



CSCI 497/597P: Computer Vision

Scott Wehrwein

Edge Detection, Continued



Reading


- Szeliski, Chapter 4.2

Announcements

Goals

- Understand the basics of edge detection:
 - The sobel operator as an approximation of the image gradient in the presence of noise.
- Understand the use of non-maximum suppression to localize edges from smoothed gradients.

Images as functions...



- Edges look like steep cliffs

Characterizing edges

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




Image derivatives

- How can we differentiate a *digital* image $F[x,y]$?
 - Option 1: reconstruct a continuous image, f , then compute the derivative
 - Option 2: take discrete derivative (finite difference)

$$\frac{\partial f}{\partial x}[x, y] \approx F[x + 1, y] - F[x, y]$$

How would you implement this as a linear filter?

$$\frac{\partial f}{\partial x} : \begin{array}{|c|c|c|} \hline & & \\ \hline & & \\ \hline & & \\ \hline \end{array}$$

H_x




$$\frac{\partial f}{\partial y} : \begin{array}{|c|c|c|} \hline & & \\ \hline & & \\ \hline & & \\ \hline \end{array}$$

H_y

Image gradient

- 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 *edge strength* is given by the gradient magnitude:



$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

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?

Image gradient




Effects of noise




Noisy input image

$$f(x)$$

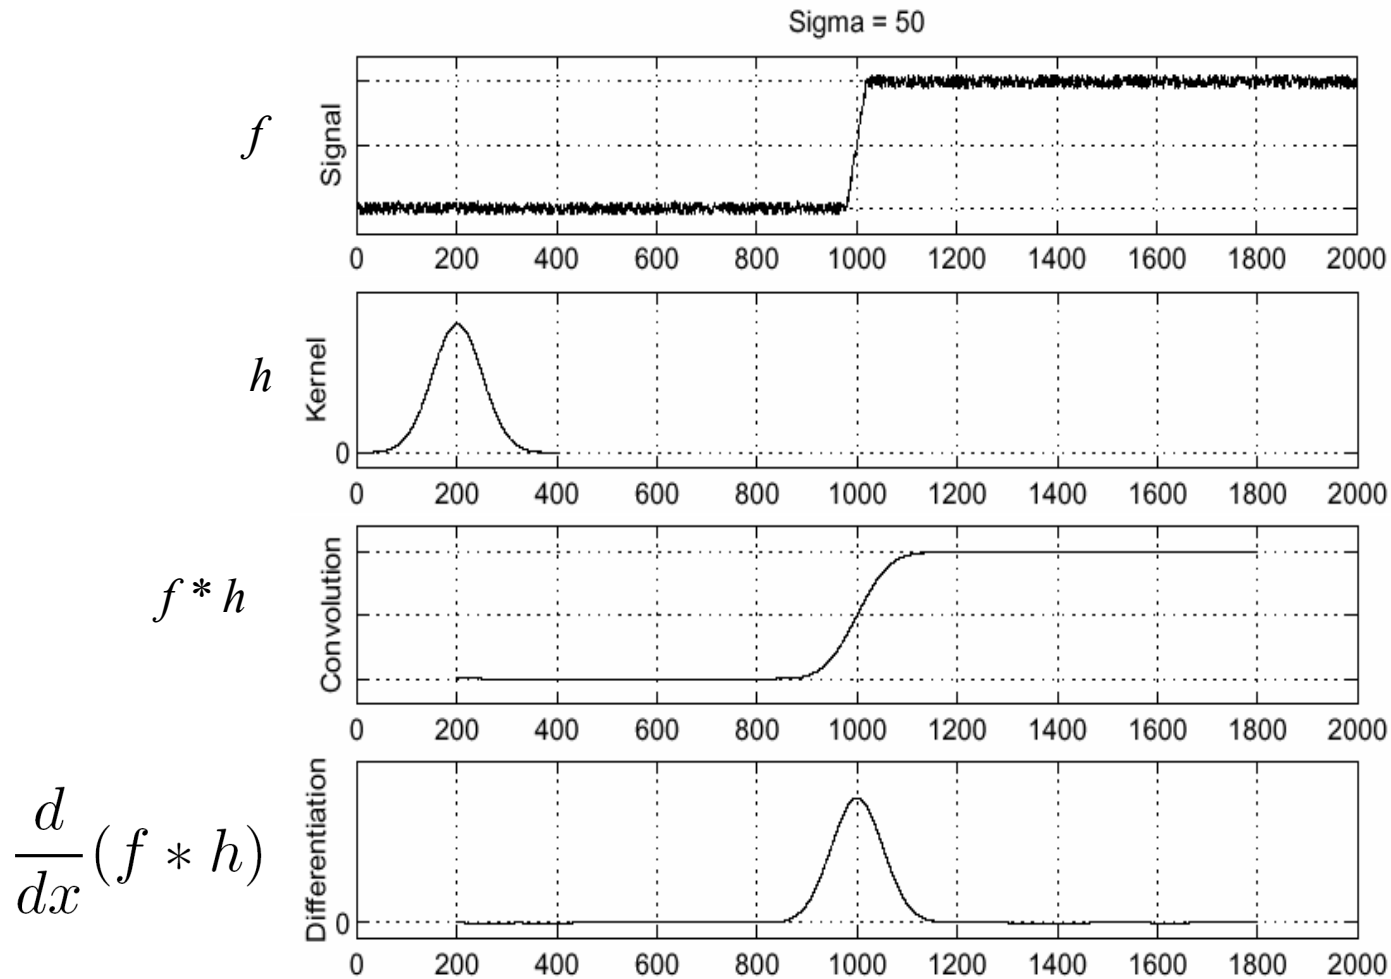


$$\frac{d}{dx}f(x)$$



Where is the edge?


Solution: smooth first



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

Associative property of convolution


- Differentiation is convolution, and convolution is associative: $\frac{d}{dx}(f * h) = f * \frac{d}{dx}h$
- This saves us one operation:




The 1D Gaussian and its derivatives

$$G_{\sigma}(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}$$

$$G'_{\sigma}(x) = \frac{d}{dx} G_{\sigma}(x) = -\frac{1}{\sigma} \left(\frac{x}{\sigma} \right) G_{\sigma}(x)$$




2D edge detection filters



Gaussian


$$h_{\sigma}(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$





derivative of Gaussian (x)

$$\frac{\partial}{\partial x} h_{\sigma}(u, v)$$


Derivative of Gaussian filter



x-direction



y-direction



The Sobel operator

- Common approximation of derivative of Gaussian

$$\frac{1}{8} \begin{array}{|c|c|c|} \hline -1 & 0 & 1 \\ \hline -2 & 0 & 2 \\ \hline -1 & 0 & 1 \\ \hline \end{array}$$

s_x

$$\frac{1}{8} \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 0 & 0 & 0 \\ \hline -1 & -2 & -1 \\ \hline \end{array}$$

s_y

- The standard defn. of the Sobel operator omits the $1/8$ term
 - doesn't make a difference for edge detection
 - the $1/8$ term **is** needed to get the right gradient magnitude

Sobel operator: example








Image with Edge



Edge Location






Image + Noise



Derivatives detect
edge *and* noise



Smoothed derivative removes
noise, but blurs edge

Criteria for a good boundary detector

- Criteria for a good boundary detector:
 - **Good detection:** Fire only on real edges, not anywhere else
 - **Good localization**
 - the edges detected must be as close as possible to the true edges
 - the detector must return one point only for each true edge point

Canny edge detector


- The classic edge detector
- Baseline for all later work on grouping
- Theoretical model: step-edges corrupted by additive Gaussian noise

J. Canny, [*A Computational Approach To Edge Detection*](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.

22,000 citations!


Source: L. Fei-Fei

Example




original image

Compute Gradients (DoG)



X-Derivative of Gaussian



Y-Derivative of Gaussian

Gradient magnitude and orientation

- Orientation is undefined at pixels with 0 gradient




Magnitude



Orientation

$\text{theta} = \text{numpy.arctan2}(\text{gy}, \text{gx})$


Non-maximum suppression for each orientation



Pixel q is a maximum if it is larger than p and r

p and r 's locations are determined by the gradient orientation


Their intensity is determined by linear interpolation.



Before Non-max Suppression



After Non-max Suppression




Hysteresis thresholding

- Threshold at low/high levels to get weak/strong edge pixels
- Do connected components, starting from strong edge pixels





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:
 - Thin multi-pixel wide “ridges” down to single pixel width
4. Thresholding and linking (hysteresis):
 - Define two thresholds: low and high
 - Use the high threshold to start edge curves and the low threshold to continue them

Does Canny always work?




The challenges of edge detection



- Texture
- Low-contrast boundaries

The challenges of edge detection



- Higher-level information

CSCI 497/597P: Computer Vision

Scott Wehrwein

Image Resampling & Interpolation



Reading

- Szeliski, Chapter 3.5


Announcements

Goals

- Know how to downsample an image naively
- Gain some intuition for why that's a bad idea
- Know how and why to build a Gaussian Pyramid
- Understand how to upsample an image naively
- Know how to use


Image

This image is too big to fit on the screen. How can we generate a half-sized version?



Source: S. Seitz

Image sub-sampling




1/4



1/8

Throw away every other row and column to create a 1/2 size image
- called *image sub-sampling*

Image sub-sampling



1/2



1/4 (2x zoom)





1/8 (4x zoom)


Why does this look so cruffy?

Source: S. Seitz


Image sub-sampling – another example



Even worse for synthetic images



Aliasing




- Occurs when your sampling rate is not high enough to capture the amount of detail in your image
- Can give you the wrong signal/image—an *alias*
- To do sampling right, need to understand the structure of your signal/image
- Enter Monsieur Fourier...
 - “But what is the Fourier Transform? A visual introduction.”
<https://www.youtube.com/watch?v=spUNpyF58BY&t=444s>
- To avoid aliasing:
 - sampling rate $\geq 2 * \text{max frequency in the image}$
 - said another way: \geq two samples per cycle
 - This minimum sampling rate is called the **Nyquist rate**

Wagon-wheel effect

Imagine a spoked wheel moving to the right (rotating clockwise).

Mark wheel with dot so we can see what's happening.

If camera shutter is only open for a fraction of a frame time (frame time = $1/30$ sec. for video, $1/24$ sec. for film):





Without dot, wheel appears to be rotating slowly backwards!
(counterclockwise)



<http://www.michaelbach.de/ot/mot-wagonWheel/index.html>

https://en.wikipedia.org/wiki/Wagon-wheel_effect

Nyquist limit – 2D example




Good sampling



Bad sampling

Gaussian pre-filtering



Gaussian 1/2



G 1/4



G 1/8

- Solution: filter the image, *then* subsample

Subsampling with Gaussian pre-filtering



Gaussian 1/2




G 1/4



G 1/8

- Solution: filter the image, *then* subsample

Compare with...



1/2





1/4 (2x zoom)




1/8 (4x zoom)

Gaussian pre-filtering




- Solution: filter the image, *then* subsample




...

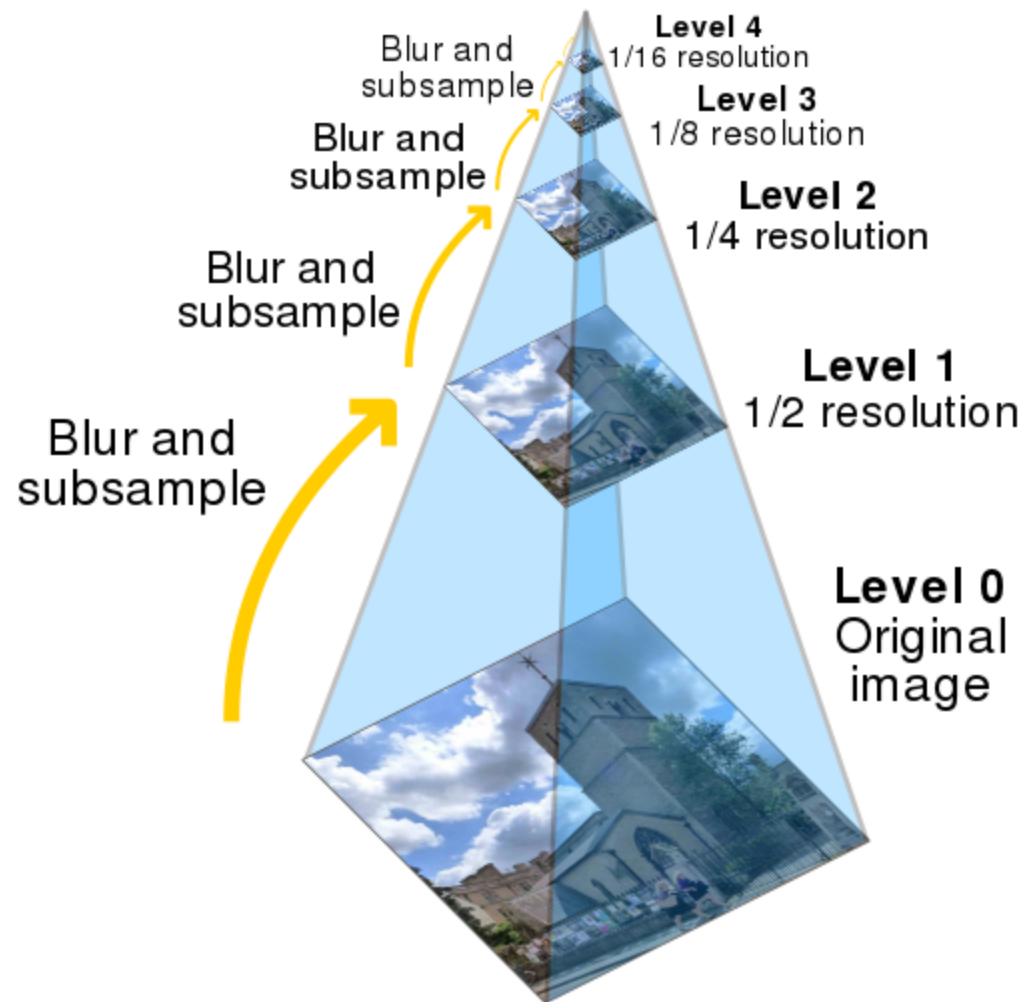


*Gaussian
pyramid*



...






Gaussian pyramids

[Burt and Adelson, 1983]

Idea: Represent $N \times N$ image as a “pyramid” of $1 \times 1, 2 \times 2, 4 \times 4, \dots, 2^k \times 2^k$ images (assuming $N=2^k$)




- In computer graphics, a *mip map* [Williams, 1983]
- A precursor to *wavelet transform*

Gaussian Pyramids have all sorts of applications in computer vision

Gaussian pyramids


[Burt and Adelson, 1983]

Idea: Represent $N \times N$ image as a “pyramid” of $1 \times 1, 2 \times 2, 4 \times 4, \dots, 2^k \times 2^k$ images (assuming $N = 2^k$)



- How much space does a Gaussian pyramid take compared to the original image?

Gaussian Pyramid




https://en.wikipedia.org/wiki/1/4_%2B_1/16_%2B_1/64_%2B_1/256_%2B_%E2%8B%AF

What are Gaussian Pyramids useful for?

- Operating at multiple **scales**



Operating at multiple scales



Operating at multiple scales

