ECE 634: Digital Video Systems Wavelets: 2/21/17

Professor Amy Reibman
MSEE 356
reibman@purdue.edu

http://engineering.purdue.edu/~reibman/ece634/index.html

A short break to discuss wavelets

- Wavelet compression
- Wavelets are another way to combine pixels to create coefficients with improved entropy and/or correlation
- Wavelets are used for compression (of images); denoising; many other applications

Reading

- R. Gonzalez, "Digital Image Processing," Chapter 7 (Wavelets)
- A. Skodras, C. Christopoulos, T. Ebrahimi, The JPEG2000 Still Image Compression Standard, IEEE Signal Processing Magazine, vol. 18, pp. 36-58, Sept. 2001.
- B.E. Usevitch, "A tutorial on modern lossy wavelet image compression: Foundations of JPEG 2000," *IEEE Signal Processing Mag.*, vol. 18, pp. 22-35, Sept. 2001.

Wavelets

- A small wave whose energy is concentrated in time
- Comes from seismography, based on a small wave that results from a short-time impact
- The tool of wavelets developed to analyze seismic "wavelets"

- Continuous Time Wavelets, Discrete Wavelets

Many ways to talk about wavelets

- Mathematical elegance
- Multi-resolution pyramid analysis tool
- Tree-structured Subband decomposition

- All have their uses; a complete overview of wavelets should incorporate all of them

Multi-resolution pyramid analysis

1. Take image, decompose into
   a. low-resolution (coarse) approximation
   b. detail info that was lost due to the approximation
2. Take (1a), decompose into
   a. low-resolution (coarse) approximation
   b. detail info that was lost due to the approximation

Multi-resolution image analysis

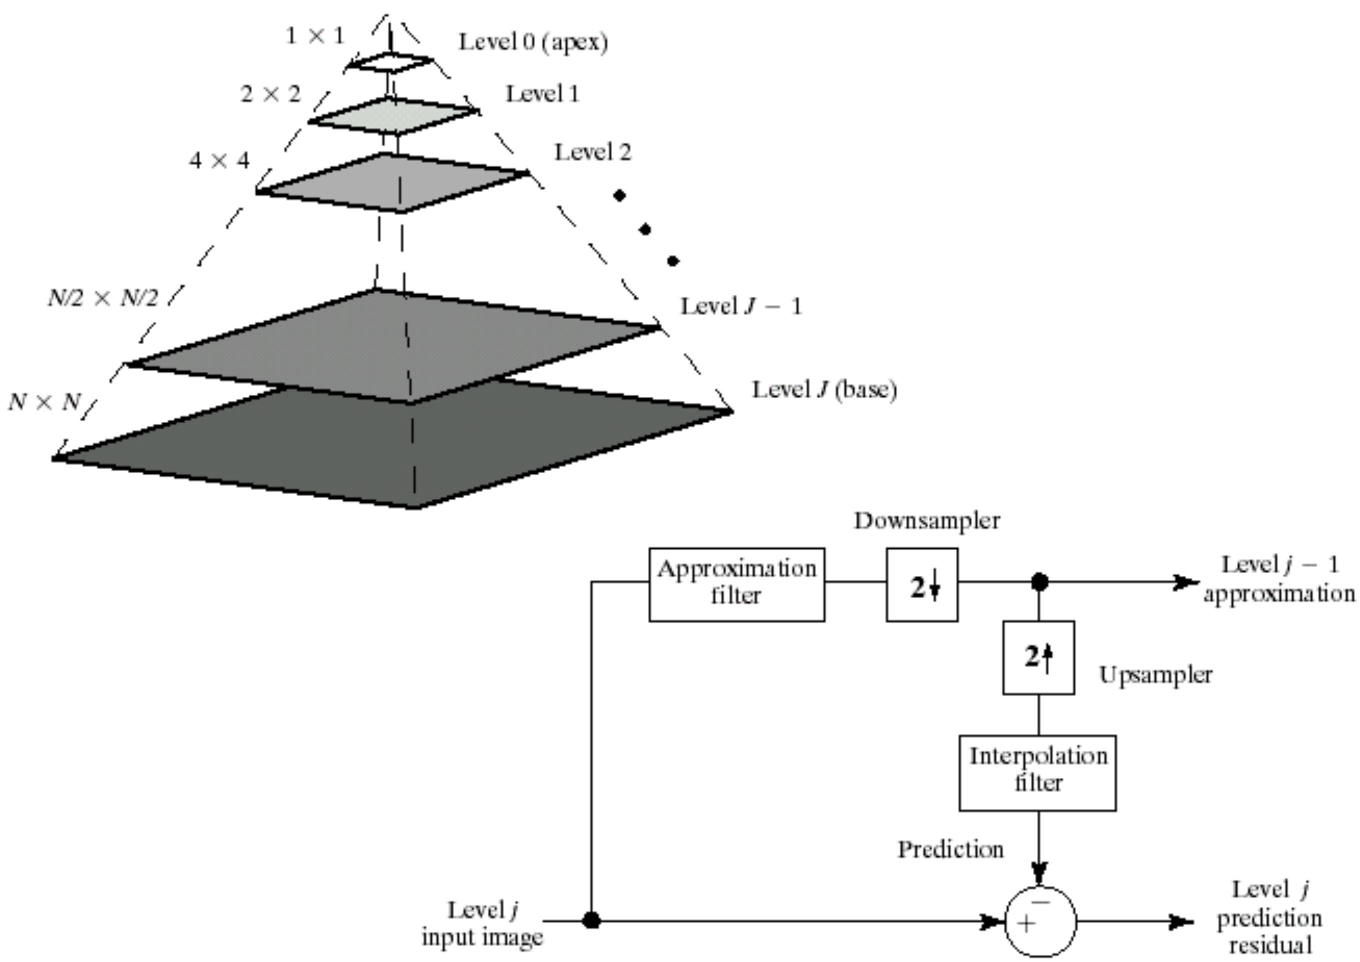

a b

FIGURE 7.2 (a) A pyramidal image structure and (b) system block diagram for creating it.

Gaussian and Laplacian Pyramids

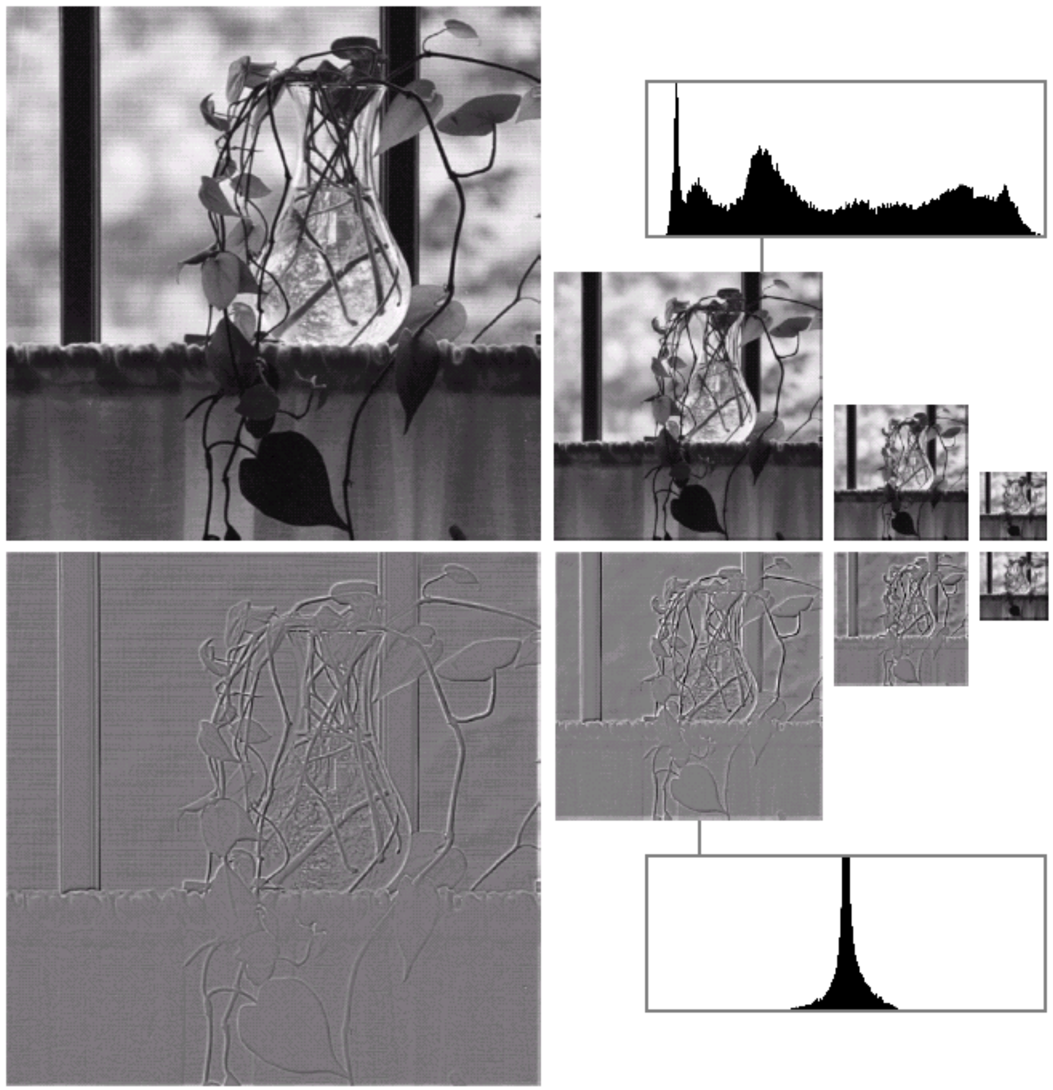

a
b

FIGURE 7.3 Two image pyramids and their statistics: (a) a Gaussian (approximation) pyramid and (b) a Laplacian (prediction residual) pyramid.

Subband decomposition

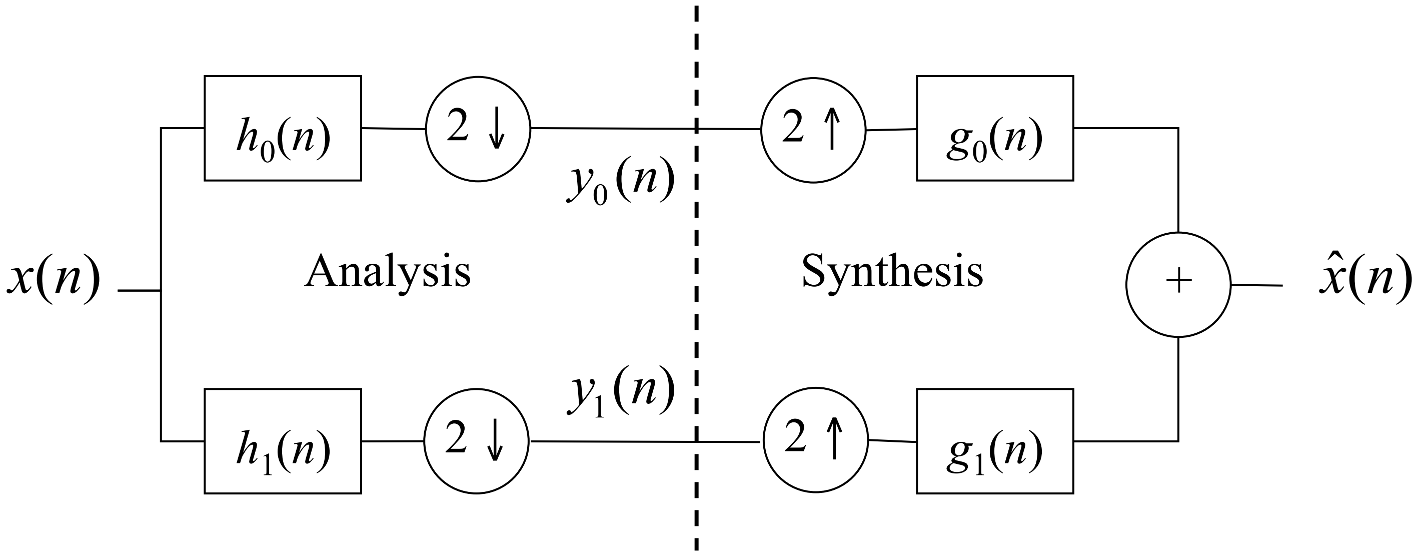

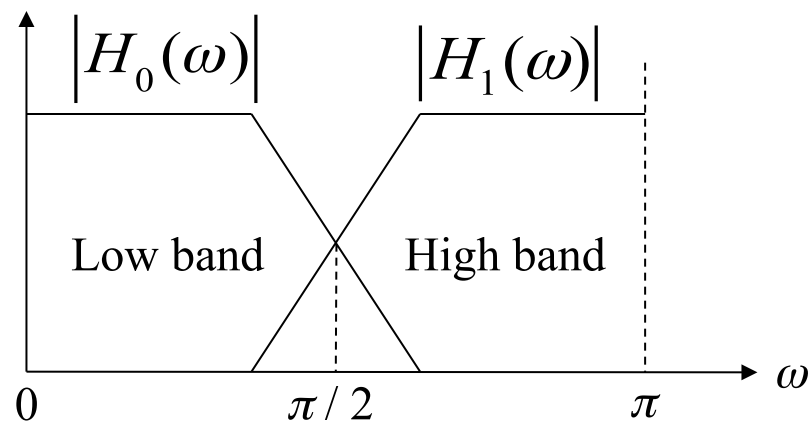

Multi-rate analysis

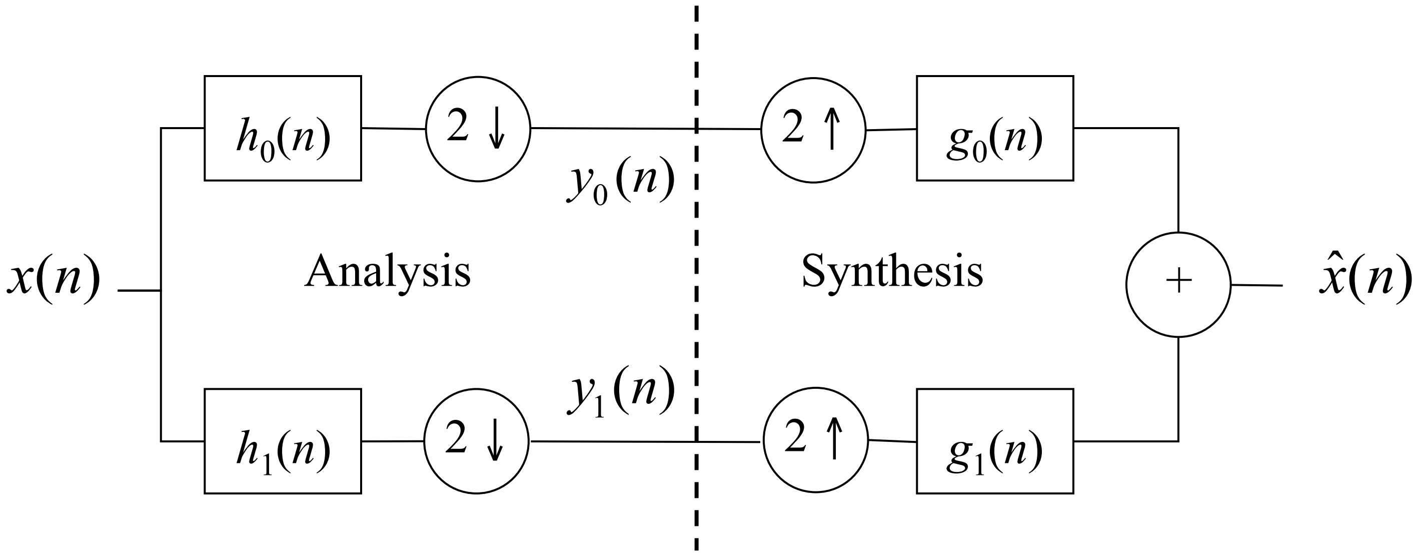

filter $X(z)H_o(z)$

decimate $$Y_0(z) = \frac{1}{2}\left[X(z^{1/2})H_o(z^{1/2}) + X(-z^{1/2})H_o(-z^{1/2})\right]$$

upsample $Y_o(z^2)$

filter $G_0(z)Y_o(z^2)$

Combining:

$$\hat{X}(z) = \tfrac{1}{2}\left[H_0(z)G_0(z) + H_1(z)G_1(z)\right]X(z)$$
$$+\tfrac{1}{2}\left[H_0(-z)G_0(z) + H_1(-z)G_1(z)\right]X(-z)$$

Subband decomposition

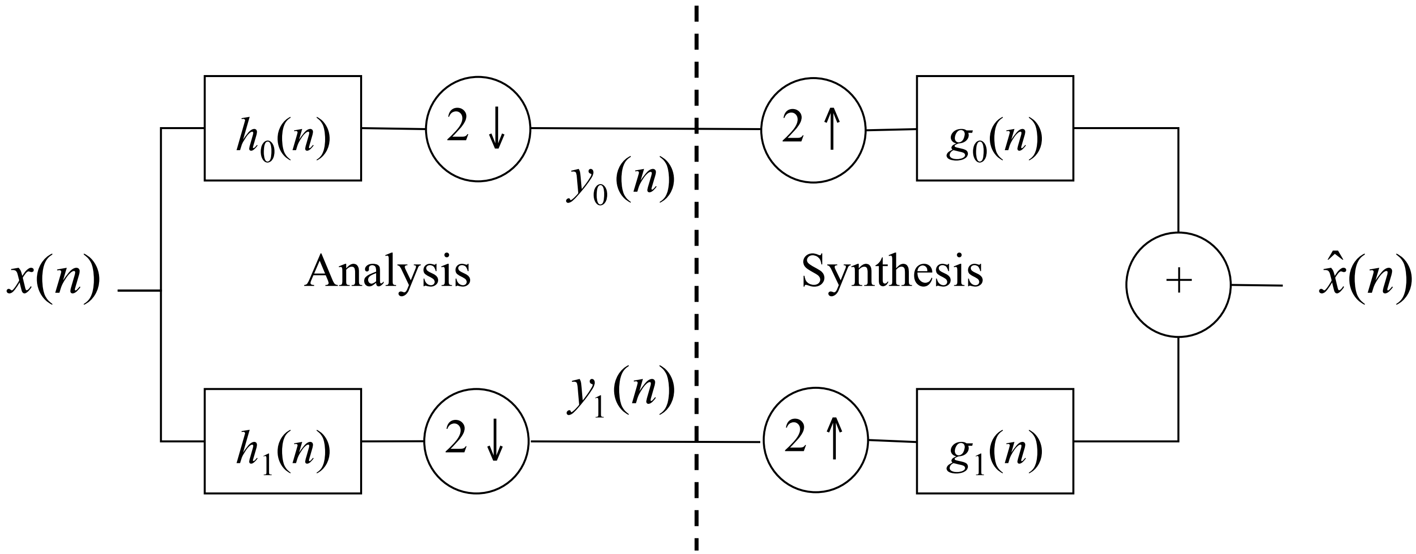

For error-free reconstruction

$$H_0(-z)G_0(z) + H_1(-z)G_1(z) = 0$$

$$H_0(z)G_0(z) + H_1(z)G_1(z) = 2$$

For finite impulse response (FIR) filters and ignoring the delay

$$g_0(n) = (-1)^n h_1(n)$$

$$g_1(n) = (-1)^{n+1} h_0(n)$$

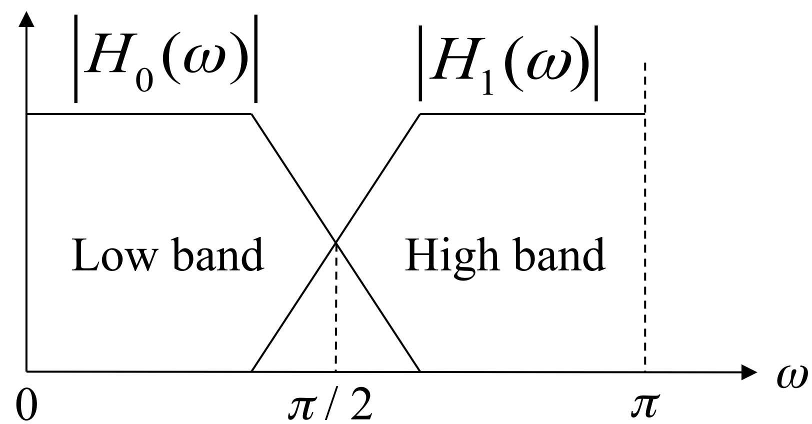

Tree-structured subband decomposition

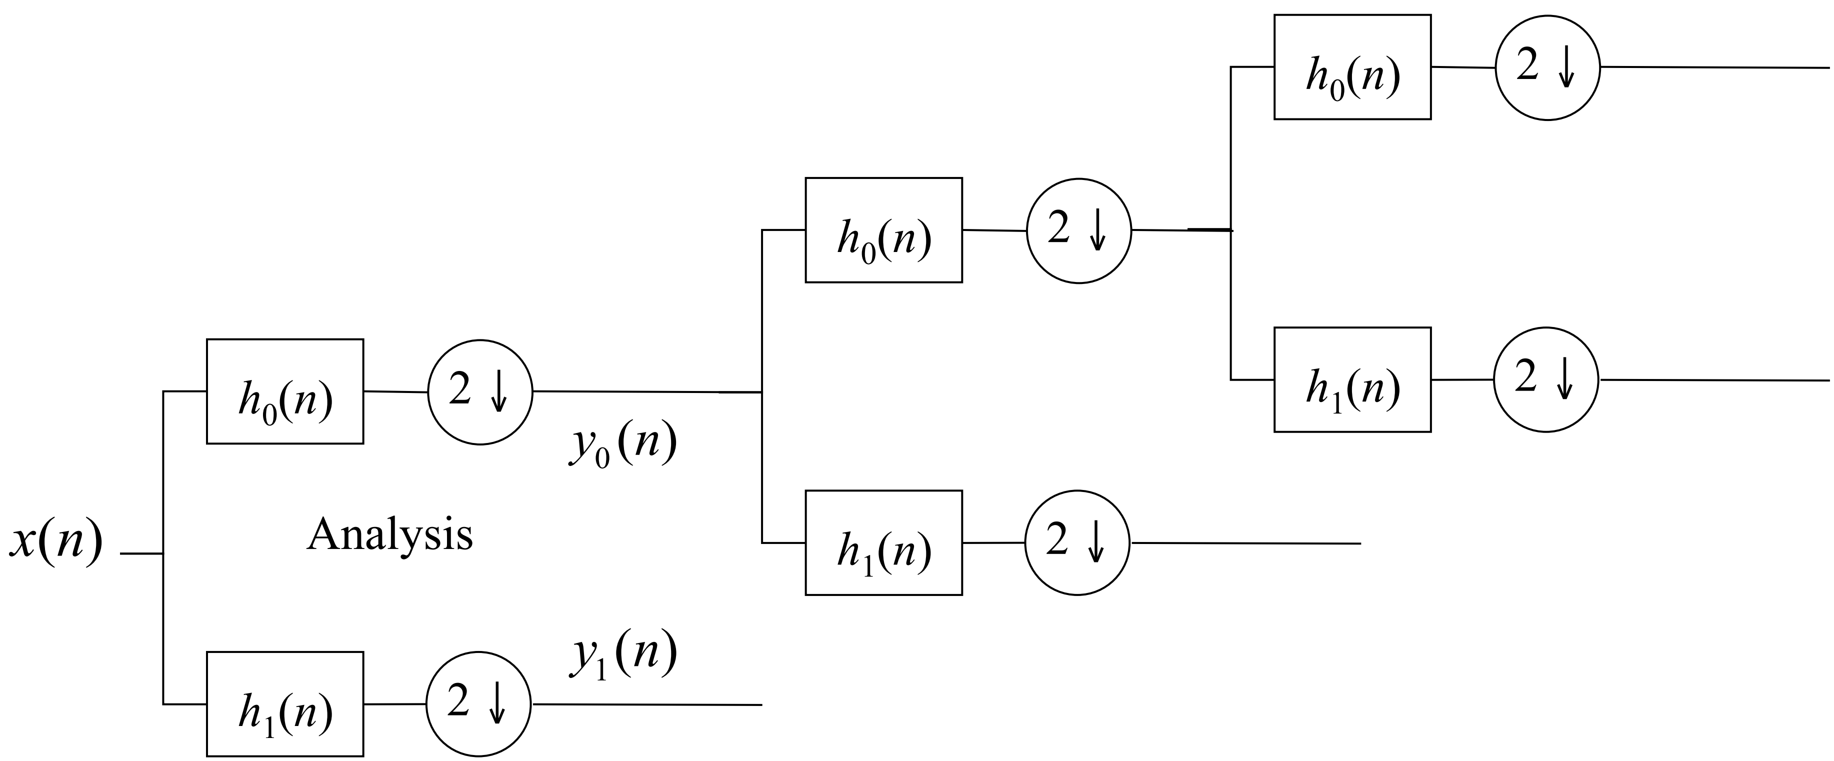

Subband Example

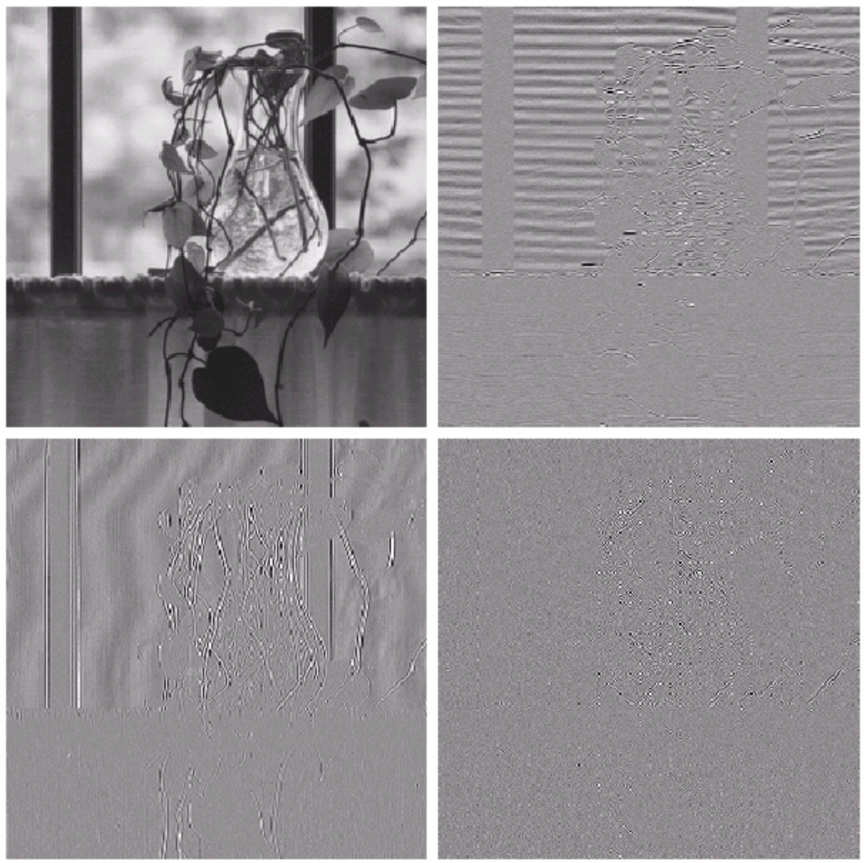

FIGURE 7.7 A four-band split of the vase in Fig. 7.1 using the subband coding system of Fig. 7.5.

LH: Low-pass Vertical,
high-pass horizontal

HL: high-pass Vertical,
low-pass horizontal

Comments

- Maximally decimated wavelet decomposition is a tree-structured subband decomposition
- The hierarchical decomposition AND the tree structure are powerful tools for many applications
- “Maximally decimated” is not required for all applications, and may cause problems

Wavelet Transform vs. Fourier Transform

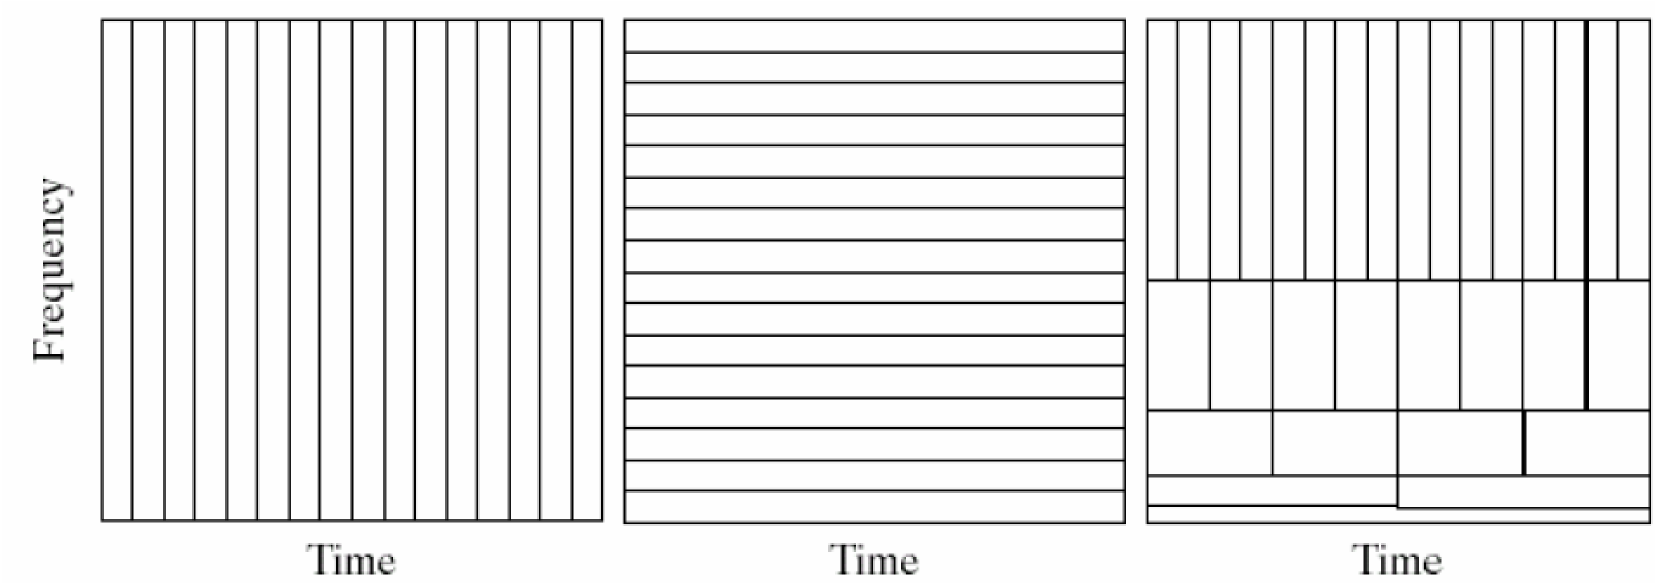

Time-frequency distribution for (a) sampled data, (b) FFT, and (c) FWT basis

Fourier transform: Basis function cover the entire signal range

Wavelet transform: Basis functions vary in frequency/scale and spatial extent

- High-frequency basis covers a smaller area
- Low-frequency basis covers a larger area

2D wavelet transform for images

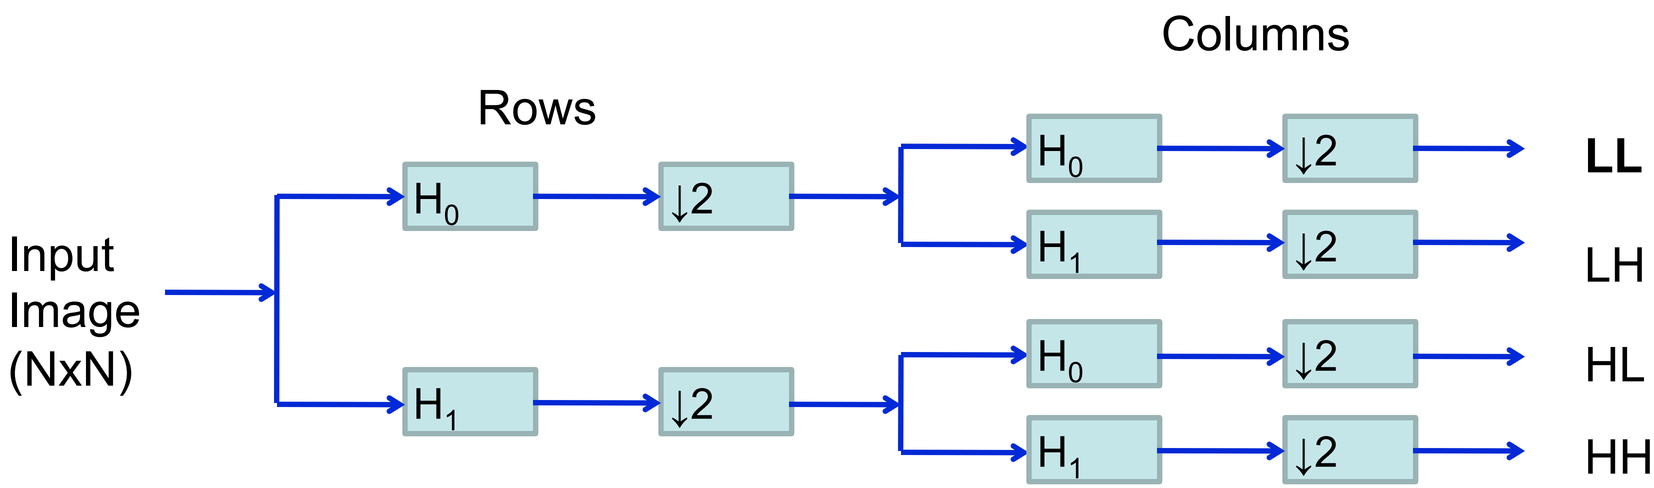

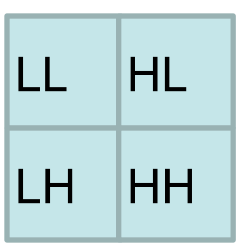

LL: low pass for both horizontal and vertical directions
LH: low pass for horizontal direction, high pass in vertical direction
HL: low pass for vertical direction, high pass in horizontal direction
HH: high pass in both horizontal and vertical directions

D. Gibbon and Z. Liu, 2013

Wavelets for images

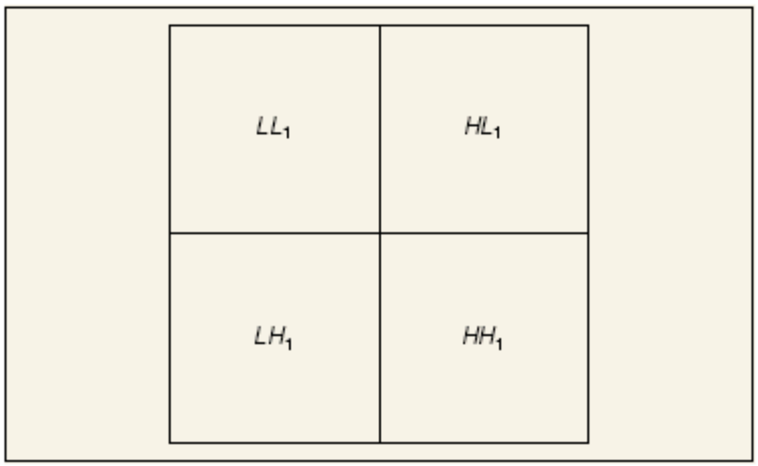

▲ 4. The subband labeling scheme for a one-level, 2-D wavelet transform.

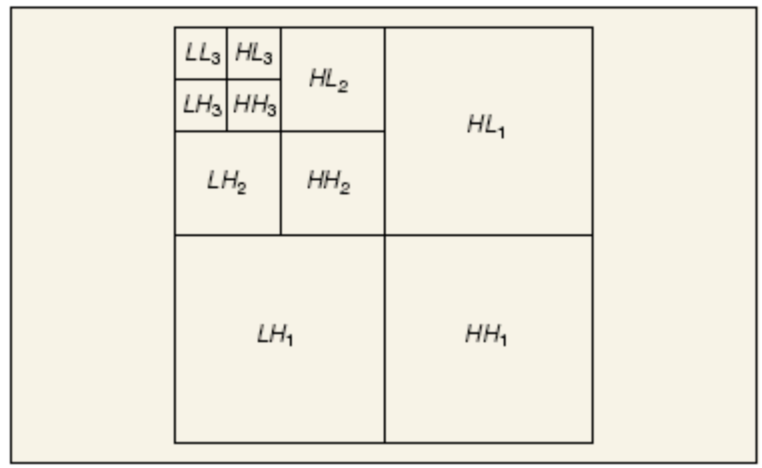

▲ 6. The subband labeling scheme for a three-level, 2-D wavelet transform.

From [Usevitch01]

2D wavelet transform is accomplished by applying the 1D decomposition along rows of an image first, and then columns.

Wavelets for images

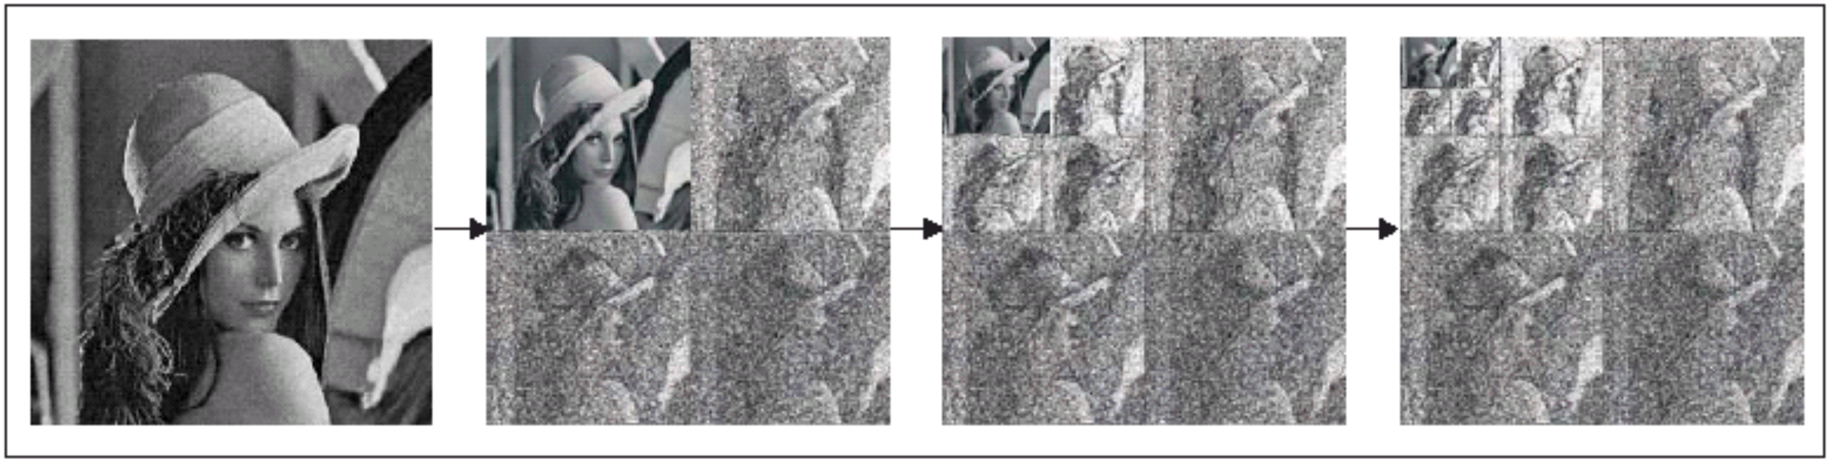

▲ *6. Three-level dyadic wavelet decomposition of the image "Lena."*

LL	HL
LH	HH

LL: low pass for both horizontal and vertical directions
LH: low pass for horizontal direction, high pass in vertical direction
HL: low pass for vertical direction, high pass in horizontal direction
HH: high pass in both horizontal and vertical directions

From [Usevitch01]

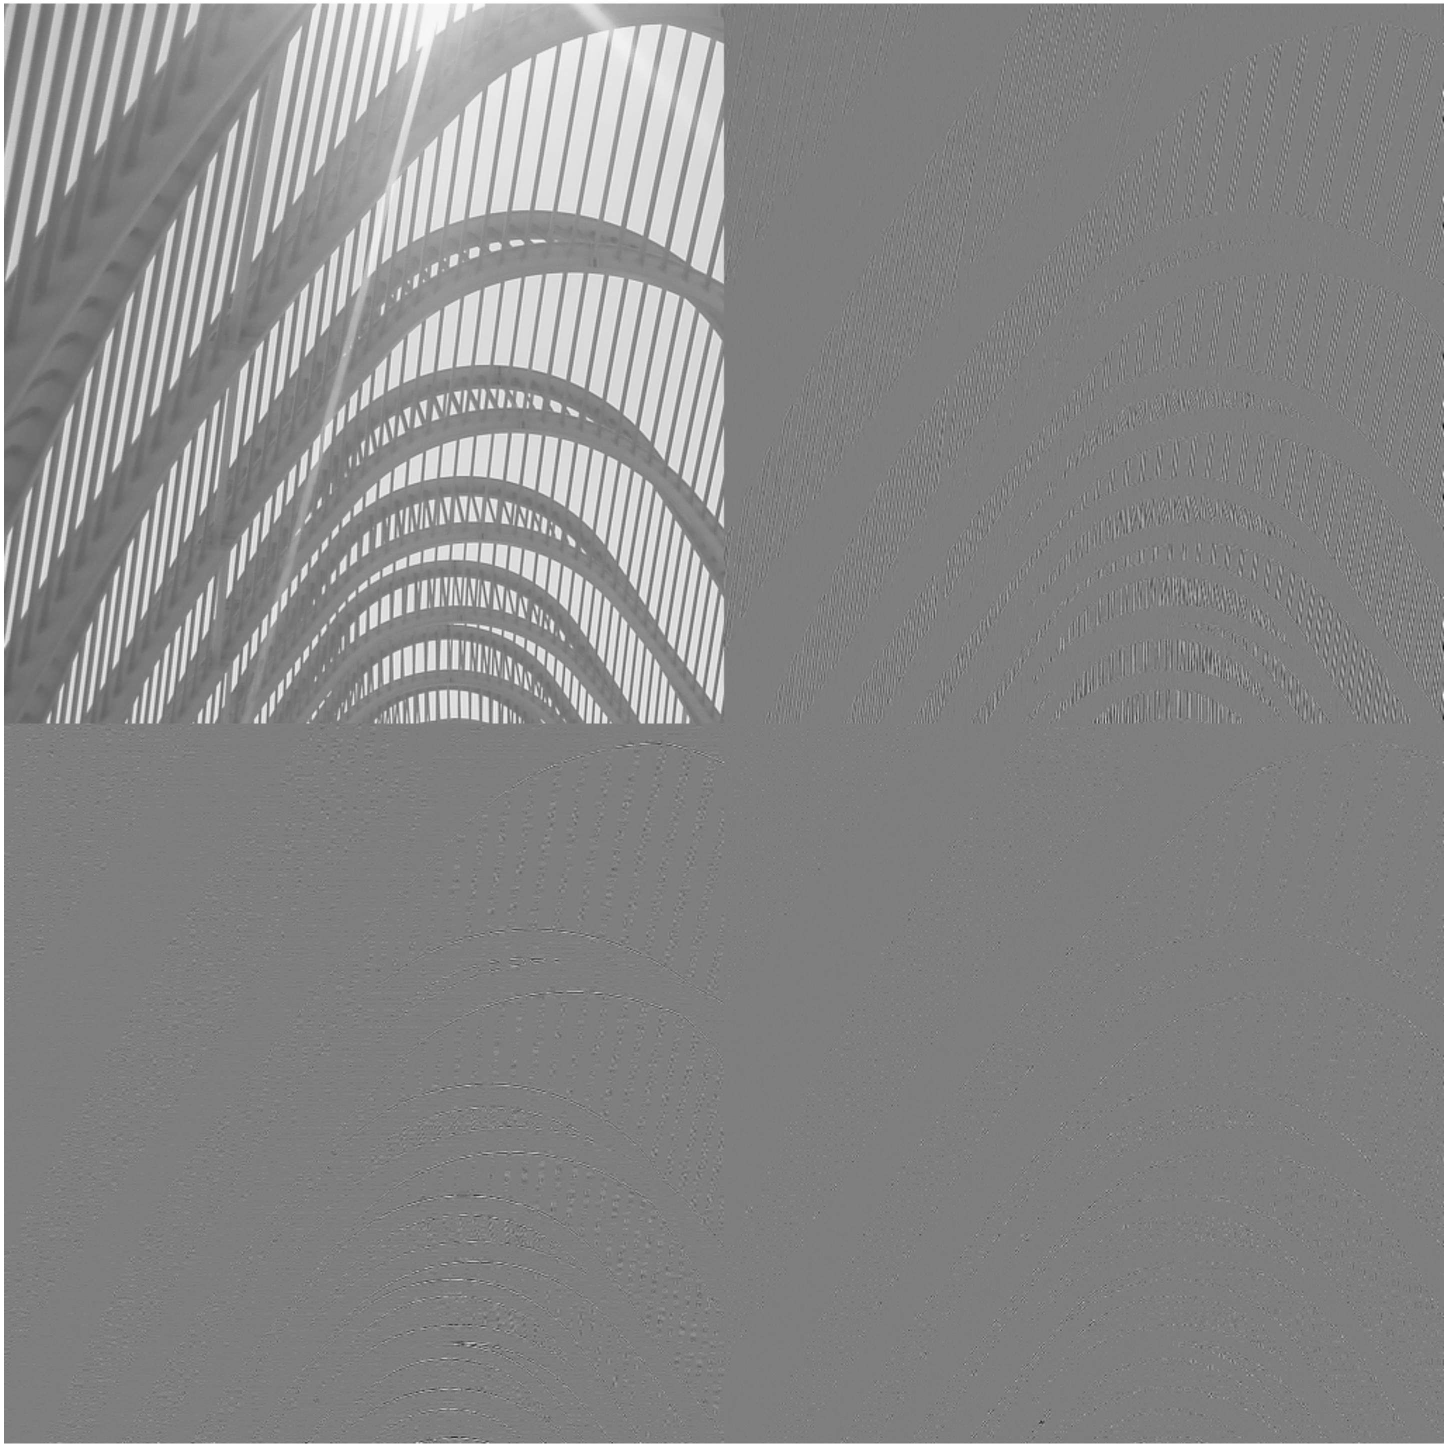

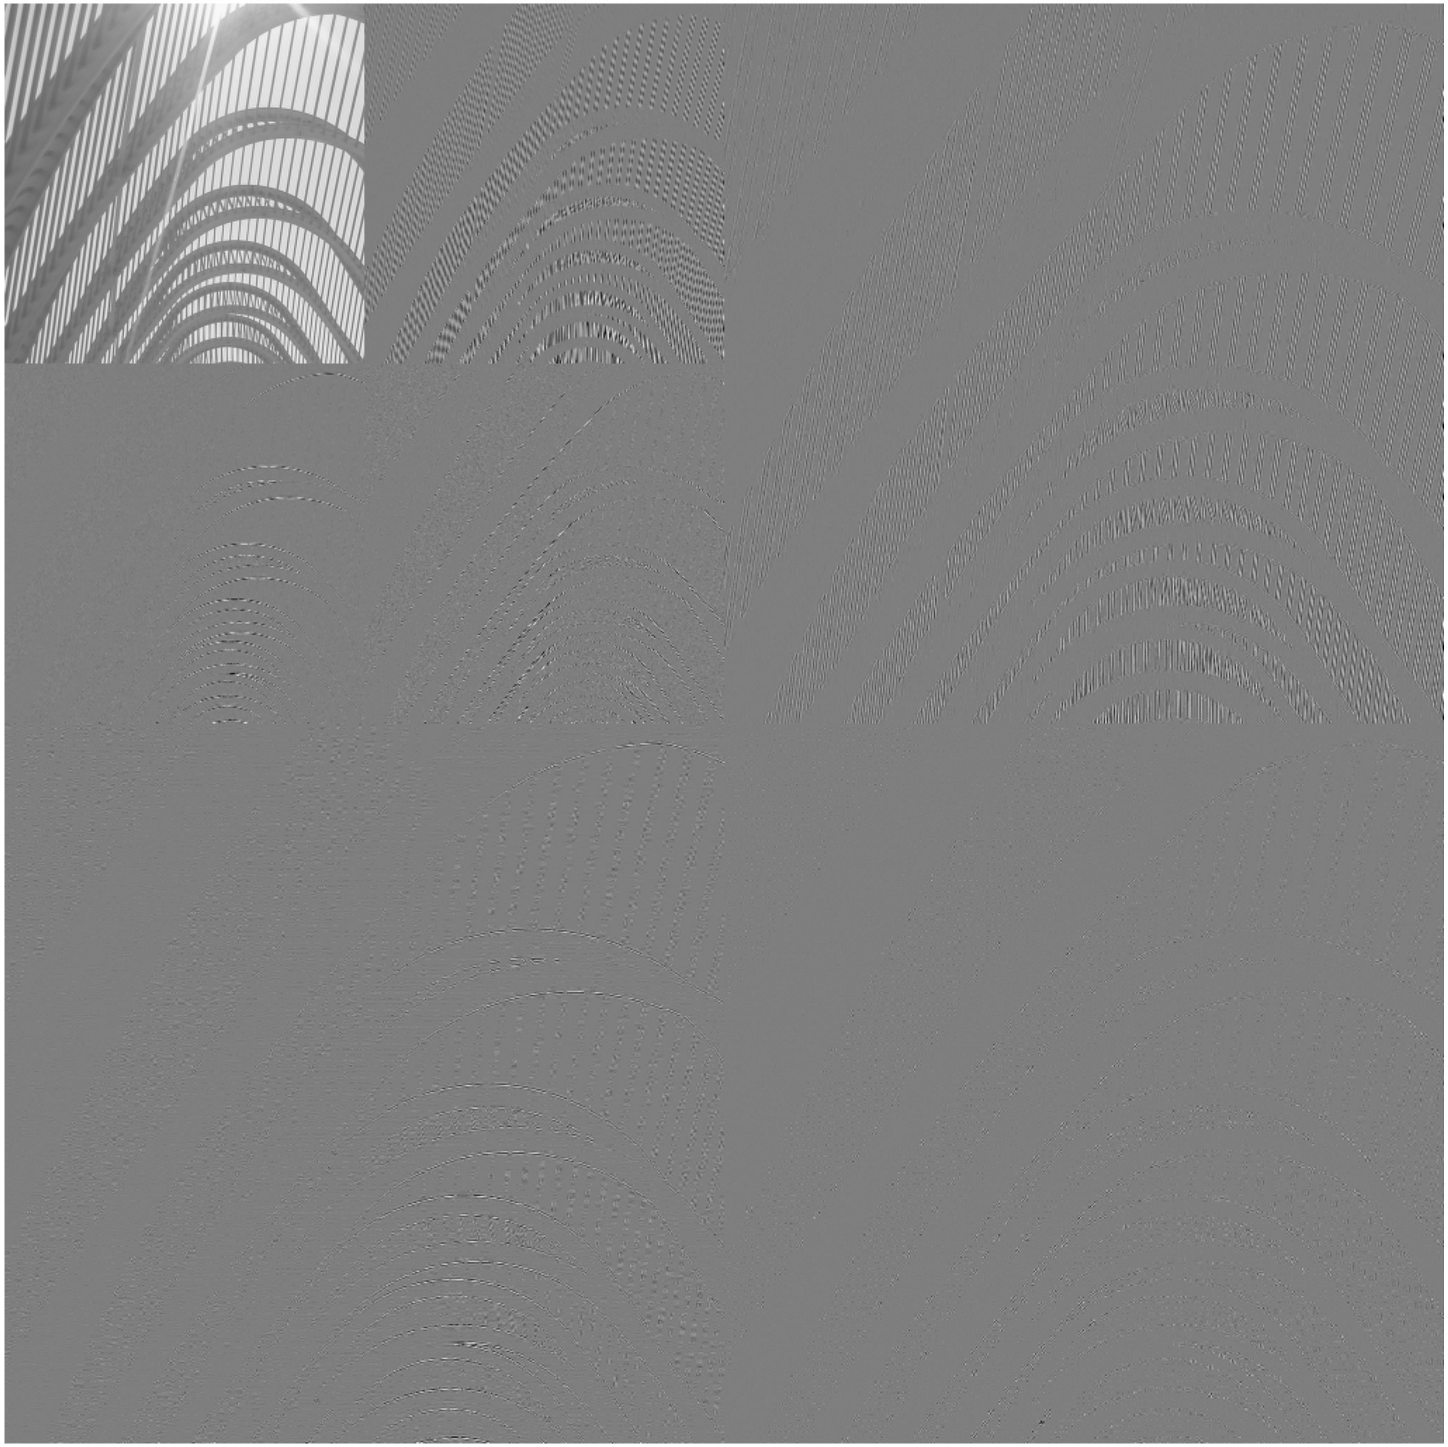

significant
horizontal
energy

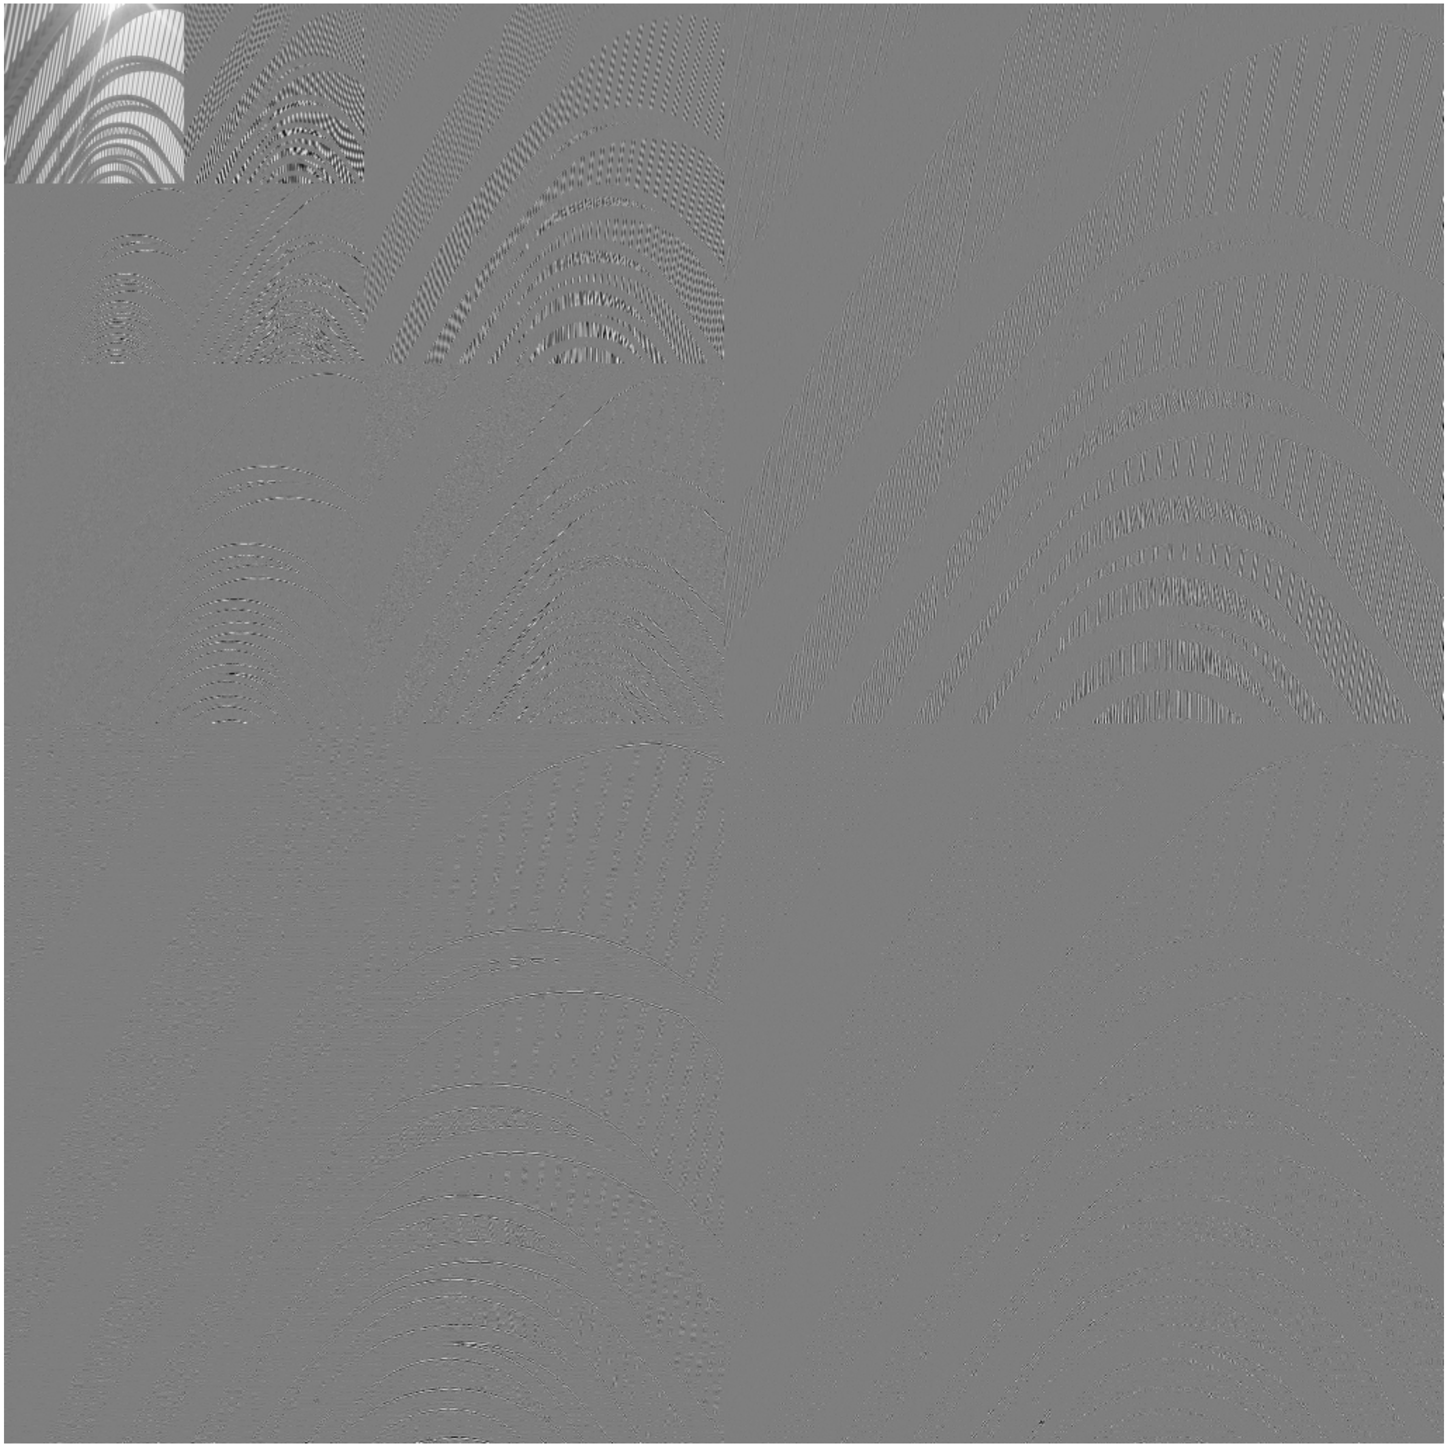

Constraints on Filters (1)

Z-transform :

$$\hat{X}(z) = \tfrac{1}{2}\left[H_0(z)G_0(z) + H_1(z)G_1(z)\right]X(z) + \tfrac{1}{2}\left[H_0(-z)G_0(z) + H_1(-z)G_1(z)\right]X(-z)$$

For perfect reconstruction $\left(\hat{X}(z) = X(z)\right)$

$$H_0(-z)G_0(z) + H_1(-z)G_1(z) = 0$$

$$H_0(z)G_0(z) + H_1(z)G_1(z) = 2$$

For finite impulse response (FIR) filters and ignoring the delay

$$g_0(n) = (-1)^n h_1(n)$$

$$g_1(n) = (-1)^{n+1} h_0(n)$$

Constraints on Filters (2)

- Biorthogonal

$$\langle g_0(k), h_0(2n-k)\rangle = \delta(n), \qquad \langle g_0(k), h_1(2n-k)\rangle = 0$$

$$\langle g_1(k), h_1(2n-k)\rangle = \delta(n), \qquad \langle g_1(k), h_0(2n-k)\rangle = 0$$

$$\Rightarrow \langle h_i(2n-k), g_j(k)\rangle = \delta(i-j)\delta(n), \quad i,j = \{0,1\}$$

The analysis and synthesis filter impulse responses of all two-band, real-coefficient, perfect reconstruction filter banks are subject to the biorthogonality constraint

Common Wavelet Filters

- Haar: simple, orthogonal, low effectiveness

$$h_0 = \frac{1}{\sqrt{2}}[1 \quad 1]; \quad h_1 = \frac{1}{\sqrt{2}}[1 \quad -1]$$

- Daubechies 9/7: bi-orthogonal
  most commonly used if numerical reconstruction errors are acceptable
- LeGall 5/3: bi-orthogonal, integer operation,
  can be implemented with integer operations only, used for lossless image coding

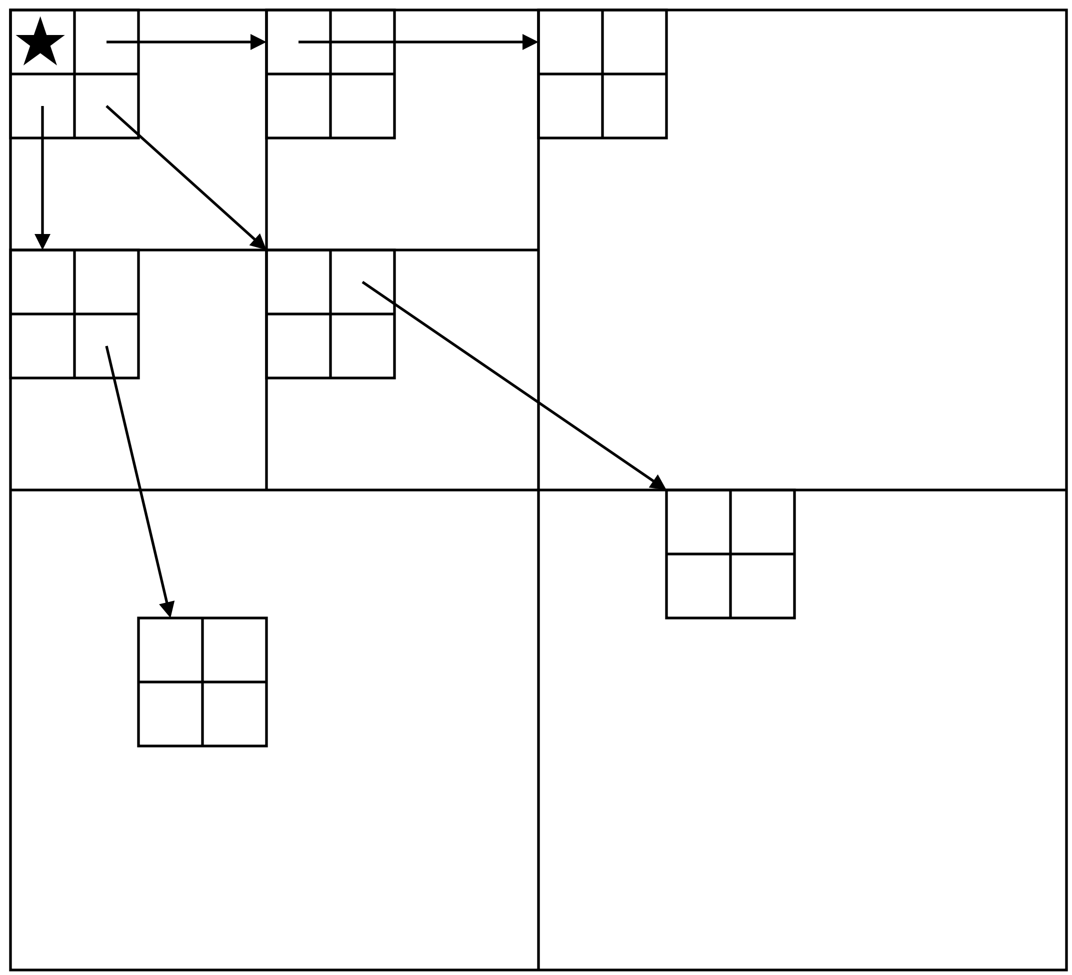

Fig.2. Examples of parent-offspring dependencies in the spatial-orientation tree.

Wavelets

- Allows for *intra-scale* prediction (like many other compression methods) – equivalently the wavelet transform is a *decorrelating transform* just like the DCT as used by JPEG

- Allows for *inter-scale* (coarse-fine scale) prediction

A reason wavelet compression is effective

- Coefficient entropies

	Entropy
Original image	7.22
1-level Haar wavelet	5.96
1-level linear spline wavelet	5.53
2-level Haar wavelet	5.02
2-level linear spline wavelet	4.57

Interscale hierarchies are powerful

from Mike Spann

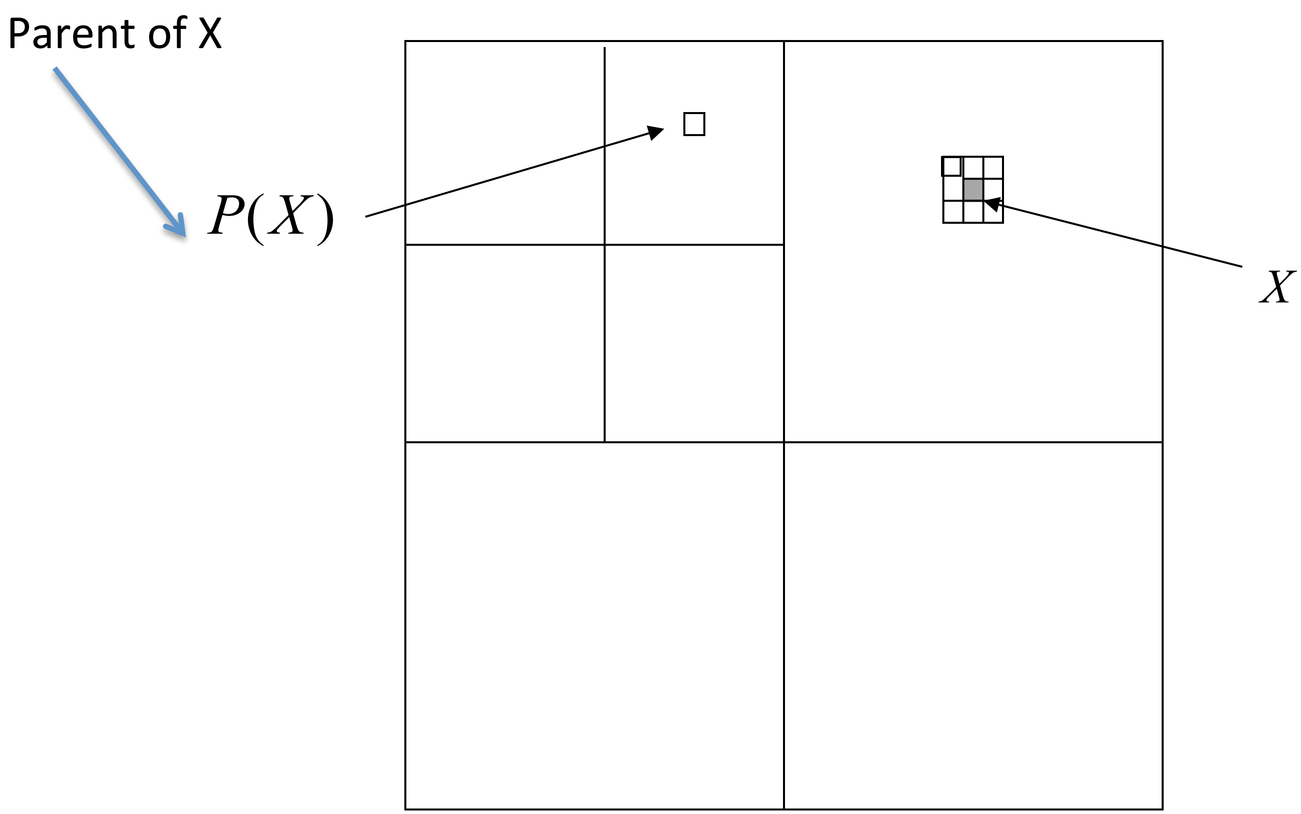

Interscale hierarchies

from Mike Spann

- Define sets *S* (small) and *L* (large) wavelet coefficients
- Without inter-scale dependencies

$$P(X \in S \mid P(X) \in S) = \frac{\#S}{N^2} \qquad P(X \in L \mid P(X) \in L) = \frac{\#L}{N^2}$$

- Measured dependencies in Lena

$P(X \in S \mid P(X) \in S)$	$P(X \in L \mid P(X) \in L)$	$\frac{\#S}{N^2}$	$\frac{\#L}{N^2}$
0.886	0.529	0.781	0.219

Intra-scale dependencies

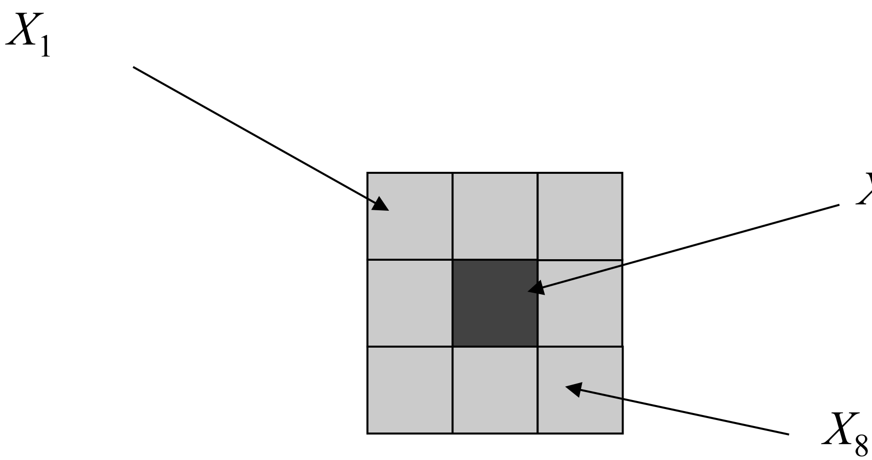

from Mike Spann

$$\bar{c} = \frac{1}{8}\sum_{n=1}^{8} c(X_n)$$

$$\{X_n\} \in S \text{ if } \bar{c} \le T$$

$$\{X_n\} \in L \text{ if } \bar{c} > T$$

$P(X \in S \mid \{X_n\} \in S)$	$P(X \in L \mid \{X_n\} \in L)$	$\frac{\#S}{N^2}$	$\frac{\#L}{N^2}$
0.912	0.623	0.781	0.219

Applications

- JPEG 2000 image compression
- (also used for digital cinema)
- Image blending
- Image denoising
- (many others)

Pyramid Blending

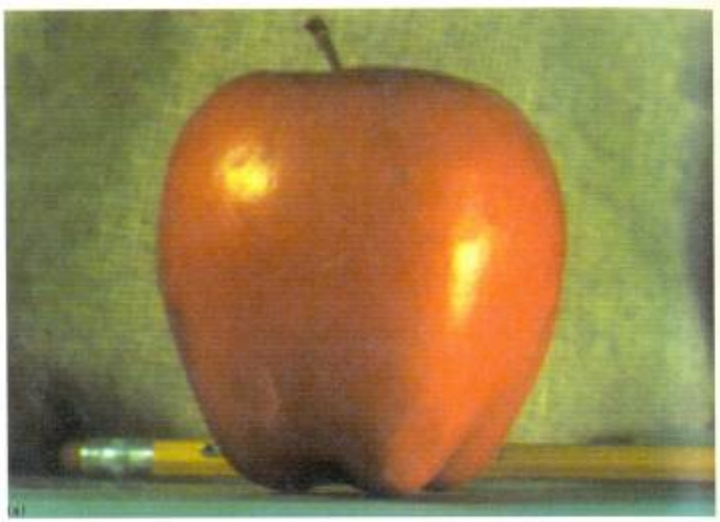

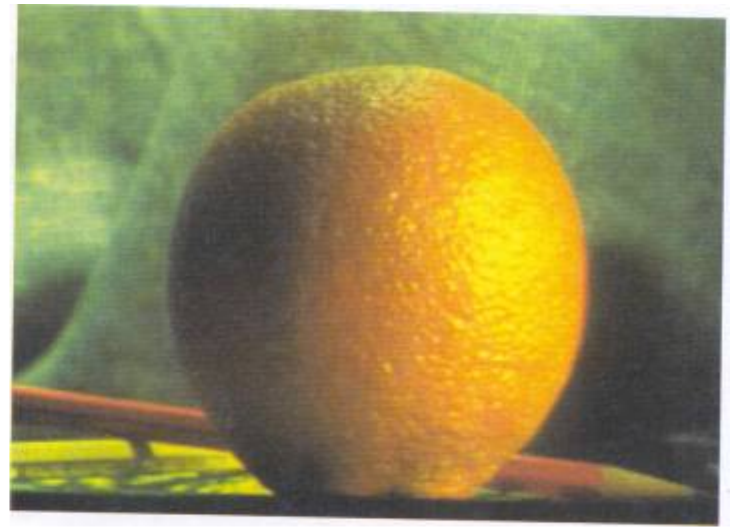

(h) (l)

A Multiresolution Spline With Application to Image Mosaics

PETER J. BURT and EDWARD H. ADELSON
RCA David Sarnoff Research Center

Application: Denoising

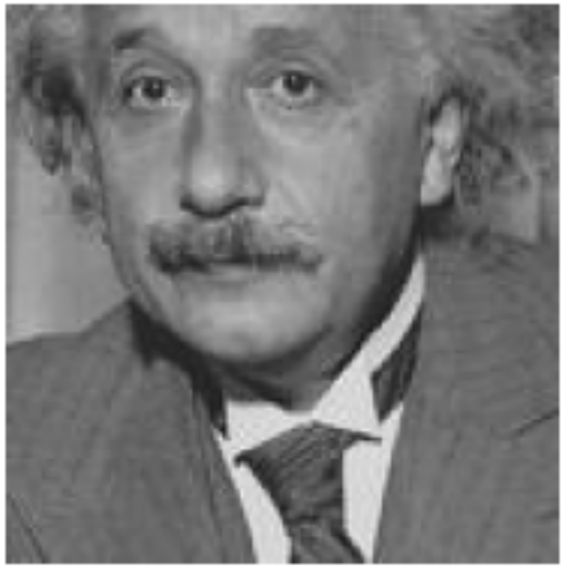

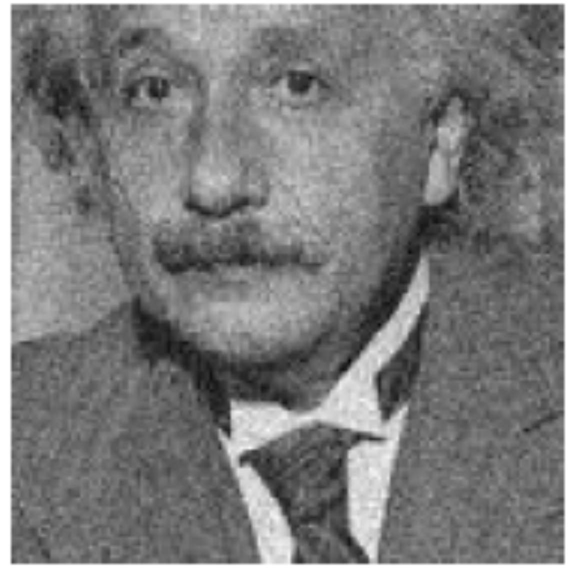

How to characterize the difference between the images?

How do we use the differences to clean up the image?

http://www.cns.nyu.edu/pub/lcv/simoncelli96c.pdf

Application: Denoising

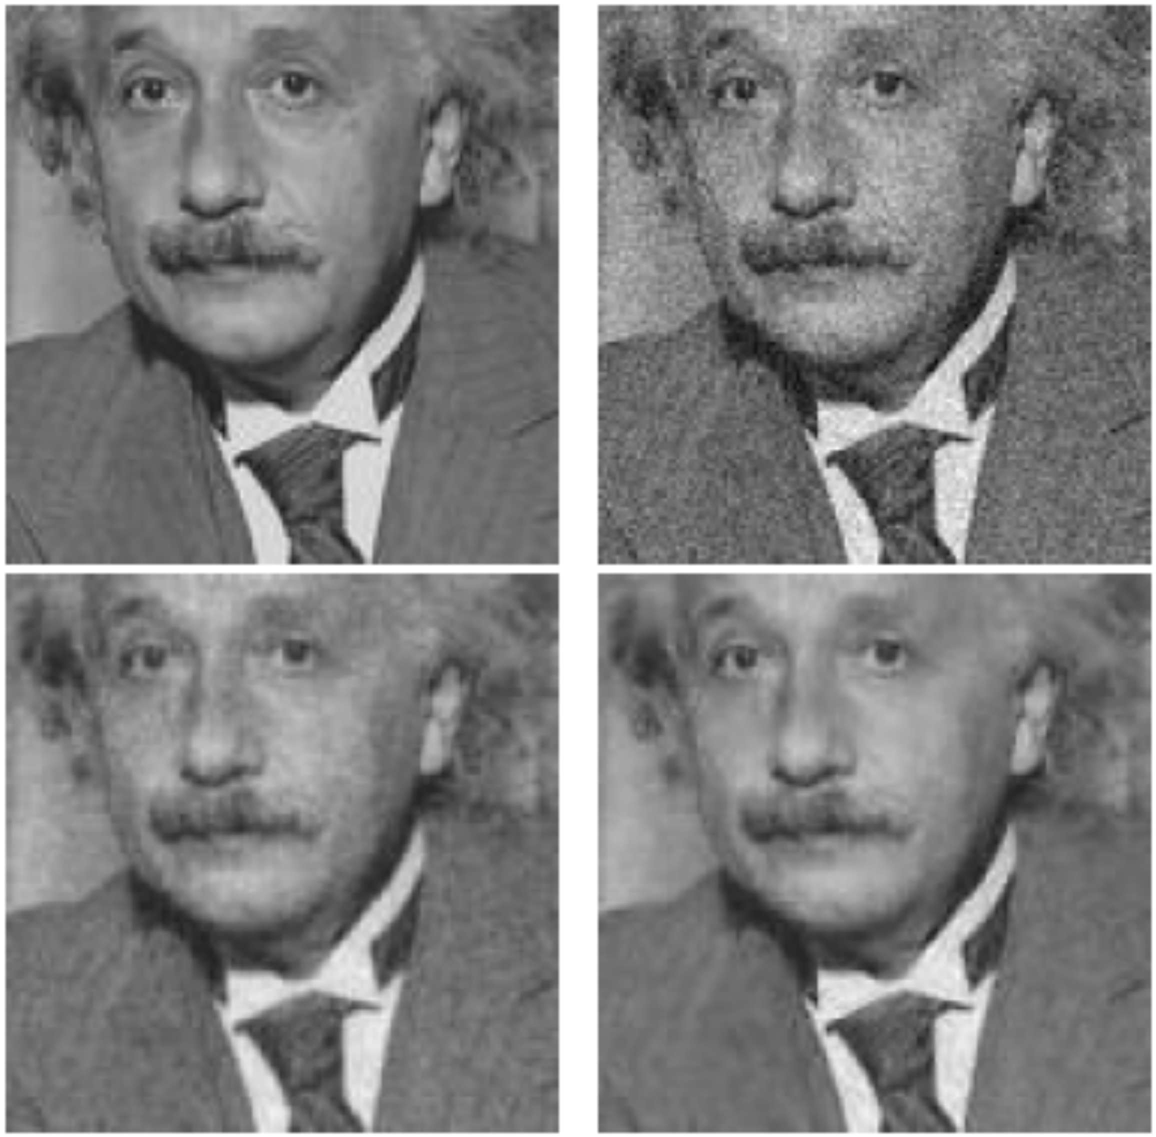

Original

Noise-corrupted

Wiener filter

Steerable pyramid coring

http://www.cns.nyu.edu/pub/lcv/simoncelli96c.pdf

JPEG

- Pros
  - Low complexity
  - Memory efficient
  - Reasonable coding efficiency
- Cons:
  - Single resolution
  - Single quality
  - No target bit-rate
  - Blocking artifacts at low rate
  - Poor error resilience
  - Low flexibility, for example, no ROI

JPEG2000 Features

- Improved coding efficiency
- Full quality scalability
  – From lossless to lossy at different bit rates
- Spatial scalability
- Improved error resilience
- Tiling
- Region of interests
- More demanding in memory and computation time

JPEG 2000: Core Processing

- Wavelet transform
  - Daubechies 9-tap/7-tap filter for irreversible transform
  - Le Gall 5-tap/3-tap filter for reversible transformation
- Quantization
  - Separate quantization step-sizes for each subband
- Entropy encoding
  - Bit plane coding - the most significant bit plane is coded first.

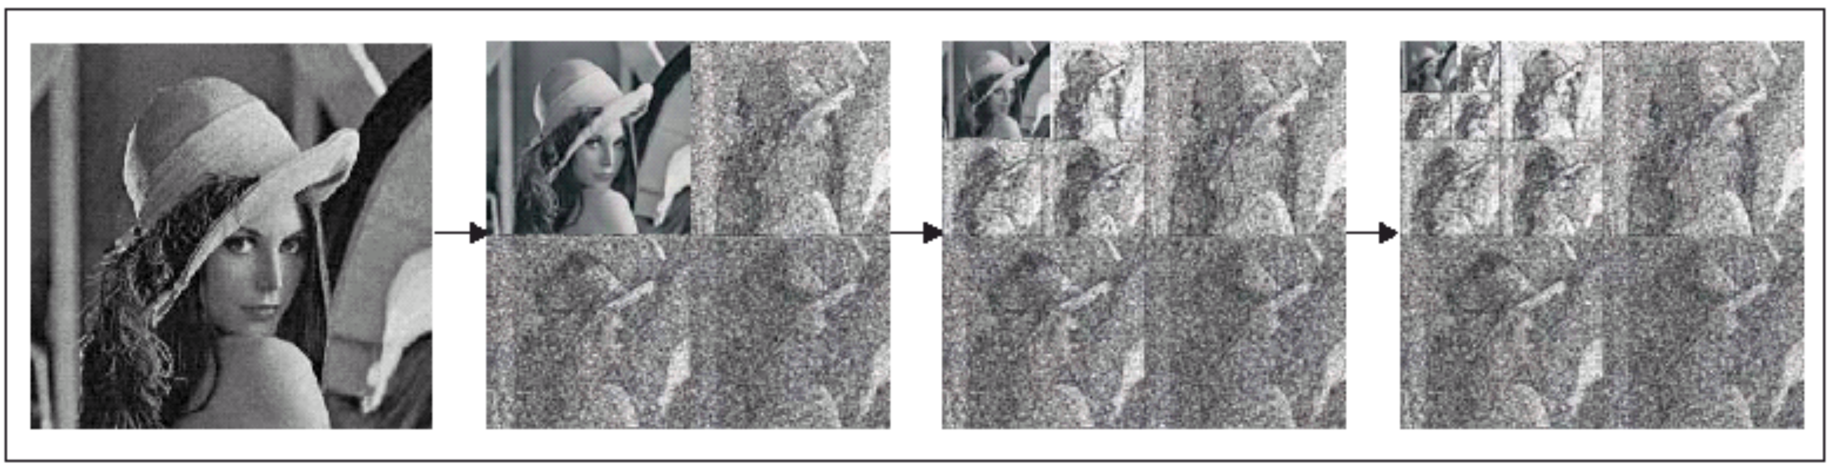

▲ *6. Three-level dyadic wavelet decomposition of the image "Lena."*

JPEG2000 vs. JPEG: Coding Efficiency

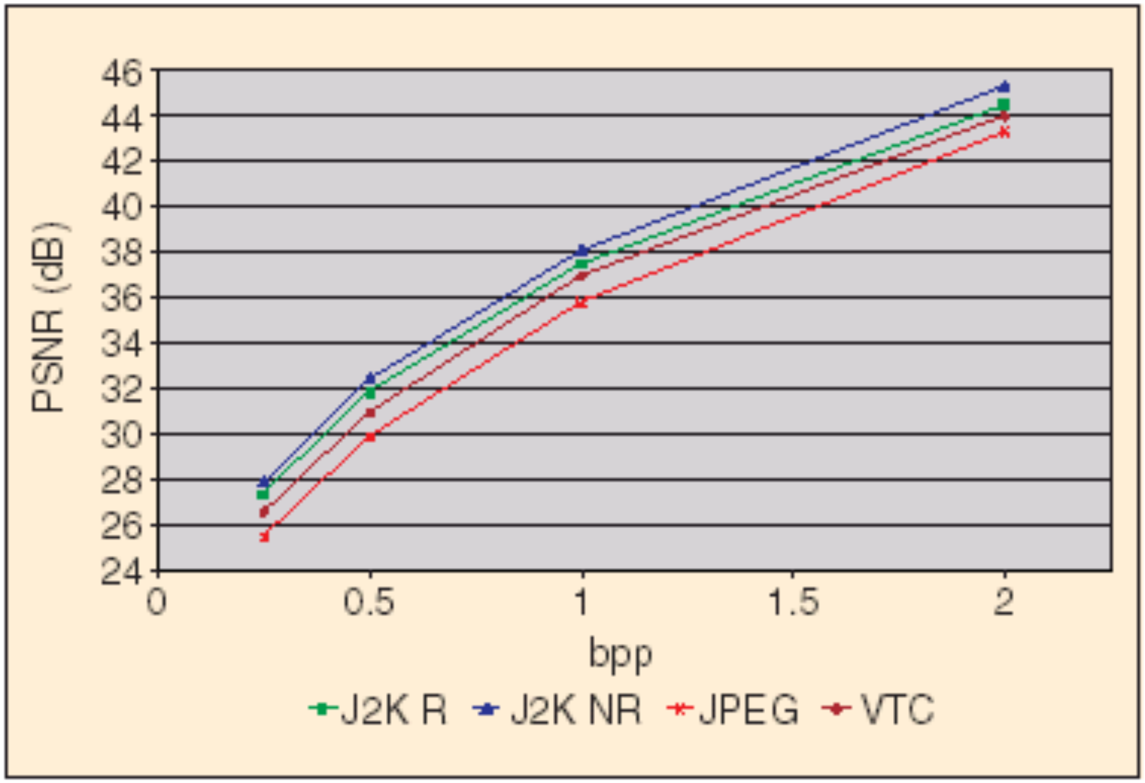

19. *PSNR results for the lossy compression of a natural image by means of different compression standards.*

J2K R: Using reversible wavelet filters;
J2K NR: Using non-reversible filter;
VTC: Visual texture coding for MPEG-4 video

PSNR: Peak signal-to-noise ratio

PSNR = $10 \log_{10}(MAX^2/MSE)$

From [skodras01]

Example Image

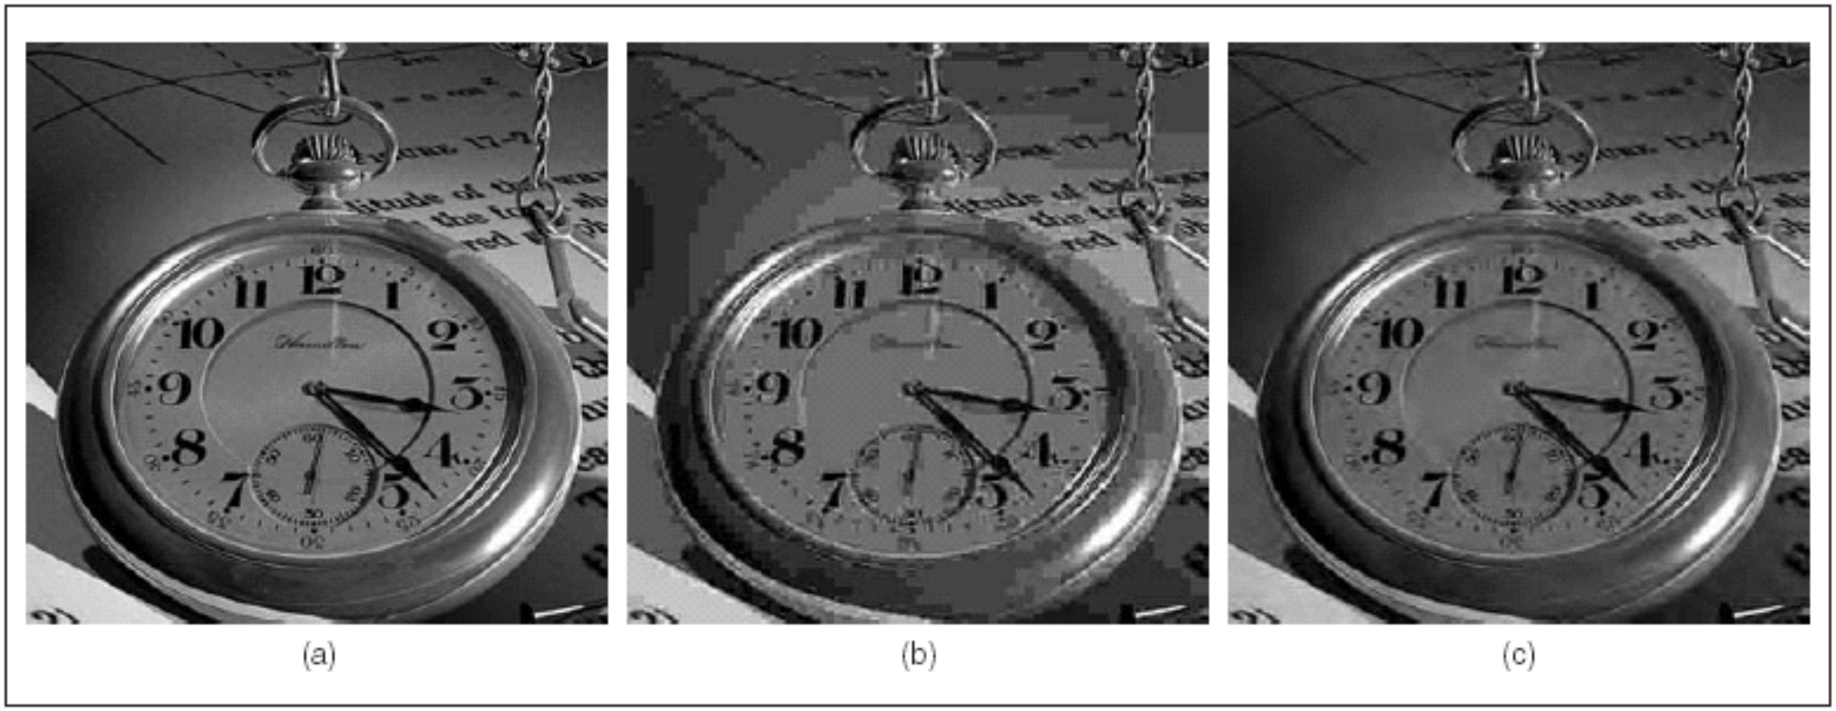

▲ 20. Image "watch" of size 512 × 512 (courtesy of Kevin Odhner): (a) original, and reconstructed after compression at 0.2 b/p by means of (b) JPEG and (c) JPEG 2000.

From [skodras01]

Another Example

(a) (b)

▲ 21. Reconstructed image "ski" after compression at 0.25 b/p by means of (a) JPEG and (b) JPEG 2000.

From [skodras01]

How J2K Achieves Scalability

- Core: Wavelet transform
  - Yields a multi-resolution representation of an original image
  - Wavelet coefficients are coded bit plane by bit plane
- Spatial scalability can be achieved by reconstructing from only low resolution wavelet coefficients
- Quality scalability can be achieved by decoding only partial bit planes

Quality Scalability of JPEG2000

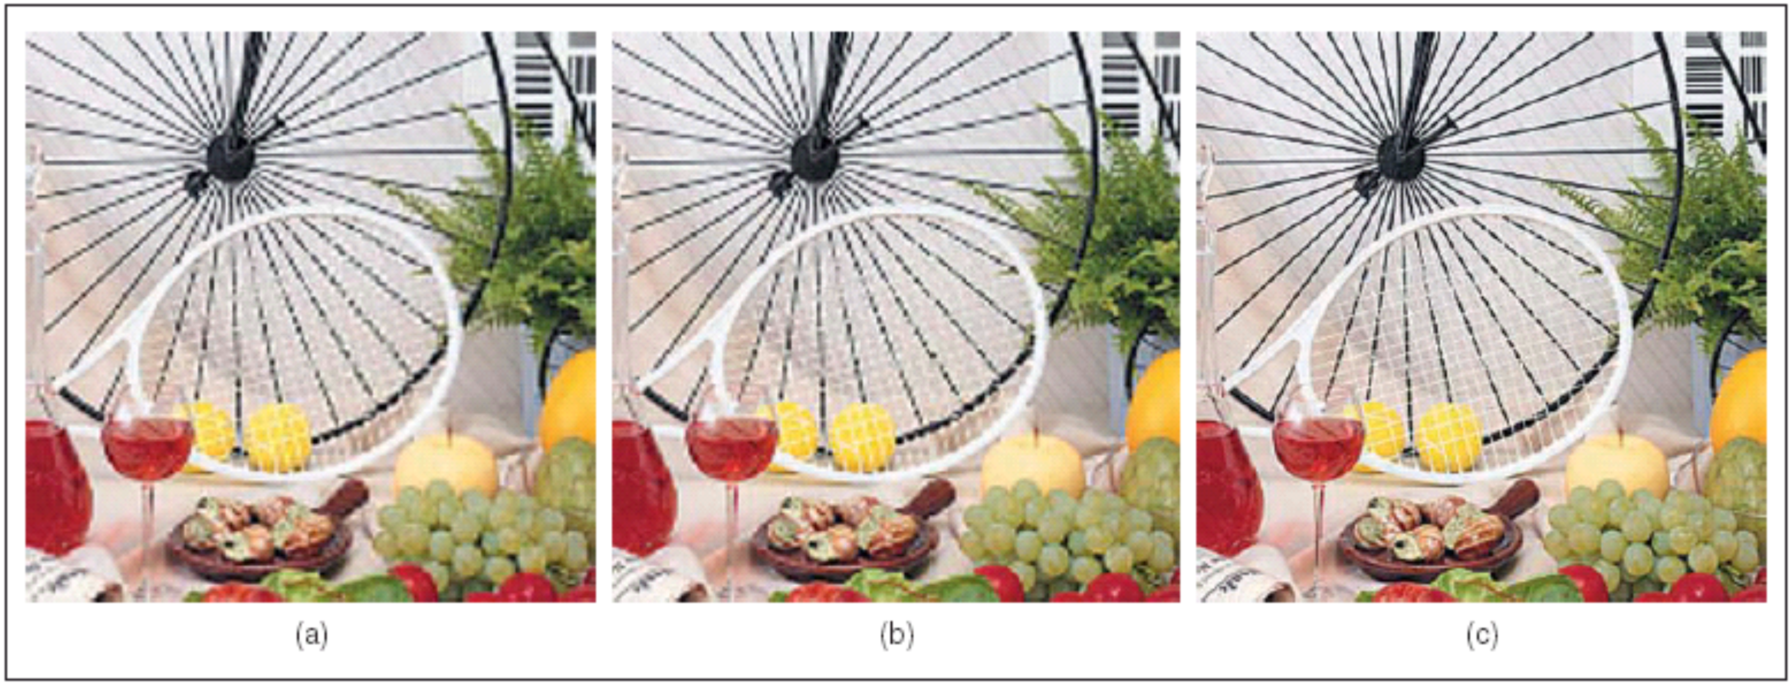

▲ *17. Example of SNR scalability. Part of the decompressed image "bike" at (a) 0.125 b/p, (b) 0.25 b/p, and (c) 0.5 b/p.*

From [skodras01]

Spatial Scalability of JPEG2000

▲ 18. *Example of the progressive-by-resolution decoding for the color image "bike."*

Wavelets have other uses

- Compression
  - Maximally decimated wavelets are most appropriate
- Signal analysis
  - Overcomplete wavelets may be more effective
  - Example: Steerable pyramids

Steerable pyramids: other wavelets

- Filters can measure local orientation direction and strength and phase at any orientation.

G2 H2