





CS376 Computer Vision

Lecture 2: Linear Filters



Qixing Huang
January 28th 2019



Announcements

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

Plan for today


- Image noise
- Linear filters
 - Examples: smoothing filters

Images as matrices

Result of averaging 100 similar snapshots




Little Leaguer



Kids with Santa




The Graduate




Newlyweds

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


Image Formation



Images as functions




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>

Digital images




a b

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





Digital images

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




10	10	10	10	24	20
21	21	70	10	30	40
00	05	10	15	30	10
05	15	10	16	17	28
75	50	30	20	25	30
98	78	77	77	75	46

2D



1D

Digital color images



Bayer filter

Digital color images

Slide Credit:
Kristen Grauman

Color images,
RGB color
space



R



G



B

Images in Matlab


- Images represented as a matrix
- Suppose we have a NxM RGB image called "I"
 - $I(1,1,1)$ = top-left pixel value in R-channel
 - $I(y, x, b)$ = y pixels down, x pixels to right in the b^{th} channel
 - $I(N, M, 3)$ = 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`



Main idea: image filtering

- Aggregate the local neighborhood at each pixel in the image
 - Function specified a pattern saying how to aggregate values from neighbors
- Uses of filtering:
 - Enhance an image (denoise, resize, level-of-details, etc)
 - Extract information (texture, edges, features, etc)
 - Detect patterns (template matching)

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

Gaussian noise



$$f(x, y) = \underbrace{\hat{f}(x, y)}_{\text{Ideal Image}} + \underbrace{\eta(x, y)}_{\text{Noise process}}$$

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?

$\sigma=1$


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.




$\sigma=1$



Effect of
sigma on
Gaussian
noise:

This shows
the noise
values
added to the
raw
intensities of
an image.


sigma=
16



Effect of
sigma on
Gaussian
noise:

Image
shows the
noise values
themselves.


$\sigma=16$



Effect of
sigma on
Gaussian
noise

This shows
the noise
values
added to the
raw
intensities of
an image.

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:




Weighted Moving Average

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



Weighted Moving Average

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



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

$G[x, y]$

	0								

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

$G[x, y]$

	0	10							

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

$G[x, y]$

	0	10	20						

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

$$G[x, y]$$

	0	10	20	30					

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

$$G[x, y]$$

	0	10	20	30	30				

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

$$G[x, y]$$

	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

Correlation filtering

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

$$G[i, j] = \underbrace{\frac{1}{(2k+1)^2}}_{\text{Attribute uniform weight to each pixel}} \underbrace{\sum_{u=-k}^k \sum_{v=-k}^k F[i+u, j+v]}_{\text{Loop over all pixels in neighborhood around image pixel } F[i,j]}$$

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?



$$G = H \otimes F$$

Boundary Issues

- What is the size of the output?

Padding Options	
numeric scalar, X	Input array values outside the bounds of the array are assigned the value X. When no padding option is specified, the default is 0.
'symmetric'	Input array values outside the bounds of the array are computed by mirror-reflecting the array across the array border.
'replicate'	Input array values outside the bounds of the array are assumed to equal the nearest array border value.
'circular'	Input array values outside the bounds of the array are computed by implicitly assuming the input array is periodic.
Output Size	
'same'	The output array is the same size as the input array. This is the default behavior when no output size options are specified.
'full'	The output array is the full filtered result, and so is larger than the input array.
Correlation and Convolution Options	
'corr'	imfilter performs multidimensional filtering using correlation, which is the same way that filter2 performs filtering. When no correlation or convolution option is specified, imfilter uses correlation.
'conv'	imfilter performs multidimensional filtering using convolution.

imfilter

Boundary Issues

- What is the size of the output?


✓ **shape — Subsection of filtered data**
'same' (default) | 'full' | 'valid'

Subsection of the filtered data, specified as one of these values:

- 'same' — Return the central part of the filtered data, which is the same size as X.
- 'full' — Return the full 2-D filtered data.
- 'valid' — Return only parts of the filtered data that are computed without zero-padded edges.


filter2

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




$\sigma = 5$ with
10 x 10
kernel




$\sigma = 5$ with
30 x 30
kernel

Gaussian filters

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



$\sigma = 2$ with
 30×30
kernel




$\sigma = 5$ with
 30×30
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
```

Slide credit:
Kristen Grauman

Keeping the two Gaussians in play straight...

Wider smoothing kernel \rightarrow

$\sigma=0.2$



no
smoothing

$\sigma=1$ pixel




$\sigma=2$ pixels





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

Predict the outputs using correlation filtering


$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} = ?$$


$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} = ?$$


$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} - \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = ?$$

Practice with linear filters



Original

0	0	0
0	1	0
0	0	0

?

Practice with linear filters



Original

0	0	0
0	1	0
0	0	0



Filtered
(no change)

Practice with linear filters



Original

0	0	0
0	0	1
0	0	0

?

Practice with linear filters



Original

0	0	0
0	0	1
0	0	0



Shifted left
by 1 pixel
with
correlation

Practice with linear filters



Original

 $\frac{1}{9}$

1	1	1
1	1	1
1	1	1

?


Practice with linear filters



Original

$$\frac{1}{9}$$

1	1	1
1	1	1
1	1	1



Blur (with a
box filter)

Practice with linear filters



Original

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



Original

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



Filtering examples: sharpening



before



after

More Examples


$$\frac{1}{9}$$

1	1	1
1	1	1
1	1	1



0	0	0
0	2	0
0	0	0

—

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

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:**

- $h * (f1 + f2) = (h * f1) + (h * f2)$

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:

$$\text{unit impulse } e = [\dots, 0, 0, 1, 0, 0, \dots]. \quad f * e = f$$

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

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

Effect of smoothing filters

5x5




Additive Gaussian noise



Salt and pepper noise


Median filter





- 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

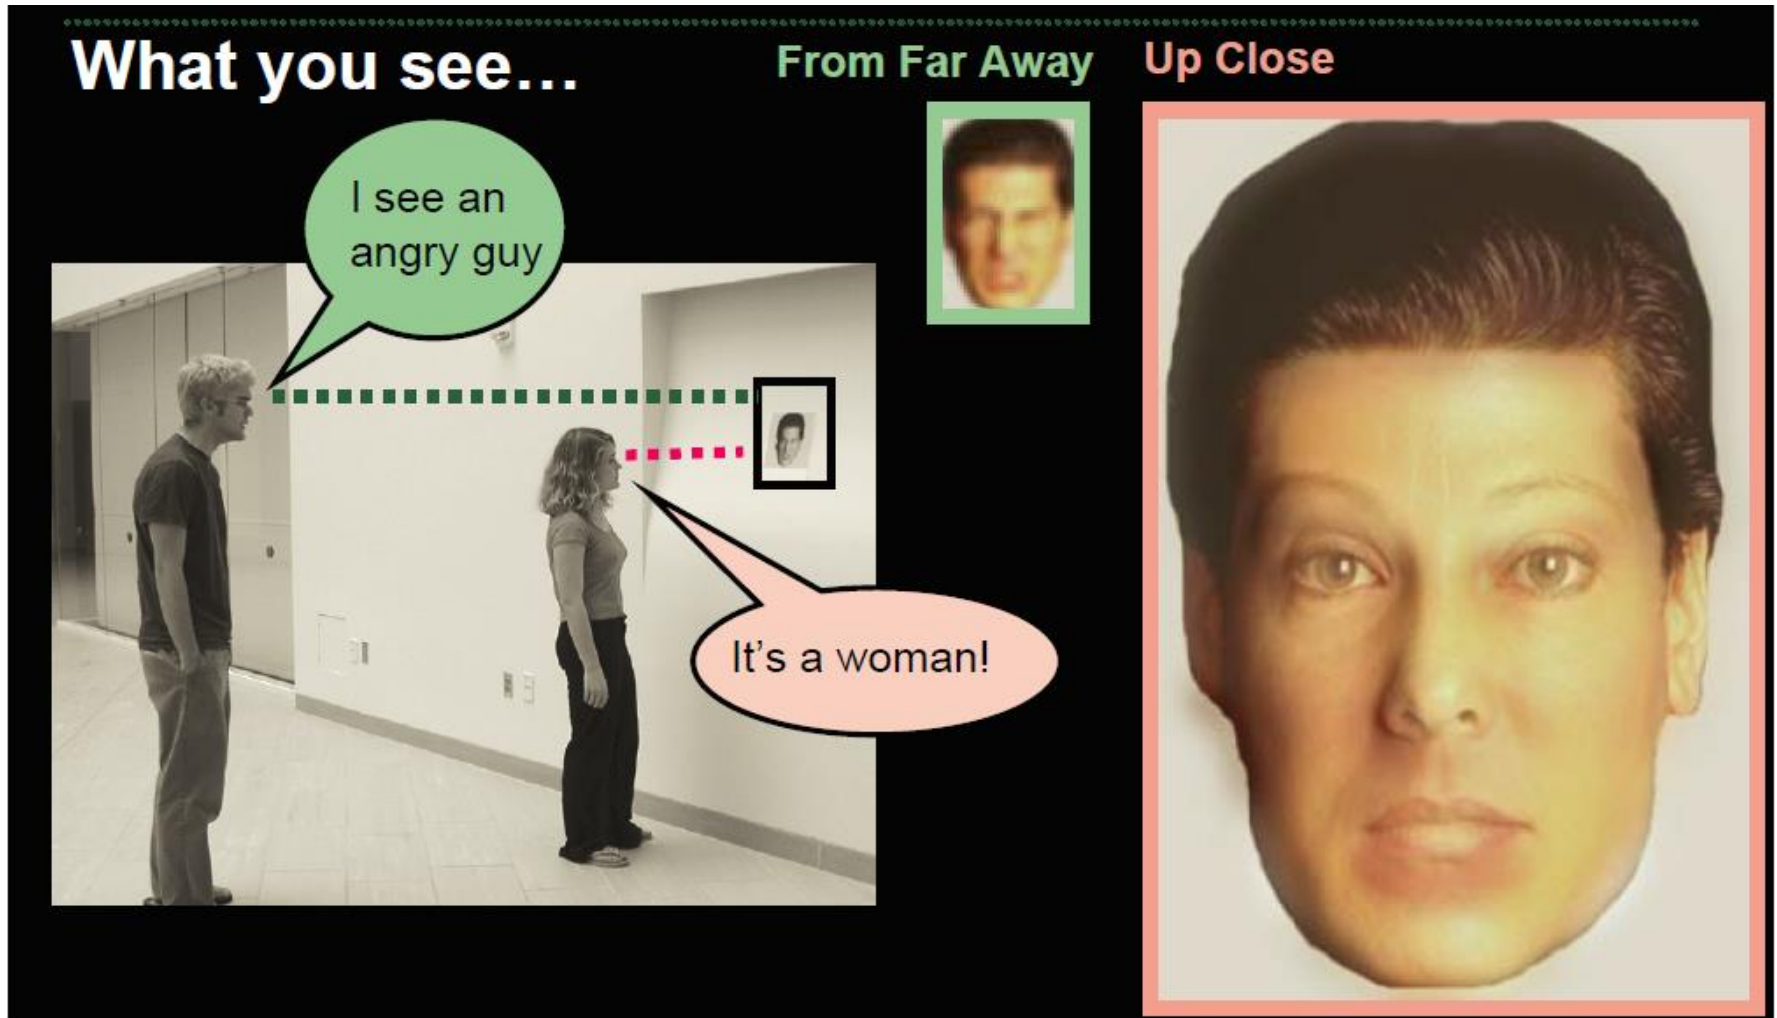
Matlab: `output im = medfilt2(im, [h w]);`

Median filter

- Median filter is edge preserving






Filtering application: Hybrid Images




Application: Hybrid Images

A. Oliva, A. Torralba, P.G. Schyns,
[“Hybrid Images,”](#) SIGGRAPH 2006

Gaussian Filter




Laplacian Filter




unit impulse

—



Gaussian

≈



Laplacian of Gaussian



Changing expression



Sad



Surprised



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 (edges, gradients, seam carving application)
- **Tomorrow:**
 - Assignment 0 is due on Canvas 11:59 PM
-