

6.869 Advances in Computer Vision:  
Learning and Interfaces

Spring 2005

Tuesday and Thursday, 2:30 to 4:00pm in 36-153

Announcements

Course Information

- Syllabus
- Problem Sets and Exams
- Grading and Requirements
- Internet Resources

Contacts <http://courses.csail.mit.edu/6.869>

Course Calendar

Lecture	Date	Description	Readings	Assignments	Materials
1	2/1	Course Introduction Cameras and Lenses	Req: FP 1.1, 2.1, 2.2, 2.3, 3.1, 3.2	FS0 out	
2	2/3	Image Filtering	Req: FP 7.1 - 7.6		
3	2/8	Image Representations: Pyramids	Req: FP 7.7, 9.2		
4	2/10	Image Statistics		FS0 due	
5	2/15	Texture	Req: FP 9.1, 9.3, 9.4	FS1 out	
6	2/17	Color	Req: FP 6.1-6.4		
7	2/22	Guest Lecture: Context in vision			
8	2/24	Guest Lecture: Medical Imaging		FS1 due	
9	3/1	Multiview Geometry	Req: Mikolajczyk and Schmid, FP 10	FS2 out	
10	3/3	Local Features	Req: Shi and Tomasi; Lowe		

Course Calendar


Lecture	Date	Description	Readings	Assignments	Materials
2	2/3	Image Filtering	Req: FP 7.1 - 7.6		
3	2/8	Image Representations: Pyramids	Req: FP 7.7, 9.2		

Today


## Reading

- Related to today's lecture:
  - Chapters 7.7, 9.2, Forsyth&Ponce..
  - Adelson article on pyramid representations, posted on web site.

## Spatial resolution and color




original




R  
G  
B


## Blurring the G component



original



processed



R  
G  
B




## Blurring the b Lab component



## Application to image compression

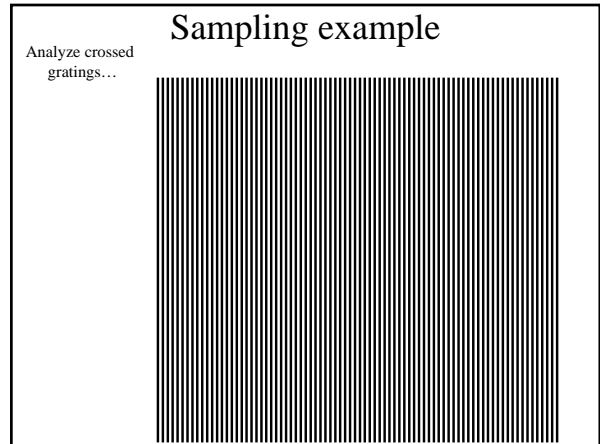
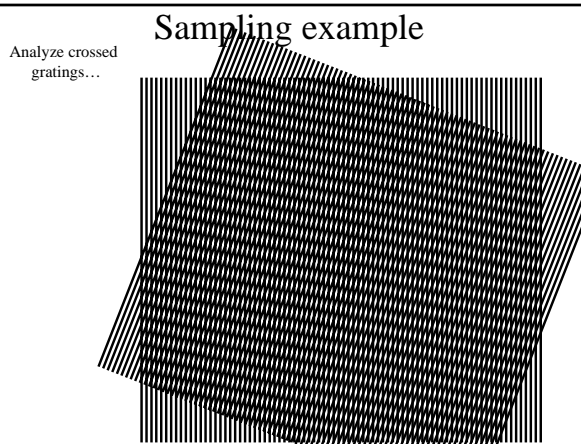
- (compression is about hiding differences from the true image where you can't see them).

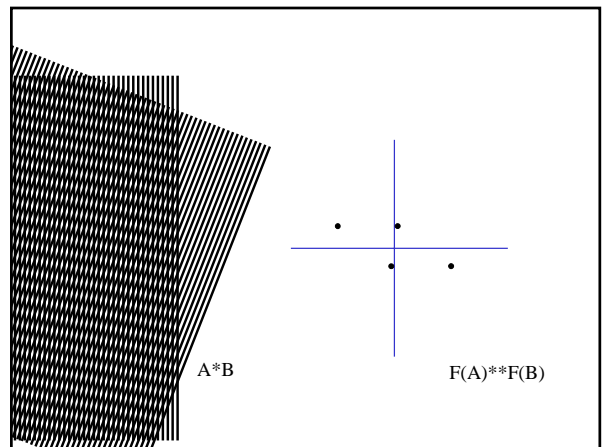


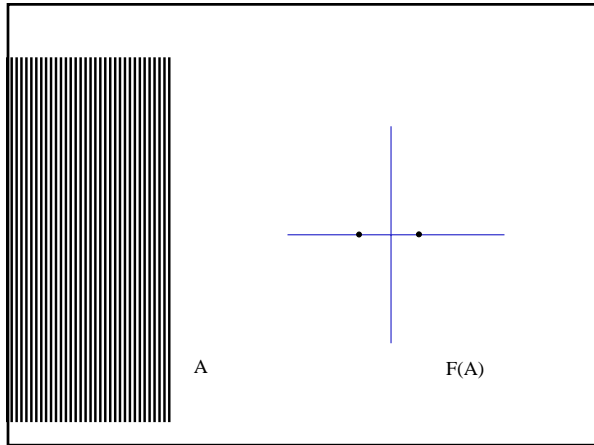


## Bandwidth (transmission resources) for the components of the television signal




Understanding image perception allowed NTSC to add color to the black and white television signal (with some, but limited, incompatibility artifacts).

## Sampling and aliasing








## Scaled representations

- Big bars (resp. spots, hands, etc.) and little bars are both interesting
  - Stripes and hairs, say
- Inefficient to detect big bars with big filters
  - And there is superfluous detail in the filter kernel
- Alternative:
  - Apply filters of fixed size to images of different sizes
  - Typically, a collection of images whose edge length changes by a factor of 2 (or root 2)
  - This is a pyramid (or Gaussian pyramid) by visual analogy

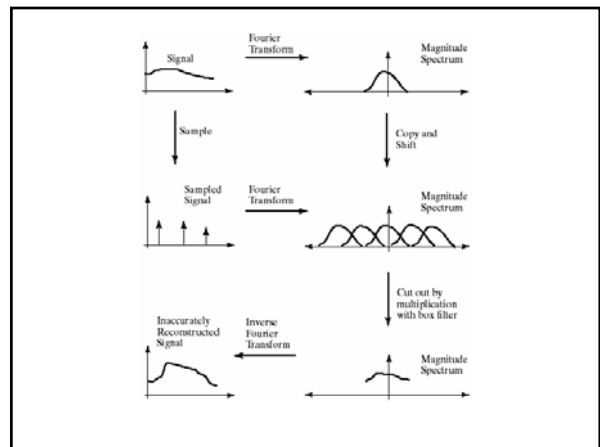

## Example application: CMU face detector



From: <http://www.ius.cs.cmu.edu/IUS/har2/har/www/CMU-CS-95-158R/>

## Aliasing


- Can't shrink an image by taking every second pixel
- If we do, characteristic errors appear
  - In the next few slides
  - Typically, small phenomena look bigger; fast phenomena can look slower
  - Common phenomenon
    - Wagon wheels rolling the wrong way in movies
    - Checkerboards misrepresented in ray tracing
    - Striped shirts look funny on colour television




## Smoothing as low-pass filtering

- The message of the FT is that high frequencies lead to trouble with sampling.
- Solution: suppress high frequencies before sampling
  - multiply the FT of the signal with something that suppresses high frequencies
  - or convolve with a low-pass filter
- A filter whose FT is a box is bad, because the filter kernel has infinite support
- Common solution: use a Gaussian
  - multiplying FT by Gaussian is equivalent to convolving image with Gaussian.


Sampling without smoothing. Top row shows the images, sampled at every second pixel to get the next; bottom row shows the magnitude spectrum of these images.



Sampling with smoothing. Top row shows the images. We get the next image by smoothing the image with a Gaussian with sigma 1 pixel, then sampling at every second pixel to get the next; bottom row shows the magnitude spectrum of these images.



Sampling with smoothing. Top row shows the images. We get the next image by smoothing the image with a Gaussian with sigma 1.4 pixels, then sampling at every second pixel to get the next; bottom row shows the magnitude spectrum of these images.



## Matlab

Subsample image in matlab.

## The Gaussian pyramid

- Smooth with gaussians, because
  - a gaussian\*gaussian=another gaussian
- Synthesis
  - smooth and sample
- Analysis
  - take the top image
- Gaussians are low pass filters, so repn is redundant

### Convolution and subsampling as a matrix multiply (1-d case)

U1 =

```

1 4 6 4 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 1 4 6 4 1 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 1 4 6 4 1 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 1 4 6 4 1 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 1 4 6 4 1 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 1 4 6 4 1 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 1 4 6 4 1 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 4 6 4 1 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 4 6 4 1
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 4 6 4 1
    
```

### Next pyramid level

U2 =

```

1 4 6 4 1 0 0 0
0 0 1 4 6 4 1 0
0 0 0 0 1 4 6 4
0 0 0 0 0 0 1 4
    
```

### b \* a, the combined effect of the two pyramid levels

>> U2 \* U1

ans =

```

1 4 10 20 31 40 44 40 31 20 10 4 1 0 0 0 0 0 0 0
0 0 0 0 1 4 10 20 31 40 44 40 31 20 10 4 1 0 0 0
0 0 0 0 0 0 0 0 1 4 10 20 31 40 44 40 30 16 4 0
0 0 0 0 0 0 0 0 0 0 0 0 1 4 10 20 25 16 4 0
    
```

### The computational advantage of pyramids




Fig. 1. A one-dimensional graphic representation of the process which generates a Gaussian pyramid. Each row of dots represents nodes within a level of the pyramid. The value of each node in the zero level is just the gray level of a corresponding image pixel. The value of each node in a high level is the weighted average of node values in the next lower level. Note that node spacing doubles from level to level, while the same weighting pattern or "generating kernel" is used to generate all levels.




Fig. 2. The equivalent weighting functions  $h_i(x)$  for nodes in levels 1, 2, 3, and infinity of the Gaussian pyramid. Note that axis scales have been adjusted by factors of 2 to aid comparison. Here the parameter  $\alpha$  of the generating kernel is 0.4, and the resulting equivalent weighting functions closely resemble the Gaussian probability density functions.

[http://www-bcs.mit.edu/people/adelson/pub\\_pdfs/pyramid83.pdf](http://www-bcs.mit.edu/people/adelson/pub_pdfs/pyramid83.pdf) IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-31, NO. 4, APRIL 1983





Fig. 4. First six levels of the Gaussian pyramid for the "Lady" image. The original image, level 0, measures 257 by 257 pixels and each higher level array is roughly half the dimensions of its predecessor. Thus, level 5 measures just 9 by 9 pixels.

[http://www-bcs.mit.edu/people/adelson/pub\\_pdfs/pyramid83.pdf](http://www-bcs.mit.edu/people/adelson/pub_pdfs/pyramid83.pdf) IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-31, NO. 4, APRIL 1983



## Image pyramids

- Gaussian
- Laplacian
- Wavelet/QMF
- Steerable pyramid

## Image pyramids

- Gaussian
- Laplacian
- Wavelet/QMF
- Steerable pyramid

## The Laplacian Pyramid

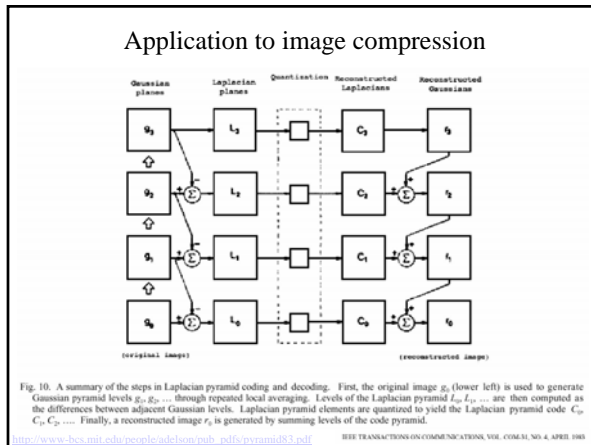


- Synthesis
  - preserve difference between upsampled Gaussian pyramid level and Gaussian pyramid level
  - band pass filter - each level represents spatial frequencies (largely) unrepresented at other levels
- Analysis
  - reconstruct Gaussian pyramid, take top layer

## Laplacian pyramid algorithm

Fig. 5. First four levels of the Gaussian and Laplacian pyramids. Gaussian images, upper row, were obtained by expanding pyramid series of Fig. 4c through Gaussian interpolation. Each level of the Laplacian pyramid is the difference between the corresponding and next higher levels of the Gaussian pyramid.

[http://www.bcs.mit.edu/people/edelson/pub\\_pdfs/pyramid83.pdf](http://www.bcs.mit.edu/people/edelson/pub_pdfs/pyramid83.pdf)
IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-31, NO. 4, APRIL 1983





Matlab manipulations with gaussian and laplacian pyramids

### Image pyramids

- Gaussian
- Laplacian
- Wavelet/QMF
- Steerable pyramid

### What is a good representation for image analysis?

(Goldilocks and the three representations)

- Fourier transform domain tells you “what” (textural properties), but not “where”. In space, this representation is too spread out.
- Pixel domain representation tells you “where” (pixel location), but not “what”. In space, this representation is too localized
- Want an image representation that gives you a local description of image events—what is happening where. That representation might be “just right”.

## Wavelets/QMF's

$$\vec{F} = U\vec{f}$$

transformed image  $\vec{F}$  ← Vectorized image  $\vec{f}$

Fourier transform, or  
Wavelet transform, or  
Steerable pyramid transform

U =

```
1  1
1 -1
```

```
>> inv(U)
```

ans =

```
0.5000  0.5000
0.5000 -0.5000
```

U =

```
1  1  0  0  0  0  0  0
1 -1  0  0  0  0  0  0
0  0  1  1  0  0  0  0
0  0  1 -1  0  0  0  0
0  0  0  0  1  1  0  0
0  0  0  0  1 -1  0  0
0  0  0  0  0  0  1  1
0  0  0  0  0  0  1 -1
```

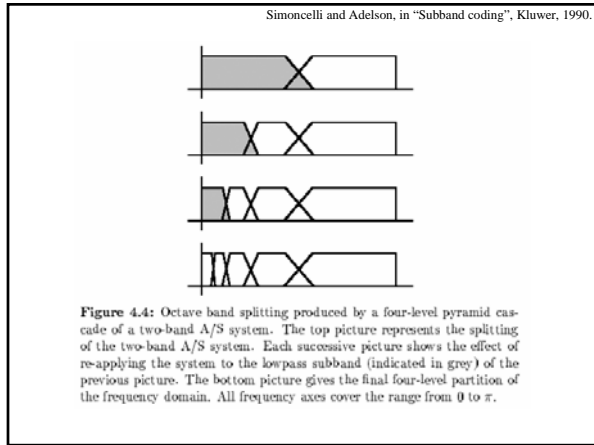
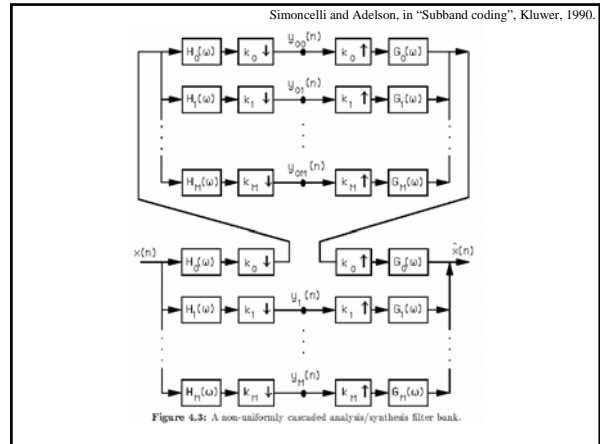
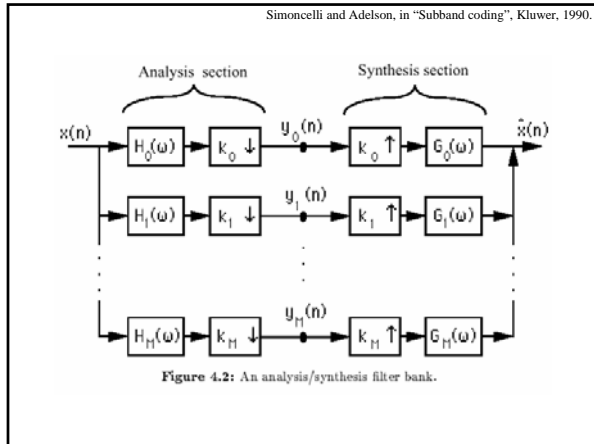
```
>> inv(U)
```

ans =

```
0.5000  0.5000  0  0  0  0  0  0
0.5000 -0.5000  0  0  0  0  0  0
0  0  0.5000  0.5000  0  0  0  0
0  0  0.5000 -0.5000  0  0  0  0
0  0  0  0  0.5000  0.5000  0  0
0  0  0  0  0.5000 -0.5000  0  0
0  0  0  0  0  0  0.5000  0.5000
0  0  0  0  0  0  0.5000 -0.5000
```

## Matlab examples of Haar wavelet representation



- Frequency characteristics of the high and low-pass representations



Simoncelli and Adelson, in "Subband coding", Kluwer, 1990.

n	QMF-5	QMF-9	QMF-13
0	0.8593118	0.7973934	0.7737113
1	0.3535534	0.41472545	0.42995453
2	-0.0761025	-0.073386621	-0.057827797
3		-0.060944743	-0.09800052
4		0.02807382	0.039045125
5			0.021651438
6			-0.014556438

Table 4.1: Odd-length QMF kernels. Half of the impulse response sample values are shown for each of the normalized lowpass QMF filters (All filters are symmetric about  $n = 0$ ). The appropriate highpass filters are obtained by delaying by one sample and multiplying with the sequence  $(-1)^n$ .



## Good and bad features of wavelet/QMF filters

- Bad:
  - Aliased subbands
  - Non-oriented diagonal subband
- Good:
  - Not overcomplete (so same number of coefficients as image pixels).
  - Good for image compression (JPEG 2000)

## Image pyramids


- Gaussian
- Laplacian
- Wavelet/QMF
- Steerable pyramid

## Steerable pyramids

- Good:
  - Oriented subbands
  - Non-aliased subbands
  - Steerable filters
- Bad:
  - Overcomplete
  - Have one high frequency residual subband, required in order to form a circular region of analysis in frequency from a square region of support in frequency.

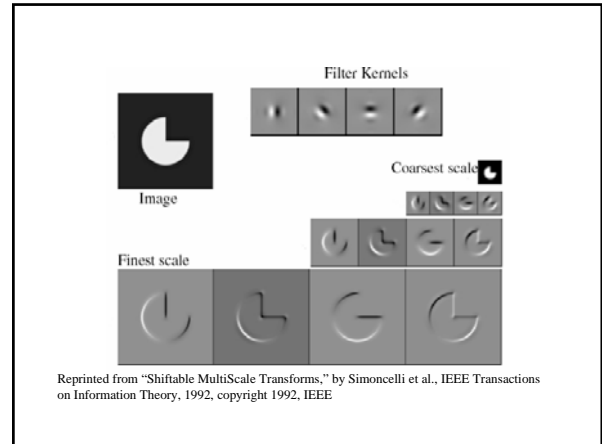

## Oriented pyramids

- Laplacian pyramid is orientation independent
- Apply an oriented filter to determine orientations at each layer
  - by clever filter design, we can simplify synthesis
  - this represents image information at a particular scale and orientation



	Laplacian Pyramid	Dyadic QMF/Wavelet	Steerable Pyramid
self-inverting (tight frame)	no	yes	yes
overcompleteness	4/3	1	4k/3
aliasing in subbands	perhaps	yes	no
rotated orientation bands	no	only on hex lattice [9]	yes

Table 1: Properties of the Steerable Pyramid relative to two other well-known multi-scale representations.



### Matlab resources for pyramids (with tutorial)

<http://www.cns.nyu.edu/~eero/software.html>

**Eero P. Simoncelli**  
Associate Investigator,  
Howard Hughes Medical Institute  
Associate Professor,  
Neural Science and Mathematics,  
New York University


### Matlab resources for pyramids (with tutorial)


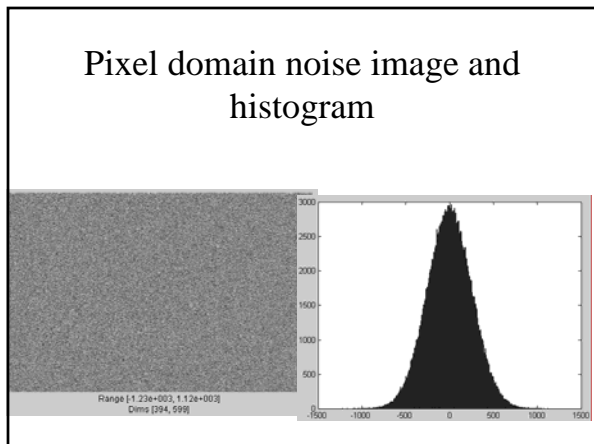
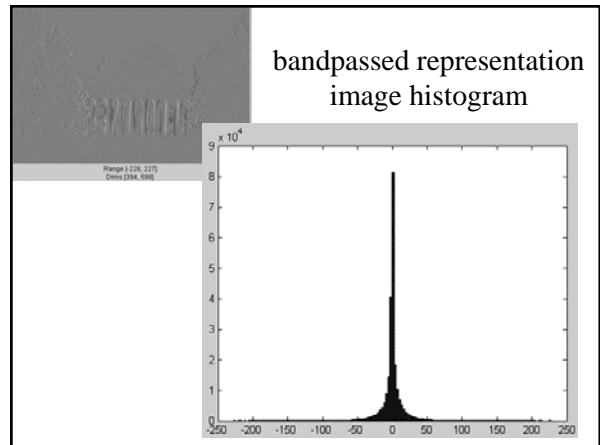

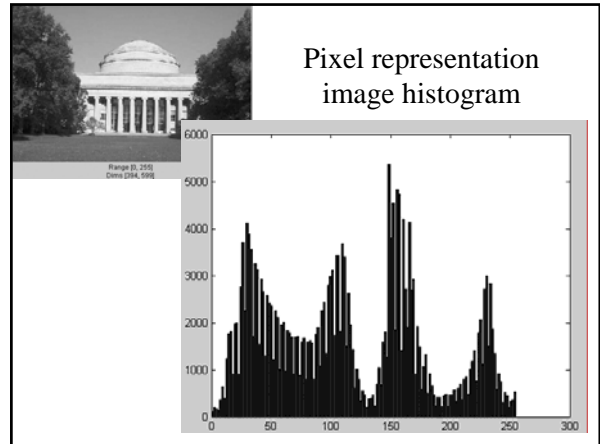

<http://www.cns.nyu.edu/~eero/software.html>

**Publicly Available Software Packages**

- **Texture Analysis/Synthesis** - Matlab code is available for analyzing and synthesizing visual textures. [README](#) | [Contents](#) | [ChangeLog](#) | [Source code](#) (UNIX/PC; gnu/par file)
- **EPWIC** - Embedded Progressive Wavelet Image Coder. C source code available.
- **matlabPyTools** - Matlab source code for multi-scale image processing. Includes tools for building and manipulating Laplacian pyramids, GMF/Wavelets, and steerable pyramids. Data structures are compatible with the Matlab wavelet toolbox, but the convolution code (in C) is faster and has many boundary-handling options. [README](#), [Contents](#), [Modification list](#), [UNIX/PC source](#) or [Macintosh source](#).
- **The Steerable Pyramid**, an (approximately) translation- and rotation-invariant multi-scale image decomposition. Matlab (see above) and C implementations are available.
- **Computational Models of cortical neurons**. Macintosh program available.
- **EPIC** - Efficient Pyramid (Wavelet) Image Coder. C source code available.
- **OBVIOUS** (Object-Based Vision & Image Understanding System). [README](#) / [ChangeLog](#) / [Disc CDROM](#) / [Source Code \(2.5MB\)](#).
- **CL-SHELL** (GNU Emacs <-> Common Lisp Interface). [README](#) / [Change Log](#) / [Source Code \(1.19B\)](#).

An application of image pyramids:  
noise removal











## CCD color sampling

## Color sensing, 3 approaches

- Scan 3 times (temporal multiplexing)
- Use 3 detectors (3-ccd camera, and color film)
- Use offset color samples (spatial multiplexing)

## Typical errors in temporal multiplexing approach


- Color offset fringes




## Typical errors in spatial multiplexing approach.

- Color fringes.

## CCD color filter pattern



## The cause of color moire




Fine black and white detail in image mis-interpreted as color information.

## Black and white edge falling on color CCD detector


Black and white image (edge)

Detector pixel colors



## Color sampling artifact

Interpolated pixel colors, for grey edge falling on colored detectors (linear interpolation).



## Typical color moire patterns

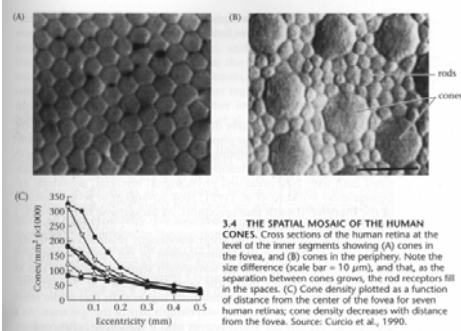


Blow-up of electronic camera image. Notice spurious colors in the regions of fine detail in the plants.

## Color sampling artifacts




## Human Photoreceptors



(From Foundations of Vision, by Brian Wandell, Sinauer Assoc.)

## Brewster's colors example (subtle).


Scale relative to human photoreceptor size: each line covers about 7 photoreceptors.




## Median Filter Interpolation

- Perform first interpolation on isolated color channels.
- Compute color difference signals.
- Median filter the color difference signal.
- Reconstruct the 3-color image.

## Two-color sampling of BW edge



## R-G, after linear interpolation



## R - G, median filtered (5x5)



## Recombining the median filtered colors

Linear interpolation

Median filter interpolation







## Didn't get a chance to show:

Local gain control.

- Summary of pyramid representations

### Image pyramids

- Gaussian  Progressively blurred and subsampled versions of the image. Adds scale invariance to fixed-size algorithms.
- Laplacian  Shows the information added in Gaussian pyramid at each spatial scale. Useful for noise reduction & coding.
- Wavelet/QMF  Bandpassed representation, complete, but with aliasing and some non-oriented subbands.
- Steerable pyramid  Shows components at each scale and orientation separately. Non-aliased subbands. Good for texture and feature analysis.

### Linear image transformations

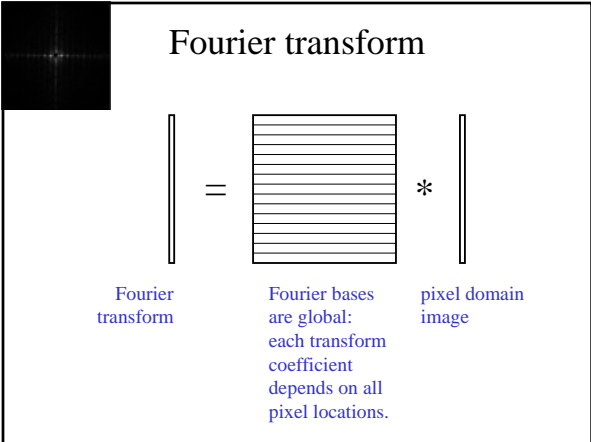
- In analyzing images, it's often useful to make a change of basis.

transformed image  $\vec{F} = U\vec{f}$  Vectorized image

Fourier transform, or Wavelet transform, or Steerable pyramid transform

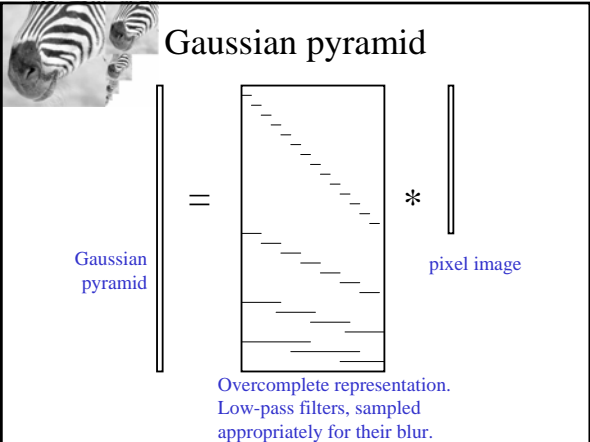
- ### Schematic pictures of each matrix transform
- Shown for 1-d images
  - The matrices for 2-d images are the same idea, but more complicated, to account for vertical, as well as horizontal, neighbor relationships.

### Fourier transform

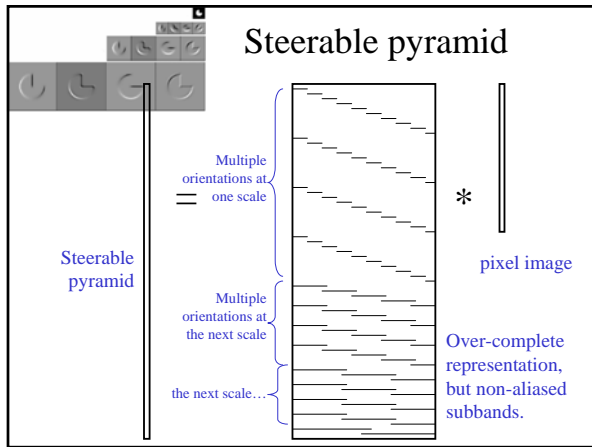
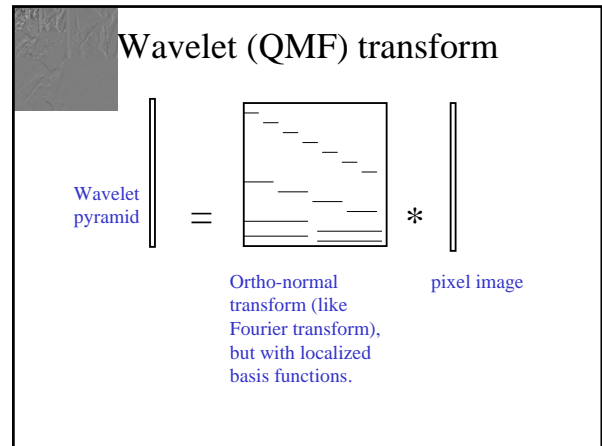



Fourier transform = Fourier bases are global: each transform coefficient depends on all pixel locations. \* pixel domain image

### Gaussian pyramid



Gaussian pyramid = Overcomplete representation. Low-pass filters, sampled appropriately for their blur. \* pixel image



### Matlab resources for pyramids (with tutorial)

<http://www.cns.nyu.edu/~eero/software.html>

**Publicly Available Software Packages**

- **Texture Analysis/Synthesis** - Matlab code is available for analyzing and synthesizing visual textures. [README](#) | [Contents](#) | [ChangeLog](#) | [Source code](#) (UNIX/PC, gzip/tar file)
- **EPWIC** - Embedded Progressive Wavelet Image Coder. C source code available.
- **matlabPyTools** - Matlab source code for multi-scale image processing. Includes tools for building and manipulating Laplacian pyramids, QMF/Wavelets, and steerable pyramids. Data structures are compatible with the Matlab wavelet toolbox, but the convolution code (in C) is faster and has many boundary-handling options. [README](#), [Contents](#), [Modification list](#), [UNIX/PC source](#) or [Macintosh source](#).
- **The Steerable Pyramid**, an (approximately) translation- and rotation-invariant multi-scale image decomposition. Matlab (see above) and C implementations are available.
- **Computational Models of cortical neurons**. Macintosh program available.
- **EPIC** - Efficient Pyramid (Wavelet) Image Coder. C source code available.
- **OBVIOUS** (Object-Based Vision & Image Understanding System). [README](#) / [ChangeLog](#) / [Disc CDROM](#) / [Source Code CD-ROM](#).
- **CL-SHELL** (Gnu Emacs <-> Common Lisp Interface). [README](#) / [Change Log](#) / [Source Code](#) (1.1.9b).

- ### Why use these representations?
- Handle real-world size variations with a constant-size vision algorithm.
  - Remove noise
  - Analyze texture
  - Recognize objects
  - Label image features

end