



Computer Vision

Lecture # 4 **Image Filtering & Its Applications**

Today: Image Filters Review



Smooth/Sharpen Images...

Find edges...

Find waldo...



Image neighborhoods

- Q: What happens if we reshuffle all pixels within the images?



- A: Its histogram won't change.
Point-wise processing unaffected.
- Need to measure properties relative to small *neighborhoods* of pixels

Images as functions





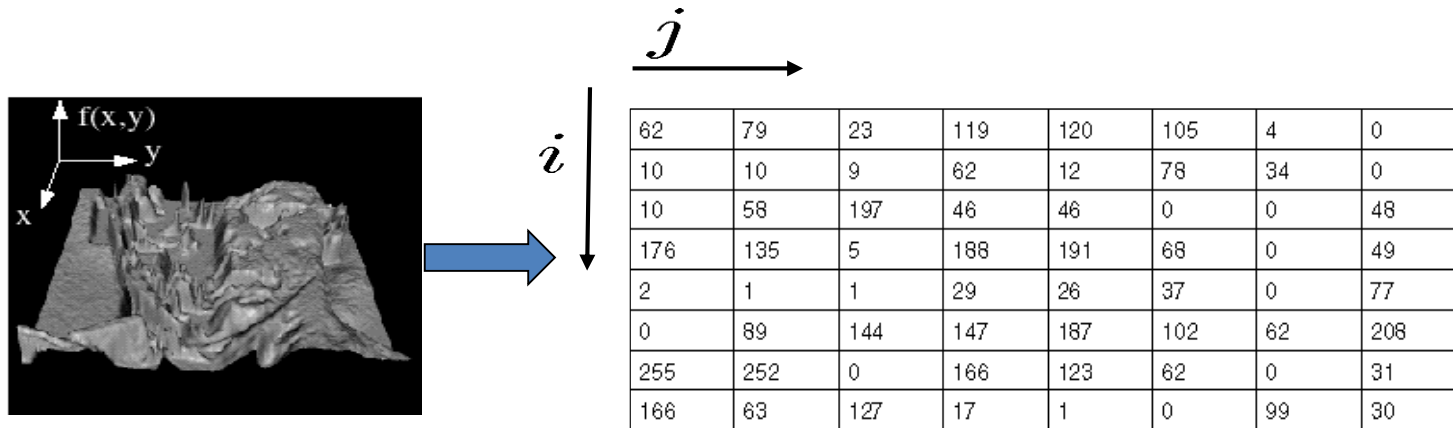
Images as functions

- 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:
 - $f: [a, b] \times [c, d] \rightarrow [0, 1.0]$
- A color image is just three functions pasted together. We can write this as a “vector-valued” function:


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

Digital images

- In computer vision we operate on **digital (discrete)** 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.



Motivation: noise reduction



- We can measure **noise** in multiple images of the same static scene.
- How could we reduce the noise, i.e., give an estimate of the true intensities?

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



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

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

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


```
>> output = im + noise;
```

$\sigma=1$

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=
16



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

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

Motivation: noise reduction




- 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 $2k+1 \times 2k+1$:

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

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?



$$G = H \otimes F$$


Smoothing by averaging



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



original



filtered

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	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

$F[x, y]$


$$\frac{1}{16}$$

1	2	1
2	4	2
1	2	1


$H[u, v]$

This kernel is an approximation of a Gaussian function:

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




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
 $\times 30$ kernel



$\sigma = 5$ with 30
 $\times 30$ kernel

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

$F[x, y]$



a	b	c
d	e	f
g	h	i

$H[u, v]$

$G[x, y]$

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

$F[x, y]$



a	b	c
d	e	f
g	h	i

$H[u, v]$

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	a	b	c	0	0
0	0	d	e	f	0	0
0	0	g	h	i	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

$G[x, y]$

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

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

?

(Note that filter sums to 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


Denoising



Additive Gaussian Noise



Reducing Gaussian noise



Smoothing with larger standard deviations suppresses noise, but also blurs the image

Reducing salt-and-pepper noise by Gaussian smoothing

3x3



5x5




7x7



Alternative idea: Median filtering

- A **median filter** operates over a window by selecting the median intensity in the window




- Is median filtering linear?

Median filter


- What advantage does median filtering have over Gaussian filtering?
 - Robustness to outliers

filters have width 5 :





Median filter

Salt-and-pepper noise



Median filtered



- MATLAB: `medfilt2(image, [h w])`

Median vs. Gaussian filtering

3x3

5x5

7x7

Gaussian



Median



Edge detection

- **Goal:** map image from 2d array of pixels to a set of curves or line segments or contours.
- **Why?**





Figure from J. Shotton et al., PAMI 2007

- **Main idea:** look for strong gradients, post-process

What can cause an edge?

Reflectance change:
appearance
information, texture

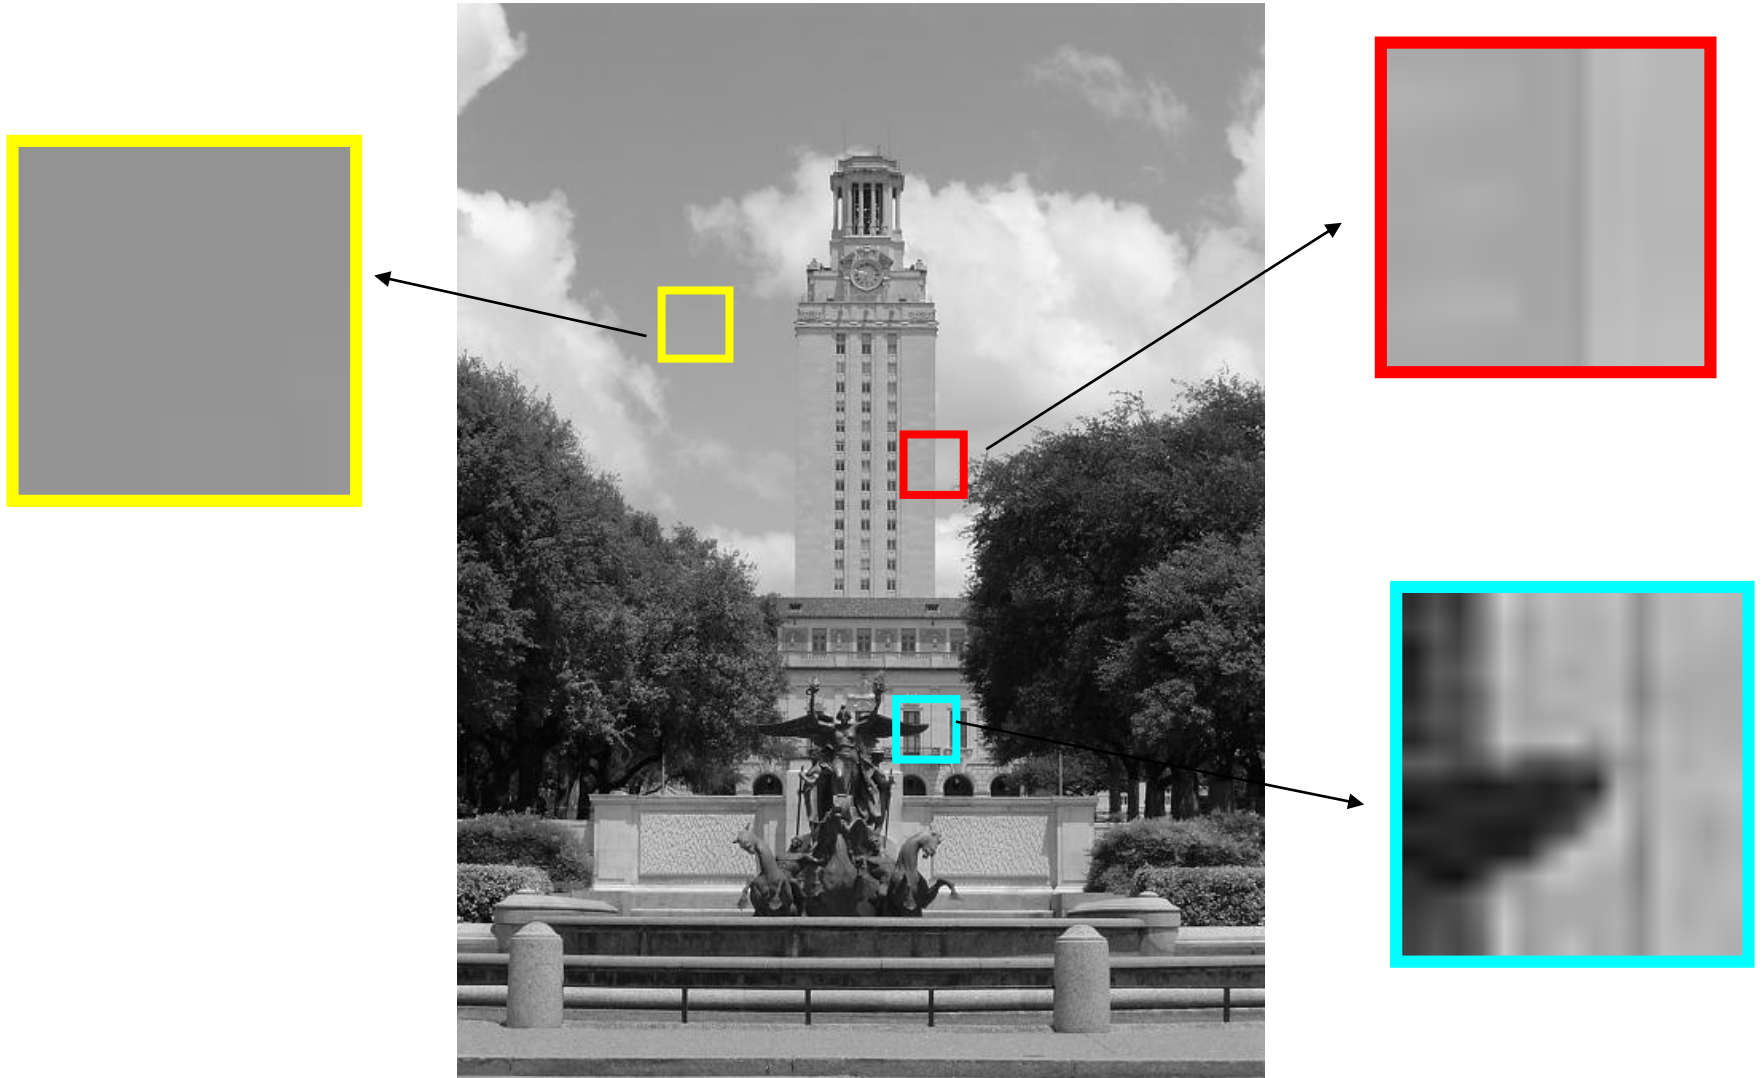
Change in surface
orientation: shape




Depth discontinuity:
object boundary

Cast shadows

Contrast and invariance




Recall : Images as functions



- Edges look like steep cliffs

Derivatives and edges

An edge is a place of rapid change in the image intensity function.



Differentiation and convolution

For 2D function, $f(x,y)$, the partial derivative is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:

$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x + 1, y) - f(x, y)}{1}$$


To implement above as convolution, what would be the associated filter?

Partial derivatives of an image




$$\frac{\partial f(x, y)}{\partial x}$$

-1	1
----	---



$$\frac{\partial f(x, y)}{\partial y}$$

-1	?	1
1	or	-1



Which shows changes with respect to x?

(showing flipped filters)

Assorted finite difference filters

Prewitt: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

Sobel: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

Roberts: $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

```
>> My = fspecial('sobel');  
>> outim = imfilter(double(im), My);  
>> imagesc(outim);  
>> colormap gray;
```







Image gradient

The gradient of an image:

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$$

The gradient points in the direction of most rapid change in intensity



The gradient direction (orientation of edge normal) is given by:

$$\theta = \tan^{-1} \left(\frac{\partial f / \partial y}{\partial f / \partial x} \right)$$

The *edge strength* is given by the gradient magnitude

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$


Thresholding

- Choose a threshold value t
- Set any pixels less than t to zero (off)
- Set any pixels greater than or equal to t to one (on)

Original image



Gradient magnitude image



Thresholding gradient with a lower threshold



Thresholding gradient with a higher threshold



Canny edge detector

- Filter image with derivative of Gaussian
- Find magnitude and orientation of gradient
- **Non-maximum suppression:**
 - Thin multi-pixel wide “ridges” down to single pixel width
- Linking and thresholding (**hysteresis**):
 - Define two thresholds: low and high
 - Use the high threshold to start edge curves and the low threshold to continue them
- MATLAB: `edge (image, 'canny') ;`
- `>>help edge`

The Canny edge detector



original image (Lena)

The Canny edge detector




norm of the gradient

The Canny edge detector



thresholding

The Canny edge detector




Problem:
pixels along
this edge
didn't survive
the
thresholding

thinning
(non-maximum suppression)

Hysteresis thresholding

- Check that maximum value of gradient value is sufficiently large
 - drop-outs? use **hysteresis**
 - use a high threshold to start edge curves and a low threshold to continue them.



Hysteresis thresholding



original image



high threshold
(strong edges)



low threshold
(weak edges)



hysteresis threshold


Object boundaries vs. edges



Background



Texture




Shadows

Edge detection is just the beginning...



image



human segmentation



gradient magnitude



Berkeley segmentation database:


<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/>

Much more on segmentation later in term...

Template matching


- Filters as **templates**:

Note that filters look like the effects they are intended to find --- “matched filters”




- Use normalized cross-correlation score to find a given pattern (template) in the image.
 - Szeliski Eq. 8.11
- Normalization needed to control for relative brightnesses.

Template matching




Scene




Template (mask)

A toy example

Template matching




Detected template




Template

Template matching



Detected template




Correlation map

Where's Waldo?




Scene



Template

Where's Waldo?




Scene




Template

Where's Waldo?



Detected template



Correlation map

Template matching




Scene



Template

What if the template is not identical to some subimage in the scene?

Template matching




Detected template




Template

Match can be meaningful, if scale, orientation, and general appearance is right.

Template matching

- Goal: find  in image
- Main challenge: What is a good similarity or distance measure between two patches?
 - Correlation
 - Zero-mean correlation
 - Sum Square Difference
 - Normalized Cross Correlation




Matching with filters

- Goal: find  in image
- Method 0: filter the image with eye patch

$$h[m,n] = \sum_{k,l} g[k,l] f[m+k,n+l]$$

f = image
g = filter



Input



Filtered Image


What went wrong?

Matching with filters


- Goal: find  in image
- Method 1: filter the image with zero-mean eye

$$h[m,n] = \sum_{k,l} (f[k,l] - \bar{f})(g[m+k,n+l])$$


← mean of f



Input





Filtered Image (scaled)



Thresholded Image

Matching with filters


- Goal: find  in image
- Method 3: Normalized cross-correlation



Input





Normalized X-Correlation



Thresholded Image

Matching with filters


- Goal: find  in image
- Method 3: Normalized cross-correlation



Input



Normalized X-Correlation




Thresholded Image


Q: What if we want to find larger or smaller eyes?

A: Image Pyramid

Review of Sampling



Gaussian pyramid




Template Matching with Image Pyramids

Input: Image, Template

1. Match template at current scale
2. Downsample image
3. Repeat 1-2 until image is very small
4. Take responses above some threshold, perhaps with non-maxima suppression

Coarse-to-fine Image Registration


1. Compute Gaussian pyramid
2. Align with coarse pyramid
3. Successively align with finer pyramids
 - Search smaller range



Why is this faster?


Are we guaranteed to get the same result?

Laplacian filter




unit impulse

—




Gaussian

≈




Laplacian of Gaussian

2D edge detection filters



Gaussian


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



derivative of Gaussian

$$\frac{\partial}{\partial x} h_{\sigma}(u, v)$$

Laplacian of Gaussian




$$\nabla^2 h_{\sigma}(u, v)$$

∇^2 is the **Laplacian** operator:

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

Laplacian pyramid



512

256


128

64

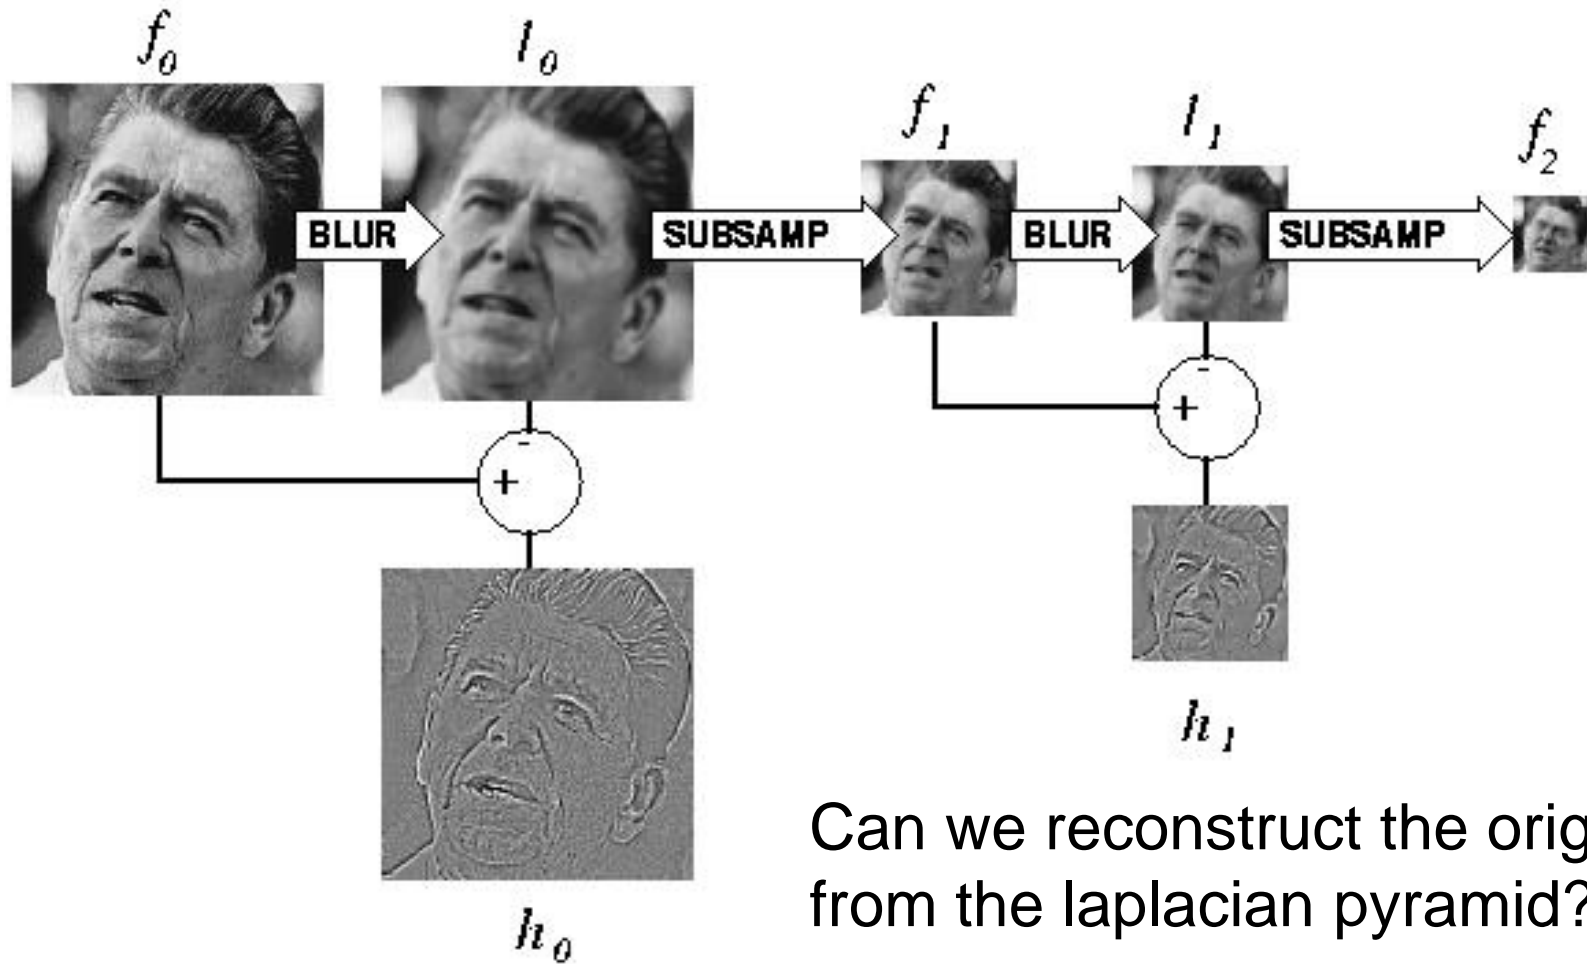
32

16

8



Computing Gaussian/Laplacian Pyramid



Can we reconstruct the original from the laplacian pyramid?

The simplest wavelet transform: the Haar transform

$$U = \begin{array}{|c|c|} \hline 1 & 1 \\ \hline 1 & -1 \\ \hline \end{array}$$

$$U^{-1} = \begin{array}{|c|c|} \hline 0.5 & 0.5 \\ \hline 0.5 & -0.5 \\ \hline \end{array}$$

The simplest set of functions:

$$\vec{F} = U\vec{f}$$

Haar transform

The simplest set of functions:

$$U = \begin{array}{|c|c|} \hline 1 & 1 \\ \hline 1 & -1 \\ \hline \end{array}$$

$$U^{-1} = \begin{array}{|c|c|} \hline 0.5 & 0.5 \\ \hline 0.5 & -0.5 \\ \hline \end{array}$$


To code a signal, repeat at several locations:

$$U = \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 1 & & & & & & & \\ \hline 1 & -1 & & & & & & & \\ \hline & & 1 & 1 & & & & & \\ \hline & & 1 & -1 & & & & & \\ \hline & & & & 1 & 1 & & & \\ \hline & & & & 1 & -1 & & & \\ \hline & & & & & & 1 & 1 & \\ \hline & & & & & & 1 & -1 & \\ \hline \end{array}$$

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

Recursive matrix construction of Haar transform

A_1



2D Haar transform

Basic elements:

1
1

1
-1

1	1
---	---

1	-1
---	----

2D Haar transform

Basic elements:

1
1

1
-1



1	1
---	---

1	-1
---	----


1
1

1	1
---	---

1	1
1	1



Low pass



2D Haar transform

Basic elements:

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$


$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$


$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$$

$$\downarrow \sqrt{2}$$



Low pass



$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

2D Haar transform


Basic elements:

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$




$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$



Low pass

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$

2D Haar transform


Basic elements:

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$




$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$




Low pass

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$




High pass vertical

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$




High pass horizontal

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1 \end{bmatrix}$$

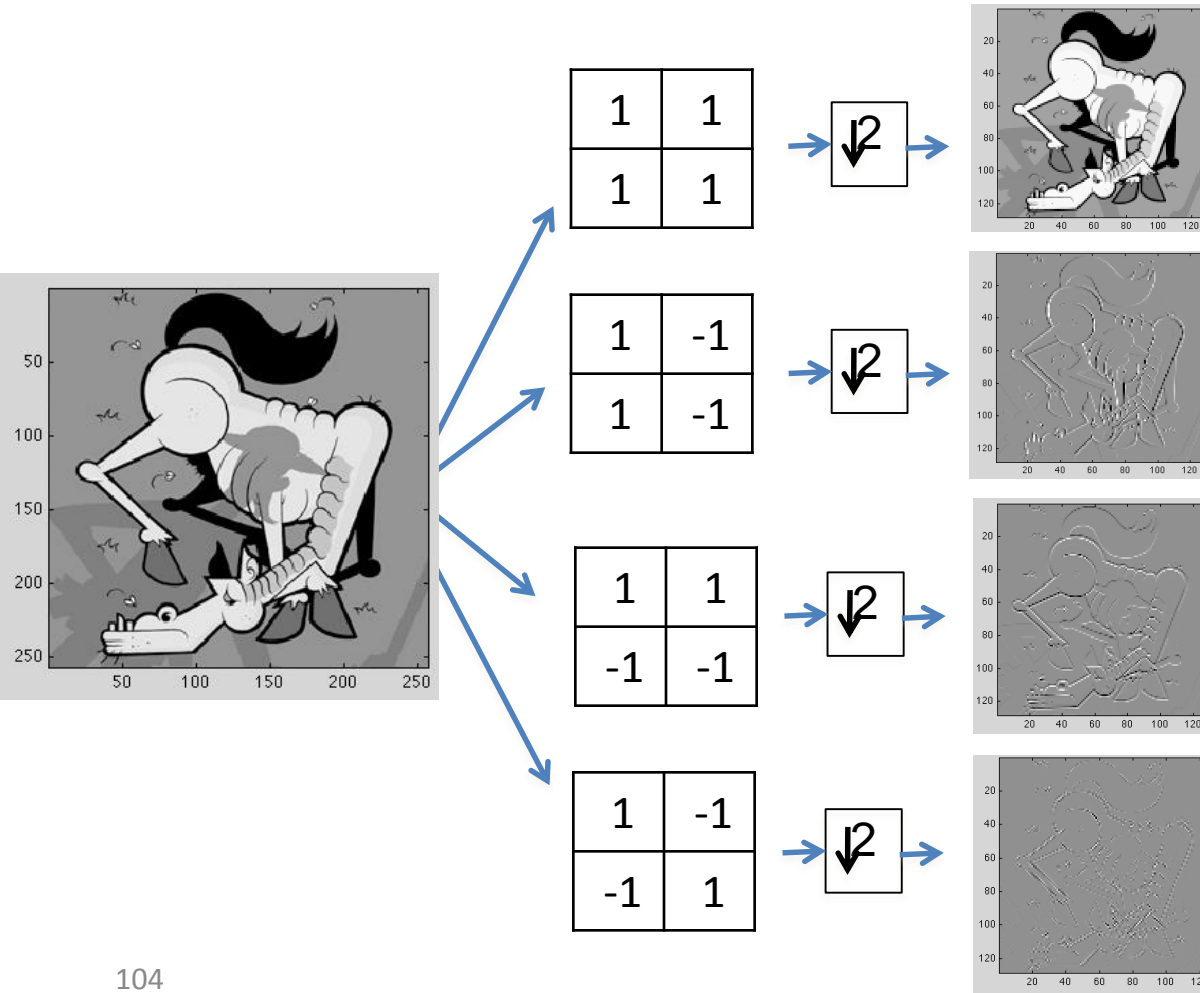
$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

$$\rightarrow \sqrt{2}$$



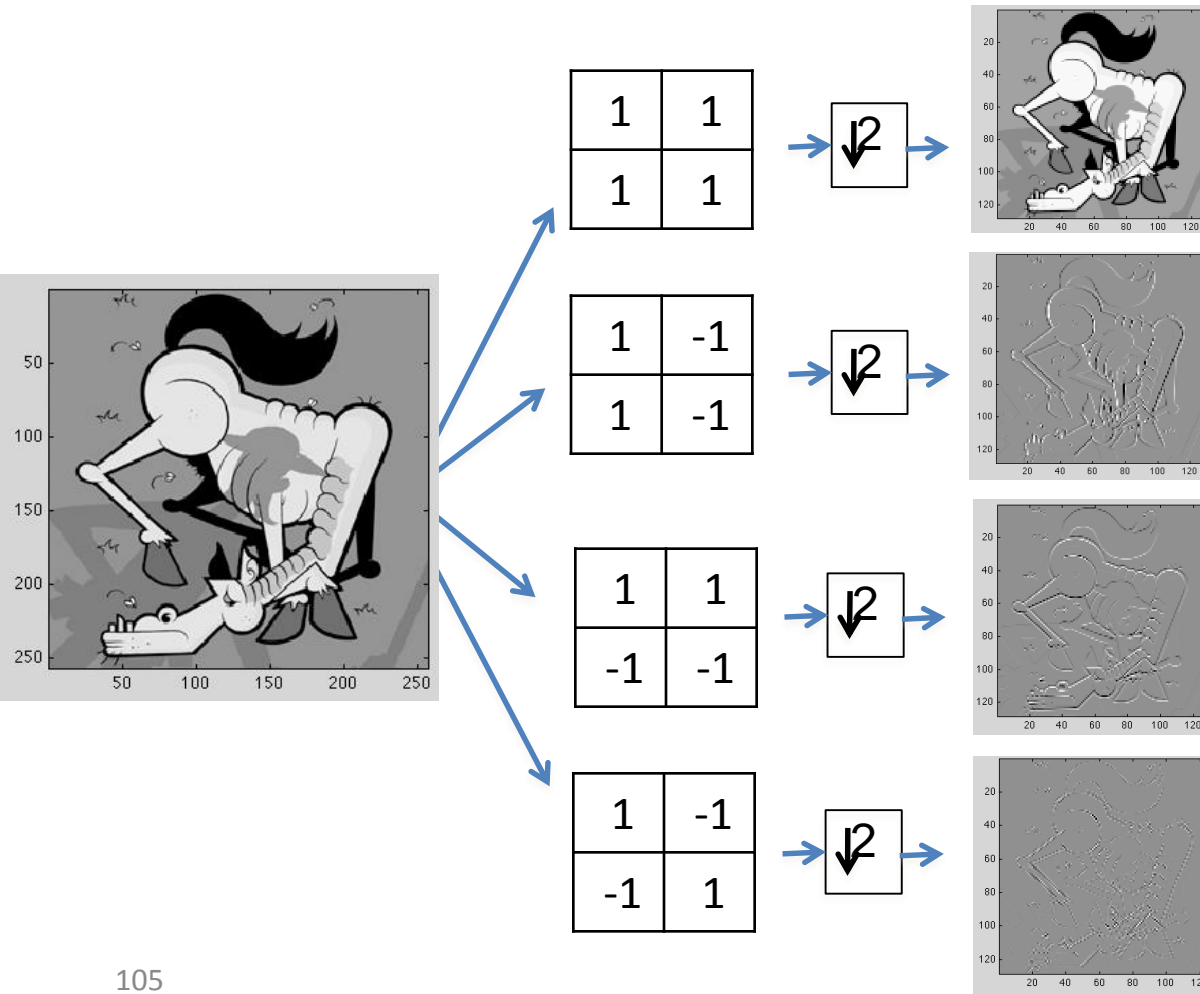
High pass diagonal

2D Haar transform




Sketch of the Fourier transform

2D Haar transform



Sketch of the Fourier transform



Pyramid cascade





Figure 4.12: Idealized diagram of the partition of the frequency plane resulting from a 4-level pyramid cascade of separable 2-band filters. The top plot represents the frequency spectrum of the original image, with axes ranging from $-\pi$ to π . This is divided into four subbands at the next level. On each subsequent level, the lowpass subband (outlined in bold) is subdivided further.

Wavelet/QMF representation



1	-1
1	-1

1	1
-1	-1

1	-1
-1	1

Same number of pixels!

Image representation

- Pixels: great for spatial resolution, poor access to frequency
- Fourier transform: great for frequency, not for spatial info
- Pyramids/filter banks: balance between spatial and frequency information

Major uses of image pyramids

- Compression
- Object detection
 - Scale search
 - Features
- Detecting stable interest points
- Registration
 - Course-to-fine

Acknowledgements

- ◆ Computer Vision A modern Approach by Frosyth
- ◆ CSCI 1430: Introduction to Computer Vision by [James Tompkin](#)
- ◆ Statistical Pattern Recognition: A Review – A.K Jain et al., PAMI (22) 2000
- ◆ Pattern Recognition and Analysis Course – A.K. Jain, MSU
- ◆ *Pattern Classification*” by Duda et al., John Wiley & Sons.
- ◆ Digital Image Processing”, Rafael C. Gonzalez & Richard E. Woods, Addison-Wesley, 2002
- ◆ Machine Vision: Automated Visual Inspection and Robot Vision”, David Vernon, Prentice Hall, 1991
- ◆ www.eu.aibo.com/
- ◆ Advances in Human Computer Interaction, Shane Pinder, InTech, Austria, October 2008
- ◆ Computer Vision A modern Approach by Frosyth