Cvision – 2
Digital Imaging

António J. R. Neves
(an@ua.pt)
&
João Paulo Silva Cunha
&
Bernardo Cunha

IEETA / Universidade de Aveiro

Outline

- Image sensors
- Camera calibration
- Sampling and quantization
- Data structures for digital images
- Histograms

Acknowledgements: Most of this course is based on the excellent courses offered by Prof. Shree Nayar at Columbia University, USA and by Prof. Srinivasa Narasimhan at CMU, USA. This was also based on Prof. Miguel Coimbra's slides. Please acknowledge the original source when reusing these slides for academic purposes.

Topic: Image Sensors

- Image sensors
- Camera Calibration
- Sampling and quantization
- Data structures for digital images
- Histograms

Image Sensors

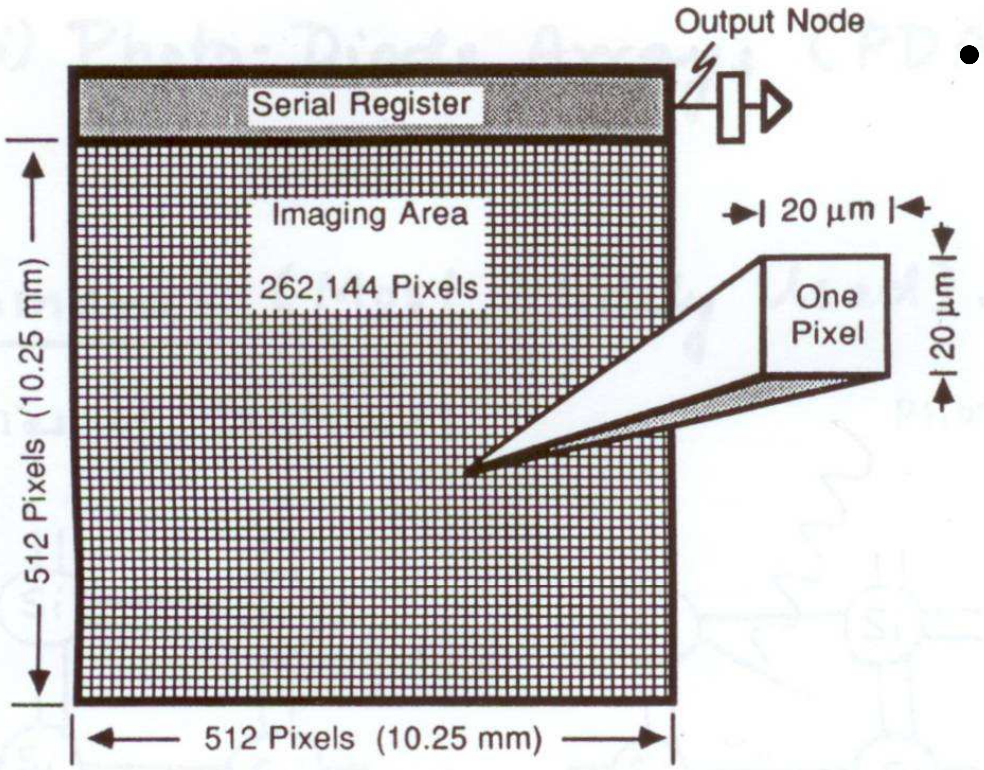

FIG. 4. Typical 512 × 512 CCD.

- Considerations
  - Speed
  - Resolution
  - Signal / Noise Ratio
  - Cost

Image Sensors

- Convert light into an electric charge

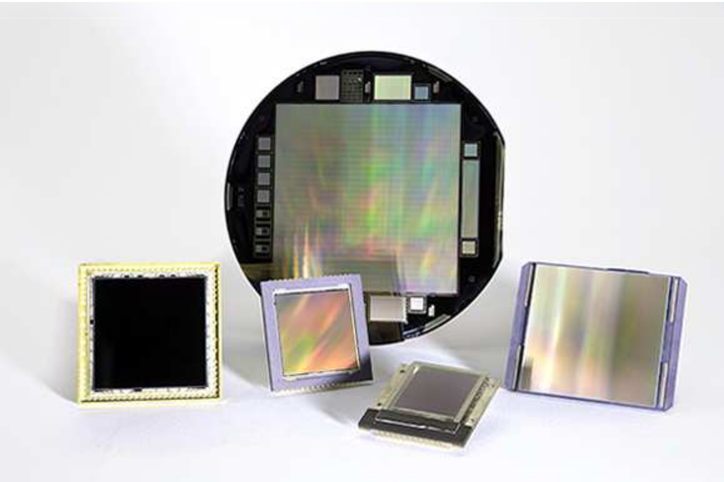

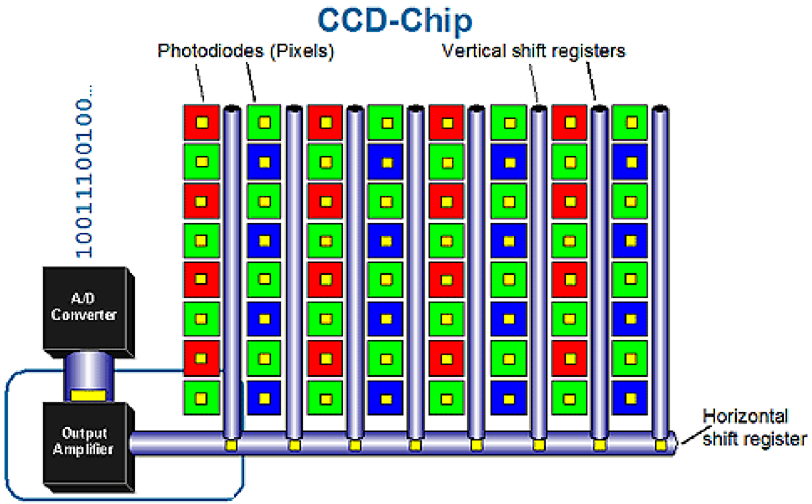

CCD (charge coupled device)

Higher dynamic range

High uniformity

Lower noise

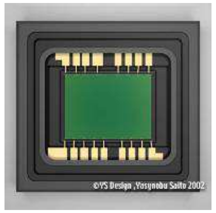

CMOS (complementary metal Oxide semiconductor)

Lower voltage

Higher speed

Lower system complexity

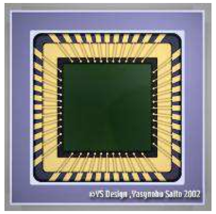

CCD Performance Characteristics

- Linearity Principle: Incoming photon flux vs. Output Signal
  - Sometimes cameras are made non-linear on purpose.
  - Calibration must be done (using reflectance charts)
- Dark Current Noise: Non-zero output signal when incoming light is zero
- Sensitivity: Minimum detectable signal produced by camera

Sensing Brightness

Incoming light has a **spectral distribution** $p(\lambda)$

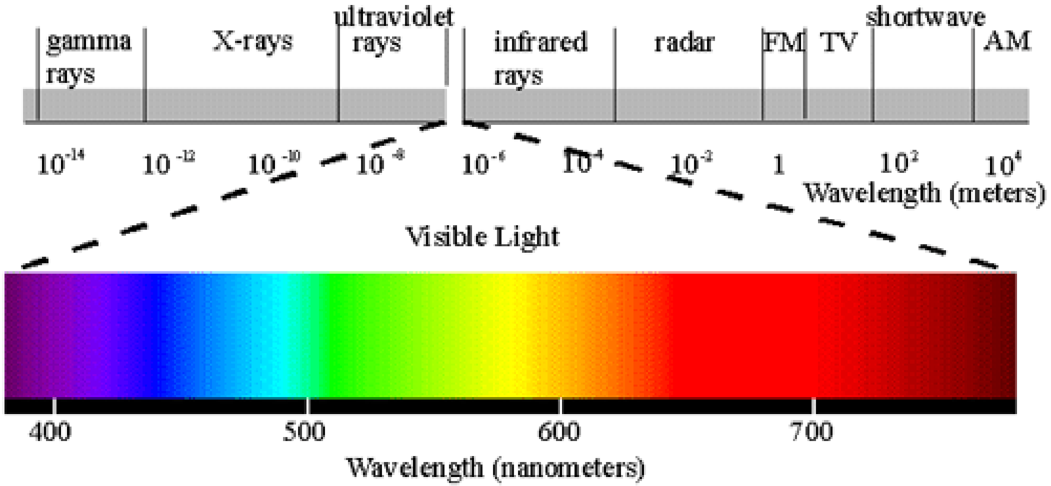

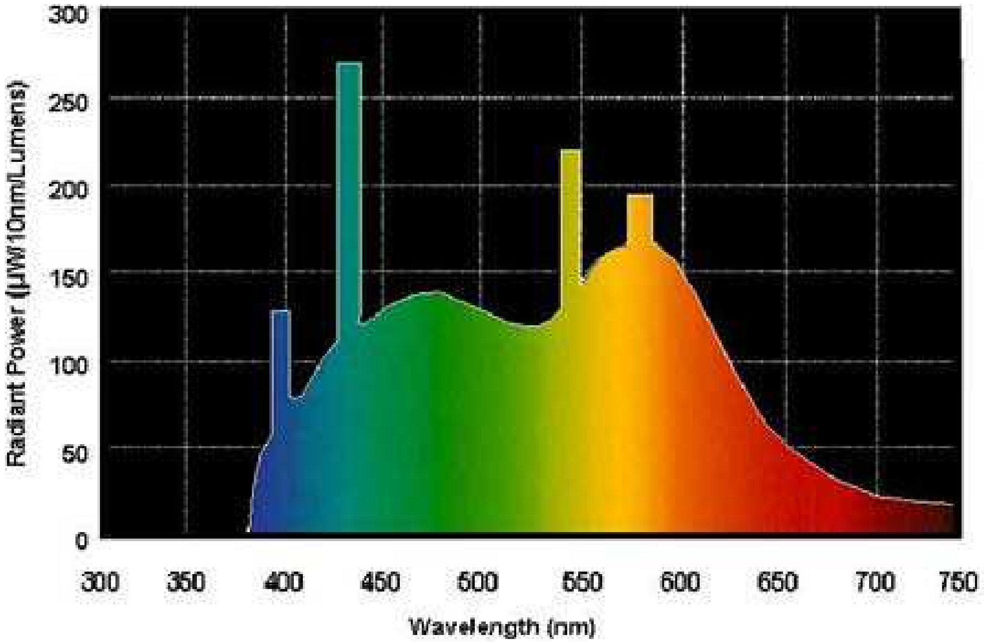

So the pixel intensity becomes

$$I = k \int_{-\infty}^{\infty} q(\lambda) p(\lambda) d\lambda$$

How do we sense colour?

- Do we have infinite number of filters?

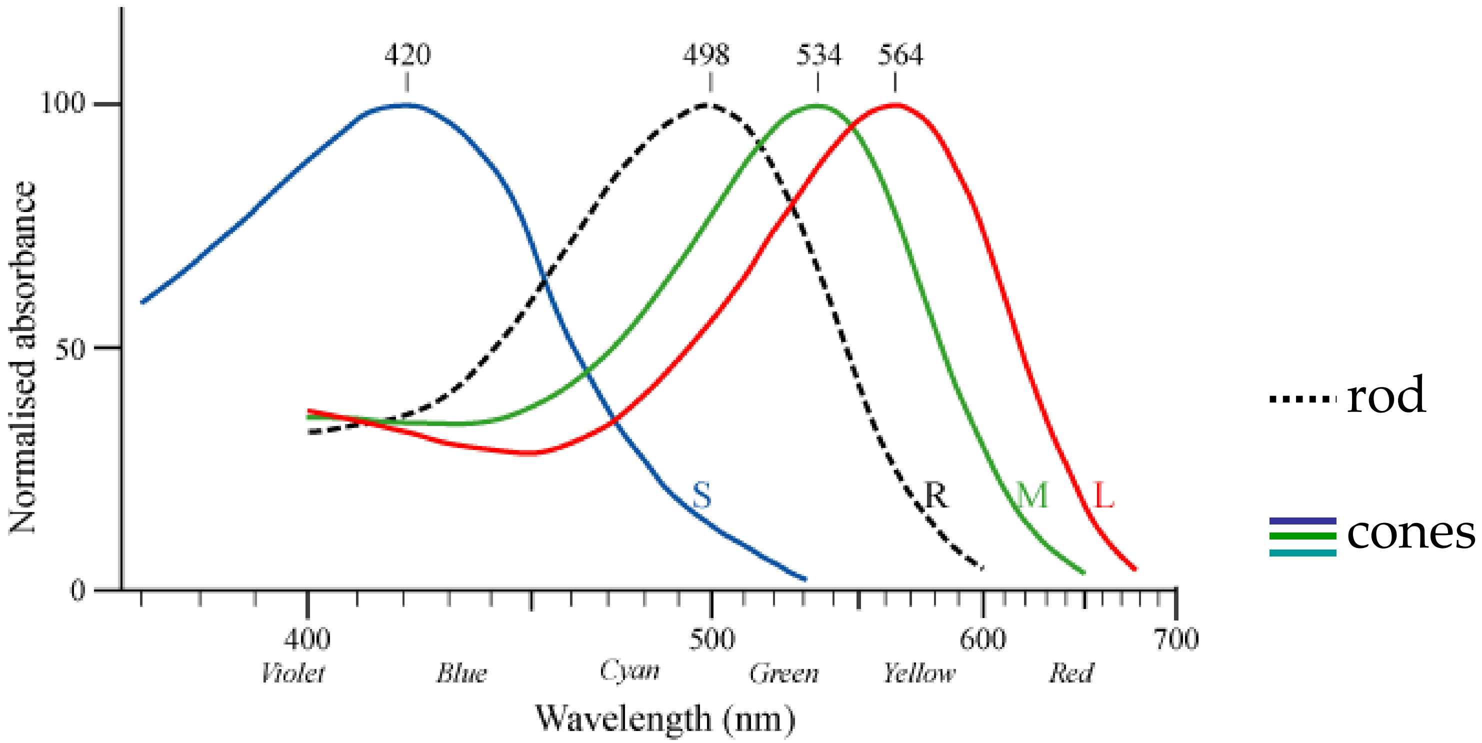

Three filters of different spectral responses

Sensing Colour

- Tristimulus (trichromatic) values (I_R, I_G, I_B)

Camera's spectral response functions: $h_R(\lambda), h_G(\lambda), h_B(\lambda)$

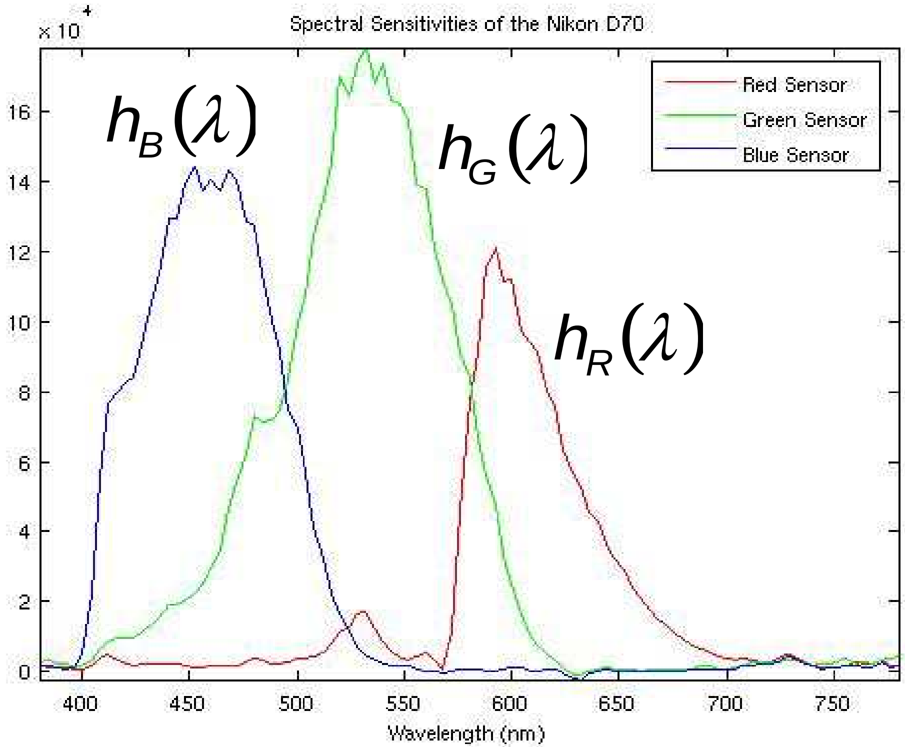

$$I_R = k \int_{-\infty}^{\infty} h_R(\lambda) p(\lambda) d\lambda$$

$$I_G = k \int_{-\infty}^{\infty} h_G(\lambda) p(\lambda) d\lambda$$

$$I_B = k \int_{-\infty}^{\infty} h_B(\lambda) p(\lambda) d\lambda$$

Sensing Colour

3 CCD

light

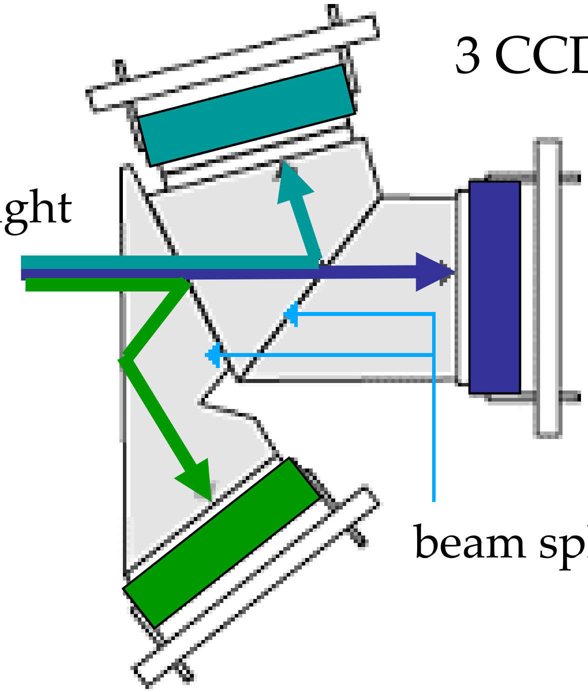

beam splitter

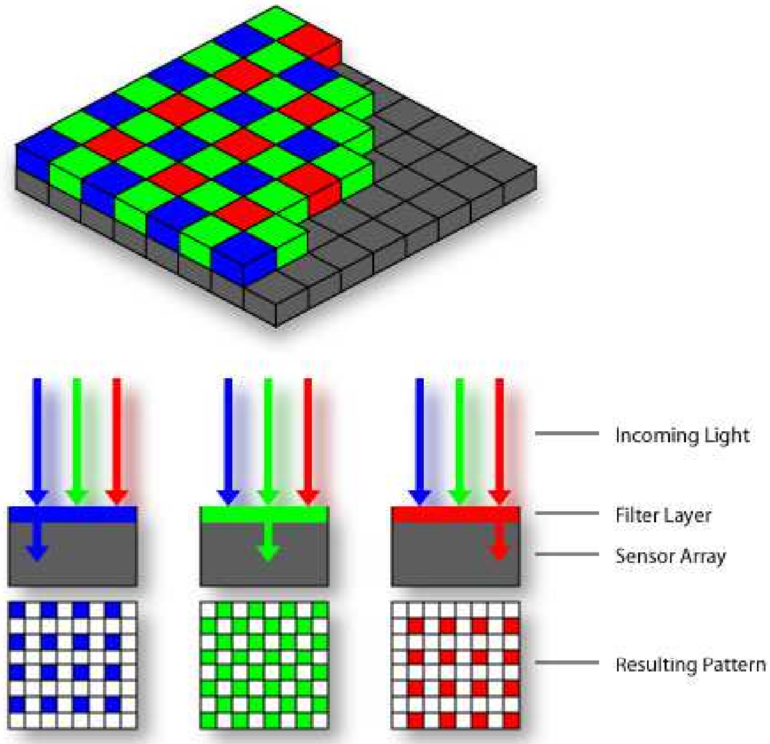

Bayer pattern

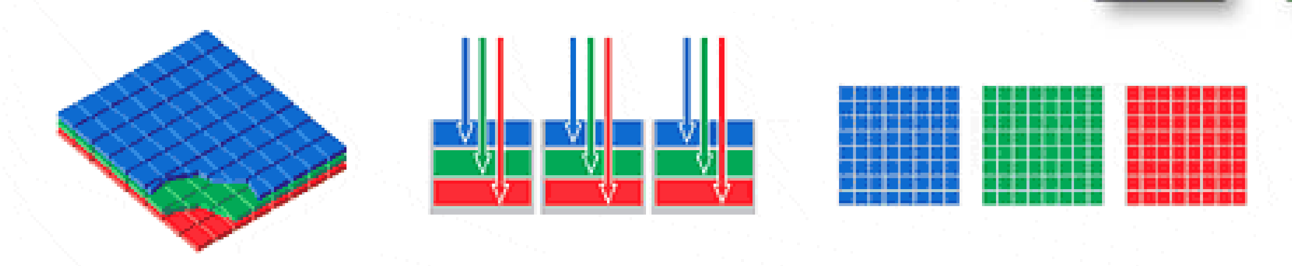

Foveon X3™

Several types of cameras

Several types of cameras

- Several interfaces (Firewire, GigE, CameraLink, USB, ...).
- Scientific usage (high resolution, long exposure time, ...).
- High speed (ex. 1000 fps).
- Linear (ex. 10000 lines per second).
- 3D
- Infrared (ex. 8 to 14 μm).
- High dynamic range (ex. using a prism and two sensors).
- Multispectral

Topic: Camera Calibration

- Image sensors
- Camera Calibration
- Sampling and quantization
- Data structures for digital images
- Histograms

Definitions - Luminance

Luminance

Luminance is normally defined as a measurement of the photometric luminous intensity per unit area of light travelling in a given direction.

Therefore it is used to describe the amount of light that goes through, or is emitted from, a particular area, and falls within a given solid angle.

The SI unit for luminance is candela per square meter (cd/m2).

The CGS unit of luminance is the *stilb*, which is equal to one candela per square centimeter or 10 kcd/m2.

Definitions - Chrominance

Chrominance

Chrominance is a numeral that describes the way a certain amount of light is distributed among the visible spectrum.

A black and white image has a balanced distribution of energy among to the visible spectrum matched to the band pass characteristics of the human visual system. This means that when viewed by a human a B&W image has no color information which means that its color information is zero.

Therefore, chrominance has no luminance information but is used together with it to describe a colored image defined, for instance, by an RGB triplet.

Any RGB triplet in which the value of R=G=B has no chrominance information.

RGB & YUV

Separating Luminance from Chrominance

Given an RGB triplet, we can define a derived triplet in which luminance and chrominance can be separated:

$$Y = W_r R + W_g G + W_b B$$

Luminance

$$U = U\max \frac{B-Y}{1-W_b} \approx 0.492(B-Y)$$

$$V = V\max \frac{R-Y}{1-W_r} \approx 0.877(R-Y)$$

Chrominance

where

$$W_r = 0.299$$
$$W_B = 0.114$$
$$W_G = 0.587$$
$$U_{\max} = 0.436$$
$$V_{\max} = 0.615$$

This values originally derivates from the general model of the human visual system and had a significant impact on the ability to develop a television color system compatible with the previous B&W television systems.

A symetric operation can be performed in order to recover the original RGB triple.

The image processing pipeline

Image processing pipeline

A typical image processing pipeline (inside the image device) for a tri-stimulus system is shown bellow. This processing can be performed on the YUV or RGB components depending on the system. This should be understood as a mere example.

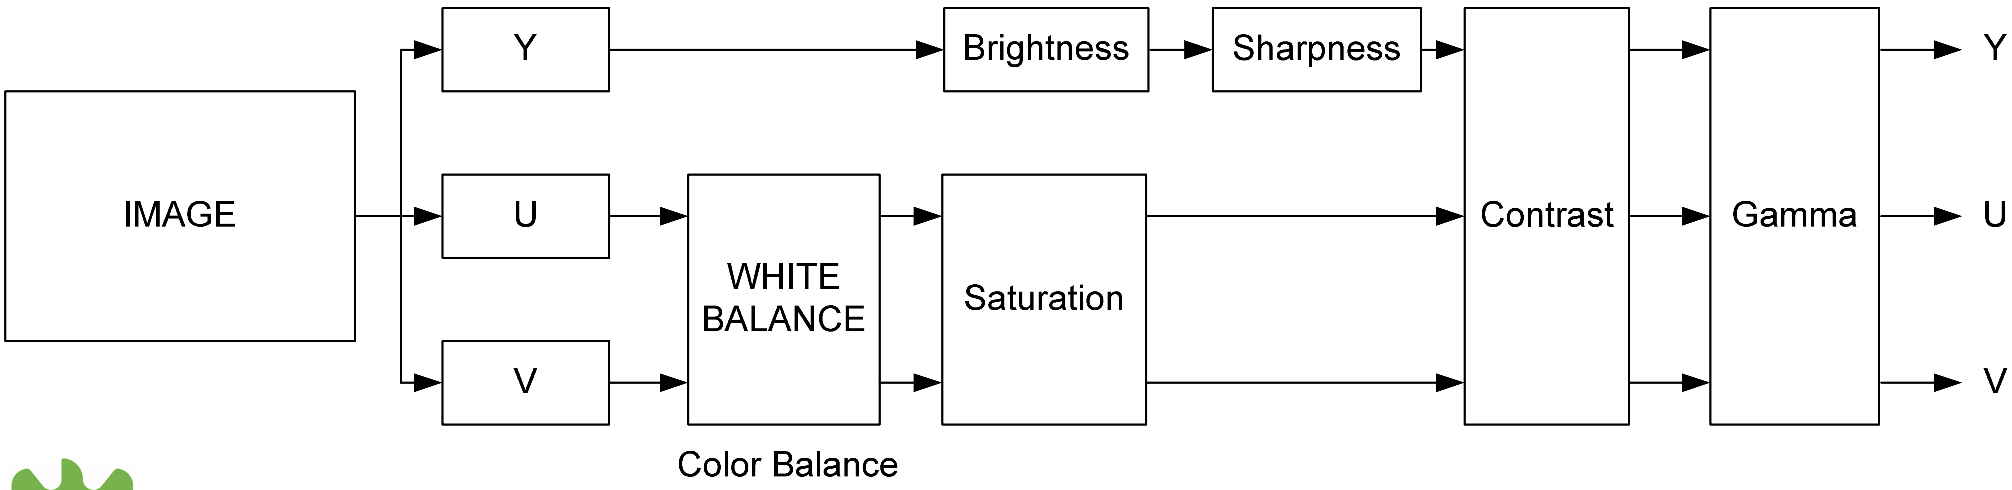

The image processing pipeline

Image processing pipeline

Depending on the system, more or less image parameters may be available for the user to control. Also, some of these parameters (namely brightness, contrast and saturation) are also intrinsic original image characteristics apart from being externally controllable parameters.

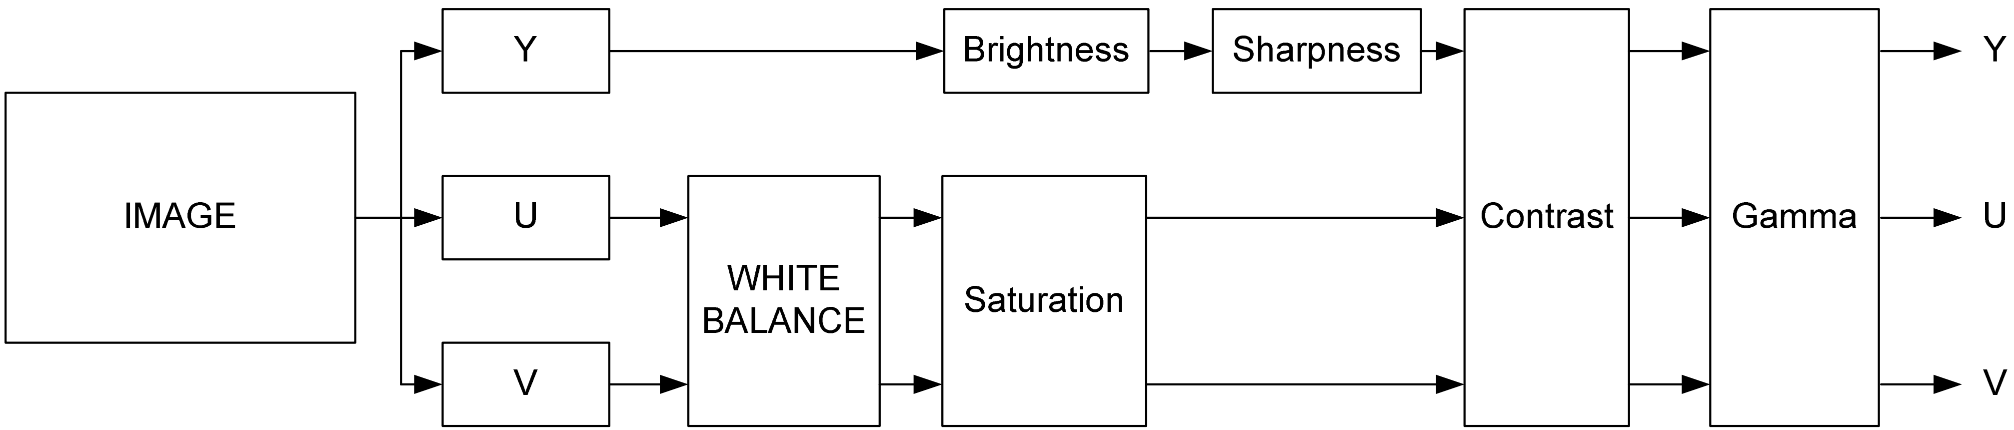

Color Balance

Brightness

Brightness (as an intrinsic image characteristic)

Brightness is one on the intrinsic original image characteristics. It represents a measure of the average amount of light that is integrated over the image during the exposure time. Exposure time (that is, the period of time during which the sensor receives light while forming the image, may or may not be a controllable parameter of the image device).

If the brightness it too high overexposure may occur which will white saturate part or the totality of the image.

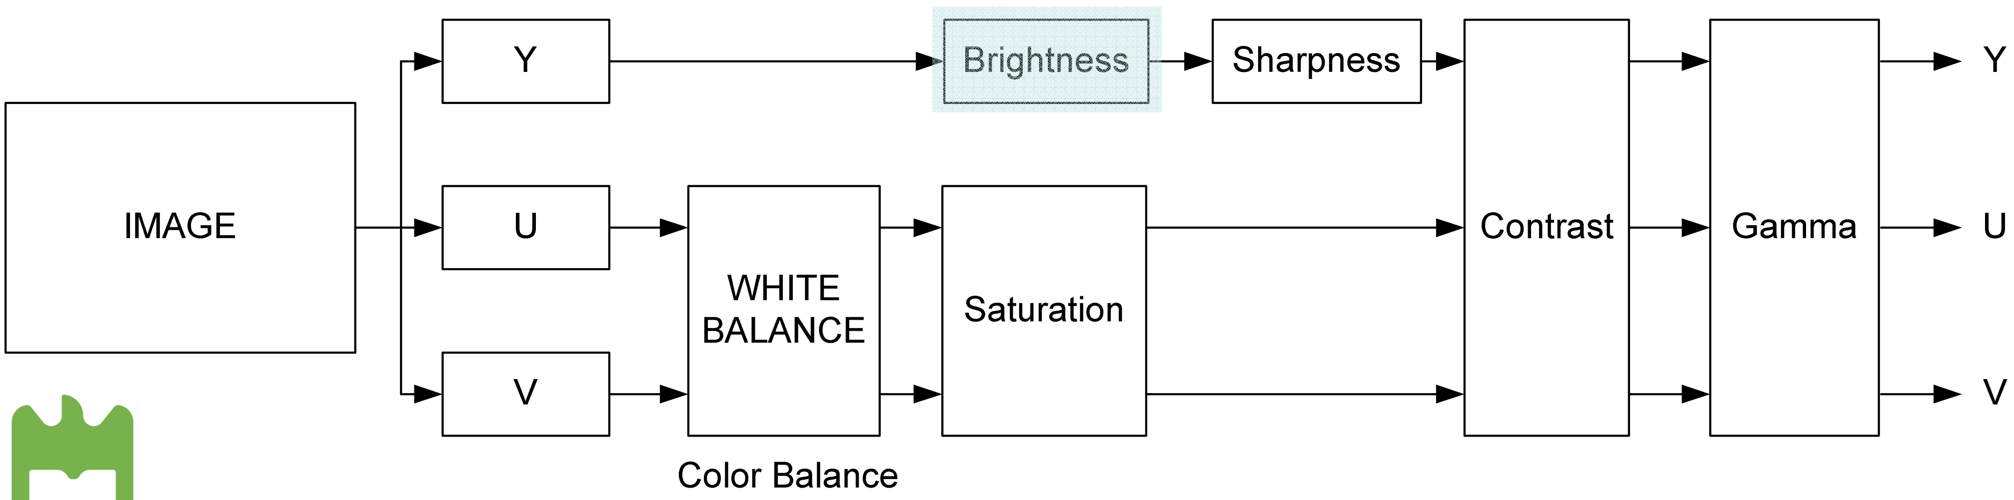

Brightness

Brightness (as a controllable parameter)

The brightness parameter is basically a constant (or offset) that can be added (subtracted) from the luminance component of the image.

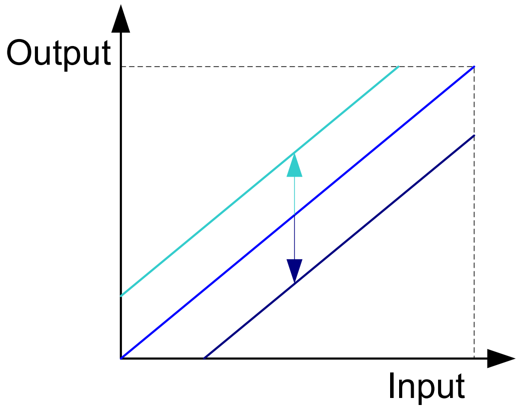

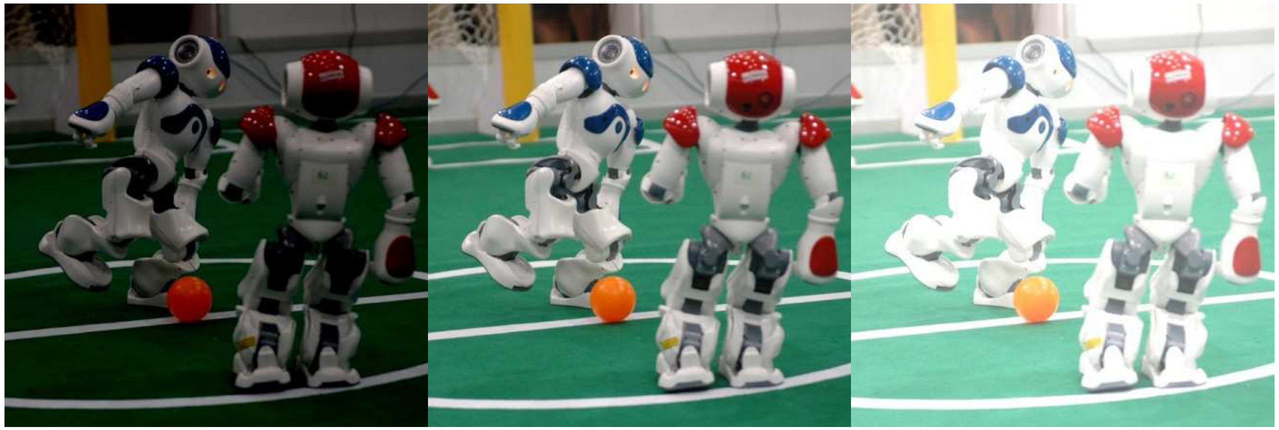

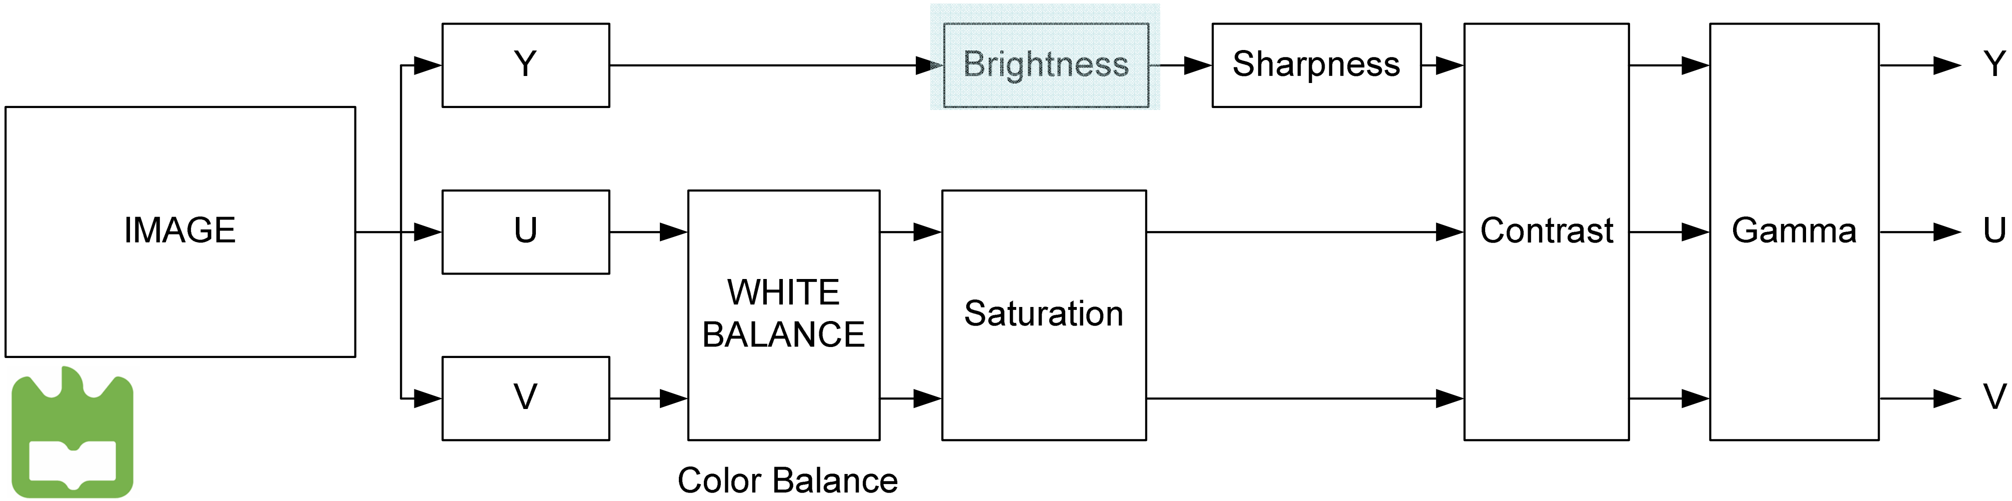

Contrast

Contrast (as an intrinsic image characteristic)

There is not a unique definition of contrast. On of the most used is that contrast is the difference in luminance (or color) along the 2D space that makes an object distinguishable. In visual perception of the real world, contrast is determined by the difference in the color and brightness of the object and other objects within the same field of view. The faster and higher the luminance (or color) changes along the space the higher the contrast is.

The maximum possible contrast of an image is also denominated contrast ratio or dynamic range.

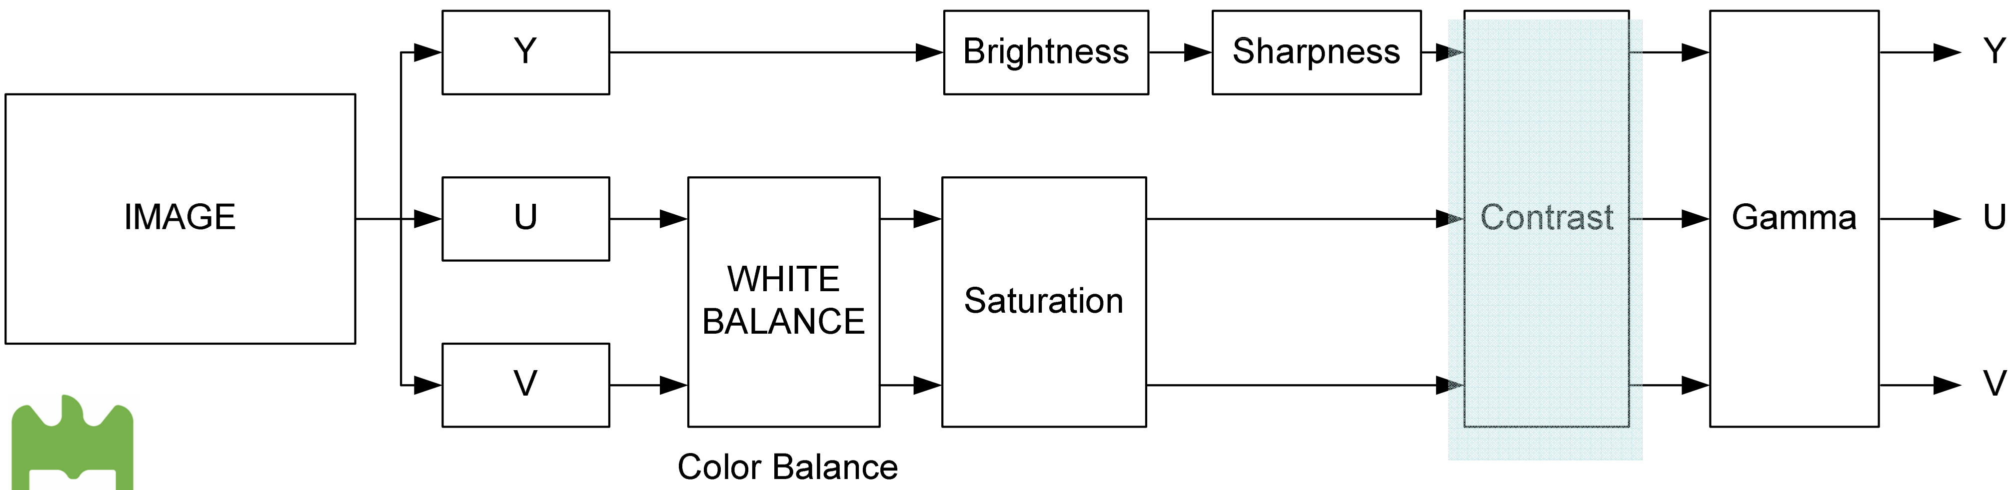

Contrast

Contrast (as an intrinsic image characteristic)

One of the possible definitions of contrast is given by the expression

$$\frac{\text{Luminance diference}}{\text{Average luminance}}$$

The human eye contrast sensitivity function is a typical band-pass filter with a maximum at around 4 cycles per degree with sensitivity reducing to both sides off that maximum. This means that the human visual system can detect lower contrast differences at 4 cycles per degree than at any other spatial frequency.

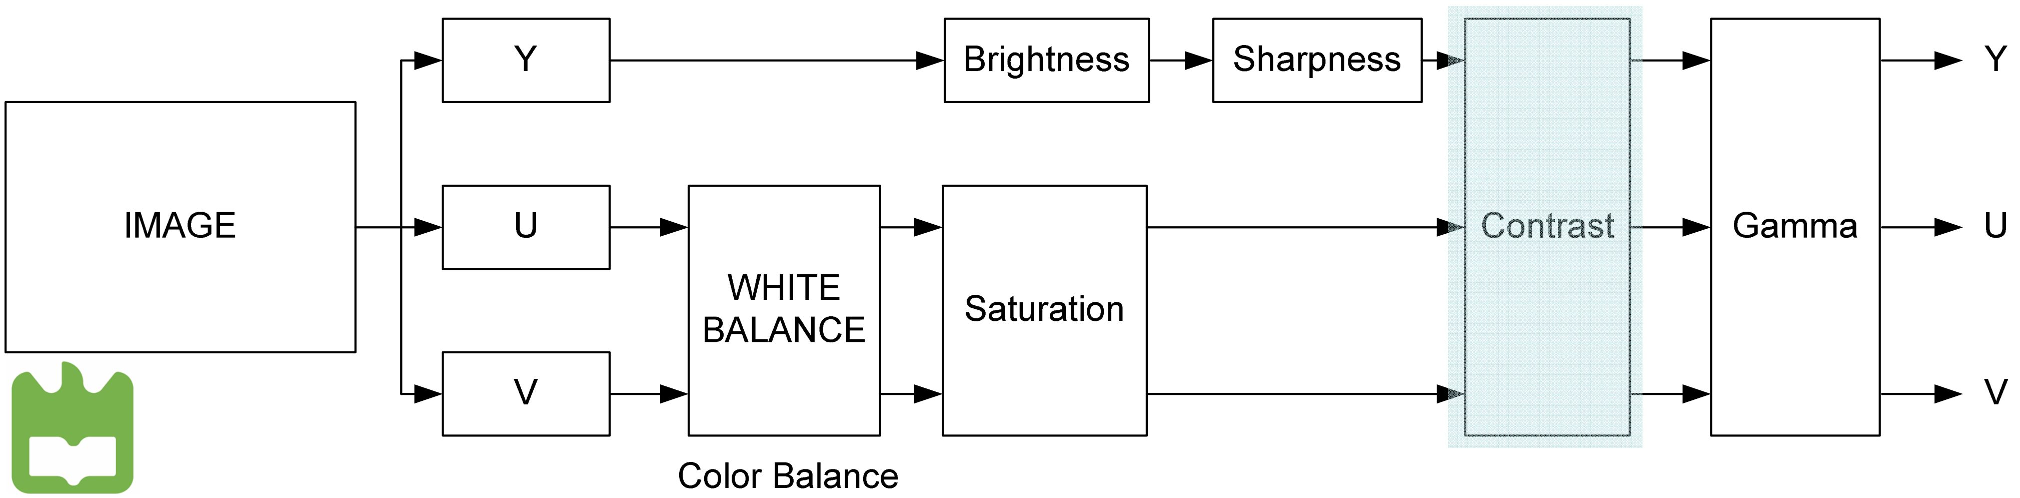

Contrast

Contrast (as a controllable parameter)

The contrast parameter is basically a variation in the gain control function of the luminance component of the image.

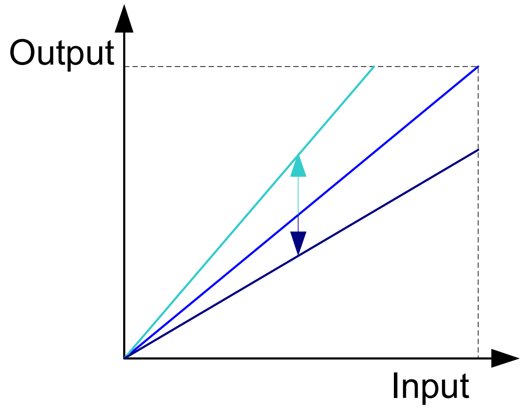

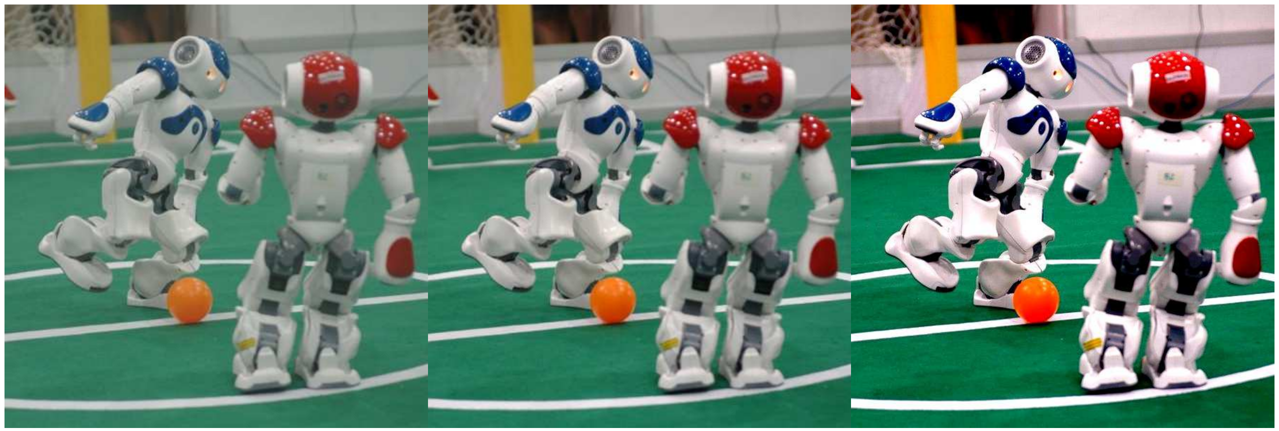

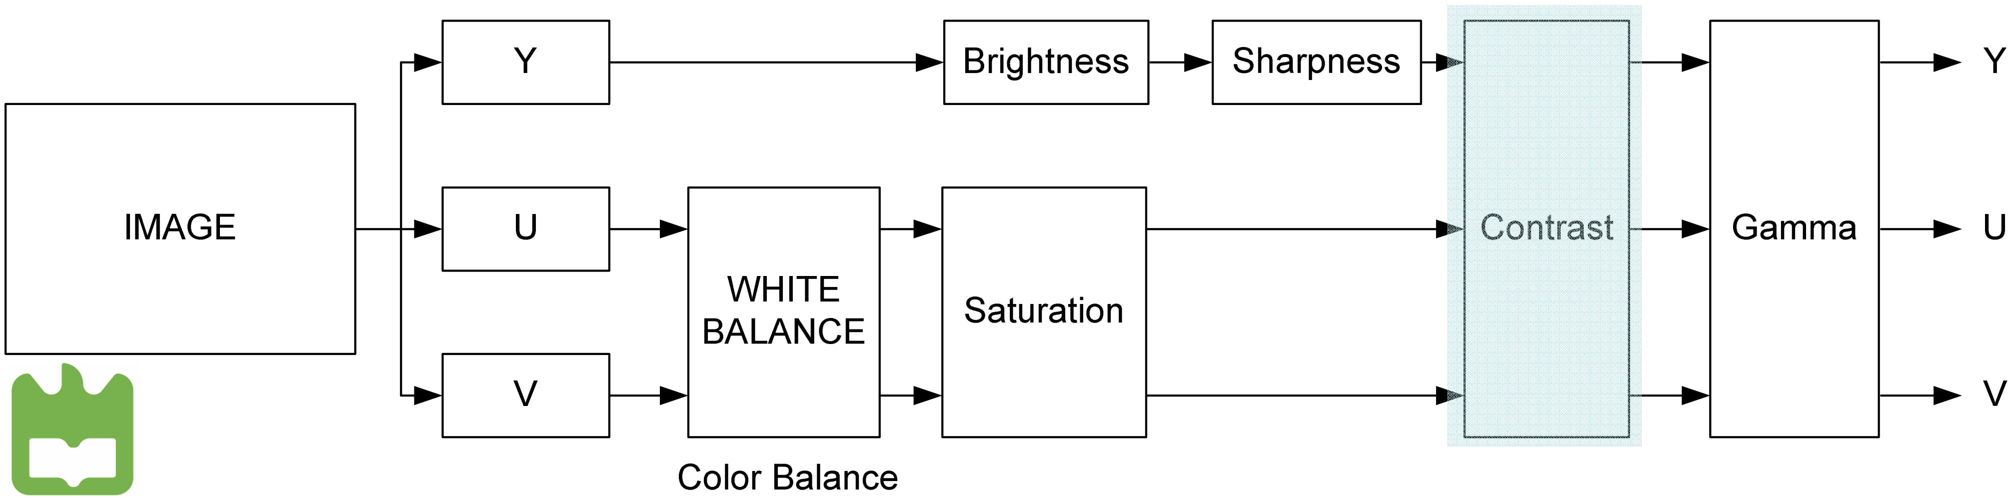

Contrast + Brightness

Contrast + Brightness (as controllable parameters)

It is common that contrat and brightness are actually a combined single transfer function.

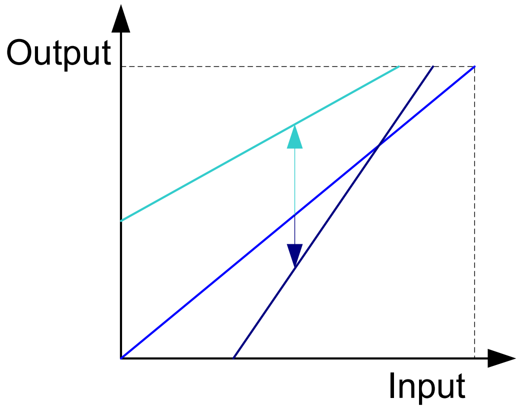

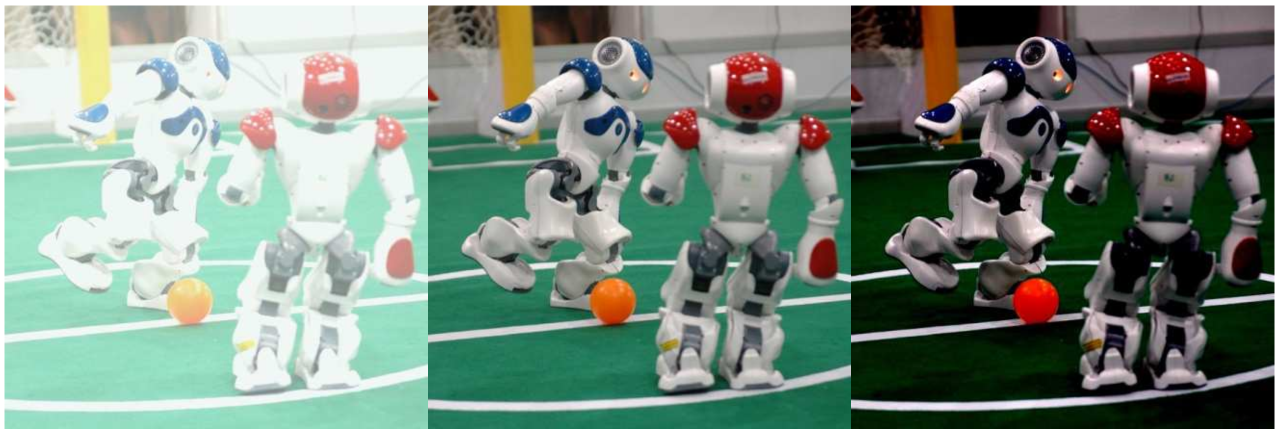

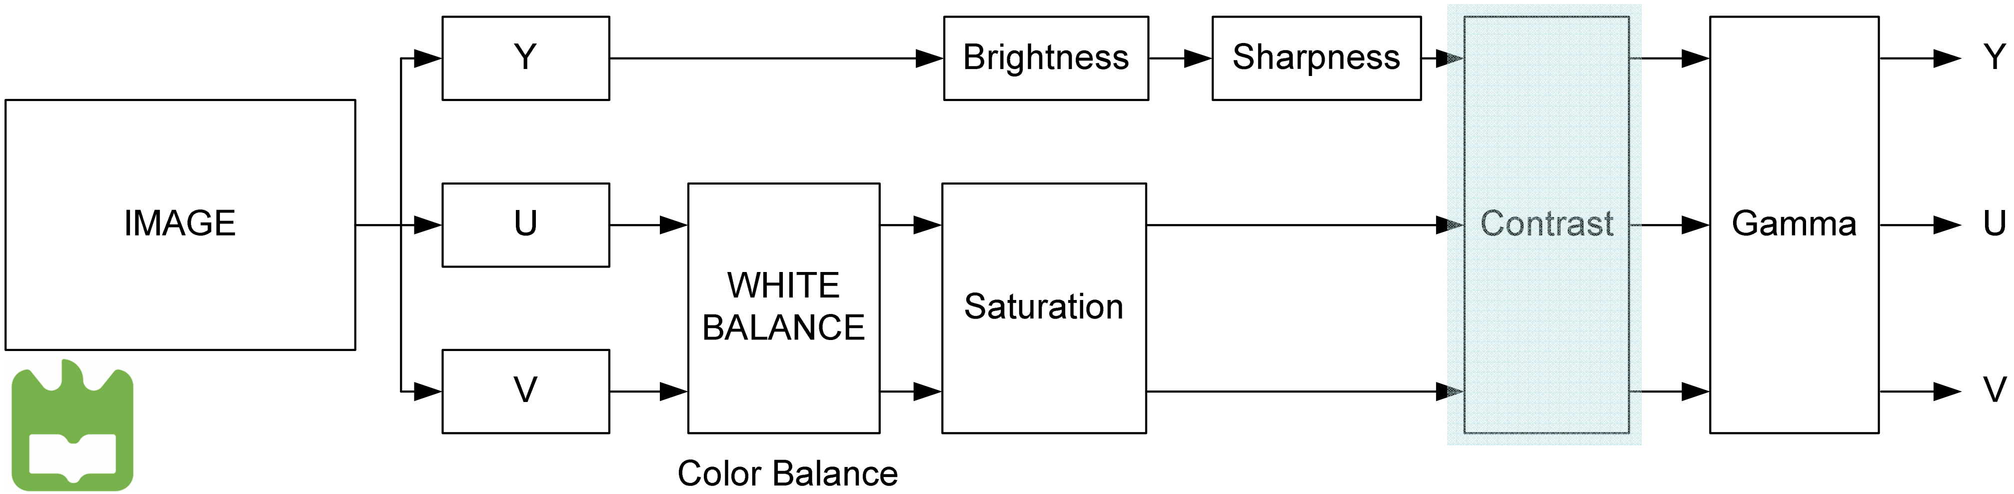

White Balance

White Balance (as controllable parameters)

White balance is the global adjustment of the intensities of the colors (typically red, green, and blue primary colors).

An important goal of this adjustment is to render specific colors – particularly neutral colors – correctly; hence, the general method is sometimes called gray balance, neutral balance, or white balance.

This balance is required because of different color spectrum energy distribution depending on the illumination source.

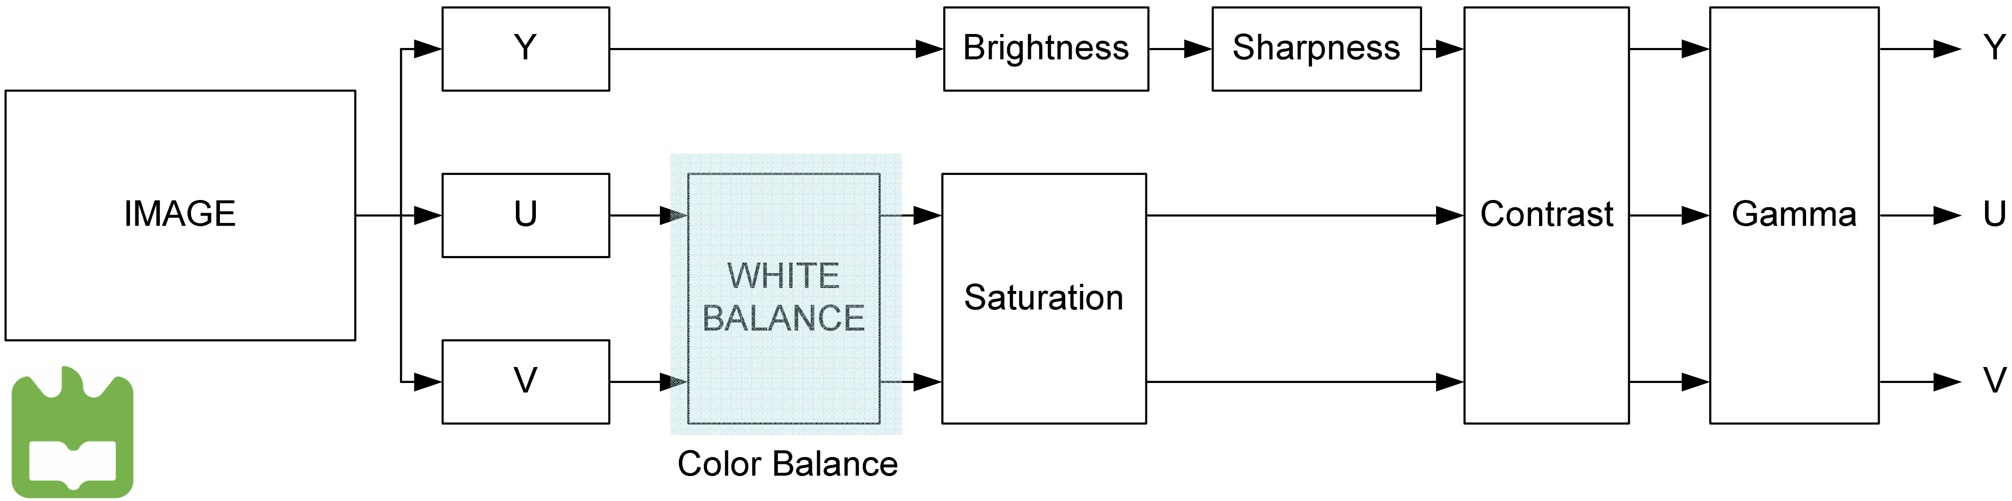

White Balance

White Balance

Examples

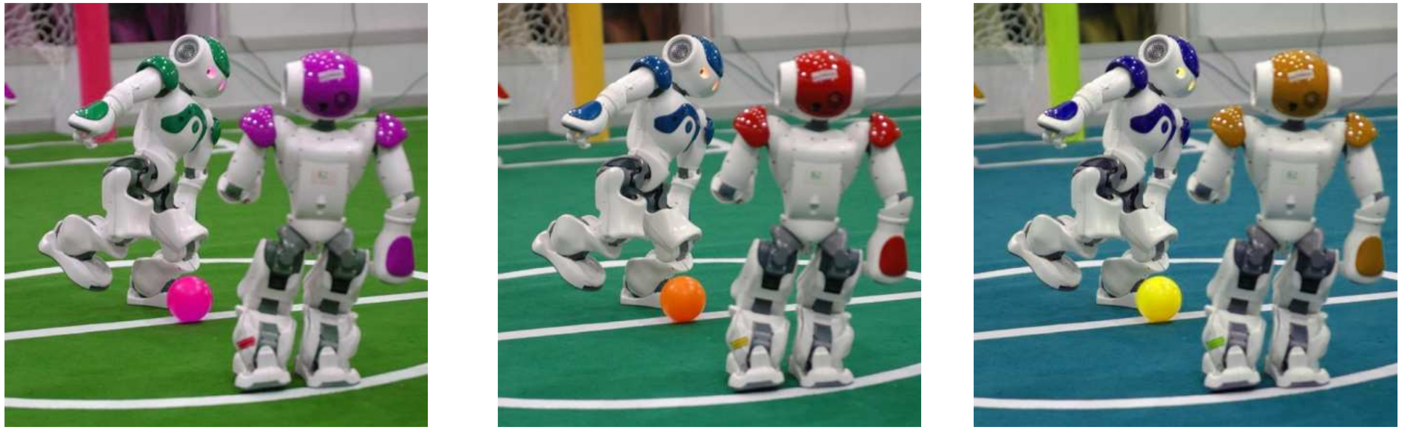

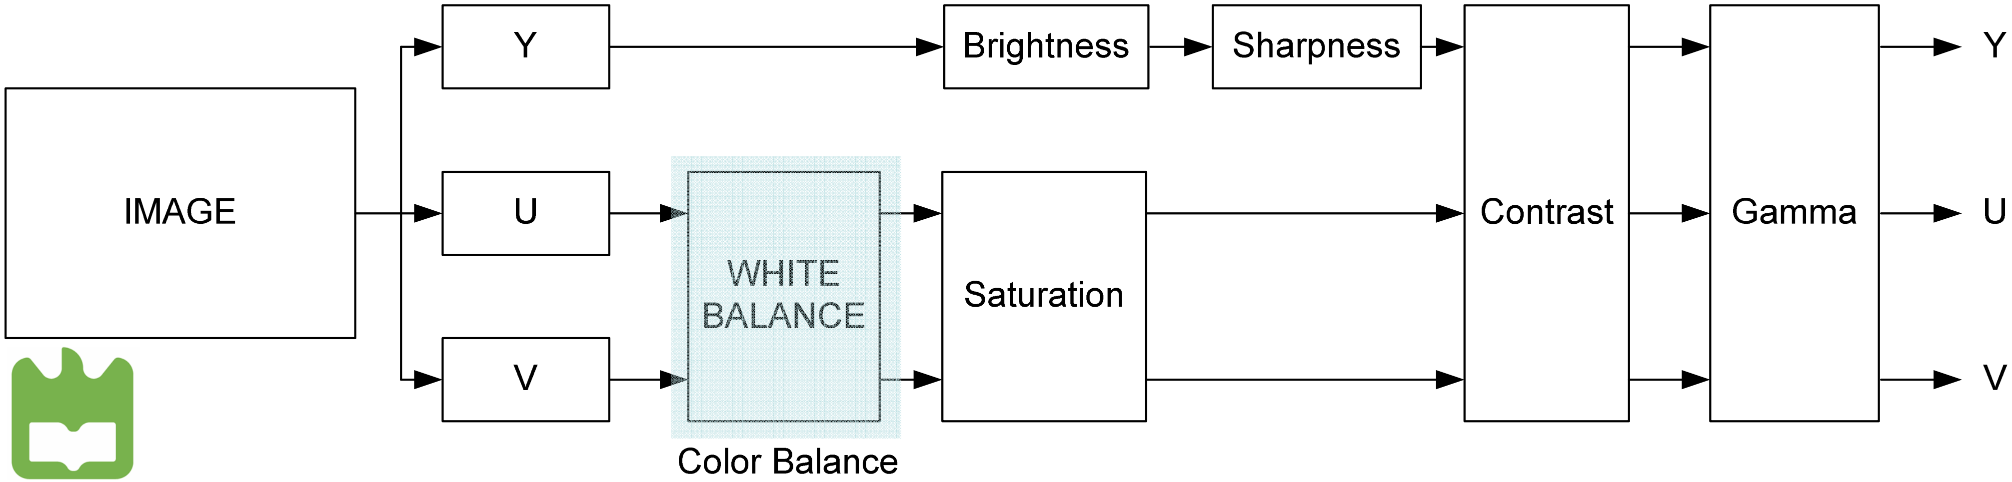

Saturation

Saturation (as an intrinsic image characteristic)

The saturation of a color is determined by a combination of light intensity that is acquired by a pixel and how much this light it is distributed across the spectrum of different wavelengths. The most purest (most saturated) color is obtained when using a single wavelength at a high intensity (laser light is a good example).

If the light intensity declines, then, as a result, the saturation also decline.

A non saturated image (B&W) has a spectrum distribution that matches the human eye spectrum sensibility. Saturation is sometimes also defined as the amount of white you have blended into a pure color.

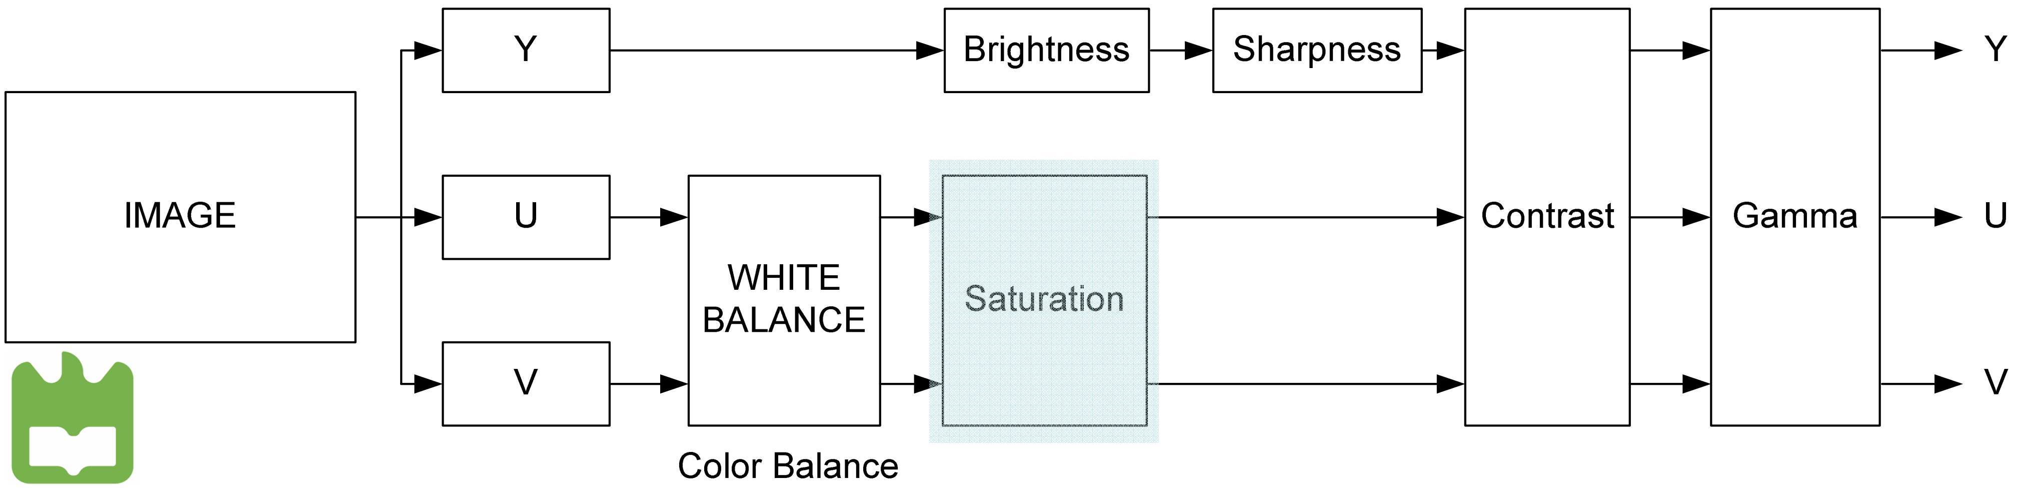

Saturation

Saturation (as a controllable parameter)

To reduce the saturation of an image we can add white to the original colors. In fact this is the same as changing the gain of the U and V chromatic components.

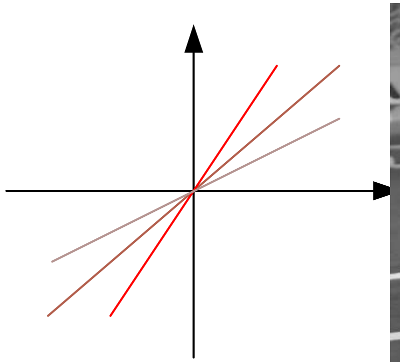

UV transfer function

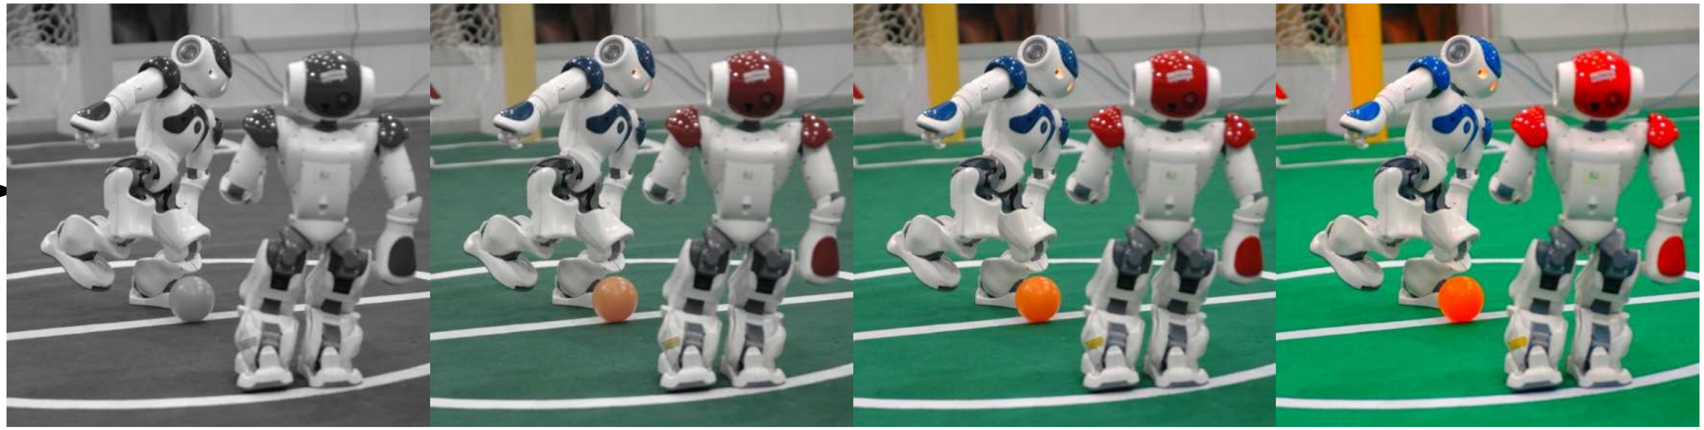

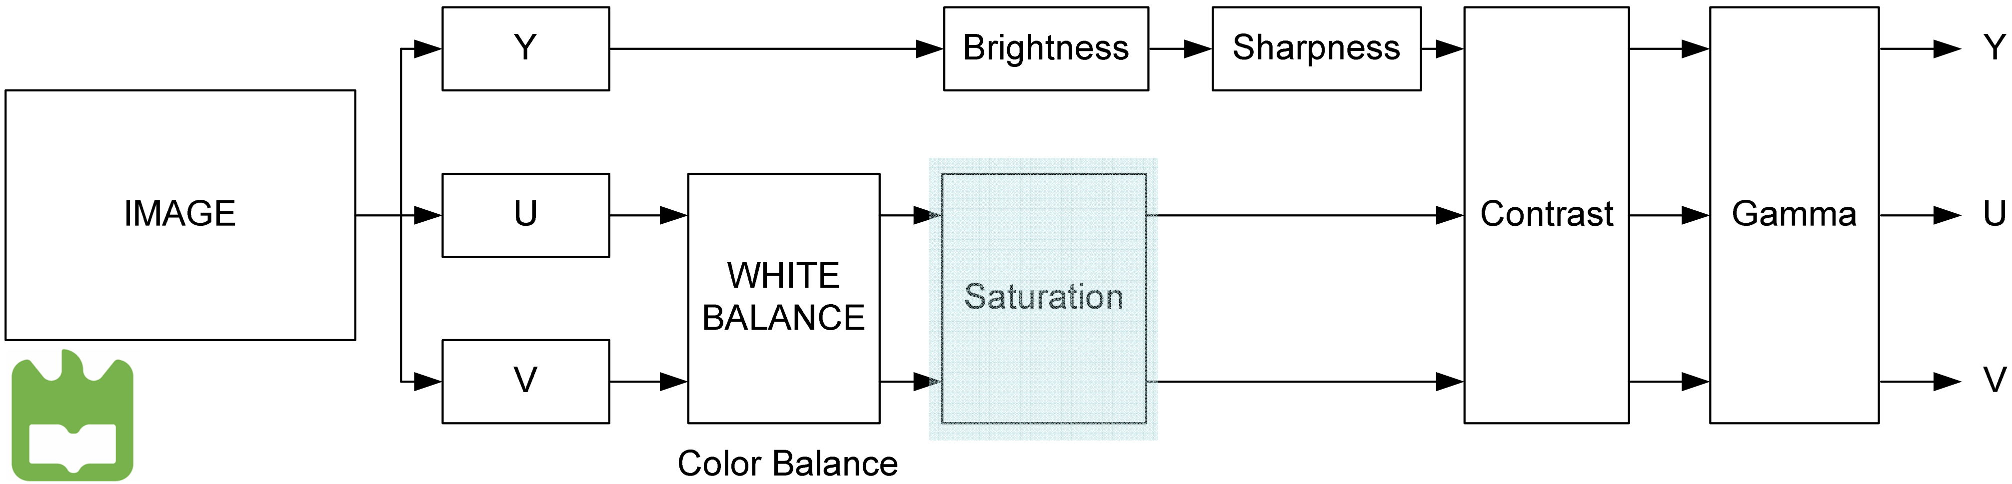

Gamma

Gamma

Gamma correction is the name of a nonlinear operation used to code and decode luminance or RGB tristimulus values. In the simplest cases gamma is defined by the power-law expression:

$$V_{out} = AV_{in}^{\delta}$$

where *A* is a constant and the input and output values are non-negative real values.

In most cases A = 1, and inputs and outputs are typically in the range 0–1.

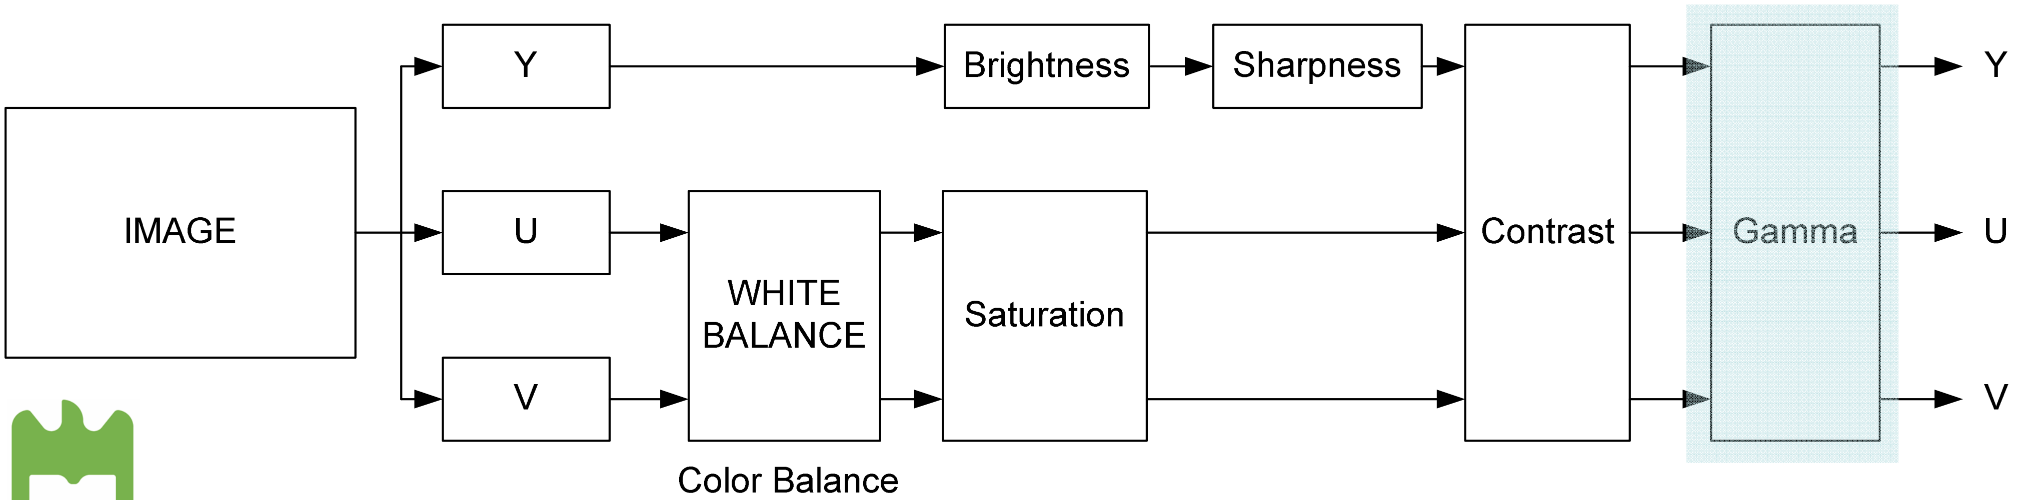

Gamma

Gamma

Examples

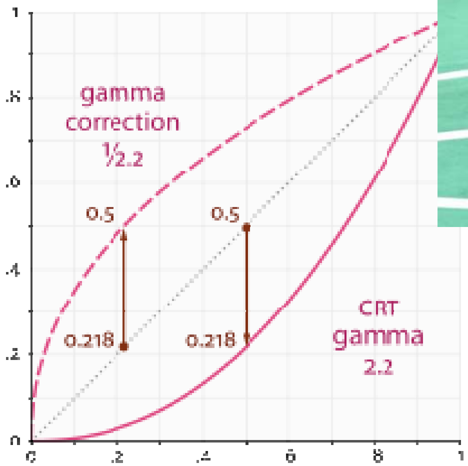

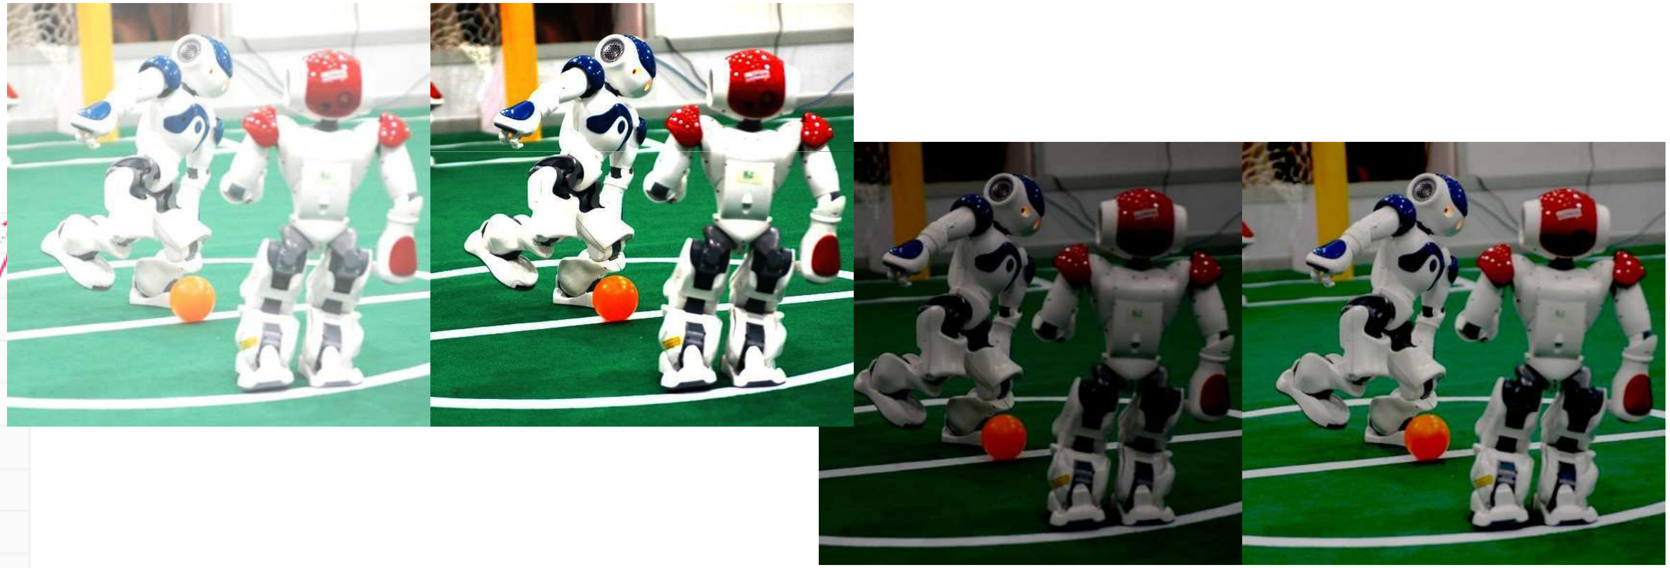

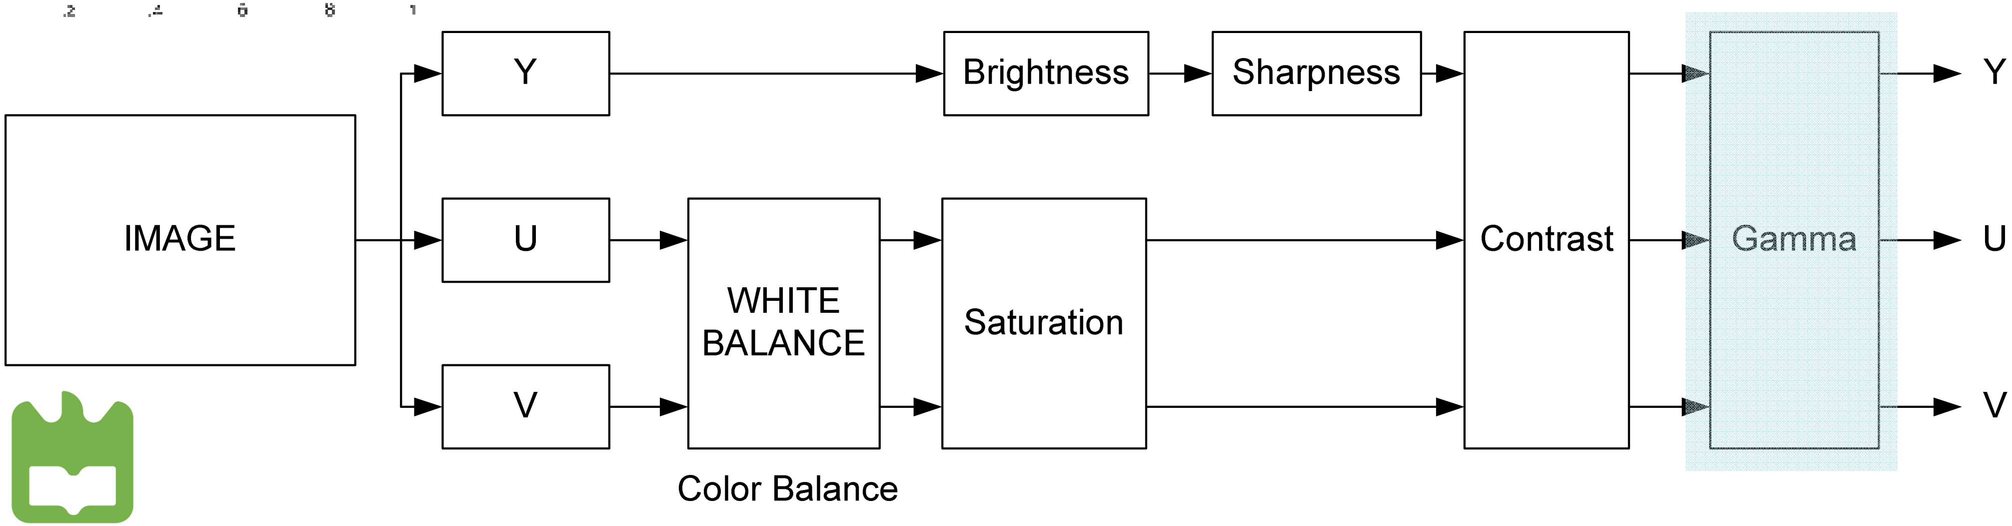

Sharpness

Sharpness (as a controllable parameter)

Sharpness is a measure of the energy frequency spatial distribution over the image.

Not all devices provide access to this parameter.

Sharpness basically allows the control of the cut-off frequency of a low pass spatial filter. This may be very useful if the image is afterward intended to be decimated, since it allows to prevent spatial aliases artifacts.

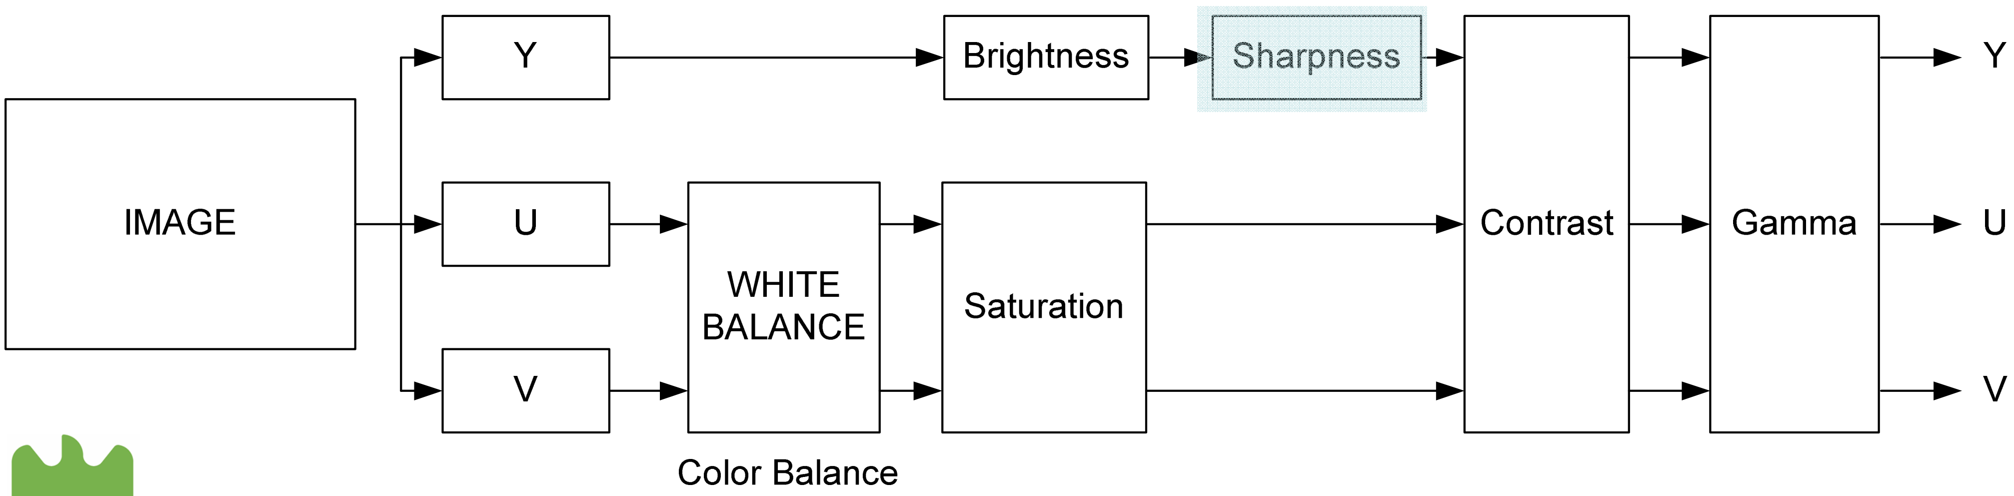

Sharpness

Sharpness (as a controllable parameter)

Examples.

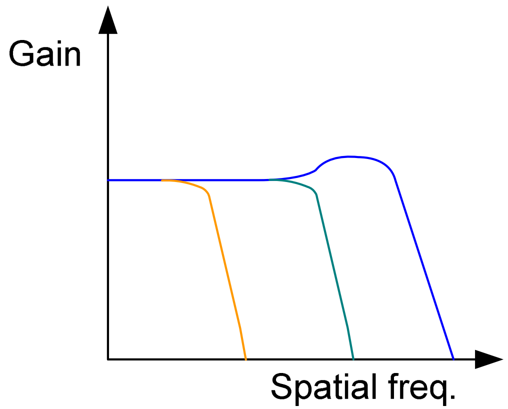

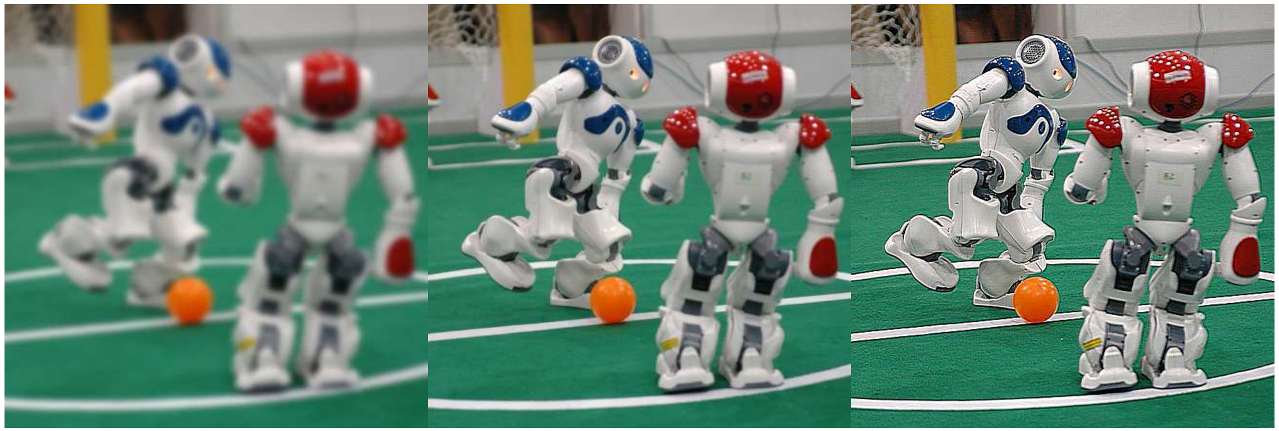

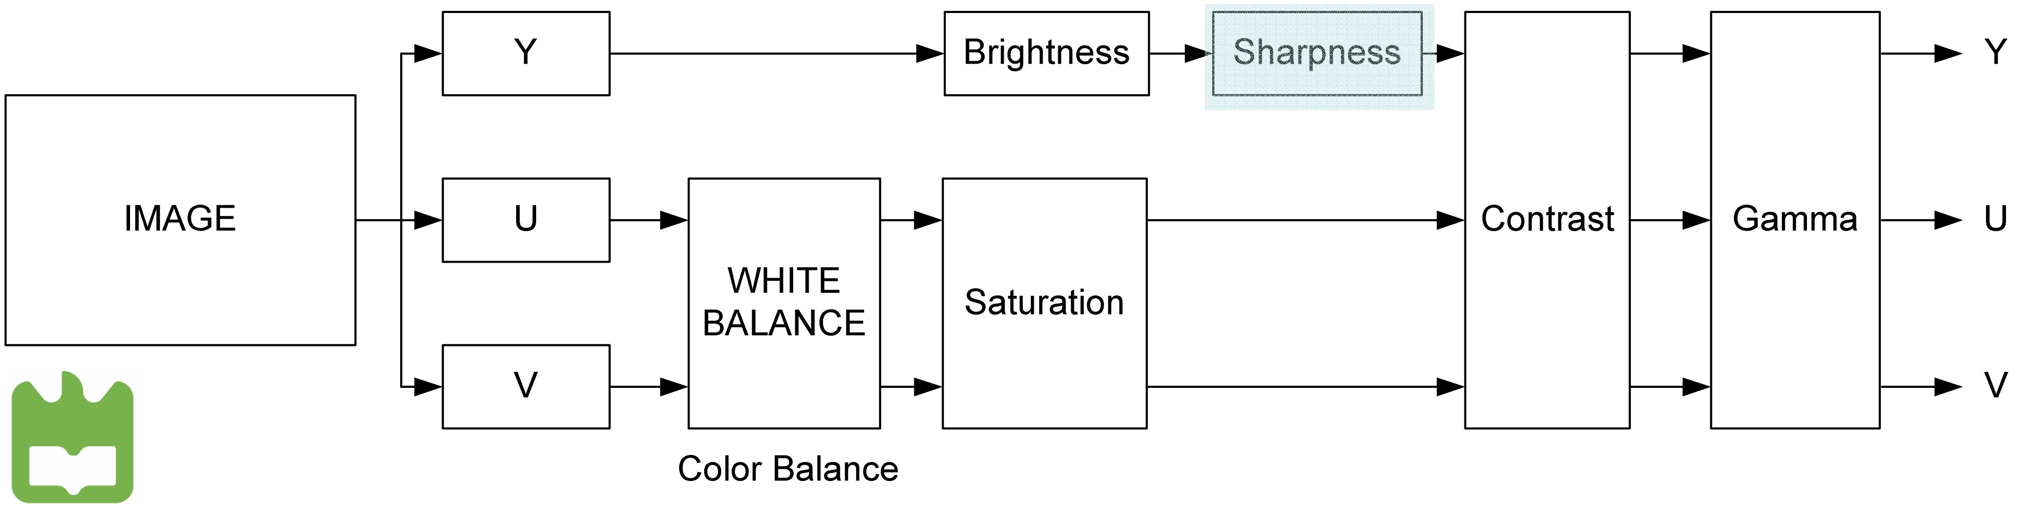

Topic: Sampling and quantization

- Image sensors
- Camera Calibration
- Sampling and quantization
- Data structures for digital images
- Histograms

Components of a Computer Vision System

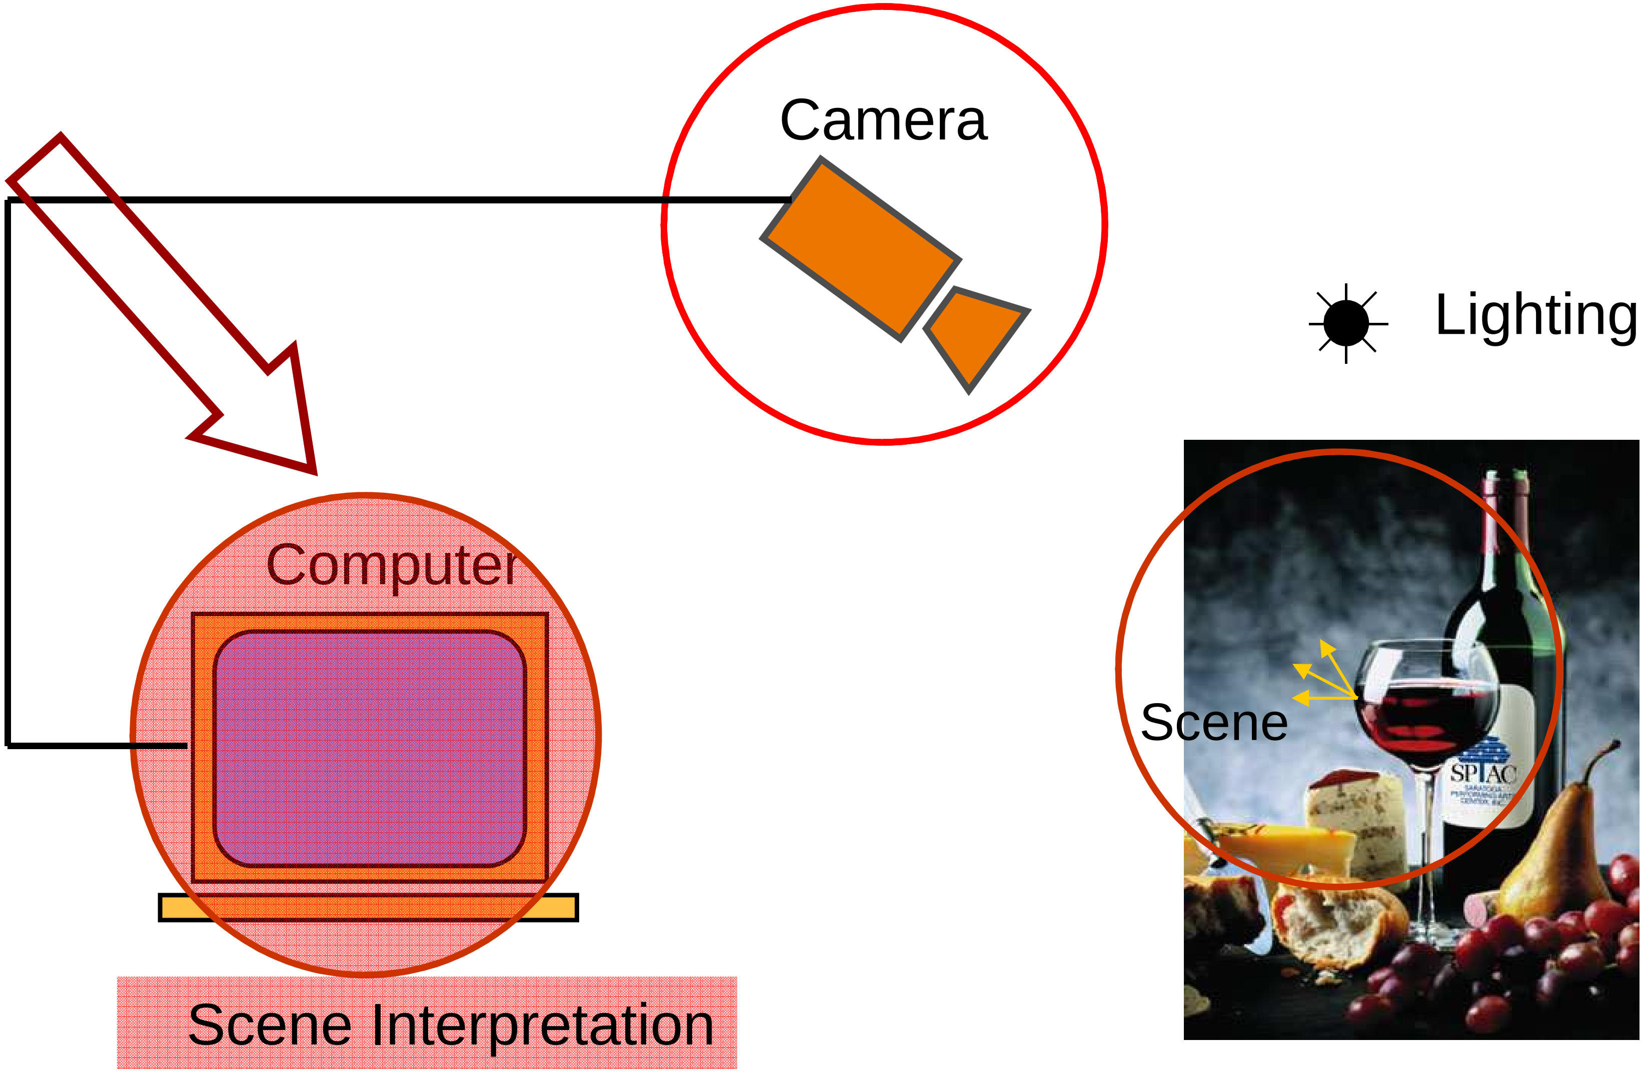

Digital Images

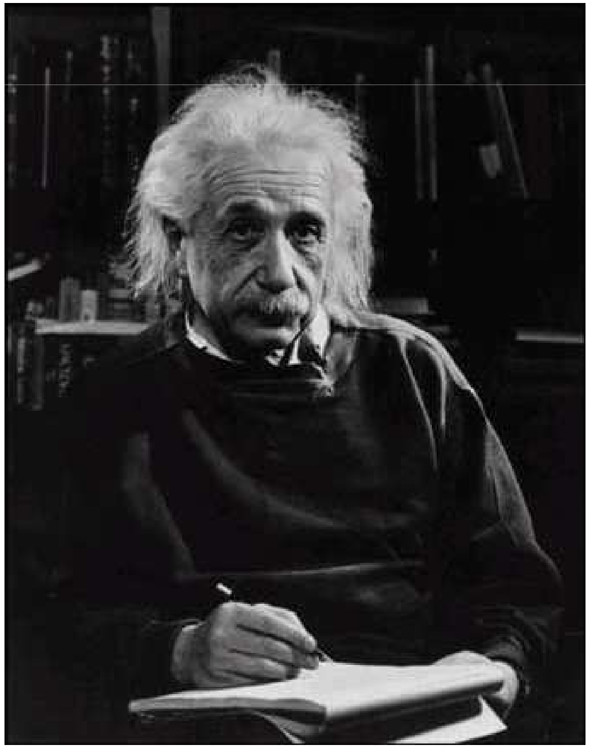

What we see

0	3	2	5	4	7	6	9	8
3	0	1	2	3	4	5	6	7
2	1	0	3	2	5	4	7	6
5	2	3	0	1	2	3	4	5
4	3	2	1	0	3	2	5	4
7	4	5	2	3	0	1	2	3
6	5	4	3	2	1	0	3	2
9	6	7	4	5	2	3	0	1
8	7	6	5	4	3	2	1	0

What a computer sees

Simple Image Model

- Image as a 2D light-intensity function

$$f(x, y)$$

- Continuous
- Non-zero, finite value

$$0 < f(x, y) < \infty$$

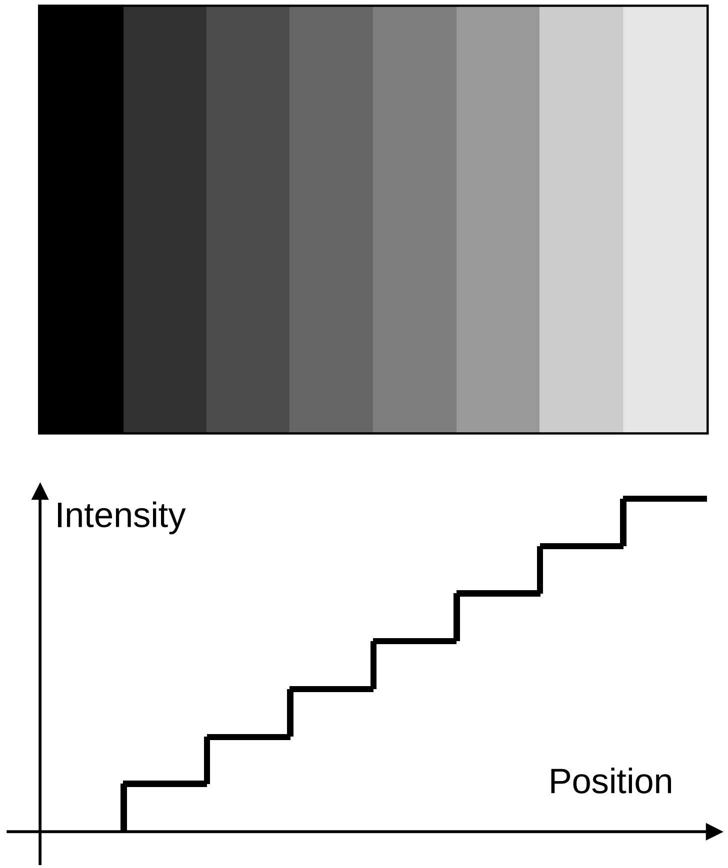

[Gonzalez & Woods]

Analog to Digital

The scene is:

– **projected** on a 2D plane,

– **sampled** on a regular grid, and each sample is

– **quantized** (rounded to the nearest integer)

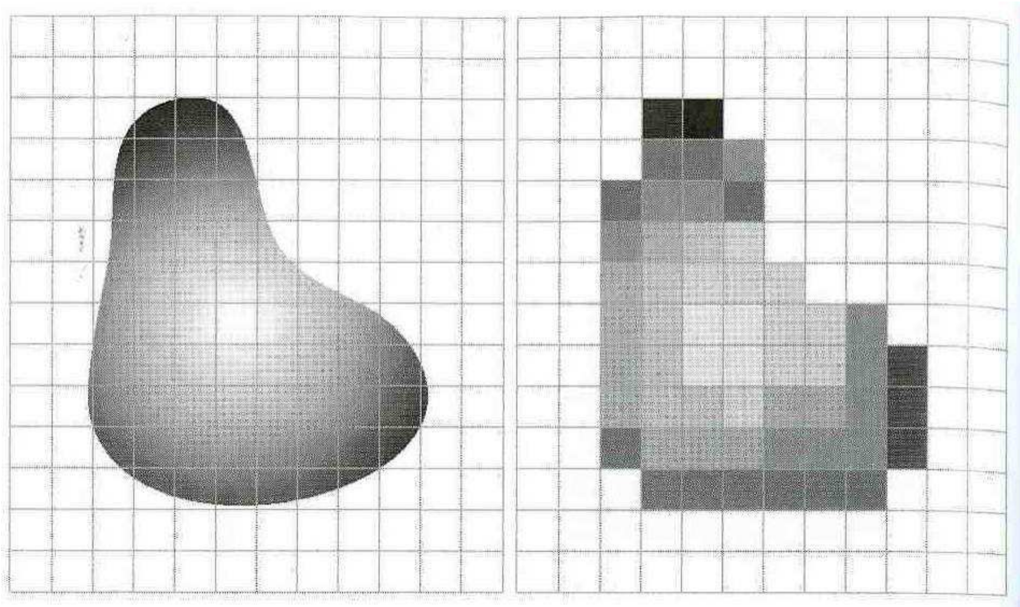

$$f(i,j) = \text{Quantize}\{f(i\Delta, j\Delta)\}$$

Images as Matrices

- Each point is a pixel with amplitude:
  - f(x,y)
- An image is a matrix with size N x M

M = [(0,0) (0,1) …

[(1,0) (1,1) …

…

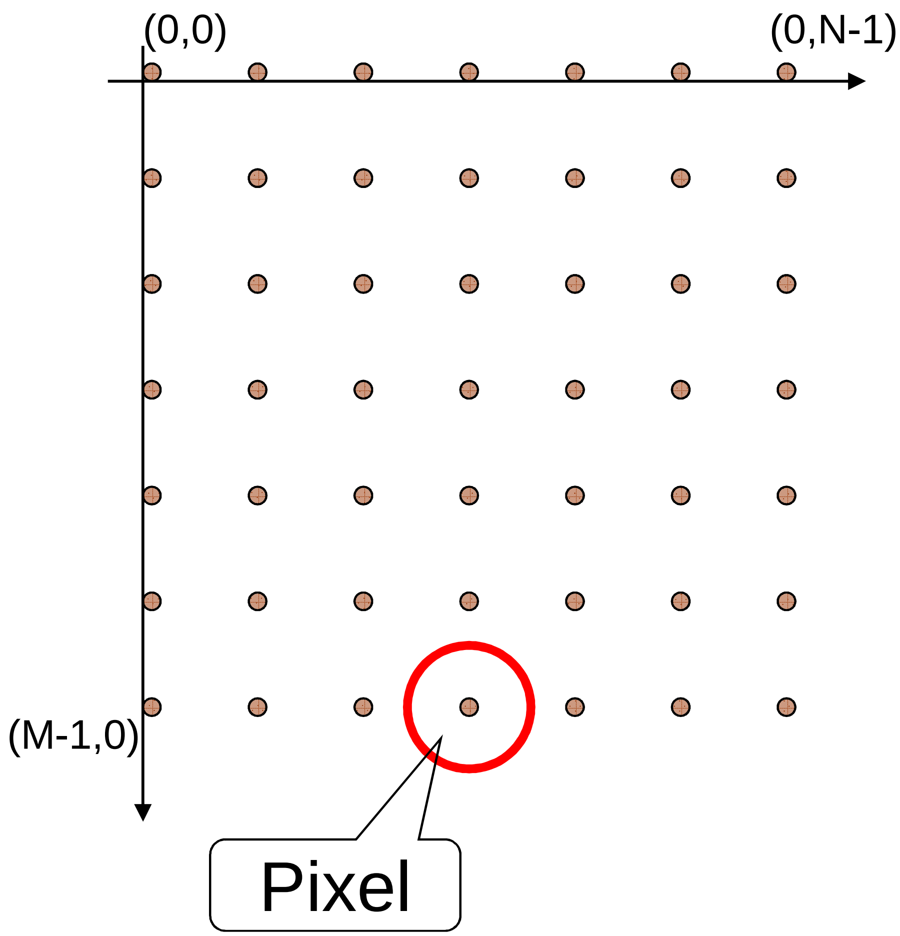

Sampling Theorem

Continuous signal:

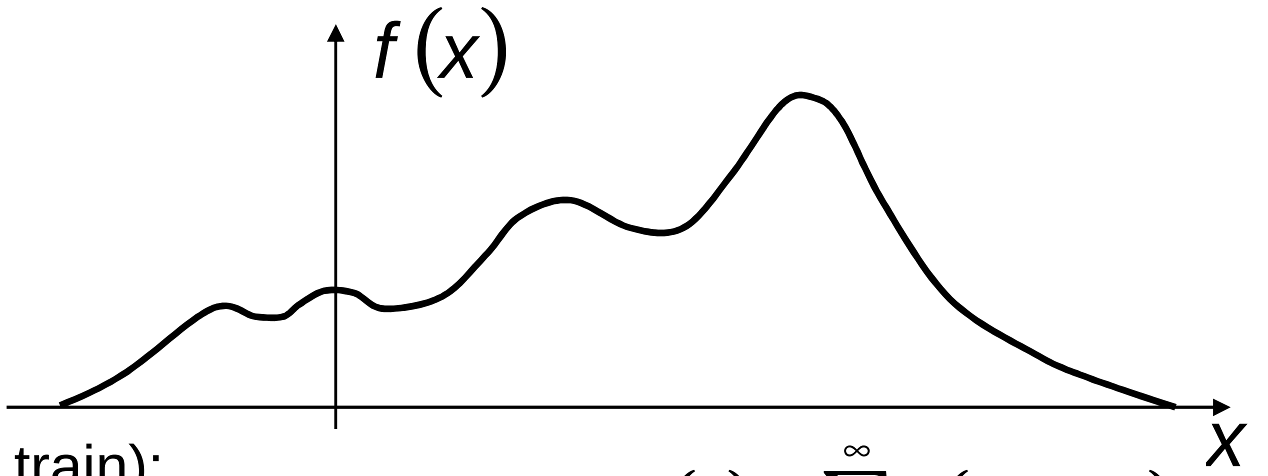

Shah function (Impulse train):

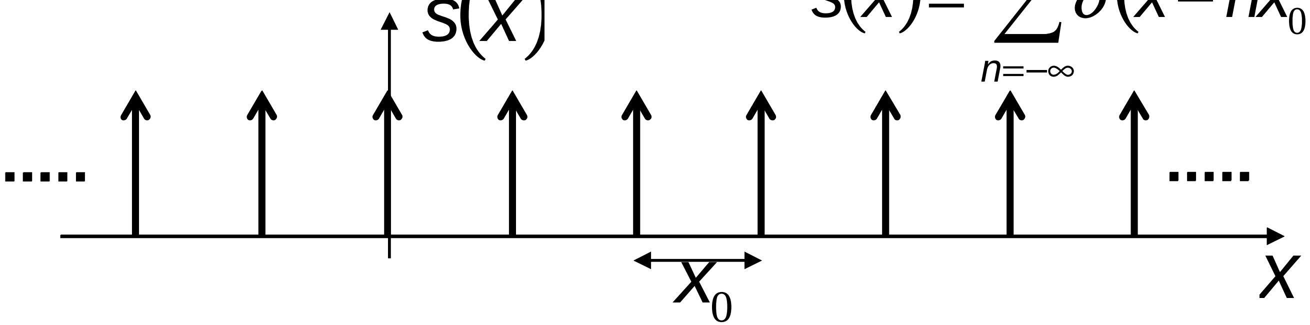

$$s(x) = \sum_{n=-\infty}^{\infty} \delta(x - nx_0)$$

Sampled function:

$$f_s(x) = f(x)s(x) = f(x) \sum_{n=-\infty}^{\infty} \delta(x - nx_0)$$

Quantization

- Analog: $0 < f(x, y) < \infty$
- Digital: Infinite storage space per pixel!
- Quantization

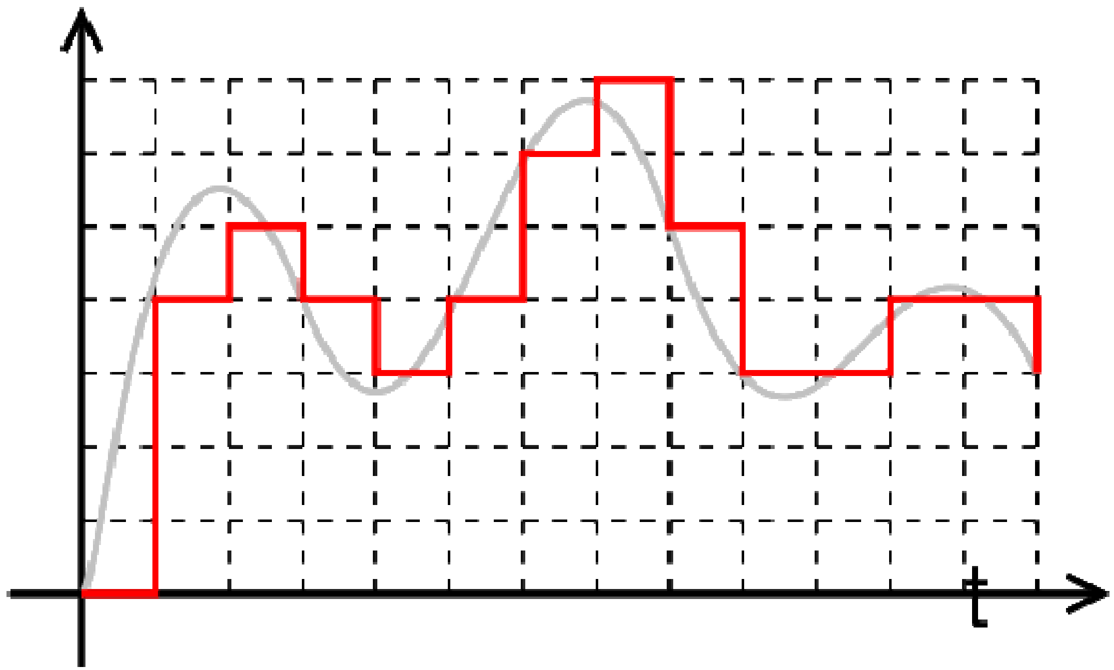

Quantization Levels

- G - number of levels
- m – storage bits
- Round each value to its nearest level

$$G = 2^m$$

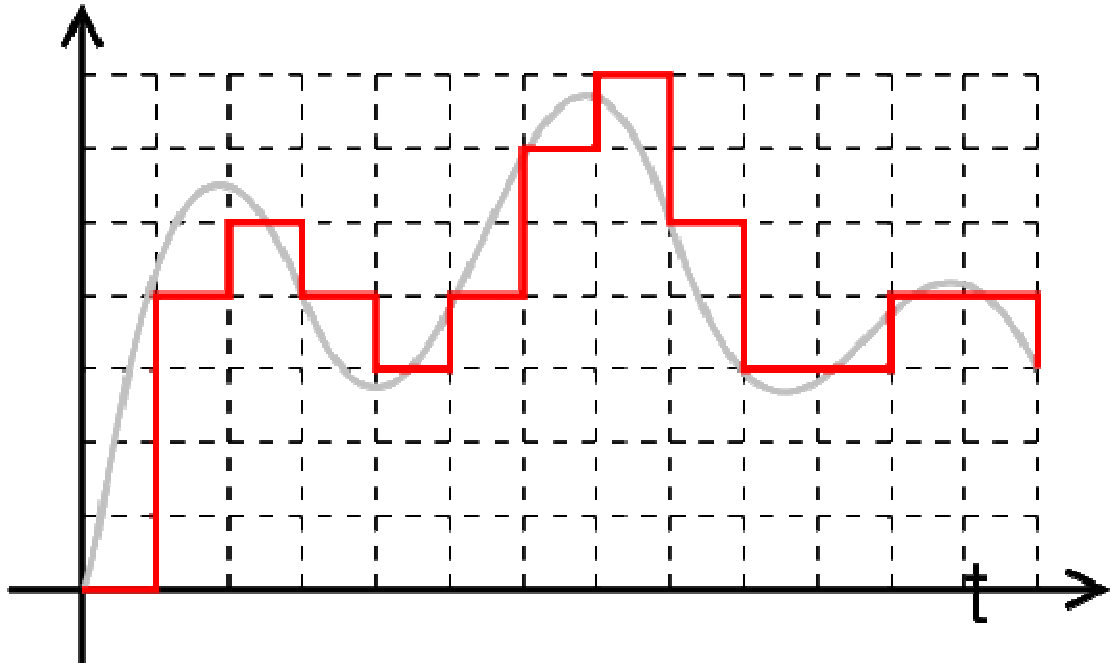

Effect of quantization

Effect of quantization

Image Size

- Storage space
  - Spatial resolution: N x M
  - Quantization: m bits per pixel
  - Required bits b:

$$b = N \times M \times m$$

- Rule of thumb:
  - More storage space means more image quality

Image Scaling

This image is too big to fit on the screen. How can we reduce it?

How to generate a half-sized version?

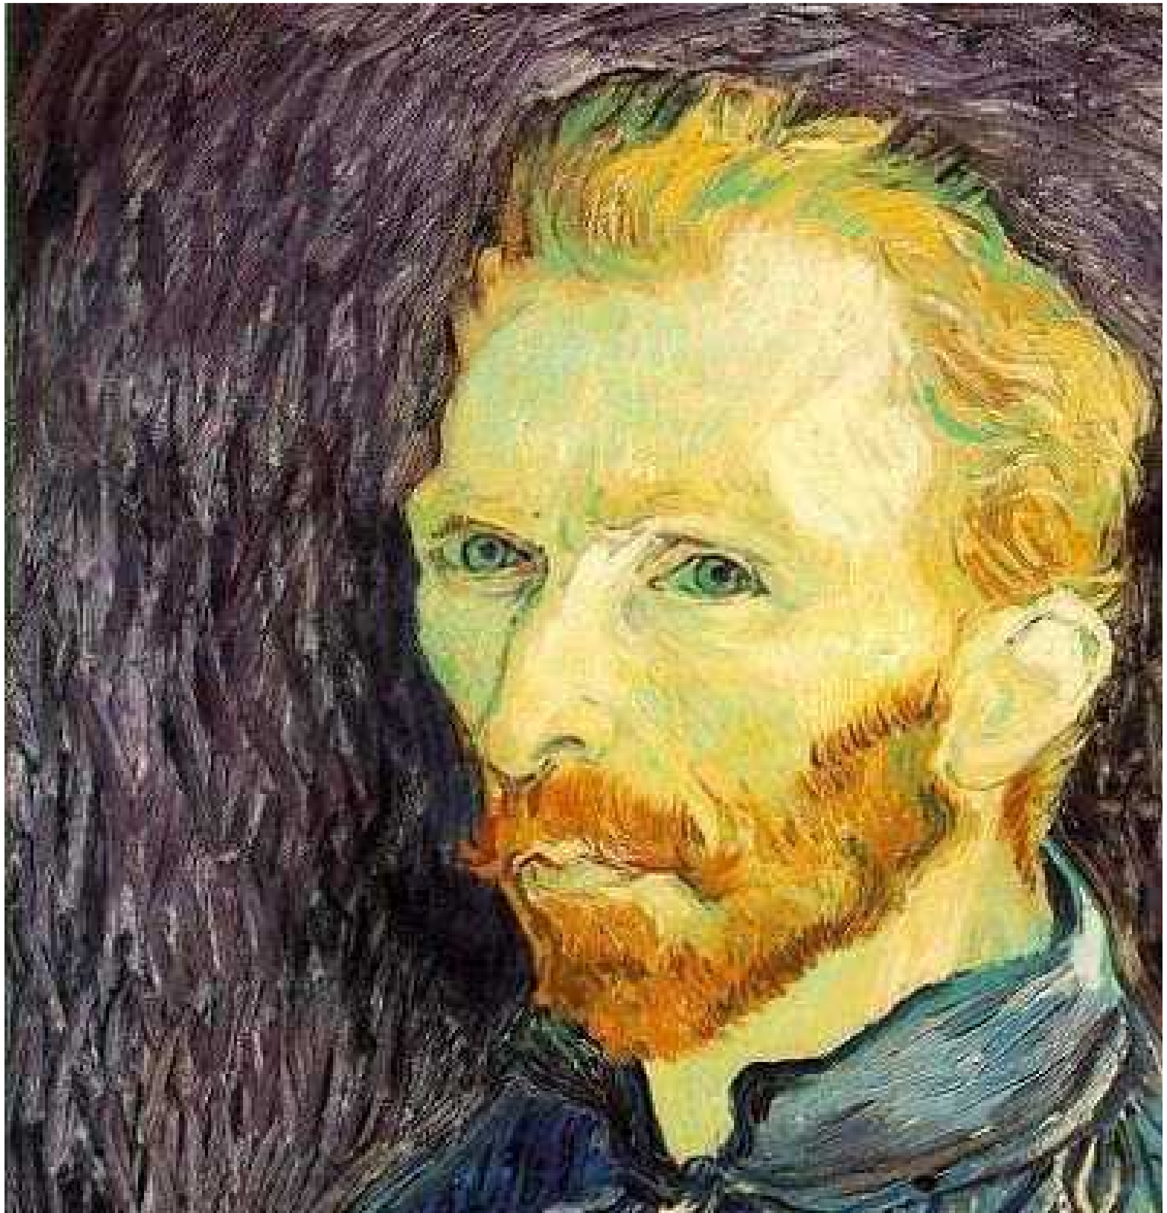

Sub-sampling

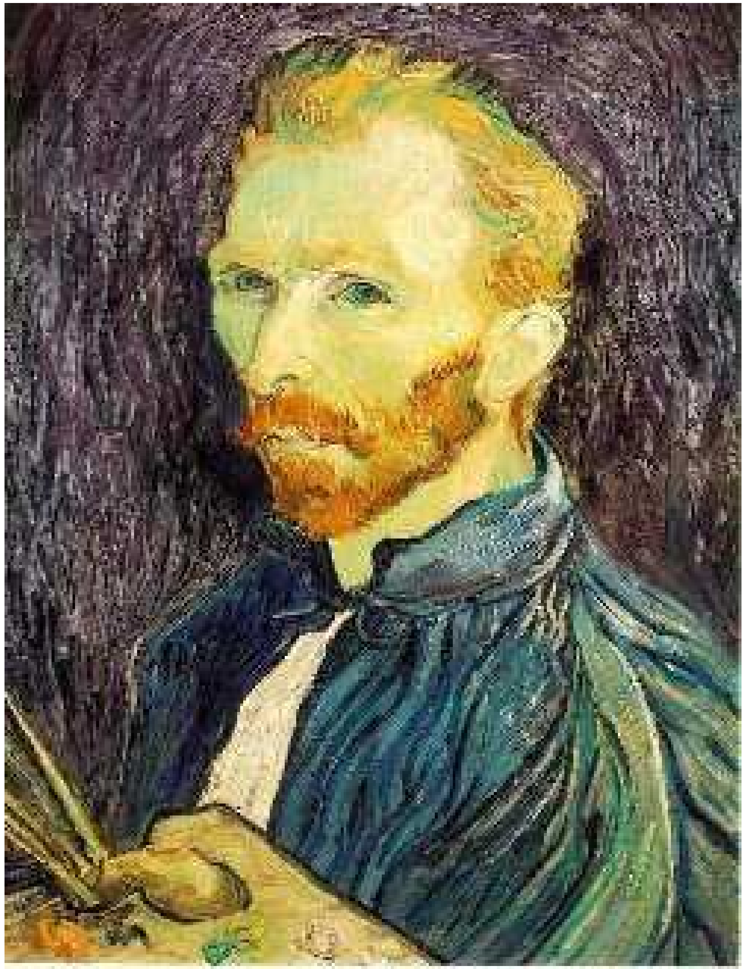

1/4

1/8

Throw away every other row and column to create a 1/2 size image
- called *image sub-sampling*

Sub-sampling

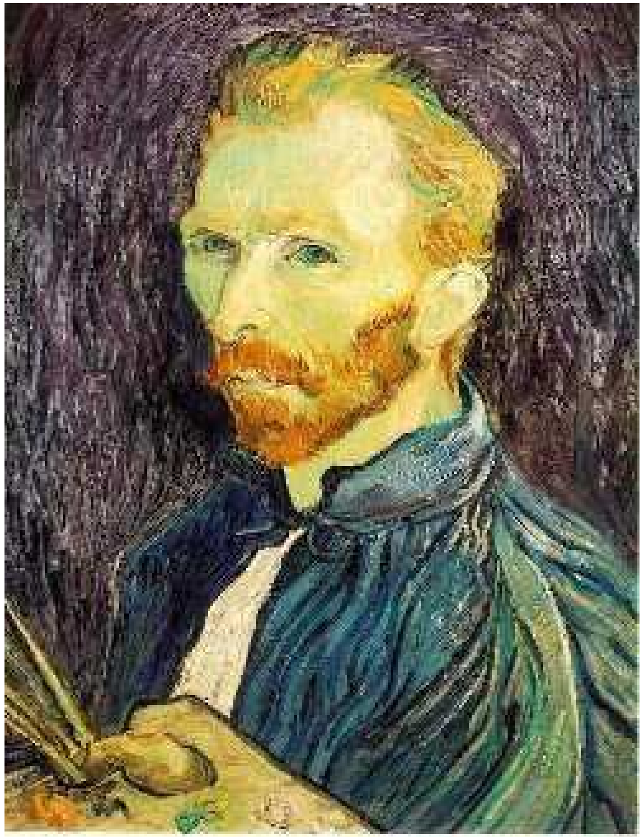

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Sub-Sampling with Gaussian Pre-Filtering

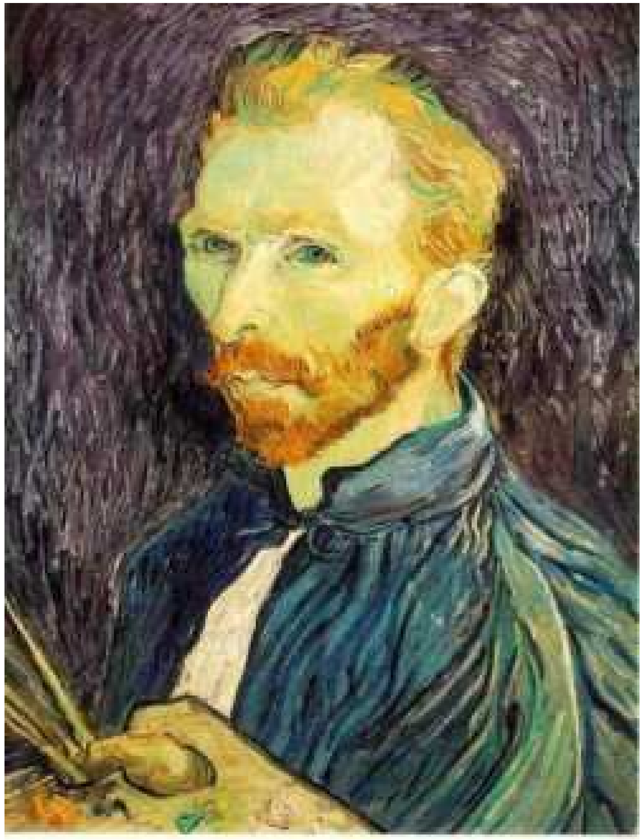

Gaussian 1/2

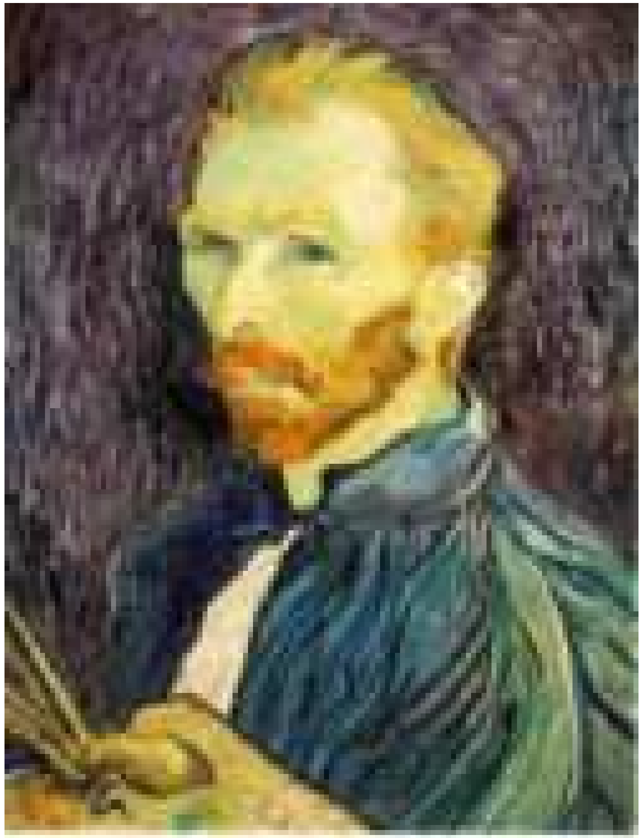

G 1/4

G 1/8

Compare with...

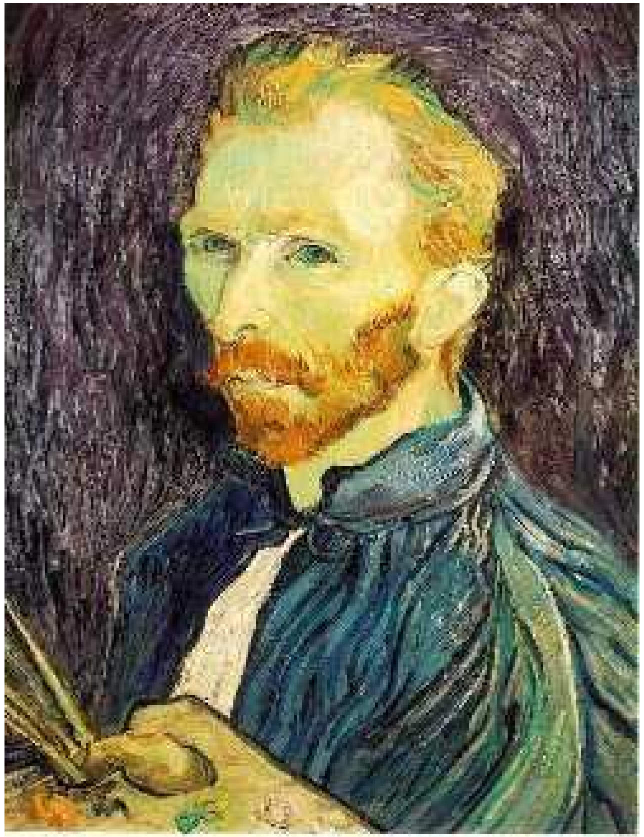

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Topic: Data structures for digital images

- Image sensors
- Sampling and quantization
- Data structures for digital images
- Histograms

Data Structures for Digital Images

- Are there other ways to represent digital images?

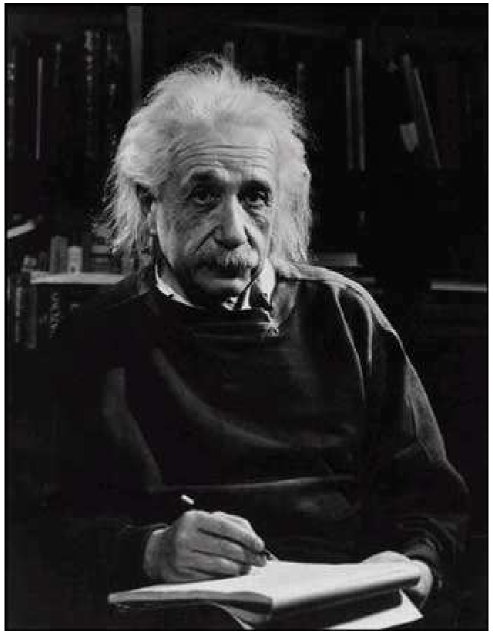

What we see

0	3	2	5	4	7	6	9	8
3	0	1	2	3	4	5	6	7
2	1	0	3	2	5	4	7	6
5	2	3	0	1	2	3	4	5
4	3	2	1	0	3	2	5	4
7	4	5	2	3	0	1	2	3
6	5	4	3	2	1	0	3	2
9	6	7	4	5	2	3	0	1
8	7	6	5	4	3	2	1	0

What a computer sees

Chain codes

- Chains represent the borders of objects.
- Coding with *chain codes*.
  - Relative.
  - Assume an initial starting point for each object.
- Needs segmentation!

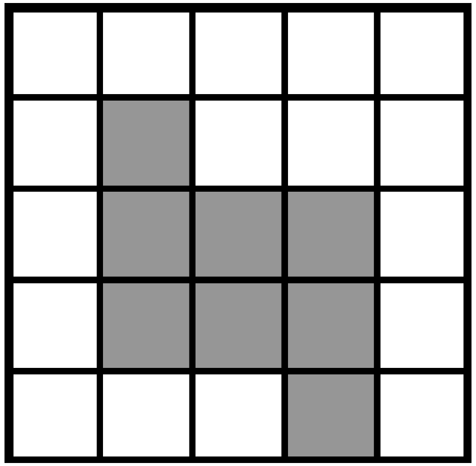

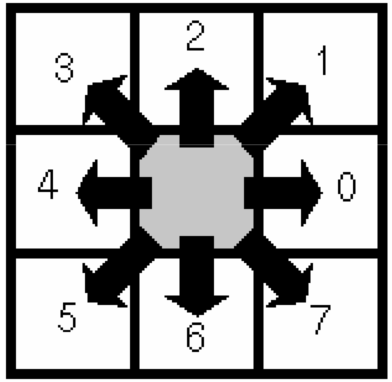

Freeman Chain Code

Using a Freeman Chain Code and considering the top-left pixel of the image as the starting point: 70663422

Topological Data Structures

- *Region Adjacency Graph*
  - **Nodes** - Regions
  - **Arcs** – Relationships
- Describes the elements of an image and their spatial relationships.
- Needs segmentation!

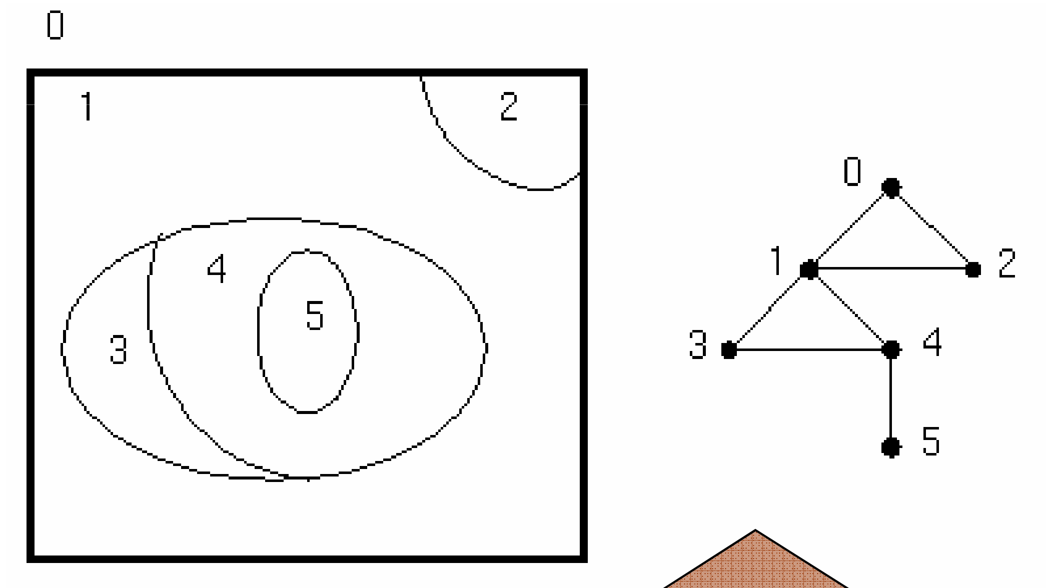

Region Adjacency Graph

Relational Structures

- Stores relations between objects.
- Important semantic information of an image.
- Needs segmentation and an image description (features)!

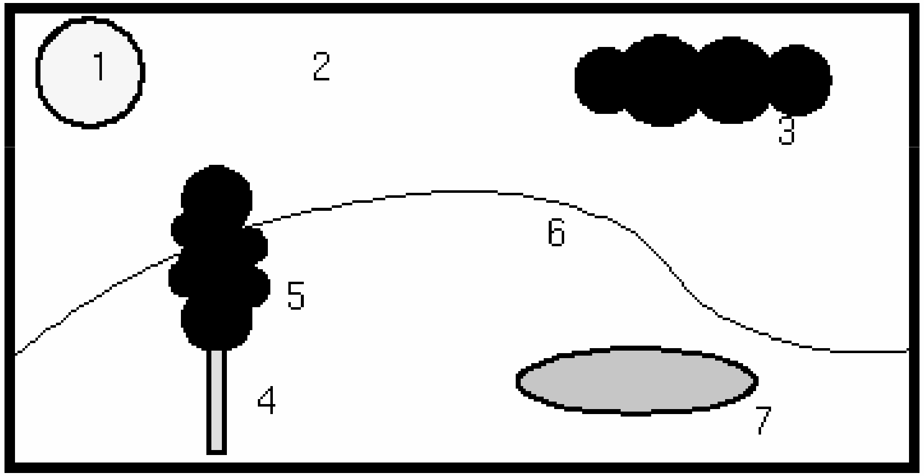

No.	Object name	Colour	Min. row	Min. col.	Inside
1	sun	white	5	40	2
2	sky	blue	0	0	-
3	cloud	grey	20	180	2
4	tree trunk	brown	95	75	6
5	tree crown	green	53	63	-
6	hill	light green	97	0	-
7	pond	blue	100	160	6

Relational Table

Topic: Histograms

- Image sensors
- Sampling and quantization
- Data structures for digital images
- Histograms

Histograms

- In statistics, a histogram is a graphical display of tabulated frequencies.

Typically represented as a bar chart:

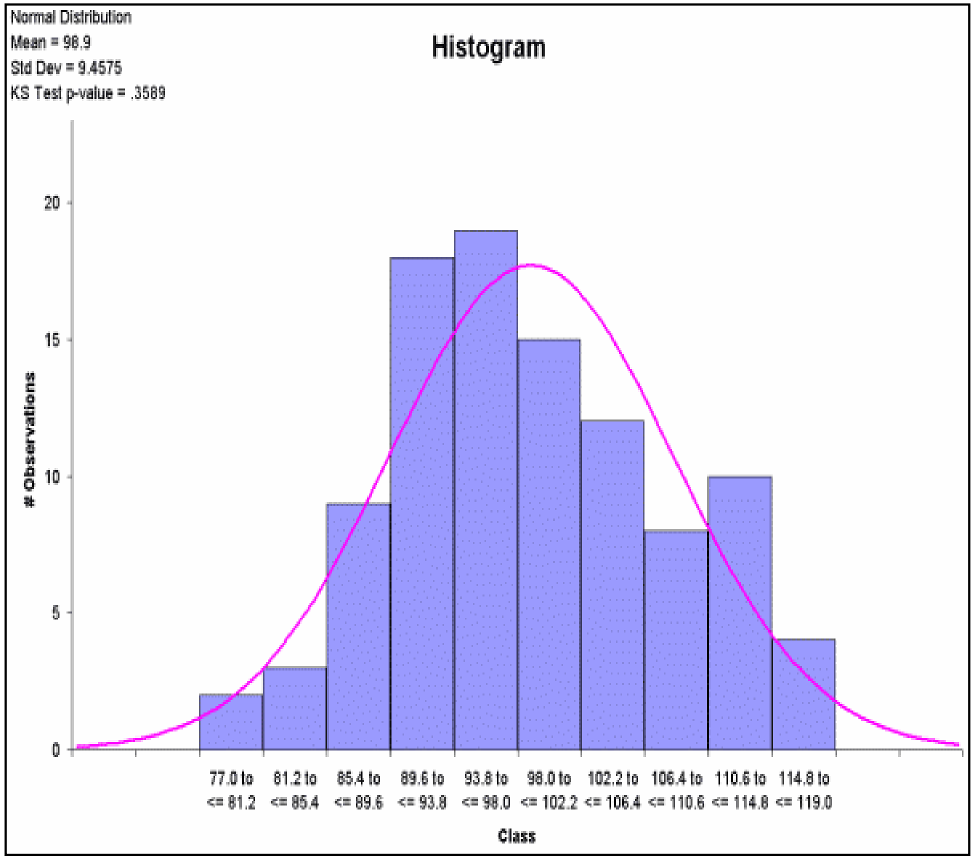

Image Histograms

- Colour or Intensity distribution.
- Typically:
  - Reduced number of *bins*.
  - Normalization.
- Compressed representation of an image.
  - No spatial information whatsoever!

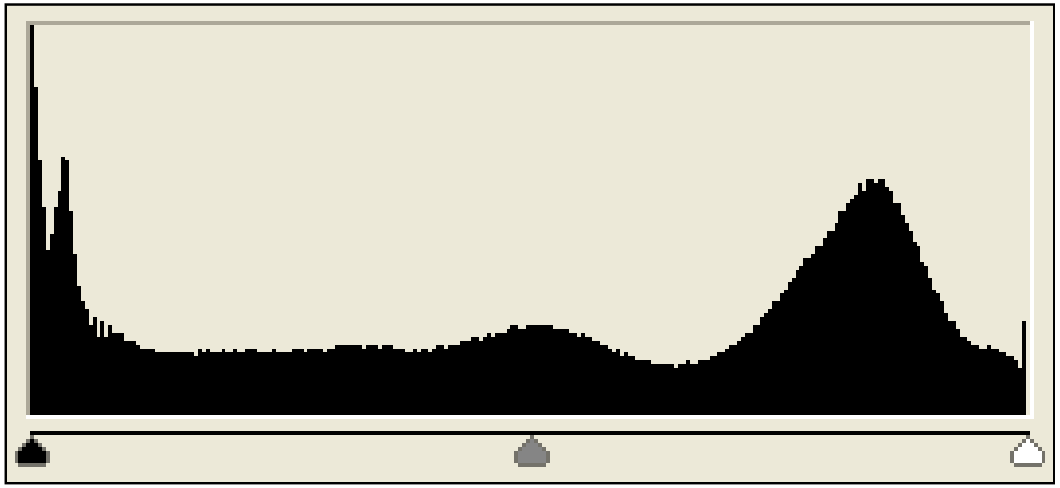

Histogram Normalization

- Improves the contrast in an image in order to stretch out the intensity range.
- The goal is to reshape the image histogram to make it flat and wide.

Color Histogram

- As many histograms as axis of the color space.

  Ex: RGB Colour space

  - Red Histogram
  - Green Histogram
  - Blue Histogram

- Combined histogram.

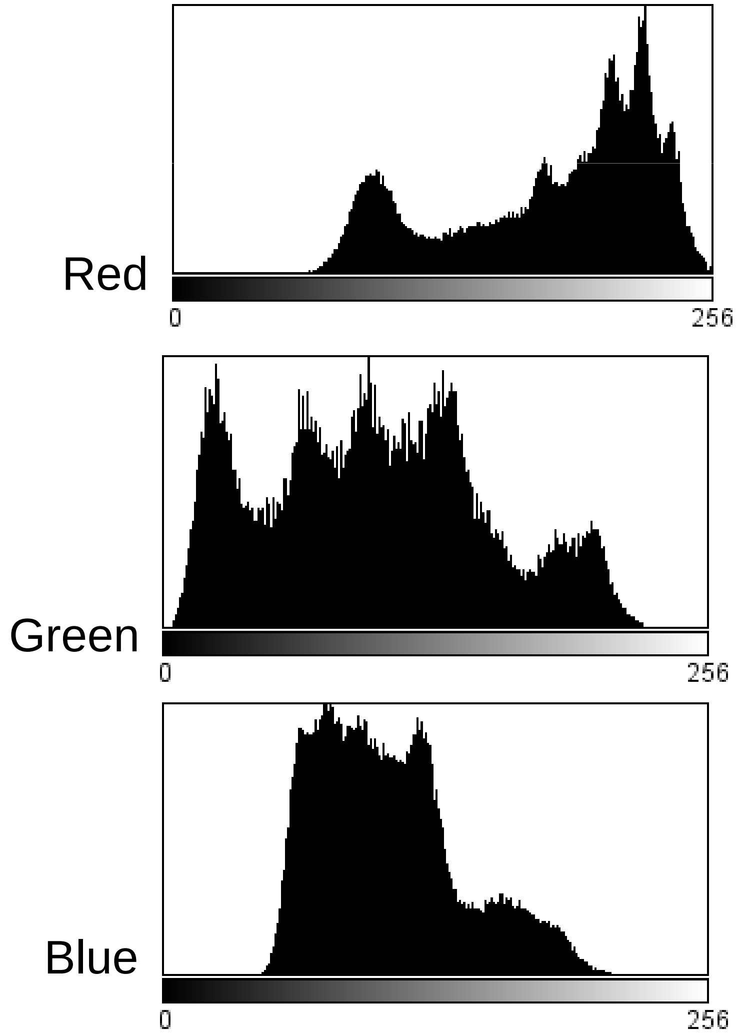

Resources

- J.C. Russ – Chapters 2
- R. Gonzalez, and R. Woods – Chapter 2