CS 369: Introduction to Robotics

Prof. Thao Nguyen

Spring 2026

HAVERFORD
COLLEGE

Admin

- Lab 2 grades posted on Moodle

Outline for today

- Robot perception
- Computer vision

Outline for today

- Robot perception
- Computer vision

Robot sensors

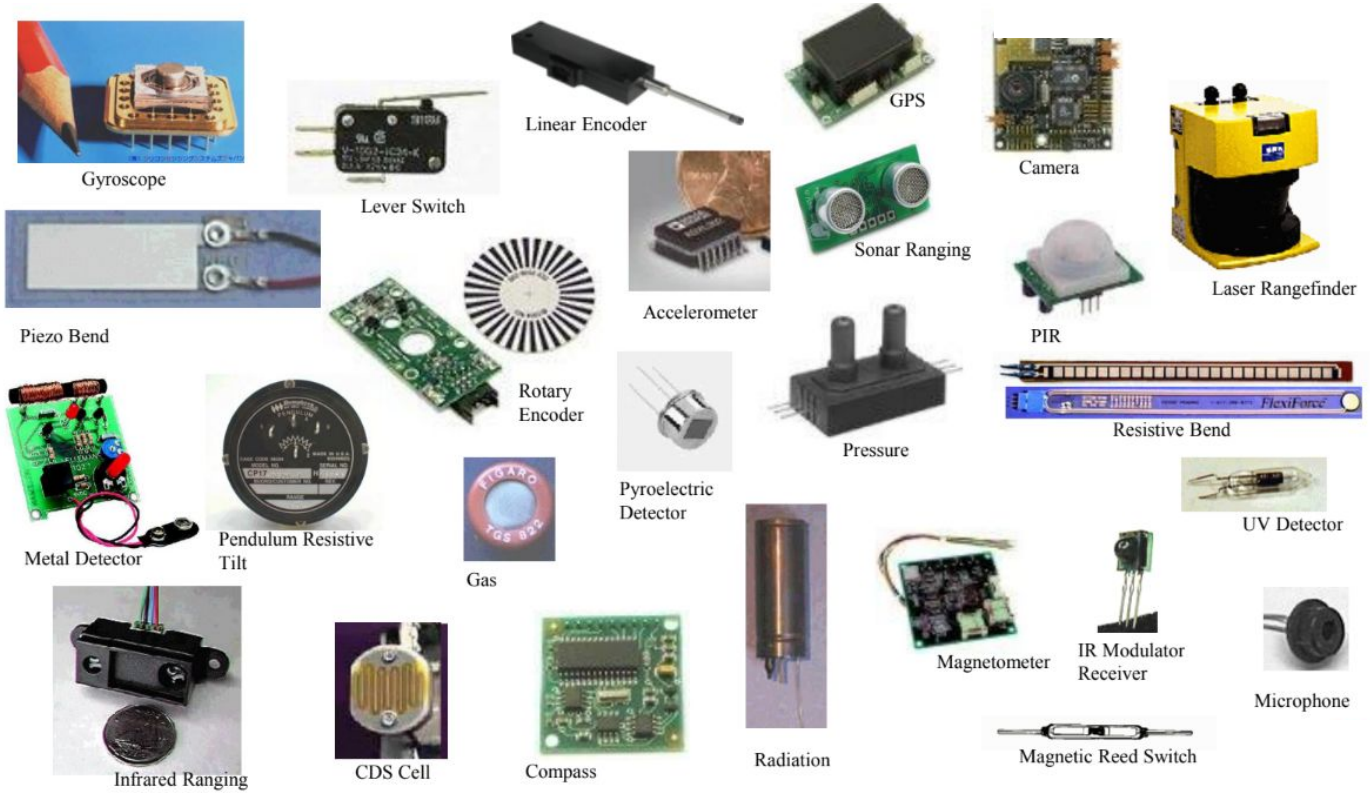

Sensor data

- Audio
- Distance
- Tactile
- Visual
- etc.

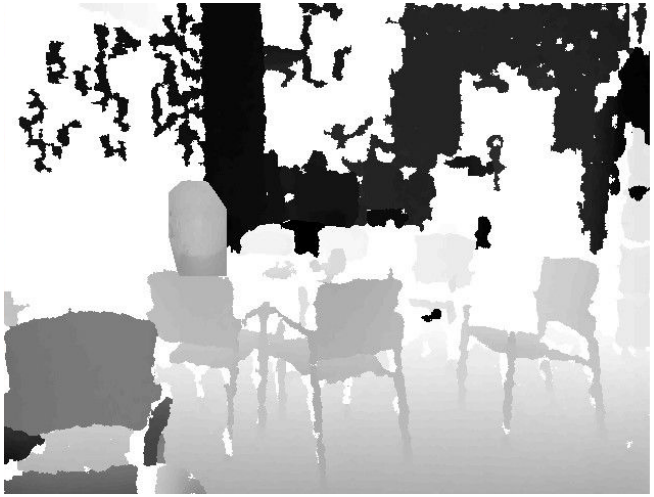

Images

- Discrete representation of a continuous function
  - An image is a two dimensional array of pixels
  - Pixel: cell with numeric value representing a uniform portion of an image

- Resolution
  - Number of pixels across the horizontal dimension
  - Number of pixels in the vertical dimension
  - Number of layers used for color

Images

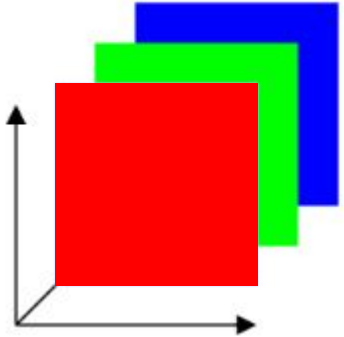

- Binary: two color image
  - Pixel is only one bit of information
  - Indicates if the intensity of color is above or below some nominal value
- Grayscale
  - All pixels represent the intensity of light in an image
  - Intensity of light in a pixel is stored as a number, generally 0 to 255 inclusive
- Color
  - Three grayscale images layered on top of eachother with each layer indicating the intensity of a specific color light — generally red, green, and blue (RGB)
  - Third dimension in a digital image

Pinhole camera model

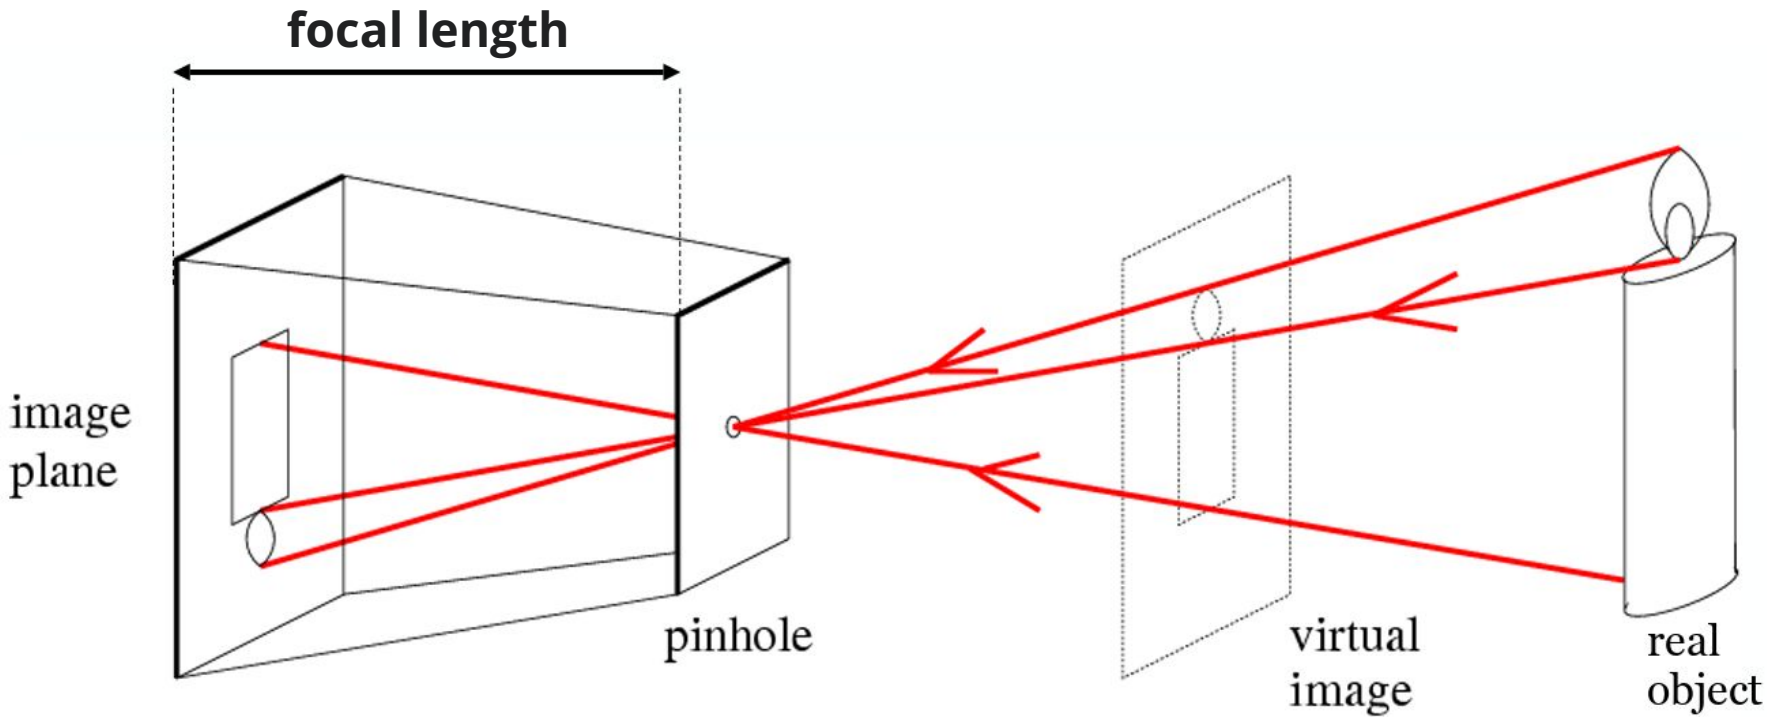

Projection on the image plane

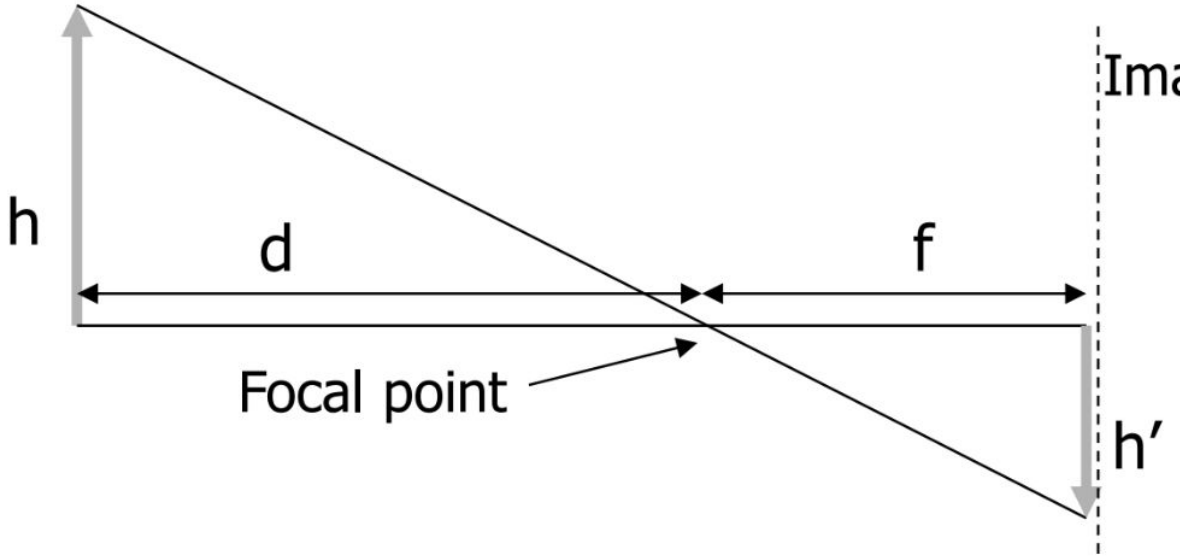

$$\frac{h}{d} = \frac{h'}{f}$$

Outline for today

- Robot perception

- Computer vision

Computer vision

- Making sense of visual data
- Typical pipeline:

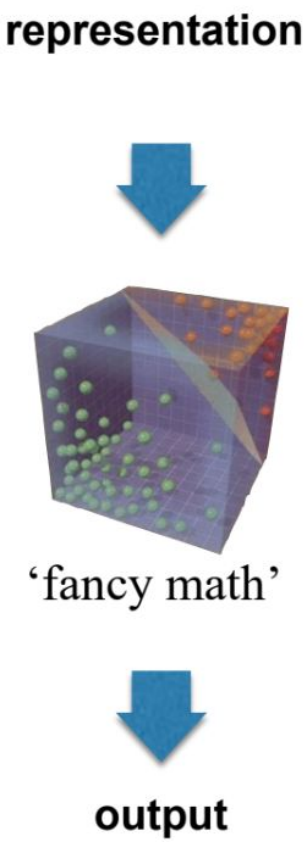

What should we look at?
(e.g. image features)

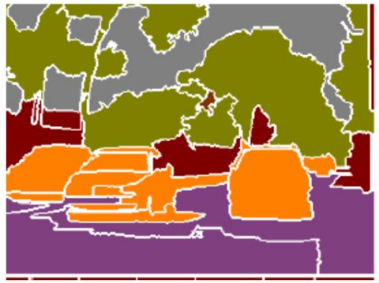

What can we understand?
(e.g. semantic segmentation)

Thresholding

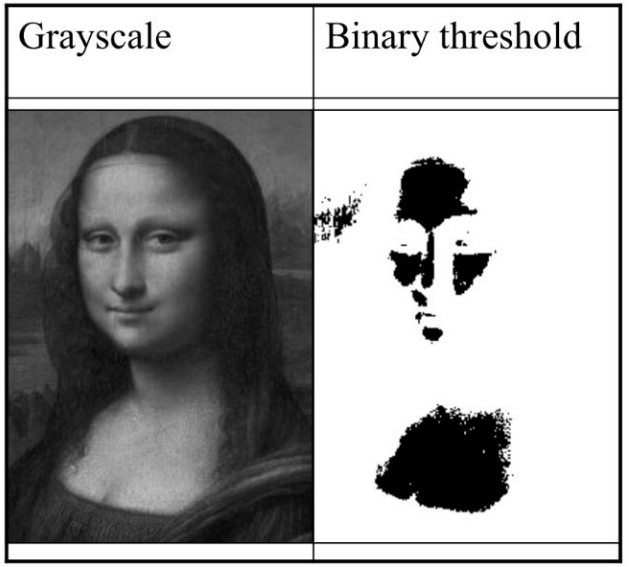

- Purpose
  - Find areas of high color intensity
  - Highlight locations of different features of the image
  - Image compression, use fewer bits to encode a pixel
- Process
  - Decide on a threshold value μ
  - Scan every pixel in the image
    - If it is greater than or equal to μ, make it 1
    - If it is less than μ, make it 0
  - Picking a good μ
    - Often 128 is a good value to start with
    - Use a histogram to determine values based on color frequency features

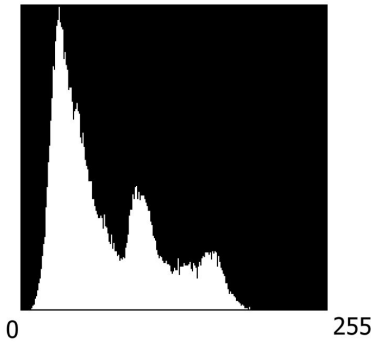

Segmentation: flood fill

Algorithm that determines and alters the area connected to a given node with some matching attribute.

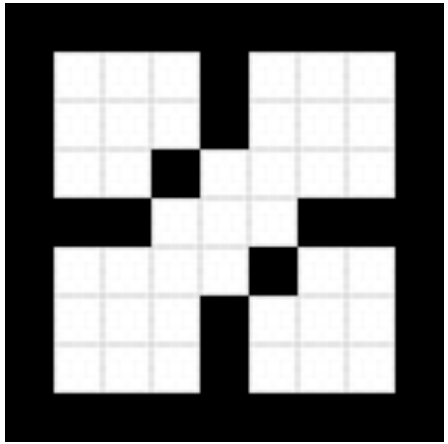

Flood fill with 4-connectivity

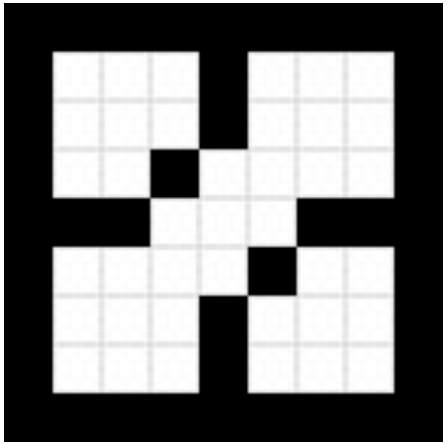

Flood fill with 8-connectivity

Wikipedia

Edge detection

The process of identifying parts of a digital image with sharp changes (discontinuities) in image intensity.

James Tompkin,
Brown University

Causes of edges

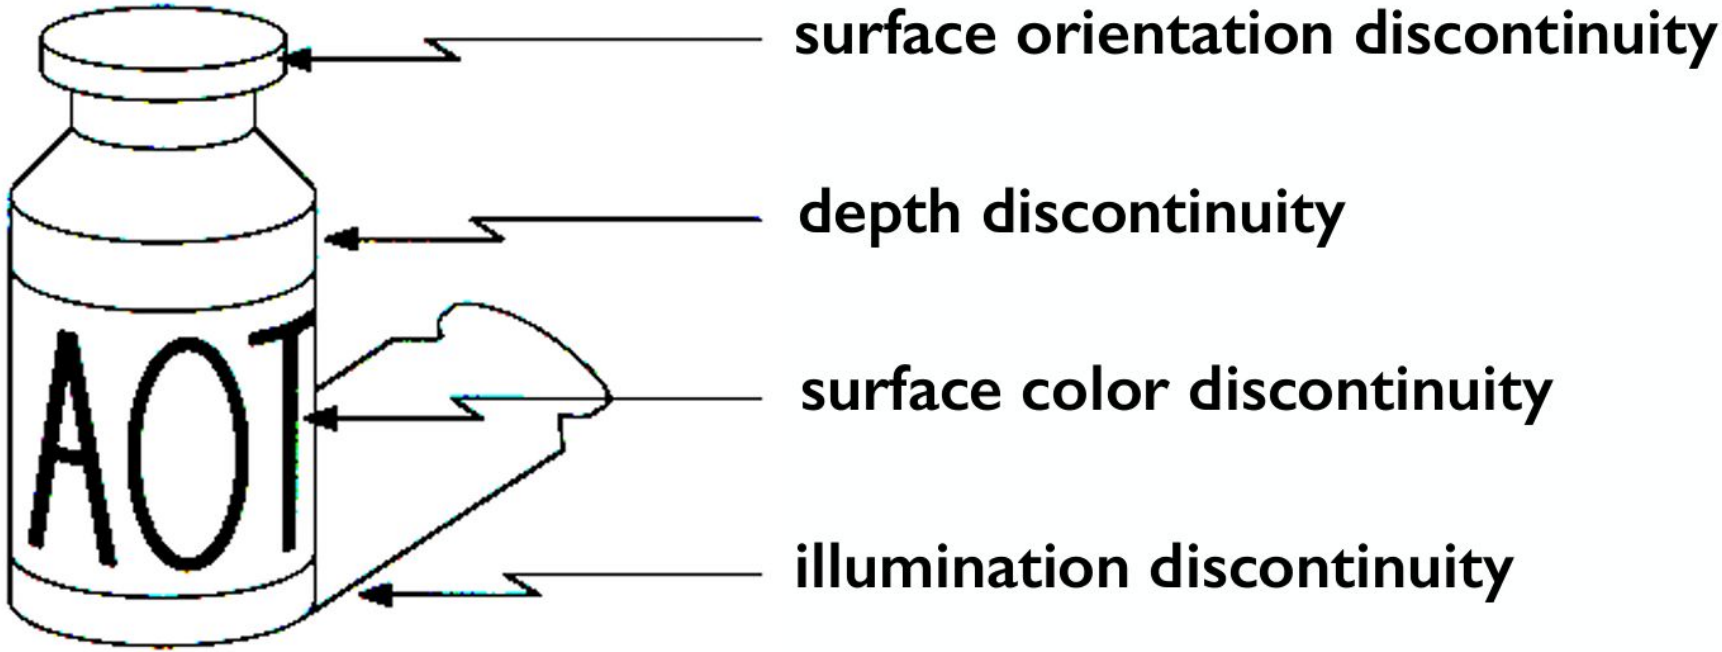

Steve Seitz, UW

Characterizing edges

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

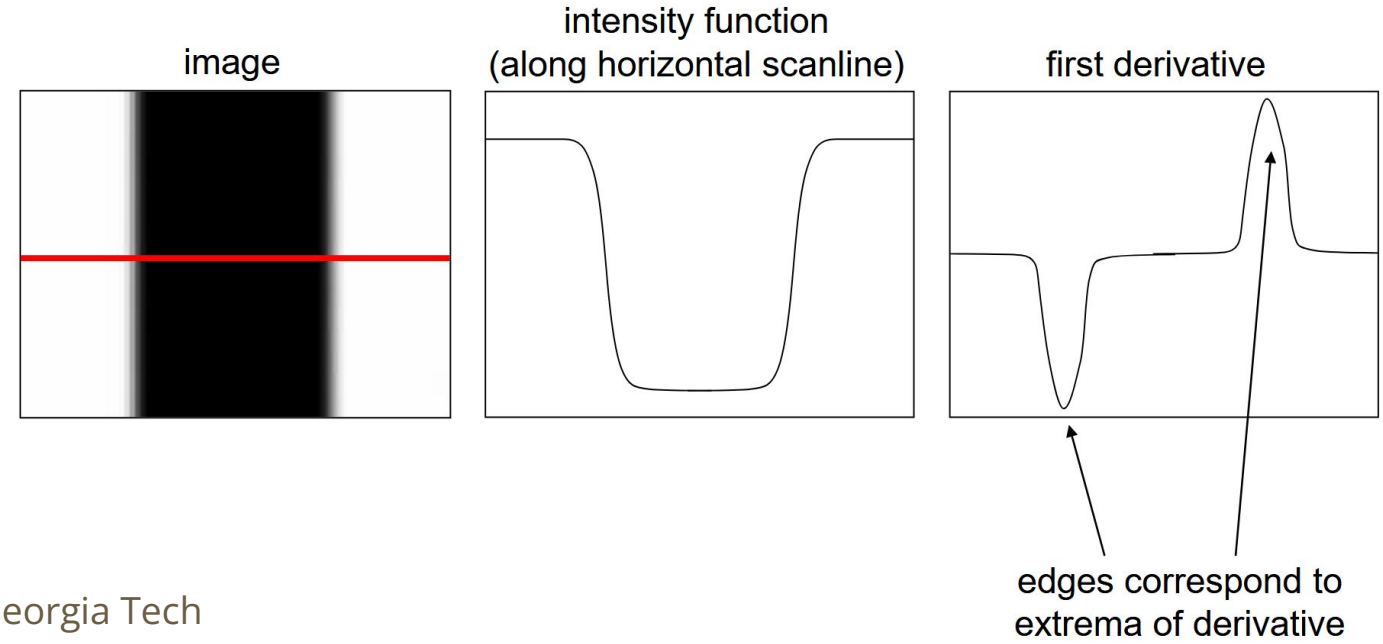

Image gradient

- Take discrete derivative (finite difference): $\frac{\partial f}{\partial x}[x, y] \approx F[x+1, y] - F[x, y]$

- The gradient of an image: $\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$

The gradient points in the direction of most rapid increase in intensity

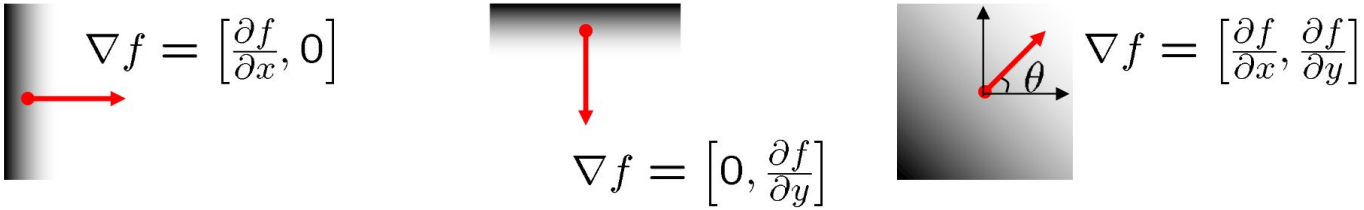

The gradient direction is given by: $\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x}\right)$

- how does this relate to the direction of the edge?

The *edge strength* is given by the gradient magnitude: $\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$

Differentiation and convolution

- Discrete derivative: $\frac{\partial f}{\partial x}[x, y] \approx F[x+1, y] - F[x, y]$
- This is the dot product of the vector [-1 1] with each pixel in the image

$$[F[x,y] \quad F[x+1,y]] \cdot [-1 \; 1]$$

- The operation of sliding a filter across an image is called convolution
- At the image boundary:
  - Extend the image using the pixel values at the boundary
  - Ignore pixels where the filter would extend beyond the image boundary
- Two-dimensional convolution: sum of element-wise products

Effects of noise

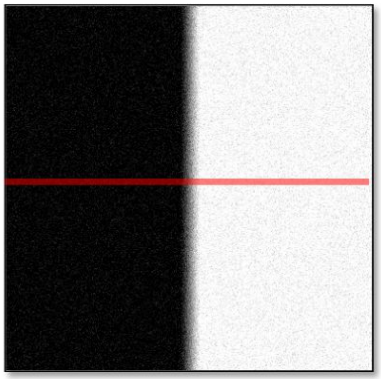

Noisy input image

$f(x)$

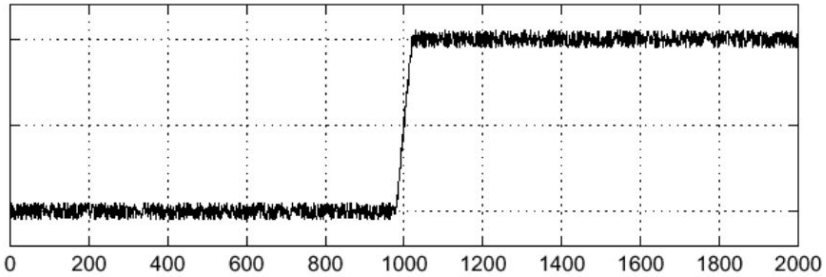

$\frac{d}{dx}f(x)$

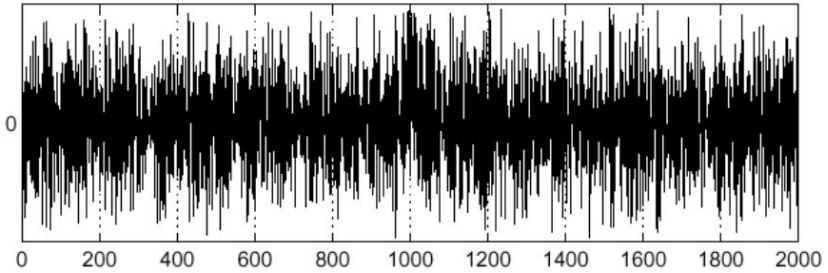

Where is the edge?

Solution: smoothing filter

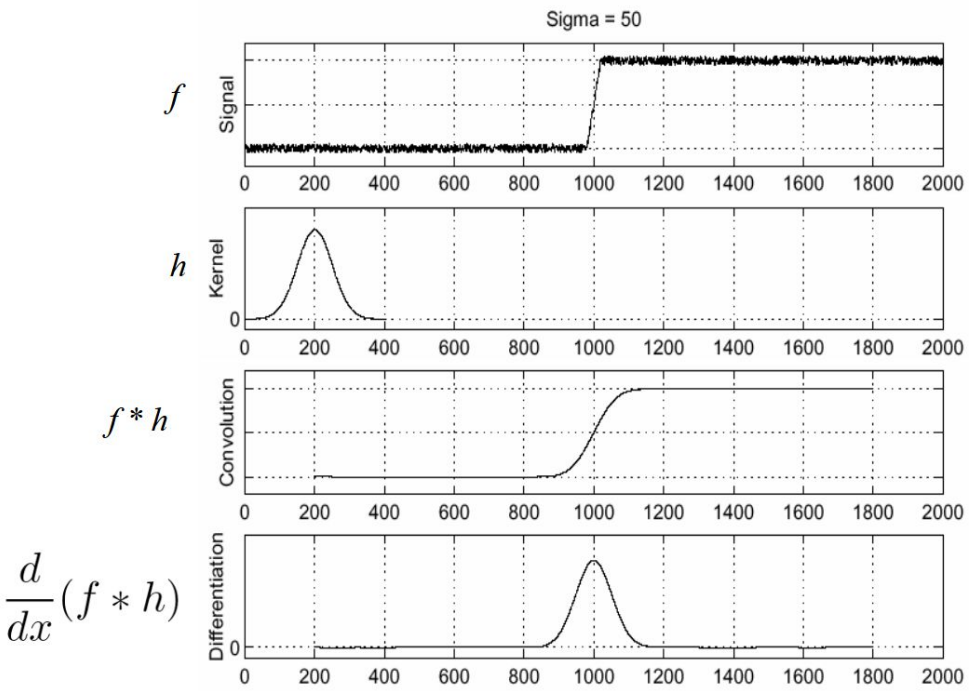

To find edges, look for peaks in $\frac{d}{dx}(f * h)$

Mask size

Wider masks lead to uncertainty about location of edge.

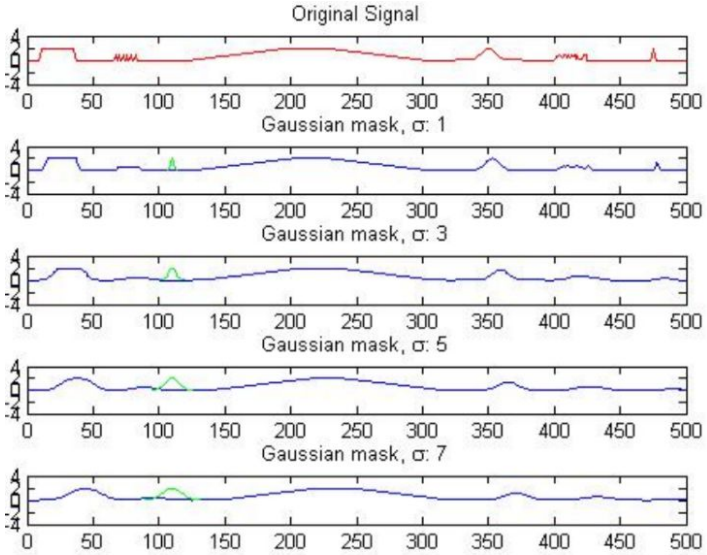

Convolution with Gaussian distribution

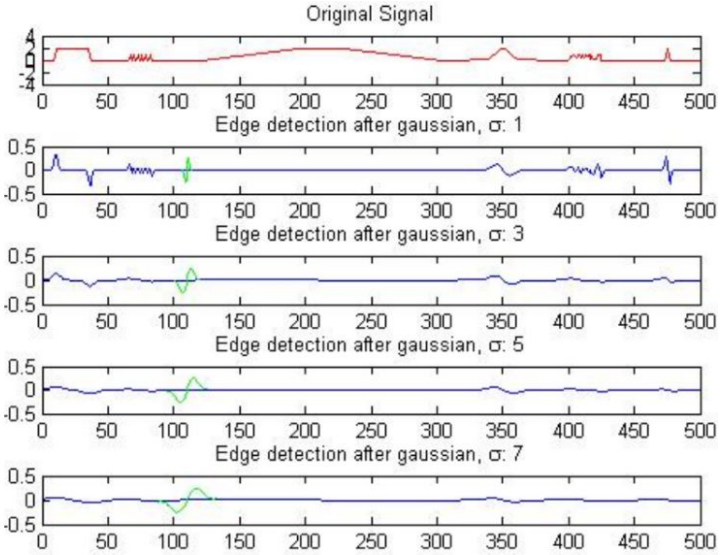

Edge detection on a signal that was convolved with a Gaussian

Canny edge detector

J. Canny, A Computational Approach To Edge Detection, IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986

- Classic and very accurate edge detector
- Process
  - Filter image with derivative of Gaussian
  - Find magnitude and orientation of gradient
  - Non-maximum suppression: thin multi-pixel 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

Hysteresis thresholding

original image

high threshold
(strong edges)

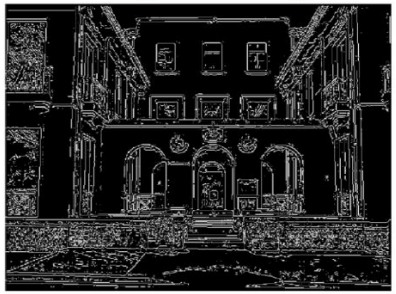

low threshold
(weak edges)

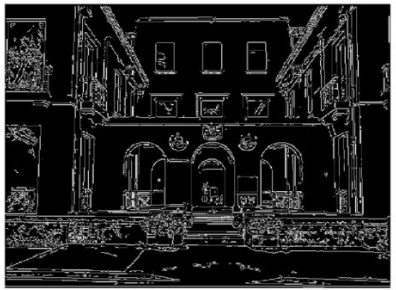

hysteresis threshold