Edge detection

Winter in Kraków photographed by Marcin Ryczek

Overview

- Motivating edge detection
- Image gradients
- Derivative of Gaussian filters
- Canny edge detector
- Role of edge detection in image understanding
- Orientations

Edge detection

- **Goal:** Identify sudden changes (discontinuities) in an image
- Intuitively, edges carry most of the semantic and shape information from the image

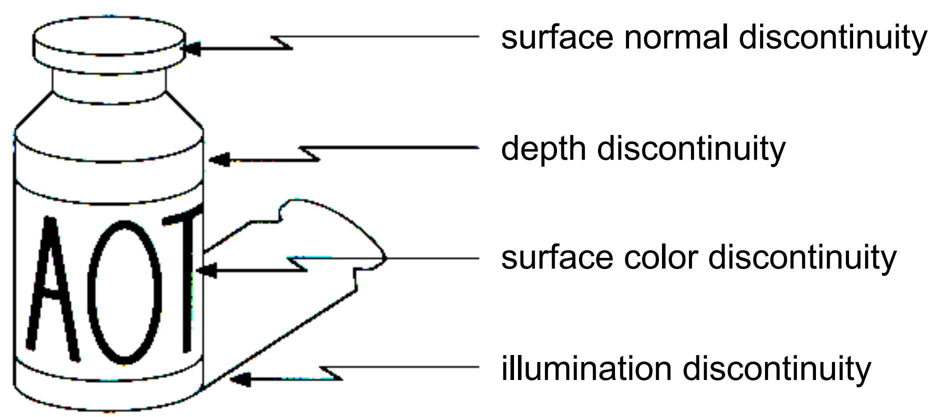

Sources: D. Lowe and S. Seitz

Edge detection

Input photo

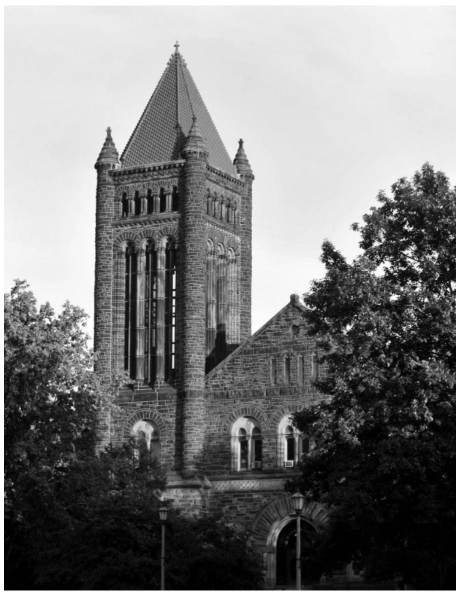

Ideal: artist's line drawing

Image source

Reality

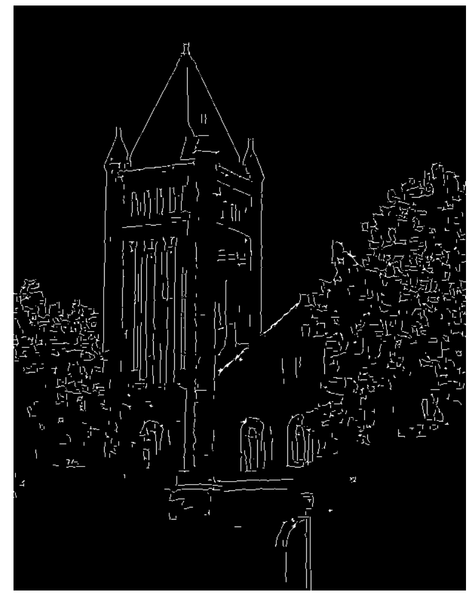

Source: S. Gupta

Edge detection

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

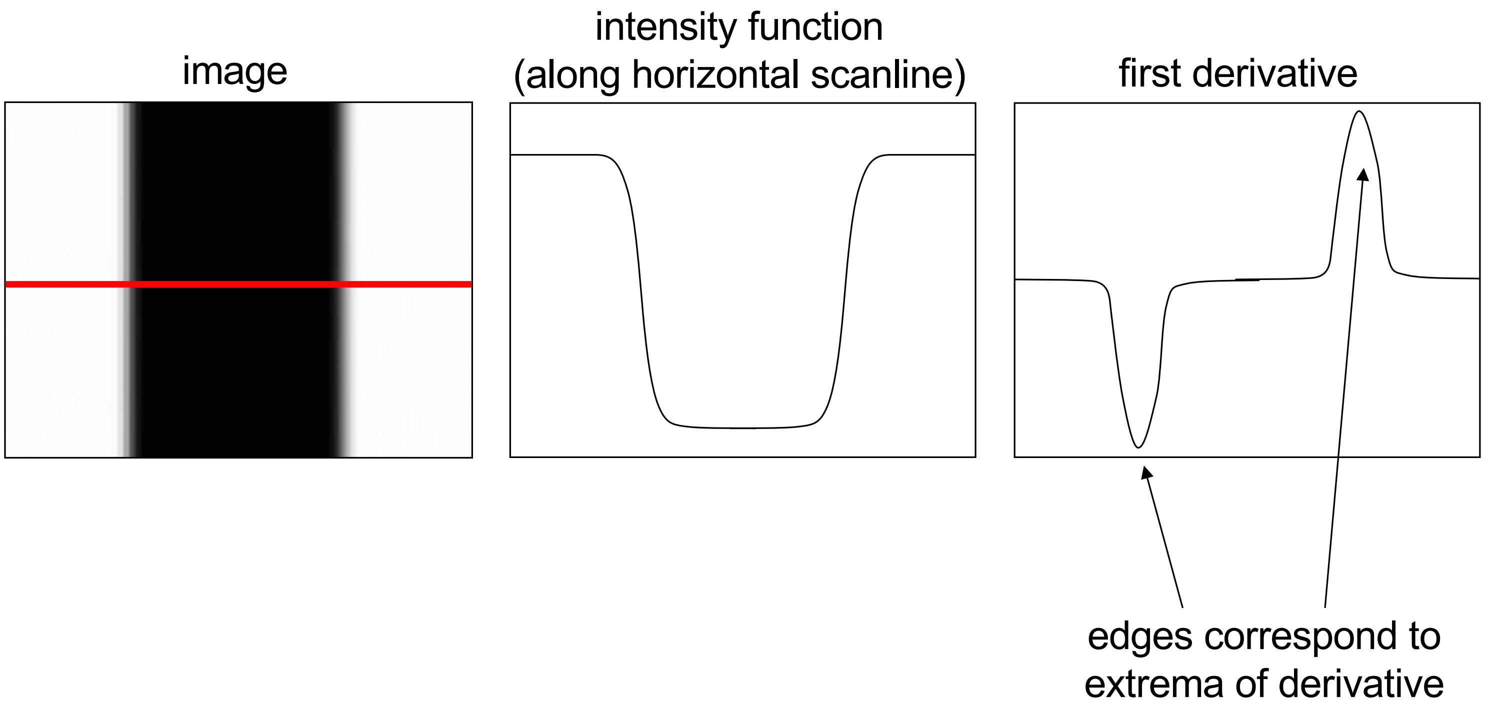

Partial derivatives of an image

- For 2D function $f(x, y)$, the partial derivative w.r.t. x is

$$\frac{\partial f(x,y)}{\partial x} = \lim_{h \to 0} \frac{f(x+h,y) - f(x,y)}{h}$$

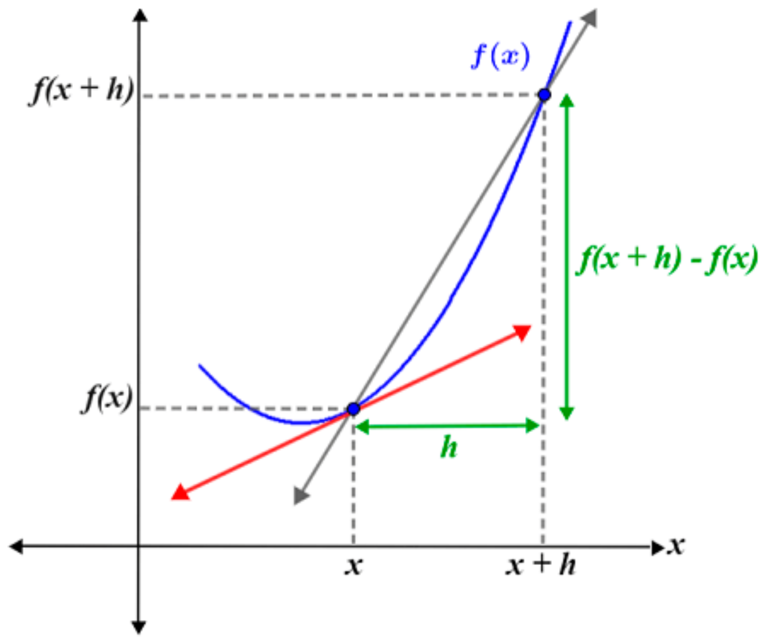

Image source

Partial derivatives of an image

- For 2D function $f(x, y)$, the partial derivative w.r.t. x is

$$\frac{\partial f(x,y)}{\partial x} = \lim_{h \to 0} \frac{f(x+h,y) - f(x,y)}{h}$$

- For discrete data, we can approximate using finite differences:

$$\frac{\partial f(x,y)}{\partial x} = \frac{f(x+1,y) - f(x,y)}{1}$$

- To implement the above as convolution, what would be the associated filter?

Source: K. Grauman

Overview

- Motivating edge detection
- Image gradients
- Derivative of Gaussian filters

“Fun” facts

Convolution is commutative

f*g=g*f

(easy proof: use the convolution theorem, multiplication commutes!)

Differentiation can be represented with convolution

there is some k so that

$$\frac{df}{dx} = k * f = f * k$$

(Won’t prove this, but we’ve sort of seen it already)

Finding noisy edges

- Consider a single row or column of the image:

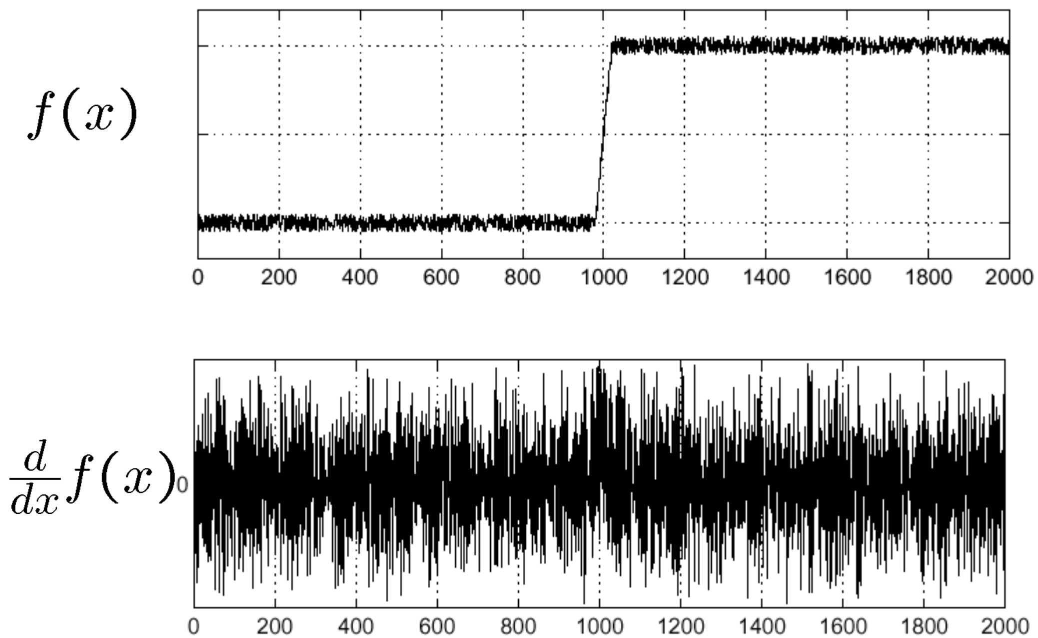

- Where is the edge?

Source: S. Seitz

Filtering with derivative of Gaussian

f

$\frac{d}{dx} g$

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

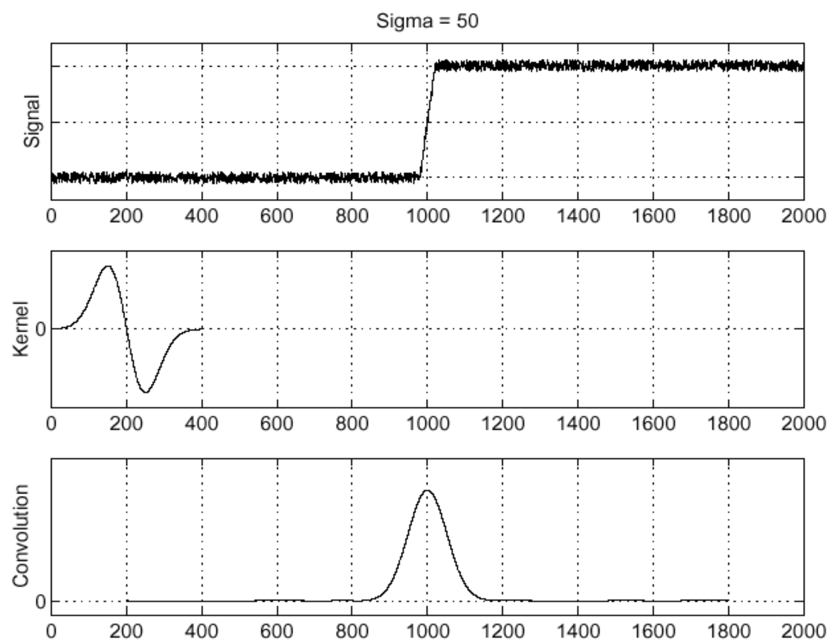

Source: S. Seitz

2D Derivative of Gaussian (d.o.g. or dog) filters

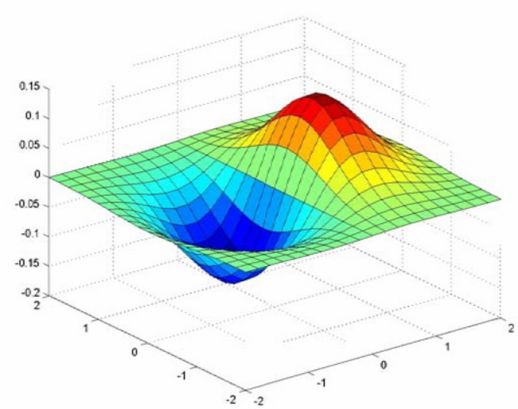

x-direction

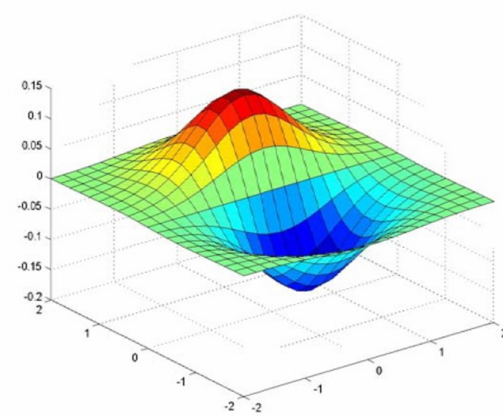

y-direction

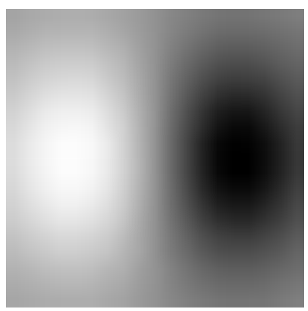

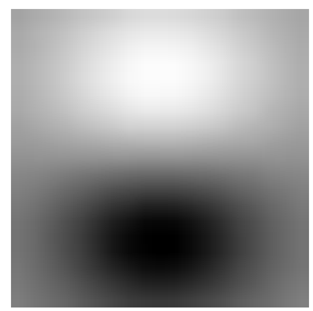

Separability of the Gaussian filter

$$\frac{1}{2\pi\sigma^2}\exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) =$$

$$\frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{x^2}{2\sigma^2}\right) \quad \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{y^2}{2\sigma^2}\right)$$

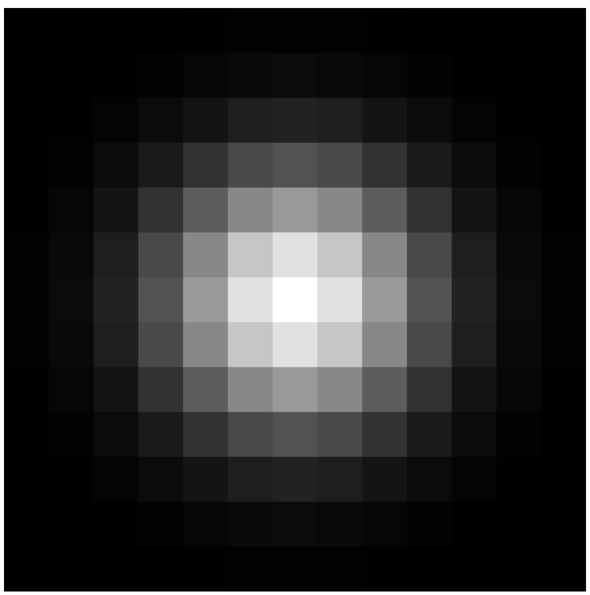

=

*

Note: matrix multiplication (outer product) and full convolution are the same in this special case

Adapted from D. Fouhey and J. Johnson

2D Derivative of Gaussian filters

- (Unnormalized) 2D Gaussian:

$$g(x,y) \propto \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) = \exp\left(-\frac{x^2}{2\sigma^2}\right)\exp\left(-\frac{y^2}{2\sigma^2}\right)$$

- (Unnormalized) Gaussian derivatives:

$$\frac{\partial g}{\partial x} \propto -x\exp\left(-\frac{x^2}{2\sigma^2}\right)\exp\left(-\frac{y^2}{2\sigma^2}\right)$$

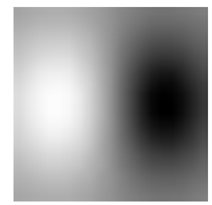

$$\frac{\partial g}{\partial y} \propto -y\exp\left(-\frac{y^2}{2\sigma^2}\right)\exp\left(-\frac{x^2}{2\sigma^2}\right)$$

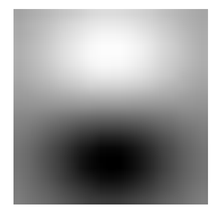

- These are products of a 1D Gaussian in one direction and 1D derivative of Gaussian in the other direction!

Derivative of Gaussian: Scale

- Using Gaussian derivatives with different values of σ finds structures at different scales or frequencies

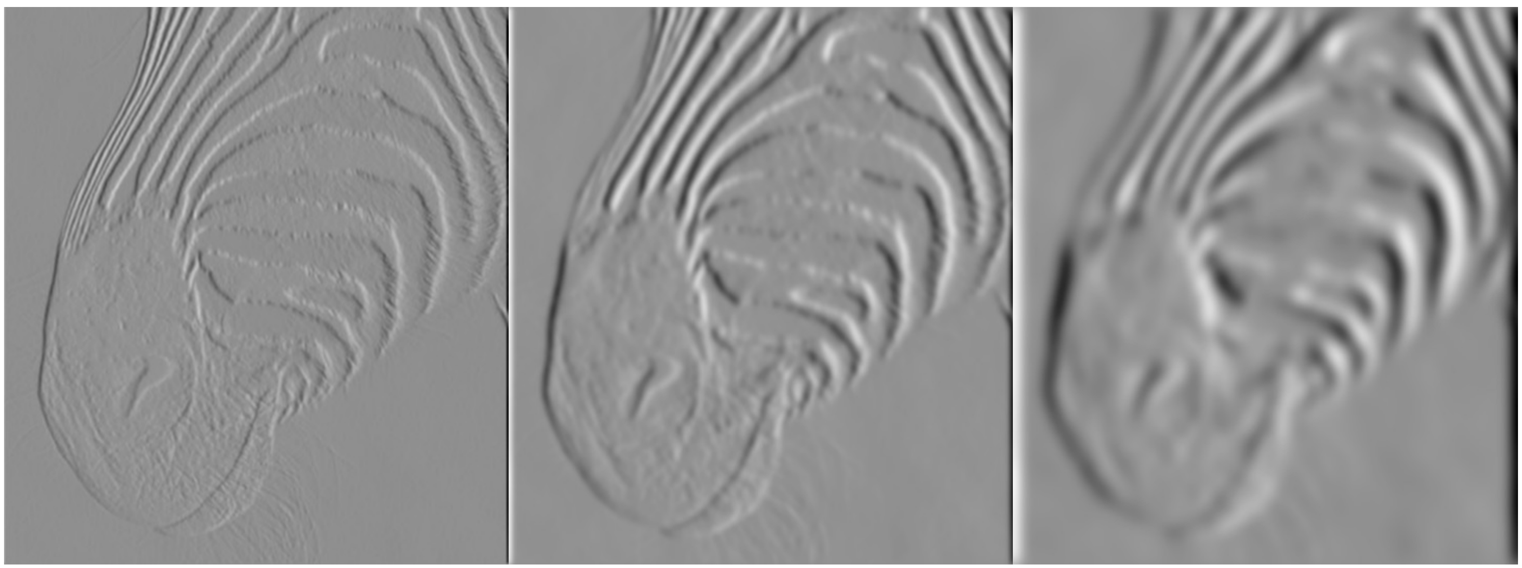

$\sigma = 1$ $\sigma = 3$ $\sigma = 7$

Source: D. Forsyth

Summing up: Types of filters

- Smoothing filters
  - Gaussian: remove "high-frequency" components; "low-pass" filter
  - Can the values of a smoothing filter be negative?
  - What should the values sum to?
    - **One:** constant regions are not affected by the filter

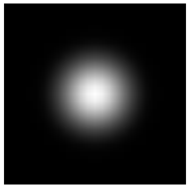

- Derivative filters
  - Derivatives of Gaussian: compute smoothed differences; "band-pass" filters
  - Can the values of a derivative filter be negative?
  - What should the values sum to?
    - **Zero:** no response in constant regions

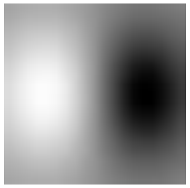

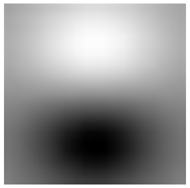

Edge detection: Overview

- Motivating edge detection
- Image gradients
- Derivative of Gaussian filters
- Canny edge detector

Building an edge detector

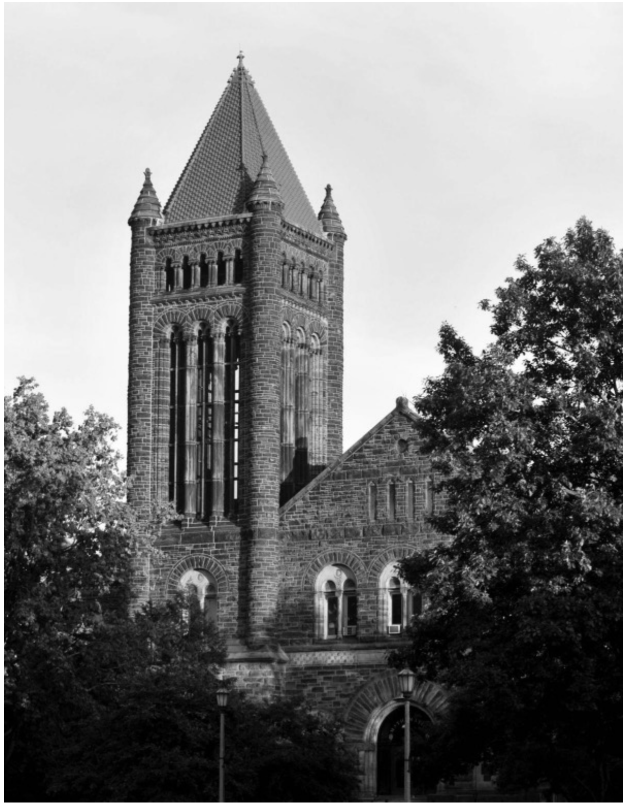

original image

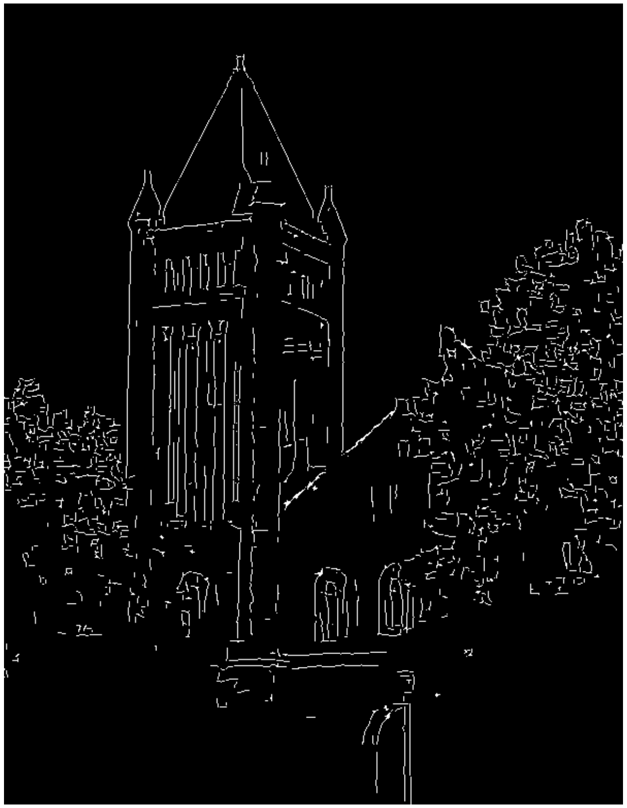

final output

J. Canny, A Computational Approach To Edge Detection, IEEE Trans. PAMI, 8:679-714, 1986

Building an edge detector

1. Compute x and y derivative images
2. Find magnitude and orientation of the gradient

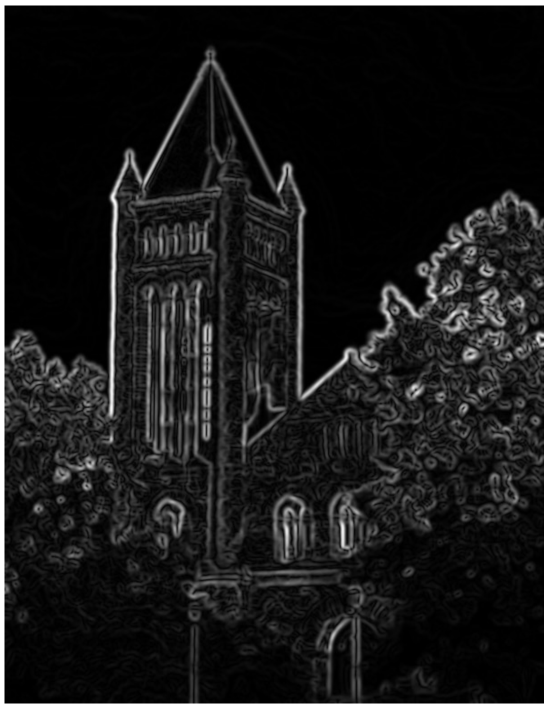

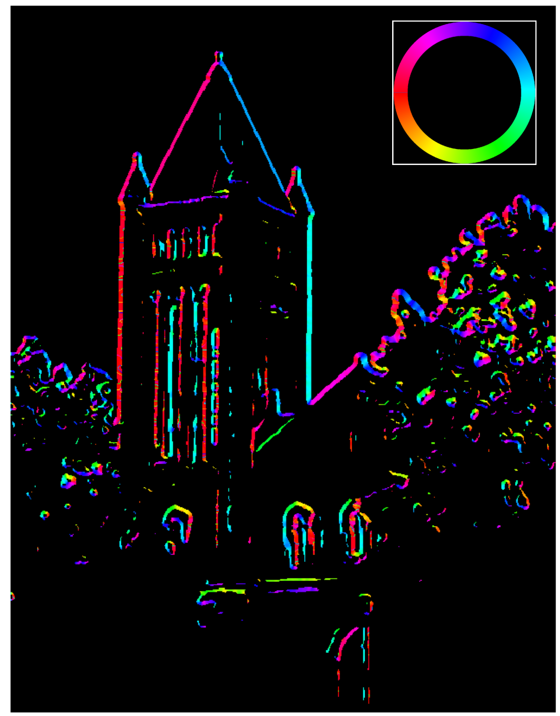

Source: S. Gupta

Building an edge detector

1. Compute x and y derivative images
2. Find magnitude and orientation of the gradient

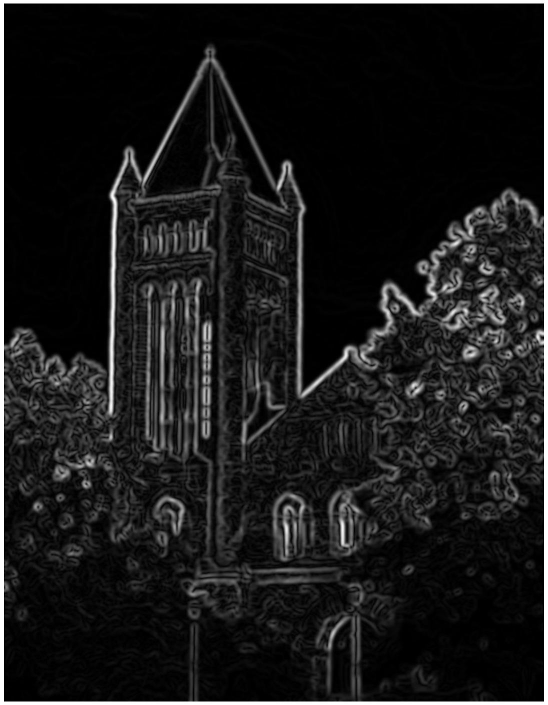

Let's threshold the gradient magnitude

Source: S. Gupta

Building an edge detector

1. Compute x and y derivative images
2. Find magnitude and orientation of the gradient

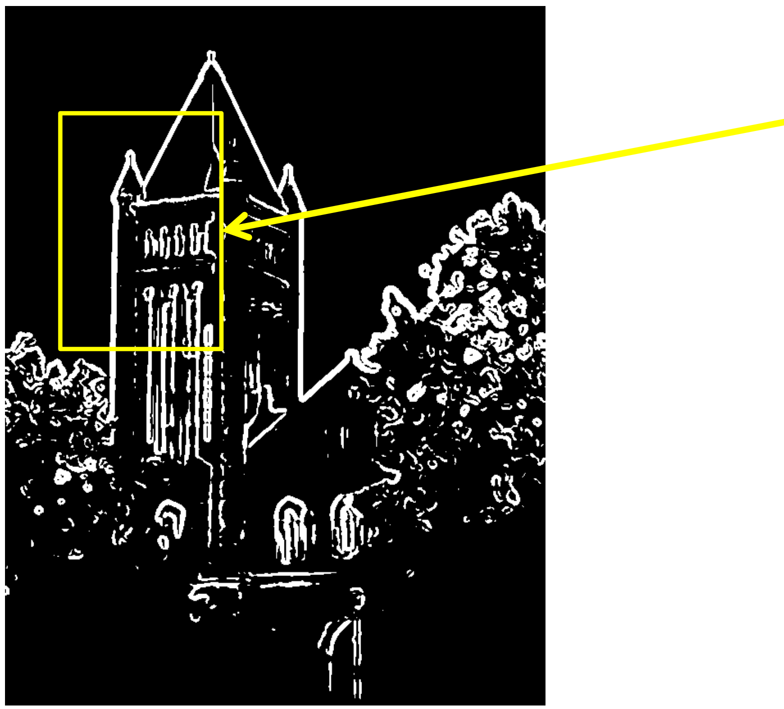

We get thick trails, not neat edge curves

Let's threshold the gradient magnitude

Source: S. Gupta

Non-maximum suppression

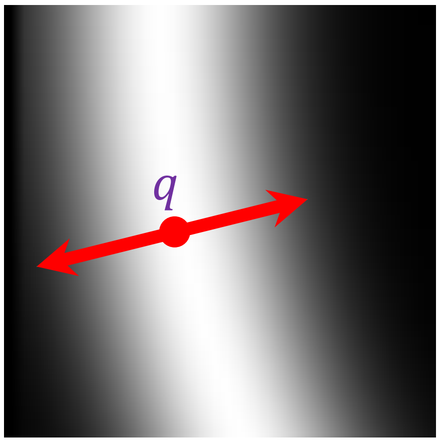

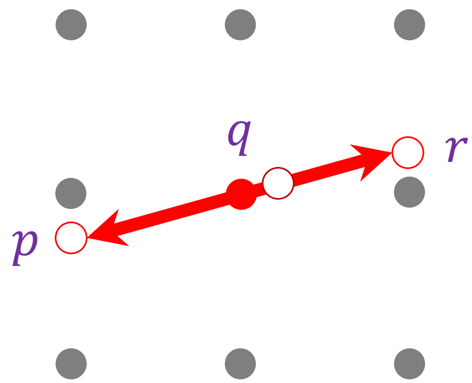

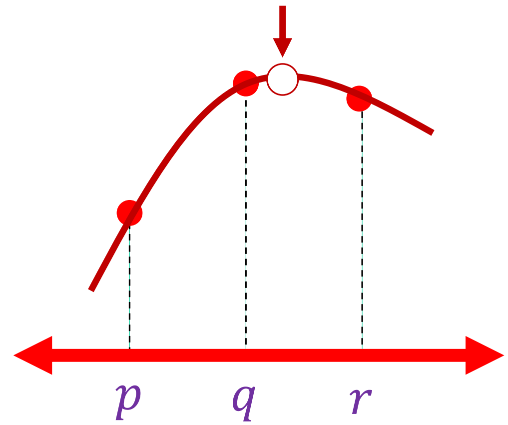

1D image "slice" normal to the edge

- For each location q above threshold, check that the gradient magnitude is higher than at adjacent points p and r along the direction of the gradient
  - Need to interpolate to get the gradient magnitude values at p and r
  - Can even use nonlinear interpolation to get sub-pixel edge localization!

Non-maximum suppression

NMS

NMS > threshold

Another problem: pixels along this edge didn't survive the thresholding

Source: S. Gupta

Hysteresis thresholding

- Use a high threshold to start edge curves, and a low threshold to continue them

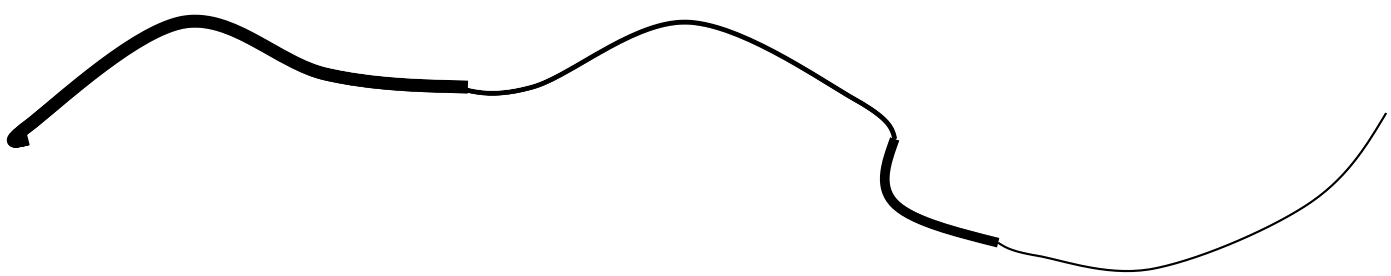

Source: Steve Seitz

Hysteresis thresholding

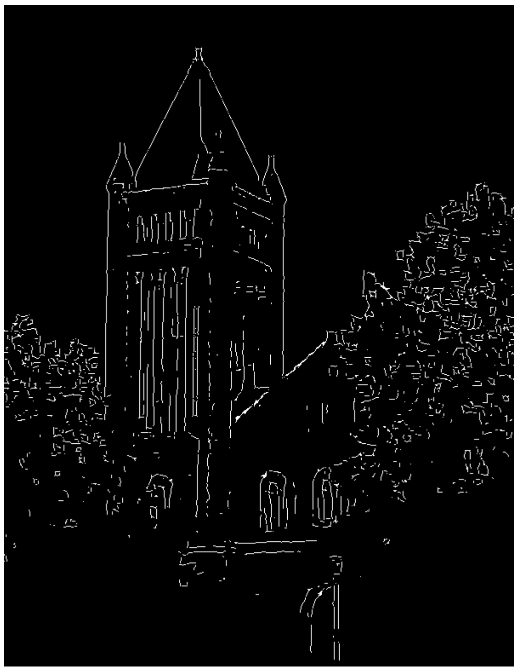

high threshold
(strong edges)

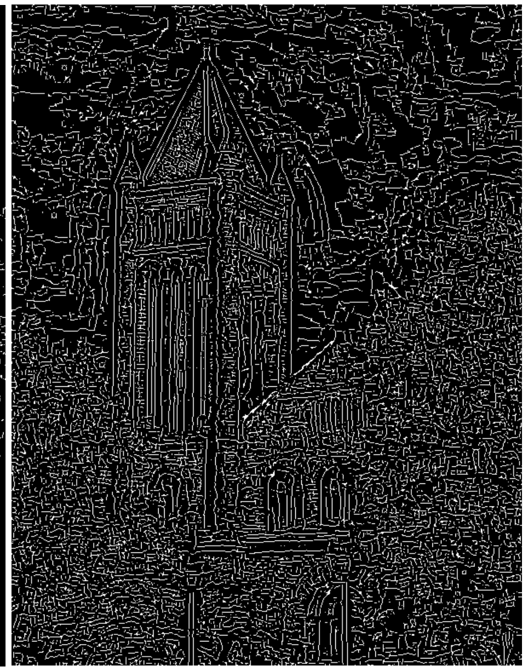

low threshold
(weak edges)

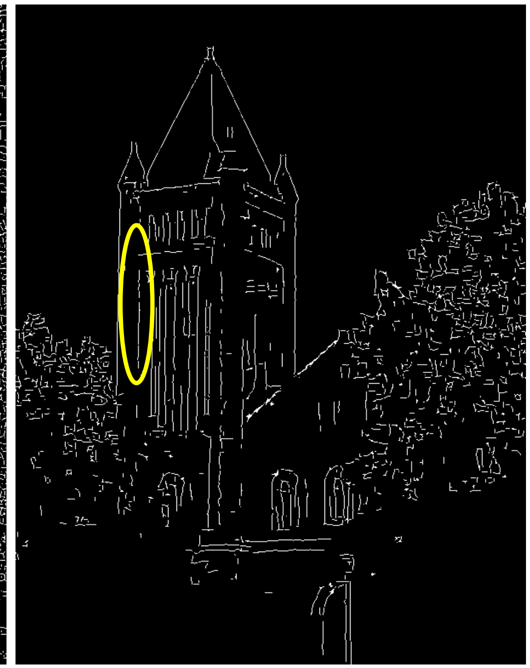

hysteresis threshold

Source: S. Gupta

Recap: Canny edge detector

1. Compute x and y derivative images
2. Find magnitude and orientation of the gradient
3. Non-maximum suppression:
   - Thin wide “ridges” down to single pixel width
4. 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

J. Canny, A Computational Approach To Edge Detection, IEEE Trans. PAMI, 8:679-714, 1986.

Overview

- Motivating edge detection
- Image gradients
- Derivative of Gaussian filters
- Canny edge detector
- What is the role of edge detection in image understanding?

Are edges an “input” or an “output”?

Figure from Marr (1982), attributed to R. C. James

Image gradients vs. meaningful contours

Image | Human segmentation | Gradient magnitude

D. Martin, C. Fowlkes, D. Tal, and J. Malik. [A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics](). ICCV 2001

Data-driven edge detection

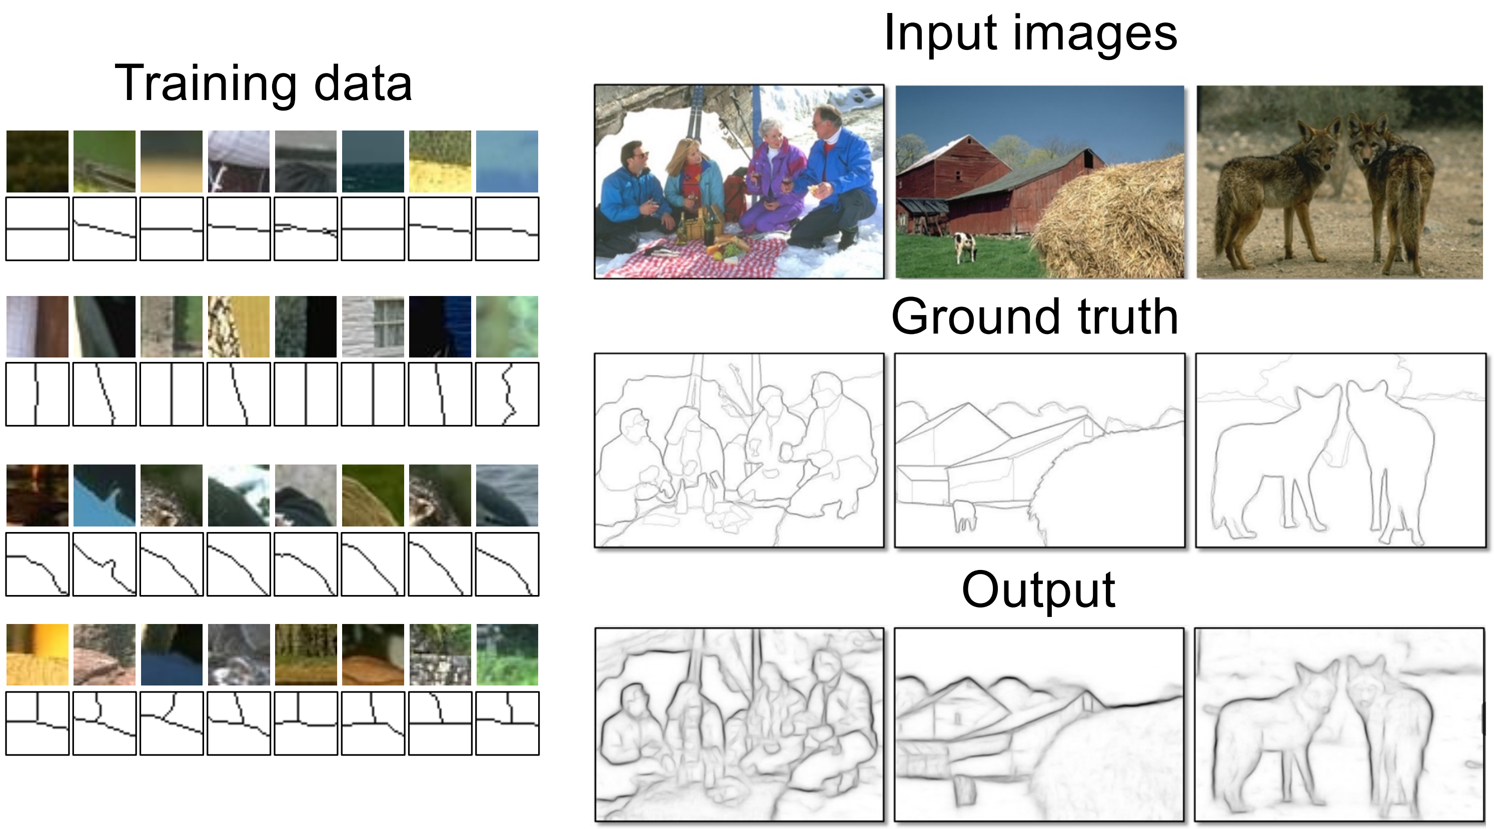

P. Dollar and L. Zitnick, [Structured forests for fast edge detection](), ICCV 2013

Data-driven edge detection

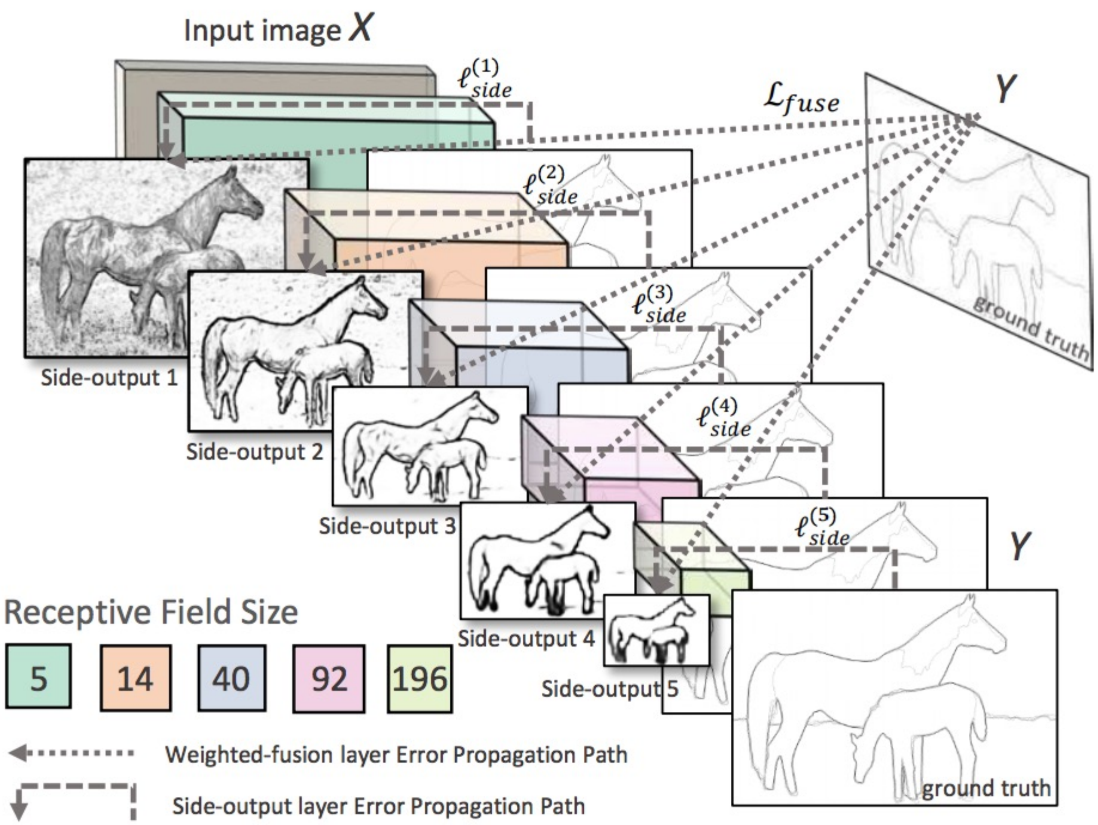

S. Xie and Z. Tu, [Holistically-nested edge detection](), ICCV 2015

Most successful approach in practice: Top-down segmentation

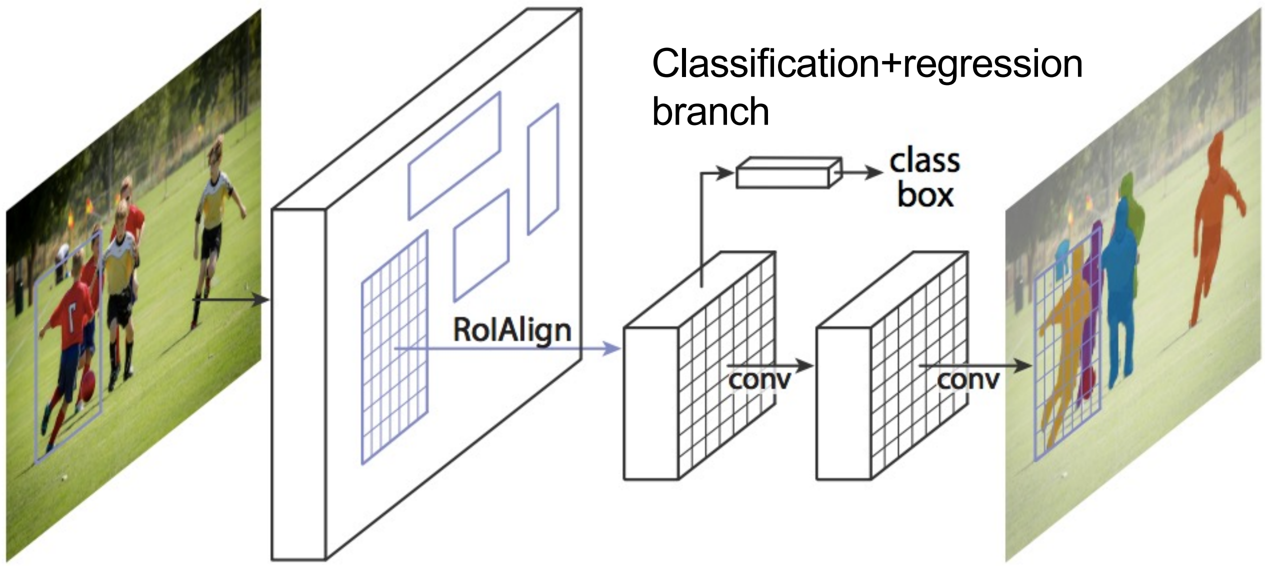

Mask branch: output binary segmentation for each instance

K. He, G. Gkioxari, P. Dollar, and R. Girshick, [Mask R-CNN](), ICCV 2017 (Best Paper Award)

Overview

- Motivating edge detection
- Image gradients
- Derivative of Gaussian filters
- Canny edge detector
- Role of edge detection in image understanding
- Orientations

Problem:

Scaling the image scales gradient magnitude

$$f \to kf \text{ implies } \sqrt{\frac{\partial f}{\partial x}^2 + \frac{\partial f}{\partial y}^2} \to k\sqrt{\frac{\partial f}{\partial x}^2 + \frac{\partial f}{\partial y}^2}$$

Which causes problems with thresholds, etc

- image gets darker, some edges disappear
- image gets lighter, some edges appear

Hysteresis helps, but doesn't cure

Gradient orientations don't change with illumination

$$f \to kf \text{ implies } \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right) \to k\left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right)$$

Q: build a representation out of orientations?

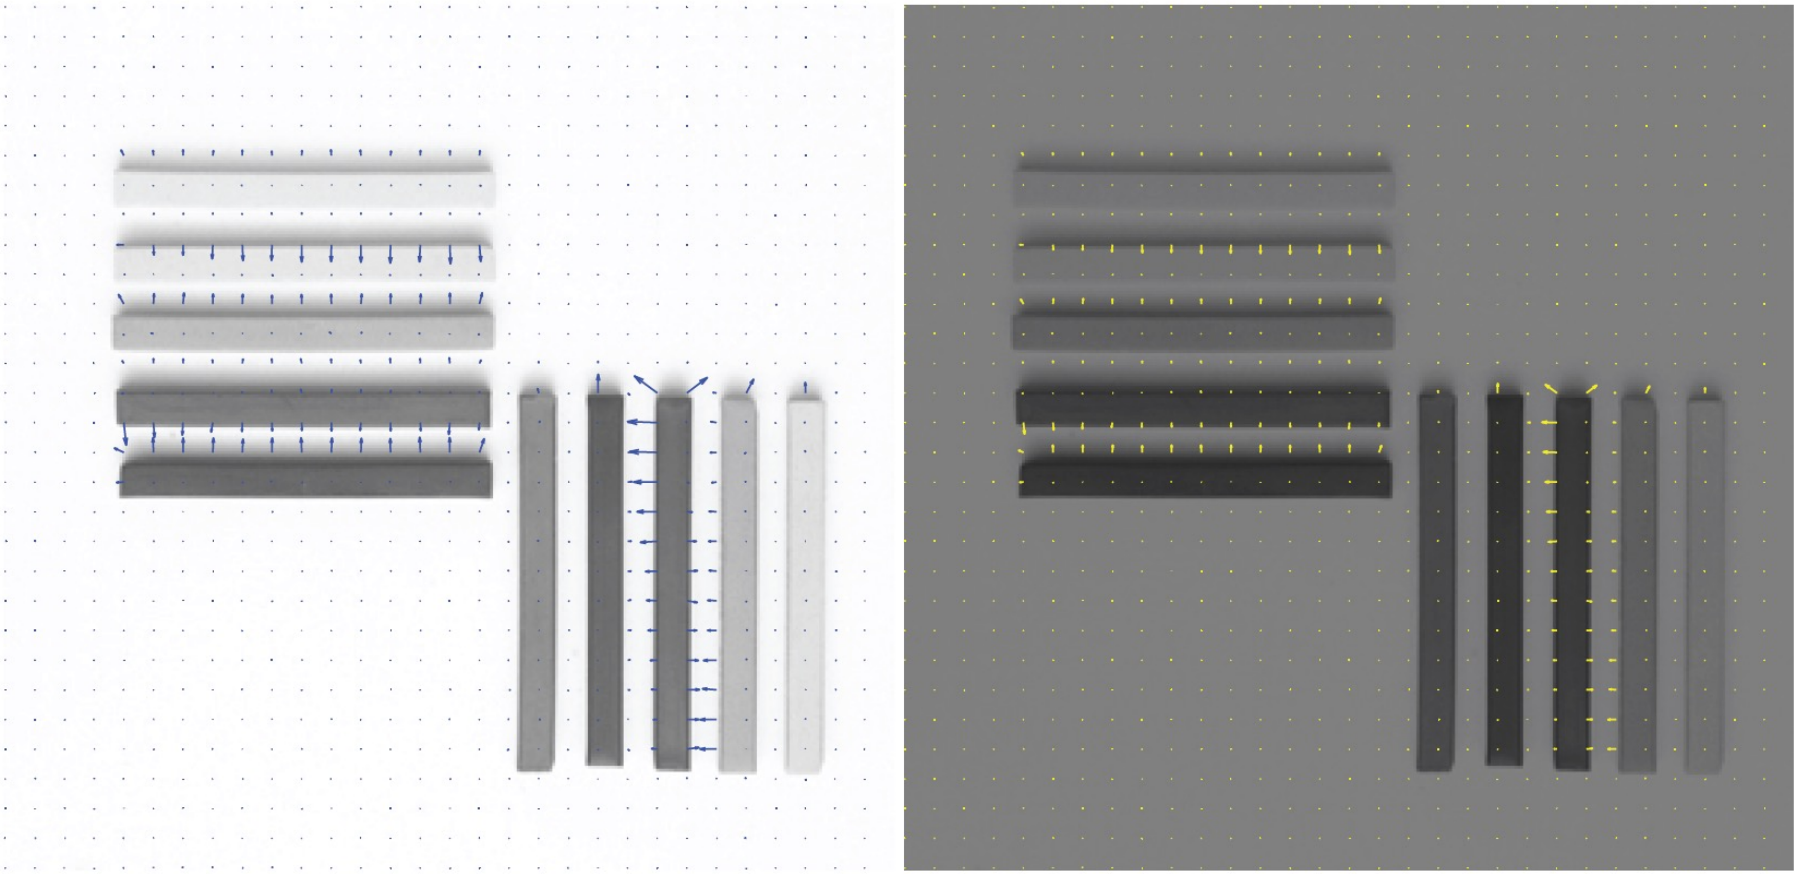

Orientation histograms

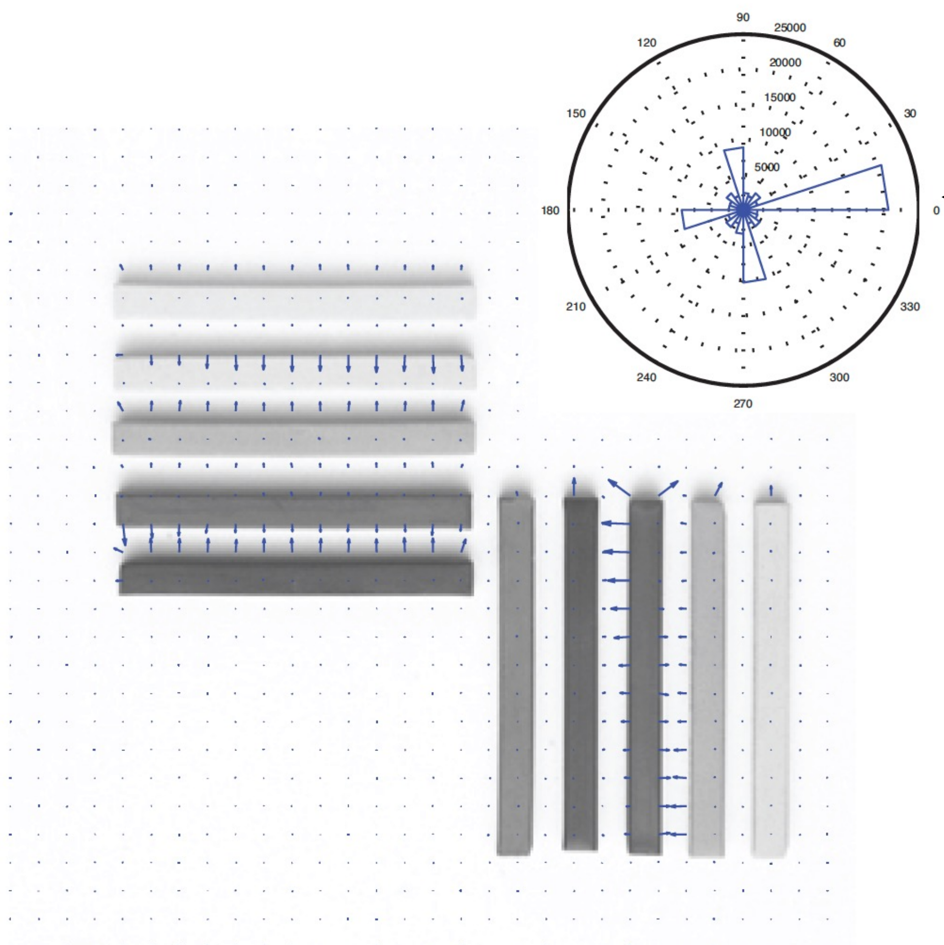

Rose plot, which shows number of vectors at each range of angles

Orientation histogram depends on scale

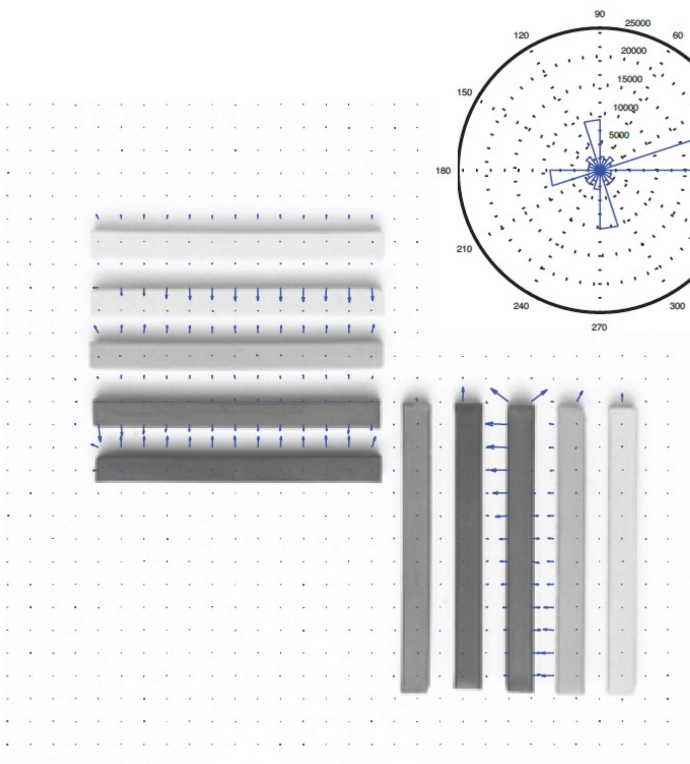

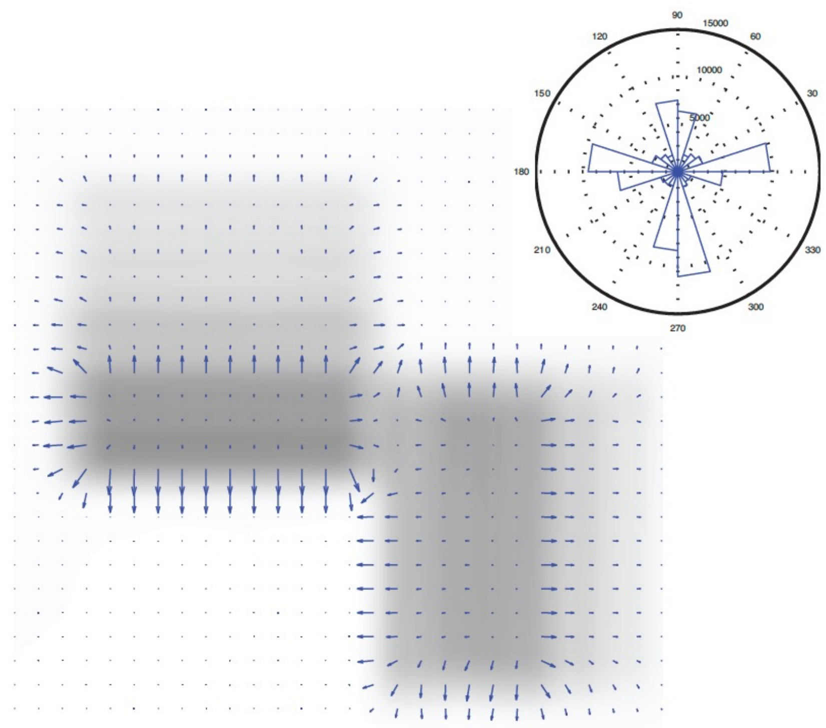

Different patterns have different orientation histograms

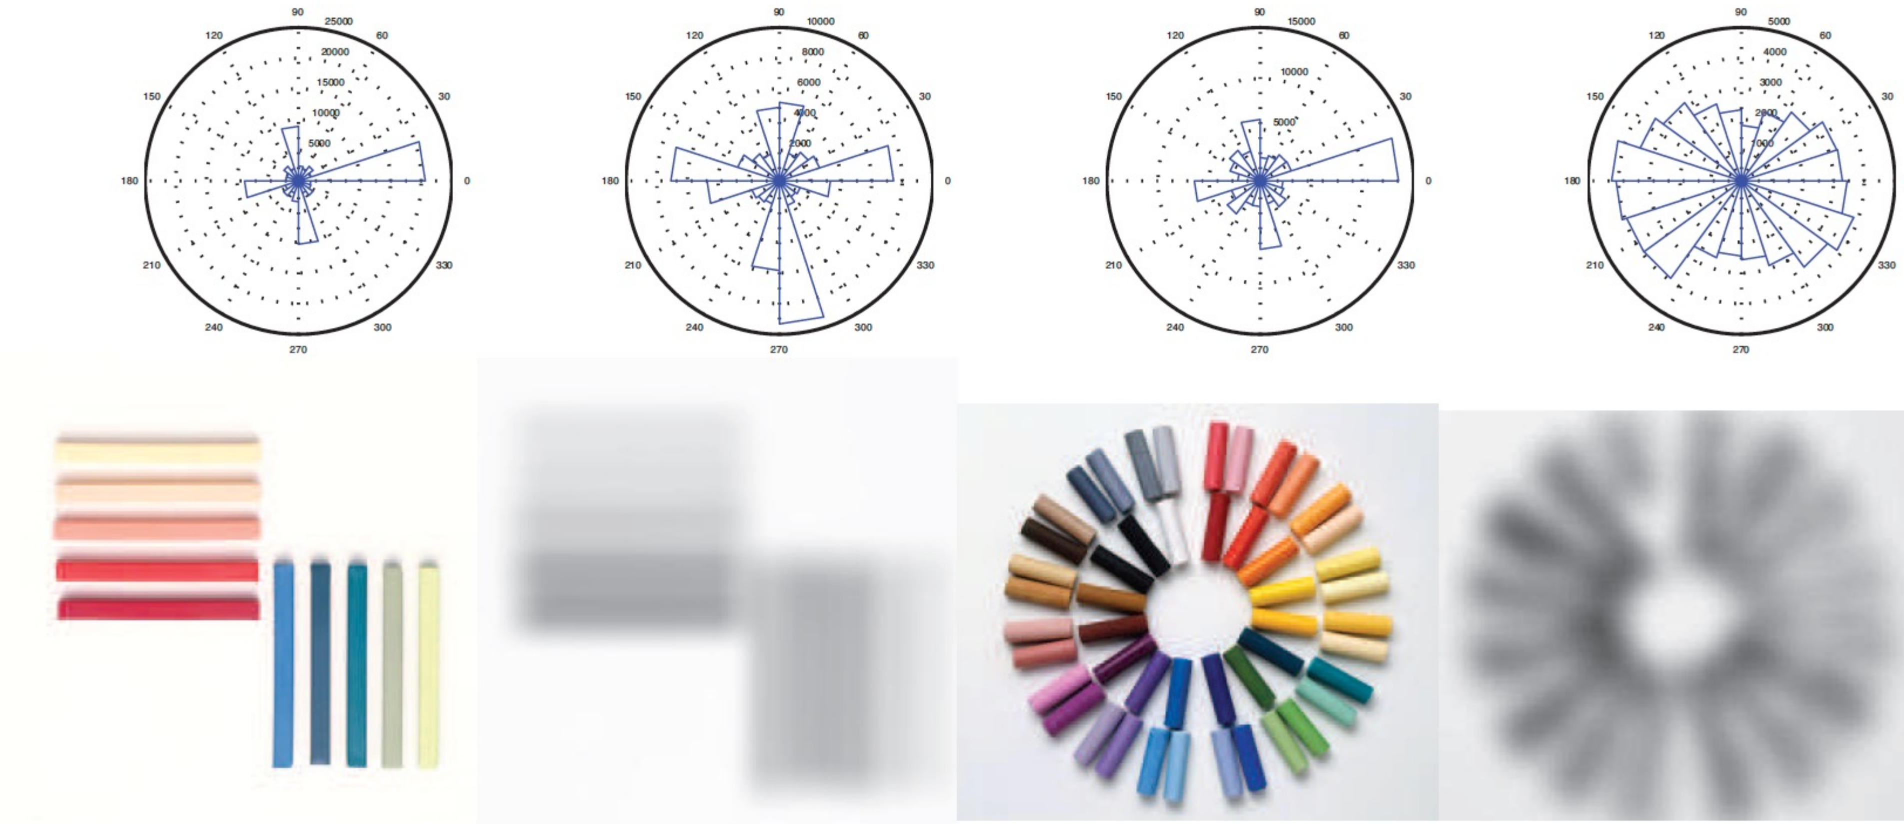

Notice something important…