




---

---

---

---

---

---

---

---

### Previously

- Filters allow local image neighborhood to influence our description and features
  - Smoothing to reduce noise
  - Derivatives to locate contrast, gradient
- Seam carving application:
  - use image gradients to measure “interestingness” or “energy”
  - remove 8-connected seams so as to preserve image’s energy

2

---

---

---

---

---

---

---

---

### Review: Partial derivatives of an image

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

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

-1

1

-1

?

1

-1

Which shows changes with respect to x?

Slide credit: Kristen Grauman (showing filters for correlation)

3

---

---

---


---

---

---

---

---



---

---

---


---

---

---

---

---



---

---

---

---

---

---

---

---

### Today

- Edge detection and matching
  - process the image gradient to find curves/contours
  - comparing contours
- Binary image analysis
  - blobs and regions

6

---

---

---

---

---

---

---

---

## Edge detection

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

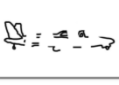
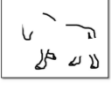




Figure from D. Lowe

Figure from J. Shotton et al., PAMI 2007

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

7

Slide credit: Kristen Grauman

---

---

---

---



---

---

---

---

## Gradients -> edges



Primary edge detection steps:

1. Smoothing: suppress noise
2. Edge enhancement: filter for contrast
3. Edge localization

Determine which local maxima from filter output are actually edges vs. noise

- Threshold, Thin

8

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

## 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)

9

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

Original image



Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

---

---

Gradient magnitude image



---

---

---

---

---

---

---

---

---

---

Thresholding gradient with a lower threshold



Slide credit: Kristen Grauman

---

---

---

---

---

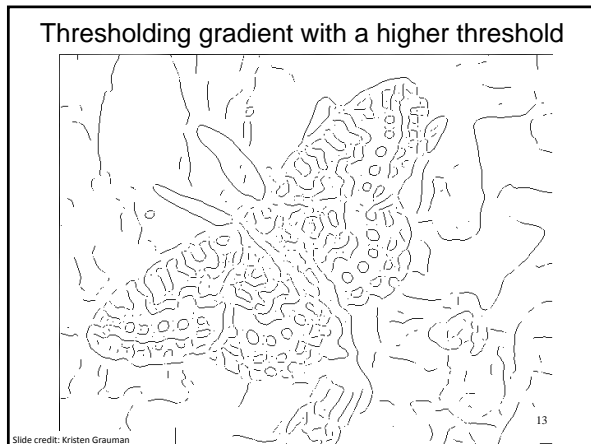
---

---

---

---

---




---

---

---

---

---

---

---

---

### Canny edge detector

- Filter image with derivative of Gaussian
- Find magnitude and orientation of gradient
- **Non-maximum suppression:**
  - Thin 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`

Slide credit: David Lowe, Fei-Fei Li

14

---

---

---

---

---

---

---

---

### The Canny edge detector

original image (Lena)

Slide credit: Steve Seltz

15

---

---

---

---


---

---

---

---

### The Canny edge detector



gradient magnitude

Slide credit: Kristen Grauman 16

---

---

---

---


---

---

---

---

### Compute Gradients



X-Derivative of Gaussian   Y-Derivative of Gaussian   Gradient Magnitude

Slide credit: Svetlana Lazebnik 17

---

---

---

---


---

---

---

---

### The Canny edge detector



gradient magnitude

Slide credit: Kristen Grauman 18

---

---

---

---


---

---

---

---

### The Canny edge detector



thresholding

Slide credit: Kristen Grauman 19

---

---

---

---

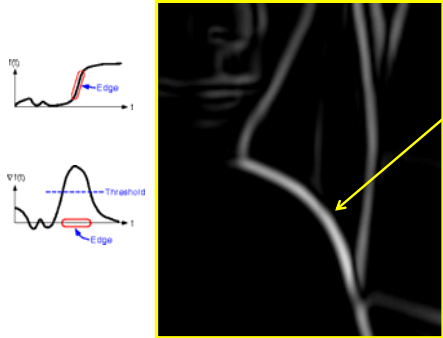
---

---

---

---

### The Canny edge detector



How to turn these thick regions of the gradient into curves?

Slide credit: Kristen Grauman 20

---

---

---

---

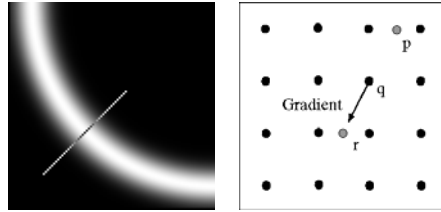
---

---

---

---

### Non-maximum suppression



Check if pixel is local maximum along gradient direction  
Select single max across width of the edge  
Requires checking interpolated pixels p and r

Slide credit: Kristen Grauman 21

---

---

---

---


---

---

---

---

### The Canny edge detector



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

thinning  
(non-maximum suppression)

Slide credit: Kristen Grauman

22

---

---

---

---

---


---

---

---

### Hysteresis thresholding

- Use a high threshold to start edge curves, and a low threshold to continue them.



Slide credit: Steve Seltz

23

---

---

---

---


---

---


---

---

### Hysteresis thresholding



original image



high threshold (strong edges)      low threshold (weak edges)      hysteresis threshold

Slide credit: Fei-Fei Li

24

---

---

---

---

---

---

---

---



## Hysteresis thresholding

[http://users.ecs.soton.ac.uk/msn/book/new\\_demo/thresholding/](http://users.ecs.soton.ac.uk/msn/book/new_demo/thresholding/)

25

---

---

---

---

---

---

---

---

## Recap: Canny edge detector

- Filter image with derivative of Gaussian
- Find magnitude and orientation of gradient
- **Non-maximum suppression:**
  - Thin 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`

26

Slide credit: David Lowe, Fei-Fei Li

---

---

---

---







---

---

---

---

## Low-level edges vs. perceived contours

		
		
<b>Background</b>	<b>Texture</b>	<b>Shadows</b>

27

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

### Low-level edges vs. perceived contours

image	human segmentation	gradient magnitude

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

Slide credit: Svetlana Lazebnik

---

---

---

---

---

---

---

---

---

---

Learn from humans which combination of features is most indicative of a "good" contour?

[D. Martin et al. PAMI 2004]

Human-marked segment boundaries

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

---

---

### pB boundary detector

Martin, Fowlkes, Malik 2004: Learning to Detect Natural Boundaries...  
<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/papers/mfm-pami-boundary.pdf>

Figure from Fowlkes

---

---

---

---

---

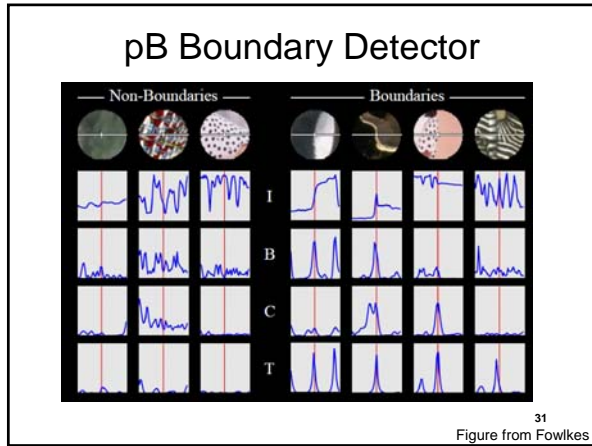
---

---

---

---

---




---

---

---

---

---

---

---

---

---

---




---

---

---

---

---

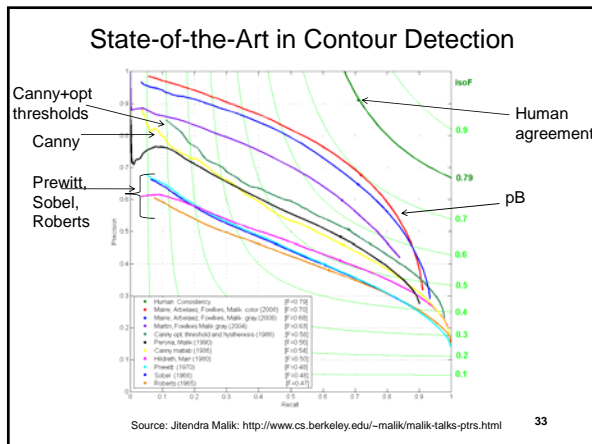
---

---

---

---

---




---

---

---

---

---

---

---


---

---

---

## Holistically-Nested Edge Detection (Xie, Tu ICCV 2015)

1. holistic image training and prediction
2. multi-scale and multi-level feature learning
3. Deeply-supervised fully-convolutional network



34

---

---

---

---

---

---


---

---

---

---

## State-of-the-Art in Contour Detection



35

---

---

---

---

---

---

---

---

---

---

## Today

- Edge detection and matching
  - process the image gradient to find curves/contours
  - comparing contours
- Binary image analysis
  - blobs and regions

Slide credit: Kristen Grauman

36

---

---

---

---

---

---

---

---

---

---




Fig. 1. Examples of two handwritten digits. In terms of pixel-to-pixel comparisons, these two images are quite different, but to the human observer, the shapes appear to be similar.

Figure from Belongie et al.

37

---

---

---

---

---

---

---

---

### Chamfer distance

- Average distance to nearest feature/edge

$$D_{chamfer}(T, I) \equiv \frac{1}{|T|} \sum_{t \in T} d_I(t)$$

$I$  = Set of edge points in image

$T$  = Set of edge points on (shifted) template

$d_I(t)$  = Minimum distance between point  $t$  and some point in  $I$

38

---

---

---

---

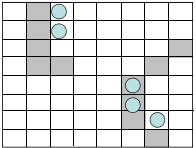
---

---

---

---

### Chamfer distance



$$D_{chamfer}(T, I) \equiv \frac{1}{|T|} \sum_{t \in T} d_I(t)$$

Slide credit: Kristen Grauman

39

---

---

---

---

---

---



---

---

## Chamfer distance

- Average distance to nearest feature

$$D_{chamfer}(T, I) \equiv \frac{1}{|T|} \sum_{t \in T} d_I(t)$$

Edge image

How is the measure different than just filtering with a mask having the shape points?

How expensive is a naïve implementation?

Slide credit: Kristen Grauman 40

---

---

---

---

---

---

---

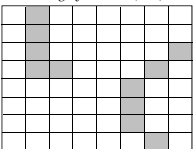
---

---

---

## Distance transform

Image features (2D)



Distance Transform

1	0	1	2	3	4	3	2
1	0	1	2	3	3	2	1
1	0	1	2	3	2	1	0
1	0	0	1	2	1	0	1
2	1	1	2	1	0	1	2
3	2	2	2	1	0	1	2
4	3	3	2	1	0	1	2
5	4	4	3	2	1	0	1

**Distance Transform** is a function  $D(x)$  that for each image pixel  $p$  assigns a non-negative number  $D(p)$  corresponding to distance from  $p$  to the nearest feature in the image  $I$

Slide credit: Yuri Boykov 41

---

---

---

---

---

---


---

---


---

---

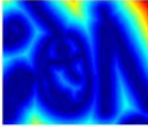
## Distance transform



original



edges



distance transform

Value at (x,y) tells how far that position is from the nearest edge point (or other binary image structure)

>> help bwdist

Slide credit: Kristen Grauman 42

---

---

---

---

---

---

---

---

---

---

### Distance transform (1D)

Two pass  $O(n)$  algorithm for 1D  $L_1$  norm

1. **Initialize:** For all  $j$   
 $D[j] \leftarrow 1_{P[j]}$  // 0 if  $j$  is in  $P$ , infinity otherwise

Image features (edges)

0 1 0 1 0 0 0 1 0

Distance transform

Slide adapted from Dan Huttonlocher

---

---

---

---

---

---

---

---

---

---

### Distance Transform (2D)

- 2D case analogous to 1D
  - Initialization
  - Forward and backward pass
    - Fwd pass finds closest above and to left
    - Bwd pass finds closest below and to right

-	1
1	0
0	1
1	-


∞	∞	∞	∞
∞	0	∞	∞
∞	∞	∞	∞
∞	∞	∞	∞

∞	∞	∞	∞
∞	0	1	∞
∞	∞	∞	∞
∞	∞	∞	∞

∞	∞	∞	∞
∞	0	1	2
∞	∞	∞	∞
∞	∞	∞	∞

2	1	2	3
1	0	1	2
1	0	1	2
2	1	2	3

Slide credit: Dan Huttonlocher

---

---

---

---

---

---

---

---




---

---

### Chamfer distance

- Average distance to nearest feature

$$D_{chamfer}(T, I) \equiv \frac{1}{|T|} \sum_{t \in T} d_I(t)$$

Edge image

Distance transform image

Slide credit: Kristen Grauman

---

---

---

---

---

---


---

---

---

---

### Chamfer distance



Edge image      Distance transform image

Fig from D. Gavrilu, DAGM 1999

46

---

---

---

---

---

---

---

---

---

---

### Chamfer distance: properties

- Sensitive to scale and rotation
- Tolerant of small shape changes, clutter
- Need large number of template shapes
- Inexpensive way to match shapes

Slide credit: Kristen Grauman

47

---

---

---

---

---

---

---

---

---

---

### Today

- Edge detection and matching
  - process the image gradient to find curves/contours
  - comparing contours
- Binary image analysis
  - blobs and regions

Slide credit: Kristen Grauman

48

---

---

---

---

---

---

---

---

---

---



### Binary images

Slide credit: Kristen Grauman

49

---

---

---

---

---

---

---

---

### Binary image analysis: basic steps

- Convert the image into binary form
  - Thresholding
- Clean up the thresholded image
  - Morphological operators
- Extract separate blobs
  - Connected components
- Describe the blobs with region properties

Slide credit: Kristen Grauman

50

---

---

---

---

---

---

---

---

### Binary images

- Two pixel values
  - Foreground and background
  - Mark region(s) of interest

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

Slide credit: Kristen Grauman

51

---

---

---

---

---

---

---

---

### Thresholding

- Grayscale -> binary mask
- Useful if object of interest's intensity distribution is distinct from background

$$F_T[i, j] = \begin{cases} 1 & \text{if } F[i, j] \geq T \\ 0 & \text{otherwise.} \end{cases}$$

$$F_T[i, j] = \begin{cases} 1 & \text{if } T_1 \leq F[i, j] \leq T_2 \\ 0 & \text{otherwise.} \end{cases}$$

$$F_T[i, j] = \begin{cases} 1 & \text{if } F[i, j] \in Z \\ 0 & \text{otherwise.} \end{cases}$$

- [Example](http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/FITZGIBBON/simplebinary.html)  
[http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL\\_COPIES/FITZGIBBON/simplebinary.html](http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/FITZGIBBON/simplebinary.html)

Slide credit: Kristen Grauman 52

---

---

---

---

---

---



---

---

### Thresholding

- Given a grayscale image or an intermediate matrix -> threshold to create a binary output.

Example: edge detection


→


Gradient magnitude
`fg_pix = find(gradient_mag > t);`

Looking for pixels where gradient is strong.

Slide adapted from Kristen Grauman 53

---

---

---

---

---

---



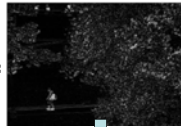
---

---

### Thresholding

- Given a grayscale image or an intermediate matrix -> threshold to create a binary output.

Example: background subtraction


-

=


Looking for pixels that differ significantly from the "empty" background.

`fg_pix = find(diff > t);`

Slide credit: Kristen Grauman 54

---

---

---

---

---

---

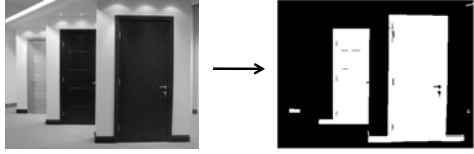
---

---

## Thresholding

- Given a grayscale image or an intermediate matrix → threshold to create a binary output.

Example: intensity-based detection



Looking for dark pixels

`fg_pix = find(im < 65);`

55

Slide credit: Kristen Grauman

---

---

---

---

---

---

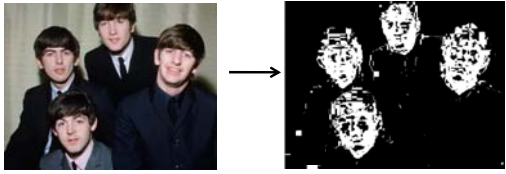
---

---

## Thresholding

- Given a grayscale image or an intermediate matrix → threshold to create a binary output.

Example: color-based detection



Looking for pixels within a certain hue range.

`fg_pix = find(hue > t1 & hue < t2);`

56

Slide credit: Kristen Grauman

---

---

---

---

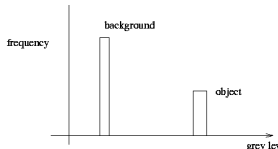
---

---

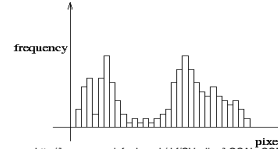
---

---

## A nice case: bimodal intensity histograms



Ideal histogram,  
light object on  
dark background



Actual observed  
histogram with  
noise

Images: [http://homepages.inf.ed.ac.uk/ibf/CVonline/LOCAL\\_COPIES/OWENS/LECT2/node3.html](http://homepages.inf.ed.ac.uk/ibf/CVonline/LOCAL_COPIES/OWENS/LECT2/node3.html)

57

---

---

---

---

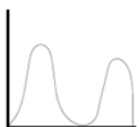
---

---

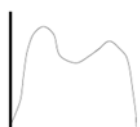
---

---


### Not so nice cases



Two distinct modes



Overlapped modes



Slide credit: Shapiro and Stockman

58

---

---

---

---

---


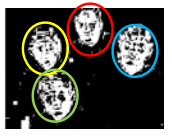
---

---

---

### Issues

- What to do with “noisy” binary outputs?
  - Holes
  - Extra small fragments
  
- How to demarcate multiple regions of interest?
  - Count objects
  - Compute further features per object

Slide credit: Kristen Grauman

59

---

---

---

---

---

---

---

---

### Morphological operators

- Change the shape of the foreground regions via intersection/union operations between a scanning structuring element and binary image
- Useful to clean up result from thresholding
- Basic operators are:
  - Dilation
  - Erosion

Slide credit: Kristen Grauman

60

---

---

---

---

---



---

---

---

### Dilation

- Expands connected components
- Grow features
- Fill holes



Before dilationAfter dilation

61

Slide credit: Kristen Grauman

---

---

---

---

---



---

---

---

### Erosion

- Erode connected components
- Shrink features
- Remove bridges, branches, noise



Before erosionAfter erosion

62

Slide credit: Kristen Grauman

---

---

---

---

---


---

---

---

### Structuring elements

- **Masks** of varying shapes and sizes used to perform morphology, for example:



- Scan mask across foreground pixels to transform the binary image

>> help strel

63

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

## Dilation vs. Erosion

At each position:

- **Dilation:** if **current pixel** is 1, then set all the output pixels corresponding to structuring element to 1.

64

---

---

---

---

---

---

---

---

## Example for Dilation

Input image 

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element 

1	1	1
---	---	---

Output Image 

1	1								
---	---	--	--	--	--	--	--	--	--

Slide credit: Adapted by Kristen Grauman from T. Moeslund

65

---

---

---

---

---

---

---

---

## Example for Dilation

Input image 

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---



Structuring Element 

1	1	1
---	---	---

Output Image 

1	1								
---	---	--	--	--	--	--	--	--	--

Slide credit: Kristen Grauman

66

---

---

---

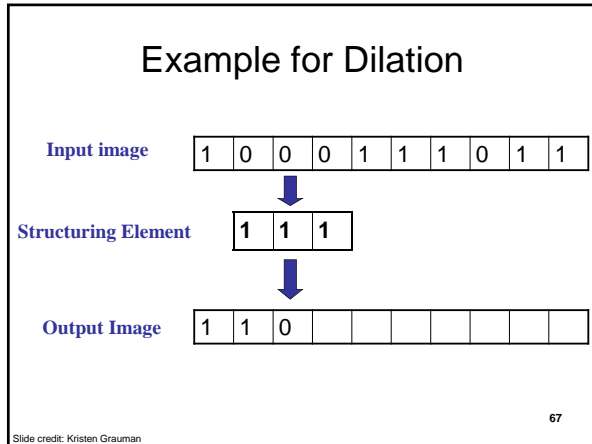
---

---

---

---

---



---

---

---

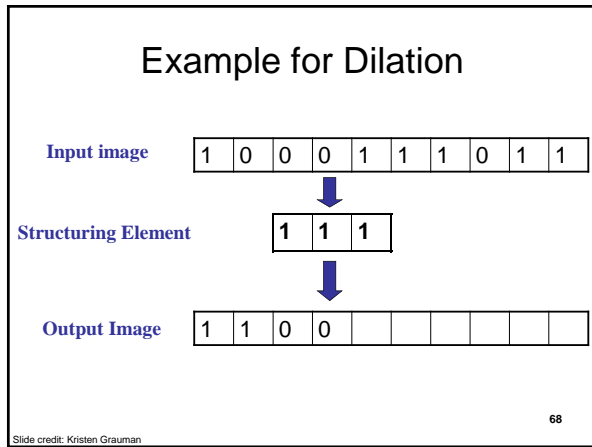
---

---

---

---

---



---

---

---

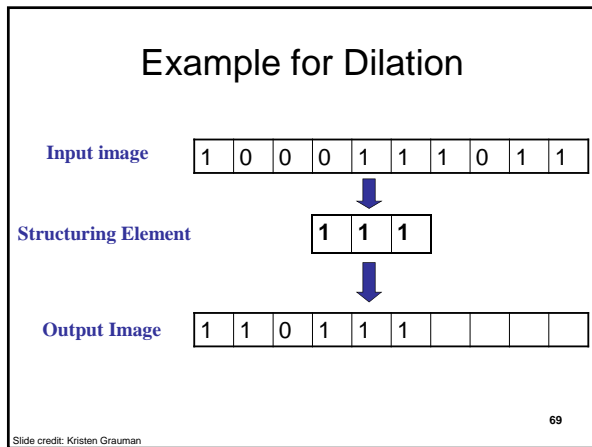
---

---

---

---

---



---

---

---

---

---

---

---

---

### Example for Dilation

**Input image**

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

↓

**Structuring Element**

1	1	1
---	---	---

↓

**Output Image**

1	1	0	1	1	1	1			
---	---	---	---	---	---	---	--	--	--

70

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

### Example for Dilation

**Input image**

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

↓

**Structuring Element**

1	1	1
---	---	---

↓

**Output Image**

1	1	0	1	1	1	1	1		
---	---	---	---	---	---	---	---	--	--

71

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

### Example for Dilation

**Input image**

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

↓

**Structuring Element**

1	1	1
---	---	---

↓

**Output Image**

1	1	0	1	1	1	1	1		
---	---	---	---	---	---	---	---	--	--

72

Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---



### Example for Dilation

**Input image**

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

↓

**Structuring Element**

1	1	1
---	---	---

↓

**Output Image**

1	1	0	1	1	1	1	1	1	1
---	---	---	---	---	---	---	---	---	---

Note that the object gets bigger and holes are filled.

>> help imdilate  
Slide credit: Kristen Grauman

---

---

---

---

---

---

---

---

### 2D example for dilation

1	1	1	1	1	1	1	1
			1	1	1	1	1
			1	1	1	1	1
			1	1	1	1	1
			1	1	1	1	1
			1	1	1	1	1
			1	1	1	1	1
			1	1	1	1	1

(a) Binary image B

1	1	1
1	1	1
1	1	1

(b) Structuring element S

1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1

(c) Dilation B @ S

Slide credit: Shapiro & Stockman

---

---

---

---

---

---

---

---

### Dilation vs. Erosion

At each position:

- **Dilation:** if **current pixel** is 1, then set all the output pixels corresponding to structuring element to 1.
- **Erosion:** if **every pixel** under the structuring element is 1, then set the output pixel corresponding to the current pixel to 1.

75

---

---

---

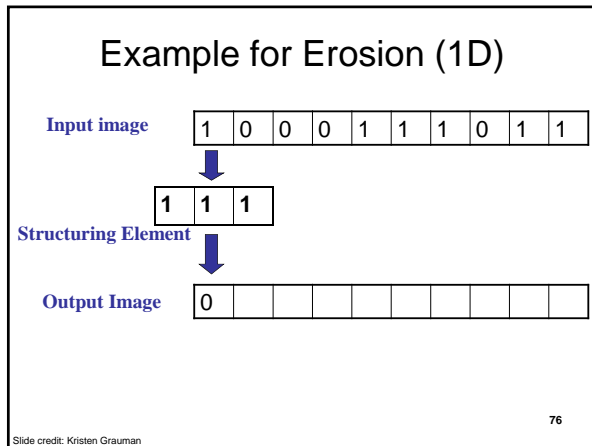
---

---

---

---

---



---

---

---

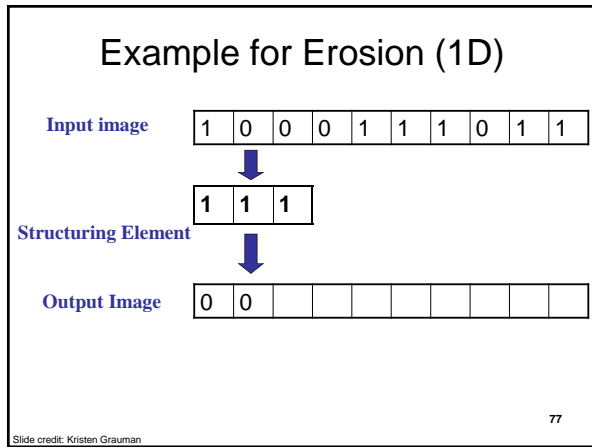
---

---

---

---

---



---

---

---

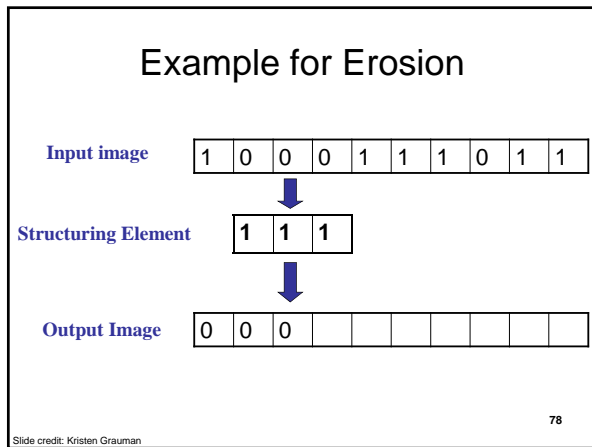
---

---

---

---

---



---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**    1 0 0 0 1 1 1 0 1 1

↓

**Structuring Element**    1 1 1

↓

**Output Image**    0 0 0 0

Slide credit: Kristen Grauman    79

---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**    1 0 0 0 1 1 1 0 1 1

↓

**Structuring Element**    1 1 1

↓

**Output Image**    0 0 0 0 0

Slide credit: Kristen Grauman    80

---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**    1 0 0 0 1 1 1 0 1 1

↓

**Structuring Element**    1 1 1

↓

**Output Image**    0 0 0 0 0 1

Slide credit: Kristen Grauman    81

---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**    1 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | 1

↓

**Structuring Element**    1 | 1 | 1

↓

**Output Image**    0 | 0 | 0 | 0 | 0 | 1 | 0 |    |    |    |

Slide credit: Kristen Grauman    82

---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**    1 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | 1

↓

**Structuring Element**    1 | 1 | 1

↓

**Output Image**    0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 |    |    |

Slide credit: Kristen Grauman    83

---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**    1 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | 1

↓

**Structuring Element**    1 | 1 | 1

↓

**Output Image**    0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 |    |

Slide credit: Kristen Grauman    84

---

---

---

---

---

---

---

---

### Example for Erosion

**Input image**

1	0	0	0	1	1	1	0	1	1
---	---	---	---	---	---	---	---	---	---

↓

**Structuring Element**

1	1
---	---

↓

**Output Image**

0	0	0	0	0	1	0	0	0	1
---	---	---	---	---	---	---	---	---	---

Note that the object gets smaller

Slide credit: Kristen Grauman >> help imerode 85

---

---

---

---

---

---


---

---


---

---

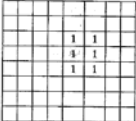
### 2D example for erosion



(a) Binary image **B**



(b) Structuring element **S**



(d) Erosion **B ⊖ S**

Slide credit: Shapiro & Stockman 86

---

---

---

---

---

---

---


---

---

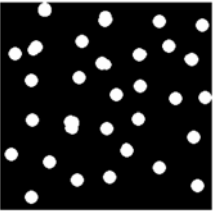
---

### Opening

- Erode, then dilate
- Remove small objects, keep original shape



Before opening



After opening

Slide credit: Kristen Grauman 87

---

---

---

---

---

---

---


---

---


---

### Closing

- Dilate, then erode
- Fill holes, but keep original shape



Before closing



After closing

Applet: <http://biqwww.epfl.ch/demo/morpho/start.php> 88

Slide credit: Kristen Grauman

---

---

---

---

---

---

---



---

---

---

### Issues

- What to do with “noisy” binary outputs?
  - Holes
  - Extra small fragments
- How to demarcate multiple regions of interest?
  - Count objects
  - Compute further features per object

Slide credit: Kristen Grauman 89

---

---

---

---

---

---

---

---

---



---

### Connected components

- Identify distinct regions of “connected pixels”

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

1	1	0	1	1	1	0	2
1	1	0	1	0	1	0	2
1	1	1	1	0	0	0	2
0	0	0	0	0	0	0	2
3	3	3	3	0	4	0	2
0	0	0	3	0	4	0	2
5	5	0	3	0	0	0	2
5	5	0	3	0	2	2	2

c) binary image and labeling, expanded for viewing

>> L = bwlabel(BW,conn) 90

---

---

---

---

---

---

---

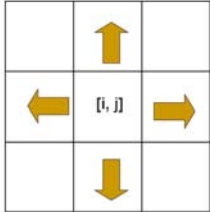
---

---

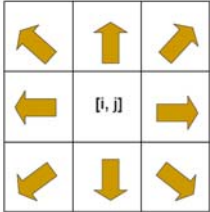
---

### Connectedness

- Defining which pixels are considered neighbors



4-connected



8-connected

Slide credit: Chaitanya Chandra 91

---

---

---

---

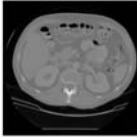
---


---

---


---

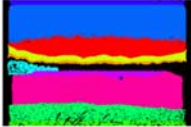
### Connected components





connected components of 1's from thresholded image





connected components of cluster labels

Slide credit: Pinar Duygulu 92

---

---

---

---

---


---

---


---


### Region properties

- Given connected components, can compute simple features per blob, such as:
  - Area (num pixels in the region)
  - Centroid (average x and y position of pixels in the region)
  - Bounding box (min and max coordinates)



A1=200  
A2=170





Slide credit: Kristen Grauman 93

---

---

---

---

---

---

---

---

### Binary image analysis: basic steps (recap)

- Convert the image into binary form
  - Thresholding
- Clean up the thresholded image
  - Morphological operators
- Extract separate blobs
  - Connected components
- Describe the blobs with region properties

Slide credit: Kristen Grauman

94

---

---

---

---

---

---

---

---

### Matlab

```
• L = bwlabel (BW,8);  
• STATS = regionprops(L,PROPERTIES) ;  
  - 'Area'  
  - 'Centroid'  
  - 'BoundingBox'  
  - 'Orientation', ...  
• IM2 = imerode(IM,SE);  
• IM2 = imdilate(IM,SE);  
• IM2 = imclose(IM, SE);  
• IM2 = imopen(IM, SE);
```

Slide adapted from Kristen Grauman

95

---

---

---

---


---

---

---

---

### Example using binary image analysis: segmentation of a liver



Slide credit: Li Shen

Application by Jie Zhu, Cornell University

96

---

---

---

---

---

---

---

---



## Binary images

- Pros
  - Can be fast to compute, easy to store
  - Simple processing techniques available
  - Lead to some useful compact shape descriptors
- Cons
  - Hard to get “clean” silhouettes
  - Noise common in realistic scenarios
  - Can be too coarse of a representation

Slide credit: Kristen Grauman

97

---

---

---

---

---


---

---

---

## Summary

- Operations, tools
  - Derivative filters
  - Smoothing, morphology
  - Thresholding
  - Connected components
  - Matching filters
  - Histograms
- Features, representations
  - Edges, gradients
  - Blobs/regions
  - Local patterns
  - Textures (next)
  - Color distributions



98

---

---

---

---

---

---

---

---

## Coming up

- Texture
  - Read Szeliski 10.5



99

---

---

---

---

---

---

---

---

Questions?  
See you Thursday!

100

---

---

---

---

---

---

---

---