Advanced Features

Jana Kosecka

Slides from: S. Thurn, D. Lowe, Forsyth and Ponce

Fitting

- We've learned how to detect edges, corners, blobs. Now what?
- We would like to form a higher-level, more compact representation of the features in the image by grouping multiple features according to a simple model

Slides S. Lazebnik

Fitting

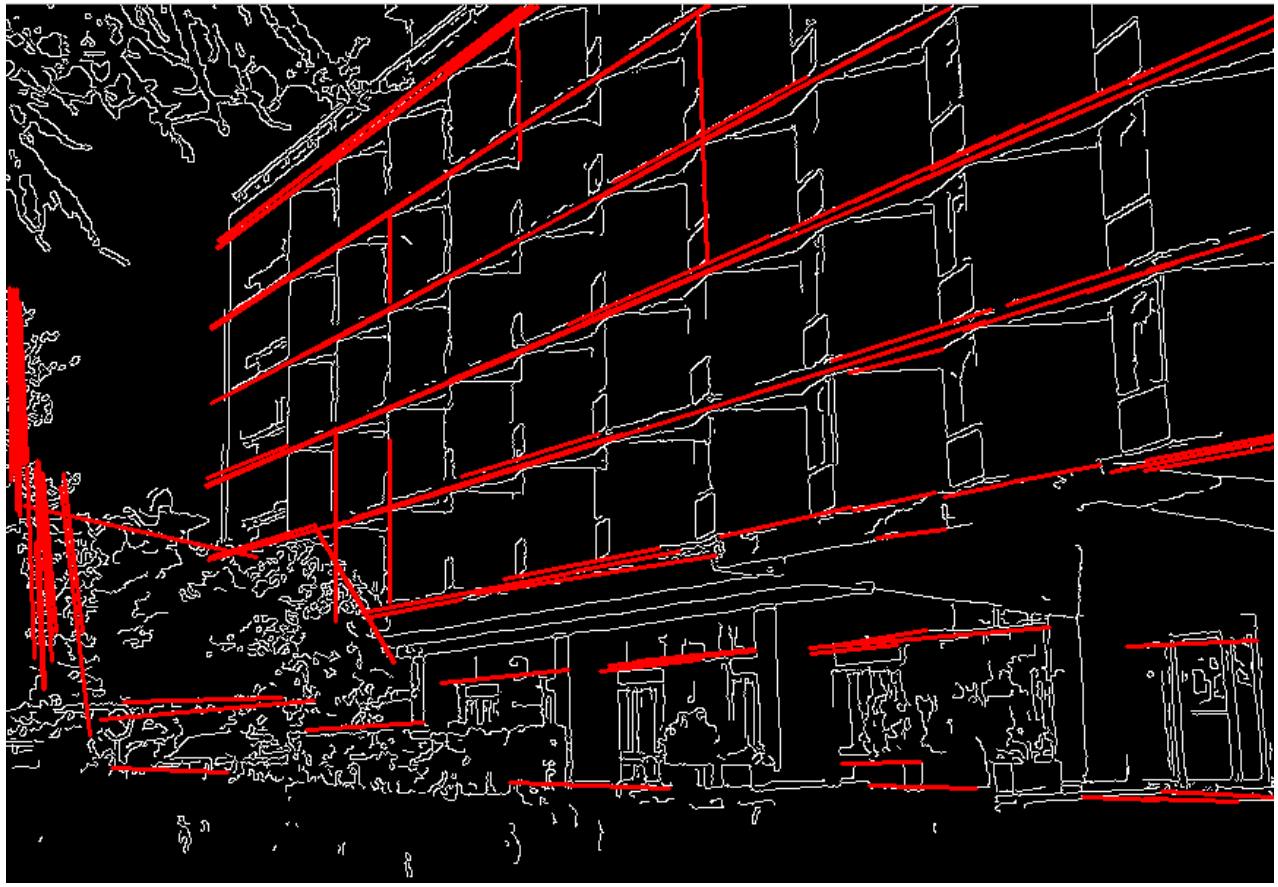

Line fitting

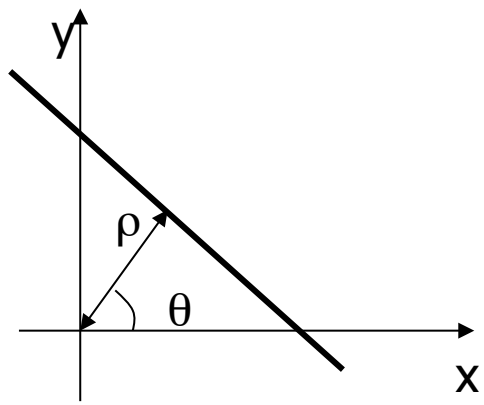

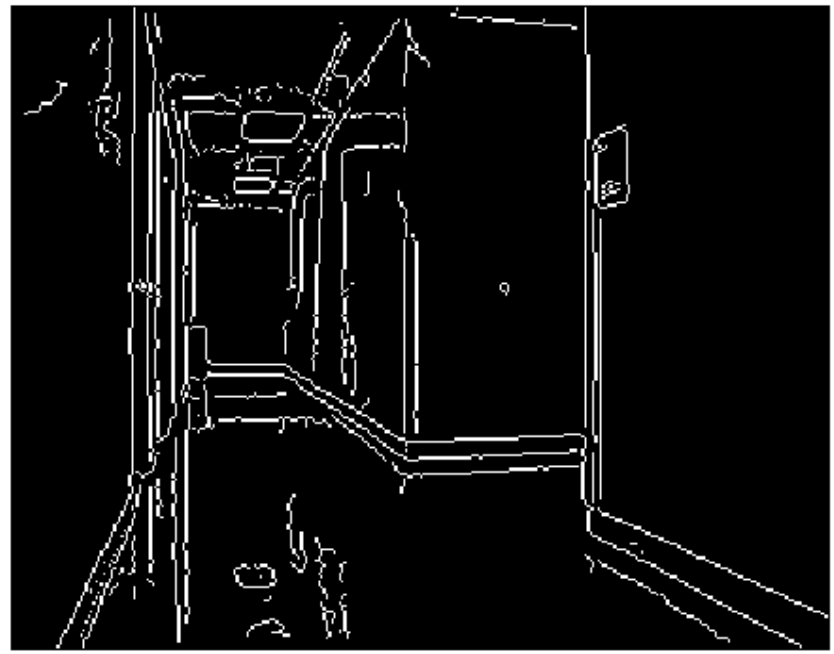

Non-max suppressed gradient magnitude

- Edge detection, non-maximum suppression
(traditionally Hough Transform – issues of resolution, threshold selection and search for peaks in Hough space)
- Connected components on edge pixels with similar orientation
- group pixels with common orientation

Line Fitting

$$A = \begin{bmatrix} \sum x_i^2 & \sum x_i y_i \\ \sum x_i y_i & \sum y_i^2 \end{bmatrix}$$

second moment matrix associated with each connected component

v_1 - eigenvector of A

$$v_1 = [\cos(\theta), \sin(\theta)]^T$$

$$\theta = \arctan(v_1(2)/v_1(1))$$

$$\rho = \bar{x}\sin(\theta) - \bar{y}\cos(\theta)$$

- Line fitting lines determined from eigenvalues and eigenvectors of A
  - Candidate line segments - associated line quality

Fitting

- Choose a parametric model to represent a set of features

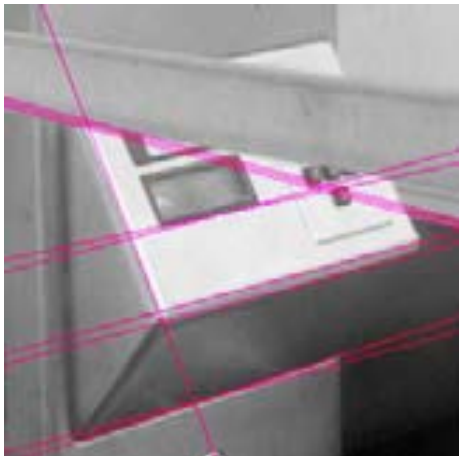

simple model: lines

simple model: circles

complicated model: car

Source: K. Grauman

Finding lines in an image

- Option 1:
  - Search for the line at every possible position/orientation
  - What is the cost of this operation?
- Option 2:
  - Use a voting scheme: Hough transform

Finding lines in an image

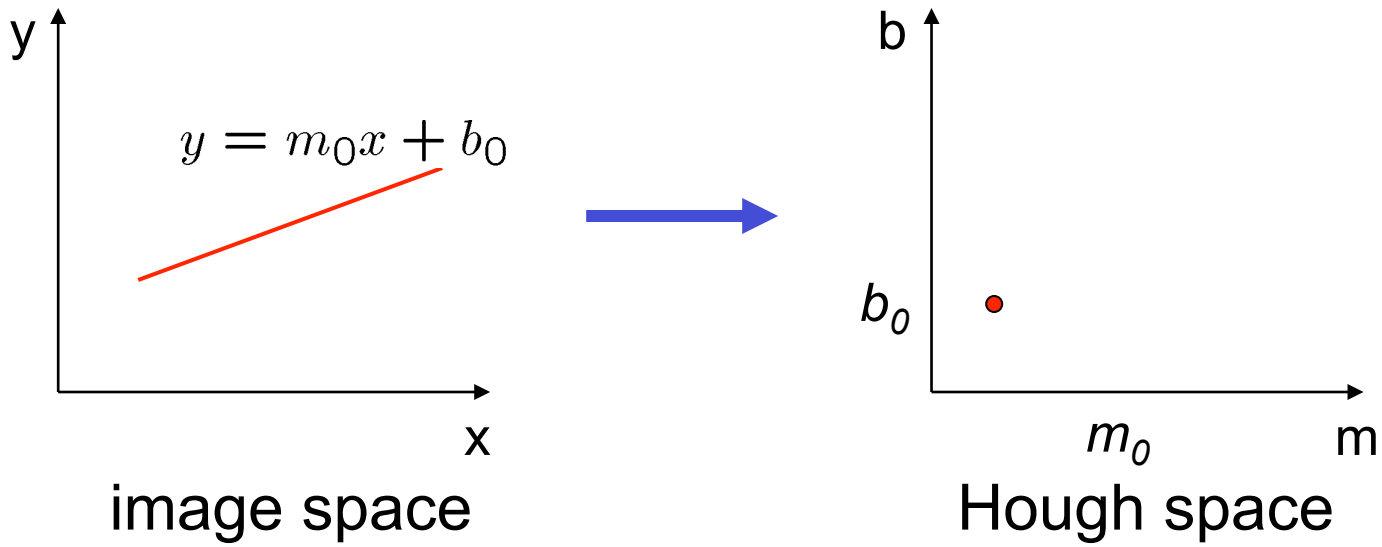

- Connection between image (x,y) and Hough (m,b) spaces
  - A line in the image corresponds to a point in Hough space
  - To go from image space to Hough space:
    - given a set of points (x,y), find all (m,b) such that y = mx + b

Finding lines in an image

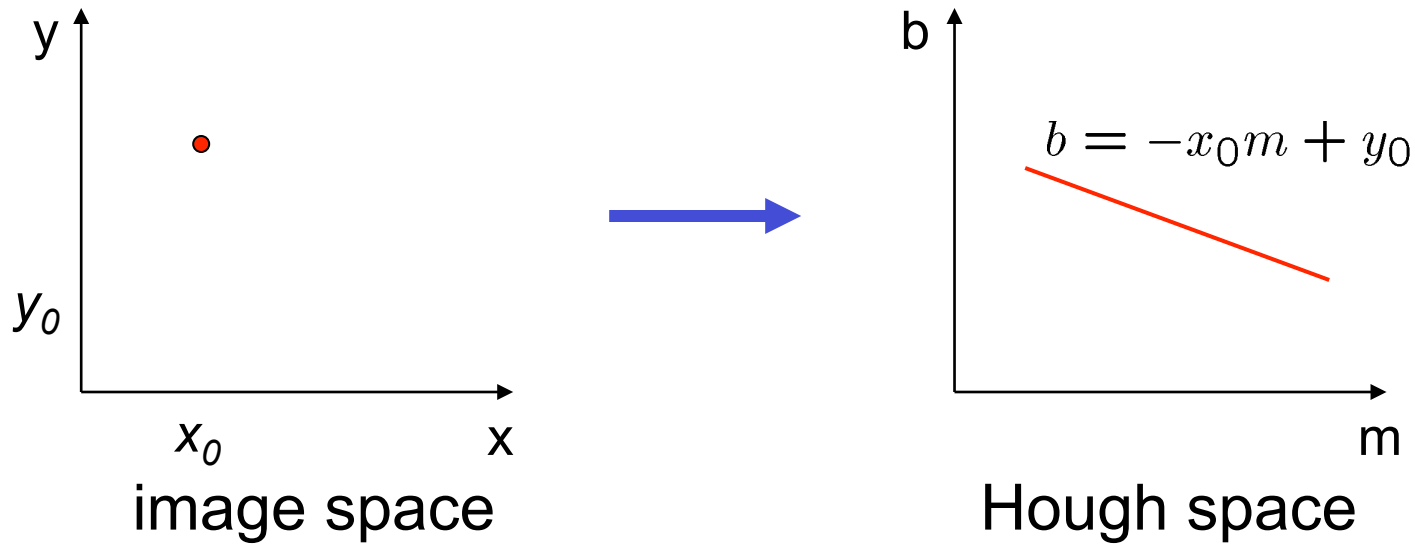

- Connection between image (x,y) and Hough (m,b) spaces
  - A line in the image corresponds to a point in Hough space
  - To go from image space to Hough space:
    - given a set of points (x,y), find all (m,b) such that y = mx + b
  - What does a point (x_0, y_0) in the image space map to?

- A: the solutions of $b = -x_0 m + y_0$
  - this is a line in Hough space

Hough transform algorithm

- Typically use a different parameterization

$$d = xcos\theta + ysin\theta$$

  - d is the perpendicular distance from the line to the origin
  - θ is the angle this perpendicular makes with the x axis
  - Why?

Idea – keep an accumulator array (Hough space) and let each edge pixel contribute to it

Line candidates are the maxima in the accumulator array

Typical Hough Transform

Basic Hough transform algorithm

1. Initialize H[d, theta]=0
2. For each edge point I[x,y] in the image
3. For theta = 0 to 180

H[$d = x cos\theta + y sin\theta$

point is now a sinusoid in Hough space

Find the value(s) of (d, theta) where H[d, theta] is maximum

The detected line in the image is given by maxima in the Hough space

What's the running time (measured in # votes)?

Radon Transform

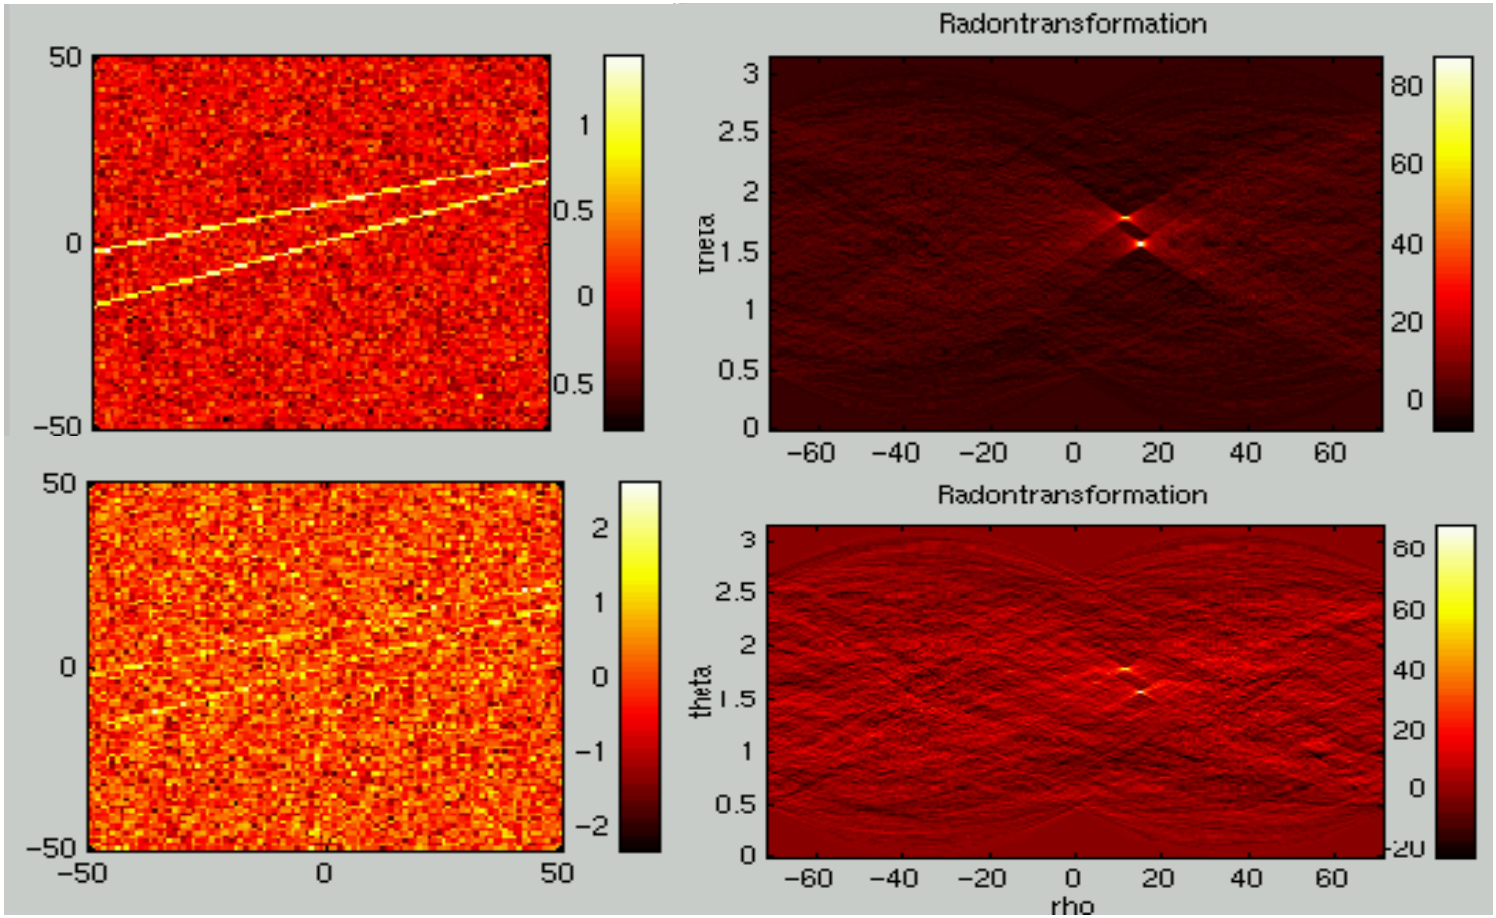

- Projection of an image along radial line

Hough Transform for Curves

- The H.T. can be generalized to detect any curve that can be expressed in parametric form:
  - $Y = f(x, a1, a2, \ldots ap)$
  - a1, a2, … ap are the parameters
  - The parameter space is p-dimensional
  - The accumulating array is LARGE!