ECG782: MULTIDIMENSIONAL DIGITAL SIGNAL PROCESSING

DIGITAL IMAGE FUNDAMENTALS

OUTLINE

- Image Formation and Models
- Pixels
- Pixel Processing
- Color

M-D SIGNALS

- Use mathematical models to describe signals
  - A function depending on some variable with a physical meaning
- 1D signal
  - E.g. speech, audio, voltage, current
    - Dependent on "time"
- 2D signal
  - E.g. image
    - Dependent on spatial coordinates in a plane
- 3D signal
  - E.g. volume in space, video
- M-D signal
  - E.g. ???

IMAGE FORMATION

- Incoming light energy is focused and collected onto an image plane

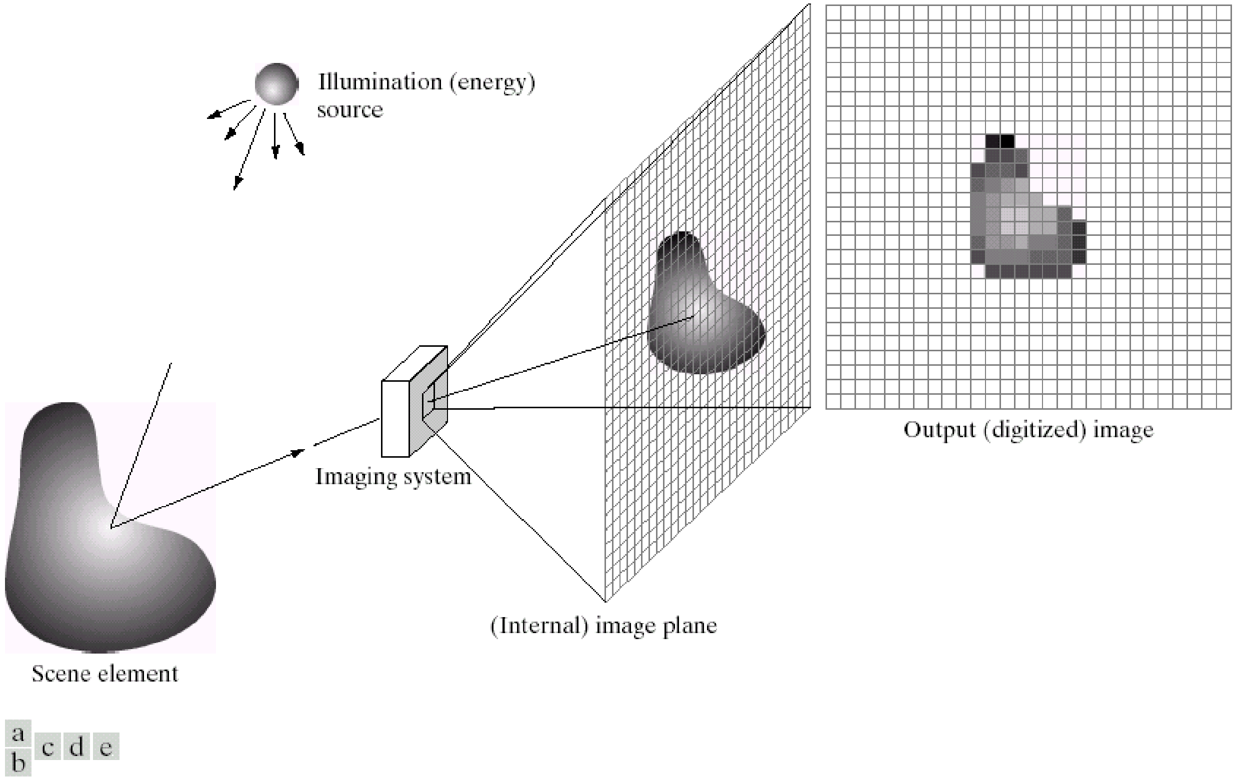

FIGURE 2.15 An example of the digital image acquisition process. (a) Energy ("illumination") source. (b) An element of a scene. (c) Imaging system. (d) Projection of the scene onto the image plane. (e) Digitized image.

IMAGE FORMATION MODEL

- Imaging takes the 3D world and projects it onto a 2D image
- Simple model for the process is called the pinhole camera

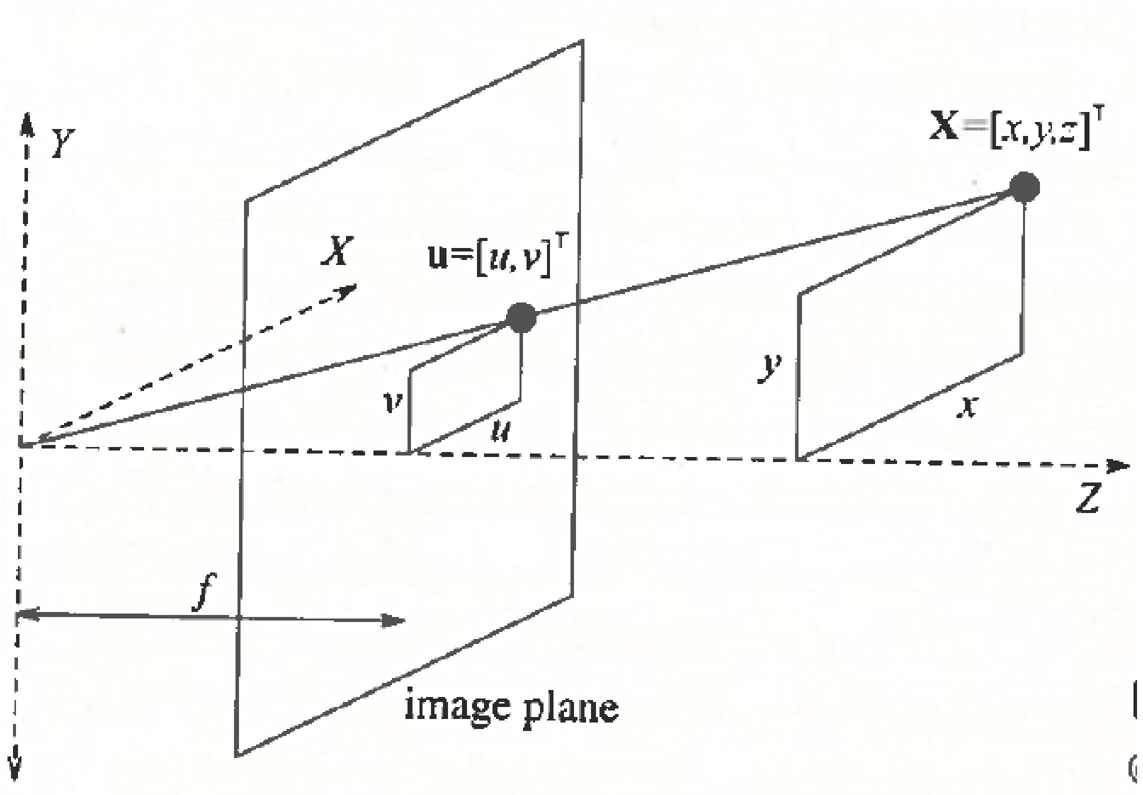

- $X = [x, y, z]^T$
  - represents point in world 3D space
- $u = [u, v]^T$
  - represents a 2D point on image plane
- f – focal length of camera
- World-image relationship
  - $u = \frac{xf}{z}$ $\qquad v = \frac{yf}{z}$

PERSPECTIVE PROJECTION

- Pinhole camera causes perspective distortion
  - Loss of information from perspective projection
  - The transform is not one-to-one
    - A line in space gets mapped to the same point
    - Need depth information to resolve ambiguity
- Orthographic (parallel) projection
  - Linear approximation with $f \to \infty$
  - This is how far away objects $z \to \infty$ are mapped onto image plane

IMAGE REPRESENTATION

- Multiple equivalent representations
- Image

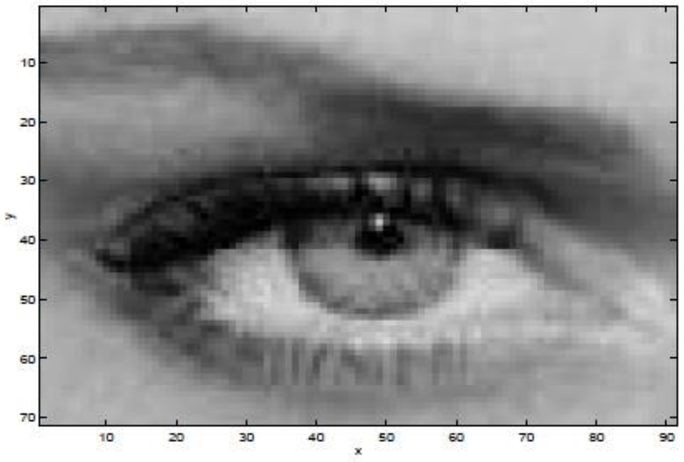

- Surface

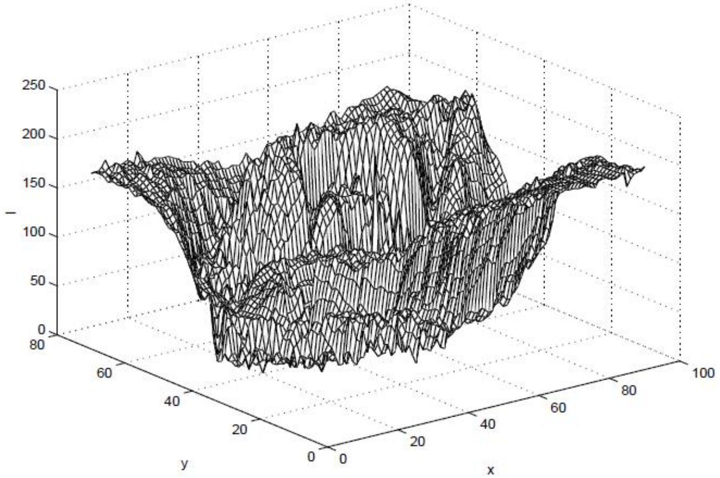

- Matrix

188	186	188	187	168	130	101	99	110	113	112	107	117	140	153	153	156	158	156	153
189	189	188	181	163	135	109	104	113	113	110	109	117	134	147	152	156	163	160	156
190	190	188	176	159	139	115	106	114	123	114	111	119	130	141	154	165	160	156	151
190	188	188	175	158	139	114	103	113	126	112	113	127	133	137	151	165	156	152	145
191	185	189	177	158	138	110	99	112	119	107	115	137	140	135	144	157	163	158	150
193	183	178	164	148	134	118	112	119	117	118	106	122	139	140	152	154	160	155	147
185	181	178	165	149	135	121	116	124	120	122	109	123	139	141	154	156	159	154	147
175	176	176	163	145	131	120	118	125	123	125	112	124	139	142	155	158	158	155	148
170	170	172	159	137	123	116	114	119	122	126	113	123	137	141	156	158	159	157	150
171	171	173	157	131	119	116	113	114	118	125	113	122	135	140	155	156	160	160	152
174	175	176	156	128	120	121	118	113	112	123	114	122	135	141	155	155	158	159	152
176	174	174	151	123	119	126	121	112	108	122	115	123	137	143	156	155	152	155	150
175	169	168	144	117	117	127	122	109	106	122	116	125	139	145	158	156	147	152	148
179	179	180	155	127	121	118	109	107	113	125	133	130	129	139	153	161	148	155	157
176	183	181	153	122	115	113	106	105	109	123	132	131	131	140	151	157	149	156	159
180	181	177	147	115	110	111	107	107	105	120	132	133	133	141	150	154	148	155	157
181	174	170	141	113	111	115	112	113	105	119	130	132	134	144	153	156	148	152	151
180	172	168	140	114	114	118	113	112	107	119	128	130	134	146	157	162	153	153	148
186	176	171	142	114	114	116	110	108	104	116	125	128	134	148	161	165	159	157	149
185	178	171	138	109	110	114	110	109	97	110	121	127	136	150	160	163	158	156	150

IMAGE REPRESENTATION

- Image $f(x, y)$ is a 2D function
  - f – amplitude, gray level, or brightness
  - (x, y) – spatial coordinates
  - Conceptually, (x, y) are continuous but are discrete in practice

- In general, the function can be vector-valued
  - E.g. color images represented by (red, green, blue)
  - $f(x, y) = [r, g, b]^T$
- The image function can be M-dimensional
  - E.g. computed tomography (CT) images are 3D
    - $f(x, y, z)$ represents x-ray absorption at point (x, y, z)

IMAGE AS FUNCTION

- Think of an image as a function, f, that maps from R^2 to R
  - $0 < f(x, y) < \infty$ is the intensity at a point (x, y)
- In reality, an image is defined over a rectangle with a finite range of values
  - $f: [a, b] \times [c, d] \rightarrow [0,1]$
- Computationally, $[0,1]$ range is convenient but usually we have an 8-bit quantized representation
  - $0 < f(x, y) < 255$
- Color image is just three separate functions pasted together
  - $f(x, y) = [r(x, y); g(x, y); b(x, y)]$

IMAGE AS MATRIX

- Images are usually represented by matrices
  - $M \times N$ dimension
- Be aware that images can have different origin definitions
  - Bottom left - typical Cartesian coordinates
  - Upper left – typical image definition (matrix or table notation)
  - Matlab uses (1,1) for origin not (0,0)

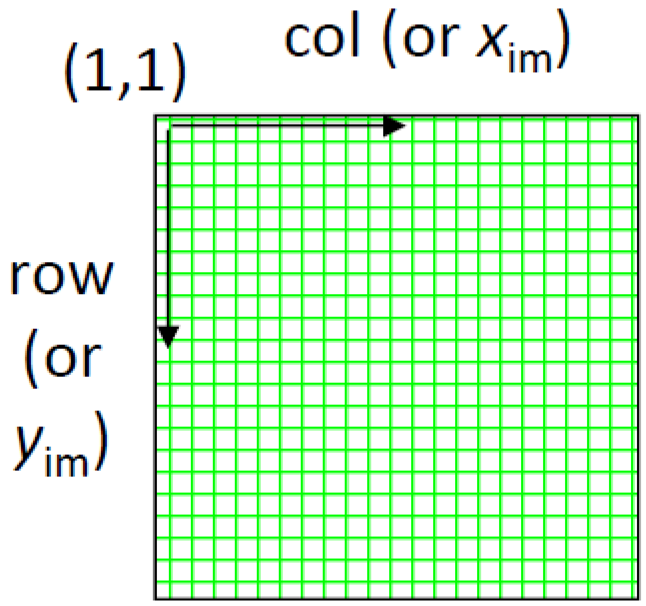

- Index an element either by
  - (x, y)
  - (row, col)

MATRIX NOTATION

- Mathematical
- Notation starts with $f(0,0)$

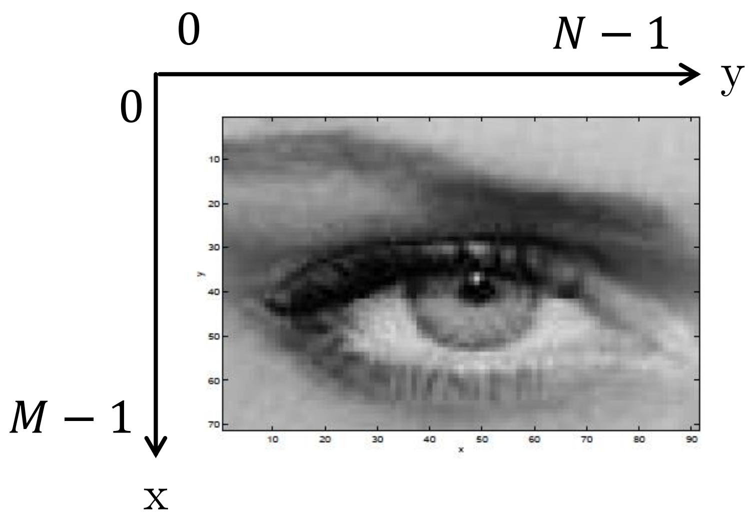

- Matlab
- Notation starts with I(1,1)
  - No zero indexing
  - Swapped axis

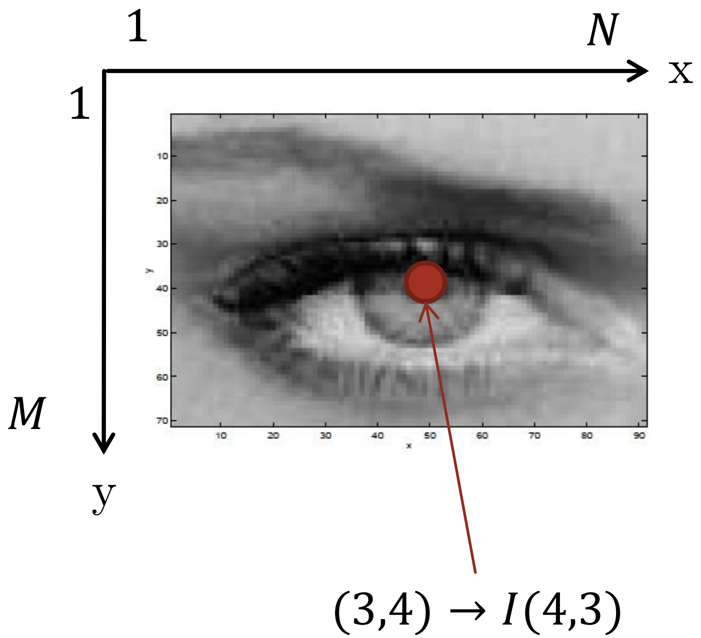

IMAGE SAMPLING

- A continuous image is sampled and ordered into a image grid

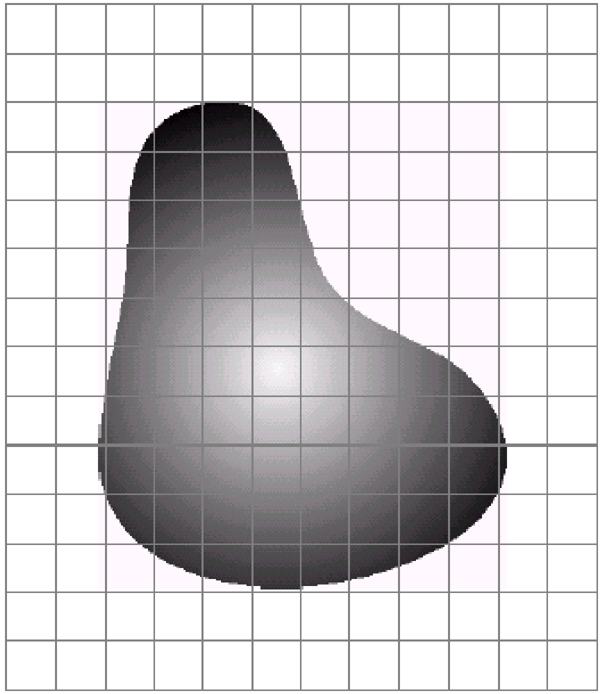

- Each grid element is known as a pixel
  - Voxel for volume element
- Consider the pixel as the smallest unit in an image
  - This is not quite a delta because it has a finite size on the CMOS sensor
  - It is possible to do sub-pixel processing (e.g. corner detection)

QUANTIZATION

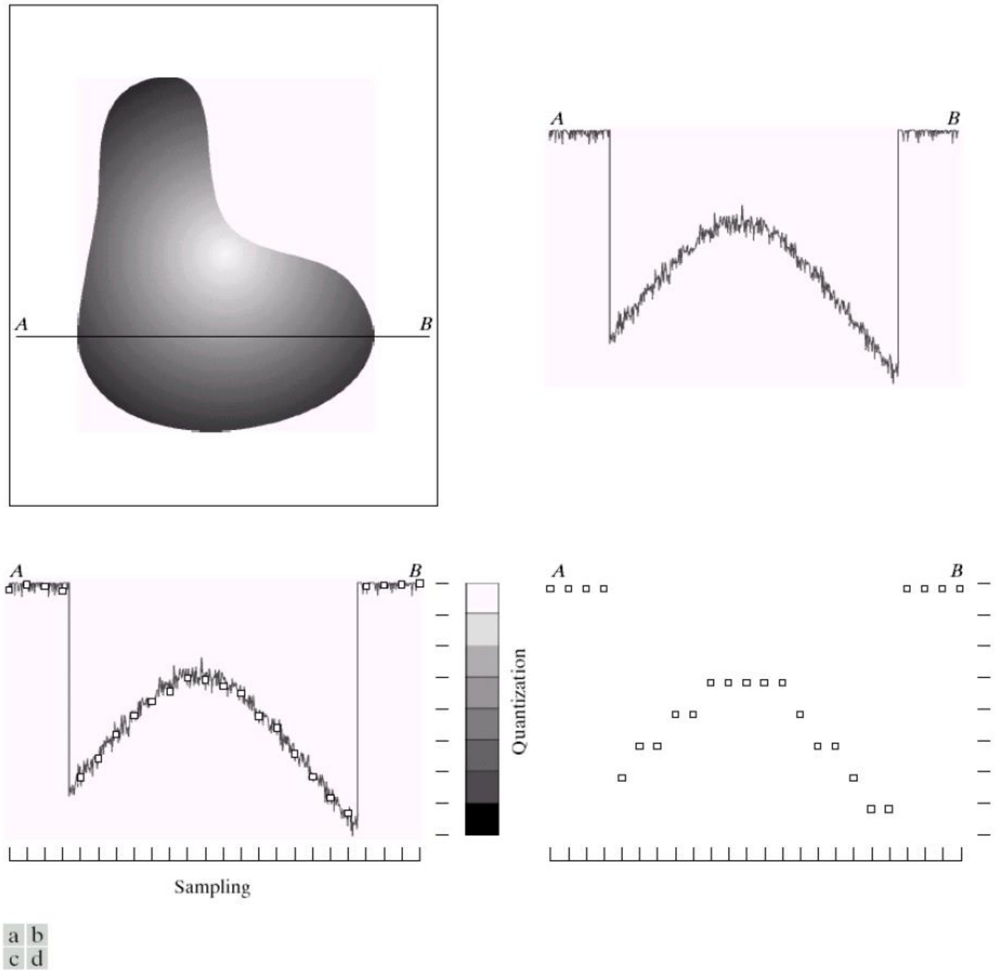

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from *A* to *B* in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

- Quantization gives the number of output levels L

a) Continuous image
b) Scan line from A to B
c) Sampling (horizontal bar) and quantization (vertical bar)
d) Digital scan line – resulting effect of sampling and quantization

QUANTIZATION LEVELS

- L = number of output levels
- k = number of bits per pixel
- Output range of image
  - $[0, L-1] = [0, 2^k - 1]$

- Image storage size
  - $b = M \times N \times k$
  - Number of bits to store image with dimensions $M \times N$

- 8-bits per channel is typical
  - Provide enough resolution to provide quality visual reproduction

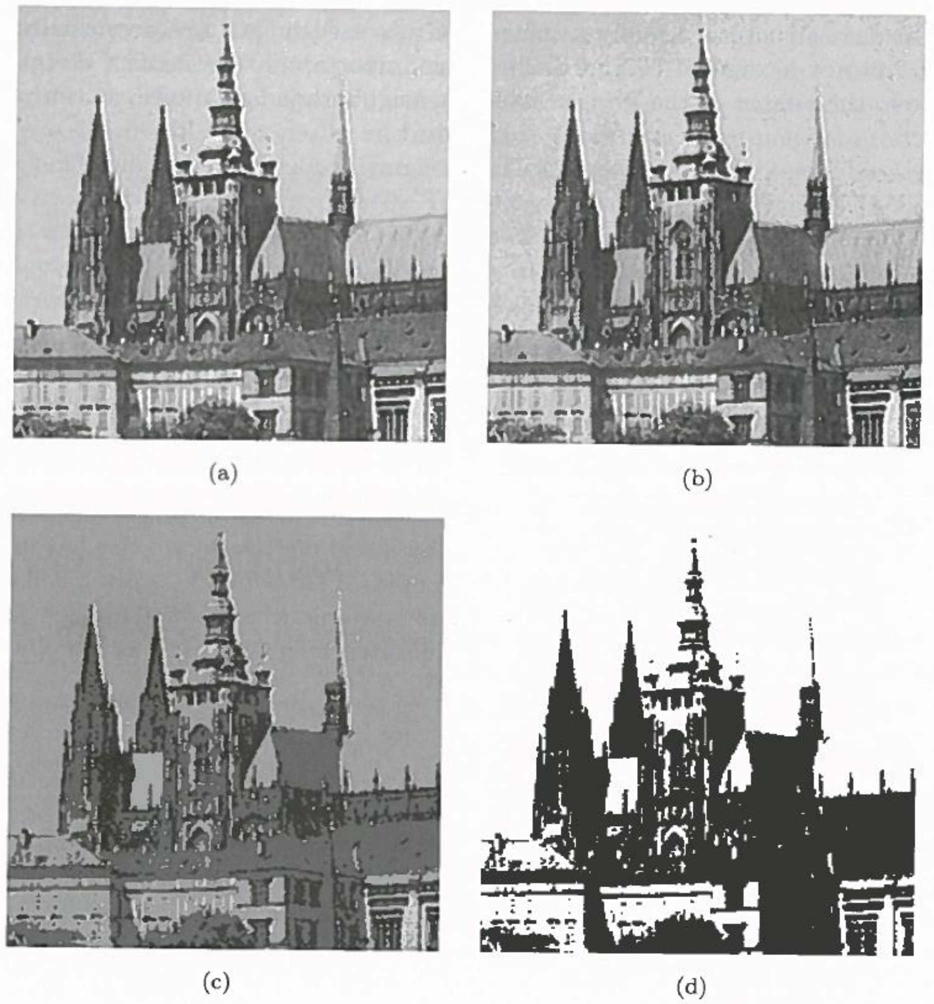

Figure 2.3: Brightness levels. (a) 64. (b) 16. (c) 4. (d) 2. © *Cengage Learning 2015.*

RESOLUTION

- Spatial resolution → smallest discernible detail in an image
- Controlled by the sampling factor (the size $M \times N$ of the CMOS sensor)

- Gray-level resolution → smallest discernible change in gray level
  - Based on number of bits for representation

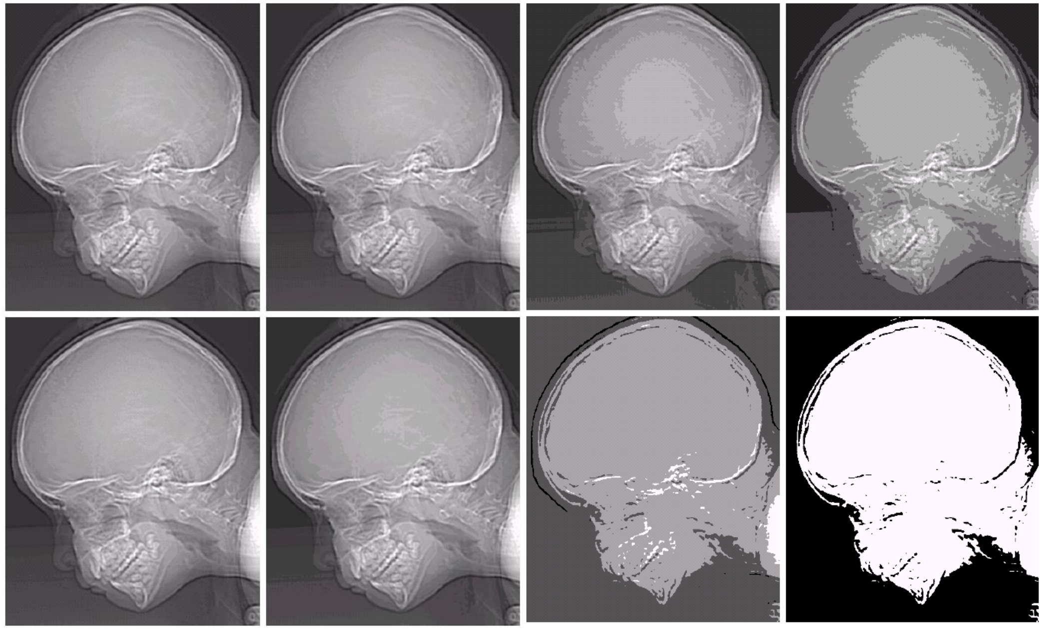

PIXEL NEIGHBORHOOD

- The pixel neighborhood corresponds to nearby pixels

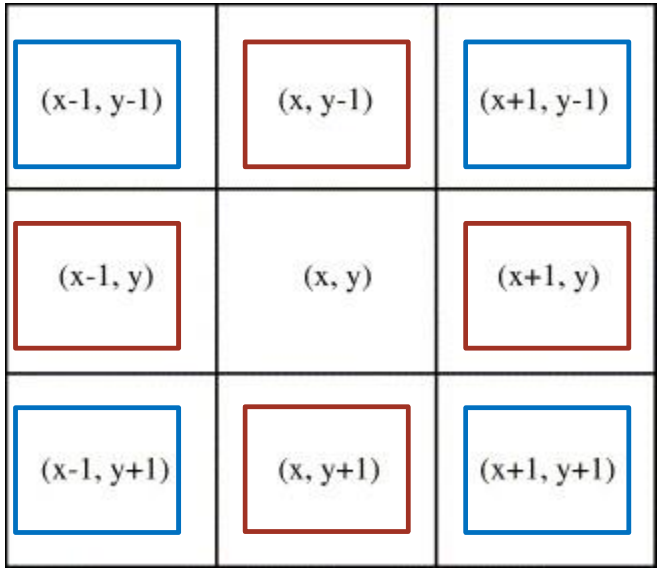

- 4-neighbors
  - Horizontal and vertical neighbors
- 8-neighbors
  - Include 4-neighbors and the diagonal pixels

CONNECTIVITY

- Path exists between pixels

- 4-connected

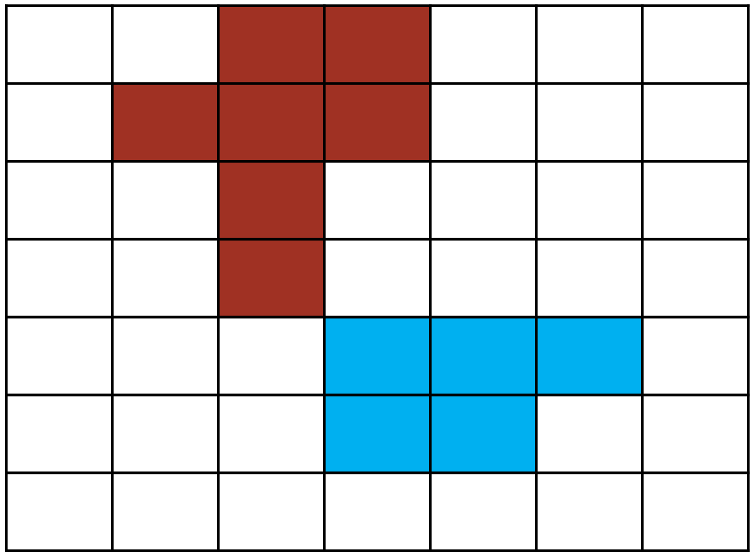

- 8-connected

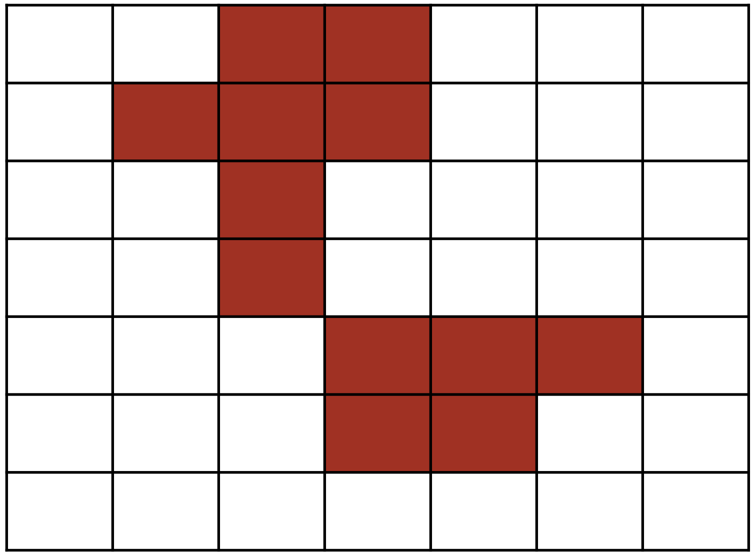

IMAGE PROCESSING

- Usually the first stage of computer vision applications
  - Pre-process an image to ensure it is in a suitable form for further analysis
- Typical operations include:
  - Exposure correction, color balancing, reduction in image noise, increasing sharpness, rotation of an image to straighten

2D SIGNAL PROCESSING

- Image processing is an extension of signal processing to two independent variables
  - Input signal → output signal
- General system

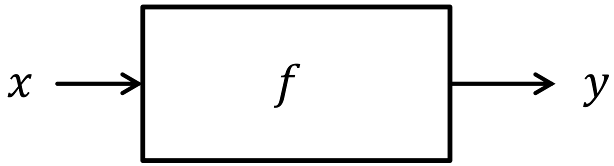

- Image processing

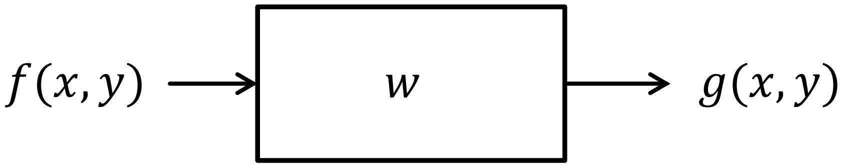

- Linear operators
  - $H(af + bg) = aH(f) + bH(g)$
  - Input is an image, output is an image
- Important class of operators for image processing because of the wealth of theoretical and practical results
  - E.g. signal processing
- However, non-linear operations can provide better performance but not always in predictable ways.

POINT OPERATORS/PROCESSES

- Output pixel value only depends on the corresponding input pixel value
- Often times we will see operations like dividing one image by another
  - Matrix division is not defined
  - The operation is carried out between corresponding pixels in the two image
  - Element-by-element dot operation in Matlab
    - `>> I3 = I1./I2`
    - Where I1 and I2 are the same size

PIXEL TRANSFORMS

- Gain and bias (Multiplication and addition of constant)
  - $g(x, y) = a(x, y)f(x, y) + b(x, y)$
  - a (gain) controls contrast
  - b (bias) controls brightness
    - Notice parameters can vary spatially (think gradients)
- Linear blend
  - $g(x) = (1 - \alpha)f_0(x) + \alpha f_1(x)$
  - We will see this used later for motion detection in video processing

COLOR TRANSFORMS

- Usually we think of a color image as three images concatenated together
  - Have a red, green, blue slice corresponding to the notion of primary colors

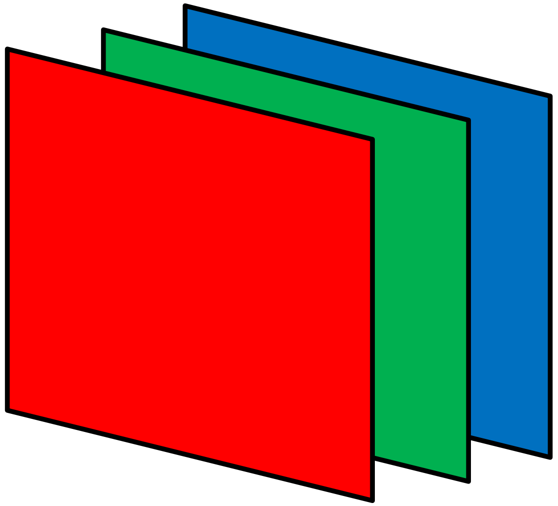

- Manipulations of these color channels may not correspond directly with desired perceptual response
  - Adding bias to all channels may actually change the apparent color instead of increasing brightness
- Need other representations of color for mathematical manipulation

COLOR IMAGES

- Color comes from underlying physical properties

- However, humans do not perceive color in the same physical process
  - There is some subjectivity (e.g. color similarity)

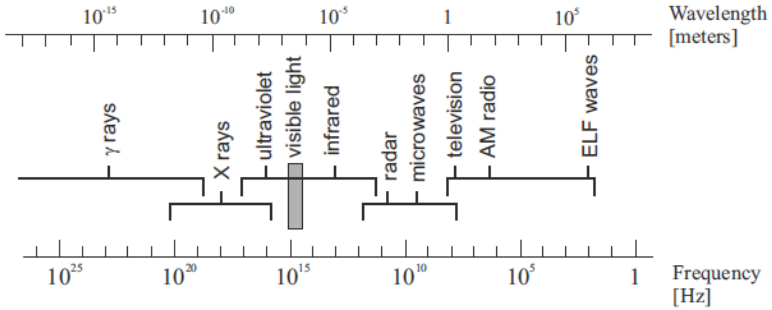

Figure 2.23: Division of the electromagnetic spectrum (ELF is Extremely Low Frequencies). *© Cengage Learning 2015.*

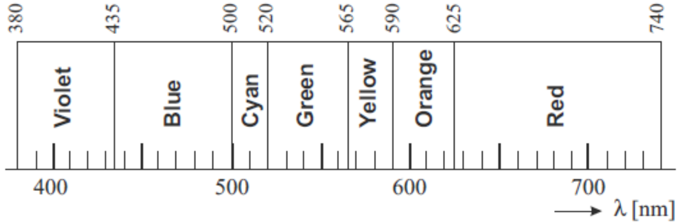

Figure 2.24: Wavelength λ of the spectrum visible to humans. *© Cengage Learning 2015.*

HUMAN COLOR PERCEPTION

- Cones in human retina are sensitive to color
  - In the center of eye
  - 3 different types for different EM frequency sensitivity
- Rods are monochromatic
  - On outside of the eye and good for low lighting and motion sensing

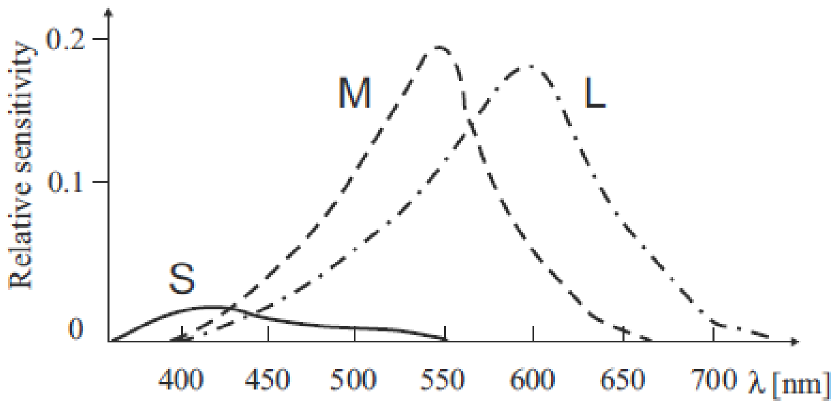

Figure 2.26: Relative sensitivity of S, M, L cones of the human eye to wavelength. © *Cengage Learning 2015.*

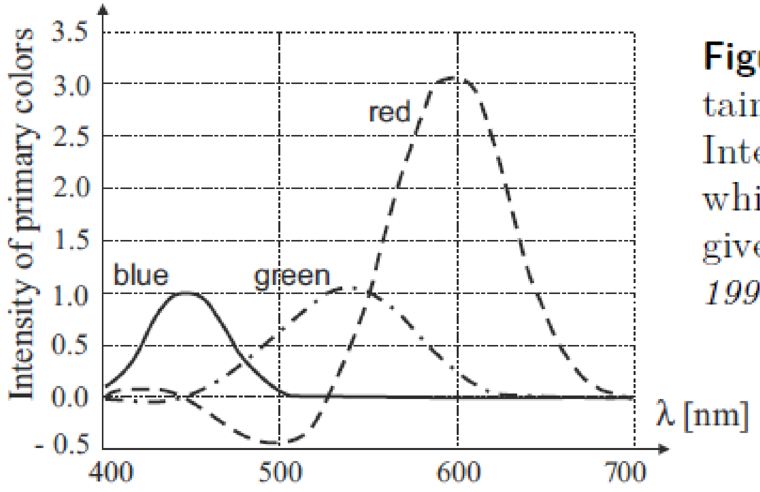

Figure 2.27: Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ. *Based on [Wandell, 1995].*

COLORSPACES

- Uniform method for defining colors
- Can transform from one to another
  - Want to take advantage of properties and color gamut
- XYZ
  - International absolute color standard
  - No negative mixing
- RGB
  - Additive color mixing for red, green, and blue
  - Widely used in computers
- CMYK
  - Cyan, magenta, yellow, black
  - Used for printers and based off of reflectivity
- HSV
  - Hue, saturation, and value = color, amount, brightness
  - Closer to human perception

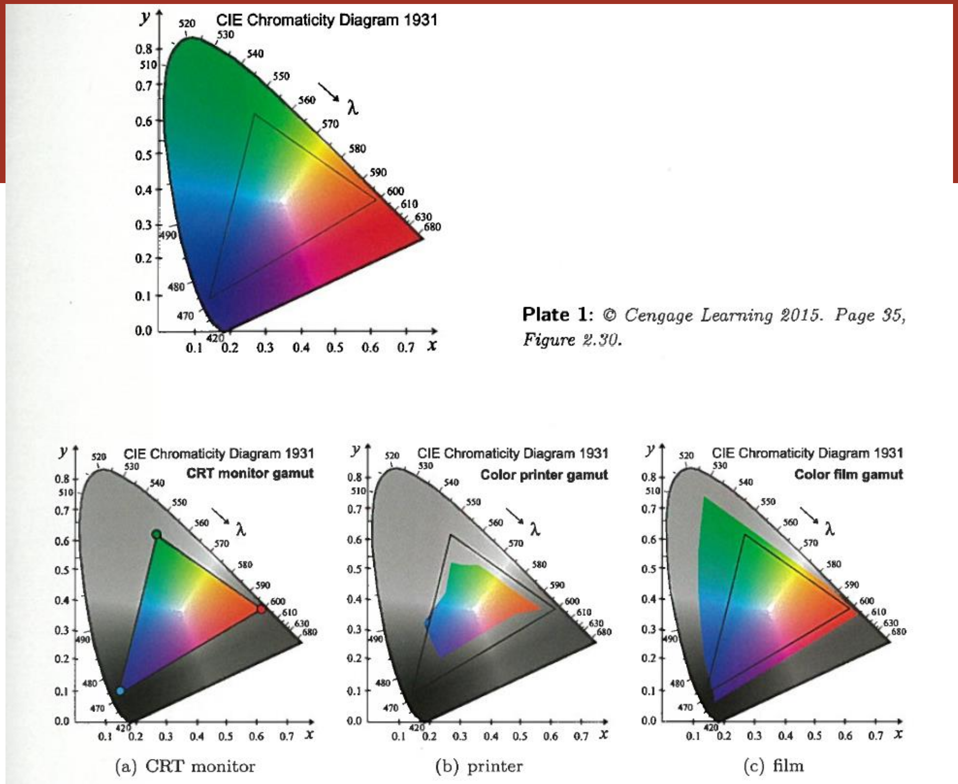

Plate 1: © *Cengage Learning 2015. Page 35, Figure 2.30.*