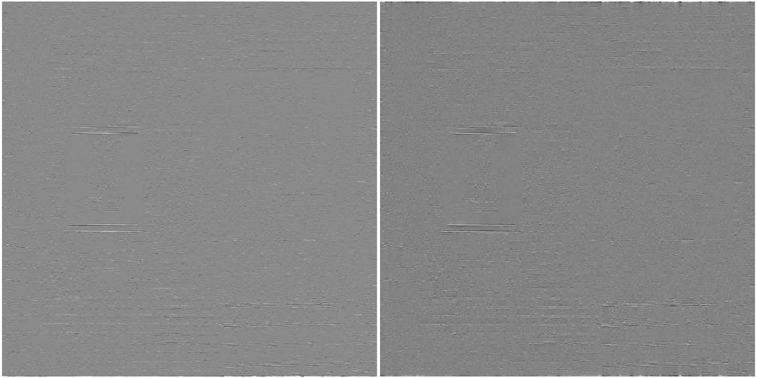

Figure : GT for horizontal subband(s): 2-band dual-tree wavelet

Oriented, ± separable

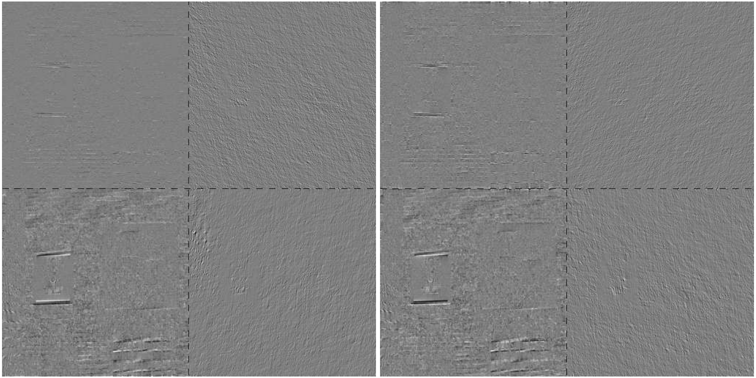

Figure : GT for horizontal subband(s): 4-band dual-tree wavelet

Directional, non-separable

Non-separable decomposition schemes, directly n-D

- non-diagonal subsampling operators & windows
- non-rectangular lattices (quincunx, skewed)
- non-MRA directional filter banks
- steerable pyramids
- M-band non-redundant directional discrete wavelets
- served as building blocks for:
  - contourlets, surfacelets
  - first generation curvelets with (pseudo-)polar FFT, loglets, directionlets, digital ridgelets, tetrolets

Directional, non-separable

Directional wavelets and frames with actions of rotation or similitude groups

$$\psi_{(\boldsymbol{b},a,\theta)}(\boldsymbol{x}) \;=\; \tfrac{1}{a}\,\psi\big(\tfrac{1}{a}\,R_\theta^{-1}(\boldsymbol{x}-\boldsymbol{b})\big),$$

where R_θ stands for the 2×2 rotation matrix

$$W_f(\boldsymbol{b},a,\theta) \;=\; \langle \psi_{(\boldsymbol{b},a,\theta)}, f\rangle$$

inverted through

$$f(\boldsymbol{x}) \;=\; c_\psi^{-1}\int_0^\infty \frac{\mathrm{d}a}{a^3}\int_0^{2\pi}\mathrm{d}\theta\int_{\mathbb{R}^2}\mathrm{d}^2\boldsymbol{b}\quad W_f(\boldsymbol{b},a,\theta)\;\psi_{(\boldsymbol{b},a,\theta)}(\boldsymbol{x})$$

Directional, non-separable

Directional wavelets and frames:

- possibility to decompose and reconstruct an image from a discretized set of parameters; often (too) isotropic
- examples: Conical-Cauchy wavelet, Morlet-Gabor frames

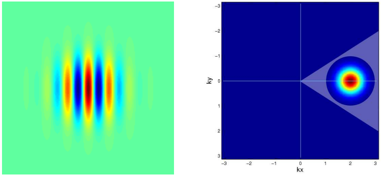

Figure : Morlet Wavelet (real part) and Fourier representation

Directional, anisotropic scaling

Ridgelets: 1-D wavelet and Radon transform $\mathfrak{R}_f(\theta, t)$

$$\mathcal{R}_f(b, a, \theta) = \int \psi_{(\boldsymbol{b},a,\theta)}(\boldsymbol{x})\, f(\boldsymbol{x})\, \mathrm{d}^2\boldsymbol{x} = \int \mathfrak{R}_f(\theta, t)\, a^{-1/2}\psi((t-b)/a)\, \mathrm{d}t$$

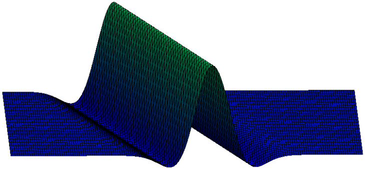

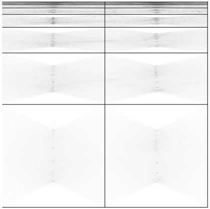

Figure : Ridgelet atom and GT decomposition

Directional, anisotropic scaling

Curvelet transform: continuous and frame

- curvelet atom: scale s, orient. $\theta \in [0, \pi)$, pos. $\boldsymbol{y} \in [0,1]^2$:

$$\psi_{s,\boldsymbol{y},\theta}(\boldsymbol{x}) = \psi_s(R_\theta^{-1}(\boldsymbol{x} - \boldsymbol{y}))$$

  $\psi_s(\boldsymbol{x}) \approx s^{-3/4}\,\psi(s^{-1/2}x_1, s^{-1}x_2)$ parabolic stretch; $(w \simeq \sqrt{l})$
  Near-optimal decay: C^2 in C^2: $O(n^{-2}\log^3 n)$

- tight frame: $\psi_{\boldsymbol{m}}(\boldsymbol{x}) = \psi_{2^j,\theta_\ell,\boldsymbol{x}_n}(\boldsymbol{x})$ where $m = (j, n, \ell)$ with sampling locations:

$$\theta_\ell = \ell\pi 2^{\lfloor j/2\rfloor - 1} \in [0, \pi) \quad \text{and} \quad \boldsymbol{x}_n = R_{\theta_\ell}(2^{j/2}n_1, 2^j n_2) \in [0,1]^2$$

- related transforms: shearlets, type-I ripplets

Directional, anisotropic scaling

Curvelet transform: continuous and frame

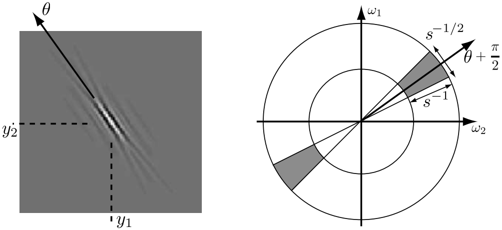

Figure : A curvelet atom and the wegde-like frequency support

Directional, anisotropic scaling

Curvelet transform: continuous and frame

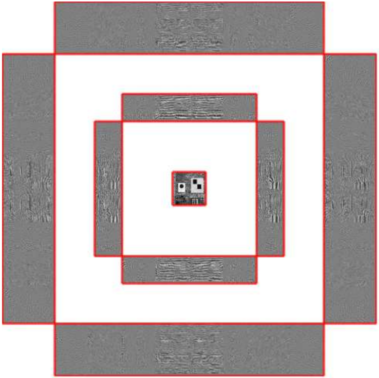

Figure : GT curvelet decomposition

Directional, anisotropic scaling

Contourlets: Laplacian pyramid + directional FB

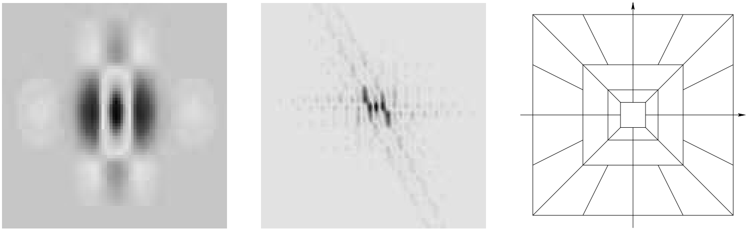

Figure : Contourlet atom and frequency tiling

from close to critical to highly oversampled

Directional, anisotropic scaling

Contourlets: Laplacian pyramid + directional FB

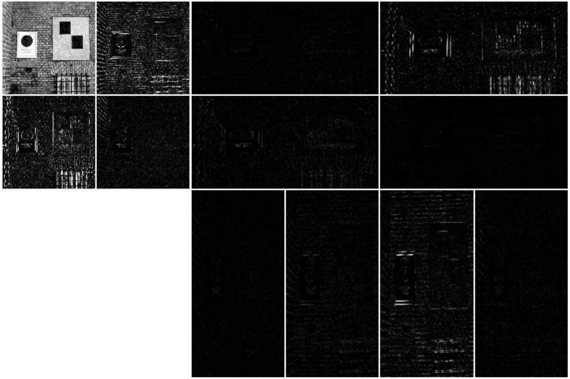

Figure : Contourlet GT (flexible) decomposition

Directional, anisotropic scaling

Shearlets

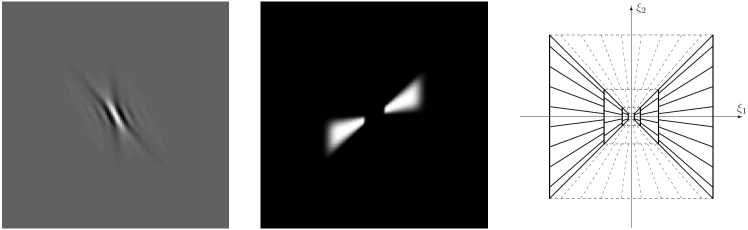

Figure : Shearlet atom in space and frequency, and frequency tiling

Do they have it all?

Directional, anisotropic scaling

Additional transforms

- previously mentioned transforms are better suited for edge representation
- oscillating textures may require more appropriate transforms
- examples:
  - wavelet and local cosine packets
  - best packets in Gabor frames
  - brushlets [Meyer, 1997; Borup, 2005]
  - wave atoms [Demanet, 2007]

Lifting representations

Lifting scheme is an unifying framework

- to design adaptive biorthogonal wavelets
- use of spatially varying local interpolations
- at each scale j, a_{j-1} are split into a_j^o and d_j^o
- wavelet coefficients d_j and coarse scale coefficients a_j: apply (linear) operators $P_j^{\lambda_j}$ and $U_j^{\lambda_j}$ parameterized by λ_j

$$d_j = d_j^o - P_j^{\lambda_j} a_j^o \quad \text{and} \quad a_j = a_j^o + U_j^{\lambda_j} d_j$$

It also

- guarantees perfect reconstruction for arbitrary filters
- adapts to non-linear filters, morphological operations
- can be used on non-translation invariant grids to build wavelets on surfaces

Lifting representations

$$d_j = d_j^o - P_j^{\lambda_j} a_j^o \quad \text{and} \quad a_j = a_j^o + U_j^{\lambda_j} d_j$$

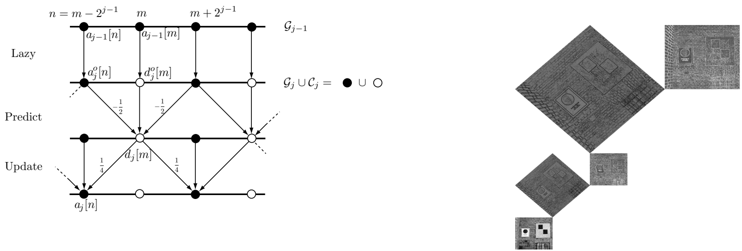

Figure : Predict and update lifting steps and MaxMin lifting of GT

Lifting representations

Extensions and related works

- adaptive predictions:
    - possibility to design the set of parameter $\lambda = \{\lambda_j\}_j$ to adapt the transform to the geometry of the image
    - λ_j is called an association field, since it links a coefficient of a_j^o to a few neighboring coefficients in d_j^o
    - each association is optimized to reduce the magnitude of wavelet coefficients d_j, and should thus follow the geometric structures in the image
    - may shorten wavelet filters near the edges
- grouplets: association fields combined to maintain orthogonality

One result among many others

Context: multivariate Stein-based denoining of a multi-spectral satellite image

Different spectral bands

One result among many others

Context: multivariate Stein-based denoining of a multi-spectral satellite image

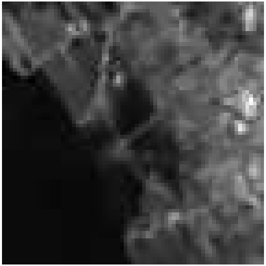

Form left to right: original, noisy, denoised

One result among many others

Context: multivariate Stein-based denoining of a multi-spectral satellite image

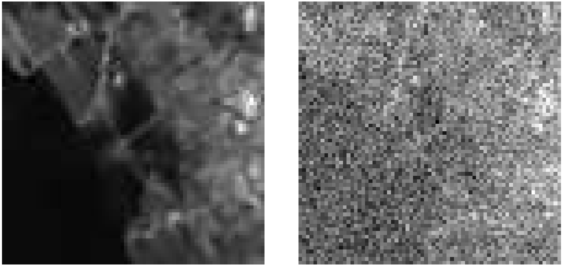

Form left to right: original, noisy, denoised

One result among many others

Context: multivariate Stein-based denoining of a multi-spectral satellite image

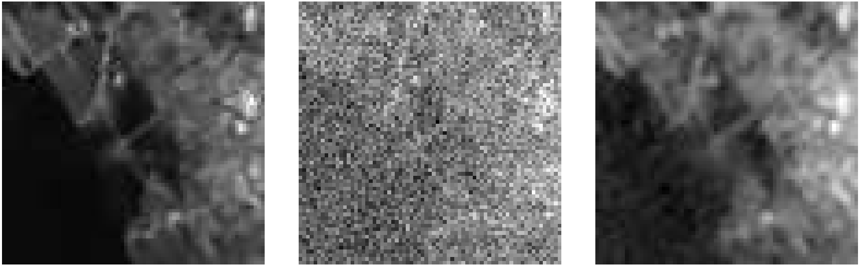

Form left to right: original, noisy, denoised

What else? Images are not (all) flat

Many multiscale designs have been transported, adapted to:

- meshes
- spheres
- two-sheeted hyperboloid and paraboloid
- 2-manifolds (case dependent)
- big deal: data on graphs

see 300+ reference list!

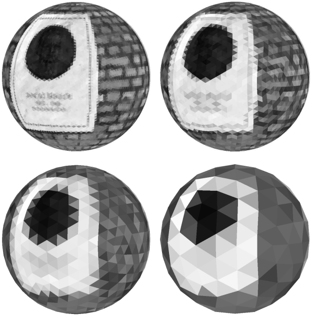

Conclusion: on a (frustrating) panorama

Take-away messages anyway?

If you only have a hammer, every problem looks like a nail

- Is there a "best" geometric and multiscale transform?
  - no: intricate data/transform/processing relationships
    - more needed on asymptotics, optimization, models
  - maybe: many candidates, progresses awaited:
    - "so ℓ_2"! Low-rank (ℓ_0/ℓ_1), math. morph. ($+, \times$ vs $\max, +$)
  - yes: those you handle best, or (my) on wishlist
    - mild redundancy, invariance, manageable correlation, fast decay, tunable frequency decomposition, complex or more

Conclusion: on a (frustrating) panorama

Postponed references & toolboxes

- A Panorama on Multiscale Geometric Representations, Intertwining Spatial, Directional and Frequency Selectivity
  Signal Processing, Dec. 2011

  Toolboxes, images, and names
  `http://www.sciencedirect.com/science/article/pii/S0165168411001356`
  `http://www.laurent-duval.eu/siva-panorama-multiscale-geometric-representations.html`
  `http://www.laurent-duval.eu/siva-wits-where-is-the-starlet.html`

Acknowledgments to:

- the many *-lets (last pick: the Gabor shearlet)