# Edge Detection



**Computer** Vision

Jia-Bin Huang, Virginia Tech

Many slides from D. Hoiem and K. Grauman

## Administrative Stuffs

- HW 1 posted, due 11:59 PM Sept 19
  - Submission through Canvas
- Questions about HW?
  - Ask me/TA after classes
  - Post questions on Piazza (no emails)
  - Attend office hours

#### Previous three classes: Image Filtering **Template matching** Spatial domain Find a template in an image Smoothing, sharpening, measuring texture FFT Inverse FFT FFT Image pyramid Coarse-to-fine search Multi-scale detection X **Frequency domain** Denoising, sampling, image compression

# Today's class

Detecting edges

• Finding straight lines

• Binary image analysis



# Why finding edges is important

- Cues for 3D shape
- Group pixels into objects or parts
- Guiding interactive image editing
- Recover geometry and viewpoint





point

point

# Origin of Edges



#### Edges are caused by a variety of factors















# Characterizing edges

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



#### Intensity profile





#### With a little Gaussian noise





# Effects of noise

Consider a single row or column of the image

 Plotting intensity as a function of position gives a signal



Where is the edge?

### Effects of noise

- Difference filters respond strongly to noise
  - Image noise results in pixels that look very different from their neighbors
  - Generally, the larger the noise the stronger the response
- What can we do about it?

#### Solution: smooth first



• To find edges, look for peaks in

 $\frac{d}{dx}(f * g)$ 

Source: S. Seitz

#### Derivative theorem of convolution

- Differentiation is convolution, and convolution is associative:  $\frac{d}{dx}(f \ast g) = f \ast \frac{d}{dx}g$
- This saves us one operation:



Source: S. Seitz

#### Derivative of Gaussian filter





• Is this filter separable?

#### Tradeoff between smoothing and localization



1 pixel

3 pixels

7 pixels

• Smoothed derivative removes noise, but blurs edge. Also finds edges at different "scales".

## Designing an edge detector

- Criteria for a good edge detector:
  - Good detection:
    - find all real edges, ignoring noise or other artifacts
  - Good localization
    - detect edges as close as possible to the true edges
    - return one point only for each true edge point
- Cues of edge detection
  - Differences in color, intensity, or texture across the boundary
  - Continuity and closure
  - High-level knowledge

## Canny edge detector

- This is probably the most widely used edge detector in computer vision
- Theoretical model: step-edges corrupted by additive Gaussian noise
- Canny has shown that the first derivative of the Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio* and localization

<u>J Canny</u> - IEEE Transactions on pattern analysis and machine ..., 1986 - ieeexplore.ieee.org Abstract-This paper describes a computational approach to edge detection. The success of the approach depends on the definition of a comprehensive set of goals for the computation of edge points. These goals must be precise enough to delimit the desired behavior of the ... Cited by 24439 Related articles All 24 versions Import into BibTeX Save More

J. Canny, <u>A Computational Approach To Edge Detection</u>, IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986. Source: L. Fei-Fei

#### Example



input image ("Lena")

#### Derivative of Gaussian filter



### Compute Gradients (DoG)



X-Derivative of Gaussian

Y-Derivative of Gaussian

Gradient Magnitude

# Get Orientation at Each Pixel

- Threshold at minimum level
- Get orientation



theta = atan2(-gy, gx)

# Non-maximum suppression for each orientation



At q, we have a maximum if the value is larger than those at both p and at r. Interpolate to get these values.



Source: D. Forsyth

#### **Bilinear Interpolation**



http://en.wikipedia.org/wiki/Bilinear\_interpolation

# Sidebar: Interpolation options

- imx2 = imresize(im, 2, interpolation\_type)
- 'nearest'
  - Copy value from nearest known
  - Very fast but creates blocky edges
- 'bilinear'
  - Weighted average from four nearest known pixels
  - Fast and reasonable results
- 'bicubic' (default)
  - Non-linear smoothing over larger area
  - Slower, visually appealing, may create negative pixel values

Examples from http://en.wikipedia.org/wiki/Bicubic interpolation



#### **Before Non-max Suppression**



#### After non-max suppression



## Hysteresis thresholding

- Threshold at low/high levels to get weak/strong edge pixels
- Do connected components, starting from strong edge pixels



# 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.



#### Final Canny Edges



## Canny edge detector

- 1. Filter image with x, y derivatives of Gaussian
- 2. Find magnitude and orientation of gradient
- 3. Non-maximum suppression:
  - Thin multi-pixel wide "ridges" down to single pixel width
- 4. Thresholding and linking (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')

#### Effect of $\sigma$ (Gaussian kernel spread/size)



The choice of  $\sigma$  depends on desired behavior

- large  $\sigma$  detects large scale edges
- small  $\sigma$  detects fine features

#### Learning to detect boundaries

image

human segmentation

gradient magnitude



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



Martin, Fowlkes, Malik 2004: Learning to Detection Natural Boundaries...

Figure from Fowlkes

### pB Boundary Detector



Figure from Fowlkes



### Results



### Human (0.95)





### Results









Human (0.95)











#### For more:

http://www.eecs.berkeley.edu/Research/Projects/CS /vision/bsds/bench/html/108082-color.html

# Global pB boundary detector



Figure from Fowlkes

http://research.microsoft.com/pubs/202540/DollarICCV13edges.pdf

Solution

- Train structured random forests to split data into patches with similar boundaries based on features
- Predict boundaries at patch level, rather than pixel level, and aggregate (average votes)

# Forests (Dollar and Zitnick ICCV 2013)

Edge Detection with Structured Random

- Goal: quickly predict whether each pixel is an edge
- Insights
  - Predictions can be learned from training data
  - Predictions for nearby pixels should not be • independent







### Edge Detection with Structured Random Forests - Algorithm

- 1. Extract overlapping 32x32 patches at three scales
- 2. Features are pixel values and pairwise differences in feature maps (LUV color, gradient magnitude, oriented gradient)
- 3. Predict *T* boundary maps in the central 16x16 region using *T* trained decision trees
- 4. Average predictions for each pixel across all patches





# Edge Detection with Structured Random Forests - Results

#### **BSDS 500**

NYU Depth dataset edges

	ODS	OIS	AP	FPS
Human	.80	.80	-	-
Canny	.60	.64	.58	15
Felz-Hutt [11]	.61	.64	.56	10
Hidayat-Green [16]	.62†	-	-	20
BEL [9]	.66†	-	-	1/10
gPb + GPU [6]	.70†	-	-	1/2 <sup>‡</sup>
gPb [1]	.71	.74	.65	1/240
gPb-owt-ucm [1]	.73	.76	.73	1/240
Sketch tokens [21]	.73	.75	.78	1
SCG [31]	.74	.76	.77	1/280
SE-SS, T=1	.72	.74	.77	60
SE-SS, $T=4$	.73	.75	.77	30
SE-MS, $T=4$	.74	.76	.78	6
	-			-

	ODS	OIS	AP	FPS
gPb [1] (rgb)	.51	.52	.37	1/240
SCG [31] (rgb)	.55	.57	.46	1/280
SE-SS (rgb)	.58	.59	.53	30
SE-MS (rgb)	.60	.61	.56	6
gPb [1] (depth)	.44	.46	.28	1/240
SCG [31] (depth)	.53	.54	.45	1/280
SE-SS (depth)	.57	.58	.54	30
SE-MS (depth)	.58	.59	.57	6
gPb [1] (rgbd)	.53	.54	.40	1/240
SCG [31] (rgbd)	.62	.63	.54	1/280
SE-SS (rgbd)	.62	.63	.59	25
SE-MS (rgbd)	.64	.65	.63	5

### Edge Detection with Structured Random Forests



### Crisp Boundary Detection using Pointwise Mutual Information (Isola et al. ECCV 2014)



Pixel combinations that are unlikely to be together are edges



http://web.mit.edu/phillipi/www/publications/crisp\_boundaries.pdf

### Crisp Boundary Detection using Pointwise Mutual Information

Algorithm	ODS	OIS	$\mathbf{AP}$
Canny [14]	0.60	0.63	0.58
Mean Shift [36]	0.64	0.68	0.56
NCuts [37]	0.64	0.68	0.45
Felz-Hutt [38]	0.61	0.64	0.56
gPb [1]	0.71	0.74	0.65
gPb-owt-ucm [1]	0.73	0.76	0.73
SCG [9]	0.74	0.76	0.77
Sketch Tokens [7]	0.73	0.75	0.78
SE [8]	0.74	0.76	0.78
$Our method-SS, \ color \ only$	0.72	0.75	0.77
Our method $-$ SS	0.73	0.76	0.79
Our method – MS	0.74	0.77	0.78



# Holistically-Nested Edge Detection



#### Xie and Tu, Holistically-Nested Edge Detection, ICCV 2015

### Holistically-Nested Edge Detection



Xie and Tu, Holistically-Nested Edge Detection, ICCV 2015

# State of edge detection

- Local edge detection is mostly solved
  - Intensity gradient, color, texture
- Work on RGB-D edge detection is currently more active
- Some methods take into account longer contours, but could probably do better
- Often used in combination with object detectors or region classifiers

# Finding straight lines





# Finding line segments using connected components

- 1. Compute canny edges
  - Compute: gx, gy (DoG in x,y directions)
  - Compute: theta = atan(gy / gx)
- 2. Assign each edge to one of 8 directions
- 3. For each direction d, get edgelets:
  - find connected components for edge pixels with directions in {d-1, d, d+1}
- Compute straightness and theta of edgelets using eig of x,y
   2<sup>nd</sup> moment matrix of their points

5. Threshold on straightness, store segment

### Canny lines $\rightarrow \dots \rightarrow$ straight edges





### **Binary** images









#### A mailwoodboll peak for the legan sea h, shape memory allogs

Yes young the

To Default Morth is the statute is previously independent of a general statute to be previously independent of the statute of

#### Marcheology,

The global cancer are an efficient of spectral spectra of the set  $\mu$  , where the domain and the spectra of spectra of the sp

### 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

### **Binary** images

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





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

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

$$F_{T}[i, j] = \begin{cases} 1 & \text{if } T_{1} \le F[i, j] \le T_{2} \\ 0 & \text{otherwise.} \end{cases}$$

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

• Example

http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL\_COPIES/FITZGIBBON/simplebin ary.html

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 credit: Kristen Grauman

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.



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

### Example: intensity-based detection







Looking for dark pixels

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

### Example: color-based detection





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

Looking for pixels within a certain hue range.

Slide credit: Kristen Grauman

# A nice case: bimodal intensity histograms



Slide credit: Kristen Grauman

### Not so nice cases



Two distinct modes





Overlapped modes

Slide credit: Shapiro and Stockman

### 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





# 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

# Dilation

- Expands connected components
- Grow features
- Fill holes



**Before dilation** 

After dilation

### Erosion

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



Before erosion

After erosion

# 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

# Dilation vs. Erosion

At each position:

• **Dilation**: if current pixel is foreground, OR the structuring element with the input image.


Slide credit: Adapted by Kristen Grauman from T. Moeslund





















Note that the object gets bigger and holes are filled.

>> help imdilate

# 2D example for dilation



# Dilation vs. Erosion

At each position:

- **Dilation**: if **current pixel** is foreground, OR the structuring element with the input image.
- Erosion: if every pixel under the structuring element's nonzero entries is foreground, OR the current pixel with S.

# Example for Erosion (1D)



# Example for Erosion (1D)



















Note that the object gets smaller

>> help imerode

### 2D example for erosion



# Opening

- First erode, then dilate
- Remove small objects, keep original shape



**Before opening** 



After opening

# Closing

- First dilate, then erode
- Fill holes, but keep original shape





**Before closing** 

After closing

Applet: <u>http://bigwww.epfl.ch/demo/jmorpho/start.php</u>

#### Morphology operators on grayscale images

- Dilation and erosion typically performed on binary images.
- If image is grayscale:
  - Diation: take the neighborhood MAX
  - Erosion: take the neighborhood MIN



original

dilated

eroded

#### 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





### 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

a) binary image

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	<b>5</b>	0	3	0	0	0	2
5	5	0	3	0	2	2	2

b) connected components labeling





c) binary image and labeling, expanded for viewing

Slide credit: Shapiro and Stockman

#### Connectedness

• Defining which pixels are considered neighbors



[i, j]

4-connected

8-connected

Slide credit: Chaitanya Chandra

# Connected components

- We'll consider a sequential algorithm that requires only 2 passes over the image.
- Input: binary image
- **Output**: "label" image, where pixels are numbered per their component
- Note: foreground here is denoted with black pixels.



#### Connected components





connected components of 1's from thresholded image





connected components of cluster labels

Slide credit: Pinar Duygulu

#### **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)
  - Circularity (ratio of mean dist. to centroid over std)



A2=170 A1 = 200





#### Circularity

a second measure uses variation off of a circle circularity(2):

$$C_2 = \frac{\mu_R}{\sigma_R}$$

where  $\mu_R$  and  $\sigma_R^2$  are the mean and variance of the distance from the centroid of the shape to the boundary pixels  $(r_k, c_k)$ . mean radial distance:

$$\mu_R = \frac{1}{K} \sum_{k=0}^{K-1} \| (r_k, c_k) - (\bar{r}, \bar{c}) \|$$

variance of radial distance:

$$\sigma_R^2 = \frac{1}{K} \sum_{k=0}^{K-1} \left[ \| (r_k, c_k) - (\bar{r}, \bar{c}) \| - \mu_R \right]^2$$



[Haralick]

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

#### Matlab

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

#### Example using binary image analysis: OCR

Digitizing Books One Word at a Time

- HOME
- WHAT IS reCAPTCHA
  - DIGITIZATION ACCURACY WHAT IS A CAPTCHA SECURITY

**Re**CAPTCHA<sup>™</sup>

- GET reCAPTCHA
- MY ACCOUNT
- EMAIL PROTECTION
- RESOURCES





The words above come from scanned books. By typing them, you help to digitize old texts.

reCAPTCHA is a free CAPTCHA service that helps to digitize books, newspapers and old time radio shows. Check out <u>our paper</u> in Science about it (or read more below).

A <u>CAPTCHA</u> is a program that can tell whether its user is a human or a computer. You've probably seen them — colorful images with distorted text at the bottom of Web registration forms. CAPTCHAs are used by many websites to prevent abuse from "bots," or automated programs usually written to generate spam. No computer program can read distorted text as well as humans can, so bots cannot navigate sites protected by CAPTCHAs.

#### [Luis von Ahn et al. http://recaptcha.net/learnmore.html]

# Example using binary image analysis: segmentation of a liver



Slide credit: Li Shen

Application by Jie Zhu, Cornell University
#### Example using binary image analysis: Bg subtraction + blob detection



University of Southern California http://iris.usc.edu/Projects/detect/detection-result.html

Slide credit: Kristen Grauman

## 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
  - Not 3d

## OpenCV Demo

# Things to remember

- Canny edge detector = smooth → derivative → thin → threshold → link
- Pb: learns weighting of gradient, color, texture differences More recent learning approaches give at least as good accuracy and are faster
- Straight line detector =

canny + gradient orientations  $\rightarrow$  orientation binning  $\rightarrow$  linking  $\rightarrow$  check for straightness

• Binary image analysis thresholding, dialation, erosion



#### Next classes: Correspondence and Alignment

- Detecting interest points
- Tracking points
- Object/image alignment and registration I(x,y,t)
  - Aligning 3D or edge points
  - Object instance recognition
  - Image stitching







