Image Processing: An Overview

Sebastiano Battiato, Ph.D.

battiato@dmi.unict.it

Program

- **Image Representation & Color Spaces**
- **Image files format (Compressed/Not compressed)**
- **Bayer Pattern & Color Interpolation**
- **Image & Video Compression**
- **Standard JPEG and JPEG2000**
- **Image Filtering**
- **Digital Enhancement, Restoration**
- **Edge Detection**

Software: Matlab, PSP 7.0.

Digital Images

An Image is a 2D function *f(x.y)* which represent a measure of some characteristics (brightness, colors) of a viewed scene.

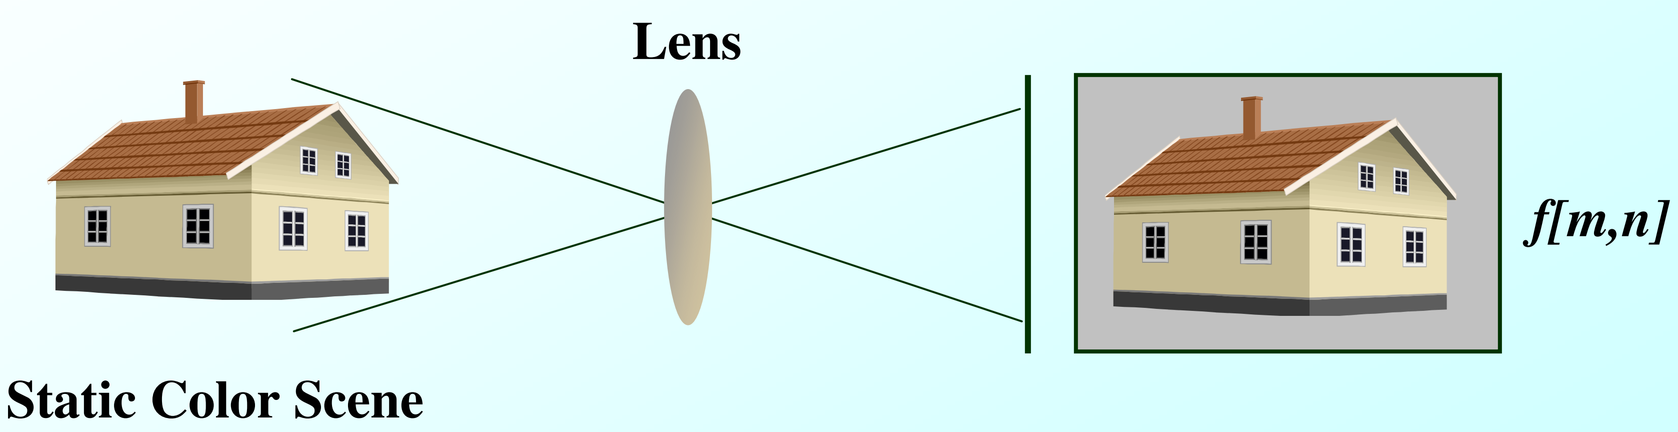

Digital images

A digital image ***f*[*m,n*]** described in a 2D discrete space is derived from an analog image ***f*(*x,y*)** in a 2D continuous space through a *sampling* process that is frequently referred to as digitization. The 2D continuous image ***f*(*x,y*)** is divided into ***N*** *rows* and ***M*** *columns*. The intersection of a row and a column is termed a ***pixel***.

The value assigned to the integer coordinates [m,n] with {m=0,1,2,...,M-1} and {n=0,1,2,...,N-1} is ***f*[*m,n*]**. In fact, in most cases ***f*(*x,y*)** --which we might consider to be the physical signal that impinges on the face of a 2D sensor--is actually a function of many variables including depth (z), color (λ), and time (t).

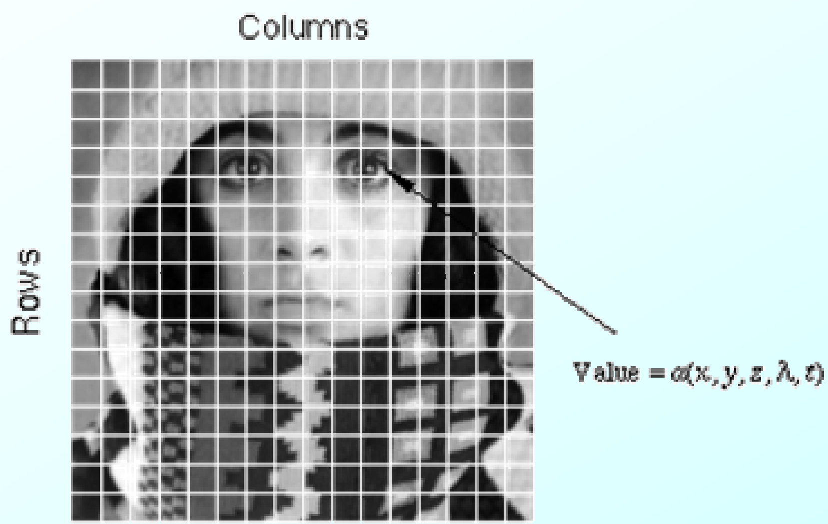

The value assigned to every pixel is the average brightness in the pixel rounded to the nearest integer value. The process of representing the amplitude of the 2D signal at a given coordinate as an integer value with *L* different gray levels is usually referred to as *amplitude quantization* **or simply *quantization*.**

Spatial Resolution

Generally speaking, spatial resolution refers to the number of specific points of picture information that are included in an image file. These points are referred to as picture elements or pixels. In a high resolution image the picture would have to be magnified to see individual pixels. In lower resolution images it may be possible to see individual pixels with the naked eye. The higher the spatial resolution, the greater the number of picture elements in the picture and correspondingly, the smaller the individual pixels will be. As a result, images having higher spatial resolution will have greater picture detail.

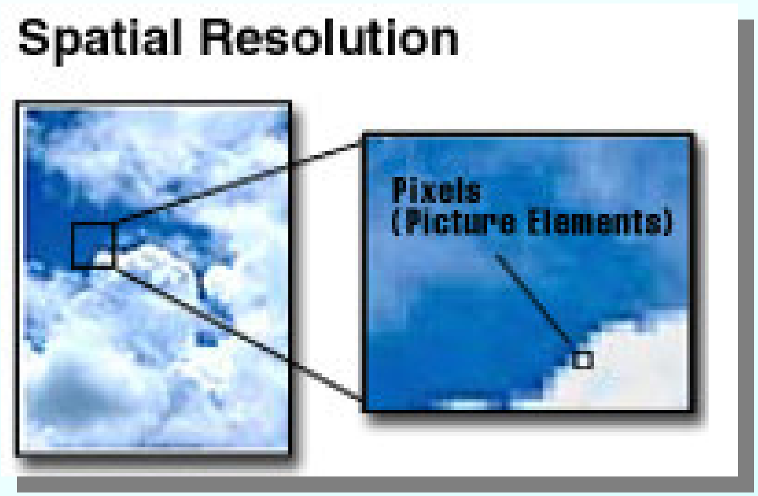

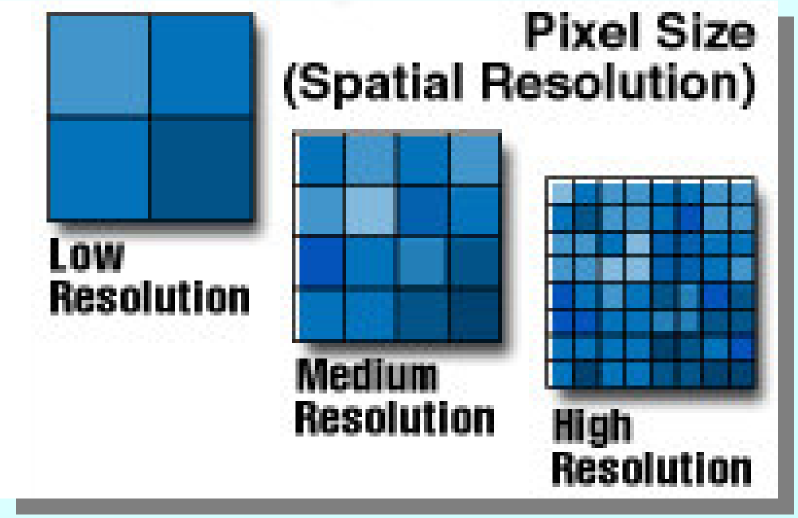

Digital pictures with higher resolution will have larger file sizes and be perceived as having better visual quality

Spatial Resolution

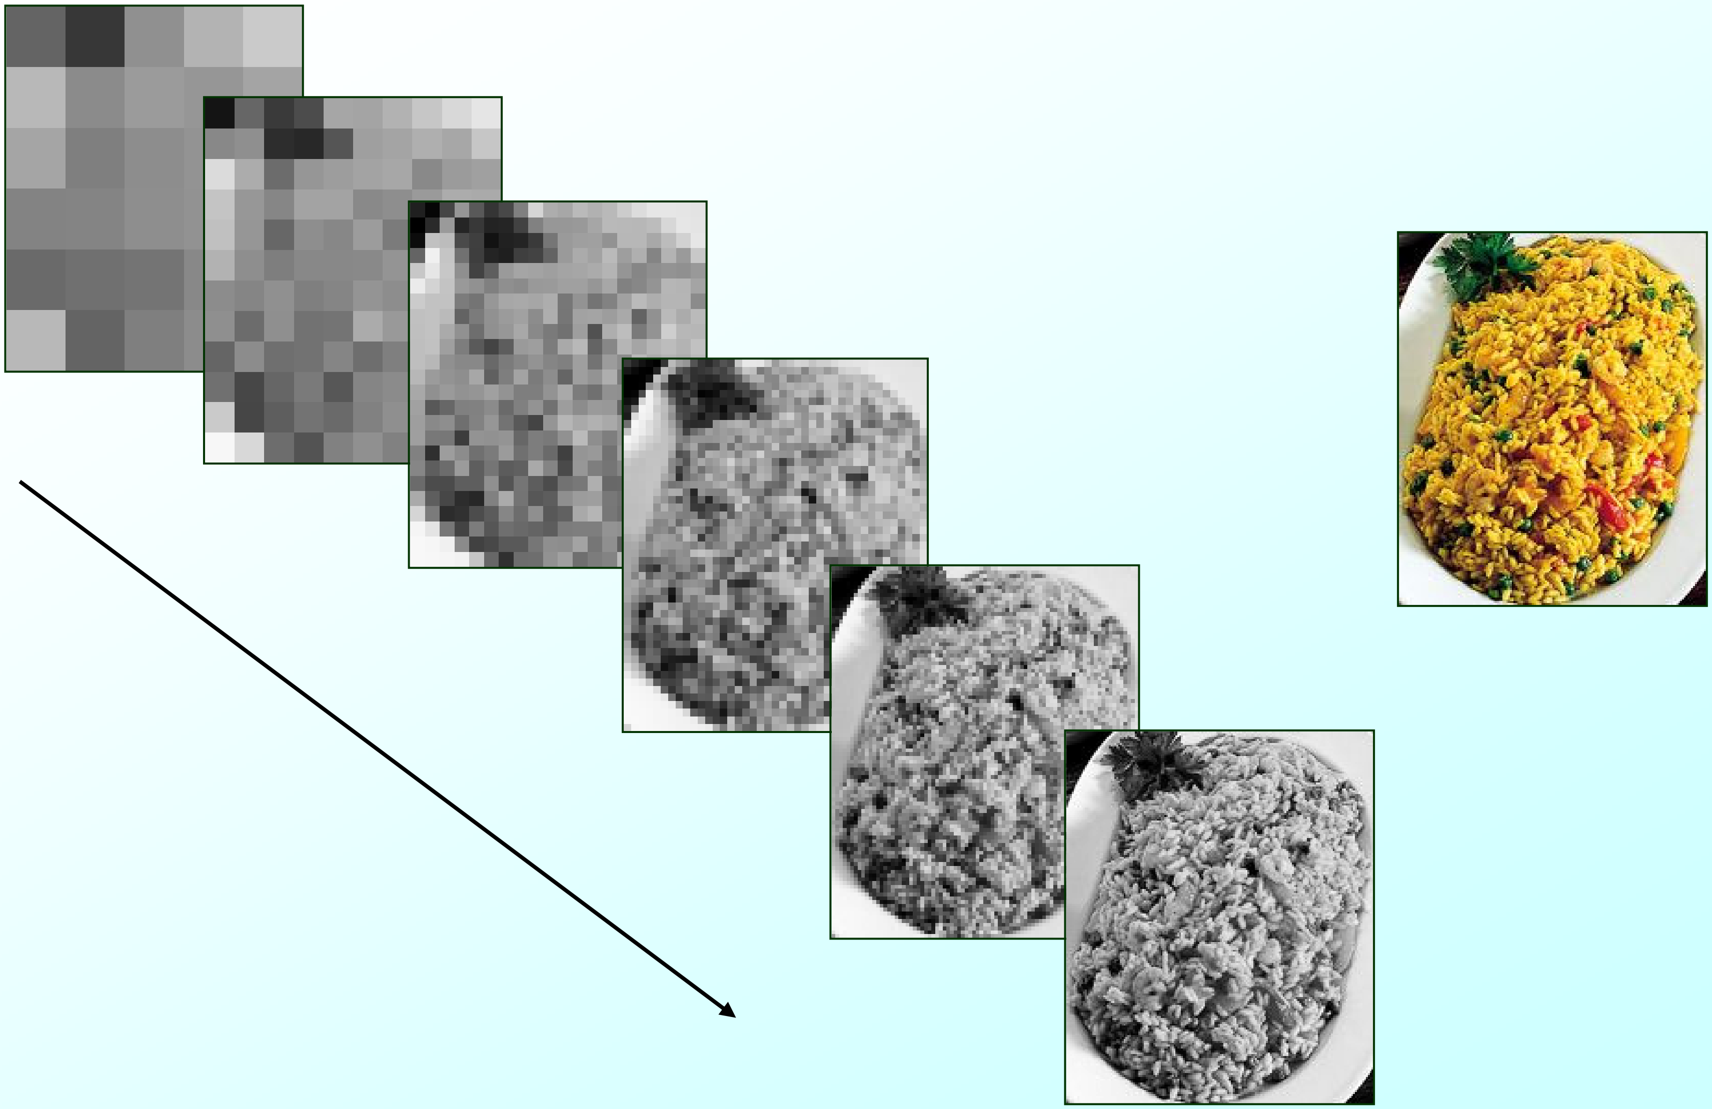

Human Vision

Light is the electromagnetic radiation to which our vision system is sensible. It is expressed as spectral energy distribution $L(\lambda)$ where the wavelength λ lies between 350nm and 780nm.

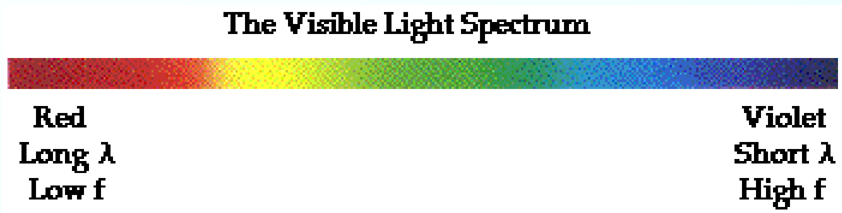

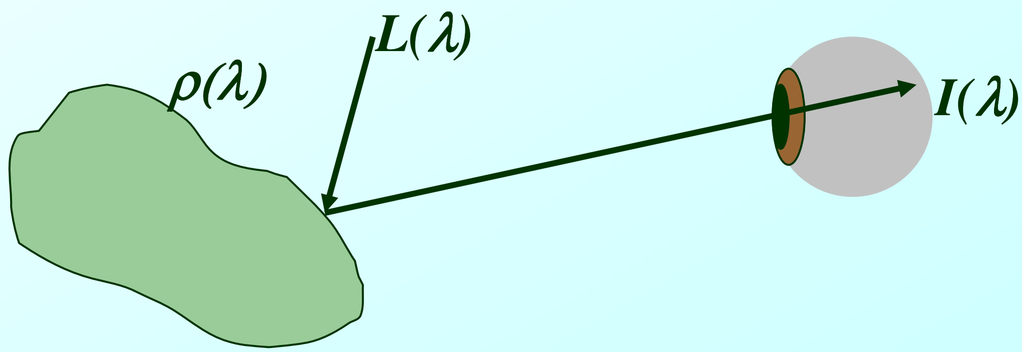

By considering the reflectivity $\rho(\lambda)$, the light $I(\lambda)$ which the eye receives from an object can be written as $I(\lambda) = \rho(\lambda)\, L(\lambda)$

IQ and HVS

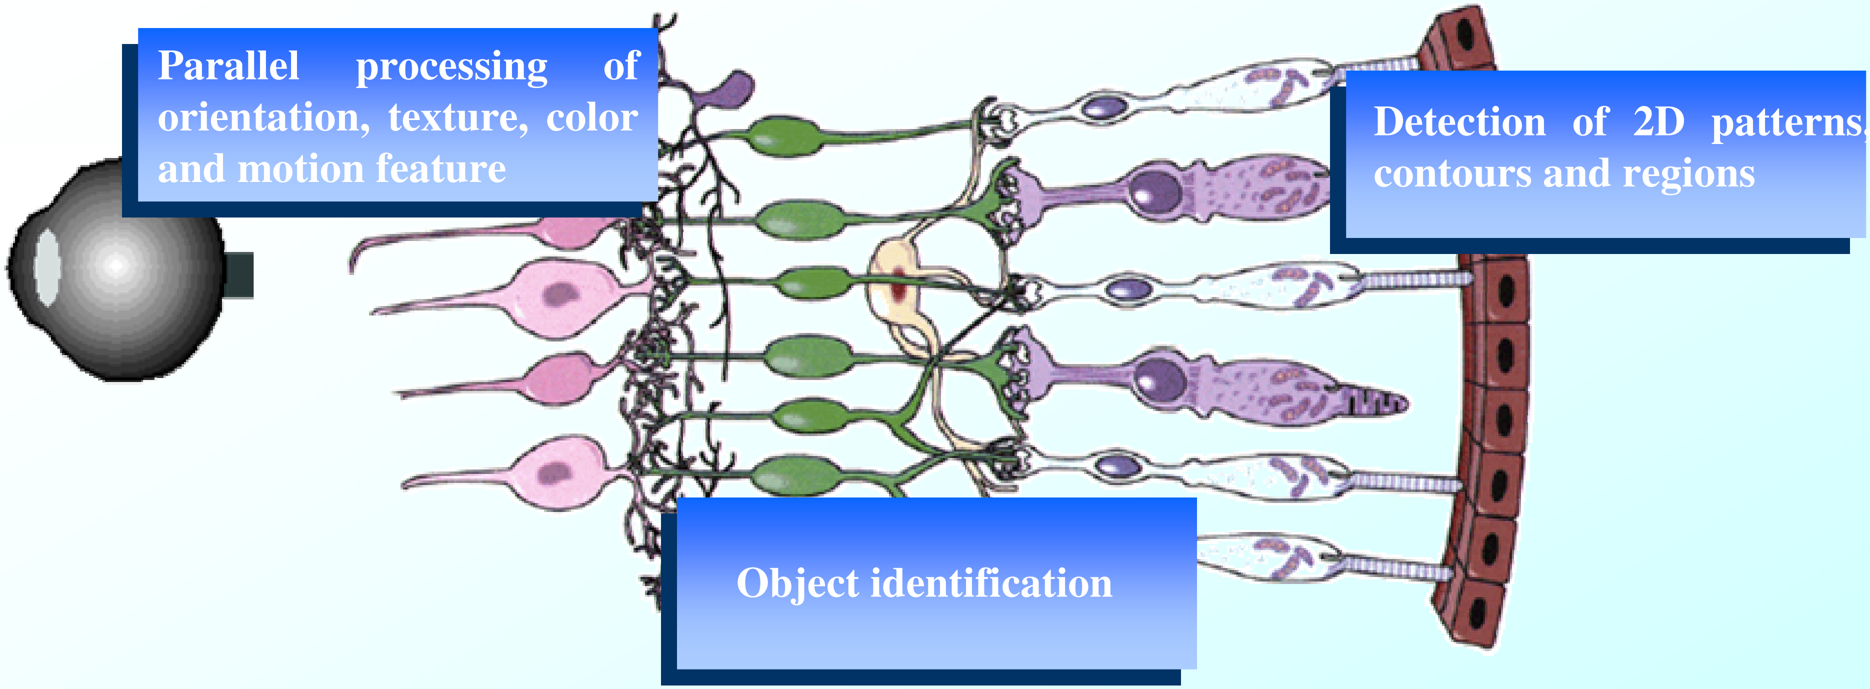

Color processing → Spatial decomposition → Local Contrast → CSF compensation → Masking effects → Pooling →

General structure of computational model of the human visual system

Contrast Sensitivity
Weber-Fecner Law

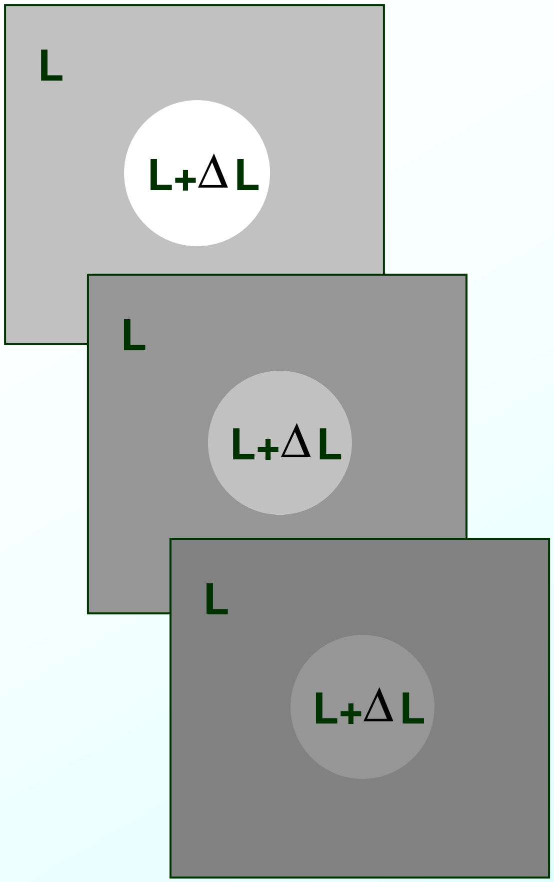

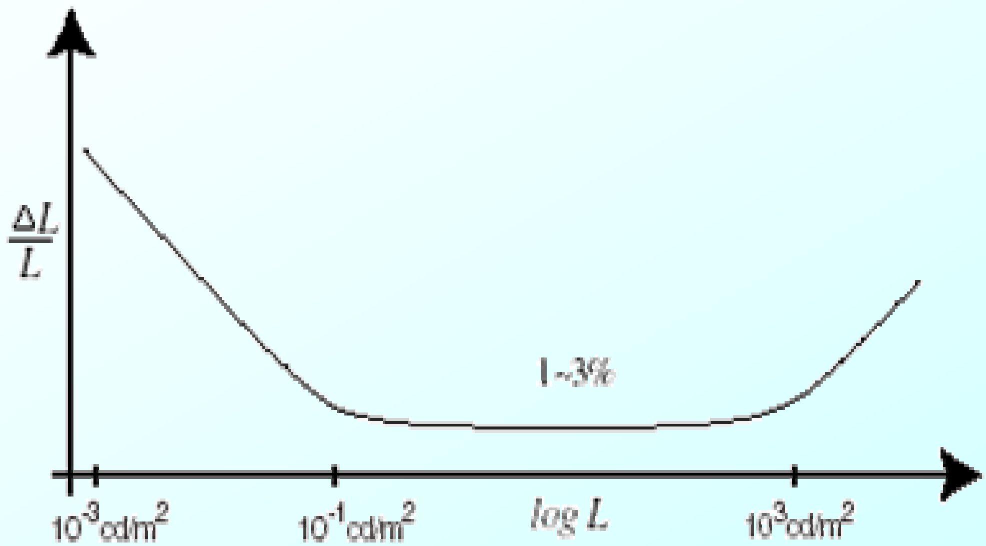

A JND (Just Noticeble Difference) of 1-3% is just sufficient to be perceived, if the background illumination is in the range between 0.1 and 1000 cd/m^2

Lightness Perception

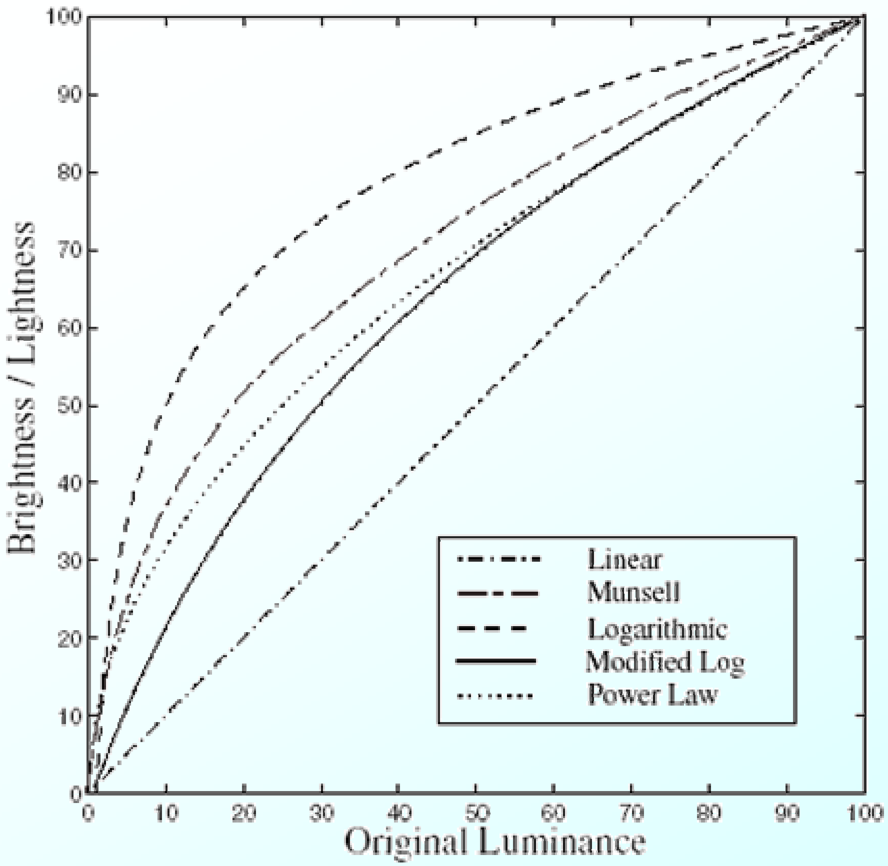

- ✓The perception of lightness is not linearly related to the real luminance of an object
- ✓A logarithmic behavior overestimates the sensitivity
- ✓A linear behavior underestimates the sensitivity
- ✓Power Low Function

$$B = a_l L^{P_l} - B_0$$

Opponent Color Perception

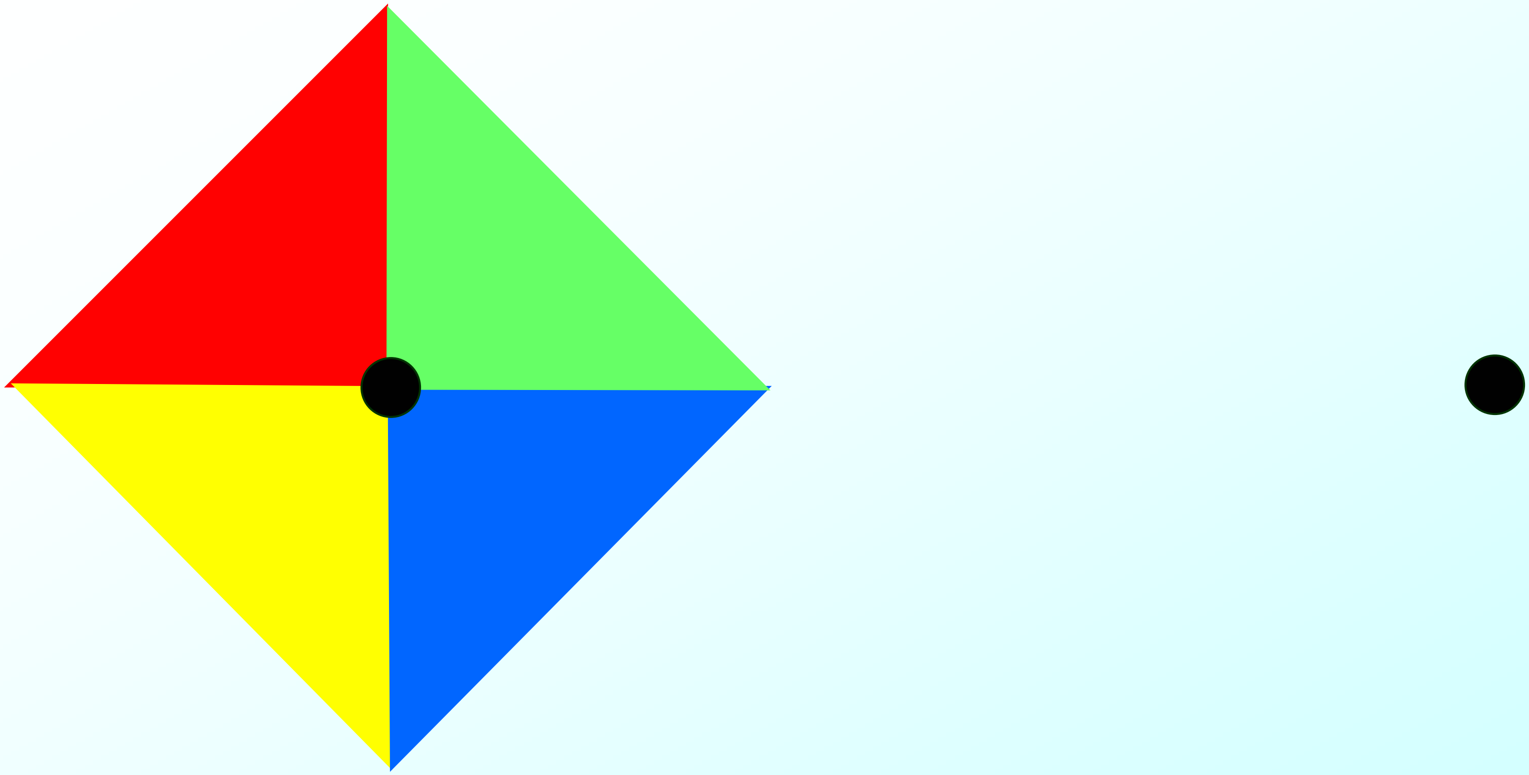

E.H. Hering introduced in 1878 the idea of *Opponent colors*

Color Encoding

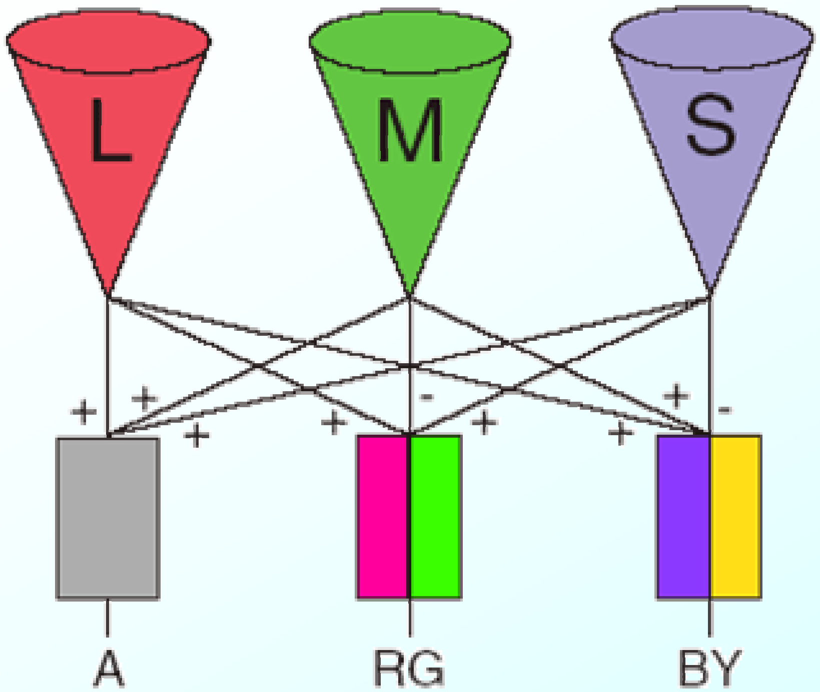

L = long wavelength cone

M = medium wavelength cone

S = short wavelength cone

Photo-receptors combined color information in one achromatic and two chromatic channels

Colors

Color is not a physical property of objects, it is a perceptual representation of the distribution of photon energy quanta within a reflectance or emission spectrum produced by an object.

Main approaches

Additive Color. Defined in terms of visible light. The additive primaries are **red**, **blue**, and **green**. As these color mix, they produce all colors in the visible spectrum.

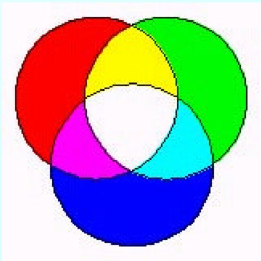

Subtractive Color. Defined in terms pigments (paints, etc.). Light waves hit an object, whose molecular structure causes it to absorbe some light wavelengths and reflect others. We see what is reflected (e.g. the remainder after some wavelengths have been subtracted)

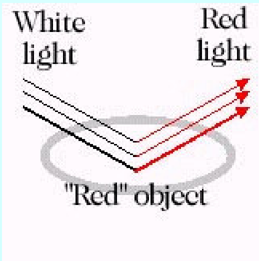

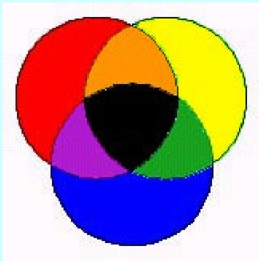

Intensity, Brightness, Luminance

Intensity is the measure over a given wavelength interval of the power of incident light.

Brightness is defined as the attribute of a visual sensation (subjective) according to which an area appears to emit more or less light. Very difficult to be measured.

Luminance is defined as the incident light power weighted by a spectral sensitivity function. Such a function is called *luminous efficiency* and is defined numerically for the standard observer.

$$Y(x, y) = \int_0^\infty V(\lambda) I(x, y, \lambda) d\lambda$$

Lightness $$L^* = 116\left(\frac{Y}{Y_0}\right)^{\frac{1}{3}} - 16$$

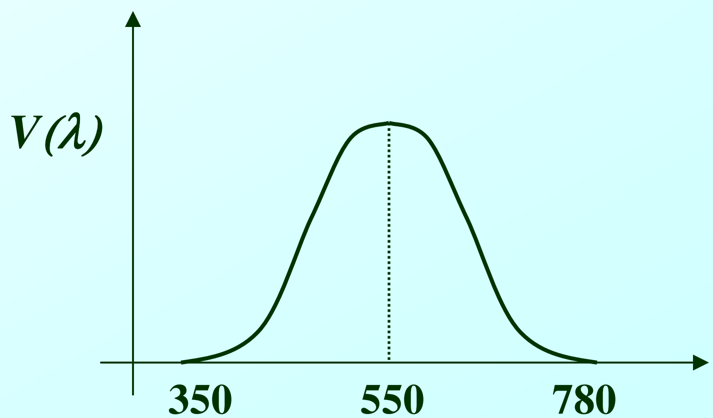

RGB

It is most common to describe color in terms of an RGB (**red**, **green**, and **blue**) color spaces. The RGB color space is based on the fact that any color can be represented by mixing percentages of the primary colors **red**, **green**, and **blue**.

RGB is used by camera and monitor suppliers since it is the easiest colorspaces to work with recording and displaying color images.

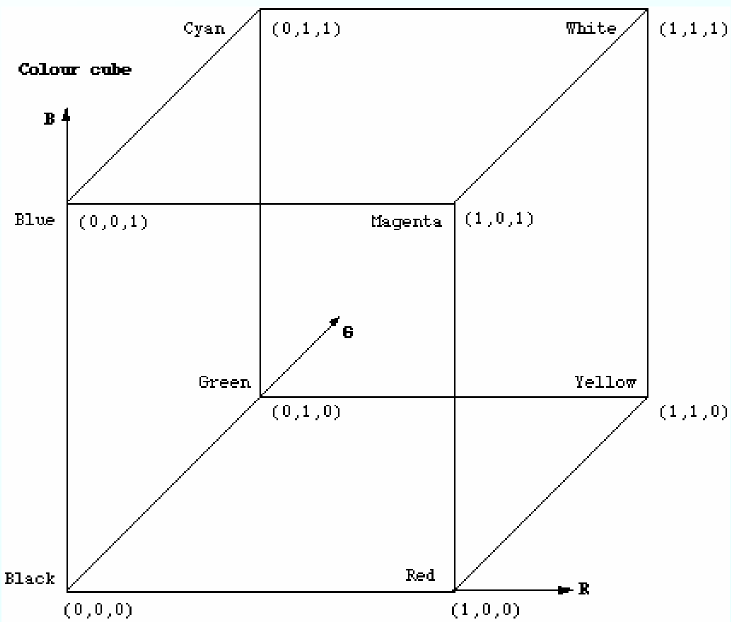

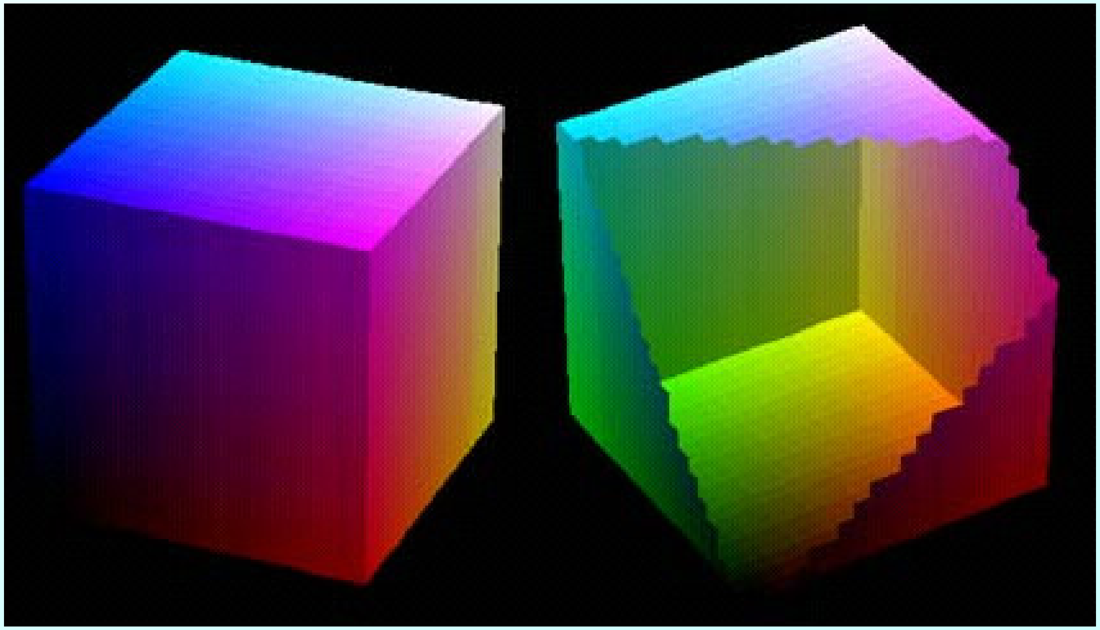

RGB Color Display

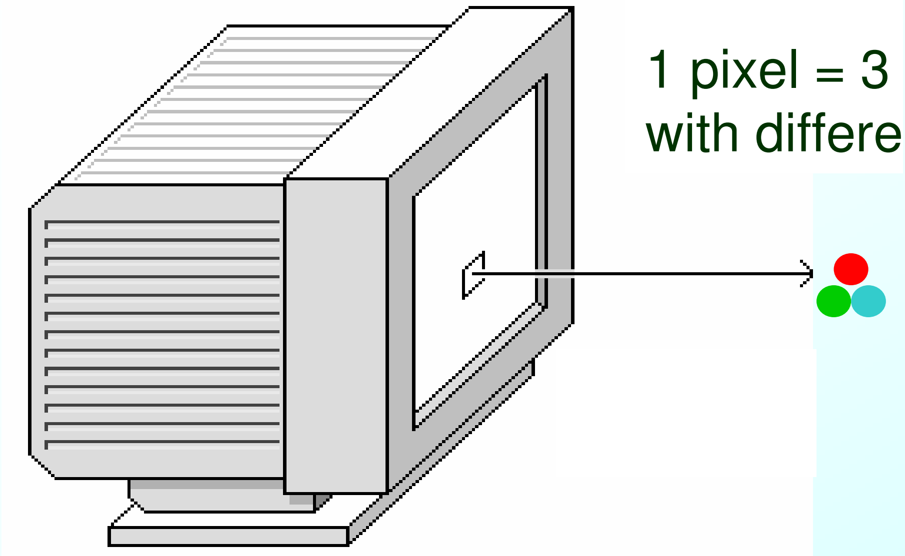

1 pixel = 3 R,G,B, phospors
with different illumination

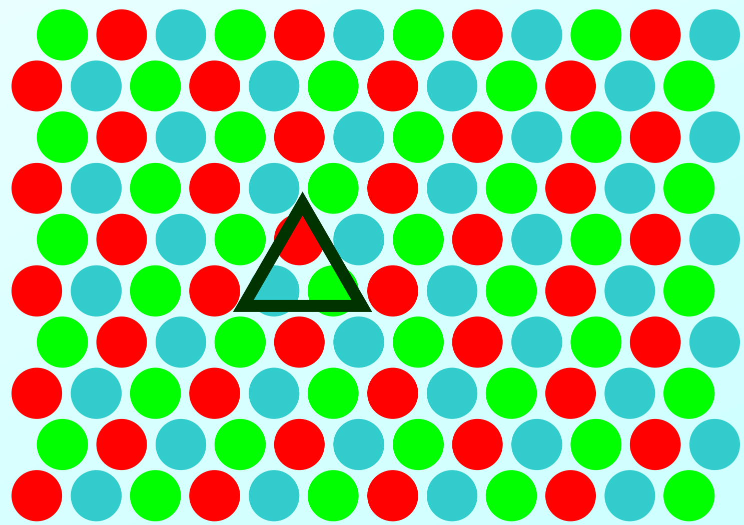

RGB Color Display
(24 bit: "TrueColor")

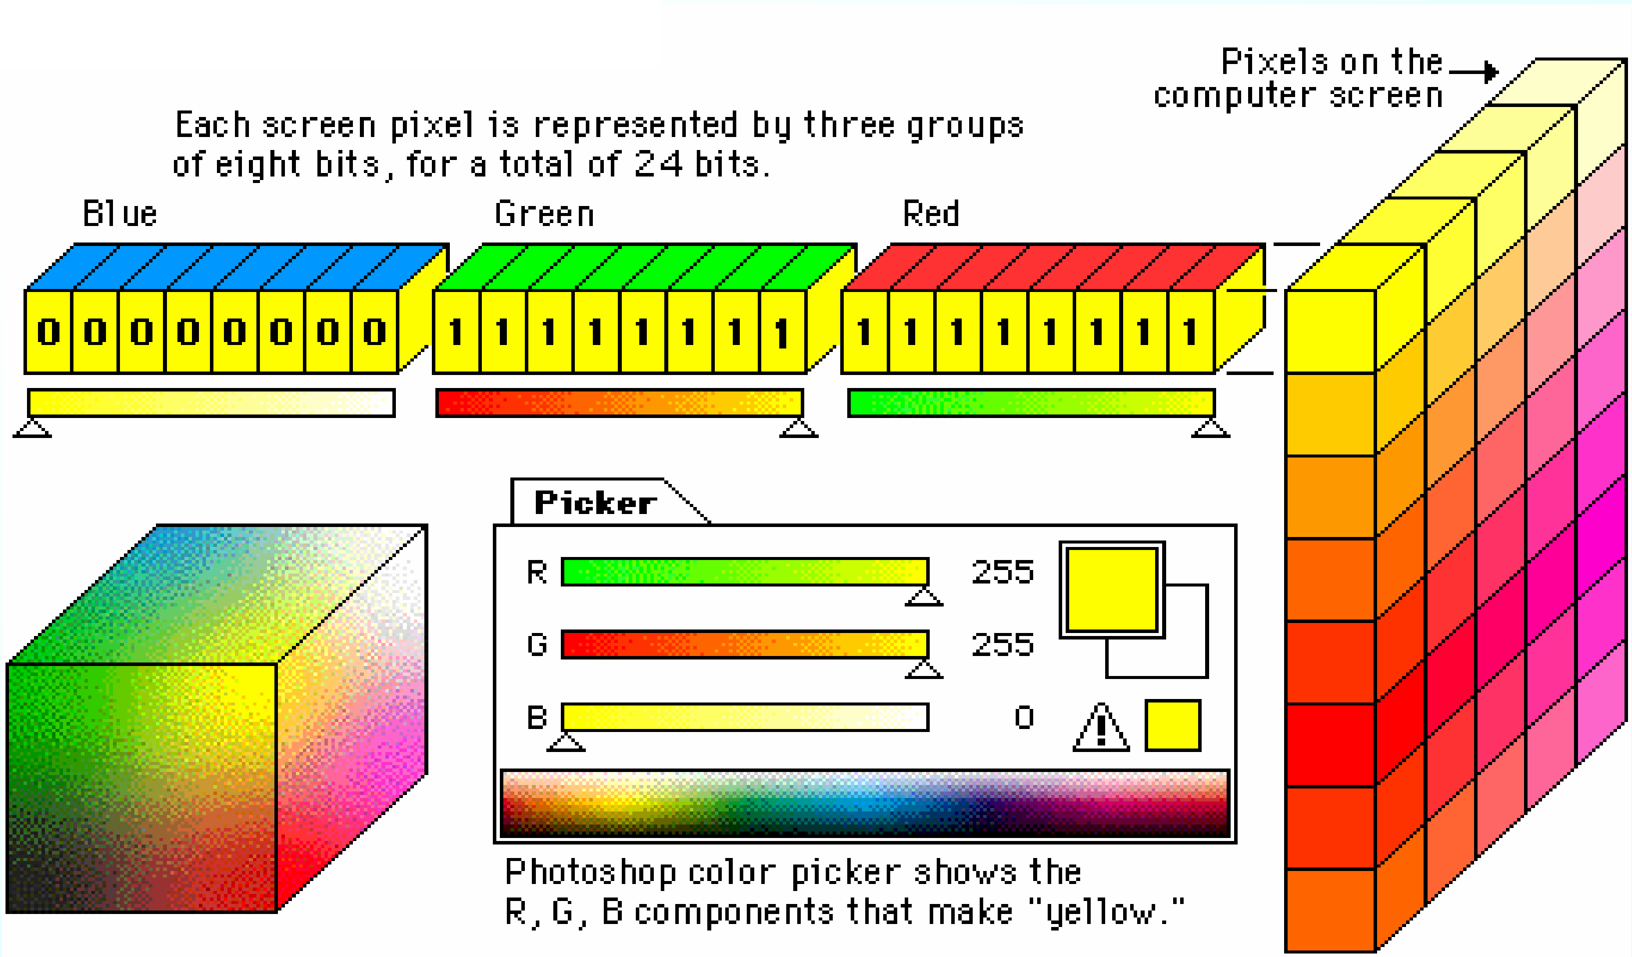

TrueColor (24 bit)

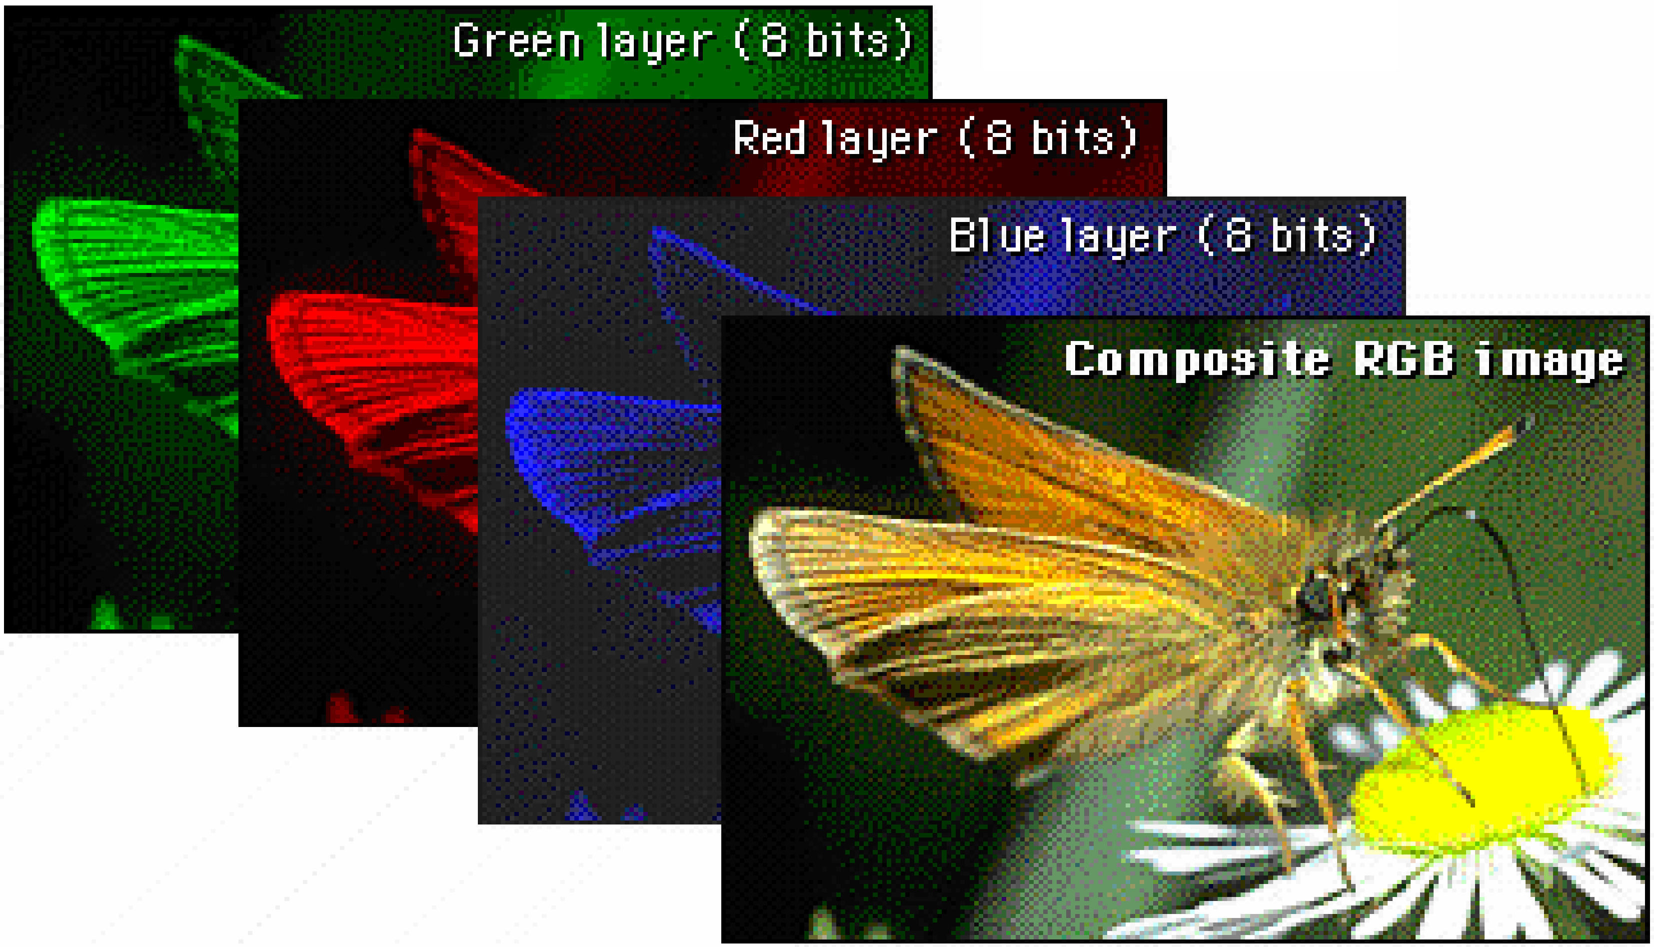

Depth Resolution

8 bit – 256 Gray levels

4 bit – 16 colors

…

24 bit True colors

Brightness resolution refers to the number of brightness levels that can be recorded in any given pixel. The greater the brightness resolution, the greater the number of levels that can be included in the picture file. In black and white images, levels are seen as shades of gray. In color images, levels will be seen as specific color hues. As a general rule, brightness resolution for black and white picture information should be at least 8bits. Brightness resolution for full color picture information should be at least 24 bits.

How many colors?

Bit	Numero di colori	
1	2^{1}	2
2	2^{2}	4
3	2^{3}	8
4	2^{4}	16
5	2^{5}	32
6	2^{6}	64
7	2^{7}	128
8	2^{8}	256
16	2^{16}	65.536 (16 bit True Color)
24	2^{24}	16.777.216 (True Color)
32	2^{32}	24 bit True-Color + 8 bit Alpha Channel

256 Color Display (8 bit)

“Palette”

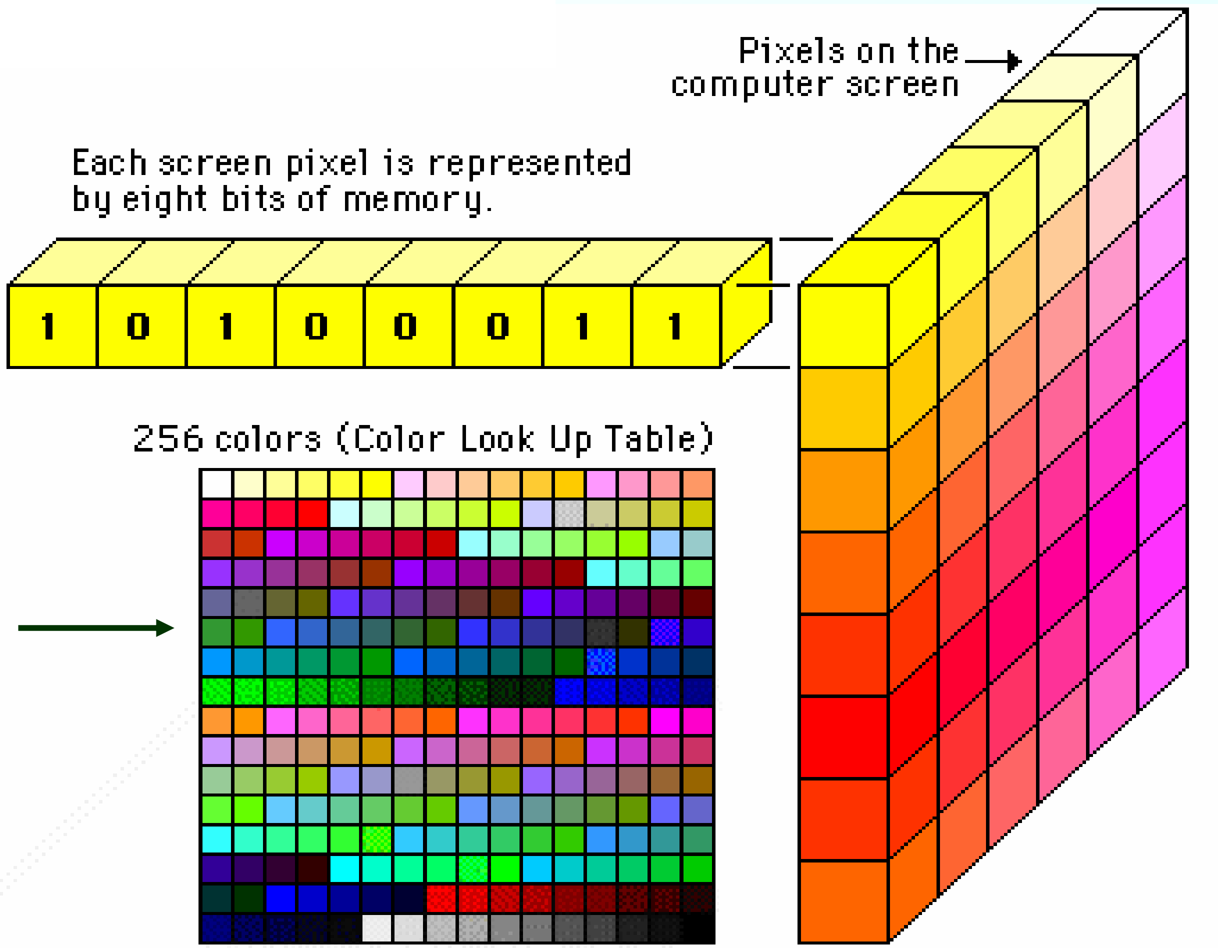

EXAMPLE 256 color

Dithering

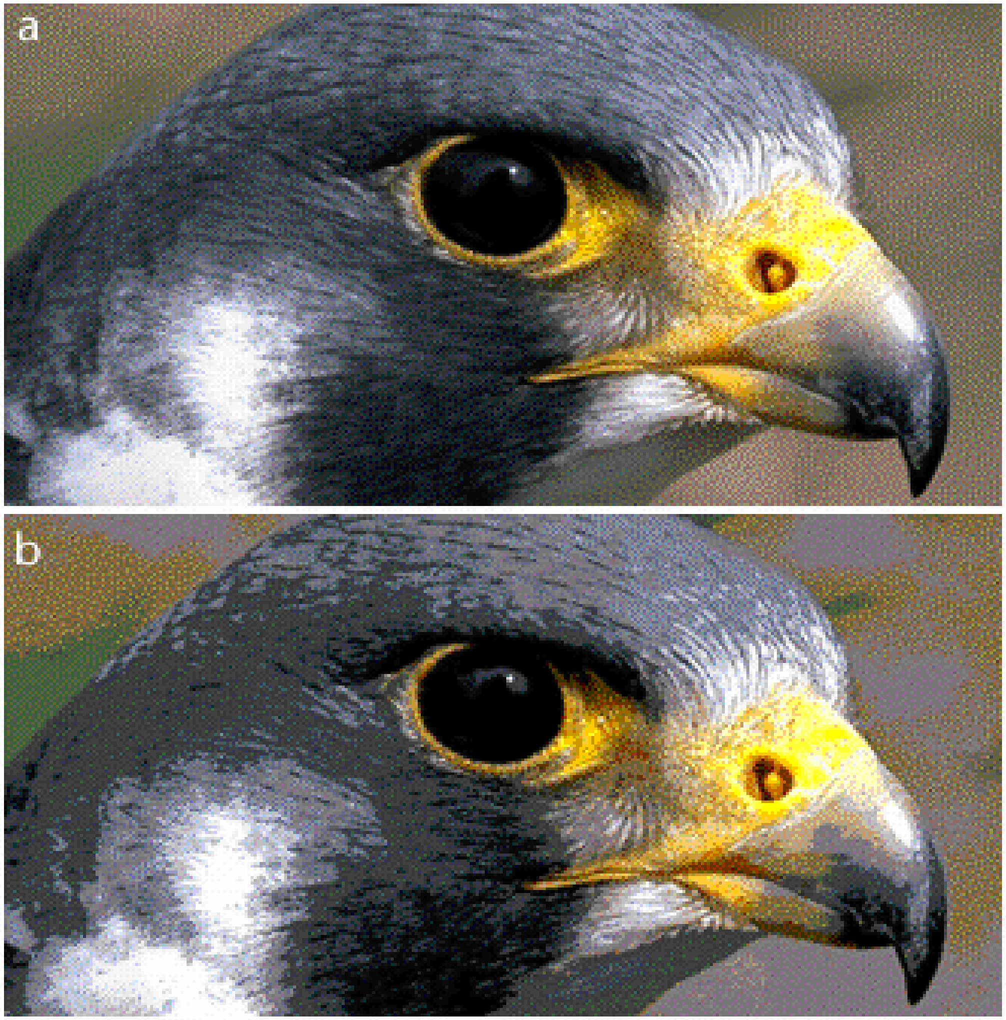

Original (TrueColor)

Dithered (256 colori)

Palettes

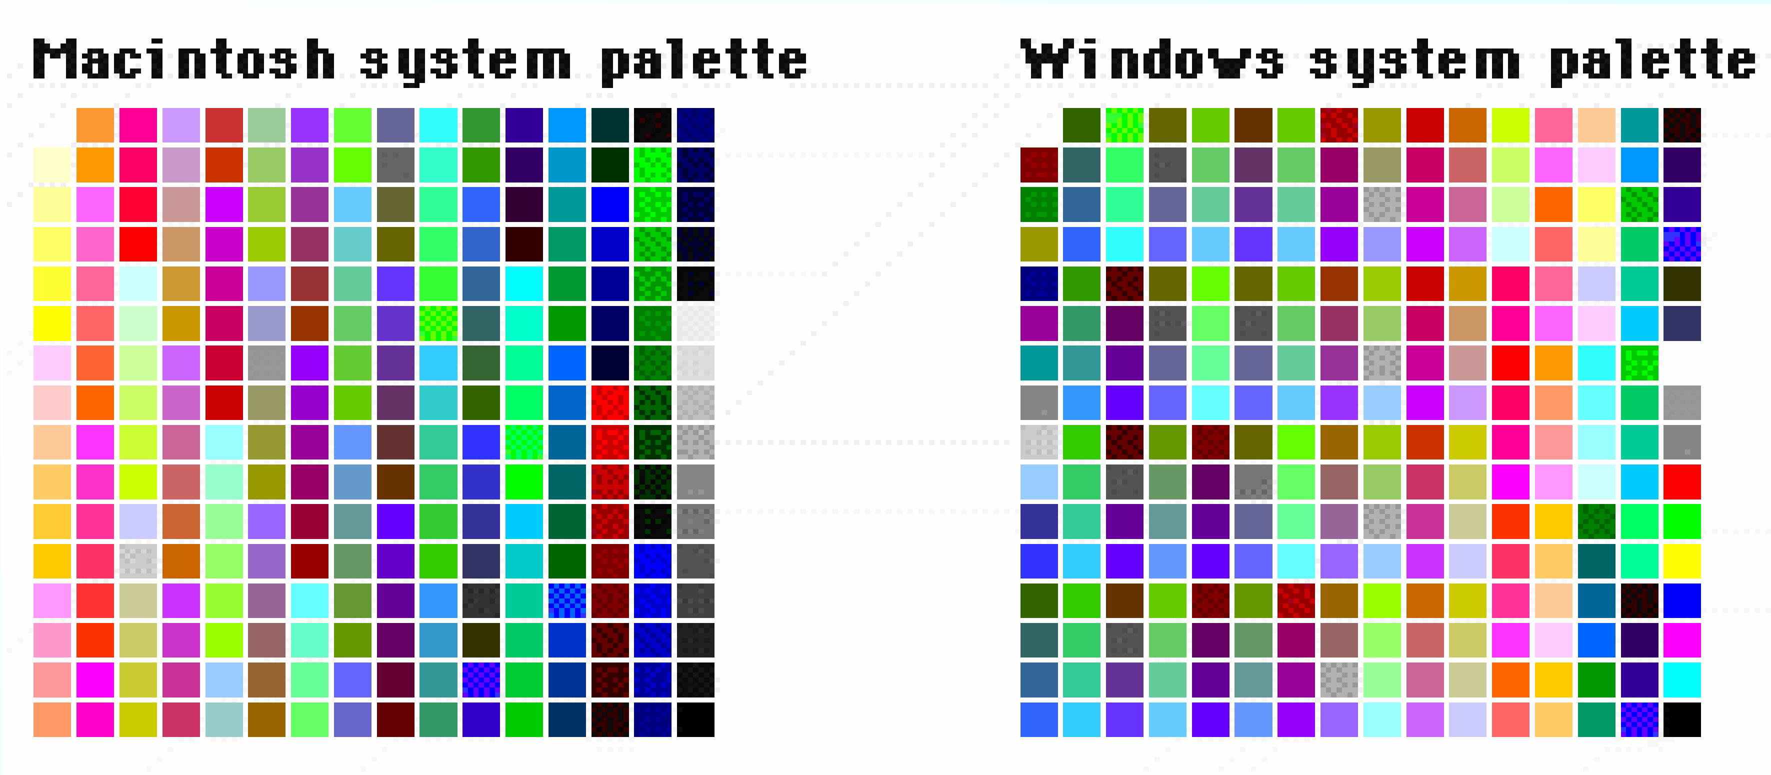

Hue, Saturation, Lightness

[**HSL (HSI)**]() is tried to better describe the way the eye perceives colors:

Hue is the attribute of a visual sensation according to which an area appears to be similar to a perceived colors.

Saturation is the colorfulness of an area judged in proportion to its depth or pureness of a given hue.

Lightness (Intensity) is the non linear perceptual response of the eye to luminance.

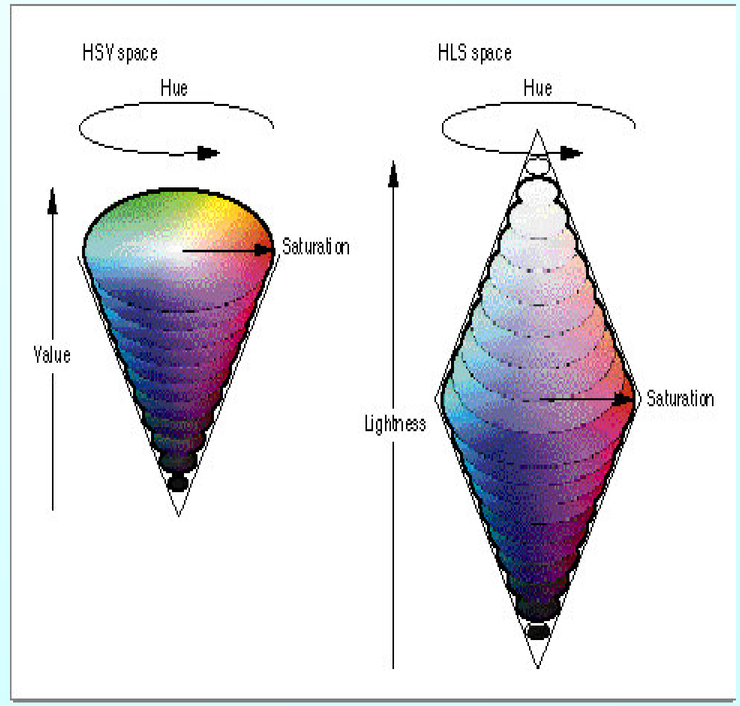

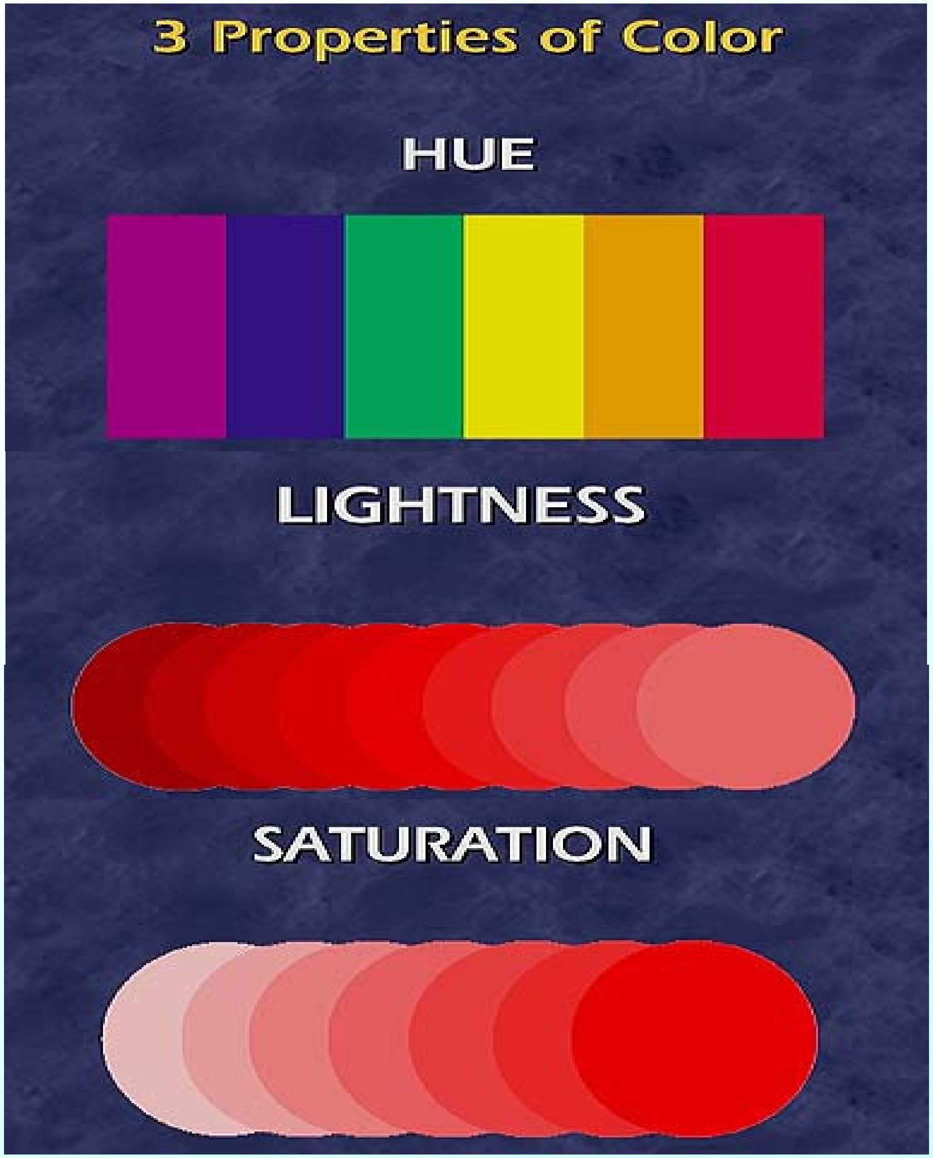
3 Properties of Color
HUE
LIGHTNESS
SATURATION

HUE

Commonly said COLOR

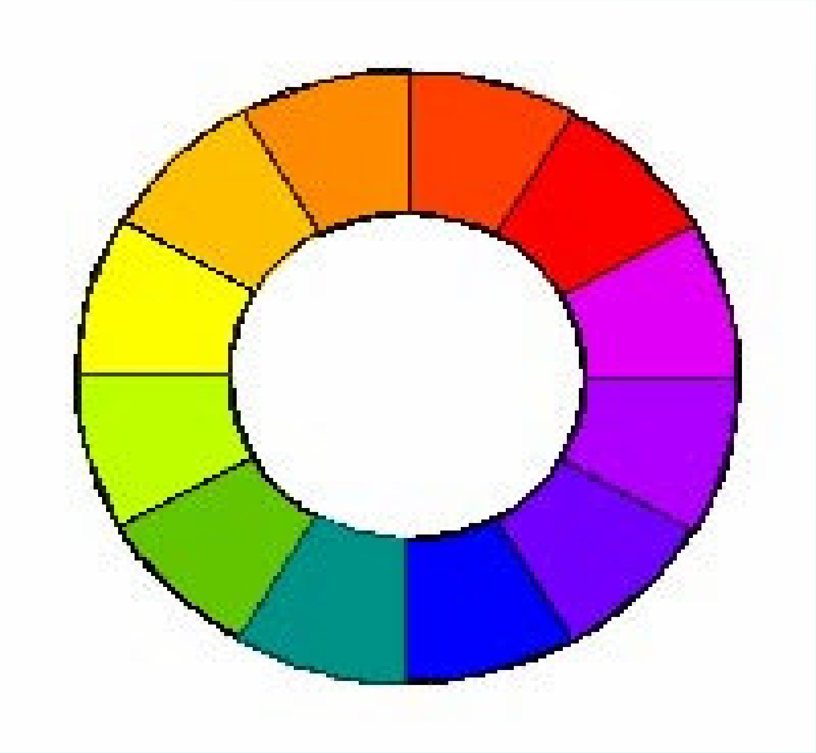

Saturation & Lightness

HUE + WHITE to obtain a perceived color

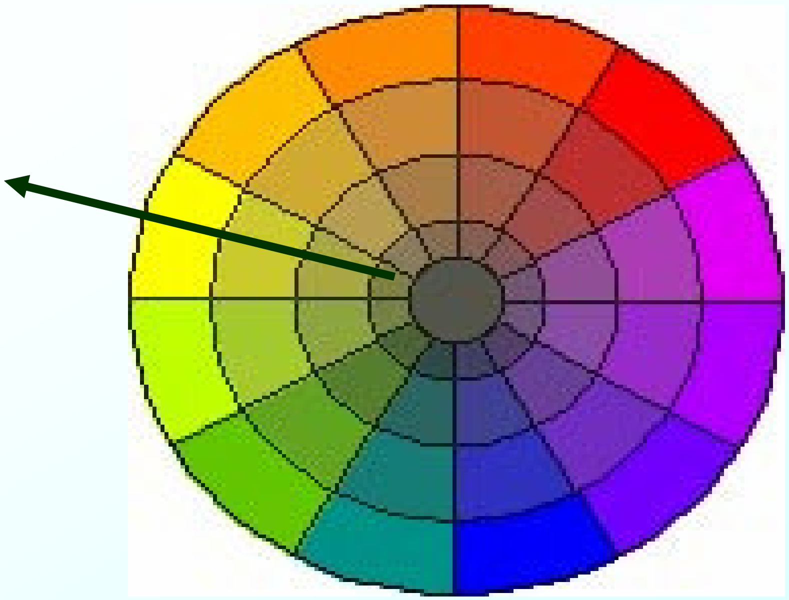

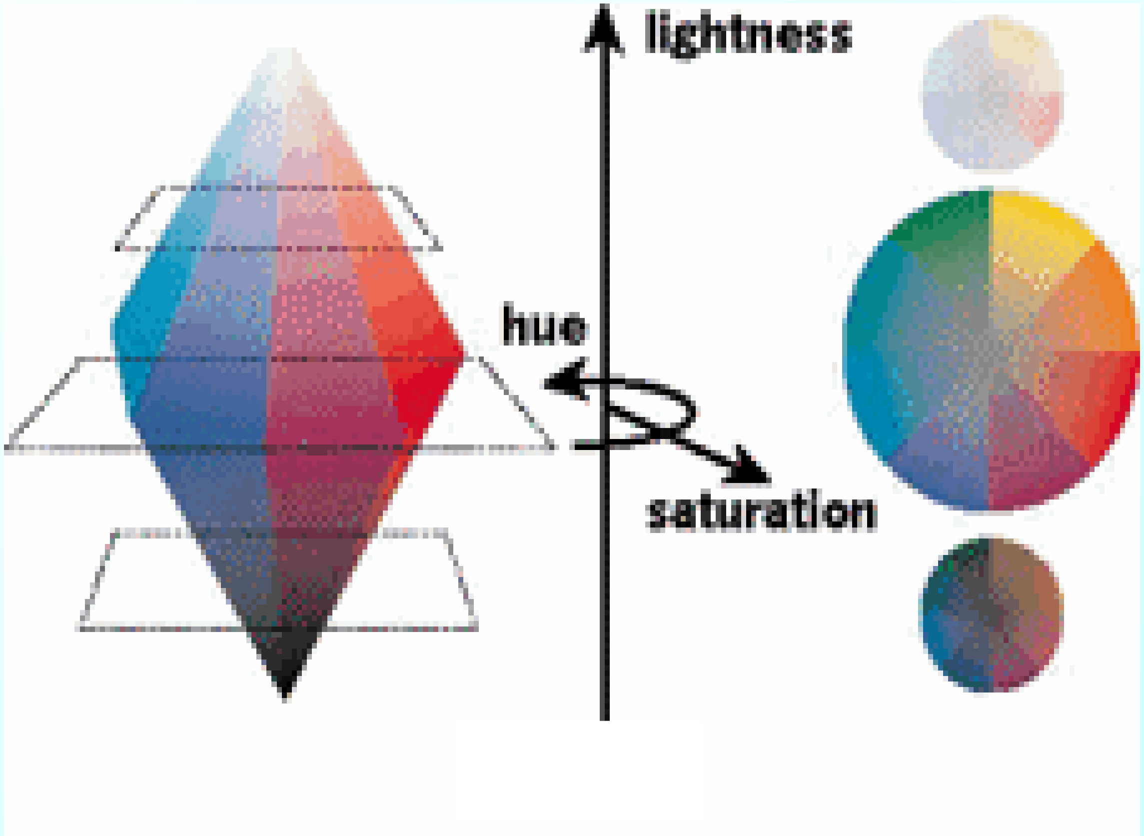

An Examples

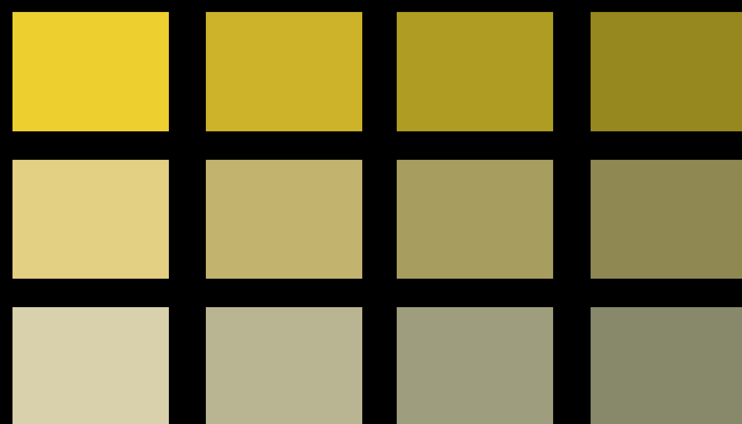

Same HUE,Decreasing Saturation

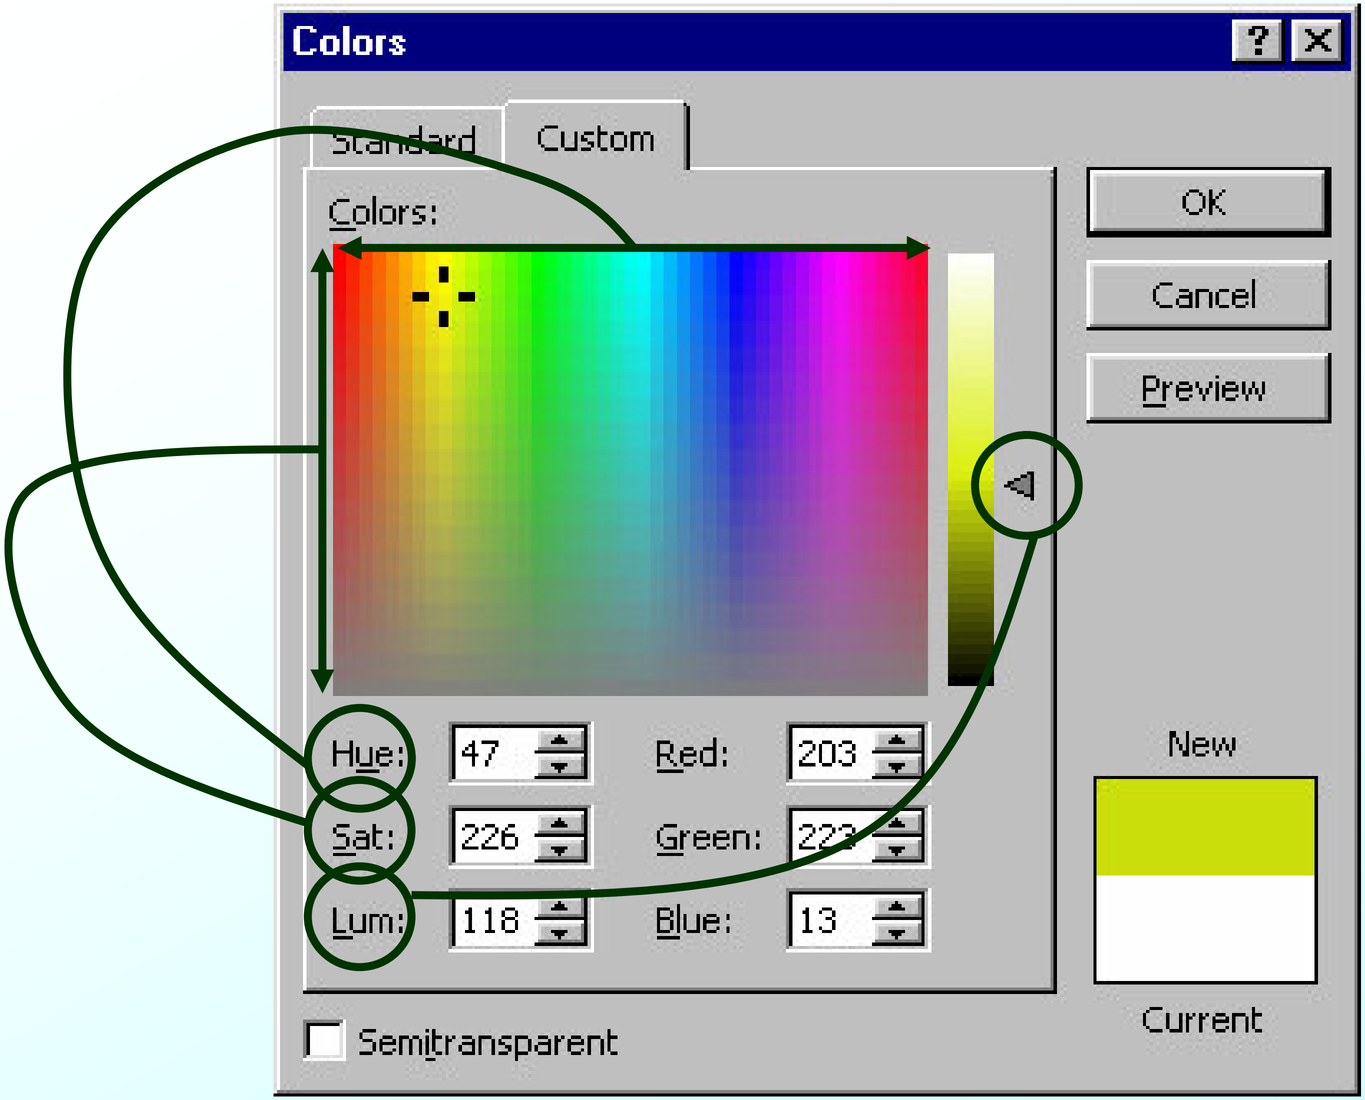
Colors
Standard
Custom
Colors:
OK
Cancel
Preview
Hue:
47
Sat:
226
Lum:
118
Red:
203
Green:
223
Blue:
13
New
Current
Semitransparent

Some notes

HSL (HSI)

Where simple operations like changing the color in an image require three separate operation on RGB channels, in the HSL (HSI) color space this is accomplished by alteration of one value *Hue*.

All traditional image processing techniques, including edge detection, filtering and histograms, require processing of just the intensity values.

RGB

The diagonal line of the cube form black (0,0,0) to white (1,1,1) represents all the greys that is, all the **red**, **green**, **blue** components are the same. In practice are used different ranges for the colours, common ones are 0-256 (8 bit – 2^8) and 0-65536 (16 bit – 2^{16}) for each component.

What Color Space? (1/2)

- **RGB**

  Used by CRT displays where proportions of excitatation of red, green and ble emitting phospors produce colorurs when visually fused. Easy to implement, non linear, device dependent, unintuitive, very common.

- **CMY (Cyan, Magenta, Yellow)**

  Subtractive colour. Used in printing and photography. Printers often include the fourth component black ink, improving black and speeding drying. Device dependent, non linear, unintuitive.

- **HSL** (HSI intensity, HSV value, HCI chroma, etc. etc.)

  Device dependent, non linear but very intuitive. Useful luminance separation.

- **YIQ,YUV, YCbCr,YCC (Luminance- Chrominance)**

  They separate luminance from chrominance (lightness from colour) and are useful in compression and image processing applications. Device dependent and un-intuitive.

What Color Space? (2/2)

- **CIE**

  There are two CIE based colour spaces, CIELUV and CIELAB. They are near linear (as close as any colour space is expected to sensibly get), device independent (unless your in the habit of swapping your eye balls with aliens), but not very intuitive to use.

  From CIELUV you can derive CIELhs or CIELhc where h is the hue (an angle), s the saturation and c the chroma. This is more intuitive to work with when specifying colours. CIELUV also has an associated chromaticity diagram, a two dimensional chart which makes additive colour mixing very easy to visualise, hence CIELUV is widely used in additive colour applications, like television.

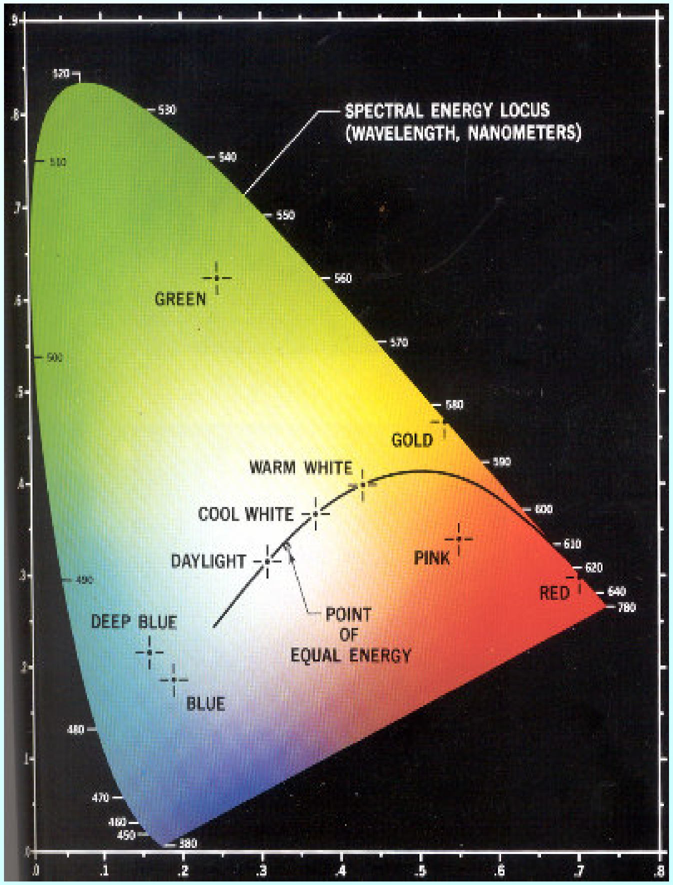

Color Transform

The human eye is more sensitive to luminance than to chrominance. Typically JPEG throw out $^3/_4$ of the chrominance information before any other compression takes place. This reduces the amount of information to be stored about the image by $^1/_2$. With all three components fully stored, 4 pixels needs 3 x 4 = 12 component values. If $^3/_4$ of two components are discarded we need 1 x 4 + 2 x 1 = 6 values.

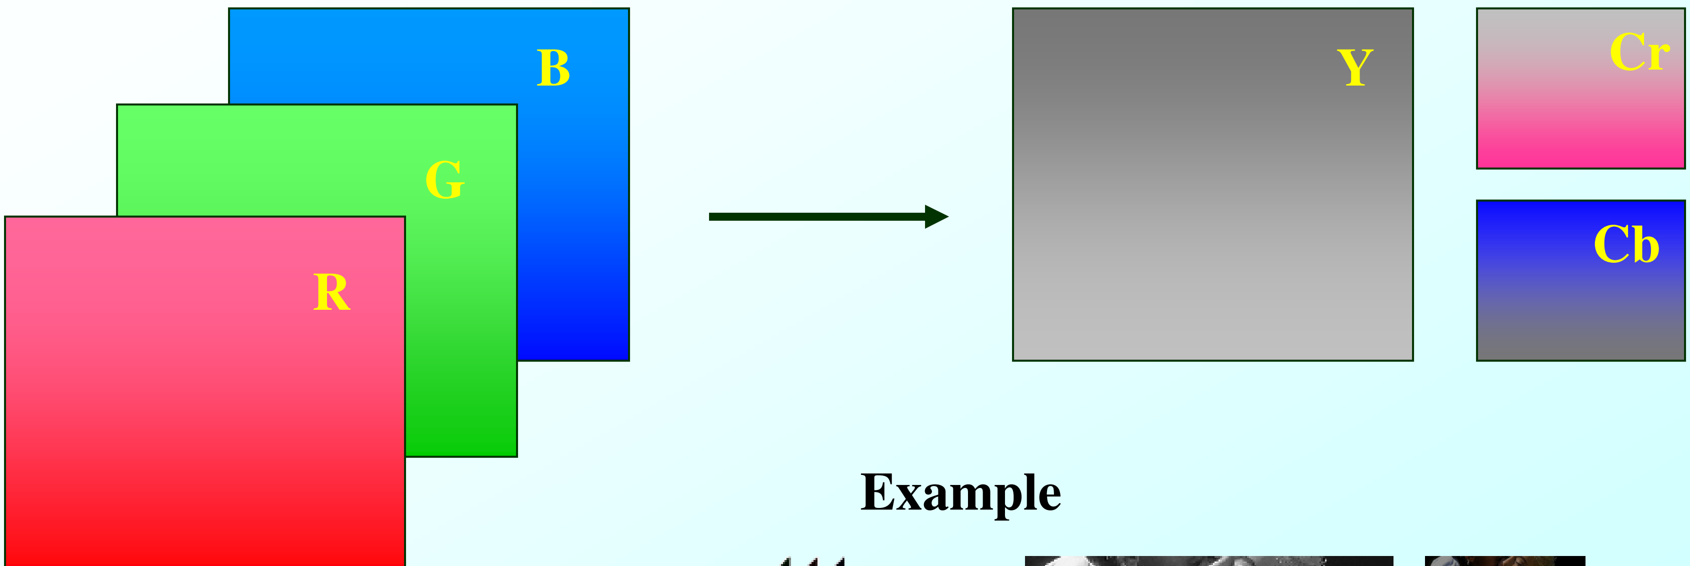

Example

RGB conversion & subsampling

RGB → CMY

$$\begin{bmatrix} C \\ M \\ Y \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

RGB → YCbCr

$$\begin{bmatrix} Y \\ C_b \\ C_r \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.523 & 0.311 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Subsampling

4:4:4 (no subsampling)

4:2:2 (Cb, Cr horizontal subsampling)

4:2:0 (Cb, Cr horizontal + vertical subsampling)

Colour Images

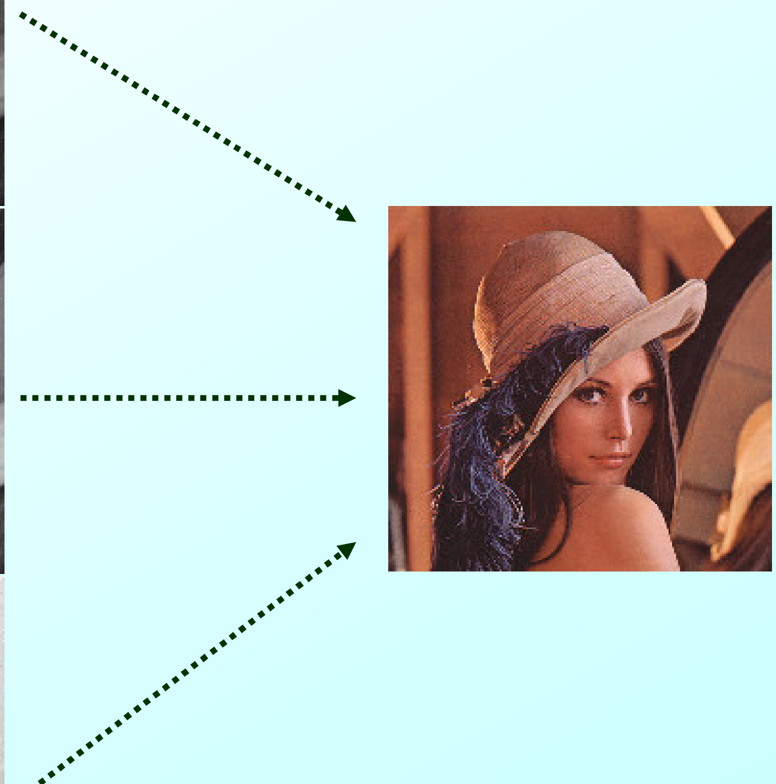

…..Demo in Psp

Zooming - Replication

An image can expanded by *zooming operation*:

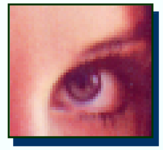

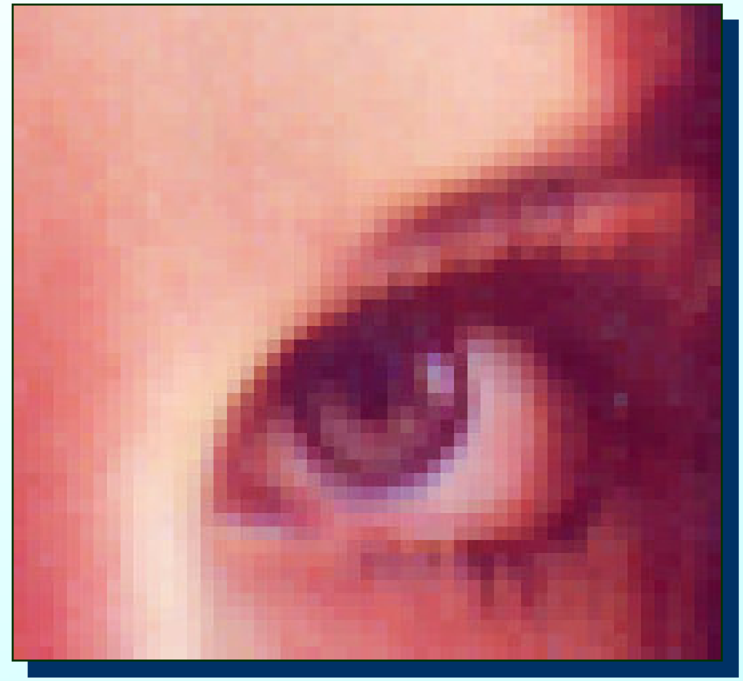

The simplest zooming operator "*Replication*" use the following steps:

- For every pixels of the image
  - Put the same value of the pixel in a grid of NxN pixels(N^2 replications of the value).

Ex: Zooming x2

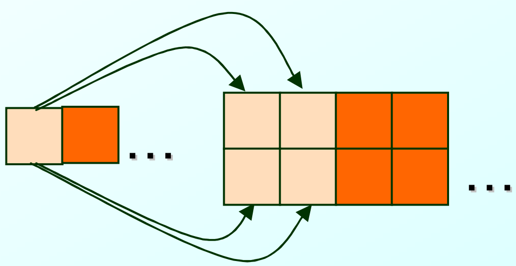

This method introduces the problem of ***Pixelization***.

Zooming - Bicubic

To remove Pixelization, we can use some intelligent and/or adaptive zooming operator.

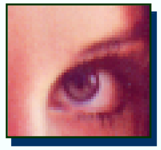

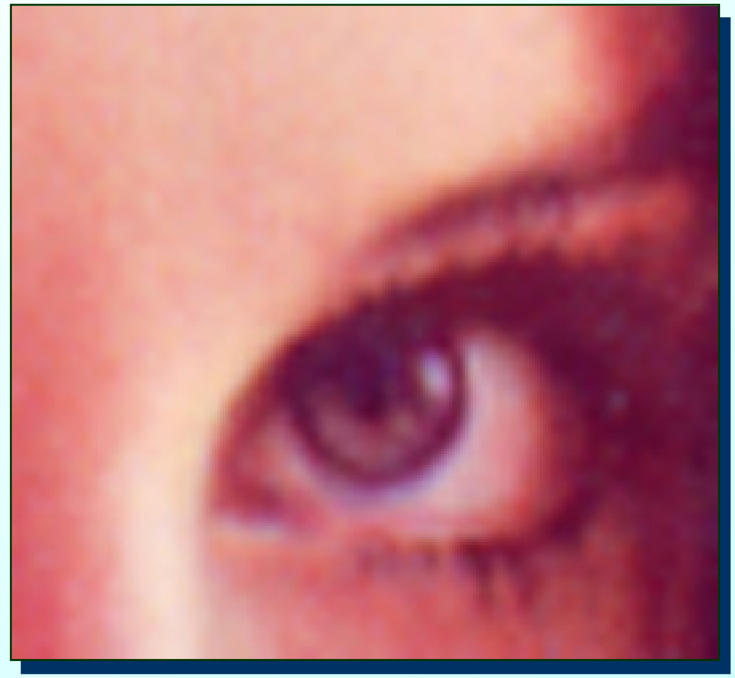

The most popular zooming method is "*Bicubic*":

- For every pixels of the image do
  - Convolve pixel value with a suitable mask to fit the interpolated information to the original information.

The ***Bicubic*** method is used for ***large enlargement*** of image size and for offline elaborations.

This method introduces the following problems:

- ***Blurring***;
- ***Computationally expensive***;

Zooming

Original

Bicubic

Zooming vs Super resolution

- Main objective of zooming methods is "***Interpolation of NEW information***".

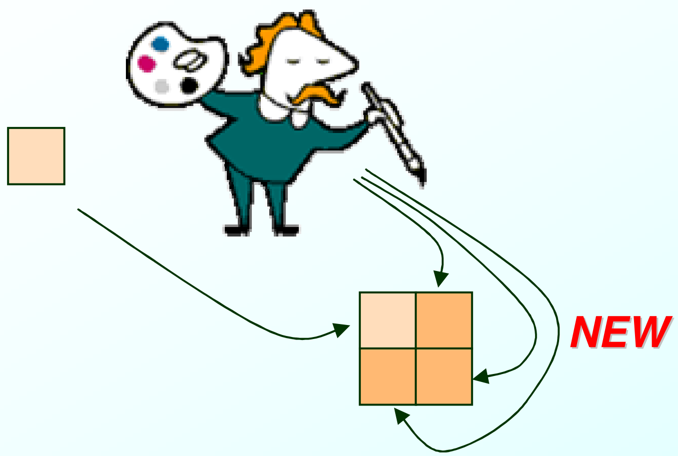

- ***Resolution Enhancement*** aims ***"To restore the REAL information"***.

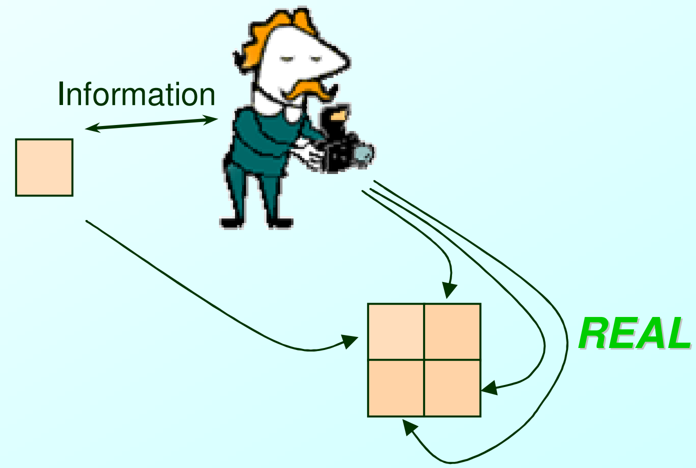

Real Experiment

Comparison between:

- **HR Image created from a sequence of frames obtained from real scene by multi-frame acquisition with CCD sensor;**
- **Upsampled frame generated with Bicubic Algorithm;**

(original size: 640x480 => upsampled : 1280x960).

LR Image

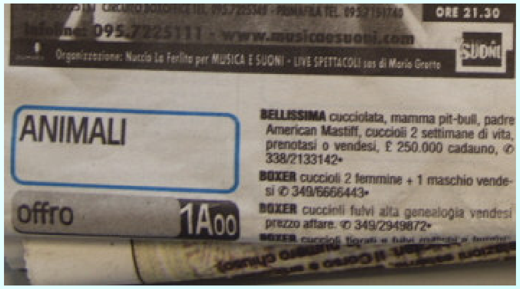

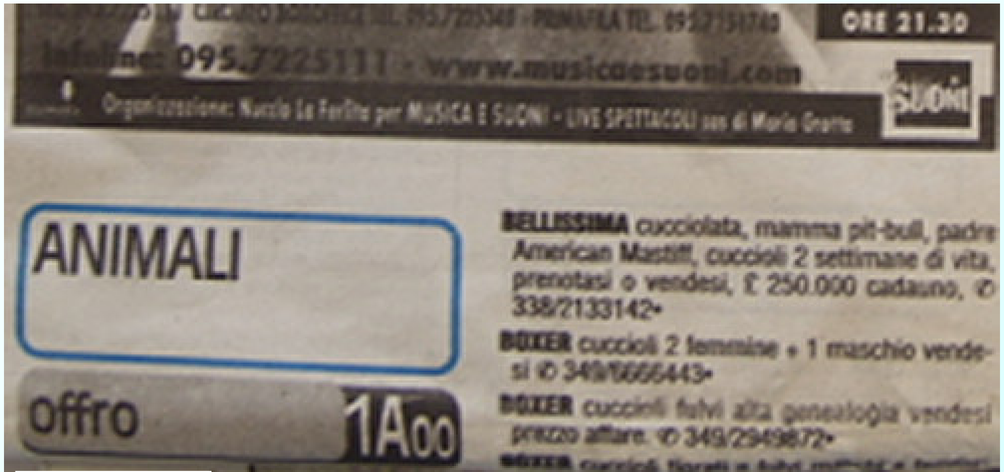

Bicubic

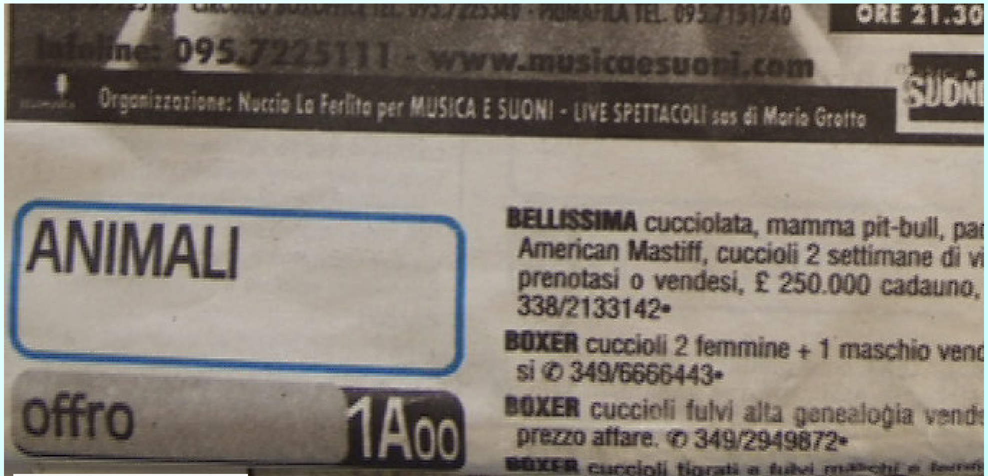

HR Image

Solid State Cameras- CCD

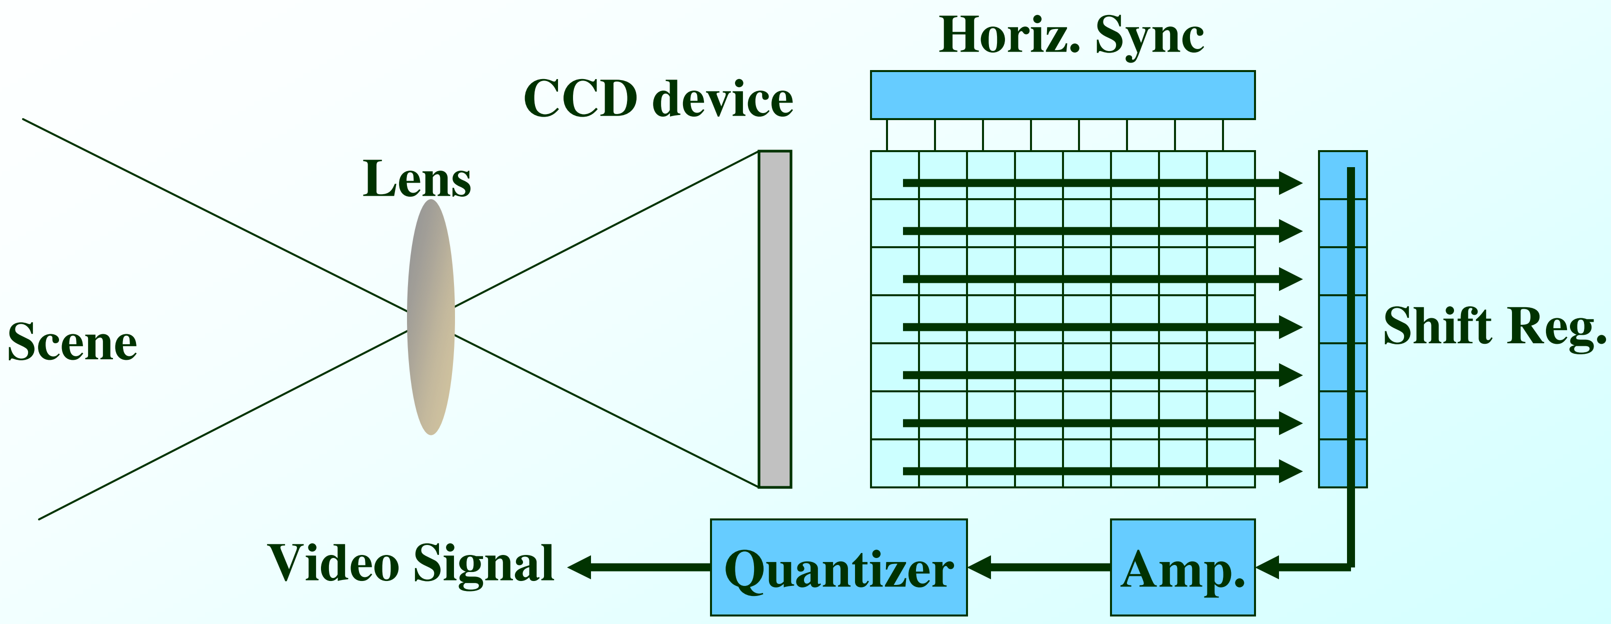

The CCD (Charge Couple Device) is basically a matrix of doped cells where electrons are liberated when light (Photons) hits. The liberated electrons are then moved (step by step) into a shift register by sync circuitry and then are read out and amplified, thus producing the brightness signal (in figure a column-wise scan is present).

Image Acquisition

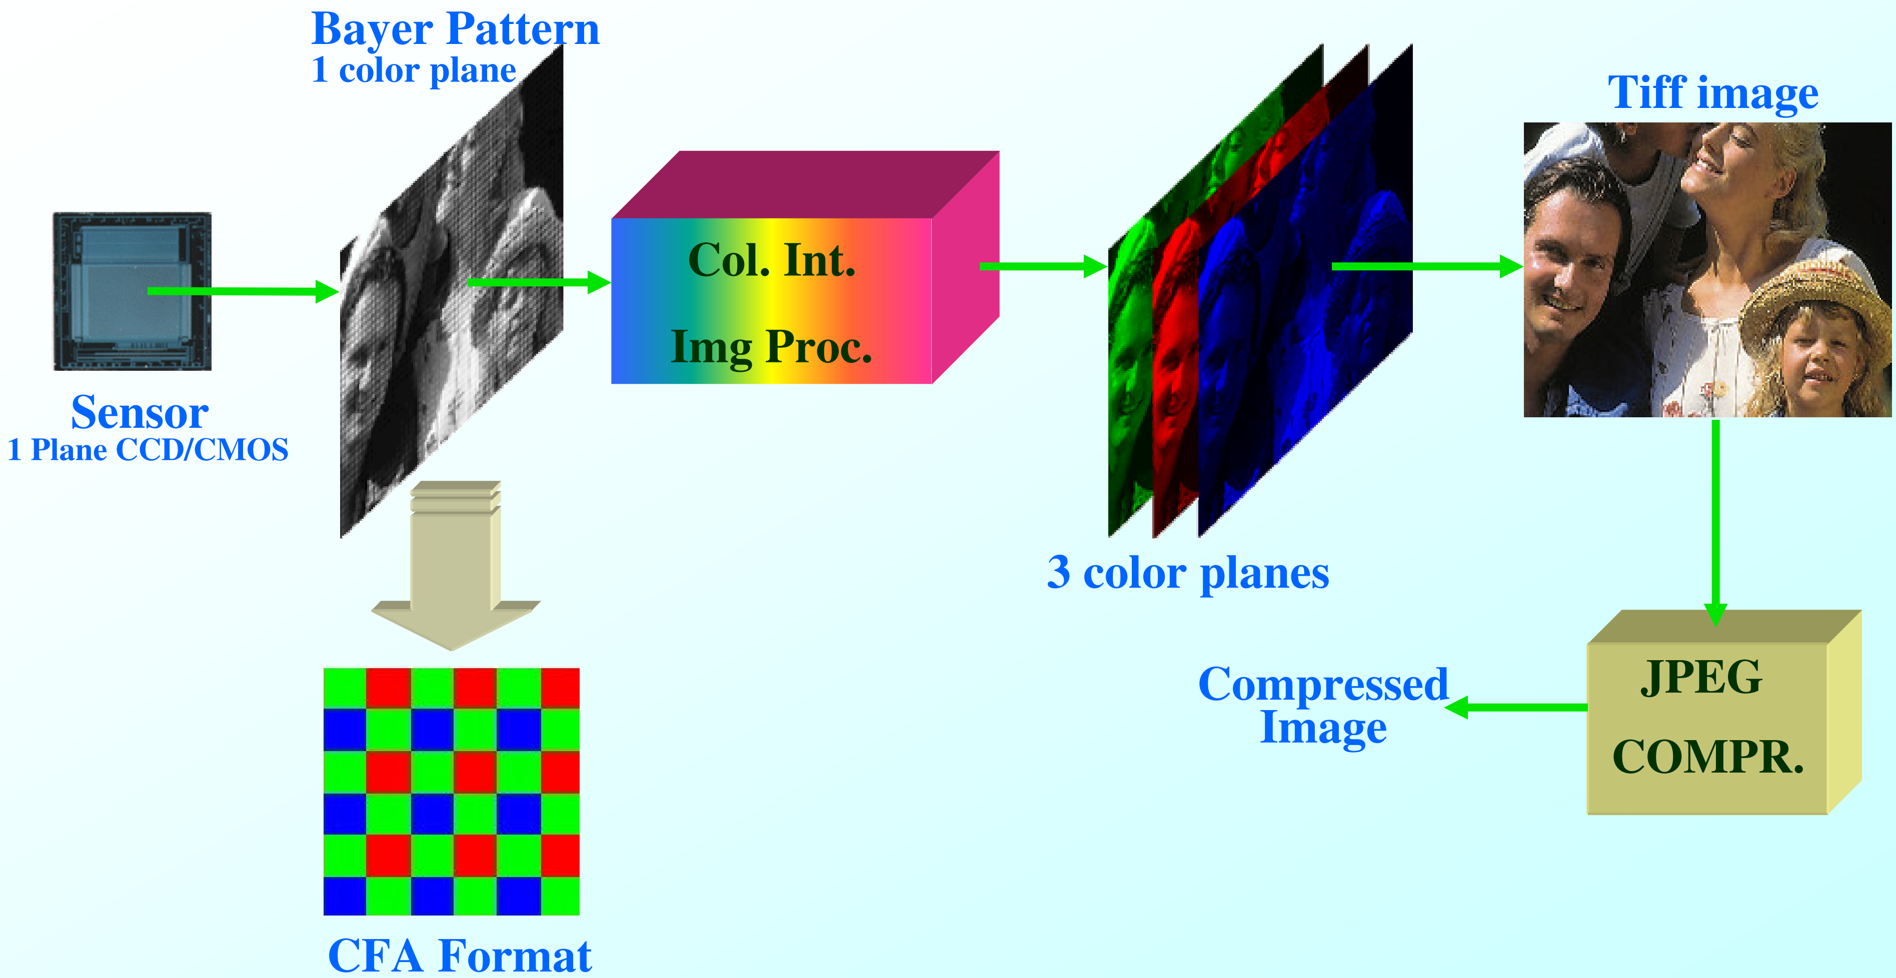

Bayer Pattern & Micro Lens

Color Sensor = Mono Sensor + RGB color filters

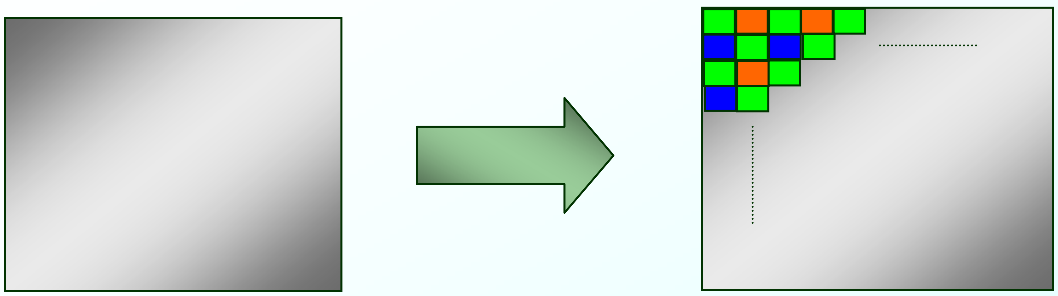

Mono Sensor

Color Sensor

Each pixel has a red, green or blue filter disposed in so called **Bayer Pattern**

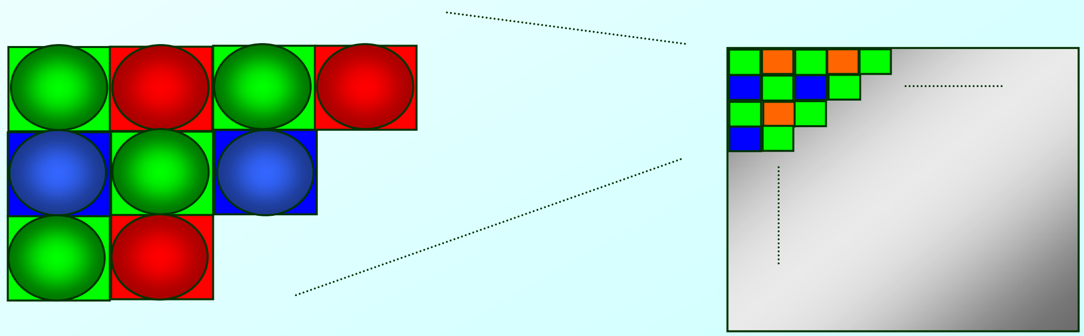

Each pixel/color filter has a small lens placed on top, improving significantly light gathering at pixel site.

A Bayer pattern image

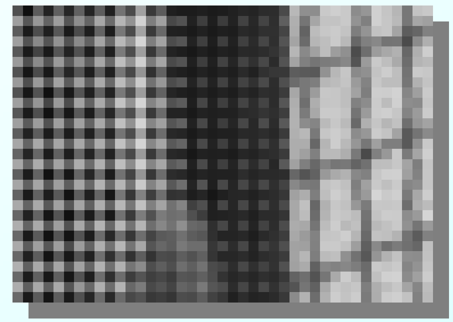

Original Bayer

Bayer Image

Original Bayer Data

Channel Splitting

Green

Red

Blue

Final True color image

Color Interpolation: some examples

Nearest Neighbor Replication

In this interpolation method each interpolated output pixel is assigned the value of the nearest pixel in the input image. The nearest neighbor can be any one of the upper, lower, left and right pixels.

An example is illustrated below for a 3x3 block in green plane. Here we assume the left neighboring pixel value is used to fill the missing ones.

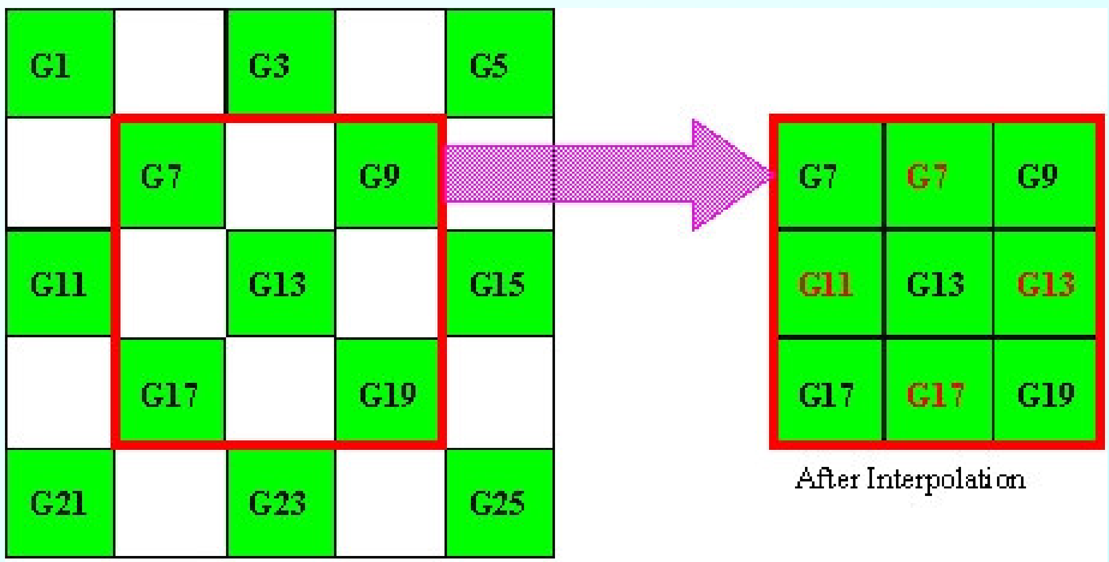

Before Interpolation

After Interpolation

Color Interpolation: some examples

Bilinear Interpolation

Interpolation of green pixels

The average of the upper, lower, left and right pixel values is assigned as the G value of the interpolated pixel. For example : G8 = (G3+G7+G9+G13) / 4

Interpolation of red/blue pixels

Interpolation of a red/blue pixel at a green position: the average of two adjacent pixel values in corresponding color is assigned to the interpolated pixel. For example : B7 = (B6+B8) / 2 ; R7 = (R2+R12) / 2

Interpolation of a red/blue pixel at a blue/red position : the average of four adjacent diagonal pixel values is assigned to the interpolated pixel. For example : R8 = (R2+R4+R12+R14) / 4 ; B12 = (B6+B8+B16+B18) / 4

G1	R2	G3	R4	G5
B6	G7	B8	G9	B10
G11	R12	G13	R14	G15
B16	G17	B18	G19	B20
G21	R22	G23	R24	G25

Color Interpolation: some examples

Edge Sensing Interpolation Algorithm (1/2)

From our description of non-adaptive algorithms, it can be seen that most of the color interpolation is done by averaging neighboring pixels indiscriminately. This causes an artifact -- the "zipper effect" in the interpolated image. To combat with this artifact, it is natural to derive an algorithm that can detect local spatial features present in the pixel neighborhood and then makes effective choices as to which predictor to use that neighborhood. The result is a reduction or elimination of "zipper-type" artifacts. And algorithms that involve this kind of "intelligent" detection and decision process are referred as ***adaptive color interpolation algorithms.***

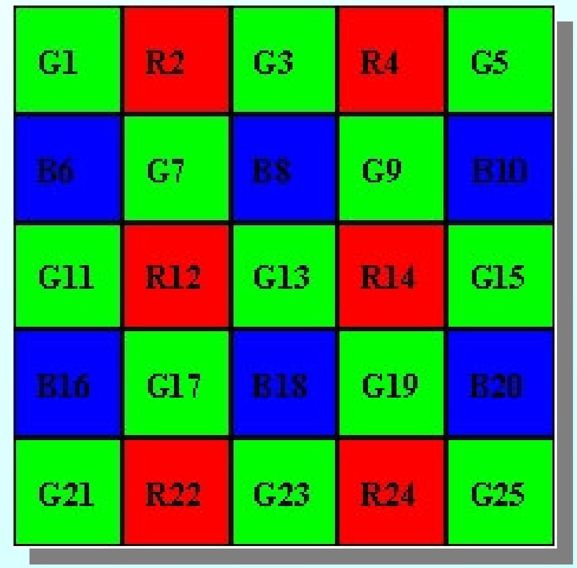

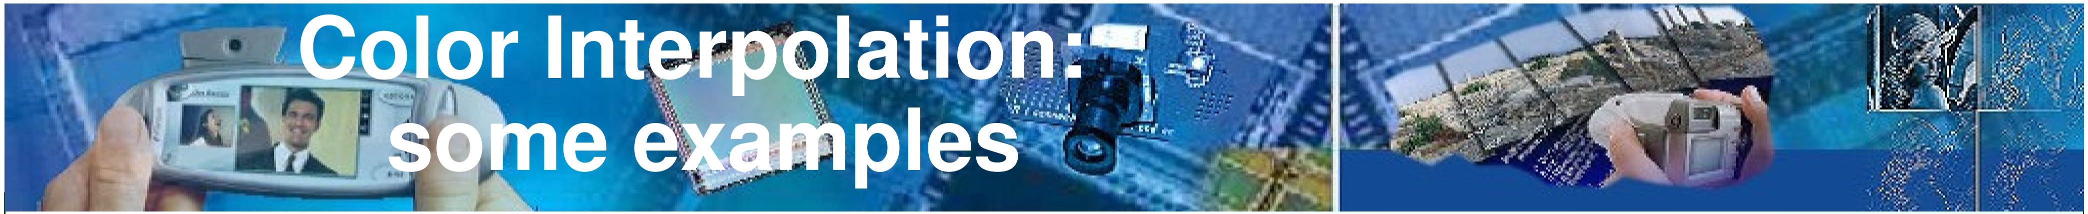

Color Interpolation: some examples

Edge Sensing Interpolation Algorithm (2/2)

Human visual systems are sensitive to edges present in the images and non-adaptive color interpolation algorithms often fail around edges since they are not able to detect "edges"

Interpolation of **green** pixels :

First, define two gradients, one in horizontal direction, the other in vertical direction, for each blue/red position. For instance, consider B8: define two gradients as $Dh = |G7-G9|$ $Dv = |G3-G13|$, where | . | denotes absolute value.Define some threshold value T

The algorithm then can be described as follows:

If Dh < T and Dv > T

G8=(G7+G9)/2;

Else if Dh > T and Dv < T

G8=(G3+G13)/2;

Else

G8=(G3+G7+G9+G13)/4;

End

The choice of T depends on the images and can have different optimum values from different neighborhoods. A particular choice of T is:

$T = (Dh+Dv)/2$

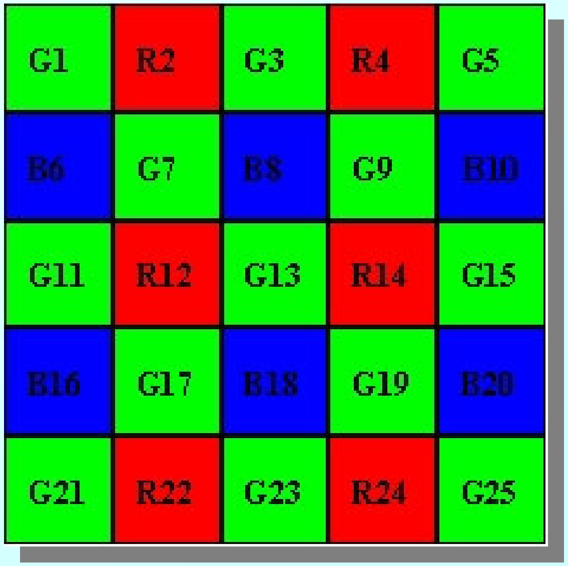

Color Interpolation: Results

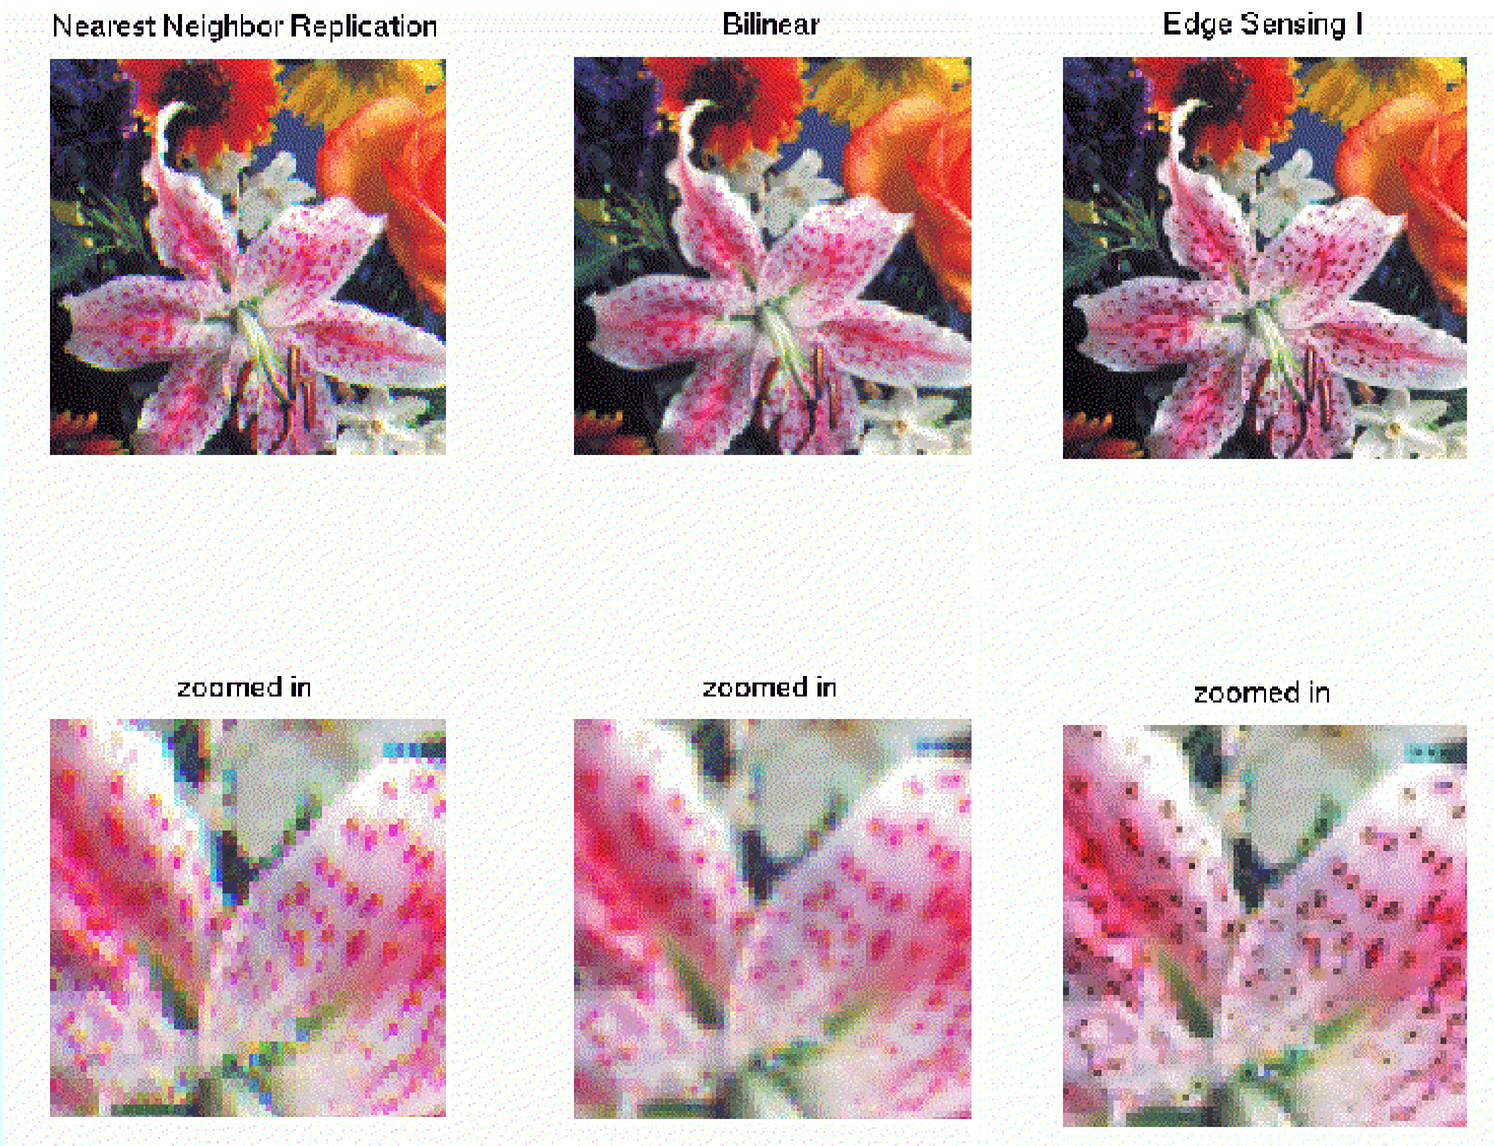

Bayer Quality

In any case there is a *drawback* using only a chromatic data for each pixel