Announcement

Quiz #9 is available in Blackboard.

Due date: 11:59pm EST, Tuesday, April 15th

Open book and open notes

Final Project Grading

- Abstract 20%
- Oral presentation 30%
- Written report 50%
- Extra credits for high quality publication-level project (Research/Hands-on project ONLY) up to 10%

Final Project Presentation

- A team project needs a combined presentation, and both members should participate to the presentation.

- Each individual presentation has 10 + 1 minutes Q&A
- Each team presentation has 13 + 2 minutes Q&A

- Absence of the final project presentation results a Zero for your final presentation grade.

Please submit your final presentation slides into Blackboard before your presentation.

What Should Be Included In the Final-Project Presentation

For a research/hands-on project

- An introduction of the background and challenges
- A brief literature review
- Methodology of your proposed method
- **Experiment**
    - **Experimental datasets**
    - **Experiment setup, evaluation metrics, implementation details, etc.**
    - **Experimental results if any**
- Conclusion and future work
- Reference

What Should Be Included In the Final-Project Presentation

For a survey project

- An introduction of the background and challenges
- **Well organized discussion, critical comments, and comparison of the papers/groups you reviewed**
  - **Organize/connect the papers in groups**
  - **Compare/Highlight the contributions/weaknesses**
  - **Do not enumerate each paper**
  - **A performance comparison on benchmark datasets is preferred**
- Conclusion
- Reference

<u>Note: You don’t need to complete the survey project by presentation date. You can cover most of the papers in your full list.</u>

Final Presentation Criteria – Research/Hands-on Project

Category	Description	Worth
Description of Problem	• Did you make clear the nature of the problem you were trying to solve? • Did you understand the challenges of the problem? • Did you show your understanding of the state-of-the-art?	15
Description of Methodology	• Did you present the methodology clearly? • Did you provide sufficient (key) information to the audience?	25
Description of Experiment	Can your experiment design – dataset, evaluation metrics, baseline approaches effectively demonstrate your proposed method?	25
Visual Aids	• Were visual aids used effectively? • Were slides clear and easy to read by the audience? (Large fonts, high contrast, consistent format, etc.)	10
Clarity & Organization	• Was the presentation easy to understand? • Did it have a logical flow and organization? • Did you include a reference list?	15
Timing	• Was the presentation well-paced? • Did it fit within the time allotted?	10

Final Presentation Criteria – Survey Project

Category	Description	Worth
Description of Problem	• Did you make clear the nature of the problem you were trying to solve? • Did you understand the challenges of the problem?	15
Organization of the surveyed papers	• Did you show a good overview of the field? • Were the papers surveyed organized in a reasonable way? • Did you provide sufficient (key) information to the audience?	25
Discussion and Comparison	• Did your discussion clearly show your understanding of the methods? • Did you make critical comments on the methods?	25
Visual Aids	• Were visual aids used effectively? • Were slides clear and easy to read by the audience? (Large fonts, high contrast, consistent format, etc.)	10
Clarity & Organization	• Was the presentation easy to understand? • Did it have a logical flow and organization? • Did you include a reference list?	15
Timing	• Was the presentation well-paced? • Did it fit within the time allotted?	10

On the Final Project Report

Written report due time: 11:59pm, Tuesday, May 6

- Report format: the same as a conference paper
  - For example, you can use a template for CVPR 2025

    Releases · cvpr-org/author-kit

    Length: around 6 pages (+ 1 page reference) double-column
- Code must be submitted with clear comments (Research project only)

Academic integrity (avoiding plagiarism)

- don’t copy other person’s work
- describe using your own words
- complete citation and acknowledgement whenever you use any other work (either published or online)

No late submission is allowed

What Should Be Included In the Final-Project Report

For a research/hands-on project

- An introduction of the background and challenges
- A brief literature review
- Methodology of your proposed method
- **Experiment**
  - **Experimental datasets**
  - **Experiment setup, evaluation metrics, implementation details, etc.**
  - **Experimental results**
- Conclusion and future work
- Reference

What Should Be Included In the Final-Project Report

For a survey project (>= 15 references)

- An introduction of the background and challenges
- **Well organized discussion, critical comments, and comparison of the papers/groups you reviewed**
  - **Organize/connect the papers in groups**
  - **Compare/Highlight the contributions/weaknesses**
  - **Do not enumerate each paper**
  - **A performance comparison on benchmark datasets is preferred**
- Conclusion
- Reference

Final Project Report Criteria – Research/Hands-on Project

Category	Description	Worth
Description of Problem	• Did you make clear the nature of the problem you were trying to solve? • Did you understand the challenges of the problem?	10
Literature Review	Did you show your understanding of the state-of-the-art?	5
Description of Methodology	• Did you present the methodology clearly? • Did you provide sufficient (key) information to the audience?	30
Description of Experiment	Can your experiment design – dataset, evaluation metrics, baseline approaches effectively demonstrate your proposed method?	30
Writing Clarity	▪ Does the report read well? ▪ Is it easy to understand?	10
Organization & Length	• Is the report well-organized? • Does it have a logical flow?	10
Reference	• Is the reference section complete and in consistent format? • Are the citations in the text in consistent format?	5

Final Project Report Criteria – Survey Project

Category	Description	Worth
Description of Problem	• Did you make clear the nature of the problem you were trying to solve? • Did you understand the challenges of the problem?	15
Organization of the surveyed papers	• Did you show a good overview of the field? • Were the papers surveyed organized in a reasonable way? • Did you provide sufficient (key) information to the audience?	25
Discussion and Comparison	• Was the problem comprehensively reviewed? Did your discussion clearly show your understanding of the methods? • Did you make critical comments on the methods you reviewed?	35
Writing Clarity	▪ Does the report read well? ▪ Is it easy to understand?	10
Organization & Length	• Is the report well-organized? • Does it have a logical flow?	10
Reference	• Is the reference section complete and in consistent format? • Are the citations in the text in consistent format?	5

Today

Image segmentation

Image Segmentation

A process that partitions R into subregions $R_1, R_2, \ldots, R_n$

Microsoft multiclass segmentation data set

Image Segmentation

(a) $\bigcup_{i=1}^{n} R_i = R$

(b) R_i is a connected set, $i = 1, \ldots, n$

(c) $R_i \cap R_j = \phi, \forall i \neq j$

(d) $Q(R_i) = TRUE$

(e) $Q(R_i \cup R_j) = FALSE$ for adjacent regions R_i and R_j

Two categories based on intensity properties:

- **Discontinuity** – edge-based algorithms
- **Similarity** – region-based algorithms

Edge-based and Region-based Segmentation

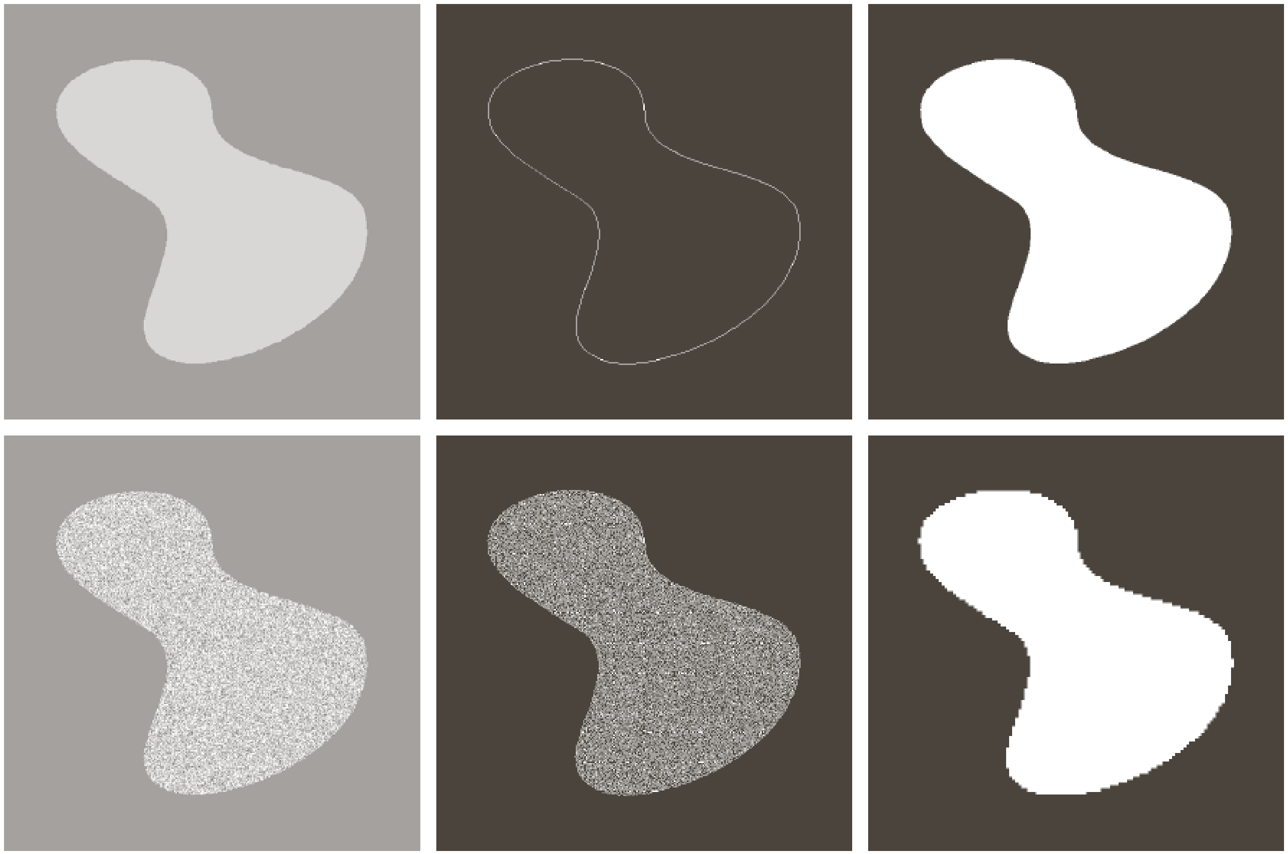

FIGURE 10.1 (a) Image containing a region of constant intensity. (b) Image showing the boundary of the inner region, obtained from intensity discontinuities. (c) Result of segmenting the image into two regions. (d) Image containing a textured region. (e) Result of edge computations. Note the large number of small edges that are connected to the original boundary, making it difficult to find a unique boundary using only edge information. (f) Result of segmentation based on region properties.

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Brief Review on Simple Edge Detectors

- **First-order derivative**
  - E.g., Roberts (2x2), Prewitt (3x3), Sobel (3x3, smooth + difference)
  - Thicker edge
  - One operator for one edge direction

- **Second-order derivative**
  - Laplacian (3x3)
  - Double edge
  - Zero-crossing

- **Common issues:**
  - Sensitive to image noise
  - Cannot deal with the scale change of the image

Advanced Edge Detection Techniques

- Deal with image noise
- Exploit the properties of image

→ Work much better for real images

Advanced edge detectors:

- Laplacian of Gaussian (LoG) $\nabla^2 G(x,y) = \left[\frac{x^2 + y^2 - 2\sigma^2}{\sigma^4} \right] e^{-\frac{x^2+y^2}{2\sigma^2}}$
- Difference of Gaussian (DoG)

$$DOG(x,y) = \frac{1}{2\pi\sigma_1^2} e^{-\frac{x^2+y^2}{2\sigma_1^2}} - \frac{1}{2\pi\sigma_2^2} e^{-\frac{x^2+y^2}{2\sigma_2^2}}, \quad \sigma_1 > \sigma_2$$

- Canny

Marr-Hildreth Detector (LoG)

Observations:

- Intensity changes are dependent on the image scale
- A sudden intensity change (step) causes a peak/trough in the1st order derivative and a zero-crossing in the 2nd order derivative
- The 2nd order derivative is especially sensitive to noise

⇒ Smooth the image using a Gaussian filter first before applying the Laplacian

Gaussian → $G(x,y) = e^{-\frac{x^2+y^2}{2\sigma^2}}$

Laplacian of a Gaussian (LoG)

$$g(x,y) = \nabla^2[G(x,y) \otimes f(x,y)] = \boxed{\nabla^2 G(x,y)} \otimes f(x,y)$$

- **Varying σ values for scale changes**
- **Rotation invariant in edge detection**

LoG Filtering

$$g(x,y) = \nabla^2 G(x,y) \otimes f(x,y)$$
$$= \nabla^2 [G(x,y) \otimes f(x,y)]$$

1. Filter the input image with an *nxn* Gaussian filter.
2. Compute the Laplacian of the intermediate image resulting from Step 1.
3. Find the zero-crossings of the image from Step 2.
   - opposite signs of the neighbors
   - the difference should be significant

Note:

- Window size $n \geq 6\sigma$ and n is an odd number

An Example – Edges are 1 Pixel Thick

Original

LoG filtering

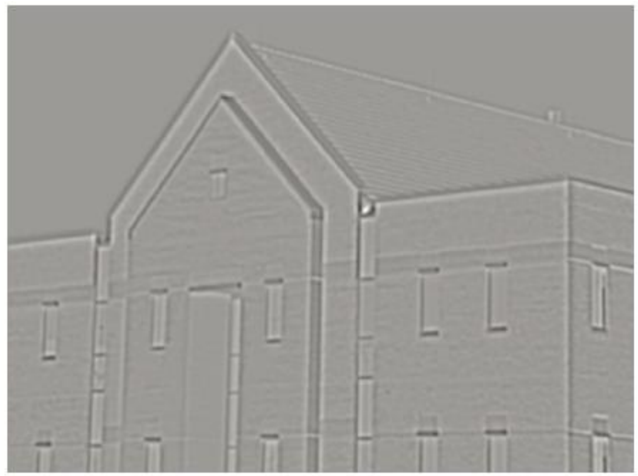

LoG filtering with T=0

Zero-crossing with T=0

Zero-crossing with T=4%max

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Approximate LoG by DoG

LoG needs multiple filters for scale variations

Difference of Gaussian (DoG)

$$DOG(x,y) = \frac{1}{2\pi\sigma_1^2} e^{-\frac{x^2+y^2}{2\sigma_1^2}} - \frac{1}{2\pi\sigma_2^2} e^{-\frac{x^2+y^2}{2\sigma_2^2}}, \quad \sigma_1 > \sigma_2$$

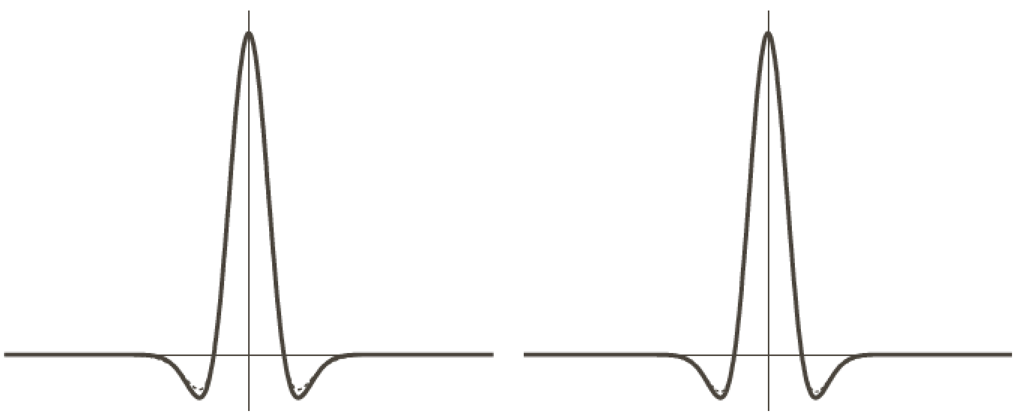

a b

FIGURE 10.23 (a) Negatives of the LoG (solid) and DoG (dotted) profiles using a standard deviation ratio of 1.75:1. (b) Profiles obtained using a ratio of 1.6:1.

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Canny Edge Detector

The general goal of edge detection:

1. Low error rate: all edges should be found and there should be no false response

2. Edge points should be well localized: the edges located must be as close as possible to the true edges

3. Single edge point response: the detector should return only one point for each true edge point

Canny edge detector: expressing these three criteria mathematically and then find optimal solutions

Canny Edge Detector

Gaussian smoothing + 1st order derivative

➡ **optimal step edge detector**

$$f_s(x,y) = G(x,y) \otimes f(x,y)$$

$$g_x = \partial f_s / \partial x \qquad g_y = \partial f_s / \partial y$$

$$M(x,y) = \sqrt{g_x^2 + g_y^2} \qquad \alpha(x,y) = \tan^{-1}\left(g_y / g_x\right)$$

Thick edge $\xrightarrow{\text{Nonmaxima suppression}}$ **Single edge**

Quantize the Edge Direction

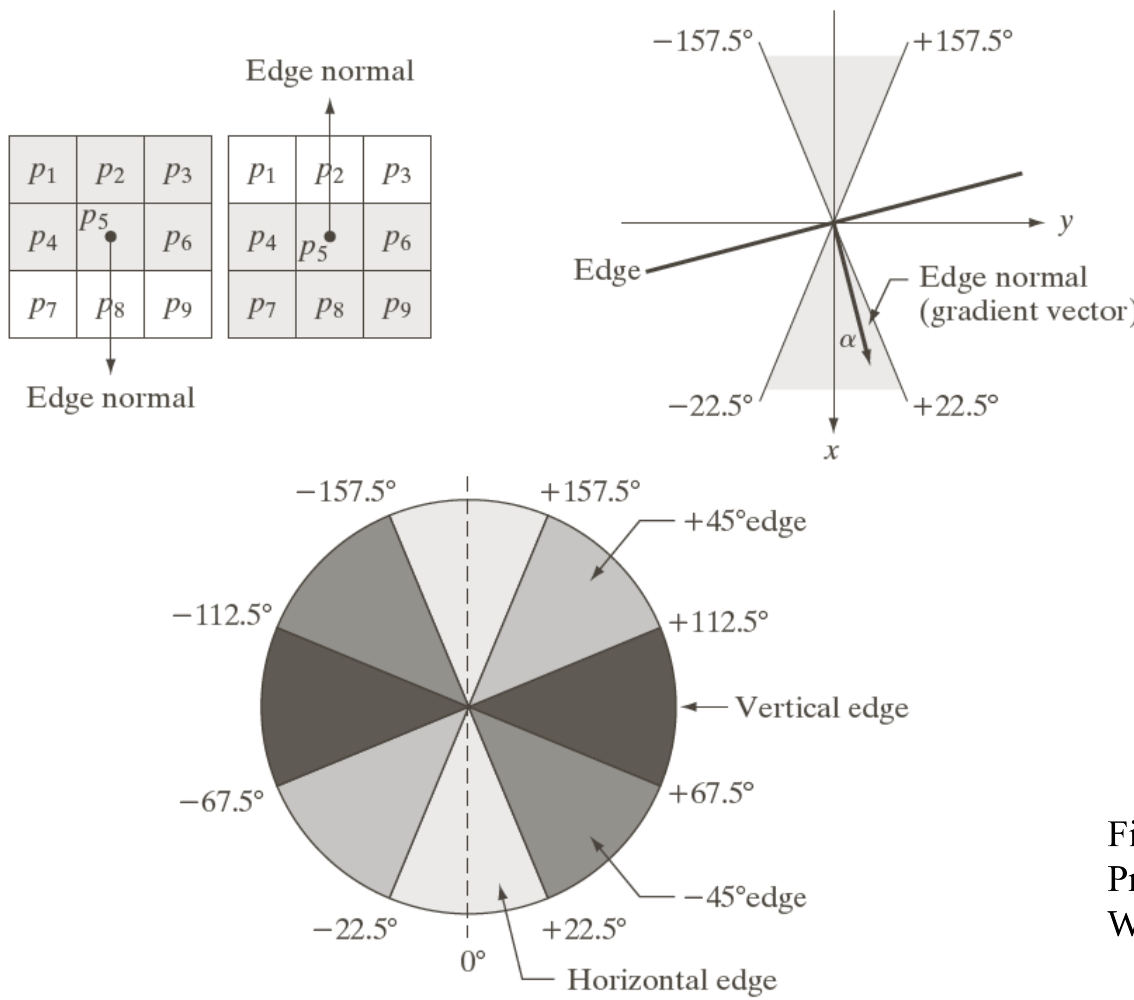

a b
c

FIGURE 10.24 (a) Two possible orientations of a horizontal edge (in gray) in a 3 × 3 neighborhood. (b) Range of values (in gray) of α, the direction angle of the *edge normal*, for a horizontal edge. (c) The angle ranges of the edge normals for the four types of edge directions in a 3 × 3 neighborhood. Each edge direction has two ranges, shown in corresponding shades of gray.

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Canny Detector -- Algorithm

1. **Smooth the input image with a Gaussian filter**
2. **Compute the gradient magnitude and angle images**
3. **Apply nonmaximum suppression on the gradient magnitude image**
   1. At (x, y), find the quantized edge normal d_k
   2. If the value $M(x, y)$ is less than at least one of its two neighbors along d_k, let $g_N(x, y) = 0$; otherwise $g_N(x, y) = M(x, y)$
4. **Reduce false edge: double thresholding and connectivity analysis to detect and link edges**
   1. High-threshold → strong edge pixels →valid edge pixels
   2. Low-threshold → weak edge pixels → valid only when connected to strong edge pixels

An Example

a b
c d

FIGURE 10.25
(a) Original image of size 834 × 1114 pixels, with intensity values scaled to the range [0, 1].
(b) Thresholded gradient of smoothed image.
(c) Image obtained using the Marr-Hildreth algorithm.
(d) Image obtained using the Canny algorithm. Note the significant improvement of the Canny image compared to the other two.

Figure from “Digital Image Processing”, 3rd edition, Gonzalez and Woods

An Example

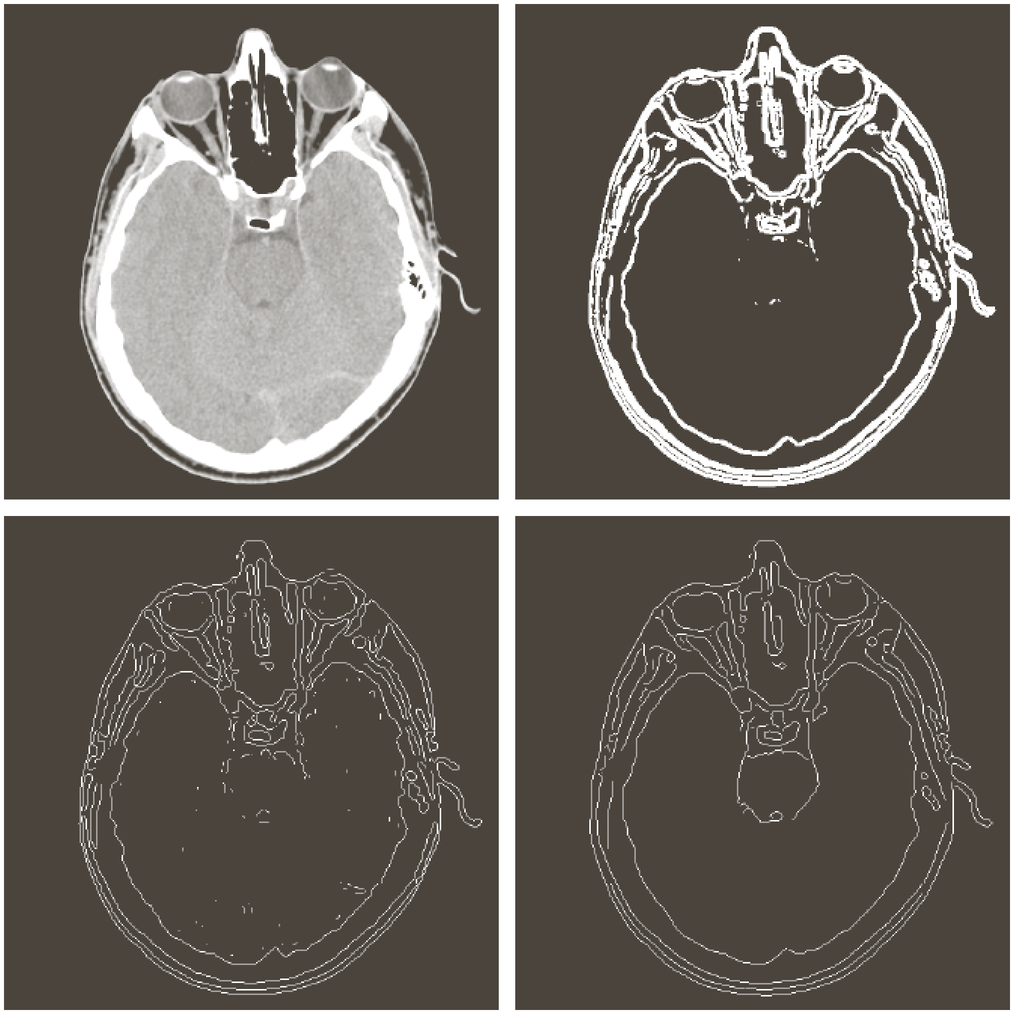

FIGURE 10.26 (a) Original head CT image of size 512 × 512 pixels, with intensity values scaled to the range [0, 1]. (b) Thresholded gradient of smoothed image. (c) Image obtained using the Marr-Hildreth algorithm. (d) Image obtained using the Canny algorithm. (Original image courtesy of Dr. David R. Pickens, Vanderbilt University.)

Figure from “Digital Image Processing”, 3rd edition, Gonzalez and Woods

Edge Linking and Boundary Detection

All edge detection algorithms can only detect fragments of boundaries, due to image noise, non-uniform illuminations, or other effects

Edge linking: link edges into longer meaningful edges or a full region boundaries

Three classic methods:

- Local processing
- Regional Processing
- Hough transform

Edge Linking – Local Processing

Link the edge points with similar properties:

- Strength
- Direction

Two edge pixels are linked if

$$|M(s,t) - M(x,y)| \le E$$

$$|\alpha(s,t) - \alpha(x,y)| \le A$$

Global Processing

A basic idea: Given n points

1. Find $n(n-1)/2$ lines between each pair of points
2. Find all subset of points that are close to particular lines. This needs $n(n-1)n/2$ comparisons

This is computationally expensive!

→ **Hough transform**

Hough Transform

A line in x-y plane is a point in the parameter plane.
A point in x-y plane is a line in the parameter plane.

X-Y plane ⟷ **Parameter plane**

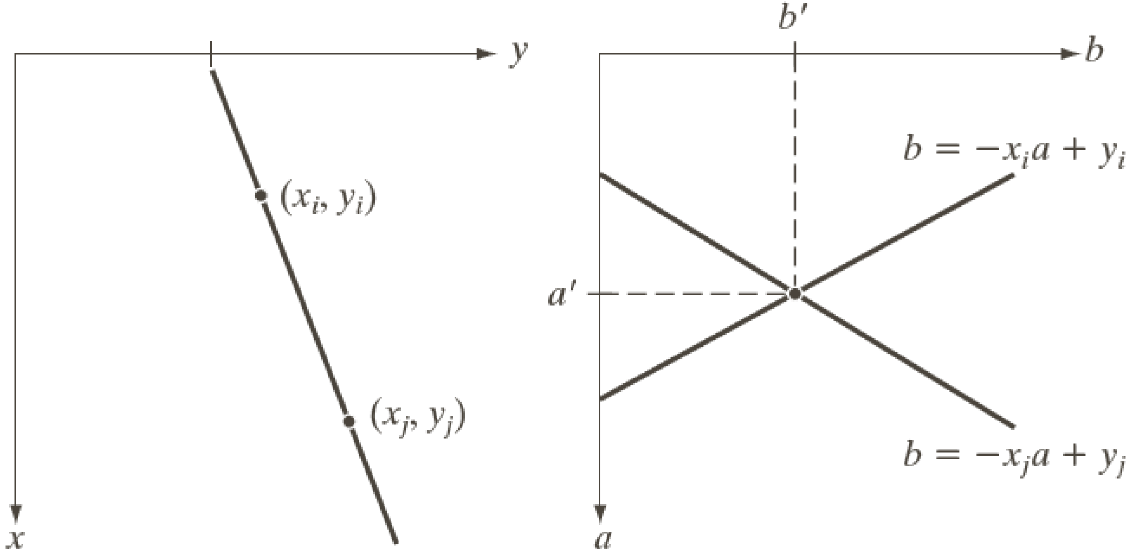

a b
FIGURE 10.31
(a) *xy*-plane.
(b) Parameter space.

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Problem of slope-intercept form: the slope *a* approaches infinity for vertical lines

Hough Transform

A line in x-y plane is a point in the parameter plane.
A point in x-y plane is a sinusoidal in the parameter plane (polar space).

$$x\cos\theta + y\sin\theta = \rho$$

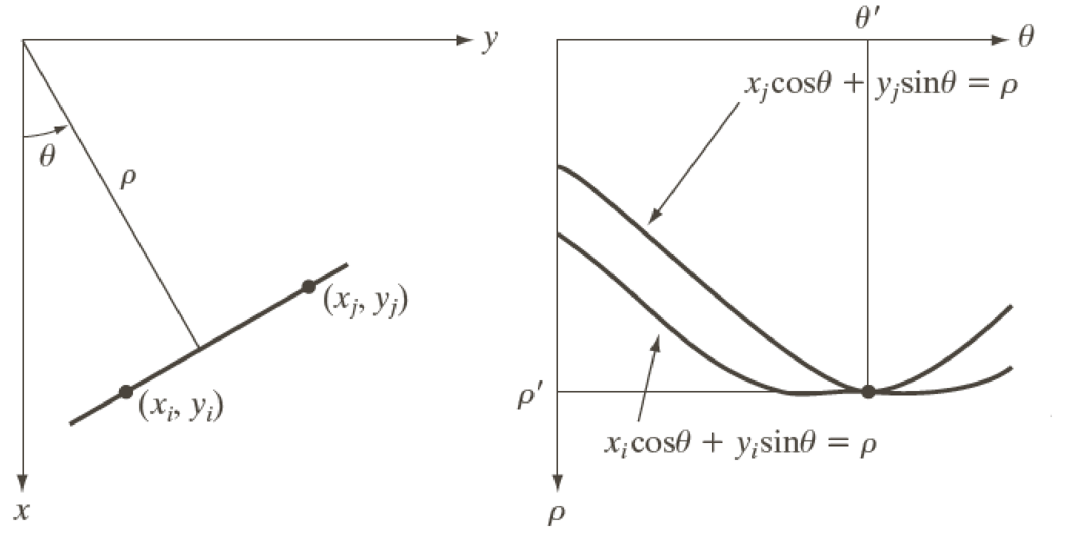

a b c

FIGURE 10.32 (a) (ρ, θ) parameterization of line in the xy-plane. (b) Sinusoidal curves in the $\rho\theta$-plane; the point of intersection (ρ', θ') corresponds to the line passing through points (x_i, y_i) and (x_j, y_j) in the xy-plane. (c) Division of the $\rho\theta$-plane into accumulator cells.

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

A Toy Example

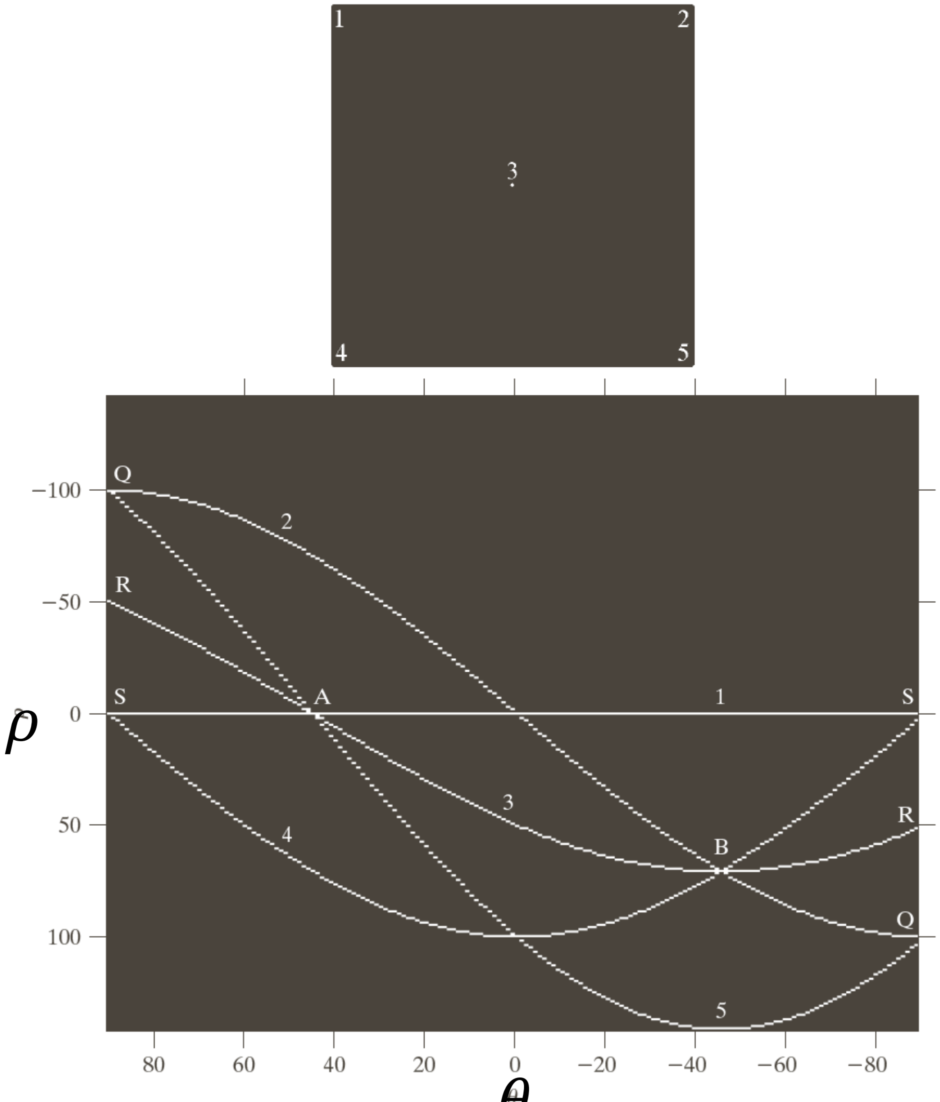

a
b

FIGURE 10.33
(a) Image of size 101 × 101 pixels, containing five points.
(b) Corresponding parameter space. (The points in (a) were enlarged to make them easier to see.)

$$x\cos\theta + y\sin\theta = \rho$$

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Hough Transform – Real Example

Find the runway

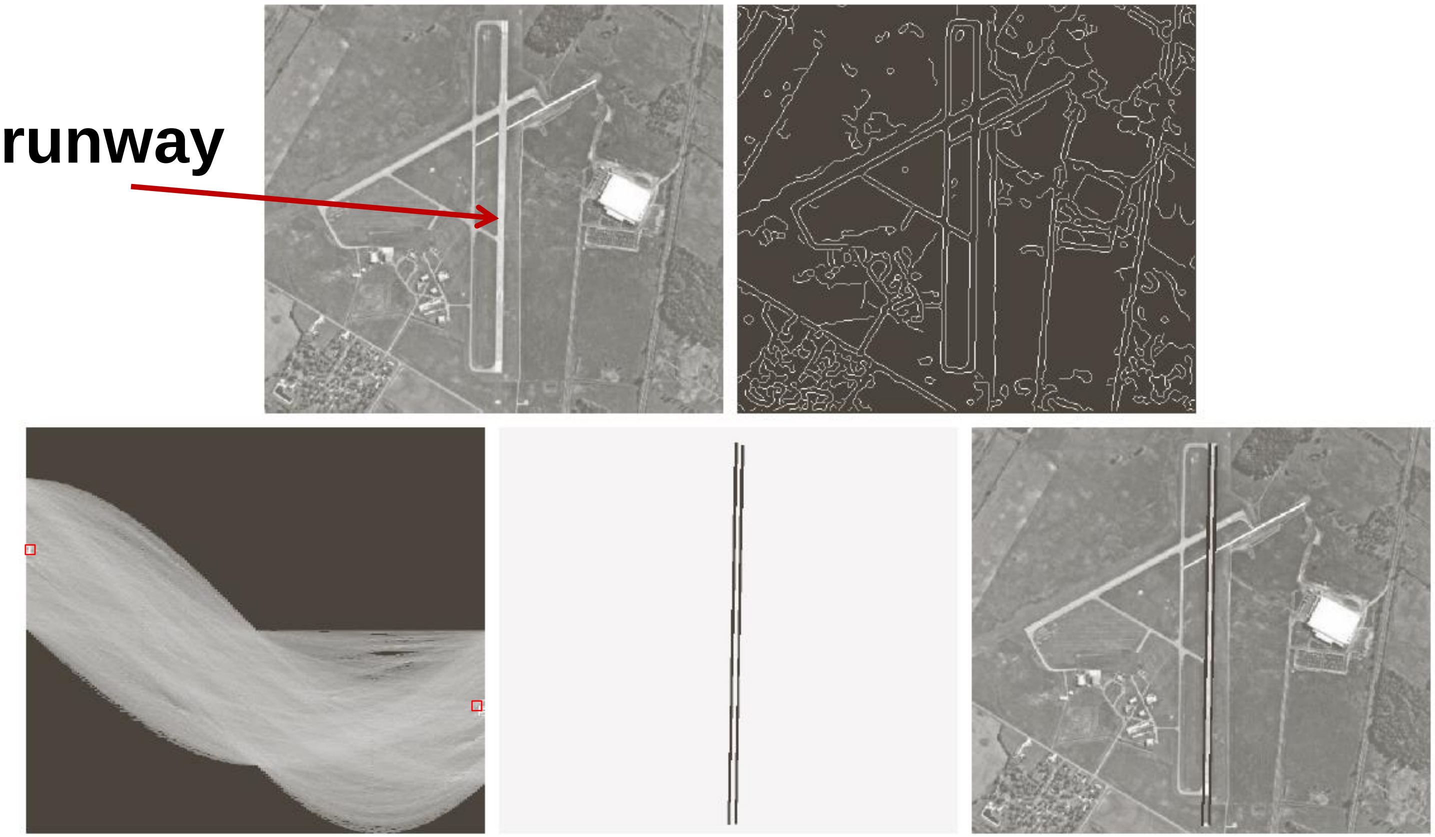

a b
c d e

FIGURE 10.34 (a) A 502 × 564 aerial image of an airport. (b) Edge image obtained using Canny's algorithm. (c) Hough parameter space (the boxes highlight the points associated with long vertical lines). (d) Lines in the image plane corresponding to the points highlighted by the boxes). (e) Lines superimposed on the original image.

Figure from "Digital Image Processing", 3rd edition, Gonzalez and Woods

Hough Transform Algorithm

- Obtain a binary edge image using any edge detector
- Specify subdivisions in the ρ-θ plane
- For each edge point

    For each θ value in the accumulator cell

        Update the accumulator cell with the corresponding ρ
- Examine the counts of accumulator cell for high pixel concentrations
- For each chosen cell, link the pixels based on the continuity

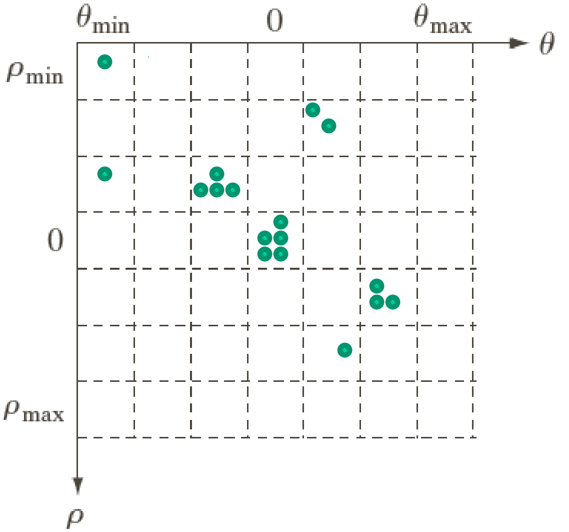

$$-90^0 \leq \theta \leq 90^0, -D \leq \rho \leq D$$

Challenges with Hough Transform

- The resolution of the cells
    - Low resolution – cannot distinguish between different lines
    - High resolution – missing lines because of noise and computation

General Ideas

Tokens

- whatever we need to group (pixels, points, surface elements, etc.)

Top-down segmentation

- tokens belong together because they lie on the same object

Bottom-up segmentation

- tokens belong together because they are locally coherent

These two are not mutually exclusive