Multi-scale decompositions

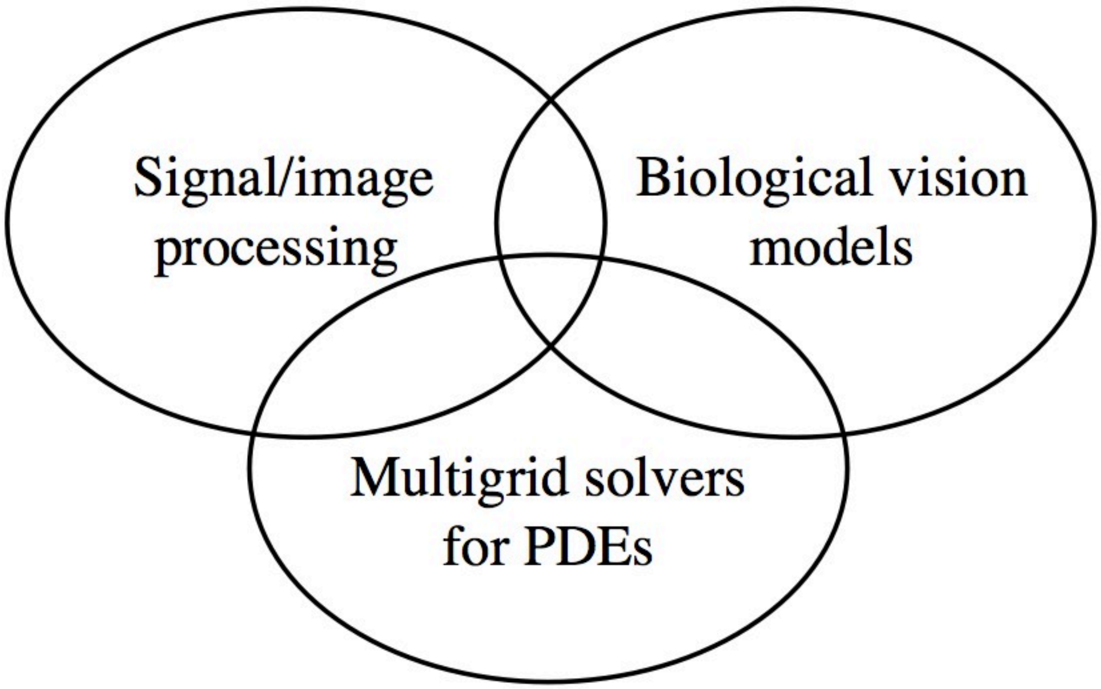

The "Wavelet revolution"

- 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 (as of 2009)
- Best results in most signal/image processing applications


image

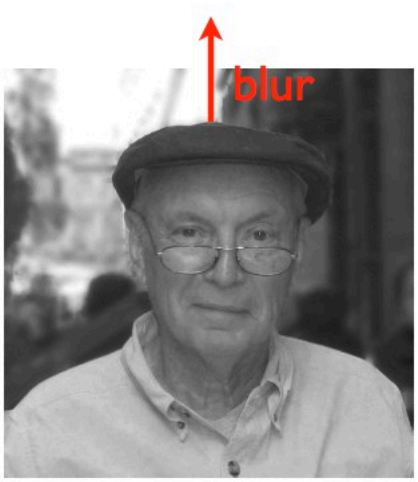

image

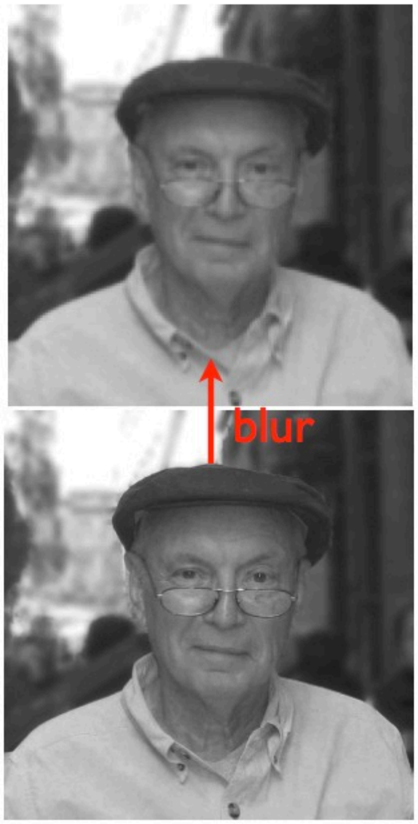

image

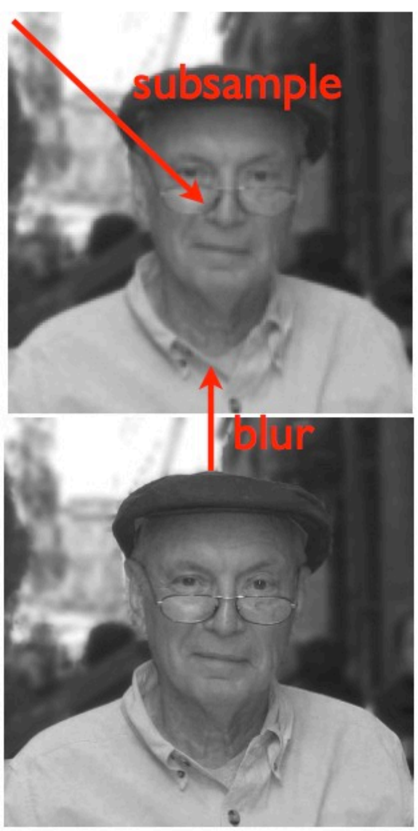

image

image

image

image

image

image

image

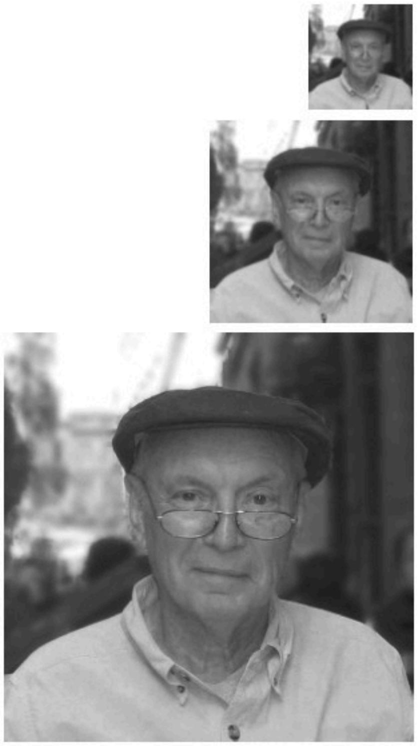

image

image

Gaussian pyramid

Why do this?

- resize image (e.g., for display)
- low-dimensional summary of data
- analyze content at different scales/ resolutions
- efficient search/matching
- efficient representation (compression)

Gaussian pyramid

"Pyramid"
data structure

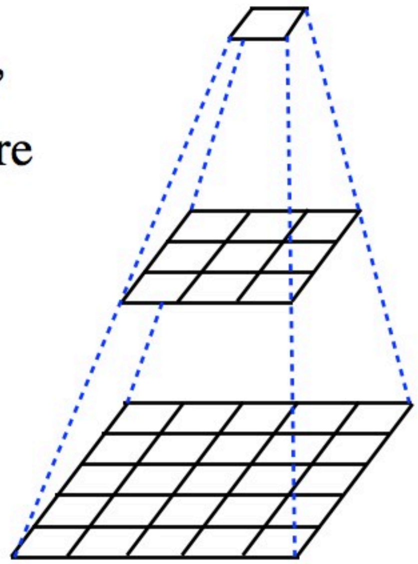

Primary operation, a filter/subsample cascade:

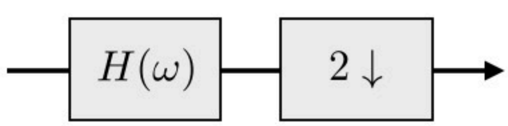

output (only keep even samples)

input

Note: should be implemented as a single operation, to avoid wasted computation!

scale space

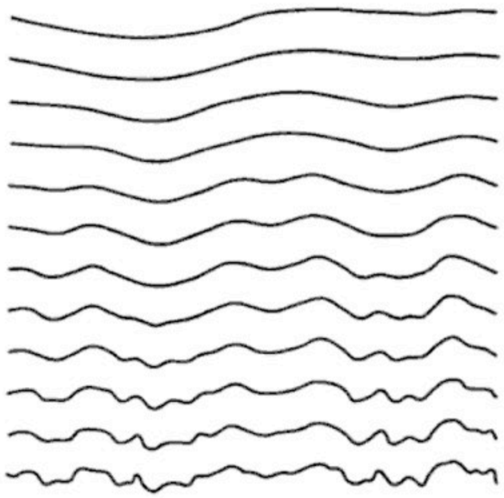

Figure 1.3. The main idea with a scale-space representation of a signal is to generate a one-parameter family of derived signals in which the fine-scale information is successively suppressed. This figure shows a signal that has been successively smoothed by convolution Gaussian kernels of increasing width. (Adapted from Witkin 1983).

[Witkin '83; Koenderink '84; fig from Lindeberg '93]

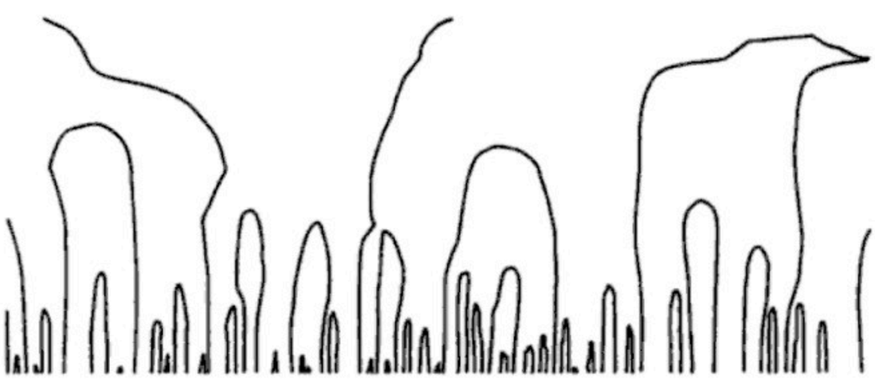

Figure 1.5. Since new zero-crossings cannot be created by the diffusion equation in the one-dimensional case, the trajectories of zero-crossings in scale-space (here, zero-crossings of the second derivative) form paths across scales that are never closed from below. (Adapted from Witkin 1983).

[Lindeberg, 93]

Coarse-to-fine Flow Estimation

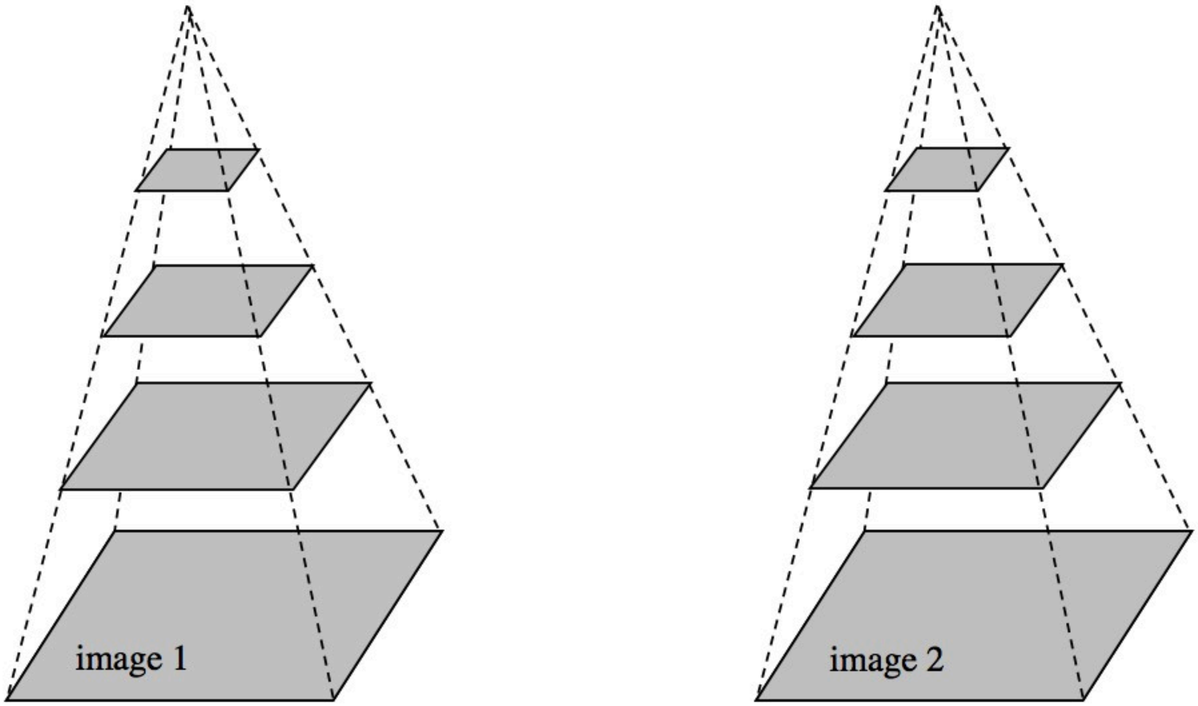

Coarse-to-fine Flow Estimation

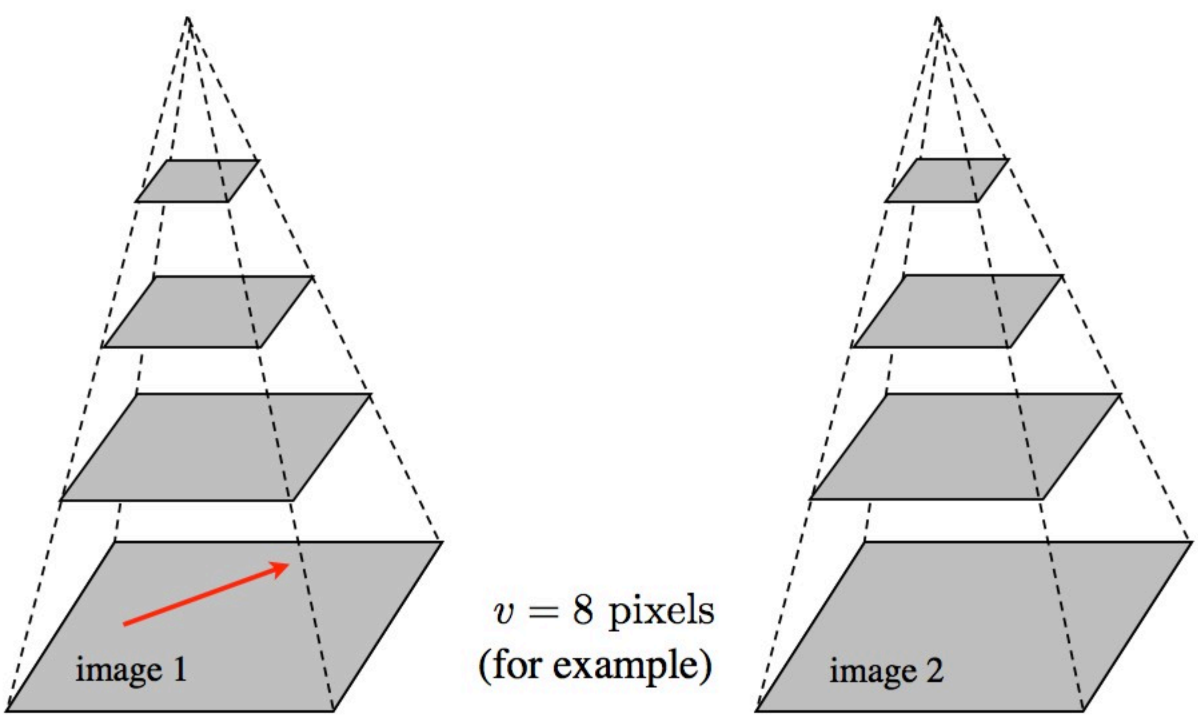

Coarse-to-fine Flow Estimation

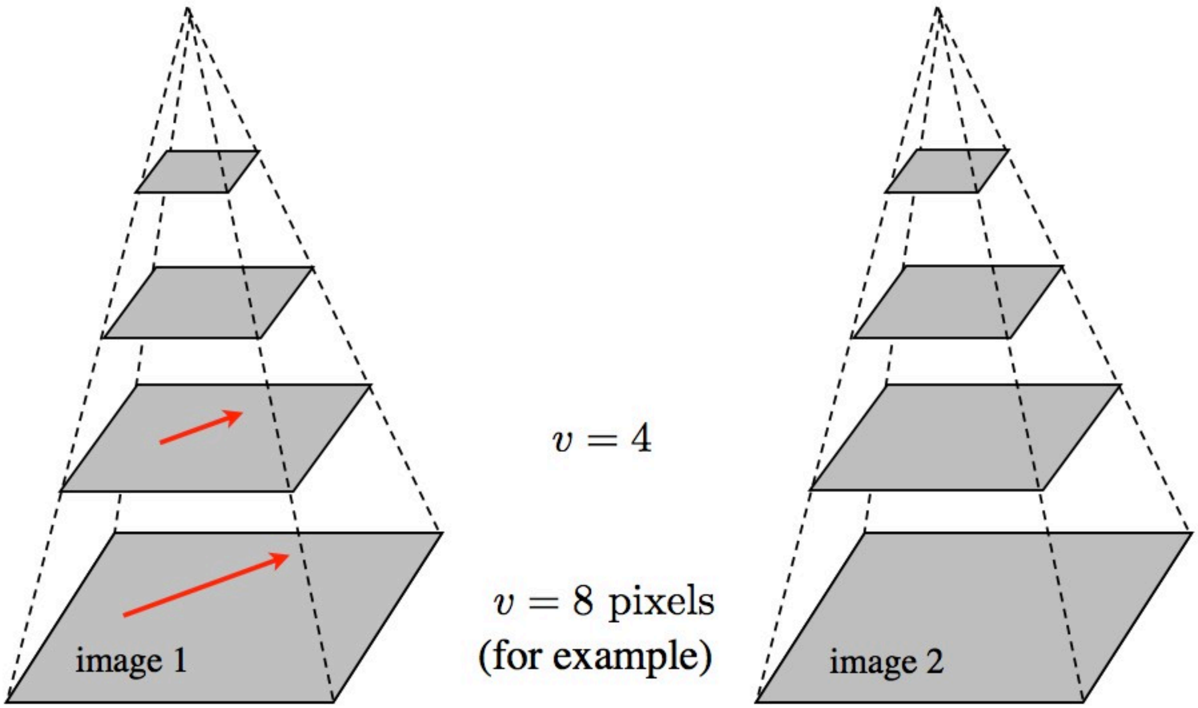

Coarse-to-fine Flow Estimation

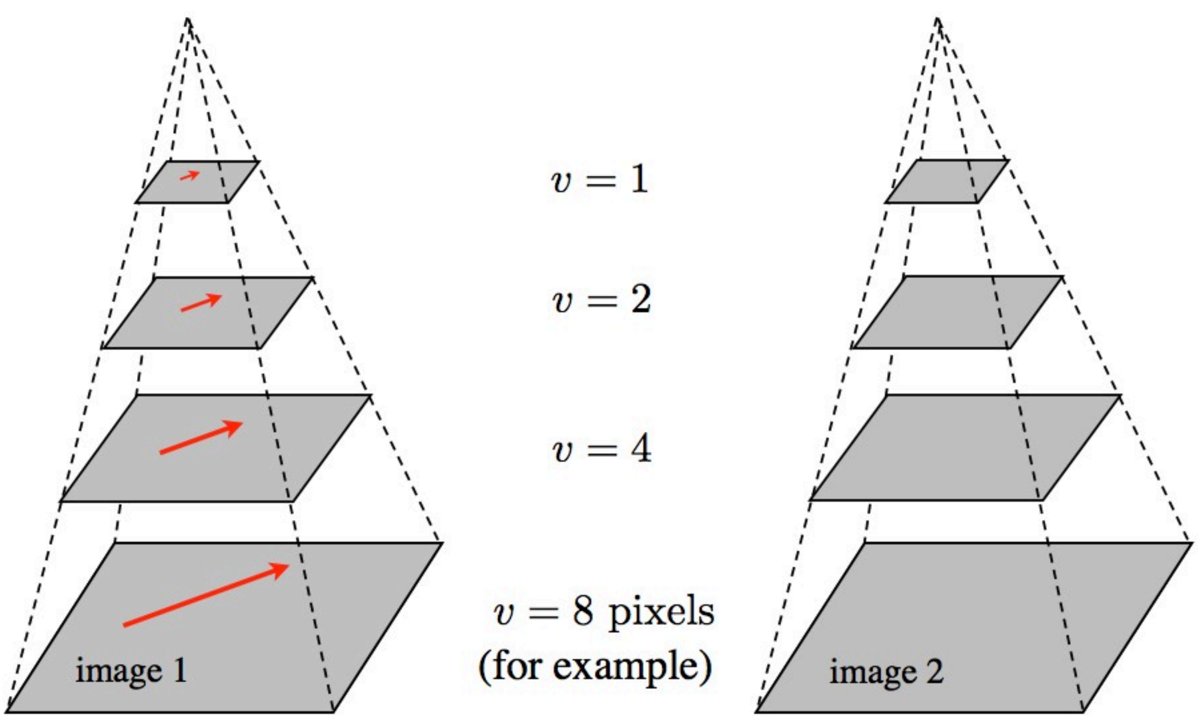

Coarse-to-fine Flow Estimation

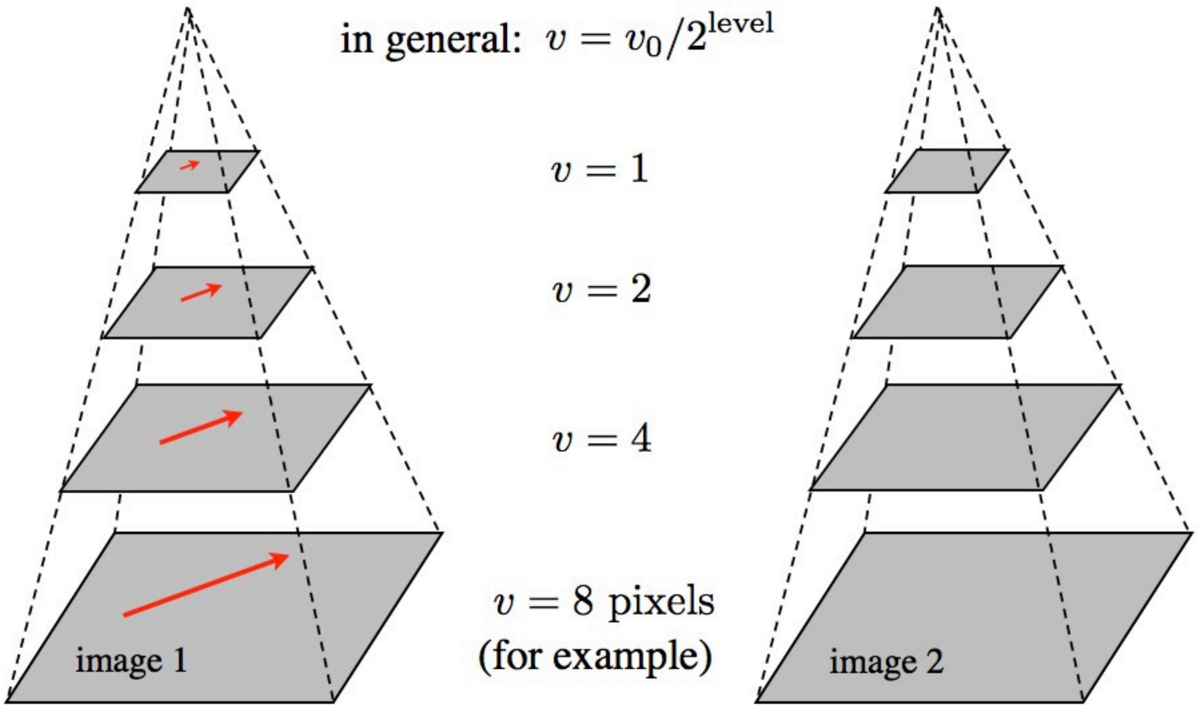

Coarse-to-fine displacement estimation (e.g., optic flow)

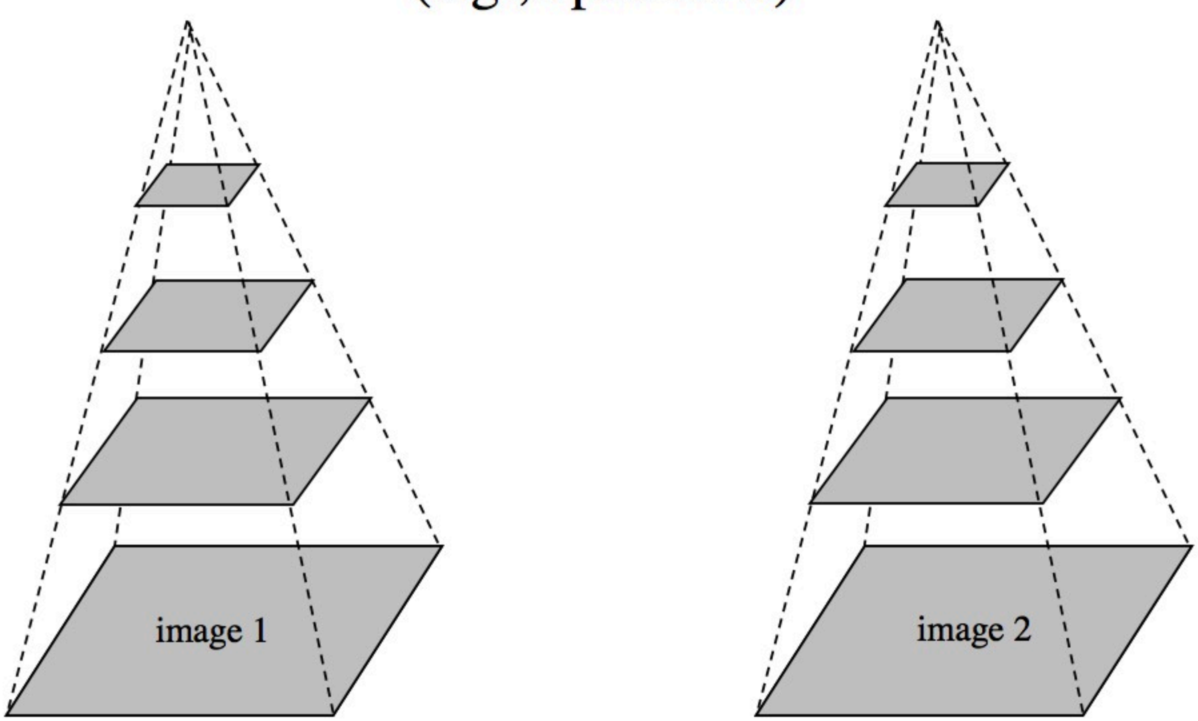

Coarse-to-fine displacement estimation
(e.g., optic flow)

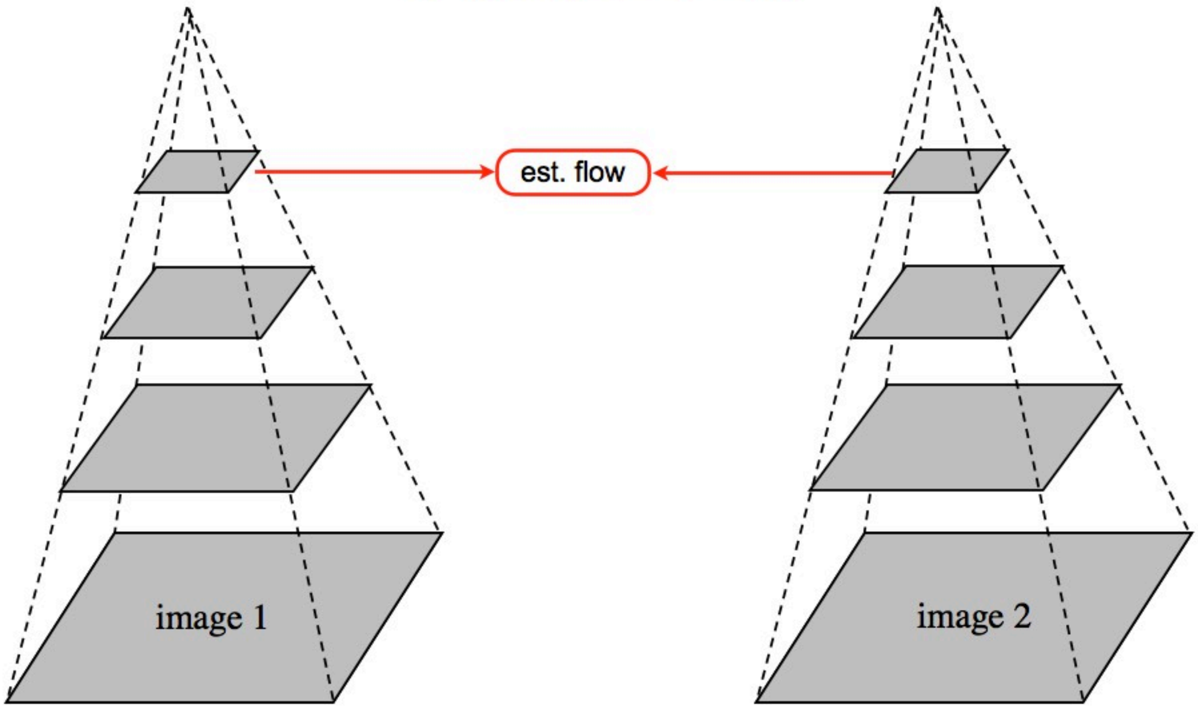

Coarse-to-fine displacement estimation
(e.g., optic flow)

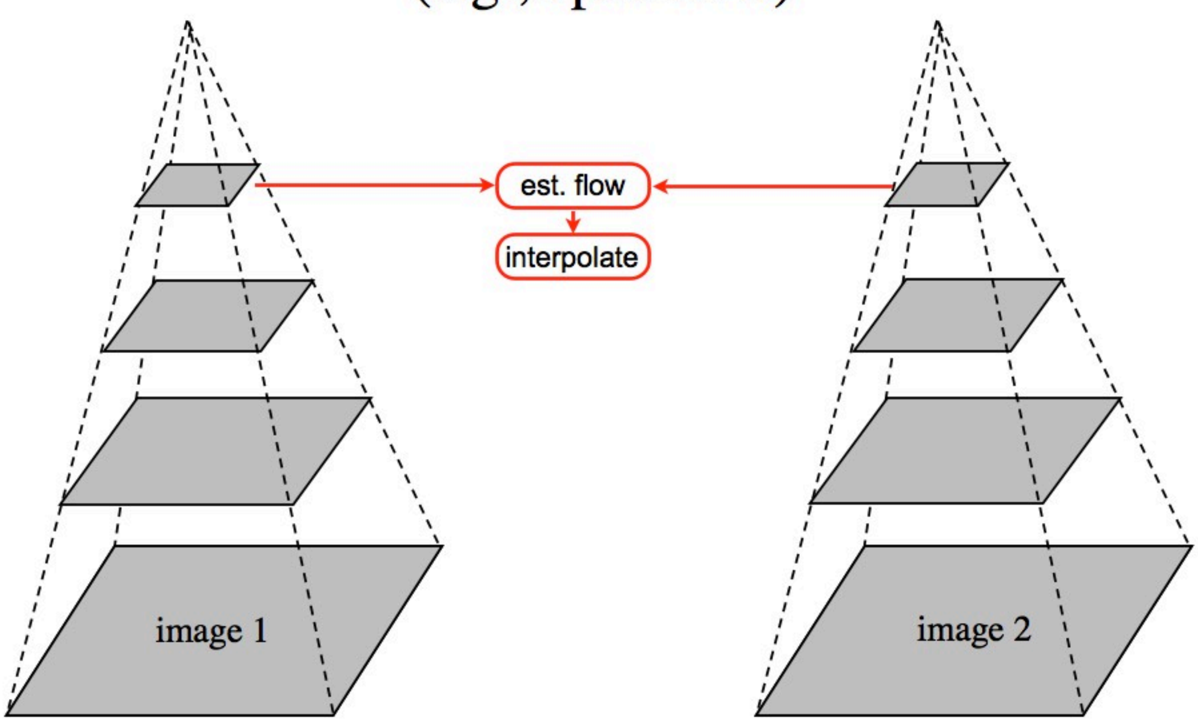

Coarse-to-fine displacement estimation
(e.g., optic flow)

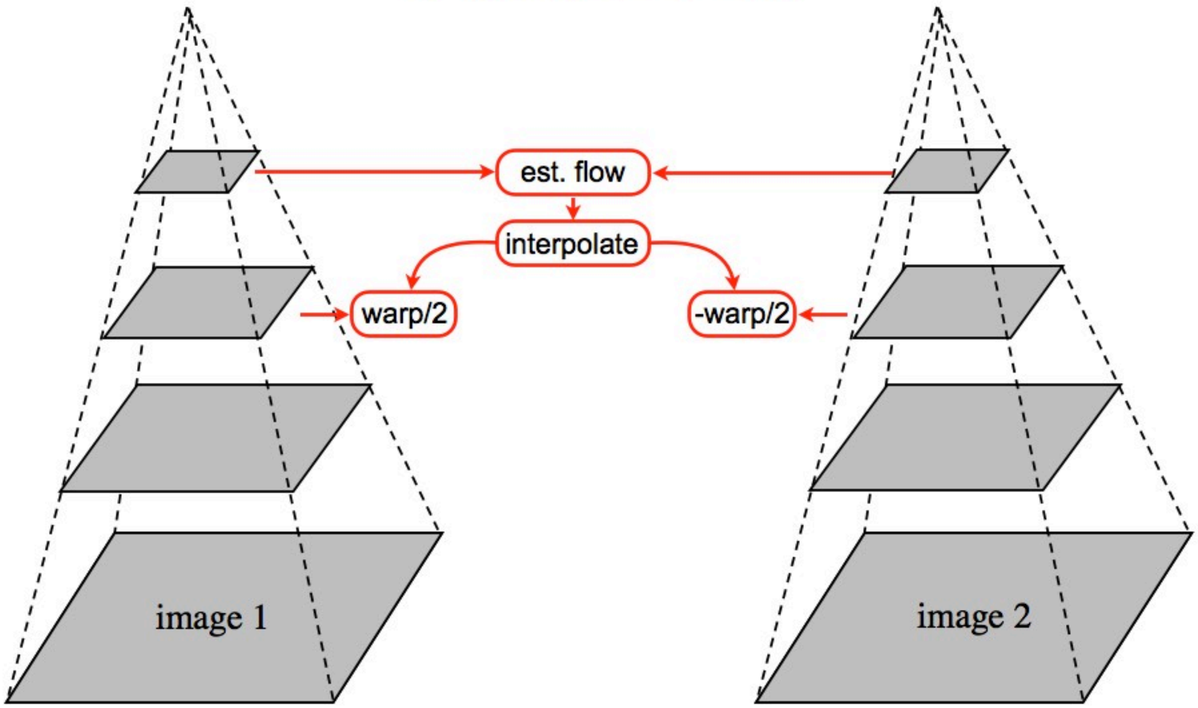

Coarse-to-fine displacement estimation
(e.g., optic flow)

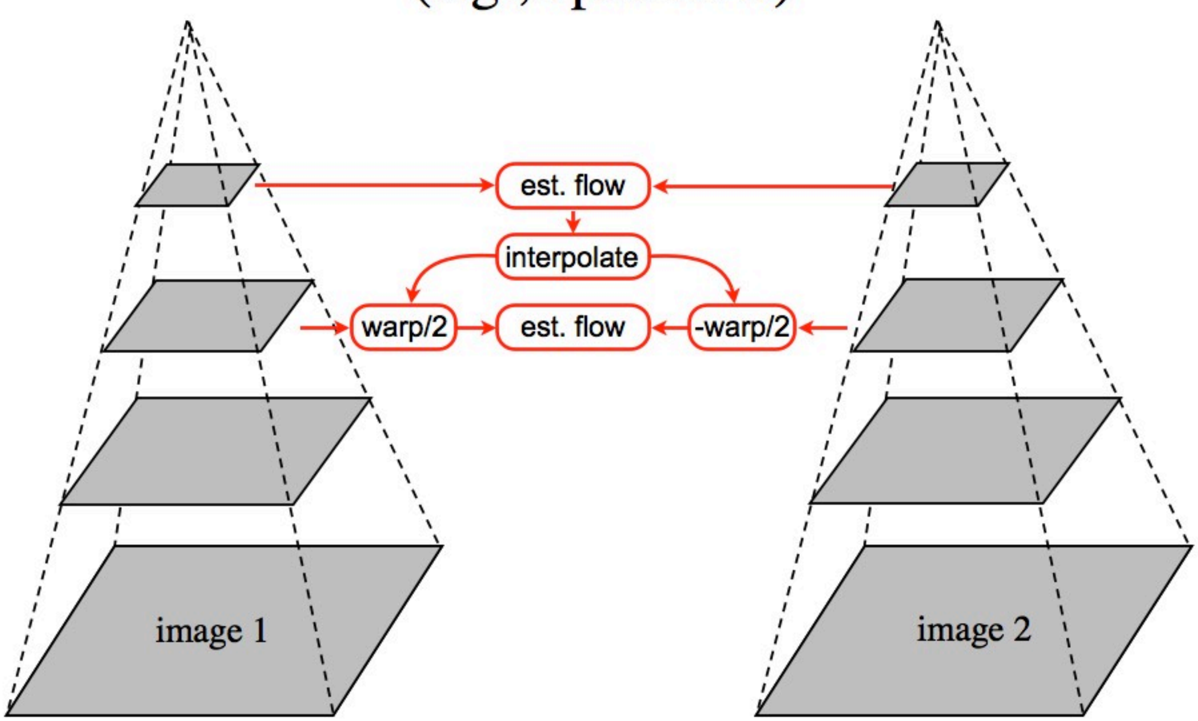

Coarse-to-fine displacement estimation
(e.g., optic flow)

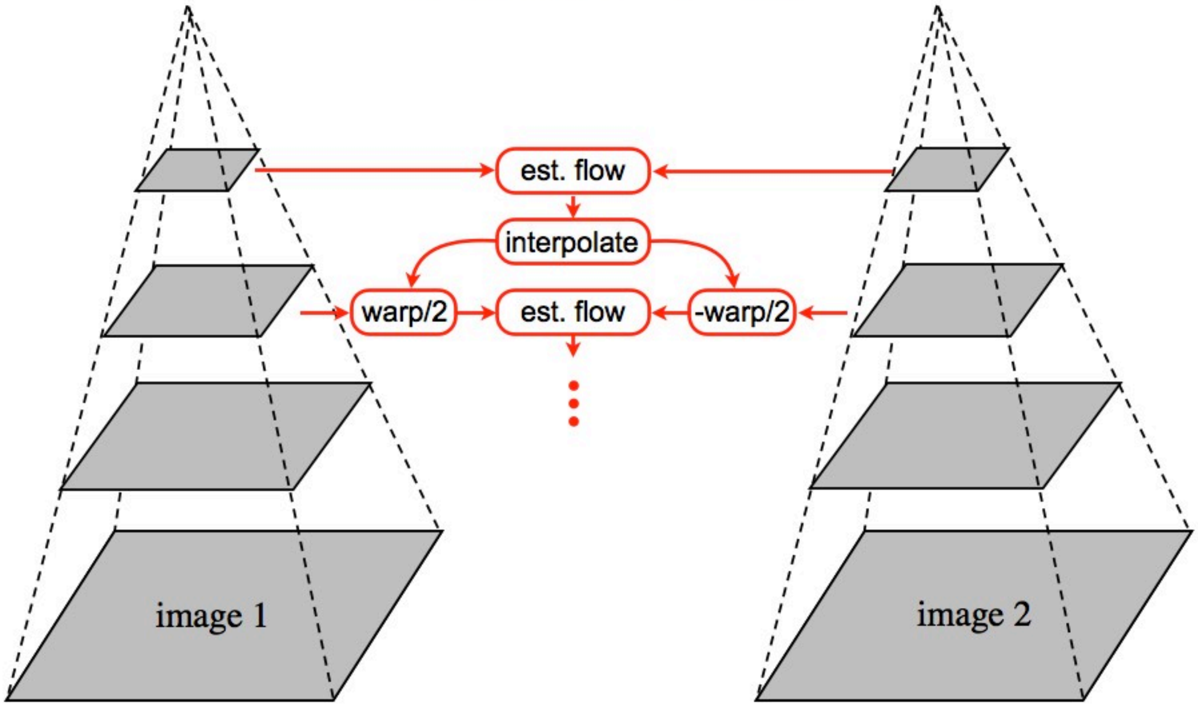

Coarse-to-fine displacement estimation
(e.g., optic flow)

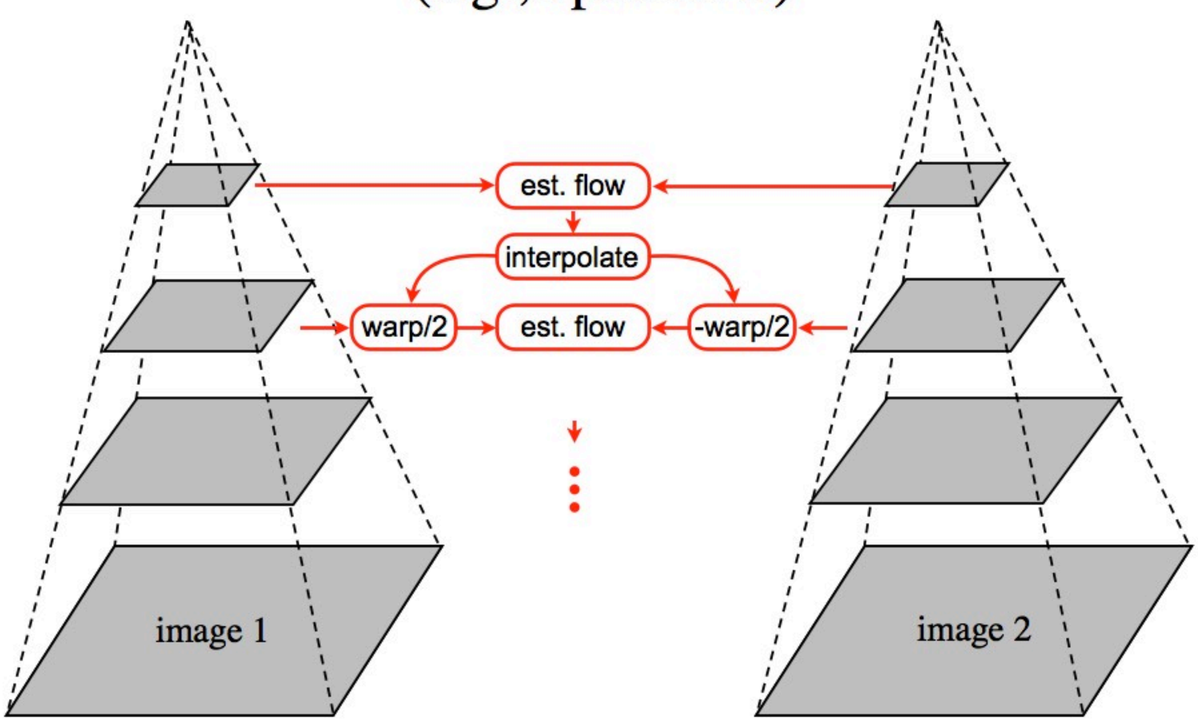

Coarse-to-fine displacement estimation (e.g., optic flow)

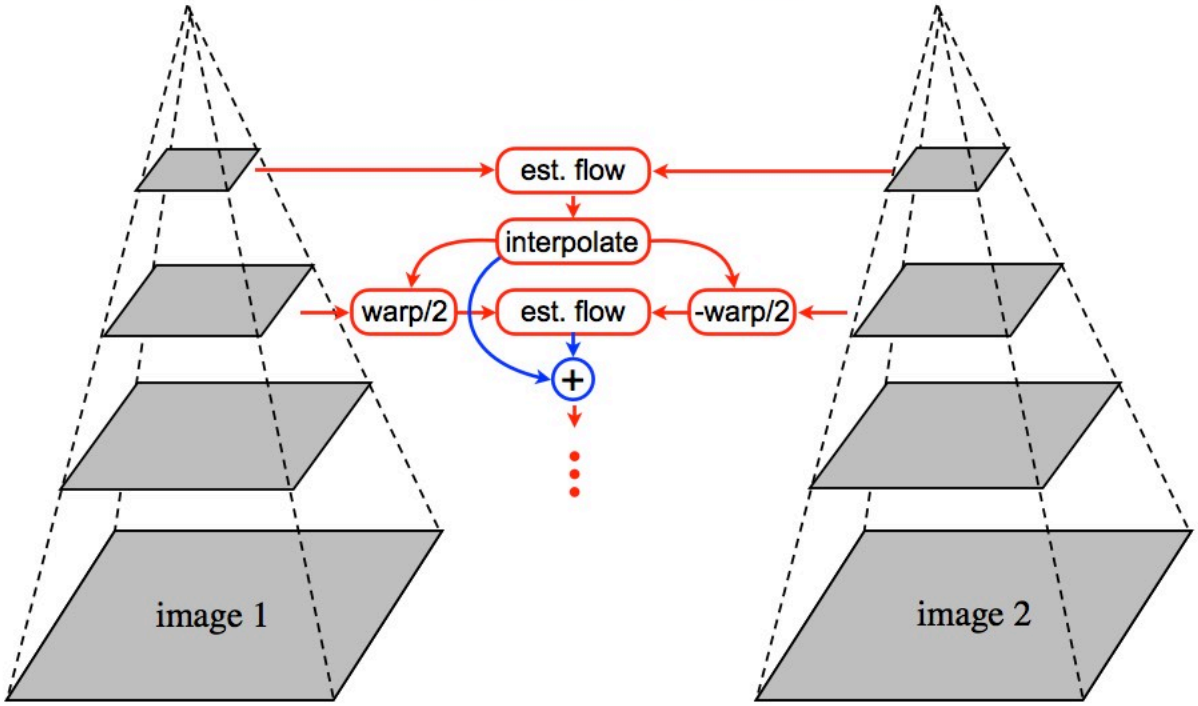

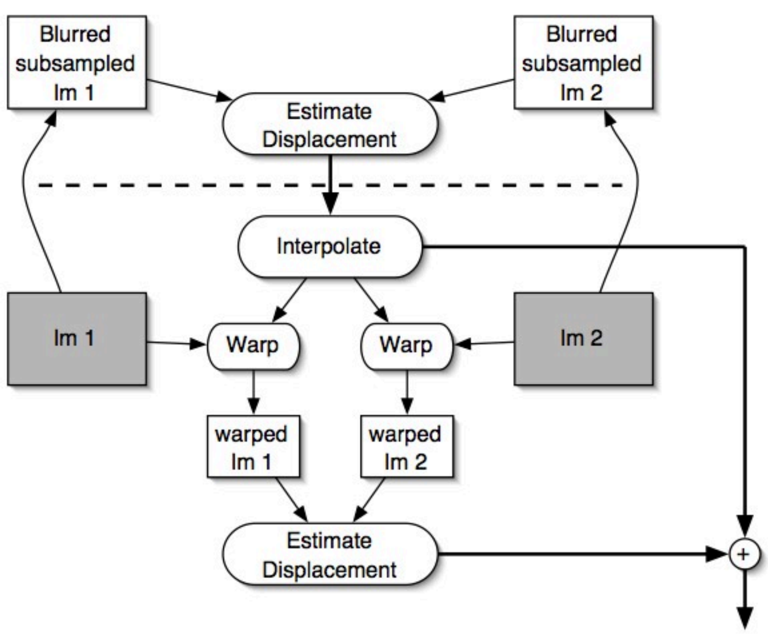

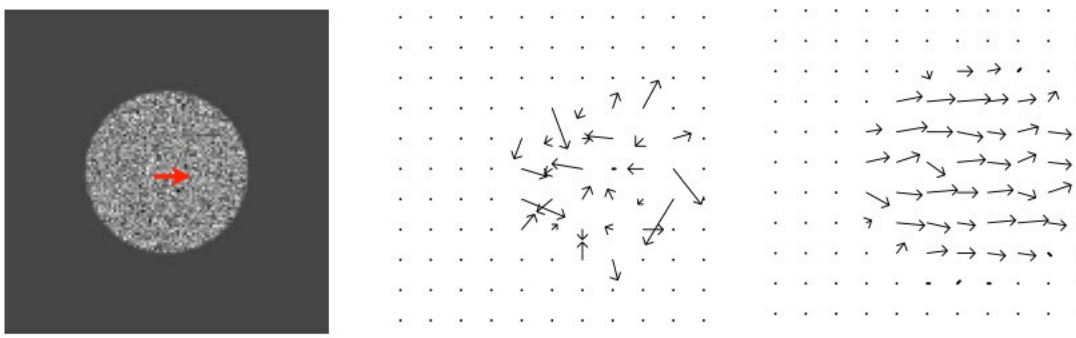

image — single scale — coarse-to-fine

Inverse operation: upsample/filter cascade:

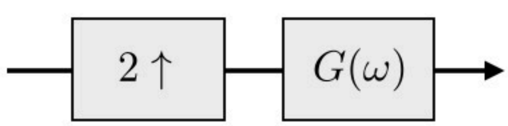

Note: should be implemented as a single operation, to avoid wasted computation!

Gaussian pyramid

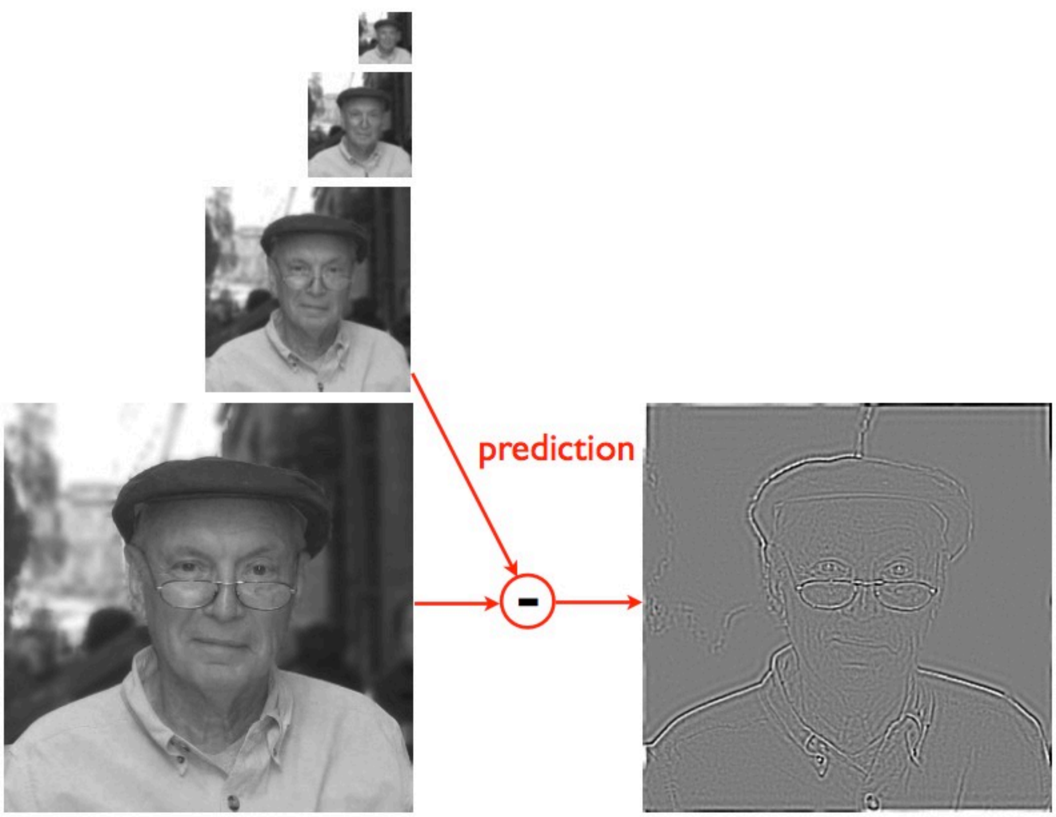

Gaussian pyramid

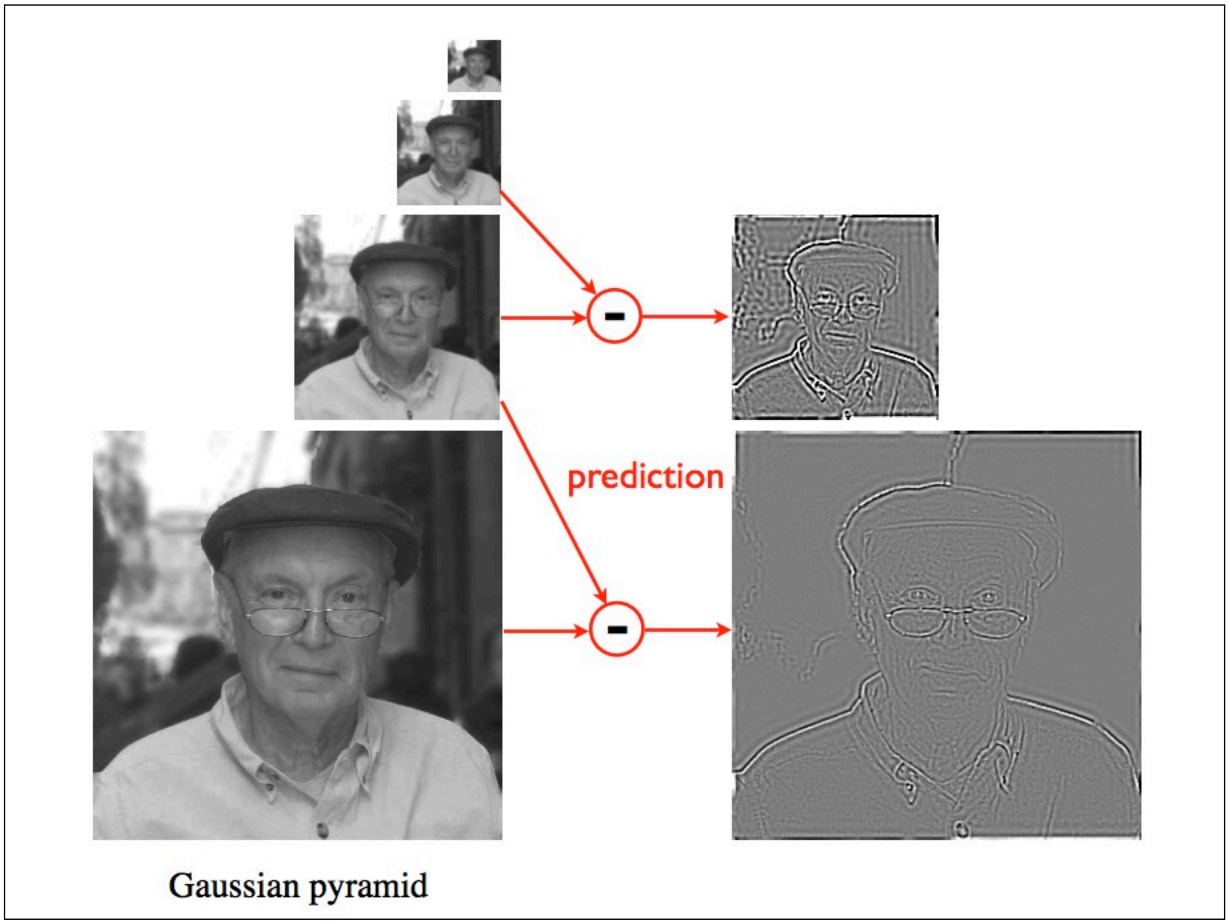
prediction
Gaussian pyramid

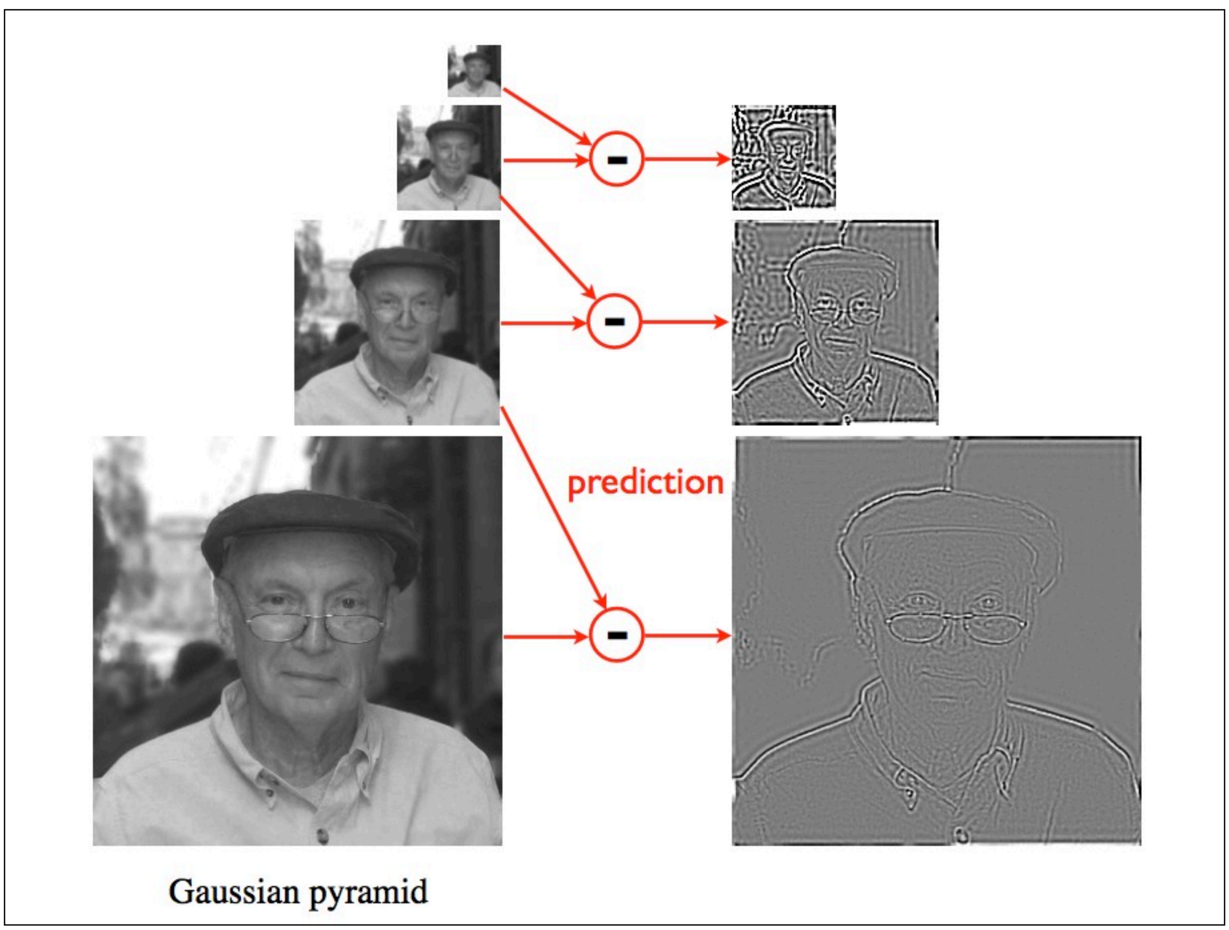
prediction
Gaussian pyramid

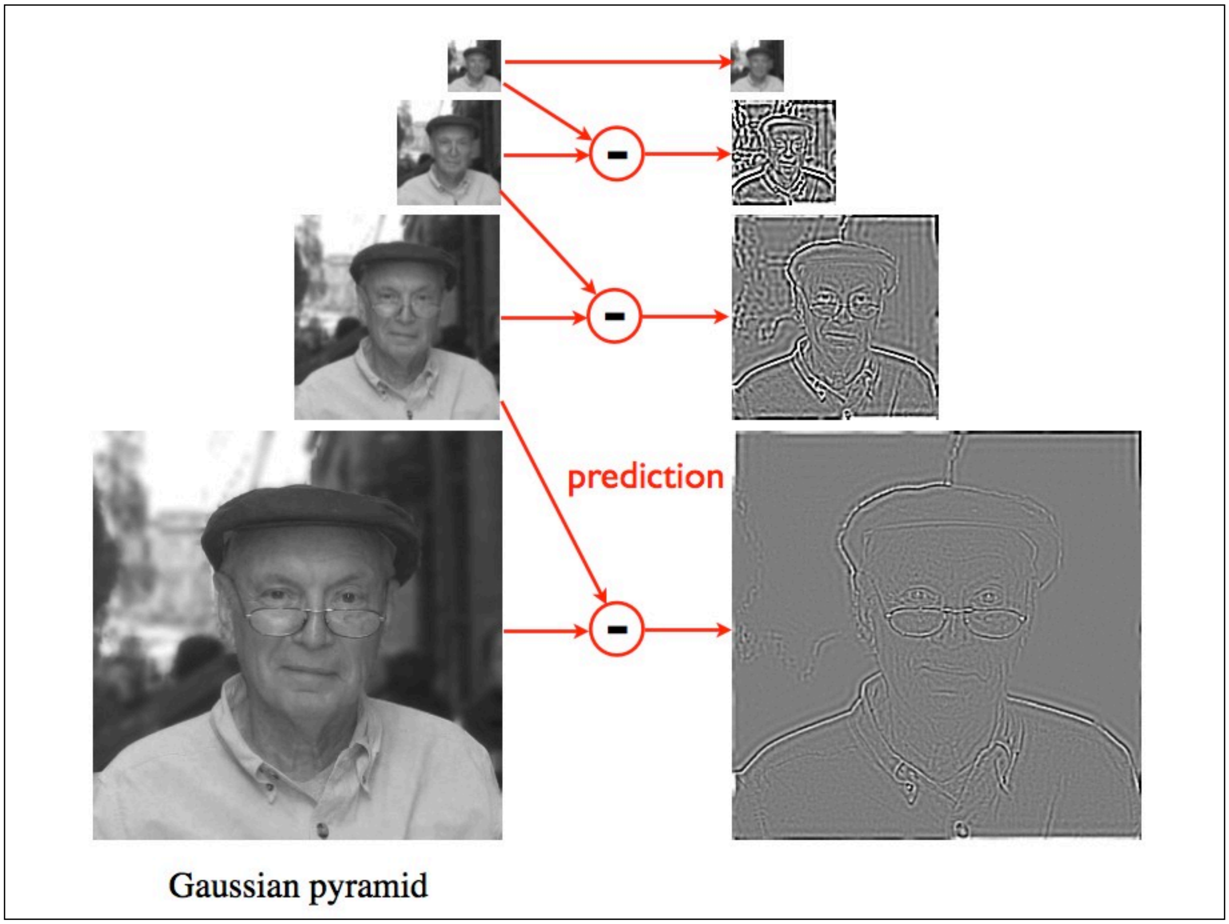
prediction
Gaussian pyramid

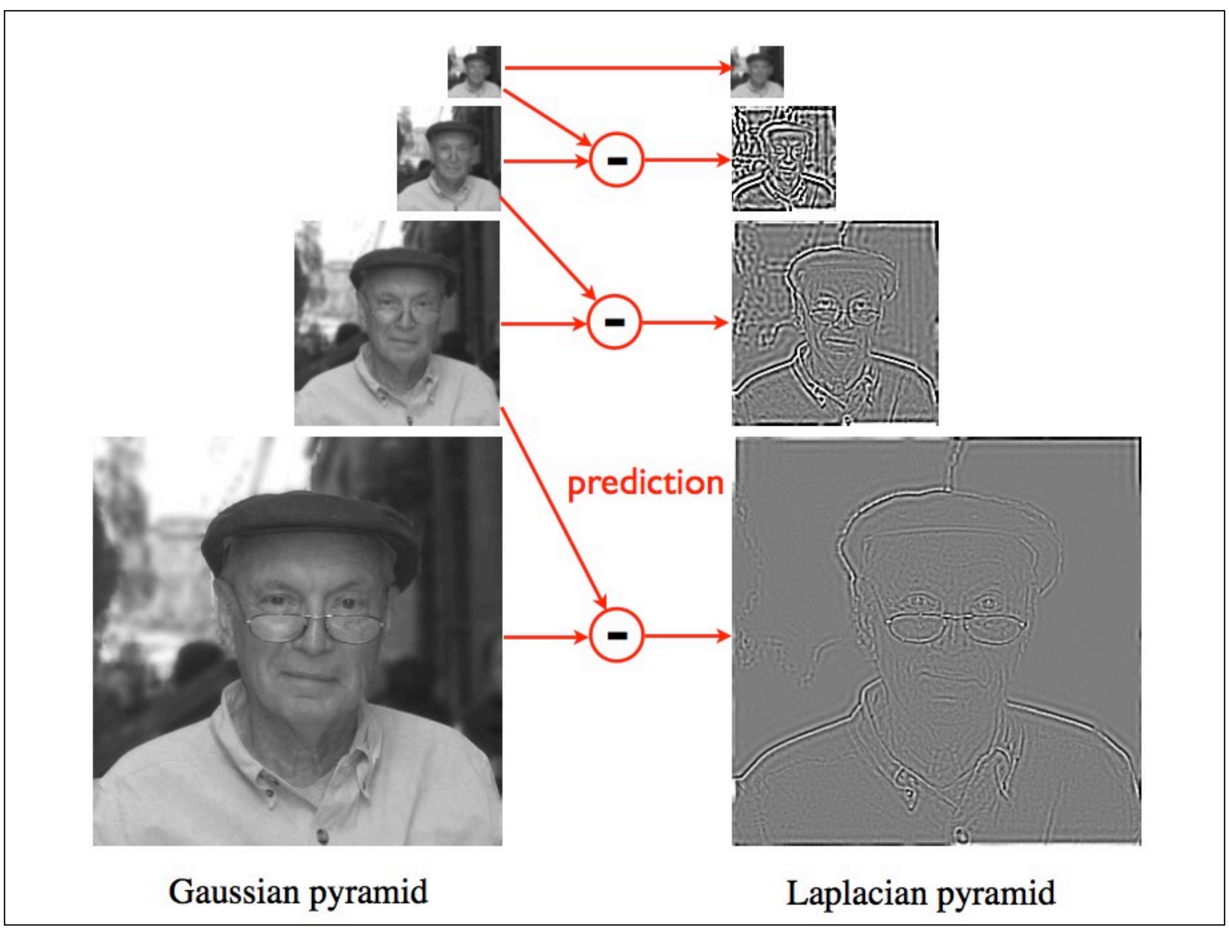
prediction
Gaussian pyramid
Laplacian pyramid

Laplacian pyramid

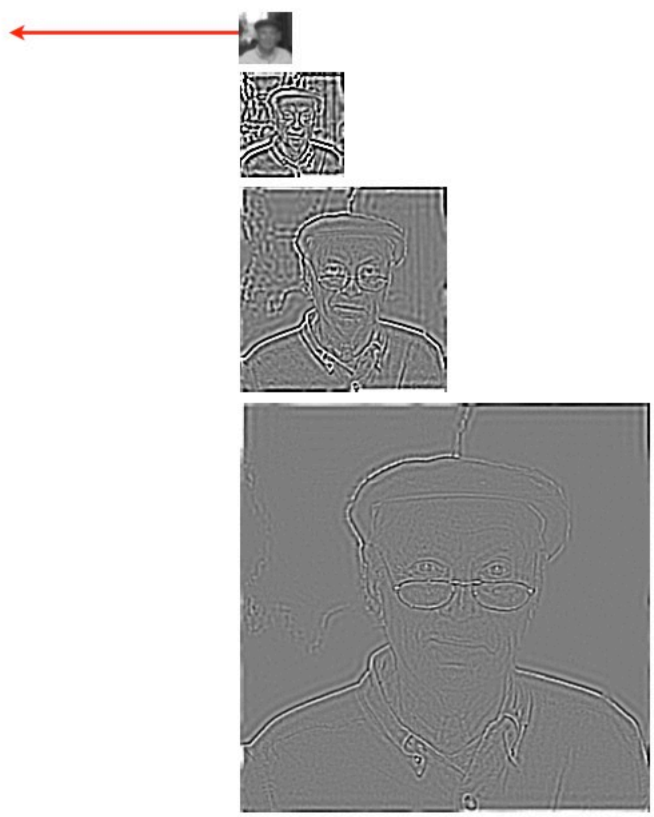

Laplacian pyramid

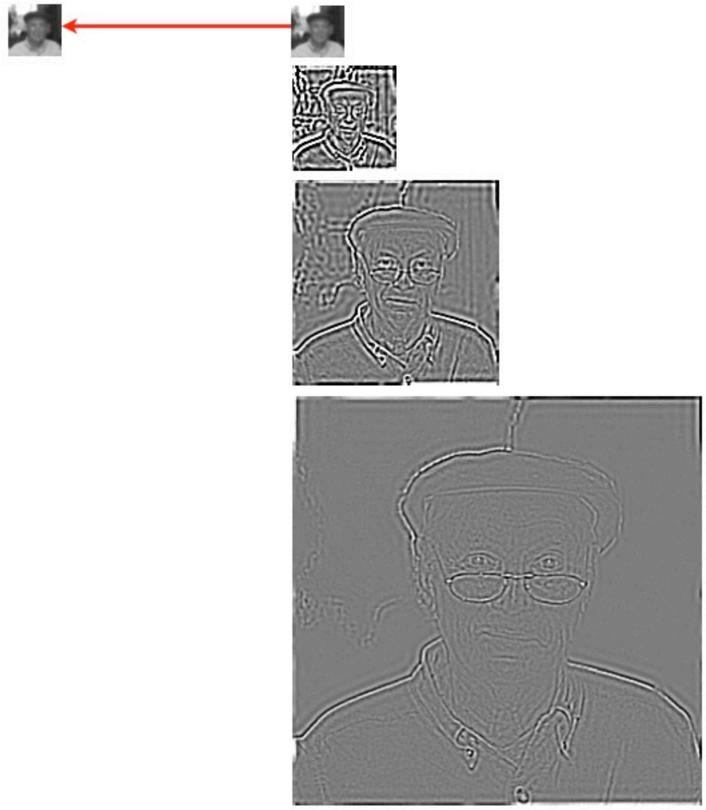

Laplacian pyramid

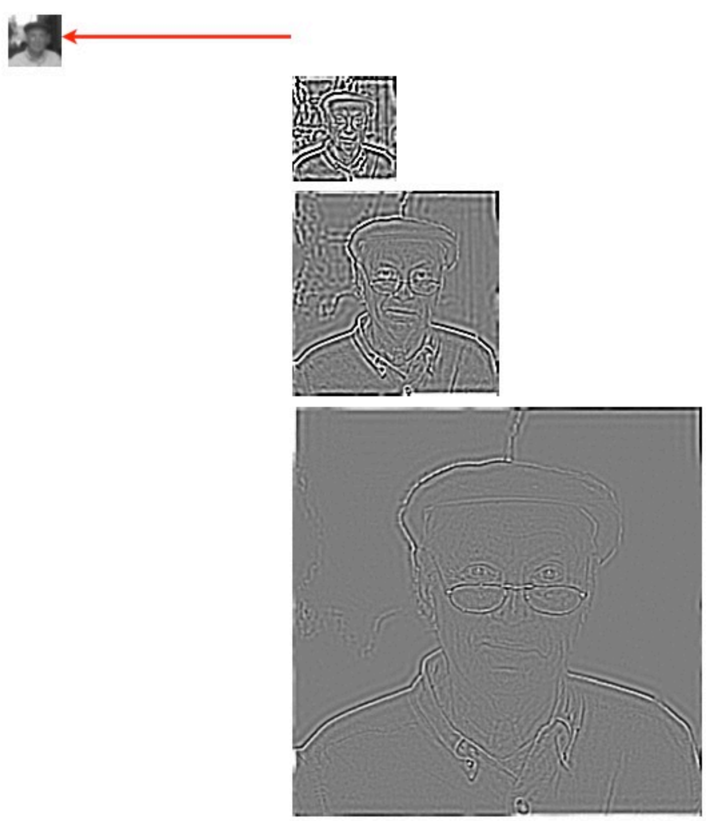

Laplacian pyramid

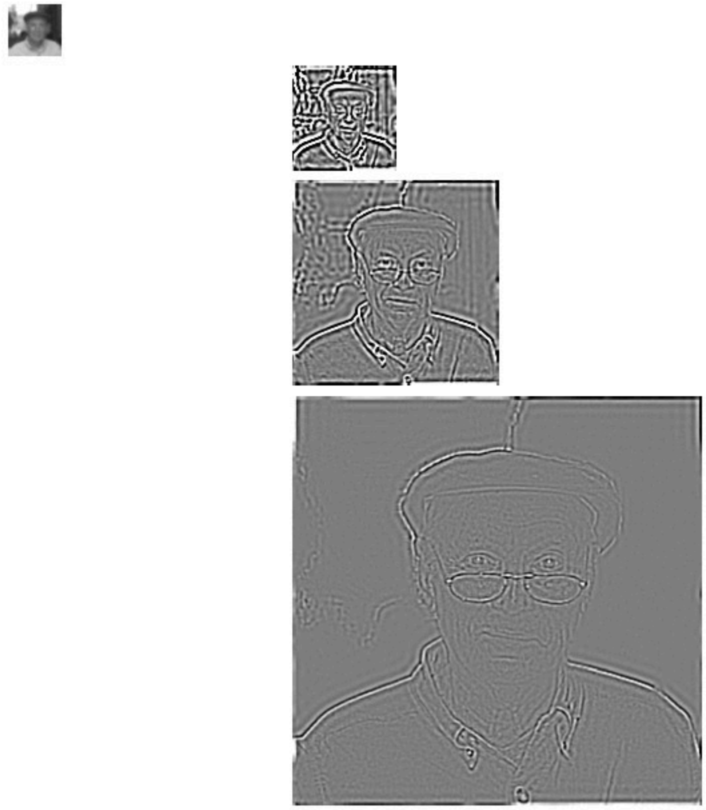

Laplacian pyramid

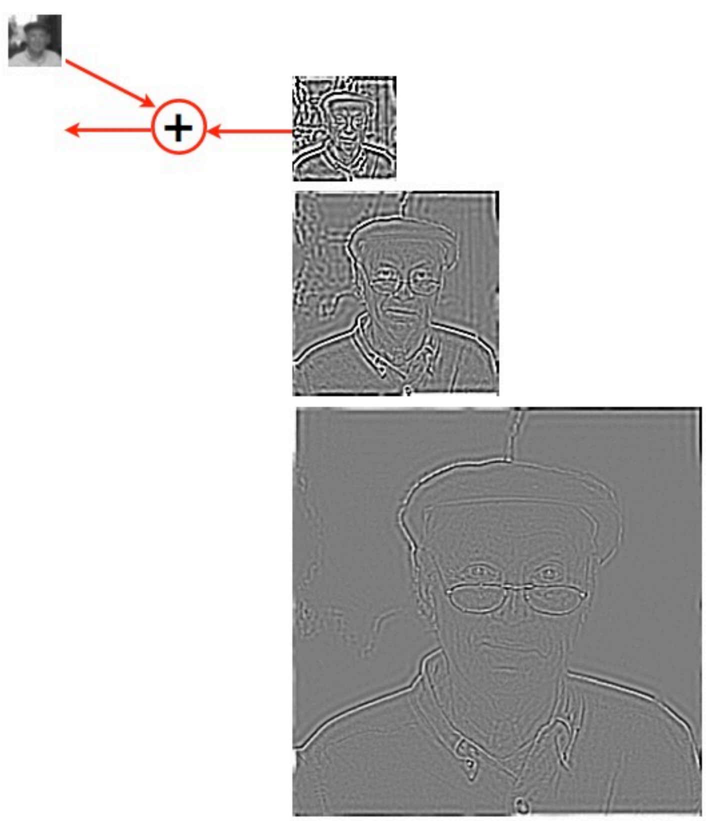

Laplacian pyramid

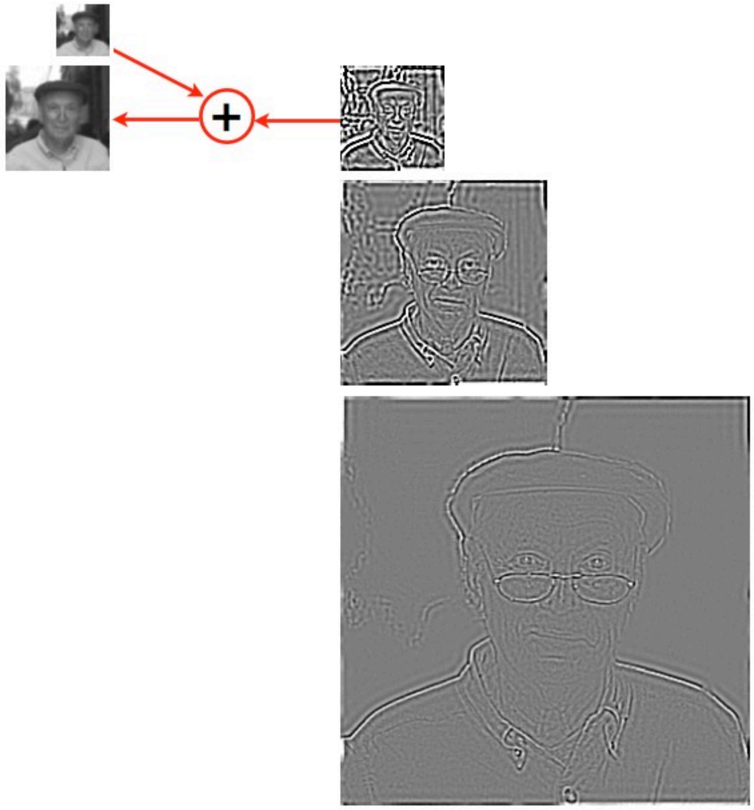

Laplacian pyramid

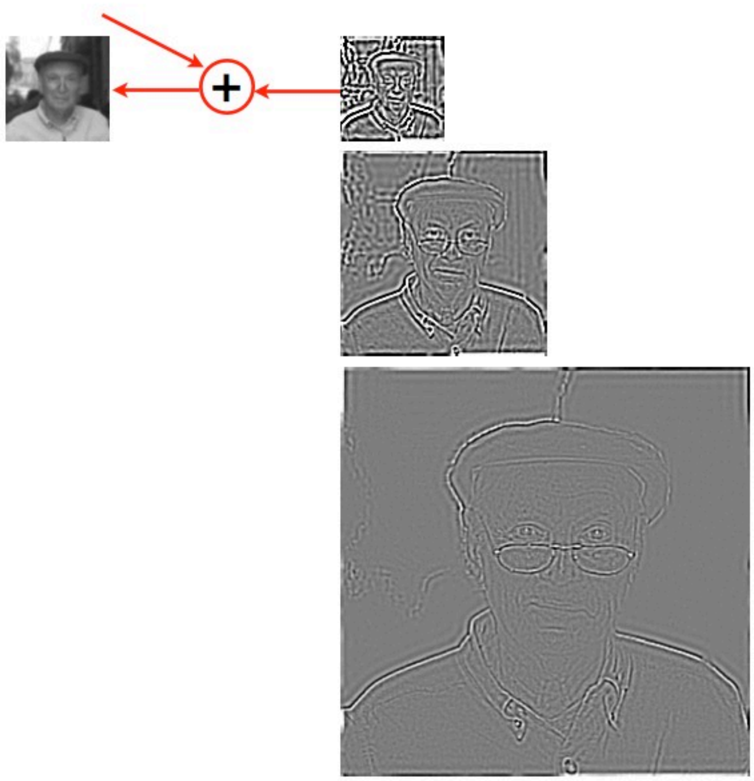

Laplacian pyramid

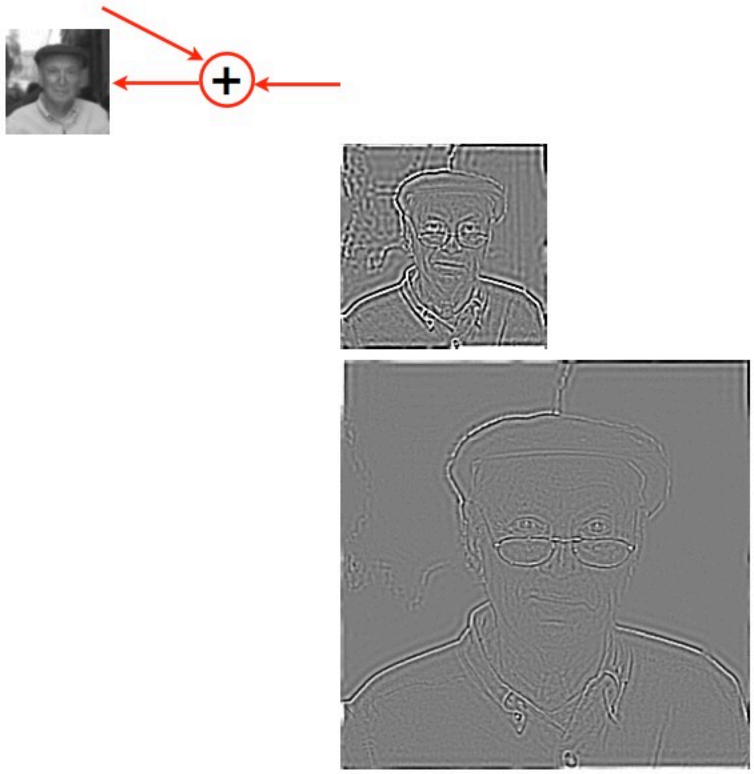

Laplacian pyramid

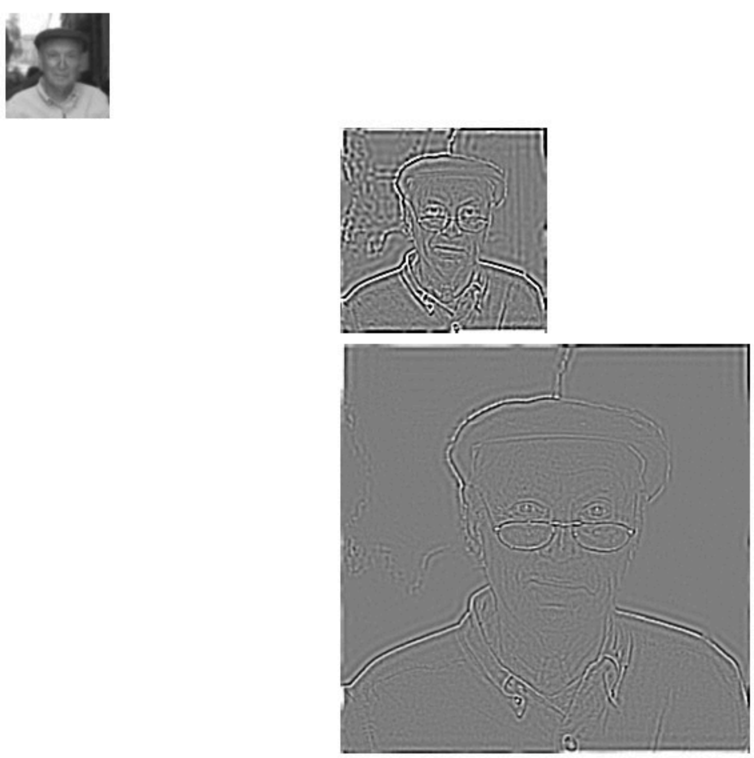

Laplacian pyramid

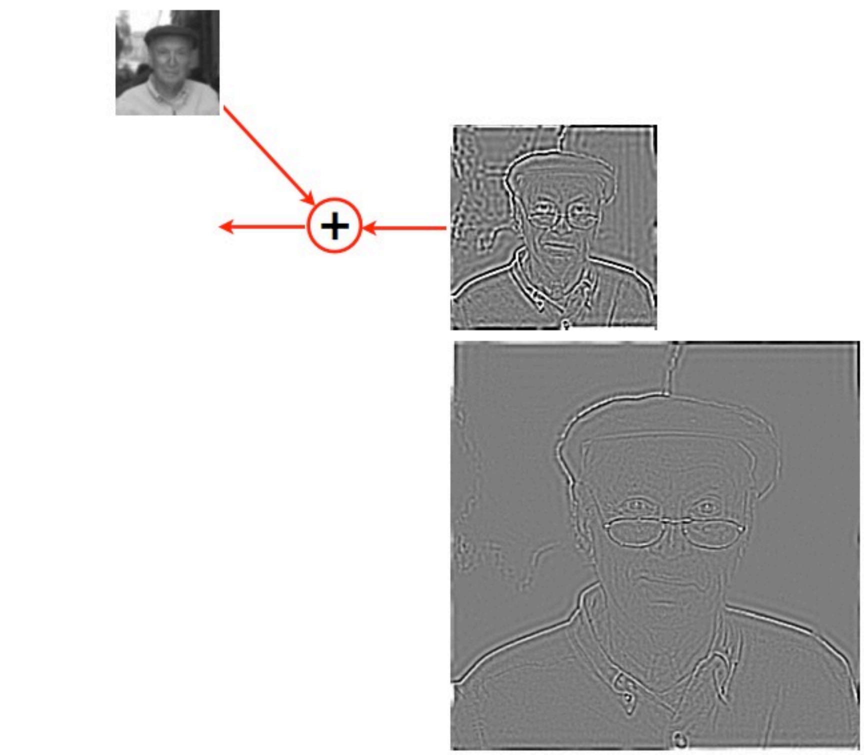

Laplacian pyramid

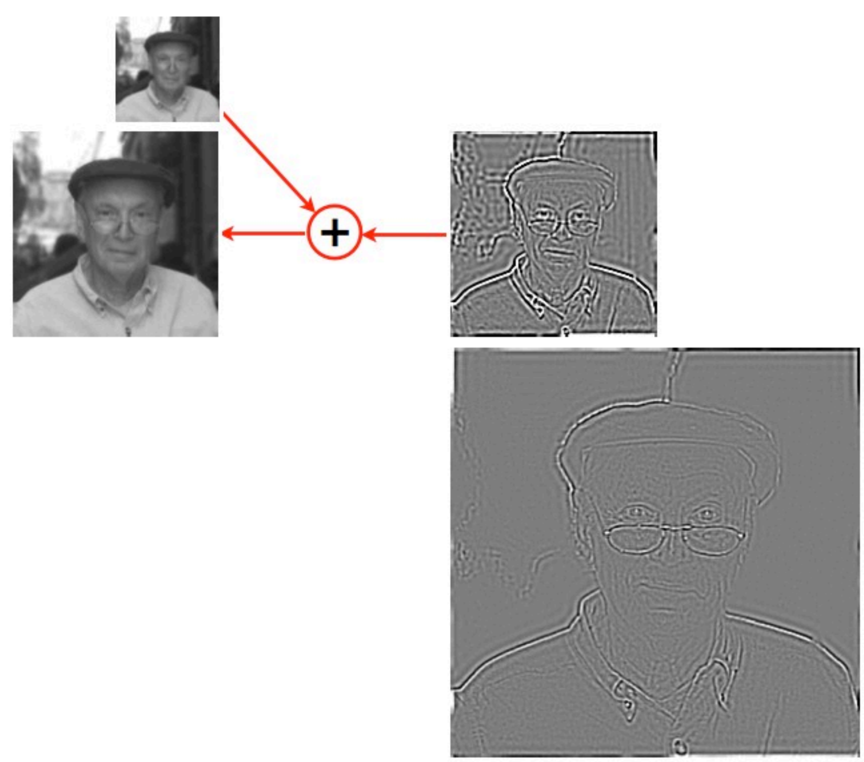

Laplacian pyramid

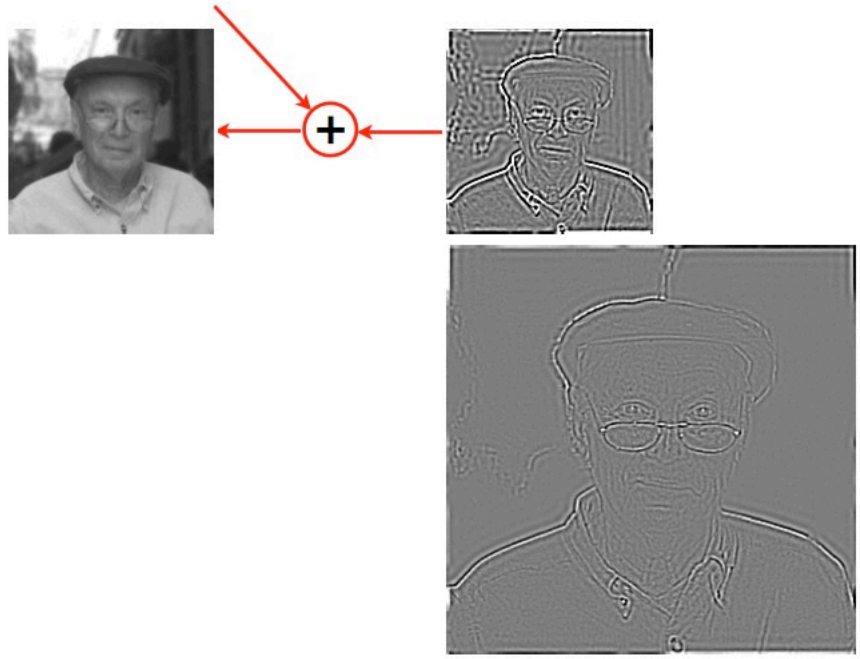

Laplacian pyramid

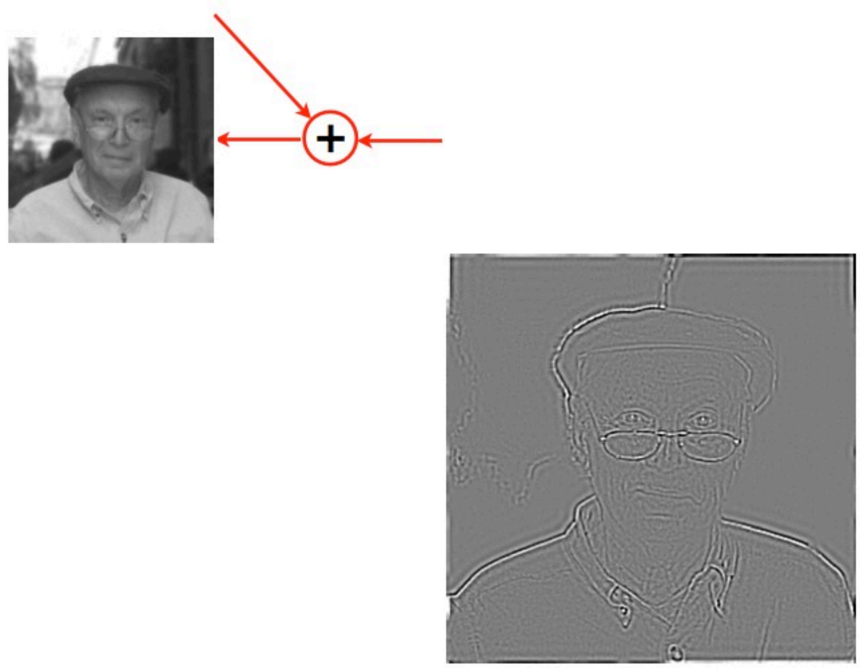

Laplacian pyramid

Laplacian pyramid

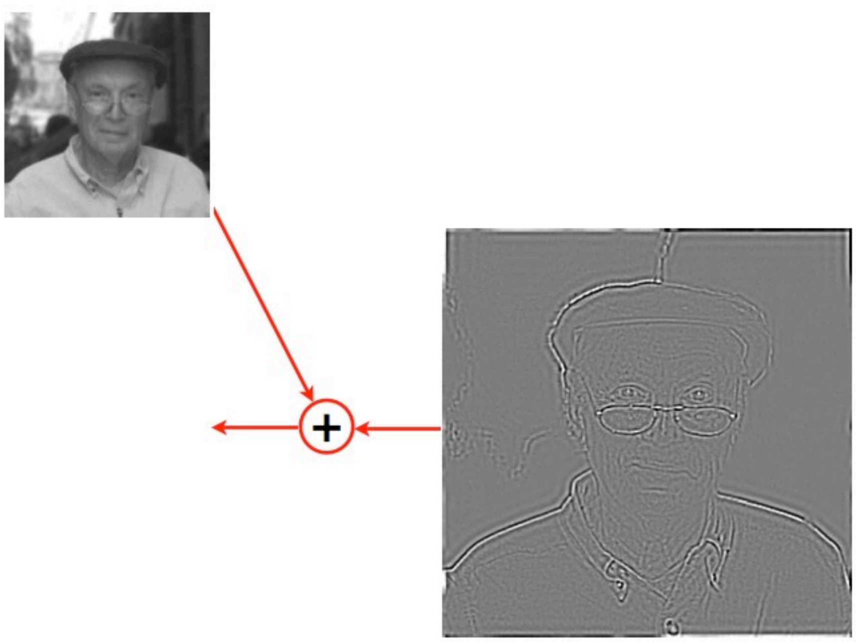

Laplacian pyramid

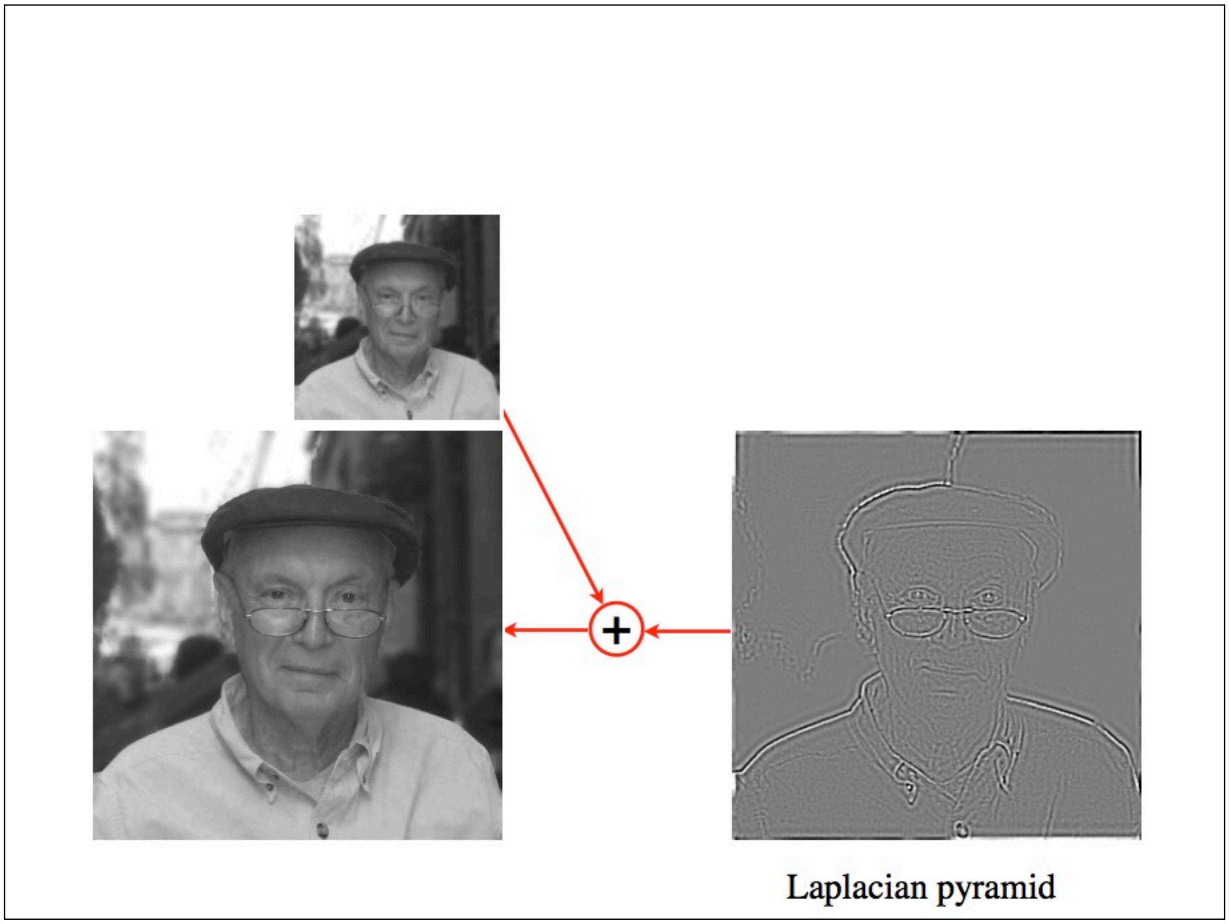
Laplacian pyramid

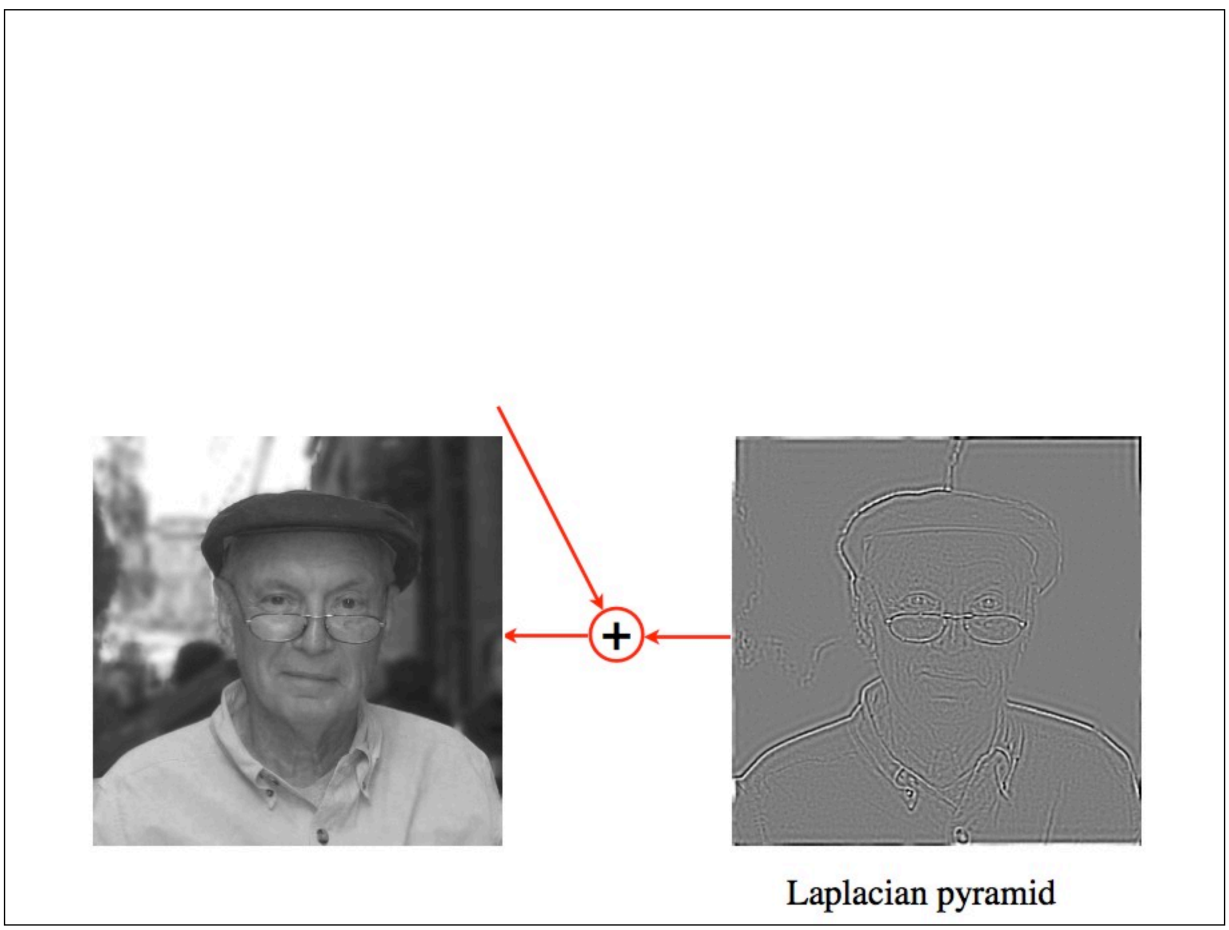
Laplacian pyramid

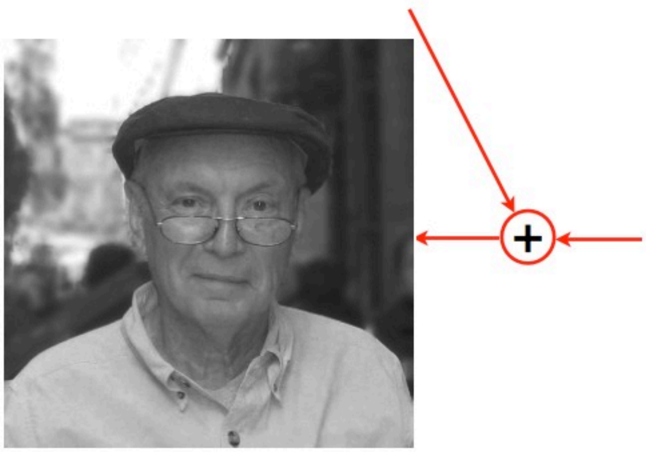

Laplacian pyramid

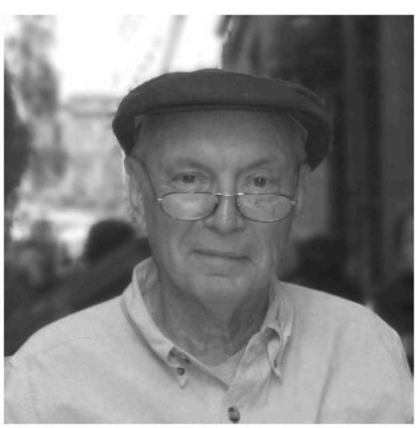

Laplacian pyramid

reconstructed image

Laplacian pyramid

Laplacian pyramid - signal processing diagram

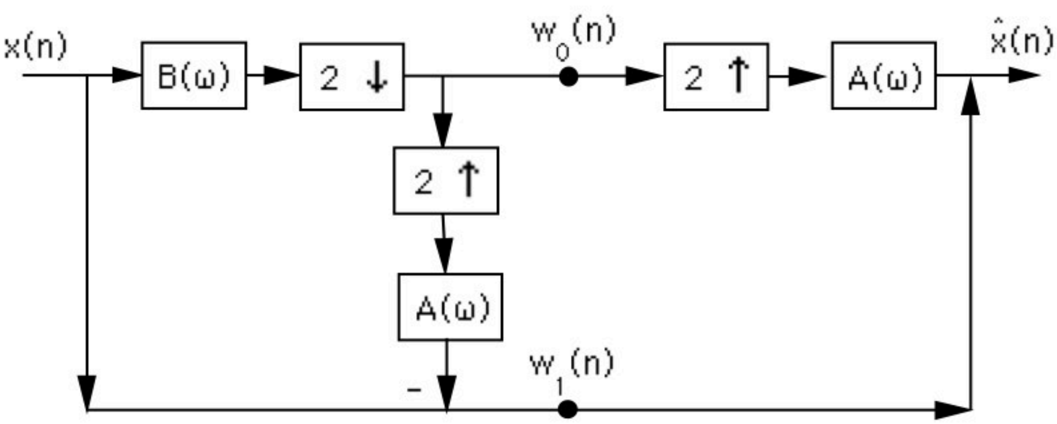

Laplacian pyramid - signal processing diagram

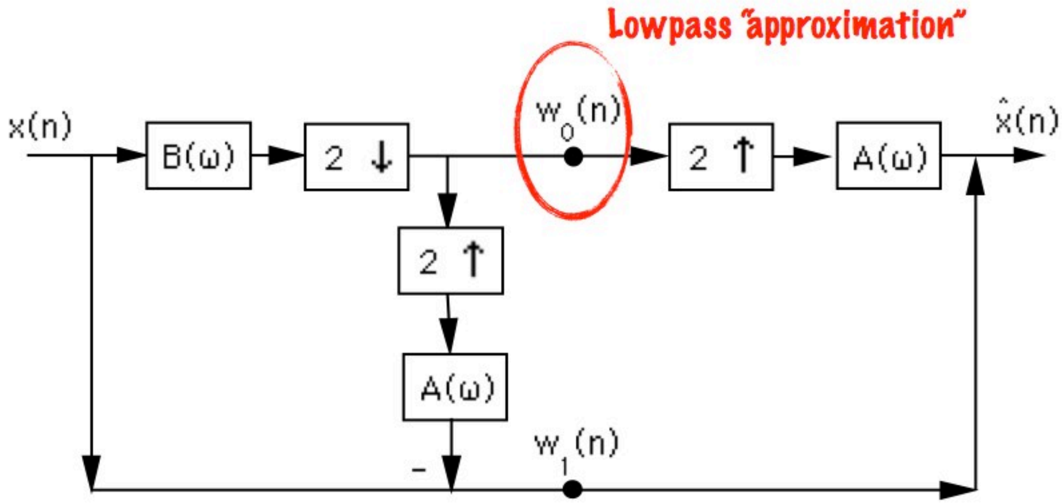

Laplacian pyramid - signal processing diagram

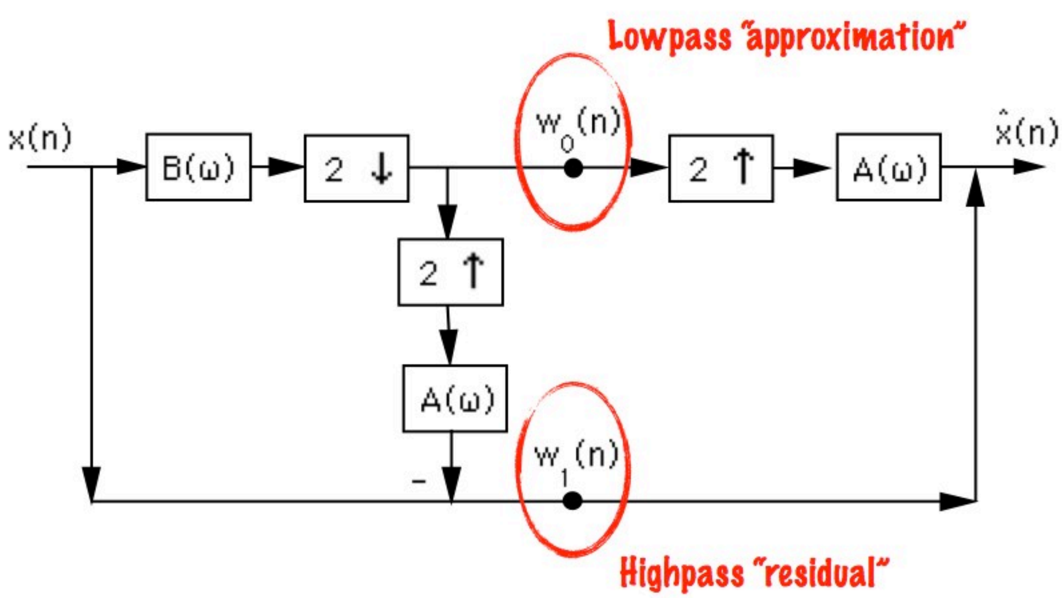

Laplacian pyramid - signal processing diagram

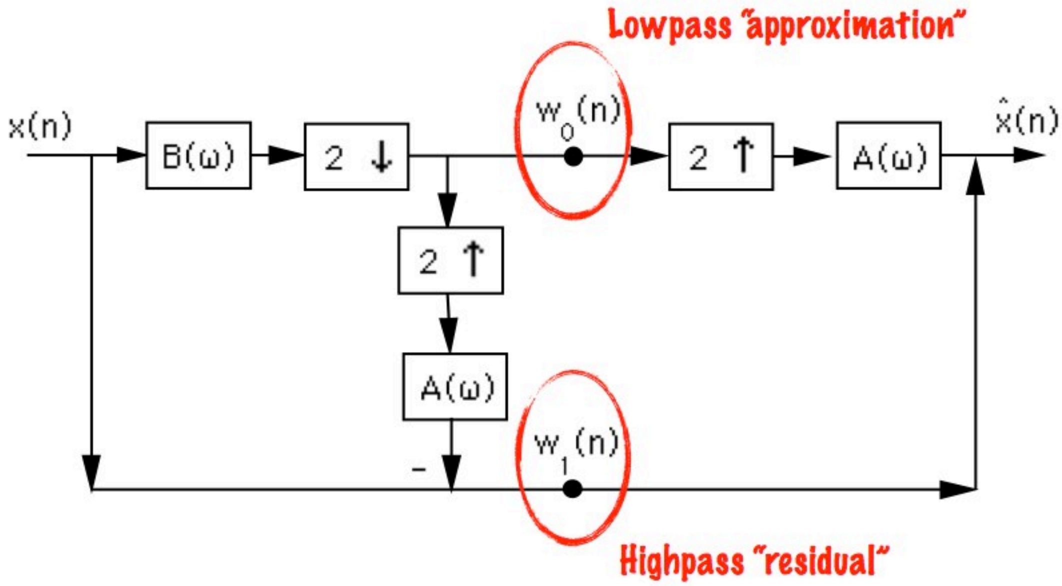

Note: perfect reconstruction for any choice of B and A!

Matrices: Laplacian pyramid + inverse

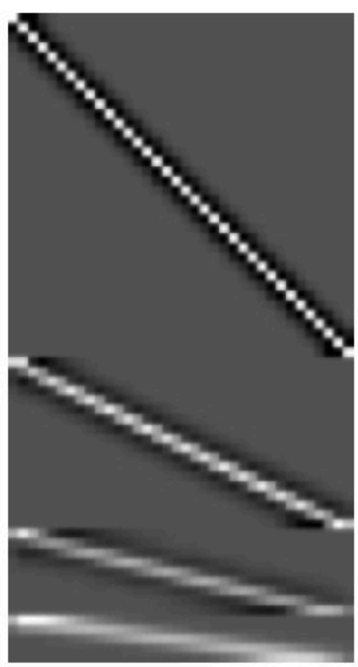

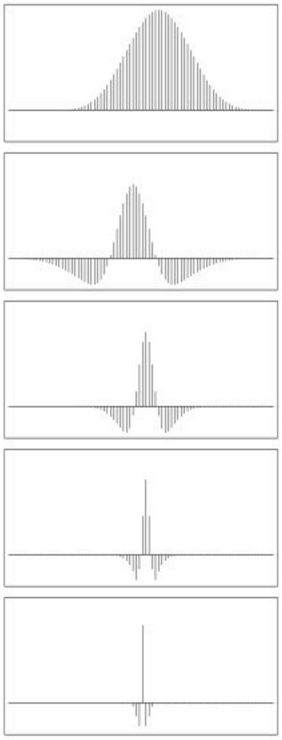

example projection functions

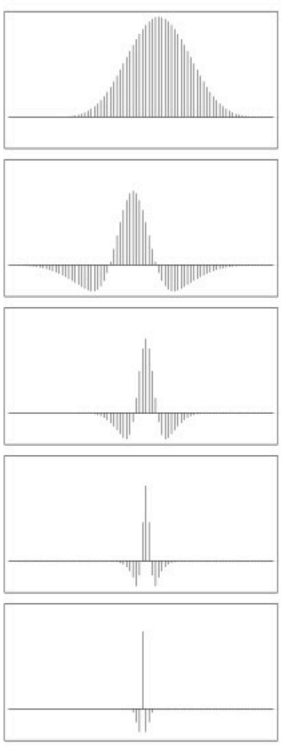

example projection functions

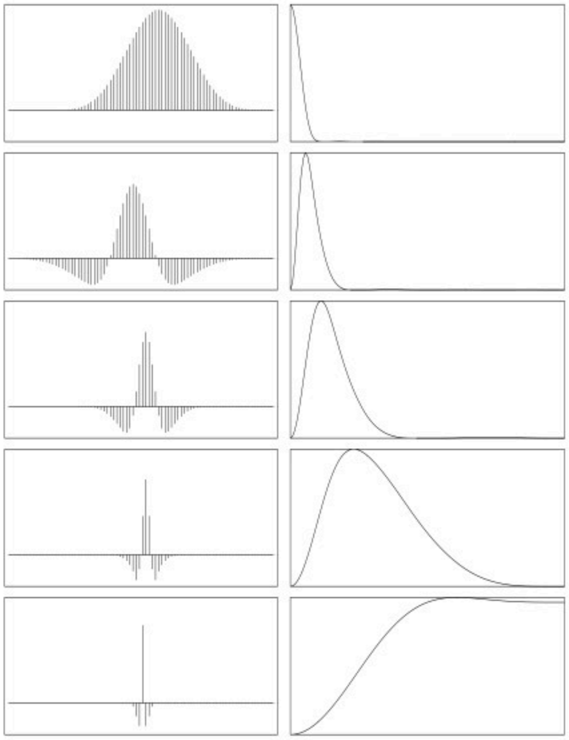

example projection functions
and their Fourier amplitudes

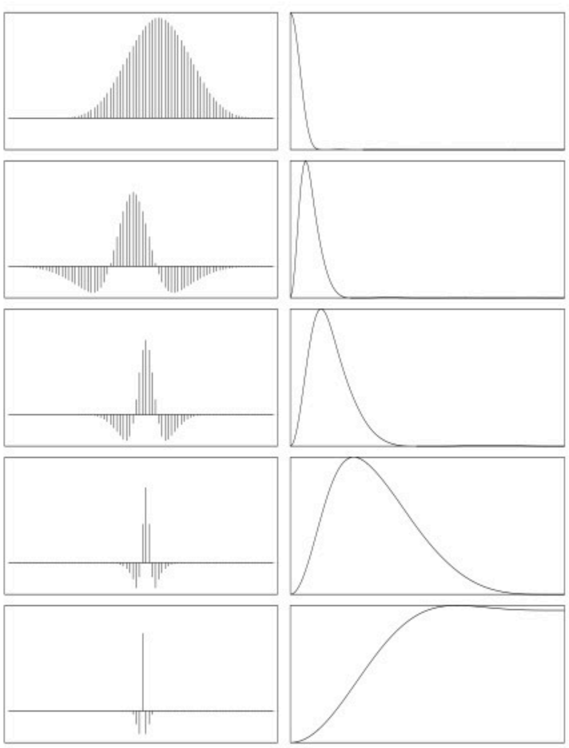

example projection functions
and their Fourier amplitudes

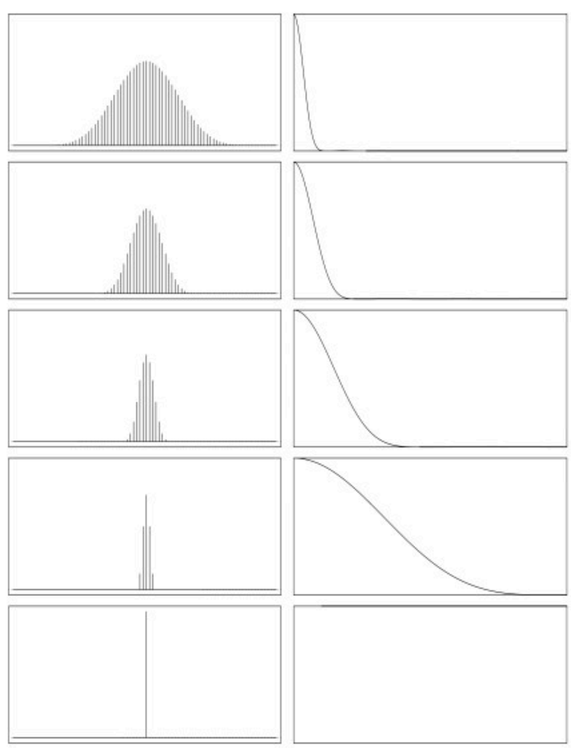

example basis functions
and their Fourier amplitudes

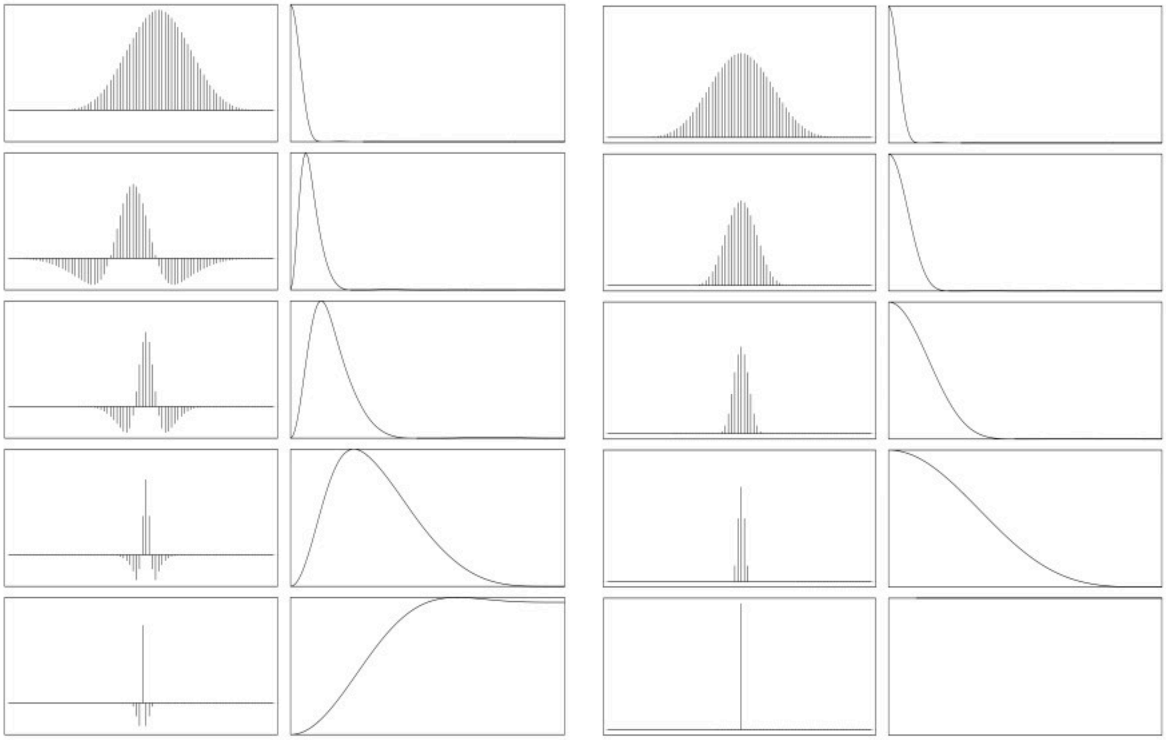

Assymmetry in forward/inverse
transform seems odd...

2D projection/basis functions

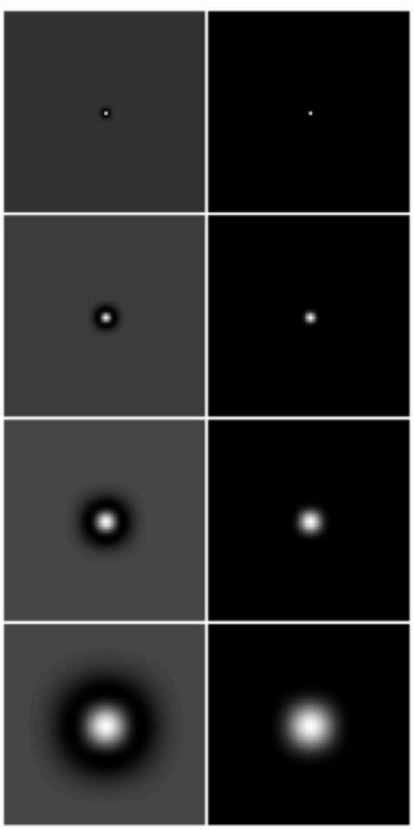

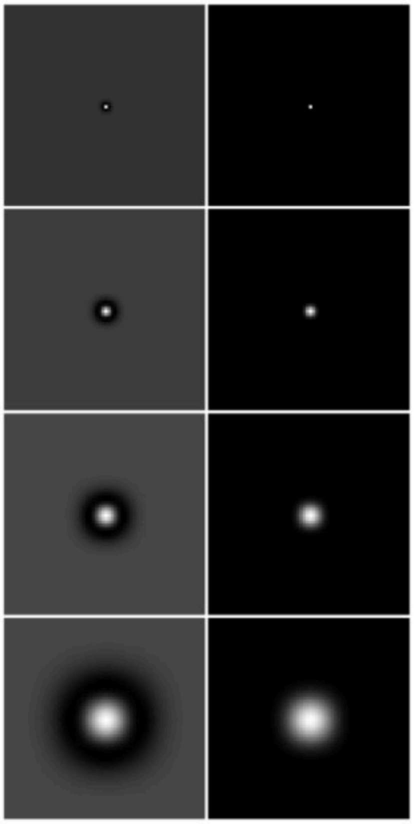

2D projection/basis functions

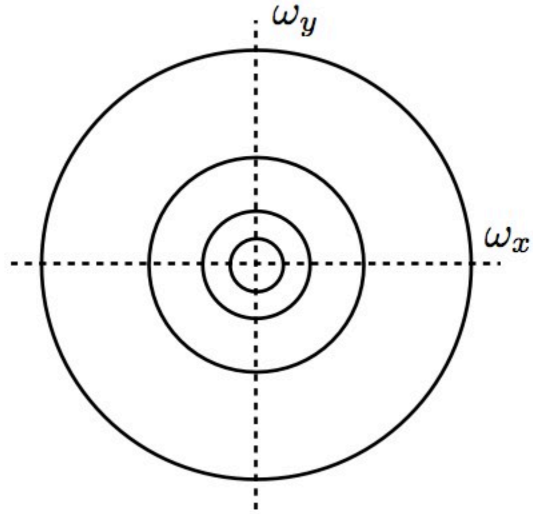

Analysis/synthesis filter bank

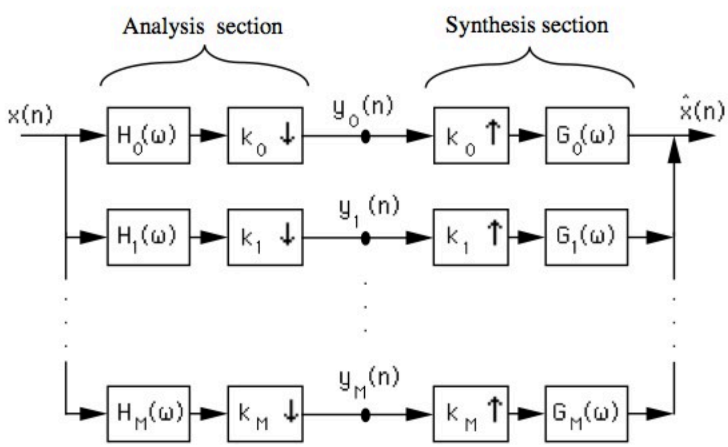

$$Y_m(\omega) = \frac{1}{k_m}\sum_{n=0}^{k_m-1} H_m\left(\frac{\omega}{k_m}+\frac{2\pi n}{k_m}\right) X\left(\frac{\omega}{k_m}+\frac{2\pi n}{k_m}\right)$$

$$\hat{X}(\omega) = \sum_{m=0}^{M-1} Y_m(k_m\,\omega)\; G_m(\omega)$$

Common frequency partitions

spatial (pixel)
basis

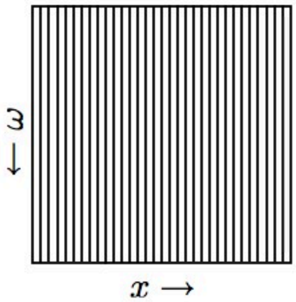

Common frequency partitions

spatial (pixel)
basis

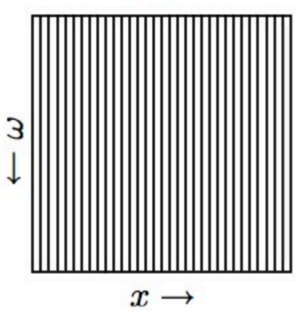

frequency (Fourier)
basis

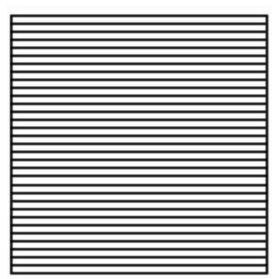

Common frequency partitions

spatial (pixel) basis

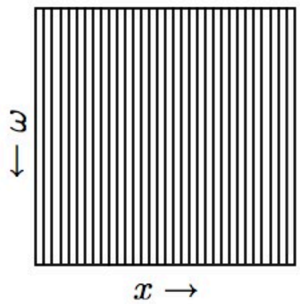

uniform subband basis

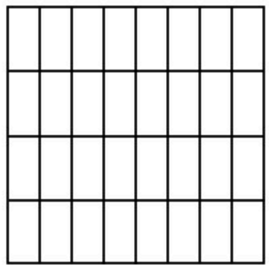

(e.g., block DCT)

frequency (Fourier) basis

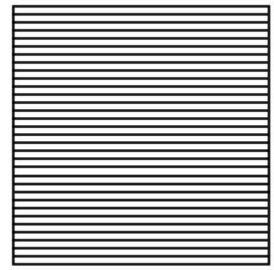

Common frequency partitions

spatial (pixel) basis

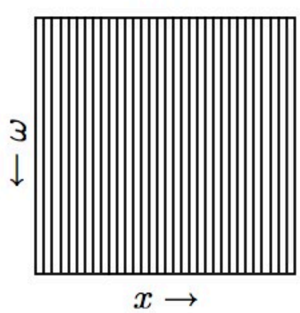

uniform subband basis

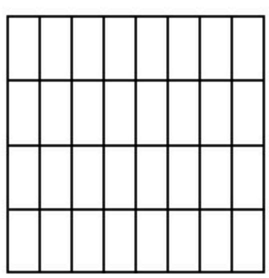

(e.g., block DCT)

octave subband basis

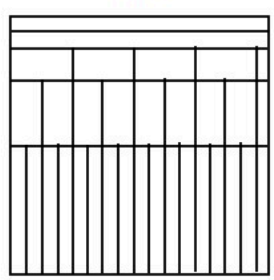

(e.g., lapl. pyramid)
(a.k.a. “constant Q”)

frequency (Fourier) basis

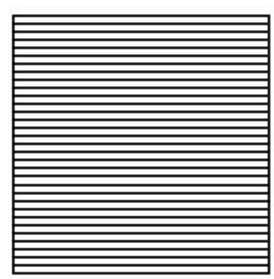

Common frequency partitions

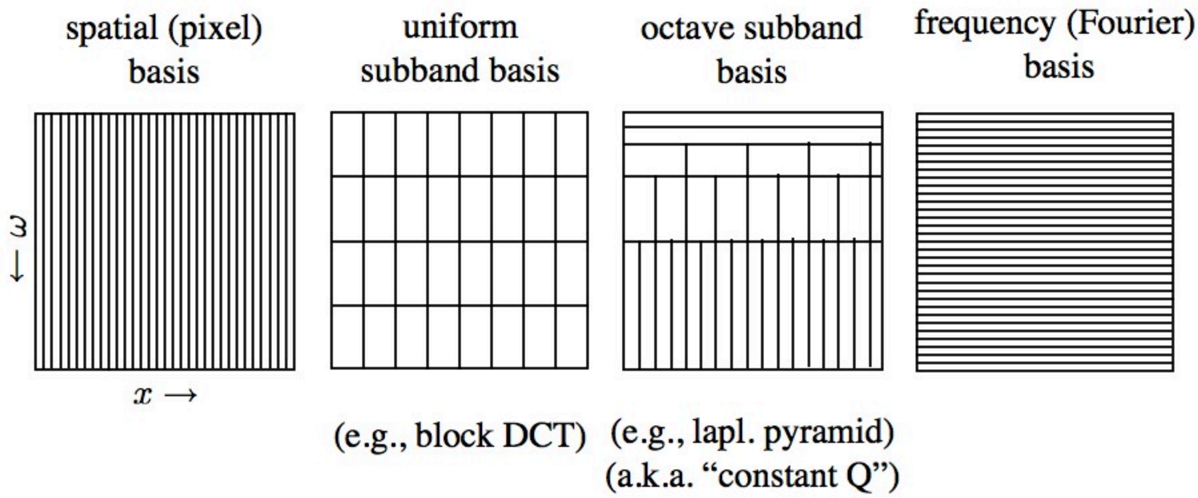

- Joint localization (product of width in space and frequency) is bounded from below [Heisenberg Uncertainty Principle].
- The bound is achieved by Gaussians, or Gaussian-windowed sinusoids ("Gabor" functions)

Desirable A/S properties

- minimal (ideally, zero) reconstruction error
- minimal aliasing within subbands
- space-frequency localization
- overcompleteness
  - special case: "critical sampling": $\sum_{m=0}^{M-1} \frac{1}{k_m} = 1$
- dis-similarity of sampling and basis functions
  - special case: self-inverting ("tight frame"): $G_m(\omega) = H_m(-\omega)$
- symmetry (or anti-symmetry) of basis functions

Some A/S examples

- identity (pixel basis)
- Fourier transform
- Block frequency (DFT, DCT) transforms
- Gabor transform
- Laplacian pyramid
- 2-band orthogonal (QMF)
- dyadic wavelets

16-point "block" DCT

- basis of cosine functions [on board]
- orthogonal:
  + perfect reconstruction
  + self-inverting
  + critical sampling
- non-overlapping blocks:
  + spatially local
  - frequency nonlocal
  - heavy aliasing

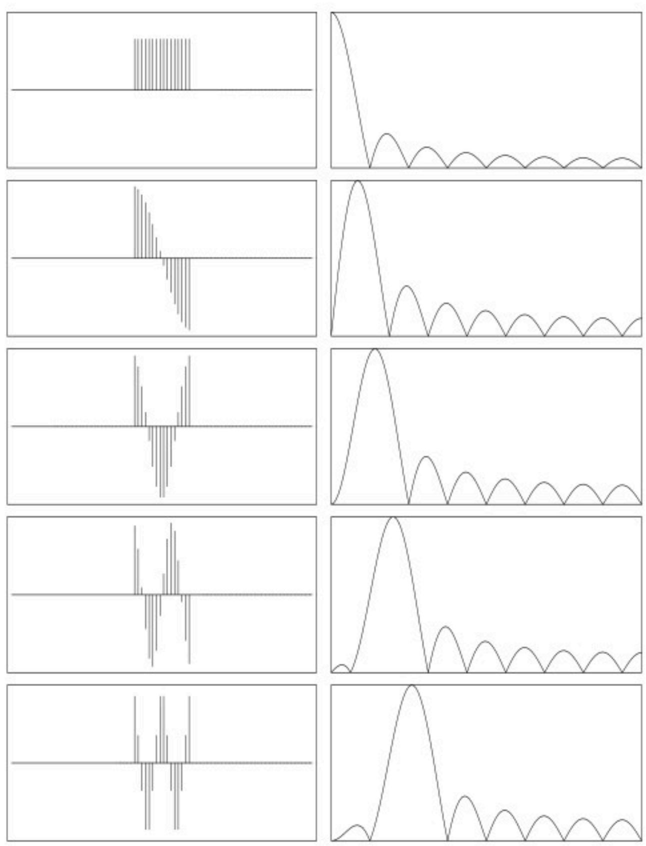

16-point "block" DCT

- basis of cosine functions [on board]
- orthogonal:
  + perfect reconstruction
  + self-inverting
  + critical sampling
- non-overlapping blocks:
  + spatially local
  - frequency nonlocal
  - heavy aliasing

1D DCT matrix

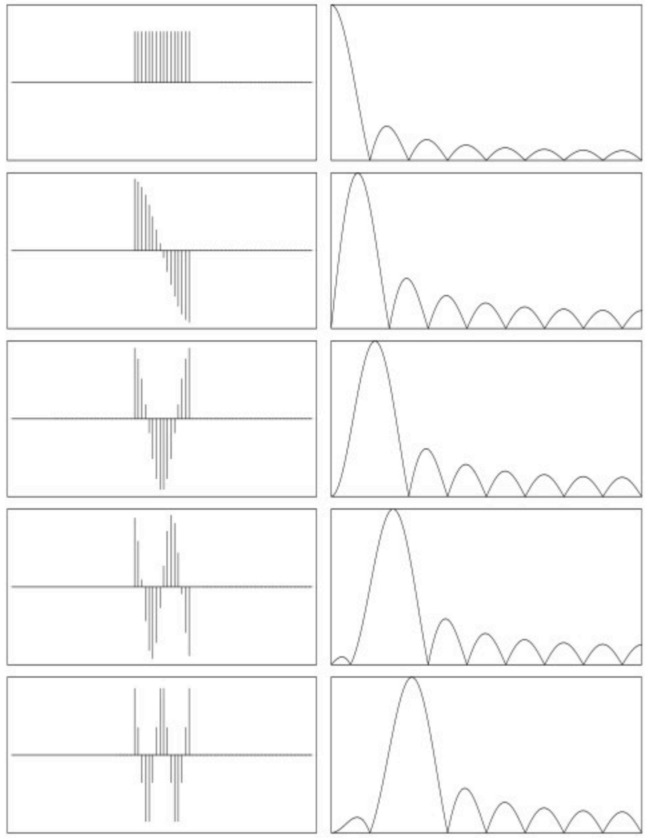

A uniform Gabor transform

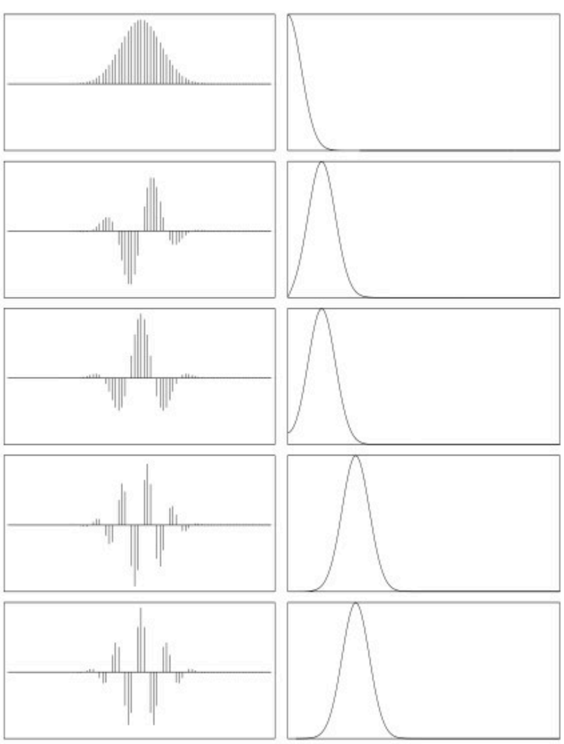

filters　　Fourier mags

A uniform Gabor transform

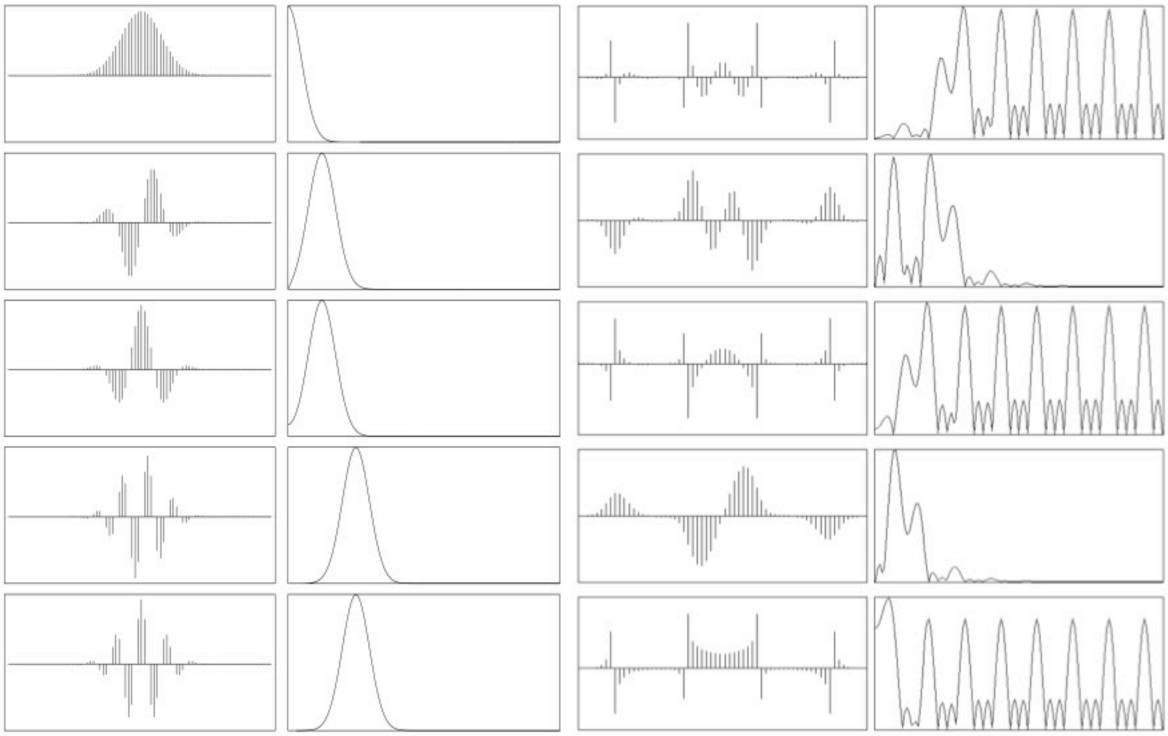

filters Fourier mags

Laplacian pyramid

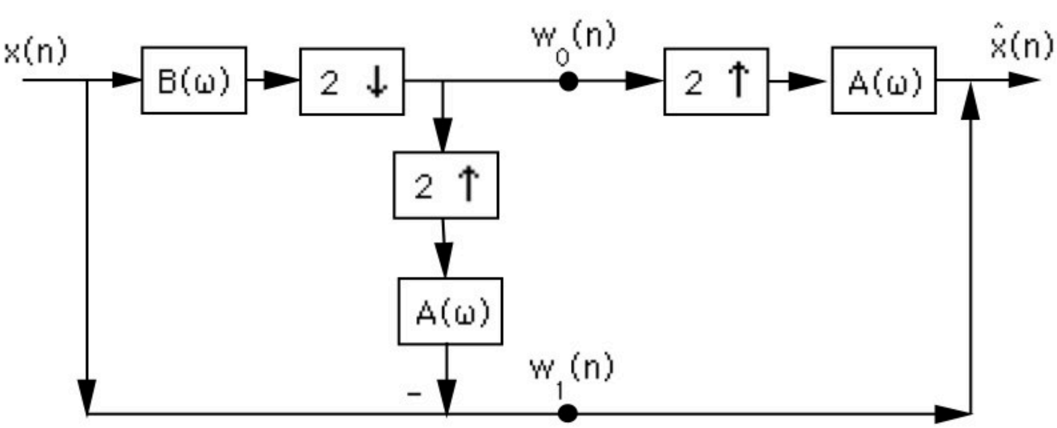

Laplacian pyramid as a 3-band A/S system

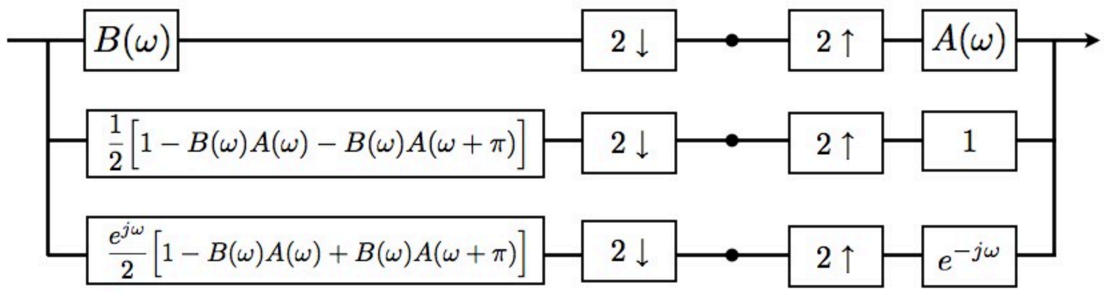

+ perfect reconstruction (for any A/B)
+ localized in space/freq
+ minimal aliasing (with proper choice of A/B)
- not self-inverting (bandpass vs. lowpass)
- overcomplete (factor of 2 in 1D, 4/3 in 2D)

Laplacian pyramid as a 3-band A/S system

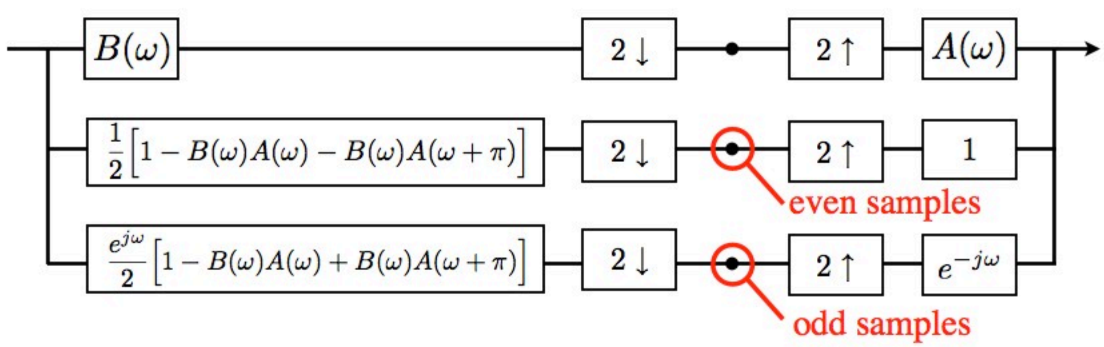

+ perfect reconstruction (for any A/B)
+ localized in space/freq
+ minimal aliasing (with proper choice of A/B)
- not self-inverting (bandpass vs. lowpass)
- overcomplete (factor of 2 in 1D, 4/3 in 2D)

2-band dyadic system

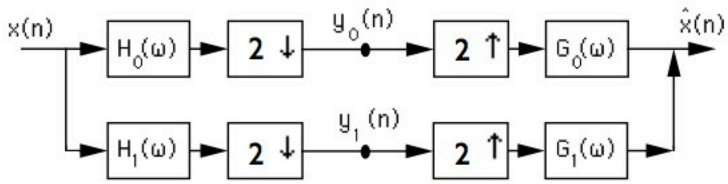

$$\hat{X}(\omega) = \frac{1}{2}\Big[H_0(\omega)G_0(\omega) + H_1(\omega)G_1(\omega)\Big]X(\omega)$$
$$+ \frac{1}{2}\Big[H_0(\omega+\pi)G_0(\omega) + H_1(\omega+\pi)G_1(\omega)\Big]X(\omega+\pi).$$

2-band dyadic system

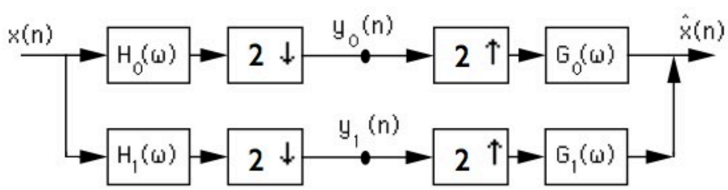

$$\hat{X}(\omega) = \frac{1}{2}\Big[H_0(\omega)G_0(\omega) + H_1(\omega)G_1(\omega)\Big]X(\omega)$$
$$+ \frac{1}{2}\Big[H_0(\omega+\pi)G_0(\omega) + H_1(\omega+\pi)G_1(\omega)\Big]X(\omega+\pi).$$

Aliasing

2-band dyadic system

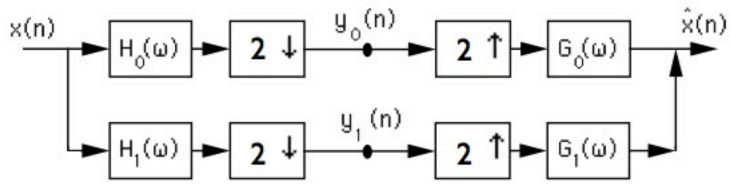

$$\hat{X}(\omega) = \frac{1}{2}\Big[H_0(\omega)G_0(\omega) + H_1(\omega)G_1(\omega)\Big]X(\omega)$$
$$+ \frac{1}{2}\Big[H_0(\omega+\pi)G_0(\omega) + H_1(\omega+\pi)G_1(\omega)\Big]X(\omega+\pi).$$

Aliasing

Choose
$$\begin{aligned} H_0(\omega) &= G_0(-\omega) = F(\omega) \\ H_1(\omega) &= G_1(-\omega) = e^{j\omega}F(-\omega+\pi) \end{aligned}$$

2-band dyadic system

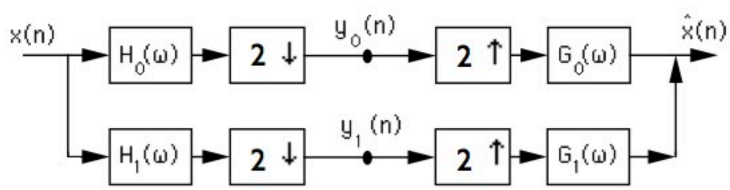

$$\hat{X}(\omega) = \frac{1}{2}\Big[H_0(\omega)G_0(\omega) + H_1(\omega)G_1(\omega)\Big]X(\omega)$$
$$+ \frac{1}{2}\Big[H_0(\omega+\pi)G_0(\omega) + H_1(\omega+\pi)G_1(\omega)\Big]X(\omega+\pi).$$

Aliasing

Choose
$$\begin{aligned} H_0(\omega) &= G_0(-\omega) = F(\omega) \\ H_1(\omega) &= G_1(-\omega) = e^{j\omega}F(-\omega+\pi) \end{aligned}$$

Then $\hat{X}(\omega) = \frac{1}{2}\Big[|F(\omega)|^2 + |F(\omega+\pi)|^2\Big]X(\omega)$

2-band dyadic system

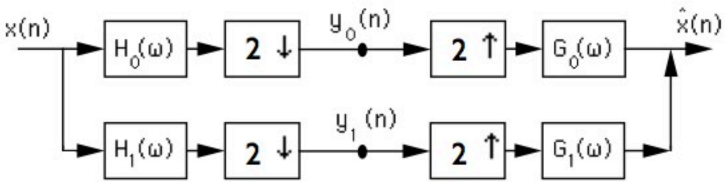

$$\hat{X}(\omega) = \frac{1}{2}\Big[H_0(\omega)G_0(\omega) + H_1(\omega)G_1(\omega)\Big]X(\omega)$$
$$+ \frac{1}{2}\Big[H_0(\omega+\pi)G_0(\omega) + H_1(\omega+\pi)G_1(\omega)\Big]X(\omega+\pi).$$

Choose
$$\begin{aligned} H_0(\omega) &= G_0(-\omega) = F(\omega) \\ H_1(\omega) &= G_1(-\omega) = e^{j\omega}F(-\omega+\pi) \end{aligned}$$

Then $\hat{X}(\omega) = \frac{1}{2}\Big[|F(\omega)|^2 + |F(\omega+\pi)|^2\Big]X(\omega)$

2-band dyadic system

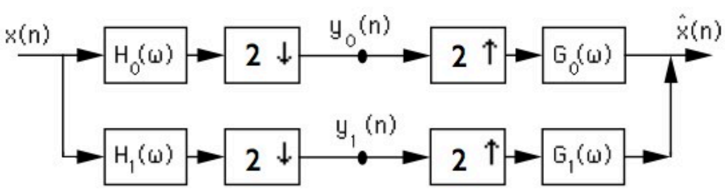

$$\hat{X}(\omega) = \frac{1}{2}\Big[H_0(\omega)G_0(\omega) + H_1(\omega)G_1(\omega)\Big]X(\omega)$$
$$+ \frac{1}{2}\Big[H_0(\omega+\pi)G_0(\omega) + H_1(\omega+\pi)G_1(\omega)\Big]X(\omega+\pi).$$

Choose
$$\begin{aligned} H_0(\omega) &= G_0(-\omega) = F(\omega) \\ H_1(\omega) &= G_1(-\omega) = e^{j\omega}F(-\omega+\pi) \end{aligned}$$

Quadrature mirror filters (QMF)

Then $\hat{X}(\omega) = \frac{1}{2}\Big[|F(\omega)|^2 + |F(\omega+\pi)|^2\Big]X(\omega)$

- Perfect reconstruction => aliasing

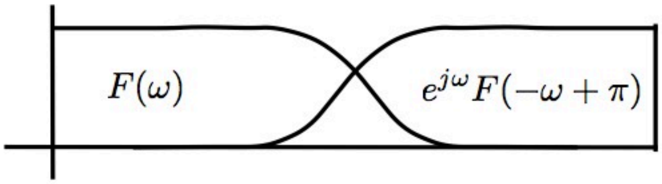

- Subbands live in orthogonal subspaces:

$e^{j\omega}F(-\omega+\pi)$

negate alternate samples

flip order

shift one sample

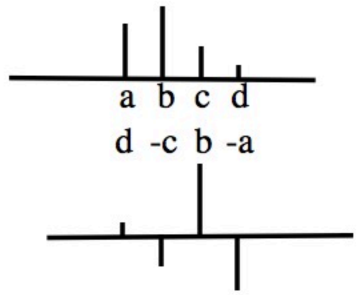

- examples: identity, sinc, Haard

Cascades

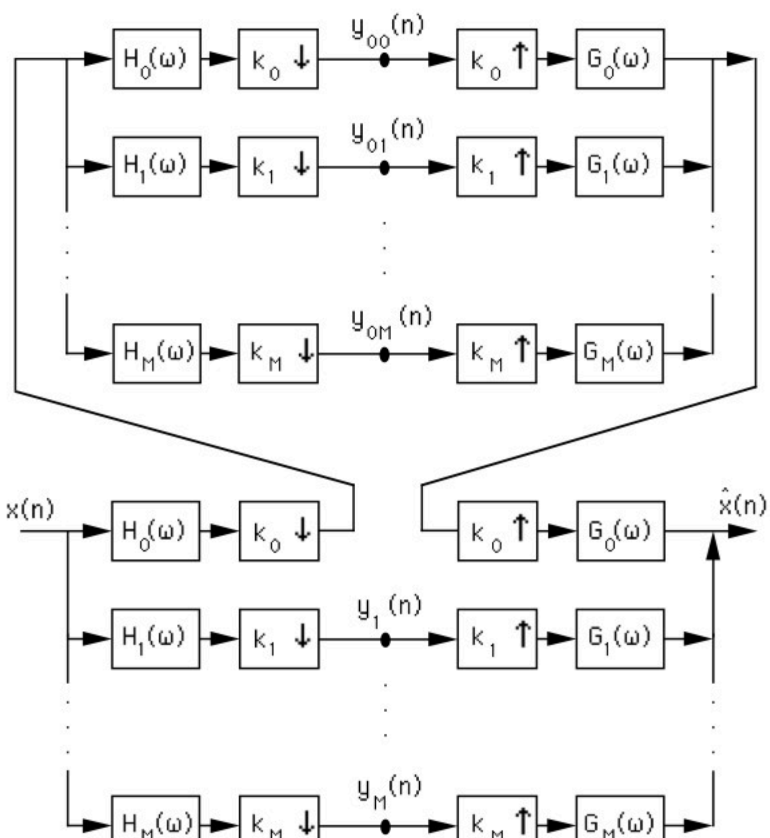

Haar (1909)

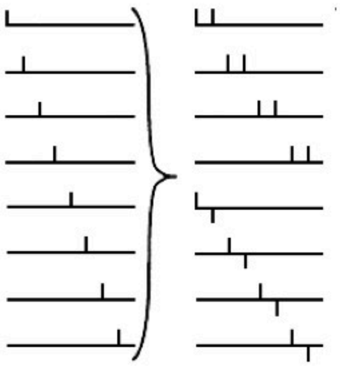

Haar (1909)

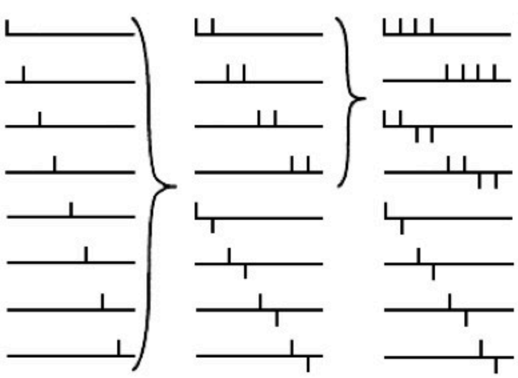

Haar (1909)

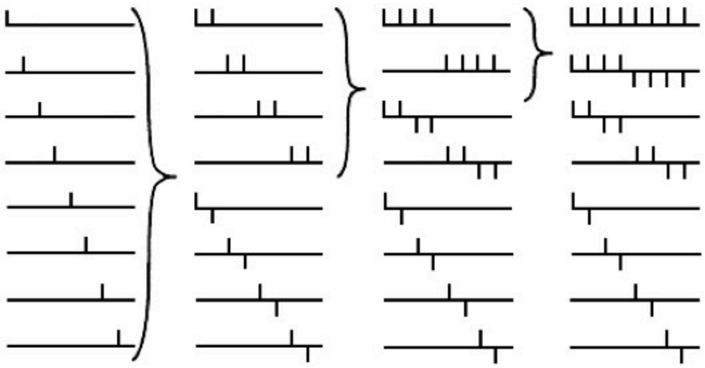

Haar (1909)

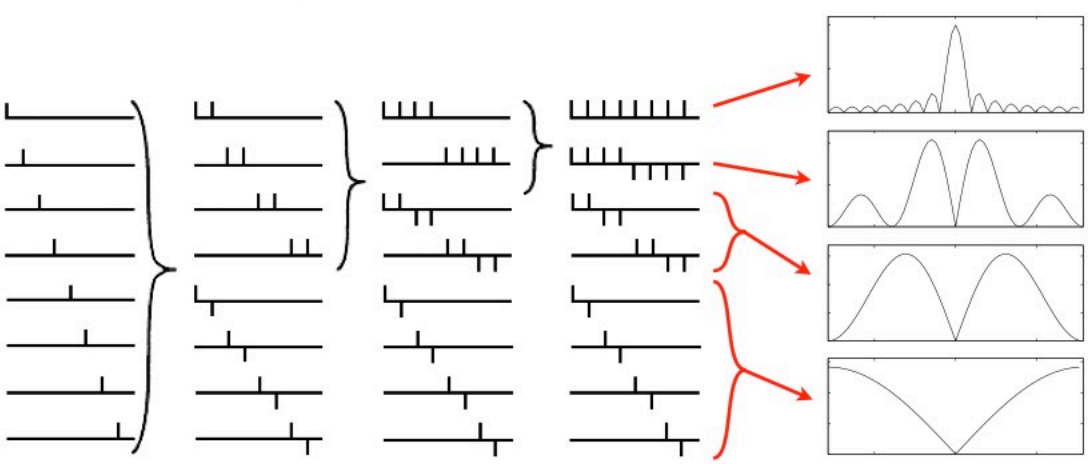

Haar (1909)

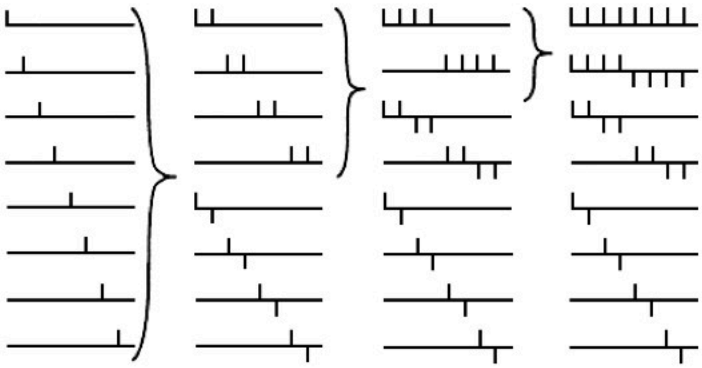

Haar, 2D

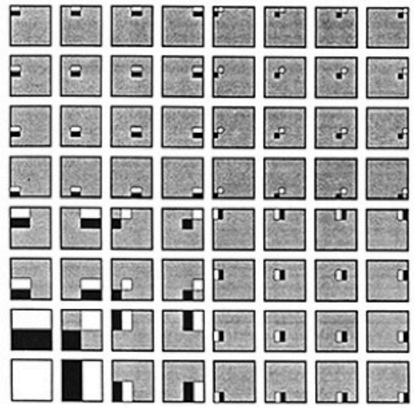

Cascaded 2-band (dyadic) system:

- octave subbands
- basis functions related by both translation *and* dilation

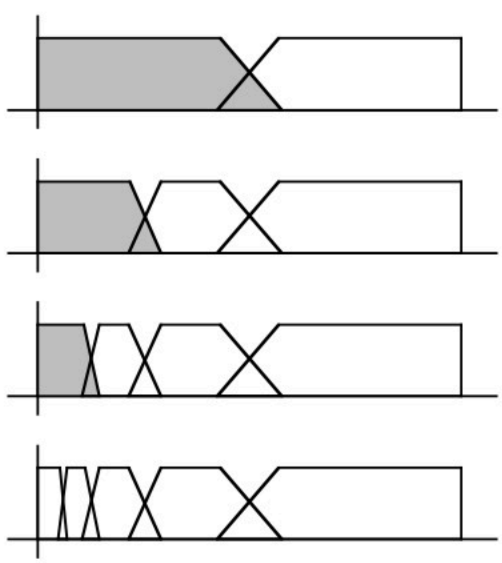

9-tap QMF pyramid basis and Fourier amplitudes

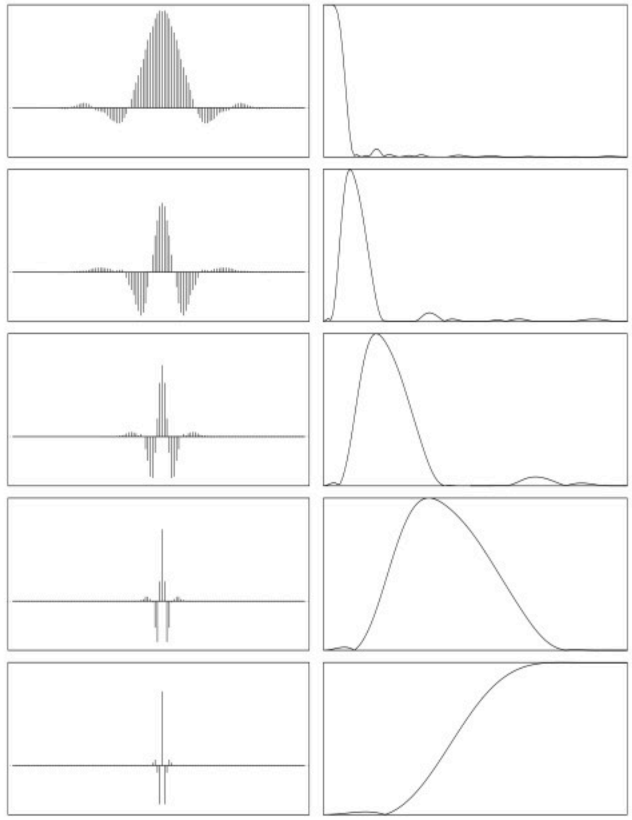

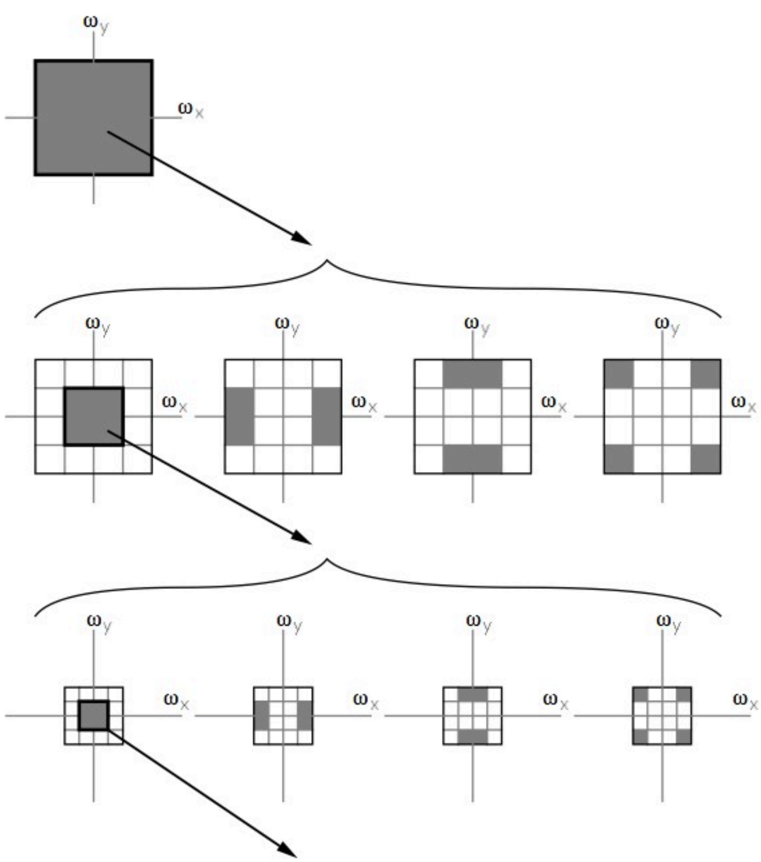

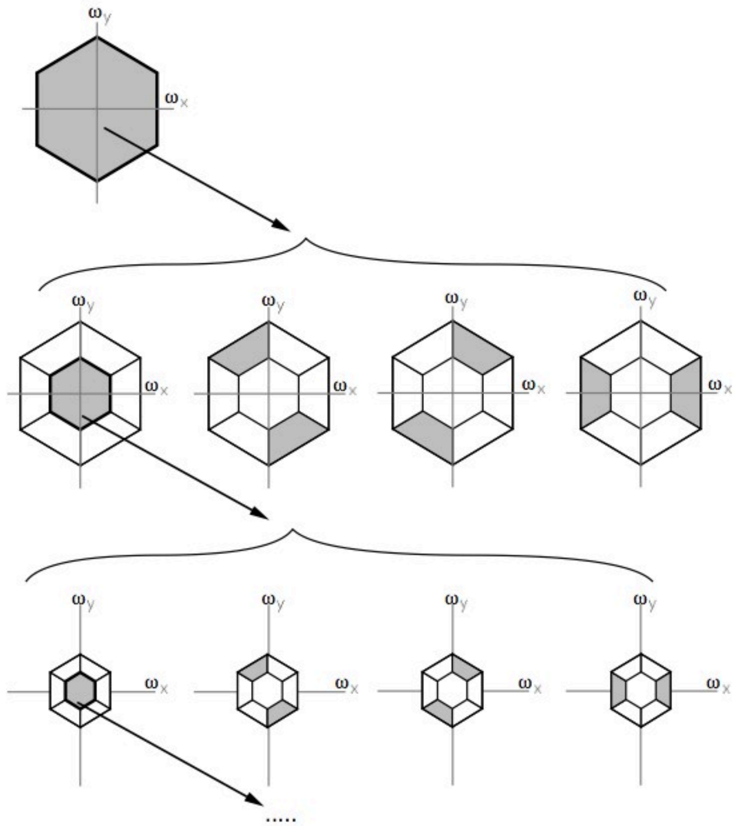

Walsh-Hadamard basis

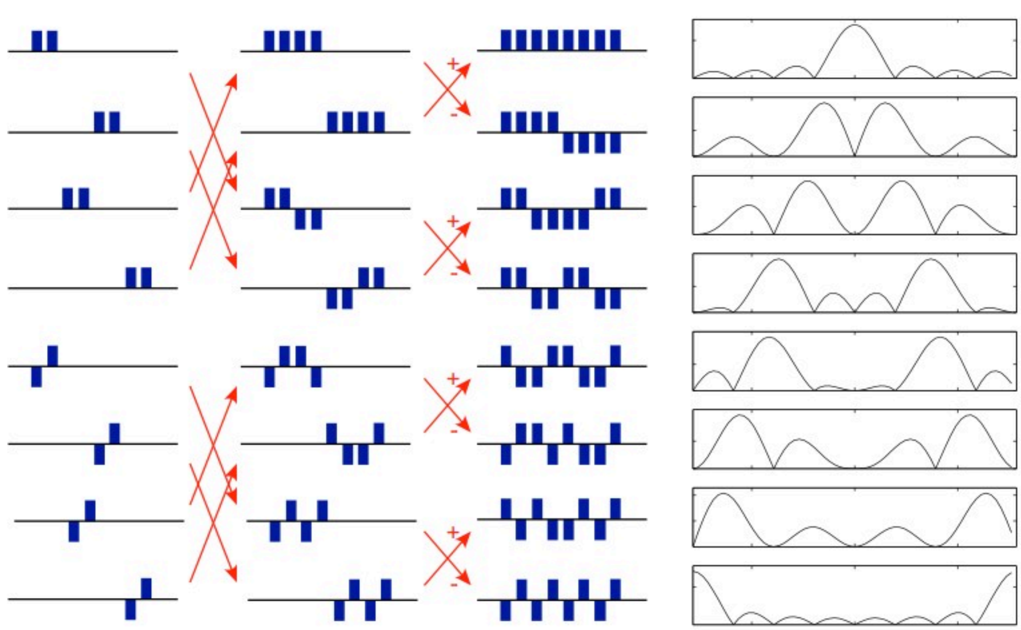

Using the same 2-band split as the Haar, recursively split all bands (result is like a "binarized" Fourier transform)

Aliasing in orthogonal wavelets

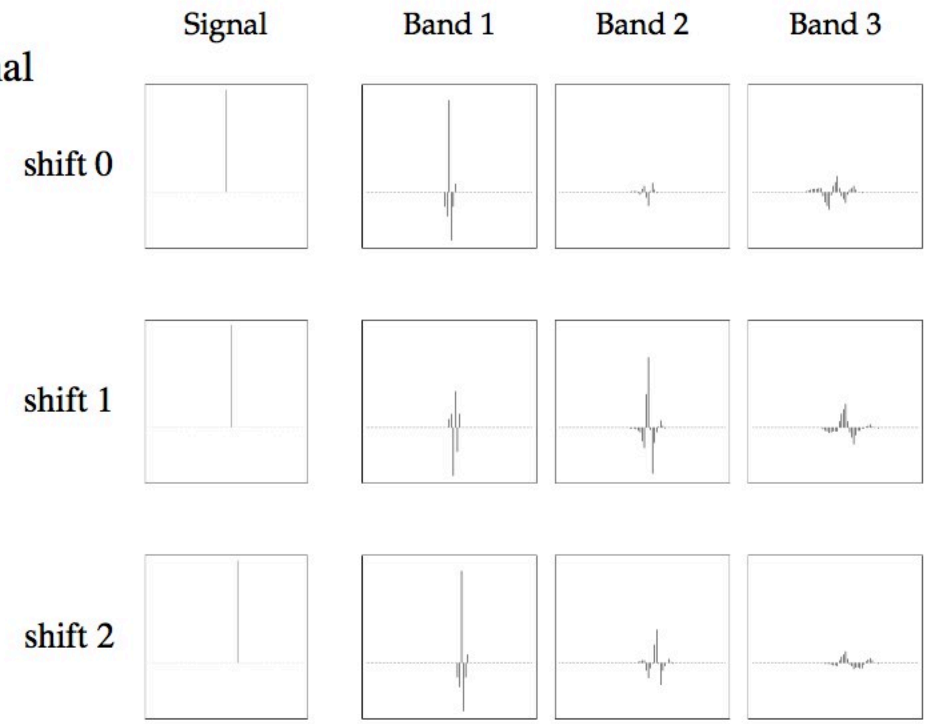

4-tap orthogonal wavelet [known as Daubechies-4]

An A/S wish-list

- basis functions related by translation, dilation (and rotation, in 2D)
- reasonably localized in both space and frequency
- minimal aliasing
- modest overcompleteness

- “self-inverting” (tight frame)
- steerable (i.e., no aliasing in orientation)
- efficient cascade implementation (pyramid)

An A/S wish-list

- basis functions related by translation, dilation (and rotation, in 2D) [no uniform subbands]
- reasonably localized in both space and frequency
- minimal aliasing
- modest overcompleteness

- "self-inverting" (tight frame)
- steerable (i.e., no aliasing in orientation)
- efficient cascade implementation (pyramid)

An A/S wish-list

- basis functions related by translation, dilation (and rotation, in 2D) [no uniform subbands] [no separable wavelets]
- reasonably localized in both space and frequency
- minimal aliasing
- modest overcompleteness

- "self-inverting" (tight frame)
- steerable (i.e., no aliasing in orientation)
- efficient cascade implementation (pyramid)

An A/S wish-list

- basis functions related by translation, dilation (and rotation, in 2D) [no uniform subbands] [no separable wavelets]
- reasonably localized in both space and frequency [no block transforms]
- minimal aliasing
- modest overcompleteness

- "self-inverting" (tight frame)
- steerable (i.e., no aliasing in orientation)
- efficient cascade implementation (pyramid)

An A/S wish-list

- basis functions related by translation, dilation (and rotation, in 2D) [no uniform subbands] [no separable wavelets]
- reasonably localized in both space and frequency [no block transforms]
- minimal aliasing [no critical sampling]
- modest overcompleteness

- "self-inverting" (tight frame)
- steerable (i.e., no aliasing in orientation)
- efficient cascade implementation (pyramid)

An A/S wish-list

- basis functions related by translation, dilation (and rotation, in 2D) [no uniform subbands] [no separable wavelets]
- reasonably localized in both space and frequency [no block transforms]
- minimal aliasing [no critical sampling]
- modest overcompleteness

- "self-inverting" (tight frame) [no Lapl. pyr, Gabor transform]
- steerable (i.e., no aliasing in orientation)
- efficient cascade implementation (pyramid)

Octave-bandwidth polar-separable representations

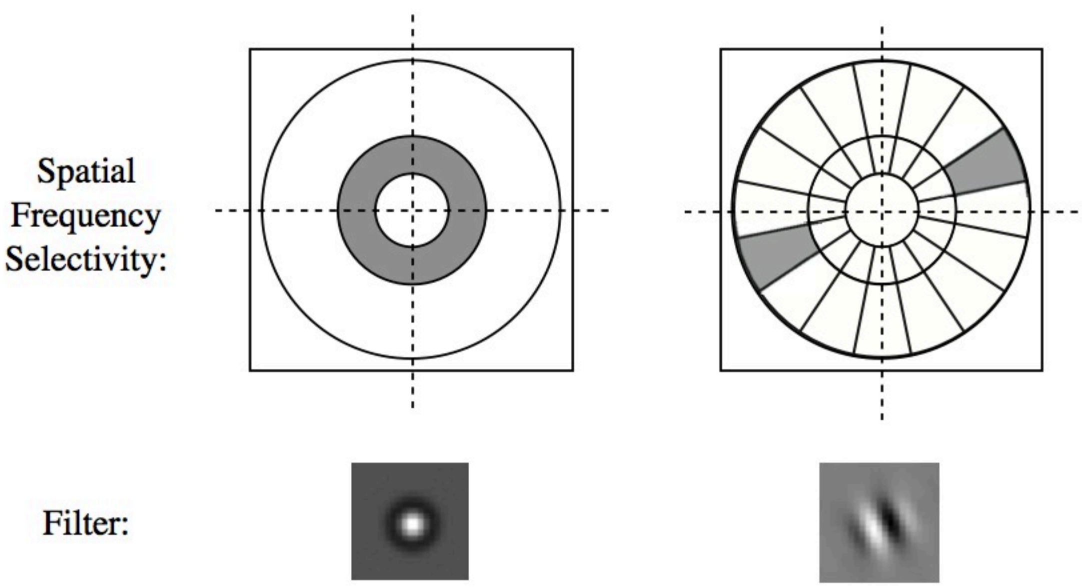

steerable pyramid
[similar to "curvelets"]

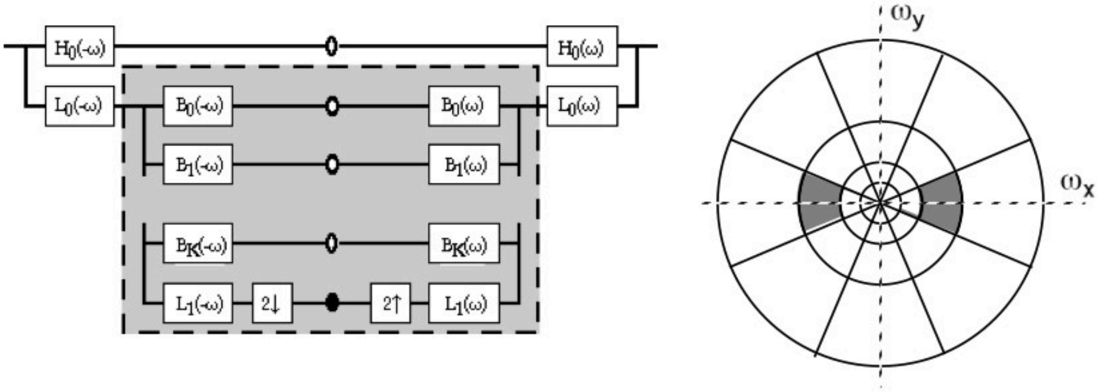

radial frequency partitions:

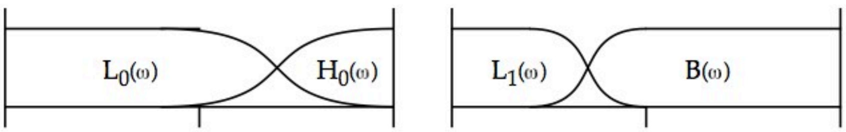

Aliasing-free ("shiftable") 1D pyramid

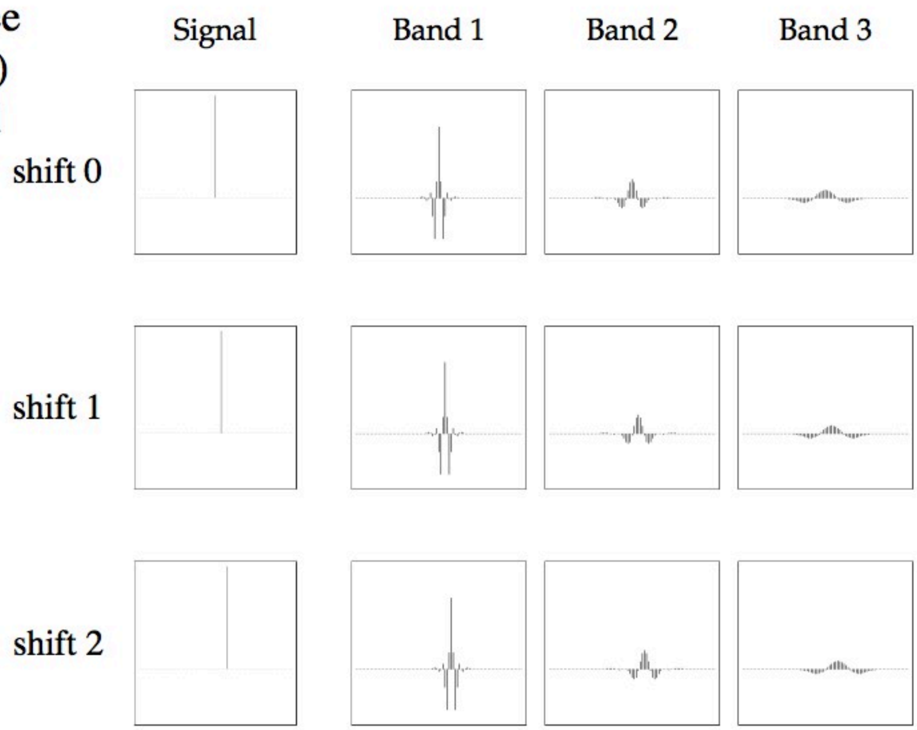

Measuring Orientation

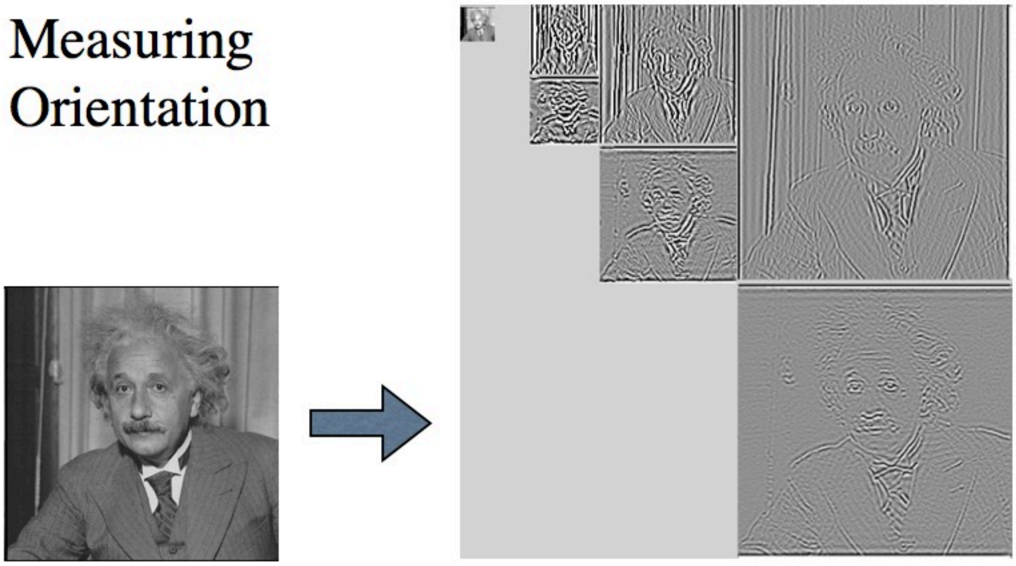

2-band steerable pyramid: Each scale contains horizontal/vertical derivatives -> multi-scale gradient
Course theme: *combines* analysis *and* representation

3-band steerable pyramid

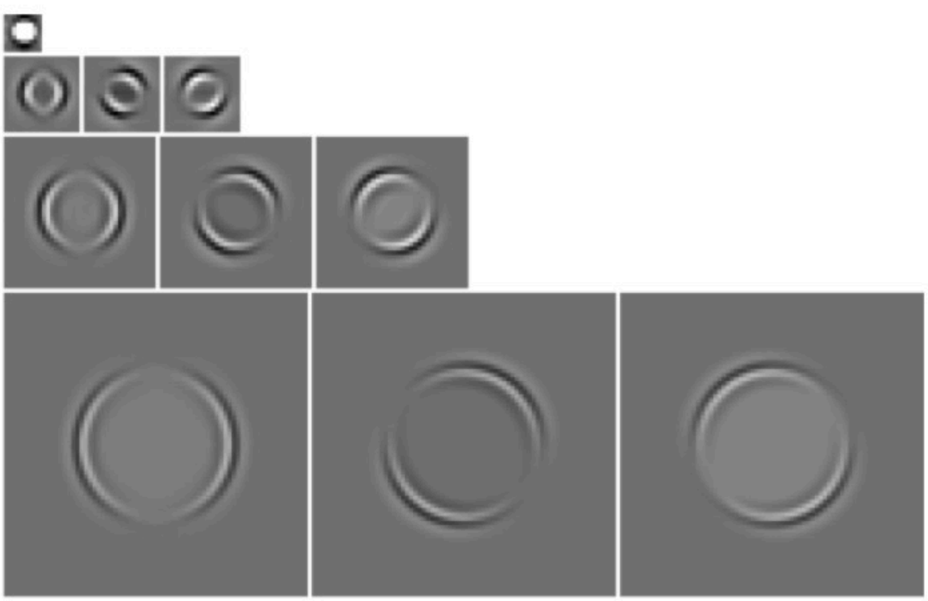

	freq. partition	tight frame	over-sampling	localization	ori/steer	shift inv.	perfect recon.
local freq	uniform	Y	large	fair	possible	possible	possible
block DCT	uniform	Y	1	fair	N	N	Y
Gabor (uniform)	uniform	N	1	p:good b:poor	Y/N	N	nearly
lapl. pyr.	octave	N	4/3	good	N/Y	possible	Y
orthog wavelet	octave	Y	1	possible	Y/N	N	Y
haar	octave	Y	1	space: good freq: poor	N	N	Y
steer pyr / curvelets	octave	Y	4K/3	good	Y	Y	nearly