Multi-scale Statistical Image Models and Denoising

Eero P. Simoncelli
Center for Neural Science, and
Courant Institute of Mathematical Sciences
New York University

http://www.cns.nyu.edu/~eero

Multi-scale roots

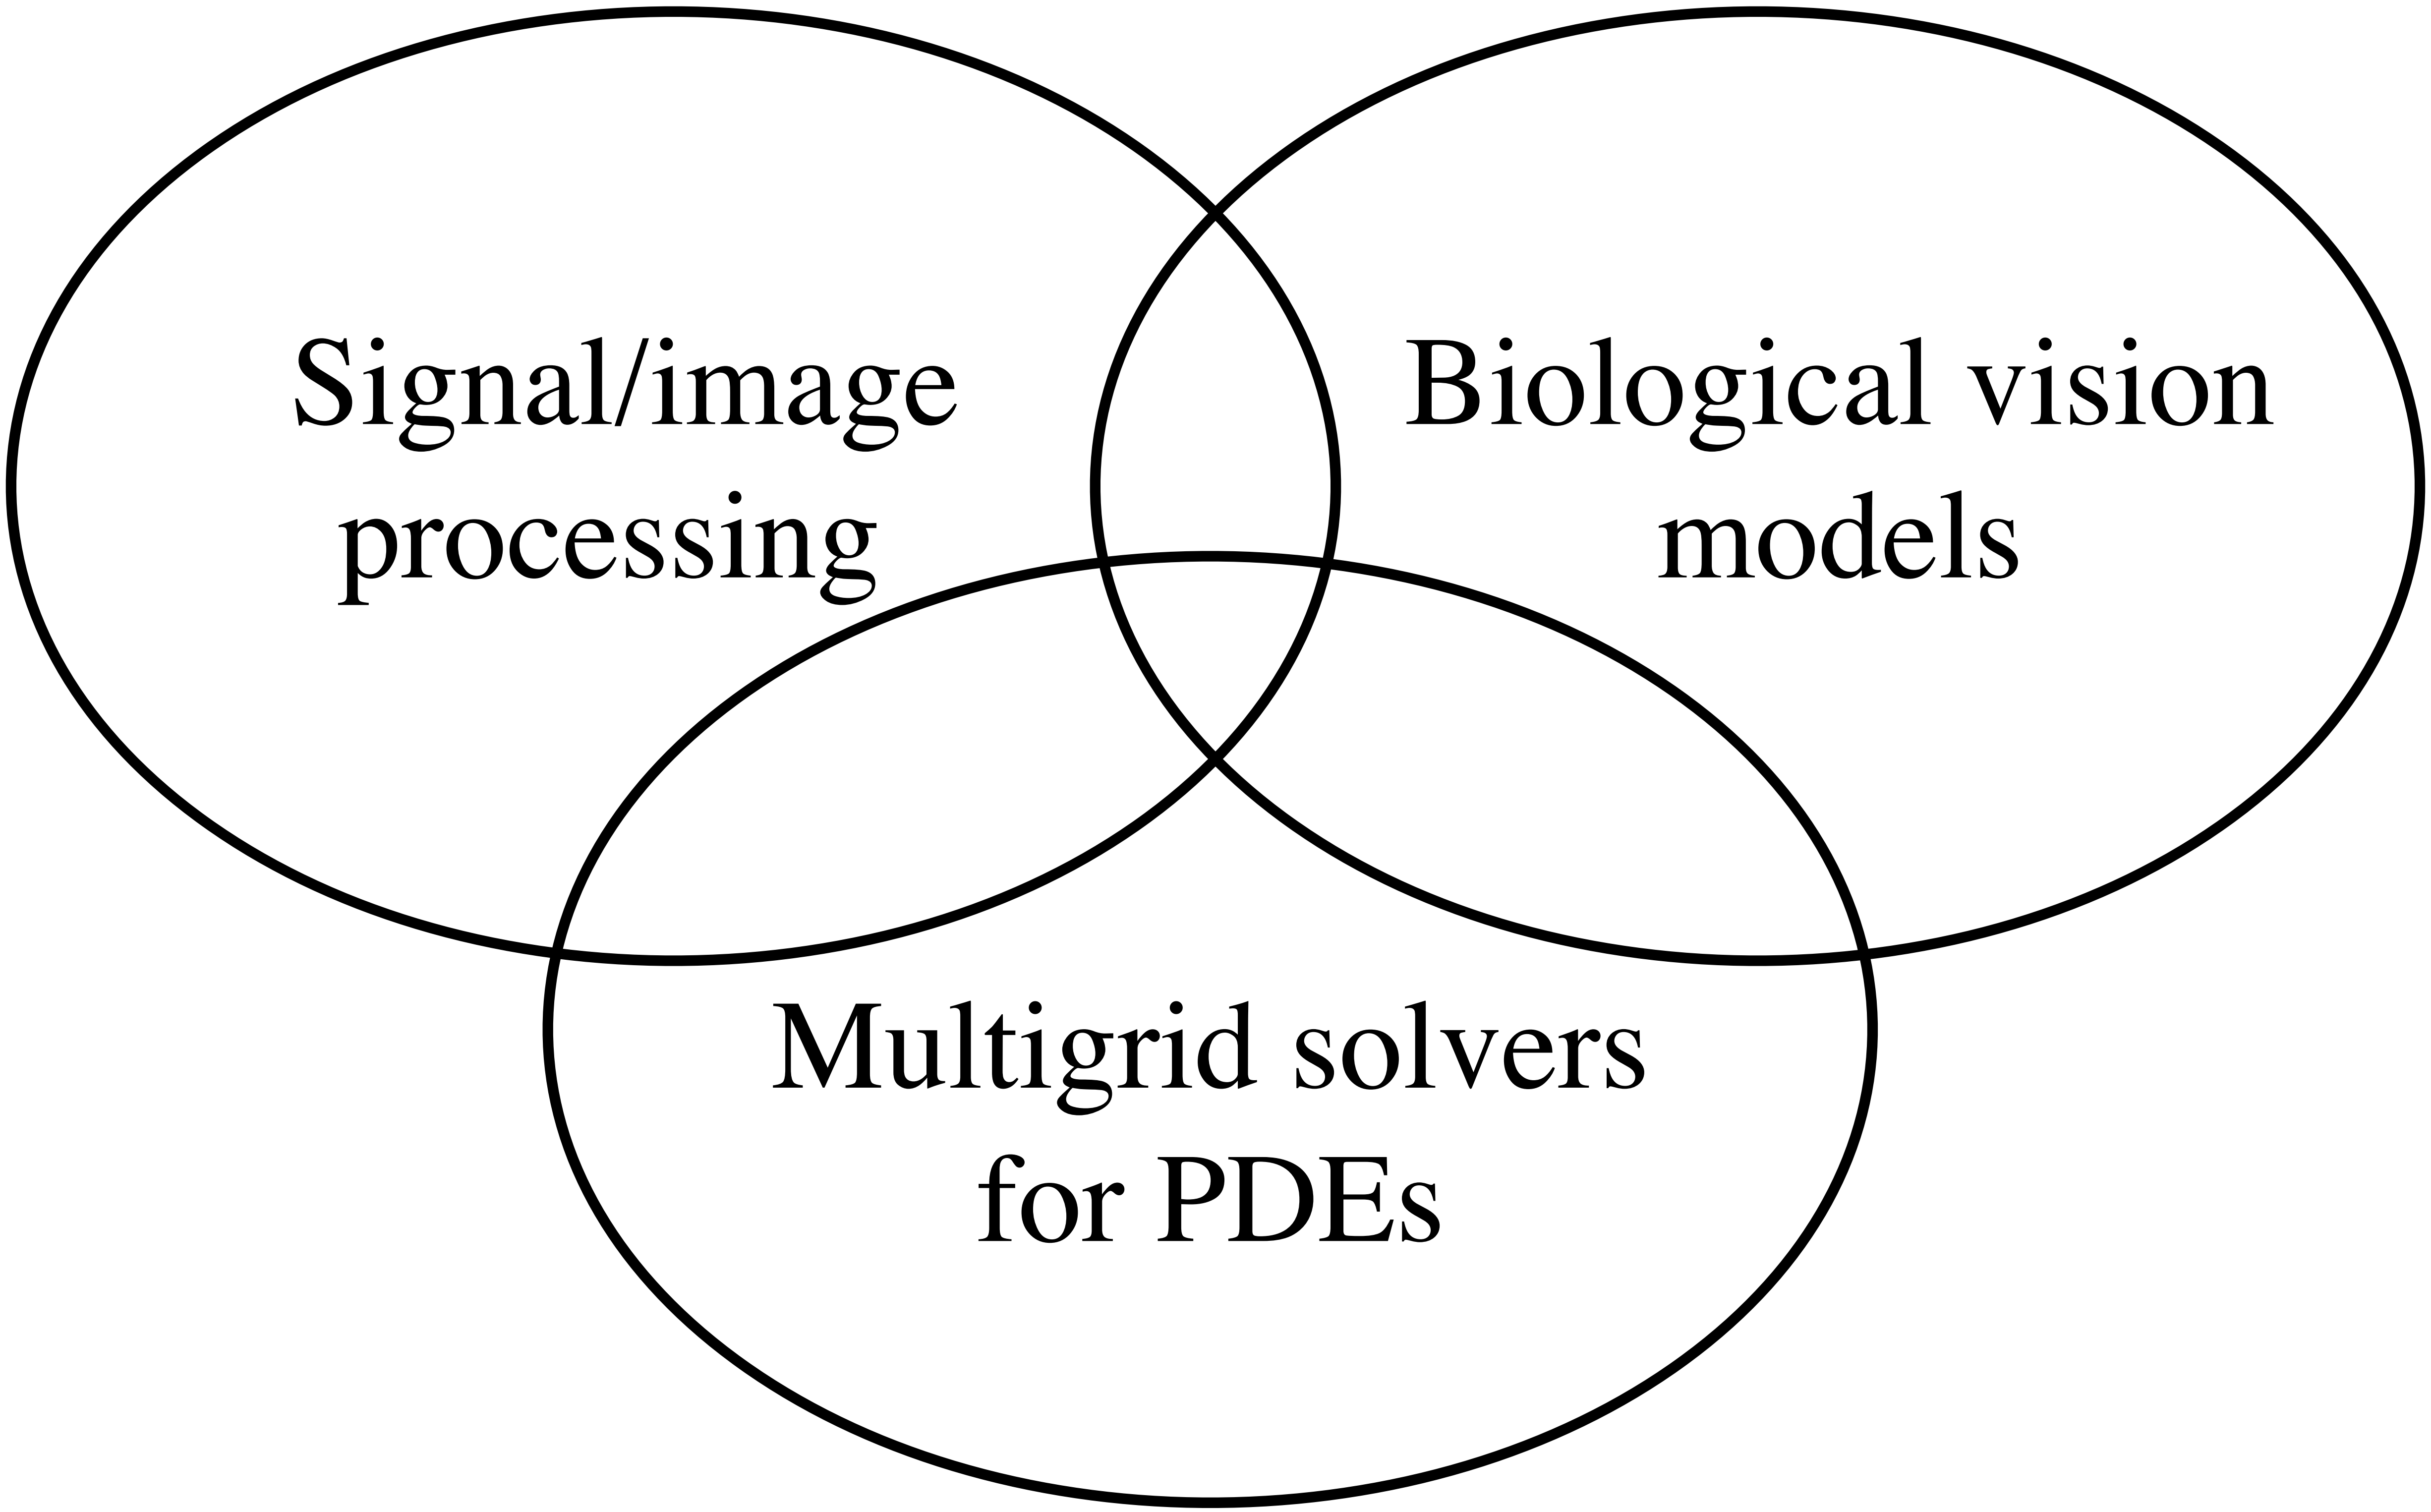

The “Wavelet revolution”

- Early 1900’s: Haar introduces first orthonormal wavelet
- Late 70’s: Quadrature mirror filters
- Early 80’s: Multi-resolution pyramids
- Late 80’s: Orthonormal wavelets
- 90’s: Return to overcomplete (non-aliased) pyramids, especially oriented pyramids
- >250,000 articles published in past 2 decades
- Best results in many signal/image processing applications

“Laplacian” pyramid

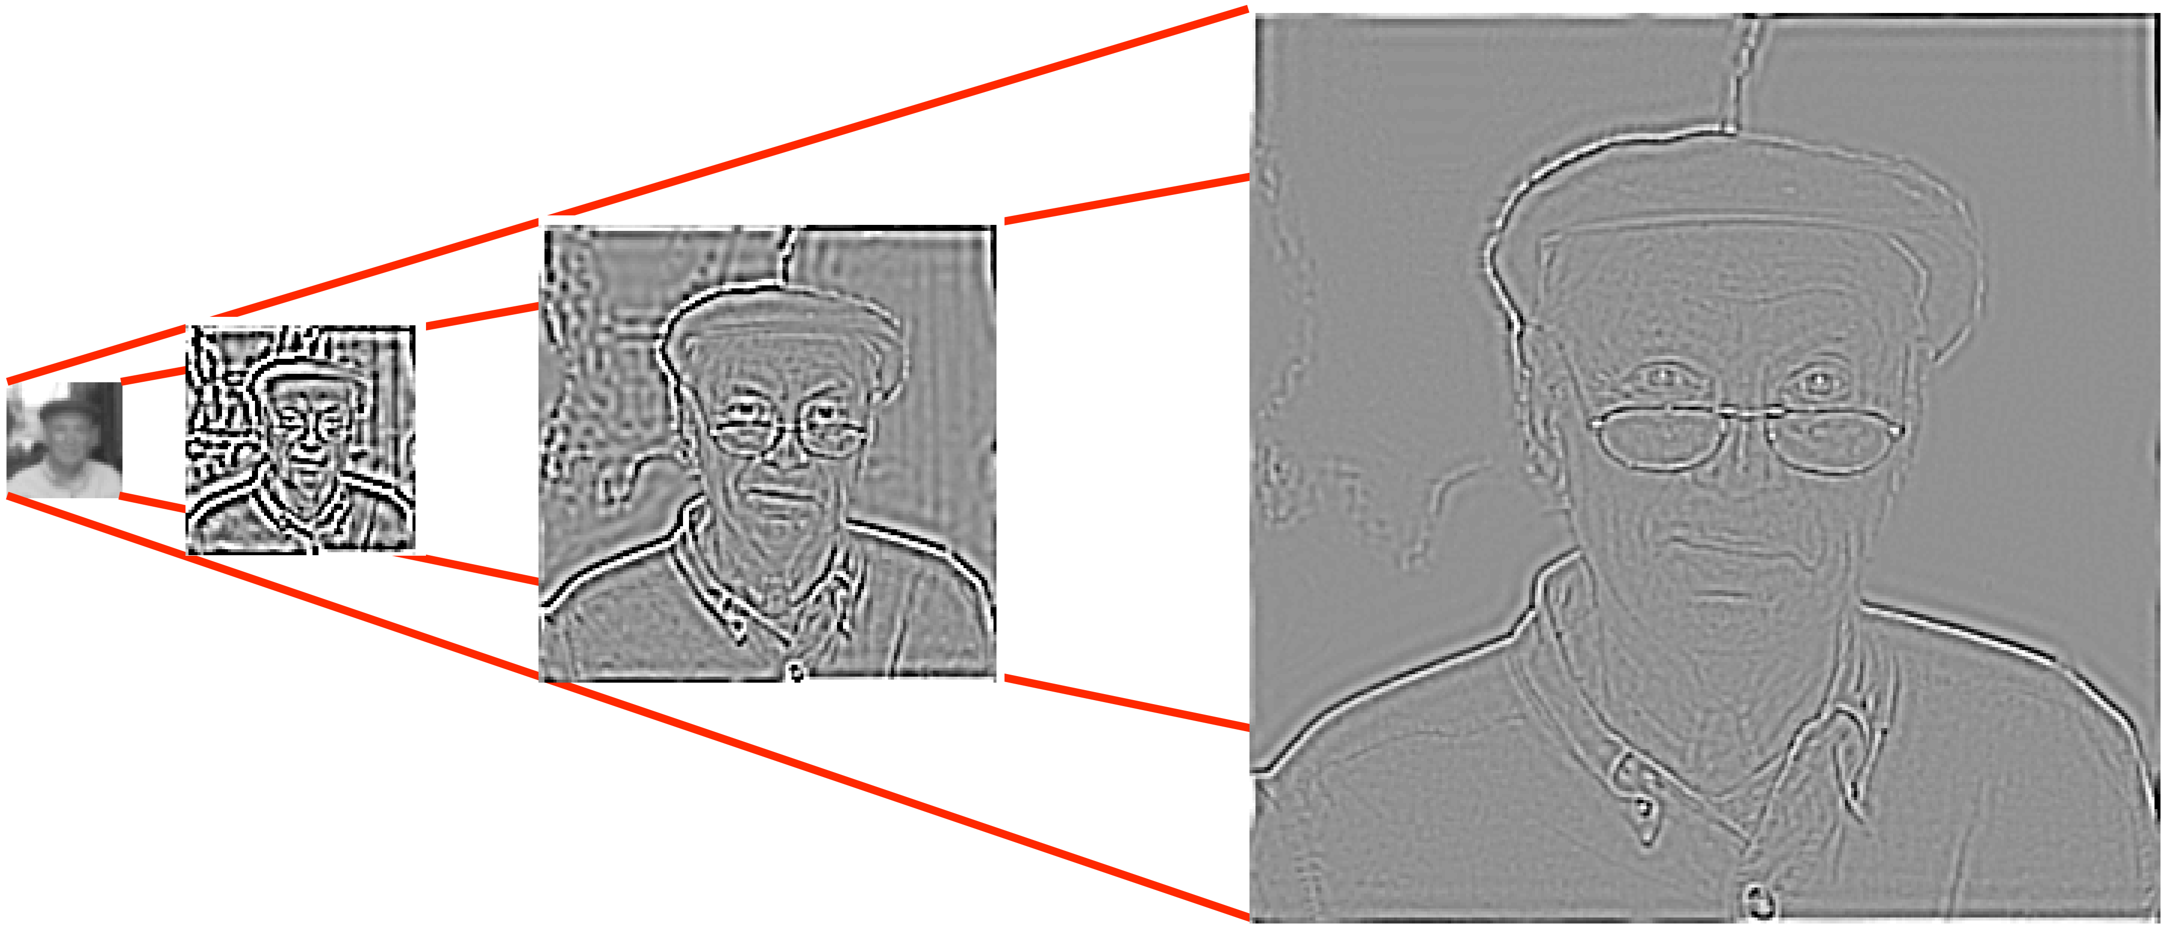

[Burt & Adelson, ‘81]

Multi-scale gradient basis

- Multi-scale bases: efficient representation
- Derivatives: good for analysis
  - Local Taylor expansion of image structures
  - Explicit geometry (orientation)
- Combination:
  - Explicit incorporation of geometry in basis
  - Bridge between PDE / harmonic analysis approaches

“Steerable” pyramid

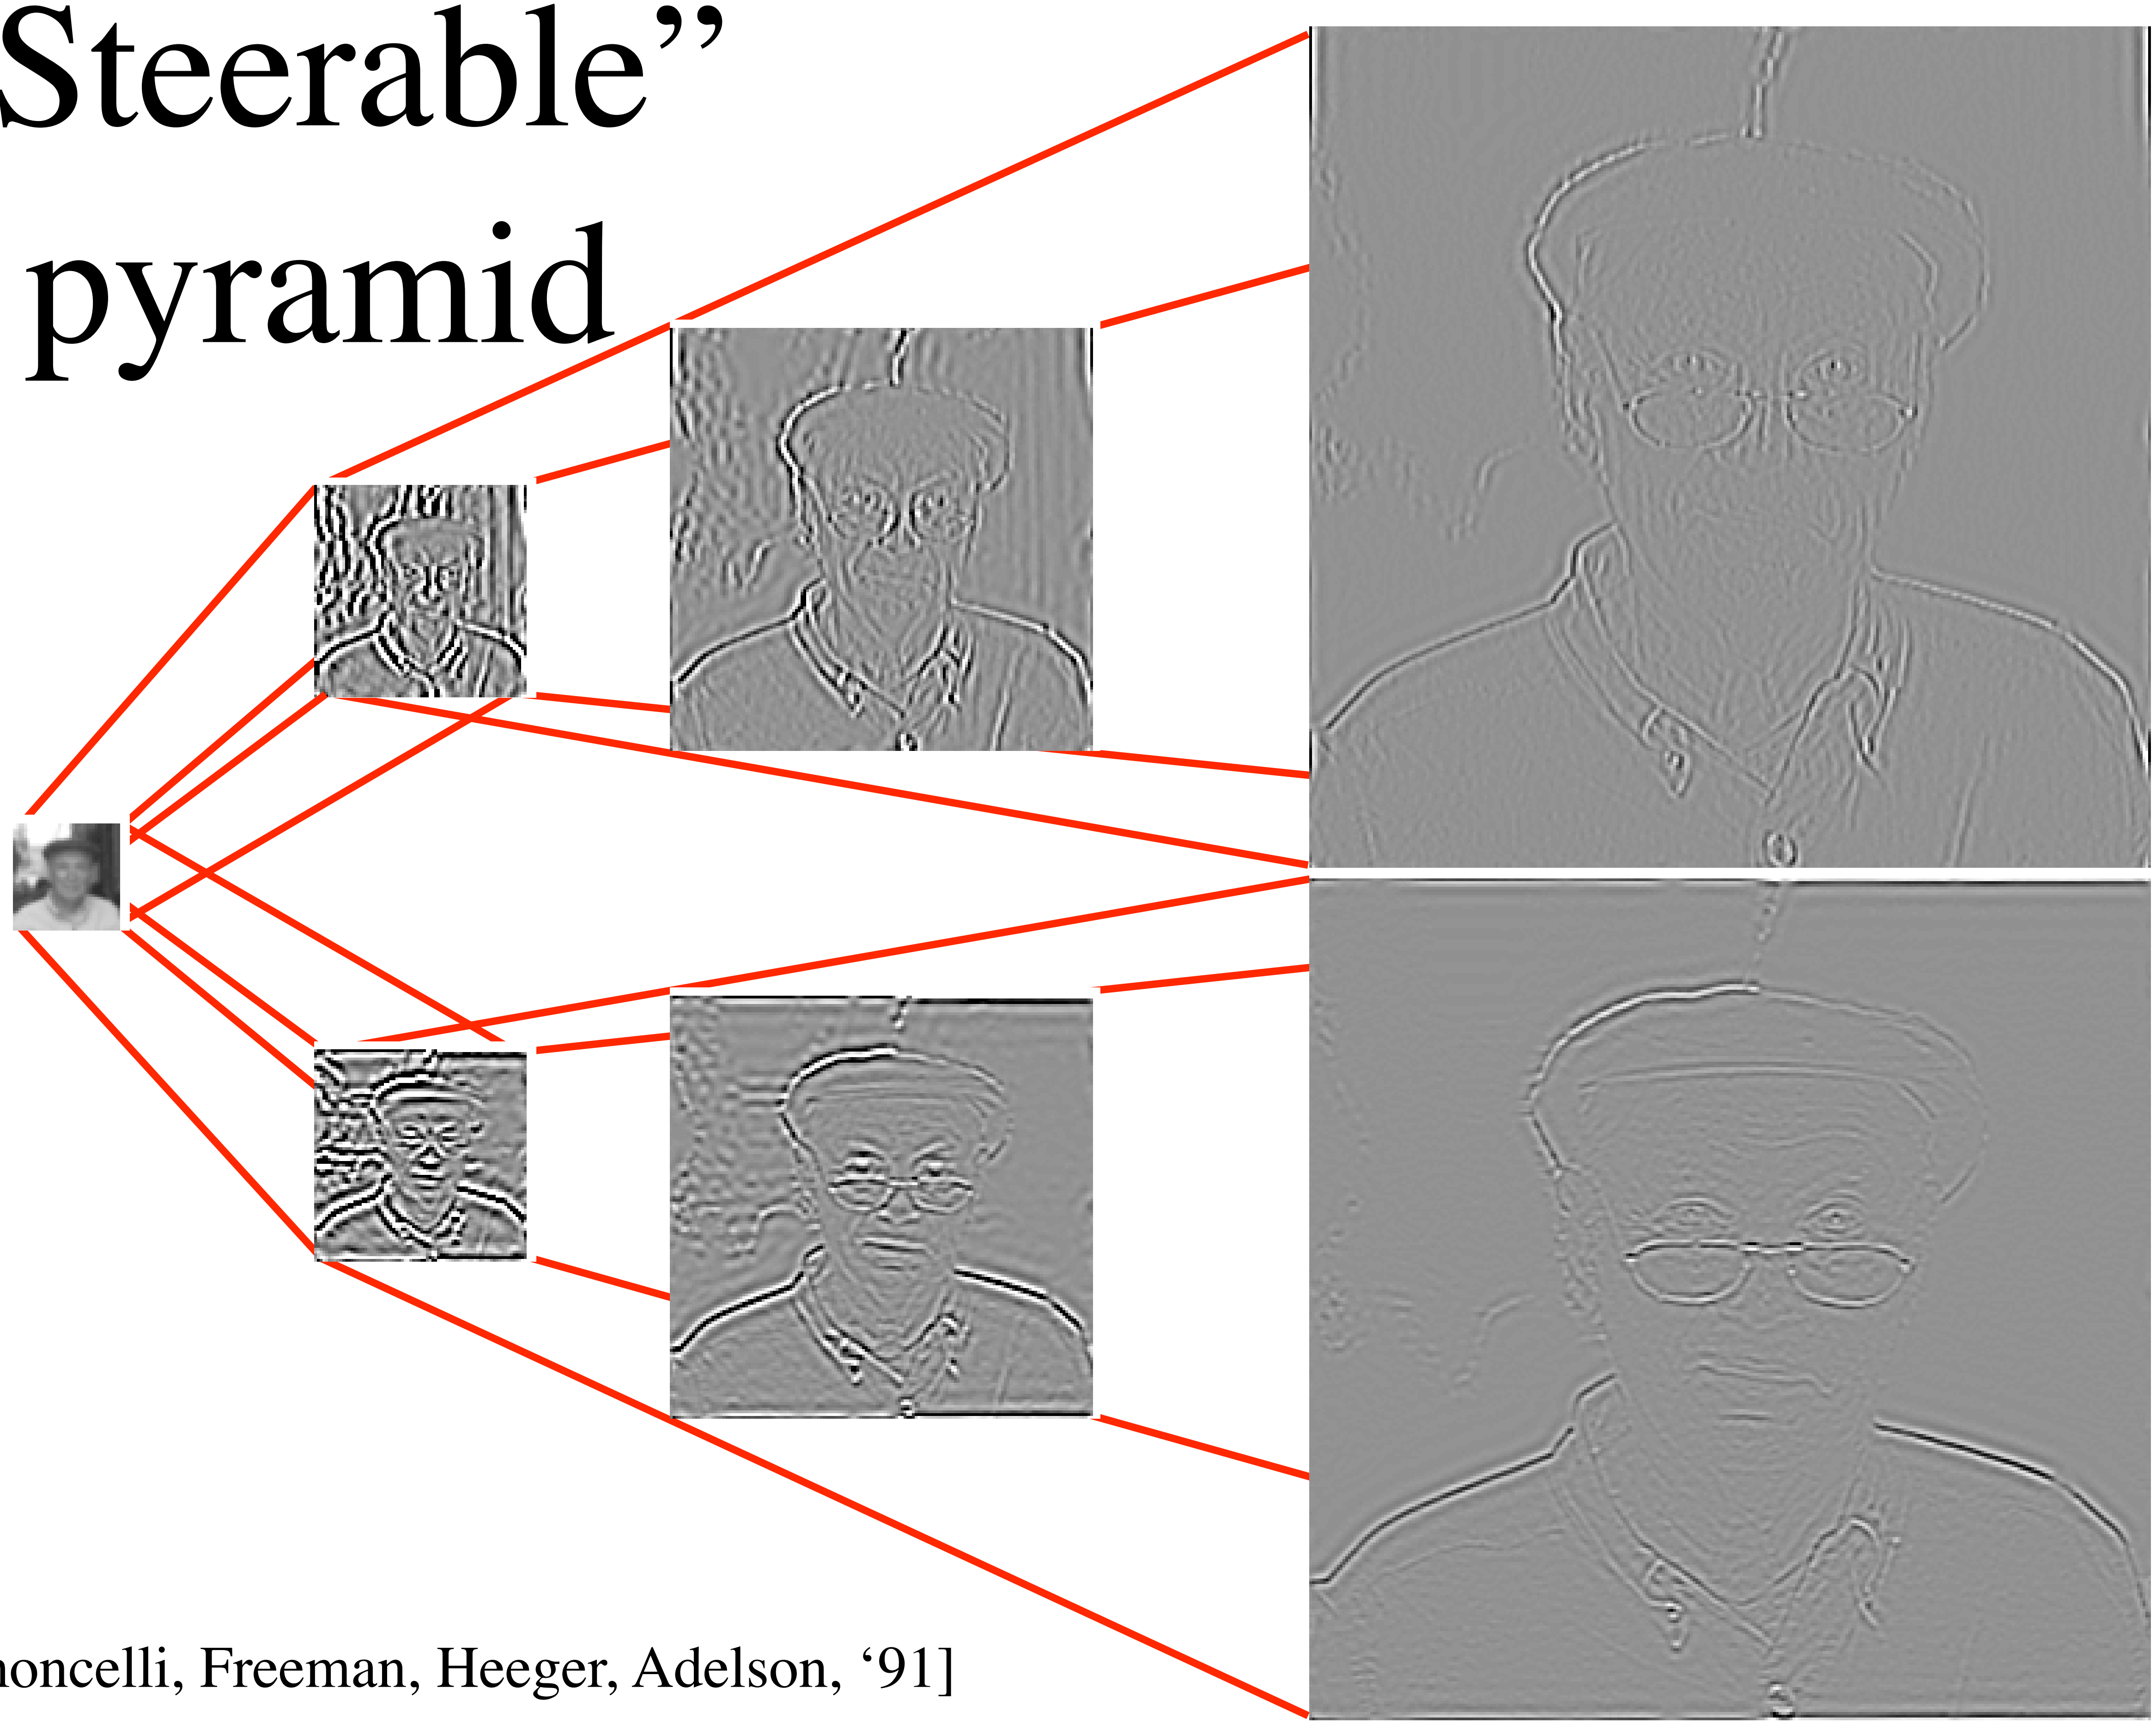

[Simoncelli, Freeman, Heeger, Adelson, ‘91]

Steerable pyramid

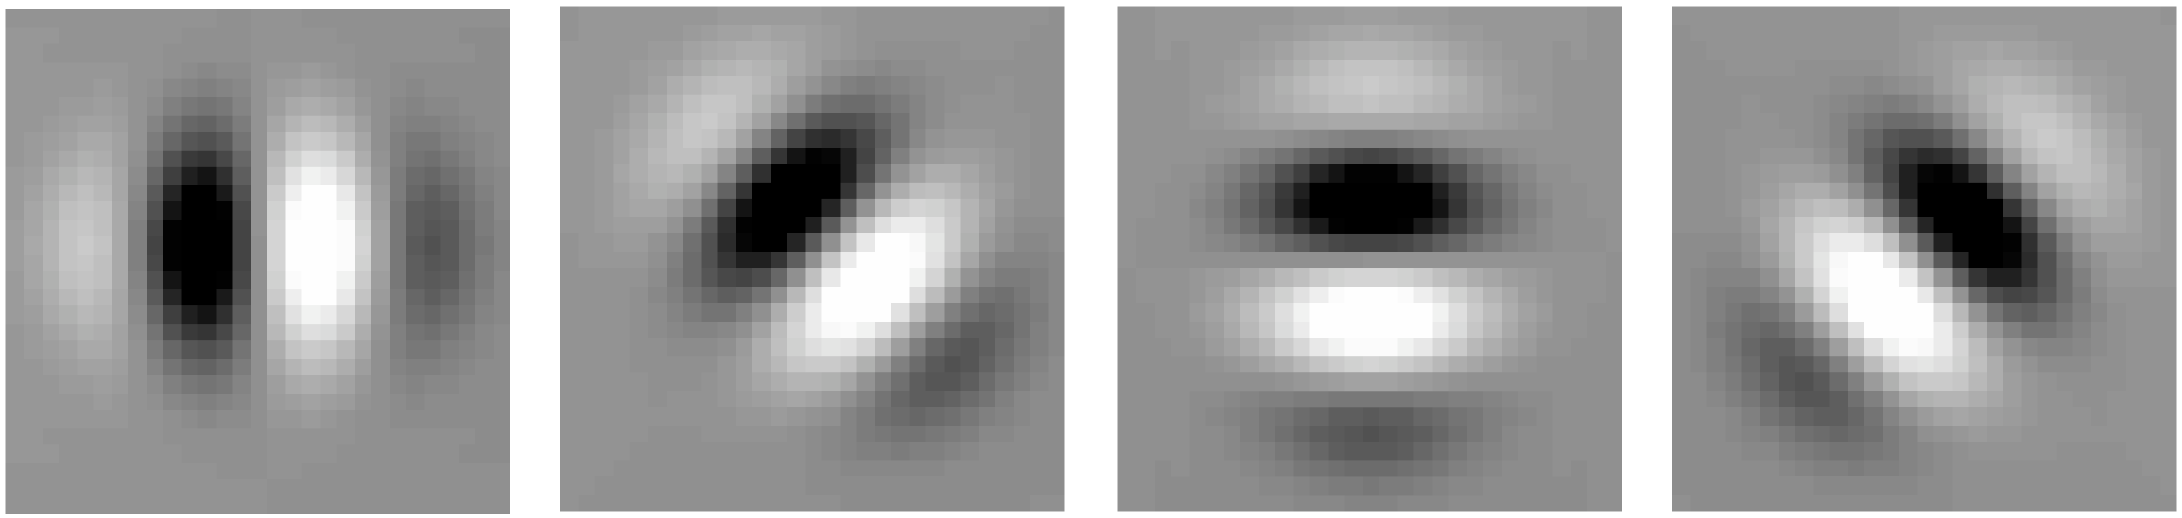

- Basis functions are *K*th derivative operators, related by translation/dilation/rotation
- Tight frame (*4(K-1)/3* overcomplete)
- Translation-invariance, rotation-invariance

[Freeman & Adelson 1991; Simoncelli et.al., 1992; Simoncelli & Freeman 1995]

Denoising

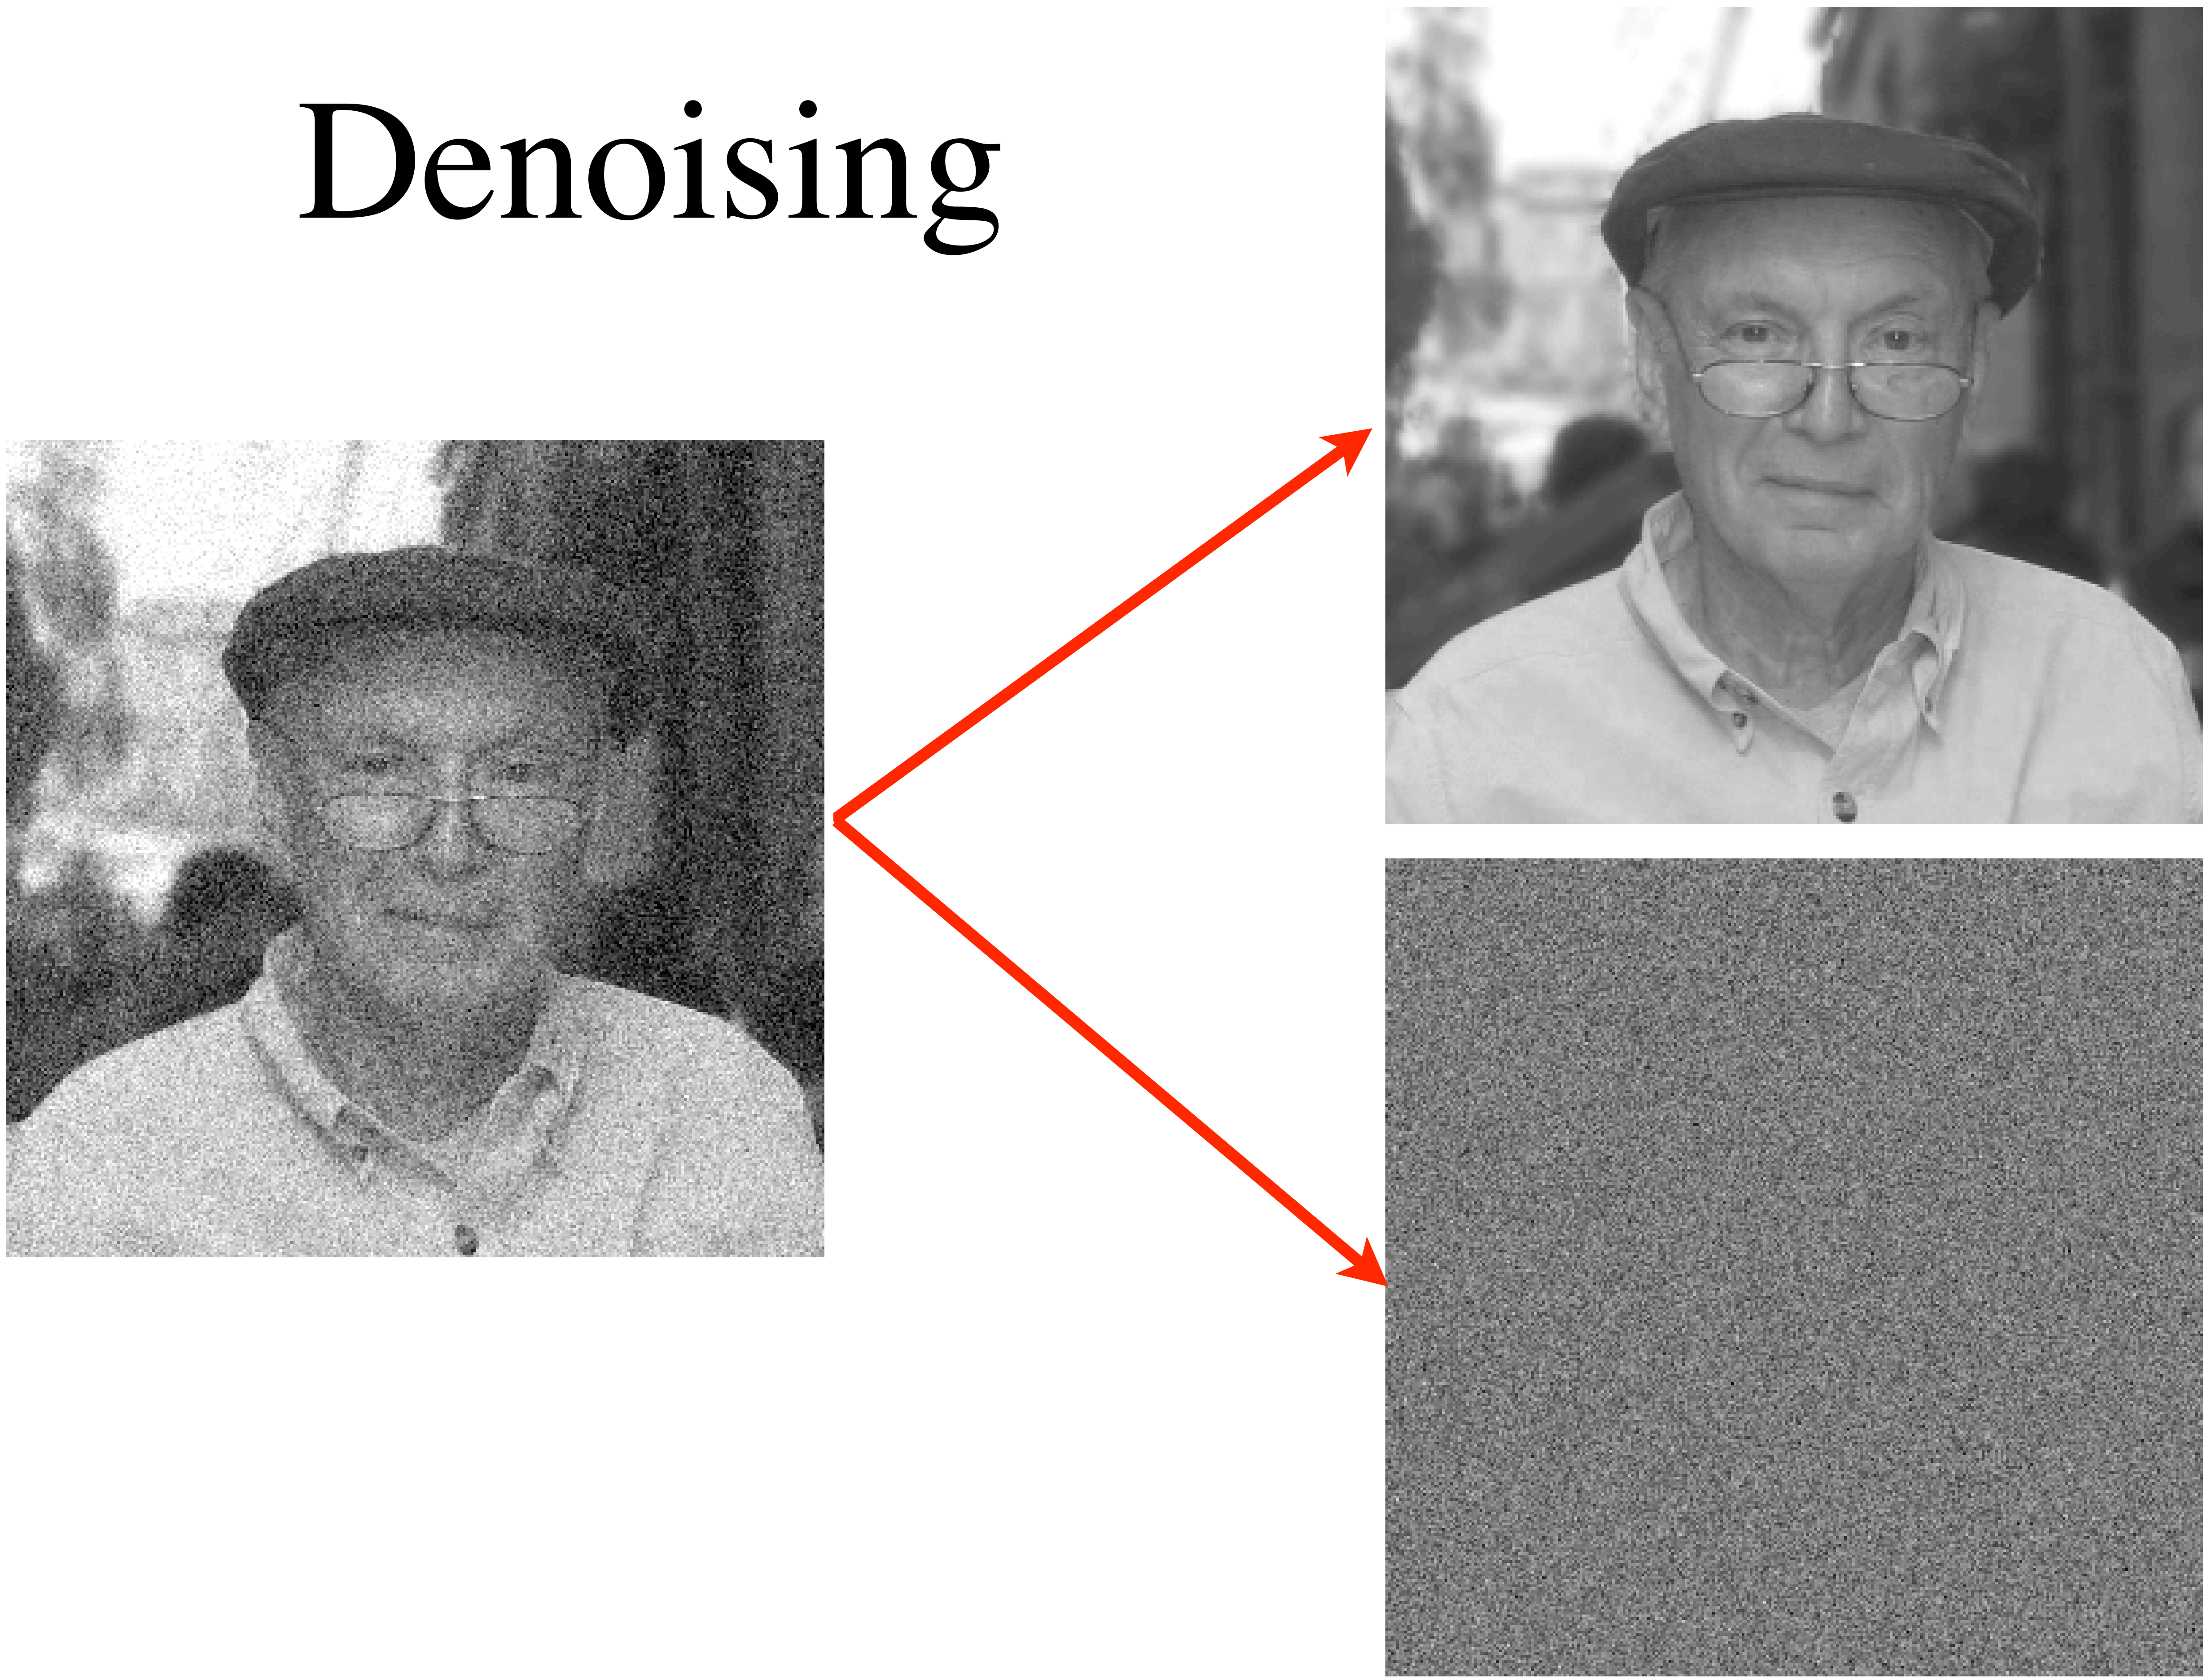

Pyramid denoising

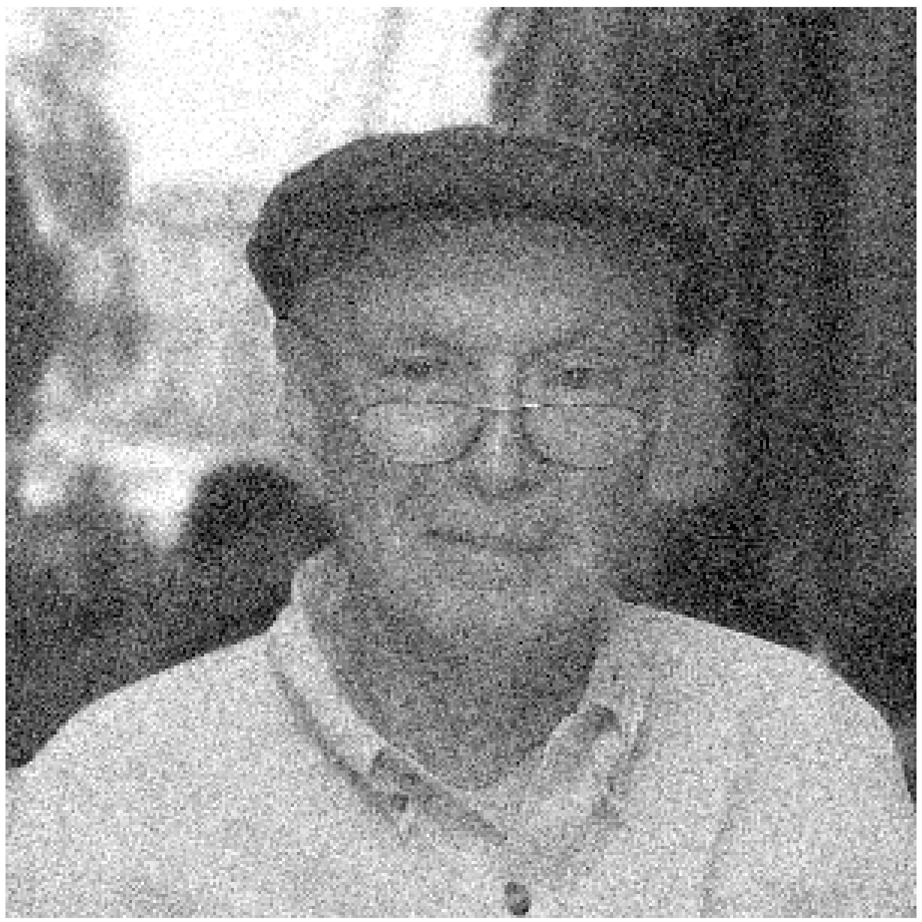

How do we distinguish signal from noise?

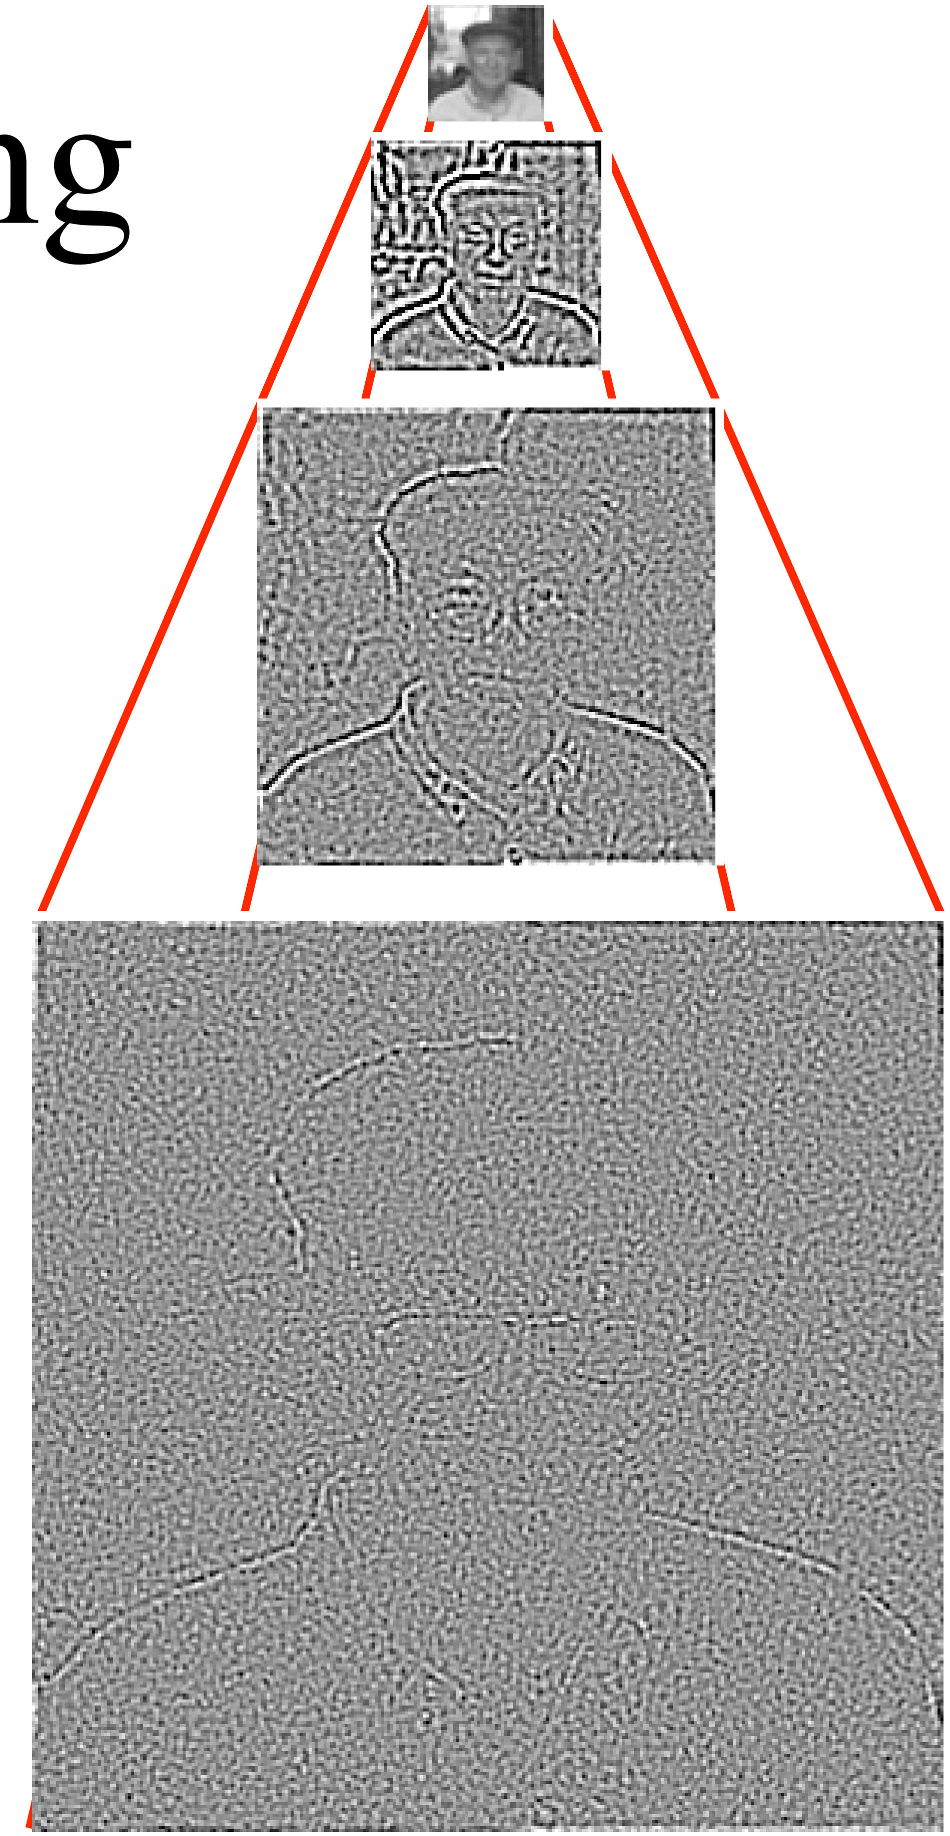

Bayesian denoising framework

- Signal: x
- Noisy observation: y
- Bayes' least squares (BLS) solution is conditional mean:

$$
\begin{aligned}
\hat{x}(y) &= \mathbb{E}(x|y) \\
&\propto \int_x x \; \mathcal{P}(y|x) \; \mathcal{P}(x)
\end{aligned}
$$

Image statistical models

I. (1950's): Fourier transform + Gaussian marginals

II. (late 80's/early 90's): Wavelets + kurtotic marginals

III. (late 90's -): Wavelets + adaptive local variance

Substantial increase in model accuracy
(at the cost of increased model complexity)

I. Classical Bayes denoising

If signal is Gaussian, BLS estimator is linear:

$$\hat{x}(y) = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_n^2} \cdot y$$

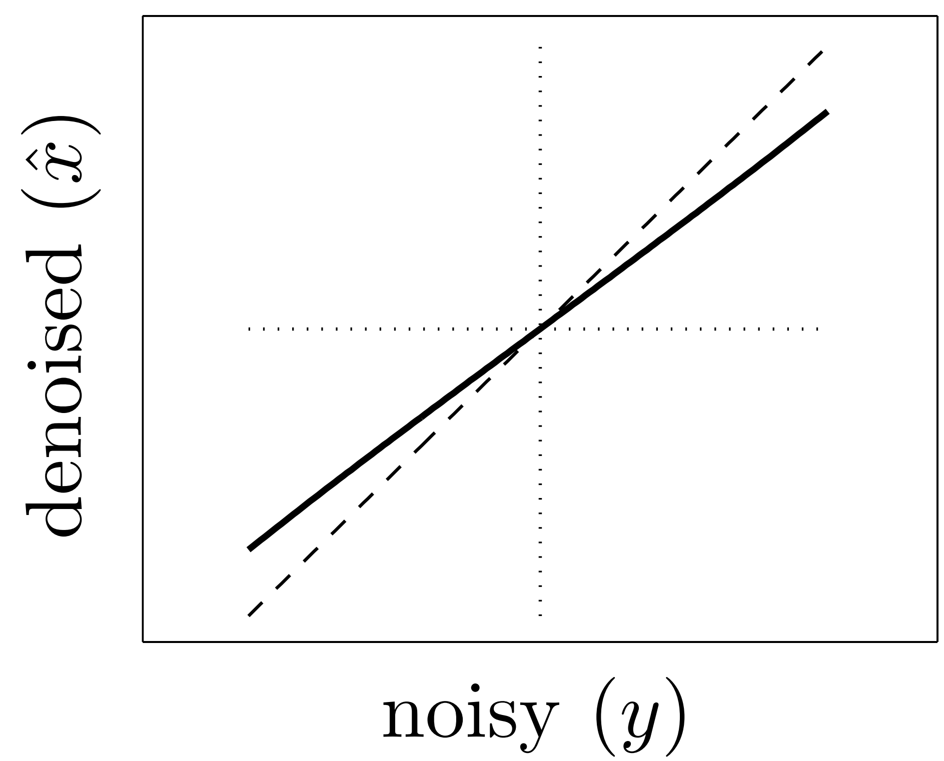

Coefficient distributions

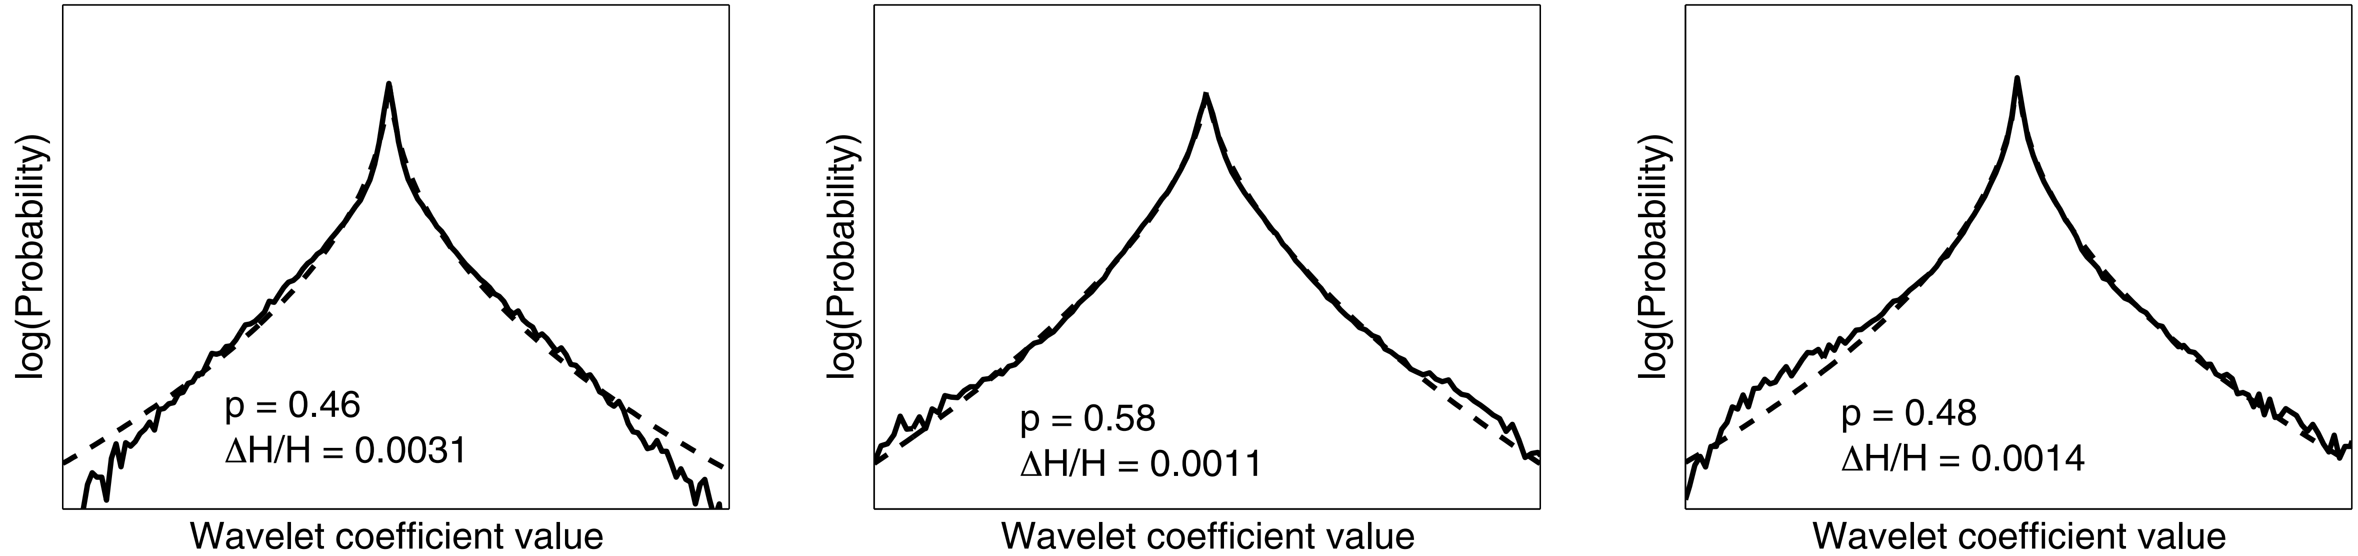

Well-fit by a generalized Gaussian:

$$P(x) \propto \exp -|x/s|^p$$

[Mallat, ‘89; Simoncelli&Adelson ‘96; Mouline&Liu ‘99; etc]

II. Bayesian coring

- Assume marginal distribution:

$$P(x) \propto \exp -|x/s|^p$$

- Then Bayes estimator is generally nonlinear:

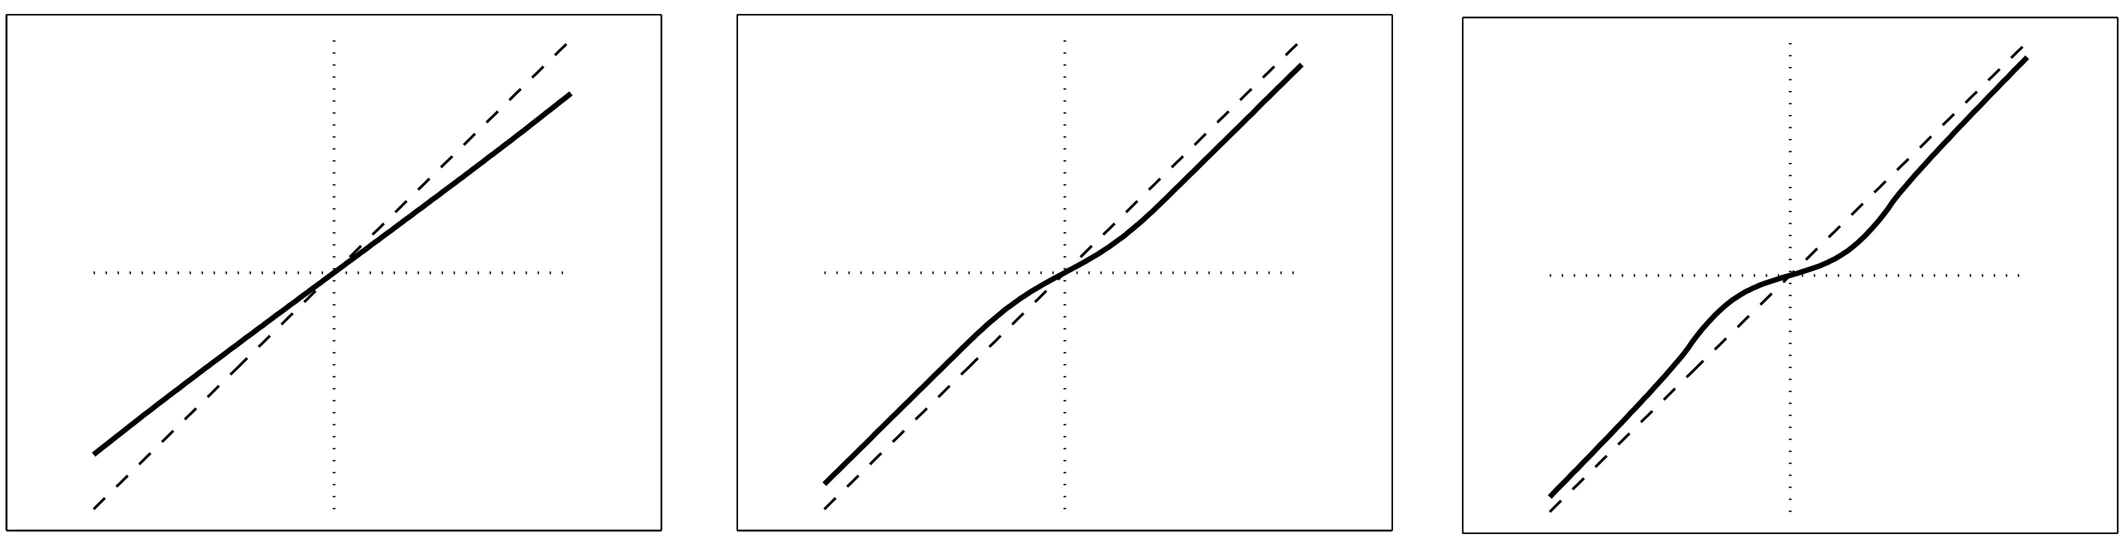

p = 2.0 *p = 1.0* *p = 0.5*

[Simoncelli & Adelson, '96]

Joint statistics

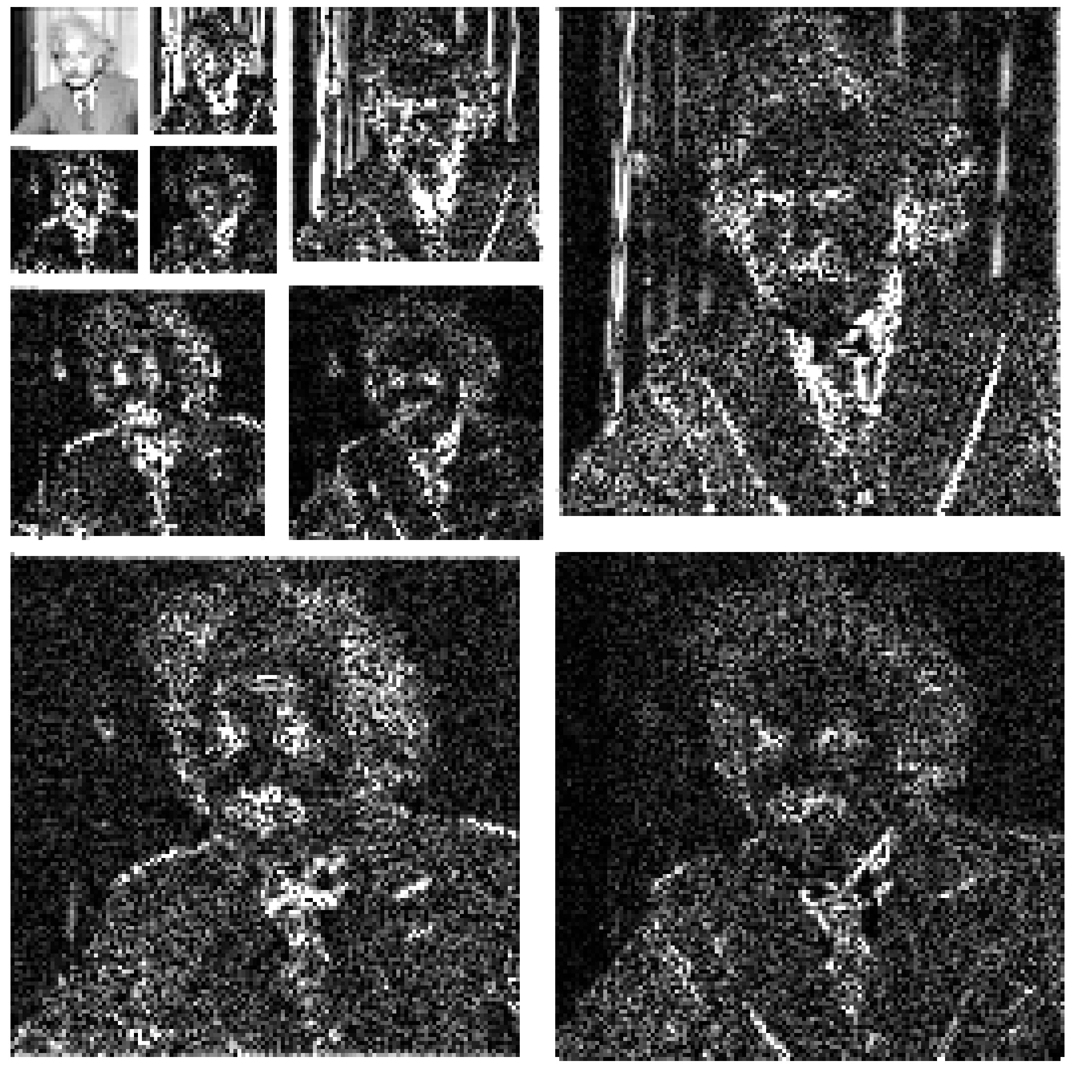

- Large-magnitude values are found at neighboring positions, orientations, and scales.

[Simoncelli, ‘97; Buccigrossi & Simoncelli, ‘97]

Joint statistics

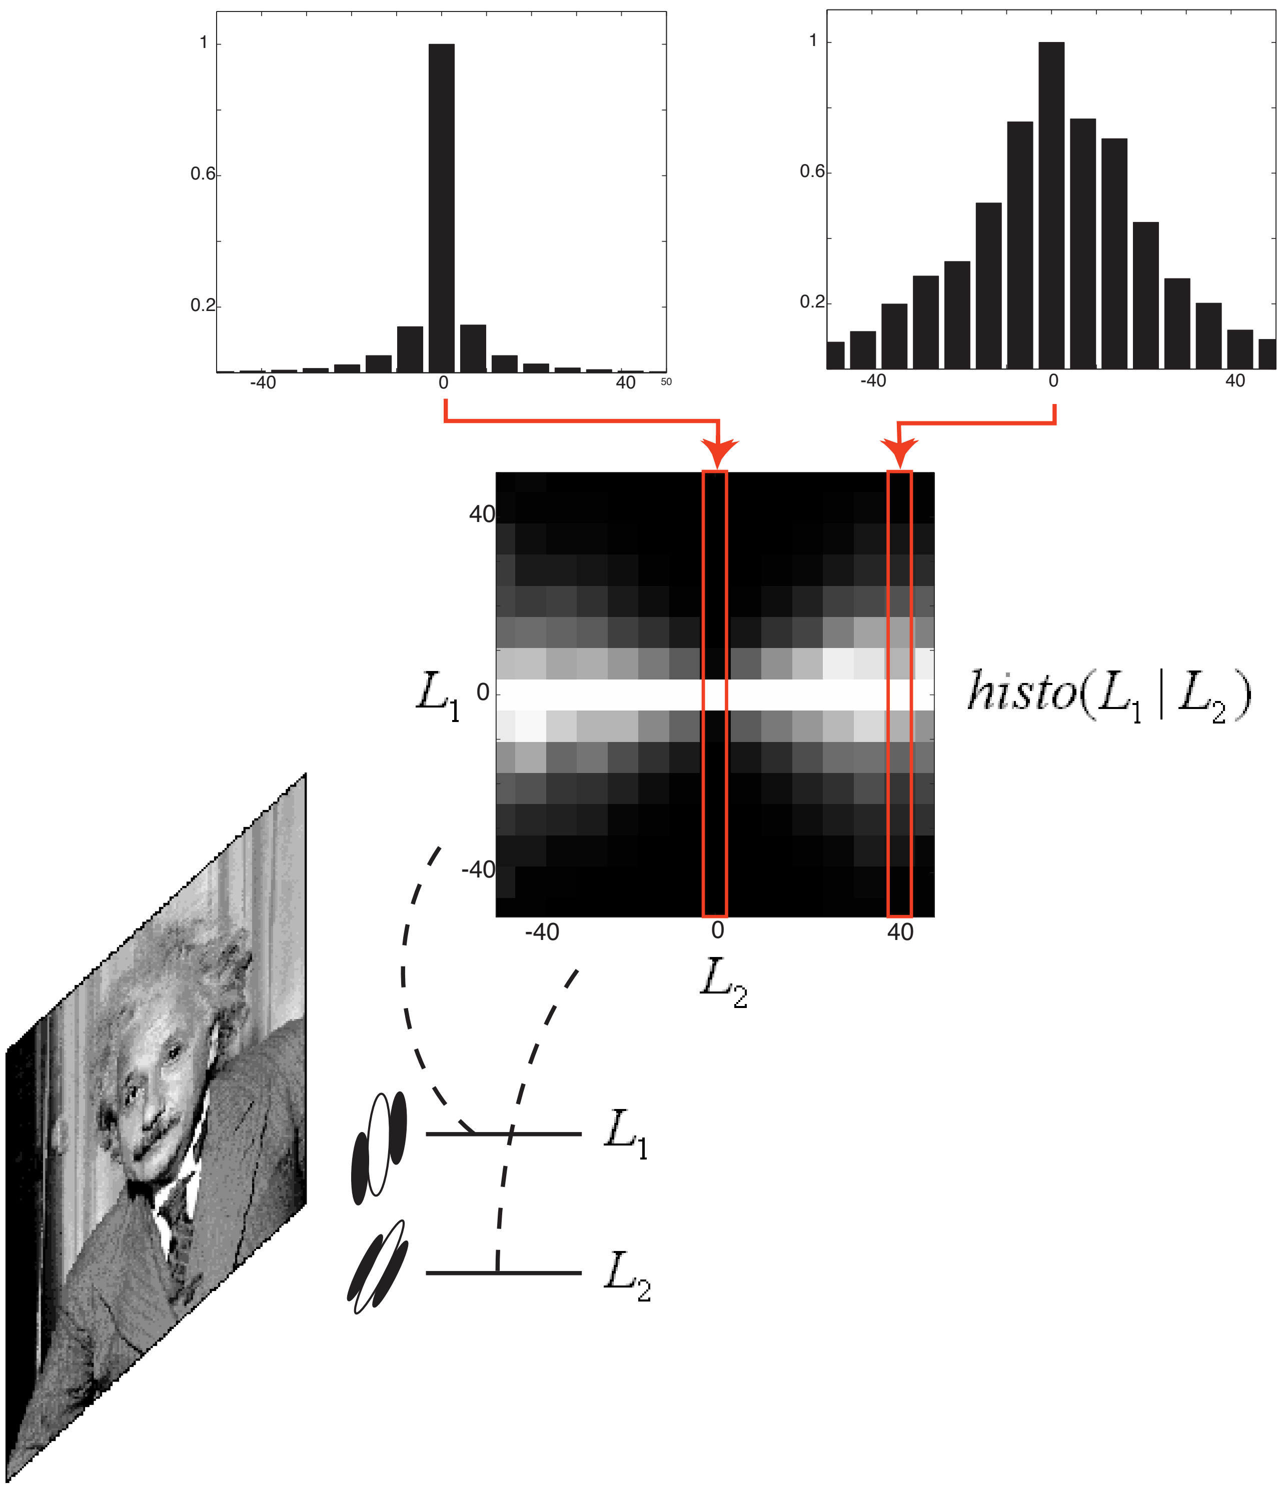

[Simoncelli, '97; Buccigrossi & Simoncelli, '97]

Joint GSM model

Model generalized neighborhood of coefficients as a Gaussian Scale Mixture (GSM) [Andrews & Mallows '74]:

$\vec{x} = \sqrt{z}\,\vec{u}$, where

- z and $\vec{u}$ are independent
- $\vec{x}|z$ is Gaussian, with covariance zC_u
- marginals are always leptokurtotic

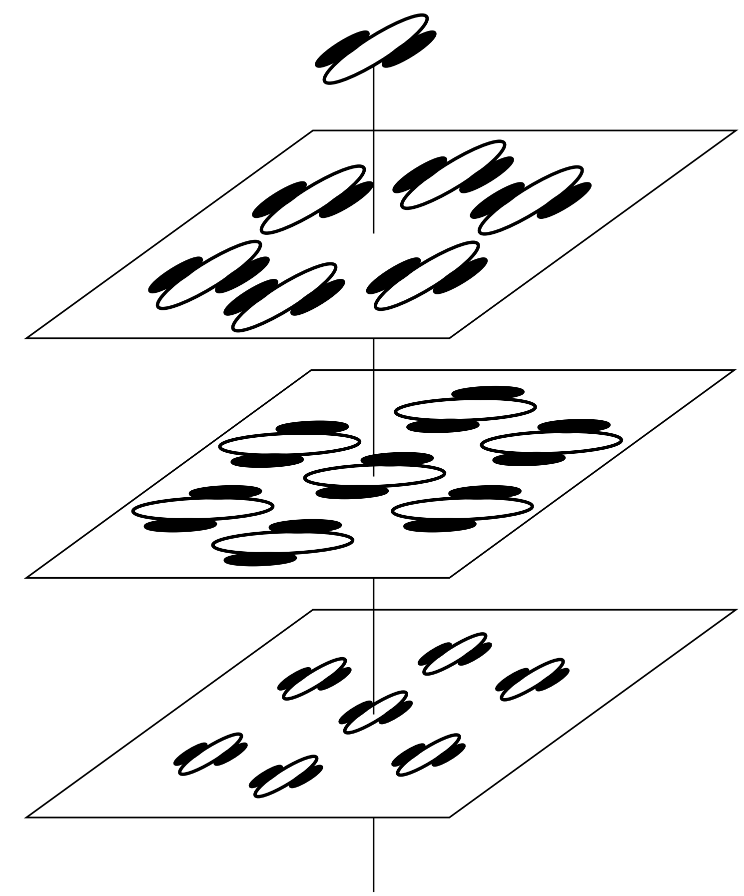

[Wainwright & Simoncelli, '99]

Simulation

Image data

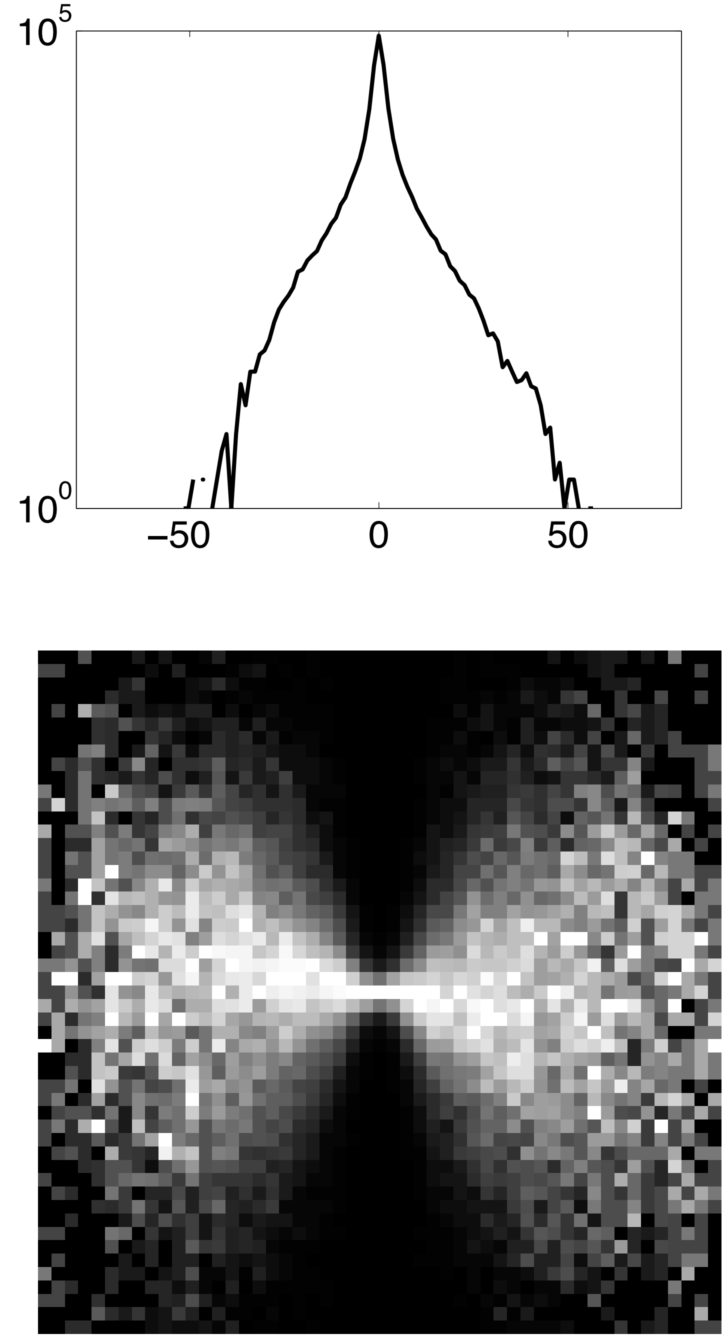

GSM simulation

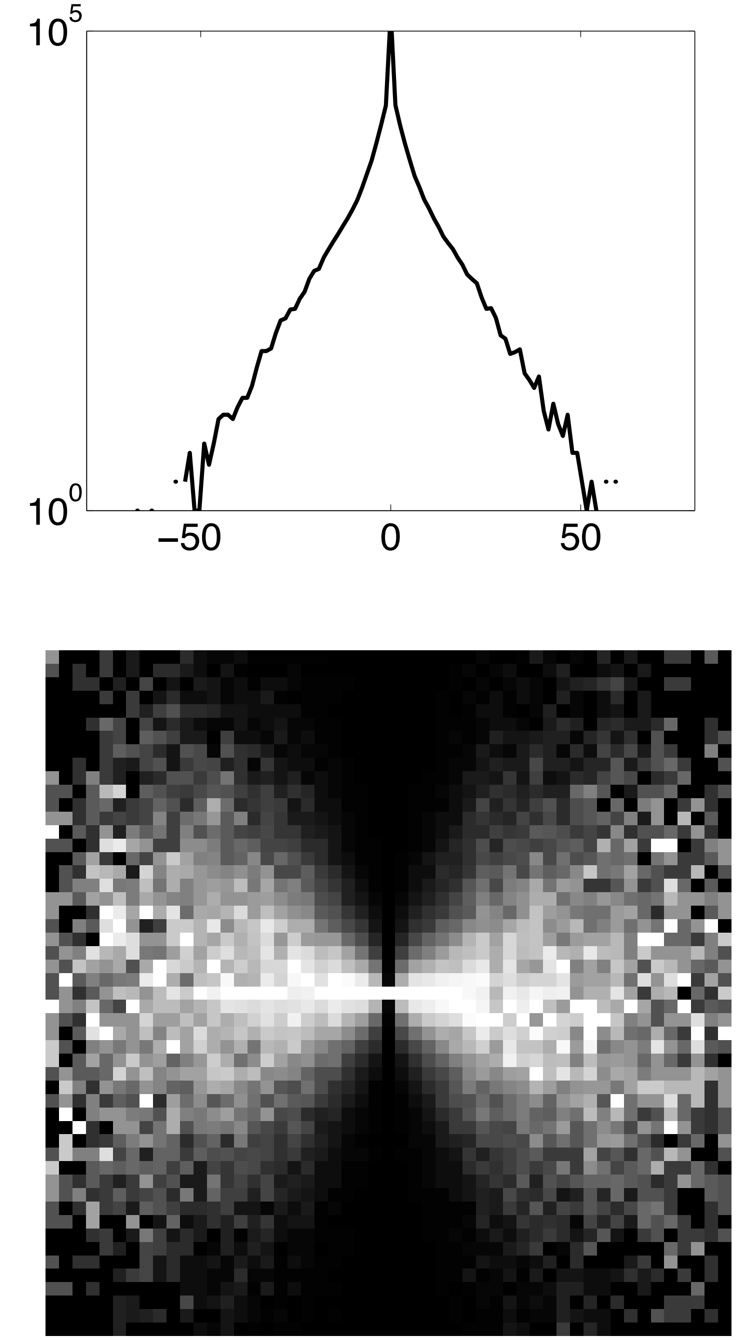

[Wainwright & Simoncelli, ‘99]

III. Joint Bayes denoising

$$\begin{aligned}\mathbb{E}(x|\vec{y}) &= \int dz\ \mathcal{P}(z|\vec{y})\ \mathbb{E}(x|\vec{y},z) \\ &= \int dz\ \mathcal{P}(z|\vec{y})\ \ \left[zC_u(zC_u+C_w)^{-1}\vec{y}\right]_{\mathrm{ctr}}\end{aligned}$$

where

$$\mathcal{P}(z|\vec{y}) = \frac{\mathcal{P}(\vec{y}|z)\ \mathcal{P}(z)}{\mathcal{P}\vec{y}}, \quad \mathcal{P}(\vec{y}|z) = \frac{\exp(-\vec{y}^T(zC_u+C_w)^{-1}\vec{y}/2)}{\sqrt{(2\pi)^N|zC_u+C_w|}}$$

Numerical computation of solution is reasonably efficient if one jointly diagonalizes C_u and C_w ...

[Portilla, Strela, Wainwright, Simoncelli, '03]

Example joint estimator

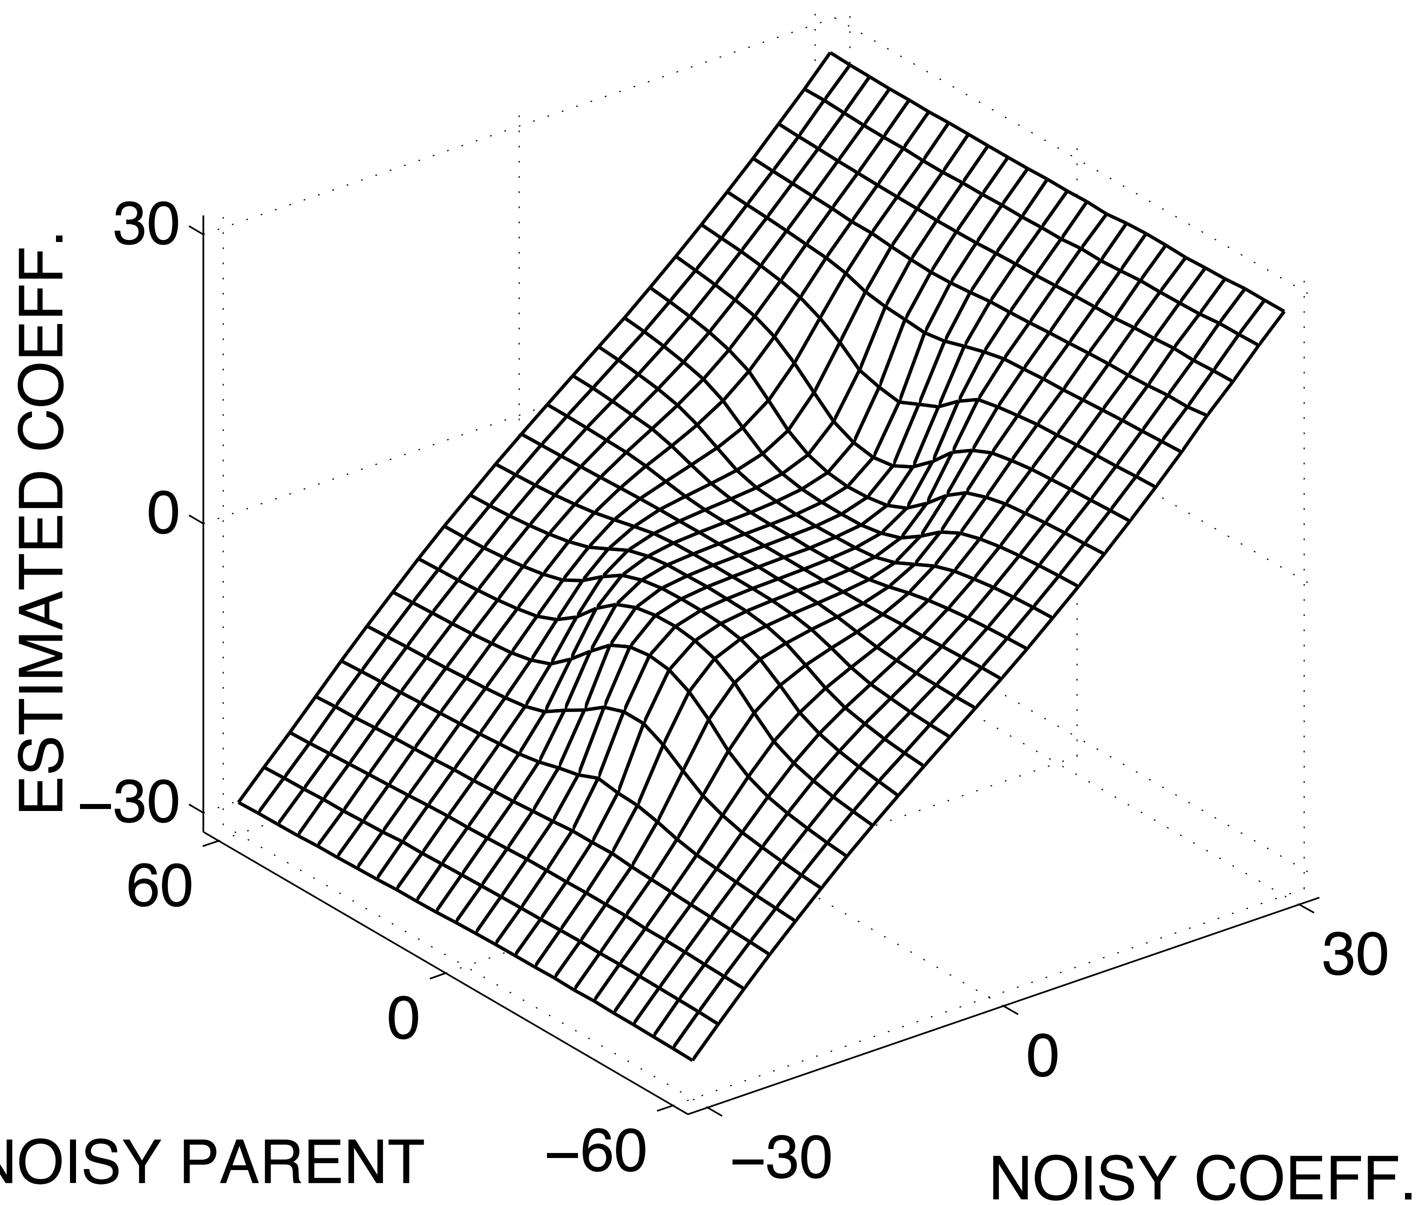

[Portilla, Wainwright, Strela, Simoncelli, ‘03;
see also: Sendur & Selesnick, ‘02]

noisy
(4.8)

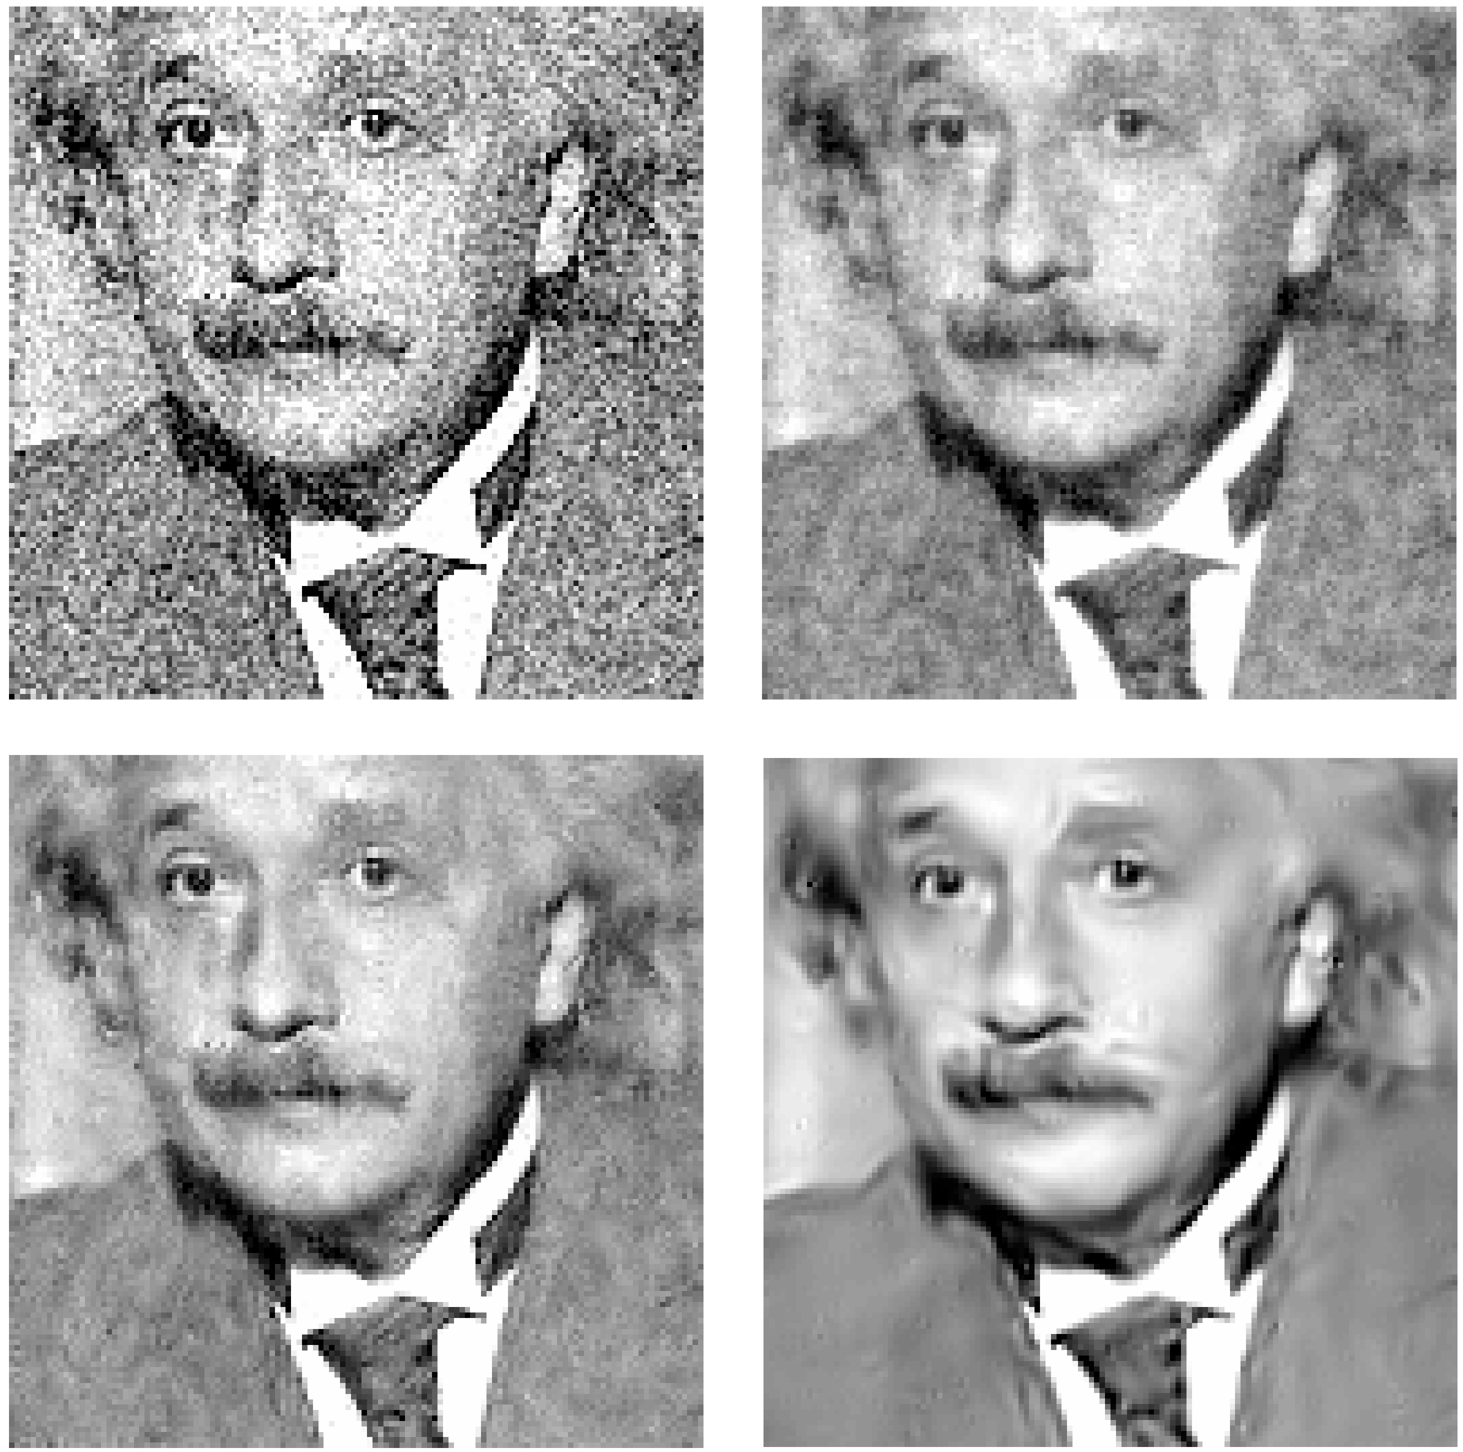

I-linear
(10.61)

II-marginal
(11.98)

III- joint
nbd: $5 \times 5 + p$
(13.60)

Original

Noisy
(22.1 dB)

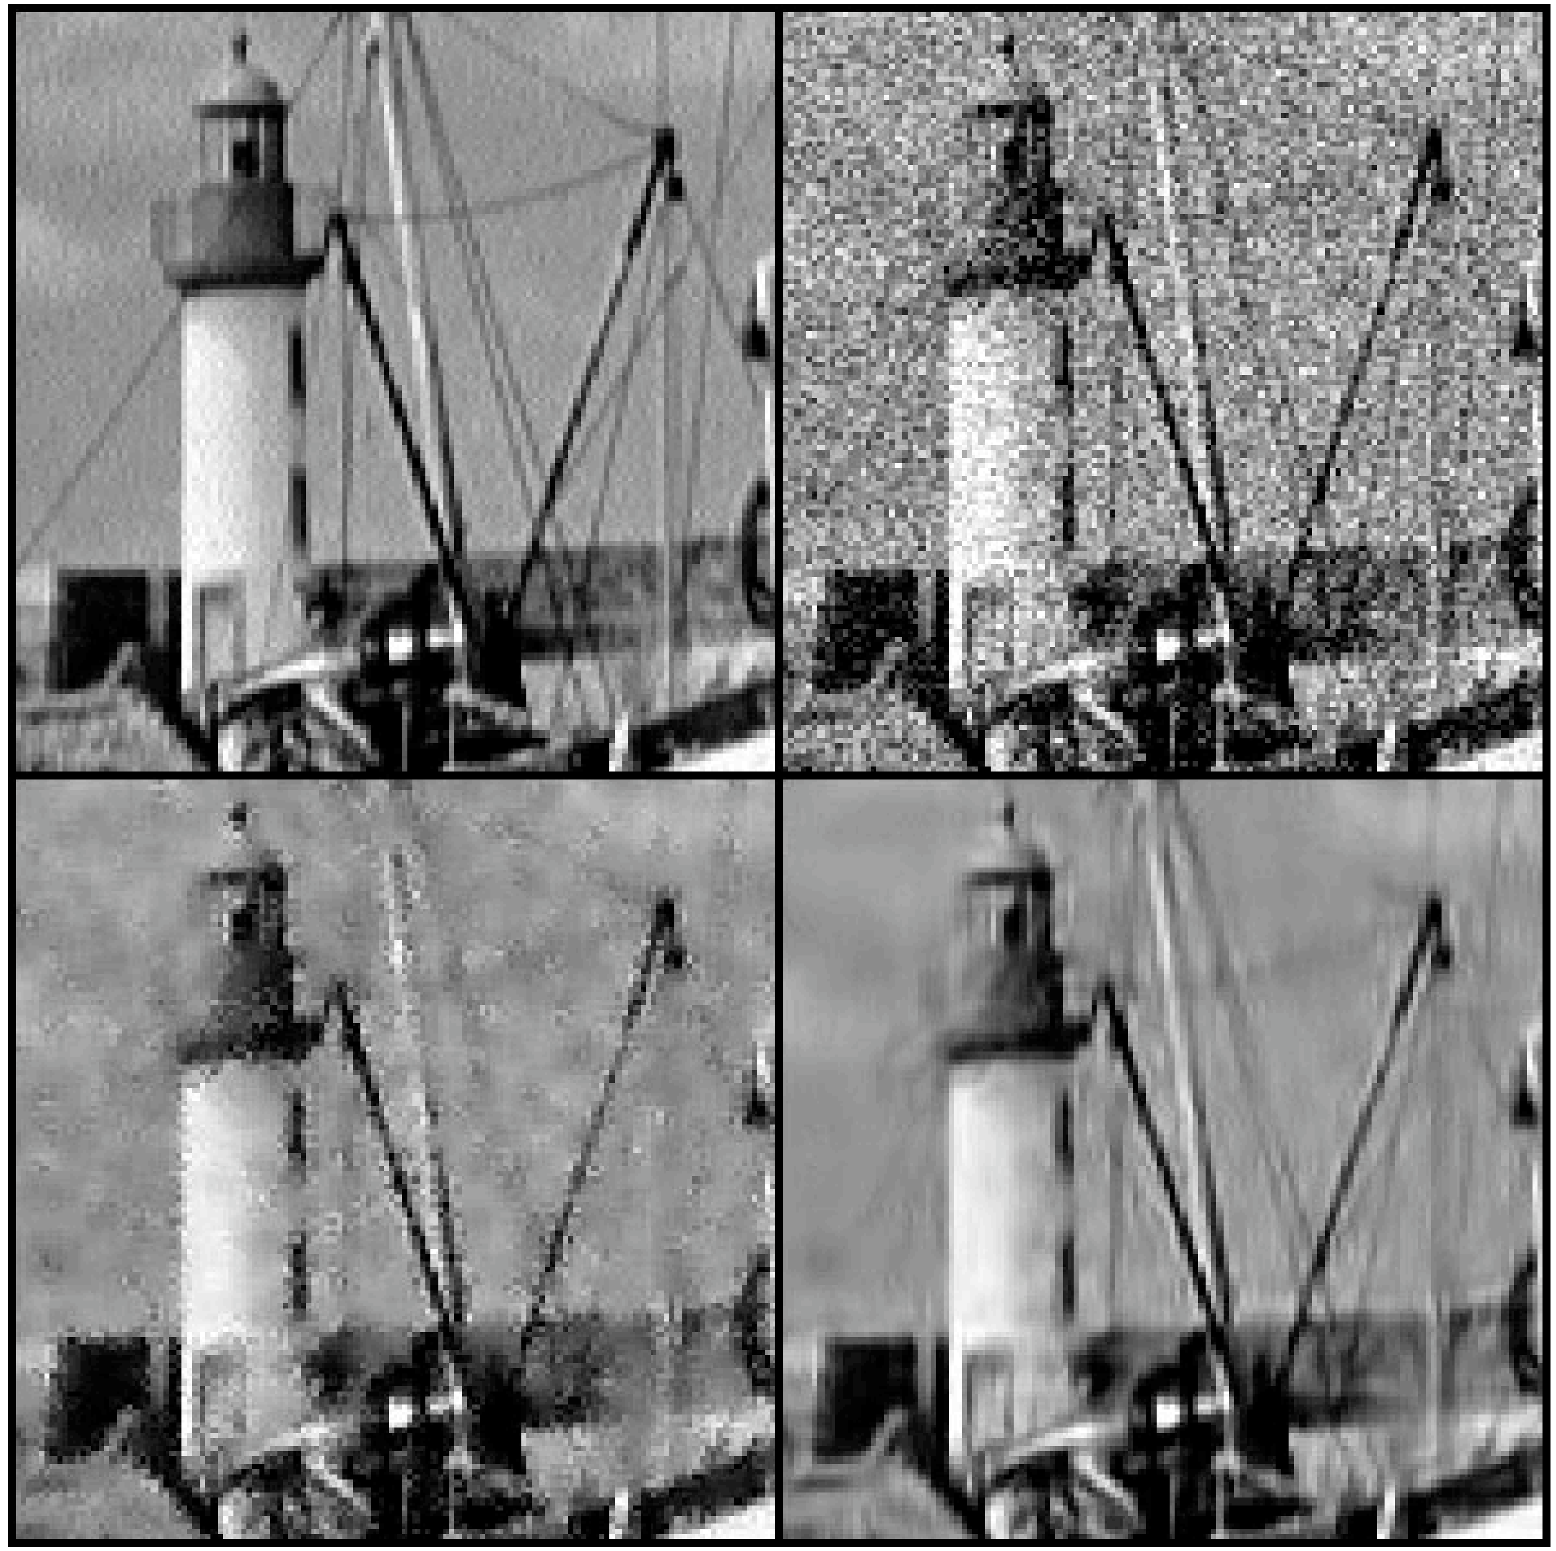

Matlab's
wiener2
(28 dB)

BLS-GSM
(30.5 dB)

Original

Noisy
(8.1 dB)

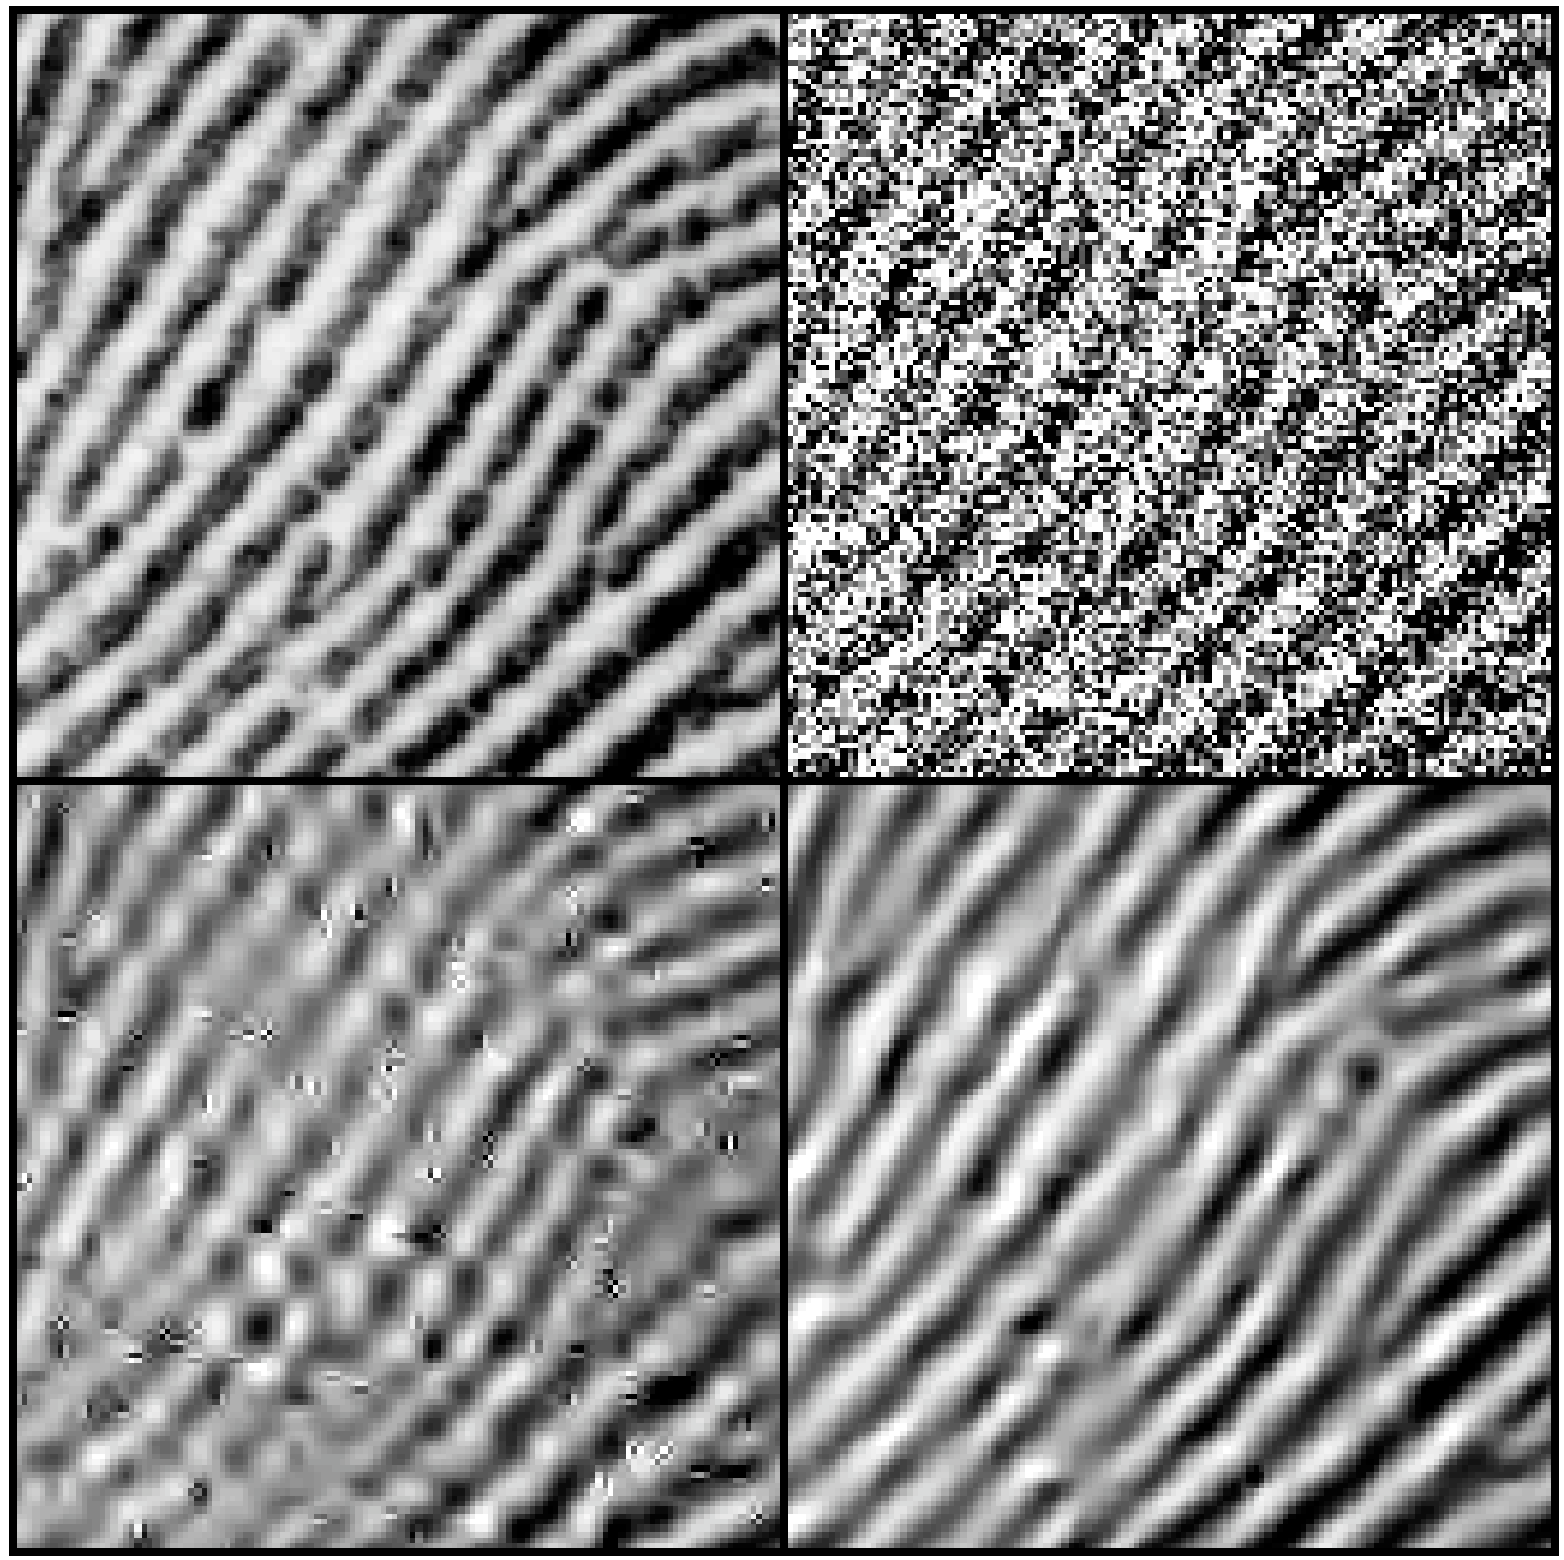

UndecWvlt
HardThresh
(19.0 dB)

BLS-GSM
(21.2 dB)

Real sensor noise

400 ISO

denoised

Comparison to other methods

Relative PSNR improvement as a function of noise level (averaged over three images):

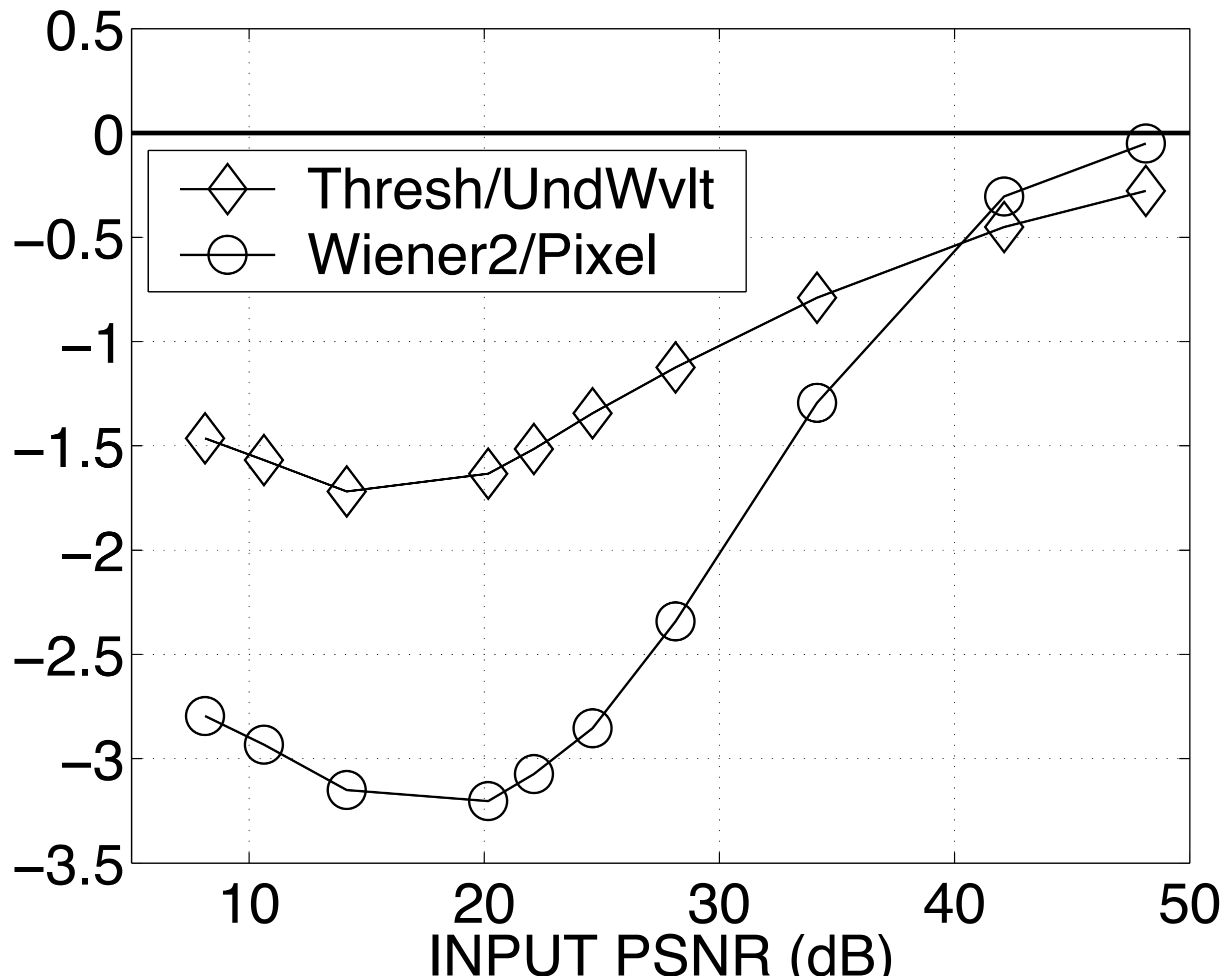

- squares: Joint model
- diamonds: soft thresholding, optimized threshold [Donoho, '95]
- circles: MatLab wiener2, optimized neighborhood [Lee, '80]

Pyramid denoising

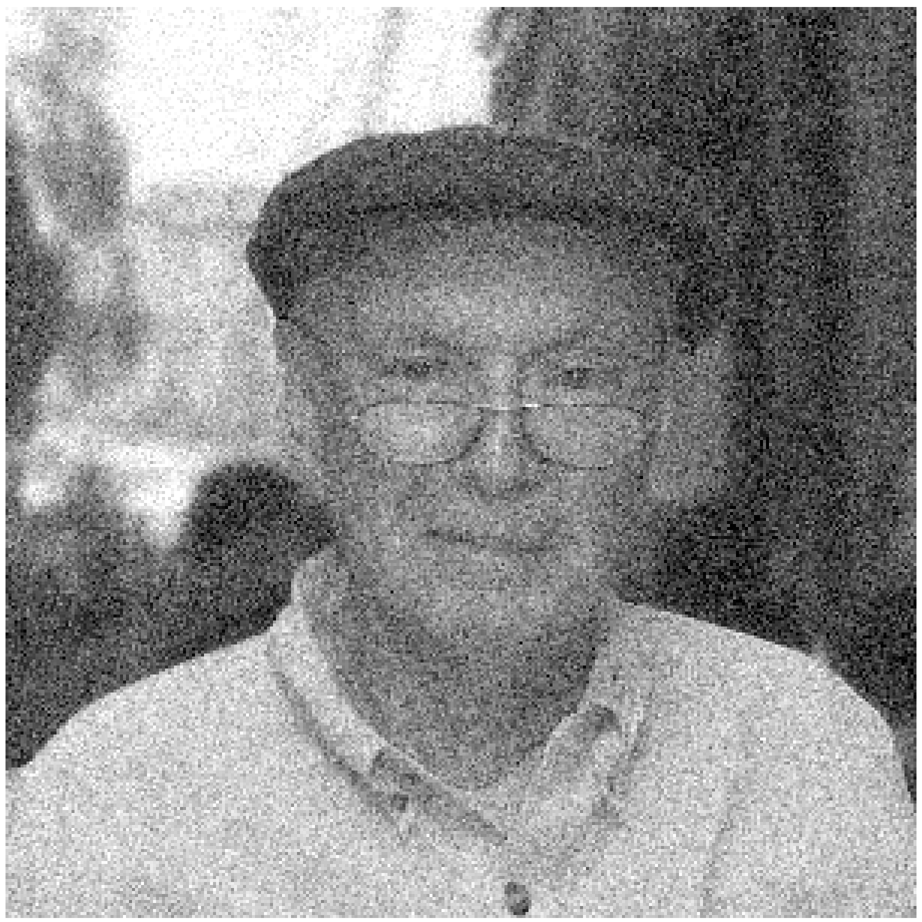

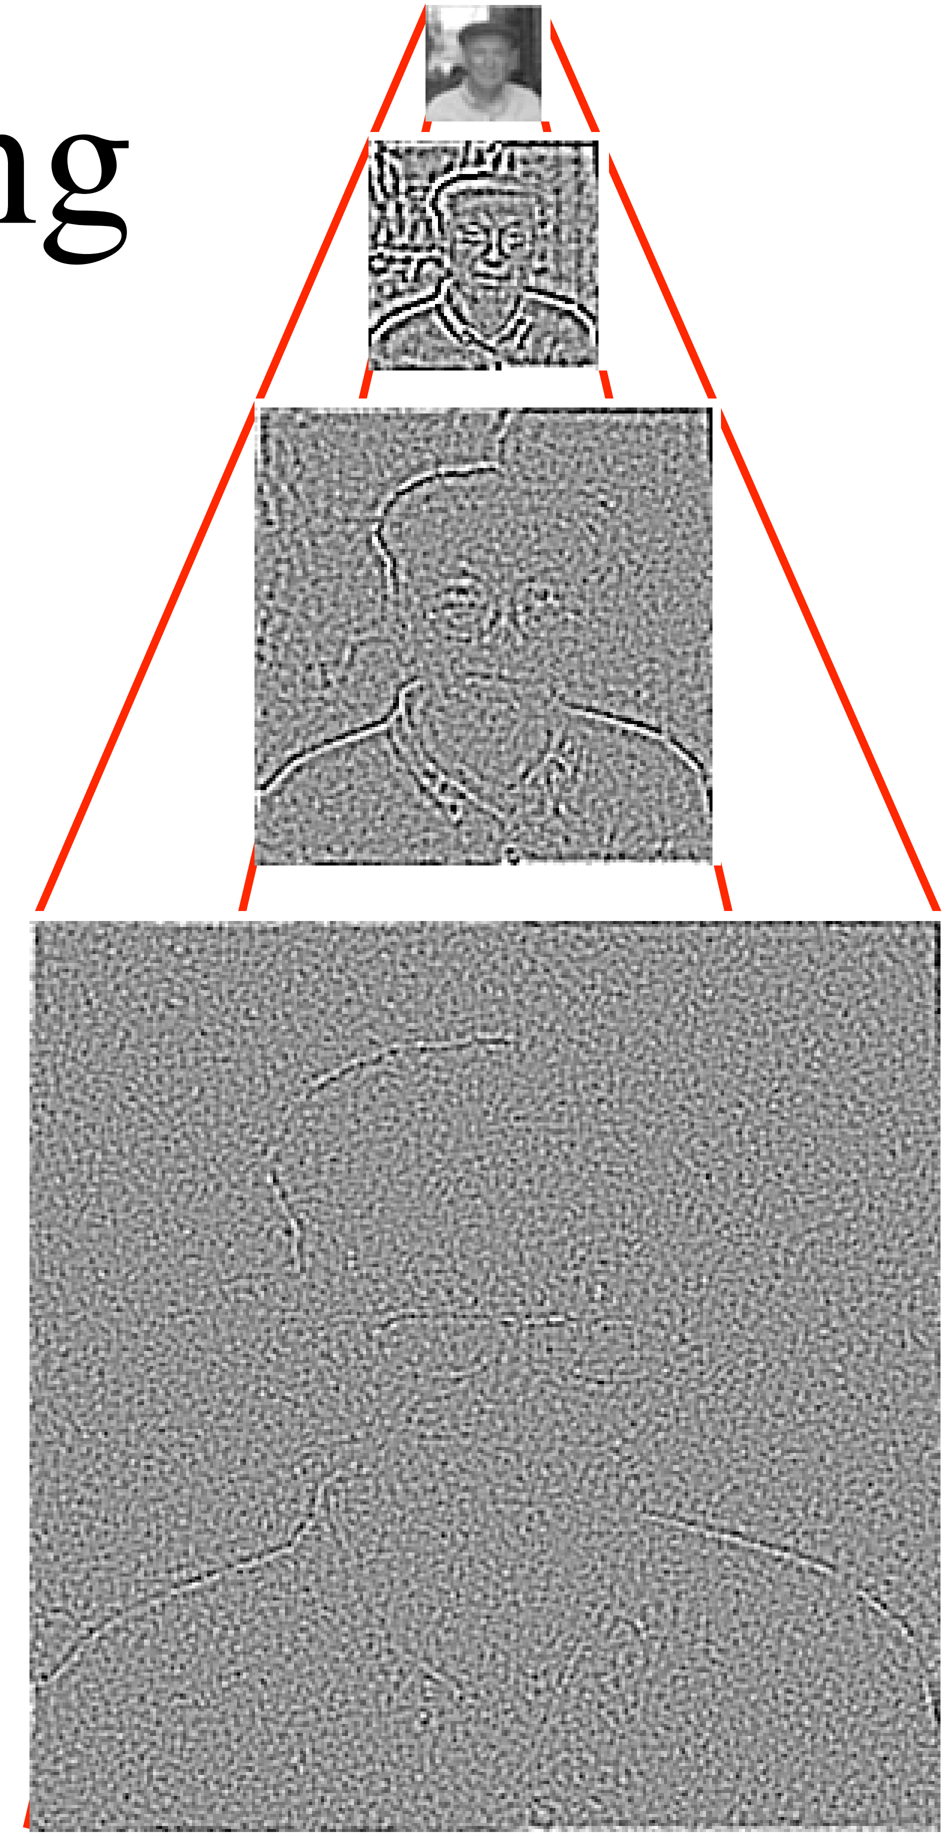

How do we distinguish signal from noise?

“Steerable” pyramid

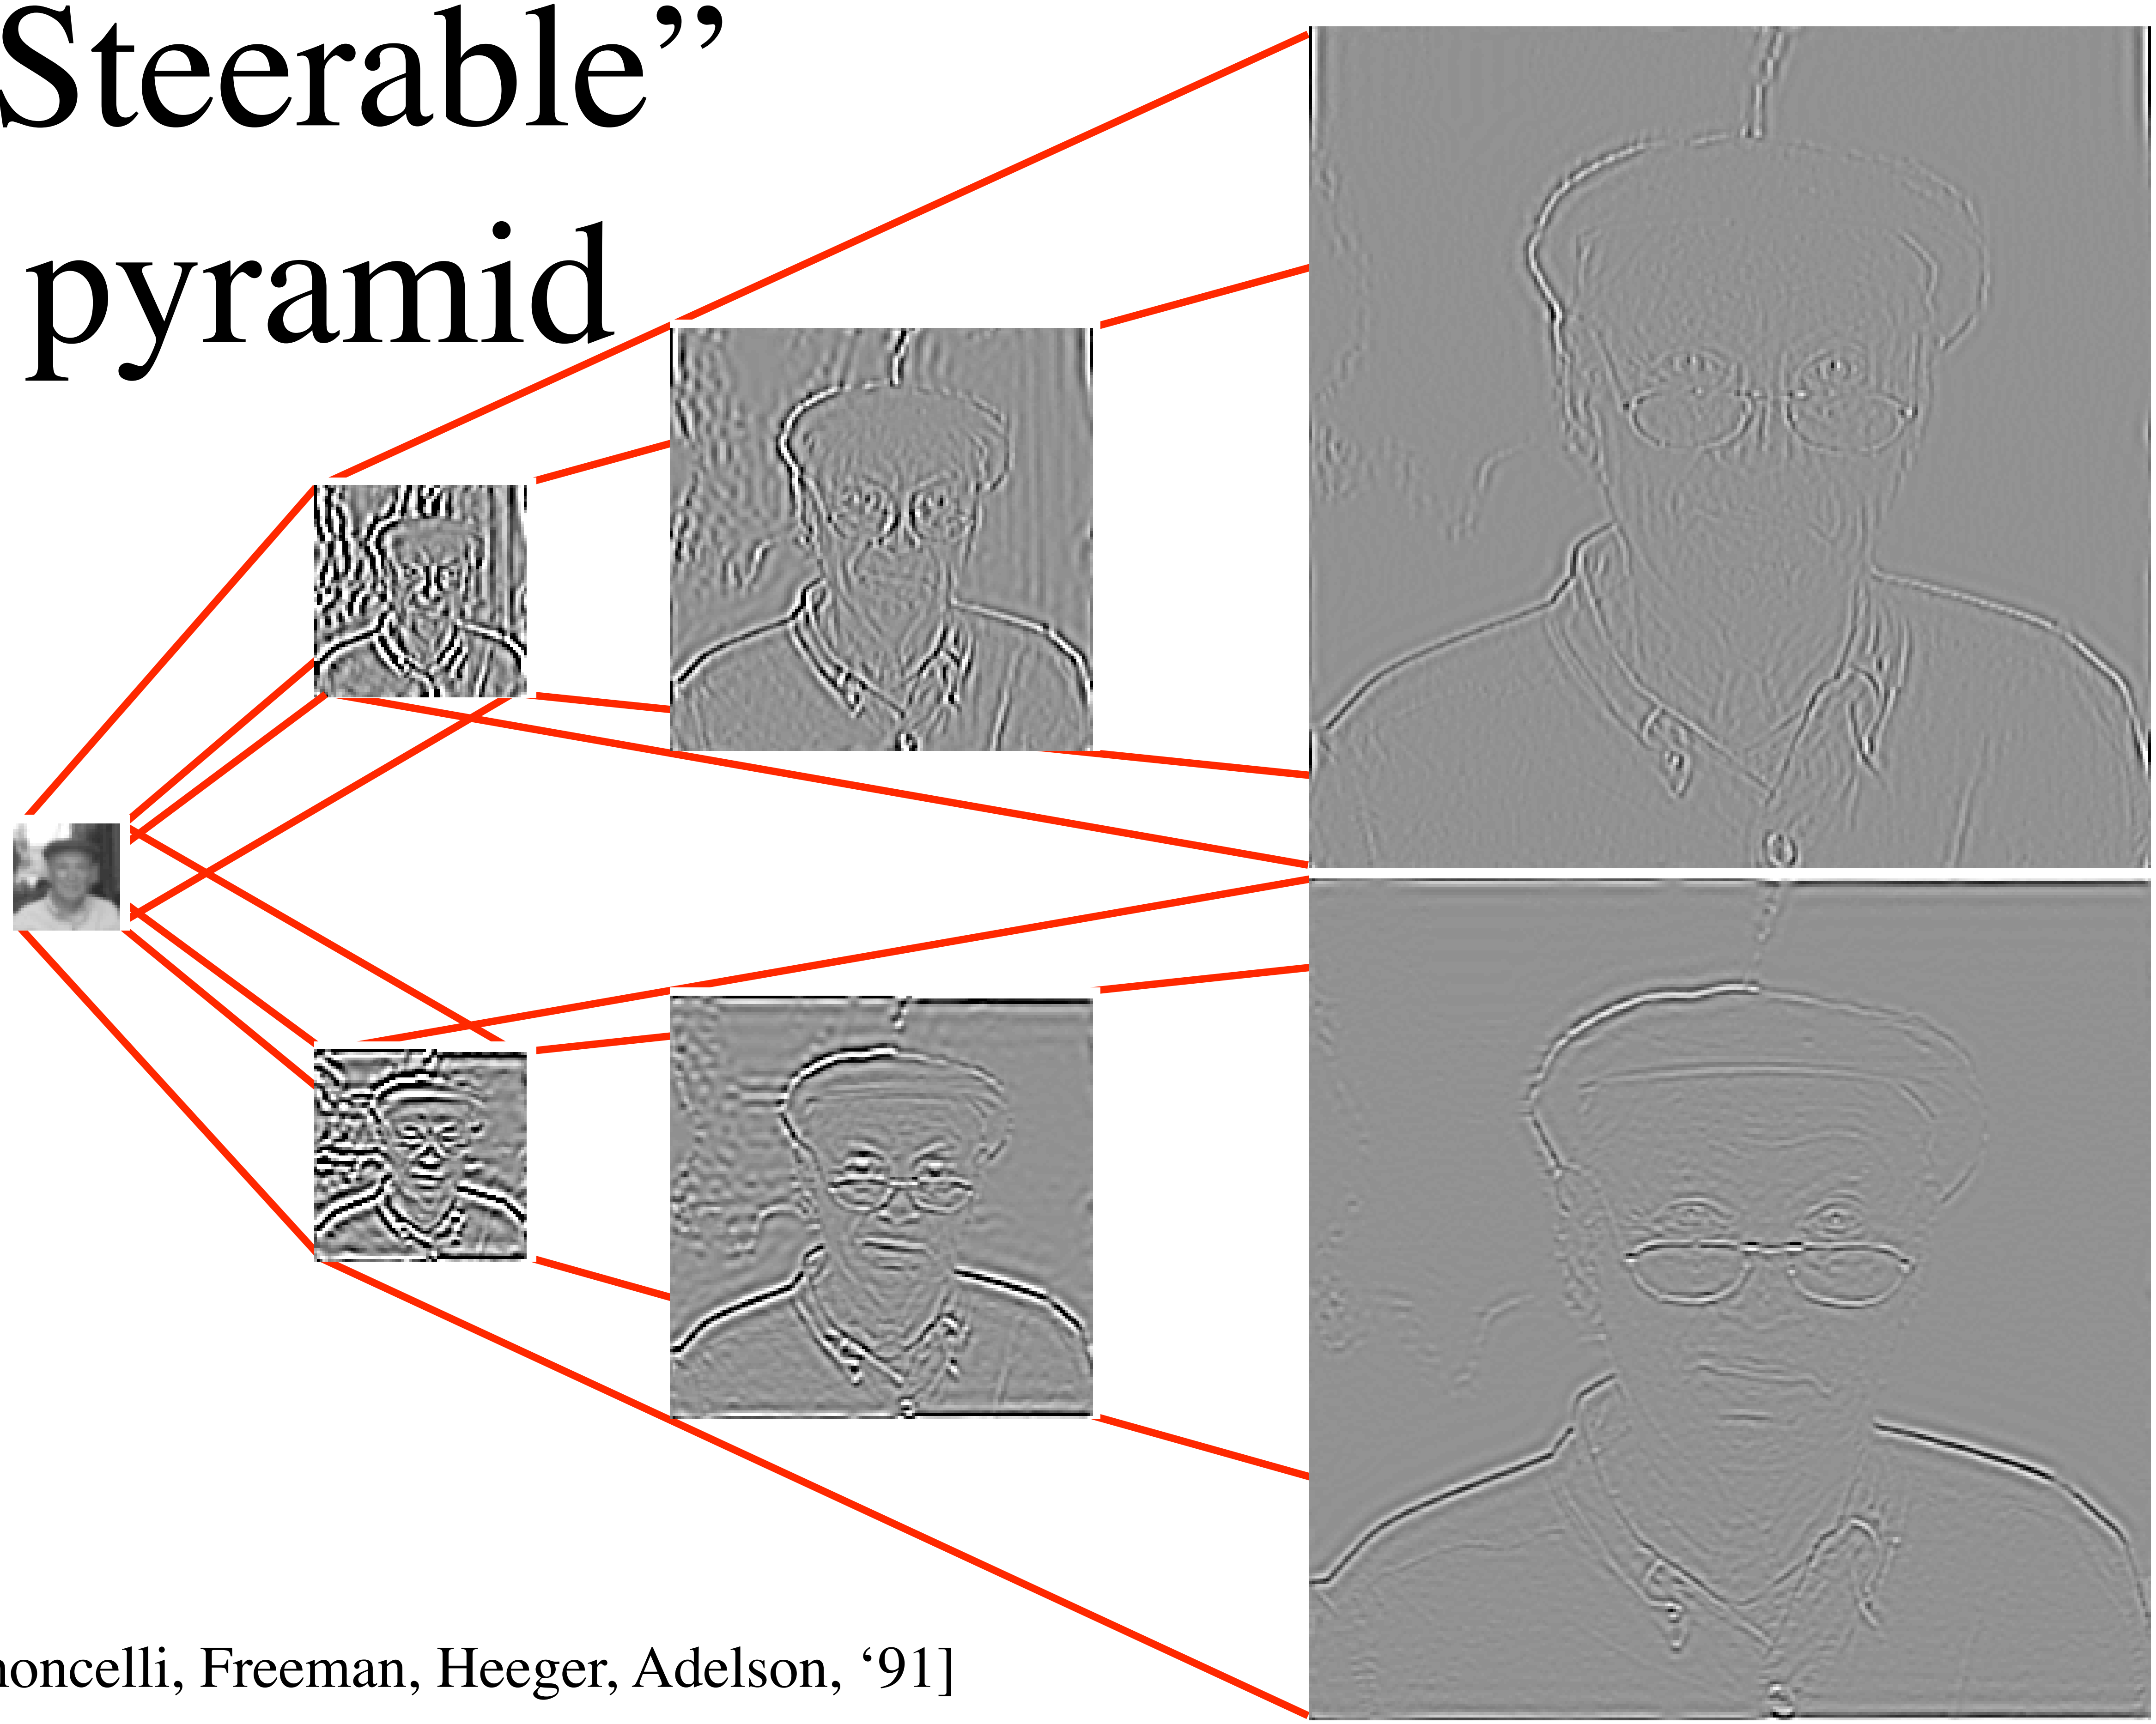

[Simoncelli, Freeman, Heeger, Adelson, ‘91]

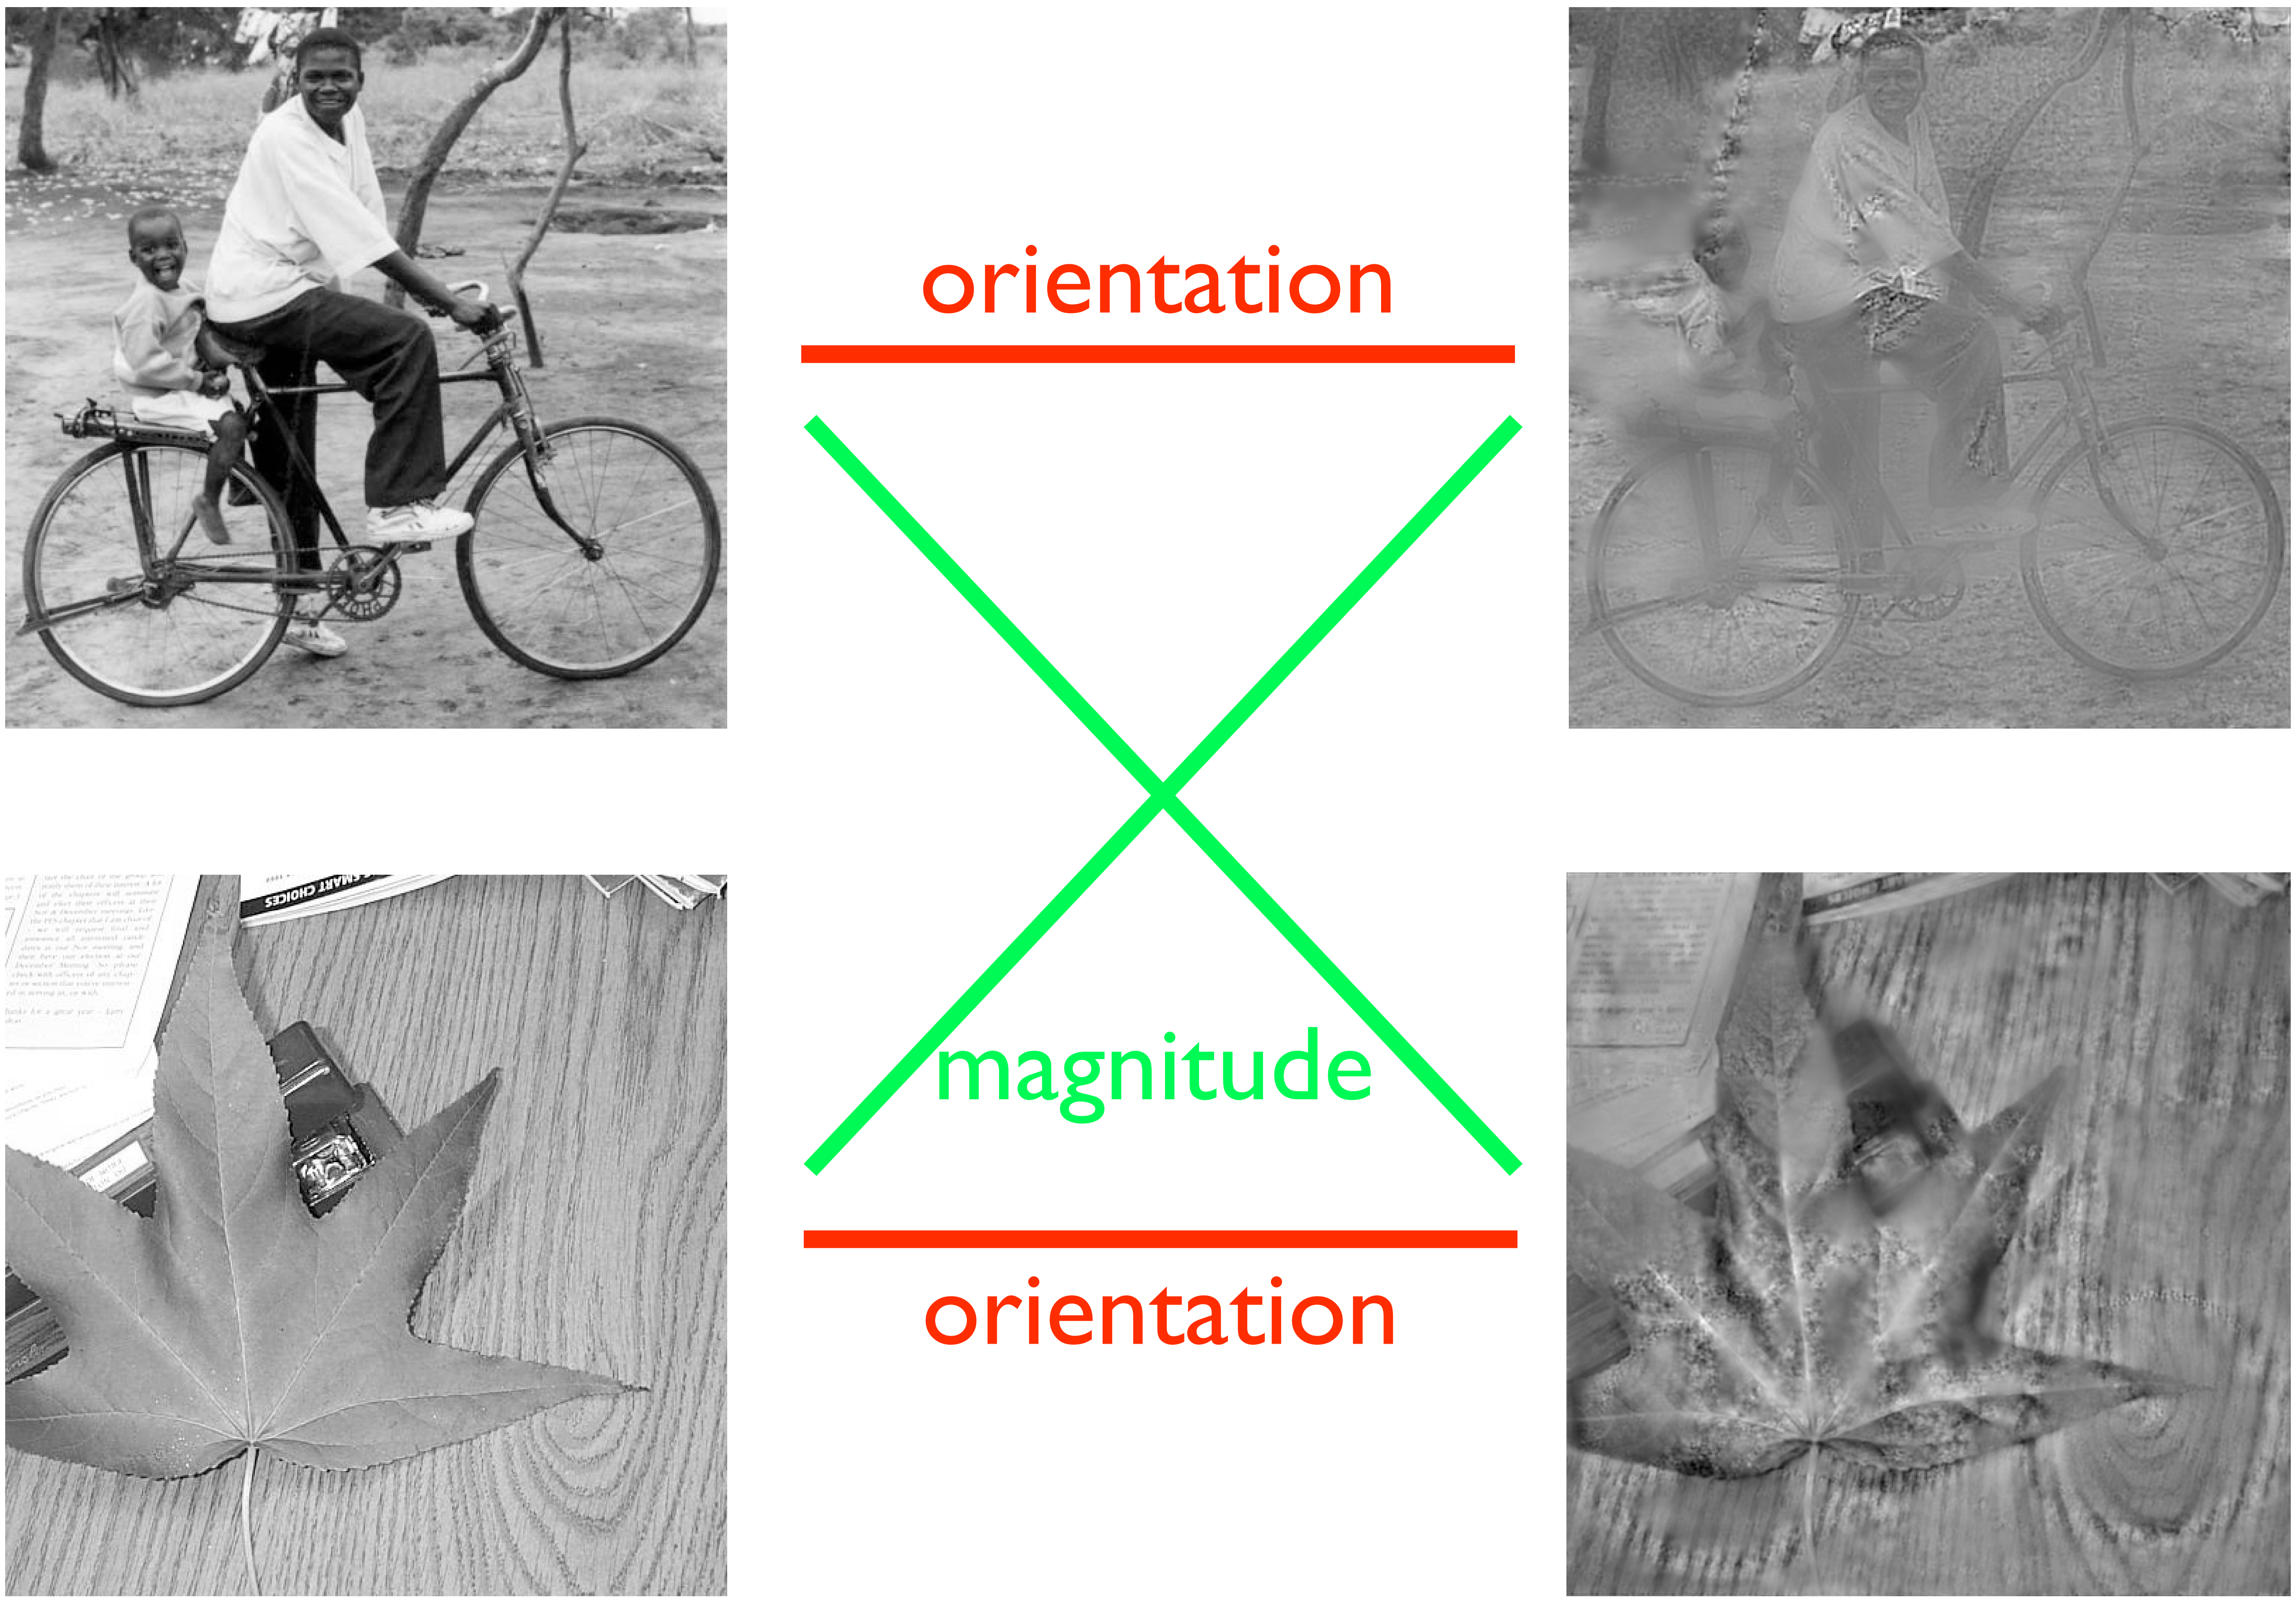

[Hammond & Simoncelli, 2005; cf. Oppenheim & Lim 1981]

Importance of local orientation

Randomized orientation

Randomized magnitude

Two-band, 6-level steerable pyramid

[with David Hammond]

Reconstruction from orientation

Original

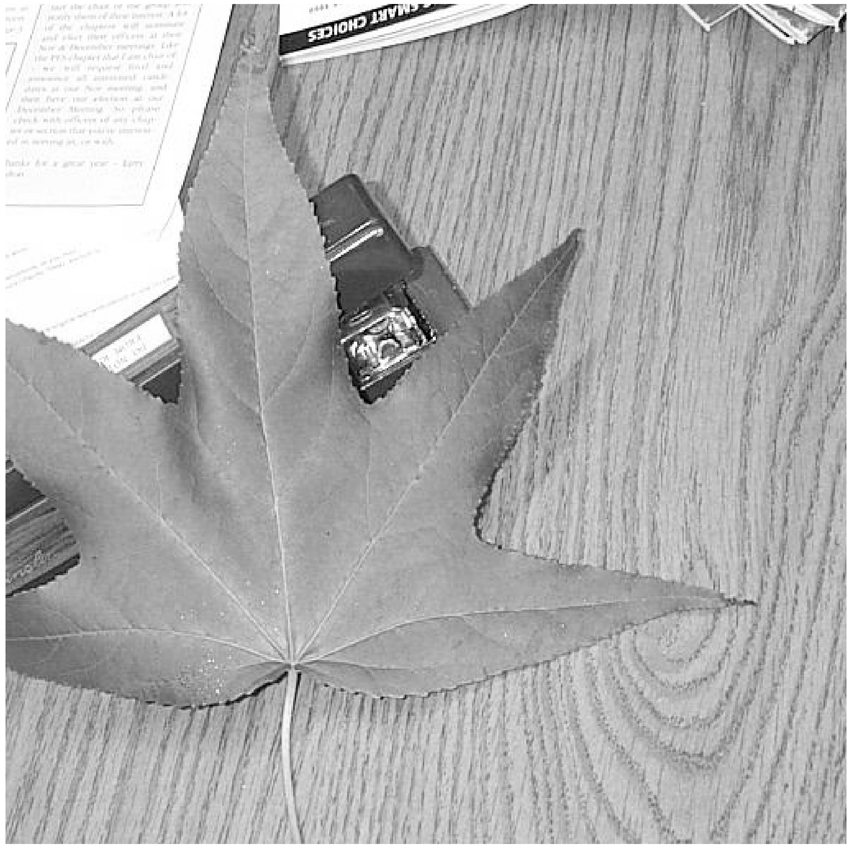

Quantized to 2 bits

- Alternating projections onto convex sets
- Resilient to quantization
- Highly redundant, across both spatial position and scale

[with David Hammond]

Spatial redundancy

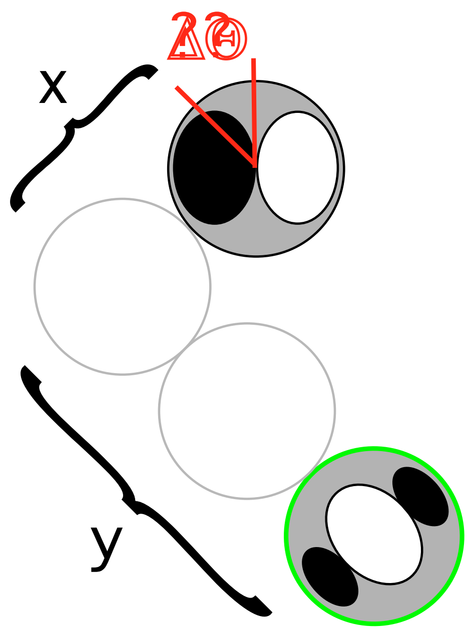

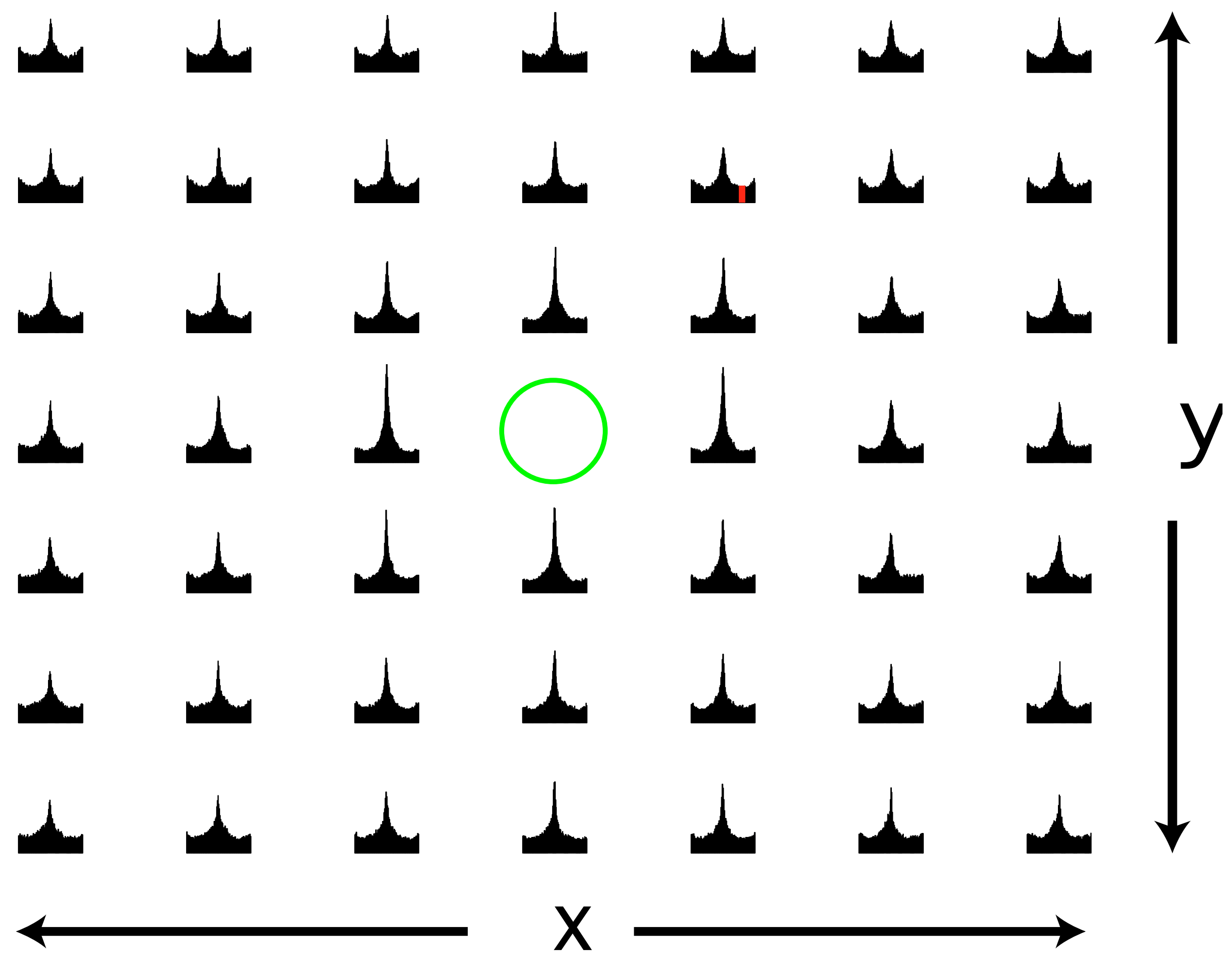

- Relative orientation histograms, at different locations
- See also: Geisler, Elder

[with Patrik Hoyer & Shani Offen]

Scale redundancy

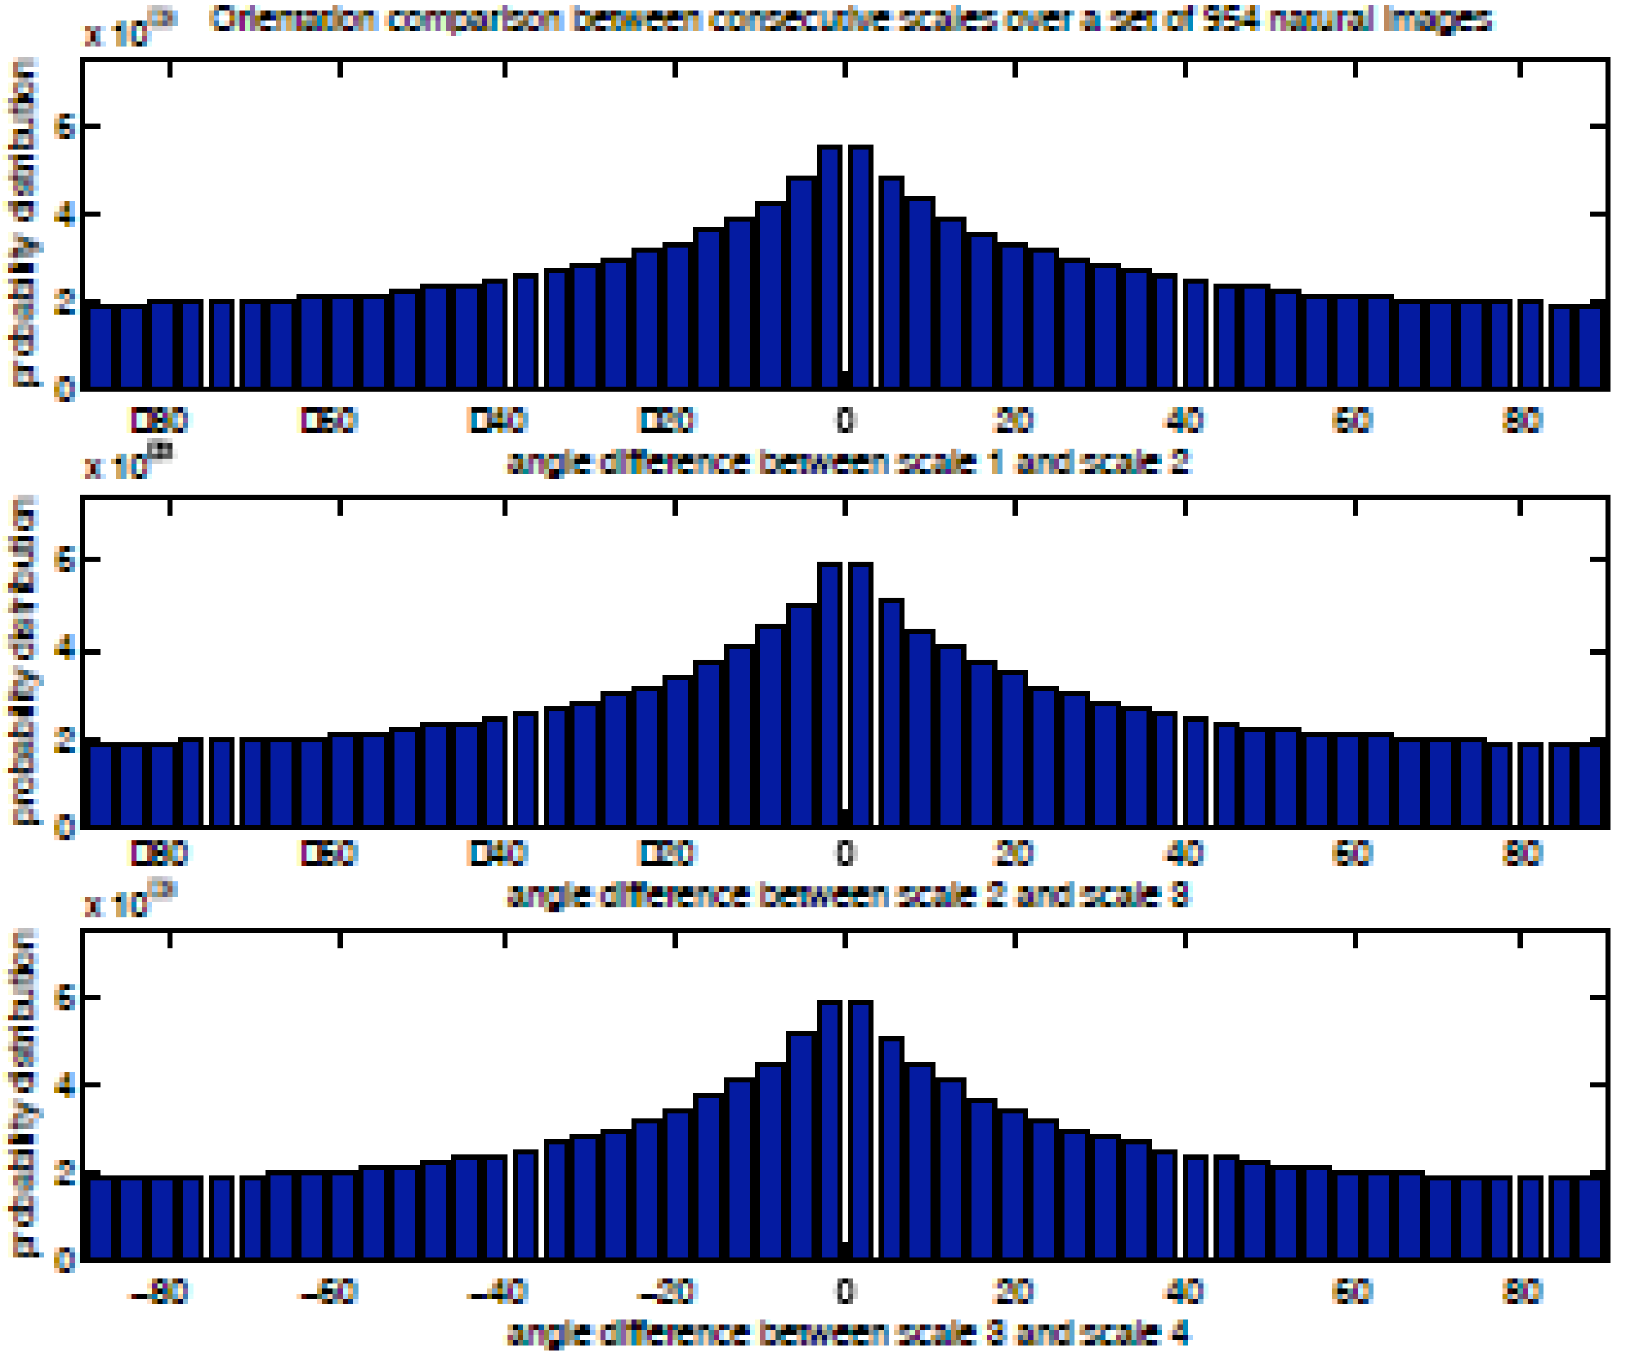

[with Clementine Marcovici]

Conclusions

- Multiresolution pyramids changed the world of image processing
- Statistical modeling can provide refinement and optimization of intuitive solutions:
  - Wiener
  - Coring
  - Locally adaptive variances
  - Locally adaptive orientation

Cast

- Local GSM model: Martin Wainwright, Javier Portilla
- Denoising: Javier Portilla, Martin Wainwright, Vasily Strela, Martin Raphan
- GSM tree model: Martin Wainwright, Alan Willsky
- Local orientation: David Hammond, Patrik Hoyer, Clementine Marcovici
- Local phase: Zhou Wang
- Texture representation/synthesis: Javier Portilla
- Compression: Robert Buccigrossi