Lecture 5 Detecting Lines







Administrative

A1 is out

- It is graded
- Due April 18th







So far: discrete derivatives in 3 ways



Raymond Yu

Lecture 5 - 3

So far: Designing filters that perform differentiation

• Using Backward differentiation:

$$g[n,m] = f[n,m] - f[n,m-1]$$

• Using Forward differentiation:

$$g[n,m] = f[n,m+1] - f[n,m]$$

Using Central differentiation:

Raymond Yu

$$g[n,m] = f[n,m+1] - f[n,m-1]$$

Lecture 5 - 4

So far: Calculating gradient magnitude and direction

Given function f[n, m]

Gradient filter
$$\nabla f[n,m] = \begin{bmatrix} \frac{df}{dn} \\ \frac{df}{dm} \end{bmatrix} = \begin{bmatrix} f_n \\ f_m \end{bmatrix}$$

Gradient magnitude $|\nabla f[n,m]| = \sqrt{f_n^2 + f_m^2}$ Gradient direction $\theta = \tan^{-1}(\frac{f_m}{f_n})$

Raymond Yu

Lecture 5 - 5

Today's agenda

- Edge detector with noisy images
- Sobel Edge detector
- Canny edge detector
- Hough Transform

Optional reading: Szeliski, Computer Vision: Algorithms and Applications, 2nd Edition Sections 7.1, 8.1.4

Raymond Yu

Lecture 5 - 6

Today's agenda

- Edge detector with noisy images
- Sobel Edge detector
- Canny edge detector
- Hough Transform

Optional reading: Szeliski, Computer Vision: Algorithms and Applications, 2nd Edition Sections 7.1, 8.1.4

Raymond Yu

Lecture 5 - 7

Characterizing edges

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



Raymond Yu

Lecture 5 - 8

Intensity profile





Raymond Yu

Q. What will happen if we use this edge detector on a noisy pixels?





Raymond Yu



• Consider a single row or column of the image

Plotting intensity as a function of position gives a signal





Lecture 5 - 1^{Spurce: S. Seitz}



Raymond Yu

• Consider a single row or column of the image

Plotting intensity as a function of position gives a signal



Lecture 5 - 12 rce: S. Seitz

- Differentiation filters respond strongly to noise
 - Image noise results in pixels that look very different from their neighbors
 - Generally, the larger the noise the larger the gradient
- Q. What is a potential quick fix for noisy images?



- Differentiation 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
- Q. What is a potential quick fix for noisy images?
- Smoothing the image should help, by forcing pixels different to their neighbors (=noise pixels?) to look more like neighbors

Raymond Yu



Smoothing with different filters

• Mean smoothing

$$\frac{1}{3} \begin{bmatrix} 1\\1\\1 \end{bmatrix} \qquad \qquad \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}$$

• Gaussian (smoothing * derivative)

T

$$\frac{1}{4} \begin{bmatrix} 1\\2\\1 \end{bmatrix} \qquad \frac{1}{4} \begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$



Lecture 5 - 15

Slide creck: Strive 14, 2025

Smoothing with different filters

Mean

Gaussian



5x5

3x3

7x7

April 14, 2025

Raymond Yu

Solution: input function

Let's look at a single image row:









Solution: smooth first



Raymond Yu

Lecture 5 -

Solution: smooth first

To find edges, look for peaks in

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



Raymond Yu

Lecture 5 -

Derivative theorem of convolution

• This theorem gives us a very useful property:



Raymond Yu

Lecture 5 - 20

Derivative of a gaussian (DoG)

• This theorem gives us a very useful property:



Raymond Yu



Derivative of a gaussian (DoG)

• This theorem gives us a very useful property:

$$\frac{d}{dx}(f*h) = f*(\frac{d}{dx}h)$$

• This saves us one operation:

We can precompute:



April 14 2025

Raymond Yu

Derivative of Gaussian filter (central derivative)



2D-gaussian

x - derivative

Raymond Yu



Derivative of Gaussian filter along x and y directions



Raymond Yu

Lecture 5 - 24

Derivative of Gaussian filter







Tradeoff between smoothing at different scales



1 pixel3 pixels7 pixels

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

Raymond Yu



Designing an edge detector

- Criteria for an "optimal" edge detector:
 - Good detection: the optimal detector must minimize the probability of false positives (detecting spurious edges caused by noise), as well as that of false negatives (missing real edges)







Designing an edge detector

- Criteria for an "optimal" edge detector:
 - Good detection: the optimal detector must minimize the probability of false positives (detecting spurious edges caused by noise), as well as that of false negatives (missing real edges)
 - Good localization: the edges detected must be as close as possible to the true edges







Designing an edge detector

- Criteria for an "optimal" edge detector:
 - Good detection: the optimal detector must minimize the probability of false positives (detecting spurious edges caused by noise), as well as that of false negatives (missing real edges)
 - Good localization: the edges detected must be as close as possible to the true edges
 - Single response: the detector must return one point only for each true edge point; that is, minimize the number of local maxima around the true edge



Raymond Yu

Lecture 5 - 29

Today's agenda

- Edge detector with noisy images
- Sobel Edge detector
- Canny edge detector
- Hough Transform

Optional reading: Szeliski, Computer Vision: Algorithms and Applications, 2nd Edition Sections 7.1, 8.1.4

Raymond Yu



Sobel Operator

- uses two 3×3 kernels which are convolved with the original image to calculate approximations of the derivatives
- one for horizontal changes, and one for vertical

$$\mathbf{G}_x = egin{bmatrix} +1 & 0 & -1 \ +2 & 0 & -2 \ +1 & 0 & -1 \end{bmatrix} \qquad \mathbf{G}_y = egin{bmatrix} +1 & +2 & +1 \ 0 & 0 & 0 \ -1 & -2 & -1 \end{bmatrix}$$

Raymond Yu

Lecture 5 - 31

Sobel Operation

• Smoothing + differentiation

$$\mathbf{G}_{x} = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \begin{bmatrix} +1 & 0 & -1 \end{bmatrix}$$

Gaussian smoothing differentiation



Sobel Operation

• Magnitude:

$$\mathbf{G}=\sqrt{{\mathbf{G}_x}^2+{\mathbf{G}_y}^2}$$

• Angle or direction of the gradient:

$$oldsymbol{\Theta} = ext{atan}igg(rac{\mathbf{G}_y}{\mathbf{G}_x}igg)$$





Sobel Filter example

Step 1: Calculate the gradient magnitude at every pixel location.

Step 2: Threshold the values to generate a binary image



Raymond Yu

Lecture 5 - 34

Sobel Filter Problems



- Poor Localization (Detects multiple adjacent edges)
- Thresholding value favors certain directions over others
 Can miss diagonal edges more than horizontal or vertical edges
 False negatives

Raymond Yu

Lecture 5 - 35

What we will learn today

- Edge detector with noisy images
- Sobel Edge detector
- Canny edge detector
- Hough Transform




So far: A simple edge detector

• This theorem gives us a very useful property:

$$\frac{d}{dx}(f*h) = f*(\frac{d}{dx}h)$$

• This saves us one operation:

We can precompute:



April 14 2025

Raymond Yu

Canny edge detector

Raymond Yu

- This is probably the most widely used edge detector in computer vision
- Theoretical model: optimal edge detection when pixels are corrupted by additive Gaussian noise
- Theory shows that first derivative of the Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio*





April 14, 2025

Canny edge detector

- 1. Suppress Noise
- 2. Compute gradient magnitude and direction
- 3. Apply Non-Maximum Suppression
 - \circ Assures minimal response
- 4. Use hysteresis and connectivity analysis to detect edges



Example



• original image

Raymond Yu

Lecture 5 - 40

Canny edge detector

- 1. Suppress Noise
- 2. Compute gradient magnitude and direction
- 3. Apply Non-Maximum Suppression
 - Assures minimal response
- 4. Use hysteresis and connectivity analysis to detect edges







Derivative of Gaussian filter



Raymond Yu

Lecture 5 - 42

Compute gradients (DoG)



X-Derivative of Gaussian

Y-Derivative of Gaussian

Gradient Magnitude

Raymond Yu



Get orientation at each pixel



$$oldsymbol{\Theta} = ext{atan}igg(rac{\mathbf{G}_y}{\mathbf{G}_x}igg)$$

Raymond Yu



Compute gradients (DoG)



Raymond Yu

Lecture 5 - 45

Canny edge detector

- 1. Suppress Noise
- 2. Compute gradient magnitude and direction
- 3. Apply Non-Maximum Suppression
 - $\circ~$ Assures minimal response
- 4. Use hysteresis and connectivity analysis to detect edges





Non-maximum suppression

- Assumption we make: An edge occurs where gradient is maximum
 - Even if their magnitude is above the threshold
- Suppress non-maxima neighboring edges
 - Only suppress edges that are in the same direction nearby
 - Don't suppress other edges







Intuition behind non-maximum suppression



Let's assume that out of the points:

p,

q,

r

r q p

q has the largest gradient. Then p and r are not real edges and should be *suppressed*

Raymond Yu

Lecture 5 - 48

What the output looks like after Non-max Suppression



Before After

Raymond Yu

Lecture 5 - 49

What if $p = [n_1, m_1]$ or $r = [n_2, m_2]$, is not a pixel location



q is a maximum if the value is larger than those at both p and at r.

How should we calculate magnitude at p and r?

April 14, 2025



Raymond Yu

What if $p = [n_1, m_1]$ or $r = [n_2, m_2]$, is not a pixel location



q is a maximum if the value is larger than those at both p and at r.

How should we calculate magnitude at p and r? Calculate p and r as averaged values of top k=8 closest pixel locations

April 14, 2025



Raymond Yu

In code, you will calculate gradient magnitudes at every q and set it to zero if it is not the local max



$$\mathbf{G}=\sqrt{{\mathbf{G}_x}^2+{\mathbf{G}_y}^2}$$

April 14, 2025

$$G[n_q, m_q] = \begin{cases} G[n_q, m_q] & \text{if } G[n_q, m_q] > G[n_p, m_p] \text{ and } G[n_q, m_q] > G[n_r, m_r] \\ 0 & \text{otherwise} \end{cases}$$

Raymond Yu

What the output looks like after Non-max Suppression



Before After

Raymond Yu

Lecture 5 - 53

Canny edge detector

- 1. Suppress Noise
- 2. Compute gradient magnitude and direction
- 3. Apply Non-Maximum Suppression
 - Assures minimal response
- 4. Use hysteresis and connectivity analysis to detect edges







Problem: if your threshold is too high (left) or too low (right), you have too many or too few edges



Problem: Also, you have too many disconnected edges

Raymond Yu



What the output of hysteresis looks like:



Before

After

April 14, 2025



Hysteresis thresholding connects edges to create long edges

- How does it work?
- Define two thresholds: Low and High

 If less than Low => not an edge
 - If greater than High => strong edge

o If between Low and High => weak edge

Raymond Yu

Lecture 5 - 57

Hysteresis thresholding

If the gradient at a pixel is between Low and High thresholds,

 Consider its neighbors iteratively then declare it an "edge pixel" if it is connected to an 'strong edge pixel' directly or via pixels between Low and High





All the white pixels are not edges (below the low threshold) The black pixels below are strong edges (above the high threshold)



Raymond Yu



Now, let's assume all the red pixels are weak edges (between low and high thresholds)



Raymond Yu

Lecture 5 - 60

Now, let's assume all the red pixels are weak edges (between low and high thresholds)



Raymond Yu

Lecture 5 - 61

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:
 - Reduce multi-pixel wide edges down to single pixel edge
- 4. Thresholding and linking (hysteresis):
 - Define two thresholds: low and high
 - Connect edges together and remove everything else

Raymond Yu

Lecture 5 - 63

Effect of σ (Gaussian kernel spread/size)



The choice of σ depends on desired behavior

- large σ detects large scale edges
- small σ detects fine features

Raymond Yu





Raymond Yu

Lecture 5 - 65

45 years of edge detection



Source: Arbelaez, Maire, Fowlkes, and Malik. TPAMI 2011 (pdf)

April 14, 2025

Lecture 5 - 66

Raymond Yu

What we will learn today

- Edge detector with noisy images
- Sobel Edge detector
- Canny edge detector
- Hough Transform





Hough transform

How Transform edge detections into lines



Raymond Yu

Lecture 5 - 68

Hough transform

- It was introduced in 1962 (Hough 1962) and first used to find lines in images a decade later (Duda 1972).
- Caveat: Hough transform can detect lines, circles and other shapes
 but only for shapes that can be expressed as a math equation.
- It gives us good detections even when the image is noisy and even if the shape is partially hidden.





Input to Hough transform algorithm

- We have performed some edge detection (Sobel filter, Canny Edge detector, etc.), including a thresholding of the edge magnitude image.
- Thus, we have some pixels that may partially describe the boundary of some objects.



Raymond Yu



Detecting lines using Hough transform

- We wish to find sets of pixels that make up straight lines.
- Instead of using [n, m], this might be easier to do with (x, y)

How do we transform [n, m] to (x, y)?

- Simple: We assume
 - n = y,
 - m = x.

Raymond Yu

- So, f[n, m] = f[y, x]



Detecting lines using Hough transform

- Consider a line that passes through two points in the image
 - \circ (x₁, y₁) and (x₂, y₂)
- Straight lines that pass that point have the form:

• How do we calculate the parameters (a, b)?

$$a = (y_2 - y_1) / (x_2 - x_1)$$

b = y_1 - a × x_1




Detecting lines using Hough transform

Lecture 5 - 73

- Consider a line that passes through two points in the image
 - $\circ~(\textbf{x}_1^{},\textbf{y}_1^{})$ and $(\textbf{x}_2^{},\textbf{y}_2^{})$
- Straight lines that pass that point have the form:

y= a*x + b

$$\begin{array}{c} x \\ (x_i, y_i) \\ (x_j, y_j) \end{array}$$

$$y = ax + b_{\bullet}$$

<u>April 14, 2025</u>

• How do we calculate the parameters (a, b)?

$$a = (y_2 - y_1) / (x_2 - x_1)$$

b = y_1 - a × x_1

Detecting lines using Hough transform

- Problem: We don't know which pairs of edge points belong to the same line.
- That's where Hough transform comes in!





- Consider a line that passes through a single point in the image $\,\circ\,\,(x_i^{},\,y_i^{})$
- All straight lines that pass that point have the form:

 $y_i = a^*x_i + b$







 $y_i = a^*x_i + b$

• This equation can be rewritten as follows:

 $\circ \mathbf{b} = -\mathbf{a}^* \mathbf{x}_i + \mathbf{y}_i$





 $y_i = a^*x_i + b$

 $\boldsymbol{\mathcal{A}}$

- This equation can be rewritten as follows:
 - b = -a*x_i + y_i
 We can now consider x and y as parameters
 - a and b as coordinates.



Raymond Yu

Lecture 5 - 77

$$y_i = a^*x_i + b$$

- $b = -a^*x_i + y_i$
- If our coordinates were (a,b) instead of (x, y):
 - We could say the above equation is a line in (a,b)-space
 - $\circ~$ parameterized by x and y.
 - So: one point (x_i,y_i) gives a line in (a,b) space.



April 14, 2025

Raymond Yu

- So: one point (x_i, y_i) gives a line in (a, b) space.
- Another point (x_i, y_i) will give rise to another line in (a, b)-space.



Raymond Yu

Lecture 5 - 79

- Doing this for 6 edge points will result in an graph like the one on the right.
- In (a,b) space these lines will intersect in a point (a', b')
 On the right, a' = 1, b' = 1
- All points on the line defined by (x_i, y_i) and (x_j, y_j) in (x, y)-space will parameterize lines that intersect in (a', b') in (a,b) space.



April 14, 2025

Raymond Yu

The need to quantize and "vote"

Not all intersections will be valid lines.

Consider two edge points that are not part of a real edge:

- They might still intersect in (a, b) space.

Problem: How do we identify intersections that are belong to the same edge versus random points?



Raymond Yu

Lecture 5 - 81

Intuition behind voting

The more lines intersect at the same (a', b') point, the more likely y=a'x + b' is a real edge in the image.

So, we need to count how many lines intersect at a point and keep the ones with high count







Counting in quantized (a, b)-space

- Quantize the parameter space (a b) by dividing it into cells
 - a. [[a_{min}, a_{max}],[b_{min},b_{max}]]
- 2. For each pair of points (x_i, y_i) and (x_j, y_j) , find the intersection (a',b') in (a,b)-space.
- Increase the value of a cell that (a', b') belongs to by 1.
- 4. Cells receiving more than a certain number of counts (also called 'votes') are assumed to correspond to lines in (x,y) space.



April 14, 2025

Raymond Yu

Output of Hough transform

• Here are the top 20 most voted lines in the image:



Raymond Yu



- We can represent lines as polar coordinates instead of y = a*x + b
- Polar coordinate representation:
 x*cosθ + y*sinθ = ρ
- We can transform points in (x, y) space to curves in (ρ θ)-space
 (x y) and (ρ θ)?





 Note that lines in (x, y)-space are not lines in (ρ, θ)-space

 Curves in (ρ, θ)-space intersect similarly like in (a, b)-space.





<u>April 14, 2025</u>

Raymond Yu

- $x^*\cos\theta + y^*\sin\theta = \rho$
- Q. For a vertical line in (x, y)-space, what are the θ and ρ values?





Raymond Yu

Lecture 5 - 87

- $x^*\cos\theta + y^*\sin\theta = \rho$
- Q. For a vertical line in (x, y)-space, what are the θ and ρ values?
 θ=0, ρ=x
- Q. For a horizontal line in (x, y)-space, what are the θ and ρ values?





Raymond Yu

Lecture 5 - 88

Hough transform remarks

• Advantages:

- Conceptually simple.
- Easy implementation
- Handles missing and occluded data very gracefully.
- Can be adapted to many types of forms, not just lines





Hough transform remarks

• Advantages:

- Conceptually simple.
- Easy implementation
- Handles missing and occluded data very gracefully.
- \circ Can be adapted to many types of forms, not just lines

• Disadvantages:

- Computationally complex for shapes with many parameters.
- Looks for only one single shape of object
- Can be "fooled" by "apparent lines".
- The length and the position of a line segment cannot be determined.
- Co-linear line segments cannot be separated.
- \circ Runs in O(N²) since all pairs of points should be considered

Raymond Yu

Lecture 5 - 90

Applications







Summary

- Edge detector with noisy images
- Sobel Edge detector
- Canny edge detector
- Hough Transform

Optional reading: Szeliski, Computer Vision: Algorithms and Applications, 2nd Edition Sections 7.1, 8.1.4

Raymond Yu



Next time

Lines and Corners





