Image formation and fundamentals

IP framework

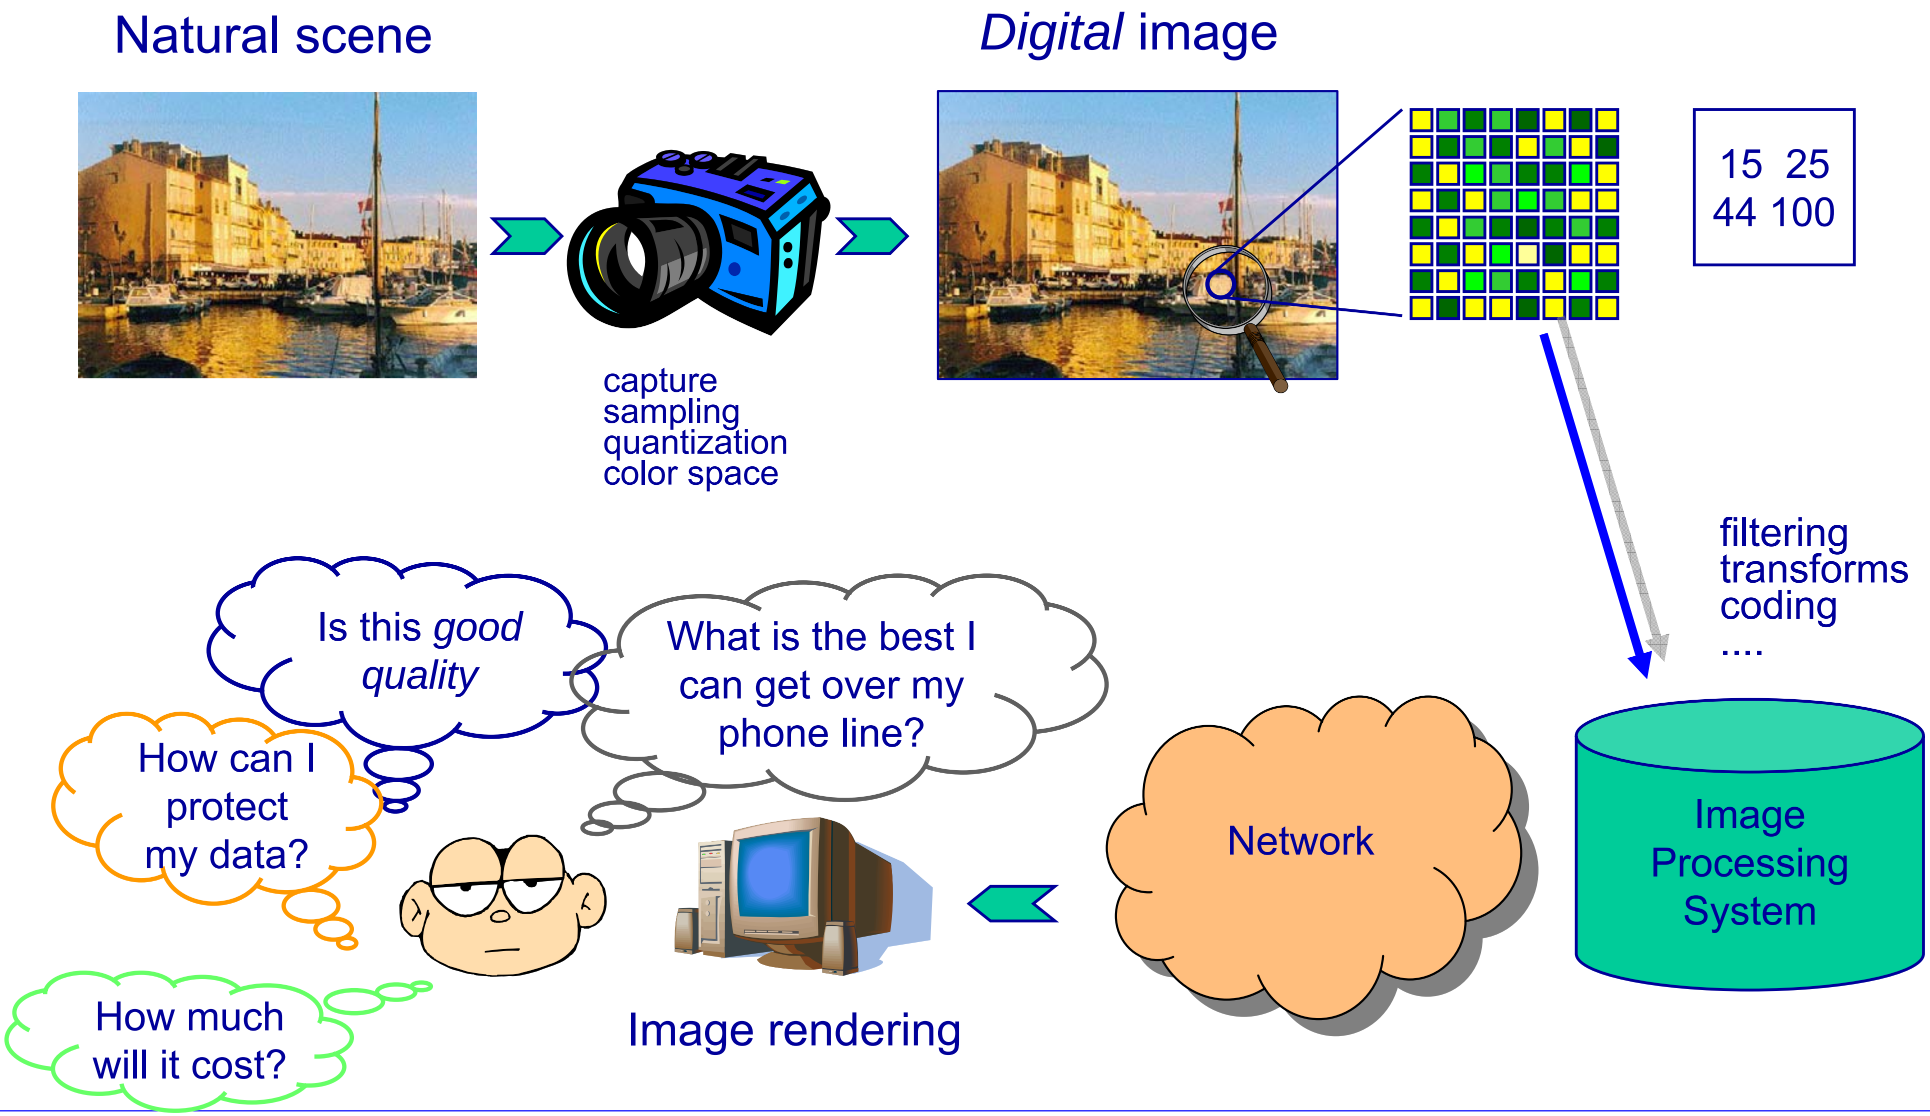

IP: basic steps

Analog image

(capturing device)

A/D conversion

Sampling (2D)

Quantization

Digital image

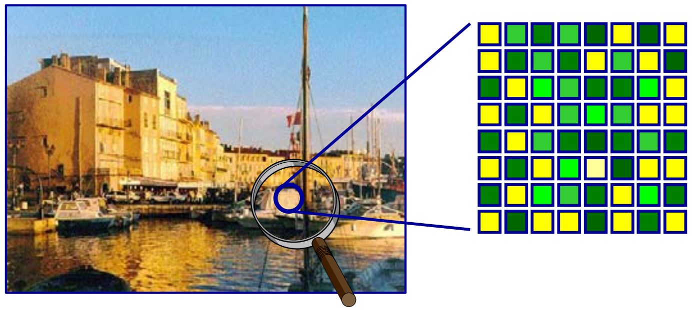

{15,1,2}
{25,44,1}
….

Digital Image Acquisition

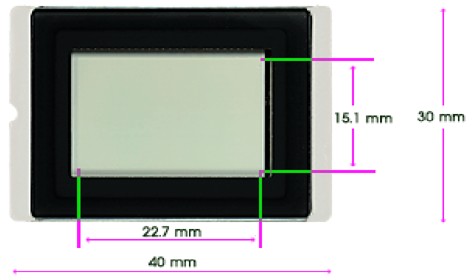

Sensor array

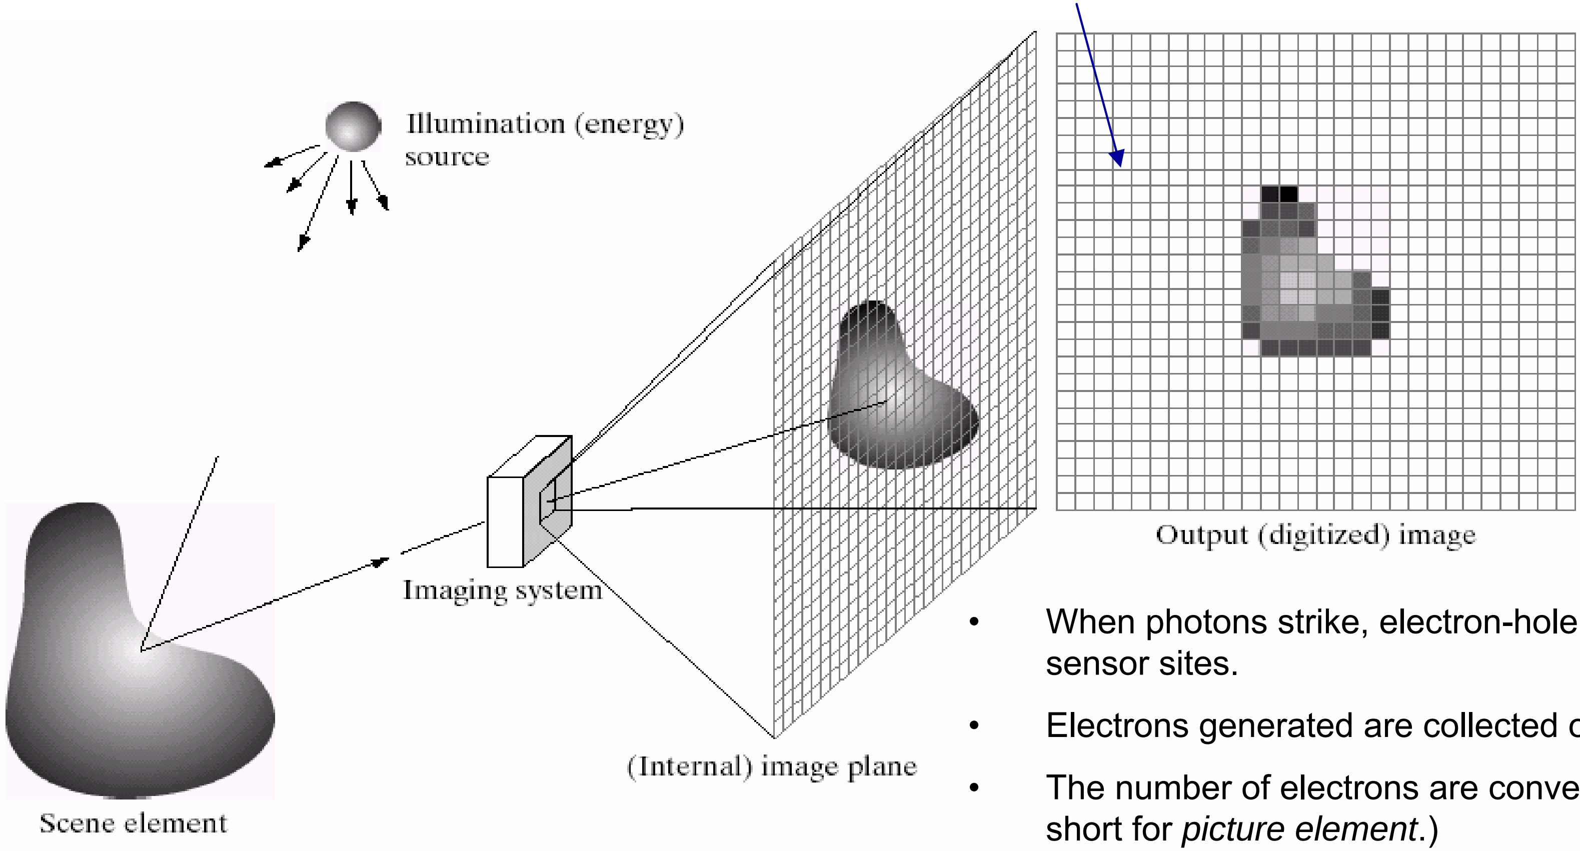

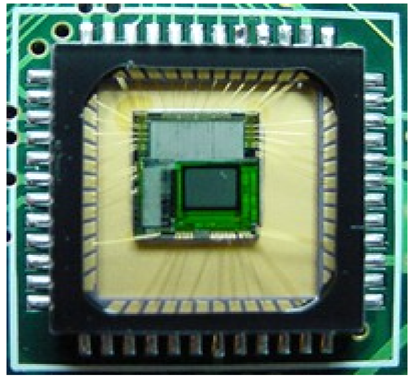

- When photons strike, electron-hole pairs are generated on sensor sites.
- Electrons generated are collected over a certain period of time.
- The number of electrons are converted to pixel values. (Pixel is short for *picture element*.)

a b c d e

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 capture

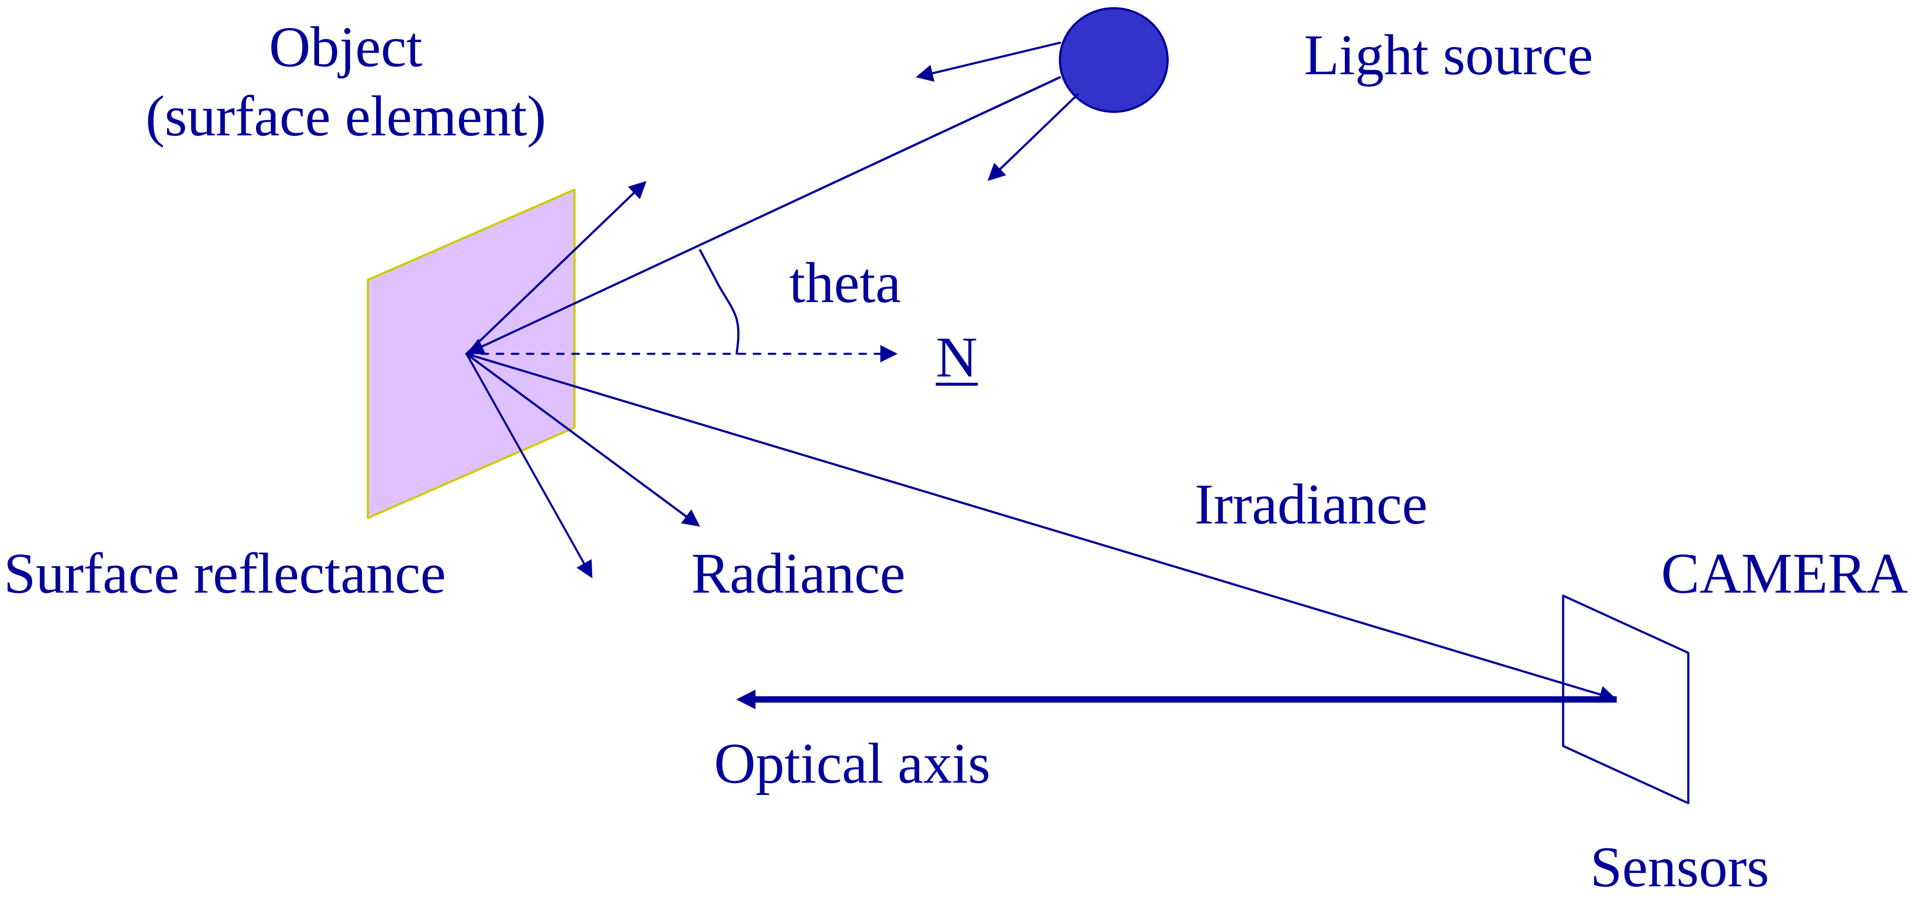

Slide 5

VM1

- radianza: energia che viene emessa dall'elemento di superficie
- irradianza: energia che colpisce la camera e dipende da lo spettro della luce, la riflettanza della superficie (che cambia lo spettro) e la sensibilità spettrale del sensore

swan; 14/01/2004

Digital Image Acquisition

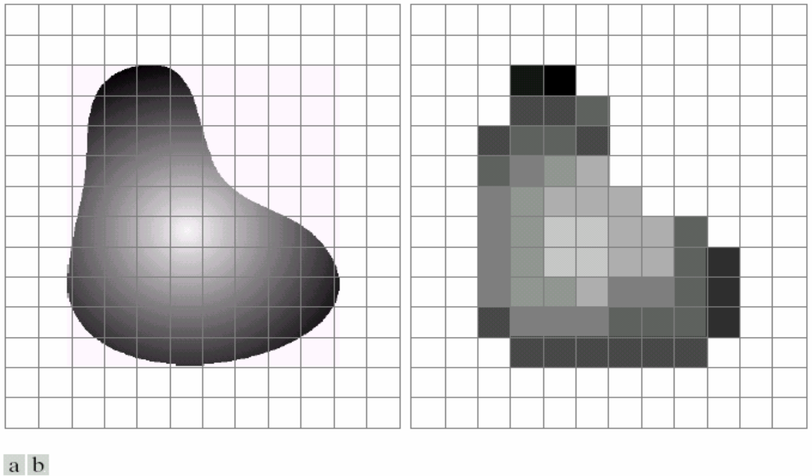

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

Two types of discretization:

1. There are finite number of pixels. (sampling → Spatial resolution)
2. The amplitude of pixel is represented by a finite number of bits. (Quantization → Gray-scale resolution)

Digital Image Acquisition

Take a look at this cross section

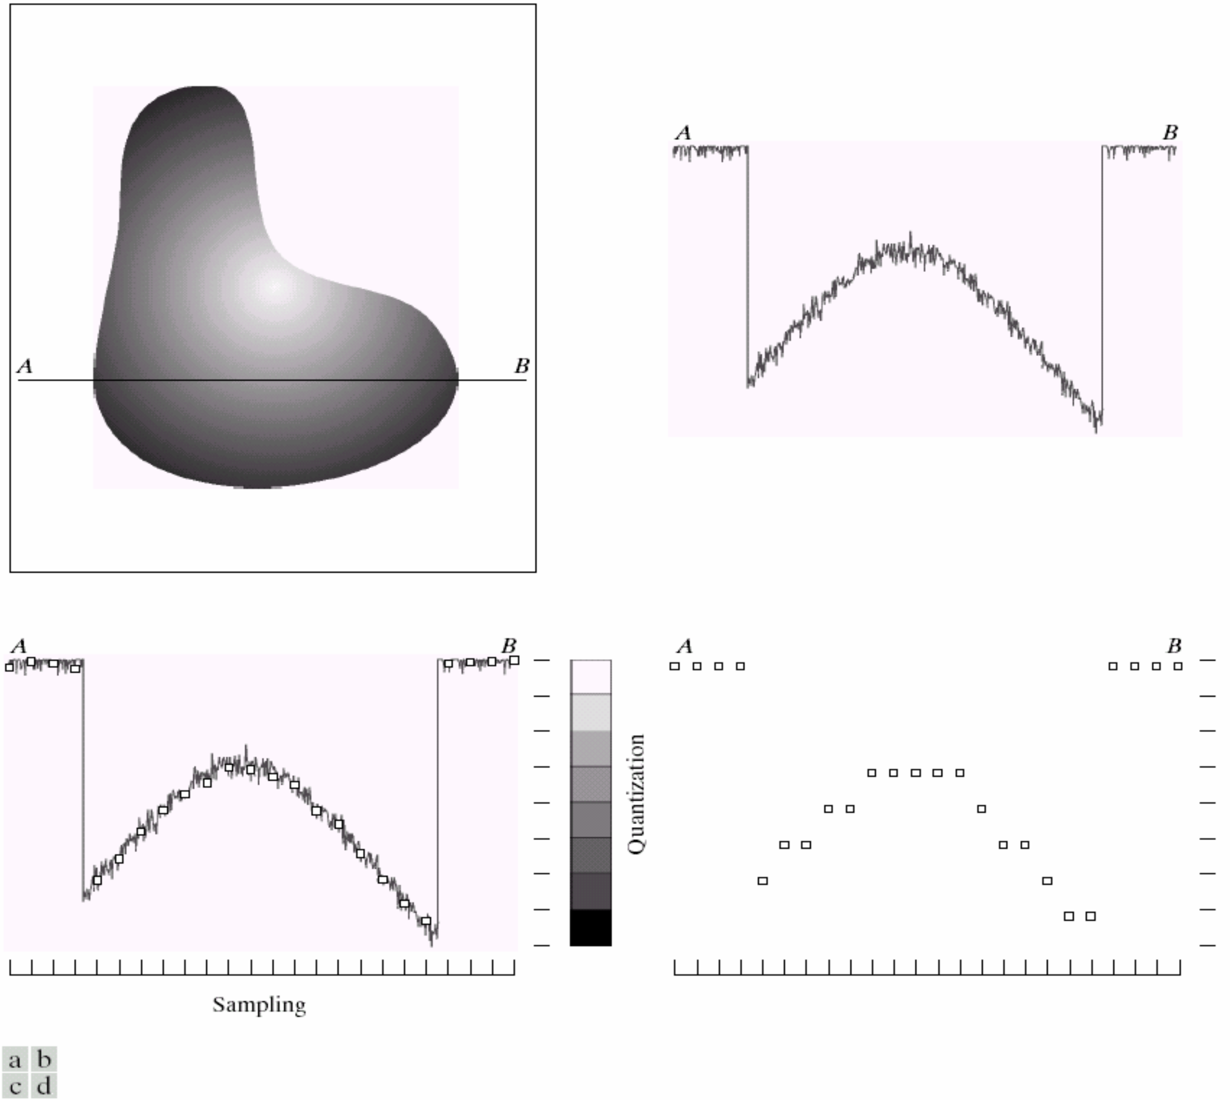

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from A to B in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

Digital Image Acquisition

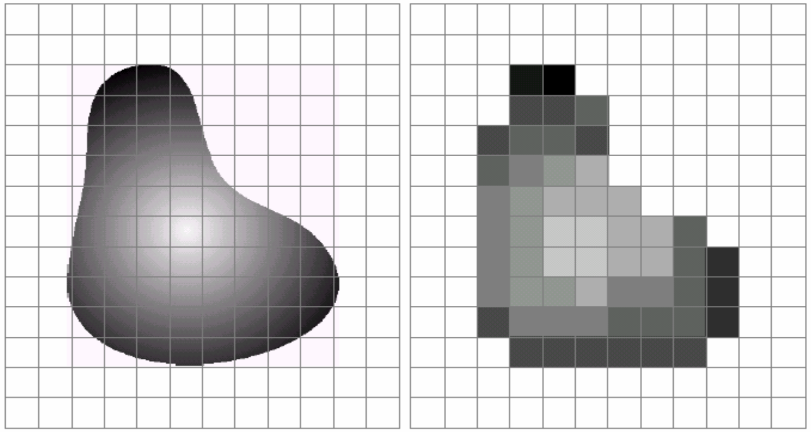

a b

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

- **256x256** - Found on very cheap cameras, this resolution is so low that the picture quality is almost always unacceptable. This is 65,000 total pixels.
- **640x480** - This is the low end on most "real" cameras. This resolution is ideal for e-mailing pictures or posting pictures on a Web site.
- **1216x912** - This is a "megapixel" image size -- 1,109,000 total pixels -- good for printing pictures.
- **1600x1200** - With almost 2 million total pixels, this is "high resolution." You can print a 4x5 inch print taken at this resolution with the same quality that you would get from a photo lab.
- **2240x1680** - Found on 4 megapixel cameras -- the current standard -- this allows even larger printed photos, with good quality for prints up to 16x20 inches.
- **4064x2704** - A top-of-the-line digital camera with 11.1 megapixels takes pictures at this resolution. At this setting, you can create 13.5x9 inch prints with no loss of picture quality.

Basics: graylevel images

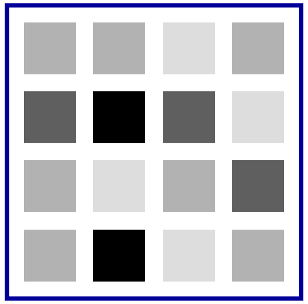

100	100	200	90
50	0	50	200
100	200	100	50
100	0	200	100

Images : Matrices of numbers
Image processing : Operations among numbers
bit depth : number of bits/pixel
N bit/pixel : 2^{N-1} shades of gray (typically N=8)

Matrix Representation of Images

- A digital image can be written as a matrix

$$x[n_1, n_2] = \begin{bmatrix} x[0,0] & x[0,1] & \cdots & x[0,N-1] \\ x[1,0] & x[1,1] & \cdots & x[1,N-1] \\ \vdots & \vdots & \ddots & \vdots \\ x[M-1,0] & \cdots & \cdots & x[M-1,N-1] \end{bmatrix}_{MxN}$$

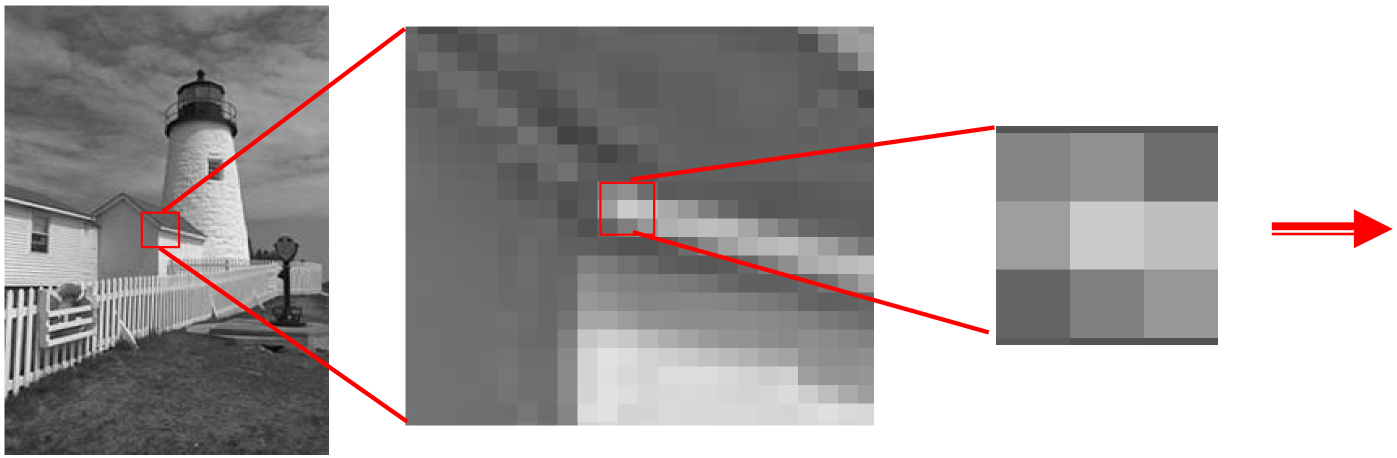

$$\begin{bmatrix} 35 & 45 & 20 \\ 43 & 64 & 52 \\ 10 & 29 & 39 \end{bmatrix}$$

Digital images acquisition

- Analog camera+A/D converter
- Digital cameras
  - CCDs (Charge Coupled Devices)
  - CMOS technology
- In both cases: optics
  - lenses, diaphrams

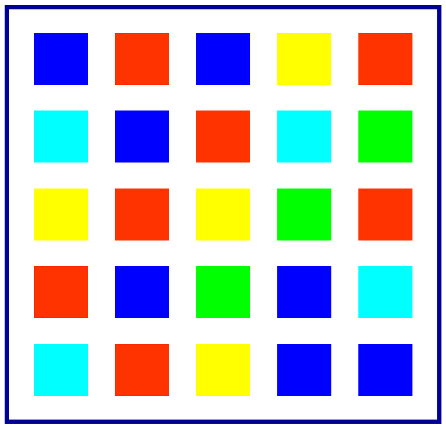

Matrices of photo sensors collecting photons of given wavelength

Features of the capture devices:

- Size and number of photosites
- Noise
- Transfer function of the optical filter

Color images

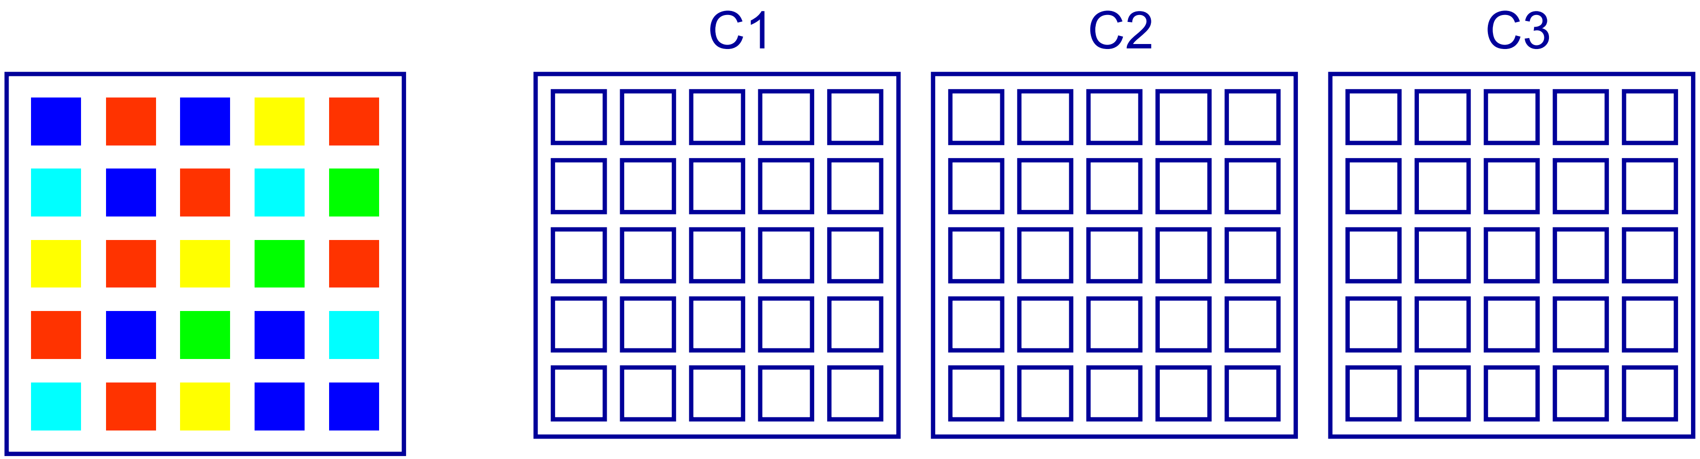

- Each colored pixel corresponds to a *vector* of three values {C1,C2,C3}
- The characteristics of the components depend on the chosen *colorspace* (RGB, YUV, CIELab,..)

Digital Color Images

- $x_R[n_1, n_2]$
  $x_G[n_1, n_2]$
  $x_B[n_1, n_2]$

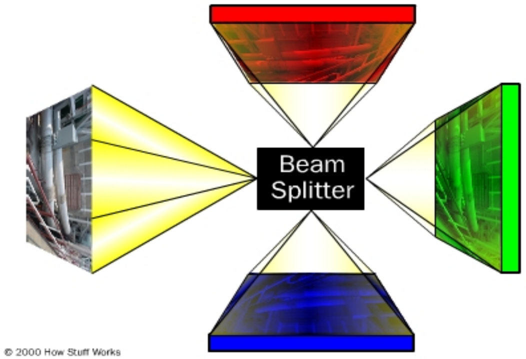

Color channels

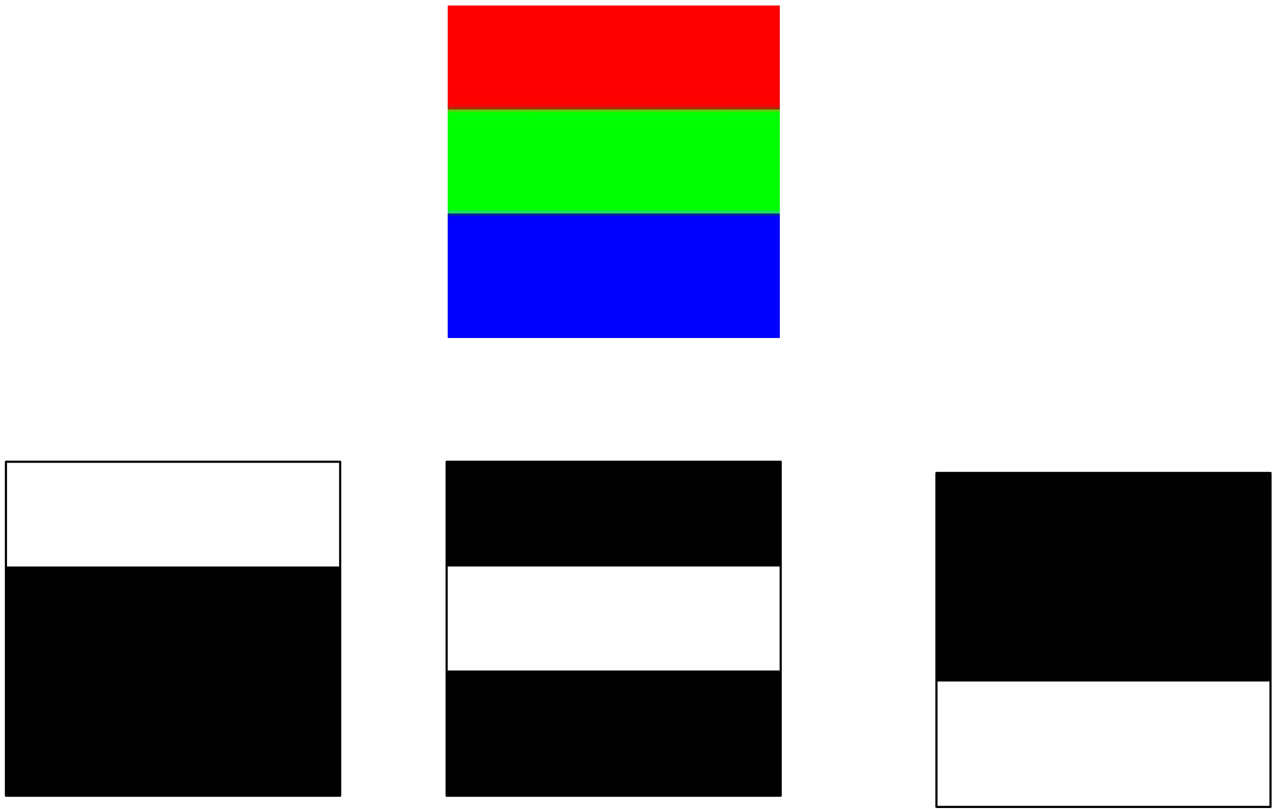

Red Green Blue

Color channels

Red

Green

Blue

The physical perspective

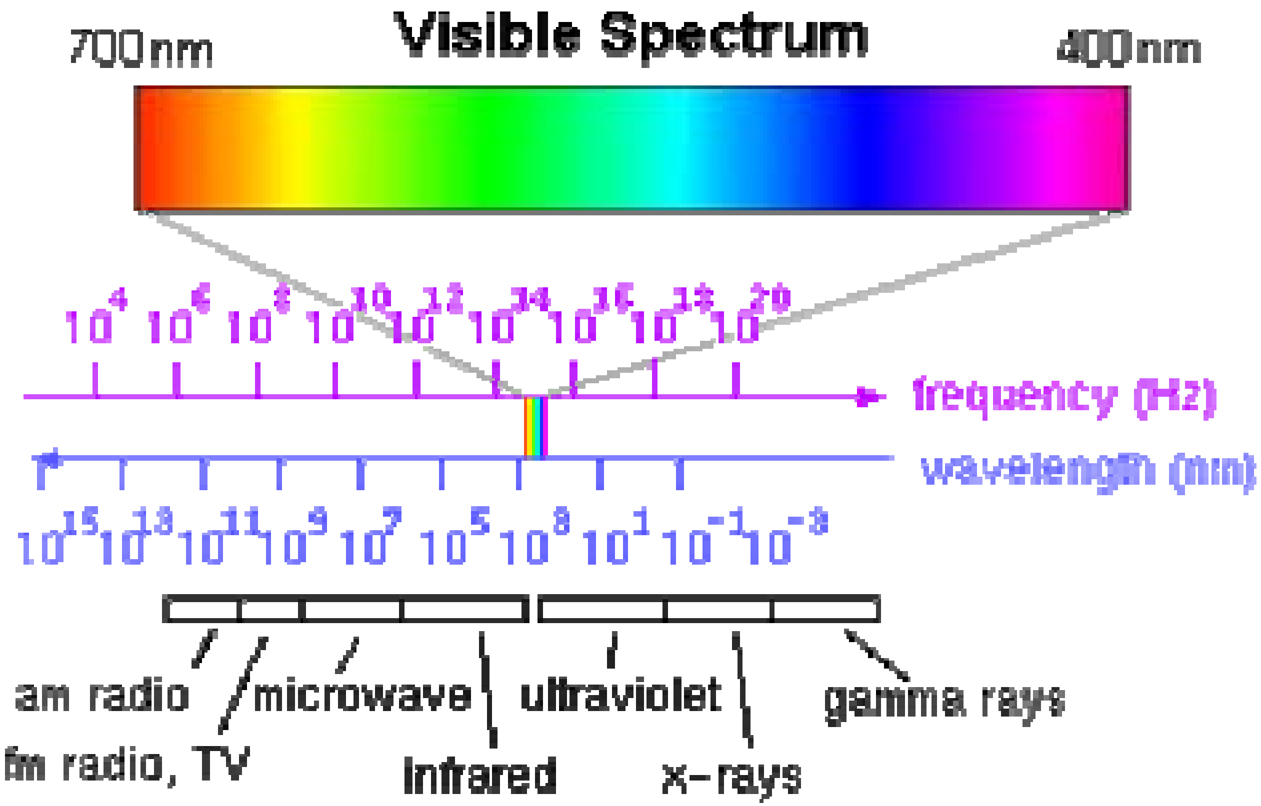

The perceptual perspective

Simultaneous contrast

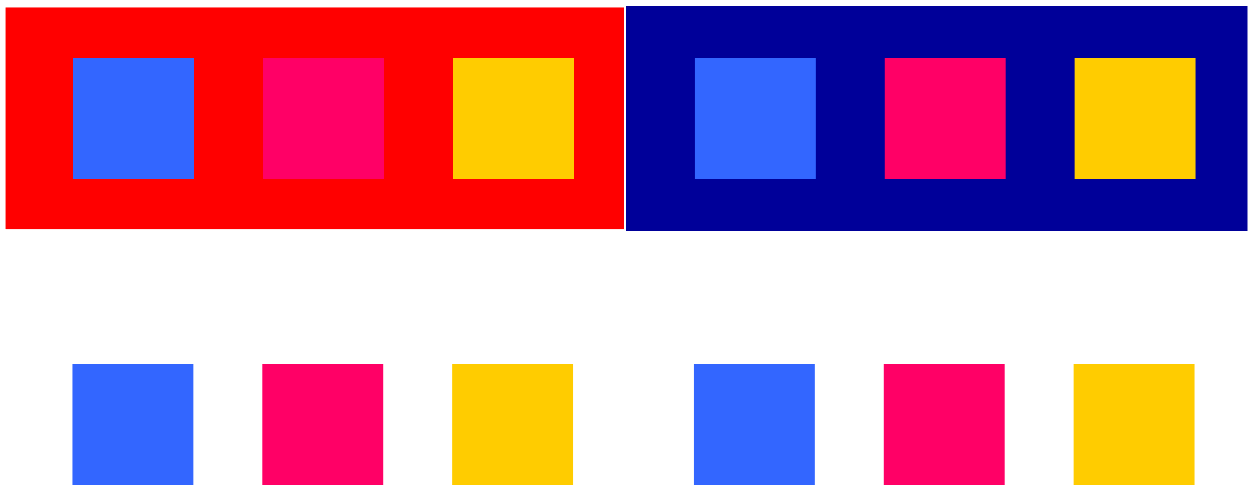

Color

- Chromatic induction

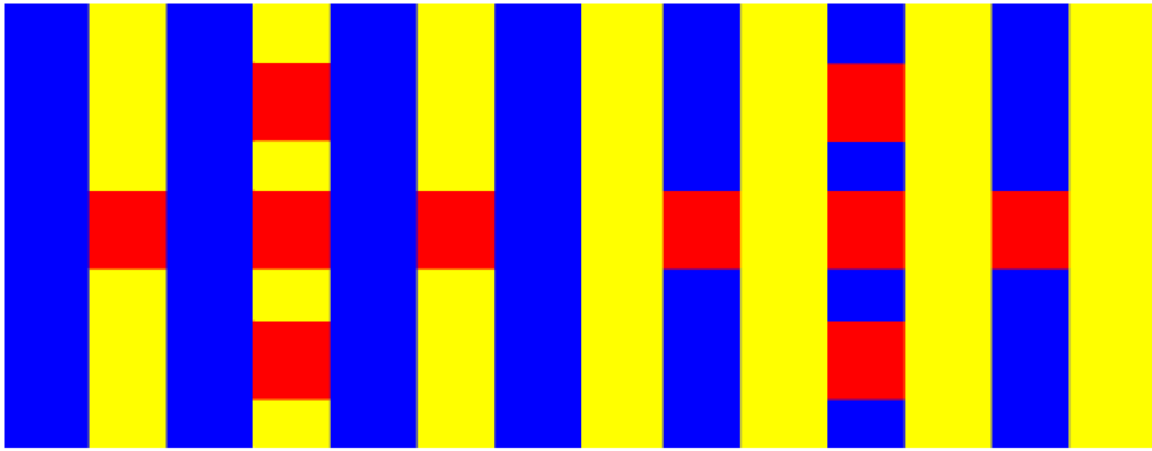

Color

- Human vision
  - Color encoding (receptor level)
  - Color perception (post-receptoral level)
  - Color semantics (cognitive level)
- Colorimetry
  - Spectral properties of radiation
  - Physical properties of materials

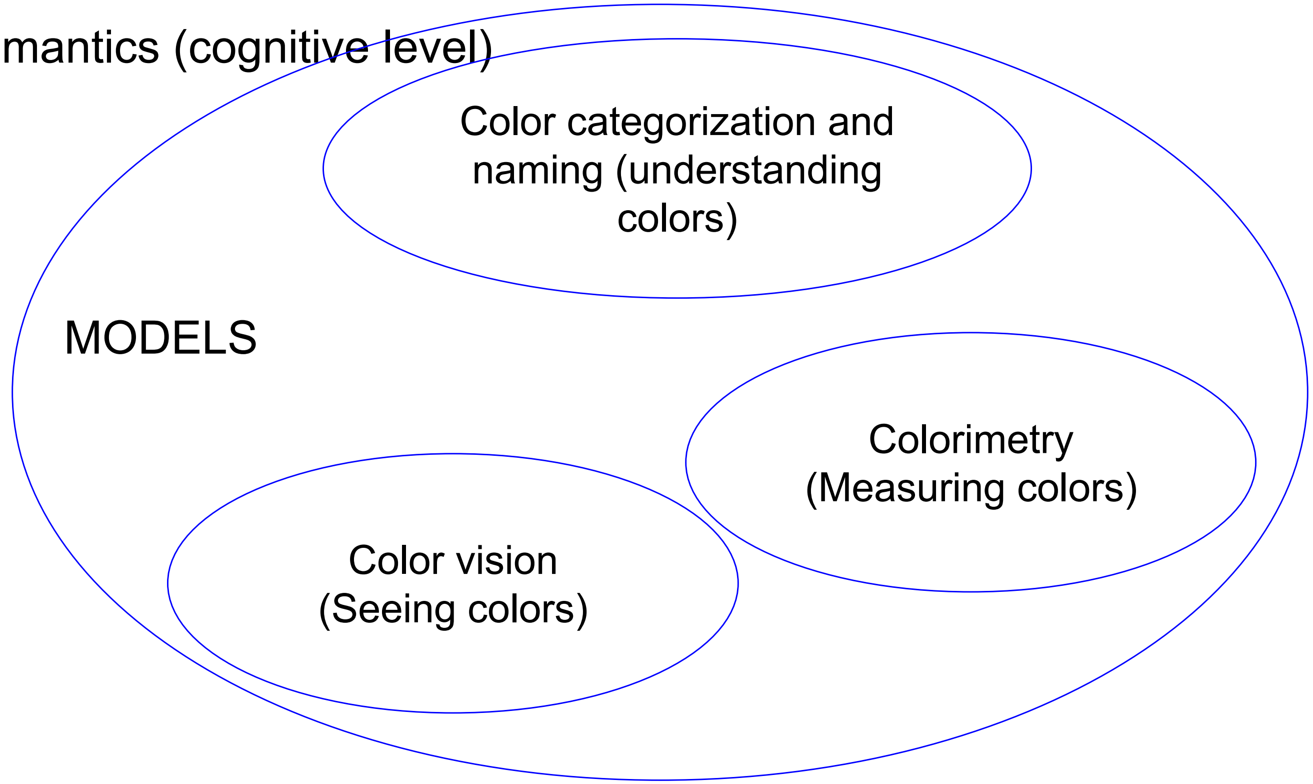

Bayer matrix

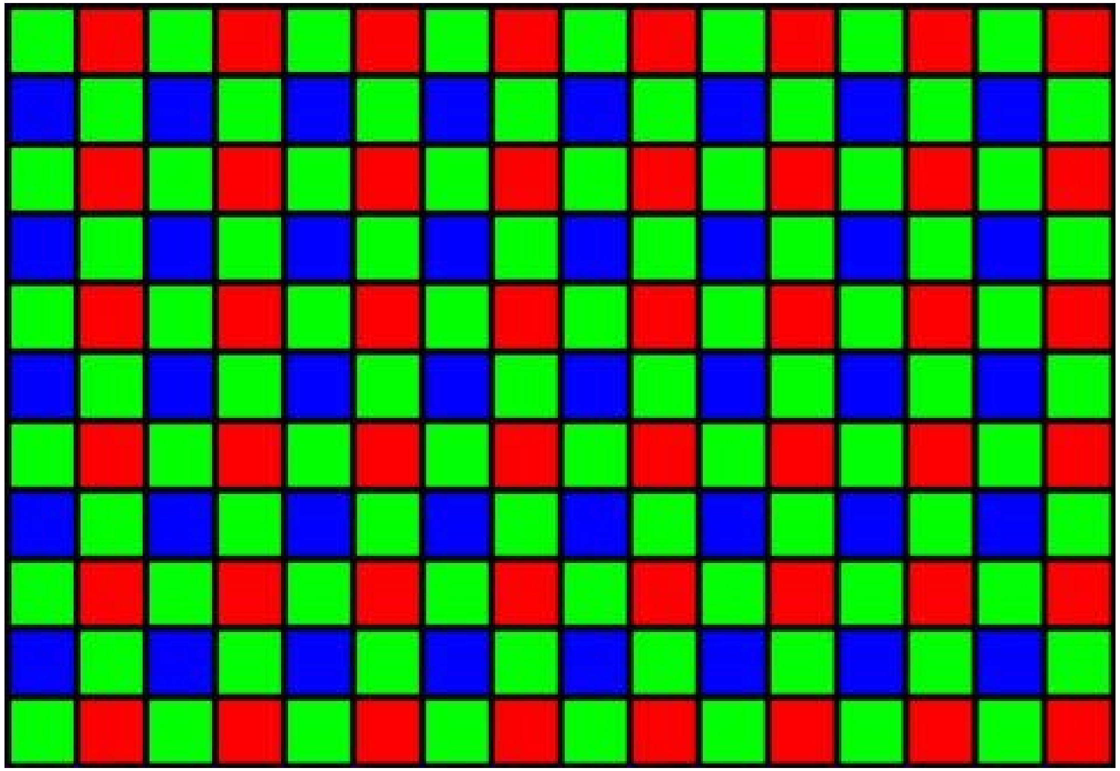

Typical sensor topology in CCD devices. The green is twice as numerous as red and blue.

Displays

LCD

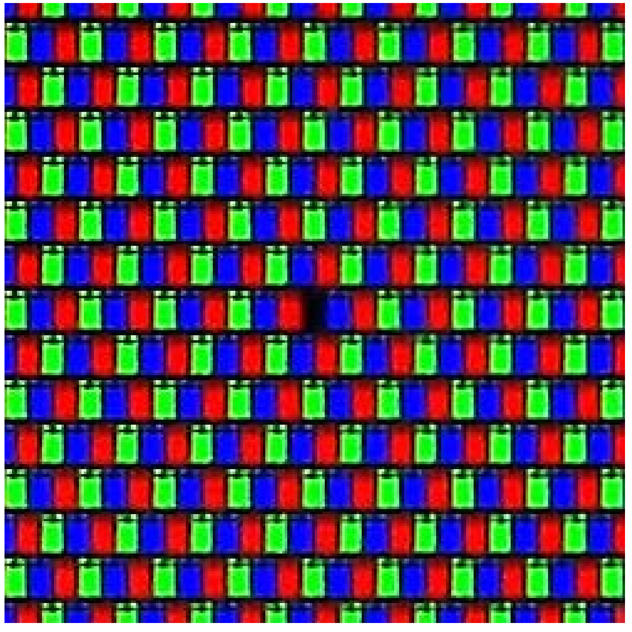

CRT

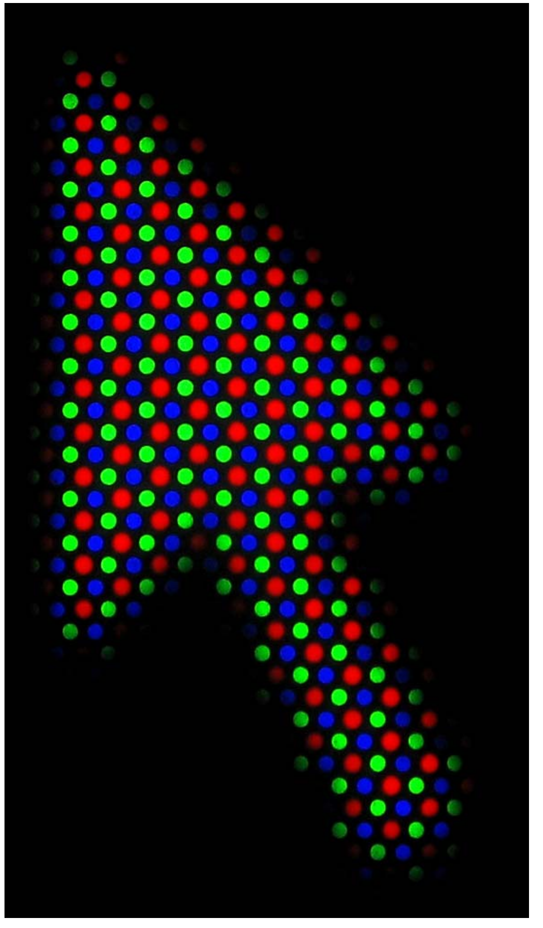

Color Displays

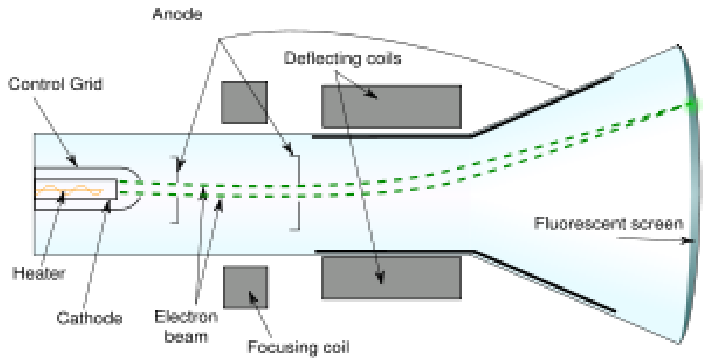

CRT

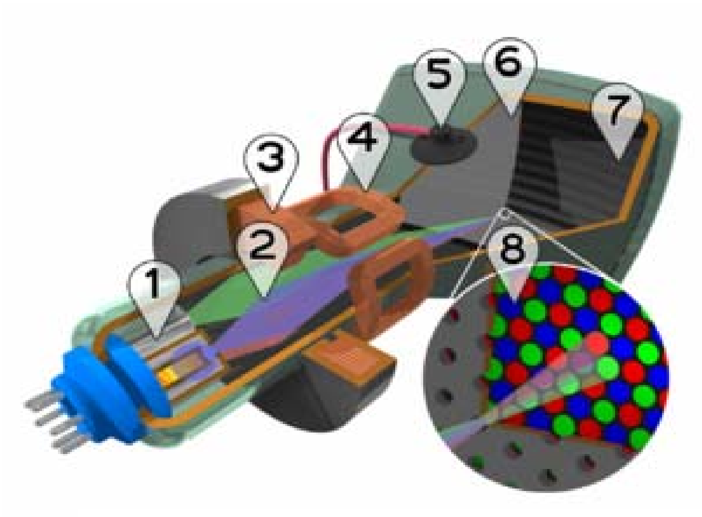

LCD

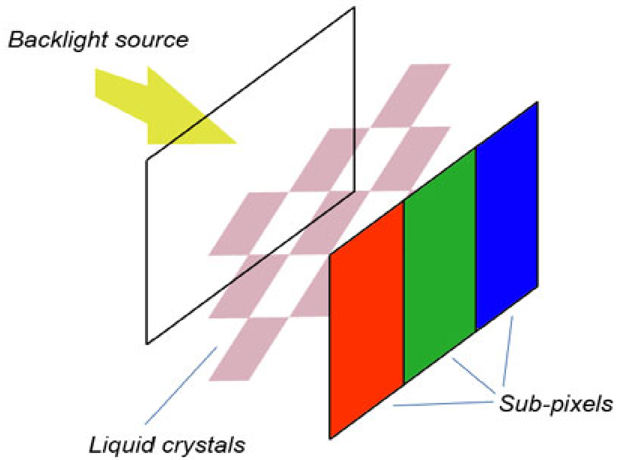

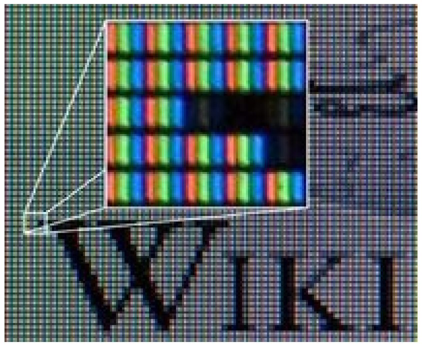

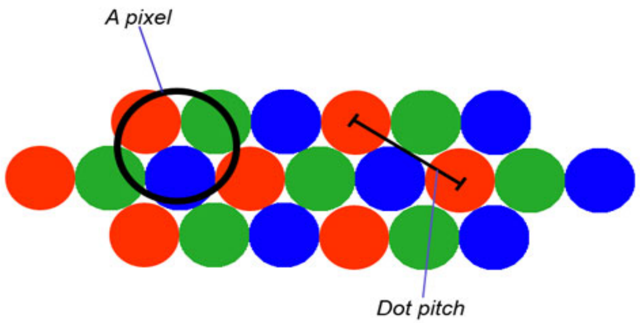

Polarize to control the amount of light passed.

Color imaging

- Color reproduction
  - Printing, rendering
- Digital photography
  - High dynamic range images
  - Mosaicking
  - Compensation for differences in illuminant (CAT: chromatic adaptation transforms)
- Post-processing
  - Image enhancement
- Coding
  - Quantization based on color CFSs (contrast sensitivity function)
  - Downsampling of chromatic channels with respect to luminance

Some definitions

- Digital images
  - Sampling+quantization
- Sampling
  - Determines the graylevel value of each pixel
    - Pixel = picture element
- Quantization
  - Reduces the resolution in the graylevel value to that set by the machine precision
- Images are stored as matrices of unisigned chars

Resolution

- Sensor resolution (CCD): Dots Per Inch (DPI)
  - Number of individual dots that can be placed within the span of one linear inch (2.54 cm)
- Image resolution
  - Pixel resolution: NxM
  - Spatial resolution: Pixels Per Inch (PPI)
  - Spectral resolution: bandwidth of each spectral component of the image
    - Color images: 3 components (R,G,B channels)
    - Multispectral images: many components (ex. SAR images)
  - Radiometric resolution: Bits Per Pixel (bpp)
    - Graylevel images: 8, 12, 16 bpp
    - Color images: 24bpp (8 bpp/channel)
  - Temporal resolution: for movies, number of frames/sec
    - Typically 25 Hz (=25 frames/sec)

VM2

Slide 25

VM2 da book Shapiro
swan; 10/04/2003

Example: pixel resolution

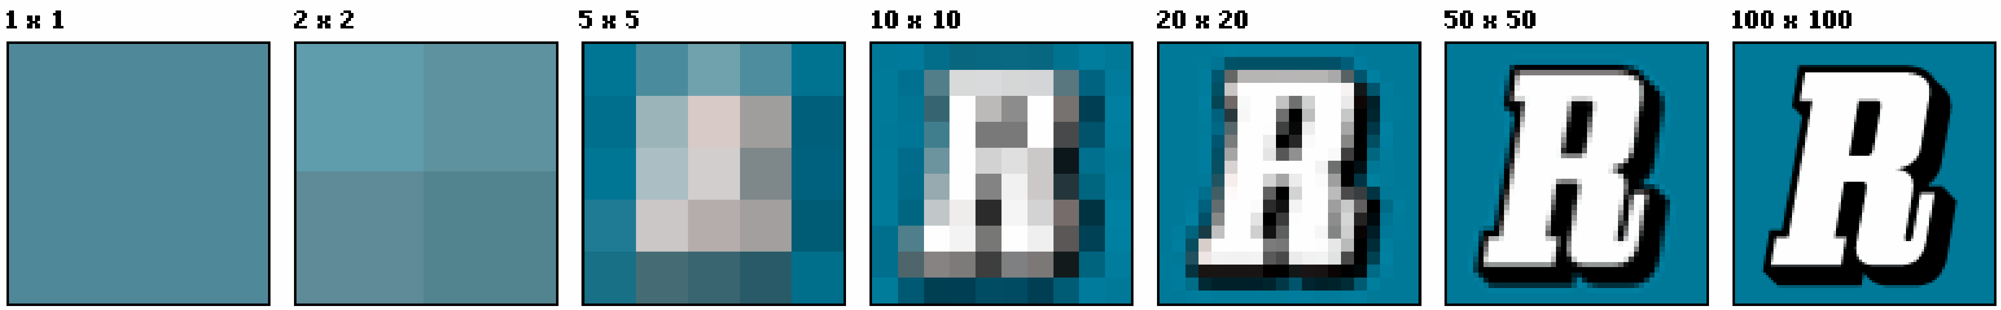

Image Resolution

Don't confuse image size and resolution.

a b c
d e f

FIGURE 2.20 (a) 1024 × 1024, 8-bit image. (b) 512 × 512 image resampled into 1024 × 1024 pixels by row and column duplication. (c) through (f) 256 × 256, 128 × 128, 64 × 64, and 32 × 32 images resampled into 1024 × 1024 pixels.

Bit Depth – Grayscale Resolution

8 bits

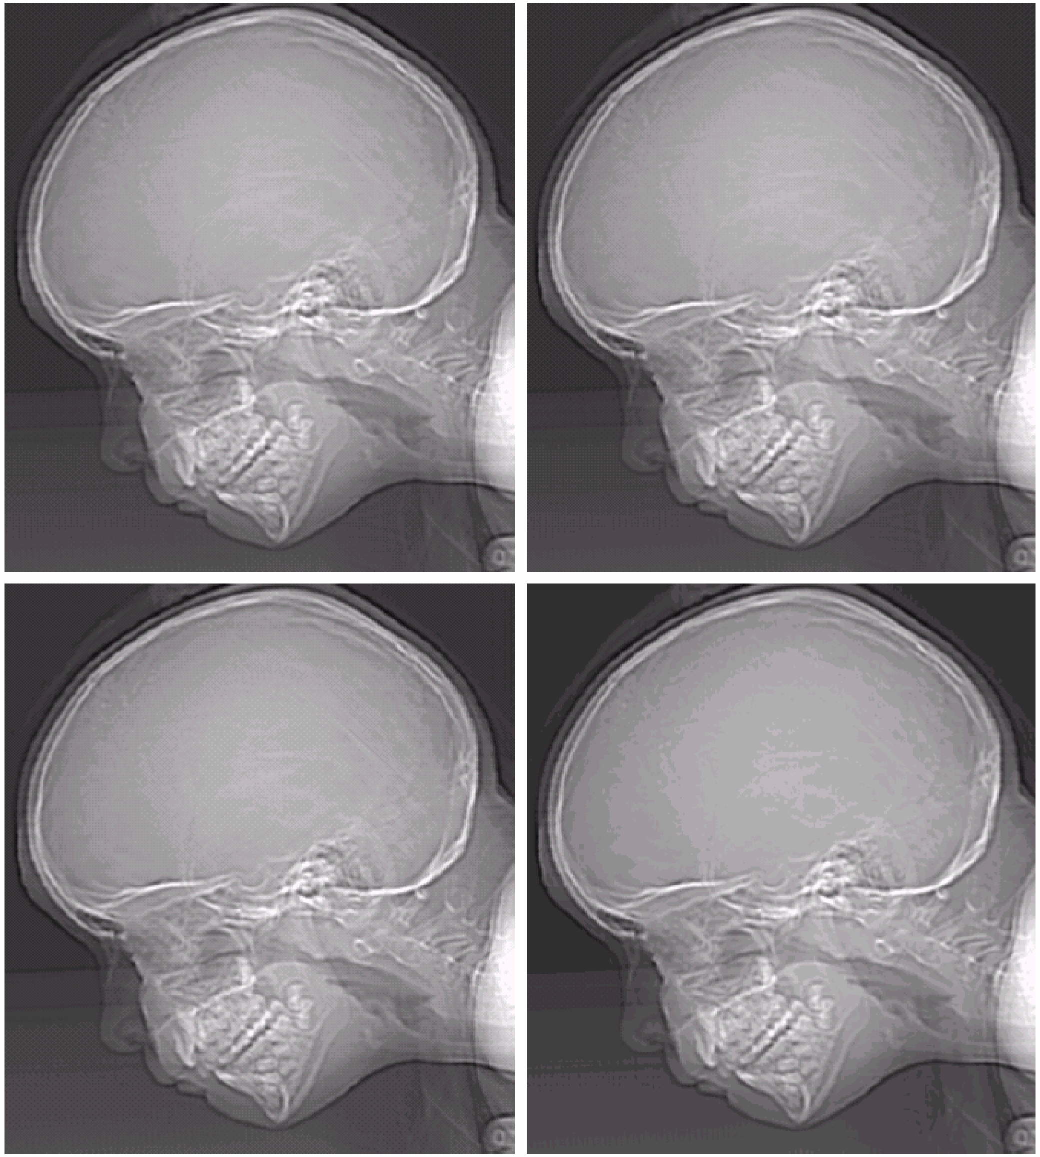

a b
c d

FIGURE 2.21 (a) 452 × 374, 256-level image. (b)–(d) Image displayed in 128, 64, and 32 gray levels, while keeping the spatial resolution constant.

7 bits

6 bits

5 bits

Bit Depth – Grayscale Resolution

e f
g h

FIGURE 2.21 *(Continued)* (e)–(h) Image displayed in 16, 8, 4, and 2 gray levels. (Original courtesy of Dr. David R. Pickens, Department of Radiology & Radiological Sciences, Vanderbilt University Medical Center.)

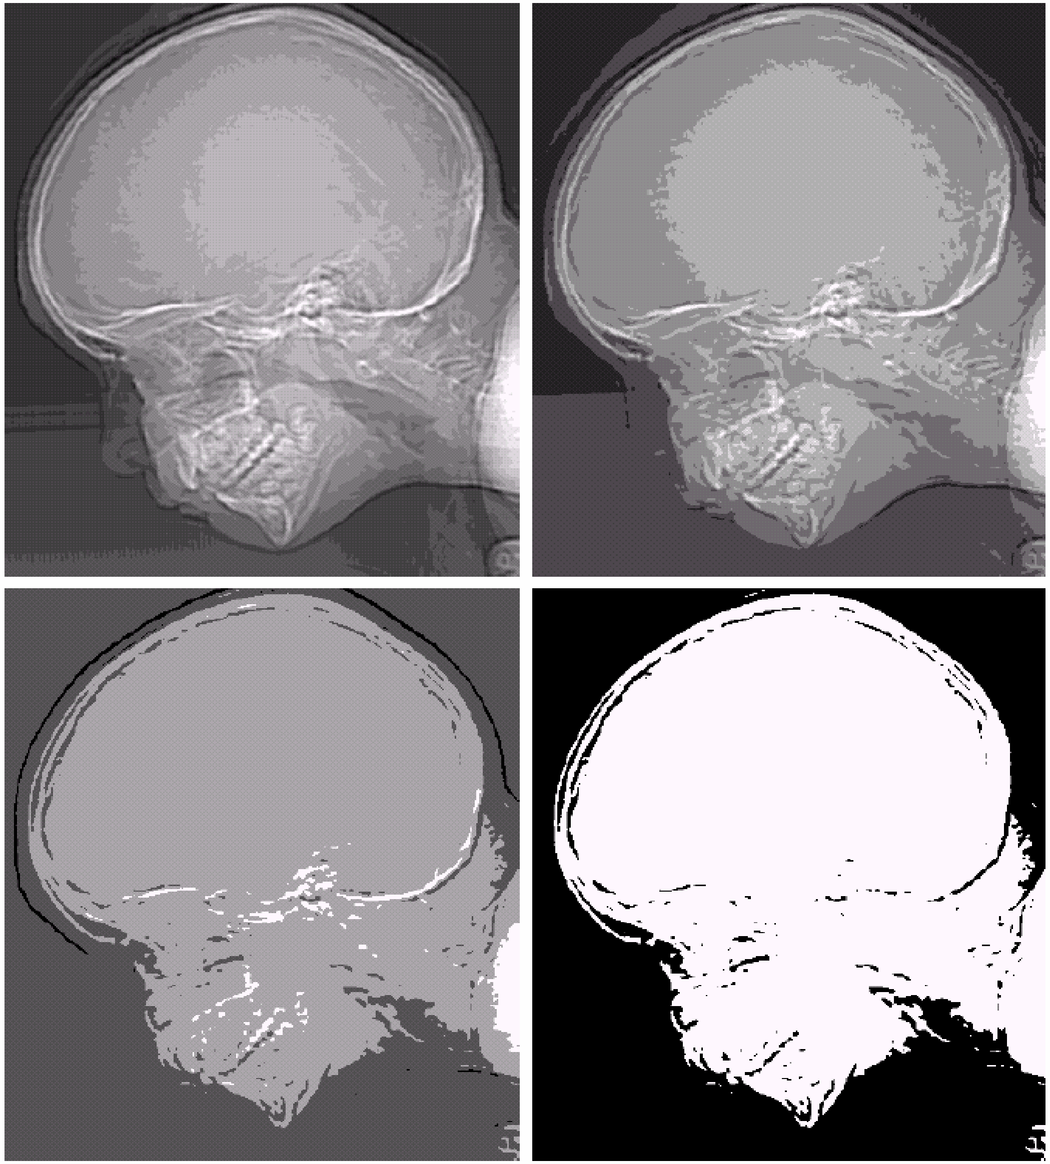

4 bits

3 bits

2 bits

1 bit

File format

- Many image formats (about 44)
- BMP, lossless
- TIFF, lossless/lossy
- GIF (Graphics Interchange Format)
  - Lossless, 256 colors, copyright protected
- JPEG (Joint Photographic Expert Group)
  - Lossless and lossy compression
  - 8 bits per color (red, green, blue) for a 24-bit total
- PNG (Portable Network Graphics)
  - Freewere
  - supports truecolor (16 million colours)

Sampling in 2D

Sampling in 1D

Continuous time signal

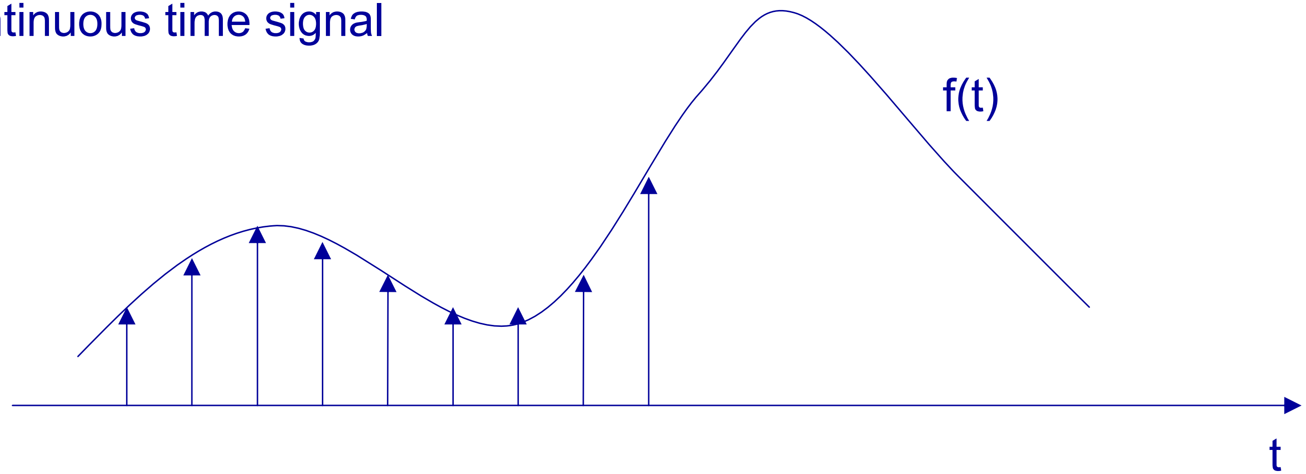

Discrete time signal

$$f[k] = f(kT_s) = f(t)\sum_k \delta(t - kT_s)$$

comb

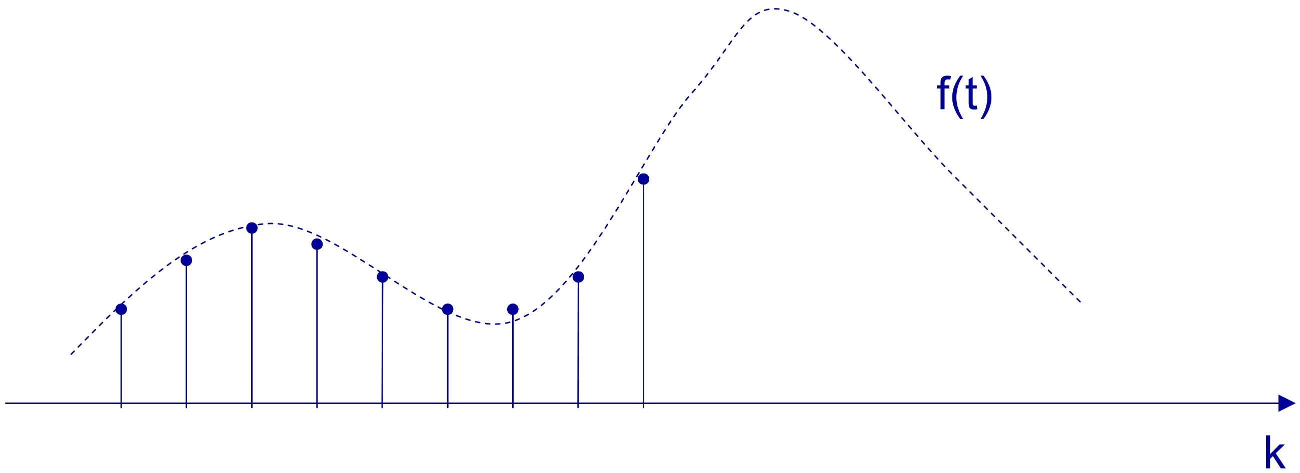

Nyquist theorem (1D)

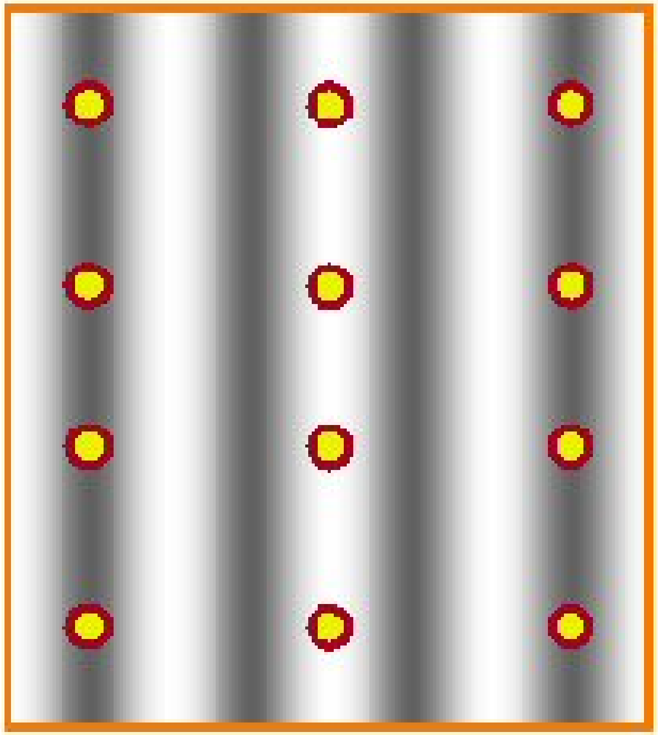

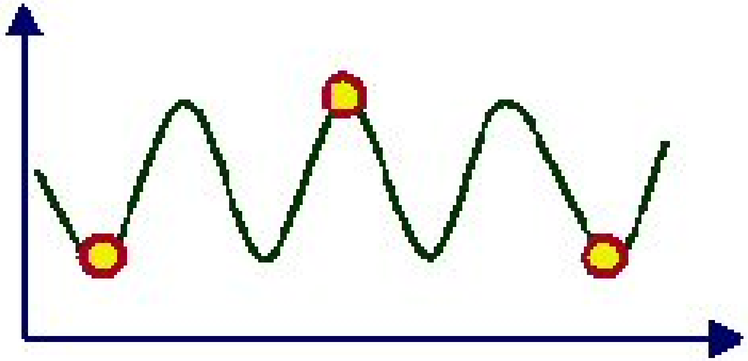

At least 2 sample/period are needed to represent a periodic signal

$$T_s \leq \frac{1}{2}\frac{2\pi}{\omega_{\max}}$$

$$\omega_s = \frac{2\pi}{T_s} \geq 2\omega_{\max}$$

Delta pulse

- 1D Dirac pulse

$$\begin{cases} \delta(x) = 1 \text{ if } x=0 \\ \delta(x) = 0 \text{ else} \end{cases}$$

- 2D Dirac pulse

$$\begin{cases} \delta(x,y) = 1 \text{ if } x=0 \text{ and } y=0 \\ \delta(x,y) = 0 \quad \text{else} \end{cases}$$

which corresponds to :

$\delta(x,y) = \delta(x)\ \delta(y)$

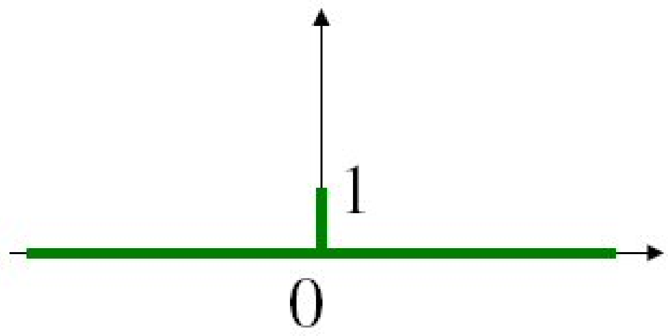

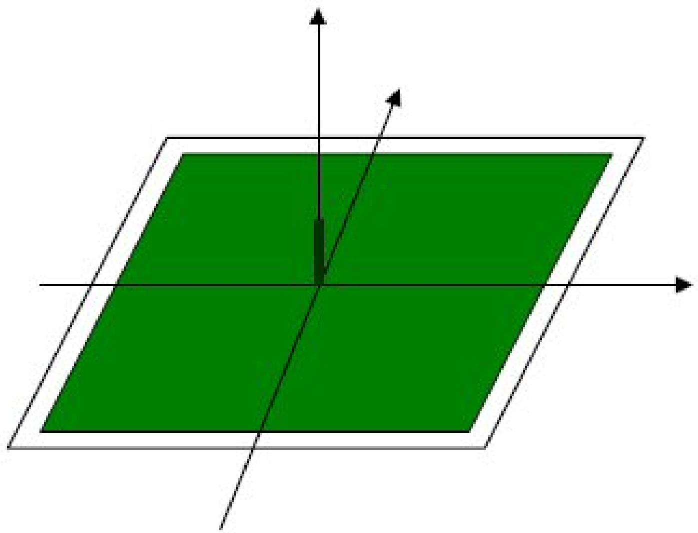

Dirac *brush*

- 1D sampling: Dirac comb (or Shah function)

- 2D sampling : Dirac « brush »

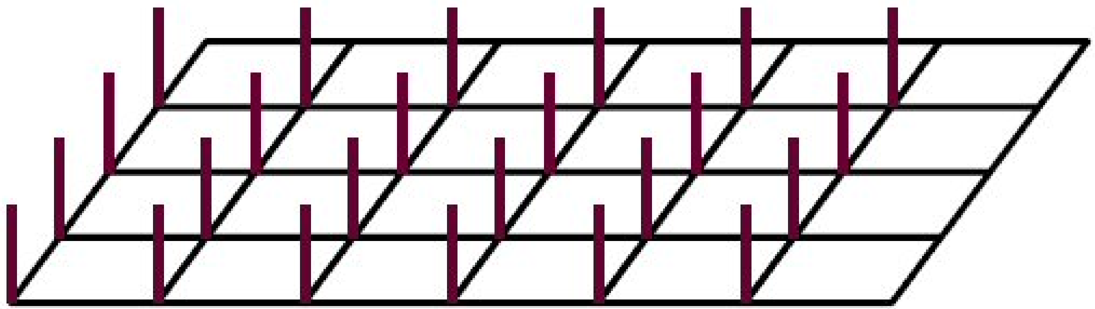

Comb

- Extended comb :

$$p_x(x,y) = \sum_{m=-\infty}^{\infty} \delta(x - m\Delta x)$$

- Comb :

$$p_x(x,y) = \delta(y) \sum_{m=-\infty}^{\infty} \delta(x - m\Delta x)$$

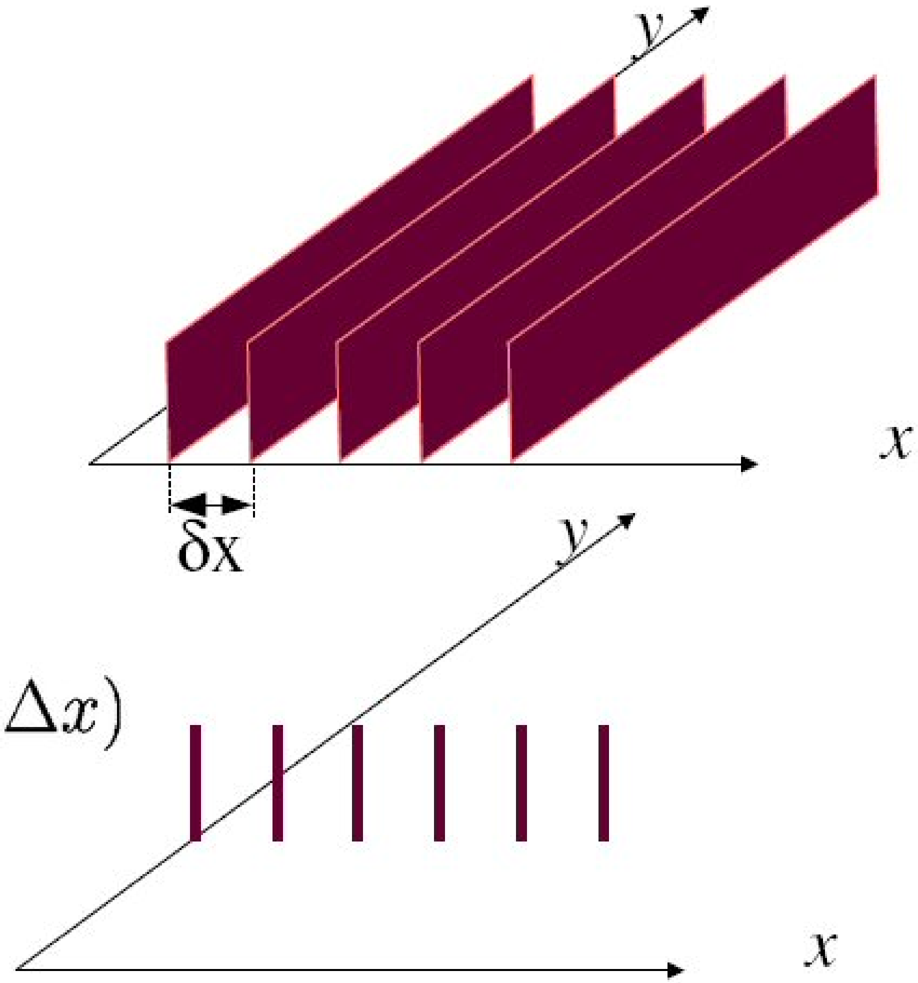

Brush

- Brush = product of 2 extended combs

$$p_x(x,y) = \sum_{m=-\infty}^{\infty} \delta(x - m\Delta x)$$

$$p_y(x,y) = \sum_{m=-\infty}^{\infty} \delta(y - n\Delta y)$$

$$b(x,y) = p_x(x,y)p_y(x,y)$$

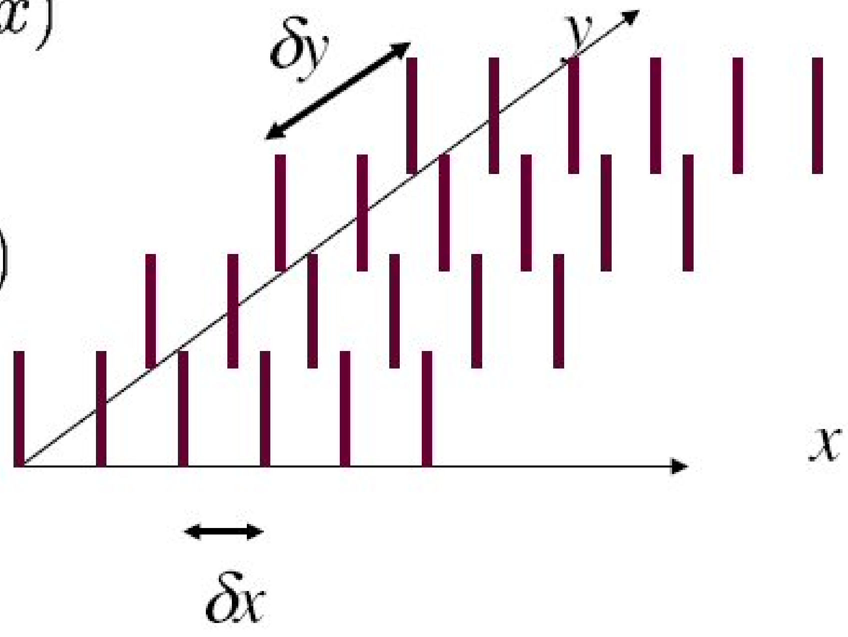

Nyquist theorem

- Sampling in p-dimensions

$$s_T(\vec{x}) = \sum_{k \in Z^p} \delta(\vec{x} - kT)$$

$$f_T(\vec{x}) = f(\vec{x}) s_T(\vec{x})$$

2D spatial domain

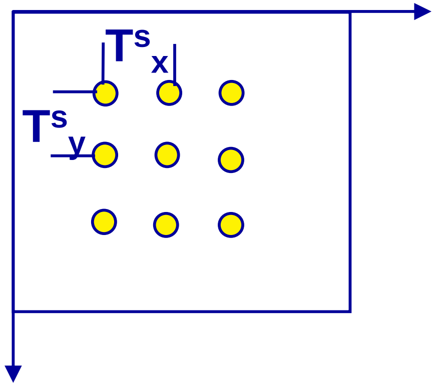

2D Fourier domain

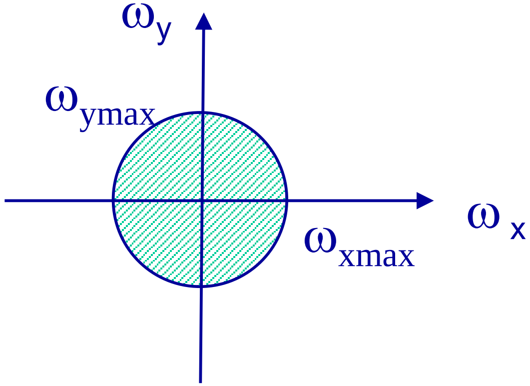

- Nyquist theorem

$$\begin{cases} \omega_x^s \geq 2\omega_{x\max} \\ \omega_y^s \geq 2\omega_{y\max} \end{cases} \Rightarrow \begin{cases} T_x^s \leq 2\pi \dfrac{1}{2\omega_{x\max}} \\ T_y^s \leq 2\pi \dfrac{1}{2\omega_{y\max}} \end{cases}$$

Spatial aliasing

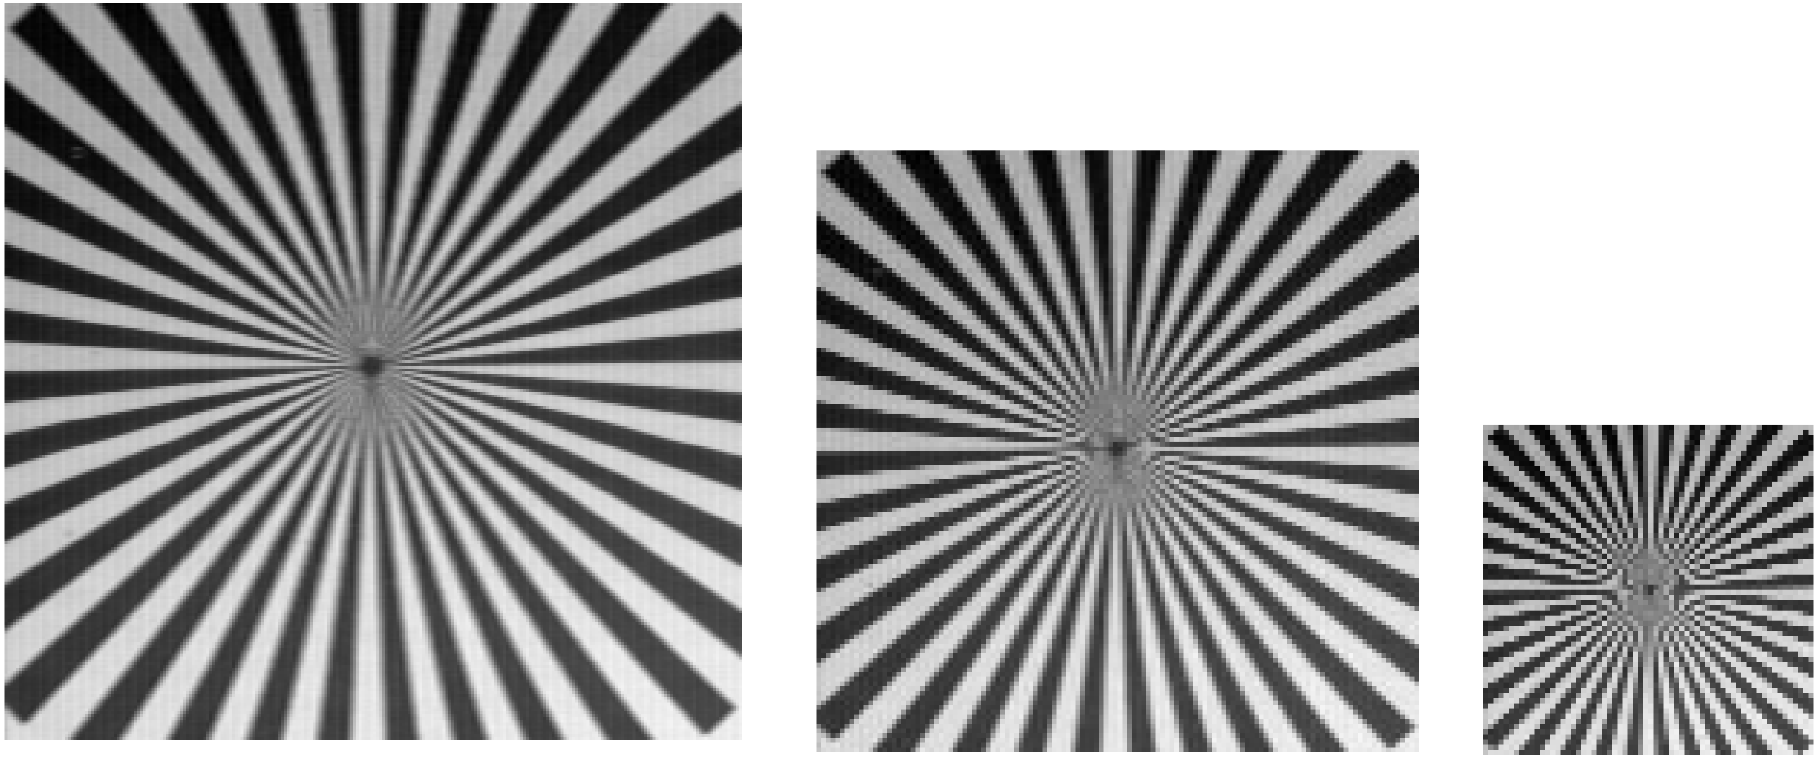

Resampling

- Change of the sampling rate
  - Increase of sampling rate: Interpolation or upsampling
    - Blurring, low visual resolution
  - Decrease of sampling rate: Rate reduction or downsampling
    - Aliasing and/or loss of spatial details

Downsampling

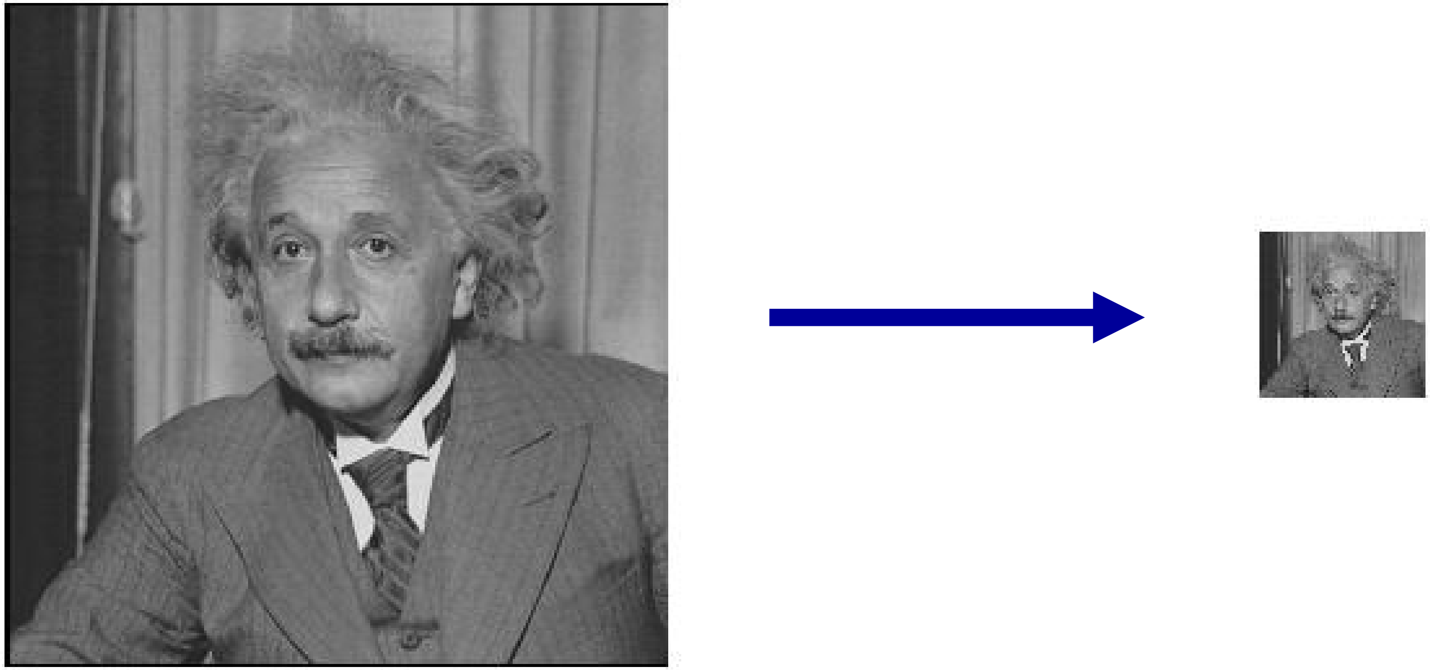

Upsampling

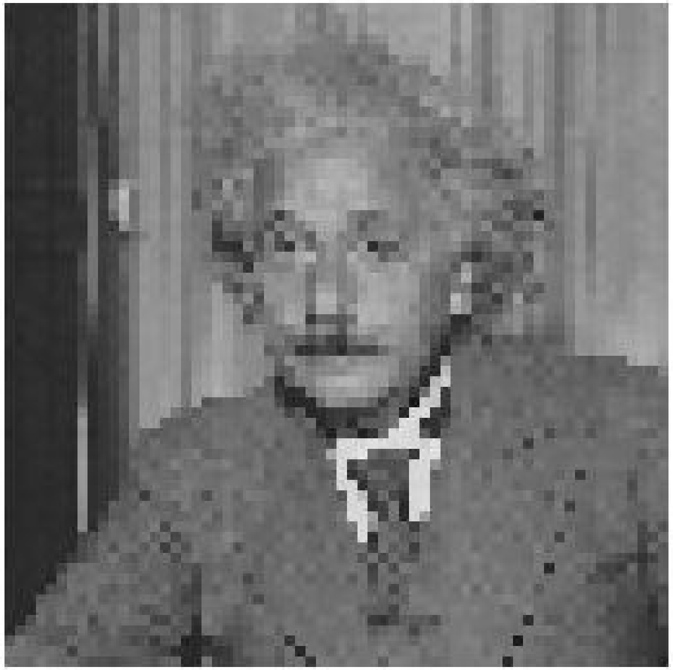

nearest neighbor (NN)

Upsampling

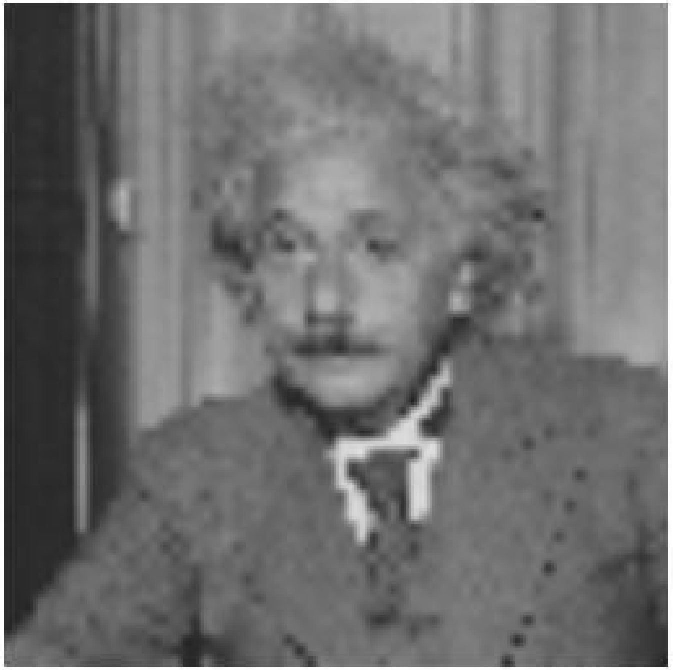

bilinear

Upsampling

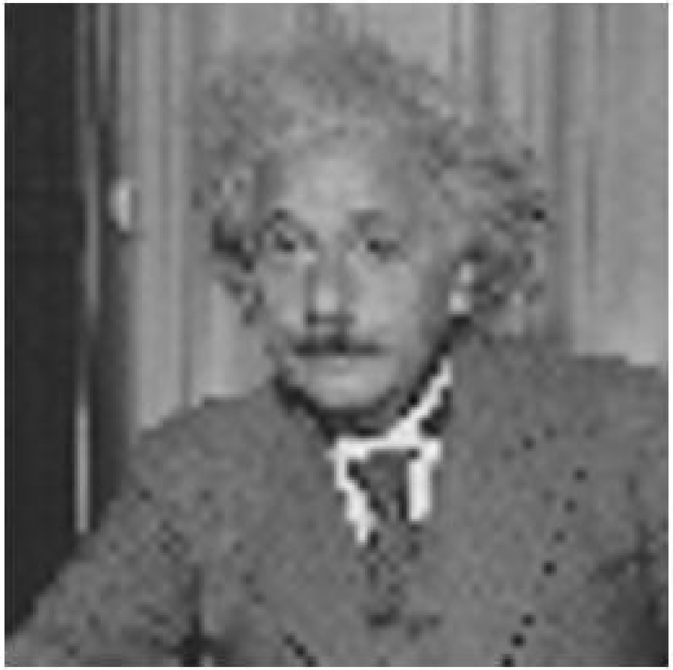

bicubic

Quantization

Quantization

- A/D conversion ⇨ quantization

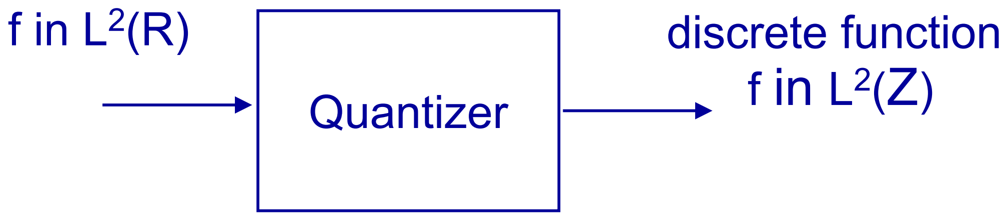

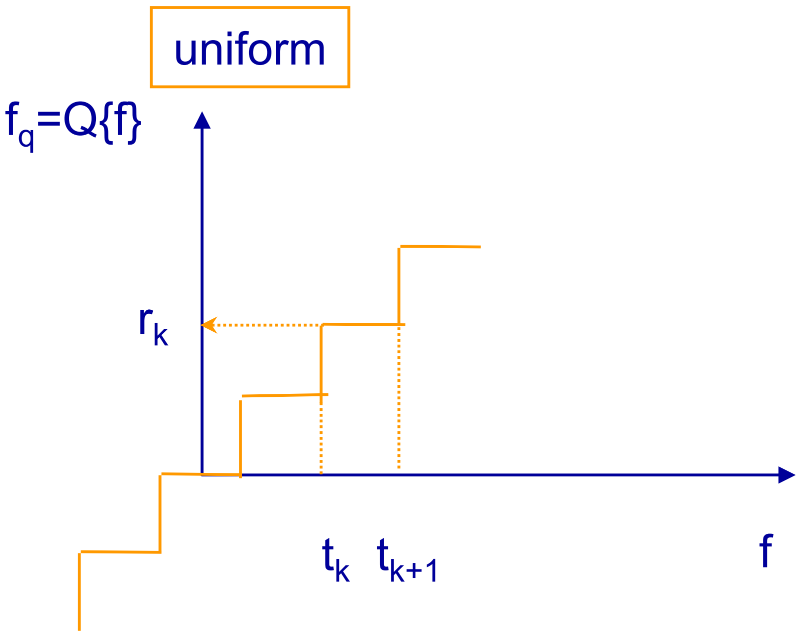

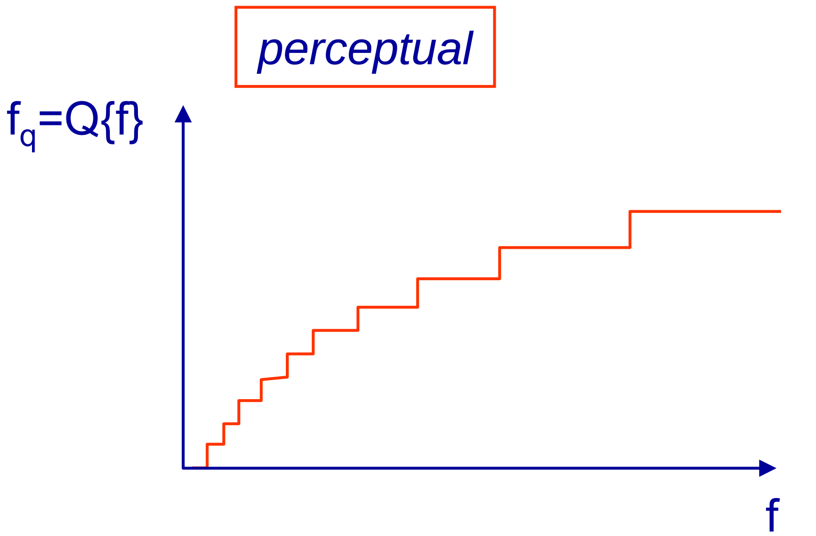

The sensitivity of the eye decreases increasing the background intensity (Weber law)

Quantization

Signal before (blue) and after quantization (red) Q

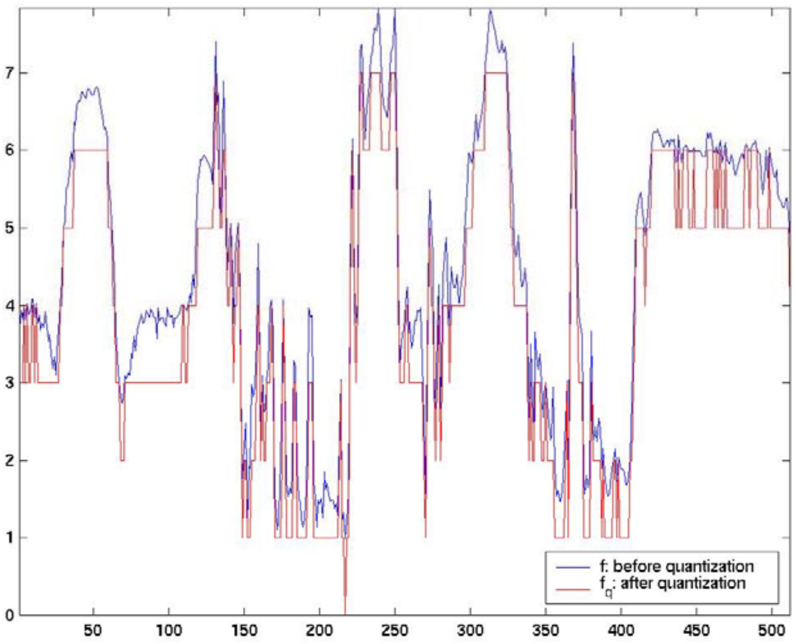

Equivalent noise: $n=f_q - f$

additive noise model: $f_q=f+n$

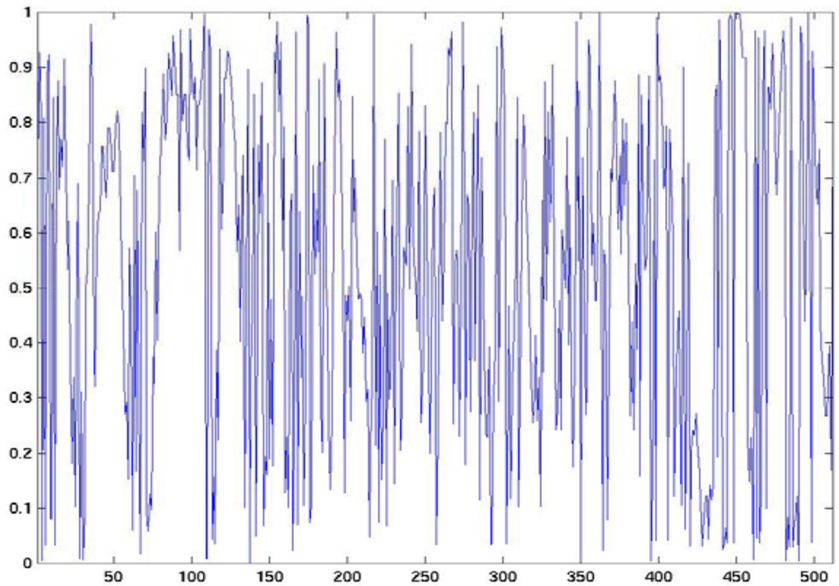

Quantization

original

5 levels

10 levels

50 levels

Distortion measure

- Distortion measure

$$D = \mathrm{E}\left[\left(f_Q - f\right)^2\right] = \sum_{k=0}^{K} \int_{t_k}^{t_{k+1}} \left(f_Q - f\right)^2 p(f)\,df$$

  - The distortion is measured as the expectation of the mean square error (MSE) difference between the original and quantized signals.
- Lack of correlation with perceived image quality
  - Even though this is a very natural way for the quantification of the quantization artifacts, it is not representative of the *visual annoyance* due to the majority of common artifacts.

- Visual models are used to define perception-based image quality assessment metrics

Example

- The PSNR does not allow to distinguish among different types of distortions leading to the same RMS error between images

- The MSE between images (b) and (c) is the same, so it is the PSNR. However, the visual annoyance of the artifacts is different