6.801/6.866 Machine Vision

Syllabus

	#	Date	Description	Readings	Assignments	Materials
	1	9/5	Course Introduction		Pset #0 (not collected)	Freeman Slides Darrell Slides Matlab Tutorial Diary
	2	9/10	Cameras, Lenses, and Sensors	Req: FP 1 Opt: H 2.1, 2.3		Freeman Slides
	3	9/12	Radiometry and Shading Models I	Req: FP 2, 5.4; H 10 Opt: FP 4	Pset #1 Assigned	Freeman Slides
[4	9/17	Radiometry and Shading Models ${\rm I\!I}$	"		Freeman Slides
	5	9/19	Multiview Geometry	Req: FP 10		Darrell Slides
	б	9/24	Stereo	Req: FP 11; H 13	Pset #1 Due	Darrell Slides
	7	9/26	Color	Req: FP 6.1-6.4	Pset #2 Assigned	Freeman Slides
	8	10/1	Shape from Shading	Req: H 11.1, 11.5-11.9 Opt: H 11.2-11.4		
•	9	10/3	Filtering I			
	10	10/8	Filtering II		Pset #2 Due	
	11	10/10	Intro to Bayesian Vision		Exam #1 Assigned	
		10/15	Columbus Day (NO LECTURE)			

Homeworks

- Problem set 1 returned today. Mean: 85. Std dev: 15
- Problem set 1 solutions posted (with password protection) by Tuesday
- Note revisions to problem set 2 matlab assignment
 - To clarify the problems and make them work out cleaner for you.
 - Modifying an existing solution to the revised version will be very simple.

Shape from shading

A shape-from-shading algorithm

$$s_{i,j} = \frac{1}{4} \left(\left(p_{i+1,j} - p_{i,j} \right)^{2} + \left(p_{i,j+1} - p_{i,j} \right)^{2} + \left(q_{i+1,j} - q_{i,j} \right)^{2} + \left(q_{i,j+1} - q_{i,j} \right)^{2} \right) \text{ surface smoothness }$$

$$r_{ij} = \left(E_{ij} - R_{s} \left(p_{ij}, q_{ij} \right) \right)^{2} \qquad \text{ local fidelity to image data } \text{ smoothness }$$

$$e = \sum_{i} \sum_{j} \left(s_{ij} + \lambda r_{ij} \right) \qquad \text{global function to minimize}$$

$$\frac{\partial e}{\partial p_{kl}} = 2(q_{kl} - \overline{q}_{kl}) - 2\lambda(E_{kl} - R_{s}(p_{kl}, q_{kl})) \frac{\partial R_{s}}{\partial q} \qquad \text{derivative w.r.t. surface parameters }$$

$$\frac{\partial e}{\partial q_{kl}} = 2(q_{kl} - \overline{q}_{kl}) - 2\lambda(E_{kl} - R_{s}(p_{kl}, q_{kl})) \frac{\partial R_{s}}{\partial q} \qquad \text{derivative w.r.t. surface parameters }$$

$$p_{kl} = \overline{p}_{kl} + \lambda(E_{kl} - R_{s}(p_{kl}, q_{kl})) \frac{\partial R_{s}}{\partial q} \qquad \text{set derivatives equal to zero.}$$

$$p_{kl}^{n+1} = \overline{p}_{kl}^{n} + \lambda\left(E_{kl} - R_{s}\left(p_{kl}^{n}, q_{kl}^{n} \right) \right) \frac{\partial R_{s}}{\partial p}$$

$$q_{kl}^{n+1} = \overline{q}_{kl}^{n} + \lambda\left(E_{kl} - R_{s}\left(p_{kl}^{n}, q_{kl}^{n} \right) \right) \frac{\partial R_{s}}{\partial q} \qquad \text{agorithm}$$



Results using this approach

Figure 11-10. Display of the image of a small resin droplet on a flower of a Cannabis sativa plant. (Reproduced by permission from the book Magnifications—Photography with the Scanning Electron Microscope by David Scharf, Schocken Books, New York, 1977.)

е	bo	0	k		М	la	qı	<i>i</i> i	-	•	•	•		•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	÷	•	
	D				0	1		c		٠	•	٠	•			•	•	•	1		•	1	٠.	•	•		•	•	1	٠.	•	•	•	•	•	•	•	•	•	•	•	•	۰.	•	
	Da	VI	ld		S	:h	a	rt.	,										١	٩.		1							¢	1	۰.	•	•	•	•	•	٠.	•	•	•	•	•	۰.	٠	
																		١	١	١	١	1	۴.		٠	1	1	1	٠	,	1	•	•	•	•	•	•	•	•	•	•	•	÷	•	
	•			i.		•	•									١	1	١	۲	١	٩		1		٠		1	1	1	,	-	-	•	•	•	•	•	•	•	•	•	•	٠	٠	
				1											1			•	1	\$				٠	٠			1	1		1	-	•	•	•	•	•	•	•	•	•	•	۰.	•	
				21										1	*	1											1	1	1	1			1	•	•	•	•	•	•					•	
	1													~													1	1			-	-	-	•	•								÷.,		
	6.3					1							~	-	-	-														-		-	-												
	12	1			20	3		0	13	0	10		-																-	-		-	-												
		1		8	1	1	1	1	13	1	- 0		1	2	1	2	1							1			-	-	-	-	1	_	_										ι.		
							6	Ċ	80					1	-	1	1	3	2	1			5	0	2	1	2	0	2	12			_												
				1		٠	• *							1						1		1					17	0	17	5		-	7	÷.,		8									
	•	•		•	•	•	•	•			`	`	•		-	٠	•	•	•		٠	•	٠	•	•	•	2	•	•	~	*	~	~		1	•	۰.	•	1	•	•	•		•	
	•	•		•	•	•	•			1	`	•	•	-	٠	٠	-	-	٠		٠	•		•	•	•	•	•	•	*	•	•	٠.	•	٠.	•	•	•	•	•	•	•	۰.	•	
	•	•	•	•	•	•	•	1	•	•	•	•	-	•	•				٠			•				•	•	•	*	*	*	1	•	٠.	1	•	•	۰.	•	•	•	•	•	•	
	•	•		•		٠	1	•	•	•	•	•											4		1		1		•	*	1		٠	•	•	•	•	•	•	•	•	•	•	۰.	
	•			•		-	-	-	-										. 1				٩		1		•		~		٠	•	•	•	•	•	•	•	•	•	•	•	•	1	
				•		•	1	-				,	,								1			\$	1		1		-	-		•	•	•	٠	٠	۰.	٠	٠	•		•		•	
				•			٠,					,		,		,	,				1				1		1	1	1		•	•	•	٠	•	•	٠.	٠	•	•	•	٠	٠	•	
										٠.	٠.	٠.	۰.			,	,		1.		1						1	1				•	•	•	٠		٠	•	•	•	•	•	•	•	
									1.			۰.		٠.	٢.	0	1		1	1		2	÷.	1	i	1			٠.														۰.	•	
	100					92	12			۰.		87.				٩.	٩.	1	1	1	2	2	2	2		2																		γ.	
	313	1			22	2	0			22					13	10				1	3		10																						
						1	1								1	1	10																												
	1	1					0	1	110			0								1				1	0		10					1													
			0.0	•		•															•		•		- 1				121	10.0			100	100					1			1.70	150		

Figure 11-11. Needle diagram calculated by the iterative scheme under the assumption that the reflectance map is $\sec \theta_i$. (The surface orientation data are actually available on a finer grid; they are sampled coarsely here for display purposes.) The needle diagram is the estimate of the shape of the surface of the resin droplet shown in the previous figure. (Figure kindly provided by Katsushi Ikeuchi.)

Horn, 1986

2-d animation



Copyrighted ... yada yada yada

E

3-d animation



Image filtering

- Reading:
 - Chapter 7, F&P
- Recommended Reading:
 - Chapter 7, 8 Horn

Oct. 3, 2002 MIT 6.801/6.866 Profs. Freeman and Darrell

Take 6.341, discrete-time signal processing

- If you want to process pixels, you need to understand signal processing well, so
 – Take 6.341
- Fantastic set of teachers:
 - Al Oppenheim
 - Greg Wornell
 - Jae Lim
- Web page: http://web.mit.edu/6.341/www/

What is image filtering?

• Modify the pixels in an image based on some function of a local neighborhood of the pixels.



Local image data

Modified image data 10

Linear functions

- Simplest: linear filtering.
 - Replace each pixel by a linear combination of its neighbors.
- The prescription for the linear combination is called the "convolution kernel".





Local image data

kernel

7

Modified image data 11

Convolution

$$f[m,n] = I \otimes g = \sum_{k,l} I[m-k,n-l]g[k,l]$$

Linear filtering (warm-up slide)



original



?

Linear filtering (warm-up slide)



original





Filtered (no change)

Linear filtering



original





shift



1.0 Definition 1.0 Definition 1.0



shifted

original

Linear filtering



original



?

Blurring

coefficient

0.3

Pixel offset



original



Blurred (filter applied in both dimensions).





Linear filtering (warm-up slide)





original

Linear filtering (no change)



original



6

Filtered (no change)

Linear filtering





original

(remember blurring)



original





Blurred (filter applied in both dimensions).

Sharpening



original





original

Sharpening example



Sharpening





before

after

Oriented filters



Gabor filters at different scales and spatial frequencies

top row shows anti-symmetric (or odd) filters, bottom row the symmetric (or even) filters.



Linear image transformations

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



Self-inverting transforms

Same basis functions are used for the inverse transform

$$\vec{f} = U^{-1}\vec{F}$$
$$= U^{+}\vec{F}$$

U transpose and complex conjugate

An example of such a transform: the Fourier transform

discrete domain

Forward transform

$$F[m,n] = \sum_{k=0}^{M-1} \sum_{l=0}^{N-1} f[k,l] e^{-\pi i \left(\frac{km}{M} + \frac{\ln}{N}\right)}$$

Inverse transform

$$f[k,l] = \frac{1}{MN} \sum_{k=0}^{M-1} \sum_{l=0}^{N-1} F[m,n] e^{+\pi i \left(\frac{km}{M} + \frac{\ln}{N}\right)}$$

To get some sense of what basis elements look like, we plot a basis element --- or rather, its real part --as a function of x,y for some fixed u, v. We get a function that is constant when (ux+vy) is constant. The magnitude of the vector (u, v) gives a frequency, and its direction gives an orientation. The function is a sinusoid with this frequency along the direction, and constant perpendicular to the direction.



Here u and v are larger than in the previous slide.



And larger still...



Phase and Magnitude

- Fourier transform of a real function is complex
 - difficult to plot, visualize
 - instead, we can think of the phase and magnitude of the transform
- Phase is the phase of the complex transform
- Magnitude is the magnitude of the complex transform

- Curious fact
 - all natural images have about the same magnitude transform
 - hence, phase seems to matter, but magnitude largely doesn't
- Demonstration
 - Take two pictures, swap the phase transforms, compute the inverse what does the result look like?


This is the magnitude transform of the cheetah pic



This is the phase transform of the cheetah pic



39



This is the magnitude transform of the zebra pic



This is the phase transform of the zebra pic



Reconstruction with zebra phase, cheetah magnitude



Reconstruction with cheetah phase, zebra magnitude



Example image synthesis with fourier basis.

• 16 images







18

#1: Range [0, 1] Dims [256, 256]



#2: Range [4.79e-007, 0.503] Dims [256, 256]



#1: Range [0, 1] Dims [256, 256]



#2: Range [8.5e-006, 1.7] Dims [256, 256]

#1: Range [0, 1] Dims [256, 256]



#2: Range [3.85e-007, 2.21] Dims [256, 256]

136

#1: Range [0, 1] Dims [256, 256]



#2: Range [8.25e-006, 3.48] Dims [256, 256]

282



#1: Range [0, 1] Dims [256, 256]



#2: Range [1.39e-005, 5.88] Dims [256, 256]

538







#2: Range [6.17e-006, 8.4] Dims [256, 256]

1088



#1: Range [0, 1] Dims [256, 256]



#2: Range [9.99e-005, 15] Dims [256, 256]

2094



#1: Range [0, 1] Dims [256, 256]



#2: Range [8.7e-005, 19] Dims [256, 256]

4052.

4052



#1: Range [0, 1] Dims [256, 256]



#2: Range [0.000556, 37.7] Dims [256, 256]

8056.



#1: Range [0, 1] Dims [256, 256]



#2: Range (0.00032, 64.5) Dims (256, 256)

15366



#1: Range [0, 1] Dims [256, 256]



#2: Range [0.000231, 91.1] Dims [256, 256]

28743



#1: Range [0, 1] Dims [256, 256]



#2: Range (0.00109, 146) Dims (256, 256)

49190.

49190



#1: Range [0, 1] Dims [256, 256]



#2: Range [0.00758, 294] Dims [256, 256]

65536.



#1: Range [0.5, 1.5] Dims [256, 256]

Fourier transform magnitude



Dims [256, 256]



Masking out the fundamental and harmonics from periodic pillars





Range [0.000551, 297] Dims [256, 256]

Name as many functions as you can that retain that same functional form in the transform domain

Linear Filters Chap. 7

TABLE 7.1 A variety of functions of two dimensions and their Fourier transforms. This table can be used in two directions (with appropriate substitutions for u, v and x, y) because the Fourier transform of the Fourier transform of a function is the function. Observant readers may suspect that the results on infinite sums of δ functions contradict the linearity of Fourier transforms. By careful inspection of limits, it is possible to show that they do not (see, e.g., Bracewell, 1995). Observant readers may also have noted that an expression for $\mathcal{F}(\frac{\partial f}{\partial y})$ can be obtained by combining two lines of this table.

Function	Fourier transform
g(x, y)	$\iint_{-\infty}^{\infty} g(x, y) e^{-i2\pi(ux+vy)} dx dy$
$\iint_{-\infty}^{\infty} \mathcal{F}(g(x, y))(u, v) e^{i2\pi(ux+vy)} du dv$	$\mathcal{F}(g(x, y))(u, v)$
$\delta(x, y)$	1
$\frac{\partial f}{\partial x}(x, y)$	$u\mathcal{F}(f)(u,v)$
$0.5\delta(x + a, y) + 0.5\delta(x - a, y)$	$\cos 2\pi a u$
$e^{-\pi(x^2+y^2)}$	$e^{-\pi(u^2+v^2)}$
$box_1(x, y)$	$\frac{\sin u}{u} \frac{\sin v}{v}$
f(ax, by)	$\frac{\mathcal{F}(f)(u/a,v/b)}{ab}$
$\sum_{i=-\infty}^{\infty} \sum_{i=-\infty}^{\infty} \delta(x-i, y-j)$	$\sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} \delta(u-i, v-j)$
(f * *g)(x, y)	$\mathcal{F}(f)\mathcal{F}(g)(u,v)$
f(x-a, y-b)	$e^{-i2\pi(au+bv)}\mathcal{F}(f)$
$f(x\cos\theta - y\sin\theta, x\sin\theta + y\cos\theta)$	$\mathcal{F}(f)(u\cos\theta - v\sin\theta, u\sin\theta + v\cos\theta)$

Forsyth&Ponce

Discrete-time, continuous frequency Fourier transform

Many sequences can be represented by a Fourier integral of the form

$$x[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} X(e^{j\omega}) e^{j\omega n} d\omega, \qquad (2.133)$$

where

$$X(e^{j\omega}) = \sum_{n=-\infty}^{\infty} x[n]e^{-j\omega n}.$$
(2.134)

Öppenheim, Schafer and Buck, Discrete-time signal processing, Prentice Hall, 1999

Discrete-time, continuous frequency Fourier transform pairs

FOURIER TRANSFORM PAIRS TABLE 2.3 Sequence Fourier Transform 1. $\delta[n]$ $-j\omega n_0$ 2. $\delta[n - n_0]$ 3. 1 $(-\infty < n < \infty)$ $\sum_{k=1}^{\infty} 2\pi \delta(\omega + 2\pi k)$ $\frac{1}{1-ae^{-j\omega}}$ 4. $a^n u[n]$ (|a| < 1) $\frac{1}{1-e^{-j\omega}}+\sum \pi\delta(\omega+2\pi k)$ 5. u[n]6. $(n+1)a^n u[n]$ (|a| < 1) $\overline{(1-ae^{-j\omega})^2}$ of obviors we call 7. $\frac{r^n \sin \omega_p(n+1)}{\sin \omega_p} u[n]$ (|r| < 1) $\frac{1}{1 - 2r \cos \omega_p e^{-j\omega} + r^2 e^{-j2\omega}}$ $X(e^{j\omega}) = \begin{cases} 1, & |\omega| < \omega_c, \\ 0, & \omega_c < |\omega| \le \pi \end{cases}$ 8. $\frac{\sin \omega_c n}{\pi n}$ 9. $x[n] = \begin{cases} 1, & 0 \le n \le M \\ 0, & \text{otherwise} \end{cases}$ $\frac{\sin[\omega(M+1)/2]}{\sin(\omega/2)}e^{-j\omega M/2}$ 10. $e^{j\omega_0 n}$ $\sum 2\pi\delta(\omega-\omega_0+2\pi k)$ $\sum \left[\pi e^{j\phi}\delta(\omega-\omega_0+2\pi k)+\pi e^{-j\phi}\delta(\omega+\omega_0+2\pi k)\right]$ 11. $\cos(\omega_0 n + \phi)$ $k = -\infty$

Oppenheim, Schafer and Buck, Discrete-time signal processing, Prentice Hall, 1999

|n| w |n| x = 1 (w |v|) = x

(MY) = A(E') D



Bracewell, The Fourier Transform and its Applications, McGraw Hill 1978

68

Why is the Fourier domain particularly useful?

• It tells us the effect of linear convolutions.

Fourier transform of convolution

Consider a (circular) convolution of g and h

$$f = g \otimes h$$

Fourier transform of convolution $f = g \otimes h$

Take DFT of both sides

$$F[m,n] = DFT(g \otimes h)$$

Fourier transform of convolution $f = g \otimes h$ $F[m,n] = DFT(g \otimes h)$

Write the DFT and convolution explicitly

$$F[m,n] = \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]h[k,l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}$$
Fourier transform of convolution $f = g \otimes h$ $F[m,n] = DFT(g \otimes h)$ $F[m,n] = \sum_{n=1}^{M-1} \sum_{n=1}^{N-1} g[u-k,v-l]h[k,l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}$

Move the exponent in

u=0 v=0 k l

$$=\sum_{u=0}^{M-1}\sum_{v=0}^{N-1}\sum_{k,l}g[u-k,v-l]e^{-\pi i\left(\frac{um}{M}+\frac{vn}{N}\right)}h[k,l]$$

Fourier transform of convolution

$$f = g \otimes h$$

$$F[m,n] = DFT(g \otimes h)$$

$$F[m,n] = \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]h[k,l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}$$

$$= \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}h[k,l]$$

Change variables in the sum

$$= \sum_{\mu=-k}^{M-k-1} \sum_{\nu=-l}^{N-l-1} \sum_{k,l} g[\mu,\nu] e^{-\pi i \left(\frac{(k+\mu)m}{M} + \frac{(l+\nu)n}{N}\right)} h[k,l]$$

Fourier transform of convolution

$$\begin{aligned} f &= g \otimes h \\ F[m,n] &= DFT(g \otimes h) \\ F[m,n] &= \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]h[k,l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)} \\ &= \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}h[k,l] \\ &= \sum_{\mu=-k}^{M-k-1} \sum_{\nu=-l}^{N-l-1} \sum_{k,l} g[\mu,\nu]e^{-\pi i \left(\frac{(k+\mu)m}{M} + \frac{(l+\nu)n}{N}\right)}h[k,l] \end{aligned}$$

Perform the DFT (circular boundary conditions)

$$=\sum_{k,l}G[m,n]e^{-\pi i\left(\frac{km}{M}+\frac{\ln}{N}\right)}h[k,l]$$

Fourier transform of convolution

$$f = g \otimes h$$

$$F[m,n] = DFT(g \otimes h)$$

$$F[m,n] = \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]h[k,l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}$$

$$= \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \sum_{k,l} g[u-k,v-l]e^{-\pi i \left(\frac{um}{M} + \frac{vn}{N}\right)}h[k,l]$$

$$= \sum_{\mu=-k}^{M-k-1} \sum_{v=-l}^{N-l-1} \sum_{k,l} g[\mu,v]e^{-\pi i \left(\frac{(k+\mu)m}{M} + \frac{(l+v)n}{N}\right)}h[k,l]$$

$$= \sum_{k,l} G[m,n]e^{-\pi i \left(\frac{km}{M} + \frac{\ln}{N}\right)}h[k,l]$$

Perform the other DFT (circular boundary conditions)

$$=G[m,n]H[m,n]$$







original

Filtered (no change)









blurred

original







original

$$F[m,n] = \sum_{k=0}^{M-1} \sum_{l=0}^{N-1} f[k,l] e^{-\pi i \left(\frac{km}{M} + \frac{\ln}{N}\right)}$$

high-pass filter
$$= 2 - \frac{1}{3} \left(1 + 2\cos\left(\frac{\pi m}{M}\right)\right)$$

$$\int_{0}^{2.3}$$

Sampling and aliasing



Sampling in 1D takes a continuous function and replaces it with a vector of values, consisting of the function's values at a set of sample points. We'll assume that these sample points are on a regular grid, and can place one at each integer for convenience.



Sampling in 2D does the same thing, only in 2D. We'll assume that these sample points are on a regular grid, and can place one at each integer point for convenience.



A continuous model for a sampled function

- We want to be able to approximate integrals sensibly
- Leads to
 - the delta function $Sample_{2D}(f(x,y)) = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} f(x,y)\delta(x-i,y-j)$
 - model on right

$$= f(x,y) \sum_{i=-\infty}^{\infty} \sum_{i=-\infty}^{\infty} \delta(x-i,y-j)$$

The Fourier transform of a sampled signal

$$F(\operatorname{Sample}_{2D}(f(x,y))) = F\left(f(x,y)\sum_{i=-\infty}^{\infty}\sum_{i=-\infty}^{\infty}\delta(x-i,y-j)\right)$$
$$= F(f(x,y)) * F\left(\sum_{i=-\infty}^{\infty}\sum_{i=-\infty}^{\infty}\delta(x-i,y-j)\right)$$
$$= \sum_{i=-\infty}^{\infty}\sum_{j=-\infty}^{\infty}F(u-i,v-j)$$





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



Resample the checkerboard by taking one sample at each circle. In the case of the top left board, new representation is reasonable. Top right also yields a reasonable representation. Bottom left is all black (dubious) and bottom right has checks that are

too big.



Constructing a pyramid by taking every second pixel leads to layers that badly misrepresent the top layer



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.

256x256 128x128 64x64 32x32 16x16





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.

256x256 128x128 64x64 32x32 16x16







What is a good representation for image analysis?

- Fourier transform domain tells you "what" (textural properties), but not "where".
- Pixel domain representation tells you "where" (pixel location), but not "what".
- Want an image representation that gives you a local description of image events what is happening where.

Image pyramids

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



512 256 128 64 32 16 8



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



512 256 128 64 32 16 8





512 256 128 64 32 16 8





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



Reprinted from "Shiftable MultiScale Transforms," by Simoncelli et al., IEEE Transactions on Information Theory, 1992, copyright 1992, IEEE