Applications and Outlook

Minh N. Do

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

www.ifp.uiuc.edu/∼minhdo
minhdo@uiuc.edu

Joint work with Arthur Cunha, Yue Lu, Jianping Zhou (UIUC), and Duncan Po (MathWorks)

Outline

1. Image modeling using the contourlet transform
2. Critically sampled (CRISP) contourlet transform
3. Image denoising and enhancement using the nonsubsampled contourlet transform
4. Outlook

1. Image Modeling using the Contourlet Transform

Contourlet Coefficient Relationships

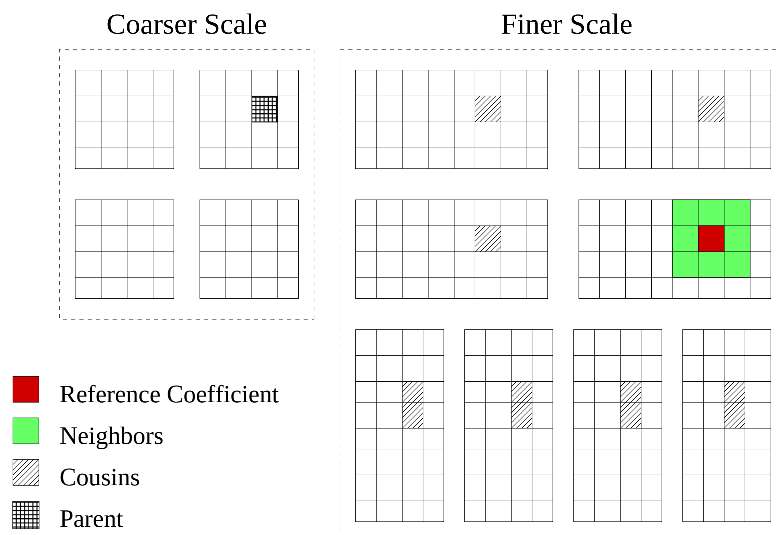

Marginal Statistics of Contourlet Coefficients

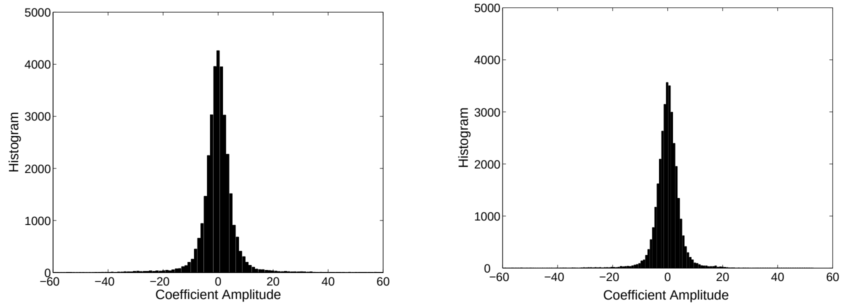

Marginal statistics of two finest subbands of the image "Peppers."

The kurtoses of the two distributions are measured at 24.50 and 19.40, showing that the coefficients are highly non-Gaussian.

Joint Statistics

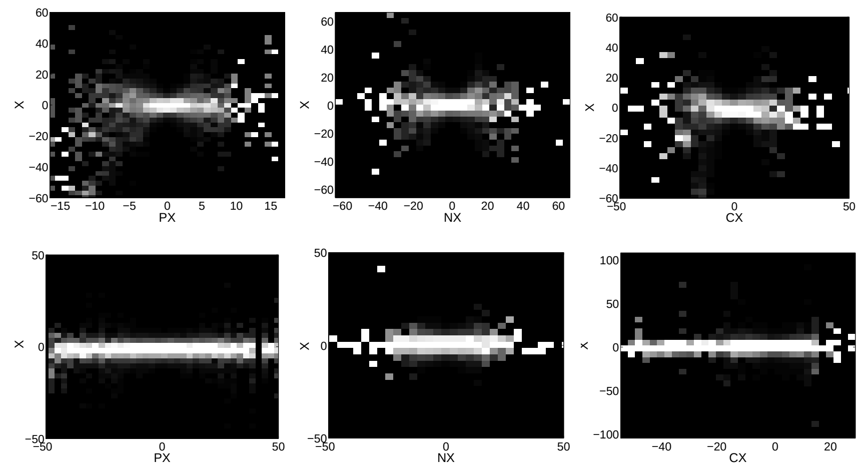

Top: Joint histograms of next neighbors
Bottom: Joint histograms of distance (three coefficients away) neighbors

Conditional Distribution

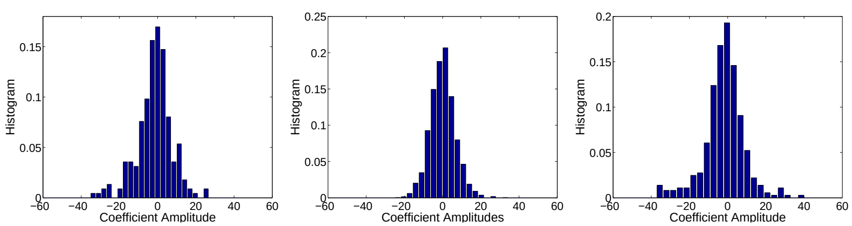

The kurtoses of the distributions are measured at 3.90, 2.90, and 2.99.

$\Rightarrow$ Contourlet coefficients are non-Gaussian but conditionally Gaussian.

Dependence Characterization using Mutual Information

Mutual information estimates between a contourlet coefficient (X) and its parent (PX), its spatial neighbors (NX), and its directional cousins (CX).

	Lena	Barbara	Peppers
$I(X;PX)$	0.11	0.14	0.10
$I(X;NX)$	0.23	0.58	0.17
$I(X;CX)$	0.19	0.39	0.14

Mutual information estimates with a **single** parent, neighbor, and cousin.

	Lena	Barbara	Peppers
$I(X;PX)$	0.11	0.14	0.08
$I(X;NX_1)$	0.09	0.31	0.07
$I(X;NX_2)$	0.07	0.27	0.05
$I(X;CX_1)$	0.08	0.20	0.06
$I(X;CX_2)$	0.06	0.17	0.05
$I(X;CX_3)$	0.06	0.20	0.04

Wavelet-domain Hidden Markov Models [CrouseNB:98]

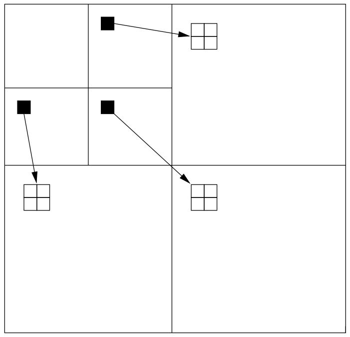

Wavelet HMT models coefficients on each direction independently.

Contourlet-domain Hidden Markov Tree Models [PoD:04]

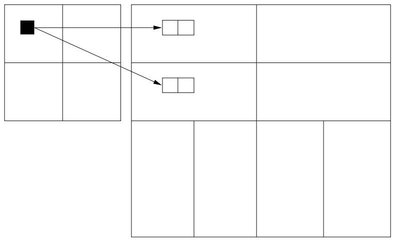

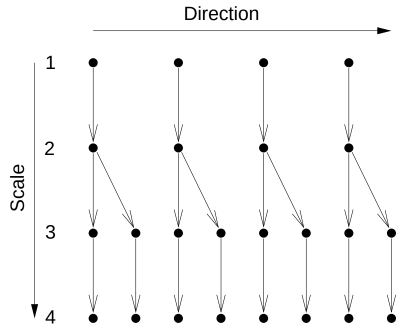

Contourlet HMT models all inter-scale, inter-direction, and inter-location independencies.

Denoising Results: *Zelda*

Left to right, top to bottom: (a) "Zelda" image, (b) noisy image ($14.61dB$), (c) wiener2 ($25.78dB$), (d) wavelet thresholding ($26.05dB$), (e) wavelet HMT ($27.63dB$), and (f) contourlet HMT ($27.07dB$).

Texture Retrieval Results: *Brodatz Database*

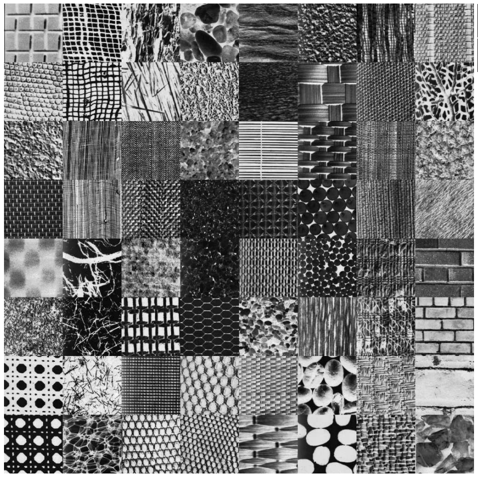

Average retrieval rates

wavelet HMT	contourlet HMT
90.87%	**93.29%**

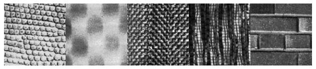

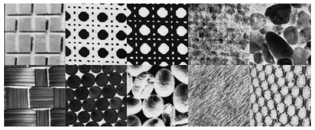

Top: Wavelets do better ($> 5\%$).
Bottom: Contourlets do better ($> 5\%$).

Outline

1. Image modeling using the contourlet transform
2. Critically sampled (CRISP) contourlet transform
3. Image denoising and enhancement using the nonsubsampled contourlet transform
4. Outlook

Frequency Partition of CRISP-Contourlets [LuD:03]

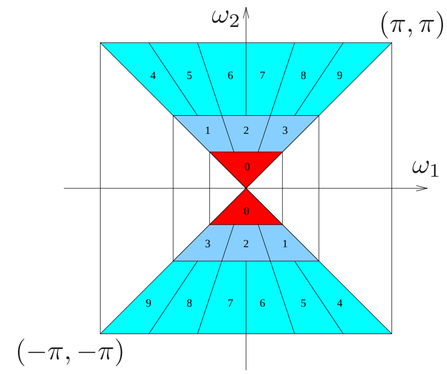

directional highpass: 3×2^n $(n = 1, 2, \ldots)$ directions at each level.
lowpass bands: 2 directional lowpass bands.

Intuition Behind the Frequency Partitioning

Proposition 1. *For a spectrum support $\mathcal{X}_{\mathcal{F}}$ to be critically sampled by a matrix $\boldsymbol{M}$, the area of the support must be $\frac{4\pi^2}{|\det(\boldsymbol{M})|}$.*

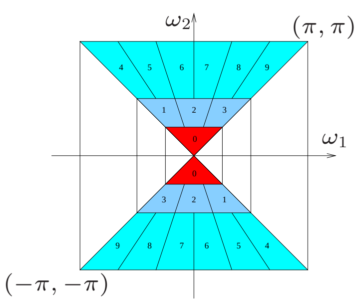

For each directional bandpass region

$$area = 4\pi^2/(4^m \cdot 2^n)$$

Proposition 2. *With $3 \cdot 2^l$ directions, the spectrum support R_k^l of bandpass directional subbands of the CRISP-contourlet transform is critically sampled by $\boldsymbol{S}_k^l = \mathrm{diag}(2^l, 4)$ (or $\mathrm{diag}(4, 2^l)$):*

$$\sum_{\boldsymbol{m} \in \mathbb{Z}^2} \mathbf{1}_{R_k^l}(\boldsymbol{\omega} - 2\pi(\boldsymbol{S}_k^l)^{-T}\boldsymbol{m}) = 1.$$

Sampling Matrices for CRISP-Contourlets

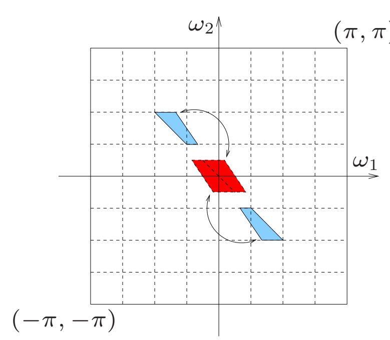

- Each directional bandpass pair can be shifted and combined to form a parallelogram support.
- The parallelogram support can be maximally decimated.
- If the shifts satisfy certain conditions, the split and shifted support can also be maximally decimated.

The decimation matrices (for $n > 1$) in CRISP-contourlets are diagonal:

$$\boldsymbol{M}_v^{(m,n)} = \begin{pmatrix} 2^{m+n+2} & 0 \\ 0 & 2^{m+2} \end{pmatrix} \quad \text{and} \quad \boldsymbol{M}_h^{(m,n)} = \begin{pmatrix} 2^{m+2} & 0 \\ 0 & 2^{m+n+2} \end{pmatrix}$$

CRISP-Contourlet Transform – Block Diagram

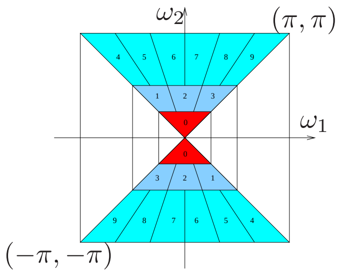

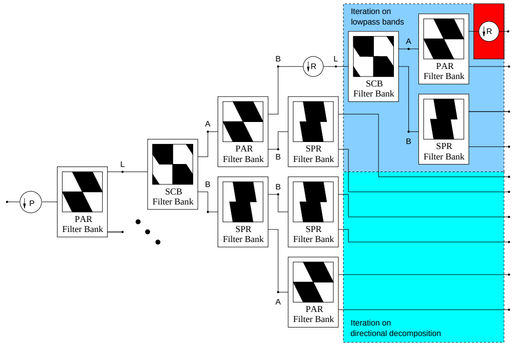

Arbitrary multiscale and multidirectional decomposition through an iterated combination of 3 filter banks:

- Type "L" outputs are followed by "SCB" filters;
- Type "A" outputs are followed by "PAR" filters.
- Type "B" outputs are followed by "SPR" filters.

Iteration for Finer Directionality

- The number of directions can be 3×2^n, for all $n \geq 1$.
- The refinement of directionality is achieved via an iteration of two filter banks.

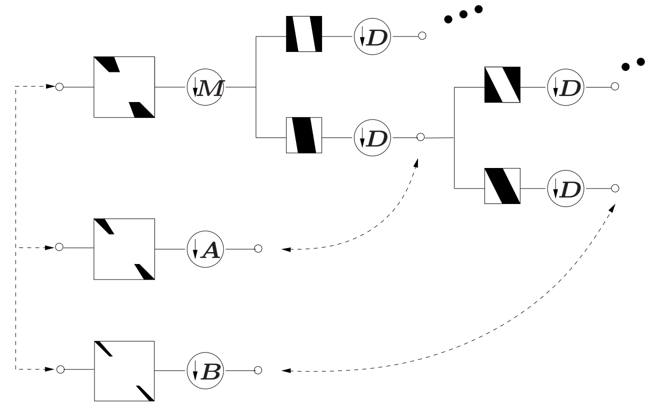

CRISP-Contourlet Transform using Nonuniform Filter Banks

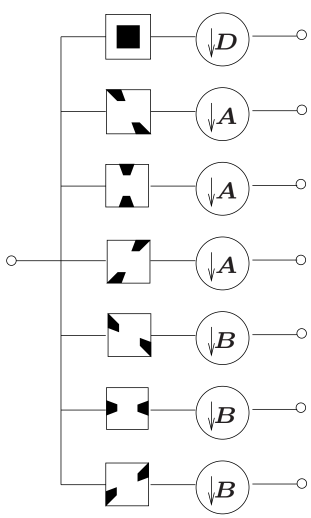

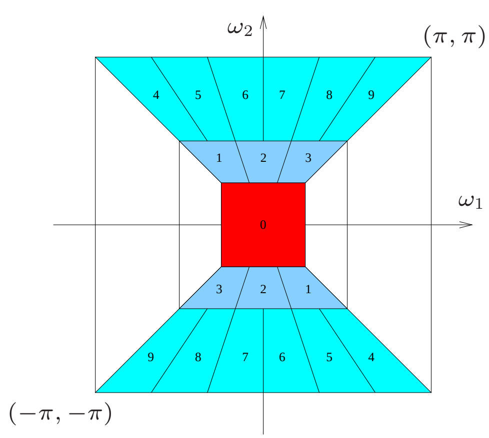

Outline

1. Image modeling using the contourlet transform
2. Critically sampled (CRISP) contourlet transform
3. Image denoising and enhancement using the nonsubsampled contourlet transform
4. Outlook

Motivation

For image-analysis applications (denoising, enhancement, ...), we want:

- Directional multiresolution shift-invariant image representation
- Structured transforms with fast algorithms
- Critical sampling is not critical

Answer: Nonsubsampled Contourlet Transform (NSCT)

Critically-Sampled Filter Banks

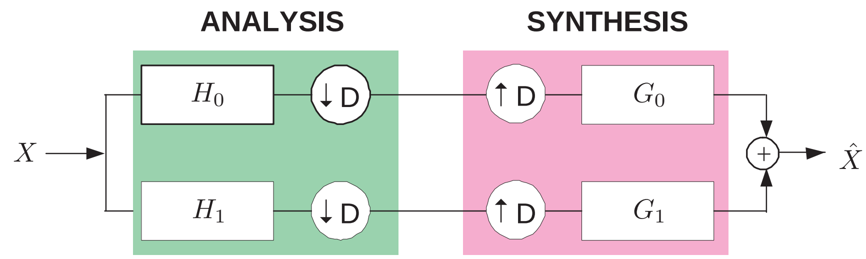

Problem: Shift-variant

Nonsubsampled Filter Banks

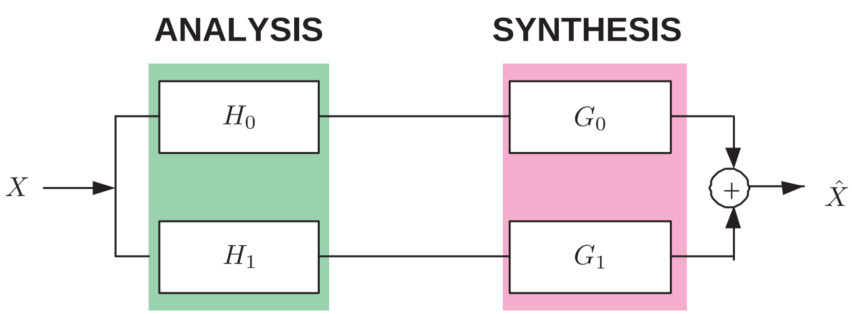

Advantages:

- Shift-invariant
- Fast transforms via the "à trous" algorithm

Iterated Nonsubsampled Contourlet Transform: Illustration

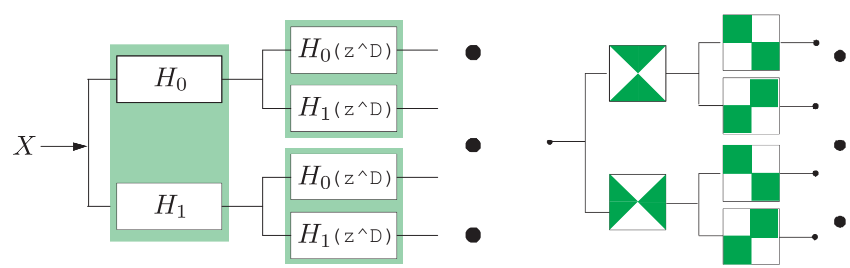

Key: Filtering "with holes" using the equivalent sampling matrices

Example of Nonsubsampled Contourlet Transform (1/2)

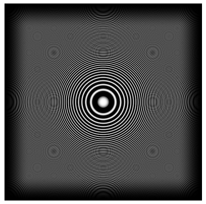

(a)

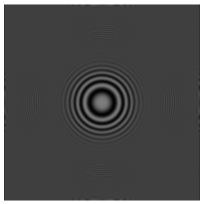

(b)

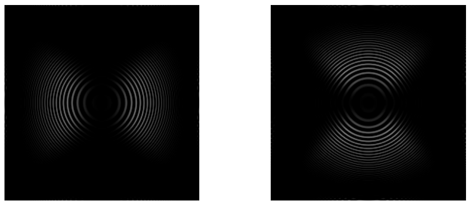

(c)

(a) Original image. (b) Lowpass image. (c) Bandpass directional images.

Example of Nonsubsampled Contourlet Transform (2/2)

Finest bandpass directional images

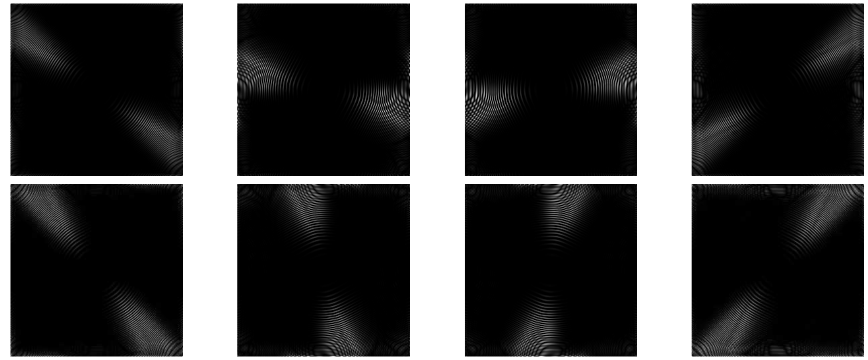

Filter Design for the Nonsubsampled Contourlet Transform (NSCT)

We have to design for two cases: pyramid filters for the LP and fan filters for the DFB

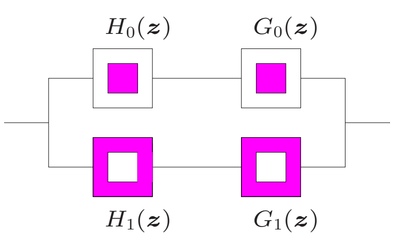

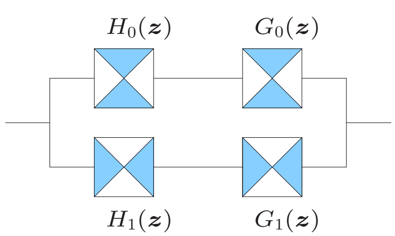

In both cases **perfect reconstruction** means

$$H_0(\boldsymbol{z})G_0(\boldsymbol{z}) + H_1(\boldsymbol{z})G_1(\boldsymbol{z}) = 1$$

This is easier than biorthogonal filter bank condition.

How to Avoid Spectral Factorization?

Answer: Use 1-D to 2-D mapping (a.k.a. McClellan transformation)

We first design a set of 1-D polynomials that solve

$$P_0(x)Q_0(x) + P_1(x)Q_1(x) = 1.$$

Then choose a mapping function $f(\boldsymbol{z})$ such that $|f(e^{j\boldsymbol{w}})| \leq 1$ and substitute

$$H_i(\boldsymbol{z}) = P_i(x)|_{x=f(\boldsymbol{z})}, \text{ and } G_i(\boldsymbol{z}) = Q_i(x)|_{x=f(\boldsymbol{z})}$$

for $i = 0, 1$.

Good Things about Mapping [Cunha]

- It preserves perfect reconstruction
- If mapping is symmetric, then resulting 2-D filter also symmetric
  $\Rightarrow$ easy linear phase FIR design
- Fast implementation through lifting factorization
- Easy control of the frequency response through the mapping filter:
  If $Q_1(x) = Q_0(-x)$ then $H_1(\boldsymbol{z})$ will have the complementary response of $H_0(\boldsymbol{z})$.
- Filters can be symmetric vertically and horizontally
  $\Rightarrow$ allows for simple symmetric extension.

How about Vanishing Moments?

Proposition 3. [CunhaD:04] *Suppose $P(x)$ and $Q(x)$ are such that*

$$P(x)Q(x) + P(-x)Q(-x) = 1$$

and further that $P(x) = (x_0 + x)^{N_P} R_P(x)$, $Q(x) = (x_0 + x)^{N_Q} R_Q(x)$.

Then if we set

$$f(x, y) = -x_0 + (x + 1)^{N_x} (y + 1)^{N_y} \tilde{f}(x, y)$$

the 2-D function $P(f(x, y))$ will have zeroes of multiplicity $N_x N_P$ and $N_y N_P$ at $x = -1$ and $y = -1$ respectively.

Similarly $Q(f(x, y))$ will have zeroes of multiplicity $N_x N_Q$ and $N_y N_Q$ at $x = -1$ and $y = -1$ respectively.

Design Mapping Function $f(x, y)$ for Pyramid Filters

A simple form of $f(x, y)$ is by using maximally flat filters:

$$P_{N,L}(x) := (1 + x)^N \sum_{l=0}^{L-1-N} \binom{N + l - 1}{l} 2^{-N-l} (1 - x)^l$$

Then set

$$f(x, y) = -1 + 2P_{N_0,L_0}(x) P_{N_1,L_1}(y)$$

so that $f(1, 1) = 1$

Design Example for Pyramid Filters

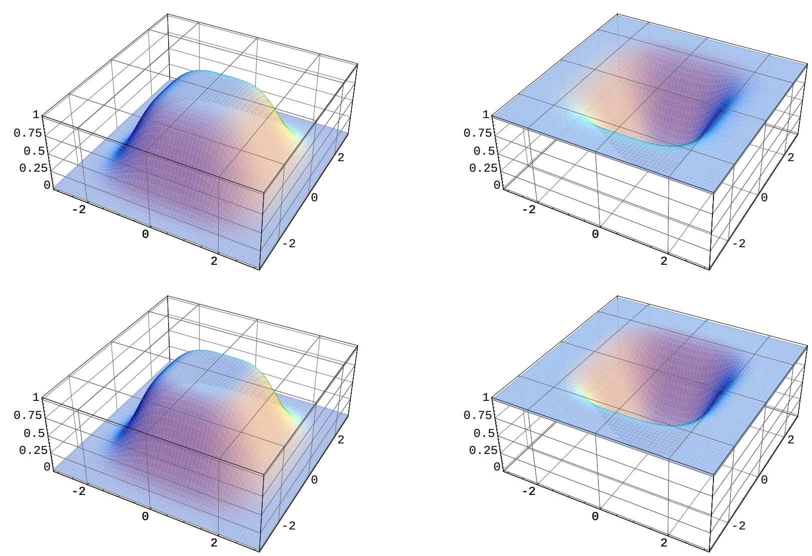

Top: Analysis; **Bottom**: Synthesis

Design Mapping Function $f(x, y)$ for Fan Filters

We can use diamond maximally flat filters and then modulate to get fan filters:

$$f_N(x,y) = 2^{-2N} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1-i} \binom{N}{i}\binom{N}{j} (1-x)^i (1+x)^{N-i} (1-y)^j (1+y)^{N-j}$$
$$+ \; 2^{-2N-1} \sum_{i=0}^{N-1} \binom{N}{i}\binom{N}{i} (1-x)^i (1+x)^{N-i} (1-y)^i (1+y)^{N-j}$$

Then set $f(x, y) = -1 + f_N(x, y)$ as before.

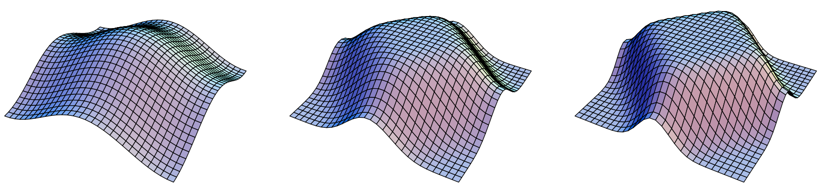

f_N: N=1 N=3 N=6

Denoising Result: *"Hat Zoom"*

Comparison against SI-Wavelet (nonsubsampled wavelet transform) methods using Bayes shrink with adaptive soft thresholding [Cunha]

Noisy Lena	SI-Wavelet	NSCT Den
PSNR 22.13db	PSNR 31.82dB	PSNR 32.14dB

Denoising Result: *"Peppers"*

Noisy Lena	SI-Wavelet	NSCT Den
PSNR 22.14db	PSNR 31.38dB	PSNR 31.53dB

Denoising Result: *"Barbara"*

Noisy "Barbara"	SI-Wavelet	NSCT Den
PSNR 22.15db	PSNR 29.34dB	PSNR 29.95dB

Comparison Against SI-Wavelet Method

"Lena"	PSNR (dB)			"Peppers"	PSNR (dB)	
Std	Noisy	SI-Wavelet	NSCT	Noisy	SI-Wavelet	NSCT
$\sigma = 10$	28.13	35.02	35.07	28.17	34.23	34.03
$\sigma = 20$	22.13	31.83	32.14	22.14	31.38	31.53
$\sigma = 30$	18.63	29.96	30.35	18.63	29.64	29.92
$\sigma = 35$	17.29	29.24	29.62	17.29	28.93	29.26

PSNR results of the tested denoising schemes for different noise levels.

Another Application: Image Enhancement [Zhou]

Image Enhancement Algorithm...

- Nonsubsampled contourlet decomposition
- Nonlinear mapping on the coefficients
  - Zero-out noises
  - Keep strong edges or features
  - Enhance weak edges or features
- Nonsubsampled contourlet reconstruction

Image Enhancement Result: *Barbara*

(a)

(b)

(c)

(a) Original image. (b) Enhanced by DWT. (c) Enhanced by NSCT.

Image Enhancement Result: *OCT Image*

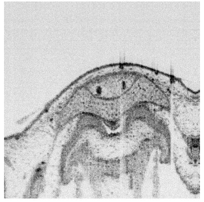

(a)

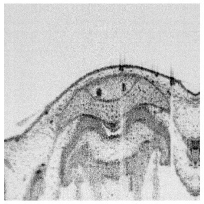

(b)

(c)

(a) Original OCT image. (b) Enhanced by DWT. (c) Enhanced by NSCT.

Outline

1. Image modeling using the contourlet transform
2. Critically sampled (CRISP) contourlet transform
3. Image denoising and enhancement using the nonsubsampled contourlet transform
4. Outlook

Speculation: Another “Wavelet” Story ?

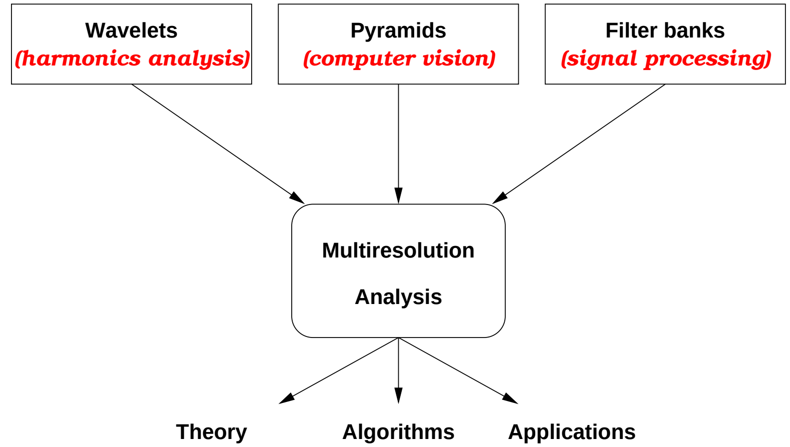

1-D Story: scale and shift

Speculation: Another “Wavelet” Story ?

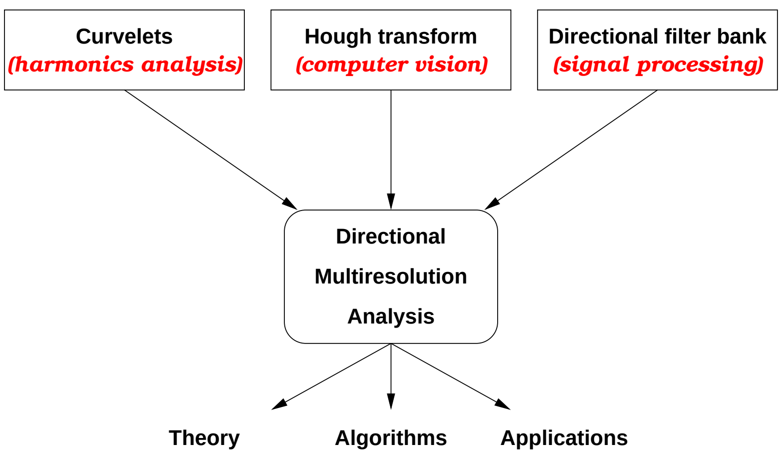

2-D Story: scale, space, and direction

Where is The Complexity in Images?

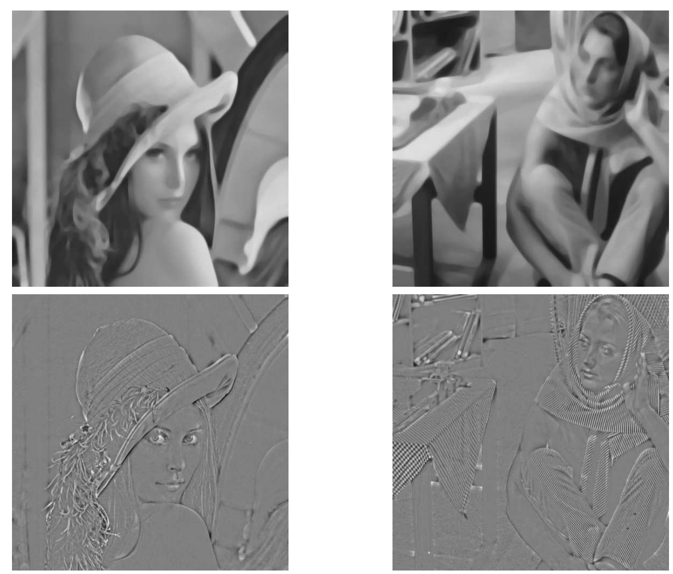

Operational Rate-Distortion by JPEG-2000

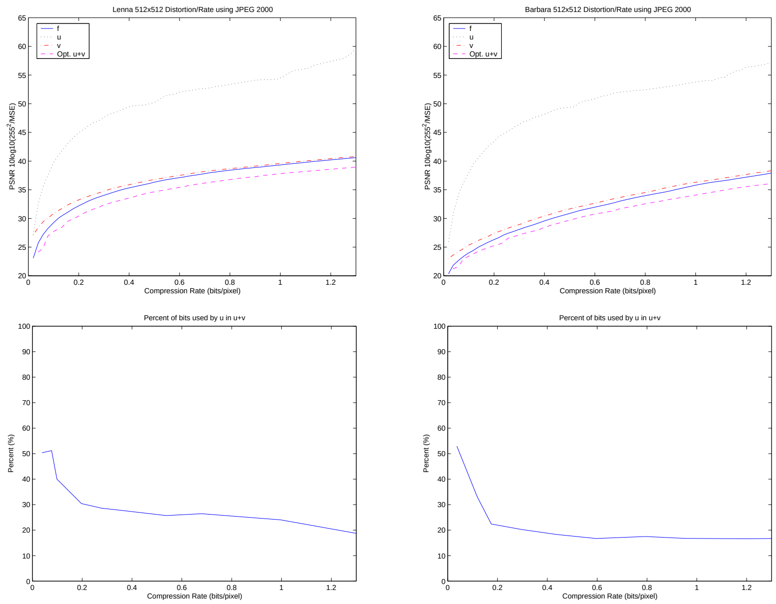

Beyond 2-D and Single Image... New Audio-Visual Paradigm

- Existing audio-visual recording and playback
  - Single camera and microphone
  - Little processing
  - Viewers: passive
- Future
  - Sensors (cameras, microphones) are cheap
  - Massive computing and storage capabilities are available
  - Viewers: interactive, immersive, remote

$\Rightarrow$ Require new signal processing theory and algorithms

Note: Video can be treated as a special case (frames in one shoot $\equiv$ multiple images of a scene)

Image-based Rendering

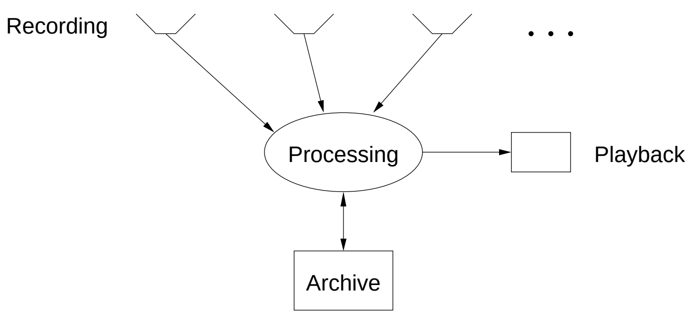

- **Goal**: Synthesize arbitrary viewpoint from a set of fixed sensors.
- Input data at a huge rate, for example: 12 M/frame $\times$ 10 frame/s $\times$ 15 views = 1.8 Gbps.
- **Challenge**: New representations for IBR data that incorporate geometric constraints.

Light-field Parameterization

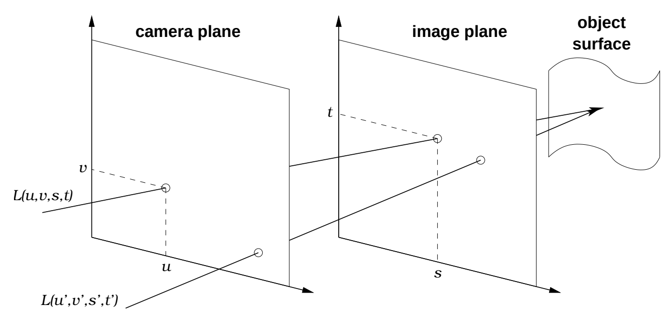

- Each light ray is addressed by a 4-D coordinate (u, v, s, t).
- Assuming the object surface is Lambertian, then $L(u, v, s, t) = L(u', v', s', t')$.

  $\Rightarrow$ The 4-D light field data $L(u, v, s, t)$ lie in lower dimensional manifolds.

Epipolar Constraint from Multiple Views

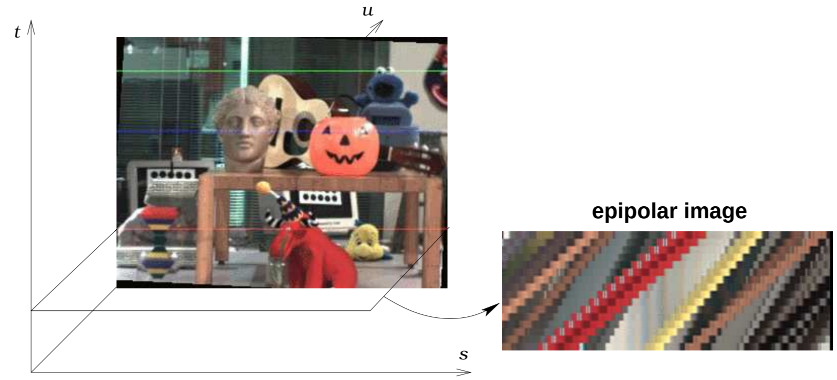

- Linear singularities in epipolar planes ... ridgelets?
- Curved singularities in image planes ... contourlets?
- **Ultimate:** High dimensional representations that can deal effectively with lower dimensional singularities.

References

- D. D.-Y. Po and M. N. Do, "Directional multiscale modeling of images using the contourlet transform," *IEEE Trans. on Image Proc.*, to appear.
- Y. Lu and M. N. Do, "CRISP-contourlets: a critically sampled directional multiresolution image representation," *Proc. of SPIE conf. on Wavelet Applications in Signal and Image Proc.*, San Diego, USA, August 2003.

- **Software and downloadable papers**: www.ifp.uiuc.edu/∼minhdo