Ondelettes et autres représentations bidimensionnelles, multi-échelles et géométriques : revue thématique

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

IFP Énergies nouvelles

2012/04/12

Laboratoire de Physique, ENS Lyon

In just one slide

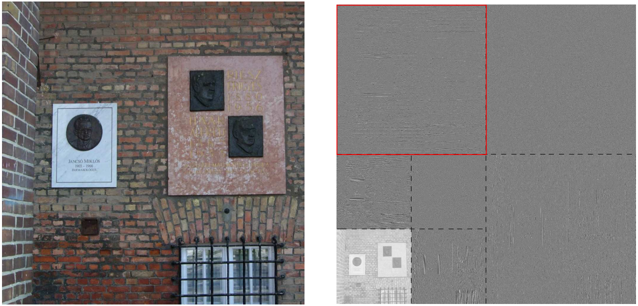

Figure: A standard separable wavelet decomposition

Where do we go from here?

Images are pixels (but...):

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

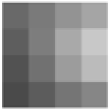

Figure: Image as a linear combination of pixels

- suffices for (simple) data (simple) manipulation
  - counting, enhancement, filtering
- very limited in higher level understanding tasks

Images are pixels (but...):

A review in an active research field:

- (partly) inspired by:
  - early vision observations
  - 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
- aimed either at pragmatic or heuristic purposes
  - known formation model *or* unknown information
- developed through a quantity of *-lets and relatives

Images are pixels, but resolution (scale?) matters

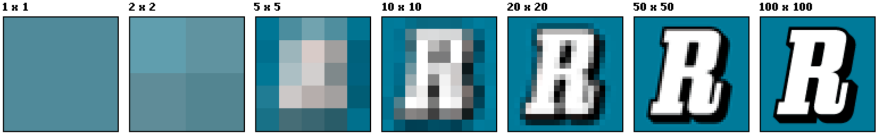

Figure: RRRrrrr: coarse to fine! [Chabat *et al.*, 2004]

- notion of sufficient resolution for understanding
- coarse-to-fine and fine-to-coarse links
- impact on more complex images?

Images are different (but...)

Figure: Different kinds of images

Images are different (but...)

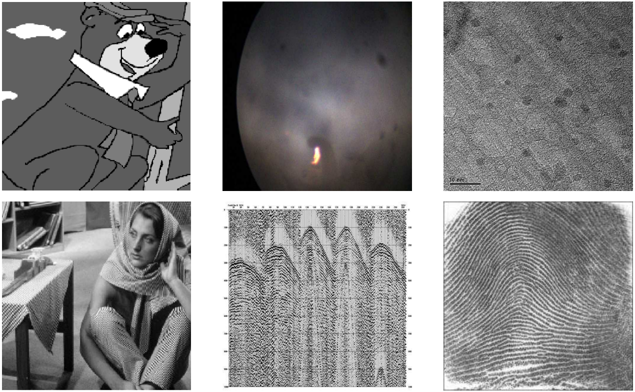

Figure: Different kinds of images

Images are different, 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]

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

- piecewise-smooth + contours + geometrical textures + unmodeled

Images are different, but might be described by models

Remember: contours and textures are usually not very well-defined

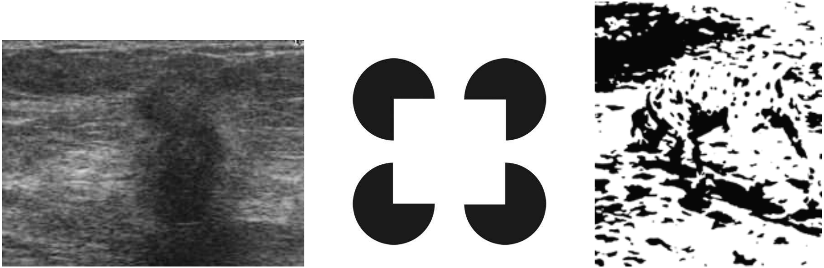

Figure: Real world image and illusion

Images are different, but might be described by models

Remember: contours and textures are usually not very well-defined

- a variety of methods for description/detection/modeling: smooth curve or polynomial fit, oriented regularized derivatives, discrete geometry, parametric curve detectors (such as the Hough transform), mathematical morphology, EMD, local *frequency estimators*, optical flow approaches, smoothed random models, etc.

Images are different, but might be described by models

Efficient descriptions needed...

- for: compression, denoising, enhancement, inpainting, restoration, fusion, super-resolution, registration, segmentation, reconstruction, source separation, image decomposition, multiple description coding, sparse sampling, etc.
- in a word: understanding (for dedicated purposes)
- first: a brief review of early approaches: from orthogonality (Fourier) to redundancy, and back (DWT)

Images are different, but might be described by models

Efficient descriptions needed... using an assumption of sparsity

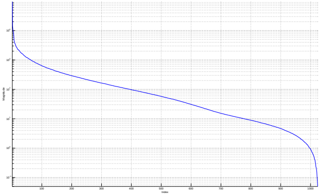

Figure: Another real world image with genuine contours and textures and its singular values

Images are different: a guiding thread

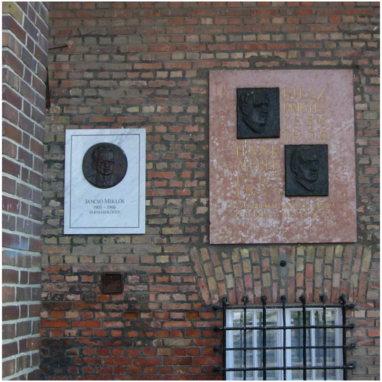

Figure: Memorial plaque in honor of A. Haar and F. Riesz: *A szegedi matematikai iskola világhírű megalapítói*, court. Prof. K. Szatmáry

Guiding thread (GT): early days

Fourier approach: critical, orthogonal

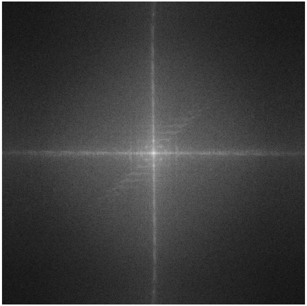

Figure: GT luminance component amplitude spectrum (log-scale)

Fast, compact, very practical but not quite informative

Guiding thread (GT): early days

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

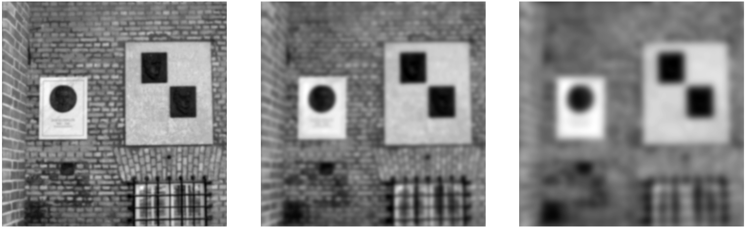

Figure: GT with Gaussian scale-space decomposition

Gaussian filters and heat diffusion interpretation
Varying persistence of features across scales

Guiding thread (GT): early days

Differences in scale-space with subsampling

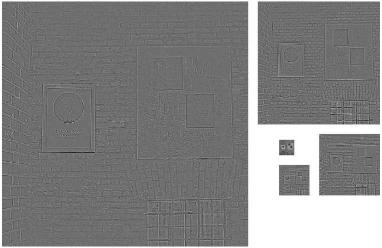

Figure: GT with Laplacian pyramid decomposition

Laplacian pyramid: complete, reduced redundancy, enhances image singularities, low activity regions with small coefficients

Guiding thread (GT): early days

Isotropic wavelets

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

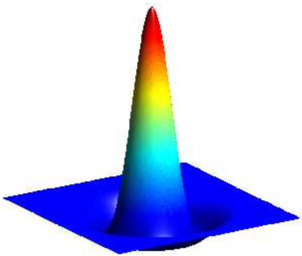

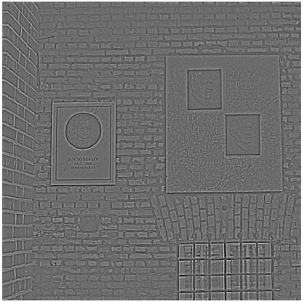

Figure: Example: Marr wavelet as a singularity detector

Multiscale edge detector, more potential wavelet shapes (DoG, Cauchy, etc.)

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. 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

DWT, back to orthogonality: fast, compact and informative, but...

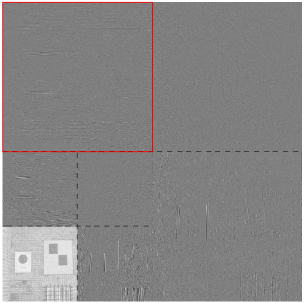

Figure: Separable wavelet decomposition of GT

is it sufficient (singularities, noise, shifts, rotations)?

Agenda

- A quindecennial panorama of improvements ($\geqslant$1998)
  - sparser representations of contours and textures through increased spatial, directional and frequency selectivity
  - from fixed to adaptive, from low to high redundancy
  - generally fast, compact, informative, practical
  - bits of hybridization and small paths
  - only a "subsampling", references postponed to the end
- Outline
  - oriented & geometrical: outcrops, non-separable directionality, anisotropic scaling
  - redundant & adaptive: pursuits, trees and lifting
  - non-Euclidian geometries
  - conclusions
  - local applications

Outcrops from 1-D separable representations

To tackle orthogonal DWT limitations

- 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 $M = p/q$
- 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

Outcrops from 1-D separable representations

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

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

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

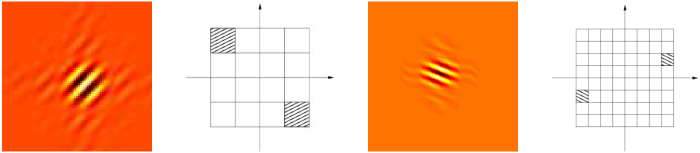

Figure: Examples of atoms and associated frequency partinioning

Outcrops from 1-D separable representations

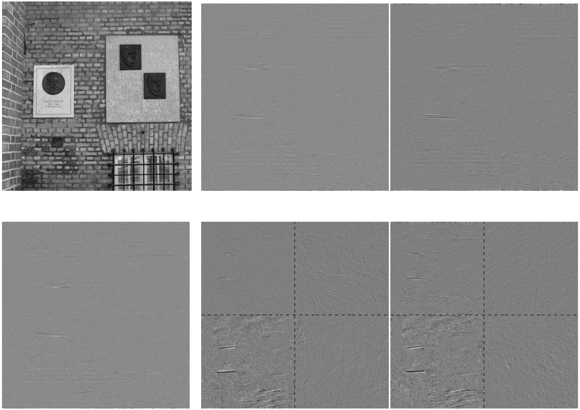

Figure: GT for horizontal subband(s): dyadic, 2-band and 4-band DTT

Non-separable directionality

Non-separable decomposition schemes, directly *n*-D

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

Non-separable directionality

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

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

Non-separable directionality

Directional wavelets and frames:

- possibility to decompose and reconstruct an image from a discretized set of parameters
- examples: Conic-Cauchy wavelet, Morlet/Gabor frames

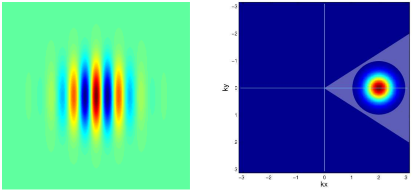

Figure: Morlet Wavelet (real part) and Fourier representation

Anisotropic scaling

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

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

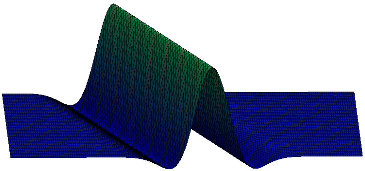

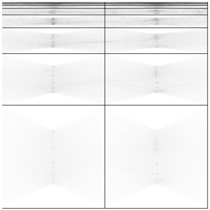

Figure: Ridgelet atom and GT decomposition

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; $(w \simeq \sqrt{l})$
  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

Anisotropic scaling

Curvelet transform: continuous and frame

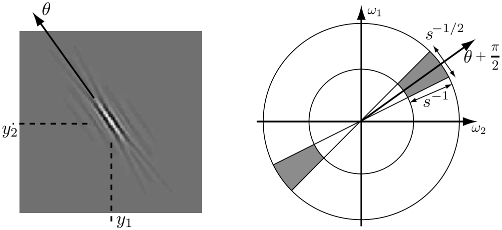

Figure: A curvelet atom and the wegde-like frequency support

Anisotropic scaling

Curvelet transform: continuous and frame

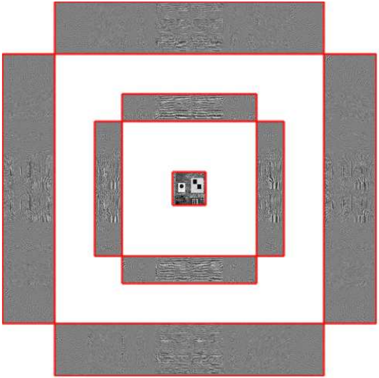

Figure: GT curvelet decomposition

Anisotropic scaling

Contourlets: Laplacian pyramid + directional FB

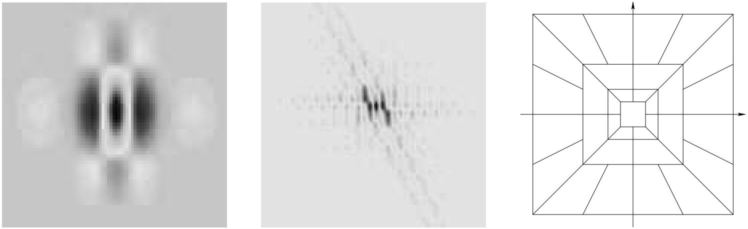

Figure: Contourlet atom and frequency tiling

Anisotropic scaling

Contourlets: Laplacian pyramid + directional FB

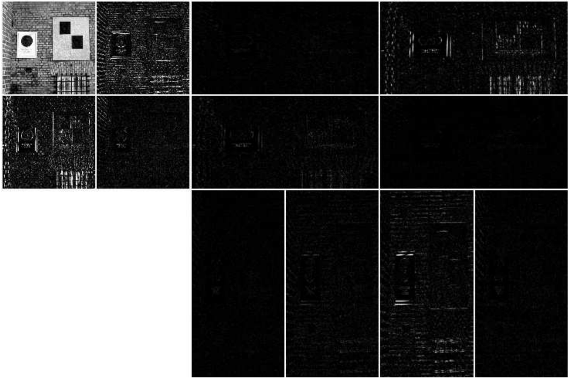

Figure: Contourlet GT (flexible) decomposition

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]

Pursuit in redundant dictionaries

Highly redundant representations improve the representation of complicated images (with edges and textures):
Approx. f_M of f with M atoms from $\mathcal{B} = \{\psi_{\boldsymbol{m}_j} : 1 \leqslant j \leqslant P\}$:

$$f_M = \Psi a = \sum_j a_j \psi_{\boldsymbol{m}_j}, \quad \text{with} \quad \|a\|_0 = \#\{j \,:\, a_j \neq 0\} \leqslant M.$$

is generally NP-hard.

- matching pursuit: greedy
- basis pursuit: sparse approximation with

$$f_M = \Psi a = \sum_j a_j \psi_{\boldsymbol{m}_j}, \; a \in \operatorname*{argmin}_{\tilde{a} \in \mathbb{R}^P} \tfrac{1}{2} \|f - \sum_j \tilde{a}_j \psi_{\boldsymbol{m}_j}\|^2 + \mu \|\tilde{a}\|_1,$$

Pursuit in redundant dictionaries

Parametric dictionaries are obtained from basic operations (like rotation, translation, dilation, shearing, modulation, etc.) applied to a continuous mother function.

- pursuits in parametric dictionaries: given a set of S transformations $T^i_{m_i}$ for $1 \leqslant i \leqslant S$ parameterized by $m_i \in \Lambda_i \subset \mathbb{R}^{n_i}$, the parametric dictionary is related to a certain discretization of $\Lambda^{\mathrm{d}} \subset \Lambda = \Lambda_1 \times \cdots \times \Lambda_S$, i.e.

$$\mathcal{B} \;=\; \{\psi_{\boldsymbol{m}}(\boldsymbol{x}) = [T^1_{m_1} \cdots T^S_{m_S} \psi](\boldsymbol{x}) \in \mathbb{L}^2(\mathbb{R}^2)$$

$$\boldsymbol{m} = (m_1, \cdots, m_S) \in \Lambda^{\mathrm{d}}\}$$

- dictionary discretization may be refined during MP iterations

Tree-structured best basis representations

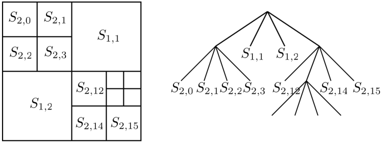

Figure: Dyadic subdivision of $[0,1]^2$ in squares $S_{j,i}$ and corresponding quad-tree λ

- quadtree-based dictionaries
- best basis selection
- adaptive approximation
  - wedgelets, platelets, bandlets
  - adaptive approximations over the wavelet domain: wedgeprints, 2nd gen. bandlets, etc.

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

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

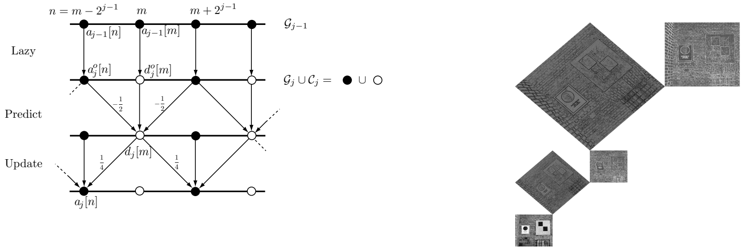

Figure: Predict and update lifting steps and MaxMin lifting of GT

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

Images are not (all) flat

On the sphere, motivated by astronomy and geosciences

- filtering (two-angle spherical param. $\boldsymbol{\xi} = (\theta, \varphi) \in S^2$)

$$(f \star h)(\boldsymbol{\rho}) \; = \int_{S^2} h(\boldsymbol{\rho}\,\boldsymbol{\xi}) f(\boldsymbol{\xi}) \,\mathrm{d}\mu(\boldsymbol{\xi}),$$

  where $\boldsymbol{\rho} \in \mathrm{SO}(3)$ is a rotation (driven by three angles) applied to the point $\boldsymbol{\xi} \in S^2$ and $\mathrm{d}\mu(\boldsymbol{\xi}) = \sin\theta \mathrm{d}\theta \mathrm{d}\varphi$.
- Fourier transform

$$\hat{f}_{\ell m} = \langle Y_{\ell m}, f \rangle = \int_{S^2} Y^*_{\ell m}(\boldsymbol{\xi})\, f(\boldsymbol{\xi}) \,\mathrm{d}\mu(\boldsymbol{\xi}), \quad f(\boldsymbol{\xi}) = \sum_{\ell,m} \hat{f}_{\ell m}\, Y_{\ell m}(\boldsymbol{\xi})$$

  with respect to orthonormal basis of *spherical harmonics* $\mathcal{Y} = \{Y_{\ell m}(\boldsymbol{\xi}) : \ell \geqslant 0, |m| \leqslant \ell\}$, i.e. the eigenvectors of the spherical Laplacian

Images are not (all) flat

- spherical Scale-Space: solution at at time $\tau > 0$ initialized to f

$$f(\xi, \tau) = \sum_{\ell, m} \hat{f}_{\ell m}(\tau) Y_{\ell m}(\xi)$$

with $\hat{f}_{\ell m}(\tau) = \hat{f}_{\ell m}\, e^{-\ell(\ell+1)\,\tau}$ and $f(\xi, 0) = f(\xi)$

- spectral Wavelets
  - spherical wavelets $\psi_a(\xi)$, $W_f(a, \xi) = (f * \psi_a)(\xi)$

$$f(\xi') \;=\; \langle f \rangle \;+\; \int_{\mathbb{R}_+} \int_{S^2} \; W_f(a, \xi)\, \psi_a(\xi' \cdot \xi) \; \frac{\mathrm{d}a}{a}\, \mathrm{d}\xi$$

  with $\langle f \rangle = \frac{1}{4\pi} \int_{S^2} f(\xi)\, \mathrm{d}\mu(\xi)$
  - MRA on the sphere with QMF
  - needlets

Images are not (all) flat

Some of the other (many) constructions

- stereographic wavelets
- Haar transforms (embedded spherical triangulations)

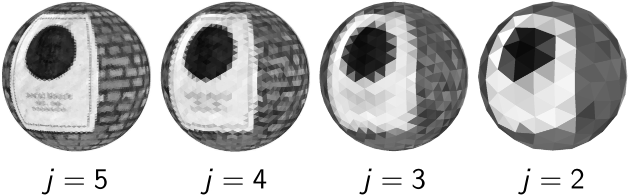

Figure: Projection on the spherical Haar multiresolution

- steerable wavelets the sphere
- HEALPix: trans. wavelets, ridgelets, curvelets on the sphere
- SOHO wavelets

Images are not (all) flat

Some of the other (many) constructions

- wavelets on general 2-manifolds
- lifting scheme on meshed surfaces

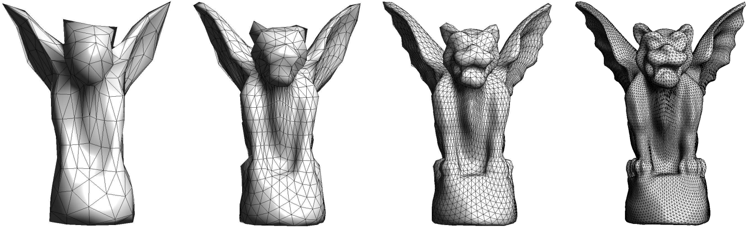

Figure: Multiresolution semi-regular mesh

Images are not (all) flat

Some of the other (many) constructions

- wavelets on general 2-manifolds
- lifting scheme on meshed surfaces

Figure: Surface approximation by thresholding (100%, 10%, 5%, 2%)

Images are not images anymore

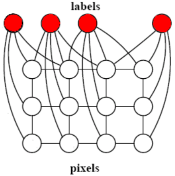

Figure: Image as a graph

Pixels in images may be viewed as *loci*:

- with connectivity information (even non-local)
- scalar, vector-valued
- labeled

Images are not images anymore

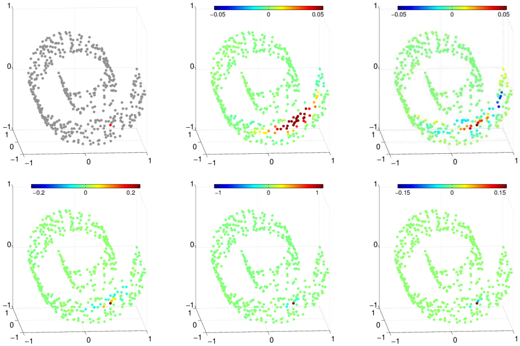

Figure: Wavelets on a swiss roll

Wavelets on graphs [Hammond *et al.*, 2011]

Conclusion: on a panorama

A panorama of 2-D images to a 3-D scene is not seamless:

- *If you only have a hammer, every problem looks like a nail*
  - no more!
- *The map is not the territory*: incomplete panorama
- from mild hybridization to GMO-let?
  - awaited progresses on asymptotics, optimization, models
  - l_0 not practical, toward structured sparsity, means efficient?
  - learned frames, robust low-rank approximations, l_0-PCA!

Conclusion: on a panorama

Take-away messages anyway?

- is there a best?
  - even worse than with (discrete) wavelets
  - intricate relationship: sparsifying transform/associated processing
- wishlist: fast, mild redundancy, some invariance, manageable correlation in transformed domain, fast decay, tunable freq. split, complex or more (phase issues)

Conclusion: on a panorama

Acknowledgments:

- L. Jacques, C. Chaux, G. Peyré
- to the many *-lets, esp. the forgotten ones

Postponed references & toolboxes & links

- A Panorama on Multiscale Geometric Representations, Intertwining Spatial, Directional and Frequency Selectivity, Signal Processing, Dec. 2011
  http://www.sciencedirect.com/science/article/pii/S0165168411001356
- more stuff:
  http://www.laurent-duval.eu/siva-panorama-multiscale-geometric-representations.html
  http://www.laurent-duval.eu/siva-wits-where-is-the-starlet.html

Local applications

Weak interest in Lenna-like images

Figure: Seismic Lenna

Local applications

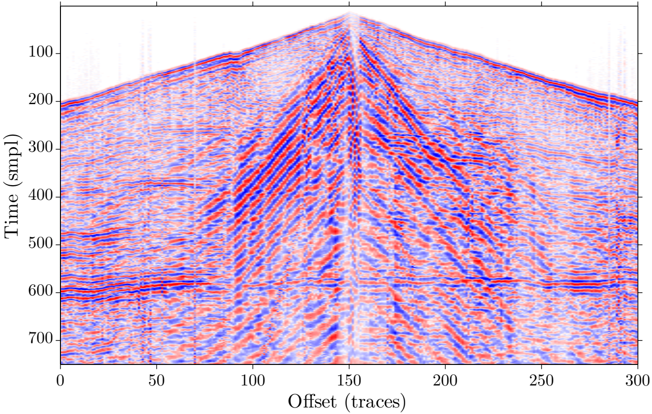

Figure: Surface wave removal (before)

Local applications

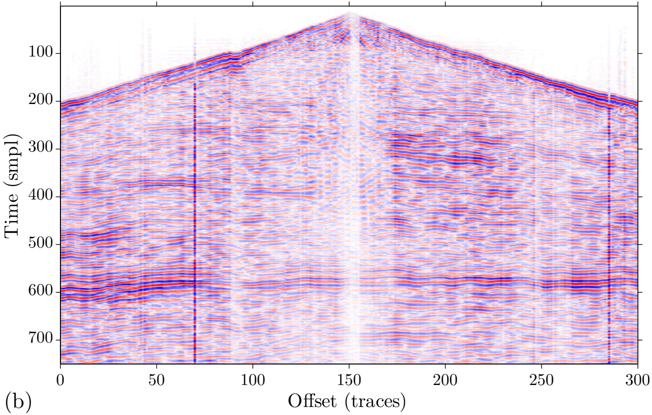

Figure: Surface wave removal (after)

Local applications

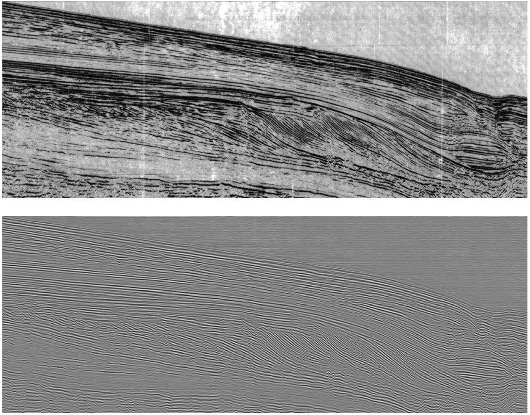

Figure: Paper to seismics

Local applications

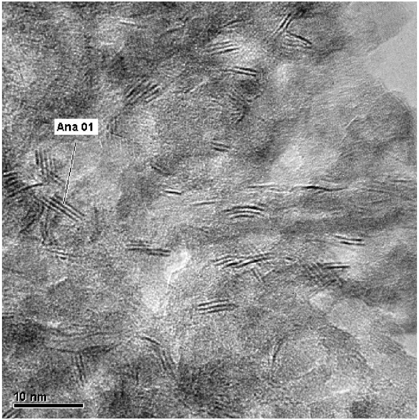

Figure: Catalyst measurements

Local applications

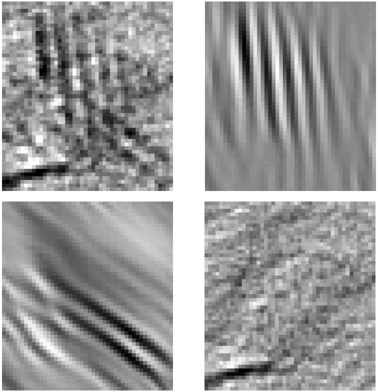

Figure: Catalyst measurements