


The slide features a large image of a tiger's face. To its left, two smaller images show the tiger's face with different filter effects. Below the main image, the text reads "Linear Filters" and "Tues, Jan 23, 2018". At the bottom, there are three small images of a panda's face, each with a different filter effect, and their corresponding filter kernels shown below them.

Announcements

- Piazza for assignment questions
- **A0** due today. Submit on Canvas.
- Office hours posted on class website


Course homepage

- <http://vision.cs.utexas.edu/376-spring2018/>

	Thurs Jan 18	Course intro	Textbook Sec 1.1-1.3 Course requirements UTCS account setup Basic Matlab tutorial Running Matlab at UT		A0 out due Tues Jan 23 See optional Latex info
	Tues Jan 23	Features and filters	Sec 3.1.1-2, 3.2	Linear filters	

Computer Vision: Algorithms and Applications

© 2010 Richard Szeliski, Microsoft Research



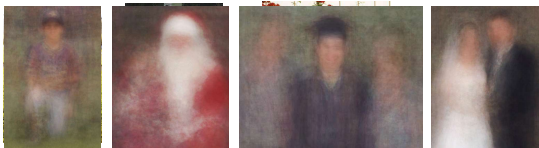
Welcome to the Web site (<http://szeliski.org/Book/>) for my computer vision textbook, which you can now purchase at a variety of locations, including [Springer](#), [Amazon](#), and [Barnes & Noble](#). This book is largely based on the computer vision courses that I have co-taught at the University of Washington (2008, 2005, 2001) and Stanford (2003) with [Steve Seitz](#) 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. As the manuscript will continue to be available even after the book is published. Note, however, that while the content of the electronic and hard-copy versions are the same, the page layout (pagination) of the electronic version is optimized for online reading. The PDFs should be enabled for continuous stretch in your viewer. Also, broken 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


Images as matrices

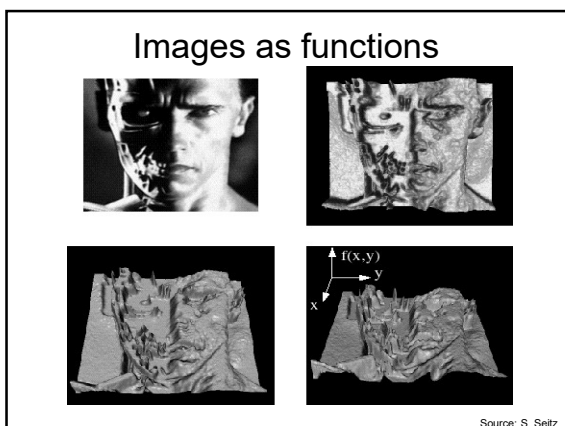
Result of averaging 100 similar snapshots



Little Leaguer *Kids with Santa* *The Graduate* *Newlyweds*

From: *100 Special Moments*, by Jason Salavon (2004)
<http://salavon.com/SpecialMoments/SpecialMoments.shtml>





Images as functions

- We can think of an image as a function, f , from \mathbb{R}^2 to \mathbb{R} :
 - $f(x, y)$ gives the **intensity** at position (x, y)
 - Realistically, we expect the image only to be defined over a rectangle, with a finite range:
 - $f: [a, b] \times [c, d] \rightarrow [0, 255]$
- A color image is just three functions pasted together. We can write this as a "vector-valued" function:

$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$

Source: S. Seltz

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

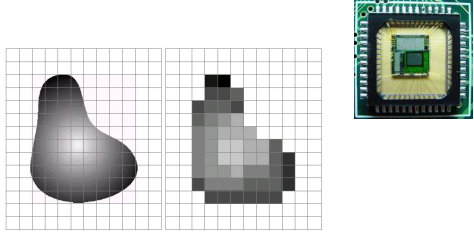
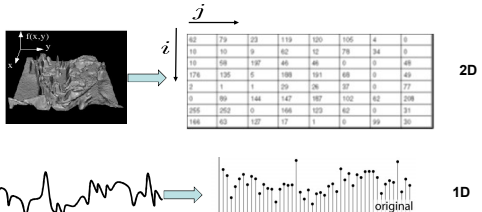


FIGURE 2.17 (a) Continuous image projected onto a sensor array. (b) Result of image sampling and quantization.

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.



82	79	23	119	100	106	4	0
118	119	9	402	12	178	214	9
10	98	192	45	46	0	5	48
176	135	5	188	191	69	0	49
3	1	1	29	26	107	0	77
0	99	144	147	187	102	403	208
262	262	0	166	123	62	0	31
166	62	127	17	1	0	99	20

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`

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

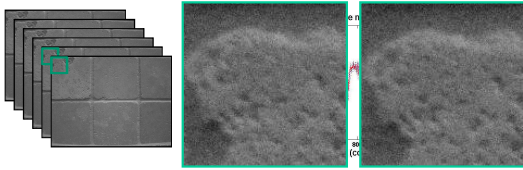
Main idea: image filtering

- 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




Original



Salt and pepper noise




Impulse noise




Gaussian noise

Source: S. Seitz

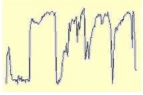
Gaussian noise

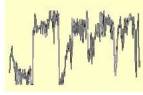


Ideal Image
 $f(x,y)$



Noise process
 $\eta(x,y)$





$f(x,y) = \overline{f(x,y)} + \overline{\eta(x,y)}$

Gaussian i.i.d. ("white") noise:
 $\eta(x,y) \sim \mathcal{N}(\mu, \sigma)$

```

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

What is impact of the sigma?

Fig. M. Hebert

sigma=1

Effect of sigma on Gaussian noise:


Image shows the noise values themselves.


sigma=4

Effect of sigma on Gaussian noise:

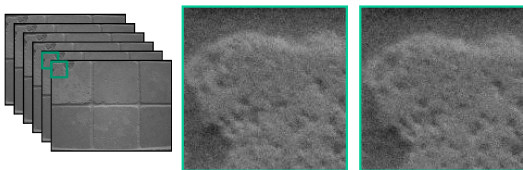
Image shows the noise values themselves.







Motivation: noise reduction



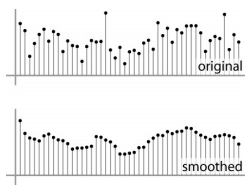
- Even multiple images of the same static scene will not be identical.
- How could we reduce the noise, i.e., give an estimate of the true intensities?
- **What if there's only one image?**

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

$F[x, y]$

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

$G[x, y]$

		0							

Source: S. Seitz

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
 $H[u, v]$
 $\frac{1}{9}$

1	1	1
1	?	1
1	1	1

"box filter"

$G[x, y]$

$G = H \otimes F$

Smoothing by averaging

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

original

filtered

What if the filter size was 5 x 5 instead of 3 x 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

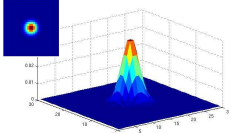
What about near the edge?

- 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')`

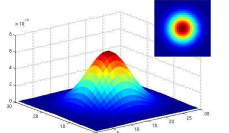
Source: S. Marschner

Gaussian 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);
    
```





→


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
    
```

Keeping the two Gaussians in play straight...

$\sigma=0.2$
 no smoothing
 $\sigma=1$ pixel
 $\sigma=2$ pixels
 Wider smoothing kernel \downarrow

Properties of smoothing filters

- Smoothing
 - Values positive
 - Sum to 1 \rightarrow 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 ?

$F[x, y]$ \otimes $H[u, v]$ $G[x, y]$

Convolution

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

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

$G = H \star F$
Notation for convolution operator

\star

H

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

*

0	0	0
0	1	0
0	0	0

= ?

*

0	0	0
0	0	1
0	0	0

= ?

*


0	0	0
0	2	0
0	0	0

=
 $\frac{1}{9}$
*

1	1	1
1	1	1
1	1	1

= ?

Practice with linear filters




0	0	0
0	1	0
0	0	0

?


Original

Source: D. Lowe

Practice with linear filters



0	0	0
0	1	0
0	0	0




Original

Filtered
(no change)

Source: D. Lowe

Practice with linear filters




0	0	0
0	0	1
0	0	0

?

Original

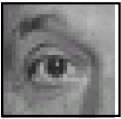
Source: D. Lowe

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



Original


$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

?

Source: D. Lowe

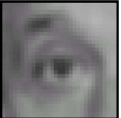
Practice with linear filters



Original

$\frac{1}{9}$


1	1	1
1	1	1
1	1	1



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



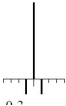
0	0	0
0	2	0
0	0	0

-
 $\frac{1}{9}$

1	1	1
1	1	1
1	1	1

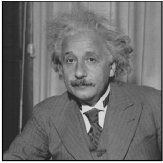
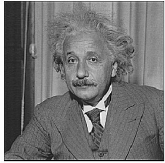


Sharpening filter:
accentuates differences
with local average



Source: D. Lowe

Filtering examples: sharpening

before
after

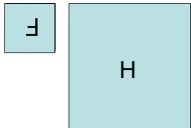
Convolution

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

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

$G = H \star F$

↑
Notation for
convolution
operator



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:**
 - $h \star (f1 + f2) = (h \star f1) + (h \star f2)$

Properties of convolution

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

Separability

- In some cases, filter is separable, and we can factor into two steps:
 - Convolve all rows with a 1D filter
 - Convolve all columns with a 1D filter

$$\begin{bmatrix} \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \end{bmatrix} = \begin{bmatrix} 0 & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & \frac{1}{3} & 0 \end{bmatrix} \circ \begin{bmatrix} 0 & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 \end{bmatrix}$$

Separability

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

g	1	2	1
---	---	---	---

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


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

Effect of smoothing filters



5x5

Additive Gaussian noise



Salt and pepper noise

Median filter

Median value →

10	15	20
23	90	27
33	31	30

10 15 20 23 27 30 31 33 90

↓ Sort

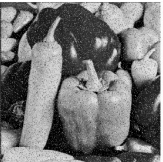
↓ Replace

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


10	15	20
23	27	27
33	31	30

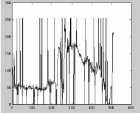
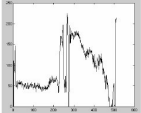
Median filter

Salt and pepper noise



Median filtered




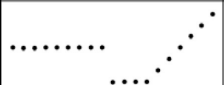




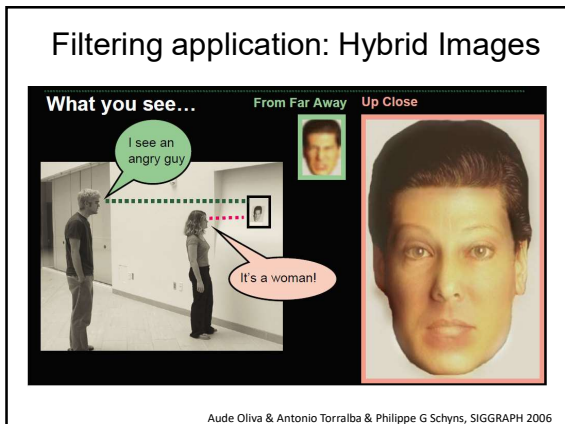
Plots of a row of the image

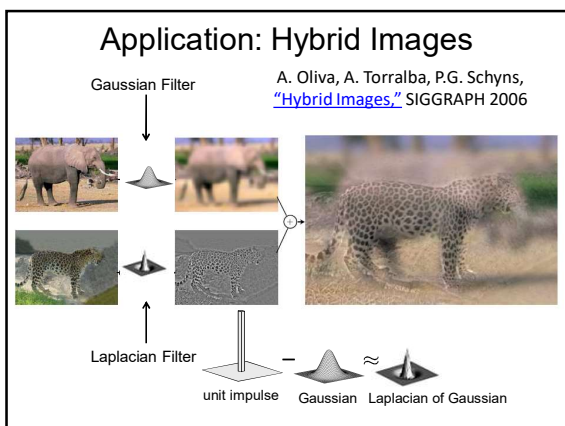
Matlab: output im = medfilt2(im, [h w]); Source: M. Hebert

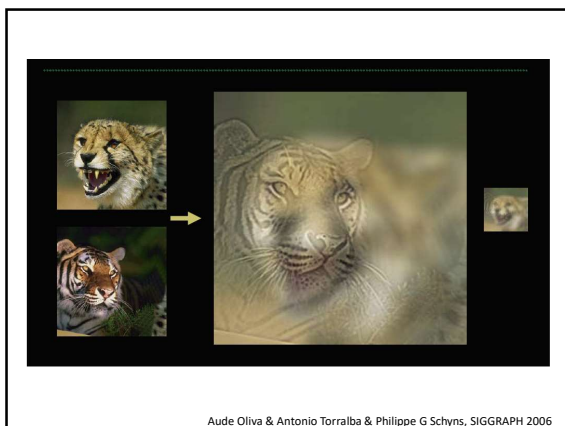
Median filter

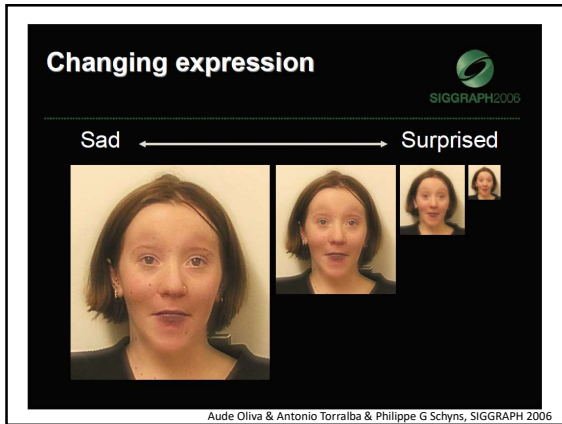
- Median filter is edge preserving

	INPUT
	MEDIAN
	MEAN









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

- **Thursday:**
 - Filtering part 2: filtering for features (edges, gradients, seam carving application)
 - See reading assignment on webpage
- **Today:**
 - Assignment 0 is due on Canvas 11:59 PM
