Wavelet X, Aug. 2003

Wavelets, Approximation and Compression: Beyond JPEG2000

Martin Vetterli   EPFL & UC Berkeley

1. Introduction: What is the Problem?

2. Wavelets and Approximation

3. Linear Versus Nonlinear Approximation

4. Compression of Piecewise Smooth/Polynomial Functions

5. New Separable Constructions
   - Footprints
   - Directional Wavelet Transforms and Frames

6. New Non-Separable Constructions
   - Contourlets
   - Tree Based Geometric Compression

7. Conclusions and Outlook

Acknowledgements

Sponsors:

- NSF Switzerland

Collaborations:

- M.Do
- P.L.Dragotti
- P.Prandoni
- R.Shukla
- V. Velisavljevic
- C.Weidmann

Discussions and Interactions:

- I. Daubechies (Princeton)
- D. Donoho (Stanford)
- J. Kovacevic and V.Goyal (Bell Labs)
- S. Mallat (Polytech. & NYU)
- K. Ramchandran (UIUC)
- M.Unser (EPFL)

How many bits for Mona Lisa?

<=> {0,1}

The problem, its history, and its importance

"... the 20 bits per second which, the psychologists assure us, the human eye is capable of taking in, ..."

D.Gabor, September 1959

How many bits to code Mona Lisa?

Index all pictures ever taken in the history of mankind

- $100 \text{ years} \cdot 10^{10} \sim 44$ bits

All pictures viewable by all humans

- $5 \cdot 10^{9}$, 100 years, $25 \cdot 3600 \cdot 24 \cdot 365 \sim 69$ bits

Huffman code Mona Lisa index

- a few bits (Lena Y/N?, Mona Lisa...)
- R(D)....

Search the Web!

http://www.paris.org/Musees/Louvre/Treasures/gifs/Mona_Lisa.jpg

JPEG

- 186K... There is plenty of room at the bottom!

Introduction

More seriously:

Parsimonious representation of visual information remains a central problem which requires good models

- storage and transmission
- indexing, searching, classification
- processing, denoising, enhancing, resolution change
- watermarking

But: it is also a fundamental question in

- information theory
- image processing
- approximation theory

New image coding standard ... JPEG 2000

Most of the best proposed coders are based on wavelets

- many variations
- all use multiresolution expansion

Final JPEG 2000 standard is strongly influenced by wavelet methods

But: do we understand images better because of this?

Other successes of wavelets in image processing:

- denoising
- enhancement
- classification

Thesis: Wavelet models for images
play an increasingly important role

Antithesis: Wavelets are just another fad!

Old Versus New JPEG: D(R) on log scale

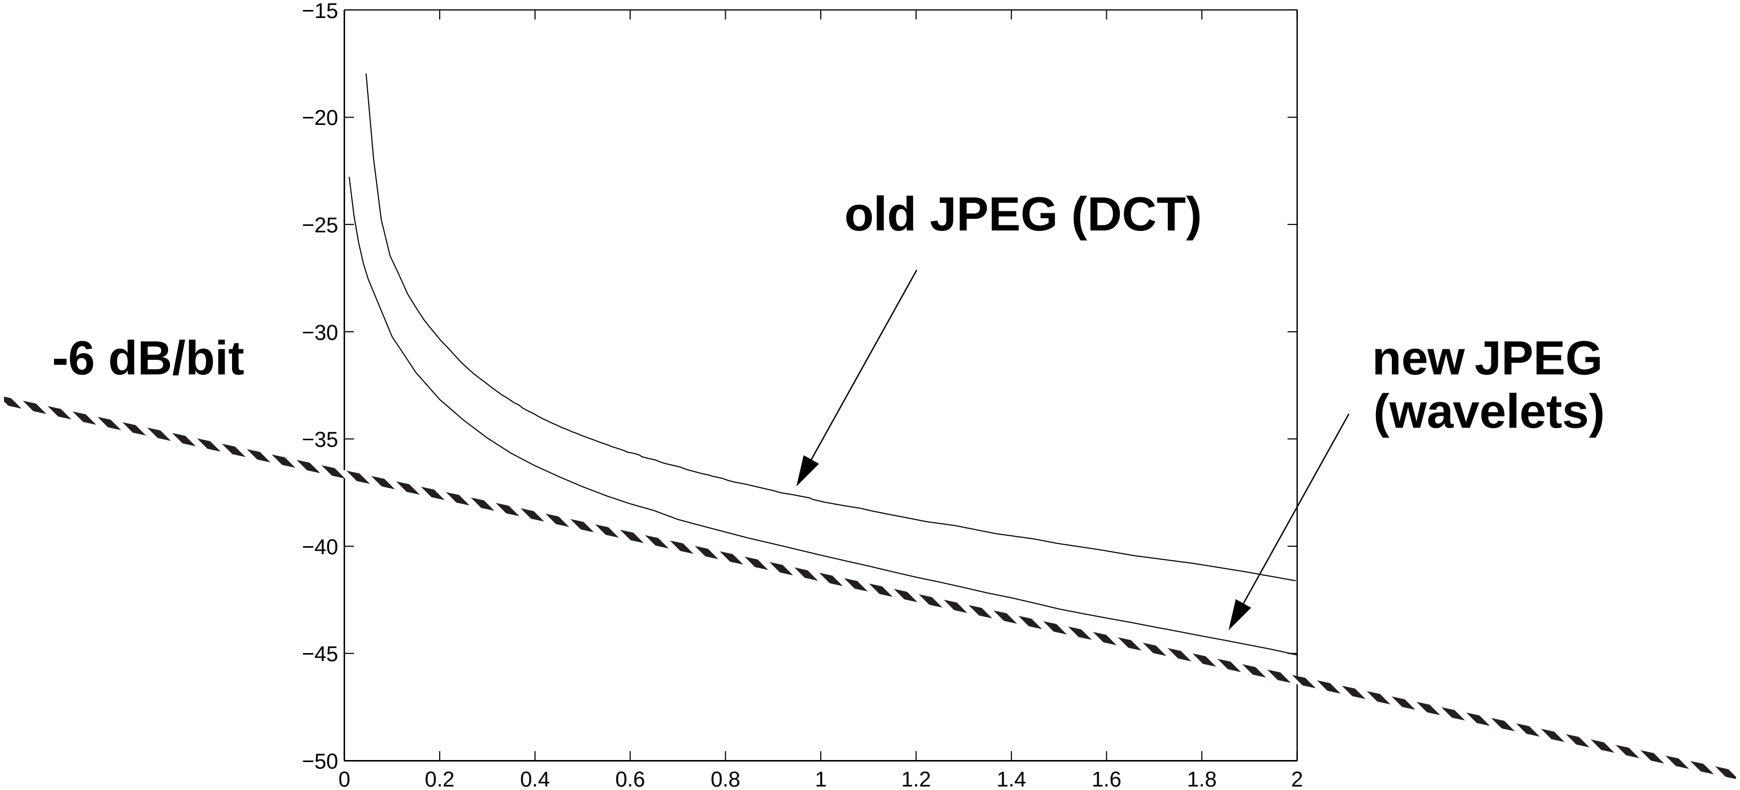

Notes

- improvement by a few dB's
- lot more functionalities (e.g. progressive download on internet)
- at high rate ~ -6db per bit: KLT behavior
- low rate behavior: much steeper: NL approxim. [Mallat et al]
- is this the limit?

Original Lena Image (256 x 256 Pixels, 24-Bit RGB)

JPEG Compressed (Compression Ratio 43:1)

JPEG2000 Compressed (Compression Ratio 43:1)

From the comparison, JPEG fails above 40:1 compression while JPEG2000 survives

Images courtesy of www.dspworx.com

2. Wavelets and approximation theory

Multirate filter banks and projections:

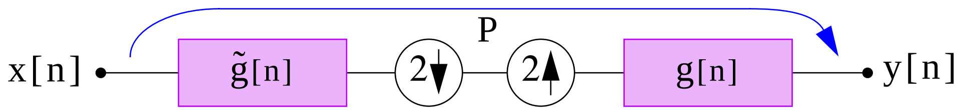

If $\{g[n-2k]\}_{k \in \mathfrak{I}}$ **is an orthonormal set, then the above operator computes an orthogonal projection**

If in addition $g[n]$ **has z-transform:**

$$G(z) = (1 + z^{-1})^N \cdot R(z)$$

then polynomials up to degree $N\text{-}1$ **are eigensignals**

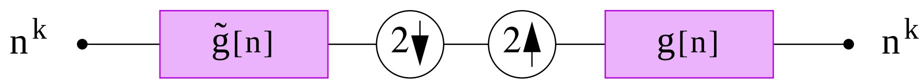

In a filter bank

- $G_0(z) = (1 + z^{-1})^N \cdot R(z)$ $\qquad$ $G_1(z) = (1 - z^{-1})^N \cdot S(z)$

- $H_0(z) = (1+z)^N \cdot \tilde{R}(z)$ $H_1(z) = (1-z)^N \cdot \tilde{S}(z)$

where ~ stands for time-reversal and $R(z) \sim S(z)$

Then

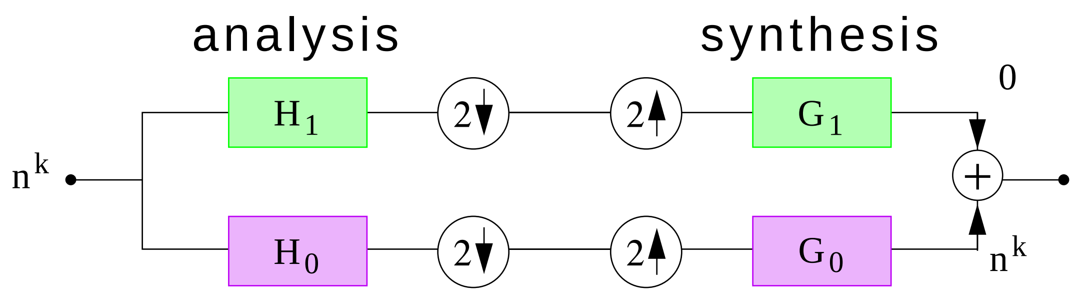

and polynomials up to degree N-1 will be cancelled by the highpass channel, but reproduced by the lowpass, or

- highpass has N zero moments
- lowpass reproduces polynomials of deg less than N (discrete-time Strang-Fix condition)

Iterated filter bank ($H_j(z) = G_j(z^{-1})$)

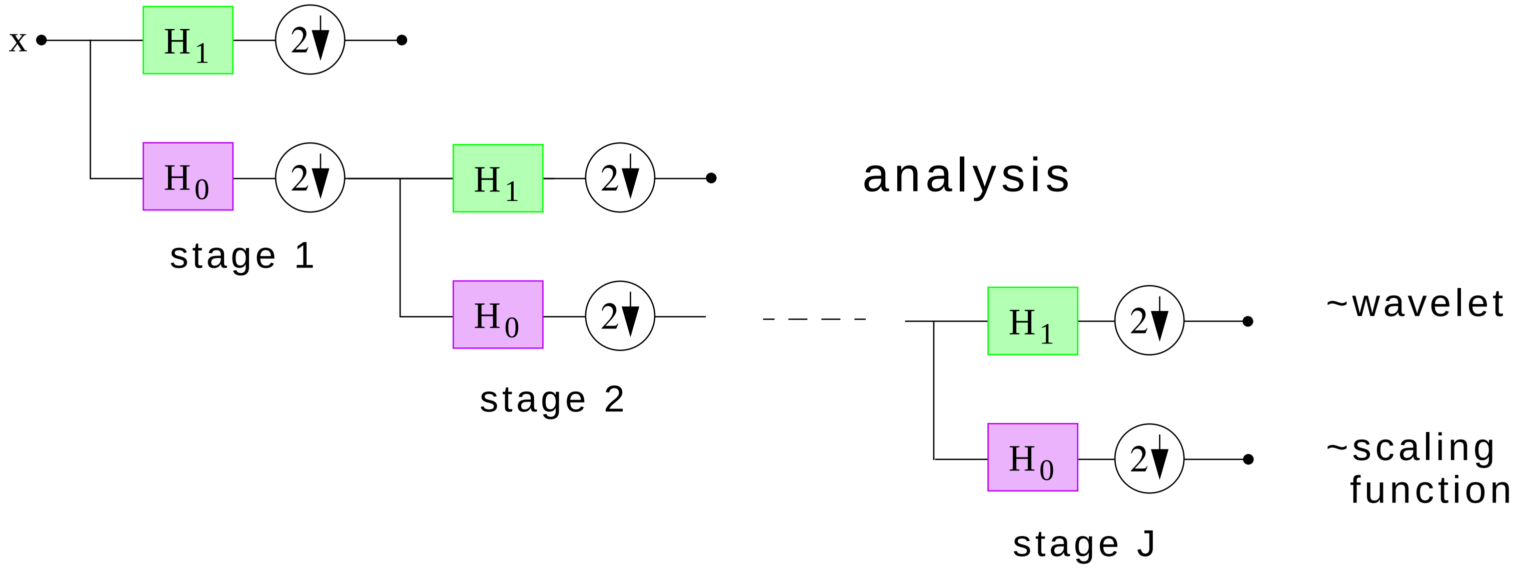

By a similar argument:

- polynomials are “eaten” in the highpass
- polynomials are reproduced by the lowpass channel

Result: Iterated synthesis filter $g^{(J)}[n]$ and its shifts by 2^J reproduce polynomials

$$n^k = \sum_m c_m \cdot g^{(J)}[n - 2^J m]$$

How about wavelet approximations?

$G_0(z)$ **having** N **zeroes at** π **implies (orthogonal case):**

- scaling function $\varphi(t)$ spans polynomials up to degree N-1

$$\sum_n c_n \cdot \varphi(t-n) = t^k \qquad k = 0, 1, \ldots N-1$$

- wavelet ψ has N zero moments
  eats polynomials up to deg. N-1
- wavelet is of length $L = 2N-1$
  $2N-1$ wavelets influenced by singularity at each scale
- wavelet/scaling function are smooth of order $0.203\ N$
  good approximation

Results carry over, with adjustments, to

- biorthogonal case
- multiwavelet case
- multidimensional case

Example: Iterate of D_4 lowpass reproduces linear ramp

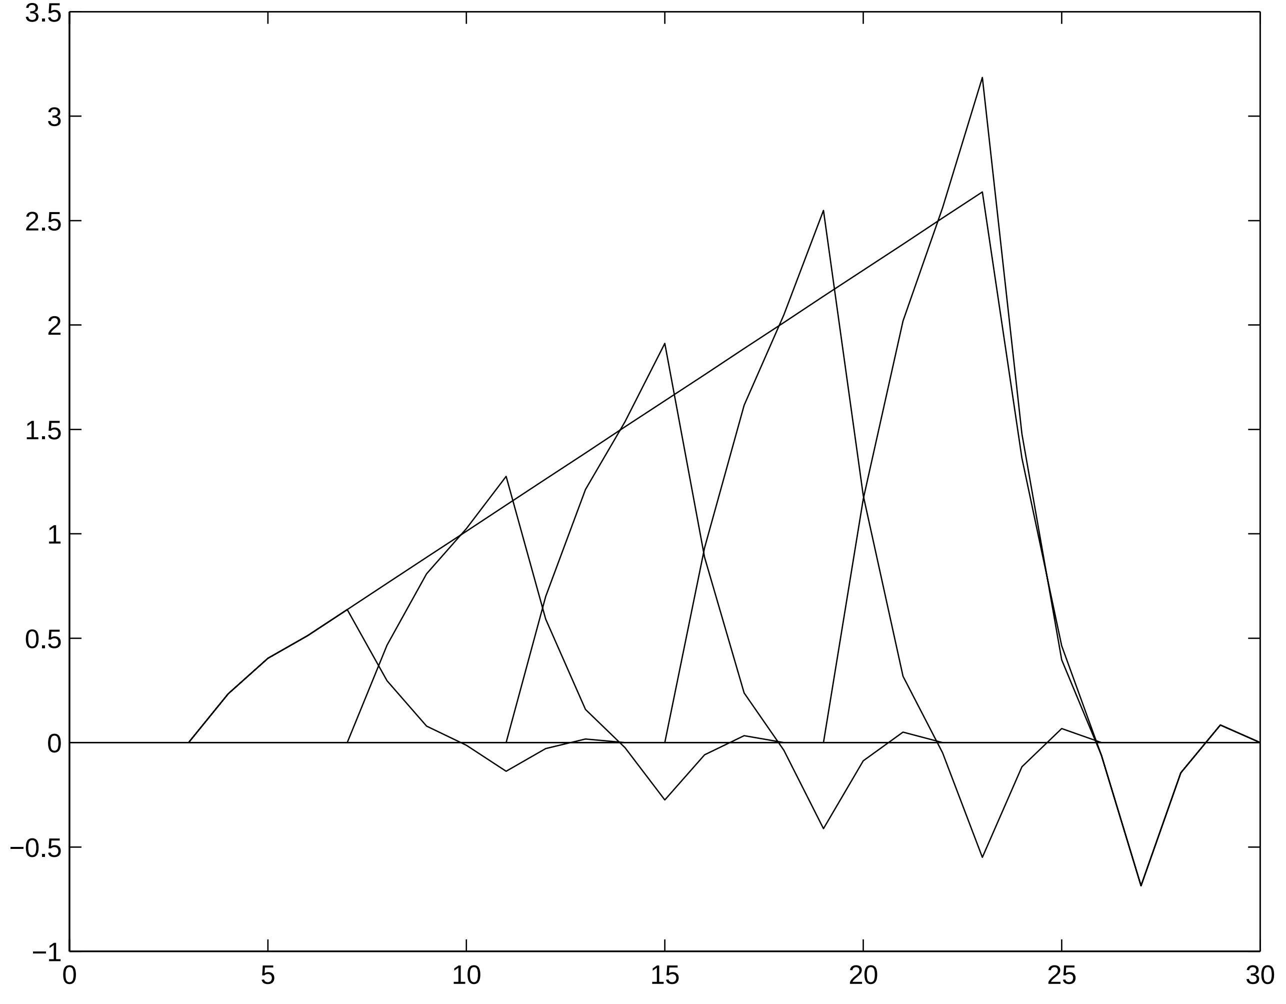

How about singularities?

If we have a singularity of order n at the origin (-1: Dirac, 0: Heaviside,...**), the CWT transform behaves as**

$$X(a, 0) = c_n \cdot a^{n/2}$$

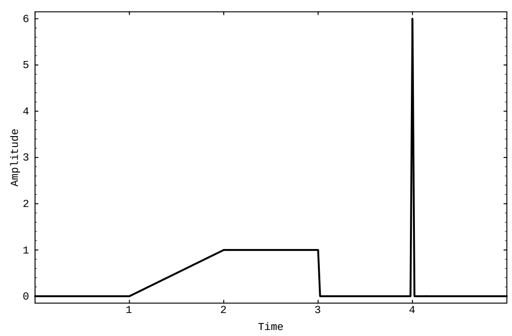

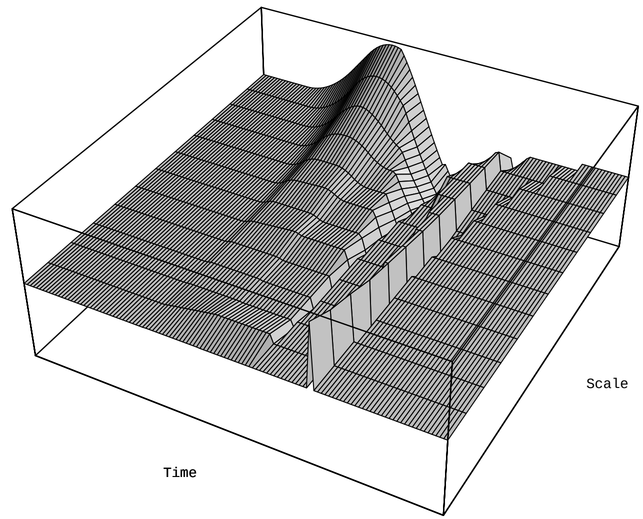

In the orthogonal wavelet series: same behavior, but only $L=2N-1$ coefficients influenced at each scale!

- e.g. Dirac/Heaviside: behavior as $2^{-m/2}$ and $2^{m/2}$

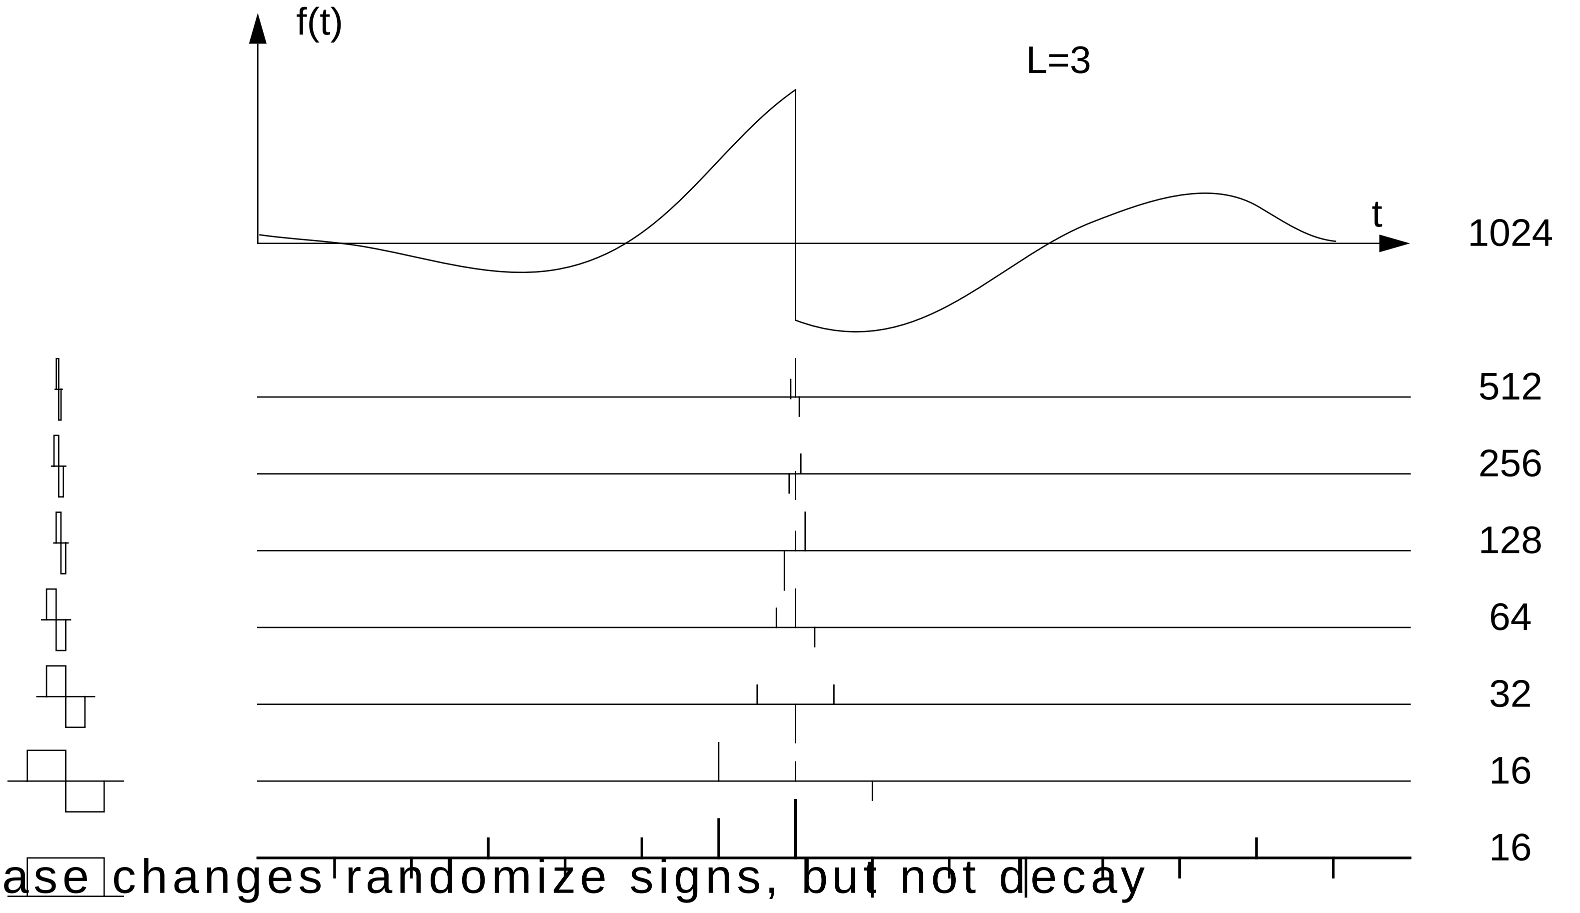

- phase changes randomize signs, but not decay
- a singularity influences only L wavelets at each scale (L=2N-1=3)

in conclusion:

Consider piecewise smooth signals

- Wavelet act as singularity detectors
- Scaling functions catch smooth parts
- “Noise” is circularly symetric

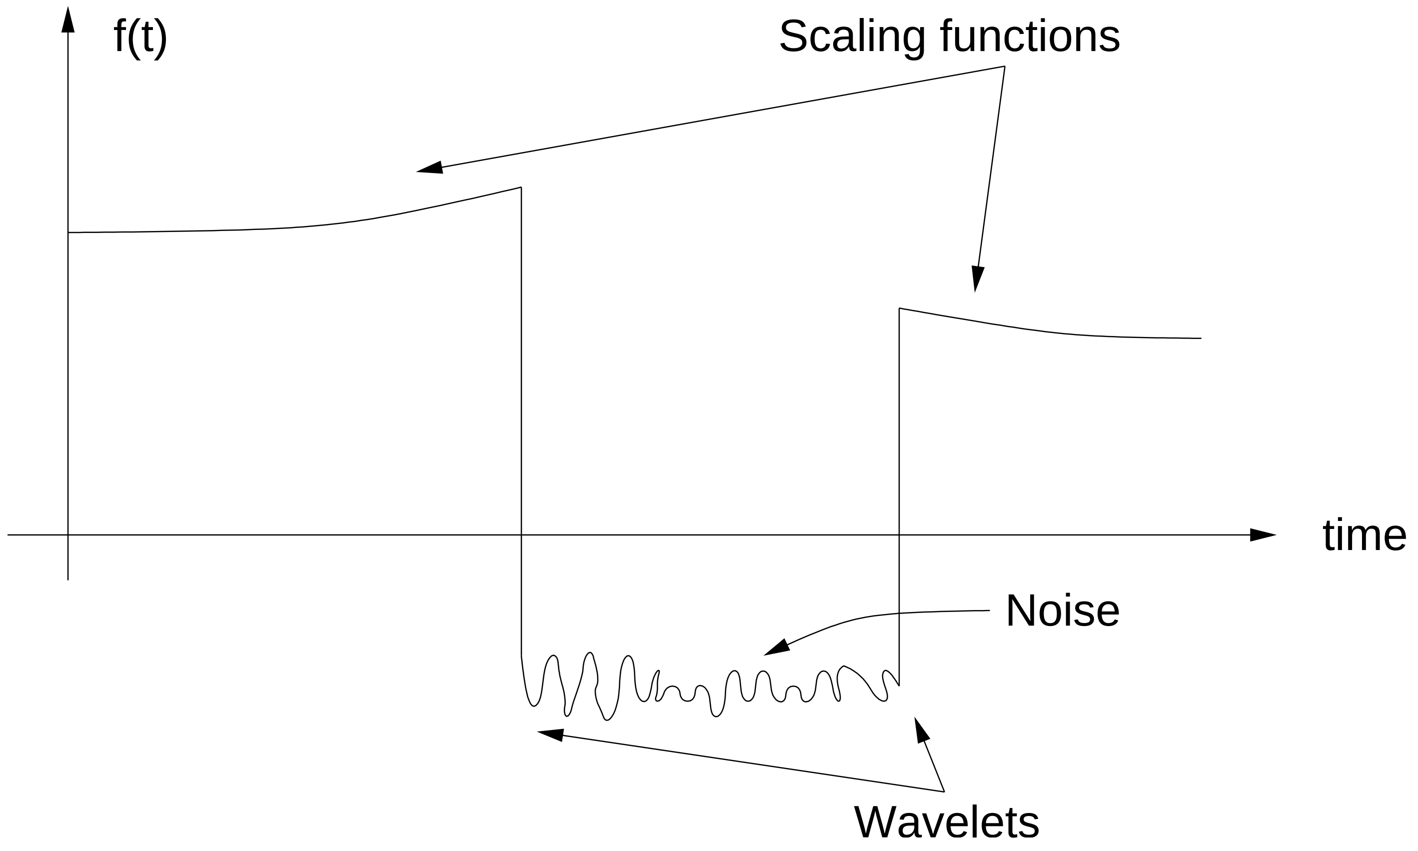

3. Linear versus non-linear approximation

The linear approximation method

Given an orthonormal basis $\{g_n\}$ for a space S and a signal

$$f = \sum_n \langle f, g_n \rangle \cdot g_n ,$$

the best linear approximation is given by the projection onto a fixed sub-space of size M (independent of f!)

$$\hat{f}_M = \sum_{n \in J_M} \langle f, g_n \rangle \cdot g_n$$

The error (MSE) is thus

$$\hat{\varepsilon}_M = \left\| f - \hat{f} \right\|^2 = \sum_{n \notin J_M} \left| \langle f, g_n \rangle \right|^2$$

The nonlinear approximation method

Given an orthonormal basis $\{g_n\}$ for a space S and a signal

$$f = \sum_n \langle f, g_n \rangle \cdot g_n ,$$

the best nonlinear approximation is given by the projection onto an adapted subspace of size M (dependent on f!)

$$\tilde{f}_M = \sum_{n \in I_M} \langle f, g_n \rangle \cdot g_n$$

$$I_M: \qquad |\langle f, g_n \rangle|_{n \in I_M} \geq |\langle f, g_m \rangle|_{m \notin I_M} \qquad \text{set of M largest } \langle \ , \ \rangle$$

The error (MSE) is thus

$$\tilde{\varepsilon}_M = \|f - \tilde{f}\|^2 = \sum_{n \notin I_M} |\langle f, g_n \rangle|^2$$

and $\tilde{\varepsilon}_M \leq \hat{\varepsilon}_M$.

Difference: take the first M coeffs (linear) or
take the largest M coeffs (non-linear)

Nonlinear approximation

- this is a simple but nonlinear scheme

Clearly, if $A(.)$ is the NL approximation scheme:

$$A(x) + A(y) \neq A(x + y)$$

in general.

This could be called ''adaptive subspace fitting''

From a compression point of view, the key is to be able to ''pay'' for the adaptivity

- in general, this will cost

$$\log\left(\begin{bmatrix} N \\ k \end{bmatrix}\right) \text{ bits}$$

which cannot be spent on coefficient representation anymore

When does it pay off?

Nonlinear approximation

Nonlinear approximation power depends on basis

Example:

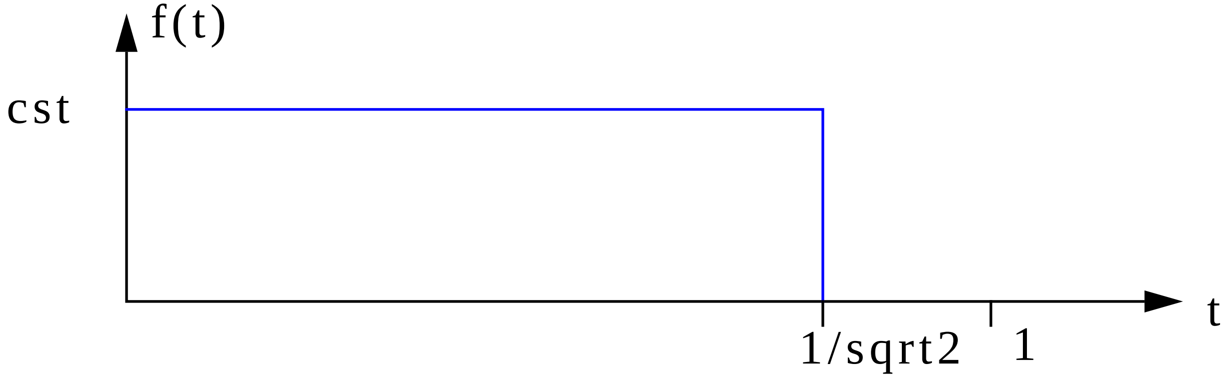

Two different bases for $[0,1]$:

- Fourier series $\{e^{j2\pi kt}\}_{k \in \mathfrak{I}}$
- Wavelet series: Haar wavelets

Linear approximation in Fourier or wavelet bases

$$\hat{\varepsilon}_M \sim 1/M$$

Nonlinear approximation in a Fourier basis

$$\tilde{\varepsilon}_M \sim 1/M$$

Nonlinear approximation in a wavelet basis

$$\tilde{\varepsilon}_M \sim 1/2^M$$

Fourier Basis: N=1024, M= 64, linear versus nonlinear

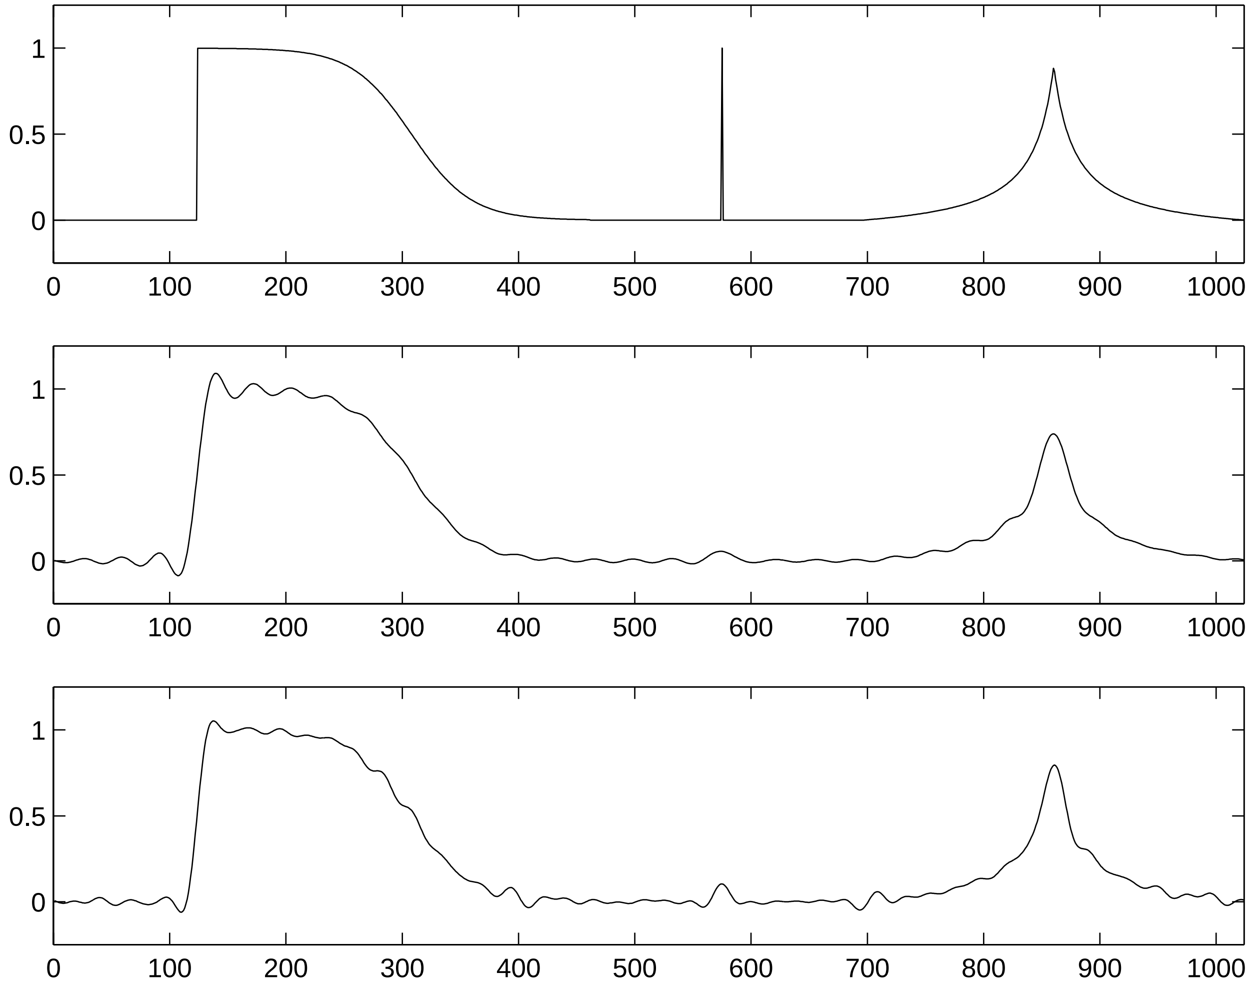

D= 2.7

D= 2.4

- nonlinear approximation is not necessarily much better!

Wavelet basis: N=1024, M= 64, J=6**, linear versus nonlinear**

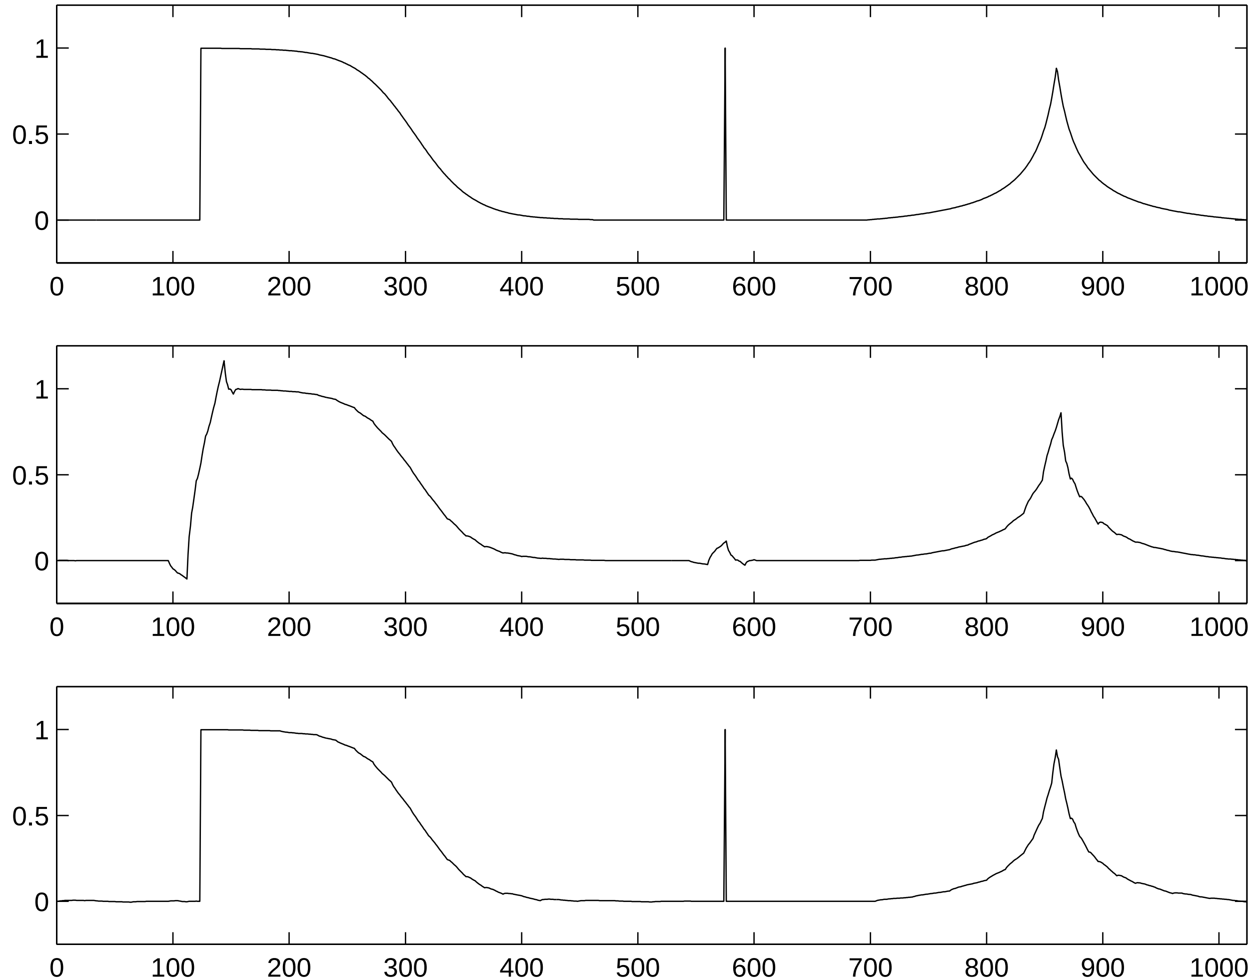

D=3.5

D=0.01

- nonlinear approximation is vastly superior!

Nonlinear approximation theory and wavelets

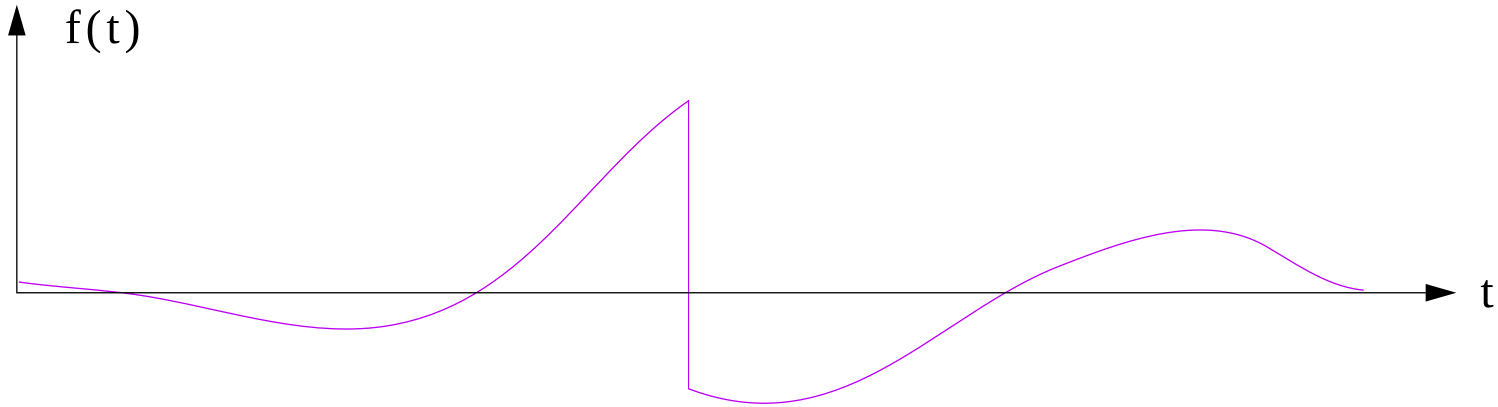

Strong approximation results for piecewise smooth fcts

- between discontinuities,
  behavior by Sobolev or Besov regularity
- K derivatives $\Rightarrow$ coeffs $\sim 2^{m(K-1/2)}$ when $m \ll 0$
- Besov spaces can be defined with wavelet bases. If

$$\|f\|_{G,p} = \left(\sum |\langle f, g_n\rangle|^p\right)^{1/p} < \infty \qquad 0 < p < 2$$

  then [DeVoreJL92]: $\tilde{\varepsilon}_M = o(M^{1-2/p})$
- limits of Besov space characterization for images...

Smooth with finite number of discontinuities

- does not change behavior fundamentally
- argument: compact support, finite effect
- dominant effect: coeffs $\sim 2^{m(K-1/2)}$

But how about $D(R)$?

We choose only the M largest coefficients, but

- how to index them?
- how to quantize them?

Method

- inspired from real coding practice!
- chain-code accross scale, to effectively achieve NLA

Results

- interesting new $D(R)$ results for various fct classes (functions of bounded variations [Cohen et al])

Connection: Kolmogorov ε-entropy, Shannon $D(R)$ [DeVore et al, Donoho]

Aproximation and Compression

A small instructive example:

Assume

- $x[n] = \alpha\, \delta[n-k]$, where signal is of length N, k is U[0,N-1] and α is N(0,1). This is a Gaussian RV at location k

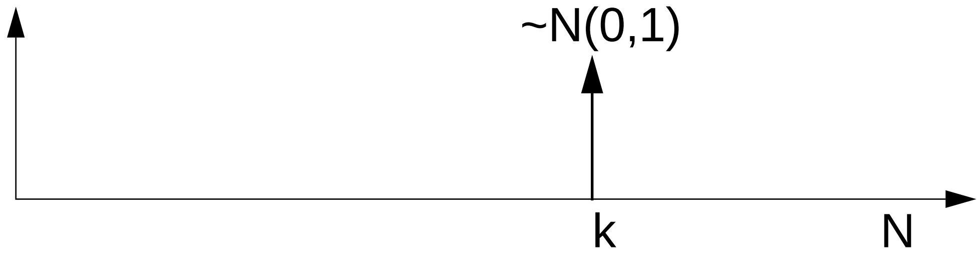

- Note: $R_x = I$!

Given budget R **for block of size** N**:**

KLT: equal distribution of R/N **bits**

$$D(R) = c \cdot \sigma^2 \cdot 2^{-2(R/N)}$$

This is the optimal linear approximation!

R(D) Analysis [WeidmannV:99]

High rate case:

- Obvious scheme: pointer + quantizer

$$D(R) = c \cdot \sigma^2 \cdot 2^{-2(R - \log N)}$$

- This is the R(D) behavior for R >> Log N
- Much better than linear approximation

Low rate case:

- trickier...
- Hamming case solved, inc. multiple spikes:
  - there is a linear decay at low rates
- L_2 case: upper bounds that beat linear approx.

4. Compression of Piecewise Smooth/Polynomial Functions

Approximation:

- we understand decay of linear approximation schemes (e.g. Sobolev regularity and Fourier series)
- we understand decay of non-linear approximation schemes (Besov regularity and wavelet series)

Compression:

- we have to pay for the representation cost!
- the ultimate bound is Shannon's rate-distortion function, but it is only known in very few cases....

Goal:

- approximate behaviors of D(R) for large R
- computationally effective algorithms

Rate-distortion bounds for piecewise polynomial functions

$R(D)$ behavior of nonlinear approximation with wavelets

Consider the simplest case

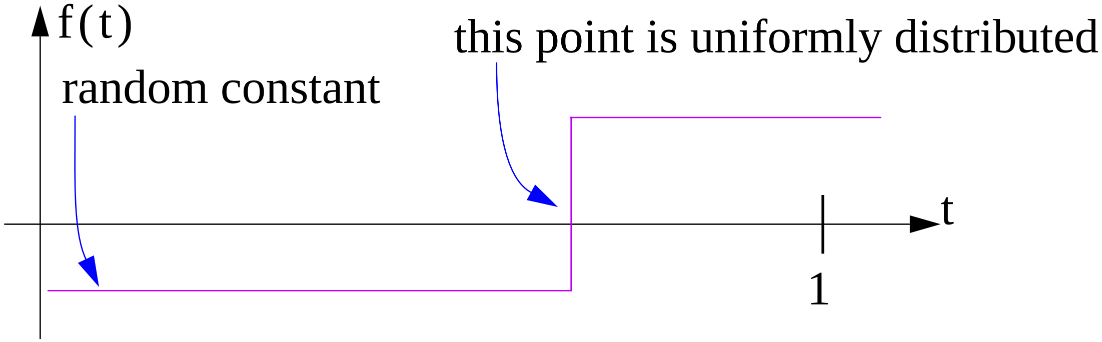

and the Haar basis. Recall that

$$\tilde{\varepsilon}_M \cong 2^{-M}$$

$$C_j \cong 2^{j/2}$$

and consider describing the significant coefficients

Quantization and rate allocation

Choose a stepsize Δ for a quantizer

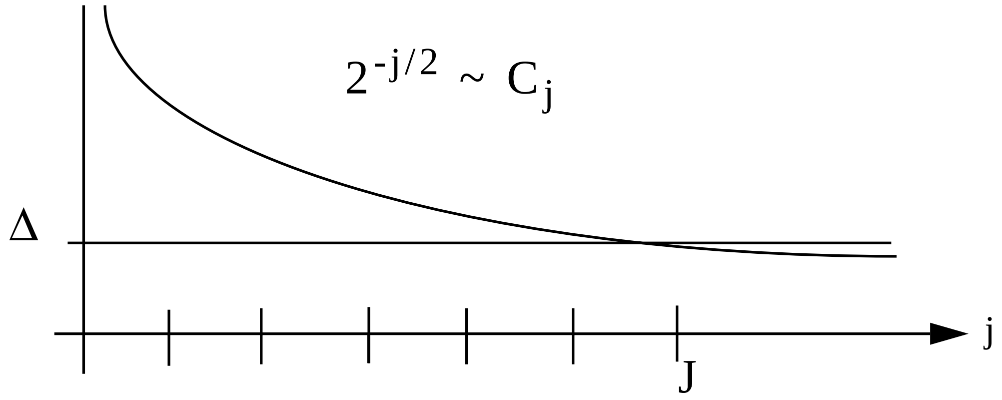

Therefore

- number of scales J before coefs set to zero $\sim \log(1/\Delta)$
- number of bits per coefficient $\sim \log(1/\Delta)$
- $R \sim J^2$

Distortion

- number of scales $\cdot \Delta^2 \sim J \cdot 2^{-J}$

Thus

$$D_w(R) = C_3 \cdot \sqrt{R} \cdot 2^{-c_2 \cdot \sqrt{R}}$$

Rate-distortion behavior using an oracle

Consider the simplest case

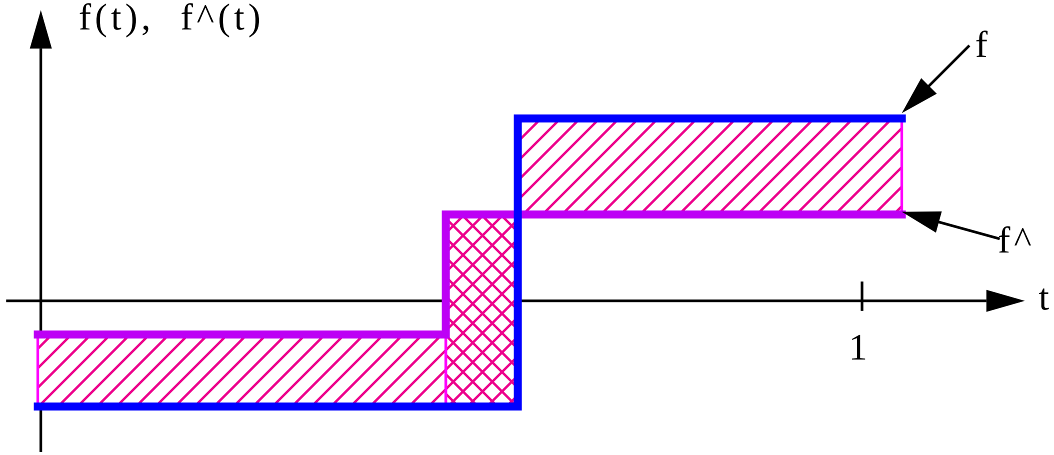

Two approximation errors

- Δ_t: quantization of step location
- Δ_a: quantization of amplitude

Rate allocation: R_t versus R_a

Result:

$$D_p(R) = C_1 \cdot 2^{-R/2}$$

General case

Piecewise polynomial, with max degree N

A. Nonlinear approximation with wavelets having $N+1$ zero moments

$$D_w(R) = C'_w \cdot (1 + \alpha\sqrt{C_w R}) \cdot 2^{-\sqrt{C_w R}}$$

B. Oracle-based method

$$D_p(R) = C'_p \cdot 2^{-(C_p \cdot R)}$$

Thus

- wavelets are a generic but suboptimal scheme
- oracle method asymptotically superior but dependent on the model

These are all high-rate results

An algorithm reaching the oracle bound

Assume discrete-time problem:

Segmentation based on Dynamic Programming

Piecewise polynomial approx.

- coding of Legendre polynomial approx. and of position
- optimal allocation of rate between position and shape

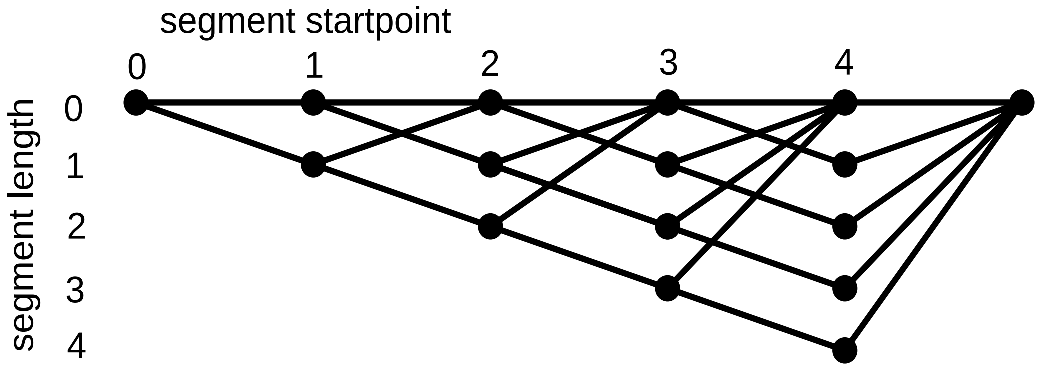

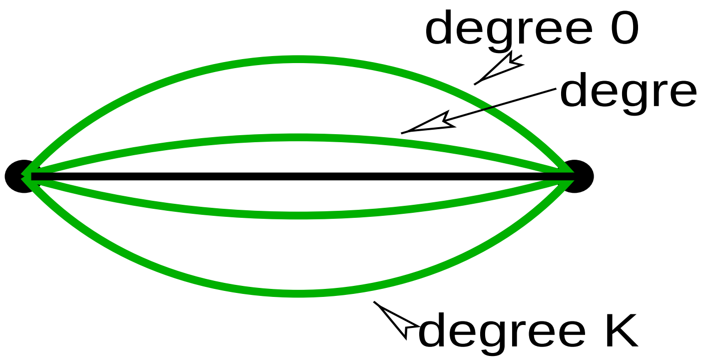

Bellmann principle: Best path uses best segments!

Algorithm: recursively find best segments

- all paths
- all possible degrees (N)
- sweep λ

R(D) performance:

$$D_{DP}(R) = C_1 \cdot 2^{-cR}$$

Example:

K = 5

N = 256

R/D curve:

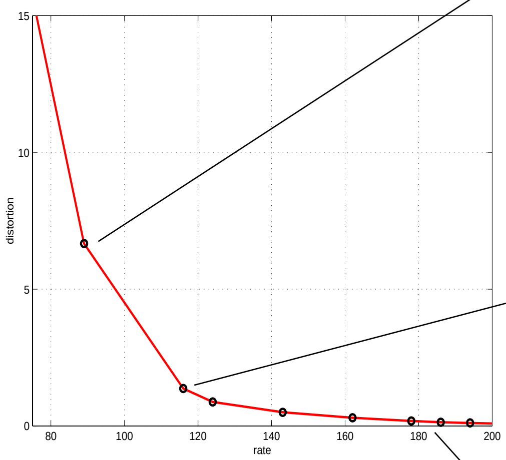

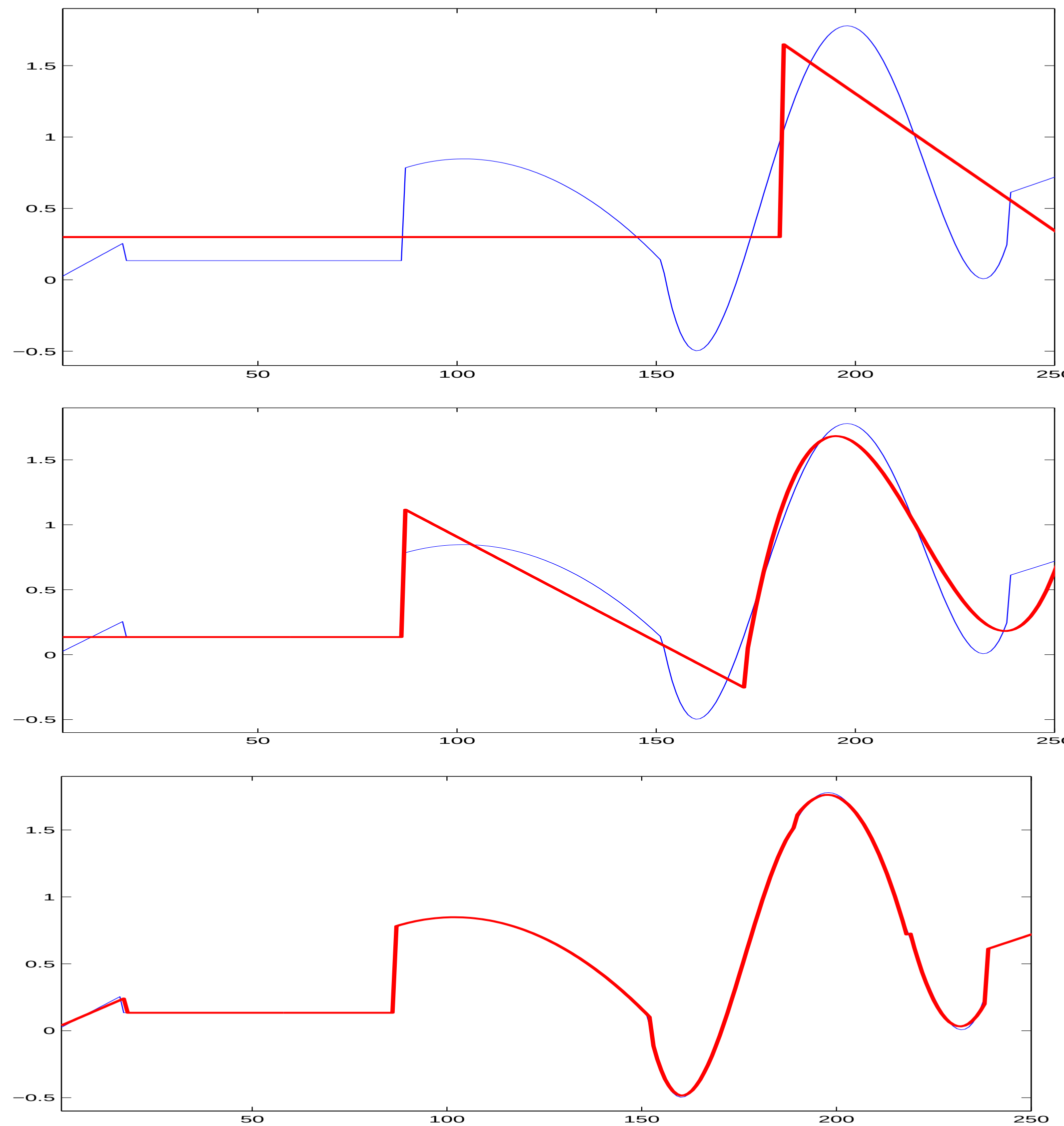

Lines from ... Lena!

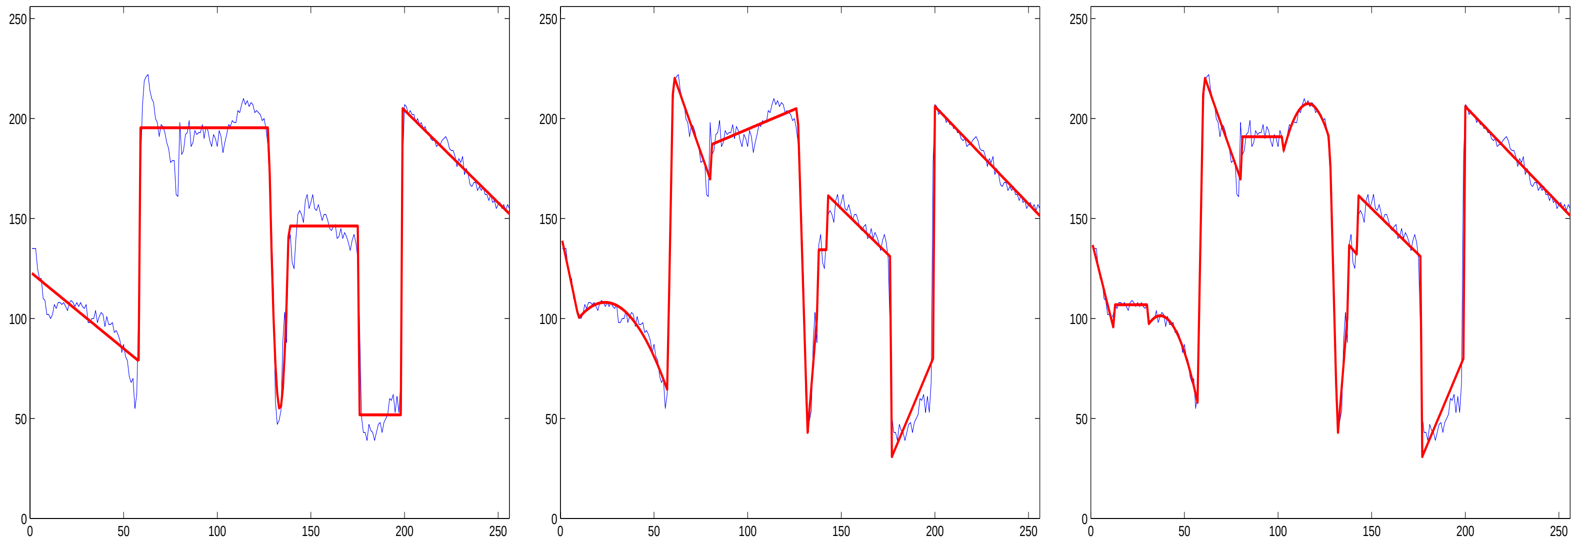

146 bits   297 bits   351 bits

Denoising

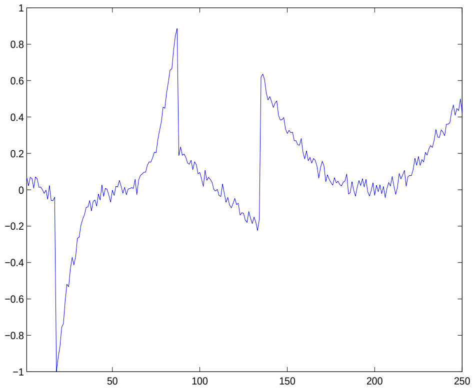

SNR 20dB

SURE soft thresholding

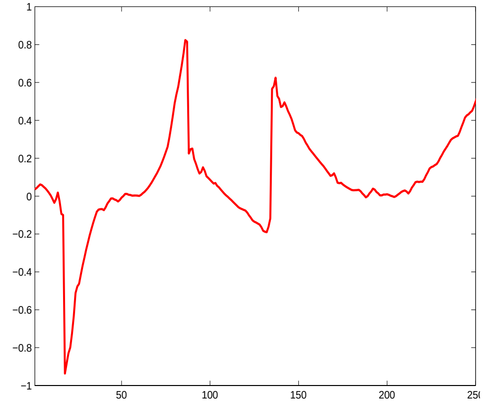

SNR 23.16dB

R/D modeling + “Occam” thrsh.

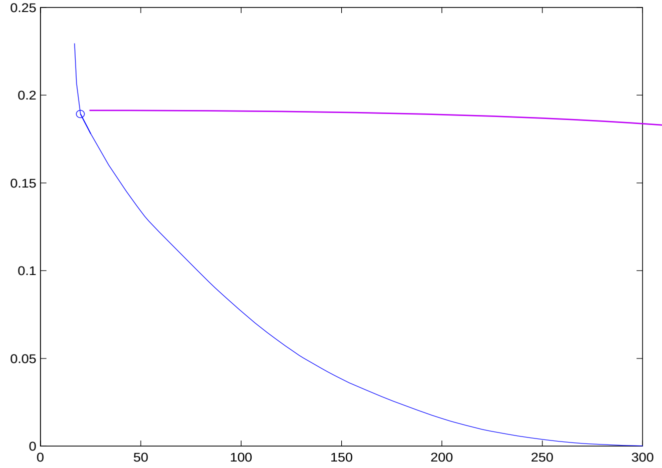

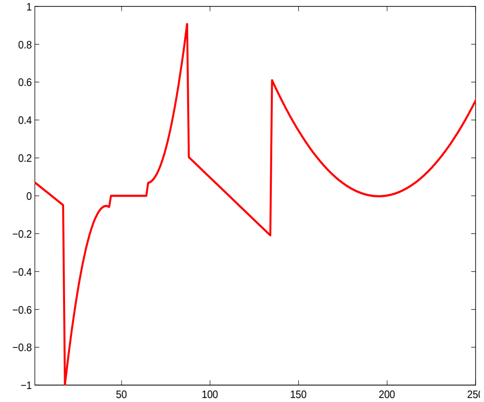

SNR 31.48dB

New Constructions

Coding beyond wavelets:

- things are not gaussian and dependencies need to be addressed
- we can squezze more out of wavelets!
- a simple data structure called **footprints** gets you there
- **directional** wavelet transforms are simple yet powerful

Coding beyond one dimension:

- piecewise smooth in 1D has point singularities
- piecewise smooth in 2D has curve singularities!
- directional analysis is the key
- there is a zoo of new constructions
- we will see **pyramidal directional filter banks** and **contourlets**
- direct approach with tree and quadtree structures that are optimal in rate-distortion sense

Therefore:

5.1 Footprints

5.2 Directional Wavelet Transforms

6.1 Contourlets

6.2 RD Optimized Tree Structures

5. New Separable Constructions

5.1 Wavelet Footprints [DragottiV:03]

Can we ‘’fix’’ the wavelet scenario?

That is, achieve the same rate-distortion performance as an oracle or a dynamic programming method but with the complexity of wavelet methods?

The structure of wavelet representation of singularities is simple:

- location: random
- structure accross scales: deterministic!

Data structure to capture discontinuities in wavelet domain

- in orthogonal expansion
- in frame

This leads to a simple and intuitive data structure

Wavelet Footprints

The wavelet footprint

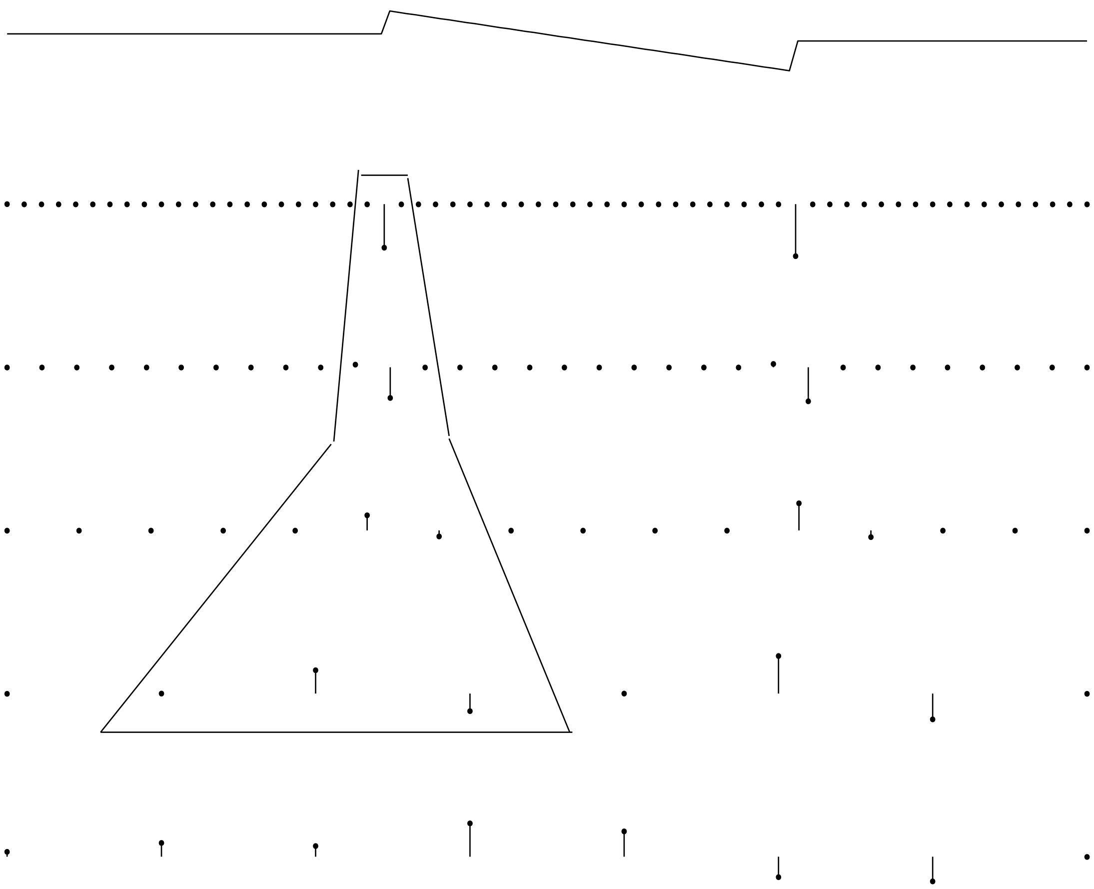

- this is the signature of the discontinuity
- behaviour well understood (classic wavelet analysis)
- vector of wavelet coefficients, or singularity subspace
- Approximation Theorem (piecewise Lipschitz α):
  Piecewise smooth = piecewise polynomial + residue in L-α!

Denoising (use coherence across scale)

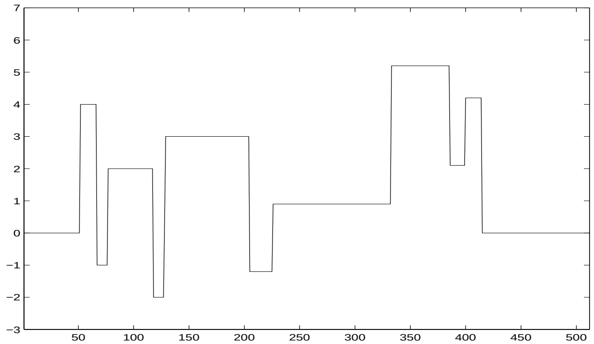

Original signal

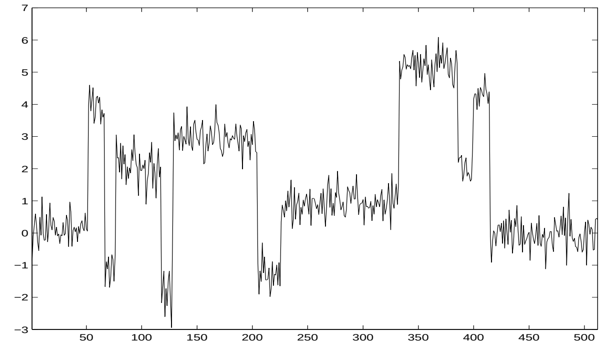

Noisy Signal (SNR=15.62dB)

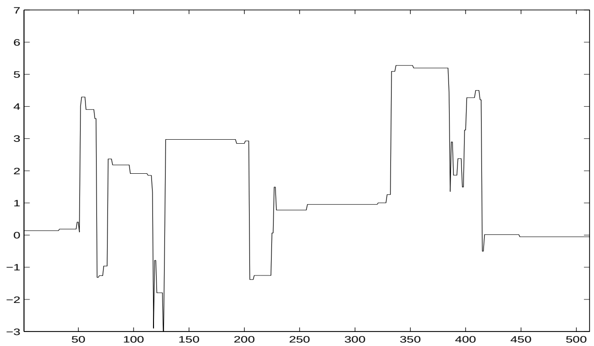

Hard-Thresholding (SNR=21.3dB)

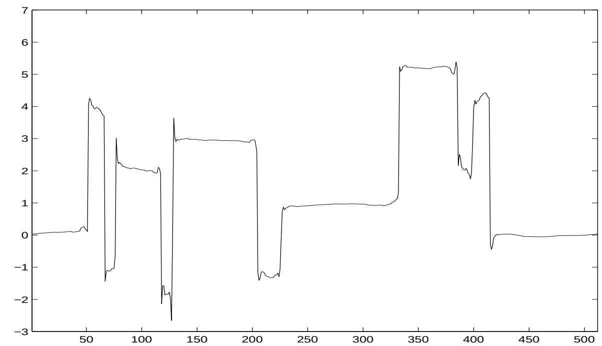

Cycle-Spinning (SNR=25.4dB)

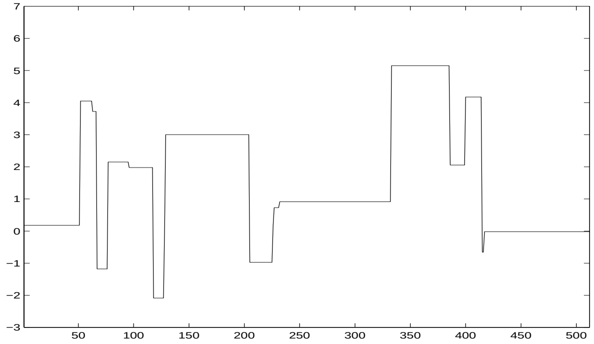

Denoising with Footprints (SNR=27.2dB)

This is a vector thresholding method adpated to wavelet singularities

Elementary compression algorithm based on footprints

- start at finest scale where discontinuity can be detected
- vector quantize corresponding footprint
- send position/footprint index

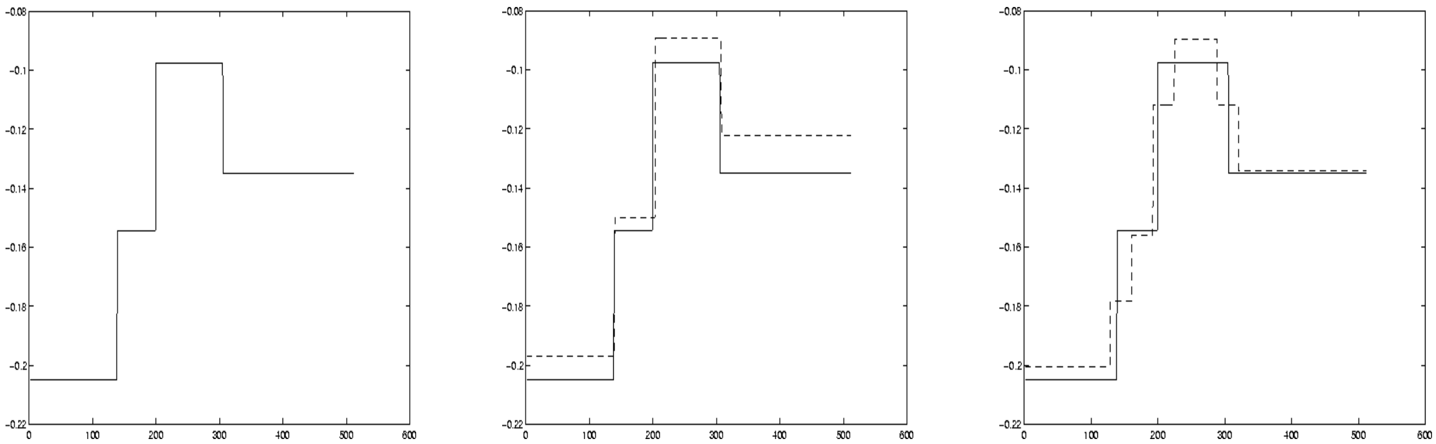

original — footprint 22.4dB, 0.1 b/p — SPIHT 23dB, 0.34 b/p

- gain: takes advantage of coherence accross scales
- initial results are encouraging!
- can be shown to attain oracle performance for piecewise polynomials

wavelet coder: $D(R) = c_0 \cdot \sqrt{R} \cdot 2^{-(C \cdot \sqrt{R})}$

footprint coder: $D(R) = c_2 \cdot 2^{-(c_3 \cdot R)}$

Compression (zerotrees versus footprints)

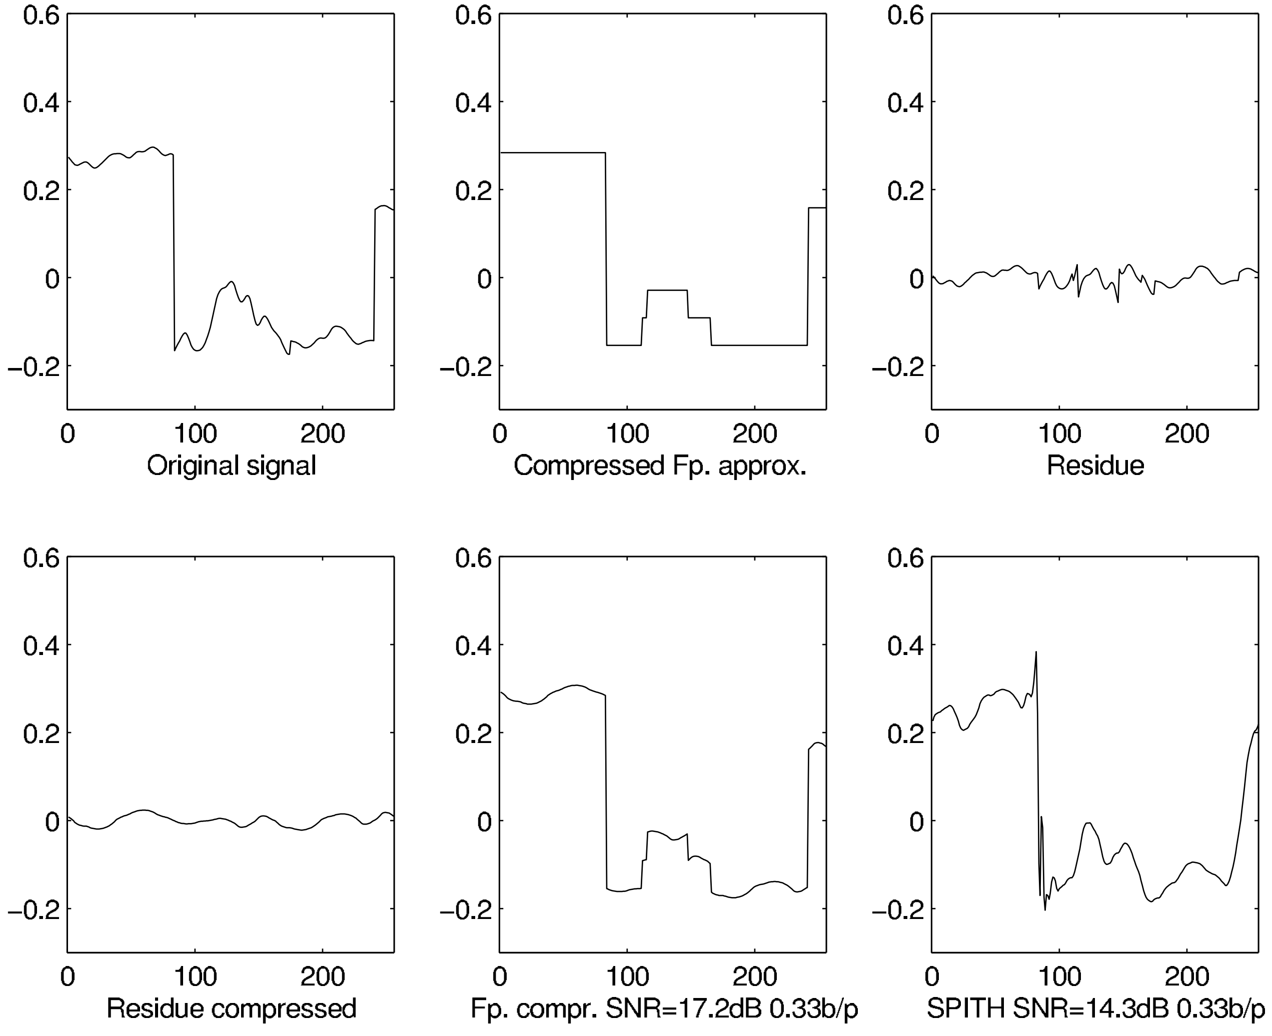

The piecewise smooth nature is well preserved

5.2 Directional Wavelet Transforms and Frames [VelisavljevicDV:03]

Classic 2D wavelet transform:

- separable horizontal-vertical transform
- good for ... horizontal and vertical directional elements

What if:

- simplicity of 2D wavelet transform
- more directionality
- adaptive directionality

=> directional wavelet transform

- discrete directions
- directional frames: good for denoising
- directional bases: good for NLA

Idea: use digital lines with ‘‘rational’’ directions

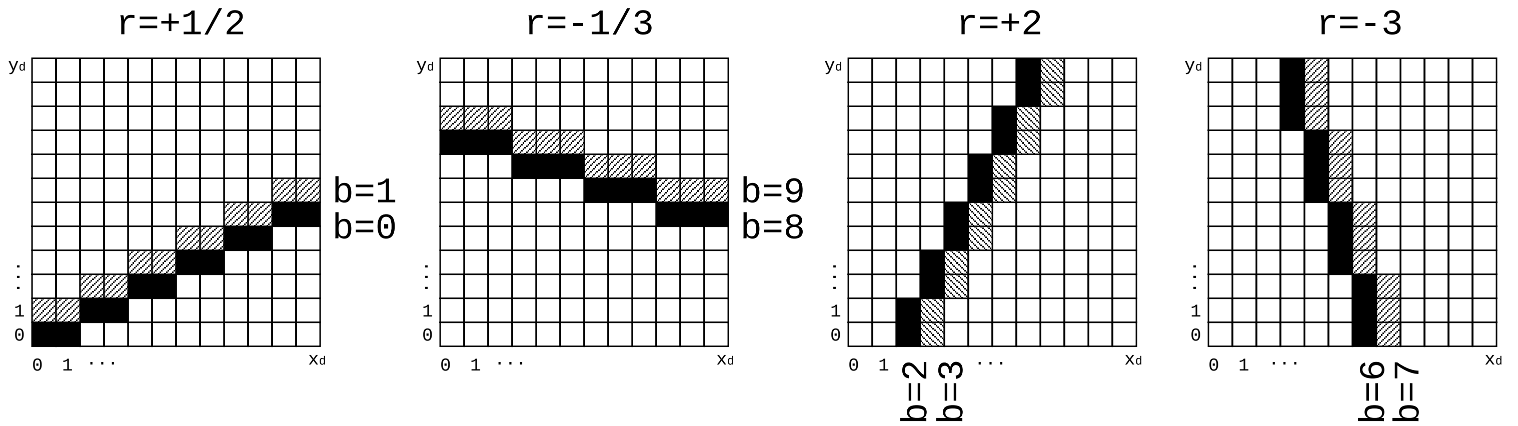

Such digital lines extend the usual ‘’horizontal-vertical’’ framework of the standard separable 2D wavelet transform

Approximation:

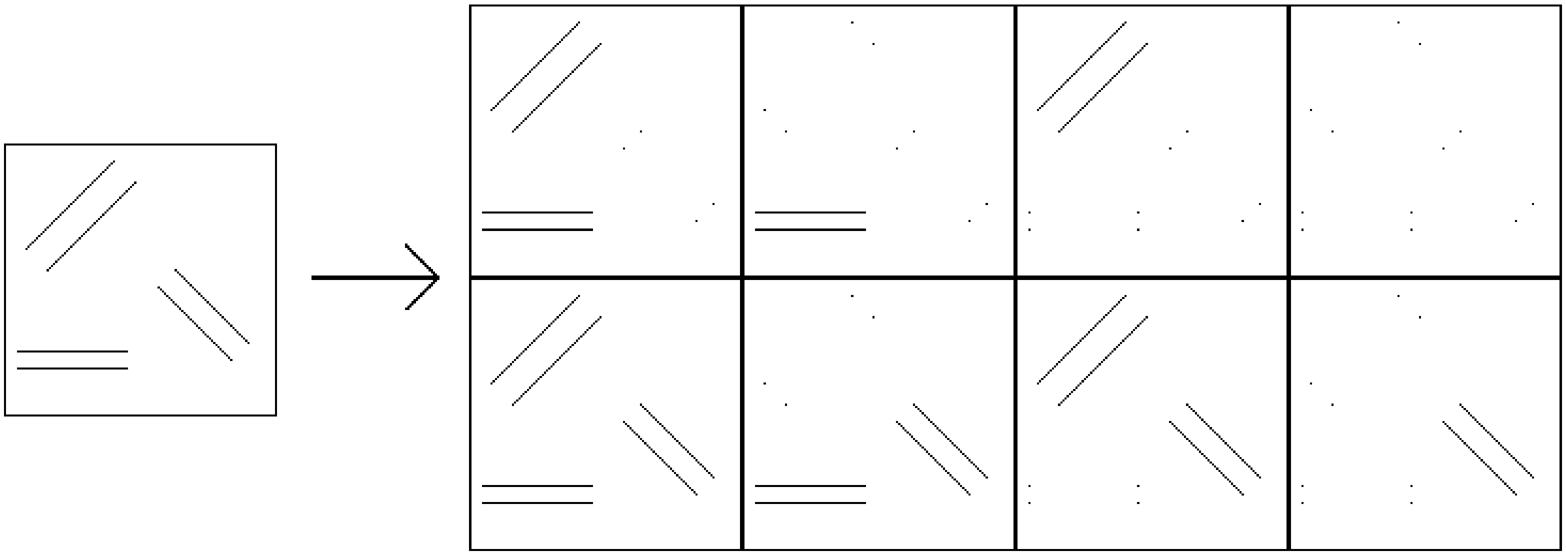

Simple example:

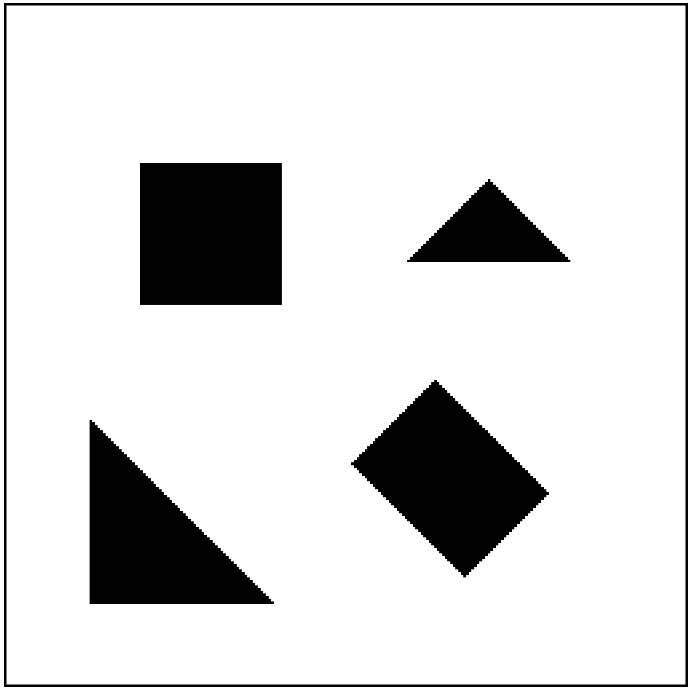

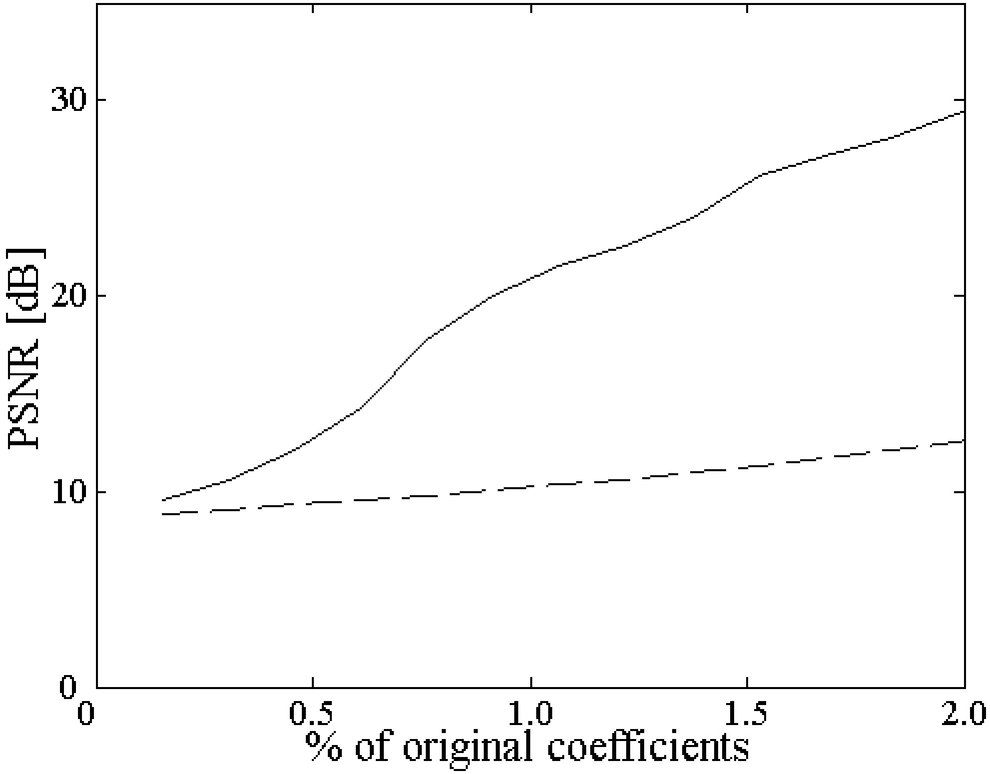

Thus, NLA in a directional bases outperforms usual wavelet NLA!

Denoising

Details:

6. New Non-Separable Constructions

Going to two dimensions requires non-separable bases

Objects in two dimensions we are interested in

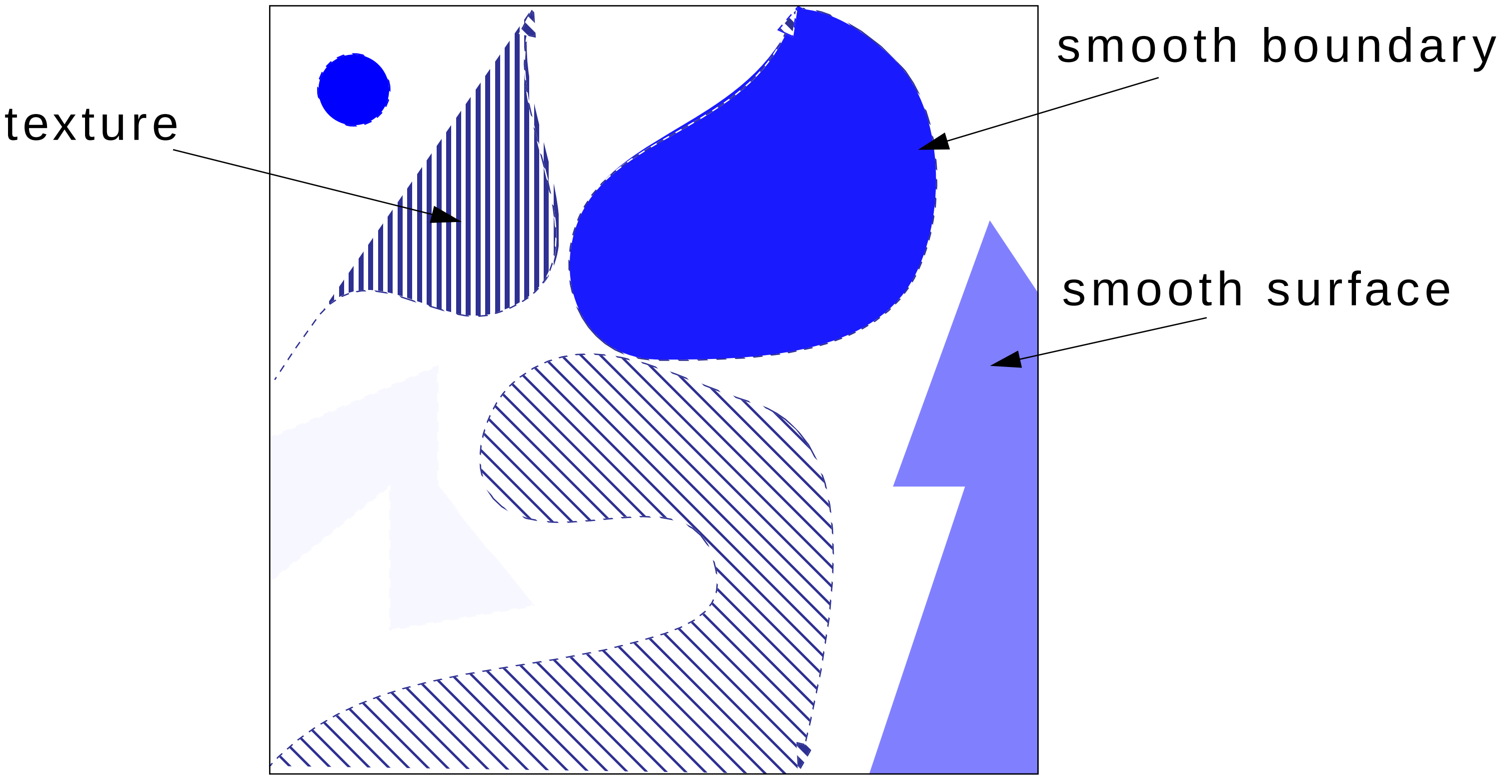

- textures: $D(R) = C_0 \cdot 2^{-2R}$ per pixel
- smooth surfaces: $D(R) = C_1 \cdot 2^{-2R}$ per object!

Current approaches to two dimensions....

Mostly separable, direct or tensor products

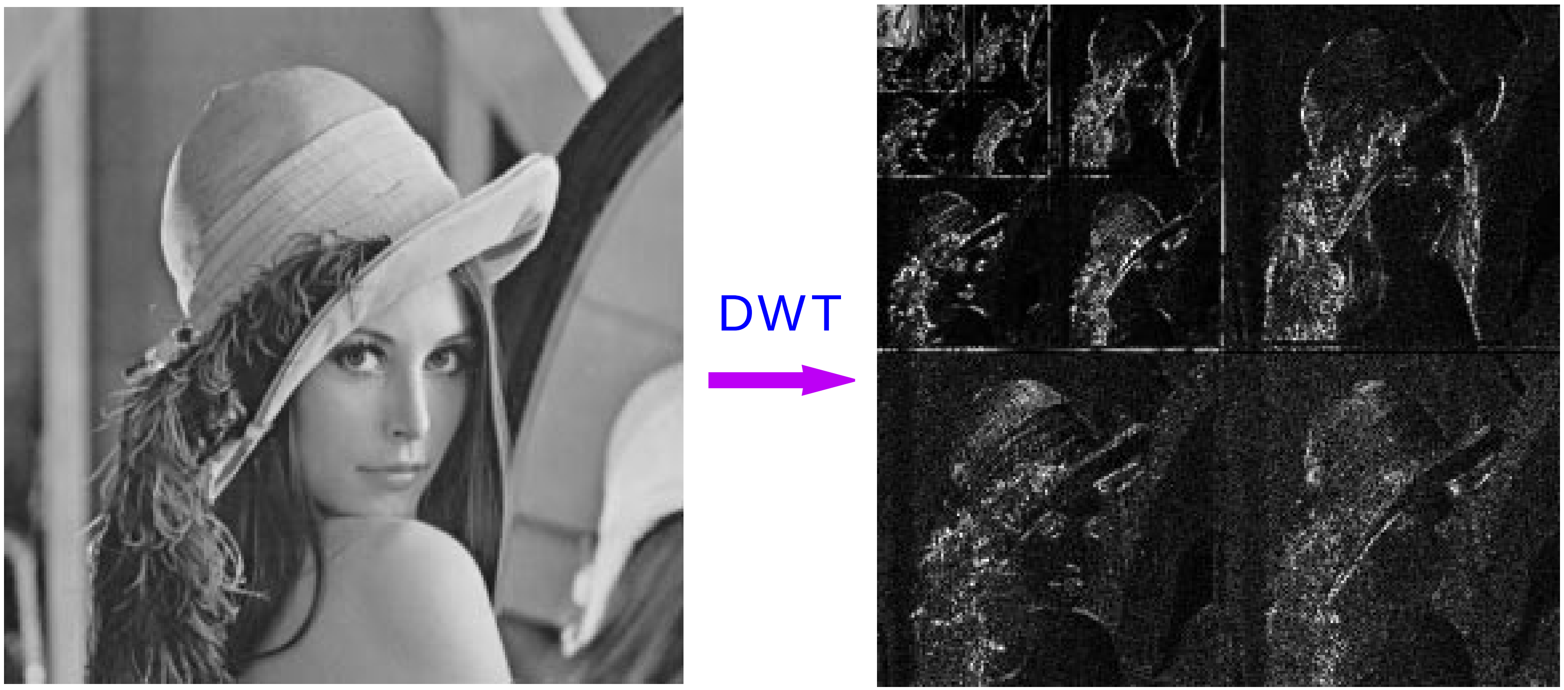

Wavelets: good for point singularities
but what is needed are sparse coding of edge singularities!

Nonseparable schemes

Approximation properties:

- wavelets good for point singularities
- ridgelets good for ridges
- curvelets good for curves

Consider object c^2 boundary between two csts

- # of wavelet coeffs: 2^j
- # of curvelet coeffs: $2^{j/2}$

Rate of approximation, M-term non-linear approximation (keeping the M largest coefficients in the expansion)

- Fourier: $O(M^{-1/2})$
- Wavelets: $O(M^{-1})$
- Curvelets: $O(M^{-2})$
- Bandelets: $O(M^{-\alpha})$

But how about:

compressionlets?

Multiresolution Contour Approximation

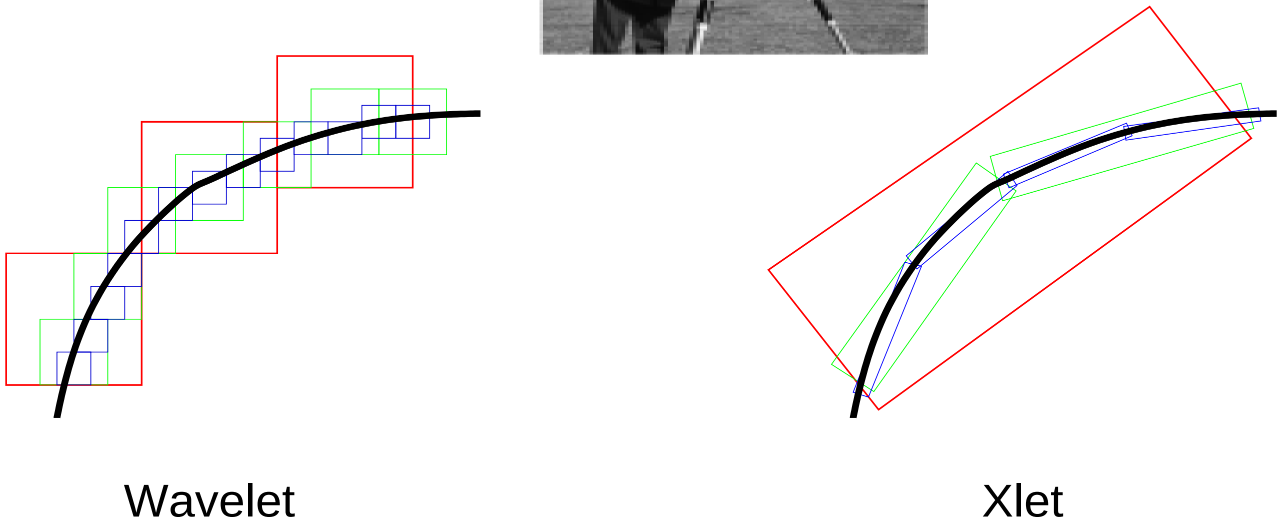

Consider object c^2 boundary between two cst

- # of wavelet coeffs: 2^j
- # of curvelet coeffs: $2^{j/2}$

Rate fo approximation, M-term non-linear approximation

- Fourier: $O(1/\sqrt{M})$
- Wavelets: $O(1/M)$
- Curvelets: $O(1/M^2)$

Objects we know how to compress....

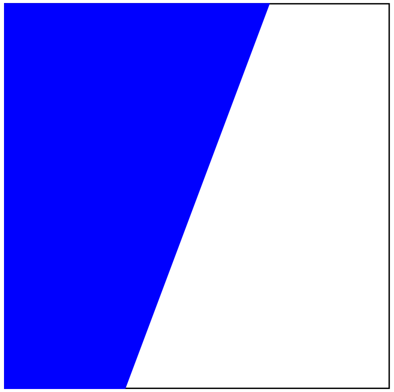

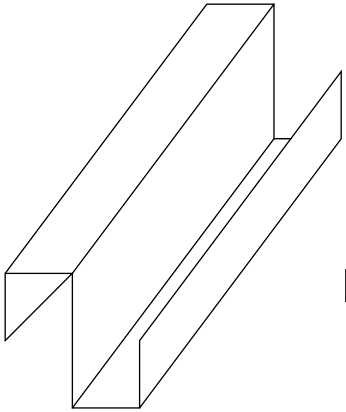

Basis element

Approximation

- Wavelets $E_M \sim 1/M$
- Ridgelets $E_M \sim 2^{-M}$

Rate/distortion

- Oracle $D(R) = C \cdot 2^{-2R}$
- Wavelets....poor
- Ridgelets....suboptimal
- adaptive schemes: close to oracle
- fixed basis: under investigation

6.1 Contourlets [DoV:03]

Idea: direct discrete-space construction that has good approximation properties for smooth functions with smooth boundaries

- directional analysis as in a Radon transform
- multiresolution as in wavelets and pyramids
- computationally easy
- bases or low redundancy frame

Directional Filter Banks [BambergerS:92, DoV:02]

- division of 2-D spectrum into fine slices using iterated tree structured filter banks

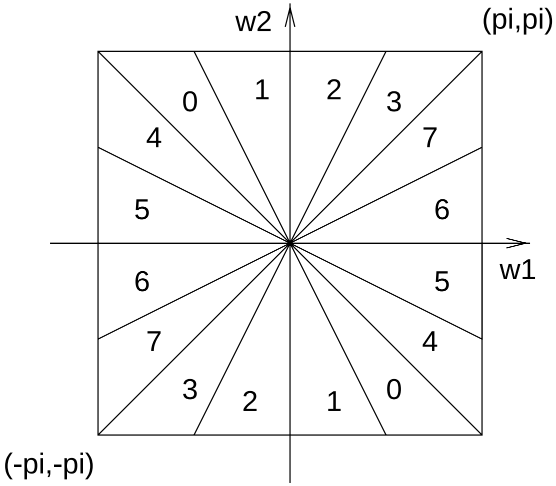

-

Iterated directional filter banks: efficient directional analysis

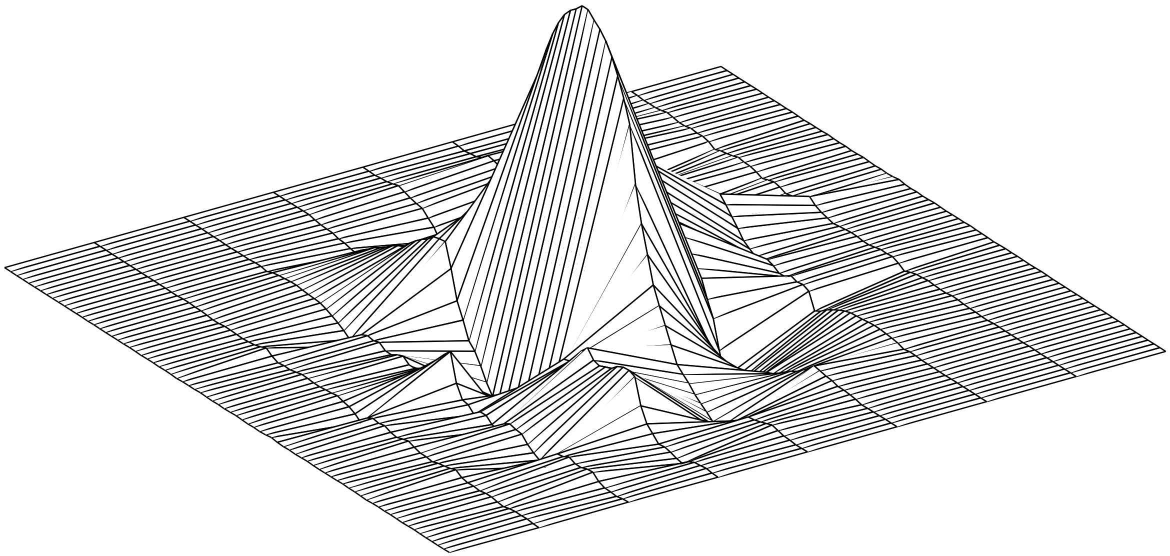

Example:

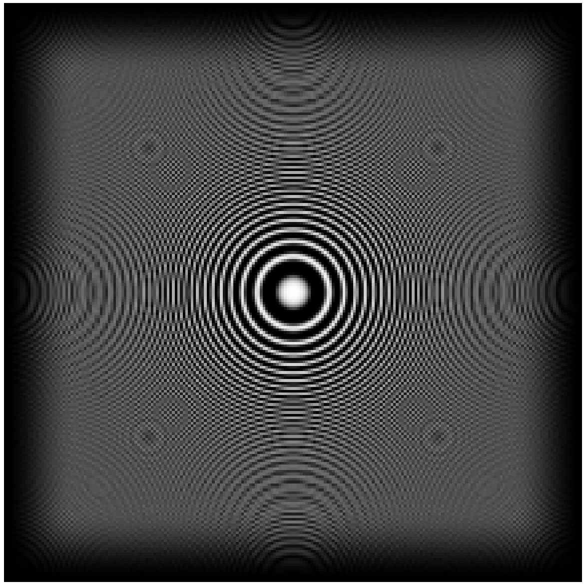

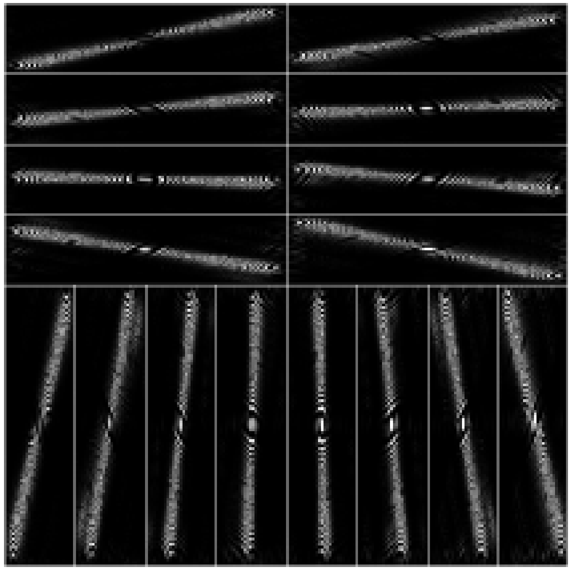

Example of basis functions [Haar]:

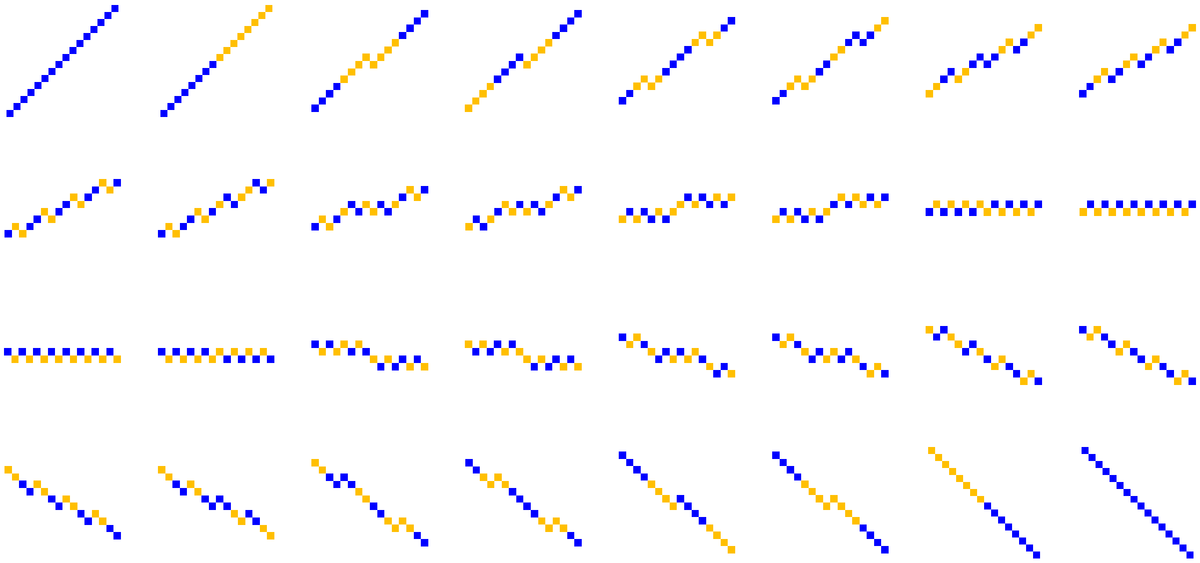

32 equivalent filters (-45..45 deg.) of a 6 level or 64 channel filter bank based on Haar filter and iterated quincunx sampling

Pyramidal Directional Filter Banks (PDFB) and contourlets

- Motivation: + add multiscale into the directional filter bank
  + improve its non-linear approximation power

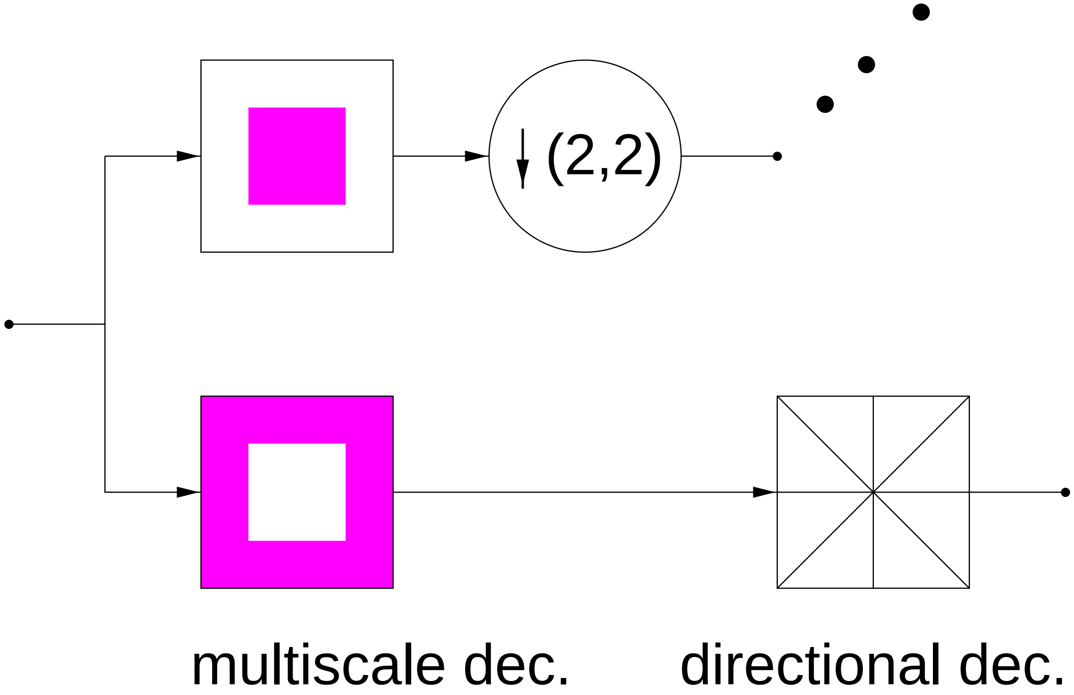

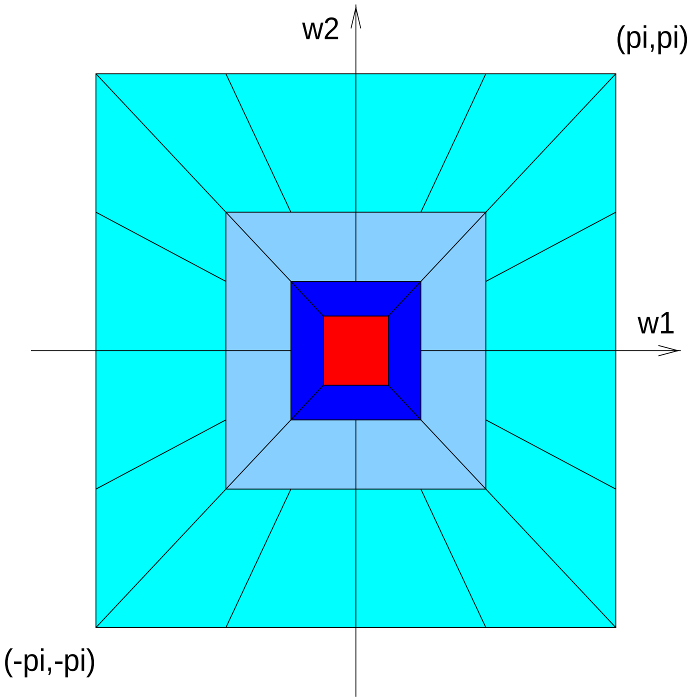

- Properties: + Flexible multiscale and directional representation for images (can have different number of direction at each scale!)

Resulting functions: contourlets!

Another view

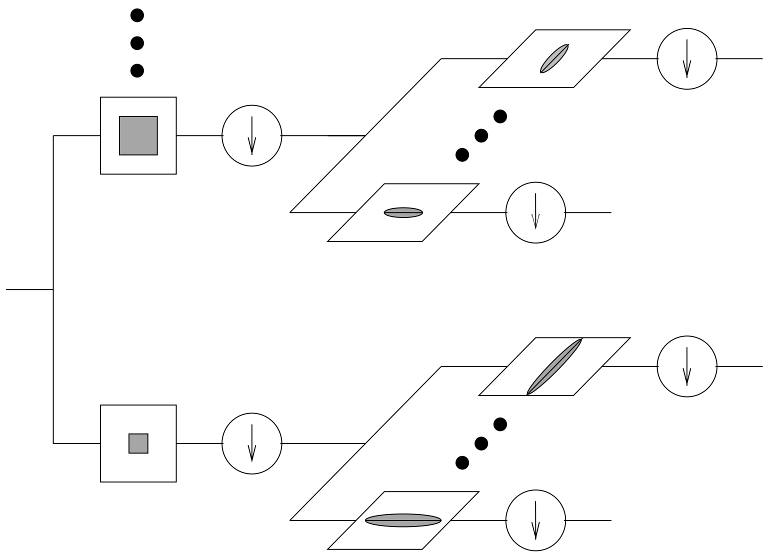

Depending on the scale (given by the pyramid) we have different number of directions

$$\text{width} = O(\text{lenth}^2)$$

LP: few directions, HP: more directions

Theorem [DoV:02]:
For a finite # of directions, this generates a tight frame for $L^2(R^2)$

A pyramidal directional filter bank

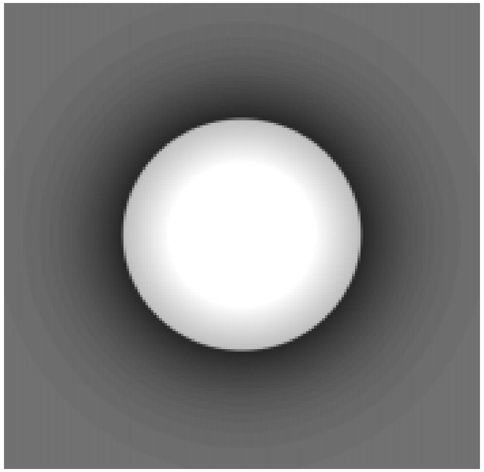

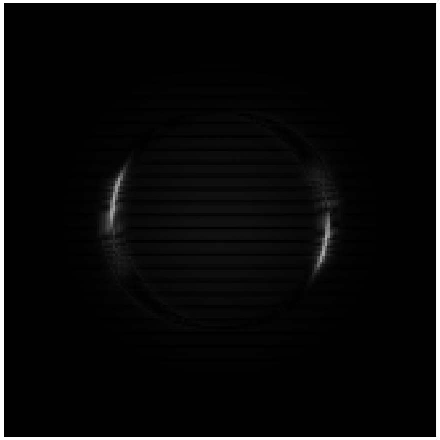

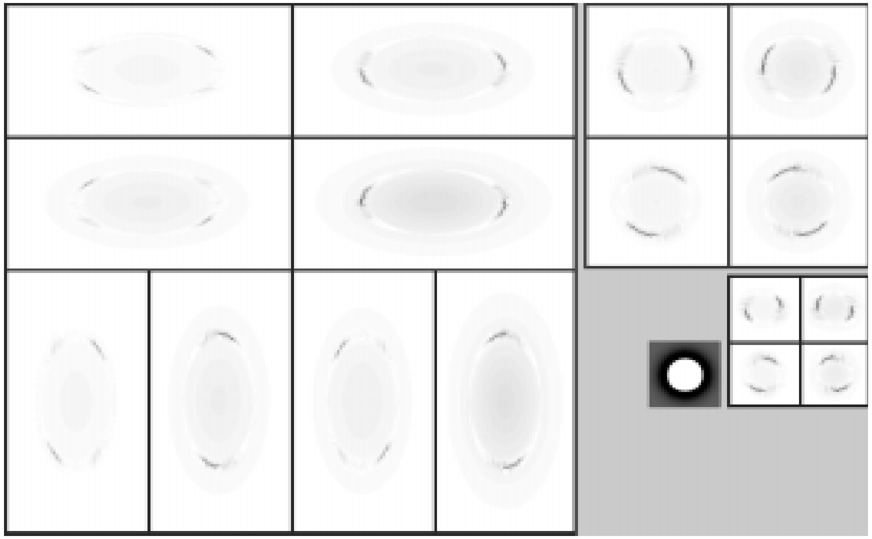

Compression, denoising, inverse problems: mostly open!

Approximation properties, Wavelets

M = 4, MSE = 1.68e−4 | M = 16, MSE = 1.66e−4 | M = 64, MSE = 1.60e−4 | M = 256, MSE = 1.44e−4

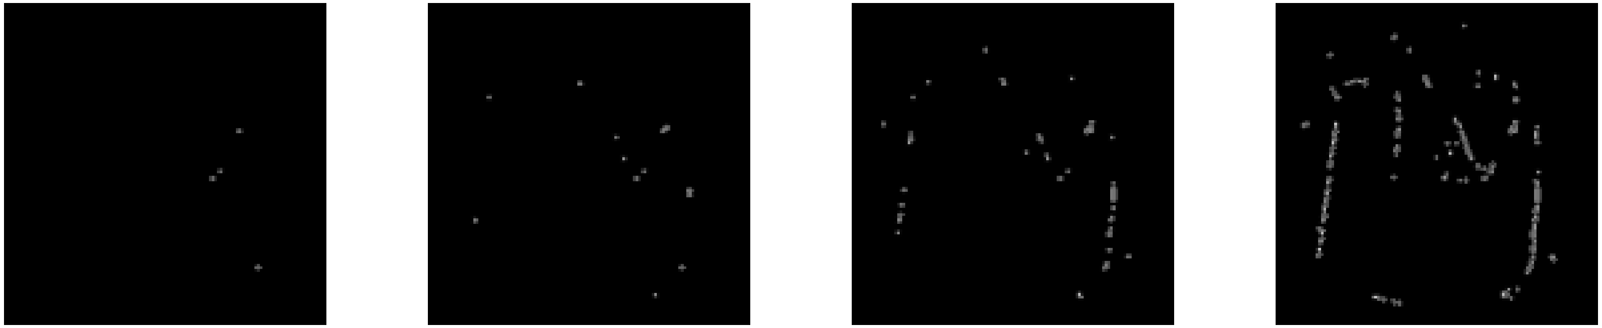

Approximation properties, Contourlets

M = 4, MSE = 1.68e−4 | M = 16, MSE = 1.63e−4 | M = 64, MSE = 1.55e−4 | M = 256, MSE = 1.43e−4

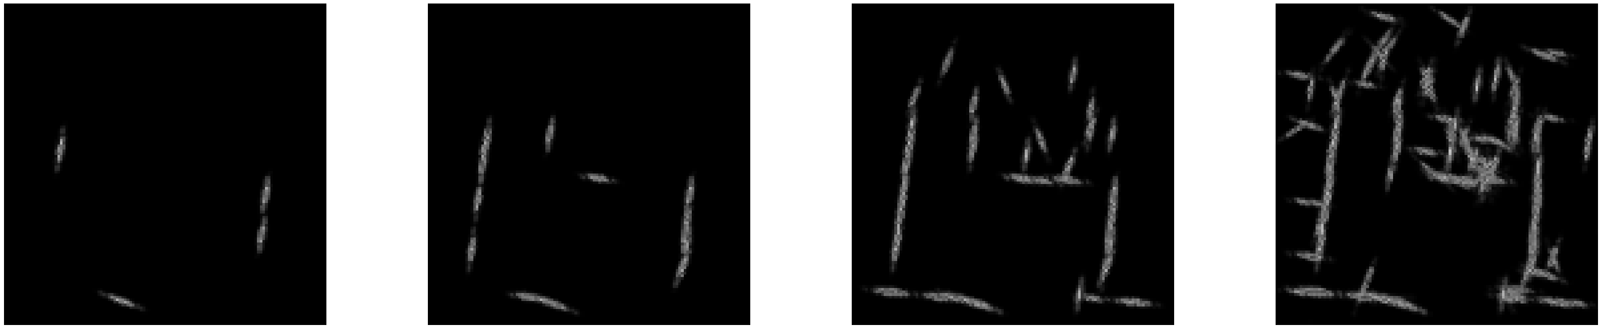

5.4 Tree Based Geometric Compression [ShuklaDDV:03]

Idea

- tree and quadtree algorithms popular
- many pruning algorithms
- concentrate on rate-distortion analysis and optimization
- new pruning and joining algorithm

Intuition:

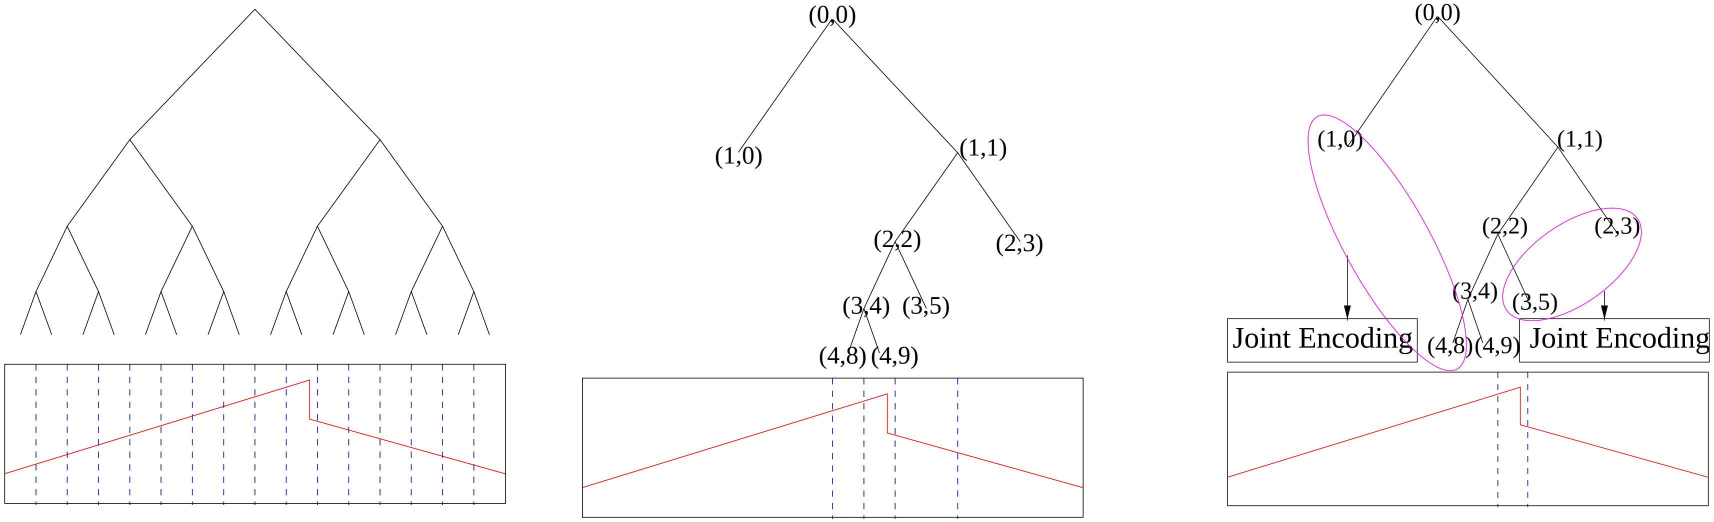

Results: Rate-distortion optimal for piecewise polynomials

$$D(R) = c_1 \cdot 2^{-(c_2 \cdot R)}$$

that is, like an oracle method (up to cosntants)

Extension to Quadtree:

- Example

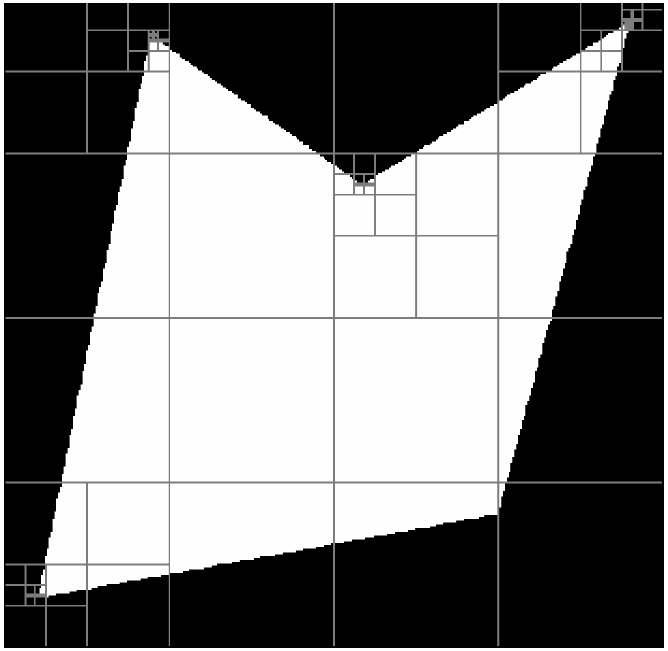

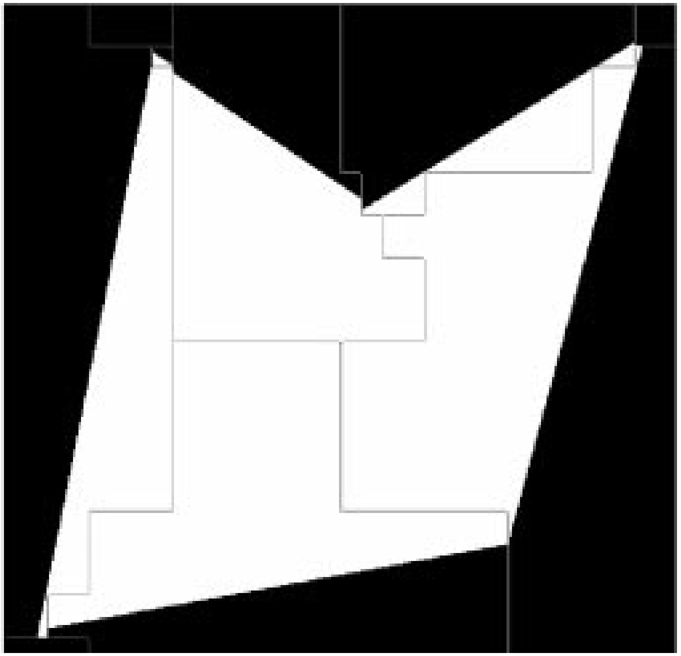

Results:

- consider a piecewise polynomial 2D signal, with polynomial boundaries, the following rate-distortion behavior is achieved

$$D(R) = c_3 \cdot 2^{-(c_4 \cdot R)}$$

- this is like an oracle method, and >> than prune algorithms which have a $\sqrt{R}$ penalty
- the method is of polynomila complexity

Bounds and experimental curves:

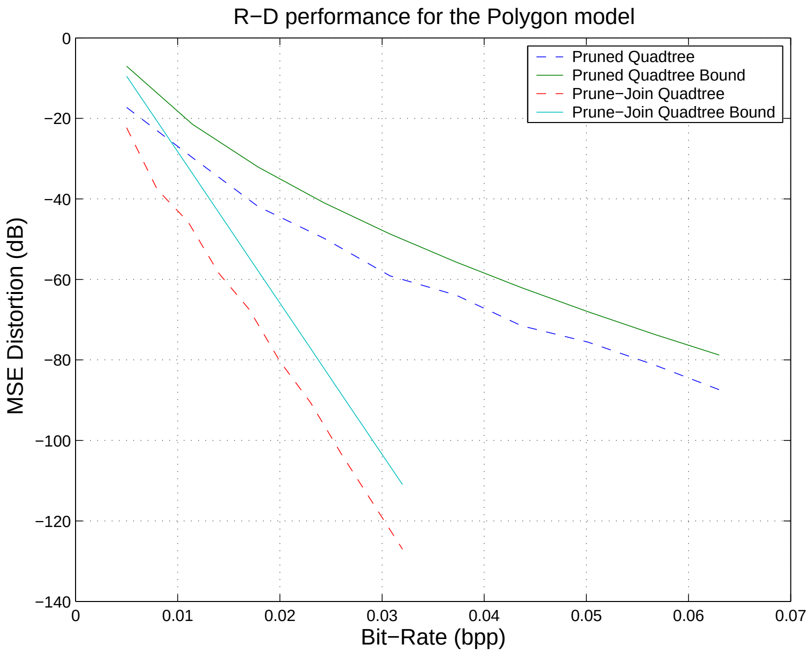

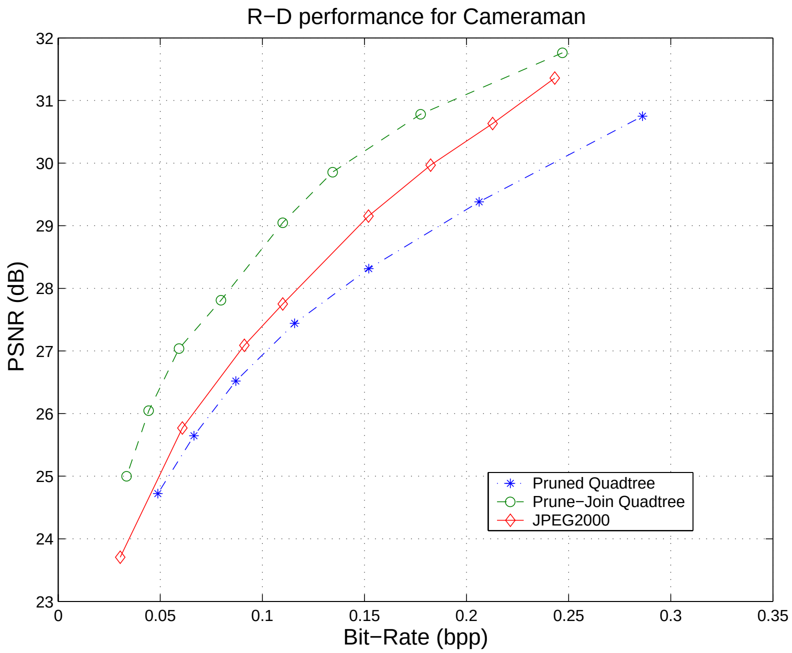

Polygon Model **Cameraman image**

Works well on images close to the model ...

Improves by 1dB at low rates on real images

Example of a geometric compression algorithm:

- **piecewise polynomial model for surface and boundaries**
- **polynomial fit to surface and to boundary on a quadtree**
- **rate-distortion optimal tree pruning and joining**

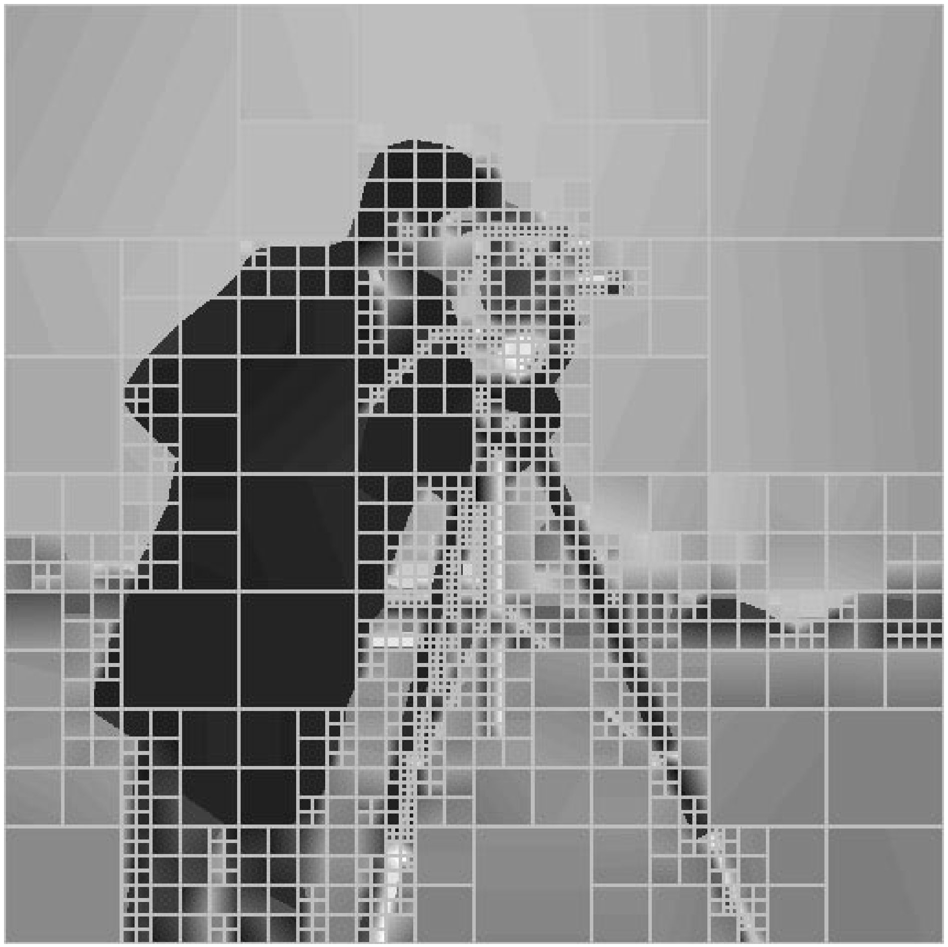

quadtree with R(D) pruning

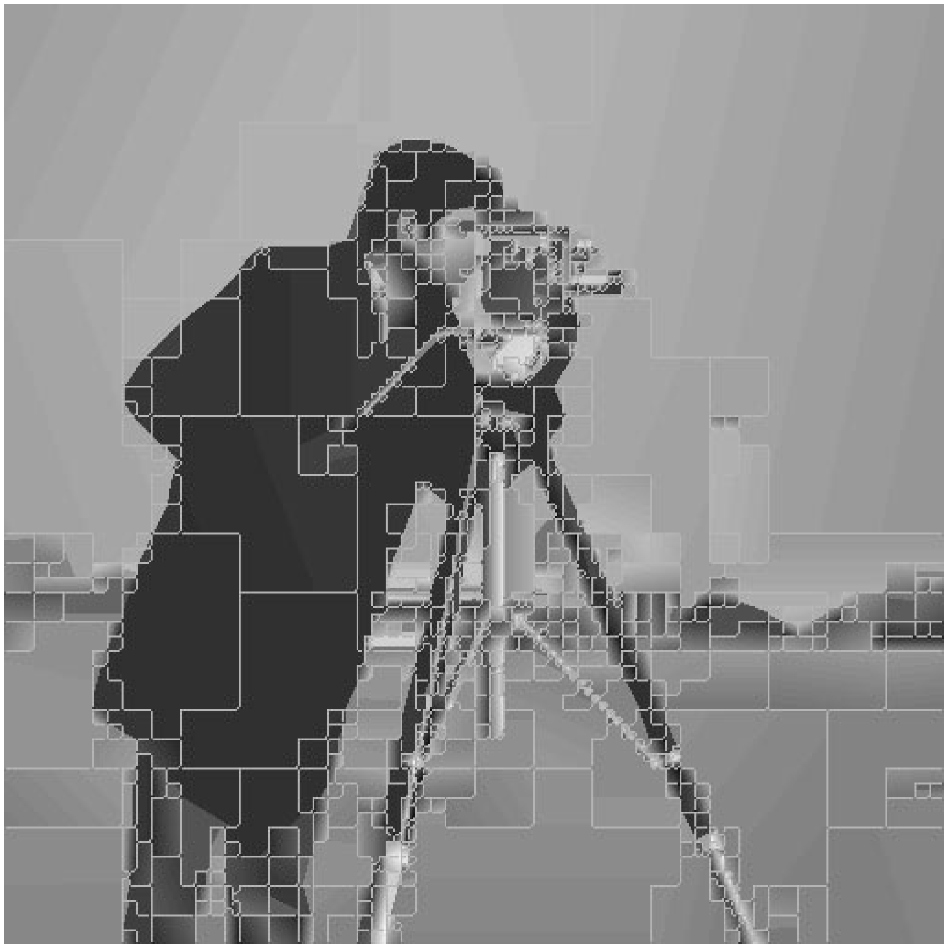

R(D) Joining of ‘’similar” leaves

Note: careful R(D) optimization!

Geometric Compression versus JPEG2000 at 0.11 bits/pixel, PSNR:

28.95

27.75

30.01

29.22

pruned-joined quadtree

JPEG2000

Details:

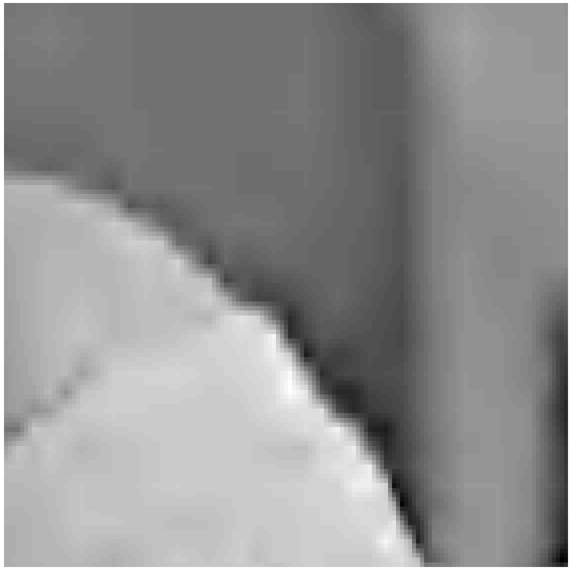

tree algorithm

JPEG2000

Conclusions on bases and approximations

Linear subspace approximation:

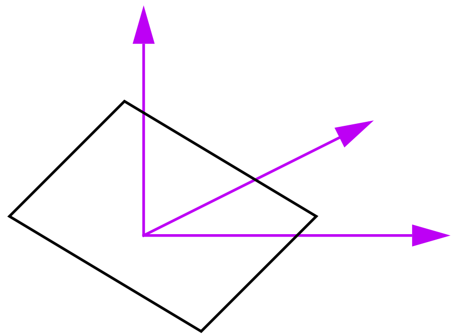

- given some prior,
  select a subspace onto which
  the data is projected
- no selection rate is spent
- example: KLT

Nonlinear subspace approximation

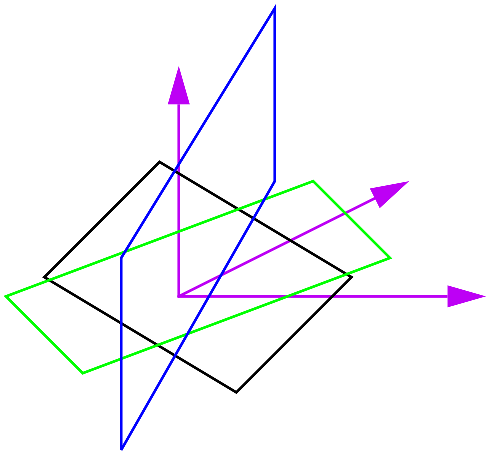

Many ways to choose

- adapted subset from a fixed basis (n term)
- subset from many possible bases
  (e.g. wavelet packets)
- subset from overcomplete set
  (e.g. matching pursuit or GSVQ)

$\Rightarrow$ **highly nonlinear approximations**
(e.g. manifold approximation)

Conclusions

Do we understand image compression better?

High rate: we know much more about certain 2D objects

- theoretical interest and for high resolution imagery

Low rate: there is still room at the bottom!

- rate-distortion analysis difficult, many dependencies still to exploit

New constructions

- **footprints**: squeezing out dependencies in wavelet coefficients
- **directional wavelet transforms**: simple but highly efficient
- **contourlets**: discrete, efficient directional multiscale bases with many propoerties of interest (e.g. NLA)
- **RD optimal quadtrees**: 1dB ahead...

Will there be a next image compression standard?

- maybe not, but there is plenty of high dimensional data out there in bad need of compression (surfaces, volumetric etc)
- if you know how to compress, you can denoise, enhance, index etc

Publications

For a tutorial:

- M. Vetterli, Wavelets, Approximation and Compression, IEEE Signal Processing Magazine, Sept. 2001

For more details

- P. L. Dragotti, M. Vetterli. Wavelets footprints: theory, algorithms and applications, IEEE Transactions on Signal Processing, to appear.
- V. Velisavljevic, P. L. Dragotti, M. Vetterli, "Directional Wavelet Transforms and Frames", Proceedings ICIP2002, Sept. 2002
- M. N. Do and M. Vetterli, Framing pyramids. IEEE Transactions on Signal Processing, to appear.
- M. N. Do and M. Vetterli, The finite ridgelet transform for image representation. IEEE Transactions on Image Processing, 2002, to appear.
- M. N. Do and M. Vetterli, Contourlets. in Beyond Wavelets, J. Stoeckler and G. V. Welland eds., Academic Press, 2002, to appear.
- R. Shukla, P. L. Dragotti, M. N. Do and M. Vetterli, Rate-distortion optimized tree structured compression algorithms for piecewise smooth images, submitted to IEEE Transactions Image Processing, 2002.