

CS 1674/2074: Filters

PhD. Nils Murrugarra-Llerena

nem177@pitt.edu



[Motivation] Filters

Sifting Sand

Imagine an image is a big pile of sand with different-sized pebbles mixed in. You want to separate the pebbles by size.

What would you do?

You wouldn't try to pick out each pebble by hand; instead, you'd use a series of **sieves** with different mesh sizes.



Filters


Topics

- Filters: motivation, math and properties
- Types of filters
 - Linear
 - Smoothing
 - Other
 - Non-linear
 - Median
- Applications of filters
 - Texture representation with filters
 - Anti-aliasing for image subsampling




How images are represented

- Color images represented as a matrix with multiple channels (=1 if grayscale)
- Suppose we have a NxM RGB image called “im”
 - `im(0,0,0)` = top-left pixel value in R-channel
 - `im(y, x, b)` = y pixels **down (rows)**, x pixels **to right (cols)** in bth channel
 - `im(N-1, M-1, 2)` = bottom-right pixel in B-channel
- `cv2.imread(filename)` returns a uint8 image (values 0 to 255)



Adapted from Derek Hoiem

Enter Noise



- The same object will look very different across images
- Even multiple images of same static scene won't be identical
- How could we reduce the noise, i.e., give an estimate of the true intensities?
- What if there's only one image?

Common types of noise

- **Impulse noise:** random occurrences of white pixels
- **Salt and pepper noise:** random occurrences of black and 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



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

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

```
>> noise = np.random.rand( *im.shape ) * sigma
>> im_noise = im + noise
```

What is impact of the sigma?


[Github: [Lab 2](#)]



Fig: M. Hebert

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

First attempt at a solution

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

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



Central pixel =
 $(10 \cdot 1 + 14 \cdot 4 + 12 \cdot 6 + 13 \cdot 4 + 20 \cdot 1) / 16$

Adapted from 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	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									

Source: S. Seitz

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								

Source: S. Seitz

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						

Source: S. Seitz

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					

Source: S. Seitz

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				

Source: S. Seitz

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	

In what ways is this output good/expected?
In what ways is it bad (what was lost)?

Source: S. Seitz

[Motivation] Interactive Filters



<https://setosa.io/ev/image-kernels/>



Explained Visually: <https://setosa.io/ev/>

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
 - Element-wise multiplication of filter and image patch
- Uses of filtering:
 - Enhance an image (denoise, resize, etc)
 - Extract information (texture, edges, etc)
 - Detect patterns (template matching)



Adapted from Derek Hoiem

Correlation Filtering

Non-weighted, averaging window size is $2k+1 \times 2k+1$:

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

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 Filtering

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



$$G = H \otimes F$$

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?

Gaussian Filter

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




$$F[x, y]$$

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


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

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



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


Smoothing with a Gaussian




Vs box filter

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 x 30
kernel



$\sigma = 5$ with
30 x 30
kernel

To join, go to: ahaslides.com/4WT8J 



Quiz question 1 of 1

0 players ready

Waiting for players to join...

Get Feedback

Start 

✓ Slide 1 selected for PowerPoint [Switch slide](#)




 0  0 / 100 

Gaussian Filters in Python

```
>> hsize = 10;
>> sigma = 5;
>> filter = fspecial_gauss(hsize, sigma)
```

```
>> plt.matshow(filter)
```



[Github: [Lab 2](#)]




```
>> filt_im = ndimage.convolve(im, filter, mode='constant') #
correlation
>> cv2.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_gaus(hsize, sigma);
    out = ndimage.convolve(im, h);
    cv2.imshow(out);
end
```

Gaussian Filters

How big should the filter be?

- Values at edges should be near zero
- Rule of thumb for Gaussian: set filter half-width ($hsize/2$) to 3σ



Source: Derek Hoiem

Properties of smoothing filters

- Smoothing
 - Values positive
 - Sum to 1 $\rightarrow$ overall intensity same as input
 - Amount of smoothing proportional to mask size
 - Remove “high-frequency” components; “low-pass” filter

Lab 2: Filters

Duration: 10 min



To join, go to: ahaslides.com/A6GP8 



Please, apply a Gaussian Filter to an image and upload your result.

Get Feedback

✓ Slide 1 selected for PowerPoint [Switch slide](#)

Group into themes

 0  0 / 100 


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*



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?

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

$$v = 0$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

$$v = 0$$

$$v = +1$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

$$v = 0$$

$$v = +1$$

$$u = 0, v = -1$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

$$v = 0$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

$$v = 0$$

$$v = +1$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

Convolution vs correlation

Cross-correlation

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

$$G = H \otimes F$$

$$u = -1, v = -1$$

$$v = 0$$

$$v = +1$$

$$u = 0, v = -1$$

F

5	2	5	4	4
5	200	3	200	4
1	5	5	4	4
5	5	1	1	2
200	1	3	5	200
1	200	200	200	1

(i, j)

Convolution

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

$$G = H \star F$$

H

.06	.12	.06
.12	.25	.12
.06	.12	.06

(0, 0)

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:

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

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

?

To join, go to: ahaslides.com/UTC5A 



Quiz question 1 of 1

0 players ready

Waiting for players to join...

Get Feedback


Start 

✓ Slide 1 selected for PowerPoint [Switch slide](#)




 0  0 / 100 

Practice with linear filters




Original

0	0	0
0	0	1
0	0	0



**Shifted left
by 1 pixel**

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

$$\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} =$$

$$\underbrace{\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array}}_{\text{Original image}} + \underbrace{\left(\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 1 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} \right)}_{\text{Image minus blur = details}}$$

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

Source: D. Lowe

Sharpening





before



after

Filters for computing *gradients*


1	0	-1
2	0	-2
1	0	-1



Adapted from Derek Hoiem


Filters for computing *gradients*

$$+1*a + 0*b + (-1)*c = a - c$$




Input image
White = 1, black = 0

Filter: $[+1 \ 0 \ -1]$



Output image


Median Filter



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

Median Filter


- Median filter is edge preserving





Adapted from Kristen Grauman

Median Filter

Salt and
pepper
noise



Median
filtered




Plots of a row of the image

https://docs.scipy.org/doc/scipy/reference/generated/scipy.ndimage.median_filter.html#scipy.ndimage.median_filter

Source: M. Hebert

Boundary Issues

- What is the size of the output?
 - 'full': output size is larger than the size of f
 - 'same': output size is same as f (**Default for Python**)




f = image
 h = filter

	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	

Adapted from S. Lazebnik

Boundary Issues

- What about near the edge?
 - the filter window might fall off the edge of the image (in 'same' or 'full')
 - need to extrapolate
 - methods:
 - clip filter (black)
 - wrap around
 - copy edge
 - reflect across edge



Source: S. Marschner

Separability

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

Separability Example

2D filtering
(center location only)

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array}$$

filter

The filter factors
into an *outer* product
of 1D filters:

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} = \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline 1 \\ \hline \end{array} \times \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array}$$

Perform filtering
along rows (*horizontal pass*):

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array} = \begin{array}{|c|c|c|} \hline & 11 & \\ \hline & 18 & \\ \hline & 18 & \\ \hline \end{array}$$

Followed by filtering
along the remaining column (*vertical pass*):

Asymptotic cost for 2D vs 1D filtering? Let image be $M \times N$, filter be $k \times k$

Separability Example

Convolution Cost Without Separability [Assume same size output]

Suppose you have a 2D image of size $M \times N$ and a $k \times k$ filter:

Naïve 2D convolution:

Each output pixel requires k^2 multiplications and additions. So total cost: $O(M \cdot N \cdot k^2)$

Convolution Cost With Separability [Assume same size output]

A filter is **separable** if it can be expressed as the outer product of two 1D filters:

$$H = f \cdot g^T$$

so instead of convolving with a full $k \times k$ kernel, you can convolve:

1. First with a $1 \times k$ filter (*horizontal pass*).
2. Then with a $k \times 1$ filter (*vertical pass*).

Cost per output pixel:

- First pass: k operations
- Second pass: another k operations
- Total: $2k$ operations. So. Total cost is: $O(M \cdot N \cdot 2k)$

Source: ChatGPT


Application: Hybrid Images



Kristen Grauman

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

Application: Hybrid Images




Kristen Grauman


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

Application: Hybrid Images

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




Application: Hybrid Images



Kristen Grauman

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




Application: Hybrid Images



Plan for next lectures

- Filters: math and properties
- Types of filters
 - Linear
 - Smoothing
 - Other
 - Non-linear
 - Median
- Applications
 - Texture representation with filters
 - Anti-aliasing for image subsampling



Texture



Due to:

Patterns, marks, etches, blobs, holes, relief, etc.

Regular (top), random (bottom) patterns



Why analyze texture?

- Important for how we perceive objects
- Can be an important appearance cue that allows us to distinguish objects, especially if shape is similar across objects

Same shape, different texture/object



Kristen Grauman

<http://animals.nationalgeographic.com/>


Same object, different texture




Kristen Grauman

Texture Representation

- Textures are made up of repeated local patterns, so:
 - Find the patterns
 - Use filters that look like patterns (spots, bars, raw patches...)
 - Consider magnitude of response
 - Describe their statistics within each local window
 - E.g. mean, standard deviation, histogram




Derivative of Gaussian filter




Figures from Noah Snavely


Texture representation: Example




Texture representation: Example



Texture representation: Example




Texture representation: Example



	<u>mean</u> <u>d/dx</u> value	<u>mean</u> <u>d/dy</u> value
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20
	⋮	

statistics to summarize patterns in small windows


Texture representation: Example




	<u>mean</u> <u>d/dx</u> <u>value</u>	<u>mean</u> <u>d/dy</u> <u>value</u>
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20

**statistics to summarize
patterns in small
windows**

Texture representation: Example



Texture representation: Example




	<u>mean</u> <u>d/dx</u> <u>value</u>	<u>mean</u> <u>d/dy</u> <u>value</u>
Win. #1	4	10
Win.#2	18	7
⋮		
Win.#9	20	20

⋮

**statistics to summarize
patterns in small
windows**


Computing Distance using Texture




$$D(a, b) = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2}$$
$$= \sqrt{\sum_{i=1}^d (a_i - b_i)^2}$$

Euclidean distance (L_2)

Texture Representation: Example



Distance reveals how dissimilar texture from window a is from texture in window b.



a: Horizontal and Vertical Edges



b: Horizontal Edges



Filter banks


- Our previous example used two filters and resulted in a 2-dimensional feature vector to describe texture in a window.
 - x and y derivatives revealed something about local structure.
- We can generalize to apply a collection of multiple (d) filters: a “filter bank”.
- Then our feature vectors will be d -dimensional.

Filter banks




- What filters to put in the bank?
 - Typically we want a combination of scales and orientations, different types of patterns.

Which filters would you use to distinguish buildings from animals? Cheetahs from tigers? Ladybugs from dalmatians?



Filter bank




Filter bank: Example

Image from <http://www.texasexplorer.com/austincap2.jpg>




Kristen Grauman

Filter bank: Example




Filter bank: Example




Kristen Grauman


Filter bank: Example




Filter bank: Example




Filter bank: Example



Kristen Grauman




Vectors of texture responses



Adapted from Kristen Grauman


38 filters



1x38

vector
representation
feature

$[r_1, r_2, \dots, r_{38}]$



Vectors of texture responses



To represent pixel, form a “feature vector” from the responses at that pixel.

To represent *image*, compute statistics over all pixel feature vectors, e.g. their mean.

$$\begin{array}{c}
 [r_{(1,1)}^1, r_{(1,1)}^2, \dots, r_{(1,1)}^{38}] \\
 \uparrow \quad \swarrow \\
 [r_{(1,2)}^1, r_{(1,2)}^2, \dots, r_{(1,2)}^{38}] \\
 \text{Pixel location} \quad \text{Filter ID} \\
 \text{(row, column)} \\
 \dots \\
 [r_{(W,H)}^1, r_{(W,H)}^2, \dots, r_{(W,H)}^{38}]
 \end{array}$$

$$[\text{mean}(r_{(:,1)}^1), \text{mean}(r_{(:,2)}^2), \dots, \text{mean}(r_{(:,38)}^{38})]$$

You try: Can you match the texture to the response?



White color
means higher
response



Derek Hoiem

To join, go to: ahaslides.com/KD1WI 



Quiz question 1 of 1

0 players ready

Waiting for players to join...

Get Feedback


Start 

✓ Slide 1 selected for PowerPoint [Switch slide](#)



 0  0 / 100 

Representing textures by mean absolute response



Classifying materials, “stuff”



Figure by Varma & Zisserman

Summary

- Filters useful for
 - Enhancing images (smoothing, removing noise), e.g.
 - Box filter (linear)
 - Gaussian filter (linear)
 - Median filter
 - Detecting patterns (e.g. gradients)
- Texture is a useful property that is often indicative of materials, appearance cues
 - Texture representations summarize repeating patterns of local structure