Digitization

World → Optics → Sensor →
Signal → Digitizer →
Digital Representation

Pre-digitization image

- What is an image before we 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 the plane.

Introduction to
Computer Vision

An image as a surface

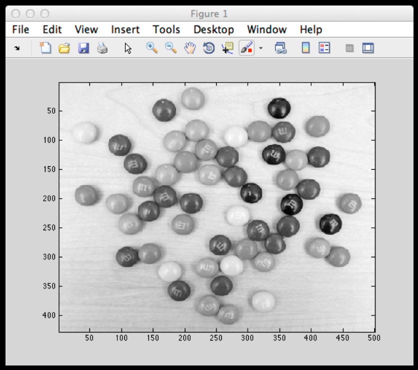

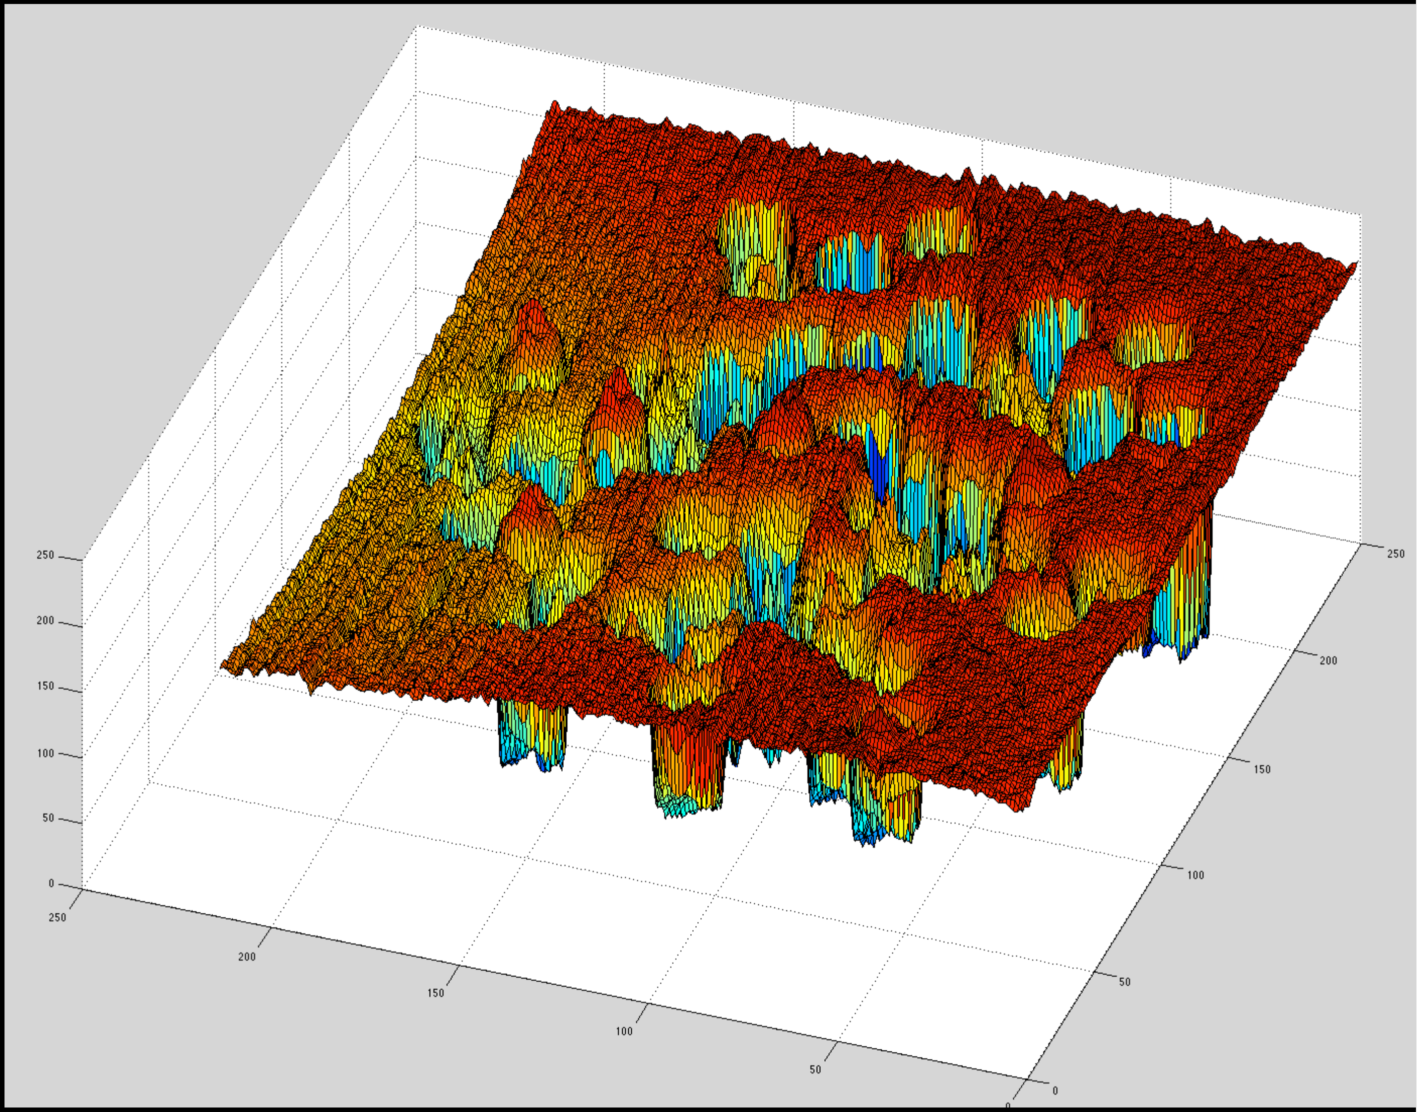

An image as a surface

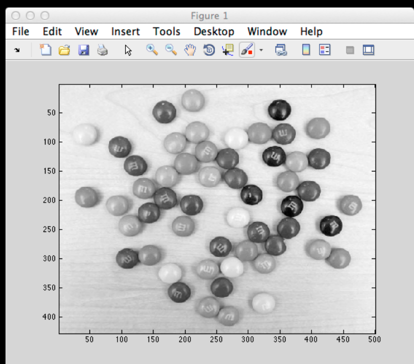

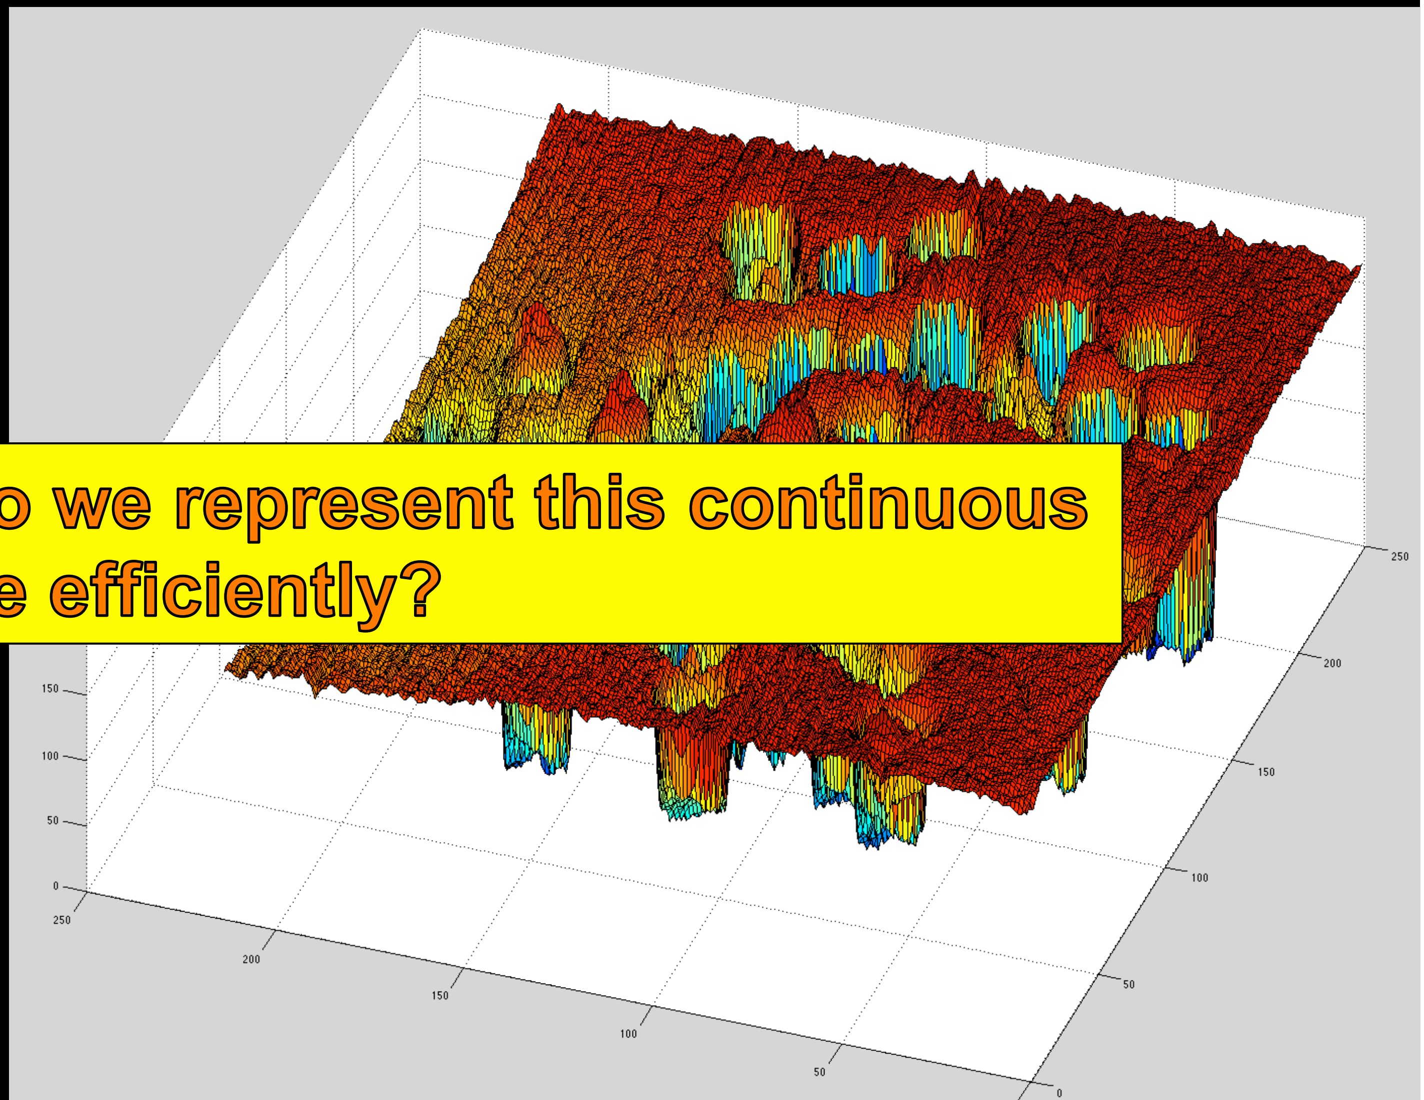

How do we represent this continuous surface efficiently?

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.

Projection through a pixel

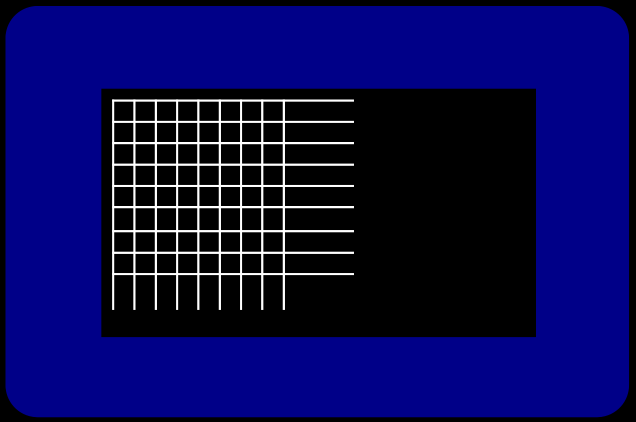

Digitized 35mm Slide or Film

Central Projection Ray

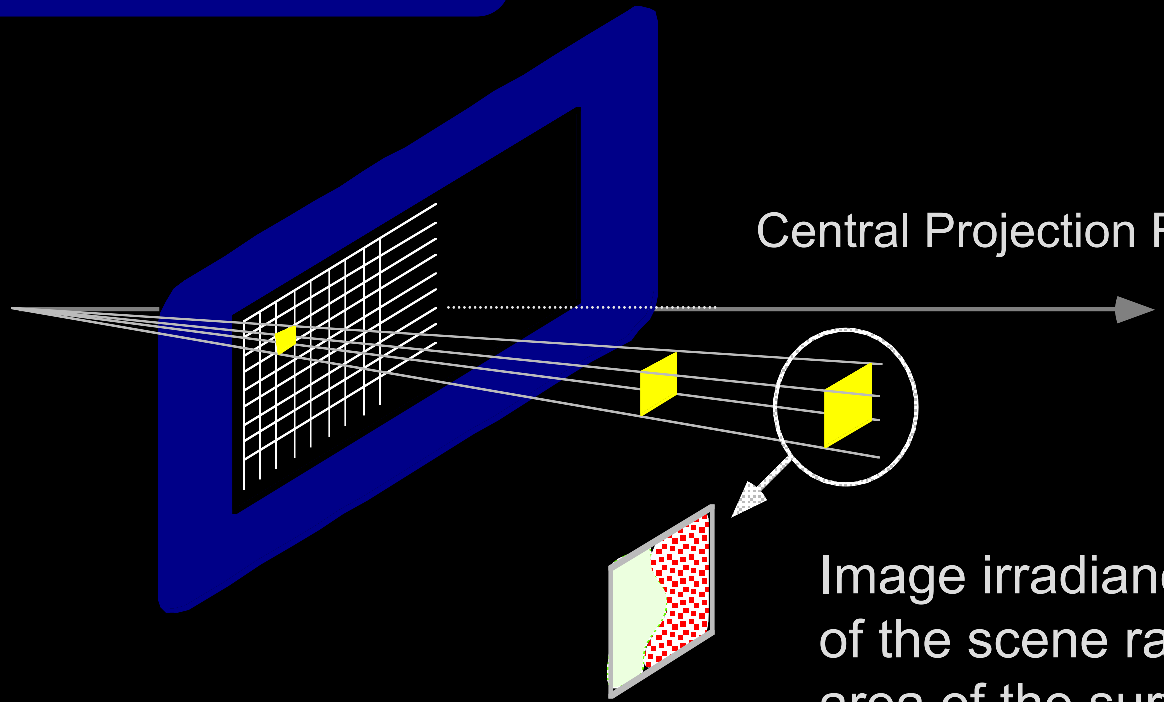

Image irradiance is the average of the scene radiance over the area of the surface intersecting the solid angle!

Signal Quantization

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

I(x,y) = .1583 volts

= ???? Digital value

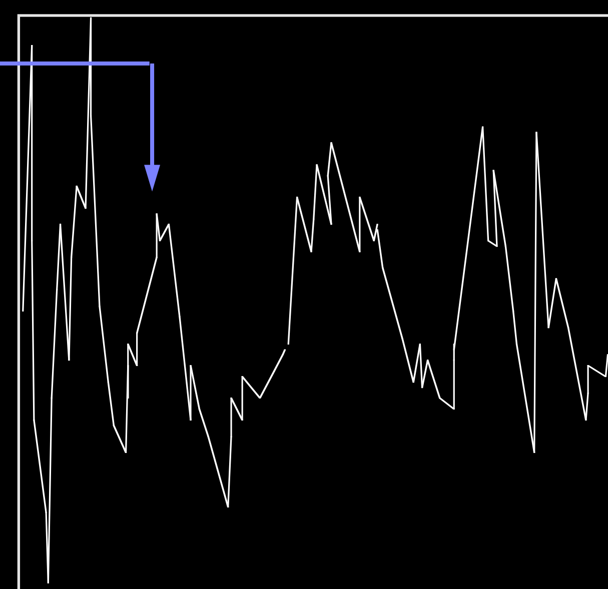

Quantization

- $I(x,y)$ = 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

Original

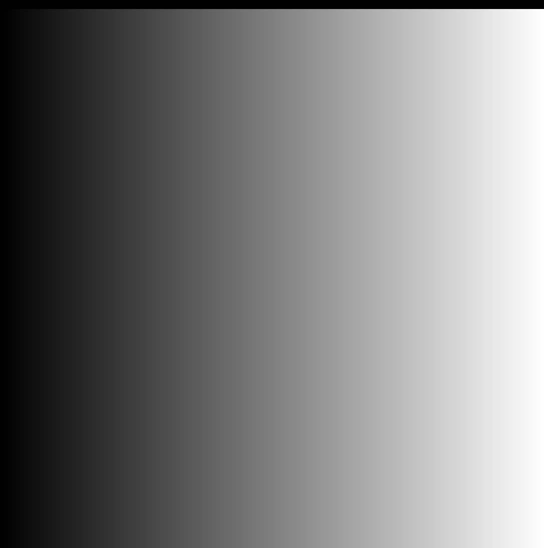

Linear Ramp

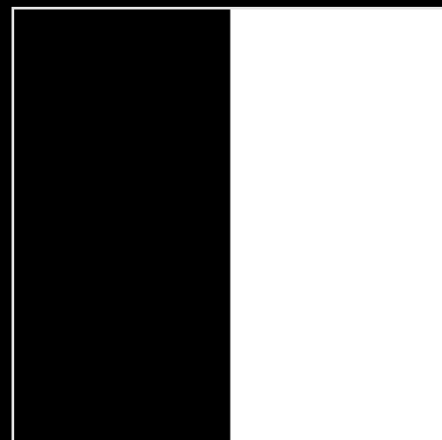

K=2

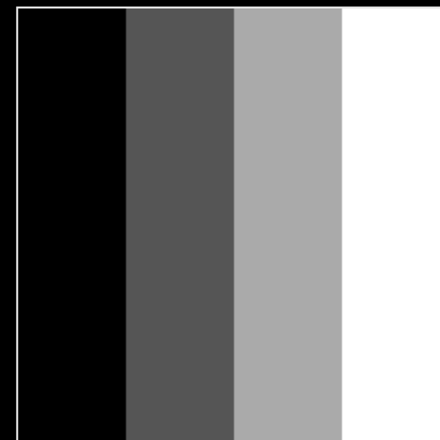

K=4

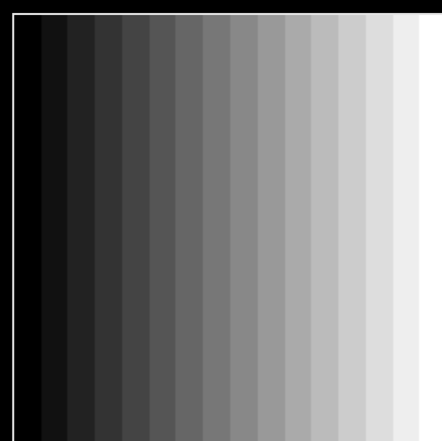

K=16

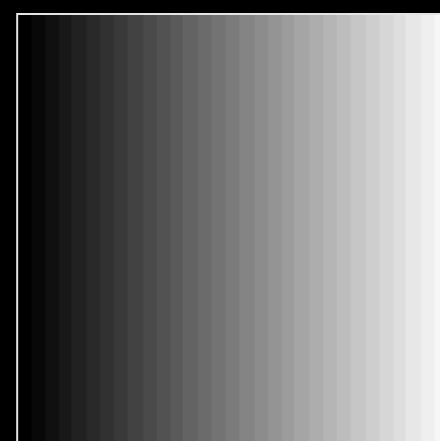

K=32

Choice of K

K=2 (each color)

K=4 (each color)

Choice of 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:

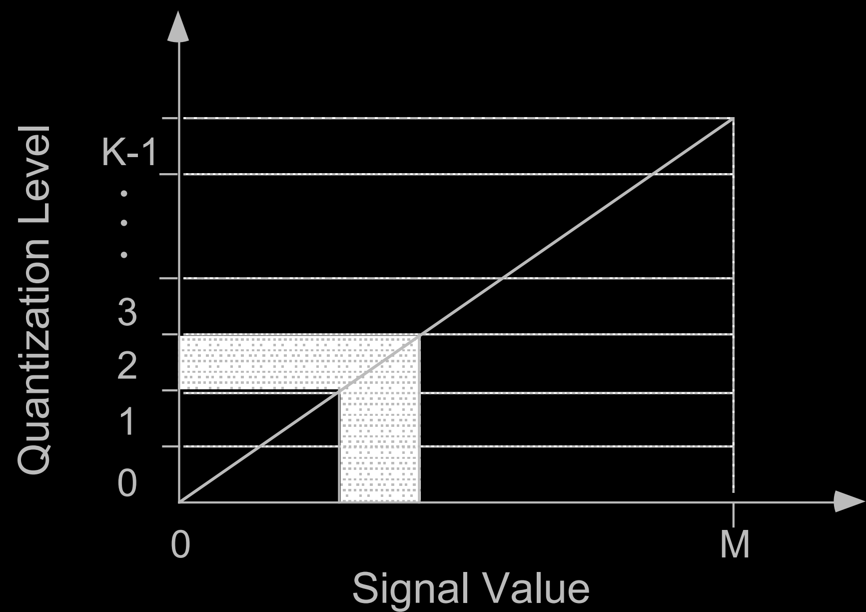

Logarithmic Quantization

- Signal is log I(x,y).
- Effect is:

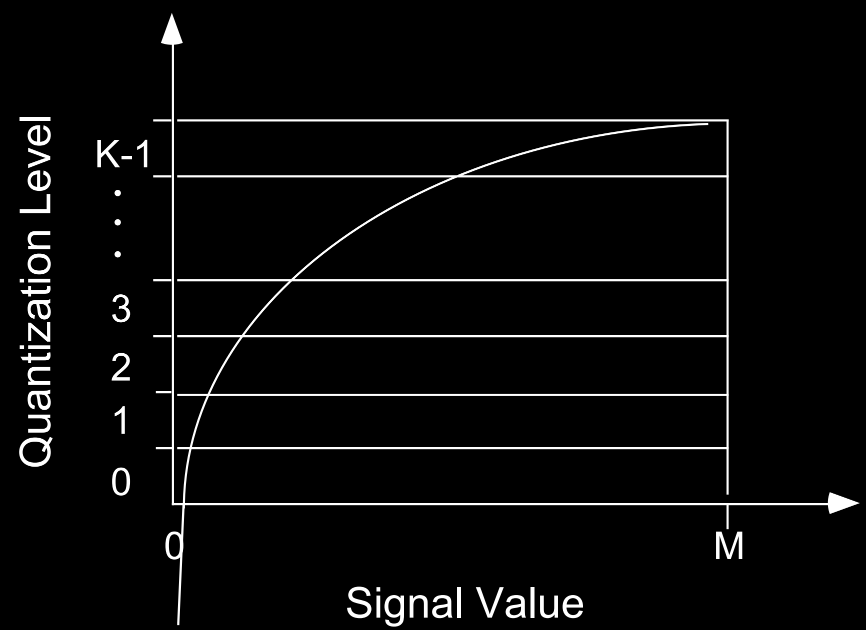

- Detail enhanced in the low signal values at expense of detail in high signal values.

Logarithmic Quantization


Original

Quantization Curve

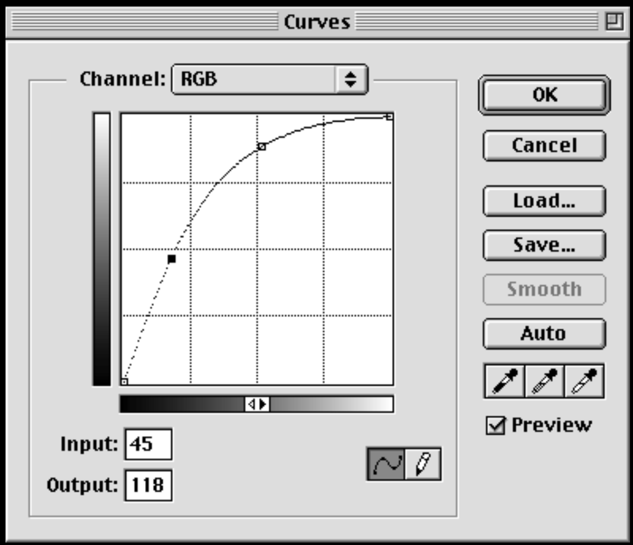


Logarithmic Quantization

Color Displays

- Given a 24-bit color image (8 bits each for R,G,B)
  - Turn on 3 subpixels with power proportional to RGB values:

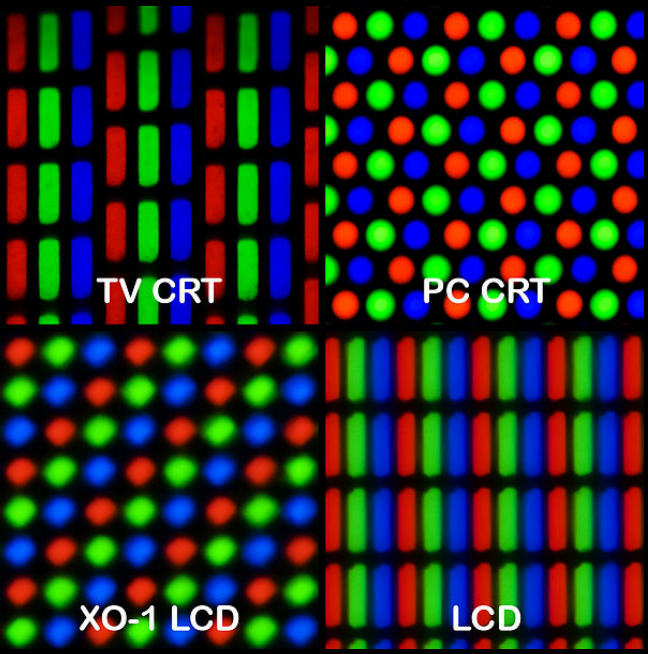

http://en.wikipedia.org/wiki/File:Pixel_geometry_01_Pengo.jpg

Color Displays

- Given a 24-bit color image (8 bits each for R,G,B)
  - Turn on 3 subpixels with power proportional to RGB values:

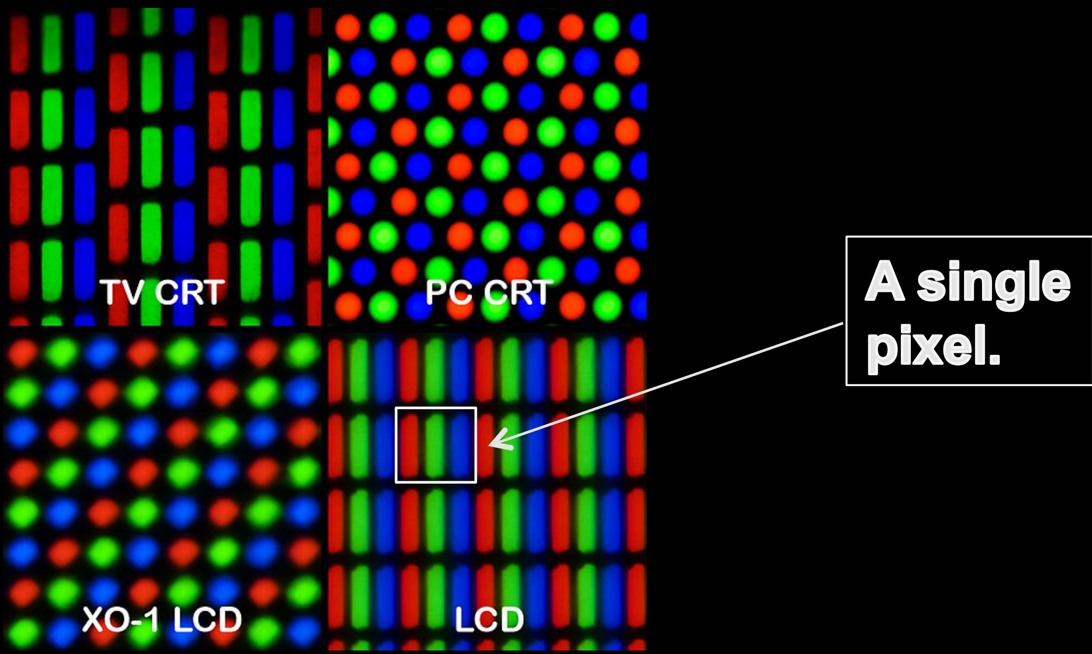

http://en.wikipedia.org/wiki/File:Pixel_geometry_01_Pengo.jpg

“White” text on a color display

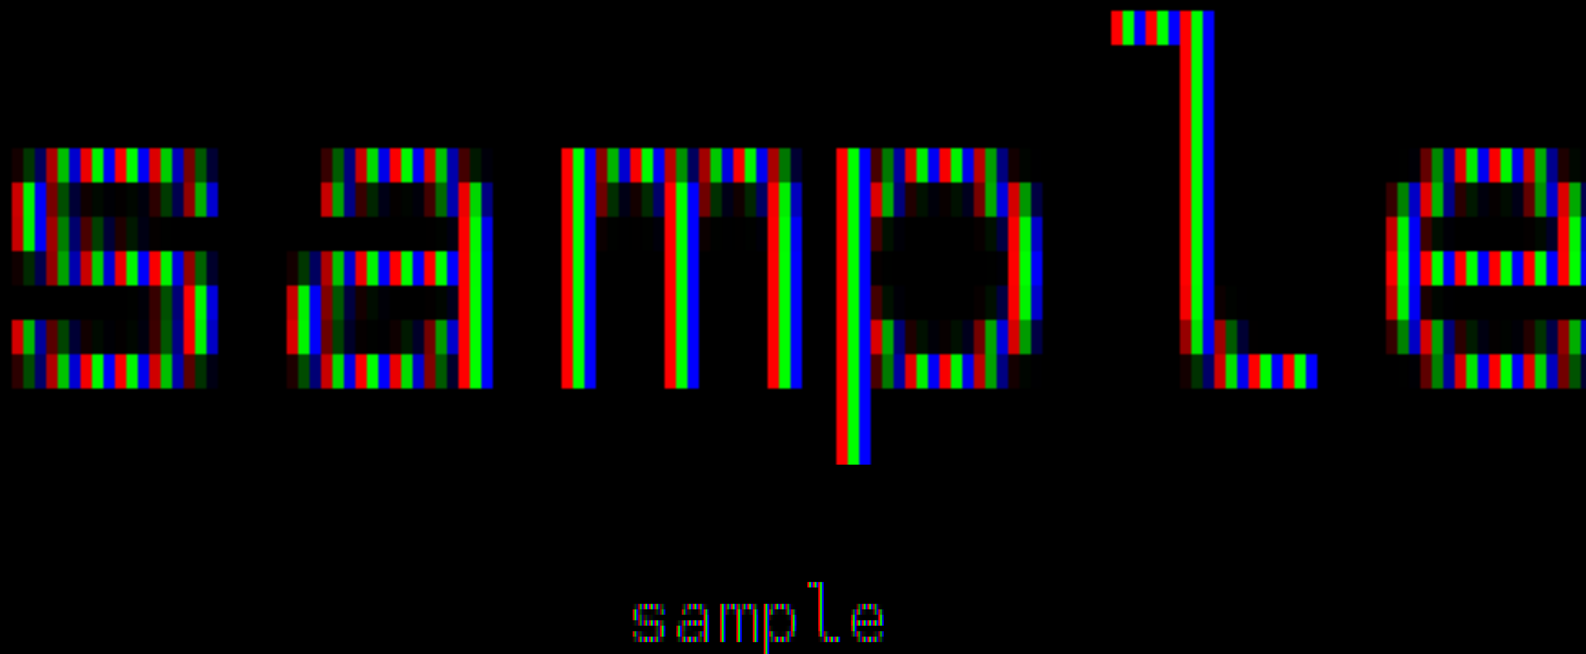

http://en.wikipedia.org/wiki/Subpixel_rendering

Color selector

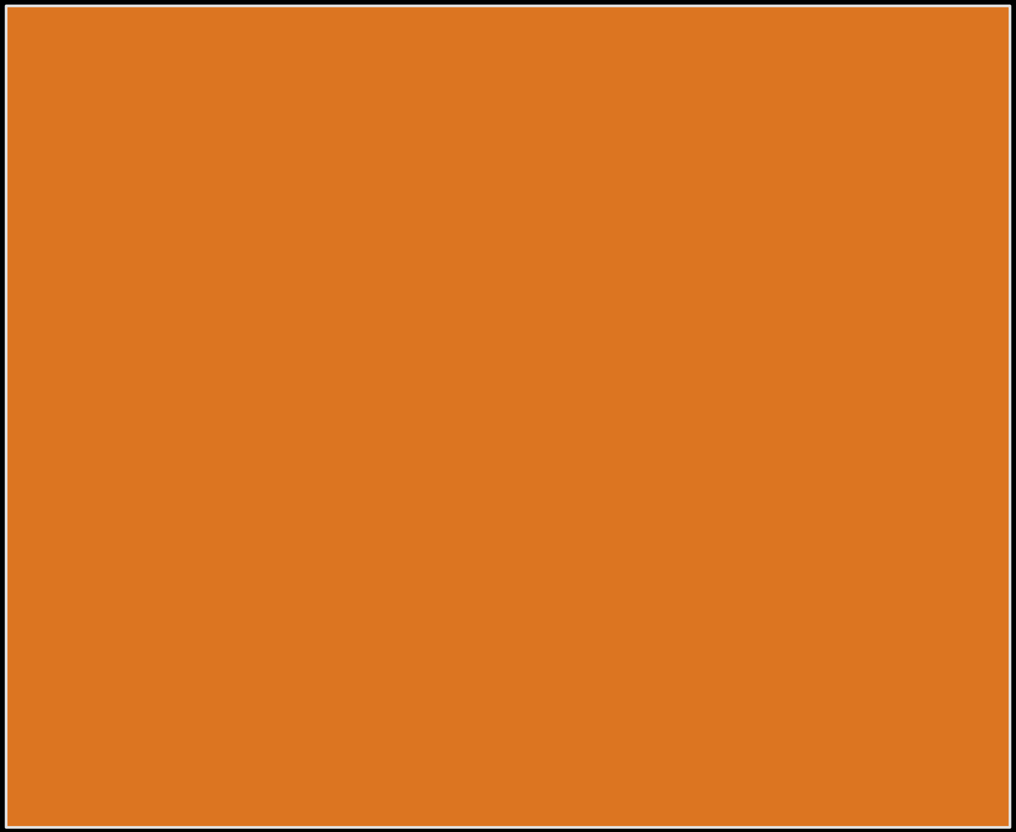

Constructing a Color LCD Display

- See movie.

Look up 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)
    - 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 image.

Lookup tables

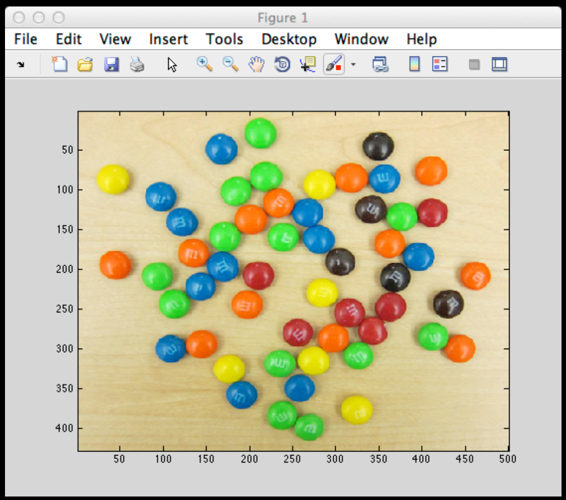

im (24 bits)
"true color"

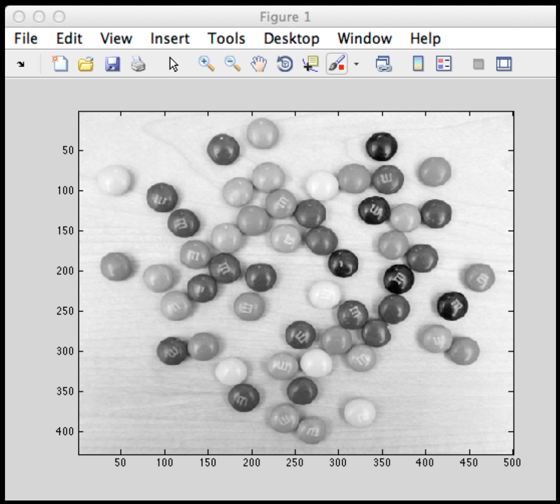

im8 (8 bits)
gray color look up table

```
>> imagesc(im);
>> size(im)

ans =

   428   500     3

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

ans =

   428   500

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

Non-gray look up tables

- We can also use other mappings:
  - 0->(17, 25, 89)
  - 1-> (45, 32, 200)
  - ...
  - 255-> (233,1,4)

  - These are called look up tables.

Introduction to
Computer Vision

More look up tables.

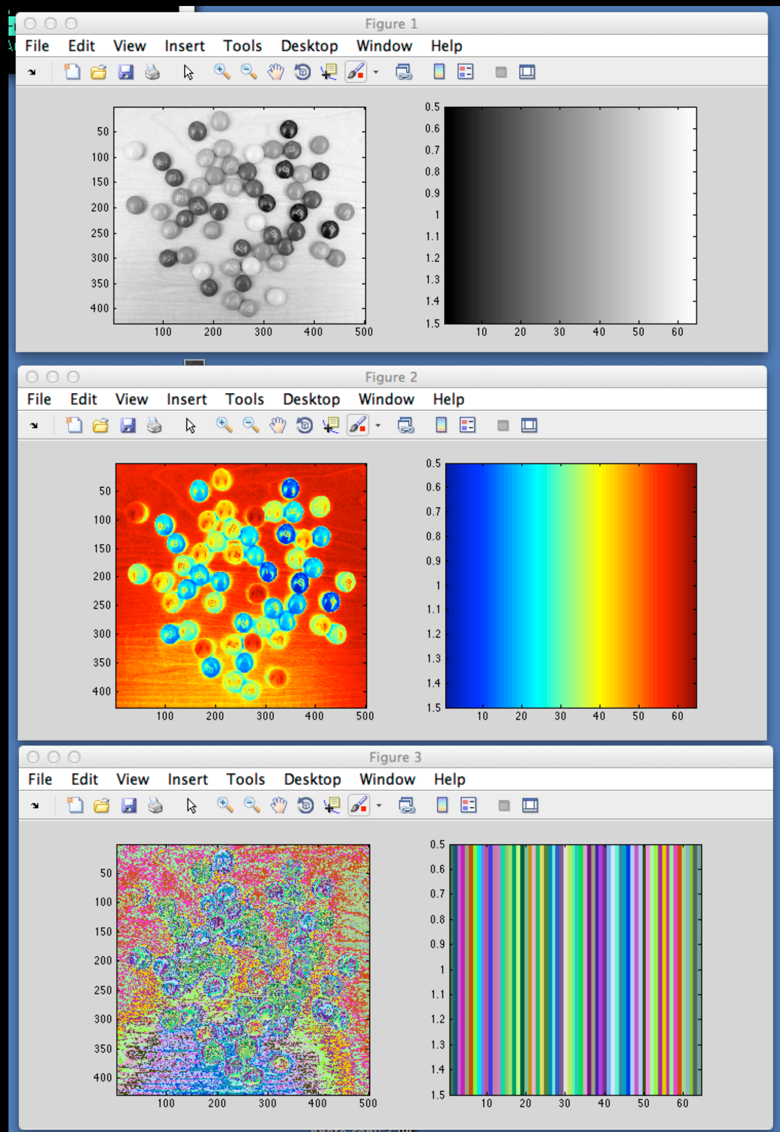

Fun with matlab

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.

Can we “enhance” an image after digitization?

Brightness Equalization

- Two methods:
  - Change the data (histogram equalization)
  - Use a look up table (brightness or color remapping)

Histogram Equalization

An unequalized image

Histogram Equalization

An unequalized image

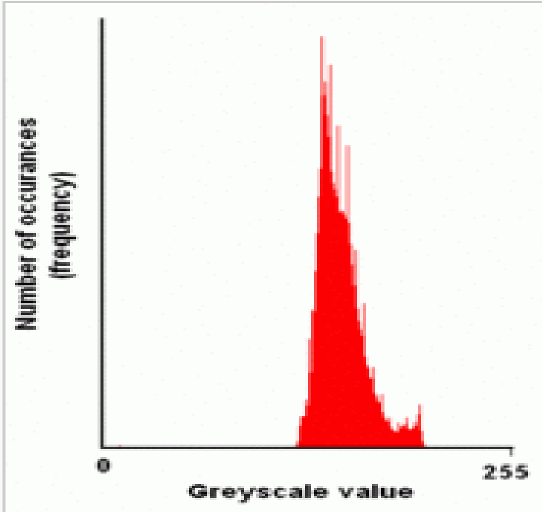

Corresponding histogram

Histogram Equalization

An unequalized image

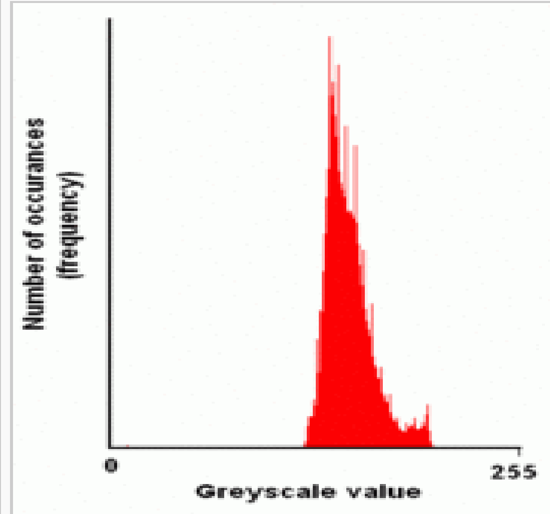

Corresponding histogram

Same image after histogram equalization

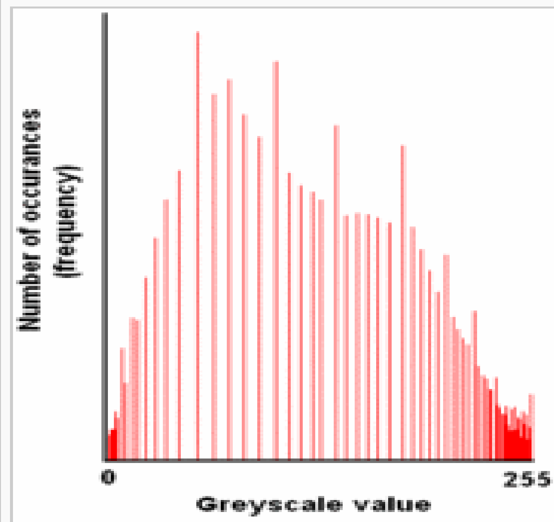

Corresponding histogram

Brightness Equalization

- Two methods:
  - Change the data (histogram equalization).
  - Use a look up table (brightness equalization).

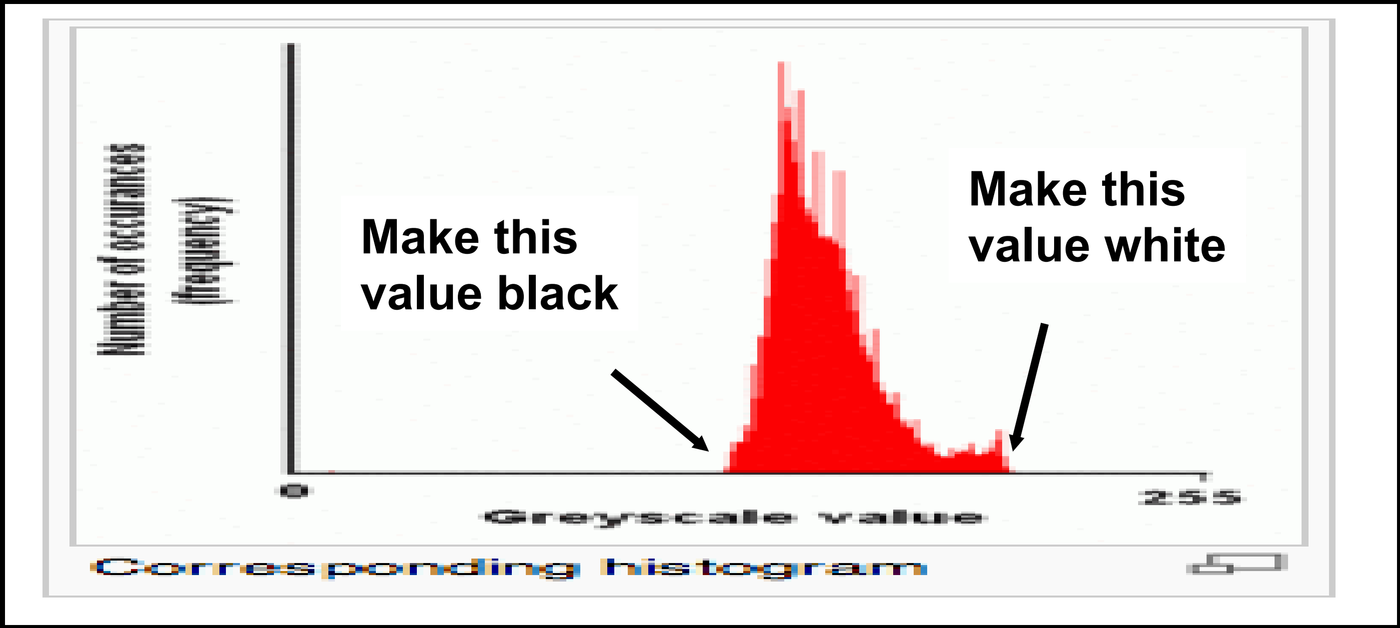

Look up tables

Map lowest value in image to black, highest value to white.

- 0 -> (0, 0, 0)
- 1 -> (0, 0, 0)
- 2 -> (0, 0, 0)
- 3 -> (0 , 0, 0)
- ...
- 130-> (0,0,0)
- 131-> (.01, .01, .01)
- 132-> (.02,.02,.02)
- ...
- 229->(1,1,1)
- 230->(1,1,1)
- ...
- 255 -> (1, 1, 1)

Brightness Equalization

An unequalized image

Mixed Pixel Problem

Mixed Pixel Problem

Pattern Recognition

- Typical recognition problems:
  - Recognize letters and words
  - Recognize people
  - Recognize classes of objects
  - Recognize places

Recognizing Text

Recognizing People

Classes of objects

Recognizing places

Recognizing Handwritten Digits

Supervised Learning

- Supervised learning:
  - Formalization of the idea of learning from examples.
- 2 elements:
  - Training data
  - Test data
- Training data:
  - Data in which the *class* has been identified.
    - Example: This is a “three”. 
- Test data:
  - Data which the algorithm is supposed to identify.
  - What is this? 

Supervised learning

- Formally:
  - n training data pairs:

$$(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), ..., (\mathbf{x}_n, y_n)$$

x's are "observations"
y's are the class labels

  - m test data samples:

$$(\mathbf{x}_{n+1}, \mathbf{x}_{n+2}, ..., \mathbf{x}_{n+m})$$