



Linear Filters

Monday, Jan 24
Prof. Kristen Grauman
UT-Austin




Announcements

- **Office hours** Mon-Thurs 5-6 pm
 - Mon: Yong Jae, PAI 5.33
 - Tues/Thurs: Shalini, PAI 5.33
 - Wed: Me, ACES 3.446



- cv-spring2011@cs.utexas.edu for assignment questions outside of office hours
- **Pset 0** due Friday Jan 28. Drop box in PAI 5.38. Attach cover page with name and CS 376




Computer Vision: Algorithms and Applications

© 2010 Richard Szeliski, Microsoft Research

Welcome to the Web site (<http://szeliski.org/Book/>) for my computer vision textbook, which you can now purchase at a variety of locations, including [Springer](#), [Amazon](#), and [Barnes & Noble](#). This book is largely based on the computer vision courses that I have co-taught at the University of Washington (2008, 2005, 2001) and Stanford (2001) with [Serra Szeliski](#) and [David Fleet](#). You are welcome to download the PDF from this Web site for personal use, but not to repost it on any other Web site. Please post a link to this URL (<http://szeliski.org/Book/>) instead. An electronic version will continue to be available even after the book is published. Note, however, that while the content of the electronic and hardcopy versions are the same, the page layout (especially electronic version) is optimized for color reading. The PDFs should be enabled for commenting directly in your viewer. Also, from links to sections, equations, and references are enabled. To get back to where you were, use Alt-Left-Arrow.

Plan for today

- Image noise
- Linear filters
 - Examples: smoothing filters
- Convolution / correlation



Digital camera



A digital camera replaces film with a sensor array

- Each cell in the array is light-sensitive diode that converts photons to electrons
- <http://electronics.howstuffworks.com/digital-camera.htm>

Slide by Steve Seitz

Digital images

FIGURE 2.17 (a) Continuous image projected onto a sensor array. (b) Result of image sampling and quantization.

Slide credit: Derek Hoiem

Digital images

- **Sample** the 2D space on a regular grid
- **Quantize** each sample (round to nearest integer)
- Image thus represented as a matrix of integer values.

62	79	23	119	100	105	4	0
10	10	9	82	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	88	0	49
2	1	1	59	26	37	0	77
0	86	144	147	187	102	62	295
255	252	0	168	123	62	0	31
168	63	127	17	1	0	99	30

Adapted from S. Seitz

Digital color images

Bayer filter

© 2000 How Stuff Works

Digital color images

Color images, RGB color space

R G B

Images in Matlab

- Images represented as a matrix
- Suppose we have a NxM RGB image called "im"
 - $im(1,1,1)$ = top-left pixel value in R-channel
 - $im(y, x, b)$ = y pixels down, x pixels to right in the bth channel
 - $im(N, M, 3)$ = bottom-right pixel in B-channel
- `imread(filename)` returns a uint8 image (values 0 to 255)
 - Convert to double format (values 0 to 1) with `im2double`

0.92	0.91	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99	
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91	
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92	
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95	
0.71	0.81	0.81	0.87	0.57	0.27	0.80	0.88	0.89	0.79	0.85	
0.49	0.62	0.69	0.98	0.50	0.60	0.56	0.50	0.61	0.45	0.33	
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74	
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93	
0.69	0.49	0.55	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99	
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97	
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93	
	0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
	0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

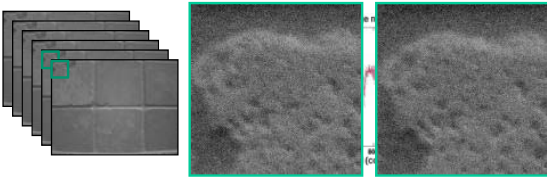
Slide credit: Derek Hoiem

Image filtering

- Compute a function of the local neighborhood at each pixel in the image
 - Function specified by a "filter" or mask saying how to combine values from neighbors.
- Uses of filtering:
 - Enhance an image (denoise, resize, etc)
 - Extract information (texture, edges, etc)
 - Detect patterns (template matching)

Adapted from Derek Hoiem


Motivation: noise reduction



- Even multiple images of the **same static scene** will not be identical.

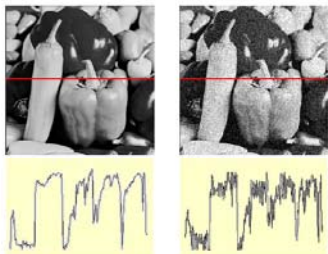
Common types of noise

- **Salt and pepper noise:** random occurrences of black and white pixels
- **Impulse noise:** random occurrences of white pixels
- **Gaussian noise:** variations in intensity drawn from a Gaussian normal distribution



Source: S. Seitz

Gaussian noise



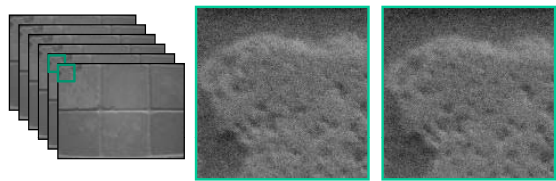
Ideal Image Noise process Gaussian i.i.d. ("white") noise:
 $f(x, y) = \hat{f}(x, y) + \eta(x, y)$
 $\eta(x, y) \sim \mathcal{N}(\mu, \sigma)$

```
>> noise = randn(size(im)).*sigma;
>> output = im + noise;
```

What is impact of the sigma?

Fig. M. Hebert

Motivation: noise reduction



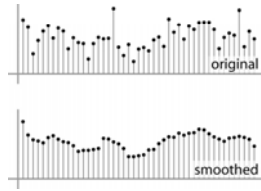
- Even multiple images of the same static scene will not be identical.
- How could we reduce the noise, i.e., give an estimate of the true intensities?
- **What if there's only one image?**

First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Assumptions:
 - Expect pixels to be like their neighbors
 - Expect noise processes to be independent from pixel to pixel

First attempt at a solution


- Let's replace each pixel with an average of all the values in its neighborhood
- Moving average in 1D:



Source: S. Marschner

Weighted Moving Average


Can add weights to our moving average
Weights [1, 1, 1, 1, 1] / 5



Source: S. Marschner

Weighted Moving Average

Non-uniform weights [1, 4, 6, 4, 1] / 16





Source: S. Marschner

Moving Average In 2D

$F[x, y]$

$G[x, y]$





Source: S. Seitz

Moving Average In 2D

$F[x, y]$

$G[x, y]$





Source: S. Seitz

Moving Average In 2D

$F[x, y]$

$G[x, y]$





Source: S. Seitz

Moving Average In 2D

$F[x, y]$

$G[x, y]$



Source: S. Seitz

Moving Average In 2D

$F[x, y]$

$G[x, y]$

Source: S. Seitz

Moving Average In 2D

$F[x, y]$

$G[x, y]$

Source: S. Seitz

Correlation filtering

Say the averaging window size is $2k+1 \times 2k+1$:

$$G[i, j] = \frac{1}{(2k+1)^2} \sum_{u=-k}^k \sum_{v=-k}^k F[i+u, j+v]$$

Attribute uniform weight to each pixel
Loop over all pixels in neighborhood around image pixel $F[i,j]$

Now generalize to allow **different weights** depending on neighboring pixel's relative position:

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k \underbrace{H[u, v]}_{\text{Non-uniform weights}} F[i+u, j+v]$$

Correlation filtering

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i+u, j+v]$$

This is called **cross-correlation**, denoted $G = H \otimes F$

Filtering an image: replace each pixel with a linear combination of its neighbors.

The filter "**kernel**" or "**mask**" $H[u, v]$ is the prescription for the weights in the linear combination.

Averaging filter

- What values belong in the kernel H for the moving average example?

$F[x, y]$

\otimes

$H[u, v]$

“box filter”

$\frac{1}{9}$

$G[x, y]$

$$G = H \otimes F$$

Smoothing by averaging

depicts box filter:
 white = high value, black = low value

original

filtered

What if the filter size was 5×5 instead of 3×3 ?

Boundary issues

What is the size of the output?

- MATLAB: output size / "shape" options
 - shape = 'full': output size is sum of sizes of f and g
 - shape = 'same': output size is same as f
 - shape = 'valid': output size is difference of sizes of f and g

Source: S. Lazebnik

Boundary issues

What about near the edge?

- the filter window falls off the edge of the image
- need to extrapolate
- methods:
 - clip filter (black)
 - wrap around
 - copy edge
 - reflect across edge

Source: S. Marschner

Boundary issues

What about near the edge?

- the filter window falls off the edge of the image
- need to extrapolate
- methods (MATLAB):
 - clip filter (black): `imfilter(f, g, 0)`
 - wrap around: `imfilter(f, g, 'circular')`
 - copy edge: `imfilter(f, g, 'replicate')`
 - reflect across edge: `imfilter(f, g, 'symmetric')`

Source: S. Marschner

Gaussian filter

- What if we want nearest neighboring pixels to have the most influence on the output?

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$$H[u, v] = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

This kernel is an approximation of a 2d Gaussian function:

$$h(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$

- Removes high-frequency components from the image ("low-pass filter").

Source: S. Seitz

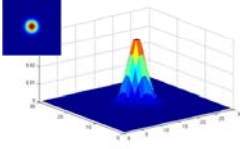
Smoothing with a Gaussian

Gaussian filters

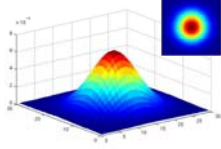
- What parameters matter here?
- Size of kernel or mask
 - Note, Gaussian function has infinite support, but discrete filters use finite kernels

Gaussian filters

- What parameters matter here?
- **Variance** of Gaussian: determines extent of smoothing



$\sigma = 2$ with
30 x 30
kernel



$\sigma = 5$ with
30 x 30
kernel



Matlab

```

>> hsize = 10;
>> sigma = 5;
>> h = fspecial('gaussian' hsize, sigma);



>> mesh(h);
>> imagesc(h);

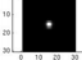
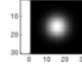
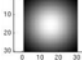
>> outim = imfilter(im, h); % correlation
>> imshow(outim);
    
```


→


Smoothing with a Gaussian

Parameter σ is the “scale” / “width” / “spread” of the Gaussian kernel, and controls the amount of smoothing.


...



...

...


```

for sigma=1:3:10
    h = fspecial('gaussian', fsize, sigma);
    out = imfilter(im, h);
    imshow(out);
    pause;
end
    
```

Properties of smoothing filters

- Smoothing
 - Values positive
 - Sum to 1 → constant regions same as input
 - Amount of smoothing proportional to mask size
 - Remove “high-frequency” components; “low-pass” filter

Filtering an impulse signal

What is the result of filtering the impulse signal (image) F with the arbitrary kernel H ?

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	1	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

⊗

a	b	c
d	e	f
g	h	i

=

?		
---	--	--

$F[x, y]$

$H[u, v]$

$G[x, y]$

Convolution

- Convolution:
 - Flip the filter in both dimensions (bottom to top, right to left)
 - Then apply cross-correlation

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i - u, j - v]$$

$$G = H \star F$$

\uparrow
 Notation for
 convolution
 operator

Convolution vs. correlation

Convolution

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i - u, j - v]$$

$$G = H \star F$$

Cross-correlation

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i + u, j + v]$$

$$G = H \otimes F$$

For a Gaussian or box filter, how will the outputs differ?
If the input is an impulse signal, how will the outputs differ?

Predict the outputs using correlation filtering

Three examples of correlation filtering on an eye image. Each example shows the eye image, a 3x3 kernel, and an equals sign followed by a question mark.

- Example 1: Eye image * $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} = ?$
- Example 2: Eye image * $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} = ?$
- Example 3: Eye image * $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} = \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = ?$

Practice with linear filters



Original

0	0	0
0	1	0
0	0	0

?

Source: D. Lowe

Practice with linear filters



Original

0	0	0
0	1	0
0	0	0



Filtered
(no change)

Source: D. Lowe

Practice with linear filters



Original

0	0	0
0	0	1
0	0	0

?

Source: D. Lowe

Practice with linear filters



Original


0	0	0
0	0	1
0	0	0



Shifted left
by 1 pixel
with
correlation

Source: D. Lowe

Practice with linear filters



$\frac{1}{9}$


1	1	1
1	1	1
1	1	1

?

Original


Source: D. Lowe

Practice with linear filters



$\frac{1}{9}$

1	1	1
1	1	1
1	1	1




Original

Blur (with a box filter)

Source: D. Lowe

Practice with linear filters



0	0	0
0	2	0
0	0	0

-

$\frac{1}{9}$


1	1	1
1	1	1
1	1	1

?

Original

Source: D. Lowe

Practice with linear filters




0	0	0
0	2	0
0	0	0

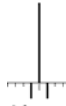
-

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1



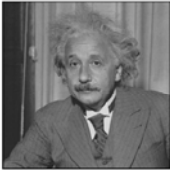
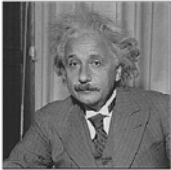
Original



Sharpening filter:
accentuates differences
with local average

Source: D. Lowe

Filtering examples: sharpening

before

after

Properties of convolution

- **Shift invariant:**
 - Operator behaves the same everywhere, i.e. the value of the output depends on the pattern in the image neighborhood, not the position of the neighborhood.
- **Superposition:**
 - $h * (f1 + f2) = (h * f1) + (h * f2)$

Properties of convolution

- Commutative:
 $f * g = g * f$
- Associative
 $(f * g) * h = f * (g * h)$
- Distributes over addition
 $f * (g + h) = (f * g) + (f * h)$
- Scalars factor out
 $kf * g = f * kg = k(f * g)$
- Identity:
unit impulse $e = [\dots, 0, 0, 1, 0, 0, \dots]$. $f * e = f$

Separability

- In some cases, filter is separable, and we can factor into two steps:
 - Convolve all rows
 - Convolve all columns

Separability

- In some cases, filter is separable, and we can factor into two steps: e.g.,

g

1	2	1
---	---	---

f

h

2	3	3
3	5	5
4	4	6


11	
18	
18	

What is the computational complexity advantage for a separable filter of size $k \times k$, in terms of number of operations per output pixel?


$f * (g * h) = (f * g) * h$

Effect of smoothing filters

5x5



Additive Gaussian noise



Salt and pepper noise

Median filter

10	15	20
23	90	27
33	31	30

Median value →

10	15	20
23	27	27
33	31	30

Sort


↓

Replace


- No new pixel values introduced
- Removes spikes: good for impulse, salt & pepper noise
- Non-linear filter

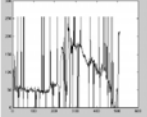
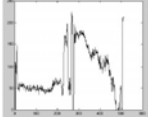
Median filter

Salt and pepper noise



Median filtered



Plots of a row of the image

Matlab: output `im = medfilt2(im, [h w]);`

Source: M. Hebert

Median filter

- Median filter is edge preserving

	INPUT
	MEDIAN
	MEAN

Filtering application: Hybrid Images

What you see...

From Far Away

Up Close

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006

Application: Hybrid Images

A. Oliva, A. Torralba, P.G. Schyns, "Hybrid Images," SIGGRAPH 2006

Gaussian Filter

Laplacian Filter

-
≈

unit impulse Gaussian Laplacian of Gaussian

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006

Changing expression

SIGGRAPH2006

Sad
←————→
Surprised

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006

Summary

- Image "noise"
- Linear filters and convolution useful for
 - Enhancing images (smoothing, removing noise)
 - Box filter
 - Gaussian filter
 - Impact of scale / width of smoothing filter
 - Detecting features (next time)
- Separable filters more efficient
- Median filter: a non-linear filter, edge-preserving

Coming up

- **Wednesday:**
 - Filtering part 2: filtering for features
- **Friday:**
 - Pset 0 is due via turnin, 11:59 PM