INTRODUCTION

Aziz M. Qaroush

Birzeit University

First Semester
2016 - 2017

Outline

- What is a digital image?
- What is digital image processing?
- State of the art examples of digital image processing
- Key stages in digital image processing
- Elements of Visual Perception
- Light and the Electromagnetic Spectrum
- Image Sensing and Acquisition
- Image Sampling and Quantization
- Image Interpolation
- Math Preliminary Related to Image Processing

What is a Digital Image?

- A **digital image** is a representation of a two- dimensional image as a finite set of digital values, called picture elements or pixels.
  - An **image can be define as a two-dimensional function f(x,y) with x and y being the spatial coordinates and f is the amplitude**
  - A **digital image is the representation of an image using finite and discrete values for x,y, and f**
  - These values are called picture elements or pixels

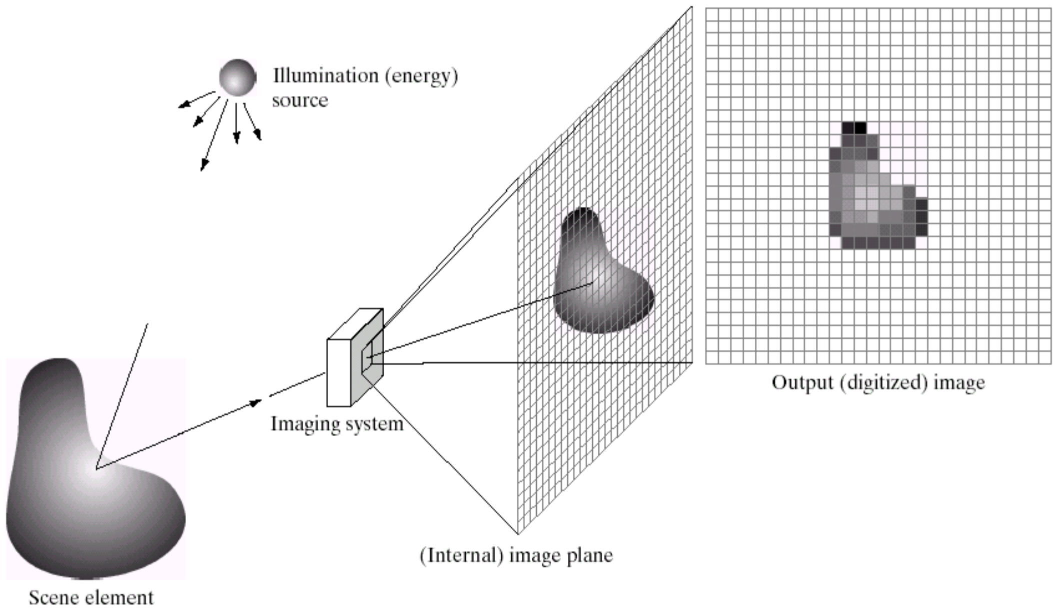

What is a Digital Image?

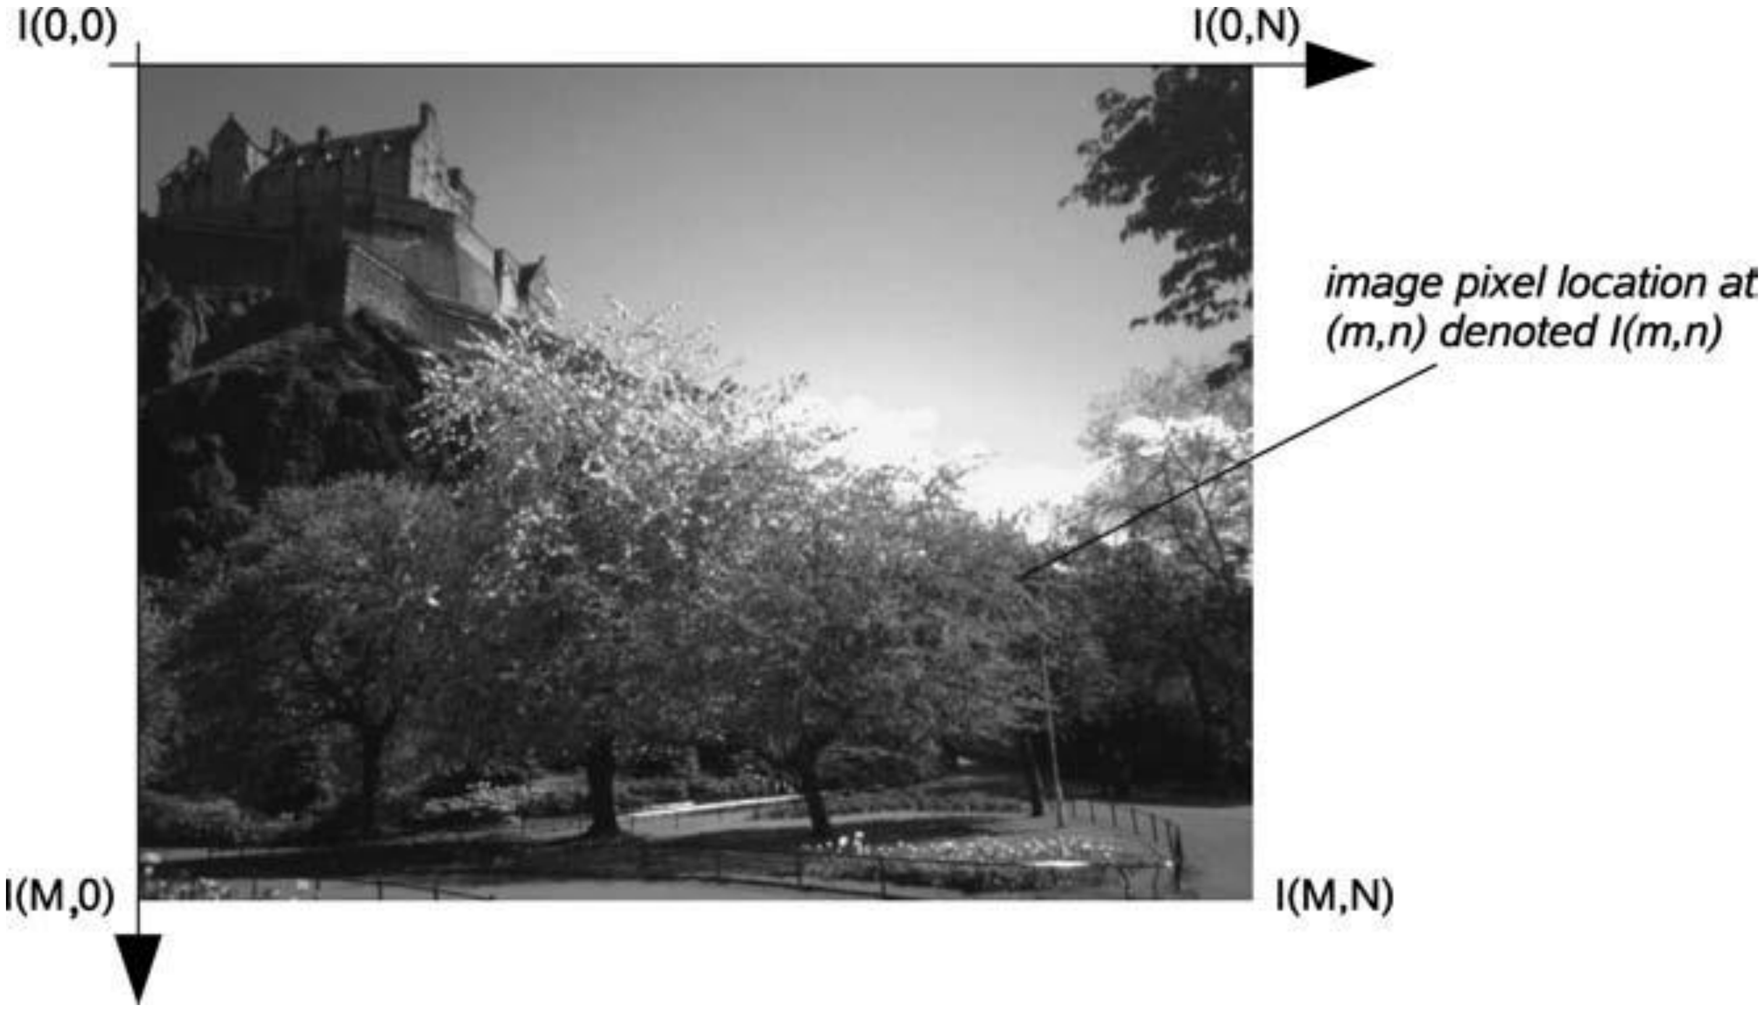

What is a Digital Image? (cont…)

- Pixel values typically represent gray levels, colors, heights, opacities etc
- **Remember** *digitization* implies that a digital image is an *approximation* of a real scene

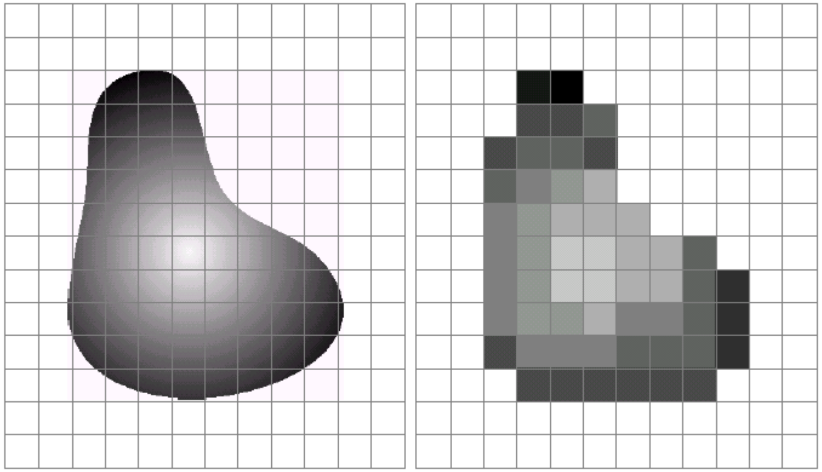

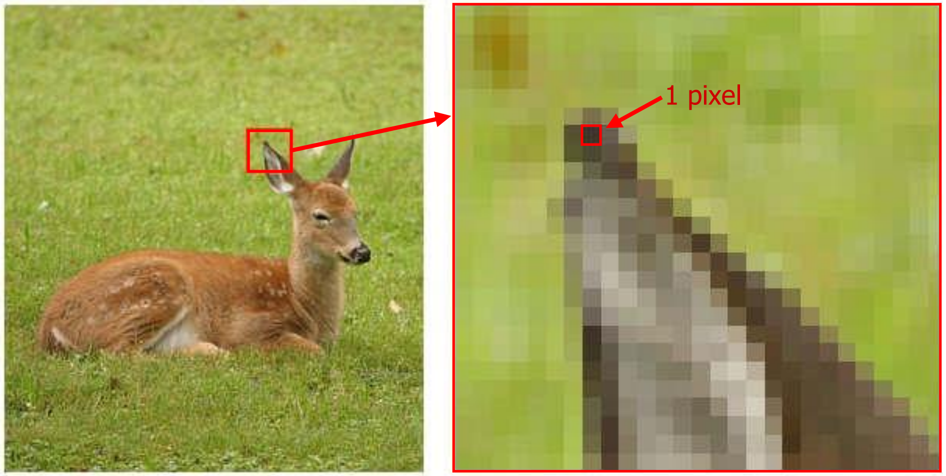

What is a Digital Image? (cont…)

- Common image formats include:
  - 1 sample per point (Grayscale)
  - 3 samples per point (Red, Green, and Blue)
  - 4 samples per point (Red, Green, Blue, and “Alpha”, a.k.a. Opacity)
  - 5 samples per point (add “Depth”)



- For most of this course we will focus on grey-scale images

Image formats

Acronym	Name	Properties
GIF	Graphics interchange format	Limited to only 256 colours (8 bit); lossless compression
JPEG	Joint Photographic Experts Group	In most common use today; lossy compression; lossless variants exist
BMP	Bit map picture	Basic image format; limited (generally) lossless compression; lossy variants exist
PNG	Portable network graphics	New lossless compression format; designed to replace GIF
TIF/TIFF	Tagged image (file) format	Highly flexible, detailed and adaptable format; compressed/uncompressed variants exist

What is Digital Image Processing?

- Digital image processing focuses on two major tasks
  - Improvement of pictorial information for human interpretation
  - Processing of image data for storage, transmission and representation for autonomous machine perception

- Some argument about where image processing ends and fields such as image analysis and computer vision start

Examples: Image Enhancement

- One of the most common uses of DIP techniques: improve quality, remove noise etc

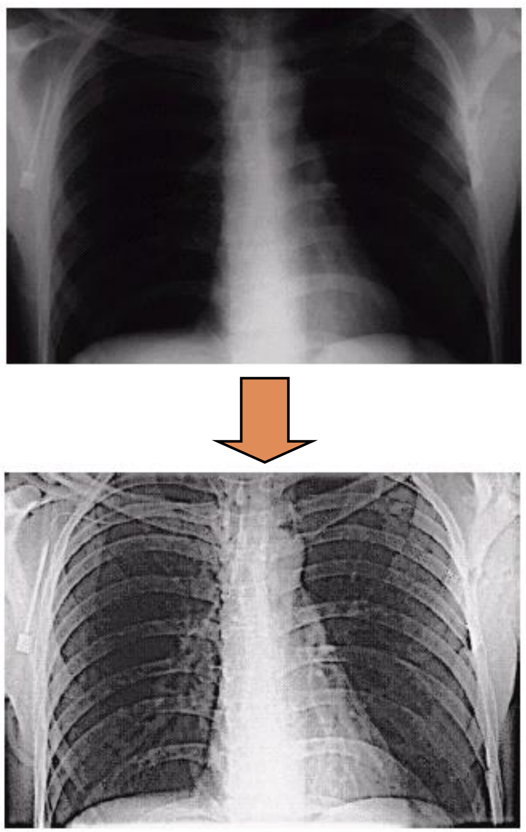

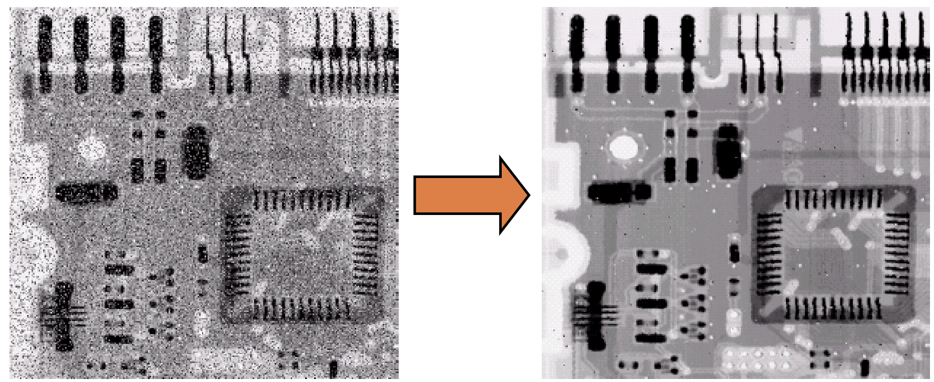

Examples: Medicine

- Take slice from MRI scan of canine heart, and find boundaries between types of tissue
  - Image with gray levels representing tissue density
  - Use a suitable filter to highlight edges

Original MRI Image of a Dog Heart

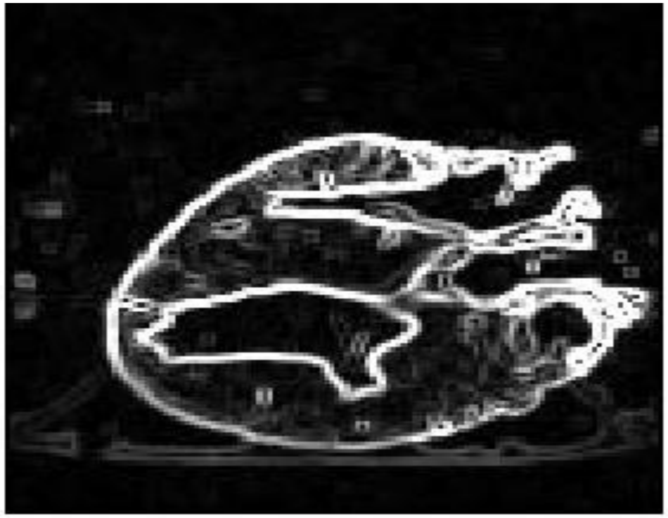

Edge Detection Image

Examples: OCR

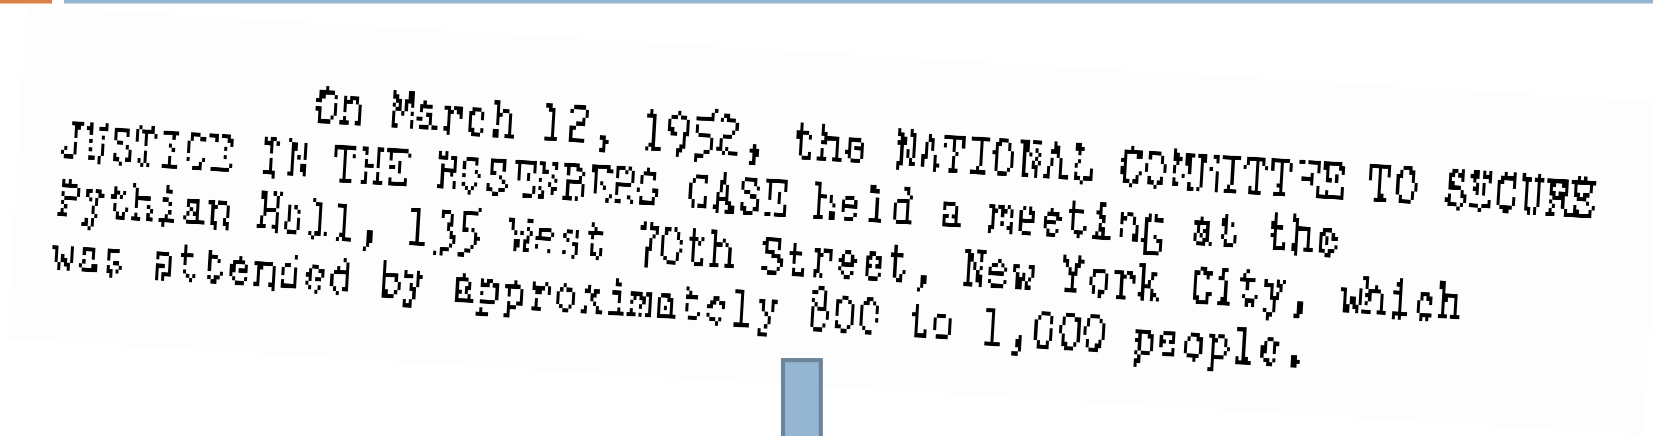
On March 12, 1952, the NATIONAL COMMITTEE TO SECURE
JUSTICE IN THE ROSENBERG CASE held a meeting at the
Pythian Hall, 135 West 70th Street, New York City, which
was attended by approximately 800 to 1,000 people.

On March 12, 1952, the NATIONAL COMMITTEE TO SECURE JUSTICE IN THE ROSENBERG CASE held a meeting at the Pythian Hall, 135 West 70th Street, New York City, which was attended by approximately 800 to 1,000 people.

Examples: The Hubble Telescope

- Launched in 1990 the Hubble telescope can take images of very distant objects
- However, an incorrect mirror made many of Hubble's images useless
- Image processing techniques were used to fix this

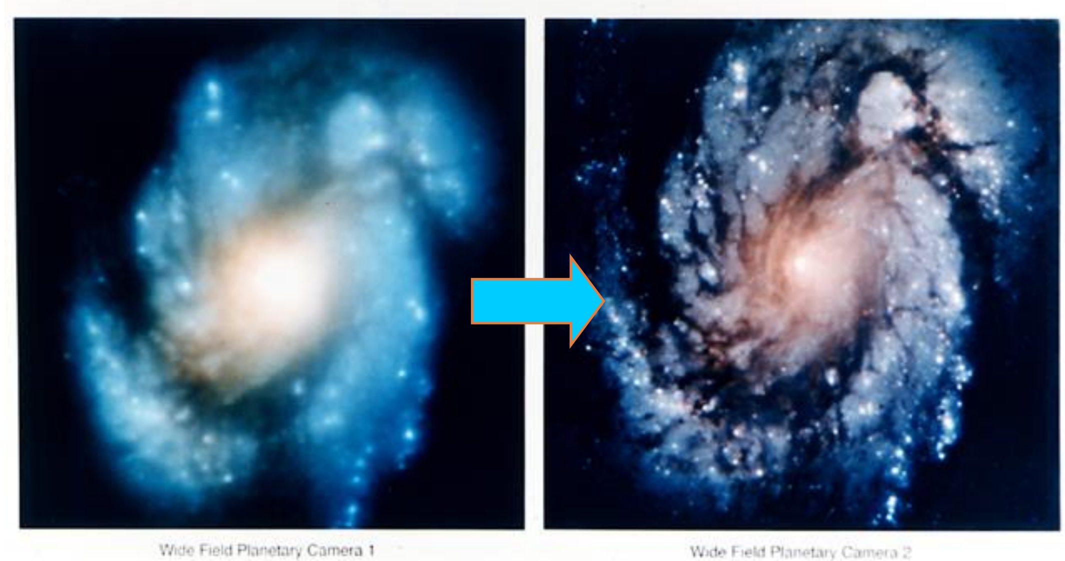

Examples: Artistic Effects

- Artistic effects are used to make images more visually appealing, to add special effects and to make composite images

Examples: GIS

- Geographic Information Systems
  - Digital image processing techniques are used extensively to manipulate satellite imagery
  - Terrain classification
  - Meteorology

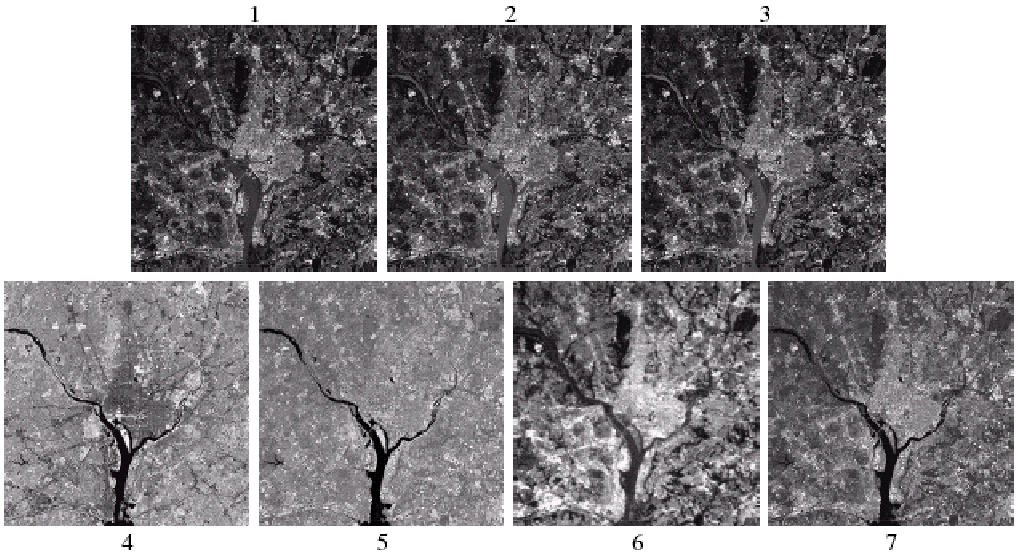

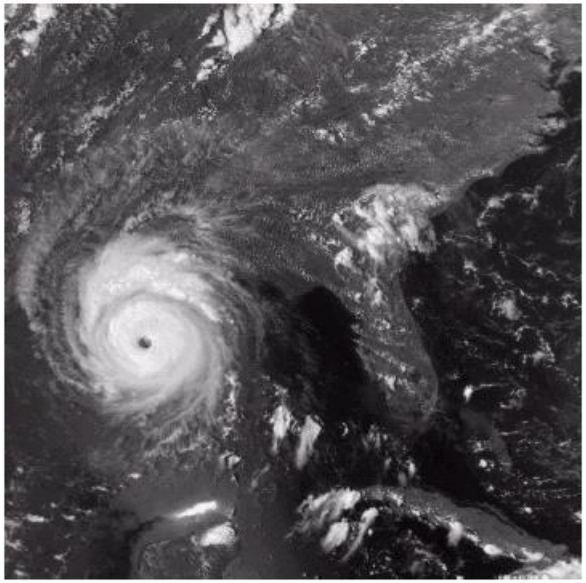

Examples: GIS (cont…)

- *Night-Time Lights of the World* data set
  - Global inventory of human settlement
  - Not hard to imagine the kind of analysis that might be done using this data

Examples: Industrial Inspection

- Human operators are expensive, slow and unreliable
- Make machines do the job instead
- Industrial vision systems are used in all kinds of industries
- Can we trust them?

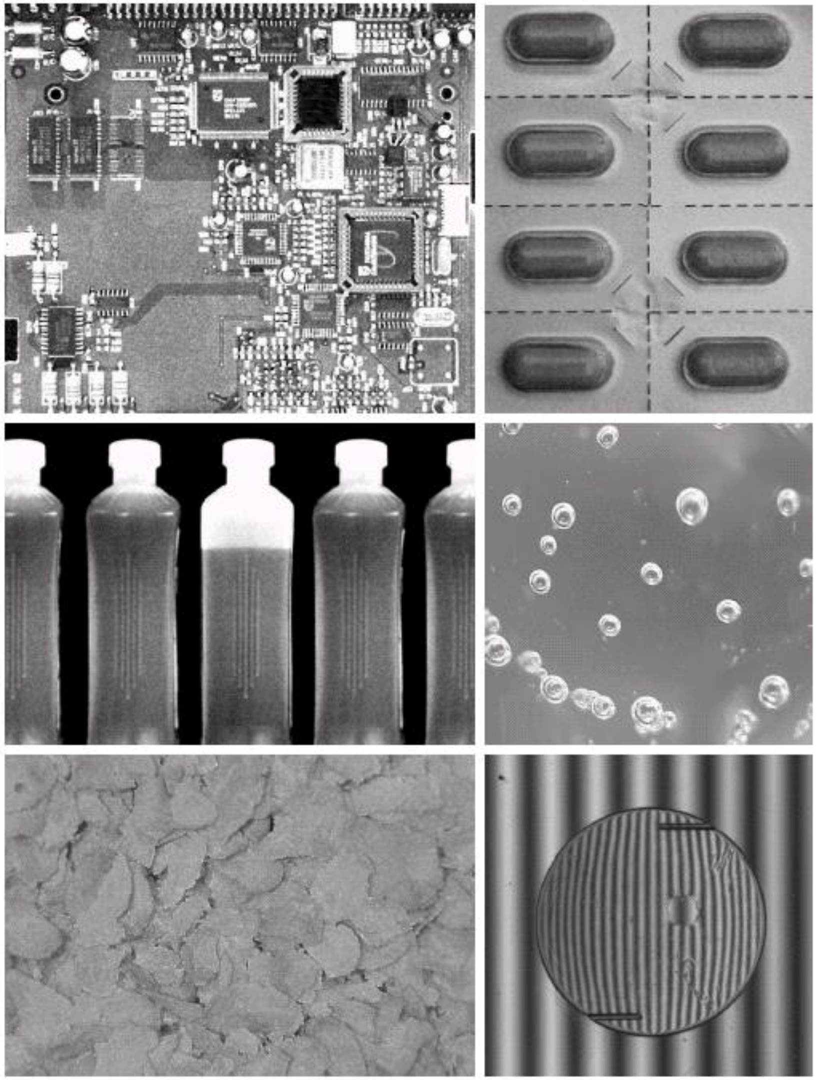

Examples: PCB Inspection

- Printed Circuit Board (PCB) inspection
  - Machine inspection is used to determine that all components are present and that all solder joints are acceptable
  - Both conventional imaging and x-ray imaging are used

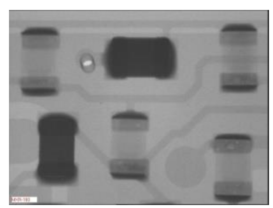

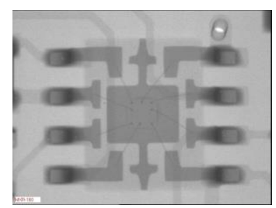

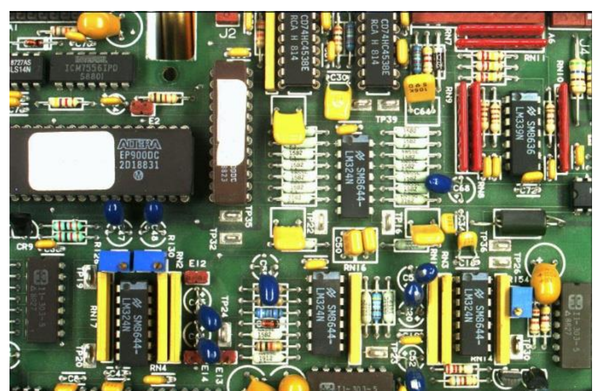

Examples: Law Enforcement

- Image processing techniques are used extensively by law enforcers
  - Number plate recognition for speed cameras/automated toll systems
  - Fingerprint recognition
  - Enhancement of CCTV images

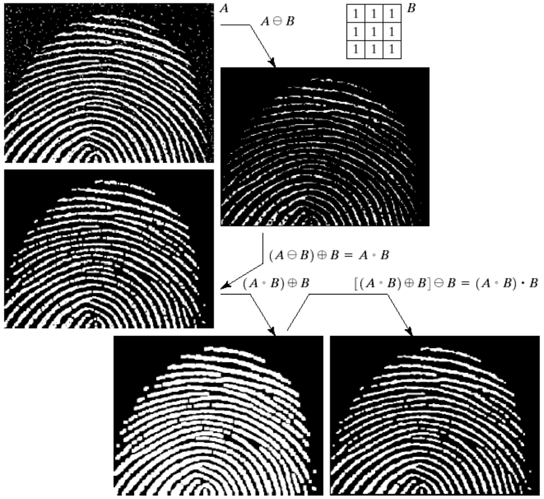

Examples: HCI

- Try to make human computer interfaces more natural
  - Face recognition
  - Gesture recognition
- Does anyone remember the user interface from "Minority Report"?
- These tasks can be extremely difficult

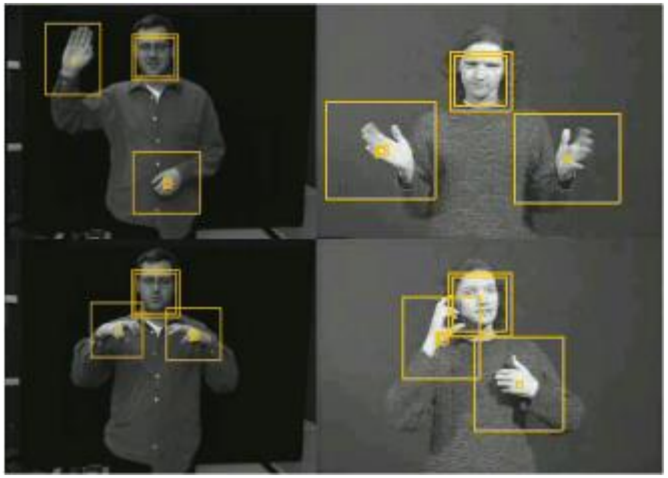

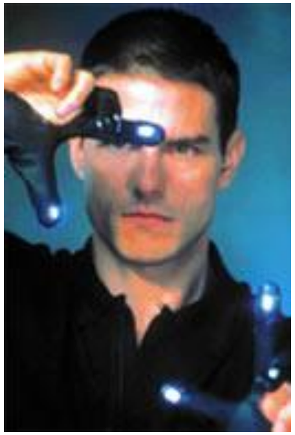

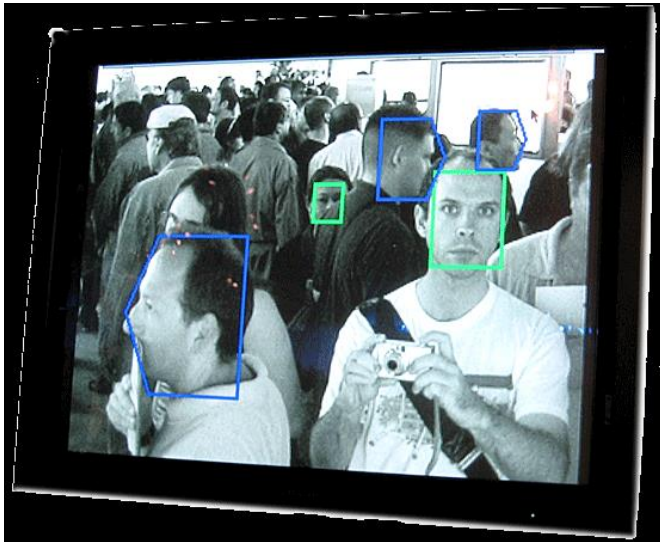

Relationship with other Fields

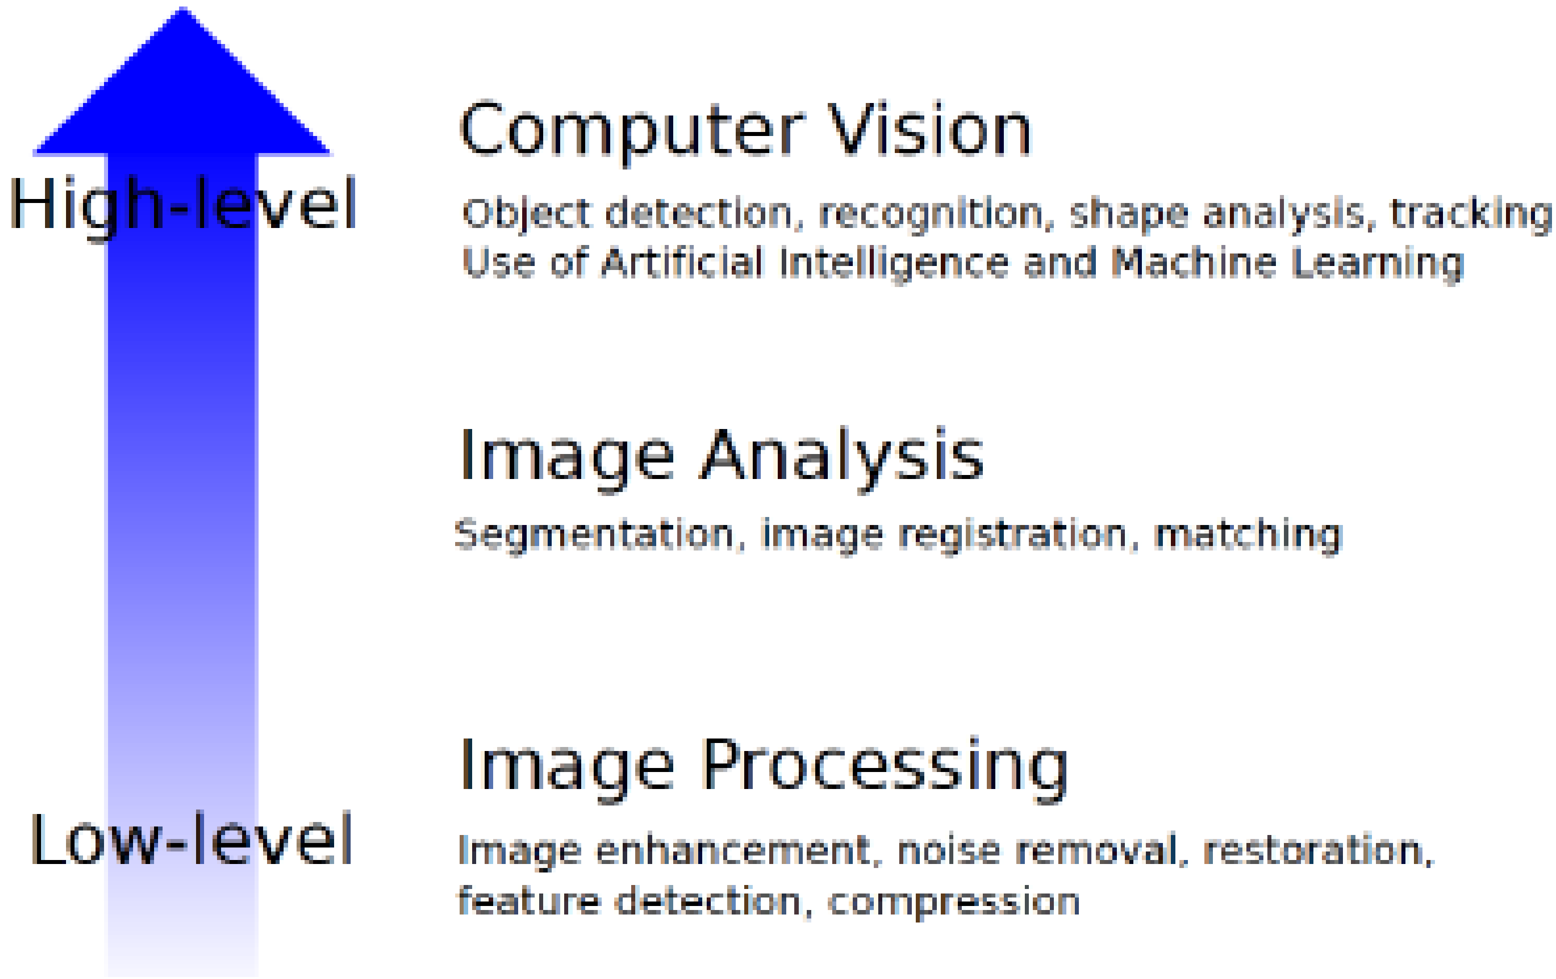

Key Stages in Digital Image Processing

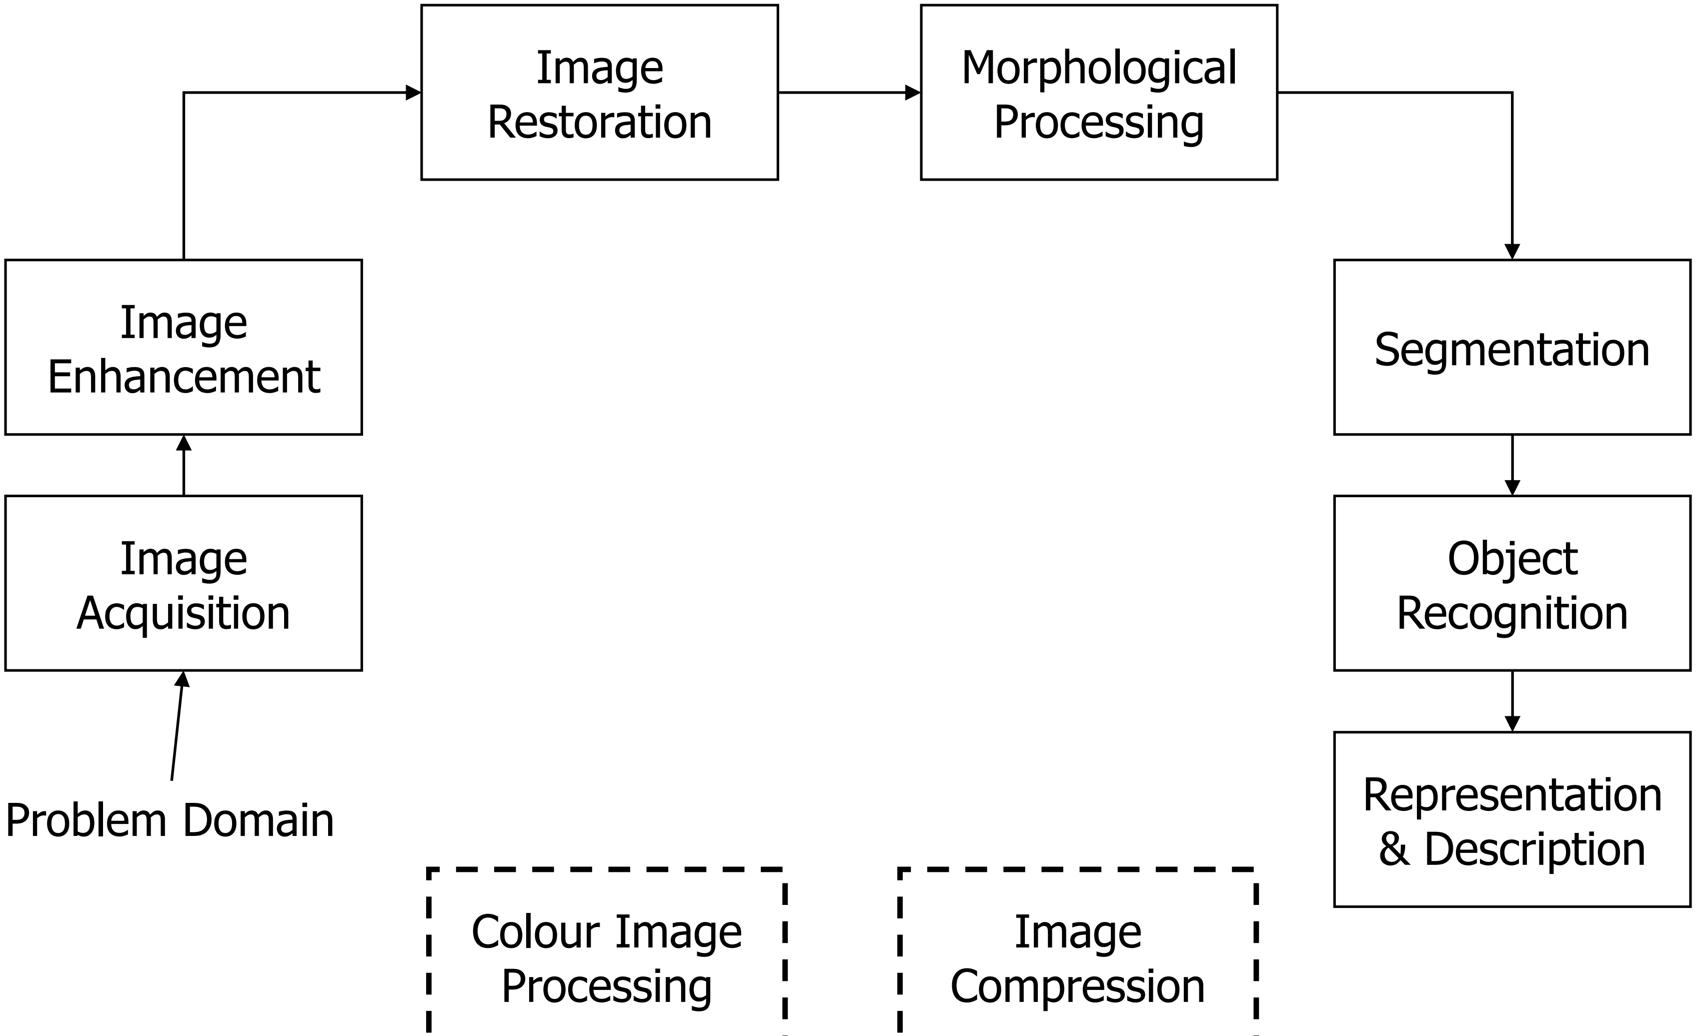

Key Stages in Digital Image Processing: Image Aquisition

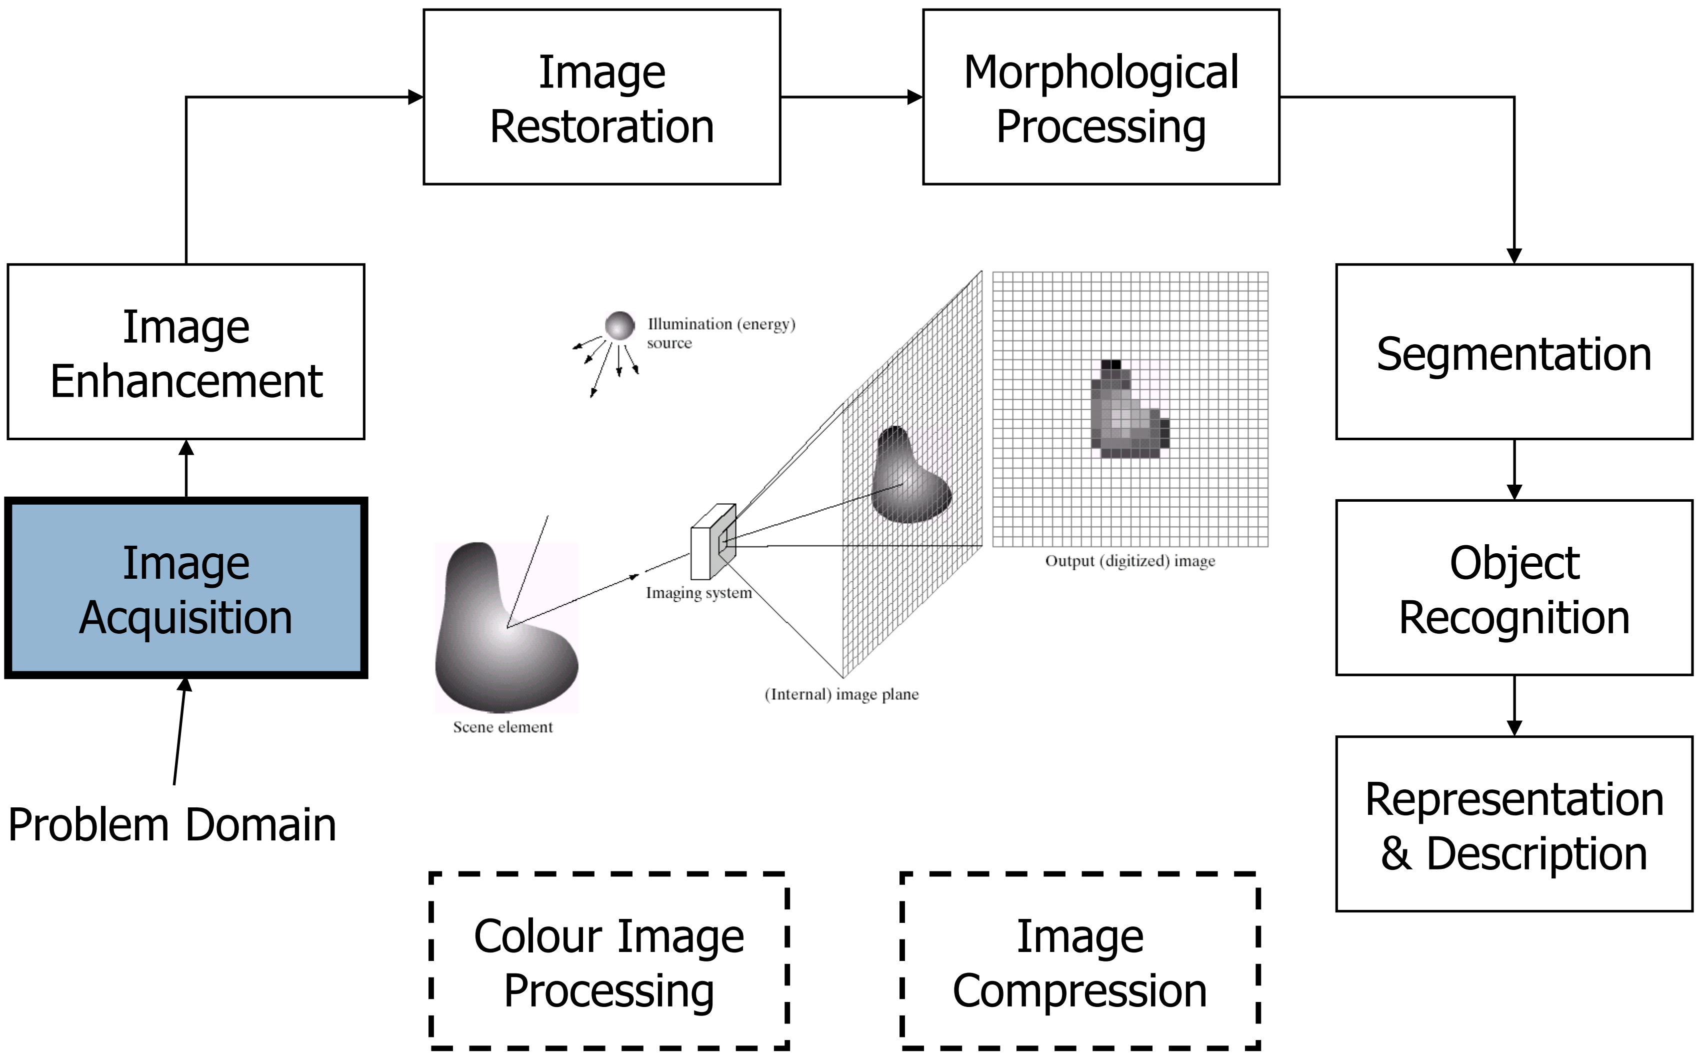

Key Stages in Digital Image Processing: Image Enhancement

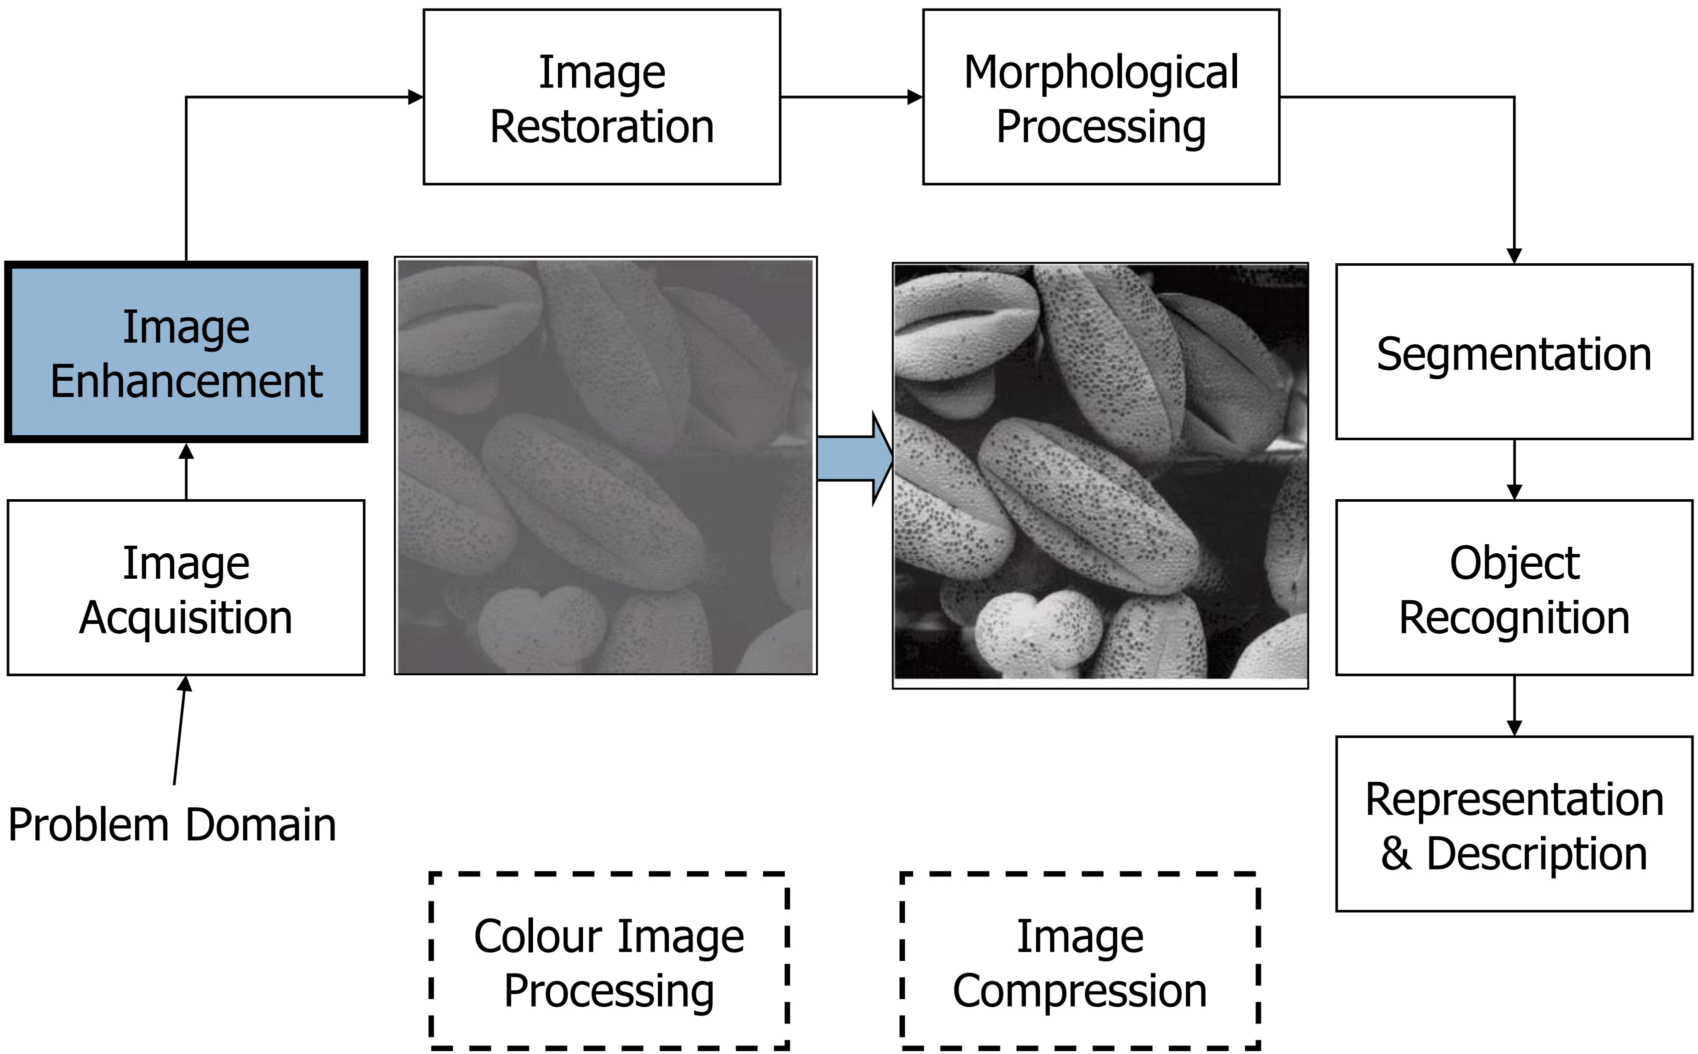

Key Stages in Digital Image Processing: Image Restoration

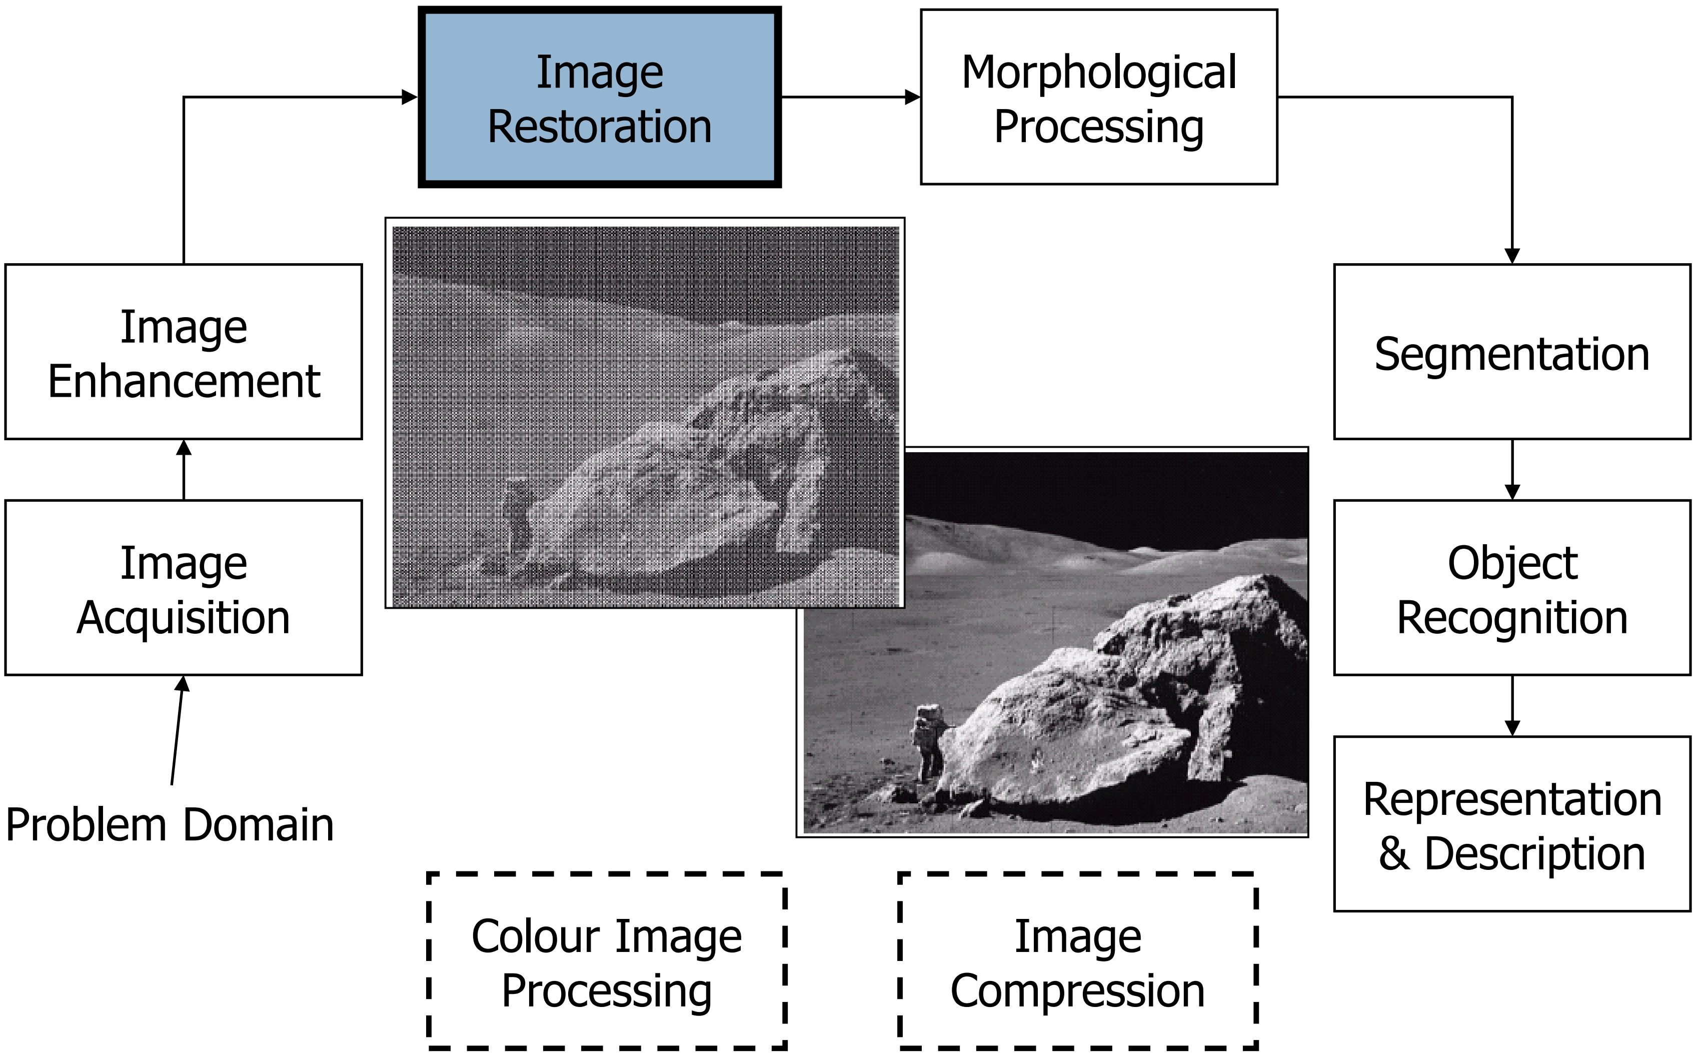

Key Stages in Digital Image Processing: Morphological Processing

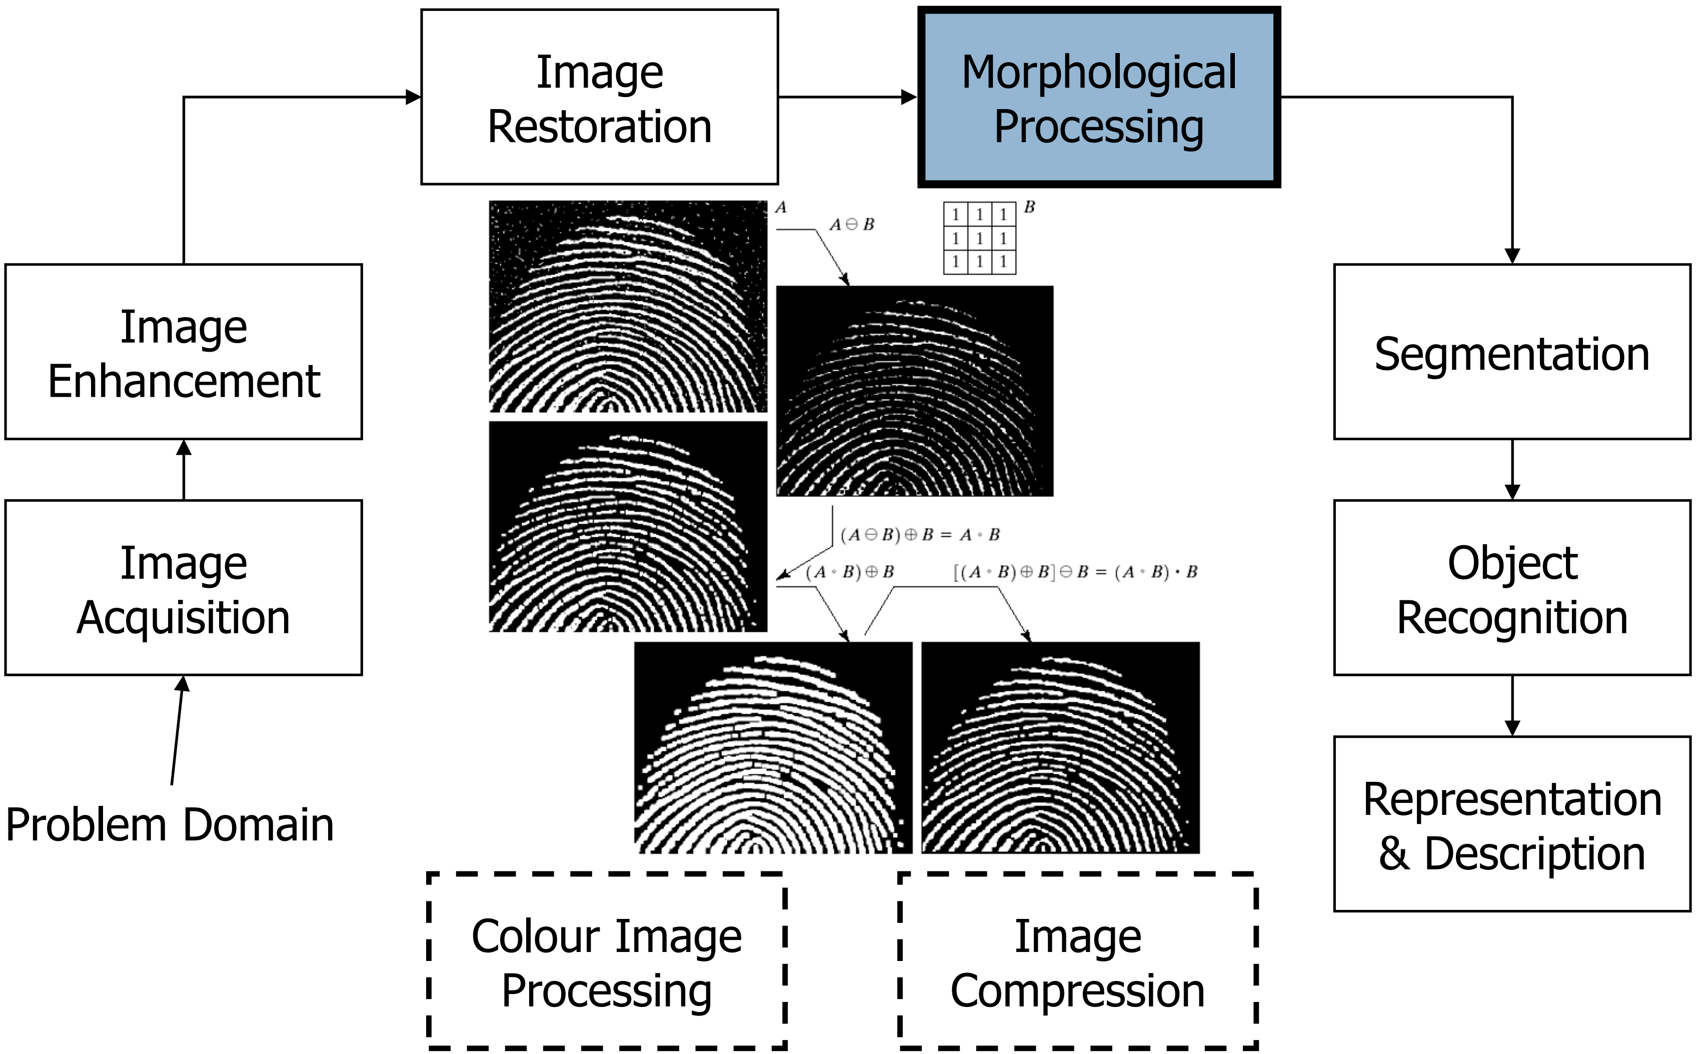

Key Stages in Digital Image Processing: Segmentation

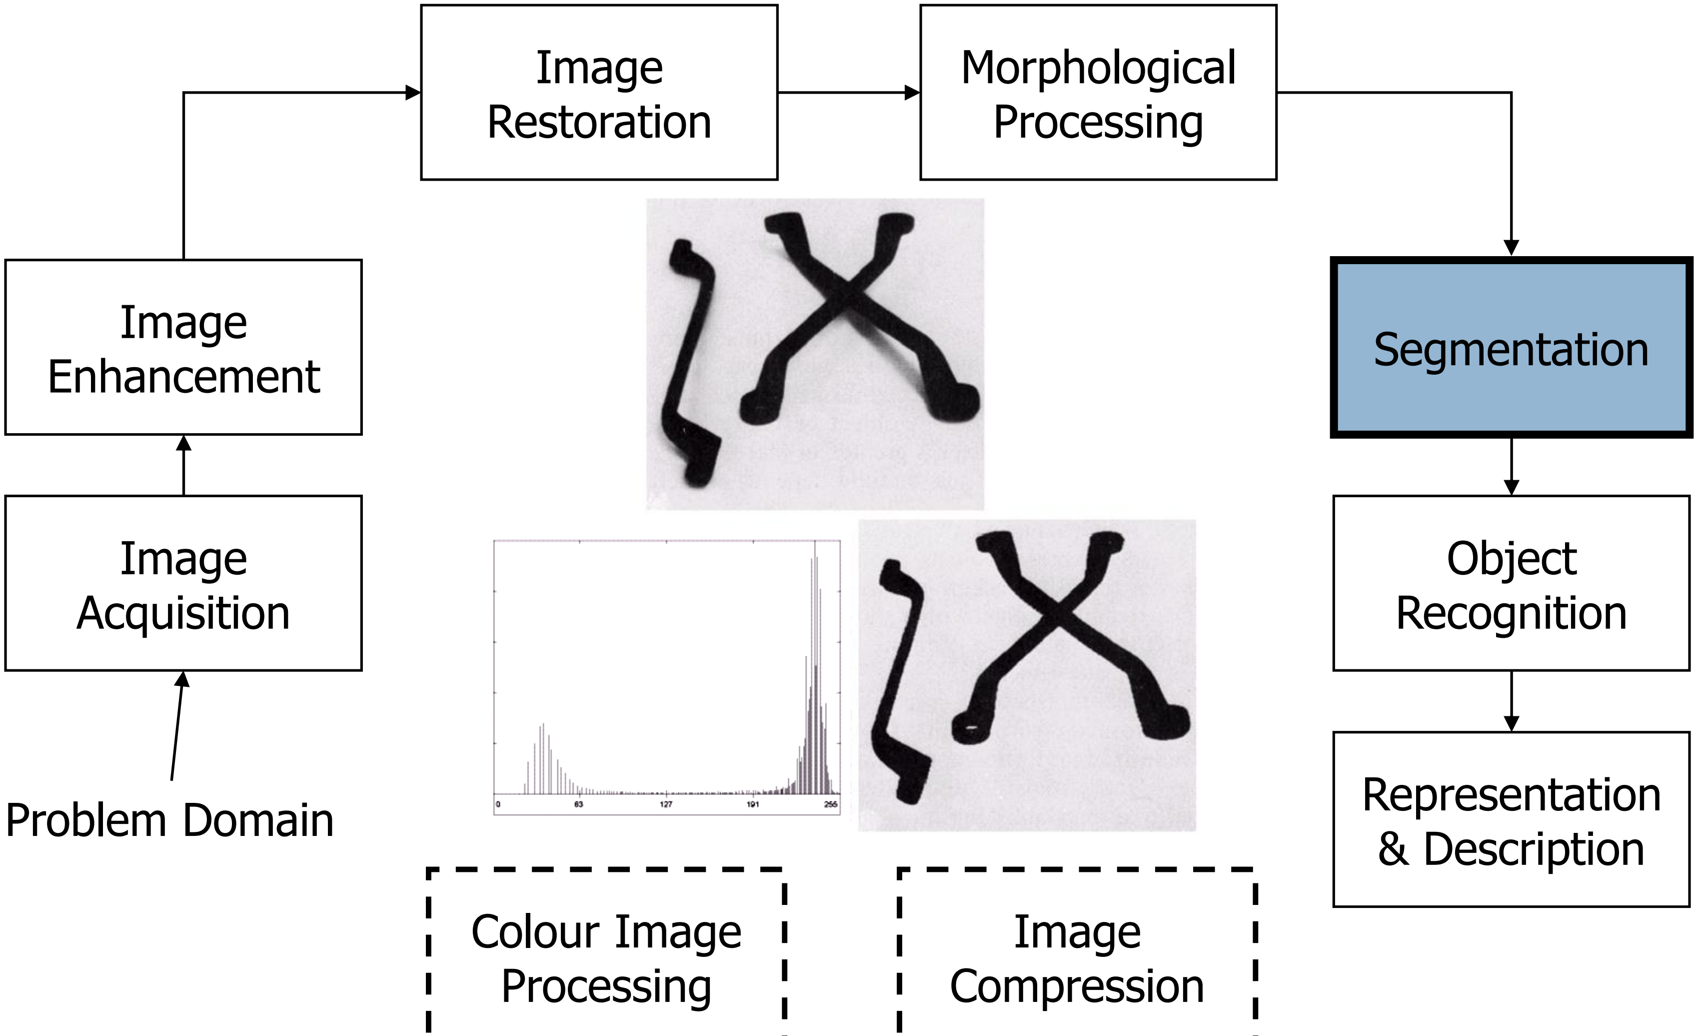

Key Stages in Digital Image Processing: Object Recognition

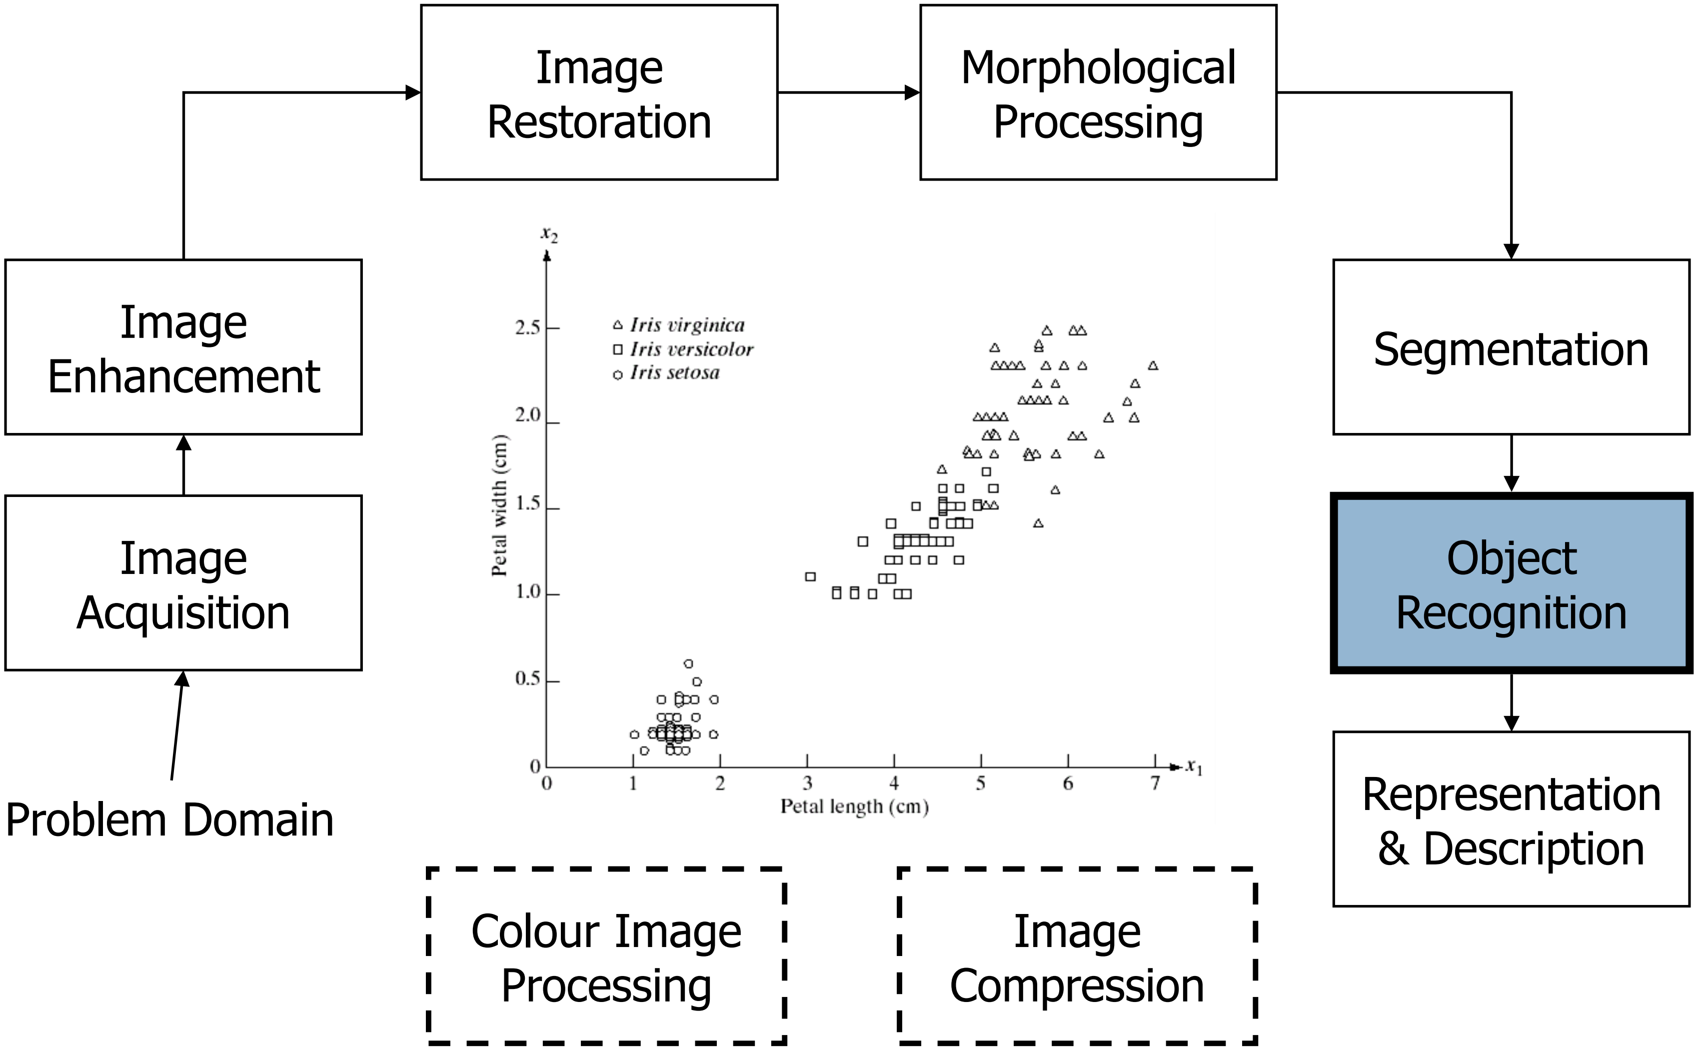

Key Stages in Digital Image Processing: Representation & Description

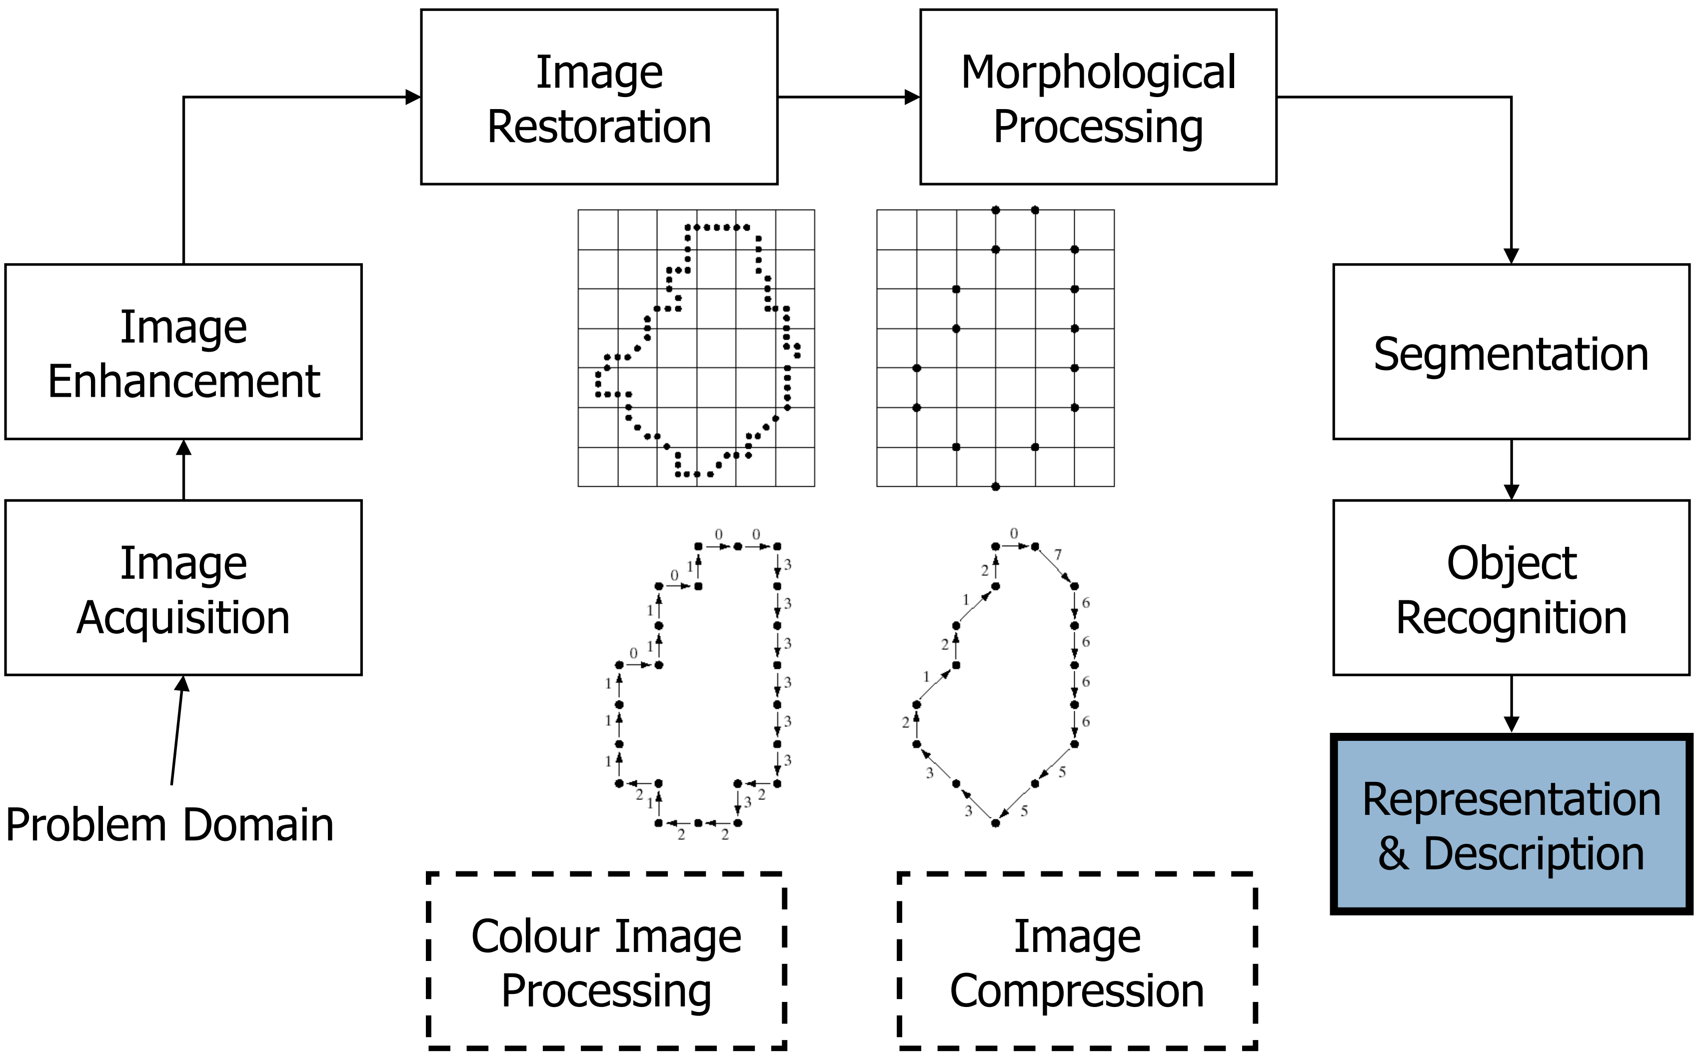

Key Stages in Digital Image Processing: Image Compression

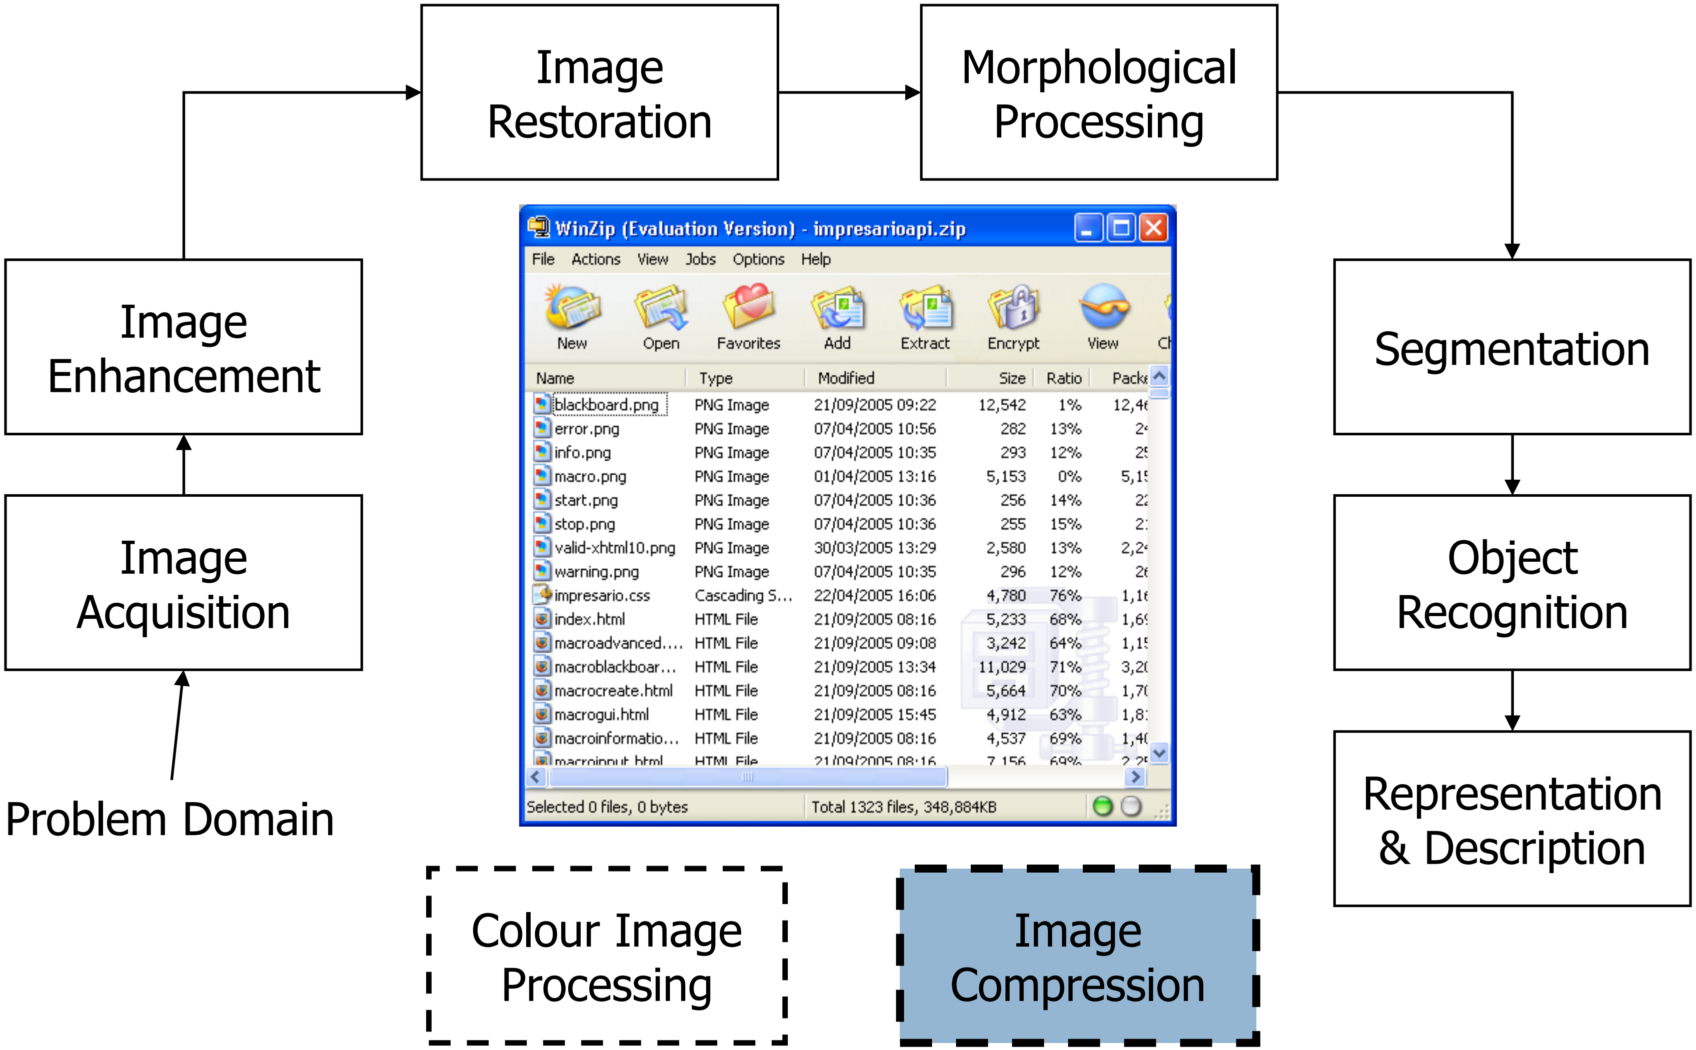

Key Stages in Digital Image Processing: Colour Image Processing

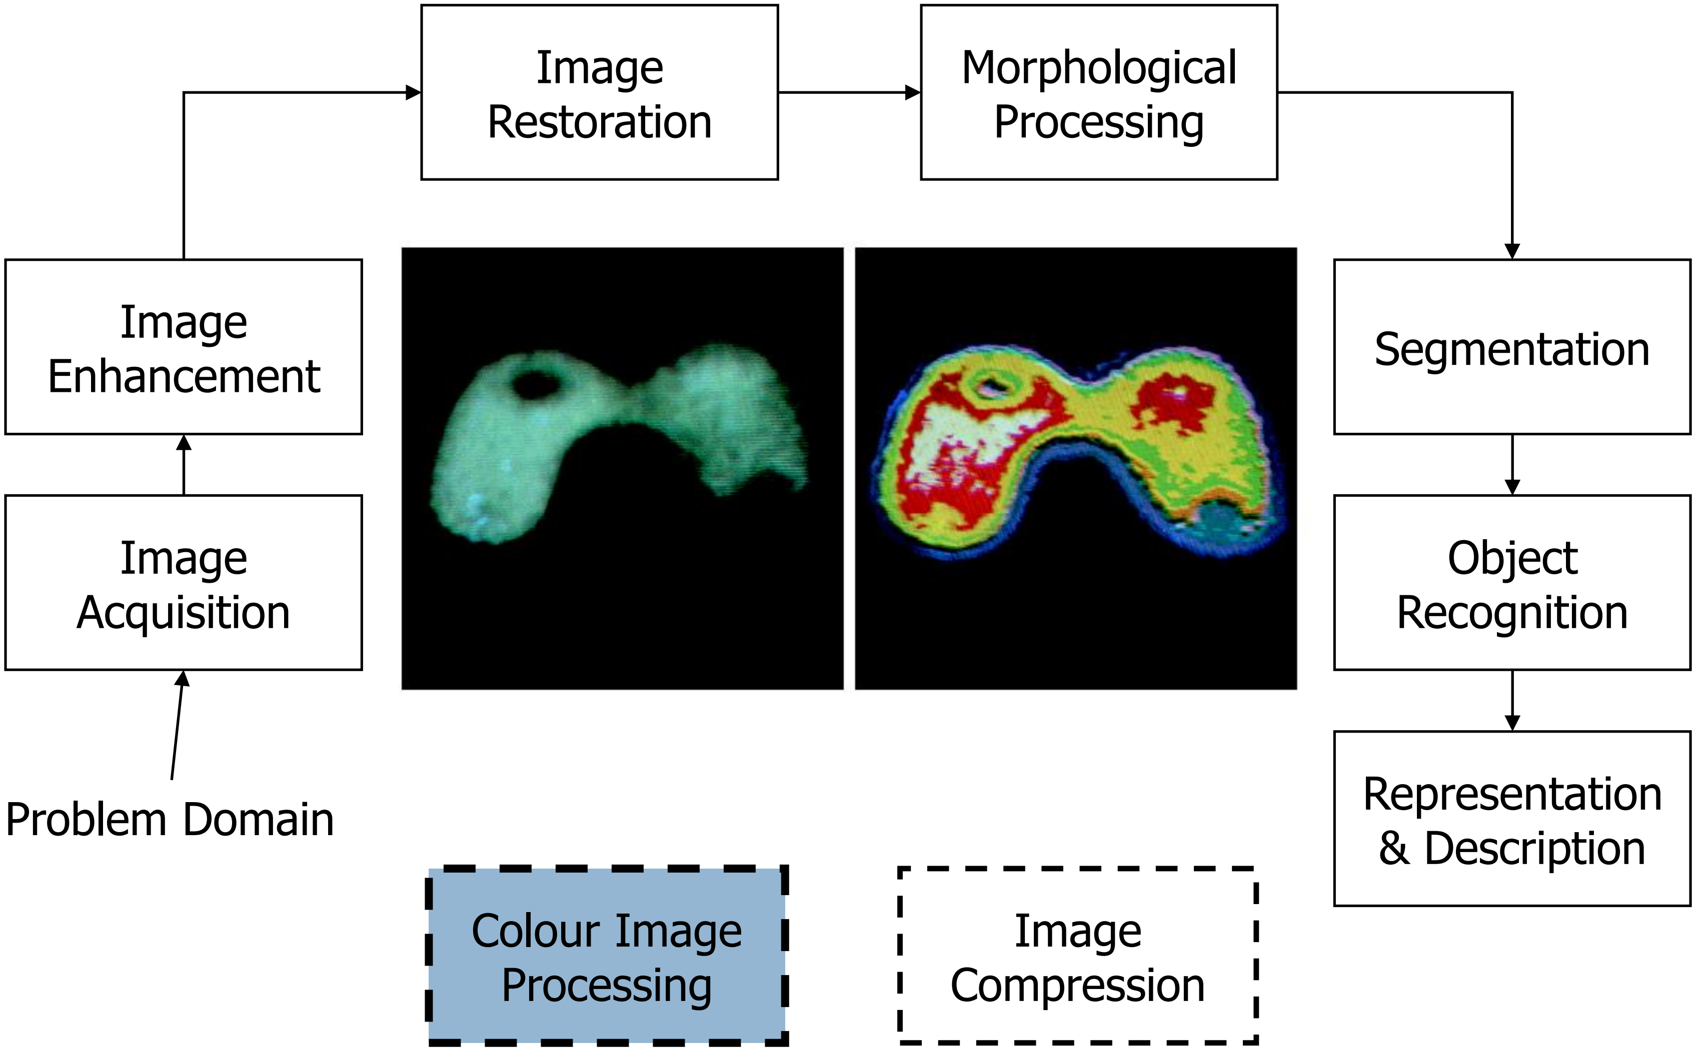

Structure of the Human Eye

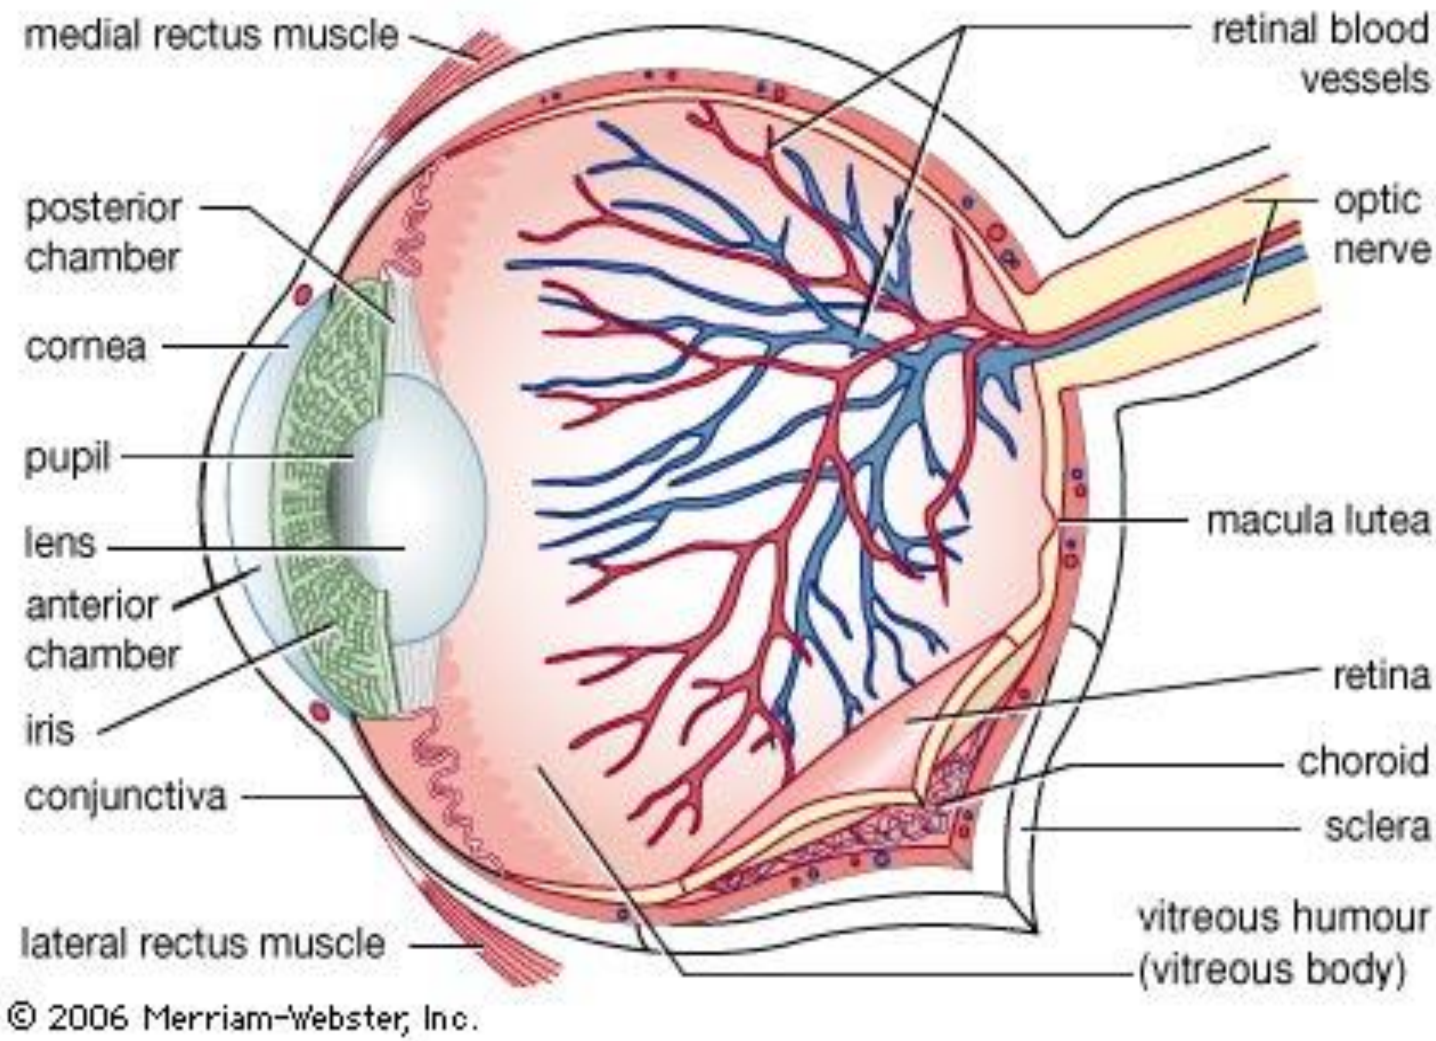

Structure of the Human Eye

- **Cornea**
  - The cornea is a strong clear bulge located at the front of the eye (where it replaces the sclera - that forms the outside surface of the rest of the eye). The front surface of the adult cornea has a radius of approximately 8mm. The cornea contributes to the image-forming process by refracting light entering the eye.
- **Fovea**
  - The fovea is a small depression (approx. 1.5 mm in diameter) in the retina. This is the part of the retina in which high-resolution vision of fine detail is possible.
- **Iris**
  - The iris is a diaphragm of variable size whose function is to adjust the size of the pupil to regulate the amount of light admitted into the eye. The iris is the coloured part of the eye (illustrated in blue above but in nature may be any of many shades of blue, green, brown, hazel, or grey).
- **Lens**
  - The lens of the eye is a flexible unit that consists of layers of tissue enclosed in a tough capsule. It is suspended from the ciliary muscles by the zonule fibers.

Structure of the Human Eye

- **Retina**
  - The retina may be described as the "screen" on which an image is formed by light that has passed into the eye via the cornea, aqueous humour, pupil, lens, then the hyaloid and finally the vitreous humour before reaching the retina. The retina contains photosensitive elements (called rods and cones) that convert the light they detect into nerve impulses that are then sent onto the brain along the optic nerve.
  - **Cones:**
    - 6-7 million primarily at fovea
    - Responsible for fine details as each cone has its own nerve
    - Highly sensitive to color
    - Photopic or bright-light vision
  - **Rods**
    - 75 to 125 million all over the retinal surface
    - Several rods share the same nerve
    - Not involved in color vision and sensitive to low illumination
    - Scotopic or dim-light vision

Structure of the Human Eye

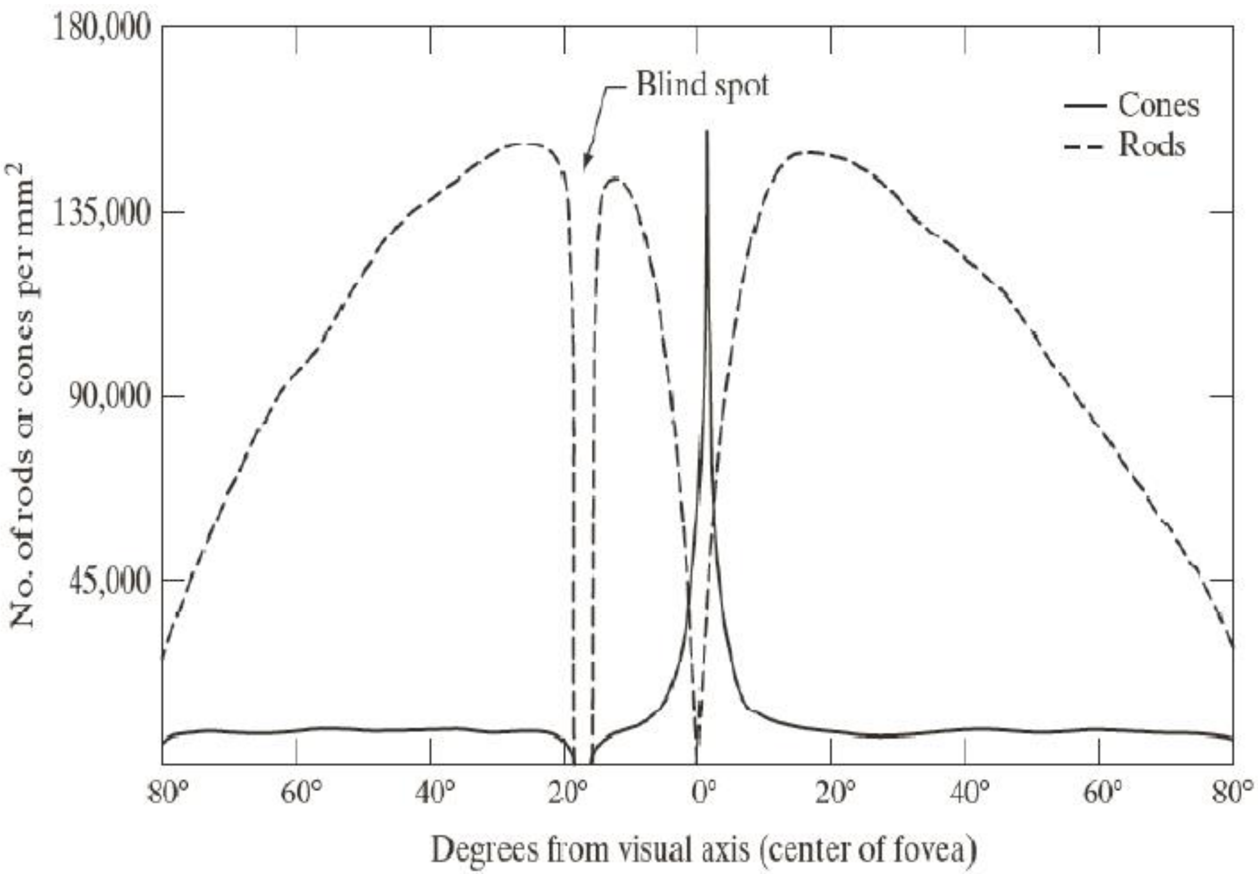

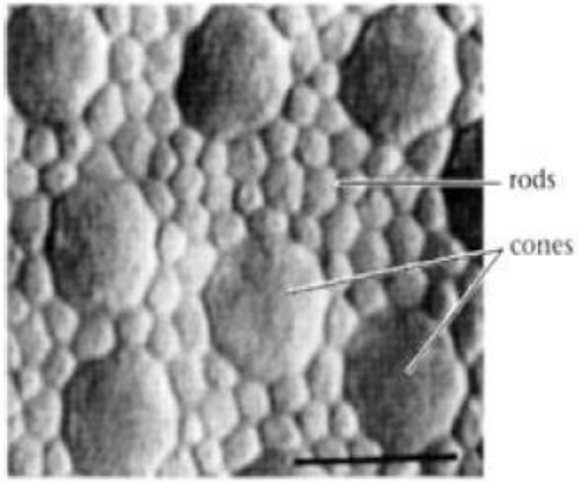

Distribution of rods and cones on the retina

Image Formation in the Eye

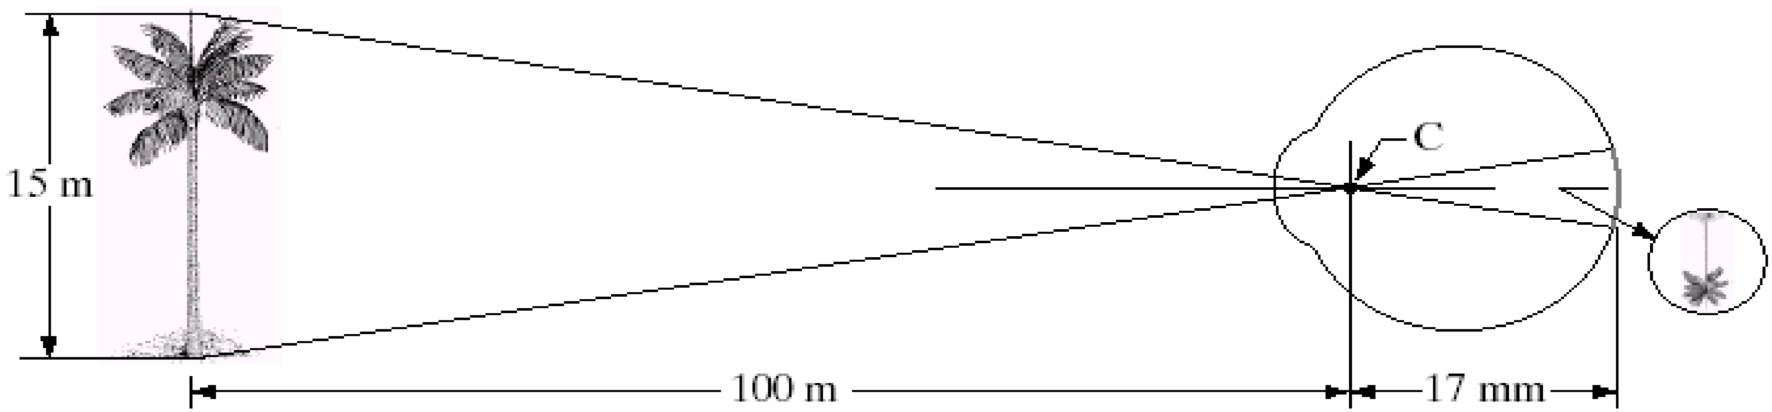

- The shape of the lens changes based on the object distance
  - Close objects; thickened with max refraction
  - Far objects; flattened with min refraction
- Size of object image on the retina
  - Focal length 14 – 17 mm

Brightness Adaptation

- The human eye can adapt to a large range of light intensity levels
- However, the eye can not operate simultaneously over this range
- Instead, based on the current lighting conditions, it adapts to a smaller range

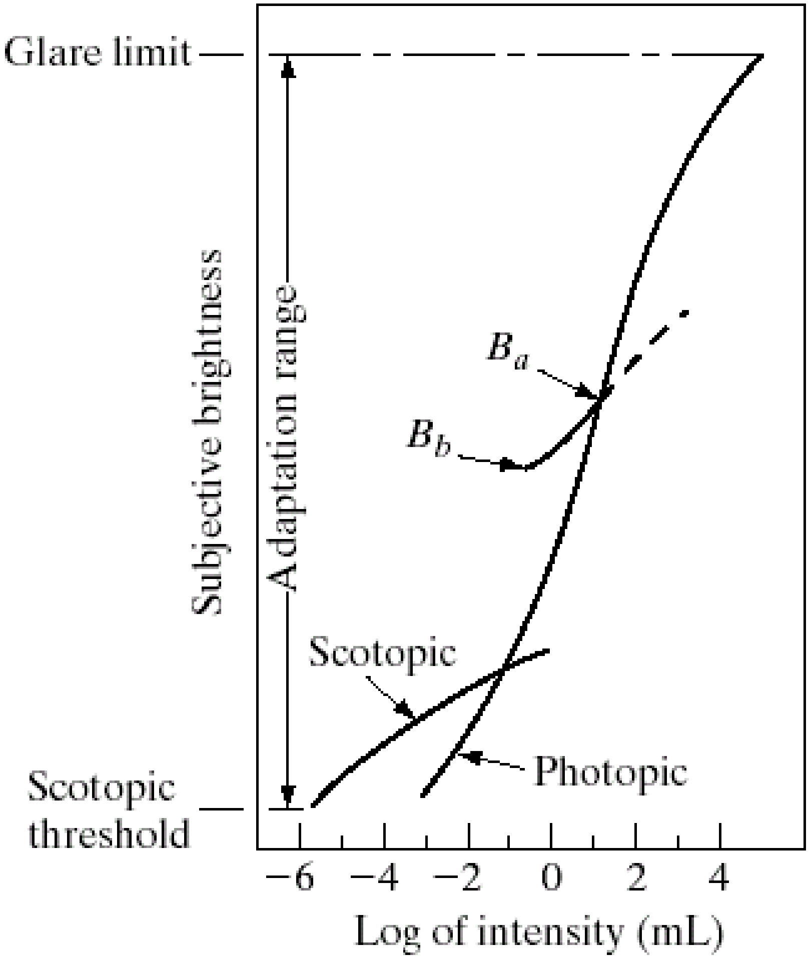

Brightness Discrimination

Simultaneous Contrast

Perceived brightness doesn't depend on its intensity only

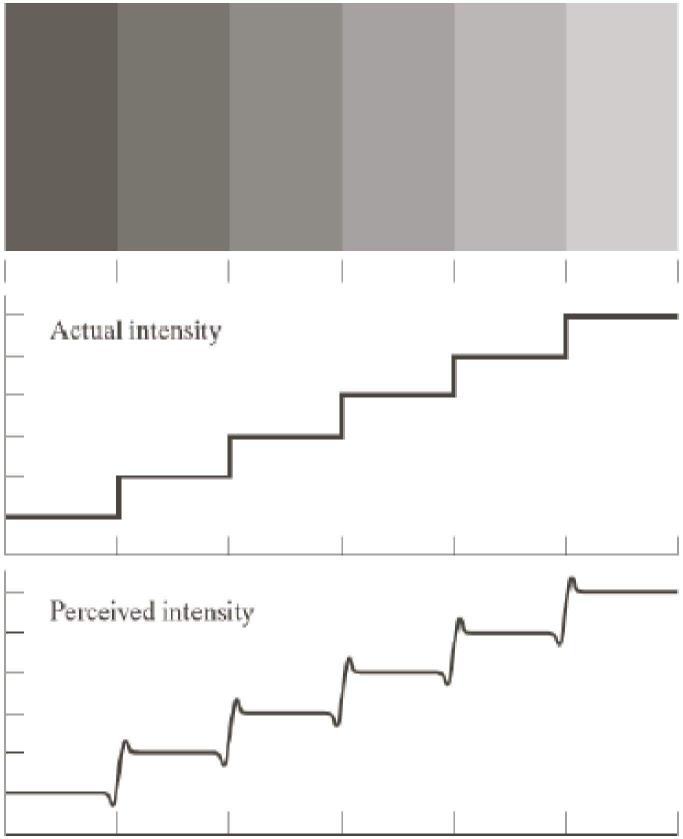

Mach bands

Visual system to overshoot or undershoot around the borders

Optical Illusions

- The eye fills nonexistent information or wrongly perceives geometrical properties

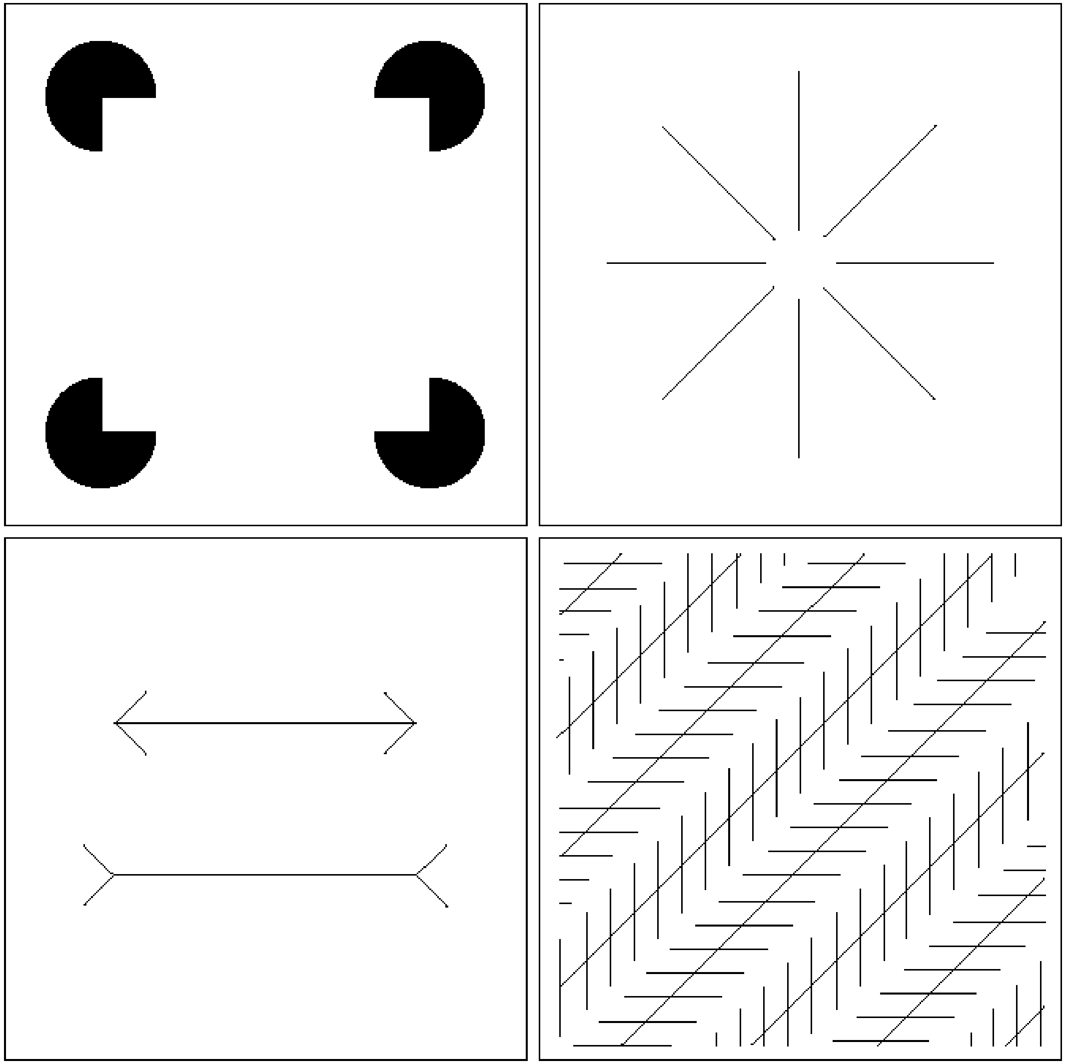

Light and Electromagnetic Spectrum

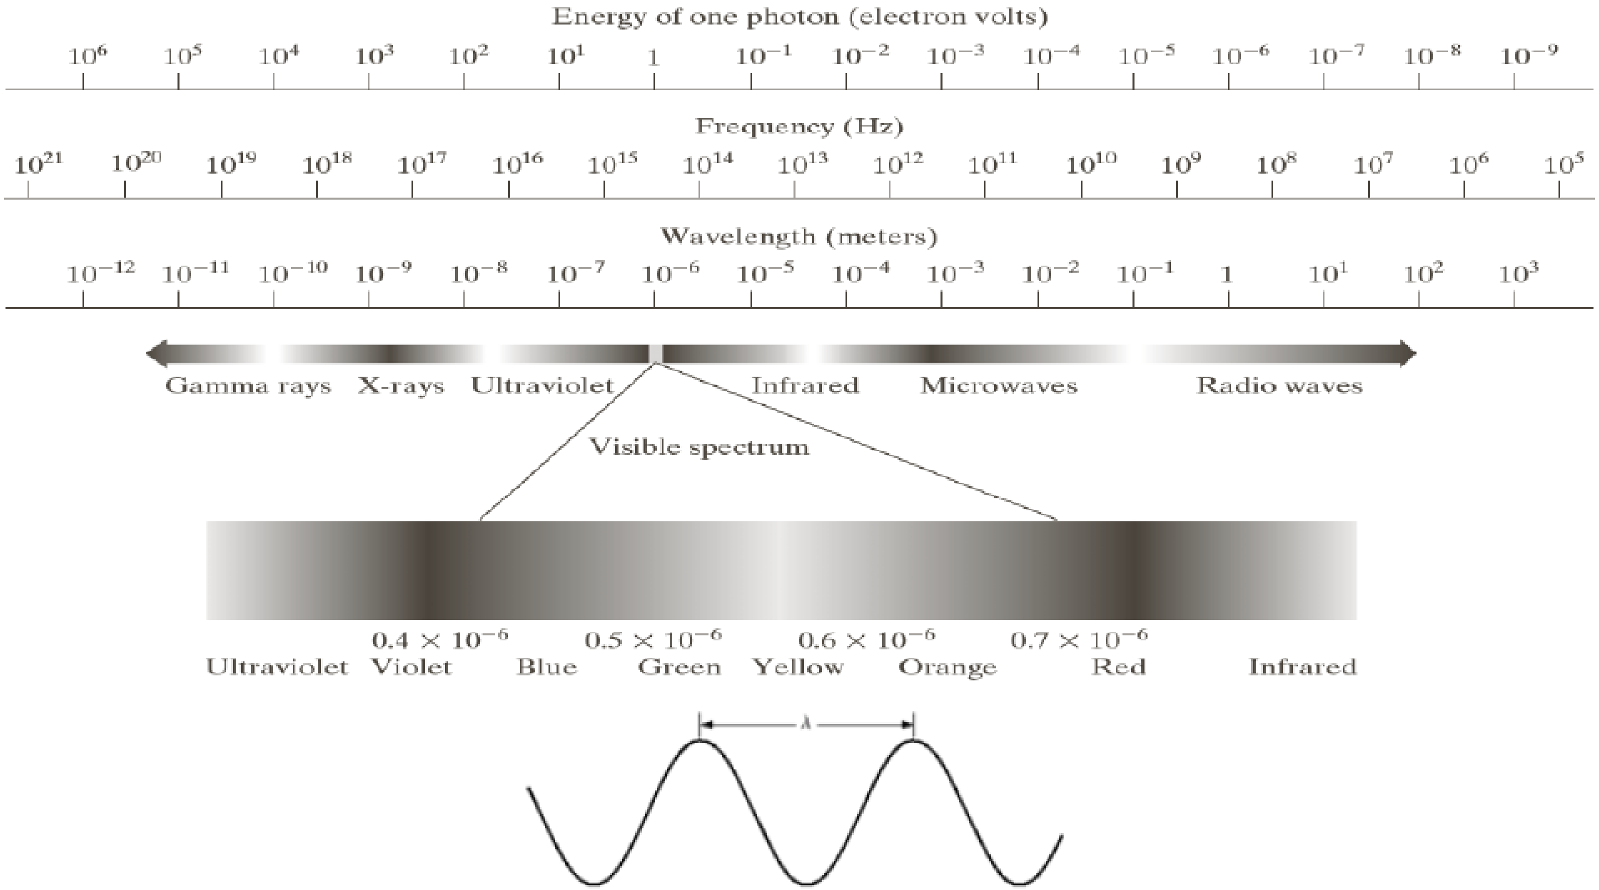

$$\lambda \times v = c$$

$$c = 2.998 \times 10^{8}\ m/s$$

$$E = h \times v$$

$$h = 6.626068 \times 10^{-34}\ m^2\ kg\ /\ s$$

Reflected Light

- The colours humans perceive are determined by nature of light reflected from an object

- For example, if white light (contains all wavelengths) is shone onto green object it absorbs most wavelengths absorbed except green wavelength (color)

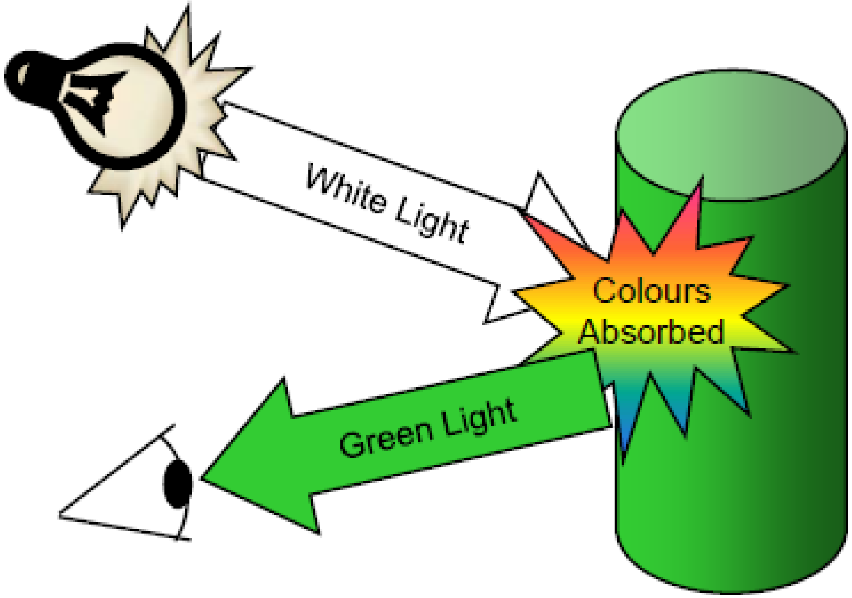

Images from Different EM Radiation

- Radar imaging (radio waves)
- Magnetic Resonance Imaging (MRI) (Radio waves)
- Microwave imaging
- Infrared imaging
- Photographs
- Ultraviolet imaging telescopes
- X-rays and Computed tomography
- Positron emission tomography (gamma rays)
- Ultrasound (not EM waves)

Light and Electromagnetic Spectrum

- Three basic quantities to describe chromatic light source
  - Radiance: total amount of energy that flows from the source
  - Luminance: the amount of energy perceived by an observer/sensor
  - Brightness: subjective property that is hard to measure. It gives the achromatic notion of intensity.
- Achromatic light / monochromatic
  - Light is void of color
  - Gray level is often the term used to describe monochromatic intensity

Image Sensing and Acquisition

- Sensing and acquisition depend on the problem domain
- Generally, we need
  - Source for illumination(gamma, x-ray, ultrasound …..)
  - Device(s) to collect the energy reflected-from/transmitted-through the objects in the scene (stars, patients, natural scenes)
- Radiance is the total energy that flows from a light source
- Luminance is the level of energy an observer perceives from a light source
- Fundamental limit:
  - To see an object the electromagnetic wavelength must be no bigger than the object
  - To image molecules far ultraviolet or soft x-ray waves must be used

Image Sensing and Acquisition

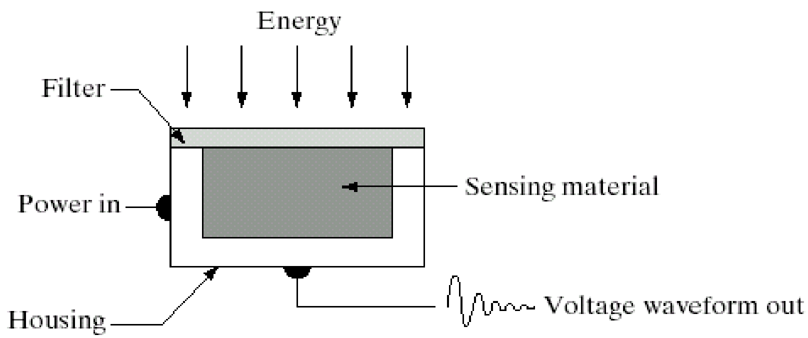

Single sensor

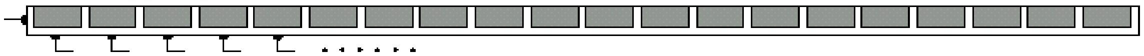

Line sensor

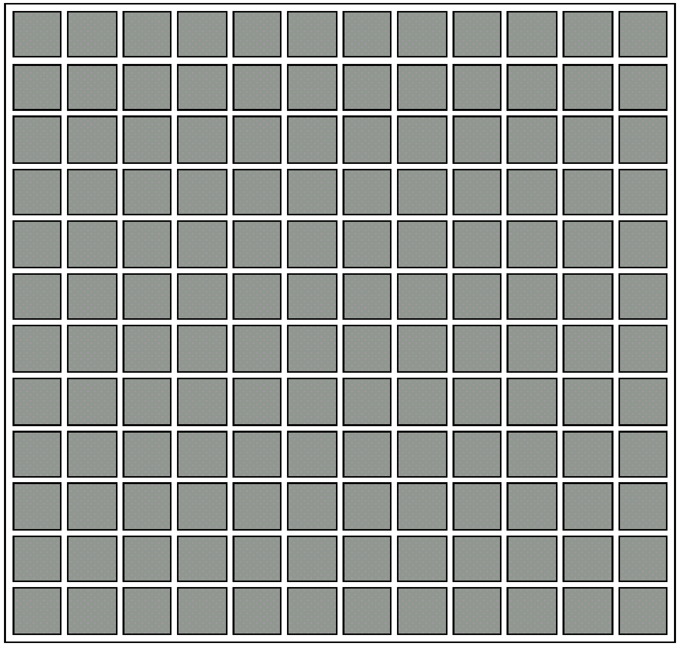

Array sensor

Image Sensing and Acquisition

- Acquisition with Sensor Arrays
  - Predominant arrangement found in many applications
  - Most expensive and usually no motion is required

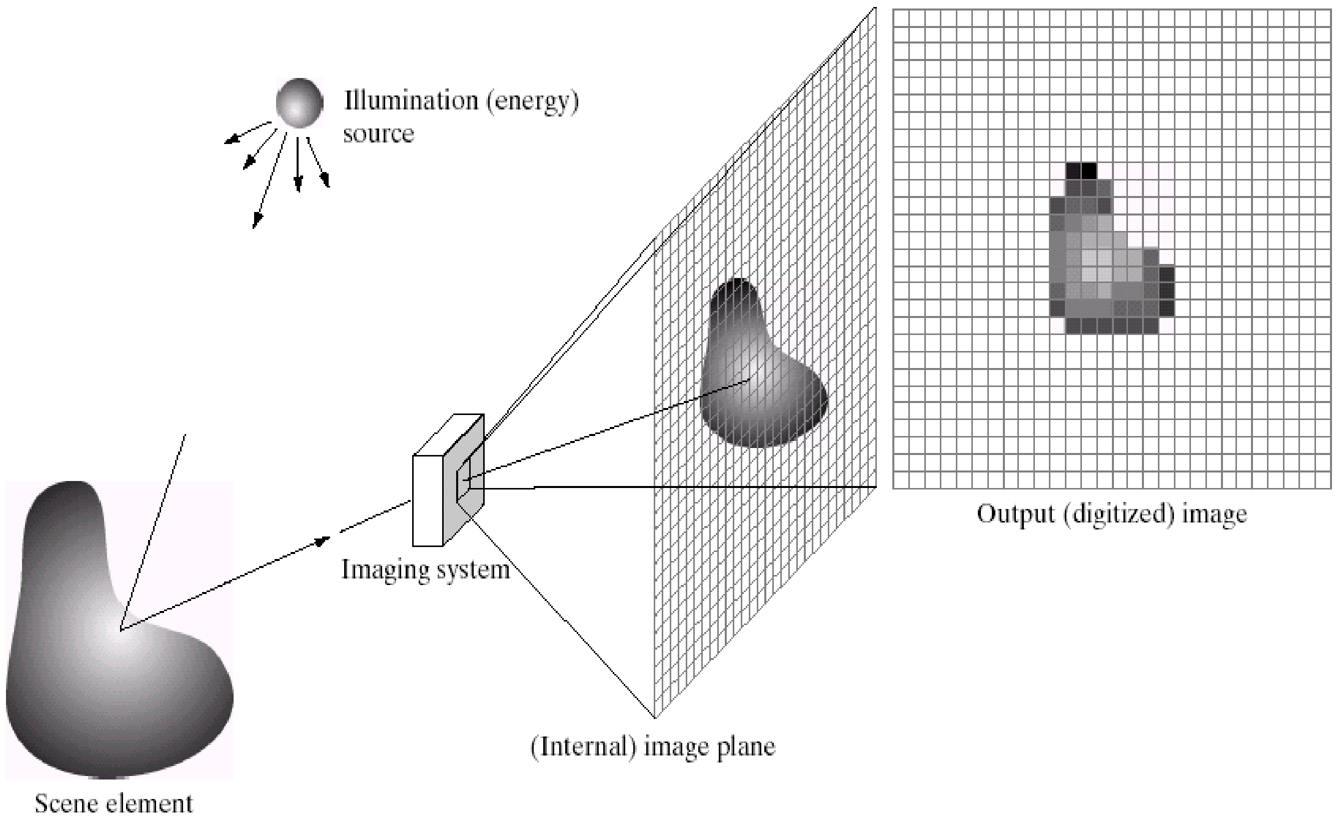

Simple Image Formation Model

- An image can represented as a two-dimensional function f(x,y)
- The amplitude of f(x,y) is a positive scalar quantity whose physical meaning depends on the source
- An image is proportional to the radiated energy

$$f(x,y) = i(x,y)r(x,y)$$

  - illumination bound: $0<i(x,y)<\infty$
  - Reflectivity bound: $0<r(x,y)<1$

- Transmission cases (x-ray): transmissivity rather than reflectivity
- Frequency dependent functions

A Simple Image Formation Model

- For monochrome images we deal with intensity or gray levels (L)

$$L = f(x,y)$$

$$L_{min} \leq L \leq L_{max}$$

- The only requirement on Lmin is to be positive and on Lmax to be finite
- Typical indoor values for Lmin and Lmax are 10 lm/m2 and 1000 lm/m2
- The interval [Lmin,Lmax] is called the gray scale and is
- usually shifted to [0,L-1] where 0 is considered black and L-1 is considered white

Image Sampling and Quantization

- Sensors used in the acquisition produce continuous voltage signal
  - In order to produce a digital form of the image, it has to go through two processes:
    - **Sampling**: digitize the spatial coordinates, x and y
    - **Quantization**: digitize the amplitude f(x,y)

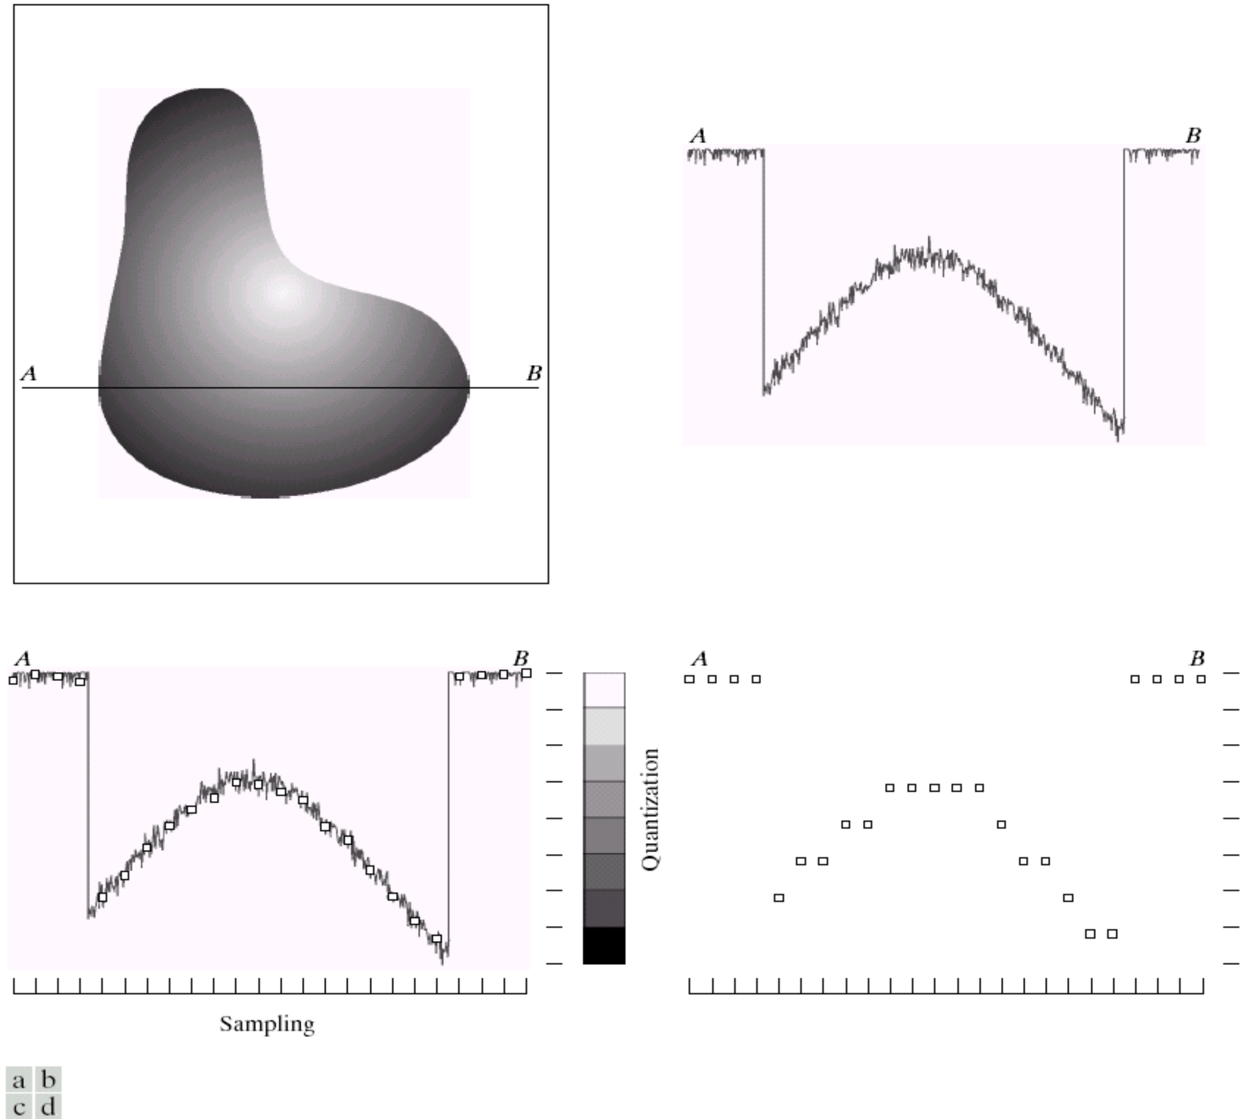

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from A to B in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

Image Sampling and Quantization

- Resolution (how much you can see the detail of the image) depends on sampling and gray levels.
- The bigger the sampling rate (n) and the gray scale (g), the better the approximation of the digitized image from the original.
- The more the quantization scale becomes, the bigger the size of the digitized image.

- The number of gray levels typically is an integer power of 2

$$L = 2^k$$

- Number of bits required to store a digitized image

$$b = M \times N \times k$$

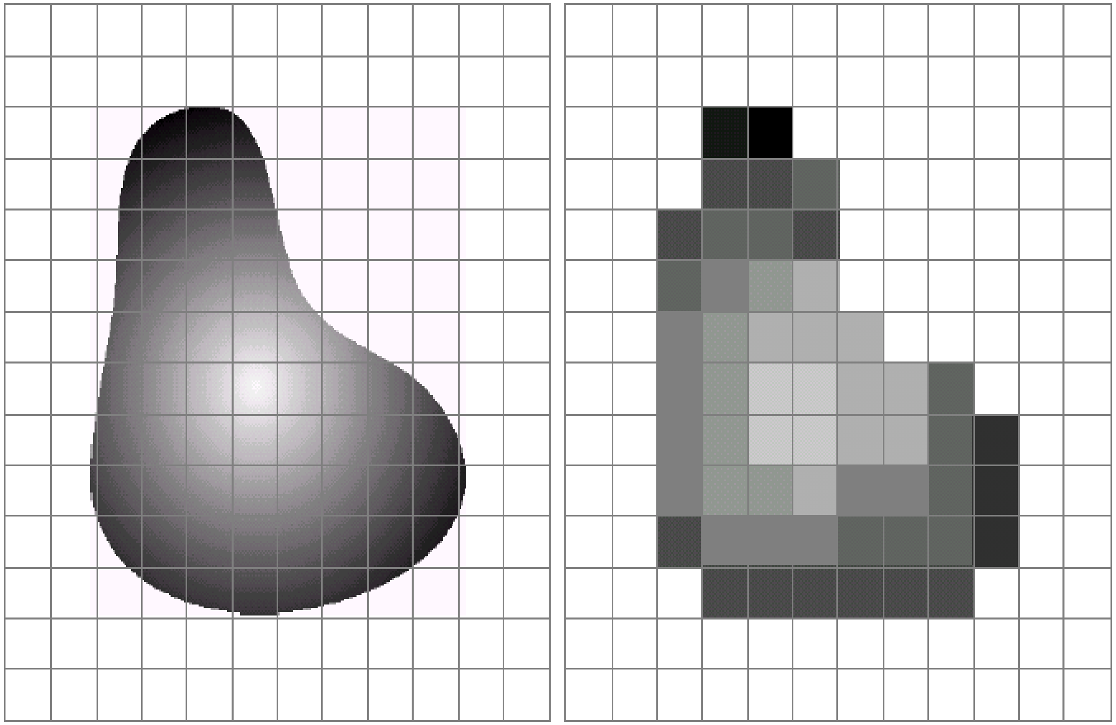

a b

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

Digital Image Representation

- Once the image is digitized, it can be represented as an array with M rows and N columns with each element called a pixel
- The values stored in the array elements represent the image values at that location

$$f(x, y) = \begin{bmatrix} f(0,0) & f(0,1) & \cdots & f(0, N-1) \\ f(1,0) & f(1,1) & \cdots & f(1, N-1) \\ \vdots & \vdots & & \vdots \\ f(M-1,0) & f(M-1,1) & \cdots & f(M-1, N-1) \end{bmatrix}$$

$$\mathbf{A} = \begin{bmatrix} a_{0,0} & a_{0,1} & \cdots & a_{0,N-1} \\ a_{1,0} & a_{1,1} & \cdots & a_{1,N-1} \\ \vdots & \vdots & & \vdots \\ a_{M-1,0} & a_{M-1,1} & \cdots & a_{M-1,N-1} \end{bmatrix}$$

Digital Image Representation

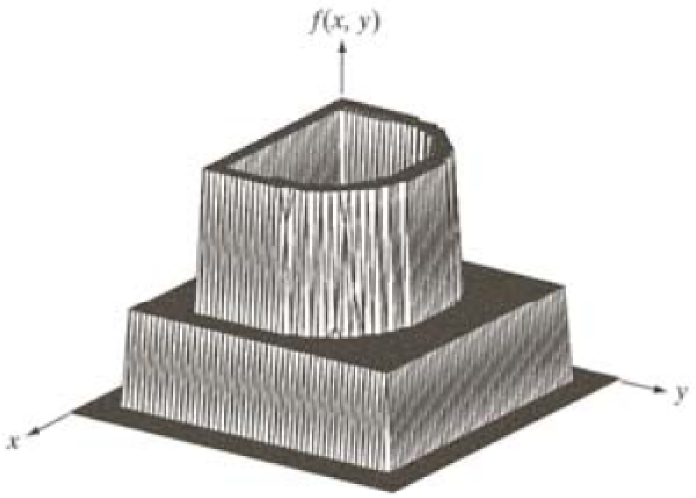

Image as
3-D surface

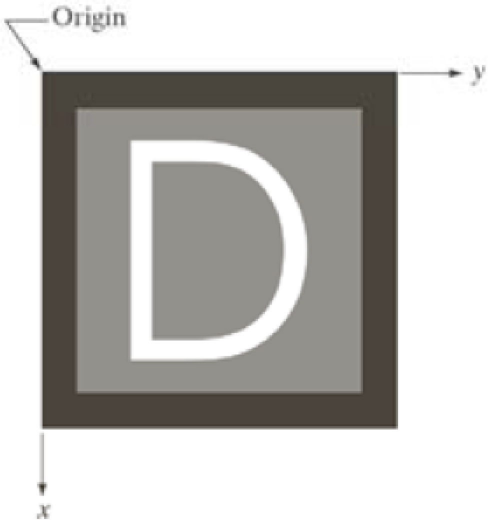

Image as visual
intensity array

Image as numerical
array

Image Sampling and Quantization

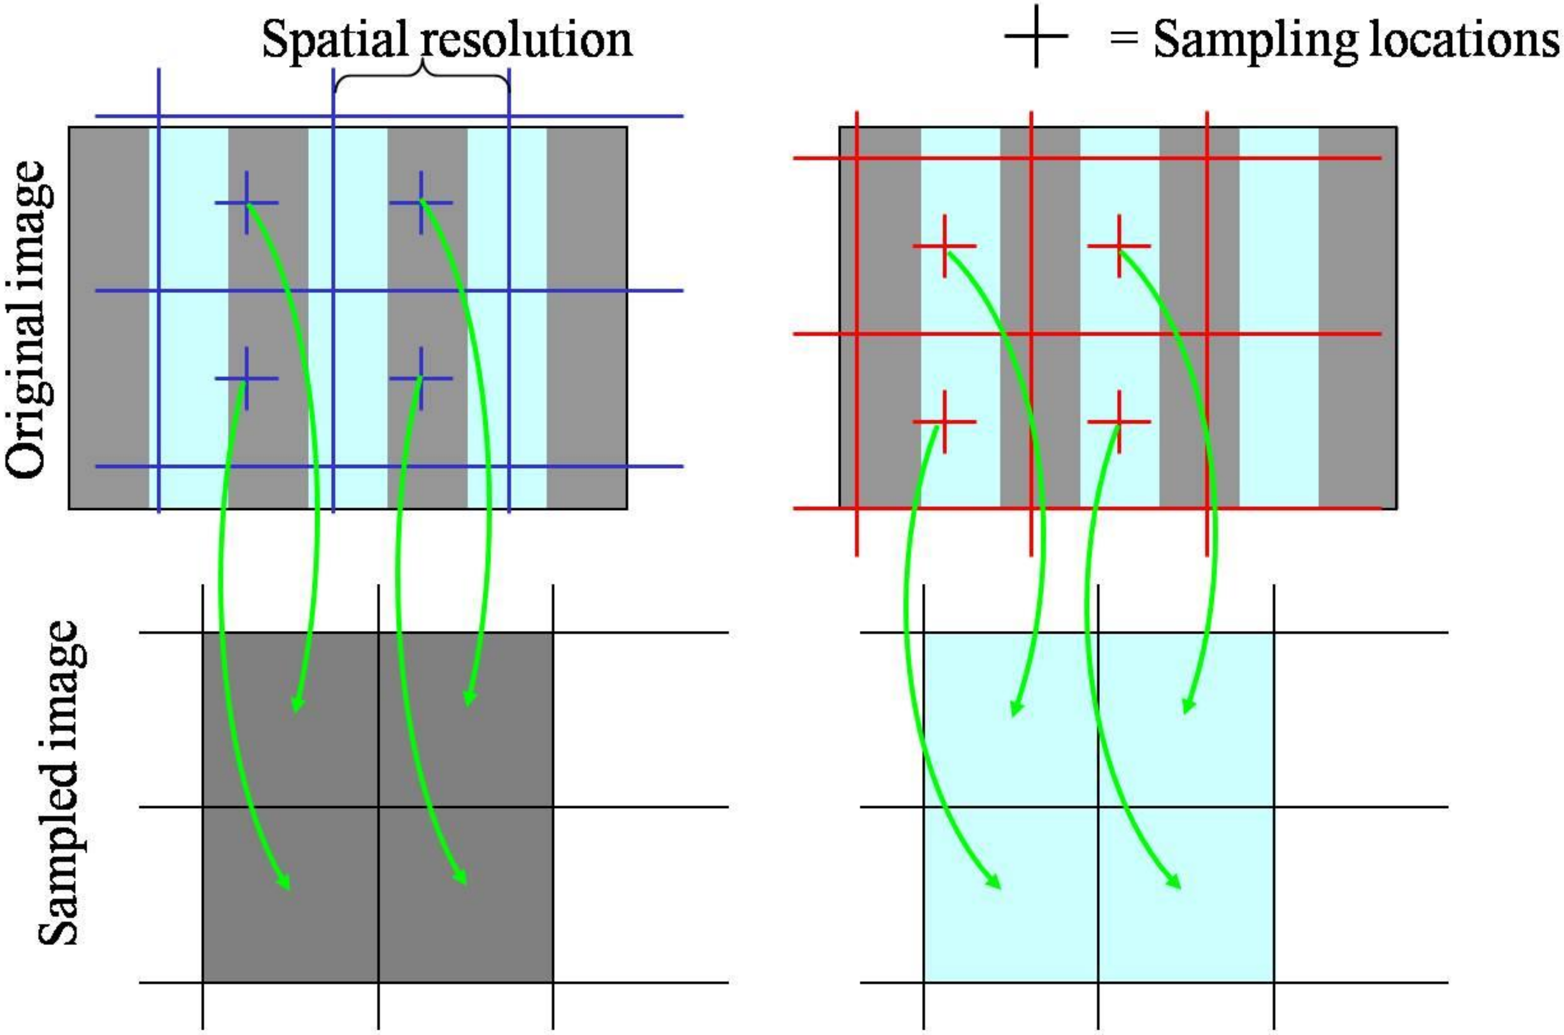

Under sampling, we lost some image details!

Image Sampling and Quantization

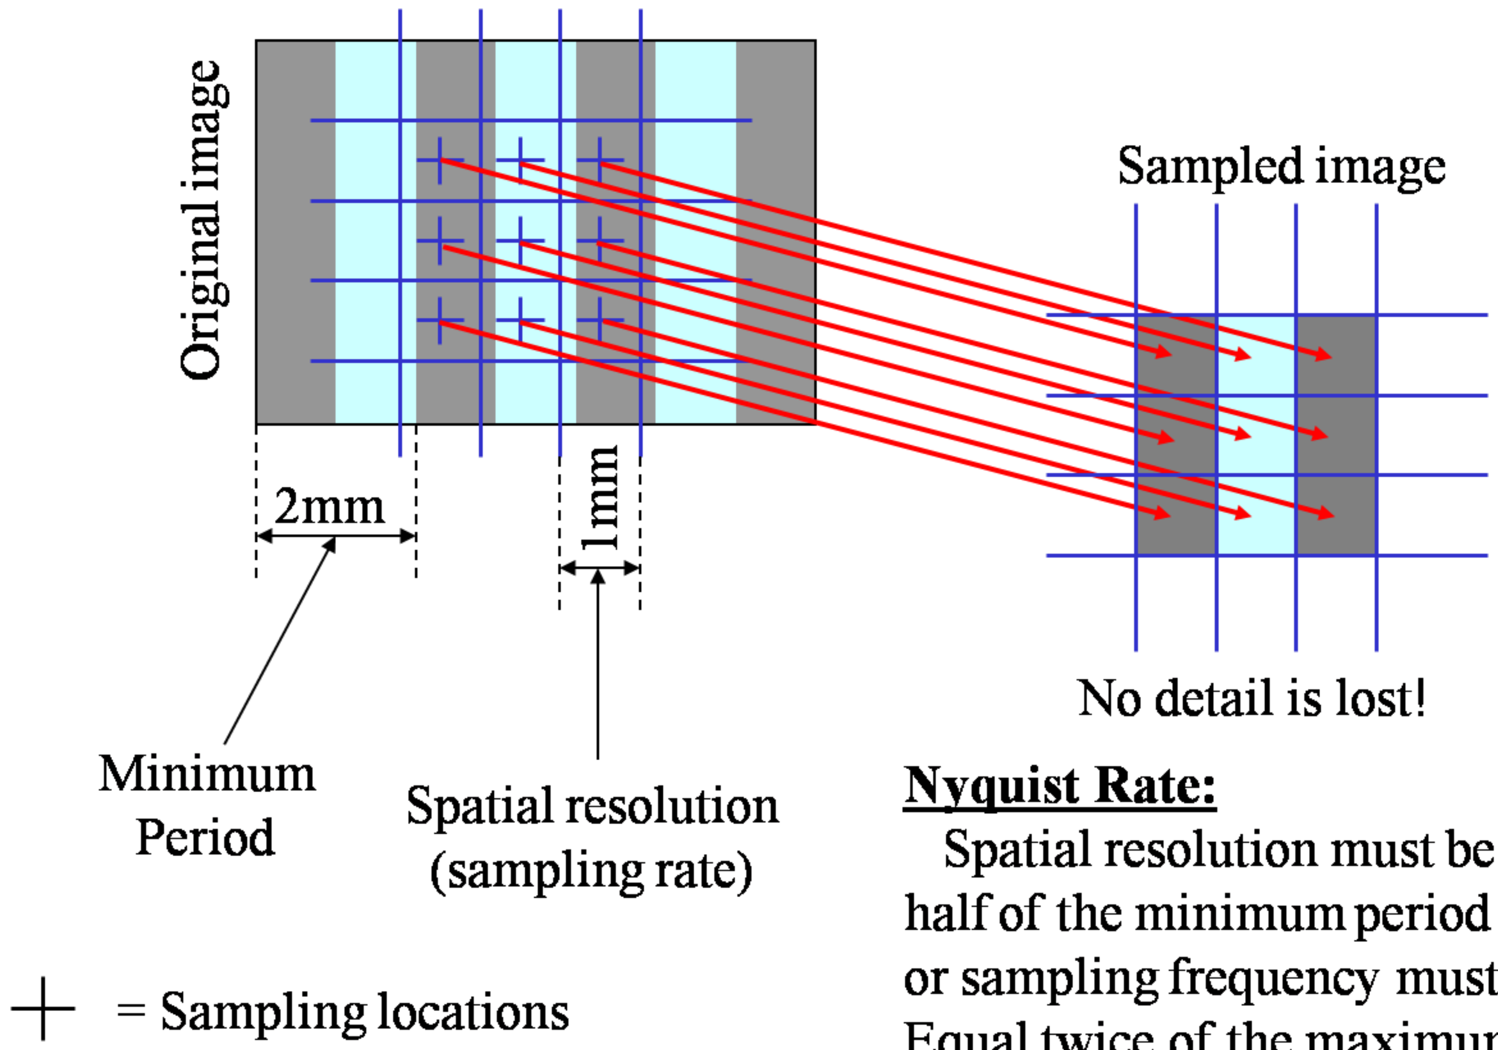

Nyquist Rate:

Spatial resolution must be less or equal half of the minimum period of the image or sampling frequency must be greater or Equal twice of the maximum frequency.

Image Sampling and Quantization

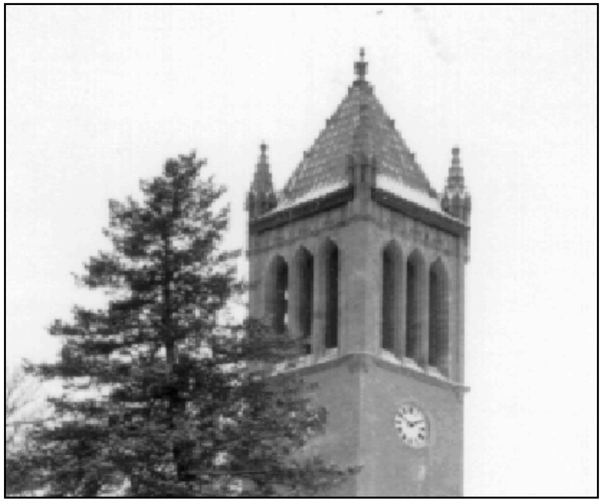

256x256 pixels

128x128 pixels

64x64 pixels

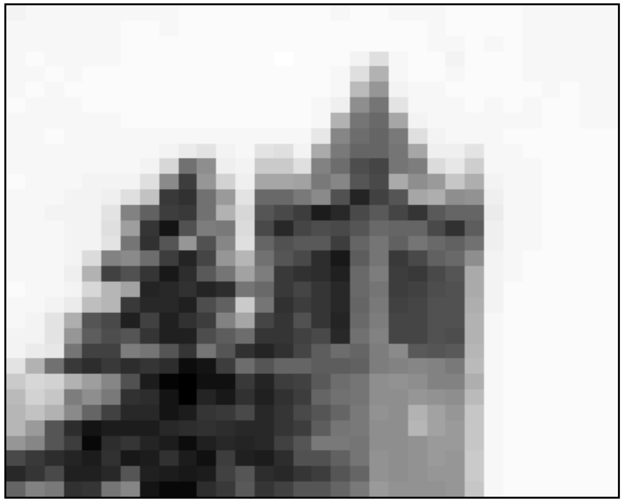

32x32 pixels

Image Sampling and Quantization

1250 pixels

300 pixels

150 pixels

75 pixels

Effect of reducing spatial resolution

Image Sampling and Quantization

- The number of quantization levels depends on
  - A/D converter specs
  - Available storage space allocated for each sample (number of bits per sample)
  - The number of available levels is an integer power of 2
- For gray level images
  - Assume that the levels are equally spaced and spans the range [0,L-1]
  - Such range is called the dynamic range
  - Typically, each gray level sample uses 8 bits, thus the number of available gray levels is 256
  - The number of gray levels specify the gray level resolution

Image Sampling and Quantization

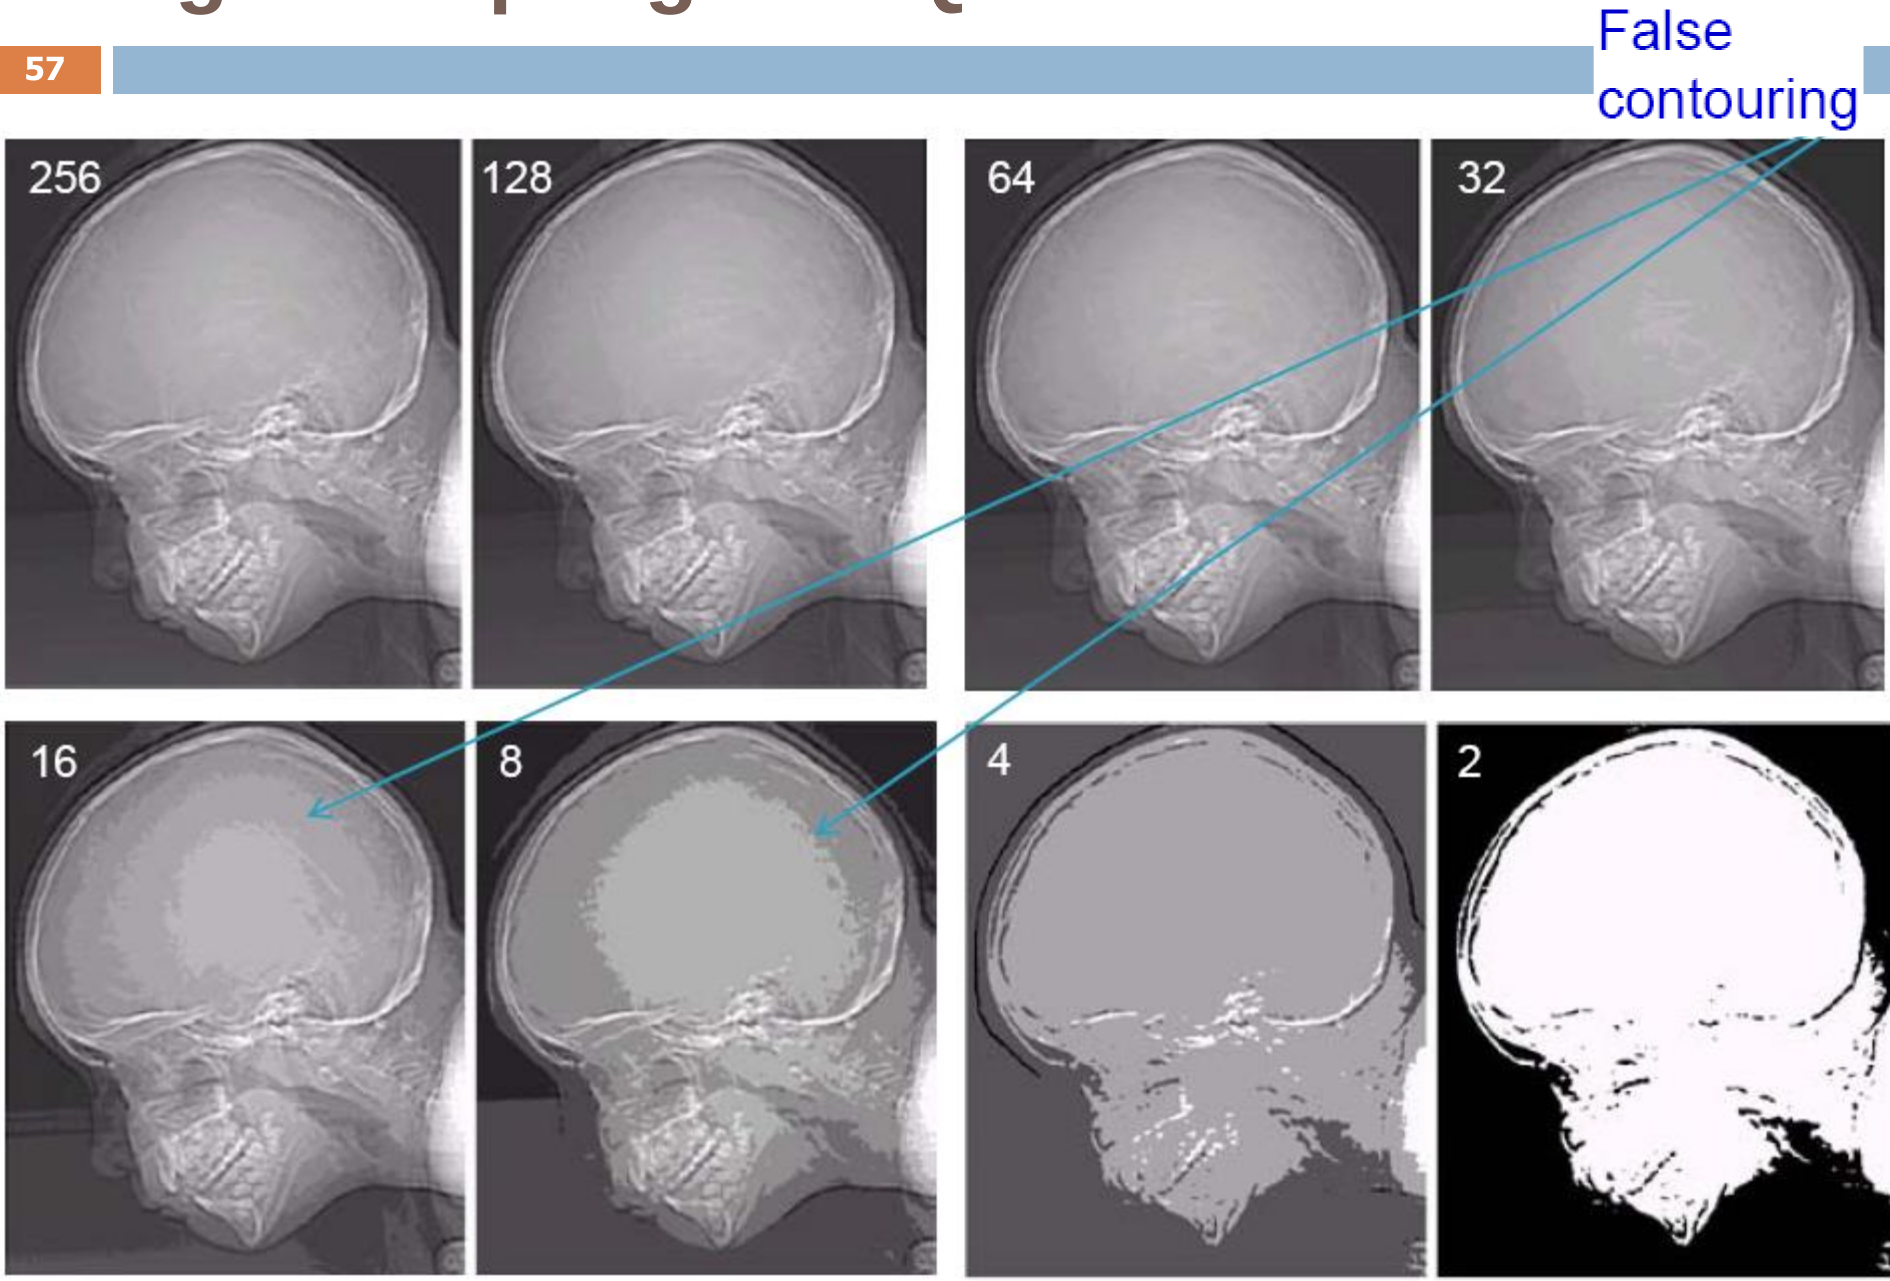

Intensity Level Resolution

- ***Intensity level resolution***: number of intensity levels used to represent the image
  - The more intensity levels used, the finer the level of detail discernable in an image
  - Intensity level resolution usually given in terms of number of bits used to store each intensity level

Number of Bits	Number of Intensity Levels	Examples
1	2	0, 1
2	4	00, 01, 10, 11
4	16	0000, 0101, 1111
8	256	00110011, 01010101
16	65,536	1010101010101010

How many Bits Per Image Element?

Grayscale (Intensity Images):

Chan.	*Bits/Pix.*	*Range*	*Use*
1	1	0...1	Binary image: document, illustration, fax
1	8	0...255	Universal: photo, scan, print
1	12	0...4095	High quality: photo, scan, print
1	14	0...16383	Professional: photo, scan, print
1	16	0...65535	Highest quality: medicine, astronomy

Color Images:

Chan.	*Bits/Pix.*	*Range*	*Use*
3	24	$[0...255]^3$	RGB, universal: photo, scan, print
3	36	$[0...4095]^3$	RGB, high quality: photo, scan, print
3	42	$[0...16383]^3$	RGB, professional: photo, scan, print
4	32	$[0...255]^4$	CMYK, digital prepress

Special Images:

Chan.	*Bits/Pix.*	*Range*	*Use*
1	16	−32768...32767	Whole numbers pos./neg., increased range
1	32	$\pm 3.4 \cdot 10^{38}$	Floating point: medicine, astronomy
1	64	$\pm 1.8 \cdot 10^{308}$	Floating point: internal processing

Image details and resolution



To satisfy human mind

1. For images of the same size, the low detail image may need more pixel depth.
2. As an image size increase, fewer gray levels may be needed.

Resolution: How Much Is Enough?

- The big question with resolution is always *how much is enough*?
  - Depends on what is in the image *(details)* and what you would like to do with it *(applications)*
  - Key questions:
    - Does image look aesthetically pleasing?
    - Can you see what you need to see in image?

Resolution: How Much Is Enough?

- **Example:** Picture on right okay for counting number of cars, but not for reading the number plate

Image Interpolation

- It is the operation in which known data is used to estimate values at unknown locations
- An important tool used in zooming, shrinking, rotation, and geometric corrections
- Zooming: enlarge a MxN image to OxP image
  - Can be simply done by row/column replication
- Shrinking: reduce a MxN image to QxR image
  - Can be simply done by row/column deletion
- Interpolation methods
  - Nearest neighbor
  - Bilinear
  - Bicubic

Image Shrinking by Pixel Deletion

Original

Image shrinked by 2x2

Image shrinked by 4x4

Image shrinked by 16x16

Zooming by Pixel Replication

Original

Image zoomed by factor of 2x2

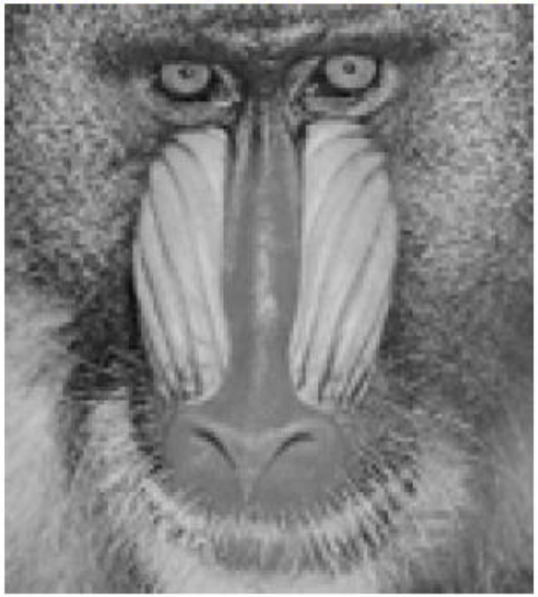

Image zoomed by factor of 4x4

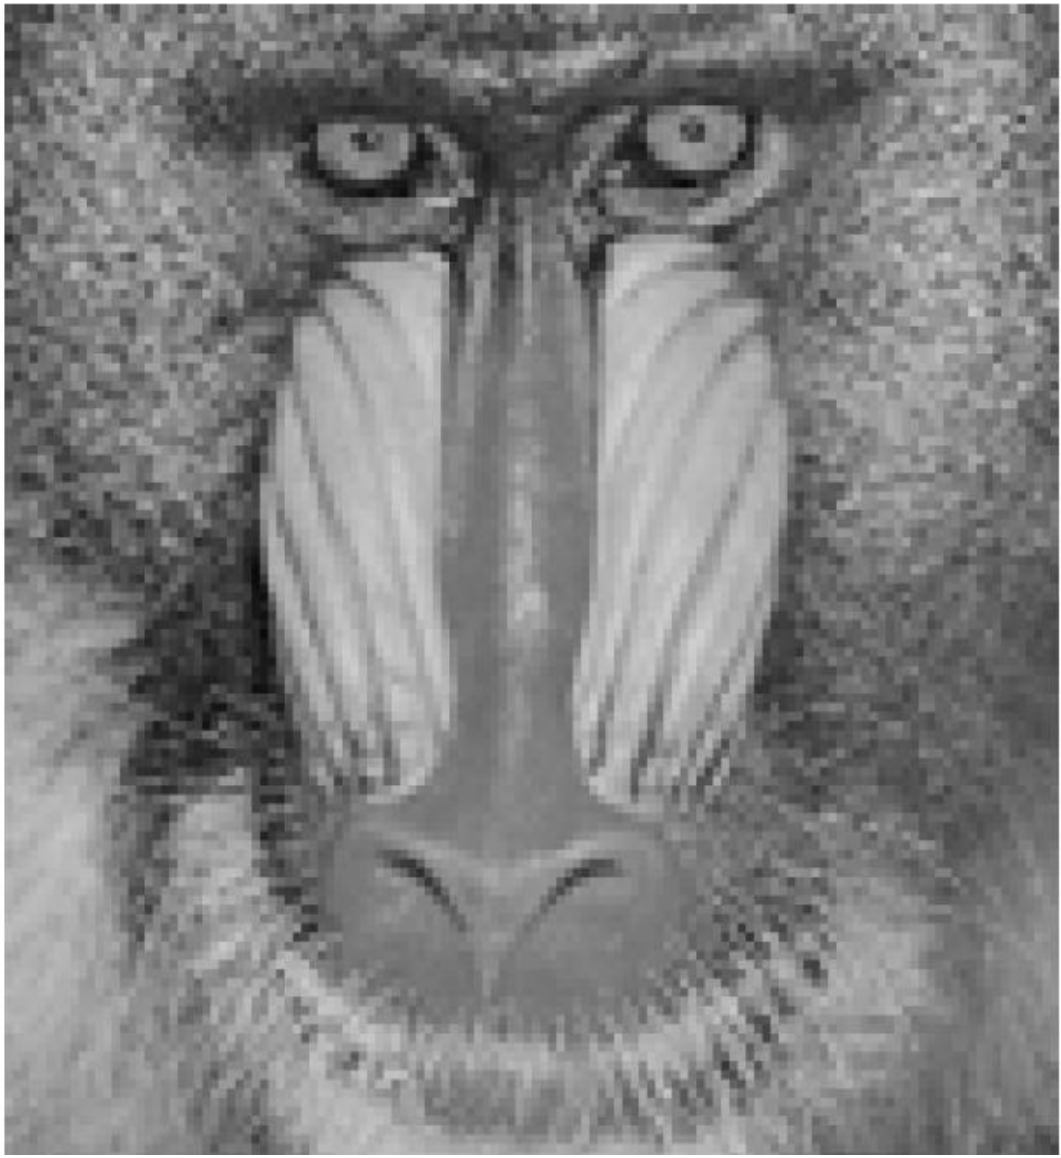

Image Interpolation

- **Zooming – Nearest Neighbor**
  - Generate a grid with the new size and assign available pixels to new locations
  - For missing pixels, assign to it the value of the closest known pixel
  - Fast but produces artifacts

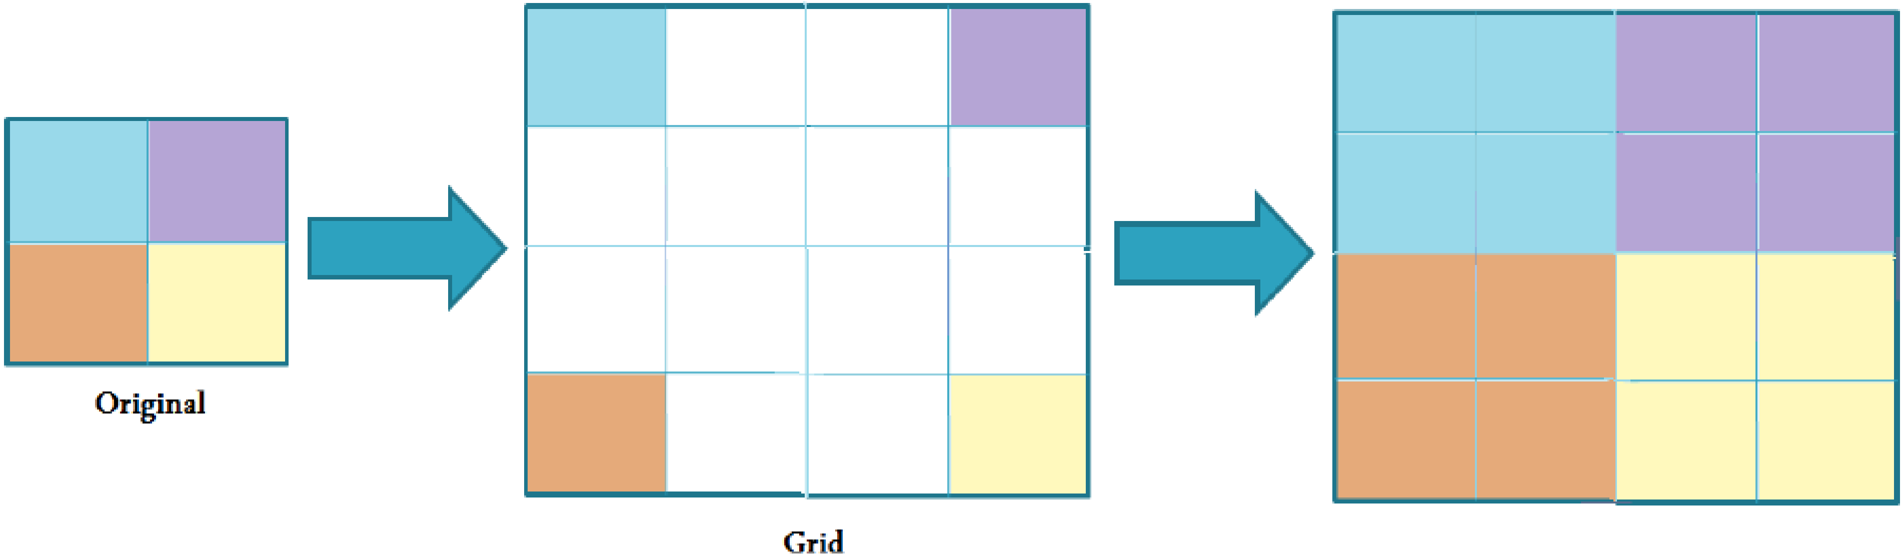

Image Interpolation

- **Zooming – Bilinear**
  - Generate a grid with the new size and assign available pixels to new locations
  - For missing pixels, assign to it the value of by

$$v(x, y) = ax + by + cxy + d$$

  - The constants are determined by solving four equations for the closest known four neighbors

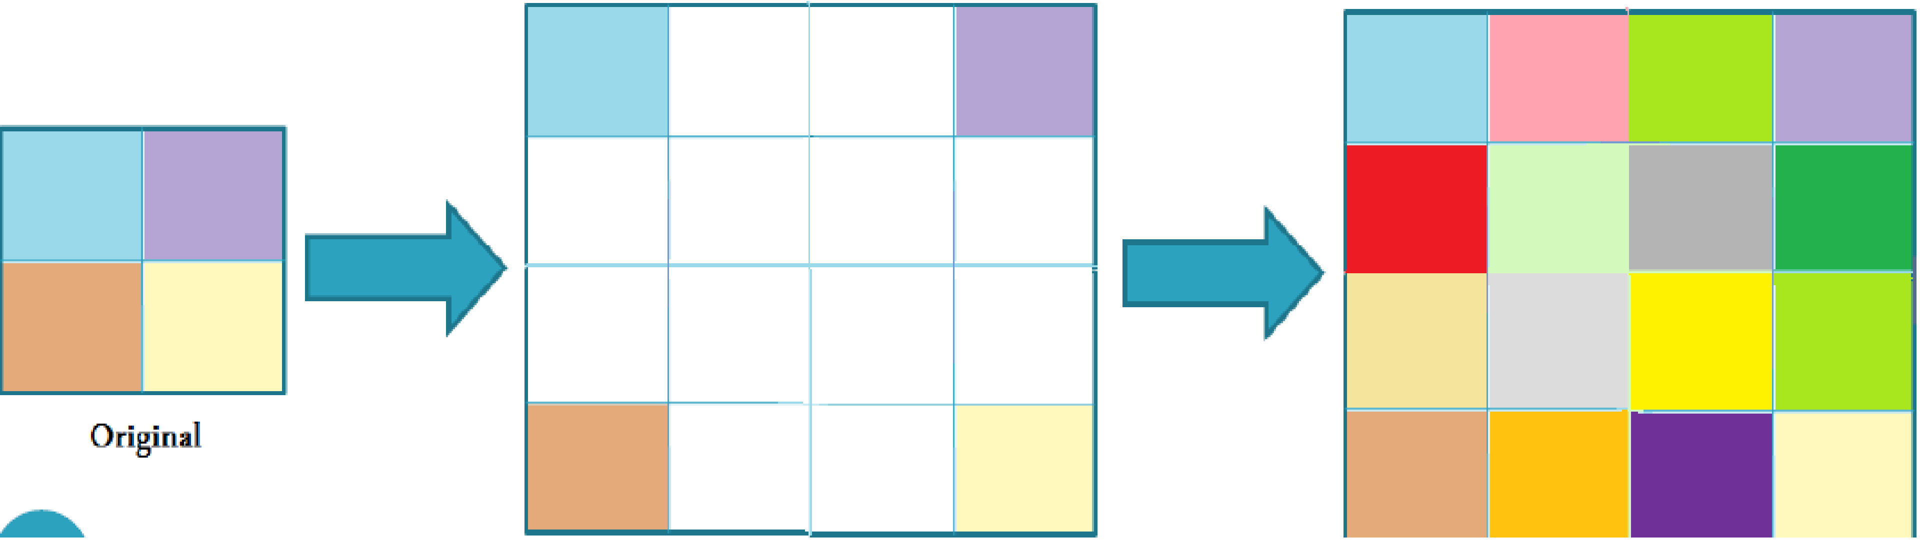

Image Interpolation

- **Zooming – Bicubic**
  - Generate a grid with the new size and assign available pixels to new locations
  - For missing pixels, assign to it the value of by

$$v(x, y) = \sum_{i=0}^{3} \sum_{j=0}^{3} a_{ij} x^i y^j$$

  - The constants are determined by solving sixteen equations for the closest sixteen neighbors

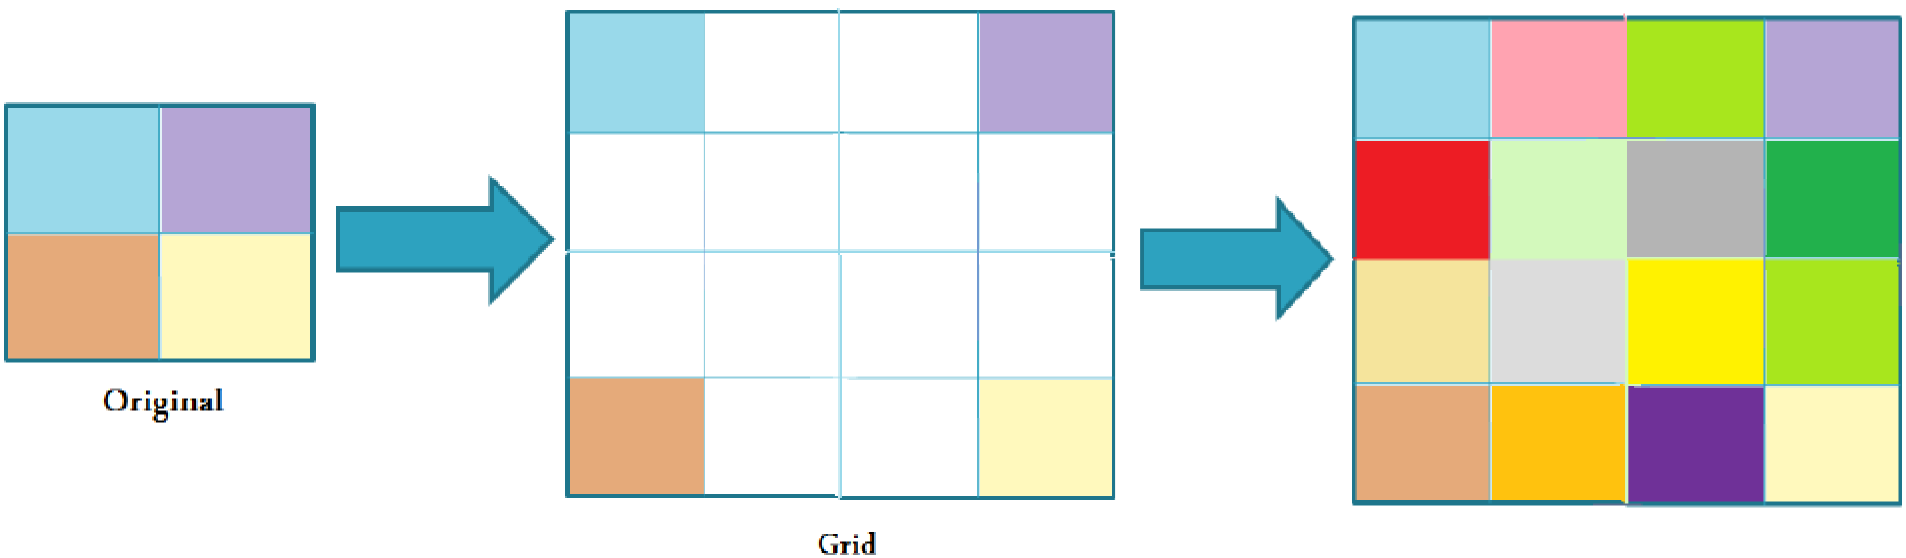

Image Interpolation

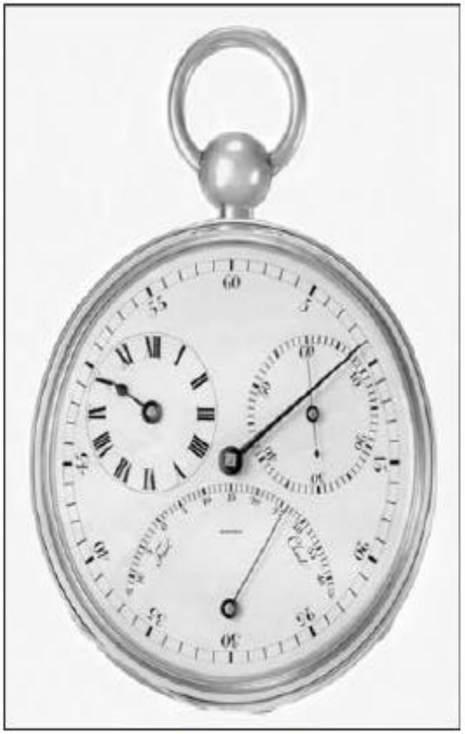

Original

Nearest Neighbor

Bilinear

Bicubic

Math Preliminaries

- **Arithmetic operations**
  - All operations are performed array-wise, i.e. element by element
  - Arrays (images) should be of the same size

$$s(x, y) = f(x, y) + g(x, y)$$

$$d(x, y) = f(x, y) - g(x, y)$$

$$p(x, y) = f(x, y) \times g(x, y)$$

$$v(x, y) = f(x, y) \div g(x, y)$$

Math Preliminaries

- **An example of addition in IP**
  - Averaging multiple images captured for the same scene reduces uncorrelated zero-mean noise

- Image Model

$$g(x, y) = f(x, y) + \eta(x, y)$$

- Averaging

$$\overline{g}(x, y) = \frac{1}{K}\sum_{i=1}^{K} g_i(x, y)$$

- If K is too large

$$E\{\overline{g}(x, y)\} = f(x, y)$$

$$\sigma^2_{\overline{g}(x,y)} = \frac{1}{K}\sigma^2_{\eta(x,y)}$$

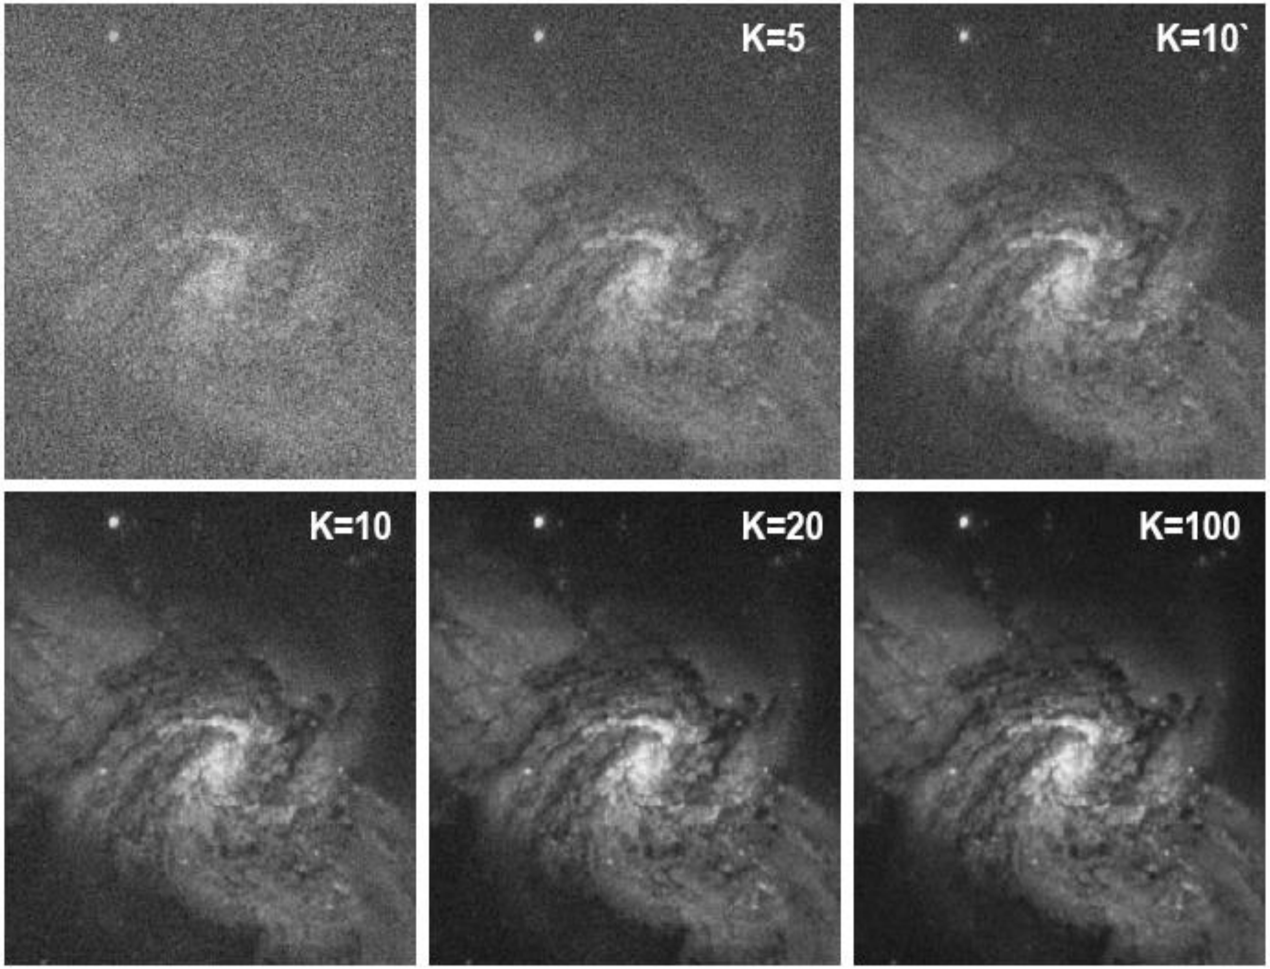

Math Preliminaries

- **An example of subtraction in IP**
  - Usually used in the enhancement of the difference between images

$$g(x, y) = f(x, y) - h(x, y)$$

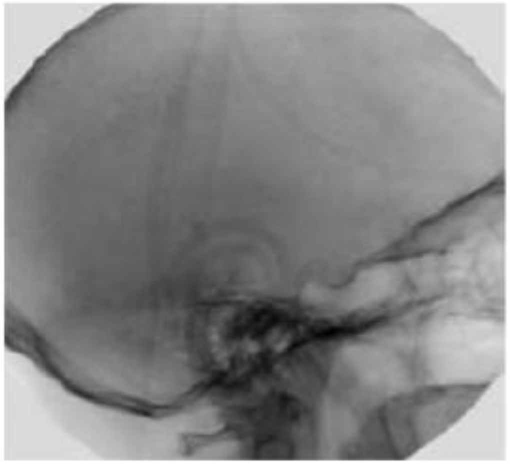

Image before Iodine injection f(x,y)

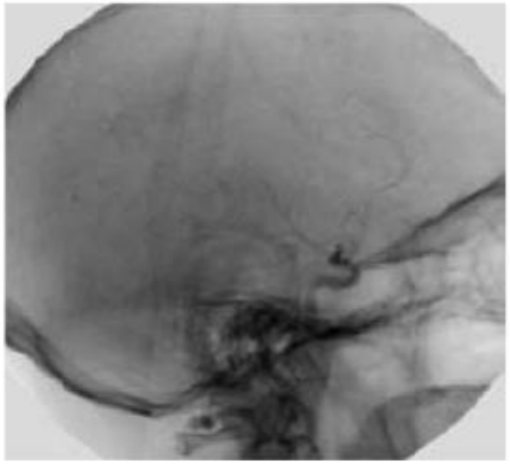

Image after Iodine injection (mask h(x,y))

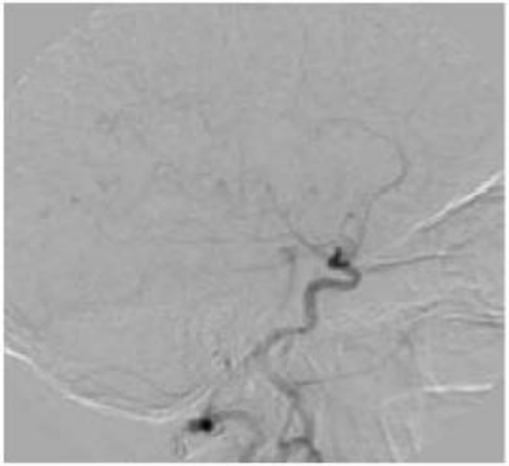

Difference image g(x,y)

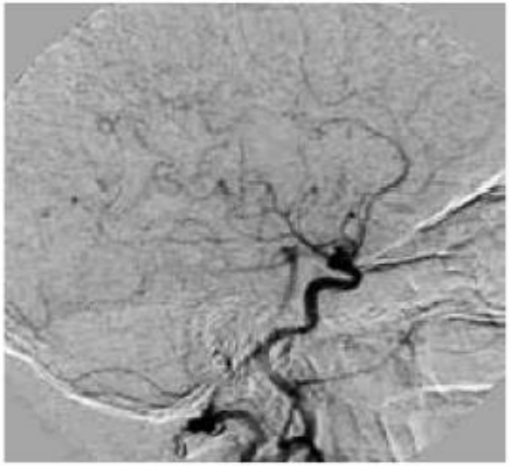

Enhanced difference image

Math Preliminaries

- **An example of multiplication in IP**
  - Region of interest (ROI) extraction by multiplying with a mask

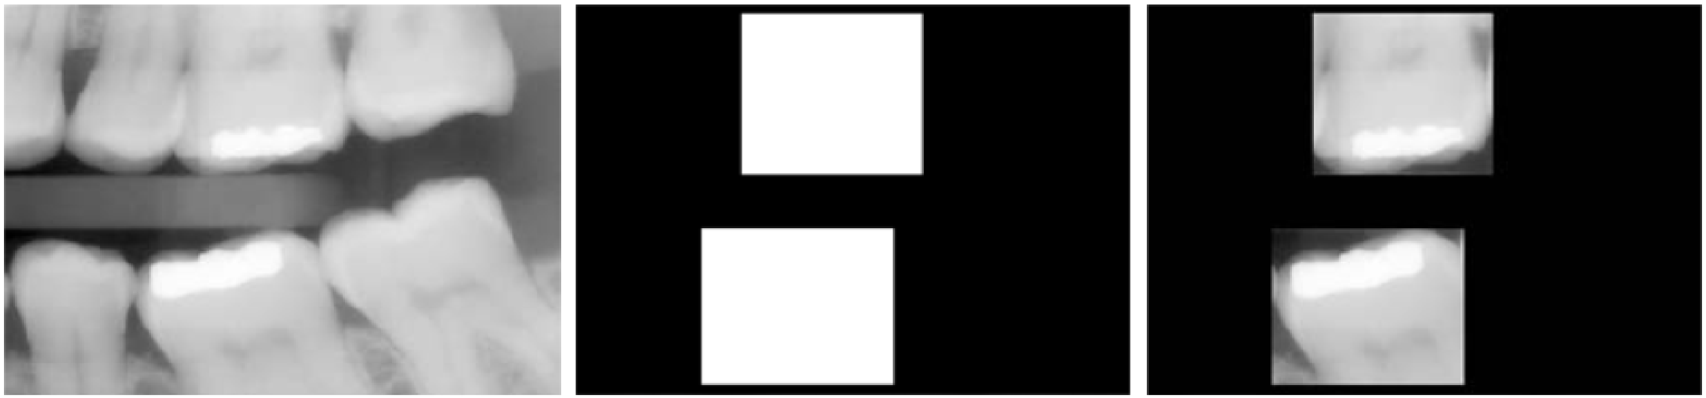

Whole image — Mask — ROI

- **An example of division in IP**
  - Shading correction

$$g(x, y) = f(x, y)h(x, y)$$

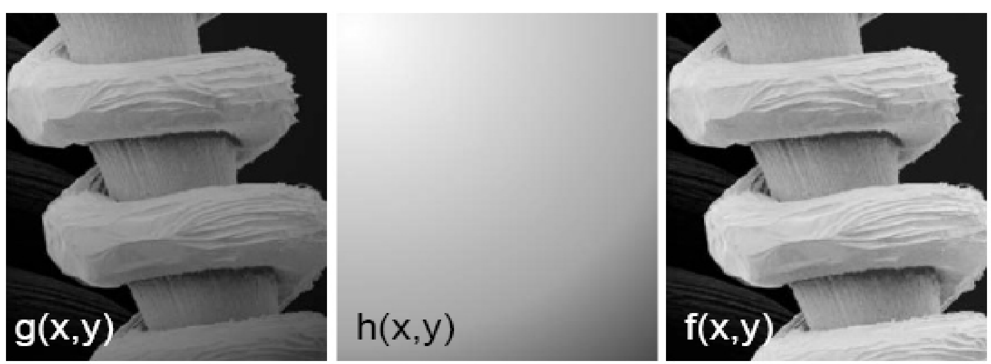

Math Preliminaries

- Spatial Operations
  - **Geometric transformations (Rubber-sheet)**
    - Modify the spatial relationships between pixels in an image
    - Geometric transformation involves two basic operations
      1. Transform the image coordinates

$$(x, y) = T\{(v, w)\}$$

Example : $(x, y) = T\{(v, w)\} = (v/2, w/2)$

- Affine transformation

$$[x \;\; y \;\; 1] = [v \;\; w \;\; 1]\,\mathbf{T} = [v \;\; w \;\; 1]\begin{bmatrix} t_{11} & t_{12} & 0 \\ t_{21} & t_{22} & 0 \\ t_{31} & t_{32} & 1 \end{bmatrix}$$

- Can be used for rotation, translation, scale, and sheer by choosing the proper values for T from the table in the following slide

Math Preliminaries

- **Spatial Operations**

Transformation Name	Affine Matrix, T	Coordinate Equations	Example
Identity	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$x = v$ $y = w$	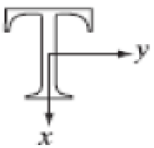
Scaling	$\begin{bmatrix} c_x & 0 & 0 \\ 0 & c_y & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$x = c_x v$ $y = c_y w$	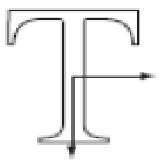
Rotation	$\begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$x = v\cos\theta - w\sin\theta$ $y = v\cos\theta + w\sin\theta$	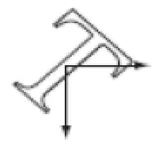
Translation	$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ t_x & t_y & 1 \end{bmatrix}$	$x = v + t_x$ $y = w + t_y$	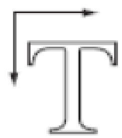
Shear (vertical)	$\begin{bmatrix} 1 & 0 & 0 \\ s_v & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$x = v + s_v w$ $y = w$	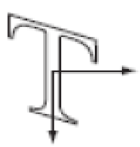
Shear (horizontal)	$\begin{bmatrix} 1 & s_h & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$	$x = v$ $y = s_h v + w$	

Math Preliminaries

- Spatial Operations
  - Geometric transformations

Original

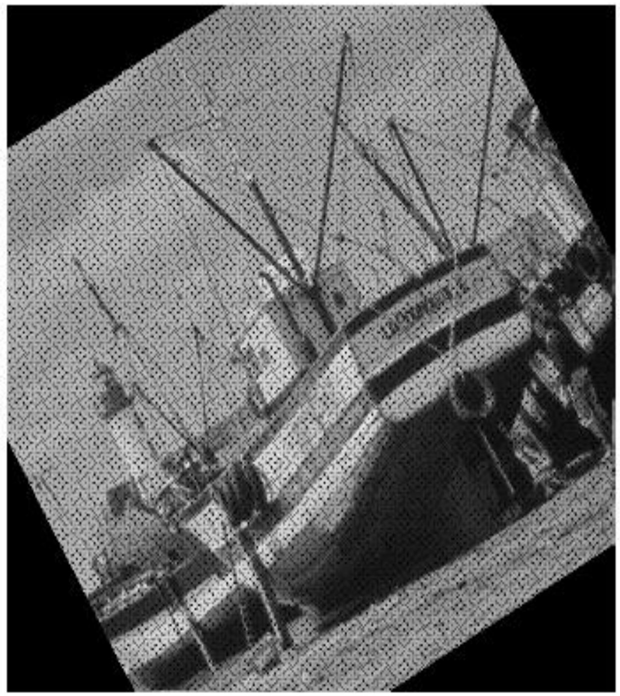

Rotation by 30° using forward transformation

Rotation by 30° using inverse transformation with nearest neighbor interpolation

Math Preliminaries

- Spatial Operations
  - Geometric transformations

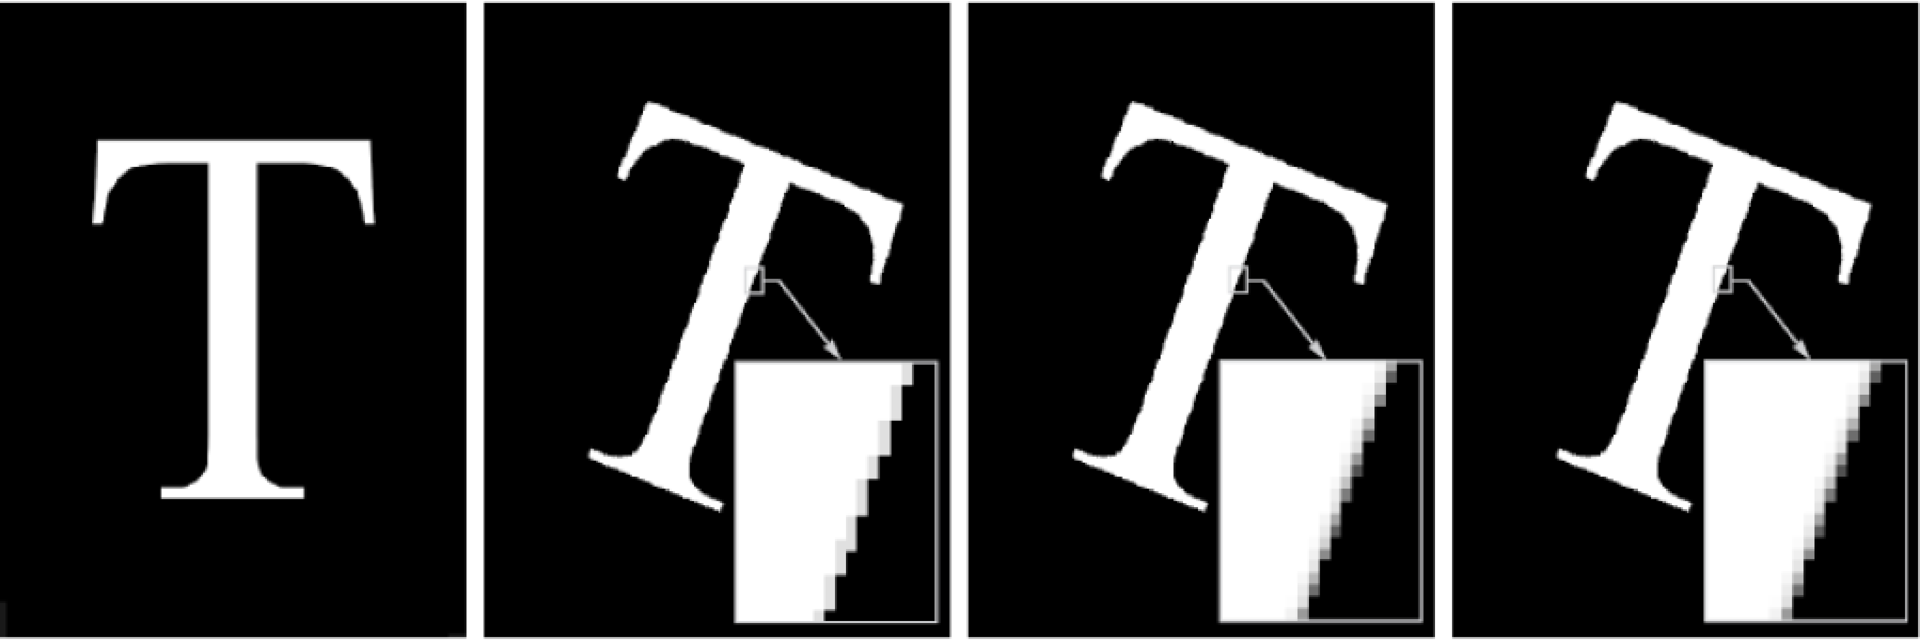

Original

Rotation by 21o using inverse mapping and nearest neighbor

Rotation by 21o using inverse mapping and bilinear interpolation

Rotation by 21o using inverse mapping and bicubic interpolation

Math Preliminaries

- **Statistical Measures**
  - Mean intensity gives the general sense of brightness

Bright
μ =192

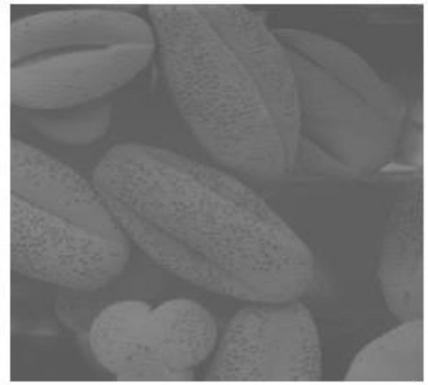

Dark
μ =102

  - Variance/standard deviation can be used to measure image contrast

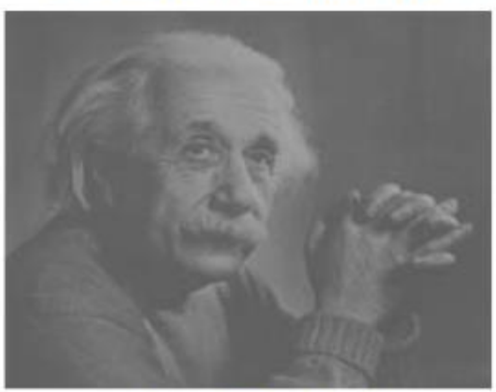

Low Contrast
σ = 14

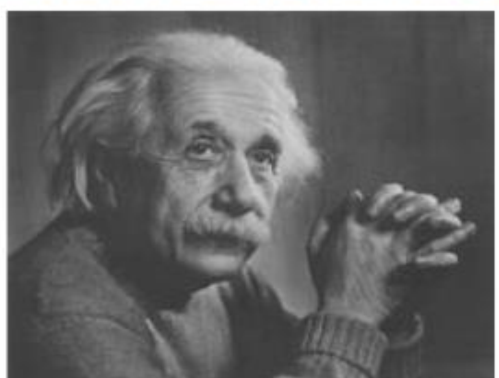

Medium Contrast
σ = 32

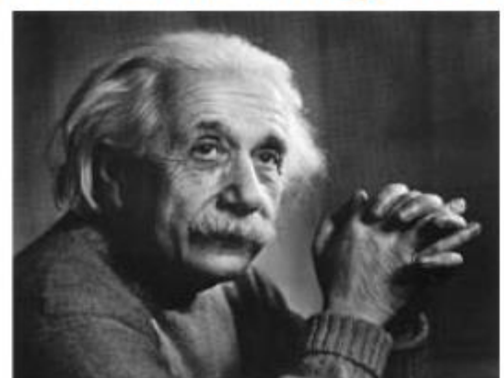

High Contrast
σ =49

Images in Matlab

- Why use Matlab ?
  - Quick-to-learn
  - Easy-to-use
  - Integrated environment
  - Platform independent
  - Scientific standard
  - Rapid Development

Images in Matlab

- See the attached documents on Ritaj
  - matlab_primer_part1
  - matlab_primer_part1