

# Visual Computing: Image Segmentation

Prof. Marc Pollefeys



But first let's finish last week's lecture

---


# Quantization

---


- Real valued function will get digital values – integer values
- Quantization is lossy!!
  - After quantization, the original signal cannot be reconstructed anymore
- This is in contrast to sampling, as a sampled but not quantized signal **can** be reconstructed.
- Simple quantization uses equally spaced levels with  $k$  intervals

$$k = 2^b$$


# Quantization



# Quantization



# Quantization



# Image Properties

---

- Image resolution
- Geometric resolution: How many pixels per area
- Radiometric resolution: How many bits per pixel



# Image resolution

---



1024x1024



512x1024



512x512

# Geometric resolution




144x144




72x72




36x36



18x18



9x9



4x4

# Radiometric resolution



256



128



64



32



16



8



4



2

# Aliasing and SNR

---

- What is the disadvantage of low sampling resolution?
- What is the disadvantage of high sampling resolution?
- Lossless vs. Lossy
  - Name some formats?

# Unassessed Assignment

---

Use python to change the geometric and radiometric quantization resolution in one of your images. For each level of sampling and quantization, plot the image function, as in slides 71 & 72, and compare the approximations to the true intensity function that you get at each level.

# Usual quantization intervals

---

- Grayscale image
  - 8 bit =  $2^8 = 256$  grayvalues
- Color image RGB (3 channels)
  - 8 bit/channel =  $2^{24} = 16.7\text{M}$  colors
- 12bit or 16bit from some sensors
- Nonlinear, for example log-scale





Photo: Paulo Barcellos Jr.

# Image Noise



- A common model is *additive Gaussian* noise:

$$I(x, y) = f(x, y) + c$$

where  $c \sim N(0, \sigma^2)$ . So that  $p(c) = (2\pi\sigma^2)^{-1} e^{-c^2/2\sigma^2}$


- Poisson noise:

(shot noise) 
$$p(k) = \frac{\lambda^k e^{-\lambda}}{k!}$$



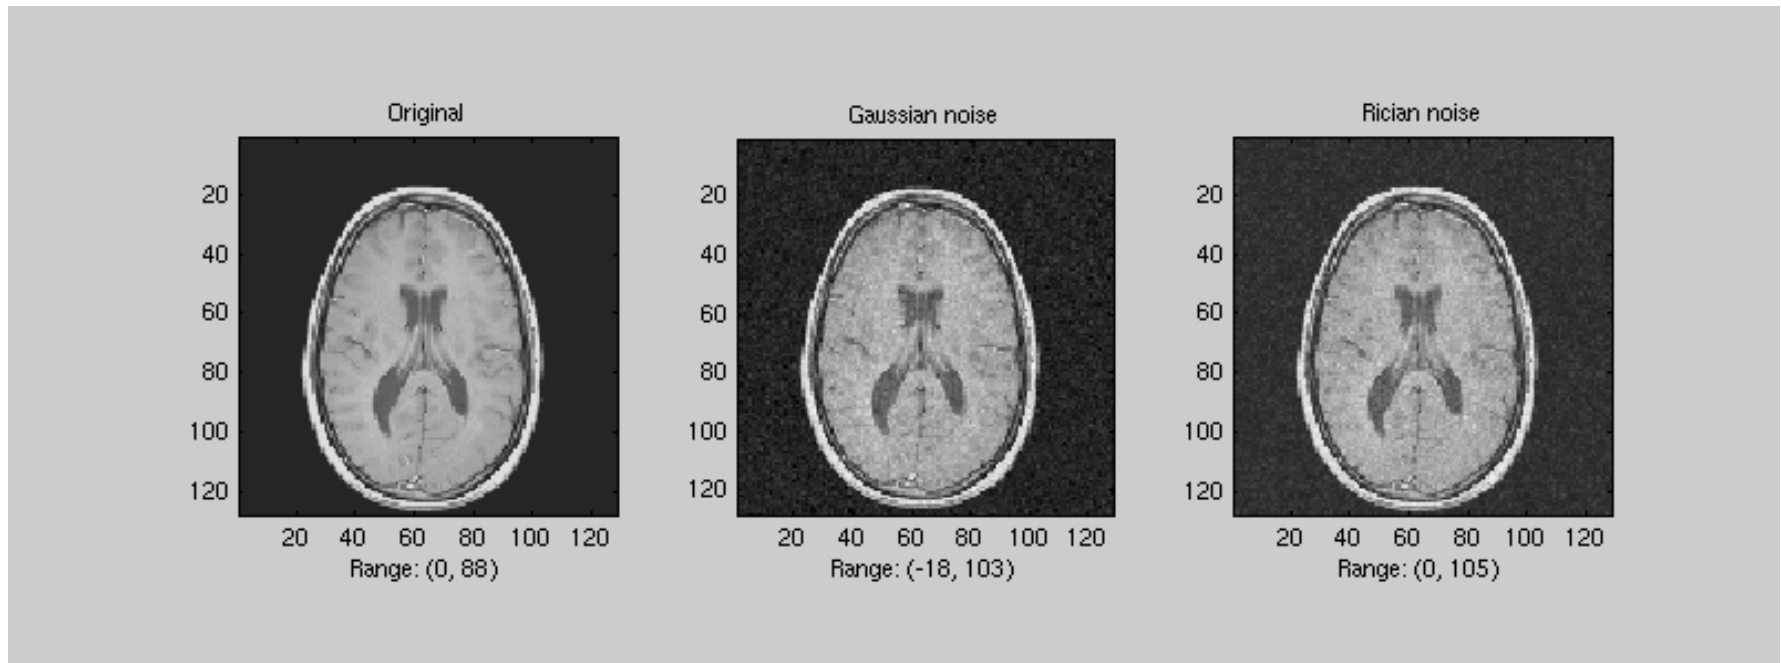


# Image Noise



- Rician noise:  
(appears in MRI)

$$p(I) = \frac{I}{\sigma^2} \exp\left(-\frac{(I^2 + f^2)}{2\sigma^2}\right) I_0\left(\frac{If}{\sigma^2}\right)$$




# Image Noise

- Multiplicative noise:

$$I = f + fc$$

- Quantization errors
- Impulse “salt-and-pepper” noise
- The *signal to noise ratio (SNR)*  $s$  is an index of image quality

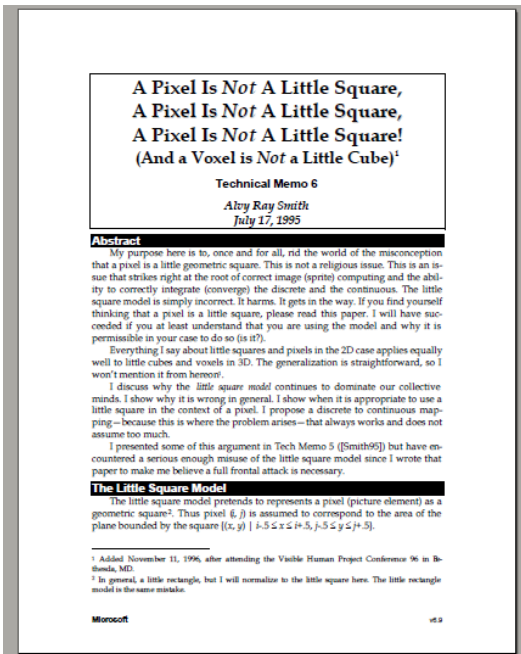


$$s = \frac{F}{\sigma}, \text{ where } F = \frac{1}{XY} \sum_{x=1}^X \sum_{y=1}^Y f(x, y)$$

Often used instead: *Peak Signal to Noise Ratio (PSNR)*  $s_{peak} = \frac{F_{max}}{\sigma}$

# Recap

## Geometric resolution




144x144



72x72



36x36



18x18



9x9



4x4

## Radiometric resolution



256



128




64



32



16



8



4



2

# Image Noise

---

- A common model is *additive Gaussian* noise:

$$I(x, y) = f(x, y) + c$$



where  $c \sim N(0, \sigma^2)$ . So that  $p(c) = (2\pi\sigma^2)^{-1} e^{-c^2/2\sigma^2}$

- The *signal to noise ratio (SNR)*  $s$  is an index of image quality

$$s = \frac{F}{\sigma}, \text{ where } F = \frac{1}{XY} \sum_{x=1}^X \sum_{y=1}^Y f(x, y)$$

Often used instead: *Peak Signal to Noise Ratio (PSNR)*  $s_{peak} = \frac{F_{\max}}{\sigma}$

## The Human Eye



Helmholtz's Schematic Eye



Visible light

Low energy

High energy

Frequency ( $s^{-1}$ )


$3 \times 10^{10}$

$3 \times 10^{12}$

$3 \times 10^{14}$

$3 \times 10^{16}$

$3 \times 10^{18}$




$10^{-1}$   $10^{-2}$   $10^{-3}$   $10^{-4}$   $10^{-5}$   $10^{-6}$   $10^{-7}$   $10^{-8}$   $10^{-9}$   $10^{-10}$   $10^{-11}$

Wavelength (m)




# The distribution of rods and cones across the retina

## color sensitivity of cones (normalized)




Bowmaker J.K. & Dartnall H.J.A. (1980). "Visual pigments of rods and cones in a human retina". *J. Physiol.* 298: 501-511.



Reprinted from *Foundations of Vision*, by B. Wandell, Sinauer Associates, Inc., (1995). © 1995 Sinauer Associates, Inc.

## Cones in the fovea




Reprinted from *Foundations of Vision*, by B. Wandell, Sinauer Associates, Inc., (1995). © 1995 Sinauer Associates, Inc.

# More eyes in nature...

## Chambered eyes



Nautilus




Octopus



Red-tailed hawk




Scallop





## Compound eyes



Sea fan



Dragonfly




Krill eye



Lobster

More info:





# Colour Images

R



+

G



+

B



=



# Color cameras

---

We consider 3 concepts:

1. Prism (with 3 sensors)
2. Filter mosaic
3. Filter wheel



... and X3

# Prism color camera


Separate light in 3 beams using dichroic prism

Requires 3 sensors & precise alignment


Good color separation




# Filter mosaic




Coat filter directly on sensor




**Bayer filter**



Demosaicing (obtain full colour & full resolution image)



More colors:



# Filter wheel

---

Rotate multiple filters in front of lens

Allows more than 3 colour bands



Only suitable for static scenes




# Prism vs. mosaic vs. wheel

<u>approach</u>	<u>Prism</u>	<u>Mosaic</u>	<u>Wheel</u>
# sensors	3	1	1
Separation	High	Average	Good
Cost	High	Low	Average
Framerate	High	High	Low
Artefacts	Low	Aliasing	Motion
Bands	3	3	3 or more


High-end cameras      Low-end cameras      Scientific applications

# color CMOS sensor

## Foveon's X3



better image quality





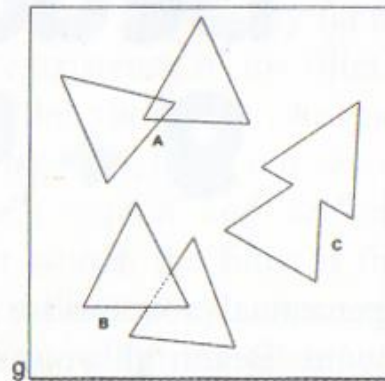







Figure 2.9(a-g)

## Gestalt Phenomena

- Figure-ground
- Proximity
- Similarity
- Continuation
- Closure
- Common fate
- Symmetry



# Young Lady or An Old Hag?

---



# What is Image Segmentation?

---

Segmentation is the ultimate classification problem.  
Once solved, Computer Vision is solved.


# What is Image Segmentation?

---

- It partitions an image into regions of interest
- The first stage in many automatic image analysis systems
- A *complete segmentation* of an image  $I$  is a finite set of regions  $R_1, \dots, R_N$ , such that

$$I = \bigcup_{i=1}^N R_i \text{ and } R_i \cap R_j = \emptyset \forall i \neq j.$$

# How should I segment this?

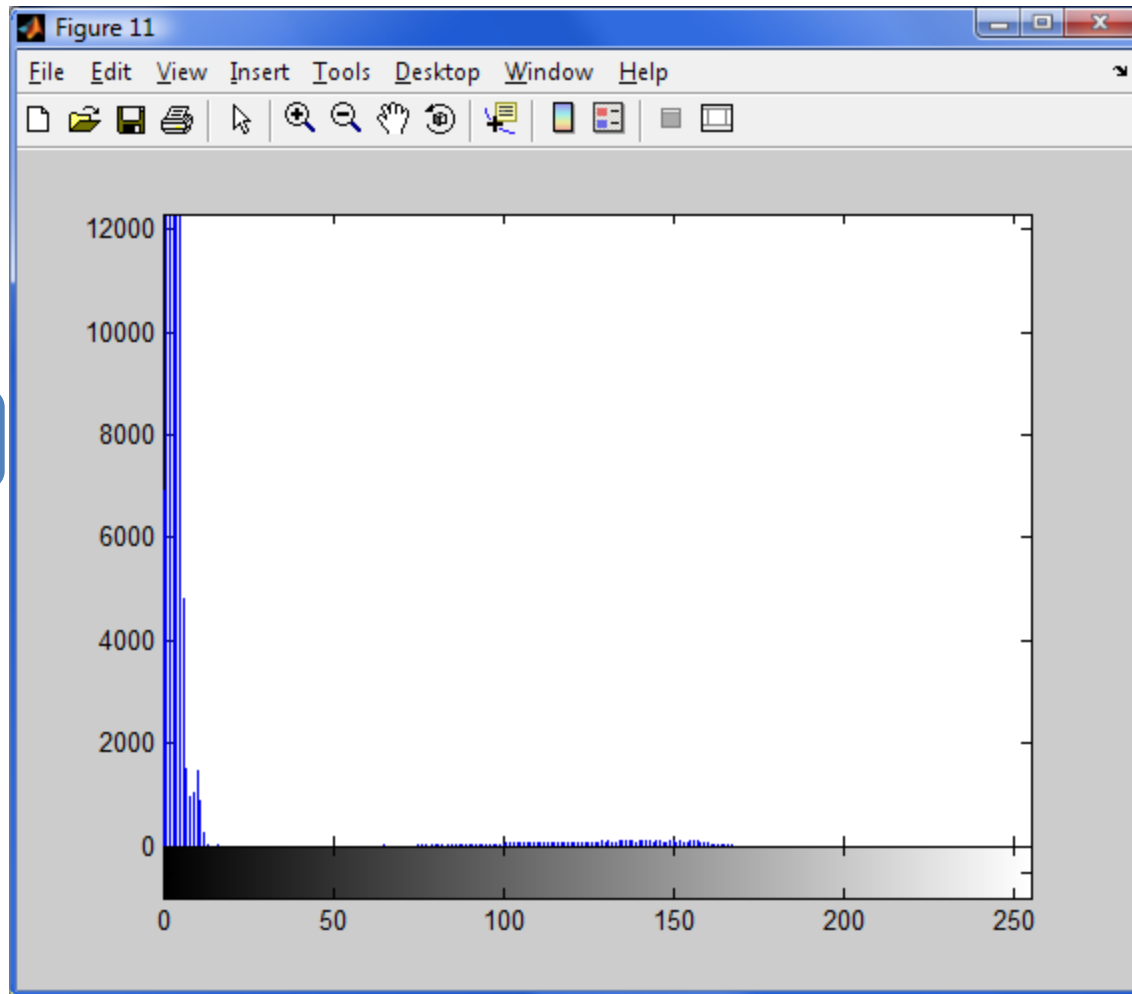


# Exclude dark pixels?

---

```
img = cv2.imread('BlobsIP.png')
cv2.imshow('BlobsIP', img)
cv2.waitKey(0)
img.shape    --> [ 244    767    3 ]
hist = np.histogram(img, bins=256)
cv2.imshow('Histogram', hist)
cv2.waitKey(0)
cv2.imshow('Mask', img[:, :, 1] > 20)
cv2.waitKey(0)
```

# Histogram



# Pixels

Grayscales [0,255]

# How should I segment this?




# Segmentation Quality

---

- The quality of a segmentation depends on what you want to do with it.
  - Segmentation algorithms must be chosen and evaluated with an application in mind.
-



# Segmentation example





# Berkeley Segmentation Database and Benchmark



<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/resources.html>



# Berkeley Segmentation Database and Benchmark



Martin et al. PAMI 2004

# Thresholding

---

- Thresholding is a simple segmentation process.
- Thresholding produces a binary image  $B$ .
- It labels each pixel **in** or **out** of the region of interest by comparison of the greylevel with a threshold  $T$ :


$$B(x, y) = \begin{cases} 1 & \text{if } I(x, y) \geq T \\ 0 & \text{if } I(x, y) < T. \end{cases}$$




# Thresholding example




T=50



$T=100$




T=150





T=200




# How do we choose T?

---

- Trial and error
- Compare results with ground truth
- Automatic methods\*

\*= We'll discuss ROC curves later

# Wouldn't it be nice...



# Planning to Segment?

---



Shraybman, HW1, 2003

# Chromakeying: Control Lighting!

---



# Chromakeying

---

- “Plain” distance measure

$$\mathbf{I}_\alpha = |\mathbf{I} - \mathbf{g}| > T$$

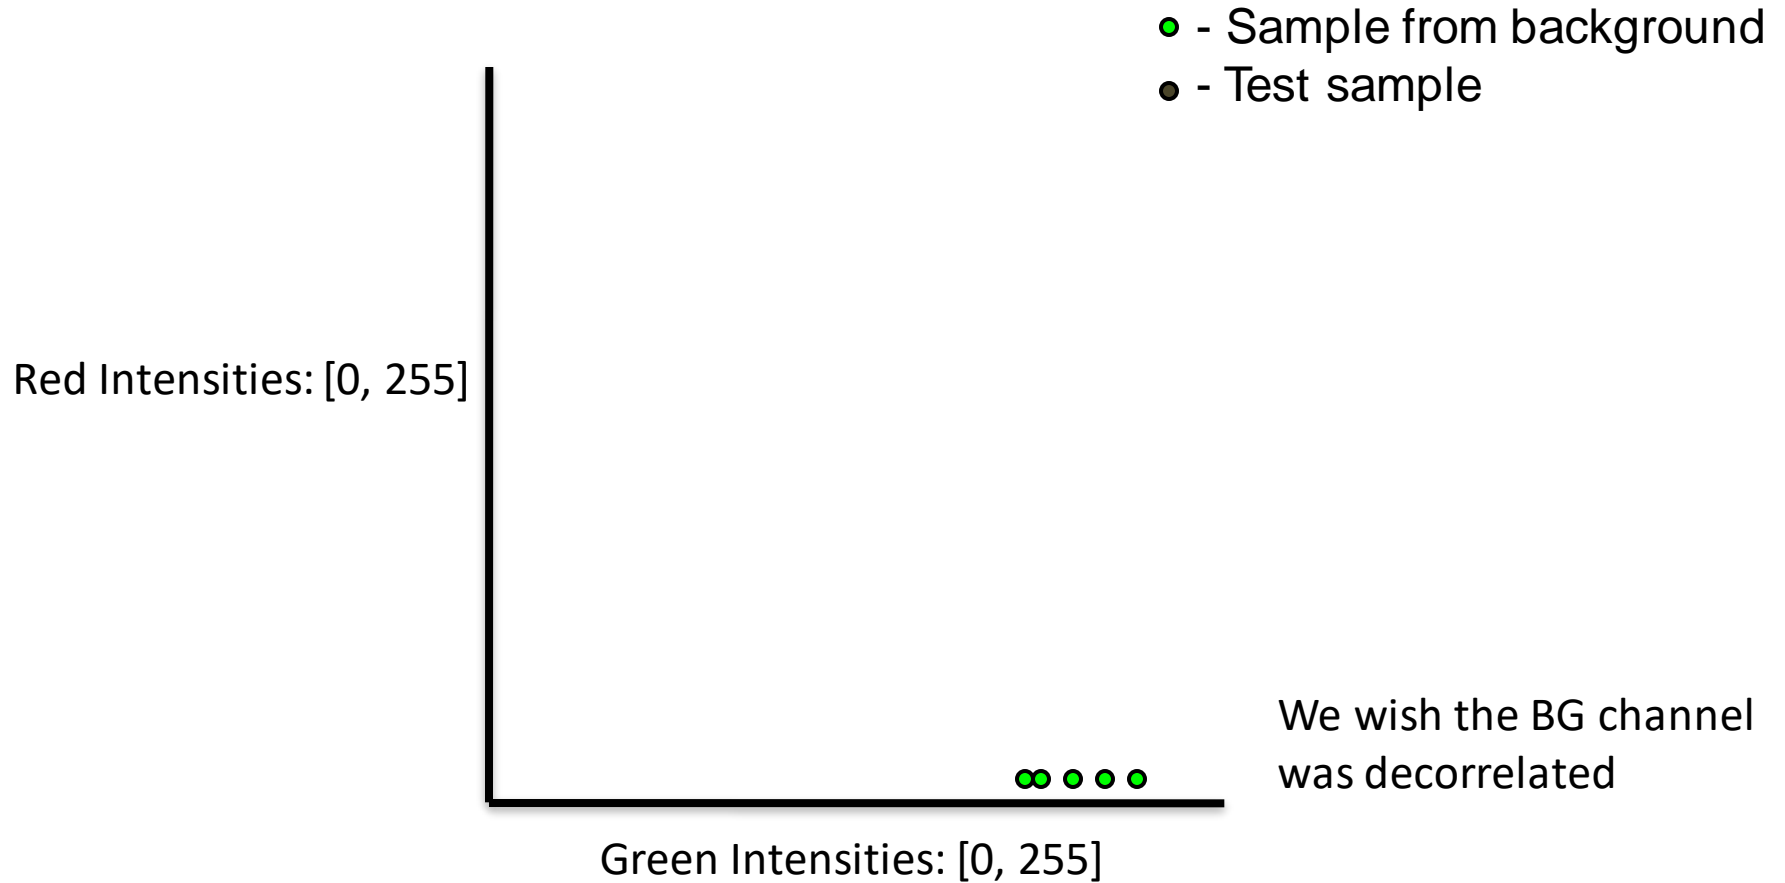
$$T = \sim 20$$

$$\mathbf{g} = ( 0 \ 255 \ 0 ) \quad \text{(for example)}$$

- Problems:
  - Variation is NOT same in all 3 channels
  - Hard alpha mask:  $\mathbf{I}_{\text{comp}} = \mathbf{I}_\alpha \mathbf{I}_a + (1 - \mathbf{I}_\alpha) \mathbf{I}_b$

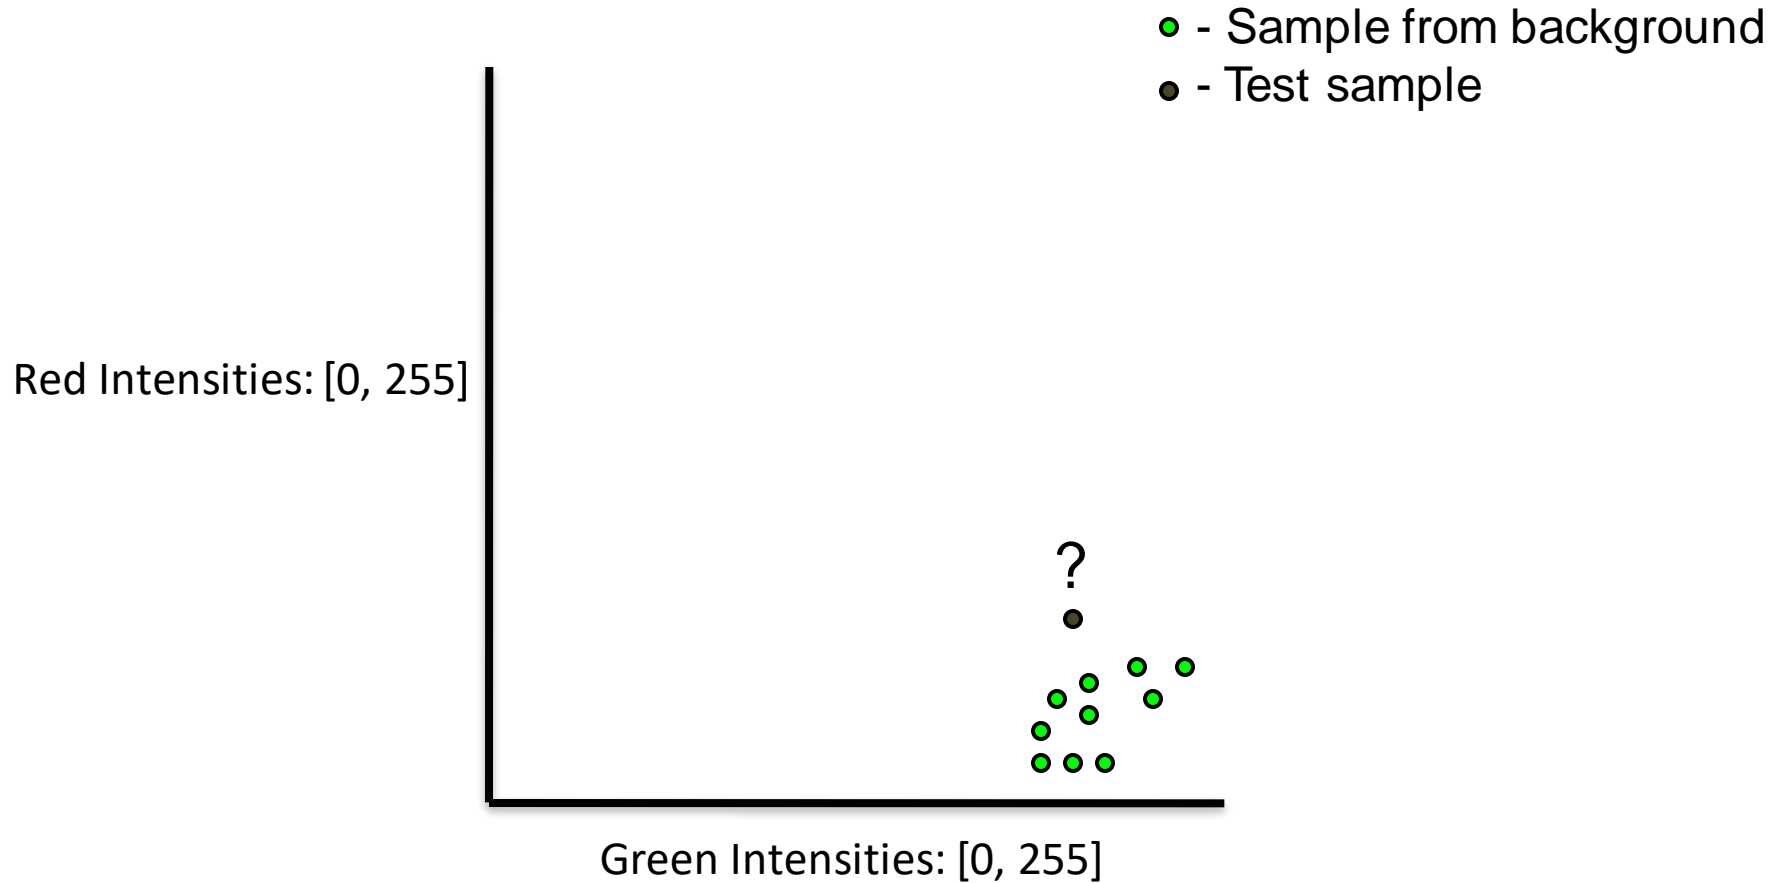
# Background Color Variation

---



# Background Color Variation


---





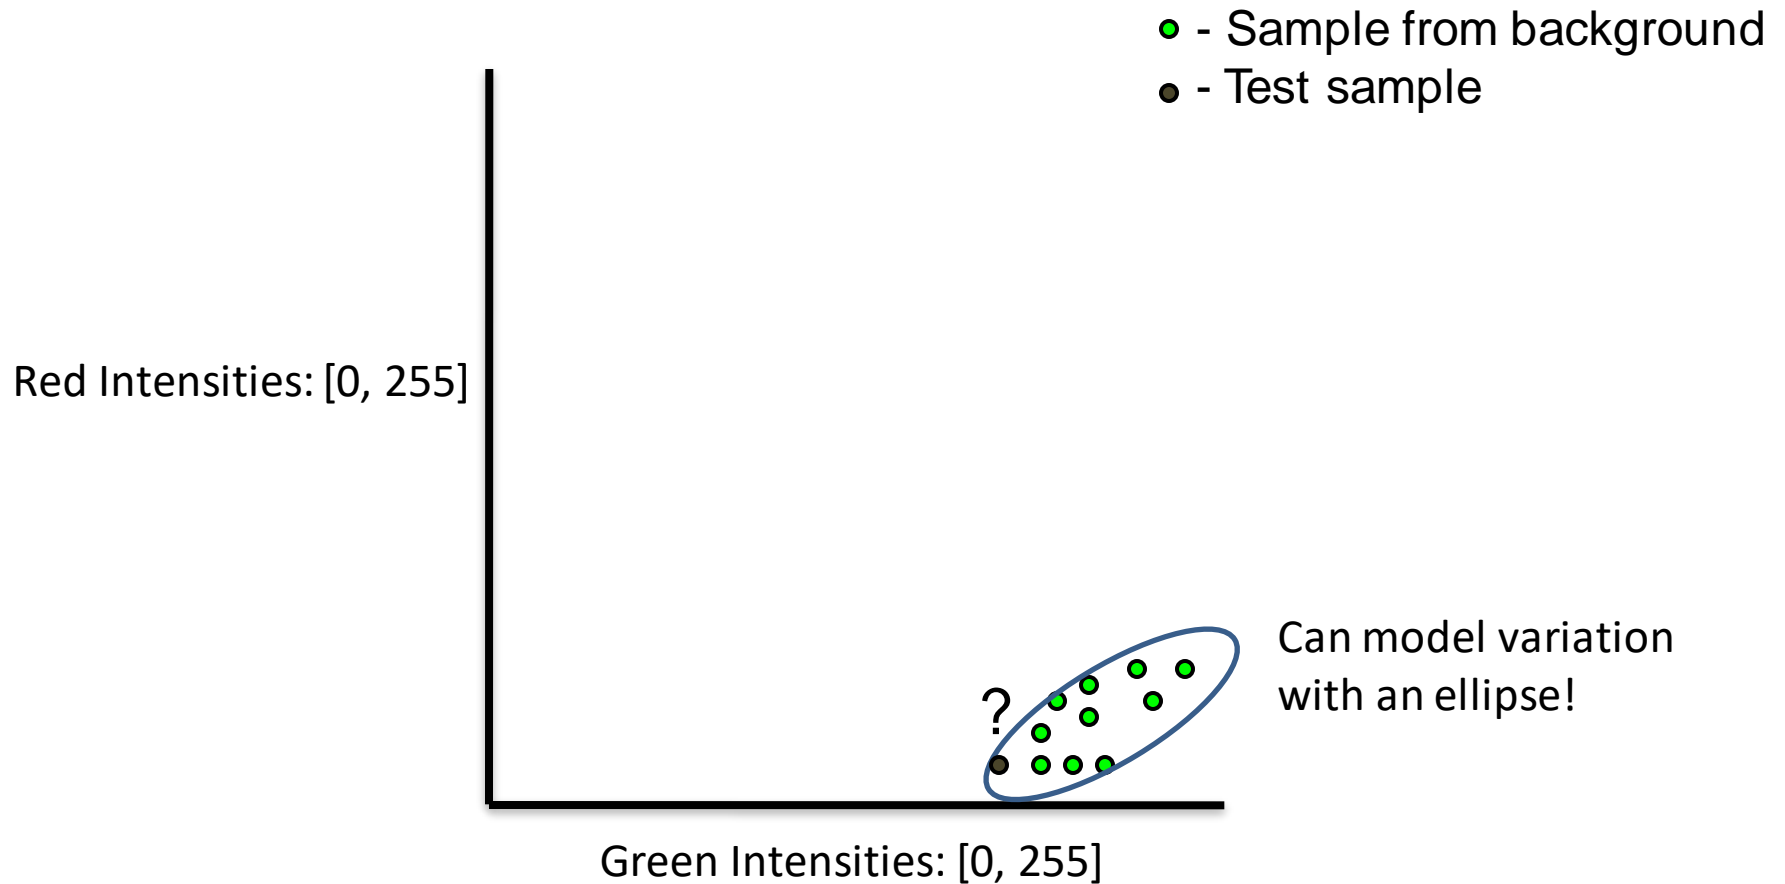
# Background Color Variation

---




# Background Color Variation

---



# Use Gaussian to Explain Most Data

---



# Green Variation in Reality



3D RGB  
plot of just  
BG

# Green Variation in Reality




3D RGB  
plot of  
geometry-  
normalized  
BG

Intensity  $I = R + G + B$

Normalized color:  $(r, g, b) = (R/I, G/I, B/I)$

# Mixed Pixels




$$I_{\text{comp}} = I_{\alpha} I_a + (1 - I_{\alpha}) I_b$$

# Segmentation Performance

---

- To use automatic analysis, one needs to know the true classification of each test
- We need to do the segmentation by hand on some example images...

Ground truth





# ROC Analysis

(ROC = Receiver operating characteristic)

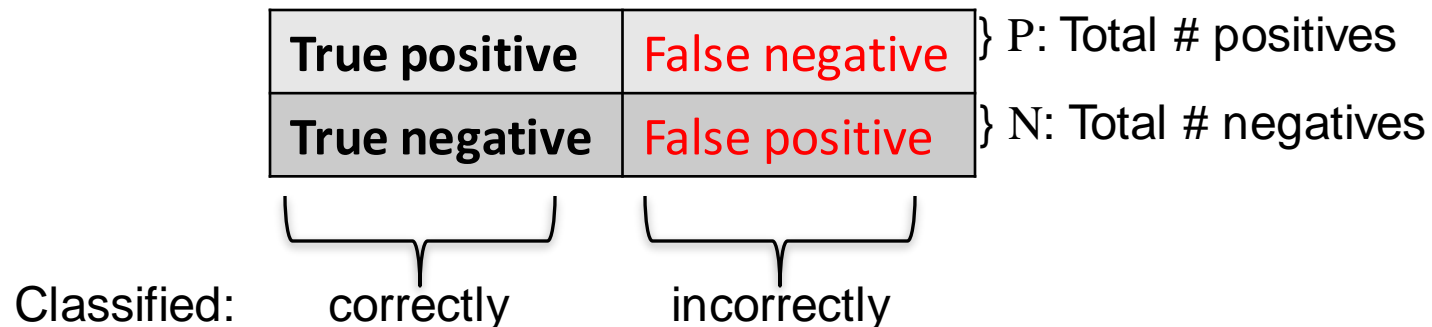
---

- An ROC curve characterizes the performance of a binary classifier.
- A binary classifier distinguishes between two different types of thing. E.g.:
  - Healthy/afflicted patients – cancer screening
  - Pregnancy tests
  - Object detection
  - Foreground/background image pixels


# Classification Error

---

- Binary classifiers make errors
- Two types of input to a binary classifier:
  - Positives
  - Negatives
- Four possible outcomes in any test:



# Classification outcomes




# The ROC Curve

---

- Characterizes the error trade-off in binary classification tasks
- It plots the TP fraction against FP fraction
- TP fraction (*sensitivity*) is  $\frac{\text{True positive count}}{P}$
- FP fraction (*1-specificity*) is  $\frac{\text{False positive count}}{N}$

# The ROC Curve




# Properties of ROC curves

---

- An ROC curve always passes through  $(0,0)$  and  $(1,1)$
- What is the ROC curve of a perfect system?
- What if the ROC curve is a straight line from  $(0,0)$  to  $(1,1)$ ?


# Threshold Sweep?

pdf



# “MAP (Maximum A Posteriori) detector”

pdf




If different outcomes are associated with different costs:  
more general “Bayes minimum risk detector”



# Operating points

- Choose an *operating point* by assigning relative costs and values to each outcome:

- $V_{TN}$  – value of true negative
- $V_{TP}$  – value of true positive
- $C_{FN}$  – cost of false negative
- $C_{FP}$  – cost of false positive




- Choose the point on the ROC curve with **gradient**

$$\beta = \frac{N V_{TN} + C_{FP}}{P V_{TP} + C_{FN}}$$


- For simplicity, we often set  $V_{TN} = V_{TP} = 0$ .

# Classification outcomes





# ROC curve



---



# Greylevel Histograms




# Positives and Negatives




# Limitations of Thresholding

---

- Why can we segment images much better by eye than through thresholding processes?
- We might improve results by considering image **context**: Surface Coherence




[Gradient.illusion.arp.jpg](#)




by [Adrian Pingstone](#), based on the [original](#) created by Edward H. Adelson






by [Adrian Pingstone](#), based on the [original](#) created by Edward H. Adelson


# Chapter 3 in Machine Vision by Jain et al.




(a)




(b)




(c)





(d)



(e)




(f)




(d)




(e)




(f)



(g)



(h)



(i)



(j)

# Note on Performance Assessment

---


- In real-life, we use two or even three separate sets of test data:
  1. A ***training set***, for tuning the algorithm
  2. A ***validation*** set for tuning the performance score
  3. An unseen ***test set*** to get a final performance score on the tuned algorithm

# Pixel connectivity

---

- We need to define which pixels are neighbors.
- Are the dark pixels in this image connected?


**Warning:**  
Pixels are samples,  
not squares.




# Pixel connectivity

---


- We need to define which pixels are neighbors.
- Are the dark pixels in this array connected?



# Pixel Neighborhoods



4-neighborhood



8-neighborhood

# Pixel paths

---

- A 4-connected path between pixels  $p_1$  and  $p_n$  is a set of pixels  $\{p_1, p_2, \dots, p_n\}$  such that  $p_i$  is a 4-neighbor of  $p_{i+1}$ ,  $i=1, \dots, n-1$ .
- In an 8-connected path,  $p_i$  is an 8-neighbor of  $p_{i+1}$ .



# Connected regions


---

- A region is 4-connected if it contains a 4-connected path between any two of its pixels.
- A region is 8-connected if it contains an 8-connected path between any two of its pixels.

# Connected regions


---

- Now what can we say about the dark pixels in this array?
- What about the light pixels?



# Connected components labelling


- Labels each connected component of a binary image with a separate number.



1	1	1	1	1	1	1	1	1
1	1	2	2	1	3	1	1	1
1	1	1	2	2	1	1	1	1
1	1	1	1	1	1	1	1	1
1	4	1	1	5	5	1	1	1
1	1	1	1	5	5	5	1	1
6	6	1	1	1	1	5	5	1
7	6	1	1	8	8	1	1	1
7	6	1	1	8	8	1	1	1

# Foreground labelling

- Only extract the connected components of the foreground




1	1	1	1	1	1	1	1	1
1	1	0	0	1	0	1	1	1
1	1	1	0	0	1	1	1	1
1	1	1	1	1	1	1	1	1
1	0	1	1	0	0	1	1	1
1	1	1	1	0	0	0	1	1
0	0	1	1	1	1	0	0	1
2	0	1	1	0	0	1	1	1
2	0	1	1	0	0	1	1	1

# Goose detector



# Goose detector



# Region Growing

---

- Start from a seed point or region.
- Add neighboring pixels that satisfy the criteria defining a region.
- Repeat until we can include no more pixels.

# Region Growing

```
def regionGrow(I, seed):  
    X, Y = I.shape  
    visited = np.zeros((X,Y))  
    visited[seed] = 1  
    boundary = []  
    boundary.append(seed)  
    while len(boundary) > 0:  
        nextPoint = boundary.pop()  
        if include(nextPoint, seed):  
            visited[nextPoint] = 2  
            for (x, y) in neighbors(nextPoint):  
                if visited[x,y] == 0:  
                    boundary.append((x, y))  
                    visited[x,y] = 1
```




# Region Growing example

---


- Pick a single seed pixel
- Inclusion test is up to you:

```
def include(p, seed):  
    test = ??  
    return test
```



Seed pixel

T=150



# Variations

---

- Seed selection
- Inclusion criteria
- Boundary constraints and snakes

# Seed selection

---

- Point and click seed point
- Seed region
  - By hand
  - Automatically, e.g., from a conservative thresholding.
- Multiple seeds
  - Automatically labels the regions

# Inclusion criteria

---

- Greylevel thresholding
- Greylevel distribution model
  - Use mean  $\mu$  and standard deviation  $\sigma$  in seed region:
  - Include if  $(I(x, y) - \mu)^2 < (n\sigma)^2$ . Eg:  $n = 3$ .
  - Can update the mean and standard deviation after every iteration.
- Color or texture information

# Snakes

---

- A snake is an *active contour*
- It is a polygon, i.e., an ordered set of points joined up by lines
- Each point on the contour moves away from the seed while its image neighborhood satisfies an inclusion criterion
- Often the contour has smoothness constraints

# Snakes

---

- The algorithm iteratively minimizes an energy function:
- $E = E_{\text{tension}} + E_{\text{stiffness}} + E_{\text{image}}$
- See Kass, Witkin, Terzopoulos, IJCV 1988



# Example





# Interim Summary

---

- Segmentation is hard
- But it is easier if you define the task carefully
  - Is the segmentation task binary or continuous?
  - What are the regions of interest?
  - How accurately must the algorithm locate the region boundaries?
- Research problems remain!

**Thursday:  
More segmentation**