

## Image processing

Brian Curless  
CSE 457  
Spring 2011

1

## Reading

Jain, Kasturi, Schunck, *Machine Vision*. McGraw-Hill, 1995. Sections 4.2-4.4, 4.5(intro), 4.5.5, 4.5.6, 5.1-5.4. [online handout]

2

## What is an image?

We can think of an **image** as a function,  $f$ , from  $\mathbb{R}^2$  to  $\mathbb{R}$ :


- $f(x, y)$  gives the intensity of a channel at position  $(x, y)$
- Realistically, we expect the image only to be defined over a rectangle, with a finite range:
  - $f: [a, b] \times [c, d] \rightarrow [0, 1]$

A color image is just three functions pasted together. We can write this as a "vector-valued" function:

$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$

3

## Images as functions



4


## What is a digital image?

In computer graphics, we usually operate on **digital (discrete)** images:

- ◆ **Sample** the space on a regular grid
- ◆ **Quantize** each sample (round to nearest integer)

If our samples are  $\Delta$  apart, we can write this as:

$$f[i, j] = \text{Quantize}\{f(i \Delta, j \Delta)\}$$



5

## Image processing

An **image processing** operation typically defines a new image  $g$  in terms of an existing image  $f$ .

The simplest operations are those that transform each pixel in isolation. These pixel-to-pixel operations can be written:

$$g(x, y) = t(f(x, y))$$

Examples: threshold, RGB  $\rightarrow$  grayscale


Note: a typical choice for mapping to grayscale is to apply the YIQ television matrix and keep the Y.

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.523 & 0.311 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

6

## Noise

Image processing is also useful for noise reduction and edge enhancement. We will focus on these applications for the remainder of the lecture...



$$v + \mathcal{N}(0, \sigma)$$

Common types of noise:

- ◆ **Salt and pepper noise:** contains random occurrences of black and white pixels
- ◆ **Impulse noise:** contains random occurrences of white pixels
- ◆ **Gaussian noise:** variations in intensity drawn from a Gaussian normal distribution

7

## Ideal noise reduction



8


## Ideal noise reduction



9

## Practical noise reduction

How can we "smooth" away noise in a single image?



Is there a more abstract way to represent this sort of operation? *Of course there is!*

10


## Discrete convolution

One of the most common methods for filtering an image is called **discrete convolution**. (We will just call this "convolution" from here on.)

In 1D, convolution is defined as:

$$\begin{aligned} g[n] &= f[n] * h[n] \\ &= \sum_n f[n'] h[n - n'] \\ &= \sum_n f[n'] \tilde{h}[n' - n] \end{aligned}$$

where  $\tilde{h}[n] = h[-n]$ .



11

## Discrete convolution

One can show that convolution has some convenient properties. Given functions  $a, b, c$ :

$$\begin{aligned} a * b &= b * a \\ (a * b) * c &= a * (b * c) \\ a * (b + c) &= a * b + a * c \end{aligned}$$

We'll make use of these properties later...

12

## Convolution in 2D

In two dimensions, convolution becomes:

$$\begin{aligned}
 g[n, m] &= f[n, m] * h[n, m] \\
 &= \sum_m \sum_n f[n', m'] h[n - n', m - m'] \\
 &= \sum_m \sum_n f[n', m'] \tilde{h}[n' - n, m' - m]
 \end{aligned}$$

where  $\tilde{h}[n, m] = h[-n, -m]$ .


13

## Convolution representation

Since  $f$  and  $h$  are defined over finite regions, we can write them out in two-dimensional arrays:

128	54	9	78	100
145	98	240	233	86
89	177	246	228	127
67	90	255	237	95
106	111	128	167	20
221	154	97	123	0

$f[n, m]$



**Note:** This is not matrix multiplication!

Q: What happens at the boundary of the image?

14

## Mean filters


How can we represent our noise-reducing averaging as a convolution filter (known as a **mean filter**)?

$$\frac{1}{N^2} \left[ \begin{array}{ccc} | & \dots & | \\ | & & | \\ | & & | \end{array} \right]$$

$N$

15

## Effect of mean filters



16

## Gaussian filters

Gaussian filters weigh pixels based on their distance from the center of the convolution filter. In particular:

$$h[n, m] = \frac{e^{-(n^2+m^2)/(2\sigma^2)}}{C}$$

This does a decent job of blurring noise while preserving features of the image.

What parameter controls the width of the Gaussian?  $\sigma$


What happens to the image as the Gaussian filter kernel gets wider? *blurrier*

What is the constant C? What should we set it to?

$$C = \sum e^{-(n^2+m^2)/2\sigma^2}$$

17

## Effect of Gaussian filters



18


## Median filters

A **median filter** operates over an  $m \times m$  region by selecting the median intensity in the region.

What advantage does a median filter have over a mean filter? *outlier removing, edge preserving*


Is a median filter a kind of convolution?

*No*




19

## Effect of median filters




20

## Comparison: Gaussian noise



21

## Comparison: salt and pepper noise



22

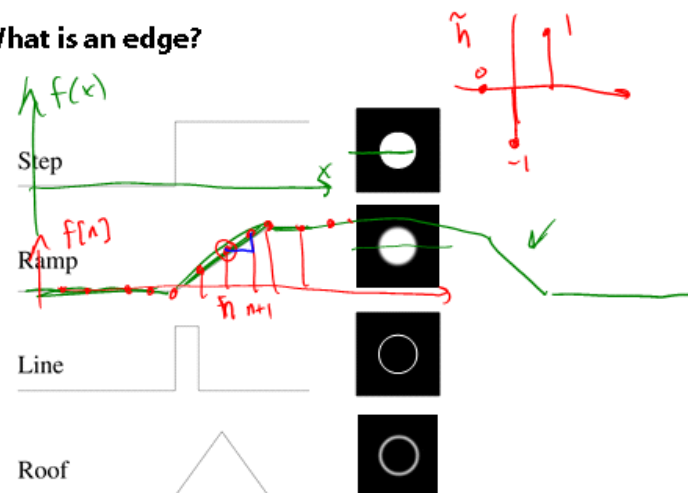
## Edge detection

One of the most important uses of image processing is **edge detection**:

- Really easy for humans
- Really difficult for computers
- Fundamental in computer vision
- Important in many graphics applications

23

## What is an edge?



Q: How might you detect an edge in 1D?

$$\left| \frac{df}{dx} \right| > \text{thresh}$$

$$\frac{df}{dx} \approx f[n+1] - f[n]$$

$$h = [1 \ -1 \ 0]$$

24

## Gradients

The **gradient** is the 2D equivalent of the derivative:

$$\nabla f(x,y) = \left( \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right) \quad \tan \theta = \frac{\partial f / \partial y}{\partial f / \partial x}$$

Properties of the gradient


- It's a vector
- Points in the direction of maximum increase of  $f$
- Magnitude is rate of increase

How can we approximate the gradient in a discrete image?

$$f_x[n,m] = f[n+1,m] - f[n,m]$$
$$f_y[n,m] = f[n,m+1] - f[n,m]$$
$$\tilde{h}_x = [0 \ -1 \ 1] \quad \tilde{h}_y = \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}$$

25

## Less than ideal edges



26

## Steps in edge detection

Edge detection algorithms typically proceed in three or four steps:

- **Filtering**: cut down on noise
- **Enhancement**: amplify the difference between edges and non-edges
- **Detection**: use a threshold operation
- **Localization** (optional): estimate geometry of edges as 1D contours that can pass between pixels

27

## Edge enhancement

A popular gradient filter is the **Sobel operator**:

$$\xi_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

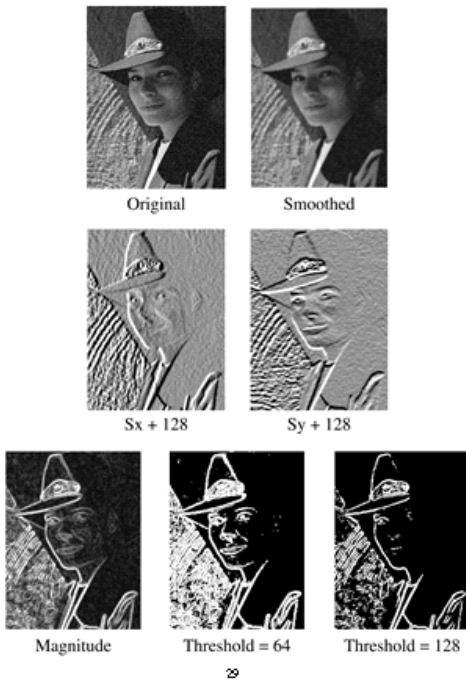
$$\xi_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

We can then compute the magnitude of the vector  $(\xi_x, \xi_y)$ .

Note that these operators are conveniently "pre-flipped" for convolution, so you can directly slide these across an image without flipping first.

28

## Results of Sobel edge detection



29



## Second derivative operators

$$h_x = [1 \ -1 \ 0]$$

$$\tilde{h}_x = [0 \ -1 \ 1]$$

$$h_{xx} = [1 \ -2 \ 1]$$

$$h_{xy} = [0 \ -1 \ 1]$$



The Sobel operator can produce thick edges. Ideally, we're looking for infinitely thin boundaries.

An alternative approach is to look for local extrema in the first derivative: places where the change in the gradient is highest.

Q: A peak in the first derivative corresponds to what in the second derivative? 0

Q: How might we write this as a convolution filter?

$$\frac{d}{dx} f \approx h_x * f$$

$$\frac{d}{dx} \frac{d}{dx} f \approx h_x * \frac{d}{dx} f$$

$$\approx h_x * (h_x * f)$$

$$\approx (h_x * h_x) * f$$

$$h_{xx}$$

30

## Localization with the Laplacian

An equivalent measure of the second derivative in 2D is the **Laplacian**:

$$\nabla^2 f(x,y) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

Using the same arguments we used to compute the gradient filters, we can derive a Laplacian filter to be:

$$\Delta = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

(The symbol  $\Delta$  is often used to refer to the *discrete* Laplacian filter.)

Zero crossings in a Laplacian filtered image can be used to localize edges.

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

$$\frac{\partial^2 f}{\partial x^2} \approx h_{xx} * f$$

$$[1 \ -2 \ 1]$$




$$\frac{\partial^2 f}{\partial y^2} \approx h_{yy} * f$$

$$\begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$$

$$\frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2} = [1 \ -2 \ 1] + \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix}$$

31

## Localization with the Laplacian



Laplacian (+128)

32



## Sharpening with the Laplacian

$$S = f - \lambda \Delta * f$$

$$\lambda \in [0, 1]$$

↓

$$\begin{bmatrix} 0 & -\lambda & 0 \\ -\lambda & 1+\lambda & -\lambda \\ 0 & -\lambda & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & -1/2 & 0 \\ -1/2 & 3 & -1/2 \\ 0 & -1/2 & 0 \end{bmatrix}$$



Original



Laplacian (+128)



Original + Laplacian



Original - Laplacian

Why does the sign make a difference?

How can you write the filter that makes the sharpened image?

33

$$S = f - \Delta * f$$

$$= 1 * f - \Delta * f$$

$$= (1 - \Delta) * f$$

$$= \left( [1] - \begin{bmatrix} 0 & 1 & 0 \\ 1 & 4 & 1 \\ 0 & 1 & 0 \end{bmatrix} \right) * f$$

$$= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 1 & 0 \\ 1 & 4 & 1 \\ 0 & 1 & 0 \end{bmatrix} * f$$

$$= \begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix} * f$$

## Summary

What you should take away from this lecture:

- The meanings of all the boldfaced terms.
- How noise reduction is done
- How discrete convolution filtering works
- The effect of mean, Gaussian, and median filters
- What an image gradient is and how it can be computed
- How edge detection is done
- What the Laplacian image is and how it is used in either edge detection or image sharpening

34