Histograms

- $h(r_k) = n_k$
  - Histogram: number of times intensity level rk appears in the image
- $p(r_k) = n_k/NM$
  - normalized histogram
  - also a probability of occurence

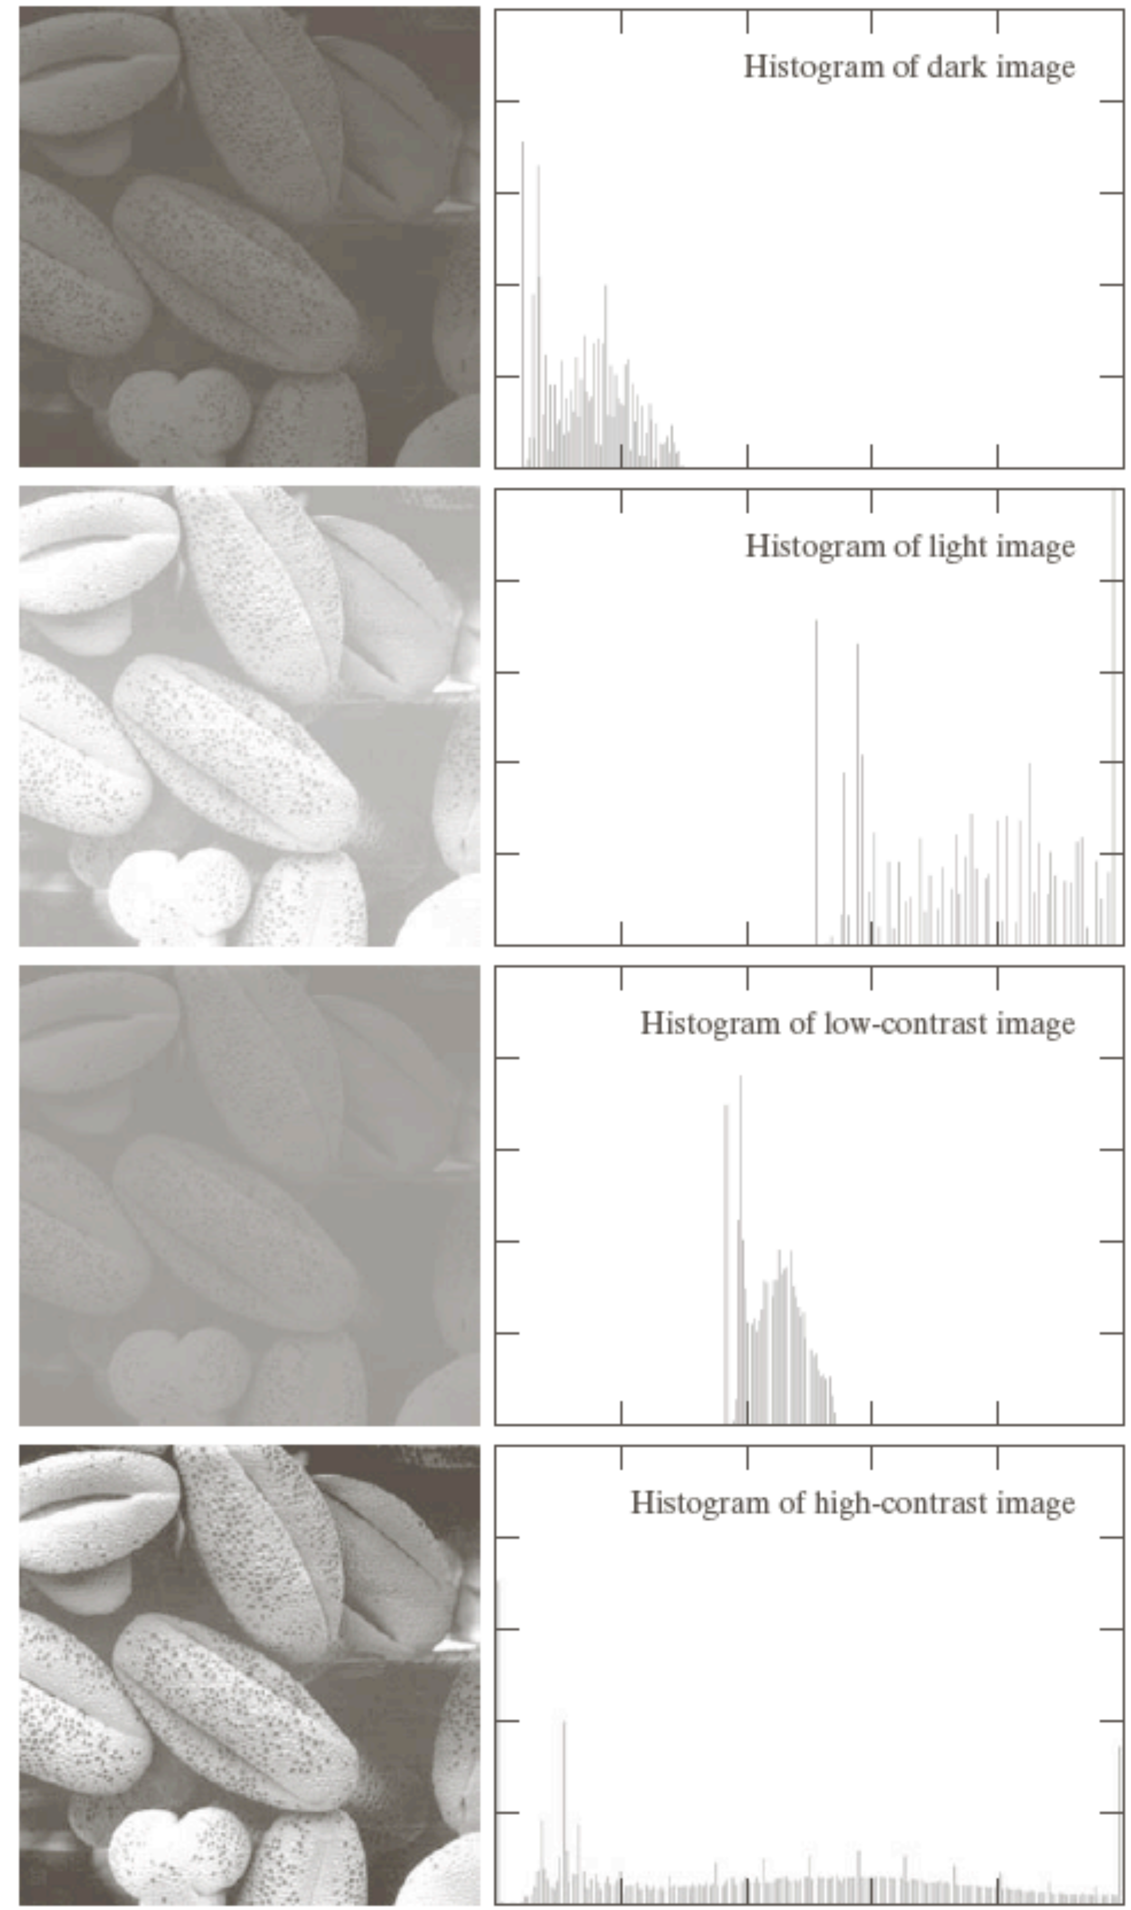

Histogram of Image Intensities

- Create bins of intensities and count number of pixels at each level
  - Normalized (divide by total # pixels)

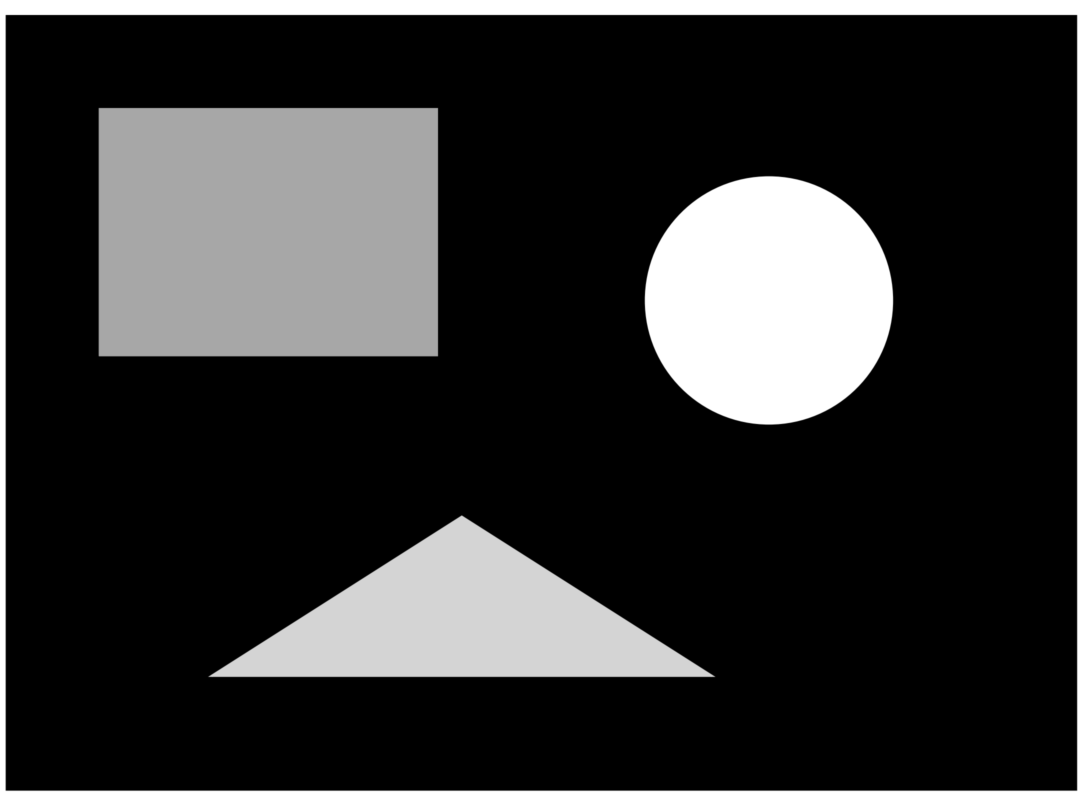

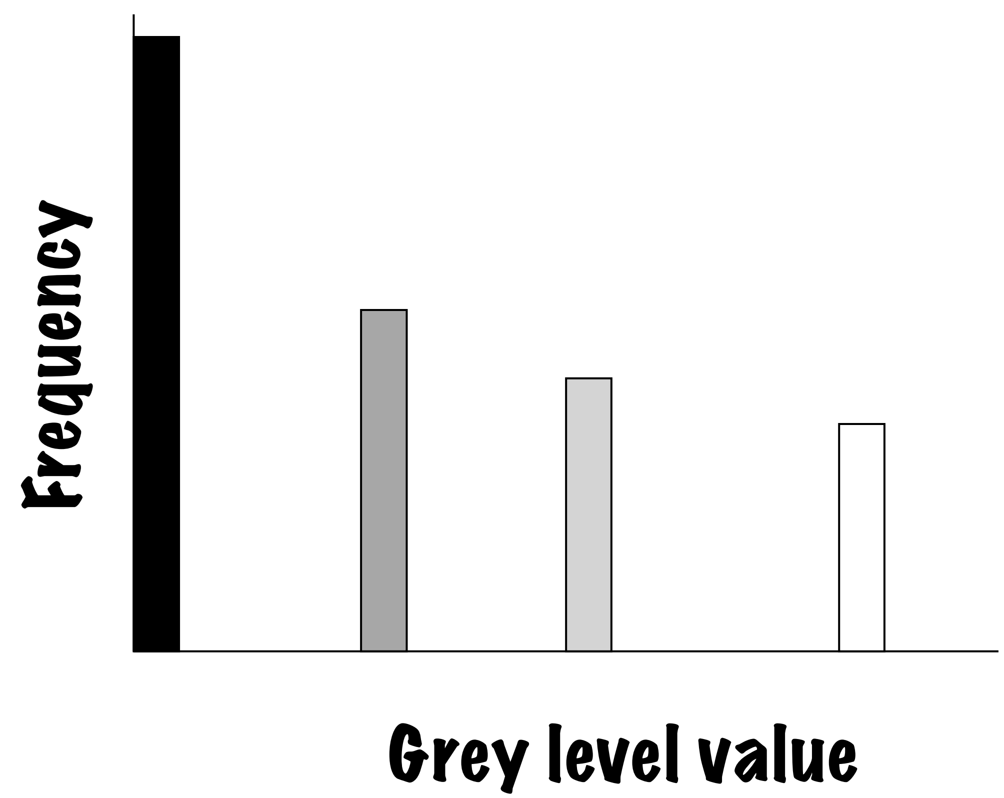

Histograms and Noise

- What happens to the histogram if we add noise?
  - $g(x, y) = f(x, y) + n(x, y)$

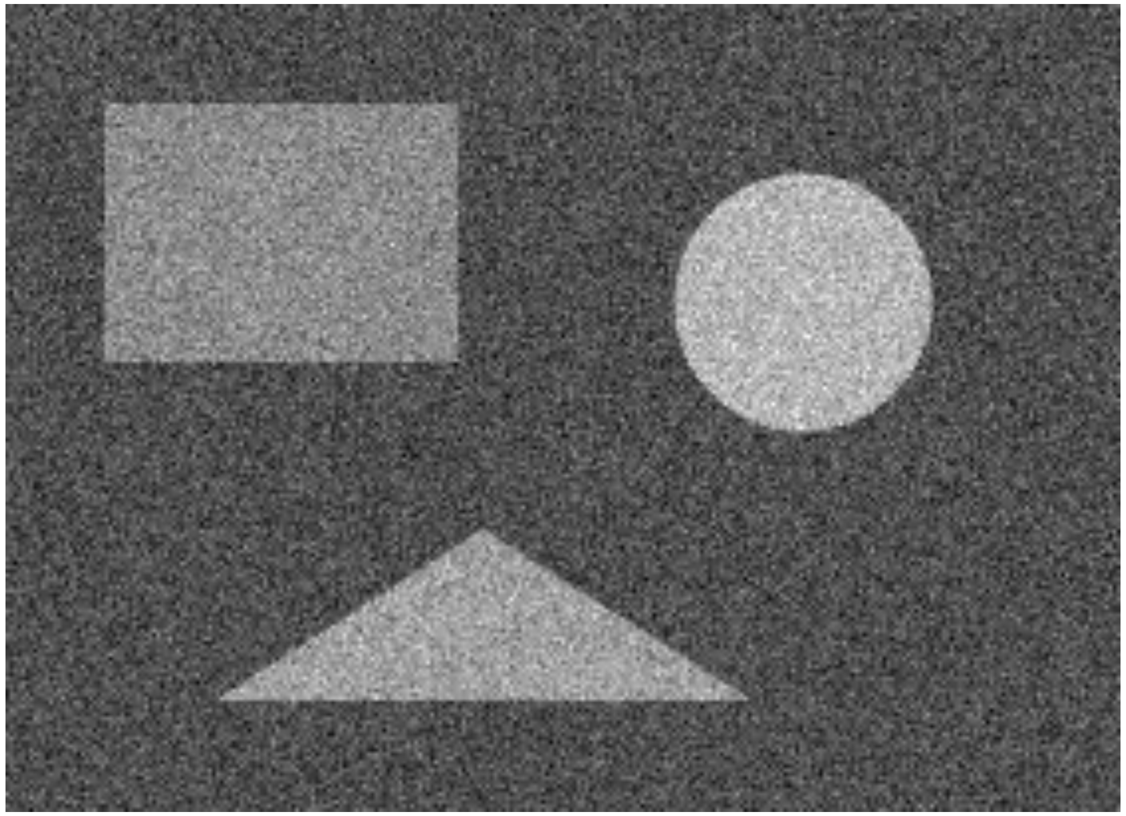

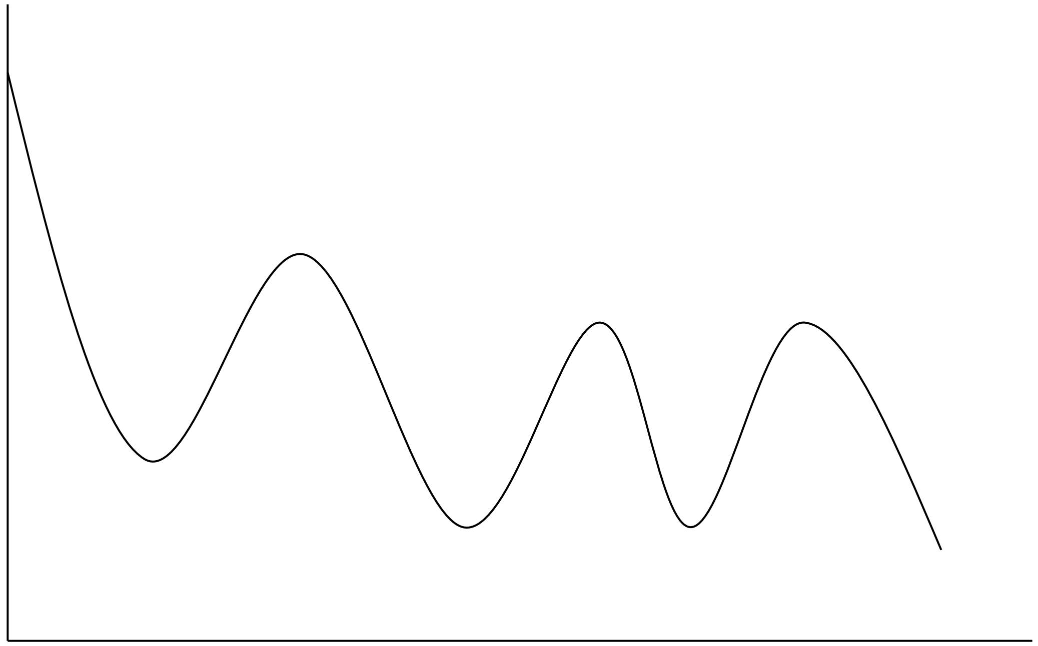

Noise reduction with pixelwise addition

- Many noisy images of the same scene
- Averaging helps remove noise
  - Why? Under what conditions?

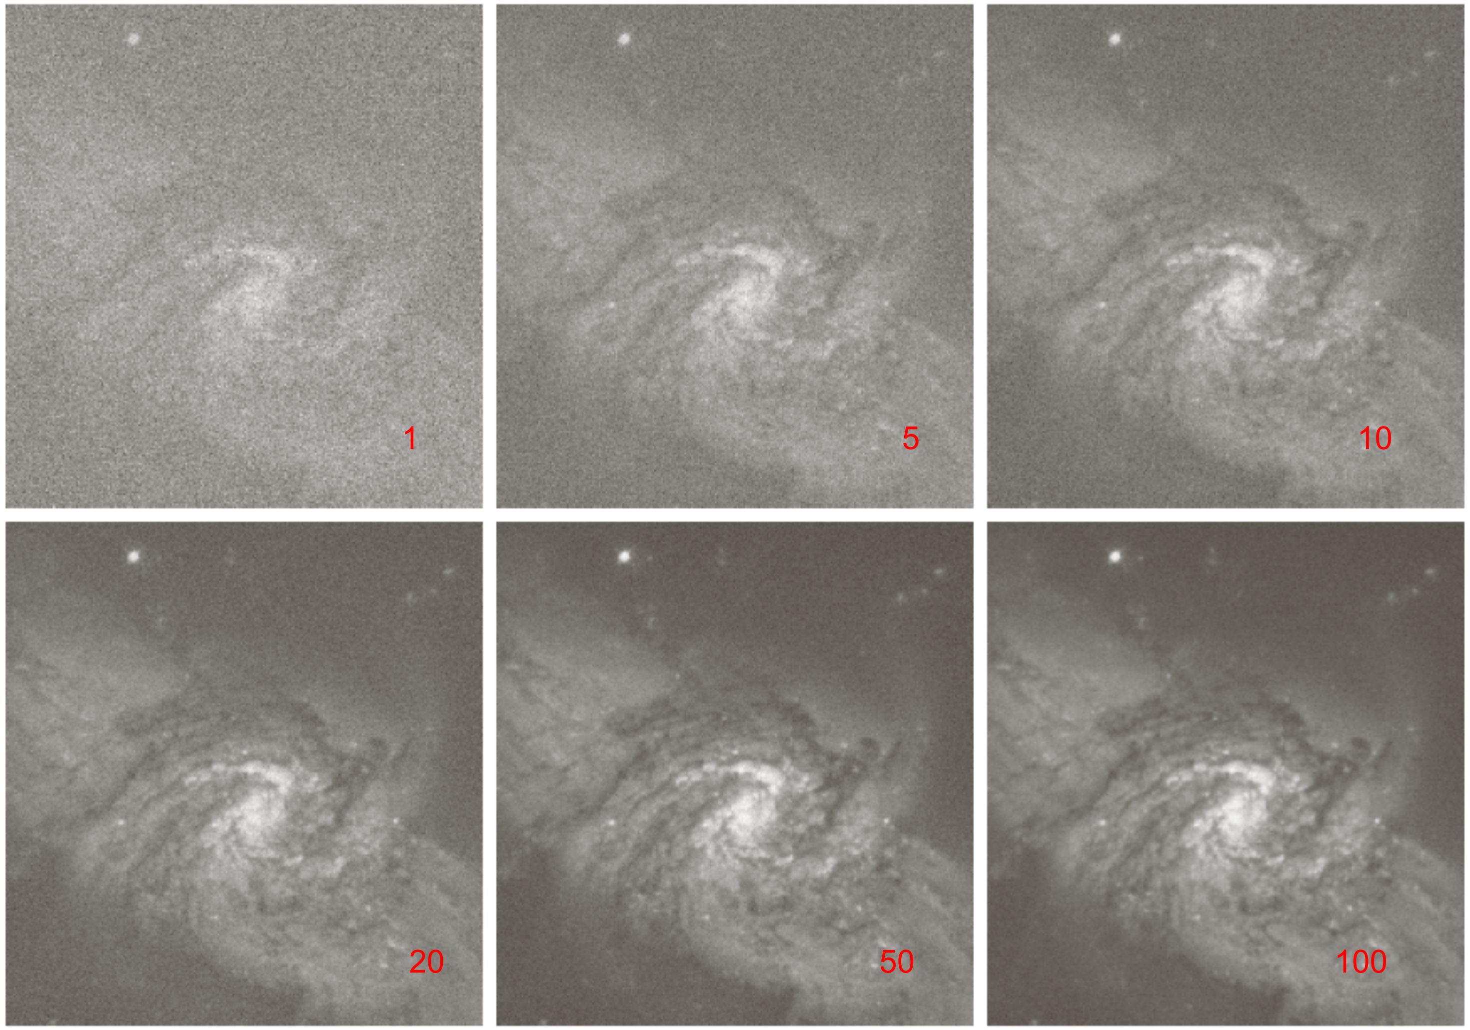

Gamma correction

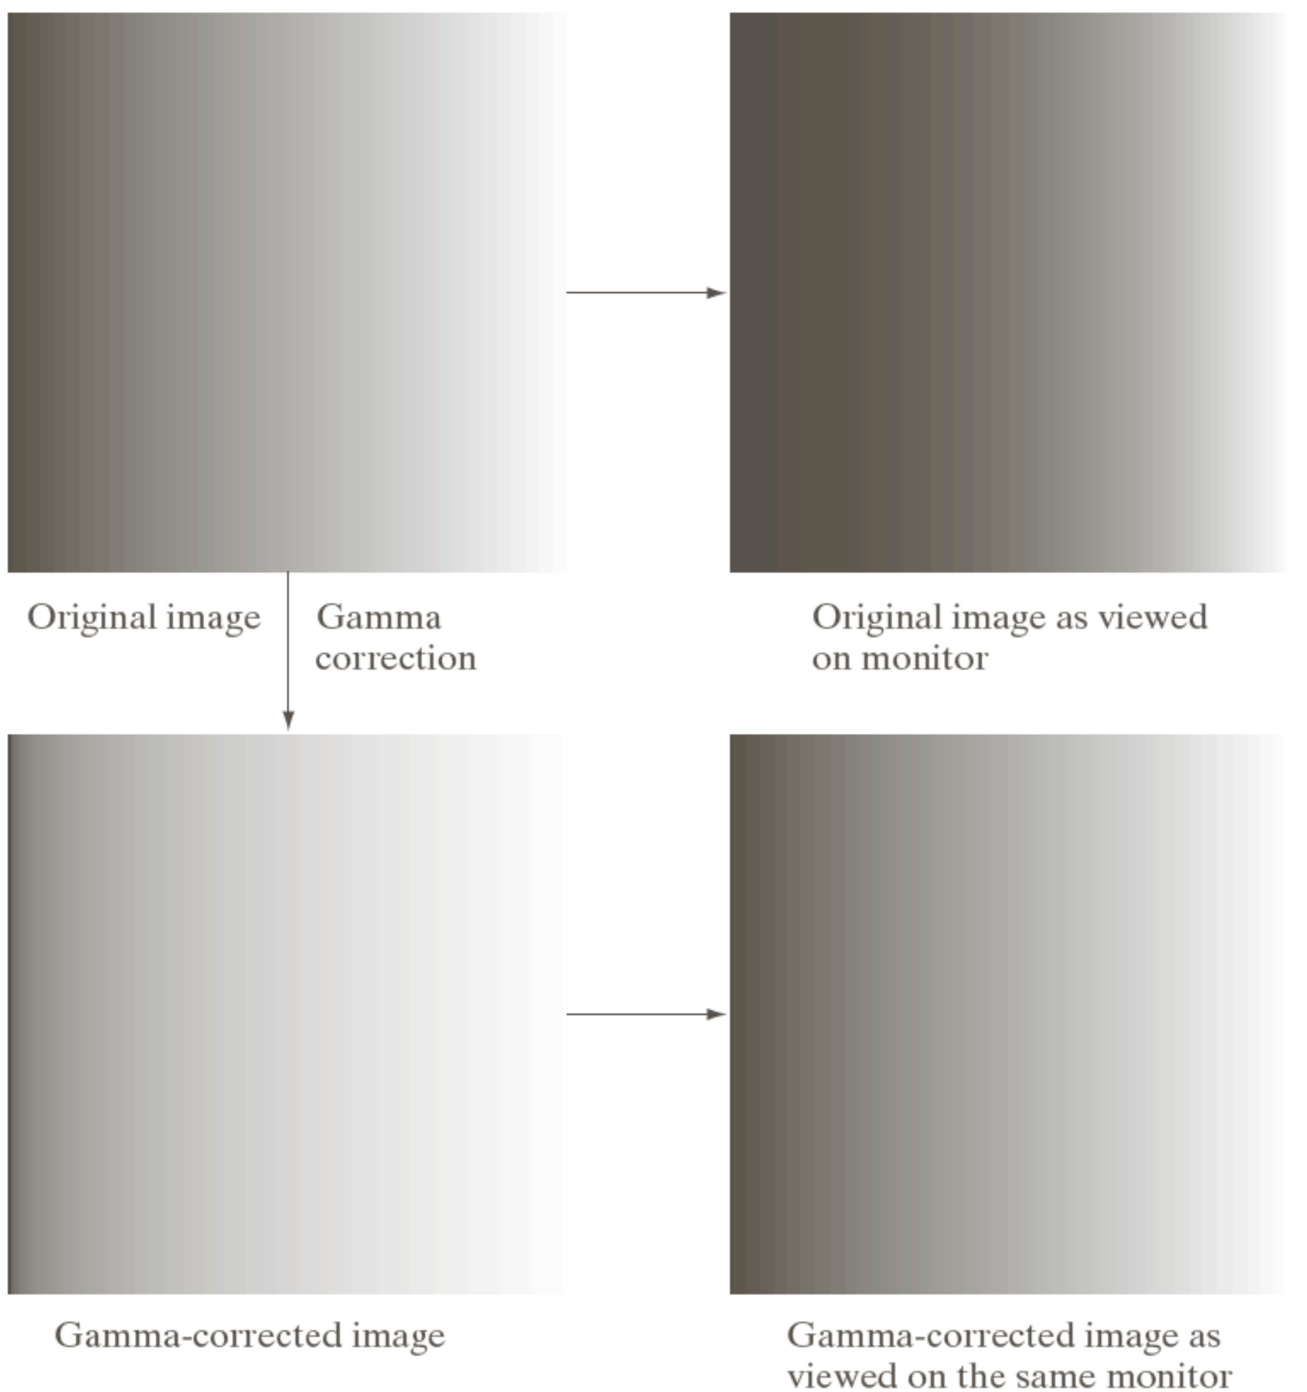

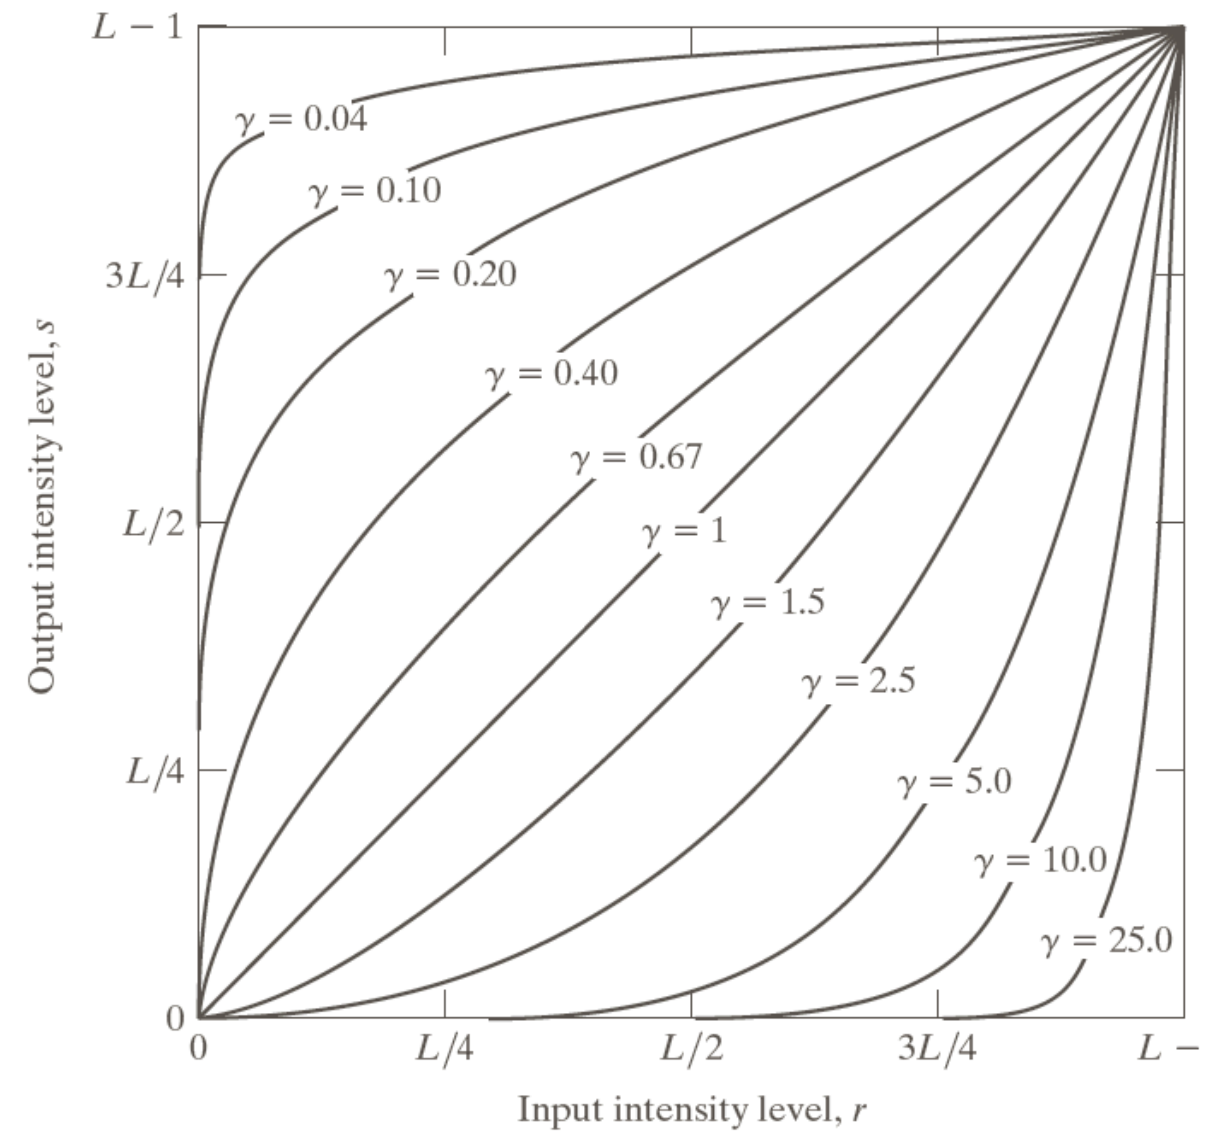

$$s = cr^{\gamma}$$

Gamma transformations

a	b
c	d

FIGURE 3.9 (a) Aerial image. (b)–(d) Results of applying the transformation in Eq. (3.2-3) with $c = 1$ and $\gamma = 3.0, 4.0,$ and 5.0, respectively. (Original image for this example courtesy of NASA.)

Gamma transformations

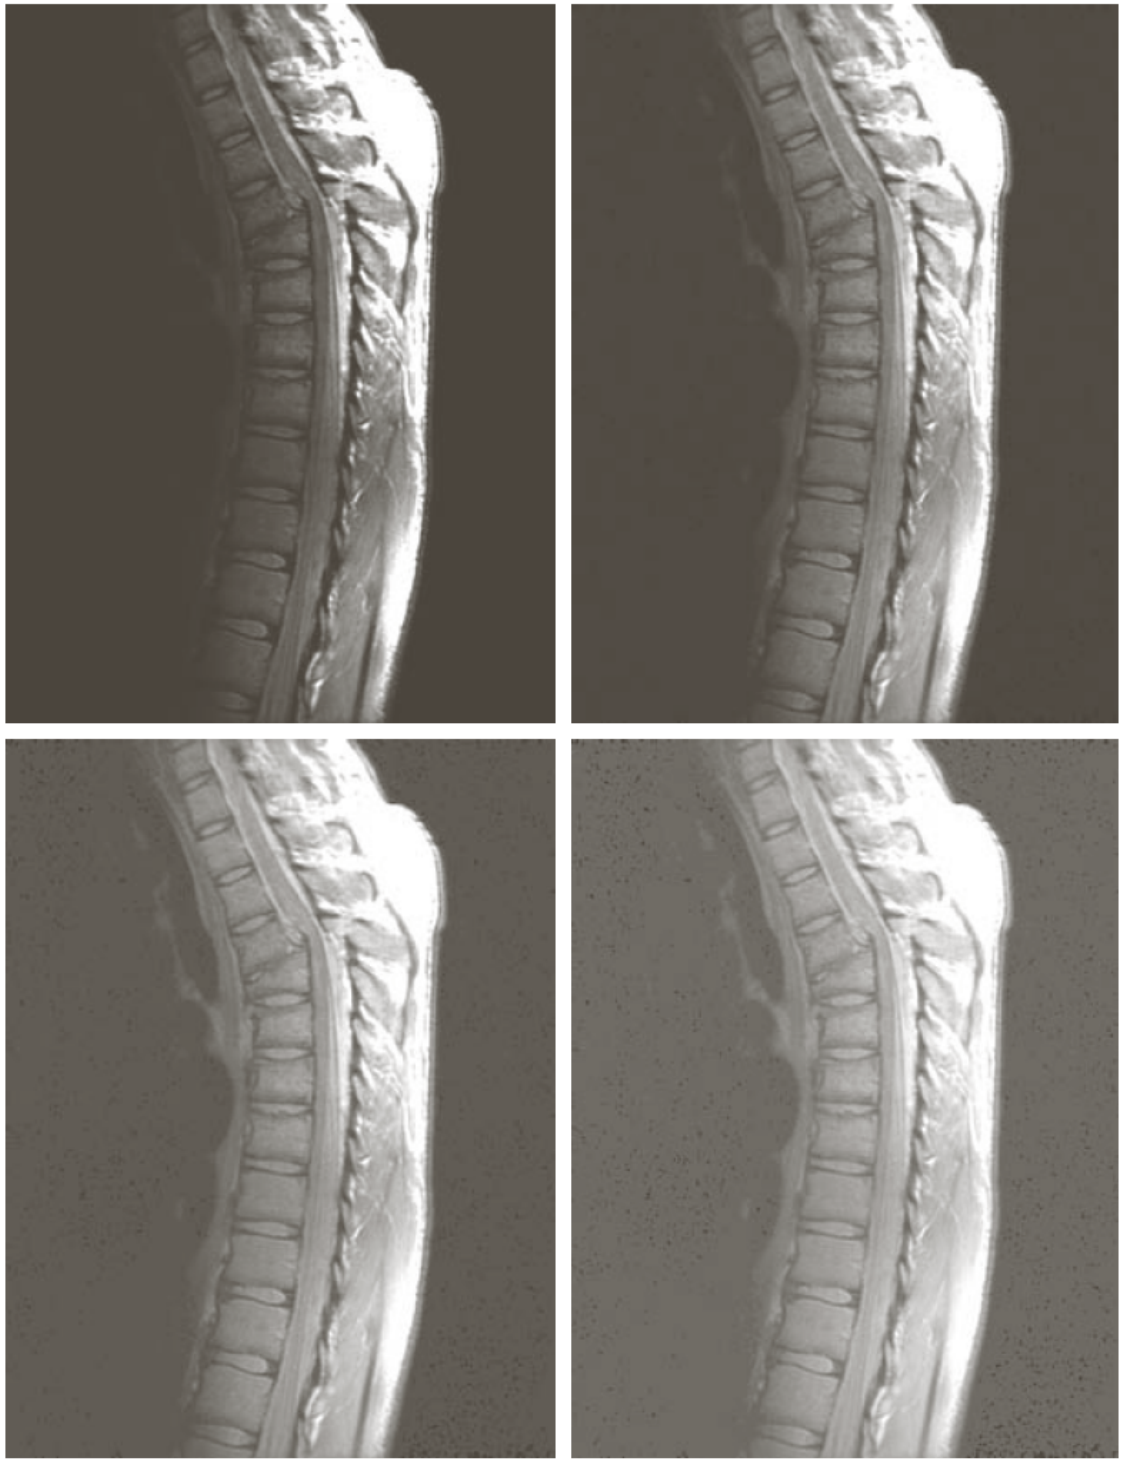

a b
c d

FIGURE 3.8 (a) Magnetic resonance image (MRI) of a fractured human spine. (b)–(d) Results of applying the transformation in Eq. (3.2-3) with $c = 1$ and $\gamma = 0.6, 0.4$, and 0.3, respectively. (Original image courtesy of Dr. David R. Pickens, Department of Radiology and Radiological Sciences, Vanderbilt University Medical Center.)

What happens to the histogram?

Automatic histogram equalization

- Cionsider image intensity as a continuous valued random variable r in the interval [0 , L-1] with pdf $p_r(r)$

- What kind of histograms do we want?

- Find an intensity transformation s=T(r) such that the pdf $p_s(s)$ is uniform in the interval [0 , L-1]

- How are the pdf's of r and s related?

$$s = T(r) \Rightarrow p_s\left(s\right) = p_r(r)\left|\frac{dr}{ds}\right|$$

Automatic histogram equalization

- Consider the transformation (range [0 , 1])

$$s = \int_0^r p_r(w)\,dw$$

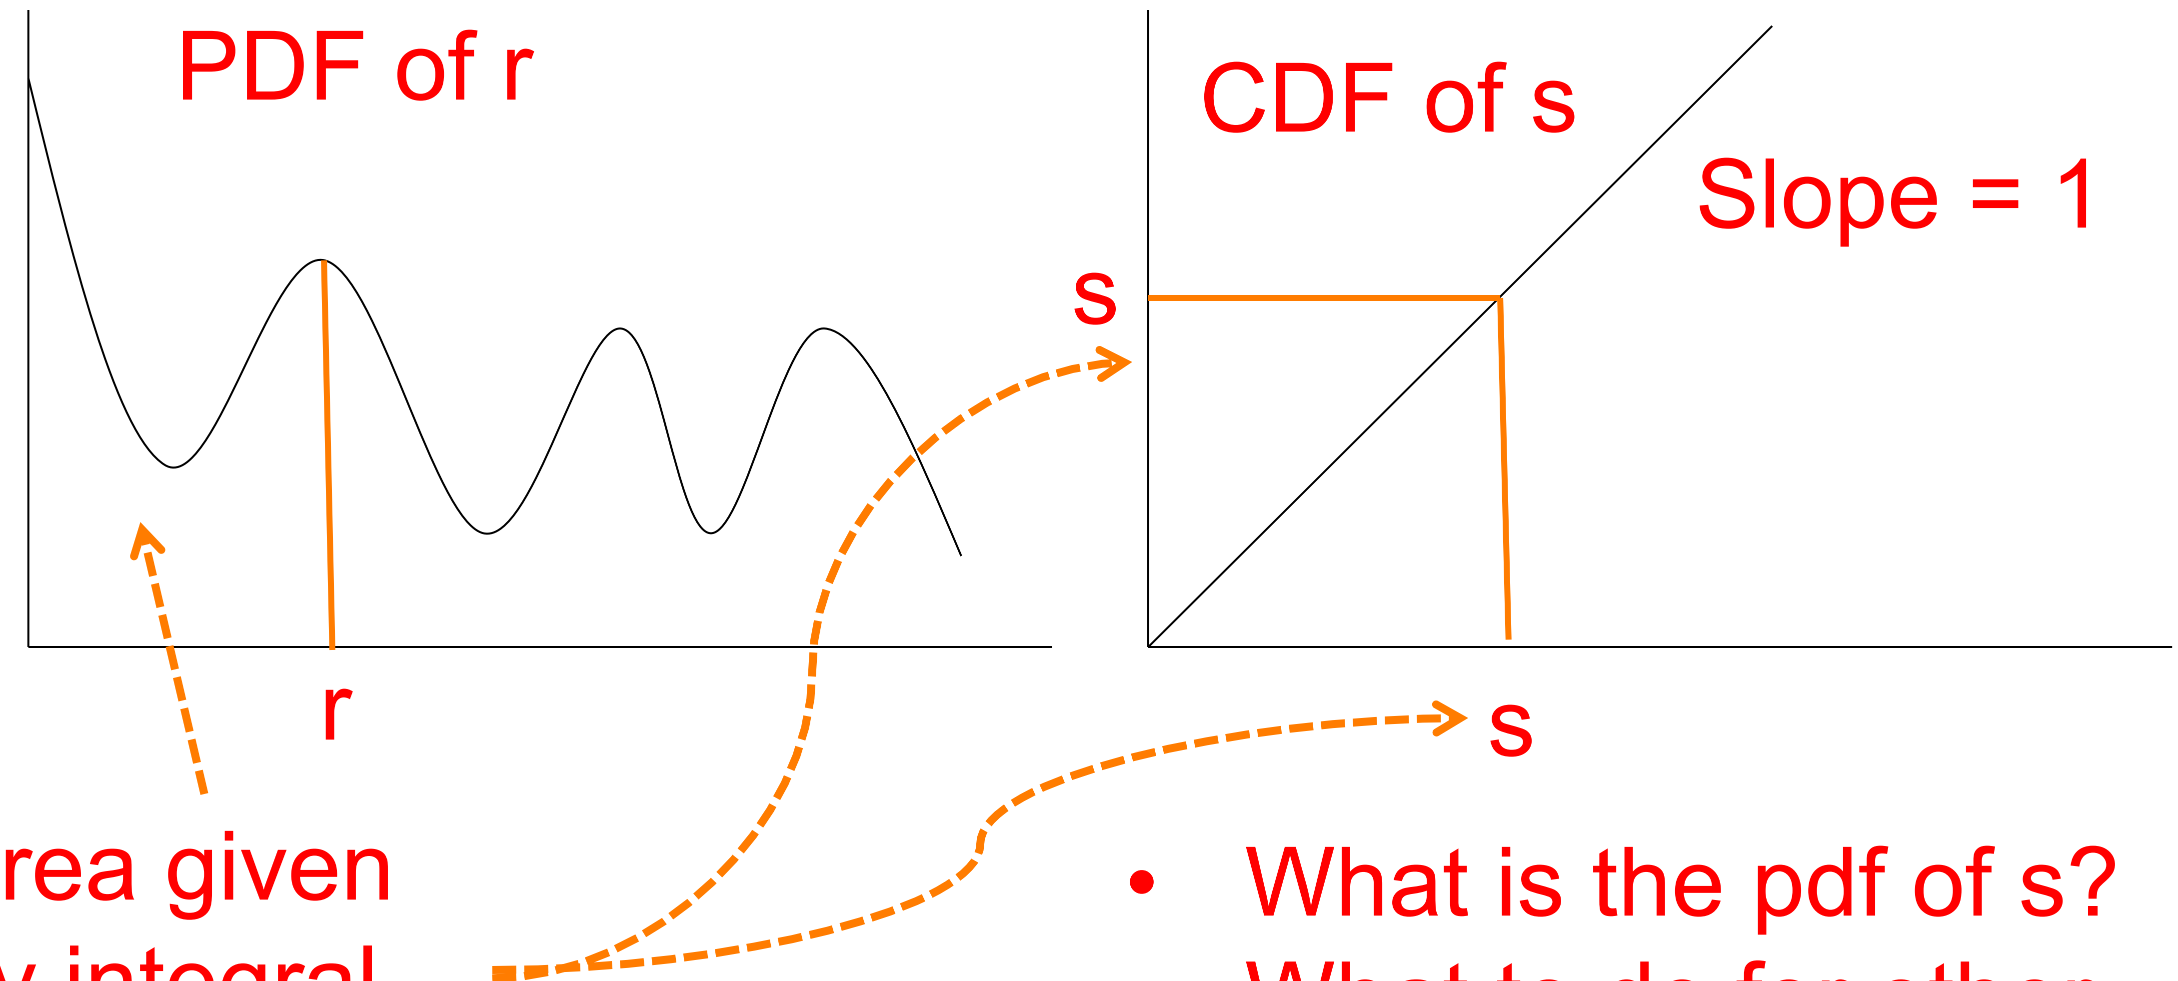

Area given by integral above

- What is the pdf of s?
- What to do for other intensity range?

Automatic histogram equalization

- Consider the transformation

$$s = (L-1)\int_0^r p_r(w)\,dw$$

- Compute ds/dr

Leibniz's rule

$$\frac{ds}{dr} = (L-1)\frac{d}{dr}\left[\int_0^r p_r(w)\,dw\right] = (L-1)\,p_r(r)$$

- Verify pdf for s is uniform

$$p_s(s) = p_r(r)\left|\frac{dr}{ds}\right| = p_r(r)\left|\frac{1}{(L-1)\,p_r(r)}\right| = \frac{1}{L-1}$$

$$0 \le s \le L-1$$

(Discrete) histogram equalization

$$s_k = (L-1)\sum_{j=0}^{k} p_r(r_j)$$

- Example with L=8

*s_0=7x0.19=1.33 -> **1***

*s_1=7x(0.19+0.25) = 3.08 -> **3***

….

r_k	n_k	$p_r(r_k) = n_k/MN$	S_k
$r_0 = 0$	790	0.19	1.33
$r_1 = 1$	1023	0.25	3.08
$r_2 = 2$	850	0.21	4.55
$r_3 = 3$	656	0.16	5.67
$r_4 = 4$	329	0.08	6.23
$r_5 = 5$	245	0.06	6.65
$r_6 = 6$	122	0.03	6.86
$r_7 = 7$	81	0.02	7.00

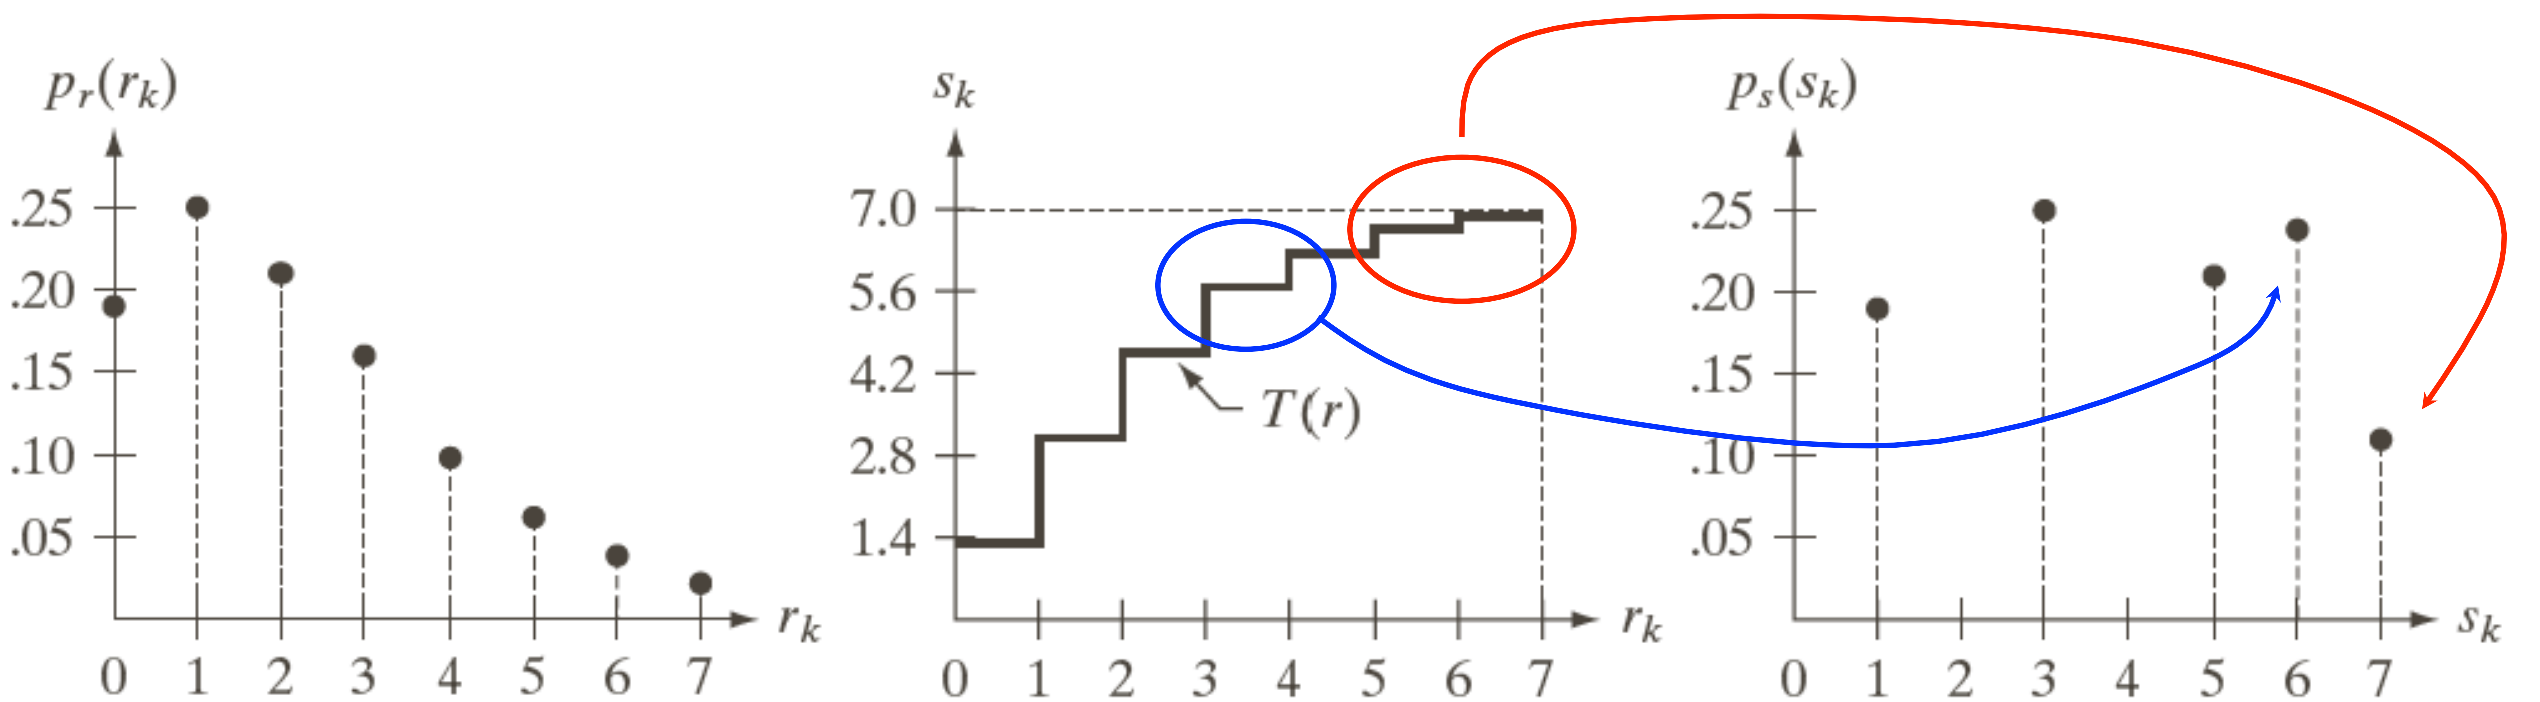

S_0=1 , S_1=3 , S_2=5 , S_3=6 , S_4=6 , S_5=7 , S_6=7 , S_7=7

Histogram equalization examples

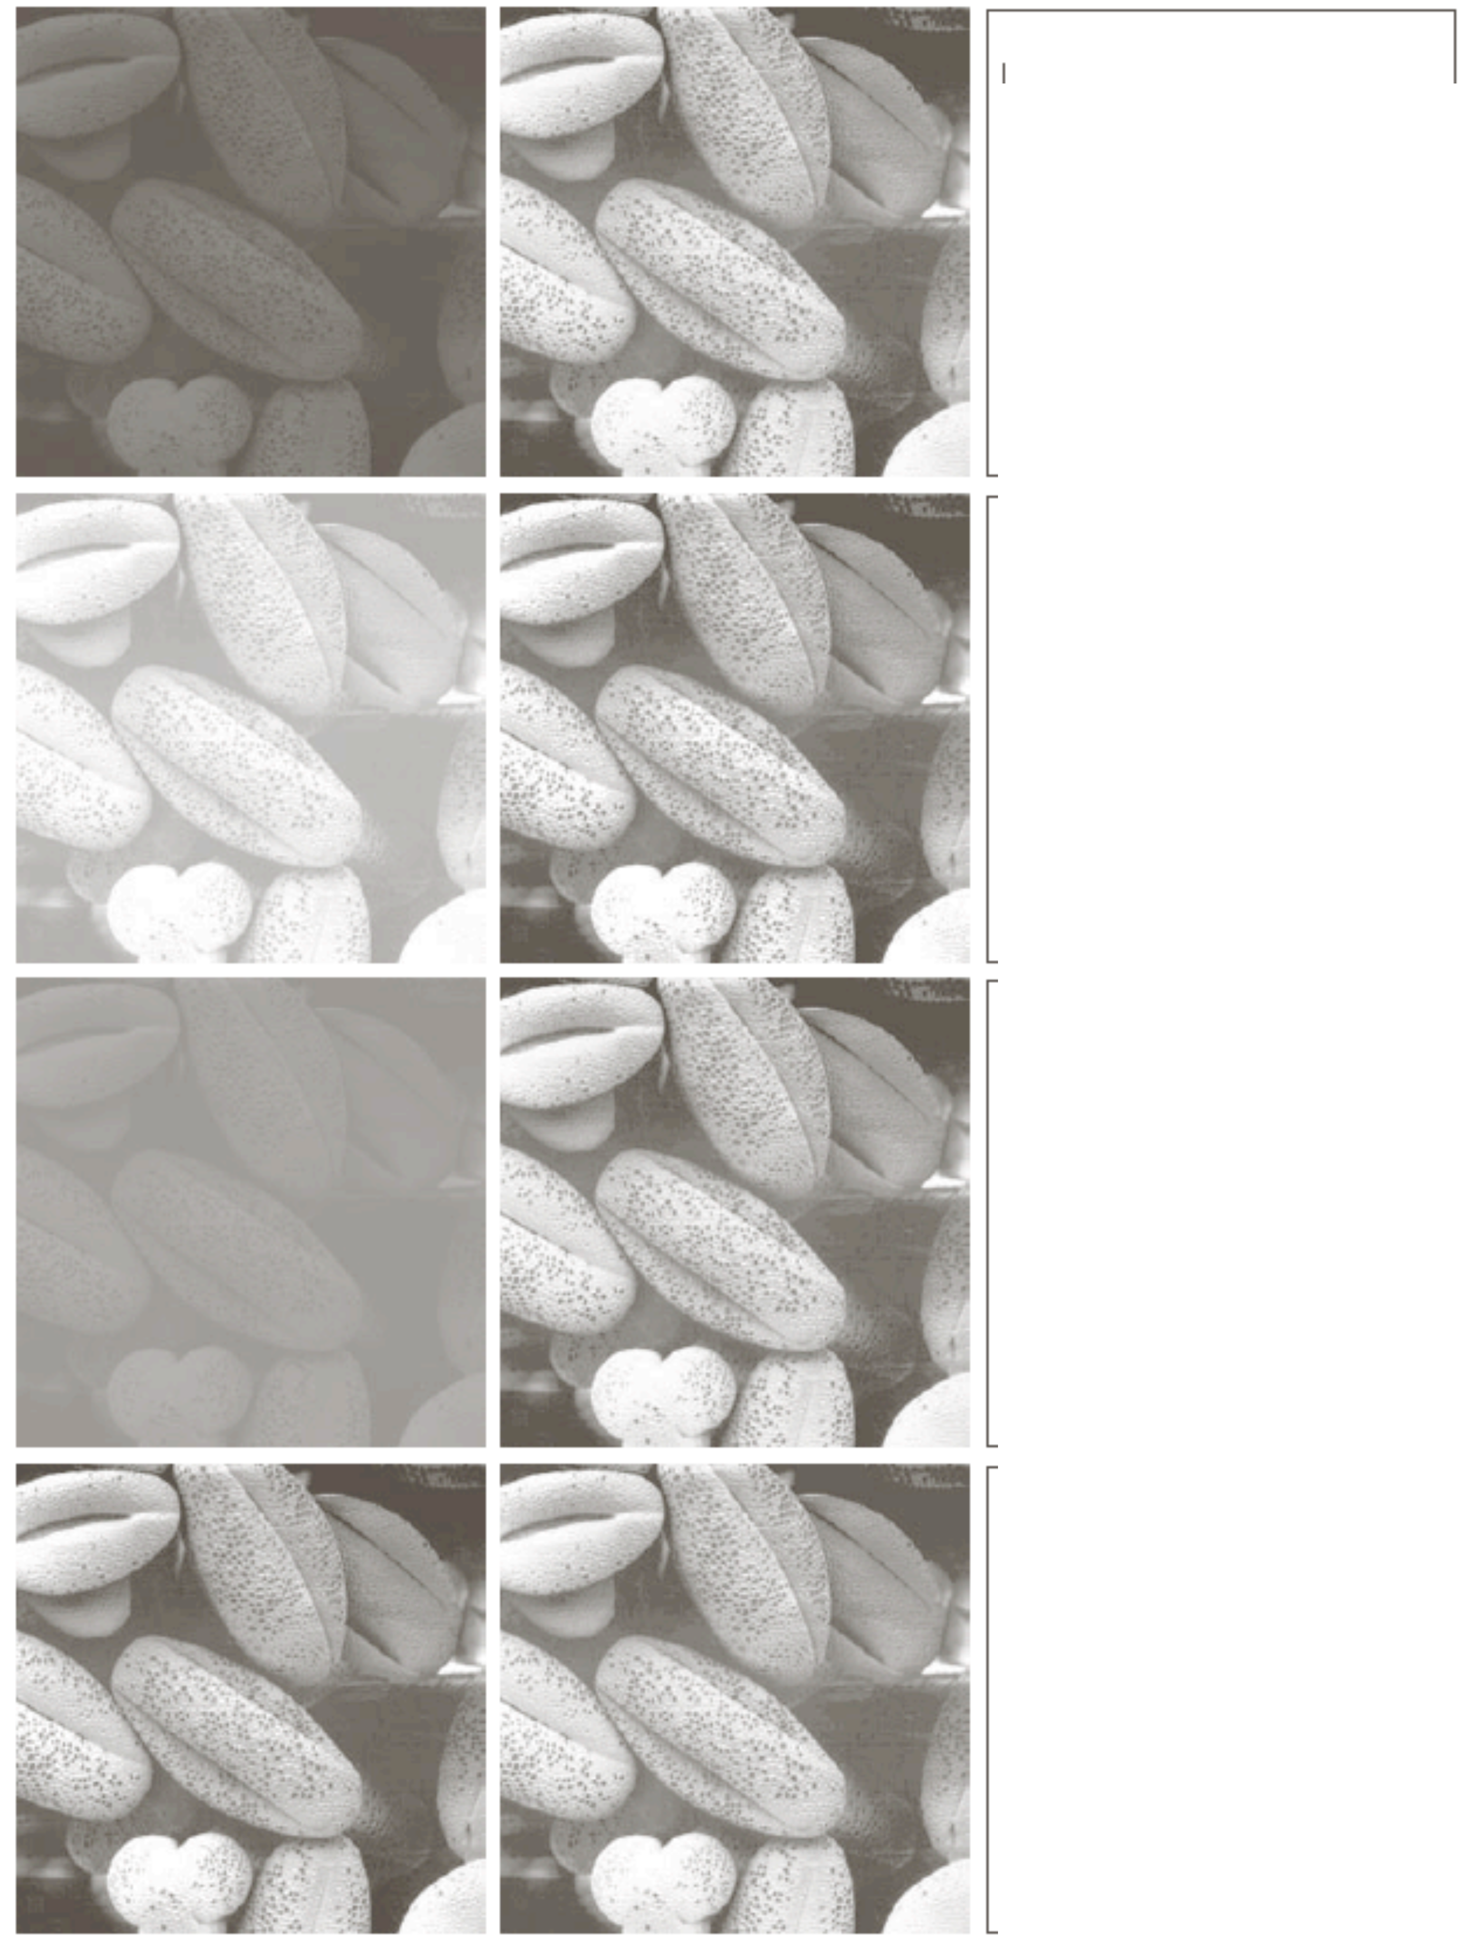

Histogram Equalization

Tuning Down Hist. Eq.

- Transformation is weighted combination of CDF and identity with parameter alpha

$$t(s) = (1 - \alpha)s + \alpha A(s)$$

α = 0.0

α = 0.2

α = 0.4

α = 0.6

α = 0.8

α = 1.0

Adaptive Histogram Equalization (AHE)

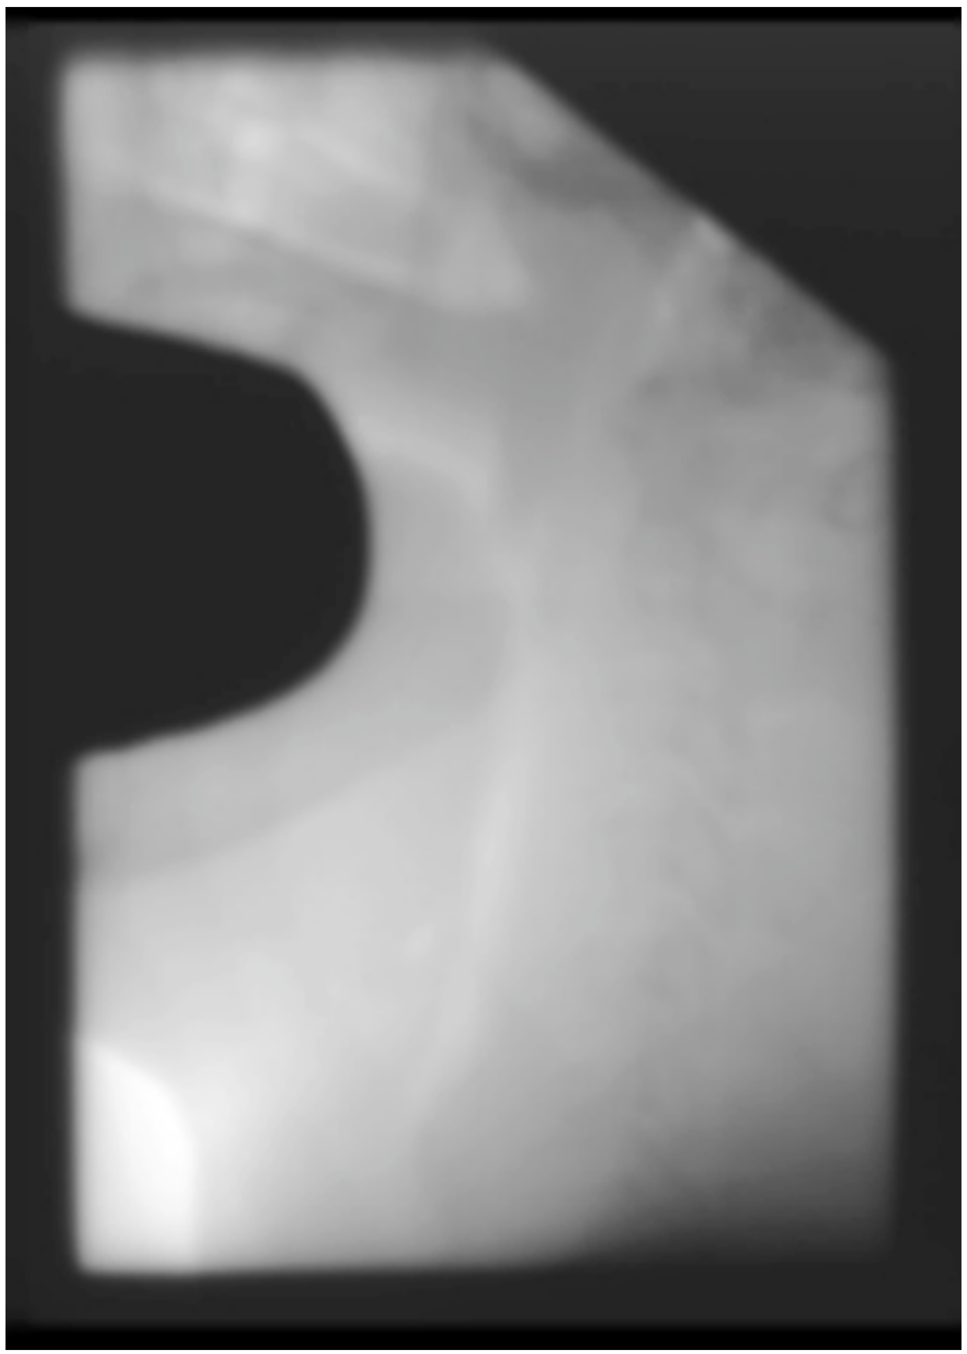

How to do adaptive histogram equalization

- Given a local image block, either
  - Do histogram equalization for that block, OR
  - Standardize the mean and variance of each block

$$g(x,y) = 128 + 50\frac{f(x,y) - m_B}{\sigma_B} \quad \text{for} \quad (x,y) \in B$$

- Repeat over all blocks (non-overlapping)
- Or, Sliding window approach

$$g(x,y) = 128 + 50\frac{f(x,y) - m_{B(x,y)}}{\sigma_{B(x,y)}}$$

Histogram statistics

- Image f(x,y), histogram $p(r_i)$
- Mean intensity

$$m = \sum_{i=0}^{L-1} r_i p(r_i) = \frac{1}{MN} \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x,y)$$

- Second central moment (intensity variance)

$$\sigma^2 = \sum_{i=0}^{L-1} (r_i - m)^2 \, p(r_i) = \frac{1}{MN} \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} \left(f(x,y) - m \right)^2$$

Local histogram statistics

- B_{xy} a neighborhood centered around (x,y)
- We can compute the intensity histogram and its statistics (mean, variance, etc.) limited to this region

$$m_{B_{xy}} = \sum_{i=0}^{L-1} r_i p_{B_{xy}}(r_i) = \frac{1}{|B_{xy}|} \sum_{(x,y)\in B_{xy}} f(x,y)$$

$$\sigma_{B_{xy}}^2 = \sum_{i=0}^{L-1} \left(r_i - m_{B_{xy}}\right)^2 p_{B_{xy}}(r_i) = \frac{1}{|B_{xy}|} \sum_{(x,y)\in B_{xy}} \left(f(x,y) - m_{B_{xy}}\right)^2$$

Local histogram statistics

- For instance, lets define S_{xy} as a 11 x 11 neighborhood centered at (x,y). Then

$$m_{B_{xy}} = \frac{1}{121}\sum_{u=-5}^{5}\sum_{v=-5}^{5} f(x+u, y+v)$$

$$\sigma_{B_{xy}}^2 = \frac{1}{121}\sum_{u=-5}^{5}\sum_{v=-5}^{5}\left(f(x+u, y+v) - m_{B_{xy}}\right)^2$$

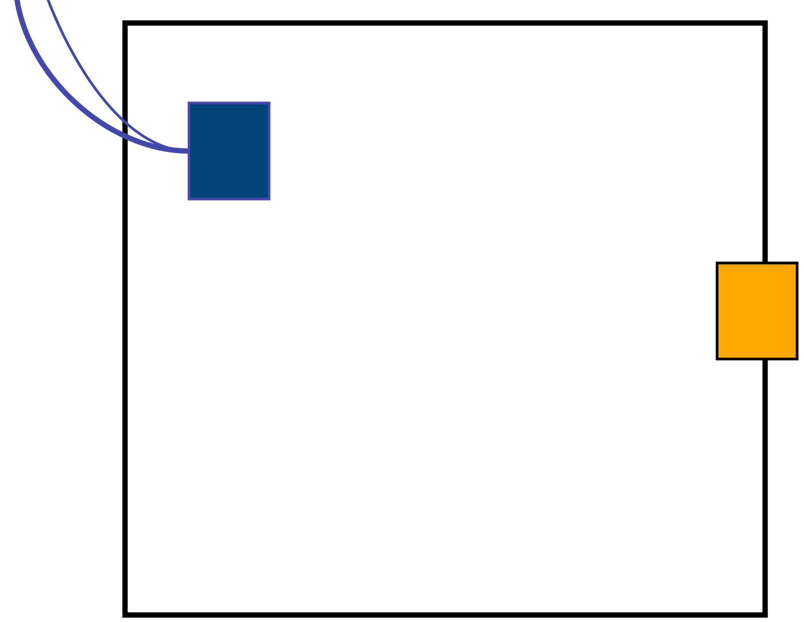

Have to take care at image boundaries

A faster alternative

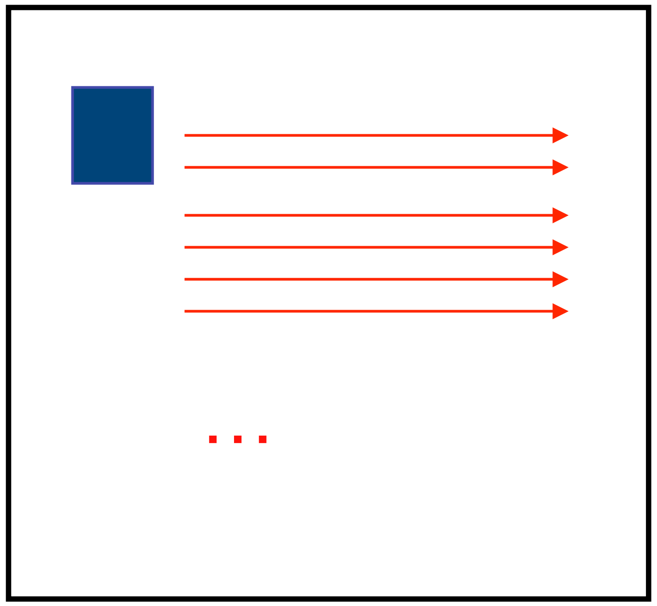

instead

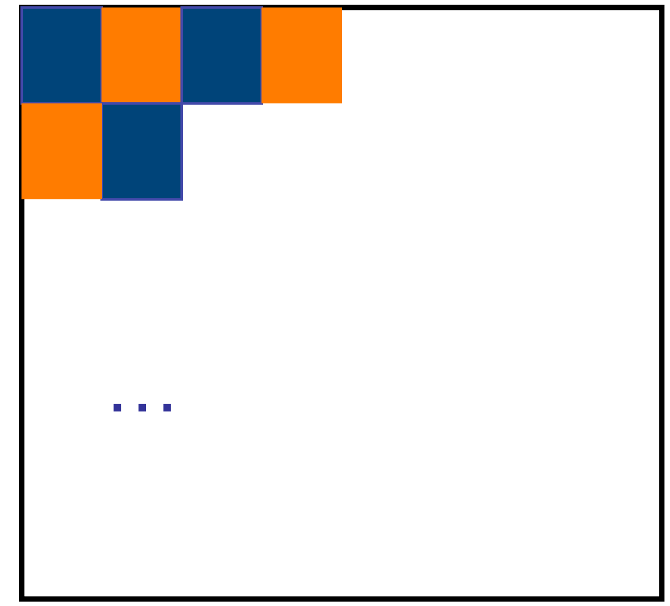

Compute statistics at every pixel

Divide image into blocks. Then compute statistics

- Much more efficient
- Can create a blocky effect in the output image

AHE Gone Bad…

What can we do about this?

Effect of Window Size

Orig

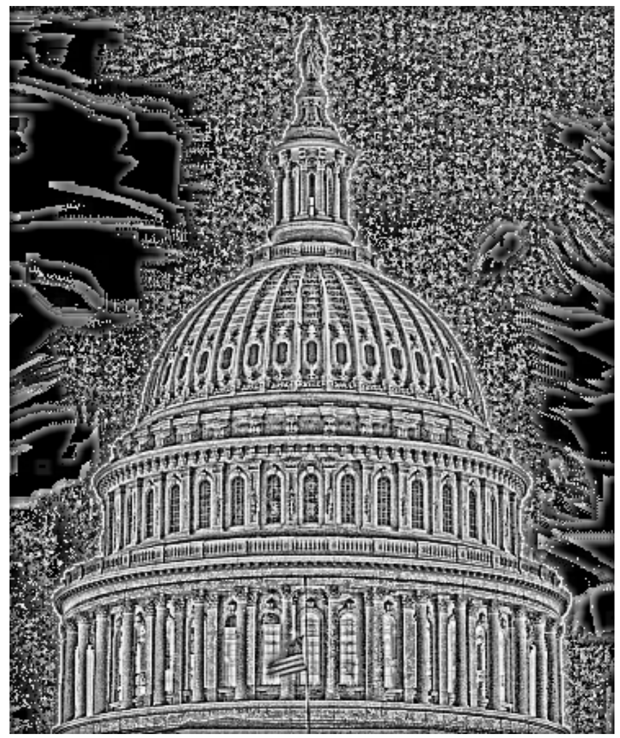

10x10

25x25

50x50

AHE Application: Microscopy Imaging

Original

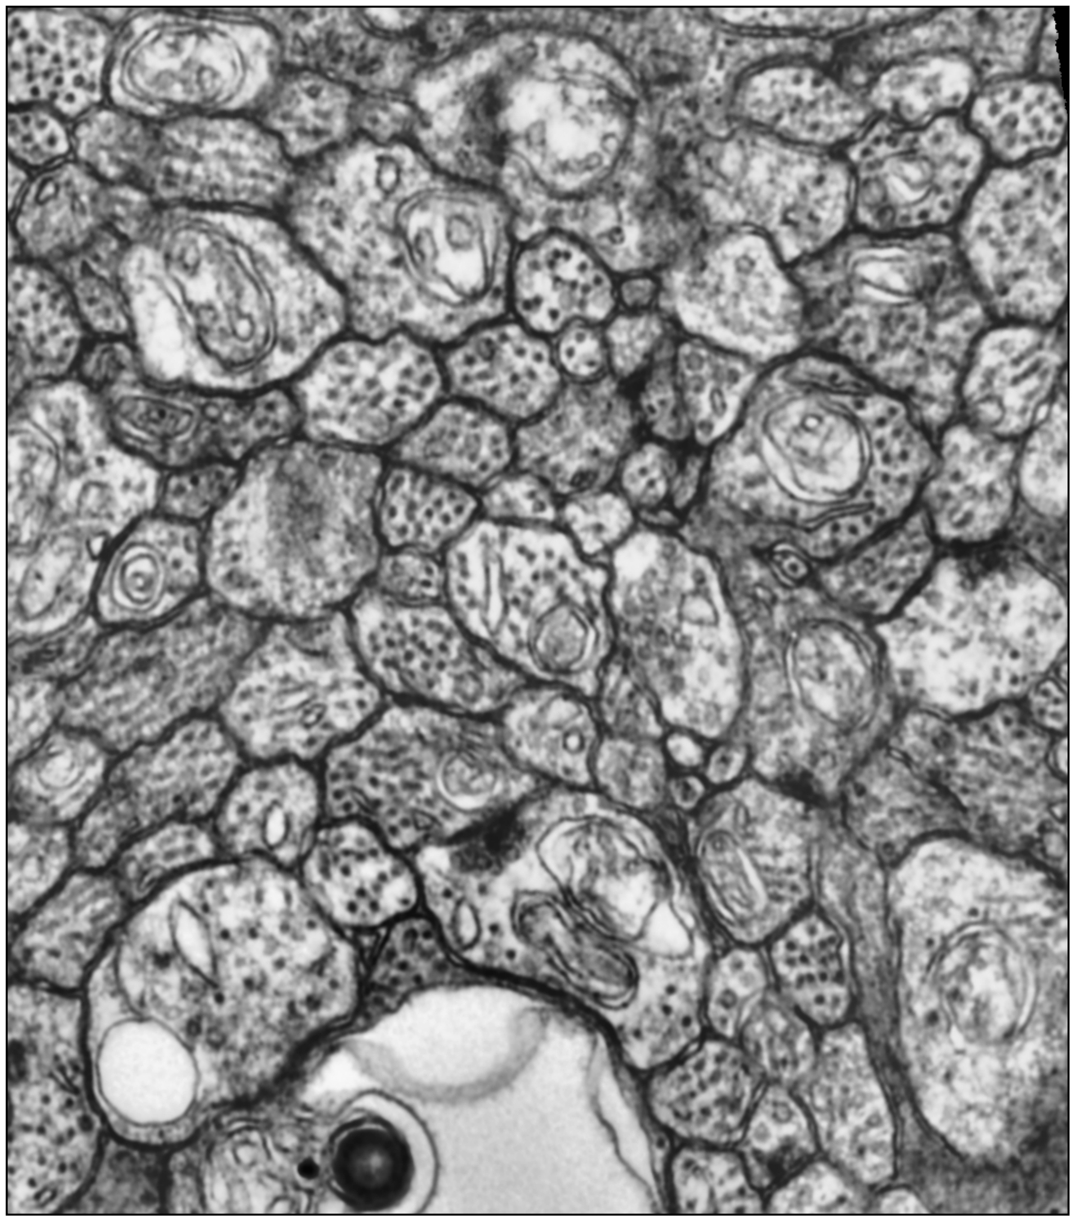

AHE

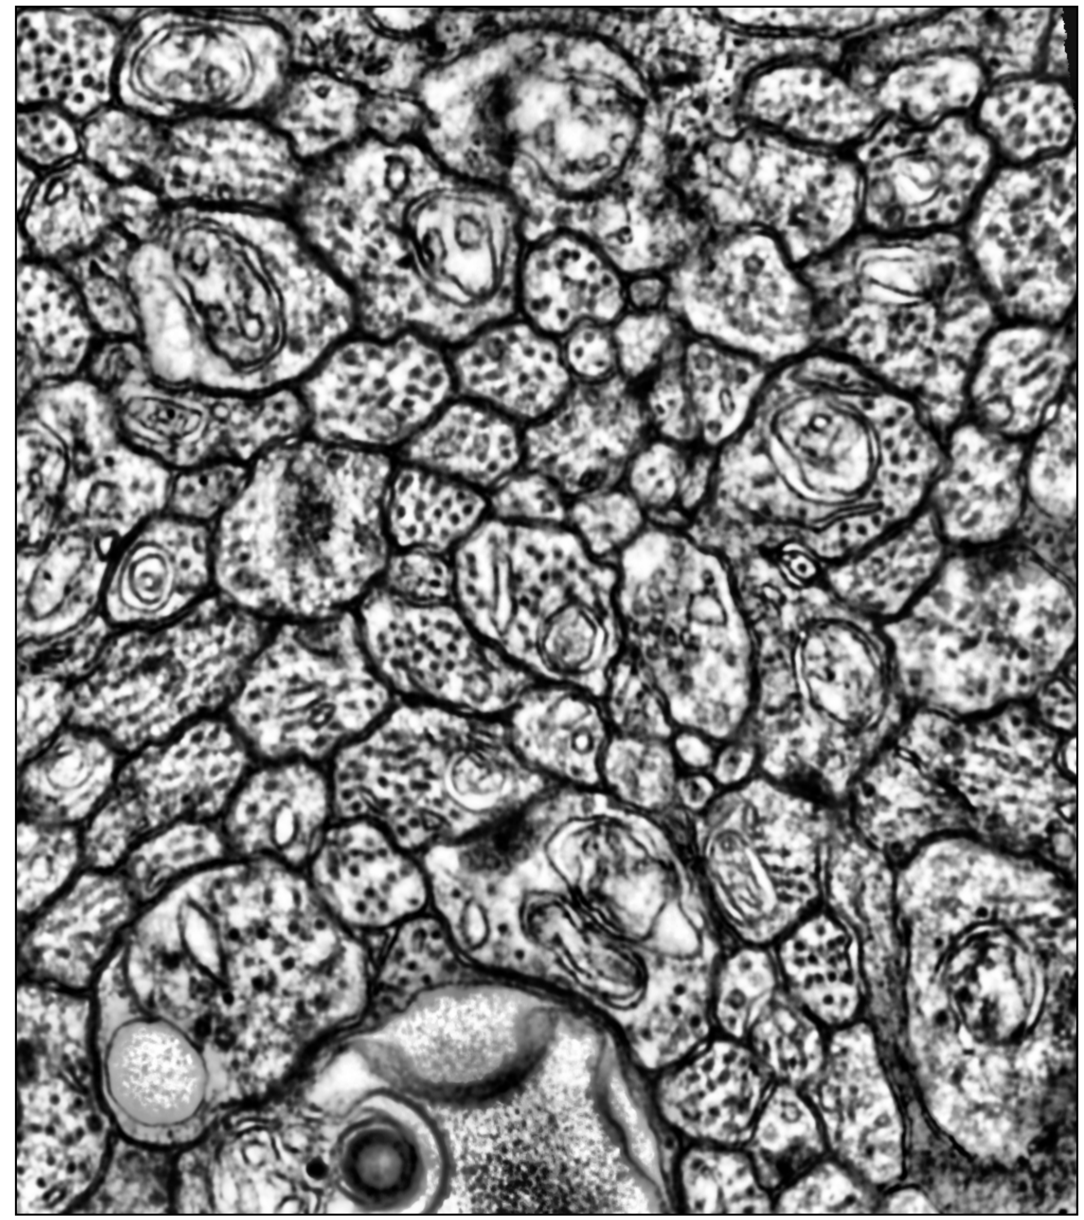

Histogram matching

- Histogram equalization aims for an uniform histogram for the output image
- Sometimes we might want to specify different histograms as the target

$$s = T(r) = (L-1)\int_0^r p_r(w)\,dw \qquad s = G(z) = (L-1)\int_0^z p_z(t)\,dt$$

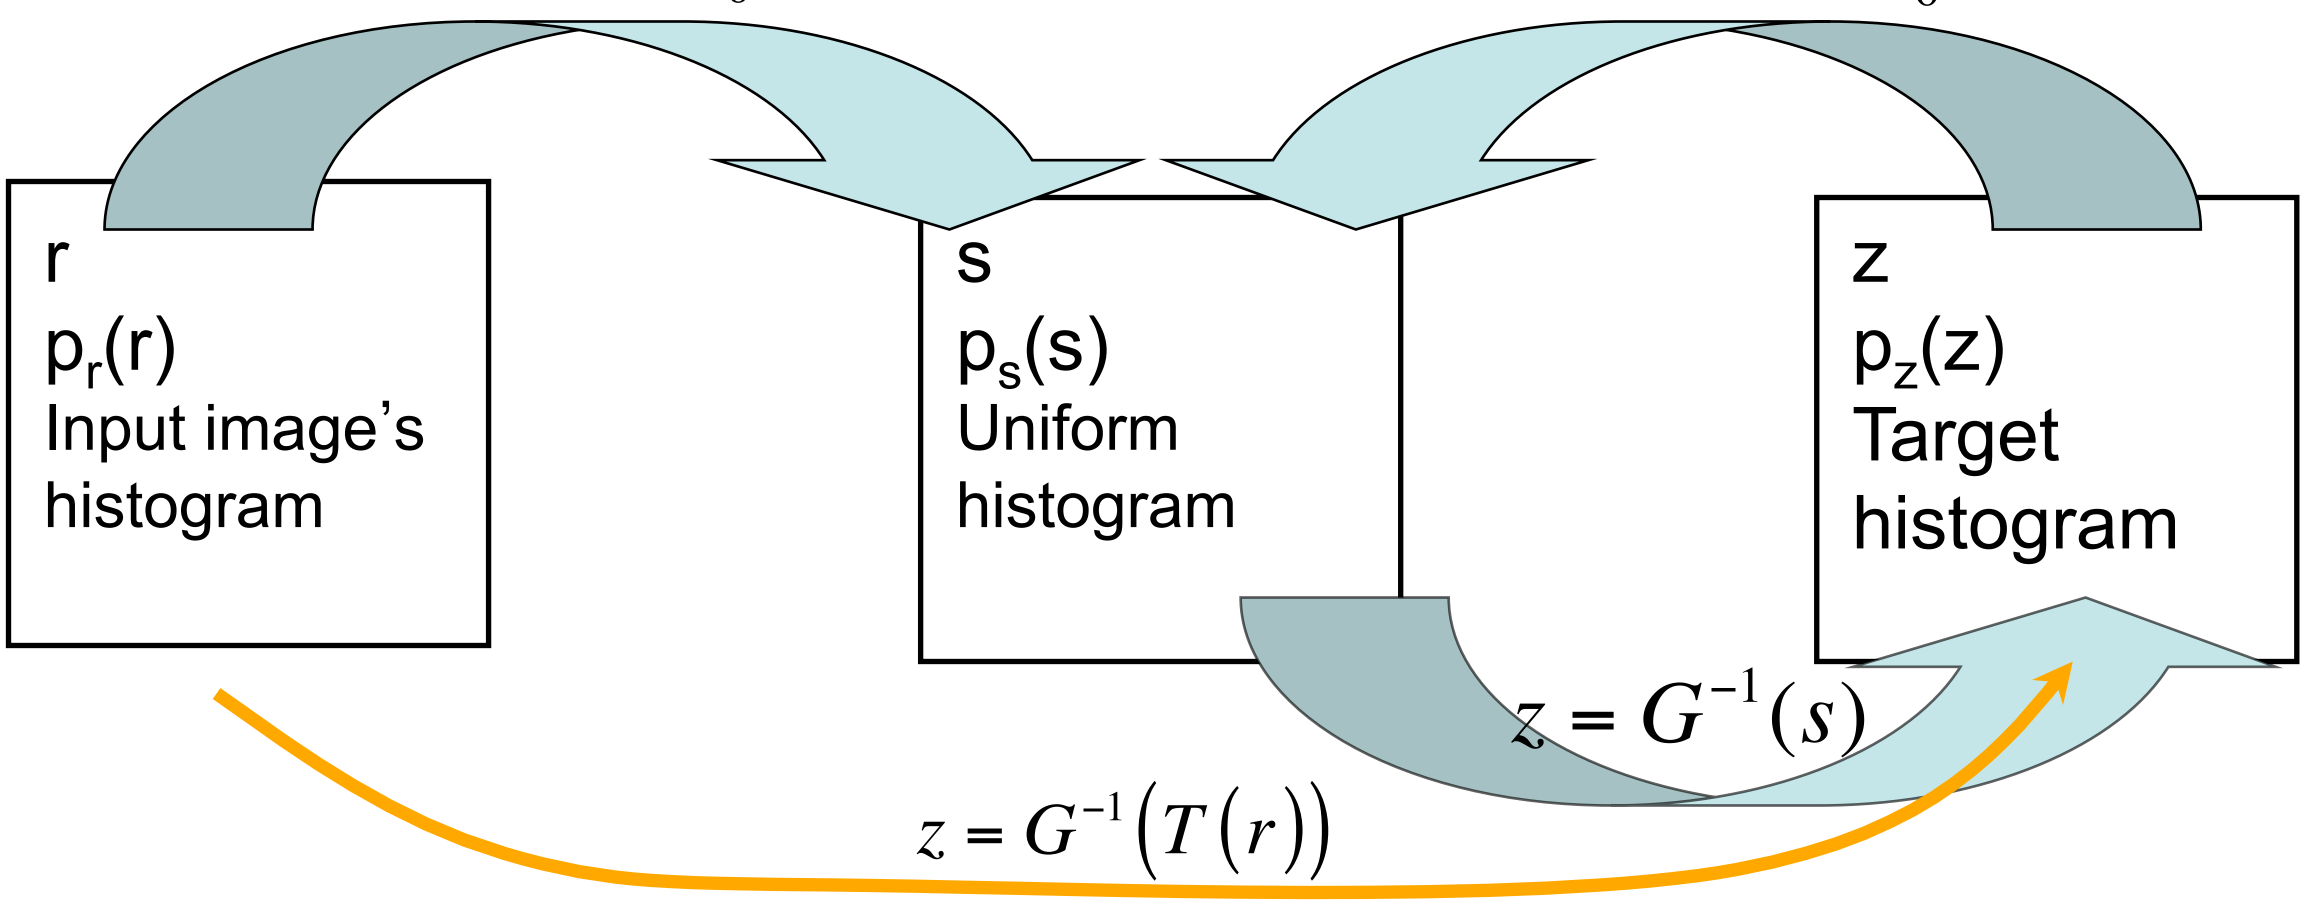

Invertible (one-to-one) transformations

- We need to be able to take the inverse of G(z)

$$s = G(z) = \left(L-1\right)\int_0^z p_z\left(t\right)dt$$

- G(z) has an inverse if it is strictly monotonically increasing as a function of z
  - This is true if $p_z(z)>0$ for all z

Histogram matching (discrete)

1. Compute $G(z_q) = (L-1)\sum_{i=0}^{q} p_z(z_i)$

   and store in a table

2. Compute $s_k = T(r_k) = (L-1)\sum_{i=0}^{k} p_r(r_j)$

3. Create mapping from s_k to z_q (inverse of G) by finding the closest value in the table stored from step 1 to all $s_{k.}$

4. Apply mapping created in step 3 to histogram equalized input image

Example

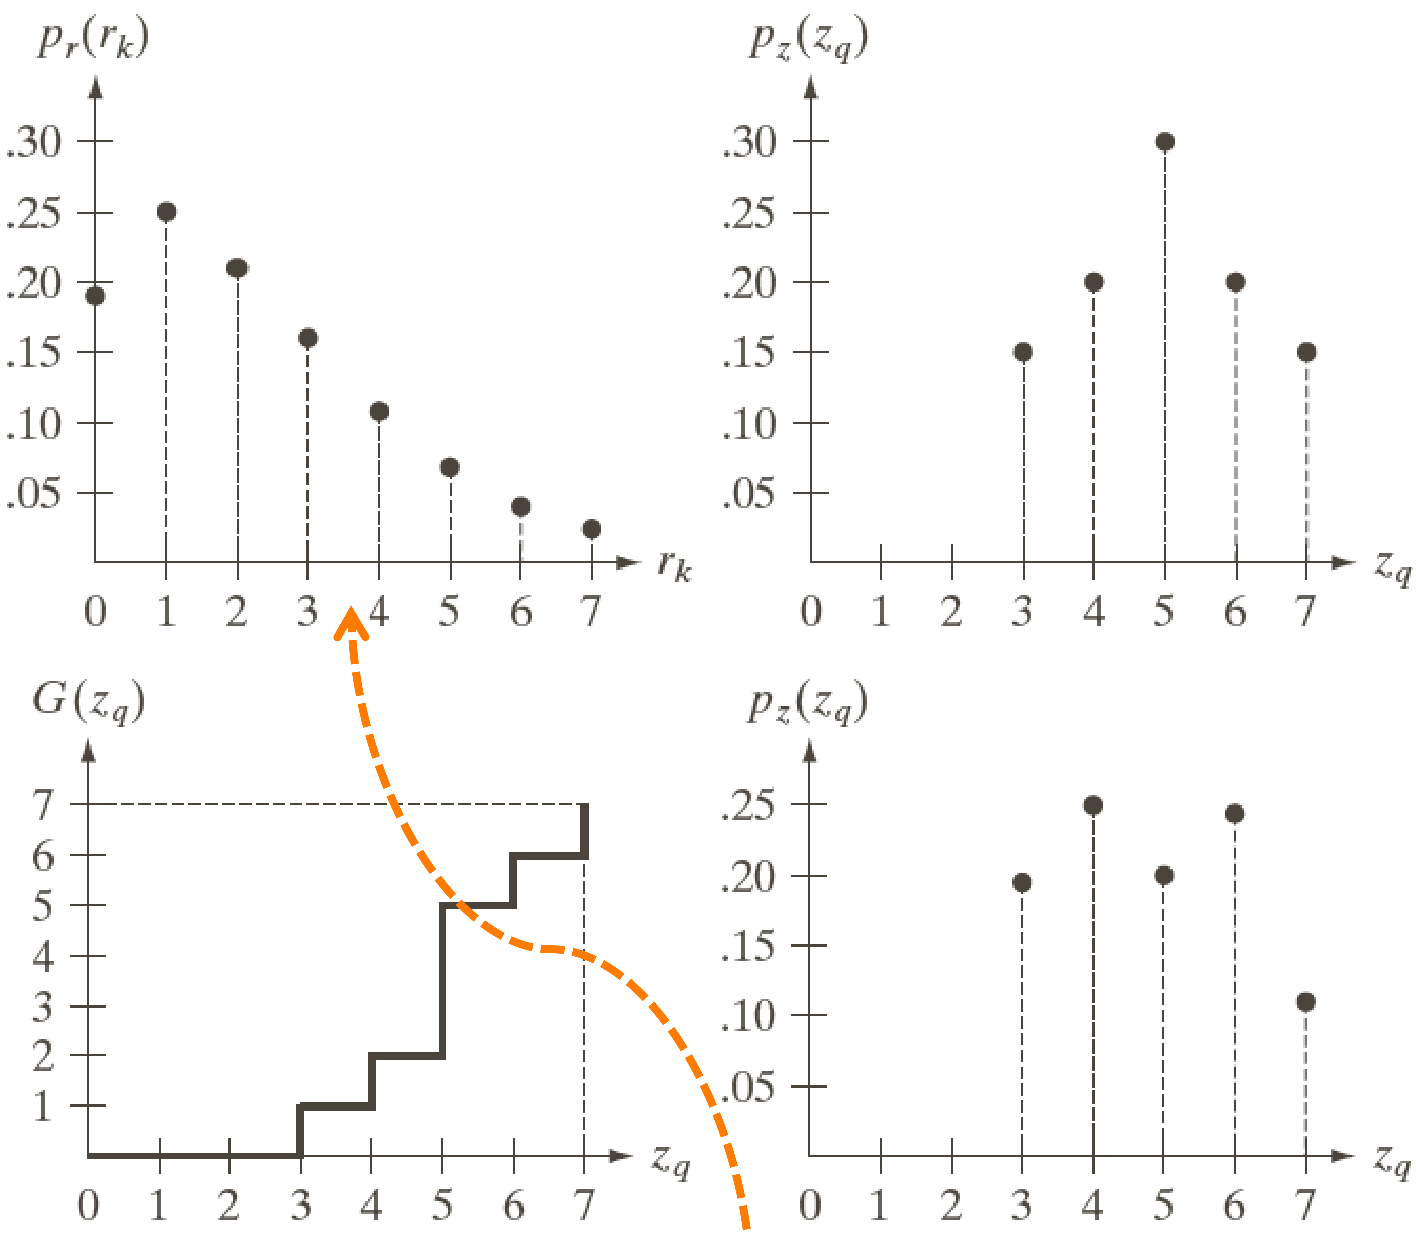

1

z_q	Specified $p_z(z_q)$	Actual $p_z(z_k)$	G
$z_0 = 0$	0.00	0.00	0.00
$z_1 = 1$	0.00	0.00	0.00
$z_2 = 2$	0.00	0.00	0.00
$z_3 = 3$	0.15	0.19	1.05
$z_4 = 4$	0.20	0.25	2.45
$z_5 = 5$	0.30	0.21	4.55
$z_6 = 6$	0.20	0.24	5.95
$z_7 = 7$	0.15	0.11	7.00

2 Did this one before

$S_0=1$, $S_1=3$, $S_2=5$, $S_3=6$,
$S_4=6$, $S_5=7$, $S_6=7$, $S_7=7$

Inverse of G

3

s_k	$\rightarrow$	z_q
1	$\rightarrow$	3
3	$\rightarrow$	4
5	$\rightarrow$	5
6	$\rightarrow$	6
7	$\rightarrow$	7

4. r_k -> s_k -> z_q

Colorization example

Match histograms of the gray scale input image to histograms of the Red, Green and Blue channels of a reference image individually to create 3 new images.

Use these new images as red, green blue channels to create the output color image.

To get realistic results input and reference image contents should be similar.

MATLAB code

- Given a gray scale image I2gray and a reference image I1color
- Uses image processing toolbox command imhistmatch
- I2red = imhistmatch(uint8(I2gray),I1color(:,:,1));
- I2blue = imhistmatch(uint8(I2gray),I1color(:,:,3));
- I2green = imhistmatch(uint8(I2gray),I1color(:,:,2));
- I2colorized(:,:,1)=I2red;
- I2colorized(:,:,2)=I2green;
- I2colorized(:,:,3)=I2blue;

Colorization example

Other uses of histograms

- So far we used histograms for image enhancement purposes
- Quantization
  - Reducing the number of gray levels or colors in an image for efficient storage
- Segmentation
  - Image segmentation is the process of subdividing an image into its constituent regions or objects.

What is image segmentation?

- Image segmentation is the process of subdividing an image into its constituent regions or objects.
- Example segmentation with two regions:

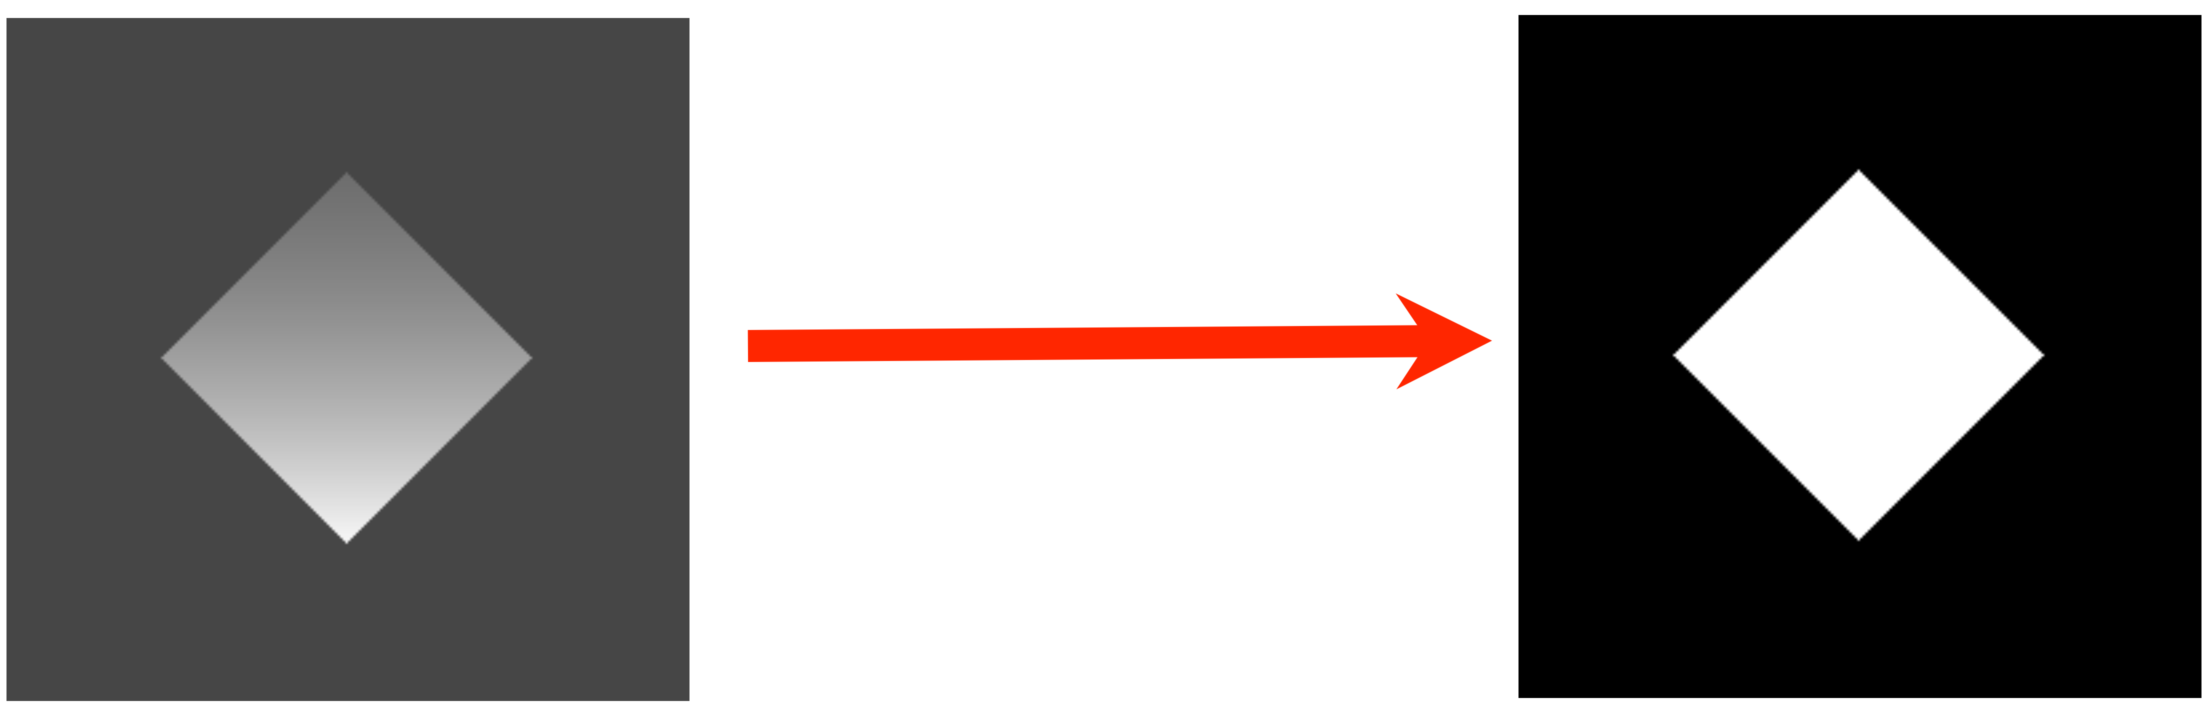

Input image
intensities 0-255

Segmentation output
0 (background)
1 (foreground)

Formal definition of segmentation

- R: set of all pixels in given image
- Segmentation into n regions $R_1, R_2, ... R_n$
- Two important properties
  - Complete

$$\bigcup_{i=1}^{n} R_i = R$$

  - Mutually exclusive

$$R_i \bigcap R_j = \emptyset \ if \ i \neq j$$

Two general approaches to segmentation

- Discontinuity based
  - Partition an image based on abrupt changes of intensity. Find pixels which correspond to these abrupt changes, these will be the boundaries between regions.
  - Edge detection is an example of this type of approach to segmentation
- Similarity based
  - Group pixels into regions of similar intensity
  - Thresholding is an example of this type of approach to segmentation

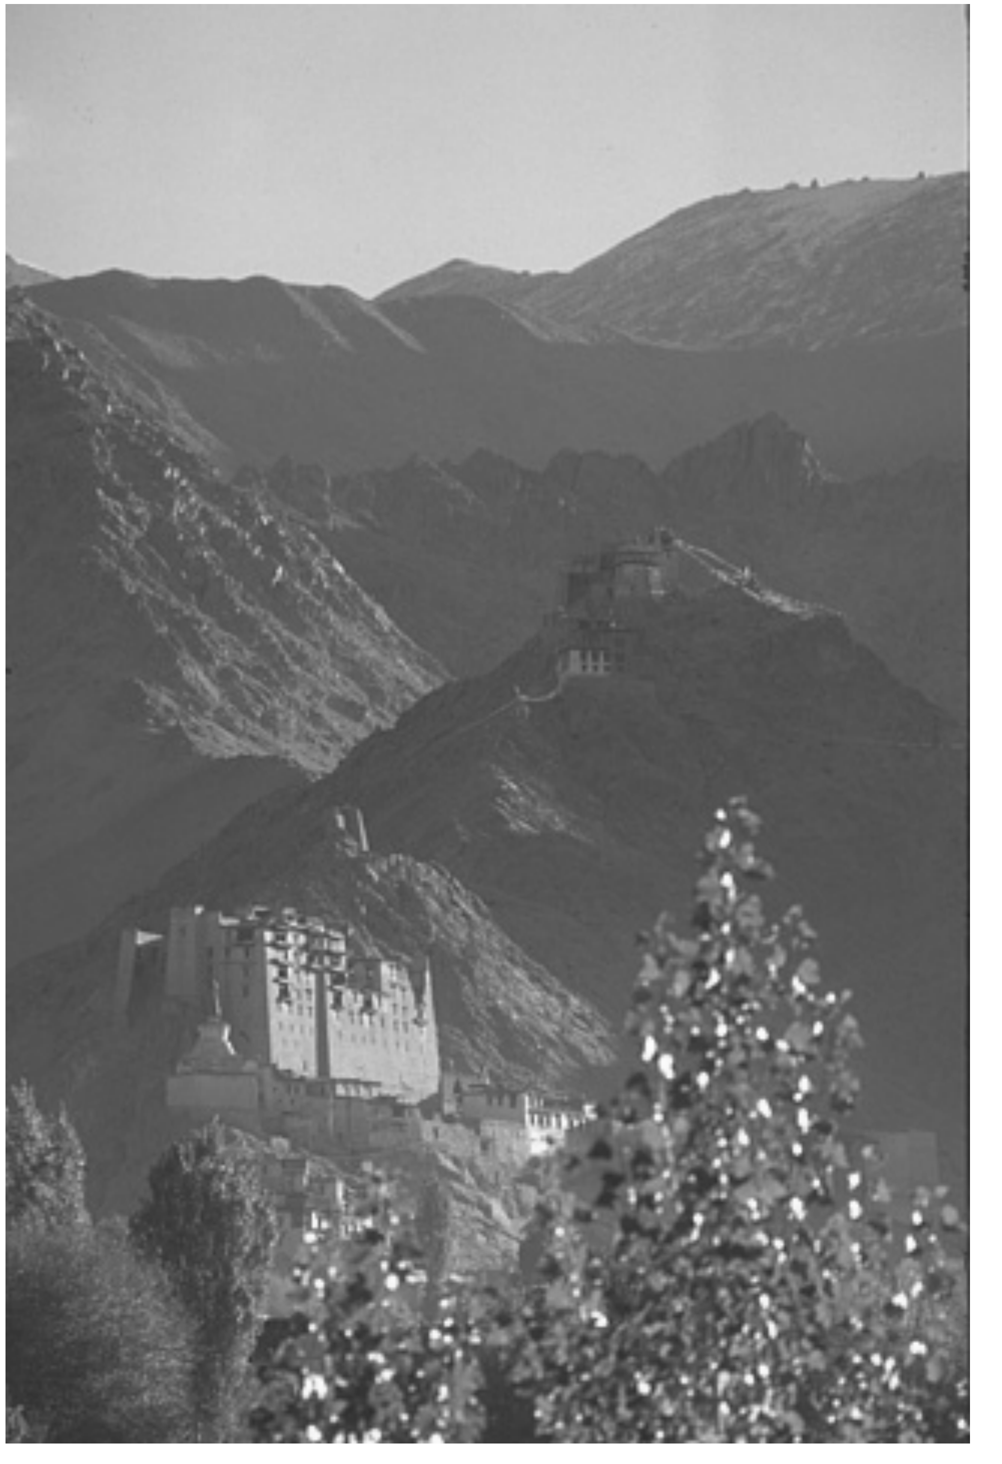

Global thresholding

$$g(x,y) = \begin{cases} 1 & if \quad f(x,y) > T \\ 0 & if \quad f(x,y) \le T \end{cases}$$

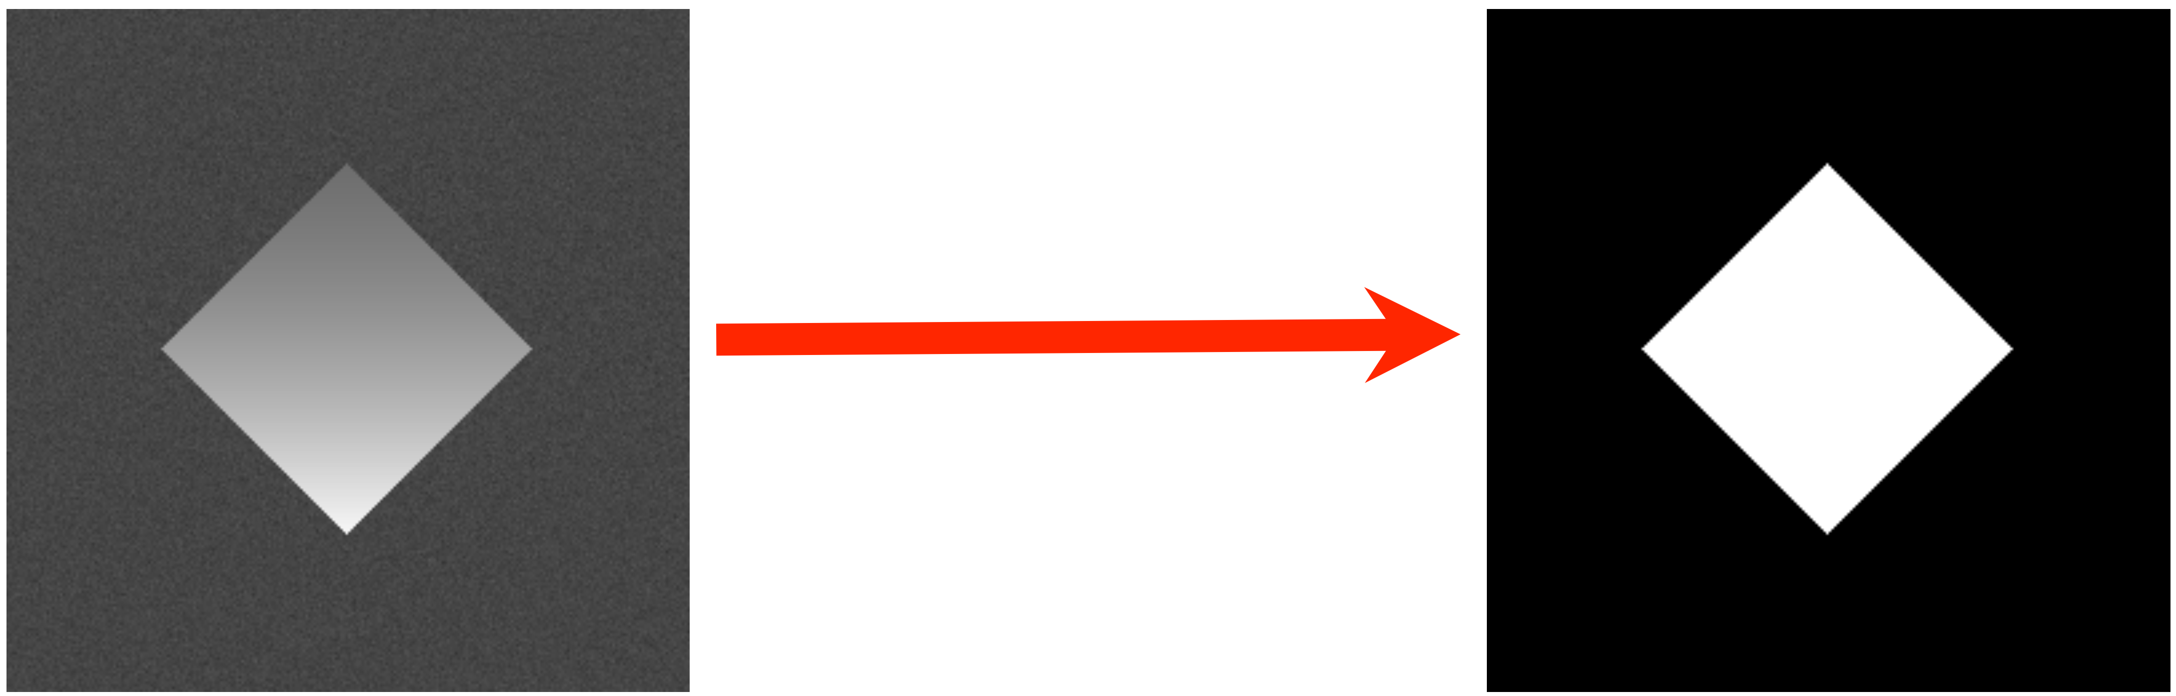

Input image f(x,y)
intensities 0-255

Segmentation output g(x,y)
0 (background)
1 (foreground)

- How can we choose T?
  - Trial and error
  - Use the histogram of f(x,y)

Role of noise

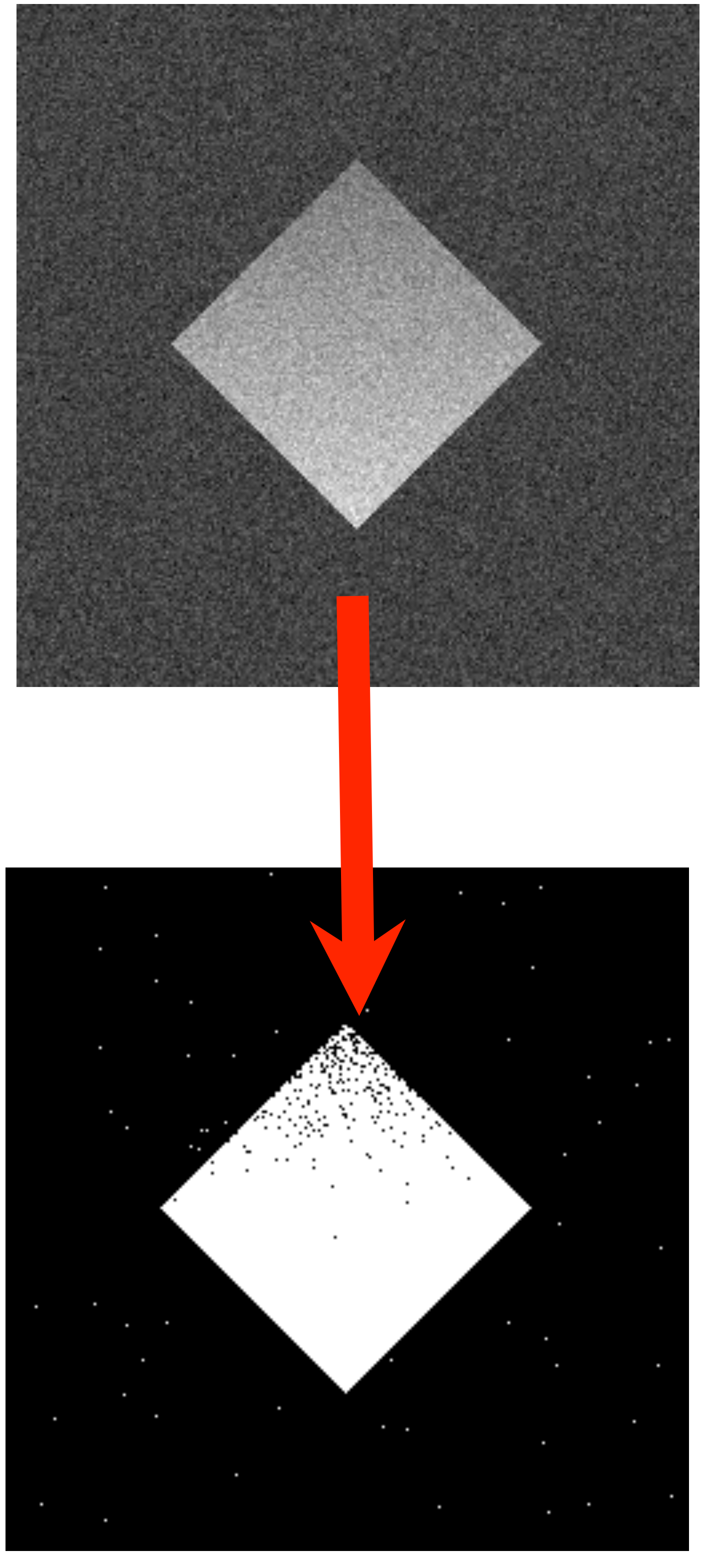

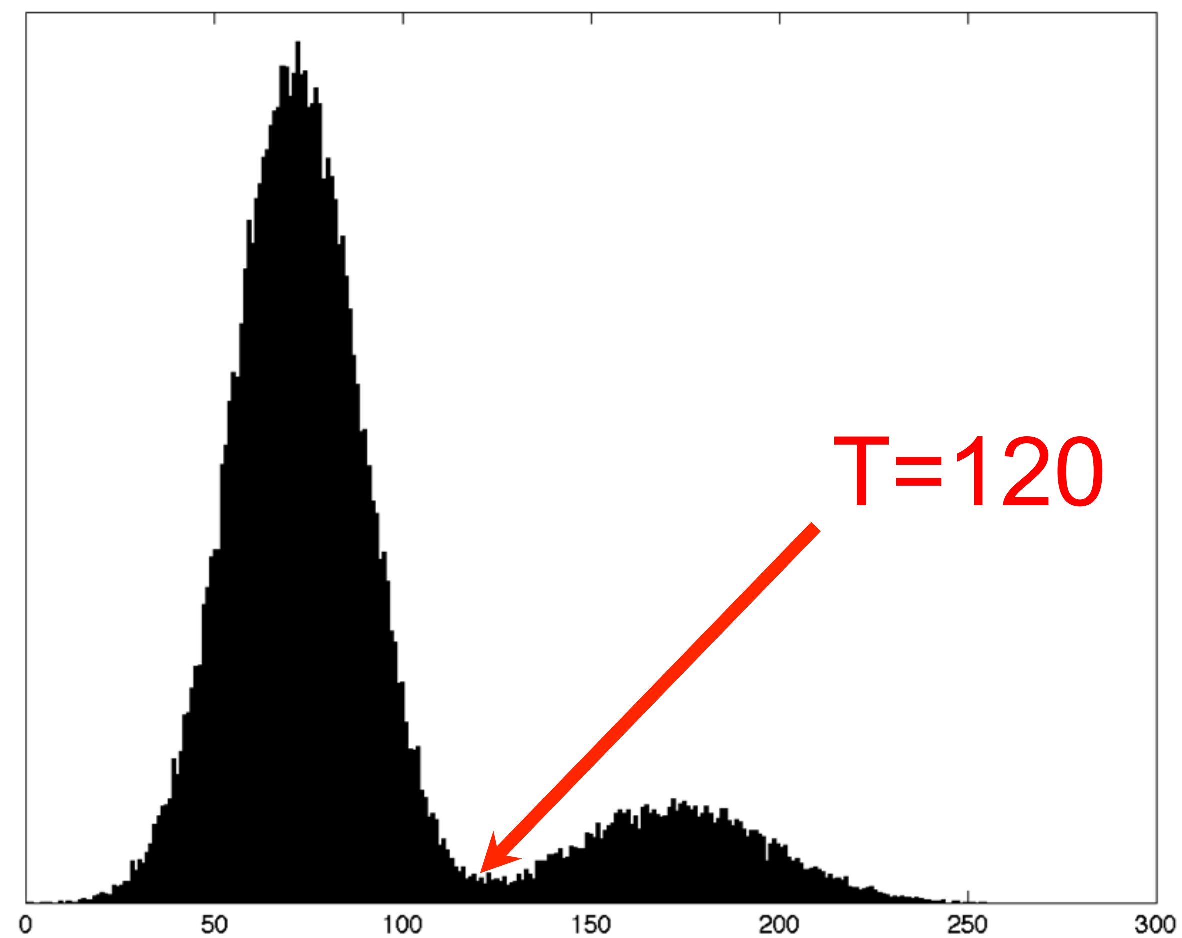

Low signal-to-noise ratio

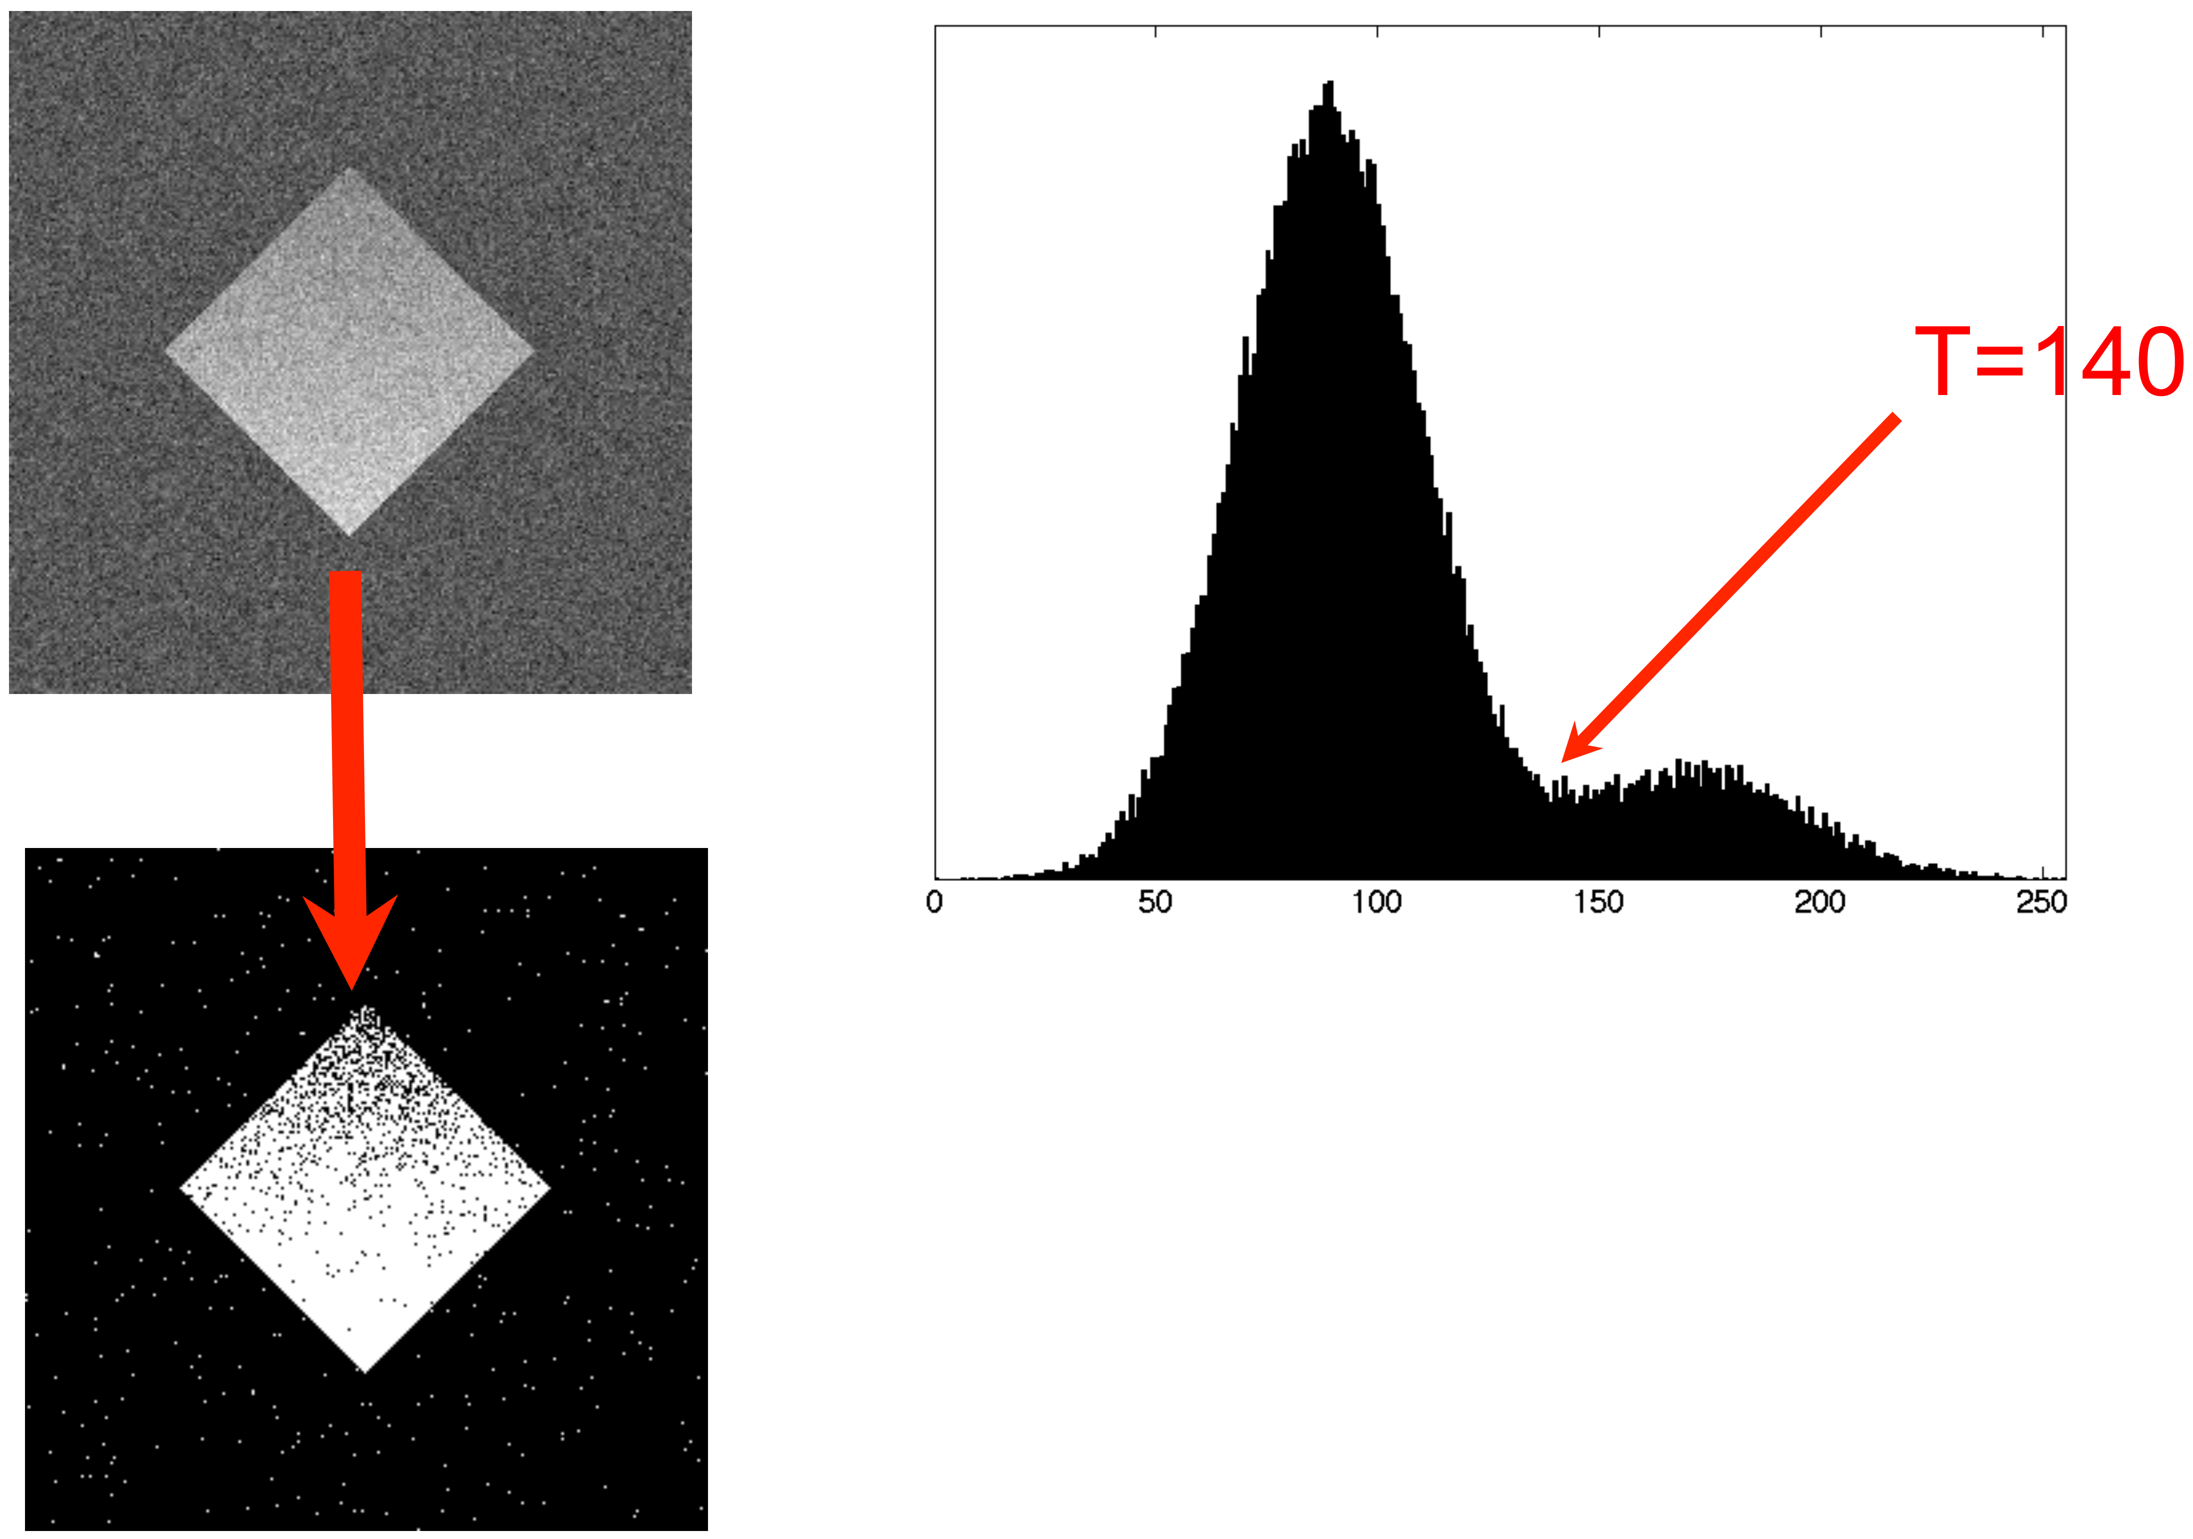

Low signal-to-noise ratio

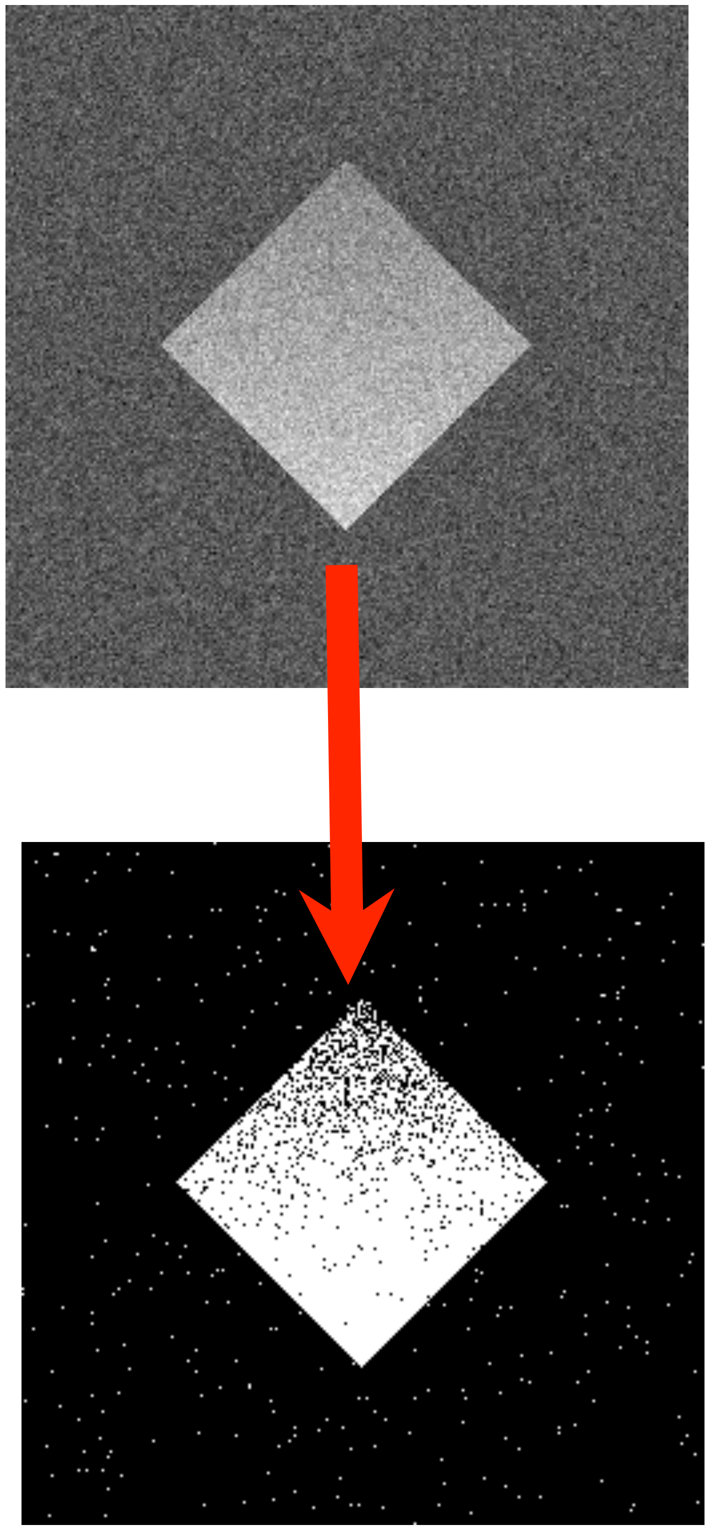

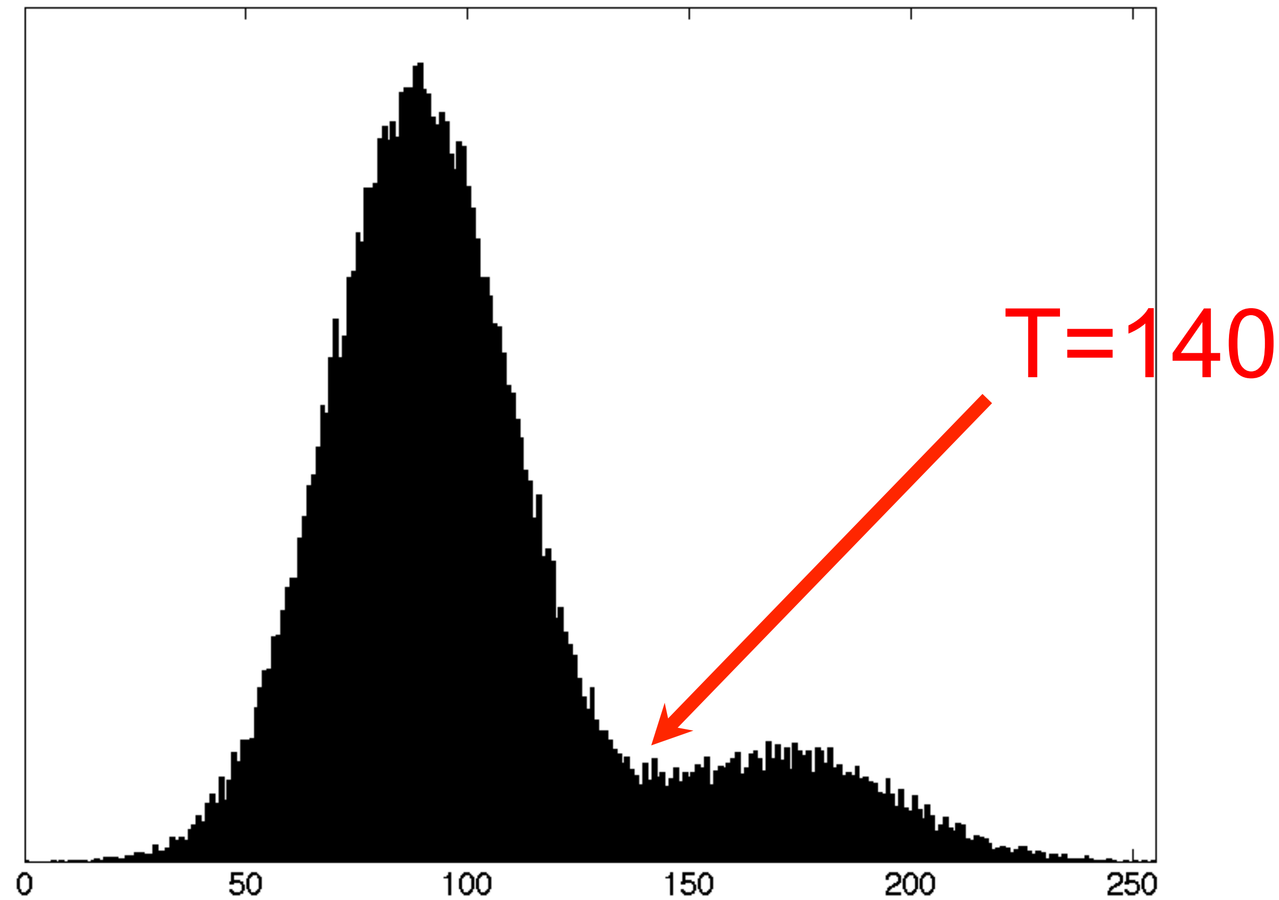

- How can we choose T?
  - Trial and error
  - Use the histogram of f(x,y)
  - Automatically
    - Otsu’s method

Effect of illumination on image histogram

Images

Histograms

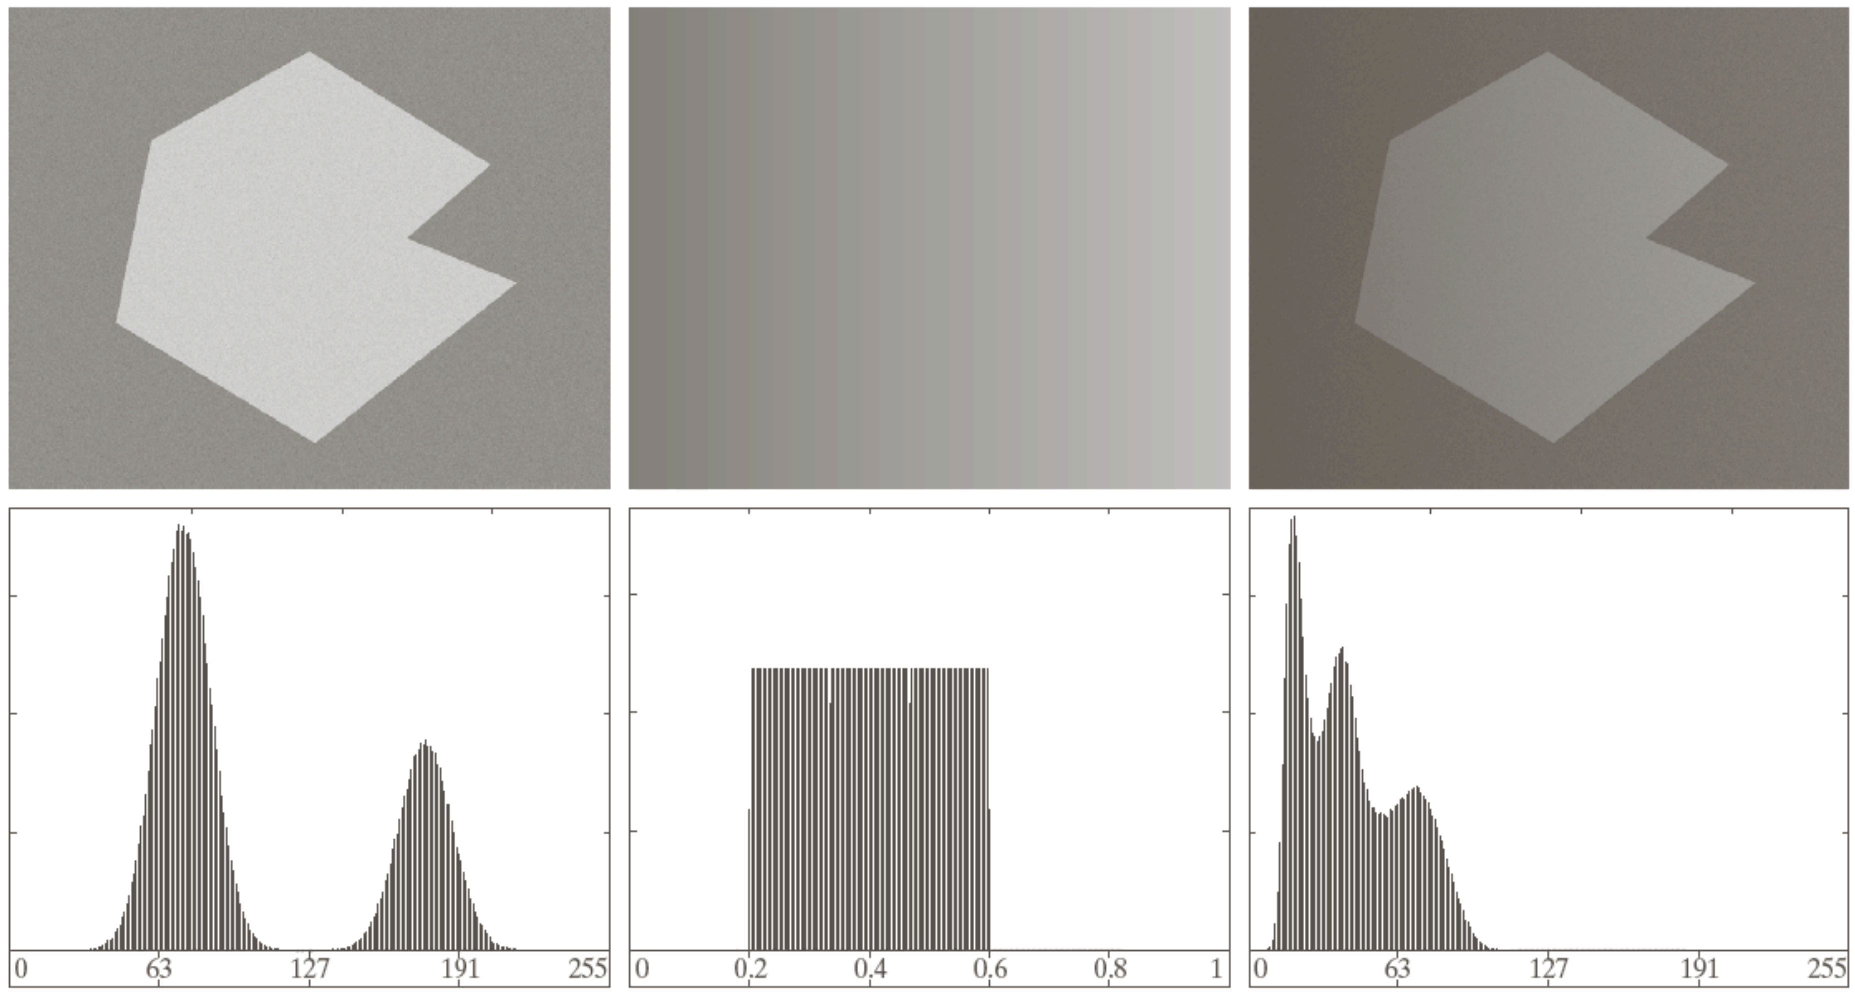

f x g = h

Original image

Illumination image

Final image

Some applications

- Microscopy image showing bright cells on a dark background
  - Find number of cells
- Time-lapse microscopy images with lighting varying over time
  - Track number of cells over time

Local thresholding

- One simple way
  - Subdivide image into blocks
  - Assumes every block has a portion of foreground and background

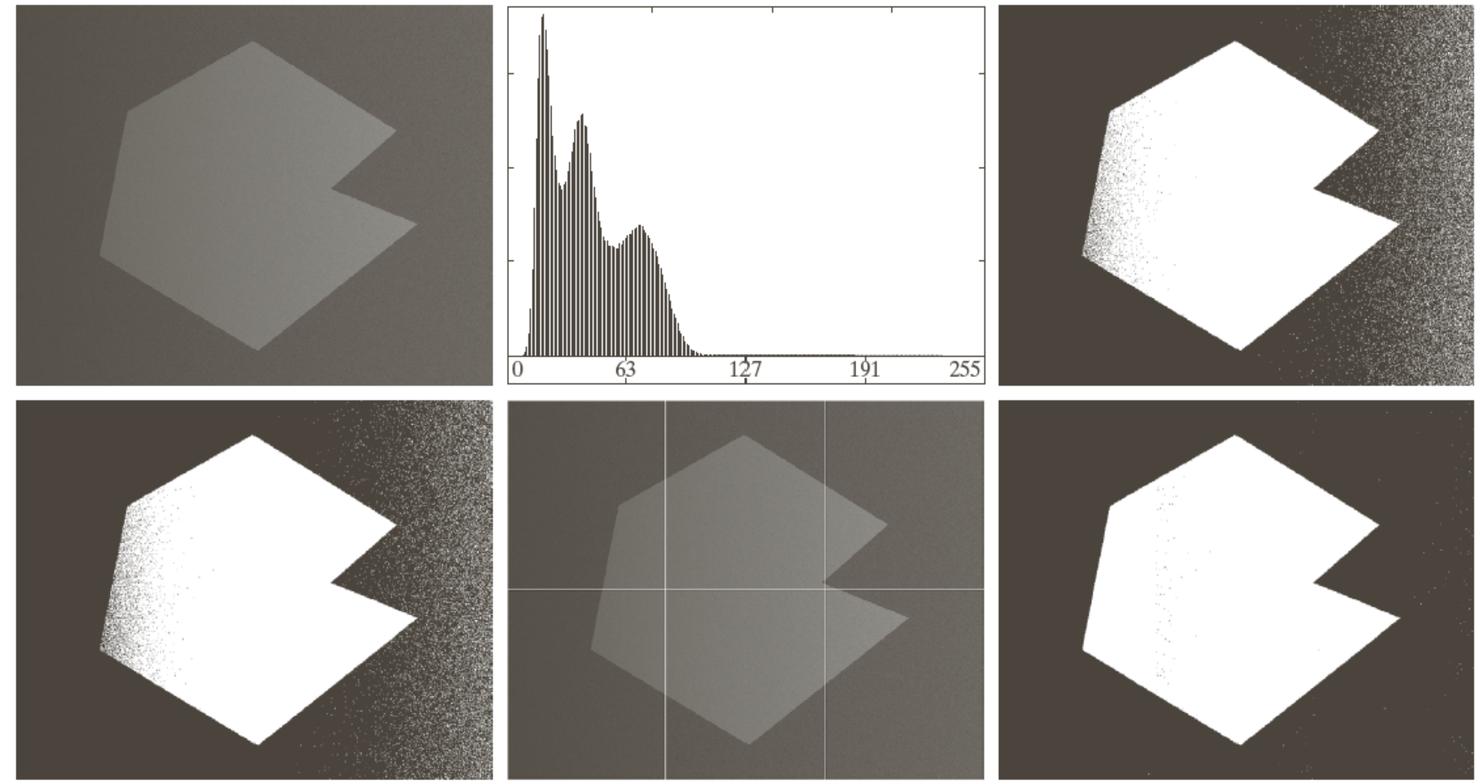

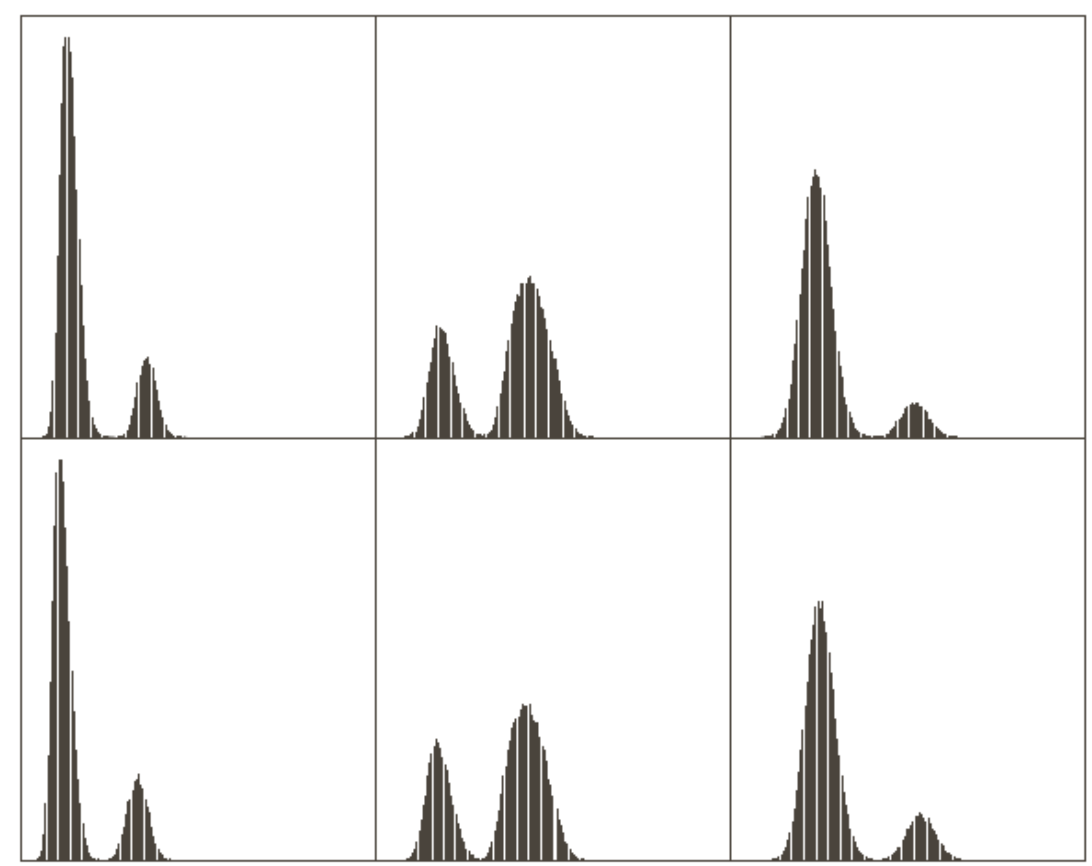

Moving mean

- At every pixel (x,y) we can choose a threshold based on the mean m(x,y) of a local window
- This is very useful for adapting to changes in illumination
  - Can be problematic if the window at (x,y) contains only foreground or only background

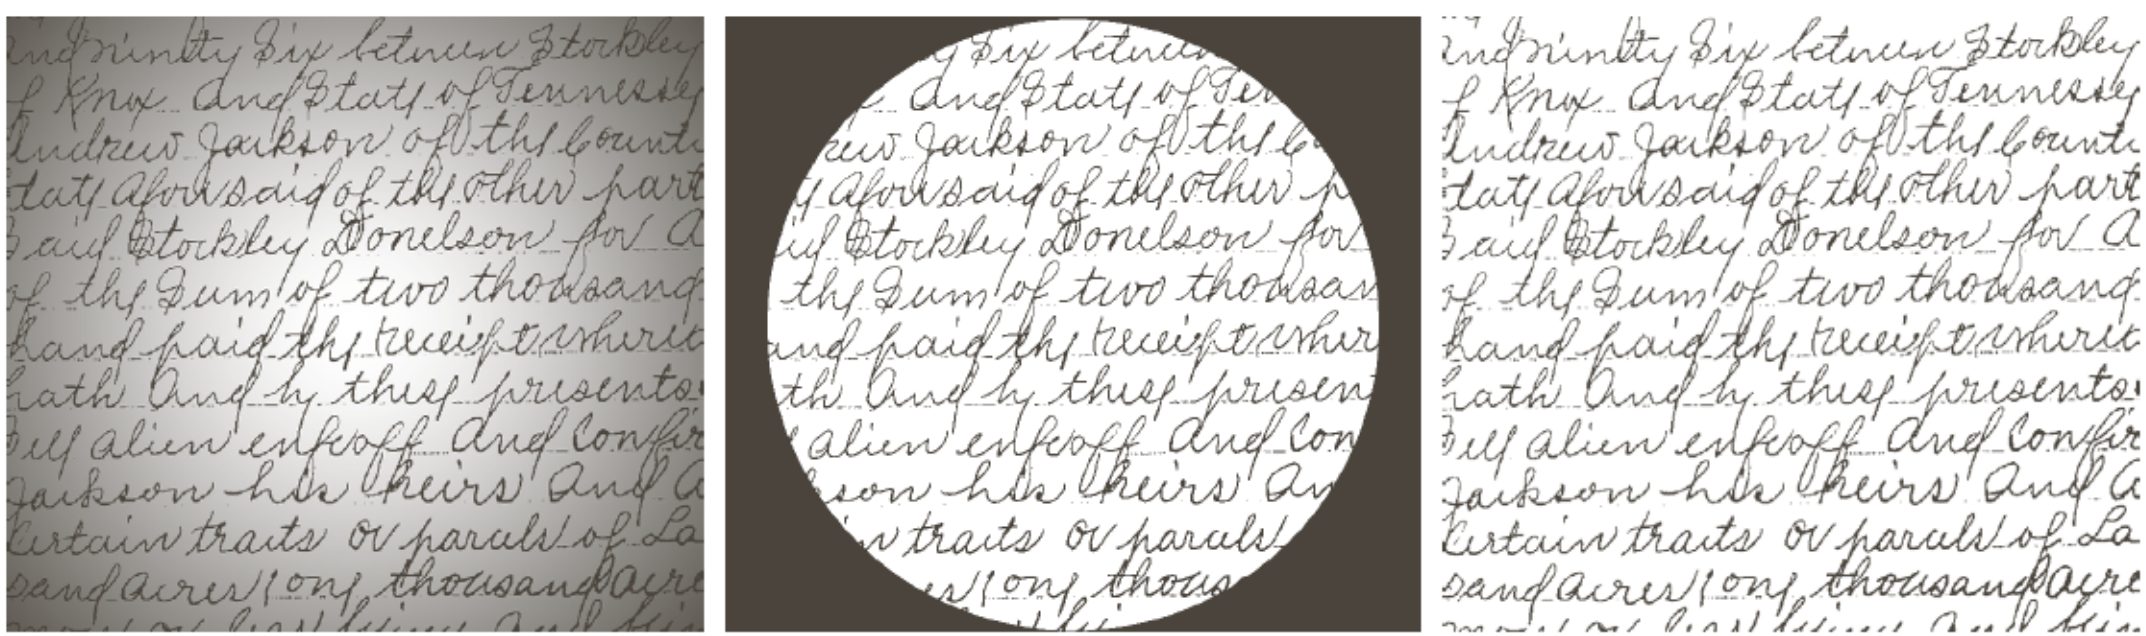

Moving mean - another example

- At every pixel (x,y) we can choose a threshold based on the mean m(x,y) of a local window
- This is very useful for adapting to changes in illumination
  - Can be problematic if the window at (x,y) contains only foreground or only background

Some Extra Things

- Gaussian/normal distribution
- Weighted means

Gaussian Distribution

- "Normal" or "bell curve"
- Two parameters
  - μ = mean, σ = standard deviation

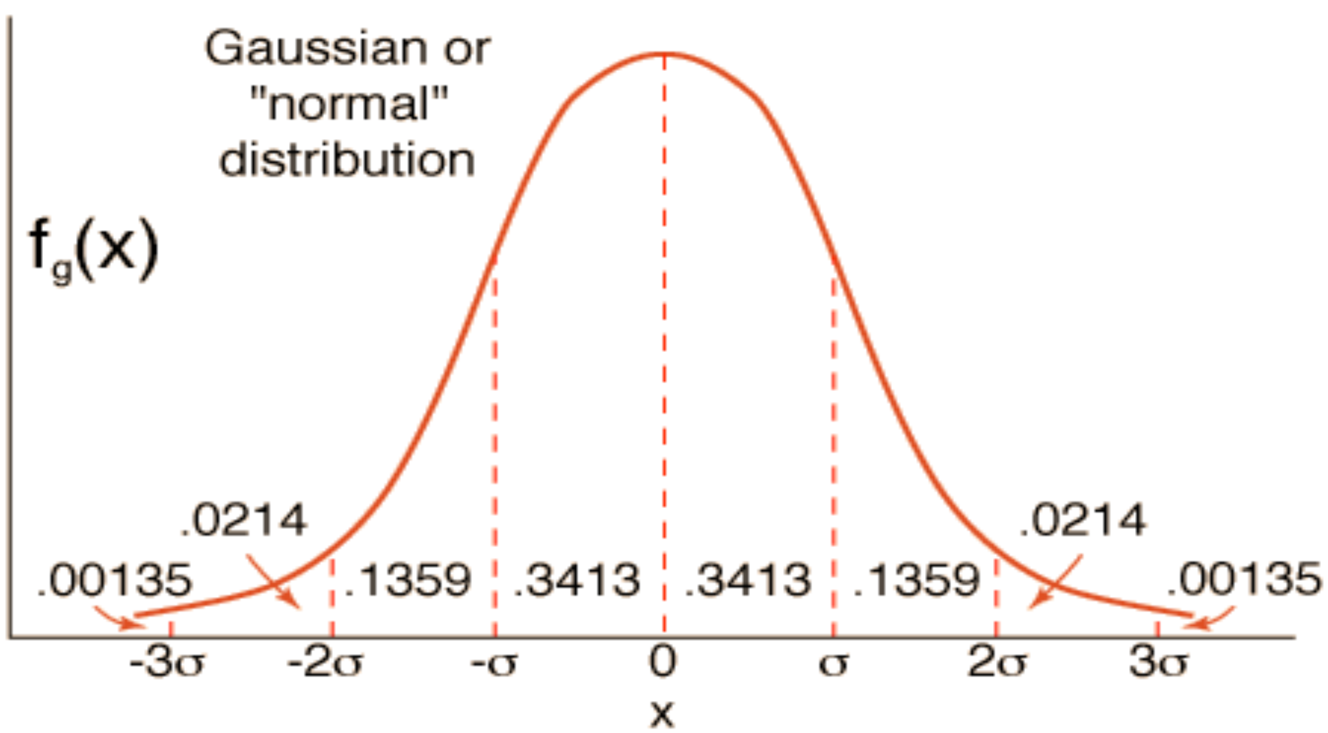

$$\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

Gaussian Properties

- Best fitting Gaussian to some data is gotten by mean and standard deviation of the samples
- Occurrence
  - Central limit theorem: mean of lots of independent & identically-distributed RVs
  - Nature (approximate)
    - Measurement error, physical characteristic, physical phenomenon
    - Diffusion of heat or chemicals

MRI Brain Tissue Segmentation

- Segment gray matter, white matter, CSF and non-brain regions
  - Manual segmentation requires expertise and is very time consuming. It is also subjective.

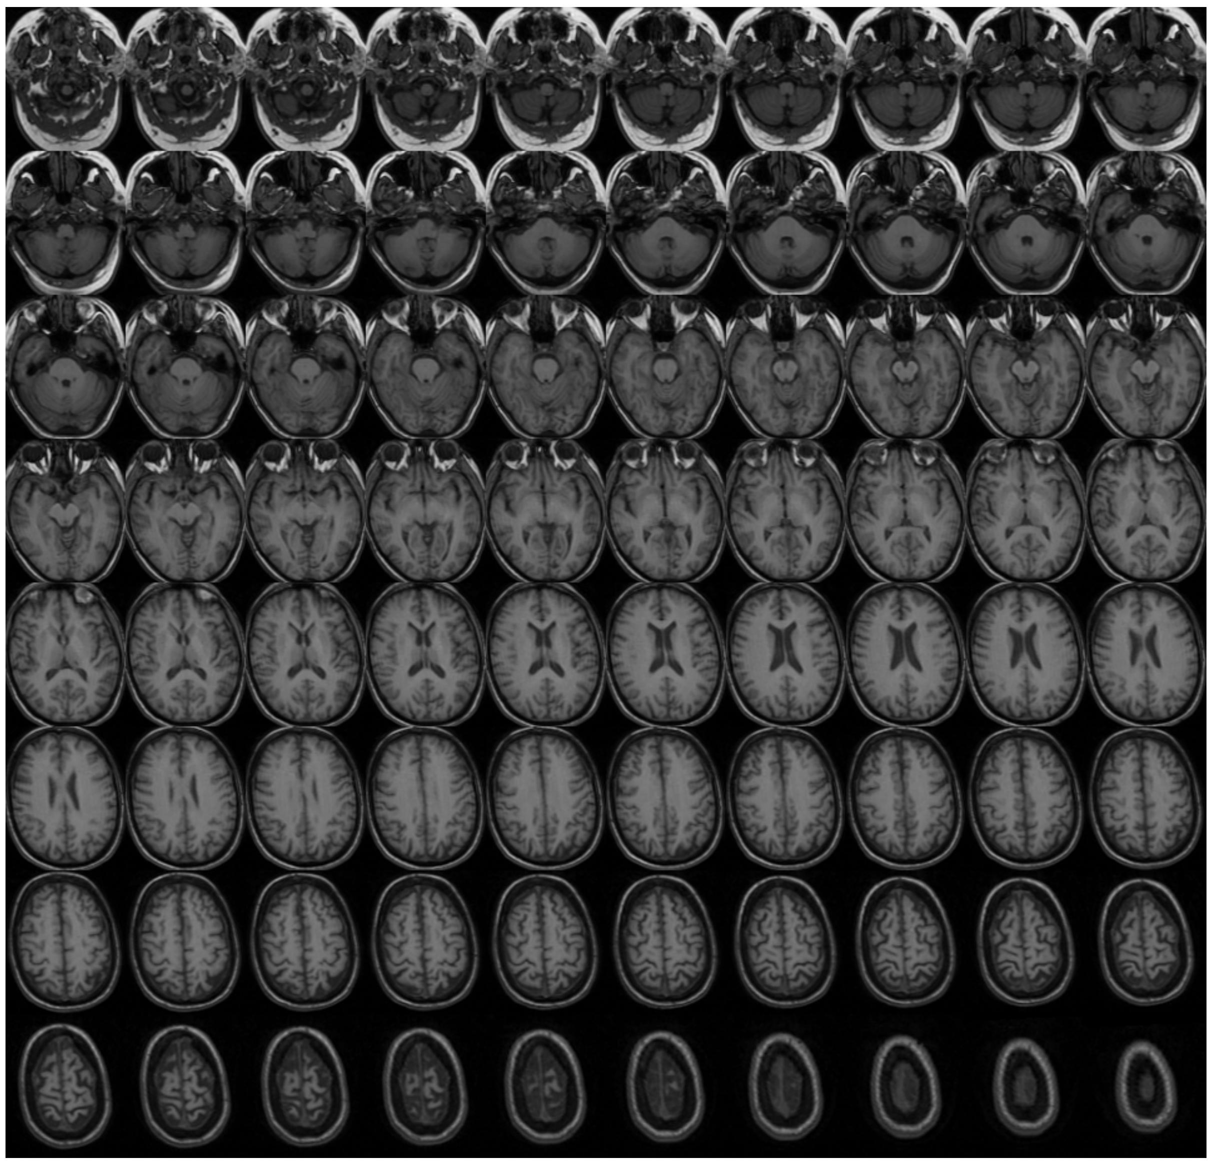

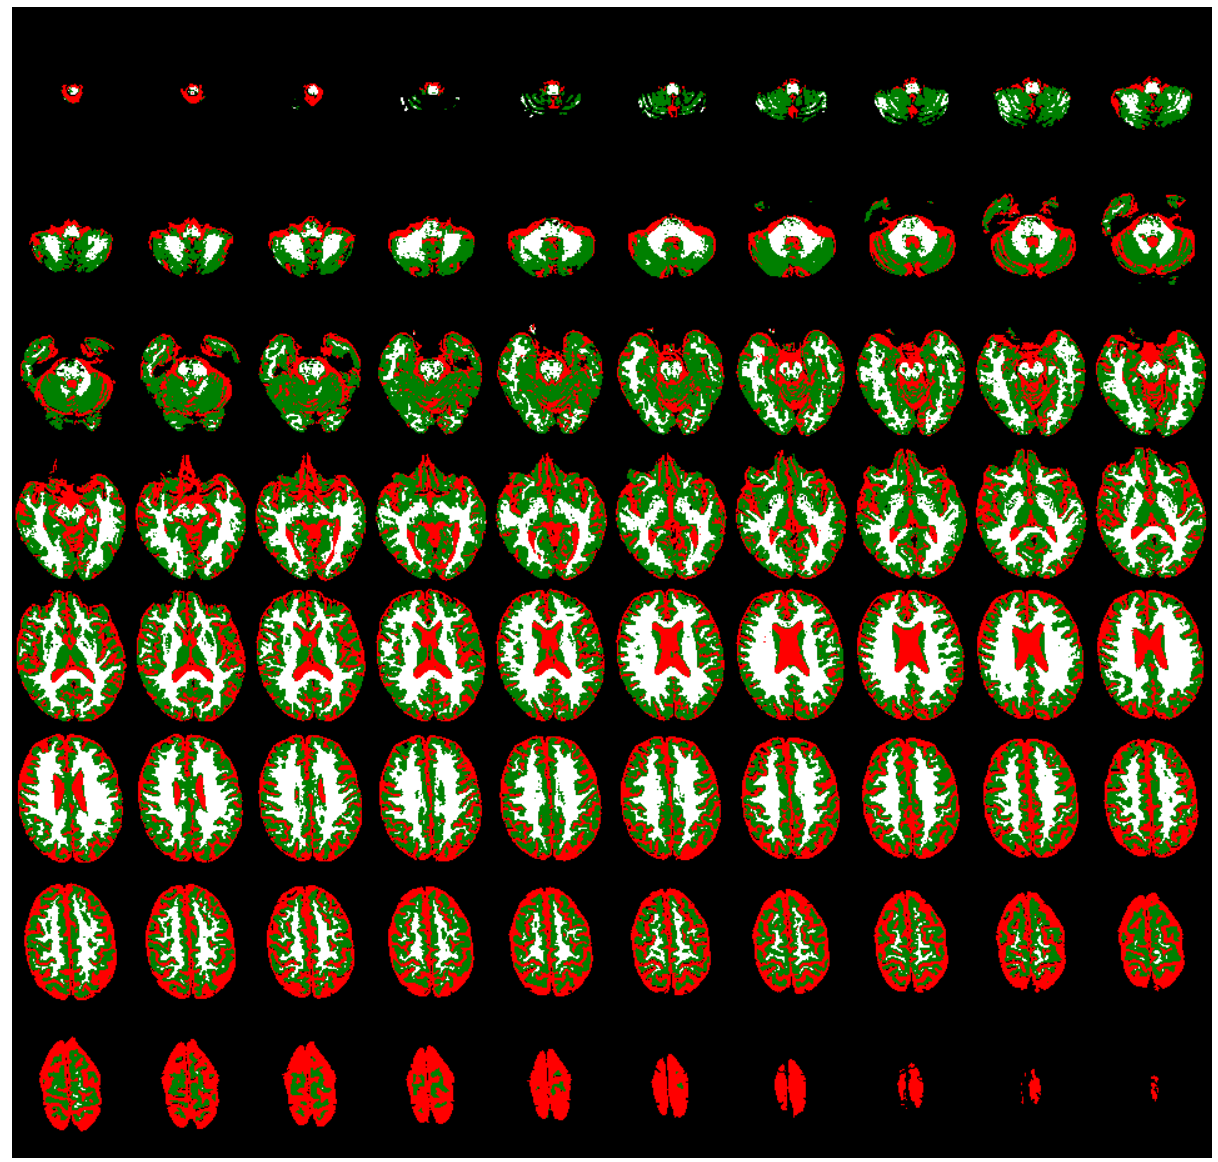

MRI Brain Example

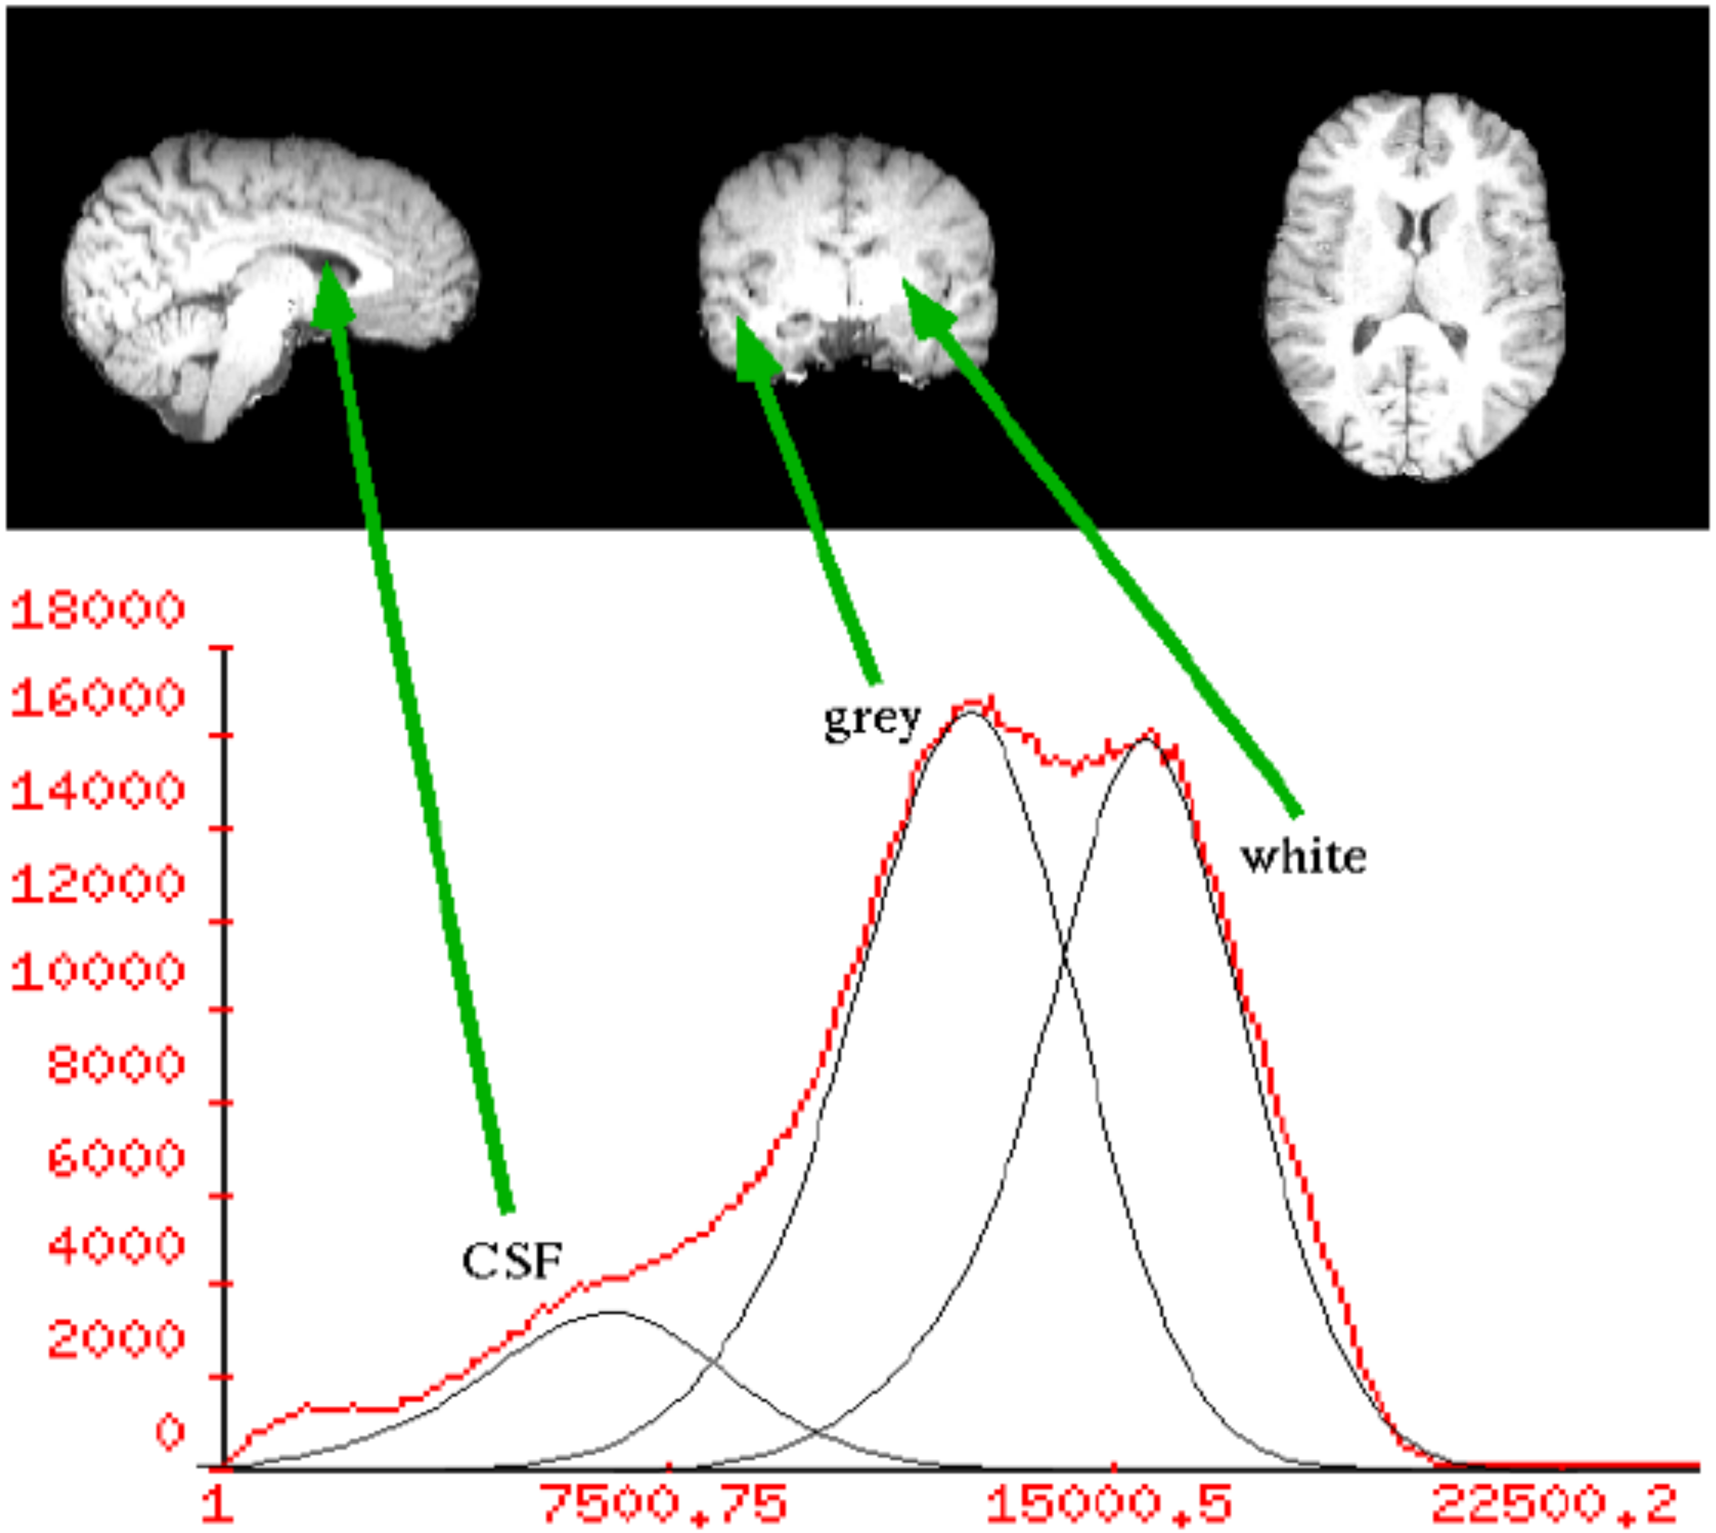

Effect of noise

- Overlap in conditional probabilities of different classes.

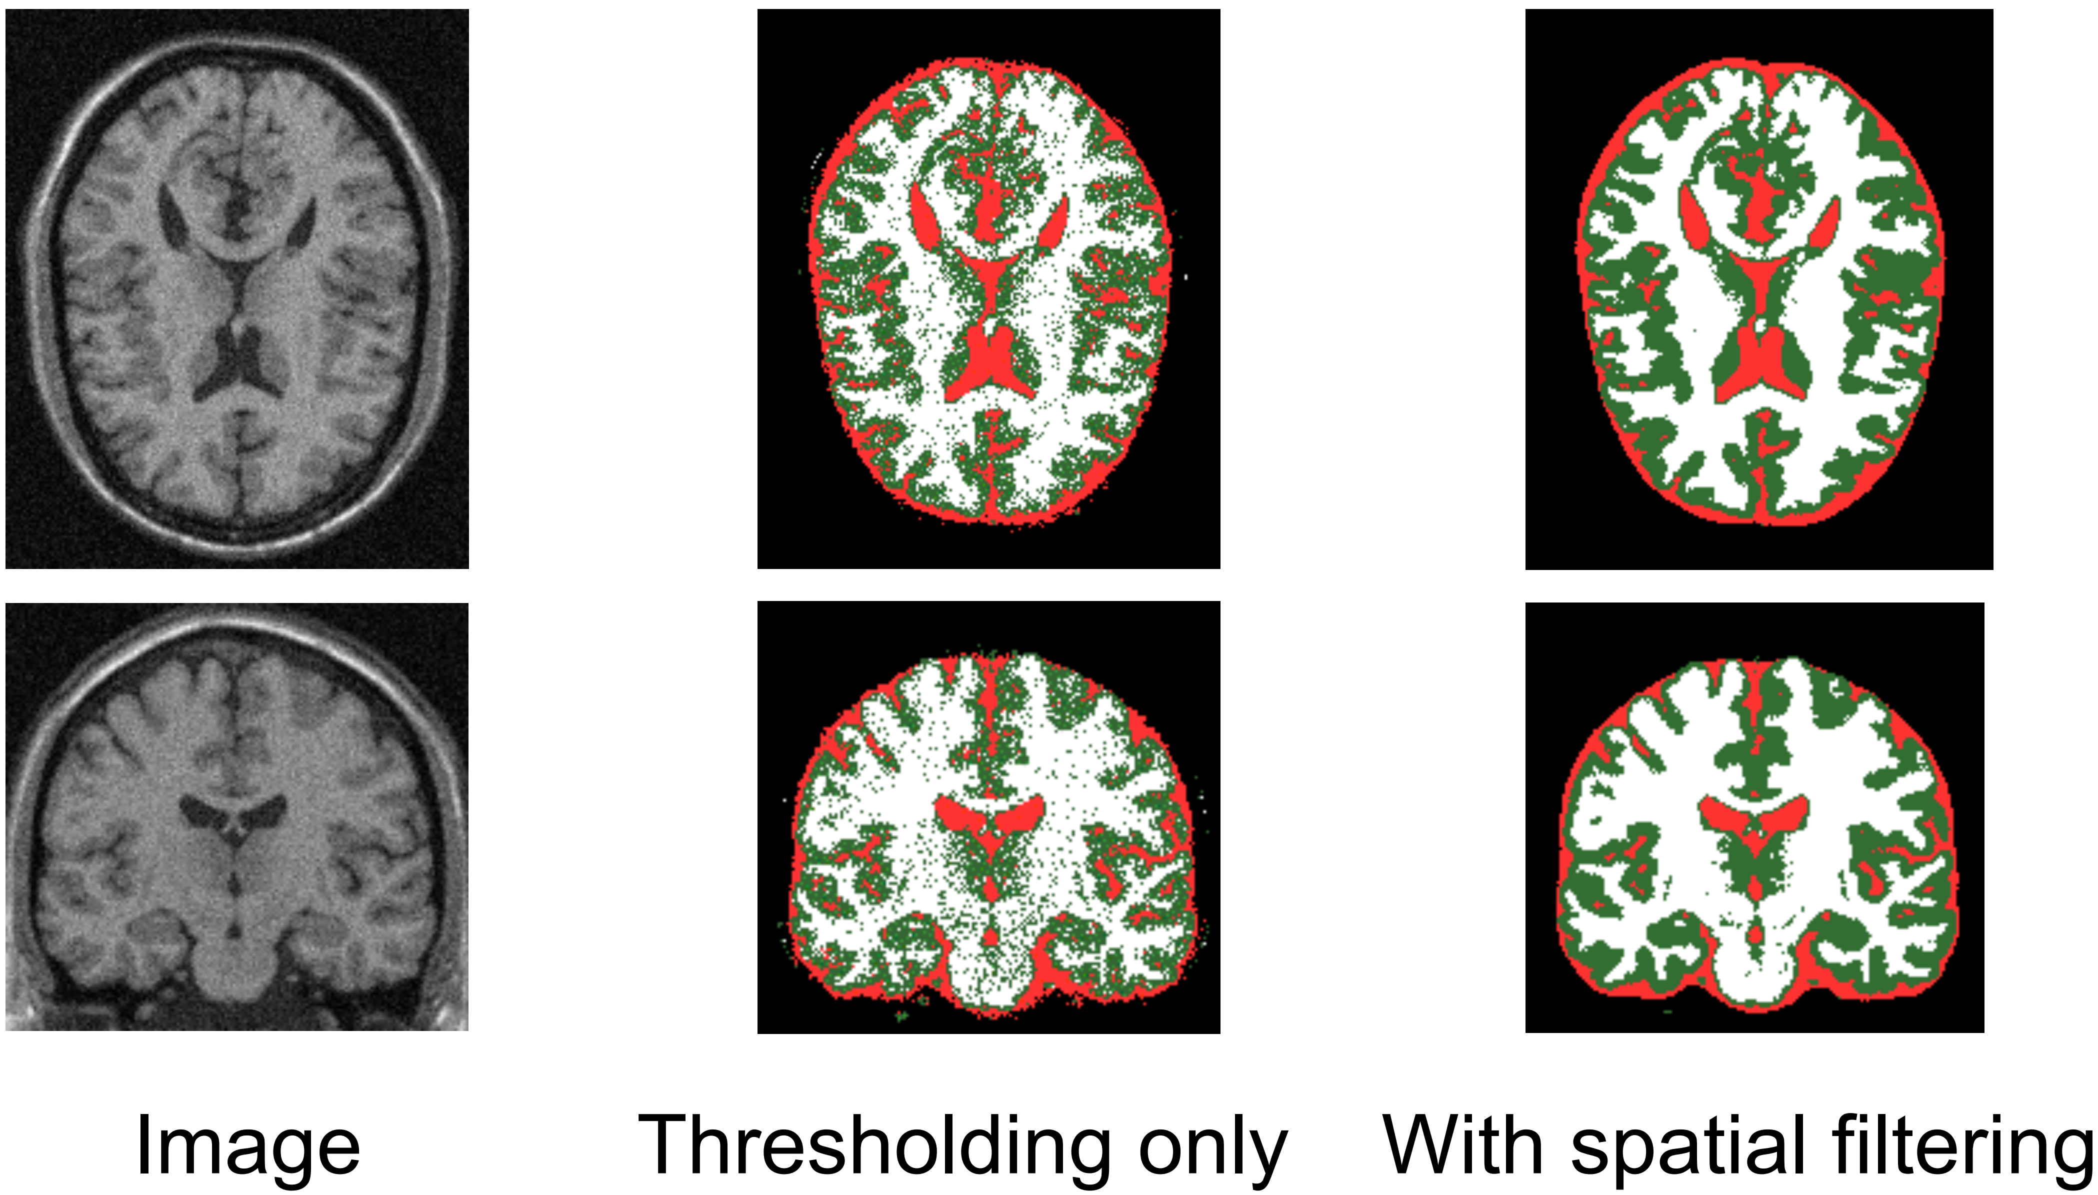

Observed image histogram

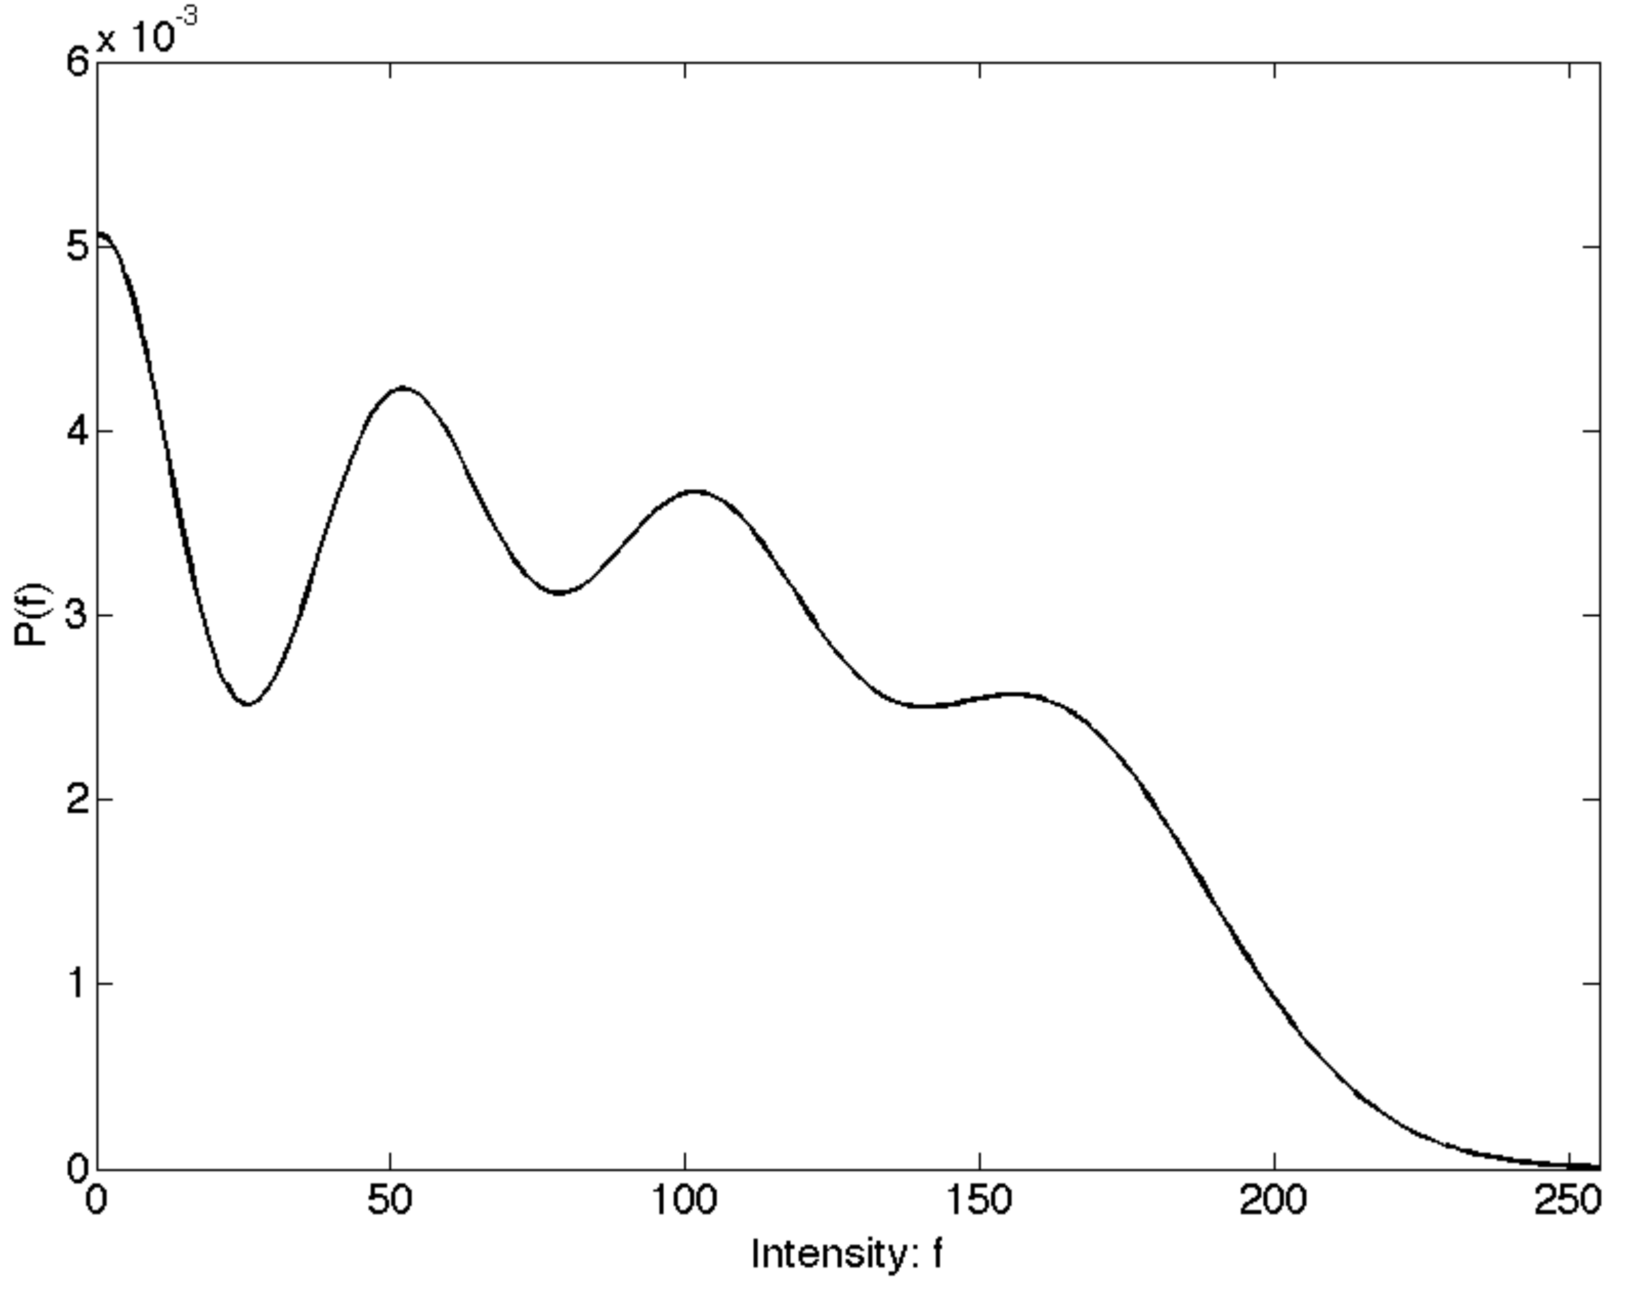

$$P(r_k) = \sum_{l \in \{csf, gm, wm, background\}} P\left(r_k \mid l\right) P(l)$$

Non-brain tissues

Large amount of overlap brain vs. non-brain

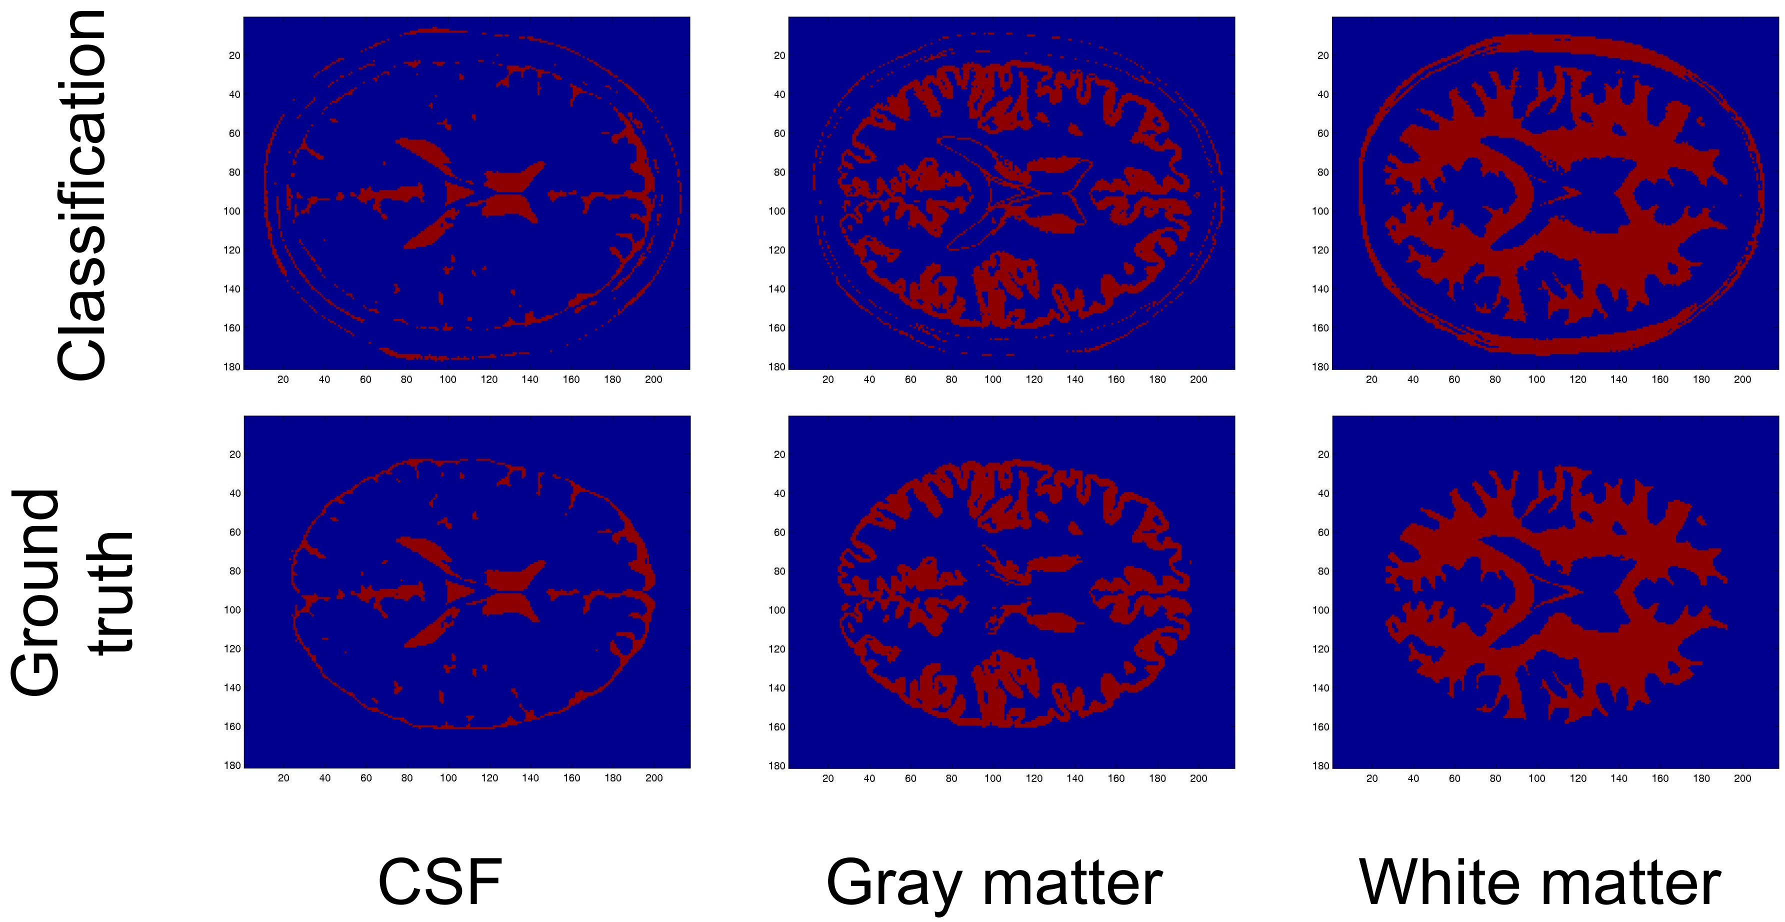

When conditional probabilities overlap this significantly, simple thresholding techniques fail. There are several possible solutions…

Multidimensional histograms

- Sometimes we have multiple images of the same object
  - Different channels in multispectral data
  - Magnetic resonance images with different pulse sequences
- Each channel provides new information
- If we have two channels, we can create a 2D histogram

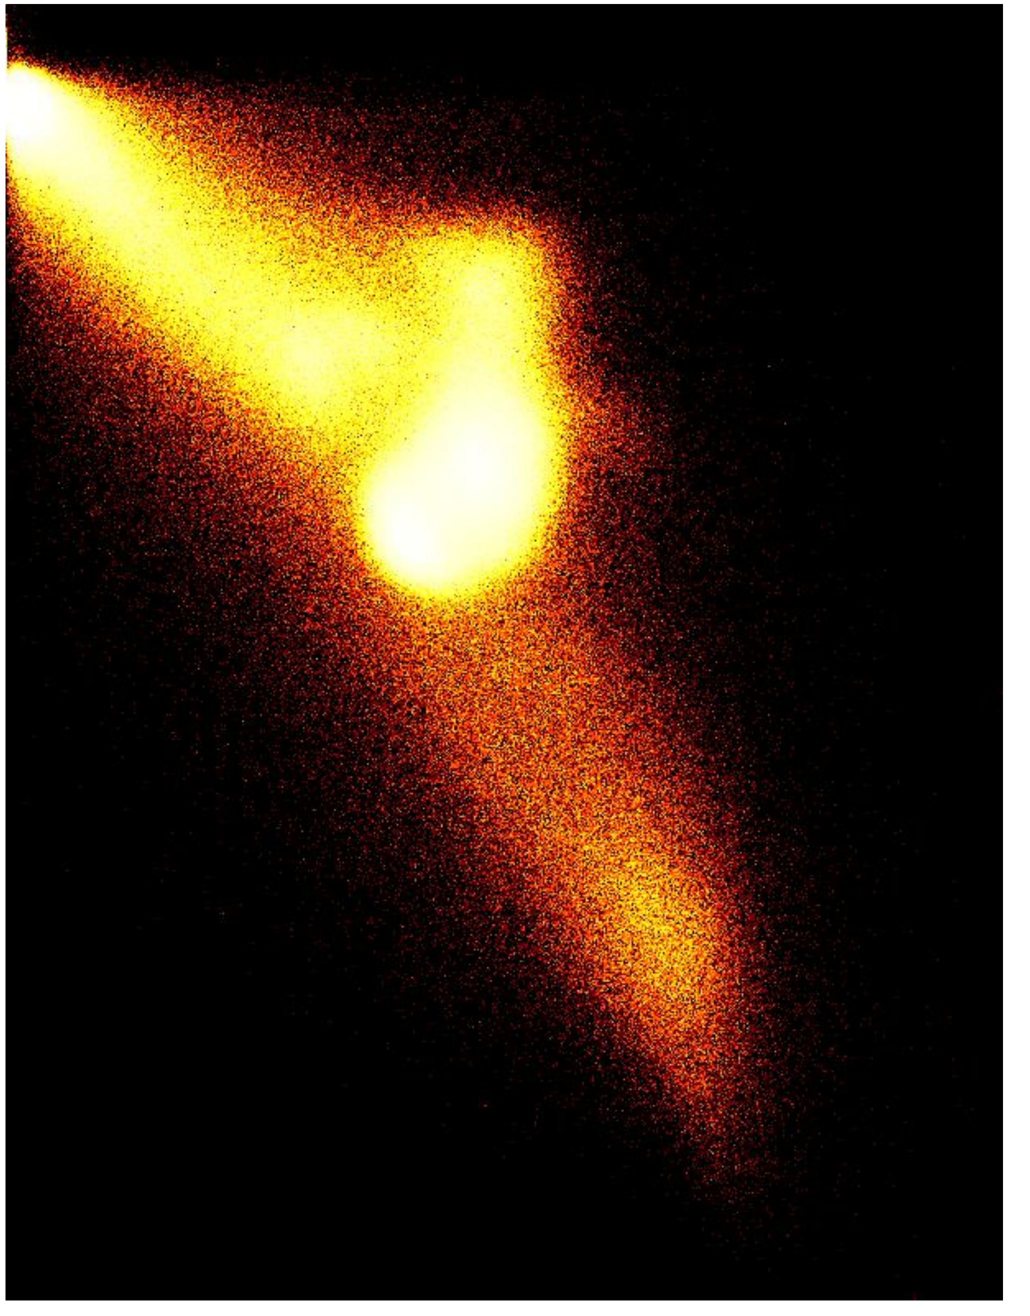

A 2D histogram of a brain image from 2 MRI channels

Non-brain tissues

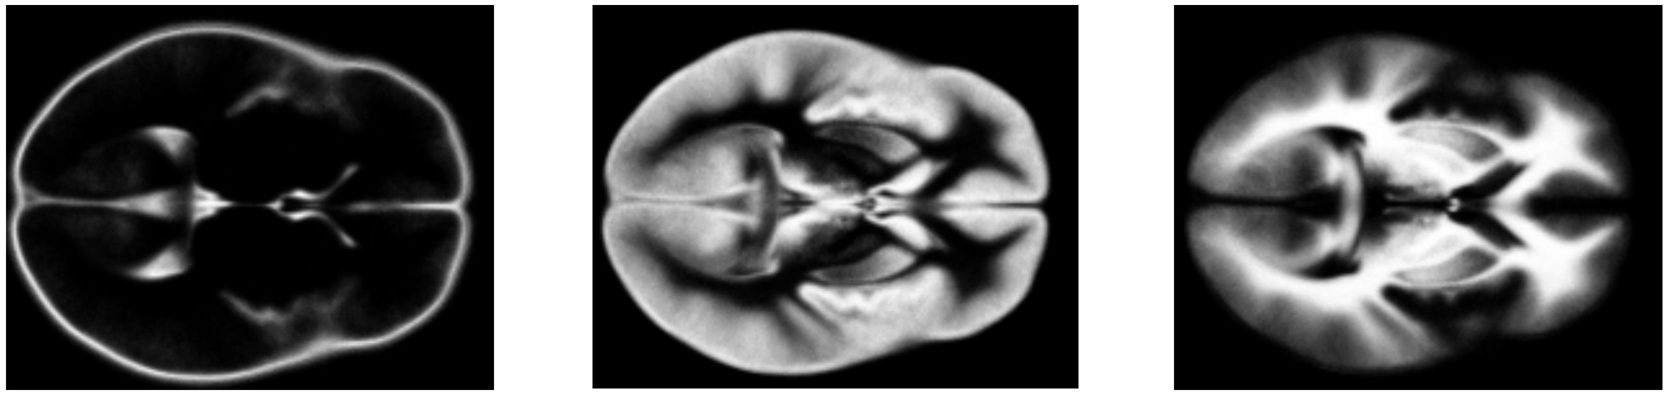

- One solution is to use atlases to guide the segmentation in addition to thresholding
- 452 subjects scanned
  - Manually segmented
  - Images aligned into common coordinate system
  - Atlas: tissue classification probability images
    - Gives a prior probability for all tissue classes at each voxel
- When we have a new MRI to segment
  - Register this atlas onto new subject's image
  - Use atlas to clarify ambiguities

Effect of non-homogeneous intensity

- Magnetic resonance images typically have a multiplicative bias field
  - This is similar to variable illumination
  - Unfortunately, the bias field can depend on the object being imaged!
  - Fortunately, it can still be estimated

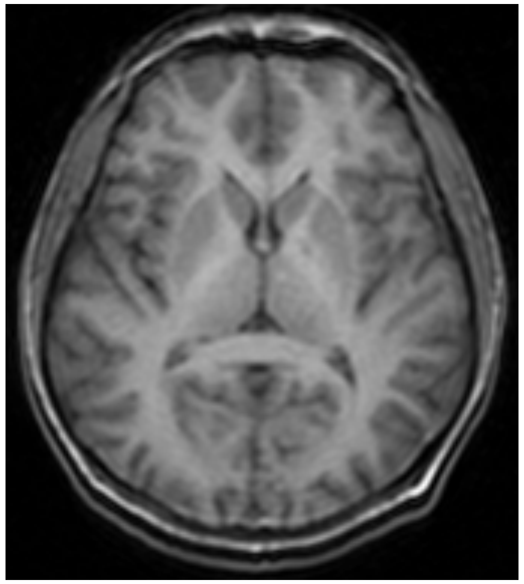

Image

Thresholding only

With bias field correction