Image processing

Subhransu Maji

CMPSCI 670: Computer Vision

September 22, 2016

Slides credit: Erik Learned-Miller and others

Image formation

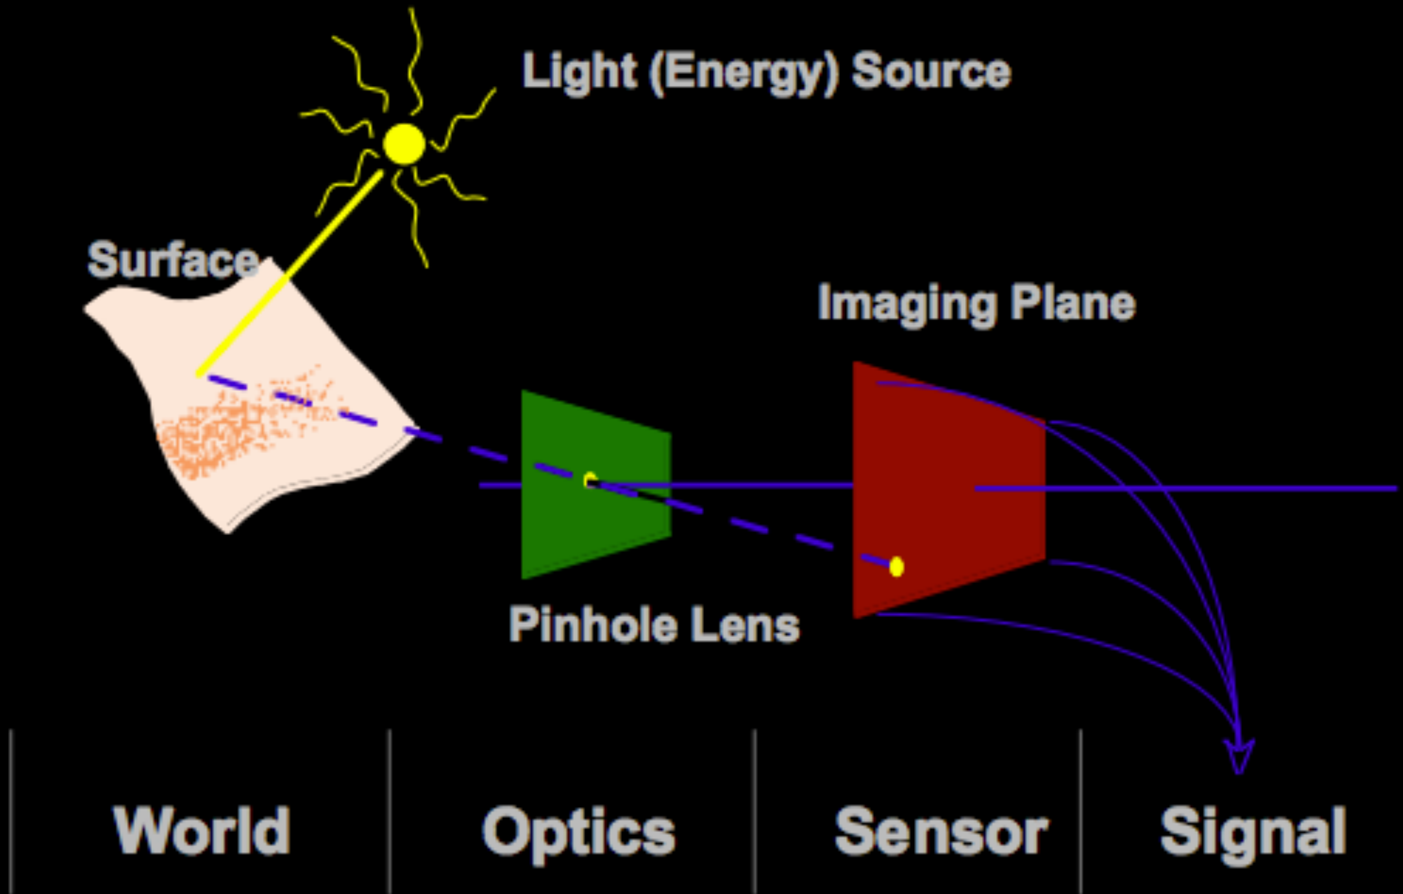

Sensor	Signal
B&W Film	Silver Density
Color Film	Silver density in three color layers
TV Camera	Electrical

Pre-digitization image

What is an image before you digitize it?

- Continuous range of wavelengths
- 2-dimensional extent
- Continuous range of power at each point

Brightness images

To simplify, consider only a brightness image

- Two-dimensional (continuous range of locations)
- Continuous range of brightness values

This is equivalent to a two-dimensional function over a plane

An image as a surface

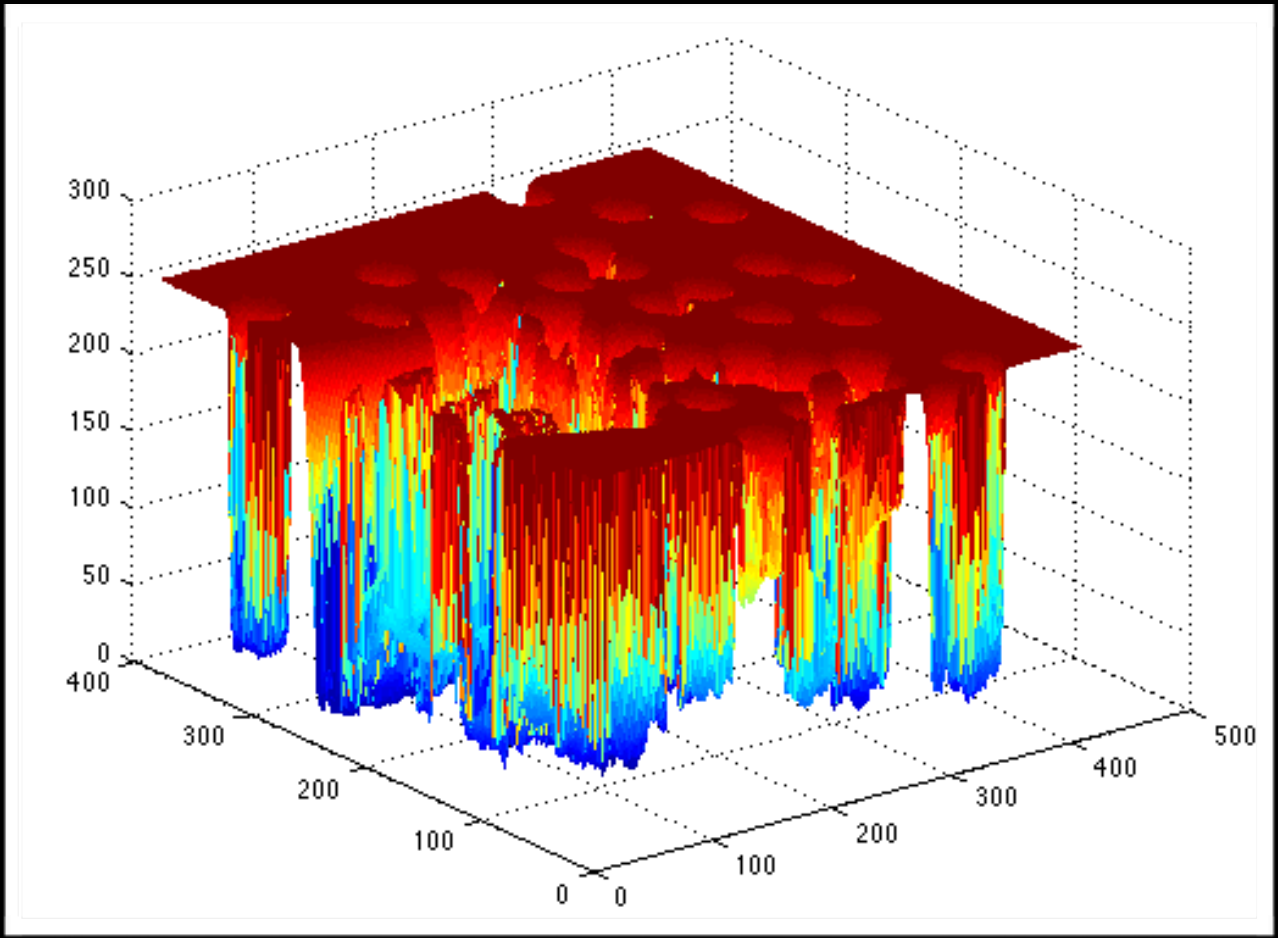

How do we represent this continuous two dimensional surface efficiently?

Discretization

Sampling strategies

- Spatial sampling
  - How many pixels?
  - What arrangement of pixels?
- Brightness sampling
  - How many brightness values?
  - Spacing of brightness values?
- For video, also the question of time sampling.

Signal quantization

Goal: determine a mapping from a continuous signal (e.g. analog video signal) to one of K discrete (digital) levels.

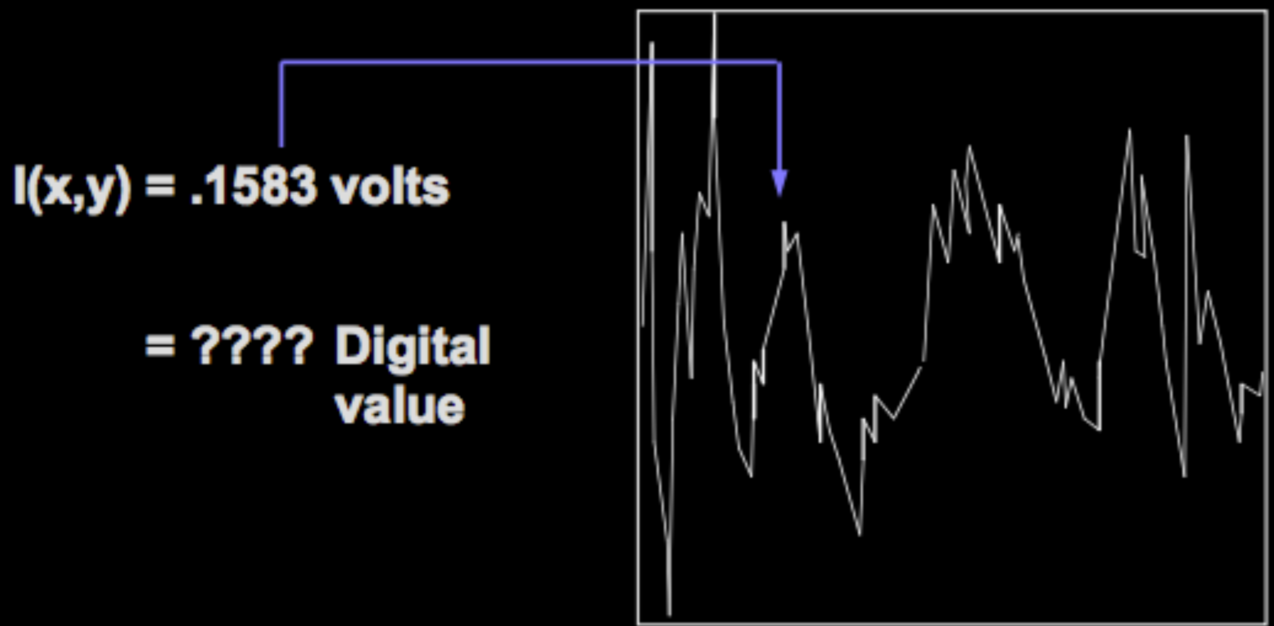

Quantization

$I(x,y)$ = continuous signal: $0 \leq I \leq M$

Want to quantize to K values 0,1,....K-1

K usually chosen to be a power of 2:

K: #Levels	#Bits
2	1
4	2
8	3
16	4
32	5
64	6
128	7
256	8

Mapping from input signal to output signal is to be determined.

Several types of mappings: uniform, logarithmic, etc.

Choice of K

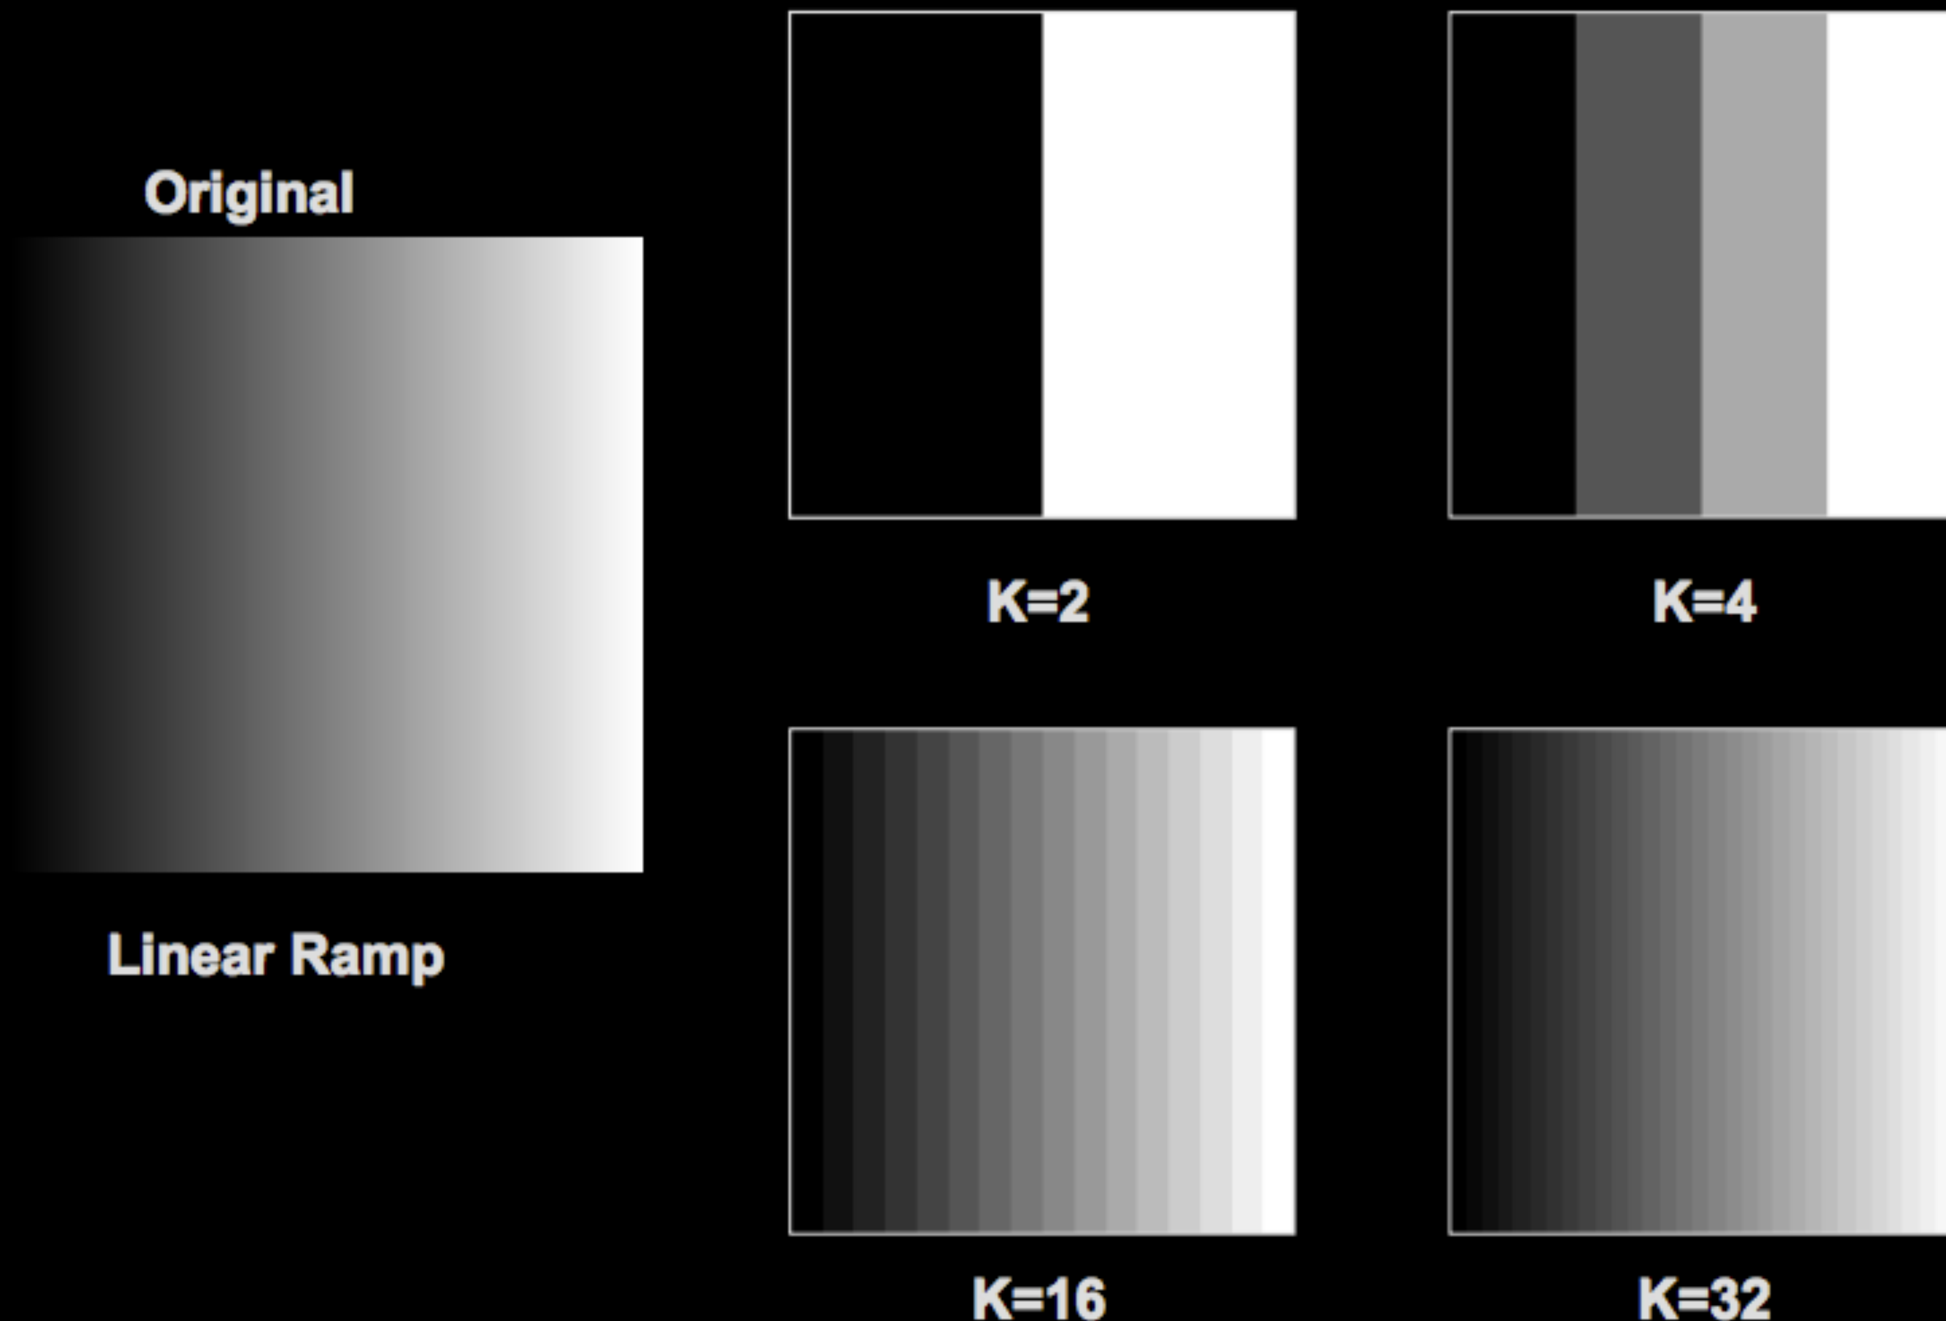

Choice of K

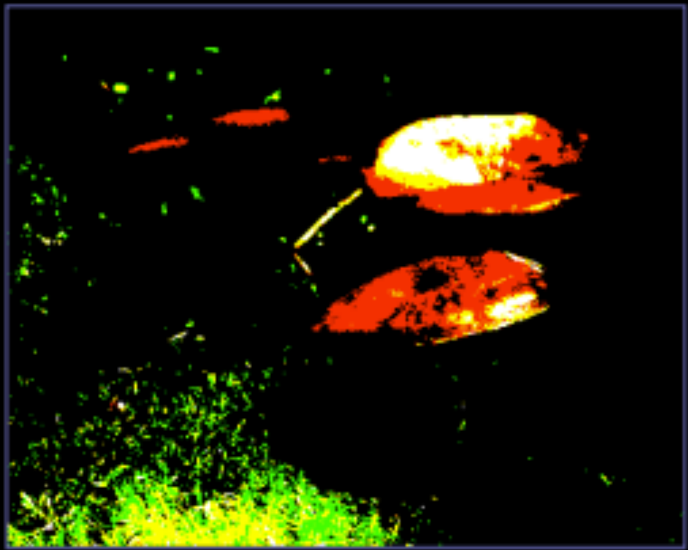

K=2 (each color)

K=4 (each color)

False contours problem

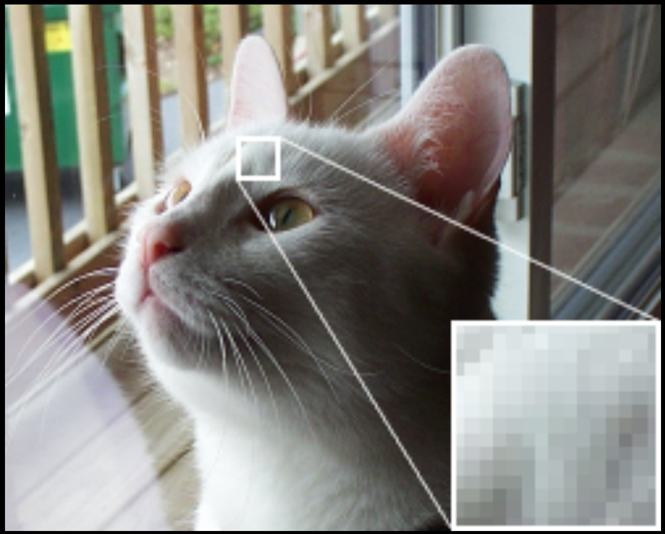

original image

16 colors

with random noise

"Dithering" adds random noise to reduce false contours

https://en.wikipedia.org/wiki/Dither

Choice of the function: uniform

Uniform sampling divides the signal range [0-M] into K equal-sized intervals.

The integers 0,...K-1 are assigned to these intervals.

All signal values within an interval are represented by the associated integer value.

Defines a mapping:

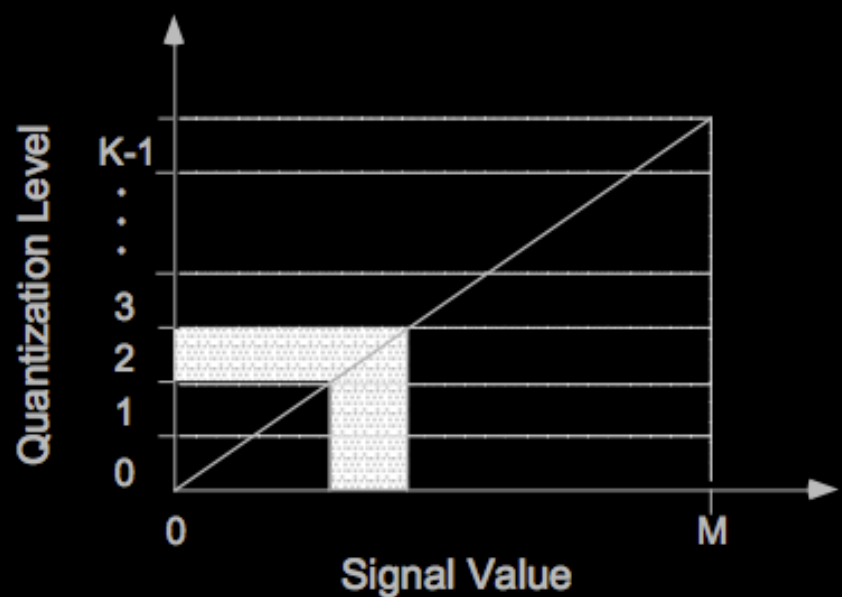

Logarithmic quantization

Signal is: log I(x,y)

Effect is:

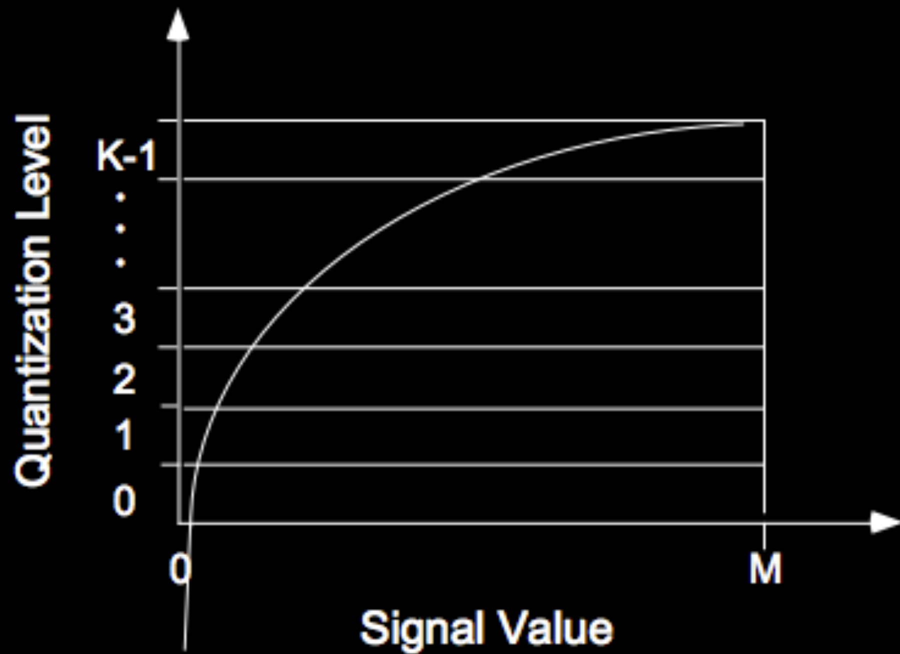

Detail enhanced in the low signal values at expense of detail in high signal values.

Logarithmic quantization

Quantization Curve

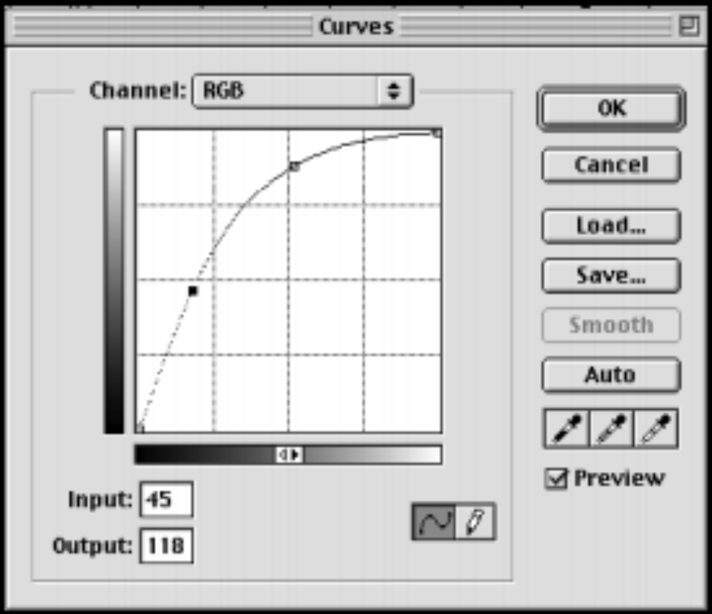

Color displays

Given a 24 bit color image (8 bits for R, G, B)

- Turn on 3 subpixels with power proportional to RGB values

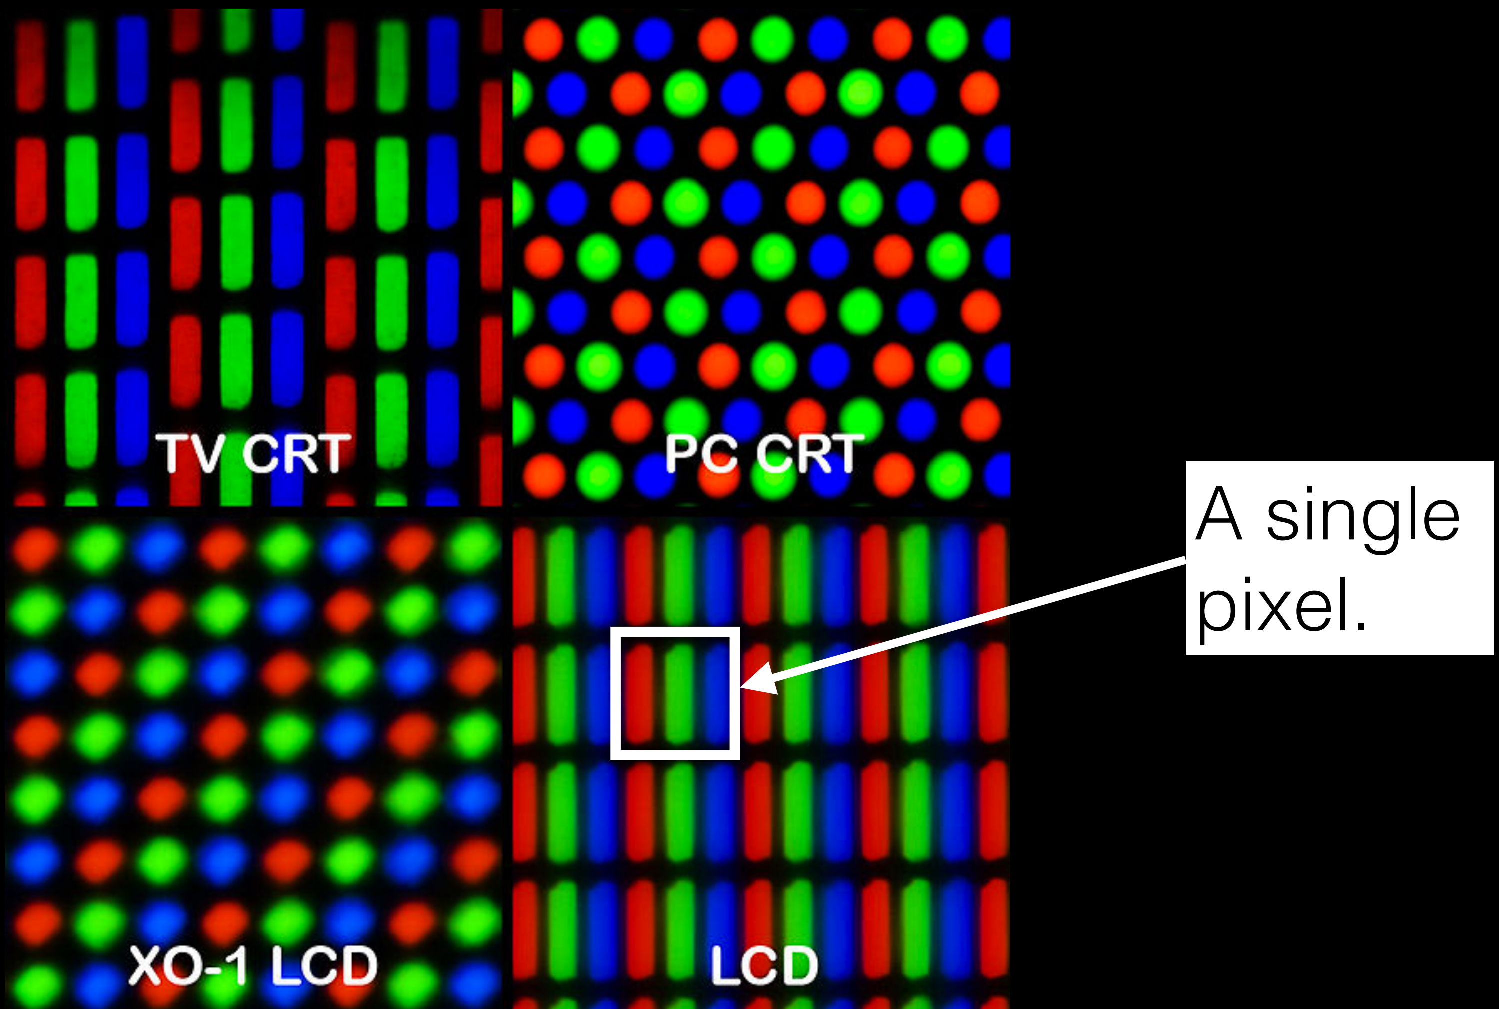

https://en.wikipedia.org/wiki/File:Pixel_geometry_01_Pengo.jpg

“White” text on color display

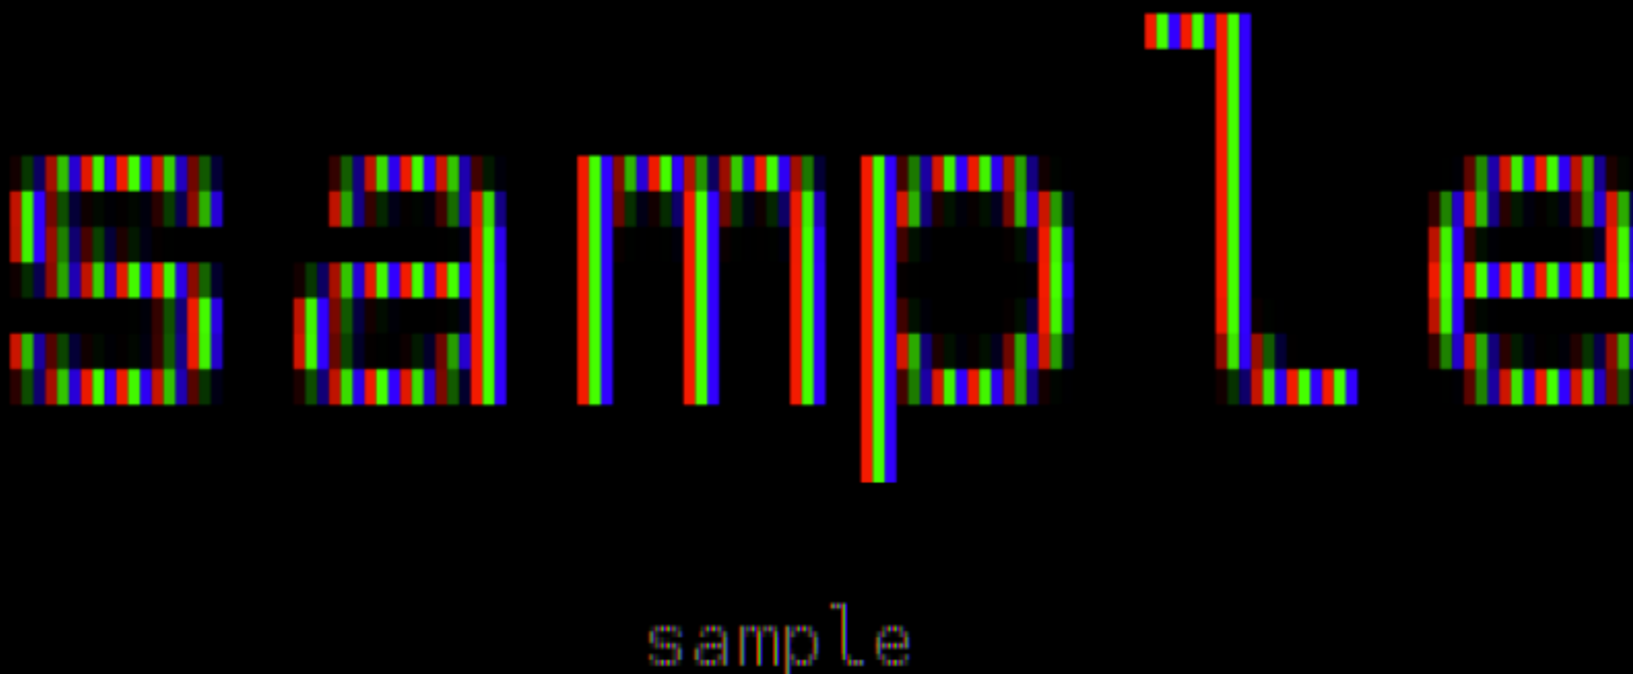

http://en.wikipedia.org/wiki/Subpixel_rendering

Lookup tables

8 bit image: 256 different values.

Simplest way to display: map each number to a gray value:

- 0 → (0.0, 0.0, 0.0) or (0,0,0)
- 1 → (0.0039, 0.0039, 0.0039) or (1,1,1)
- 2 → (0.0078, 0.0078, 0.0078) or (2,2,2)
- ...
- 255 → (1.0, 1.0, 1.0) or (255,255,255)

This is called a grayscale mapping.

Lookup tables

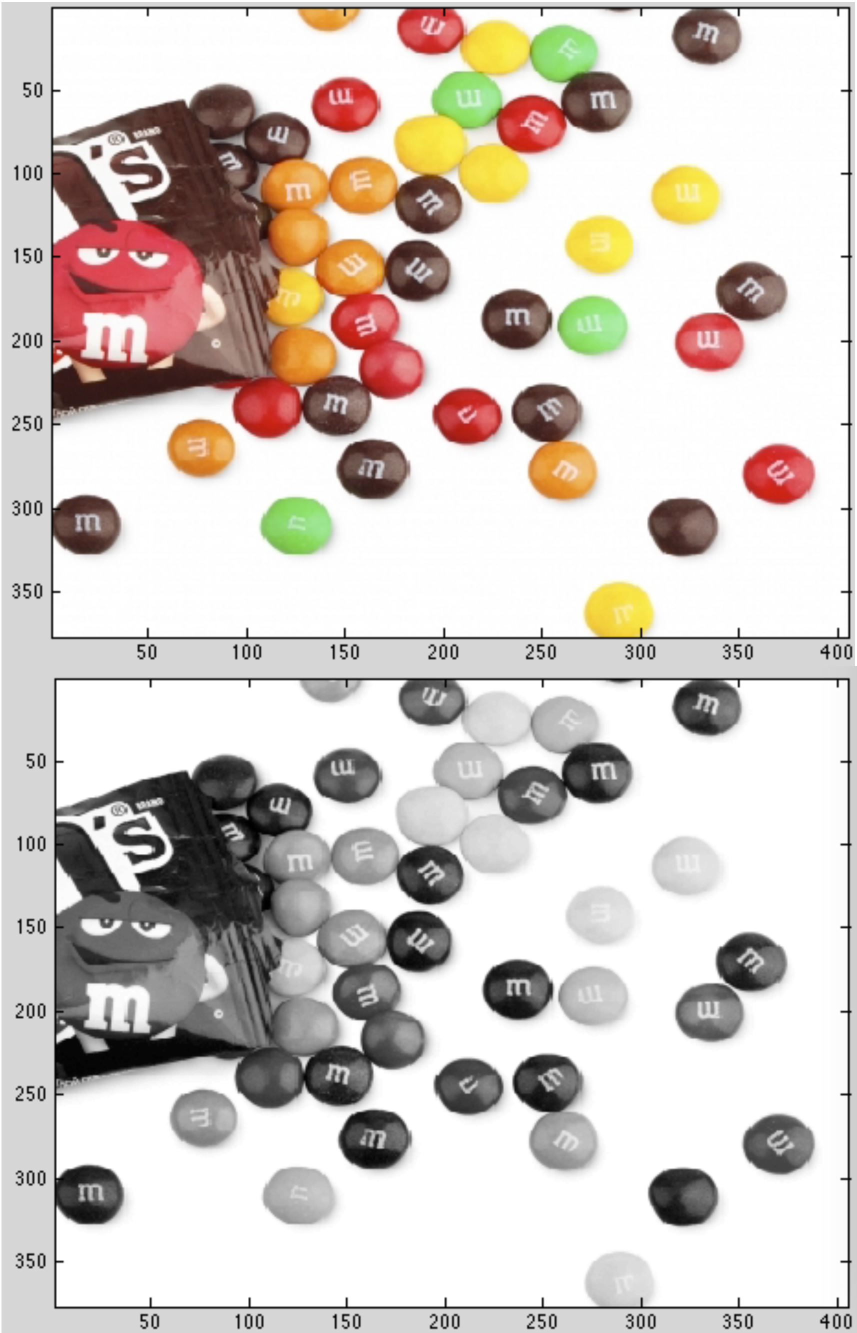


```
>> im = imread('mnms.jpeg');
>> figure;
>> imagesc(im);
>> size(im)

ans =

   376   406     3

>> im8 = rgb2gray(im);
>> size(im8)

ans =

   376   406

>> figure;
>> imagesc(im8); colormap(gray);
>>
```

Non-gray lookup tables

We can also use other mappings:

- 0 → (17, 25, 89)
- 1 → (45, 32, 200)
- ...
- 255 → (233,1,4)

These are called lookup tables.

More colormaps

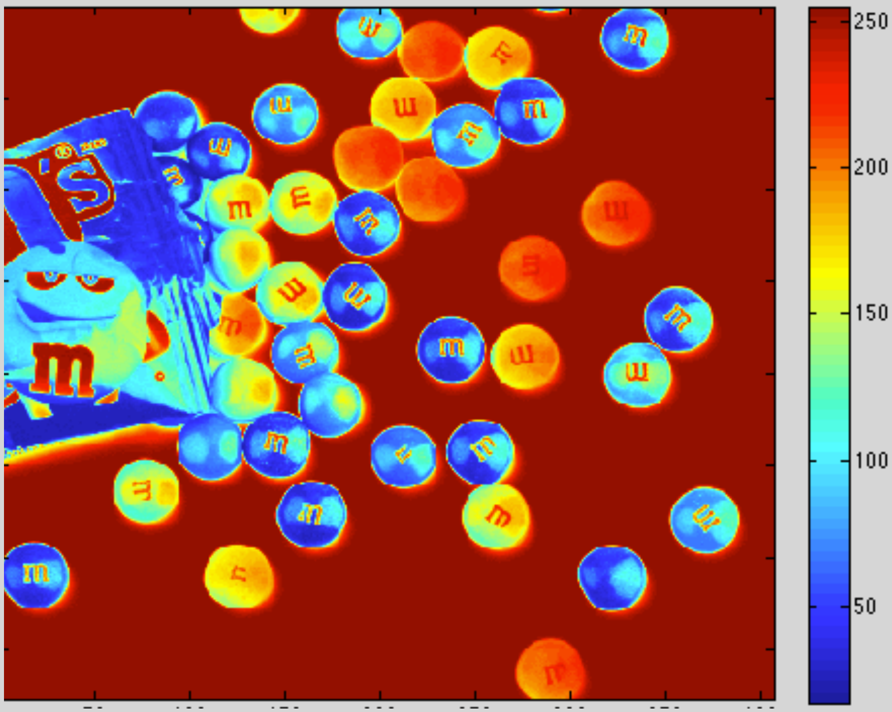

colormap jet;

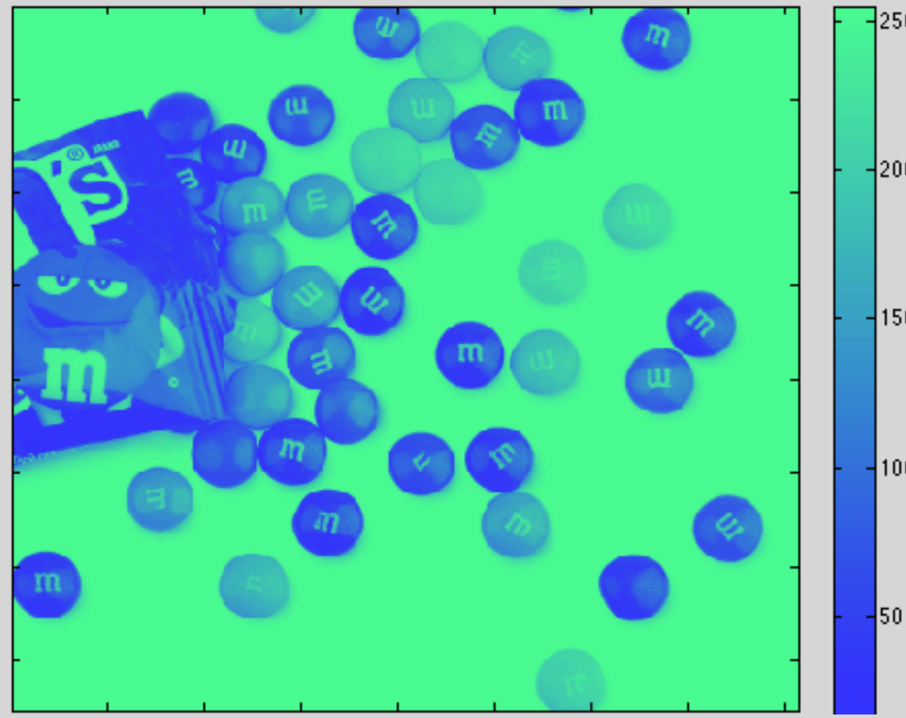

colormap winter;

Fun with Matlab

Colormap Name	Color Scale
parula	
jet	
hsv	
hot	
cool	
spring	
summer	
autumn	
winter	
gray	
bone	
copper	
pink	
lines	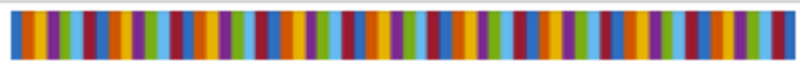
colorcube	
prism	
flag	
white	

Enhancing images

What can we do to “enhance” an image after it has already been digitized?

- We can make the information that is there easier to visualize.
- We can guess at data that is not there, but we cannot be sure, in general.

contrast enhancement

deblurring

Contrast enhancement

Two methods:

- Normalize the data (contrast stretching)
- Transform the data (histogram equalization)

Contrast stretching

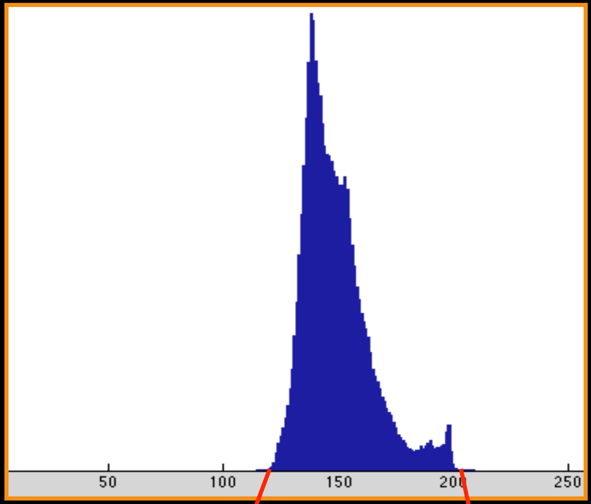

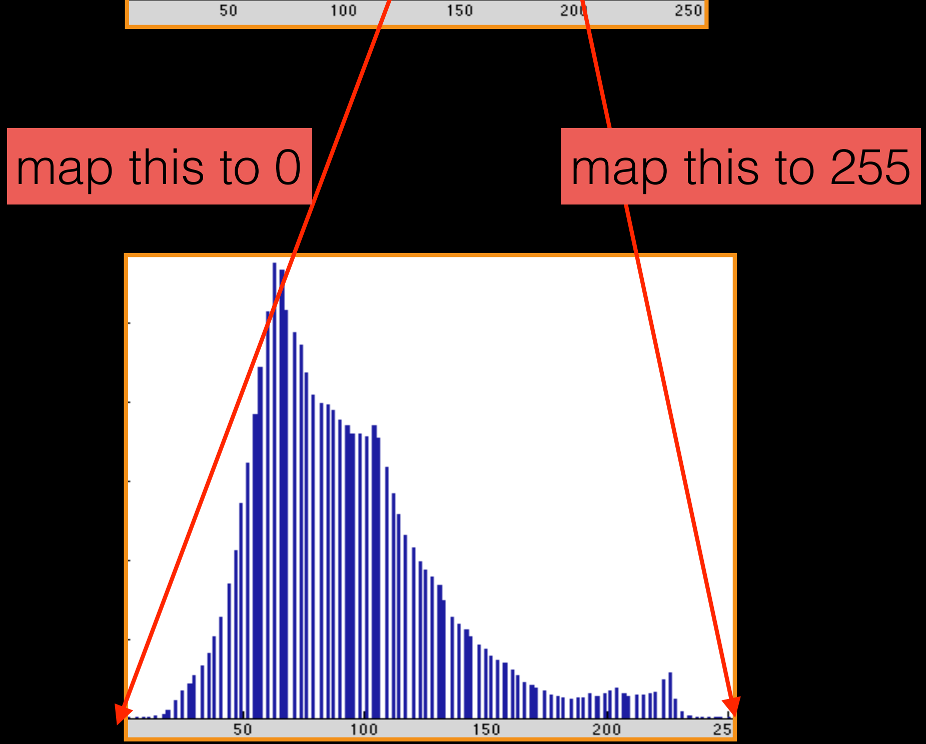

Contrast stretching

- Basic idea: scale the brightness range of the image to occupy the full range of values

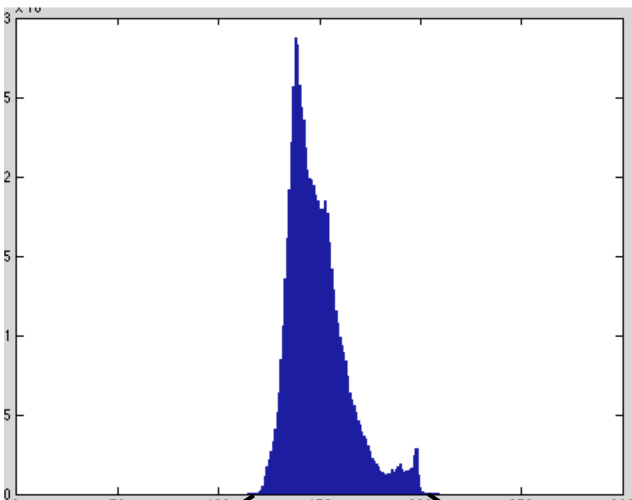

map this to 0 map this to 255

$$I \leftarrow floor\left(\frac{I - \min(I)}{\max(I) - \min(I)} \times 255\right)$$

- Issues:
  - What happens if there is one bright pixel?
  - What happens if there is one dark pixel?

Matlab demo

- imcontrast() — contrast stretching

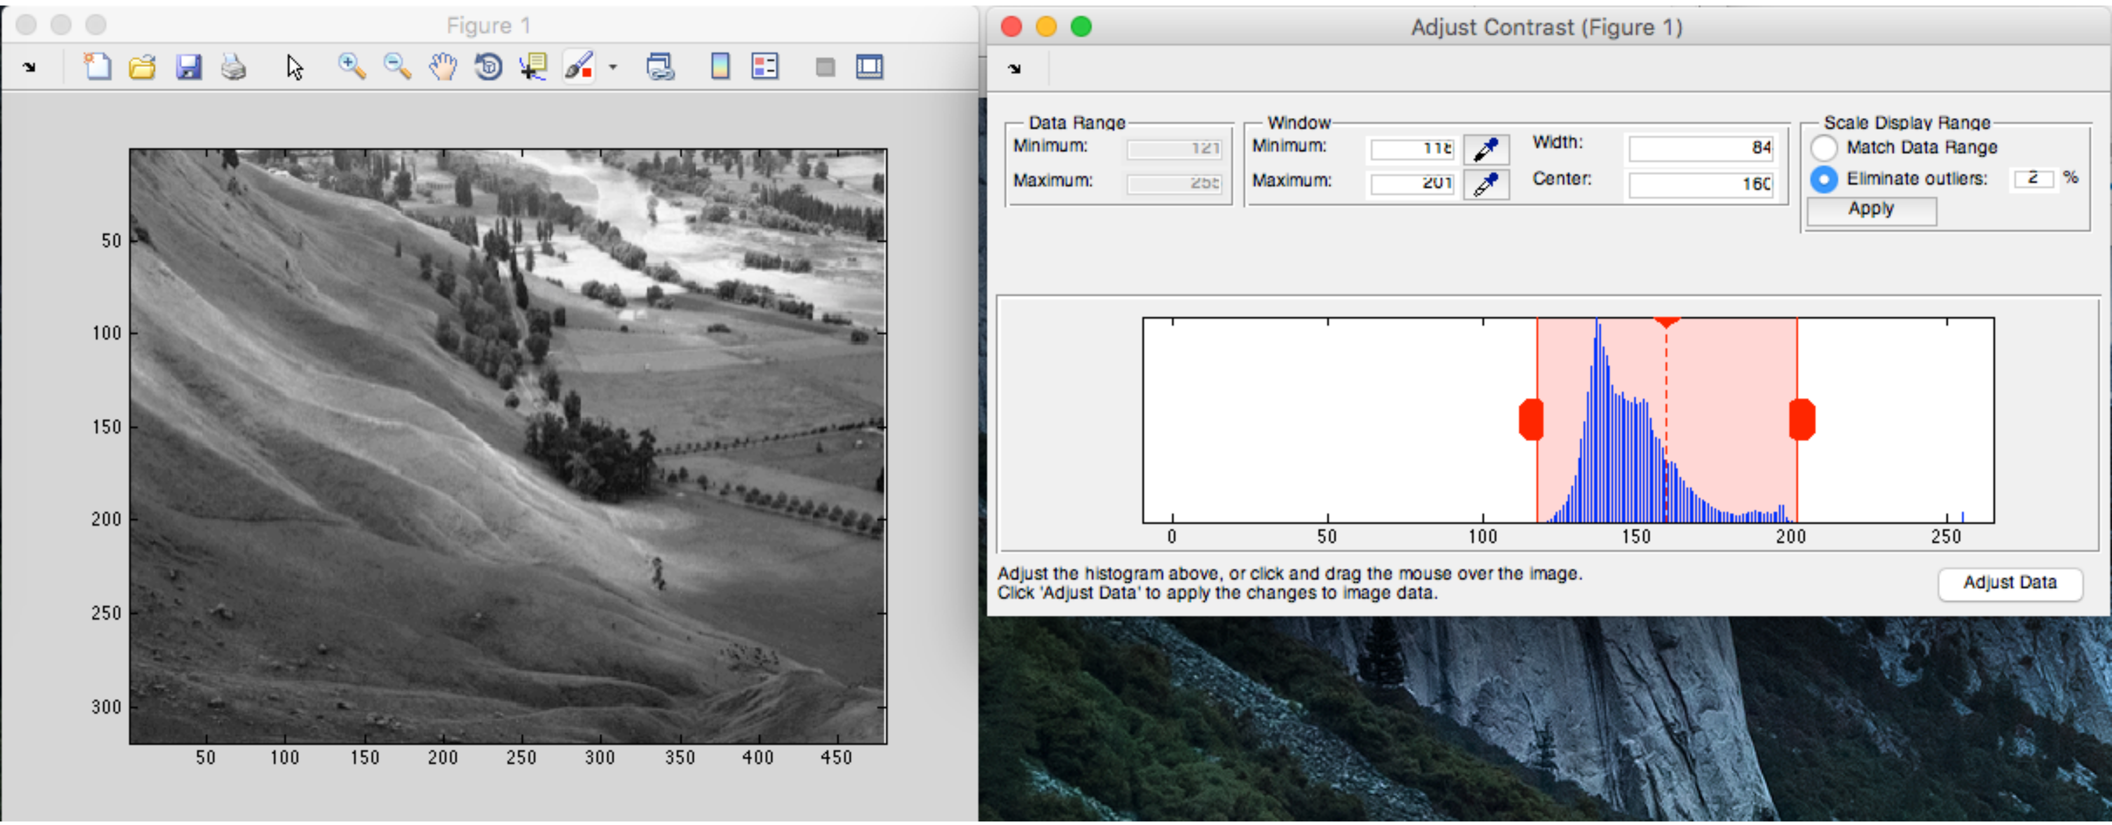

Histogram equalization

histogram

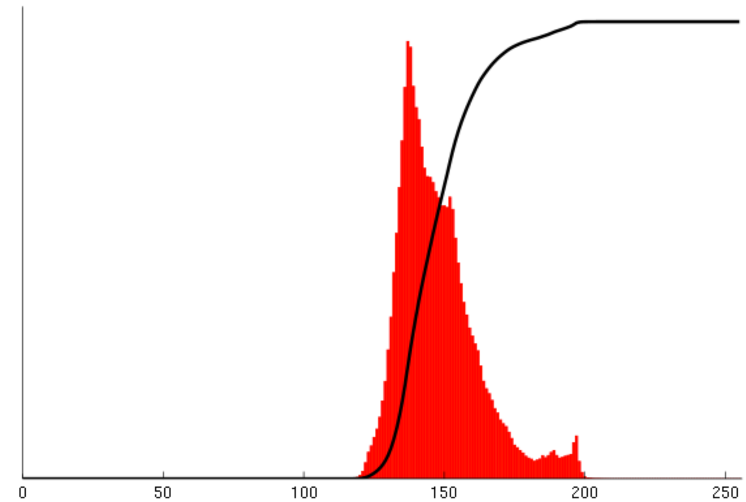

make the distribution close to the uniform distribution

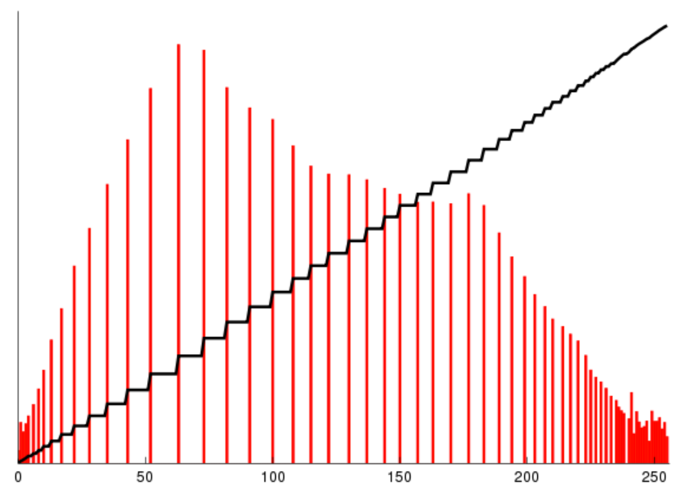

Histogram equalization

Let,

$$p_n = \frac{\text{number of pixels with intensity n}}{\text{total number of pixels}} \qquad n \in \{0, L-1\}$$

If each intensity value *k* is mapped to *T(k)*

$$\boxed{T(k) = \text{floor}\left((L-1)\sum_{n=0}^{k} p_n\right)}$$

Then *T(k)* is roughly a uniform distribution (why?)

https://en.wikipedia.org/wiki/Histogram_equalization

Histogram equalization

original image

original histogram

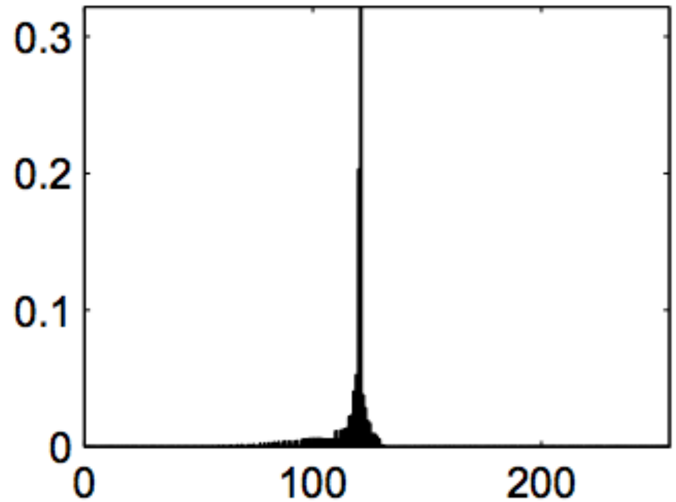

transformed image

transformed histogram

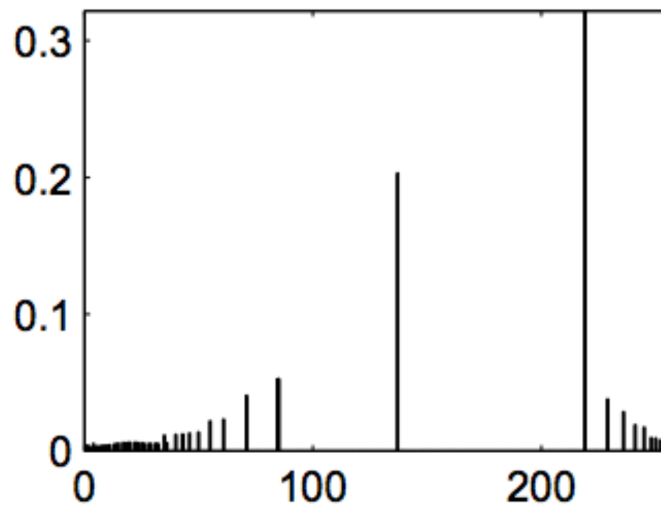

approximately uniform

source: http://www.math.uci.edu/icamp/courses/math77c/demos/hist_eq.pdf