Computer Vision

Doctoral Program in Computer Science (MAPi)

Hélder Filipe Pinto de Oliveira

Faculdade de Ciências da Universidade do Porto

Departamento de Ciência de Computadores

DOCTORAL PROGRAMME
IN COMPUTER SCIENCE

Teacher Background

Academic Background

- BSc (Licenciatura) - Eng. Eletrotécnica e de Computadores, FEUP (1999 – 2004)
- MSc - Automação, Instrumentação e Controlo, FEUP (2005 – 2008)
  - Robotics
- PhD - Eng. Eletrotécnica e de Computadores, FEUP (2008 – 2013)
  - Medical Image Analysis (Breast Cancer)

Teacher Background

Faculty Experience

- Faculdade de Ciências da Universidade do Porto
  - Departamento de Ciência de Computadores
  - Invited Assistant Professor (2016 – …)
- Faculdade de Engenharia da Universidade do Porto
  - Departamento de Engenharia Informática
  - Invited Assistant Professor (2014 – 2016)
  - Invited Assistant (2008 – 2011)

Teacher Background

Research Experience

- INESC TEC (2008 - …)
- Politecnico di Milano (april – june 2011)
- FEUP (2004 – 2005) (2007 – 2008)
- ISR-Porto (2005 – 2007)

- Projects
  - BCCT.plan (PT 2020) – Breast Cancer (2016 – 2019)
  - NanoSTIMA (PT 2020) – Medical Image Analysis (2015 – 2018)
  - PICTURE (UE) – Breast Cancer (2013 – 2016)
  - 3dBCT (FCT) – Breast Cancer (2011 – 2014)
  - 5DPO (FCT) – Robotics (2005 – 2007)
  - IVLAB (ADI) – Cork stoppers image analysis (2004 – 2005)

Teacher Background

<u>VISUM Summer School</u>

- http://visum.inesctec.pt/

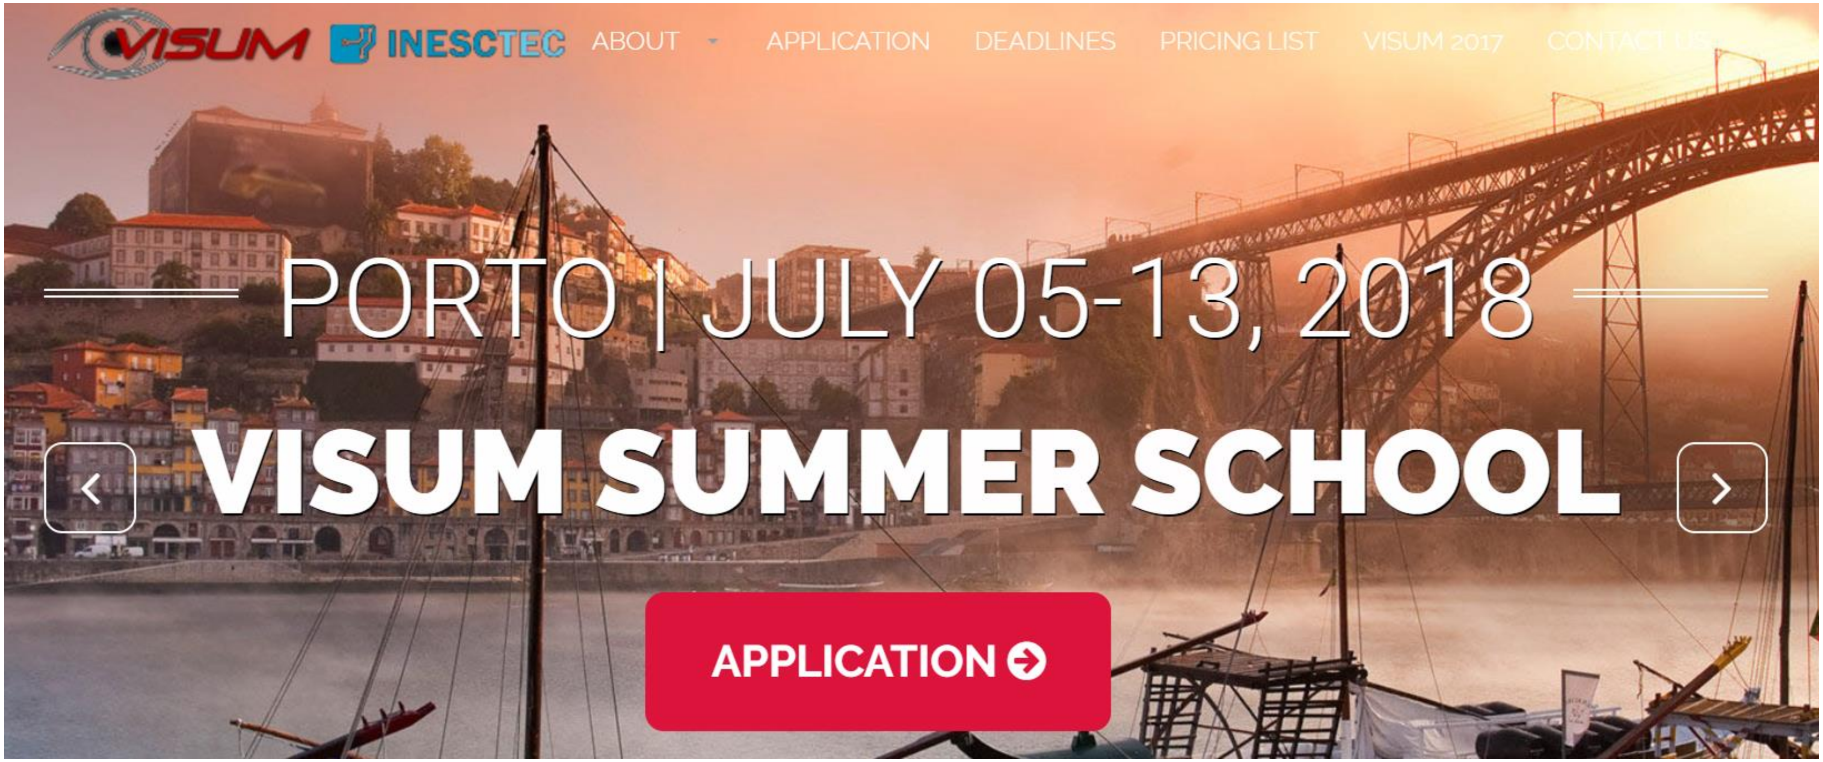

Teacher Background

Personal webpage

- http://www.inescporto.pt/~hfpo/

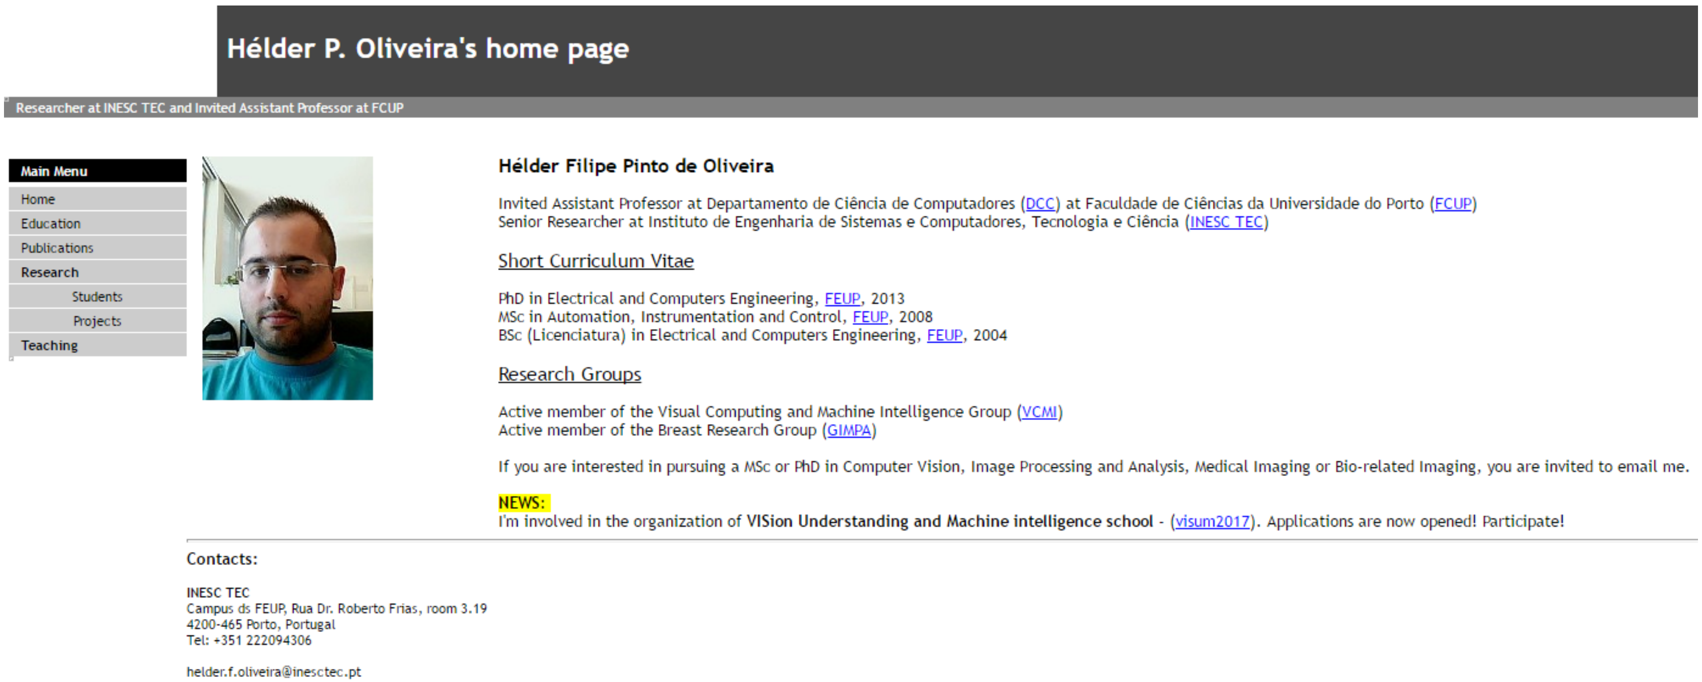

Hélder P. Oliveira's home page

Researcher at INESC TEC and Invited Assistant Professor at FCUP

Main Menu

- Home
- Education
- Publications
- Research
  - Students
  - Projects
- Teaching

Hélder Filipe Pinto de Oliveira

Invited Assistant Professor at Departamento de Ciência de Computadores (DCC) at Faculdade de Ciências da Universidade do Porto (FCUP)
Senior Researcher at Instituto de Engenharia de Sistemas e Computadores, Tecnologia e Ciência (INESC TEC)

Short Curriculum Vitae

PhD in Electrical and Computers Engineering, FEUP, 2013
MSc in Automation, Instrumentation and Control, FEUP, 2008
BSc (Licenciatura) in Electrical and Computers Engineering, FEUP, 2004

Research Groups

Active member of the Visual Computing and Machine Intelligence Group (VCMI)
Active member of the Breast Research Group (GIMPA)

If you are interested in pursuing a MSc or PhD in Computer Vision, Image Processing and Analysis, Medical Imaging or Bio-related Imaging, you are invited to email me.

NEWS:
I'm involved in the organization of **VISion Understanding and Machine intelligence school** - (visum2017). Applications are now opened! Participate!

Contacts:

INESC TEC
Campus ds FEUP, Rua Dr. Roberto Frias, room 3.19
4200-465 Porto, Portugal
Tel: +351 222094306

helder.f.oliveira@inesctec.pt

Motivation

Every image tells a story...

Motivation

- Goal of computer vision: perceive the "story" behind the picture.

- Compute properties of the world: 3D shape, names of people or objects, what happened?

Can the computer match human perception?

- Yes and no (mainly no): computers can be better at easy things; humans are much better at hard things.

- But huge progress has been made in the last years: what is considered hard keeps changing.

Compute 3D shape of the world

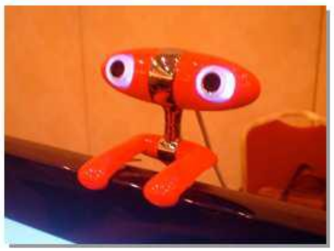

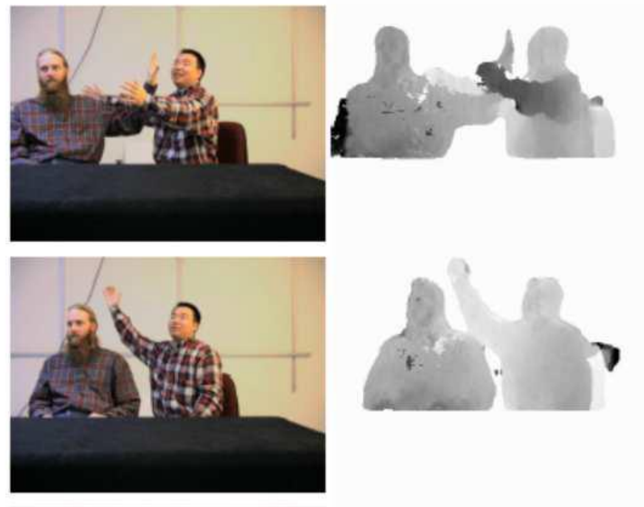

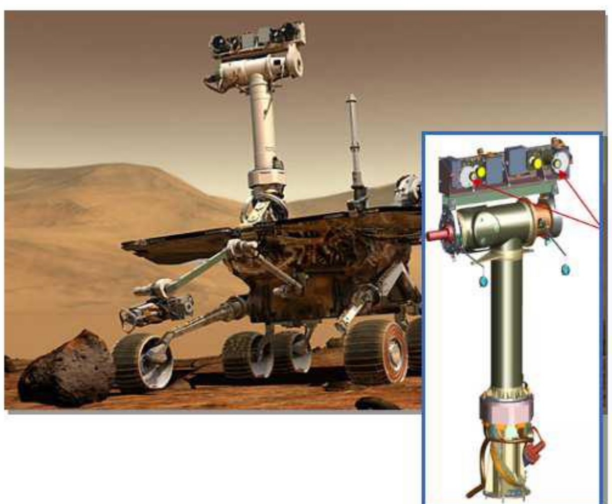

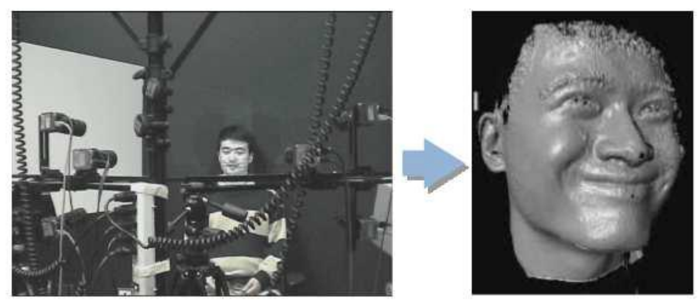

Compute 3D shape of the world

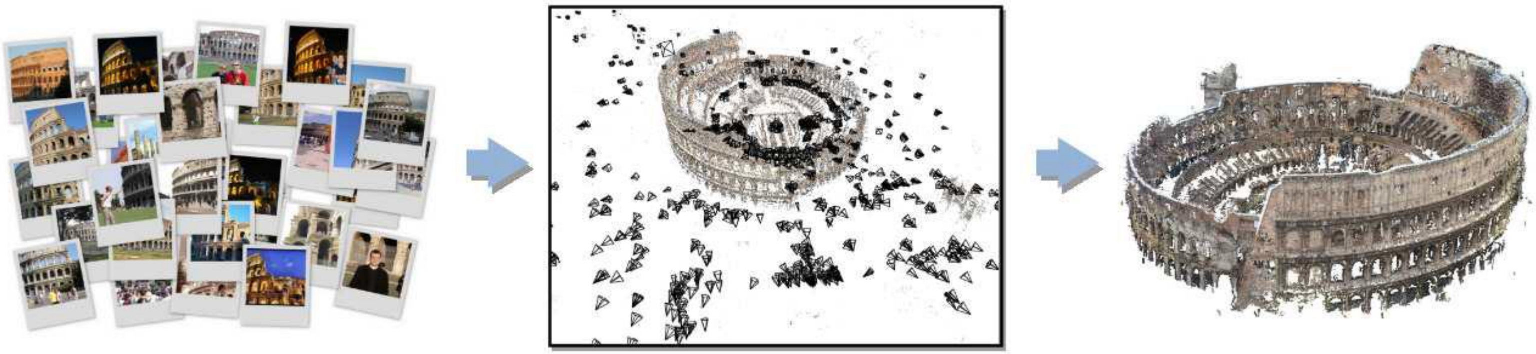

Internet Photos ("Colosseum") | Reconstructed 3D cameras and points | Dense 3D model

Optical character recognition (OCR)

Digit recognition, AT&T labs

License plate readers

Automatic check processing

Sudoku grabber

Forensics

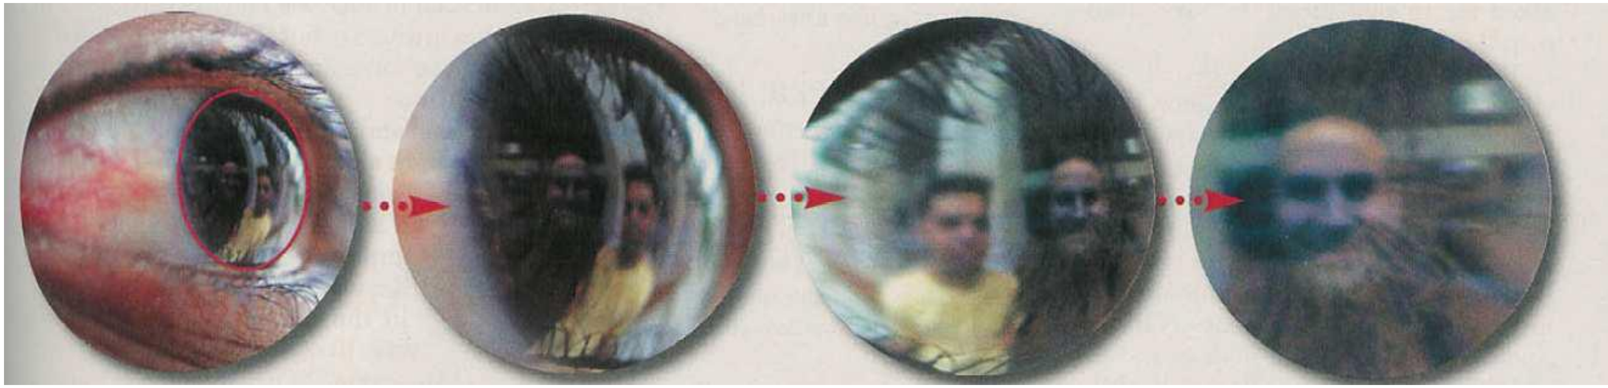

Face detection / recognition

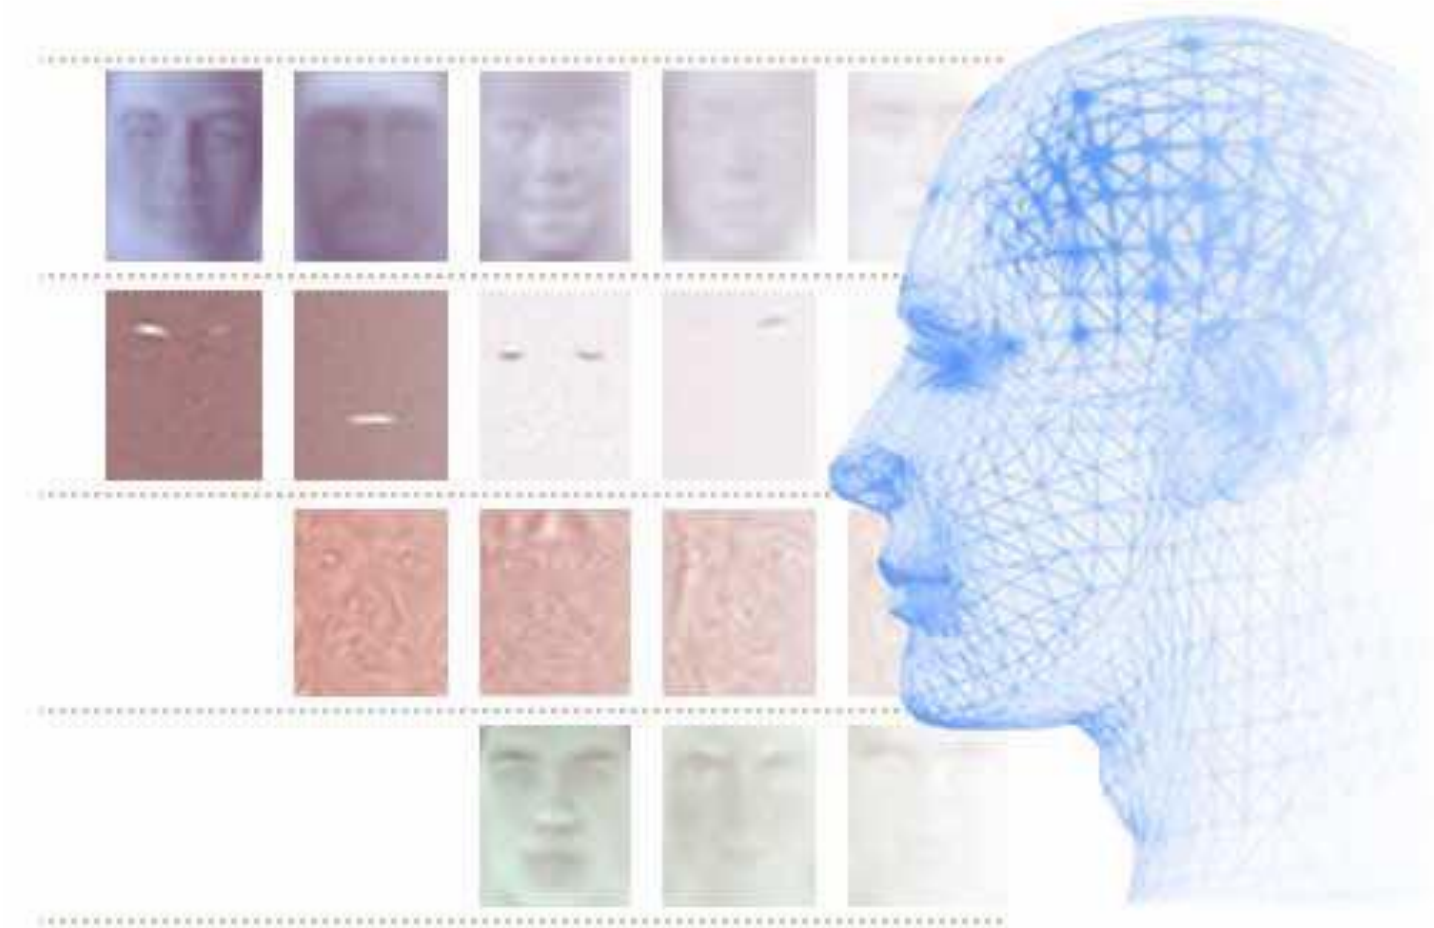

http://www.face-rec.org/

Medical image analysis

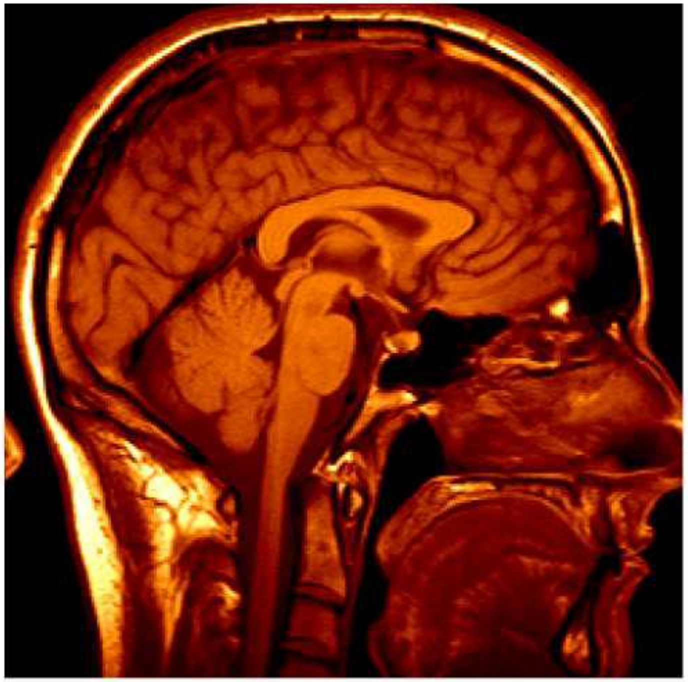

3D imaging
MRI, CT

Image guided surgery
Grimson et al., MIT

Space

The Heights of Mount Sharp
http://www.nasa.gov/mission_pages/msl/multimedia/pia16077.html
Panorama captured by Curiosity Rover, August 18, 2012 (Sol 12)

Vision systems (JPL) uses for several tasks

- Panorama stitching
- 3D terrain modeling
- Obstacle detection, position tracking

Smart Cars

- Mobileye
  - Vision systems currently in high-end models

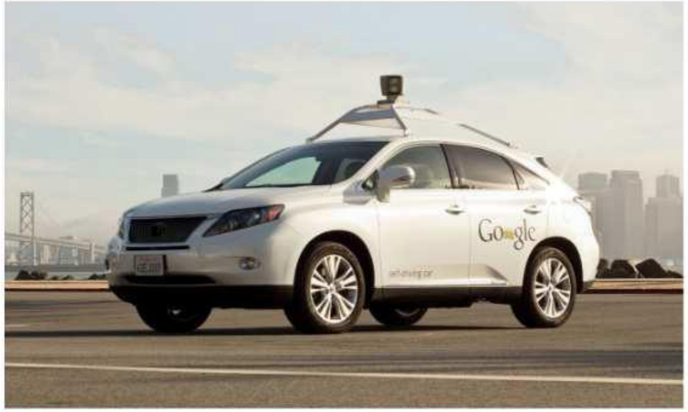

Vision based interaction

Ex: camera-based IR tracking.

Assistive technologies

Sports

OAK 7 4TH
TEN 14 11:04

Day 5: Swimming - Men's 4X200M Final
UNITED STATES
WR Split 5:41.33 -3.32
650M
5:38.01
07:13 / 14:25
Highlights of the men's 4x200m relay final on Day 5.

Shape and motion capture

Object detection

Vision as sensor in robotics

Human Vision

- Vision is a complex physical and intellectual human task that stands as a primary interaction tool with the world.
- It is a complex process not completely understood, even after hundreds of years of research.
- The visualization of a physical process involves an almost simultaneous interaction of the eyes and the brain.
- This interaction is performed by a network of neurons, receptors and other specialized cells.

Human Vision

- The human eye is equipped with a variety of optical elements, including the cornea, iris, pupil, a variable lens and the retina.
- Can do amazing things like:
  - Recognize people and objects
  - Navigate through obstacles
  - Understand mood in the scene
  - Imagine stories
- But:
  - Suffers from illusions
  - Ignores many details
  - Ambiguous description of the world
  - Doesn’t care about accuracy of world

Computer Vision

- Vision is a complex physical and intellectual human task that stands as a primary interaction tool with the world.
- Computer vision is a field that includes methods for acquiring, processing, analyzing, and understanding images . . .
- Computer vision applications are increasing:
  - surveillance;
  - machine inspection;
  - medicine;
  - robotics;
  - entertainment;
  - media.
- The main goal: make computer vision converge towards human vision. Can we ever accomplish that?

Why is computer vision difficult?

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

- Objective of Computer Vision:
  - Understand the meaning of the image taken into account the value of the pixels

Viewpoint variation

Michelangelo 1475-1564

slide credit: Fei-Fei, Fergus & Torralba

Ilumination

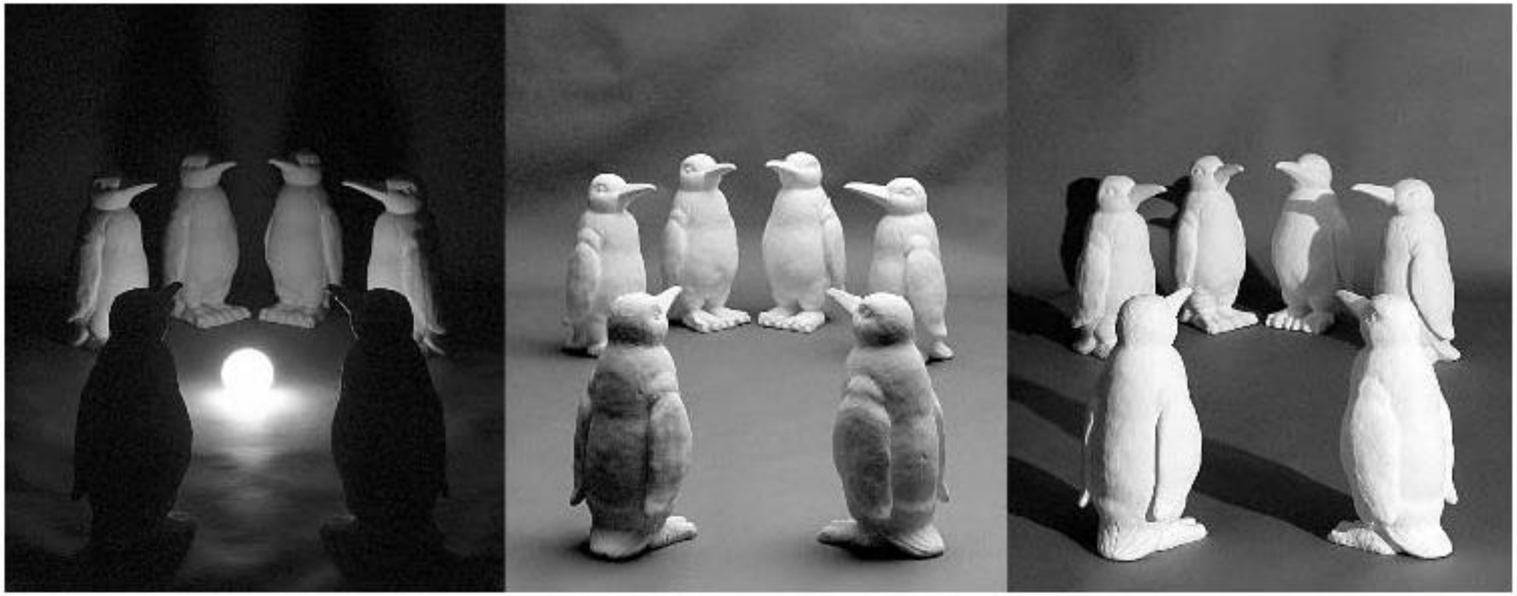

image credit: J. Koenderink

Scale

Shape and deformations

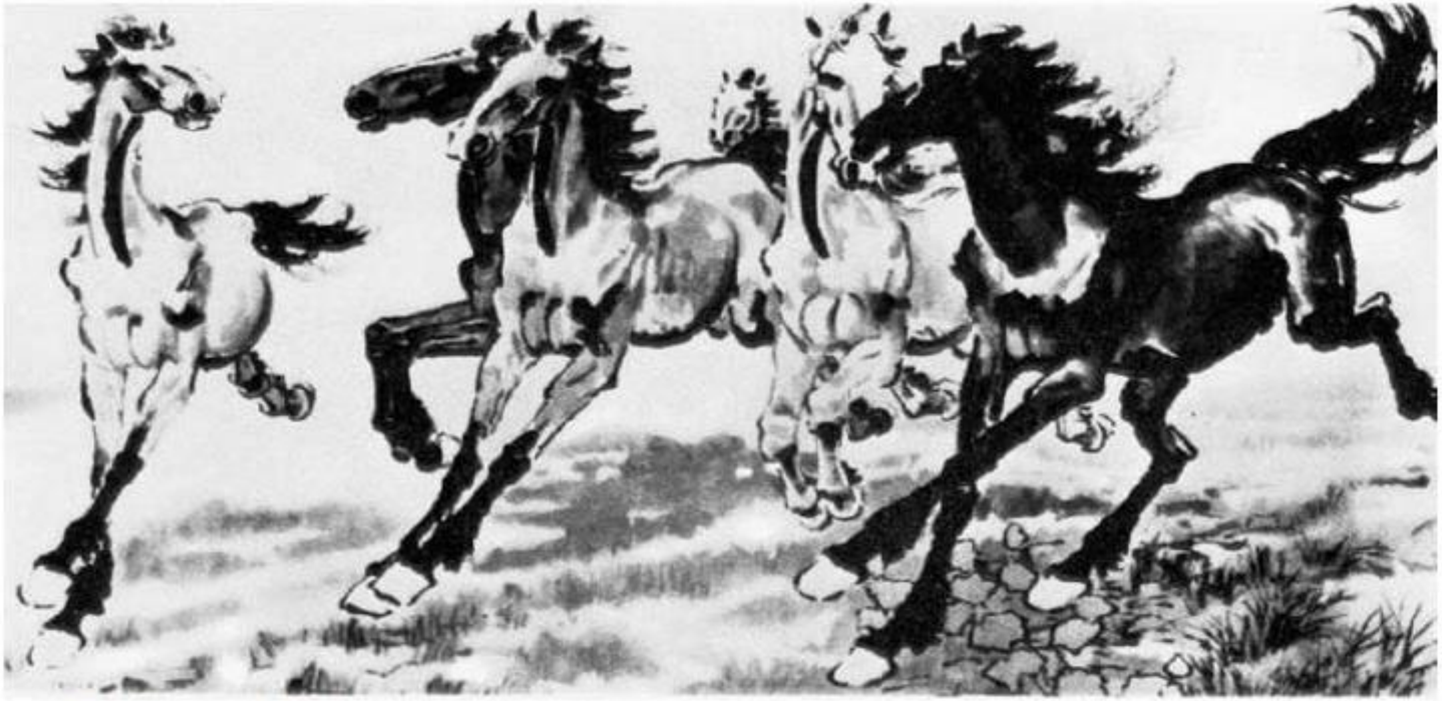

Xu, Beihong 1943

slide credit: Fei-Fei, Fergus & Torralba

Occlusions

Magritte, 1957

Background clutter

slide credit: Svetlana Lazebnik

Motion

slide credit: Svetlana Lazebnik

Intra-class variation

Need to understand the context

Need to understand the context

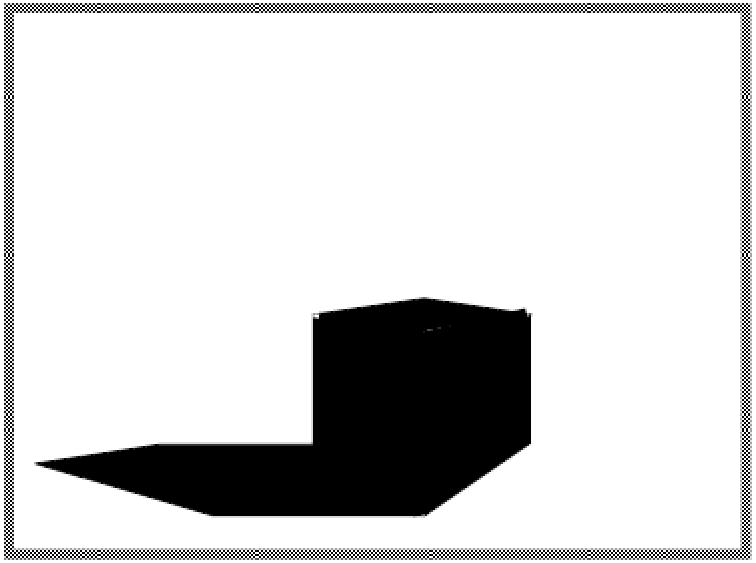

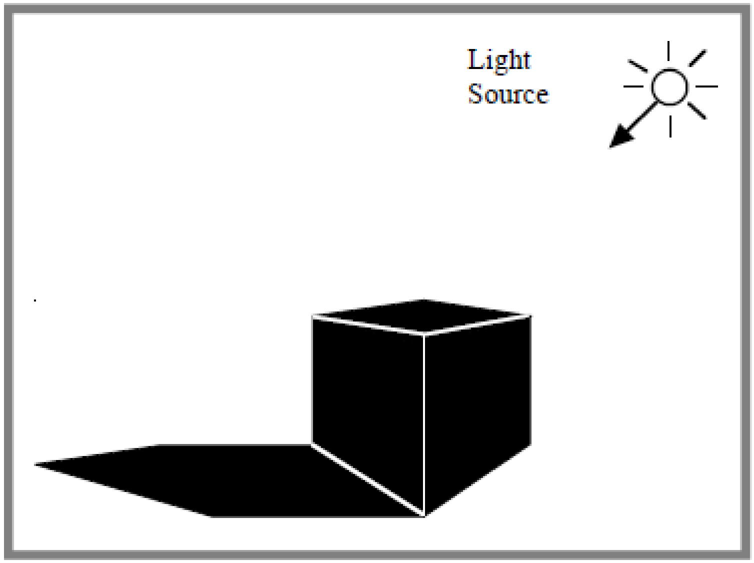

Acknowledgements: Most of this course is based on the excellent courses offered by Prof. Shree Nayar at Columbia University, USA and by Prof. Srinivasa Narasimhan at CMU, USA. Please acknowledge the original source when reusing these slides for academic purposes.

Outline

- Image Formation
  - The Human Visual System
  - Image Capturing Systems
- Digital Images
  - Sampling
  - Data Structures
  - Histograms
- Colour and Noise
  - Colour spaces
  - Colour processing
  - Noise

Outline

- Image Formation
  - The Human Visual System
  - Image Capturing Systems
- Digital Images
- Colour and Noise

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. Please acknowledge the original source when reusing these slides for academic purposes.

Components of a Computer Vision System

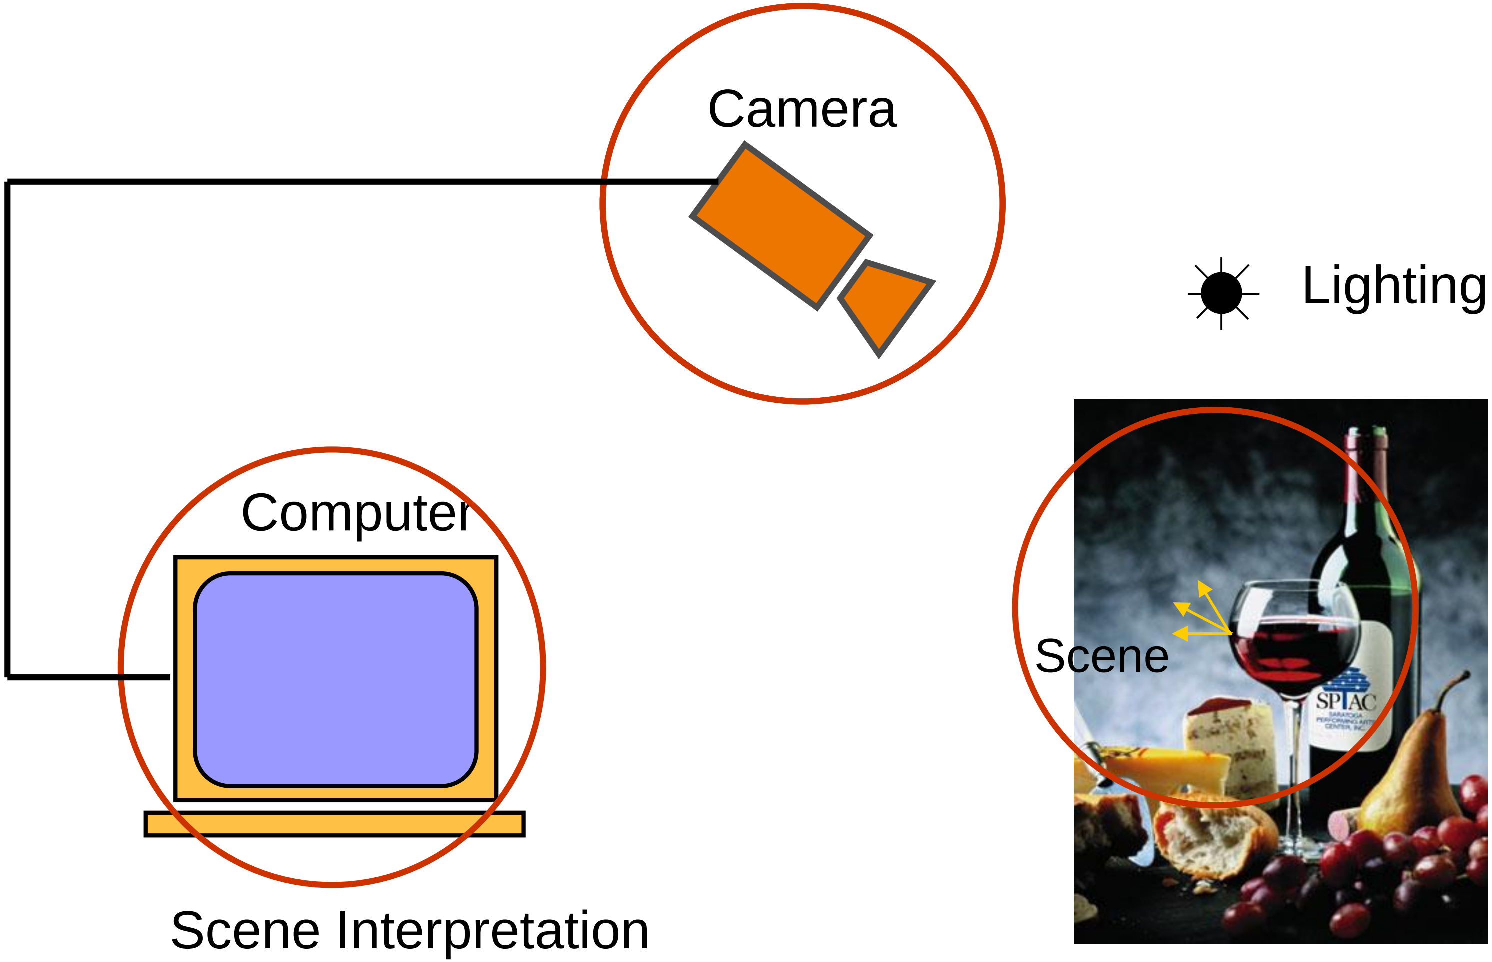

Our Eyes

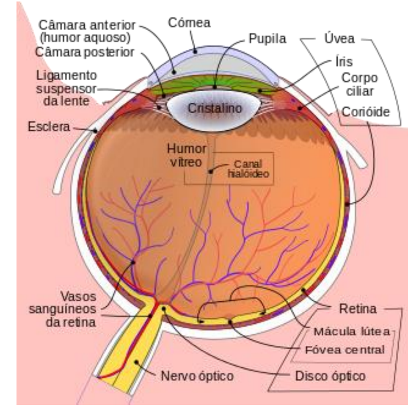

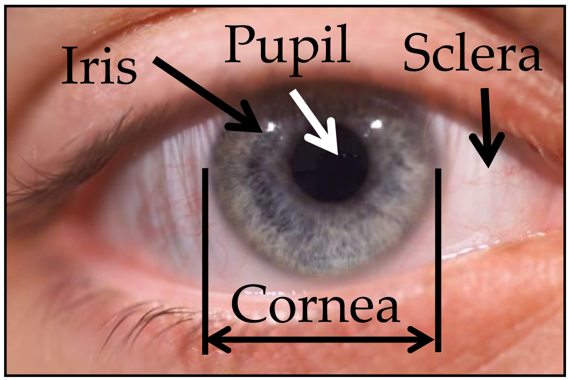

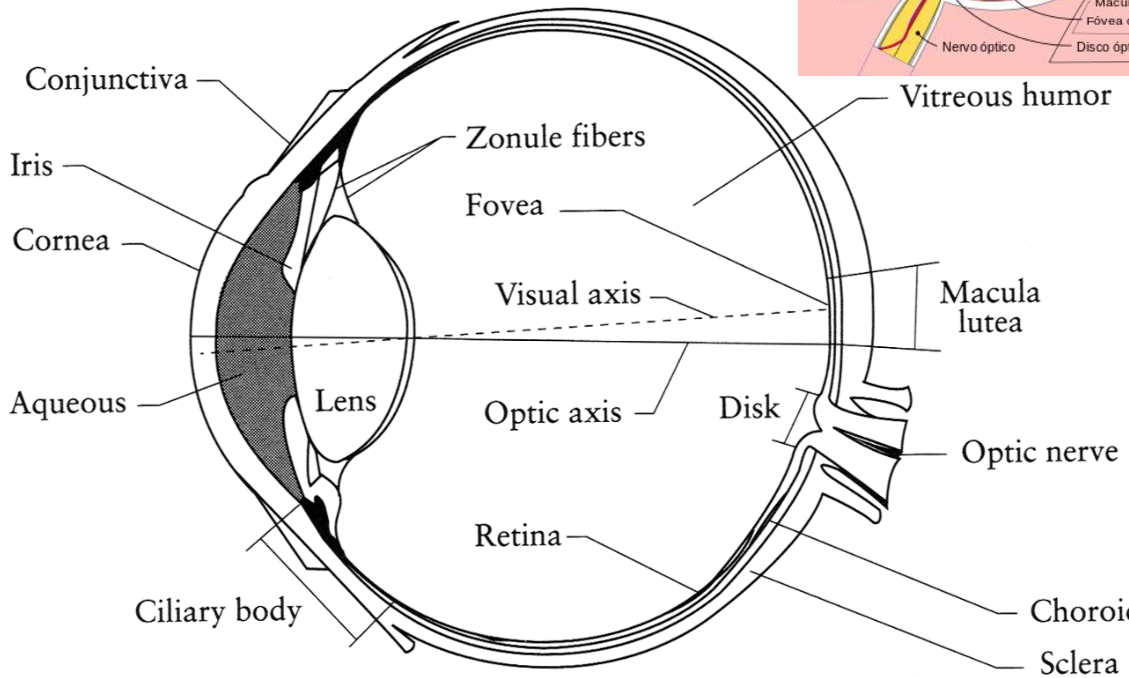

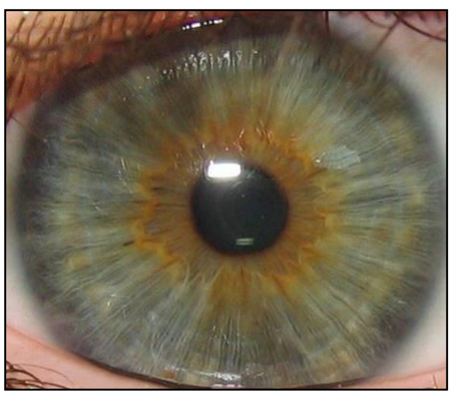

-Iris is the diaphragm that changes the aperture (pupil)

-Retina is the sensor where the fovea has the highest resolution

Colour

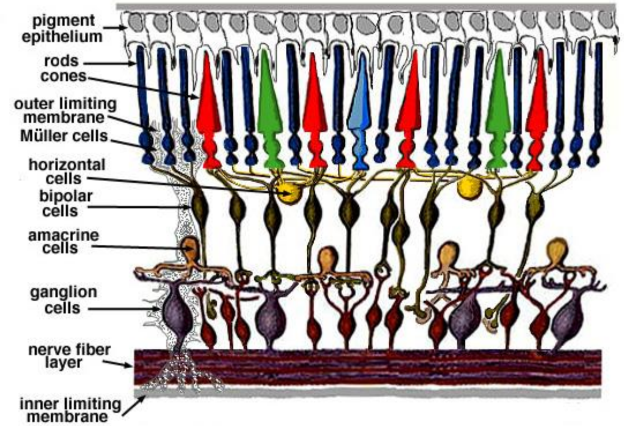

- Our retina has:
  - **Cones** – Measure the frequency of light (colour)
    - 6 to 7 millions
    - High-definition
    - Need high luminosity
  - **Rods** – Measure the intensity of light (luminance)
    - 75 to 150 millions
    - Low-definition
    - Function with low luminosity

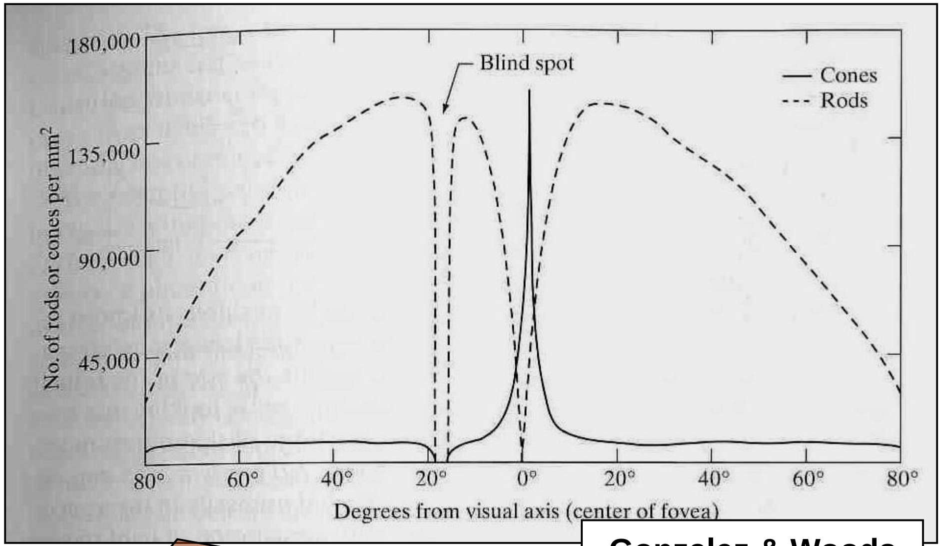

Gonzalez & Woods

We only see colour in the center of our retina!

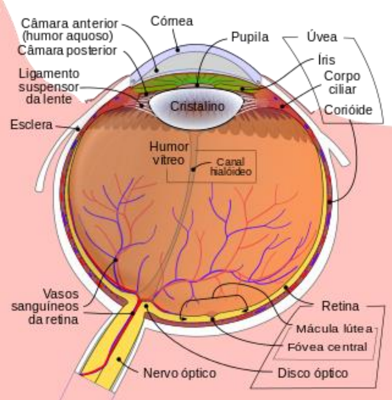

Focusing

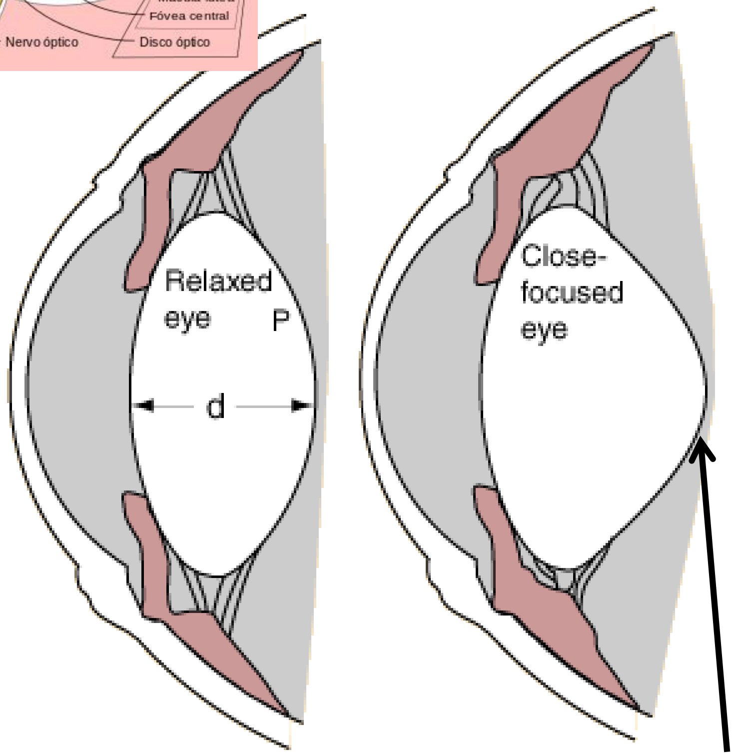

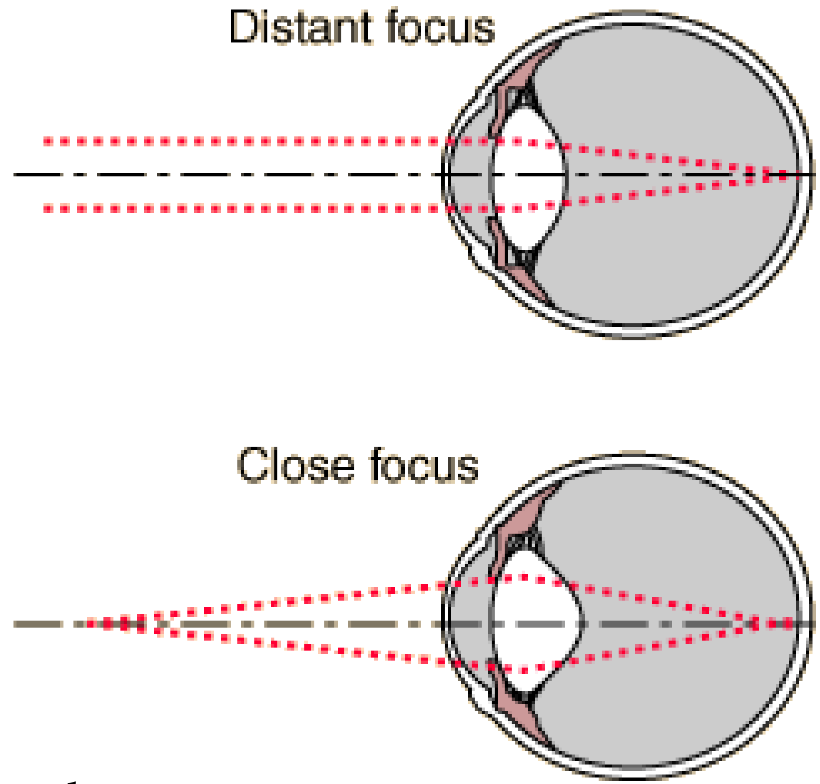

shorter focal length

Changes the focal length of the lens

Myopia and Hyperopia

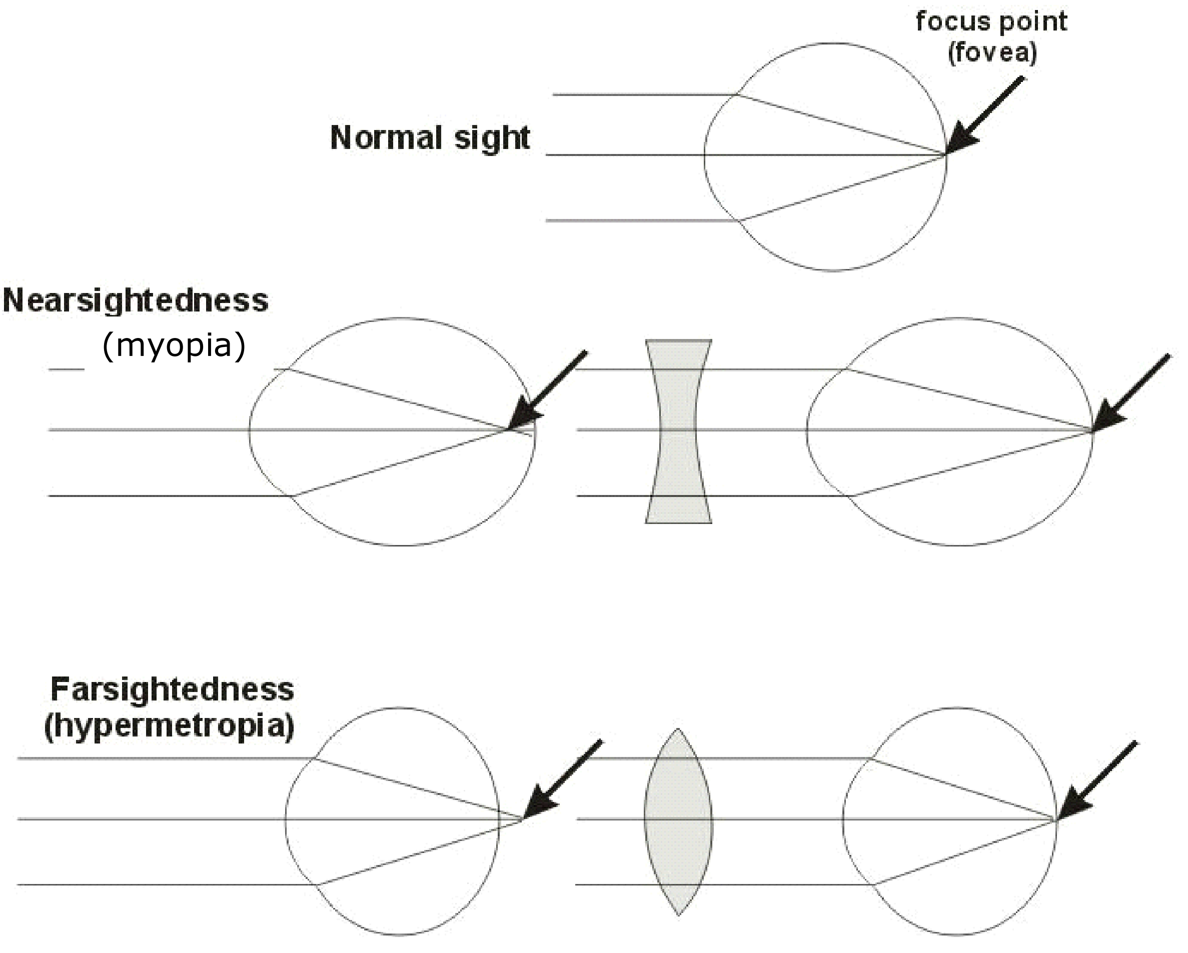

Astigmatism

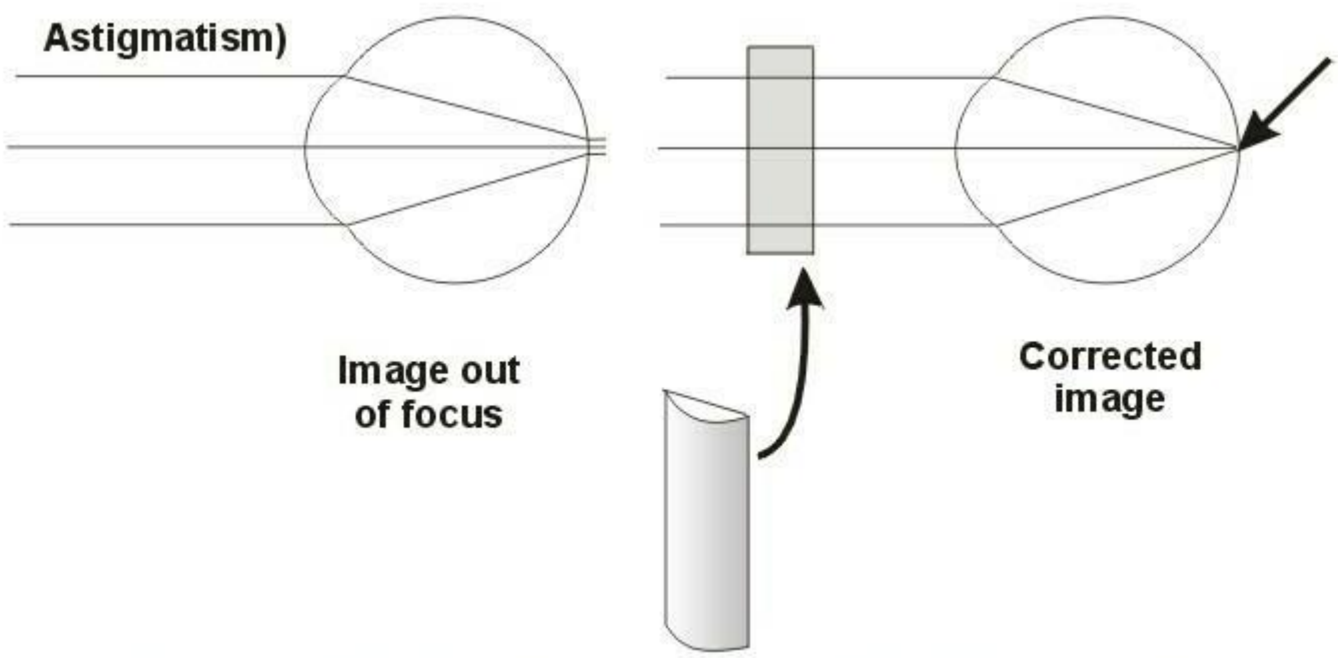

The cornea is distorted causing images to be un-focused on the retina.

Blind Spot in the Eye

Close your right eye and look directly at the “+”

https://en.wikipedia.org/wiki/Blind_spot_(vision)

More Illusions

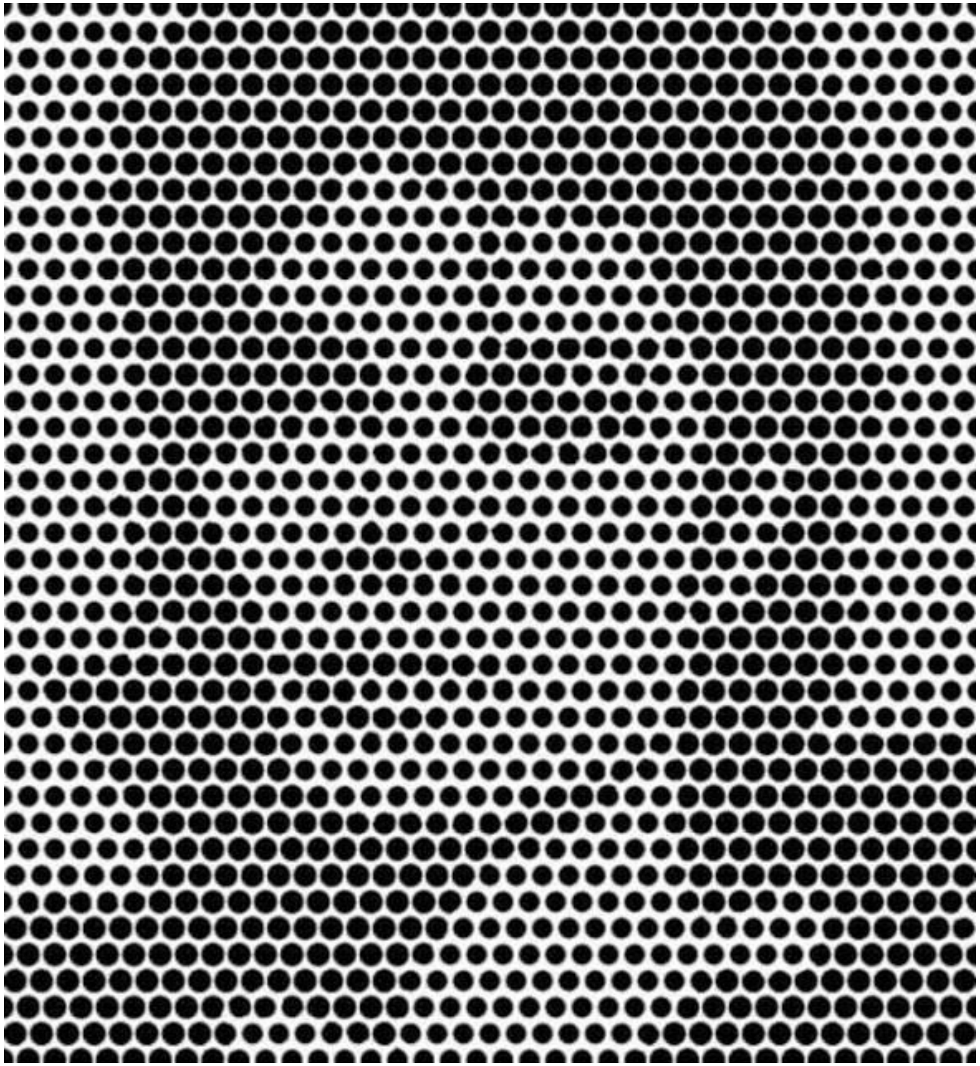

More Illusions

More Illusions

More Illusions

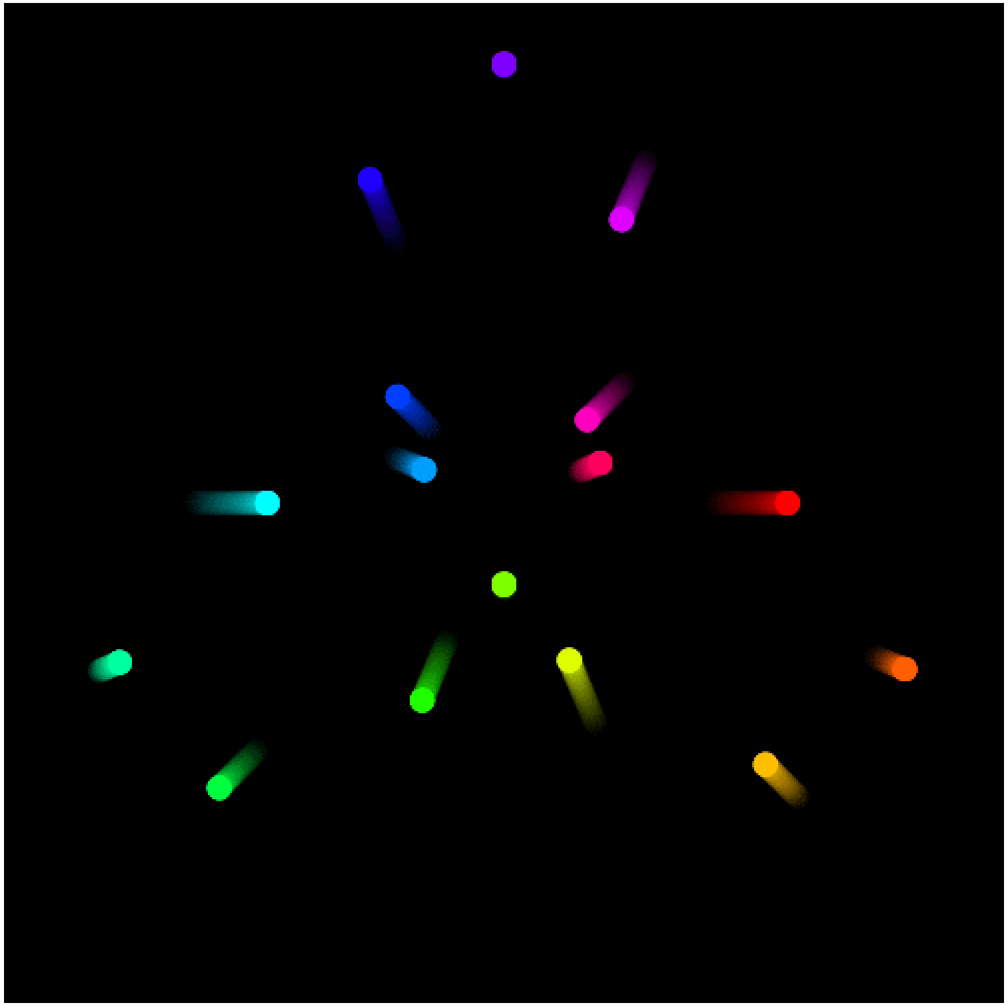

More Illusions

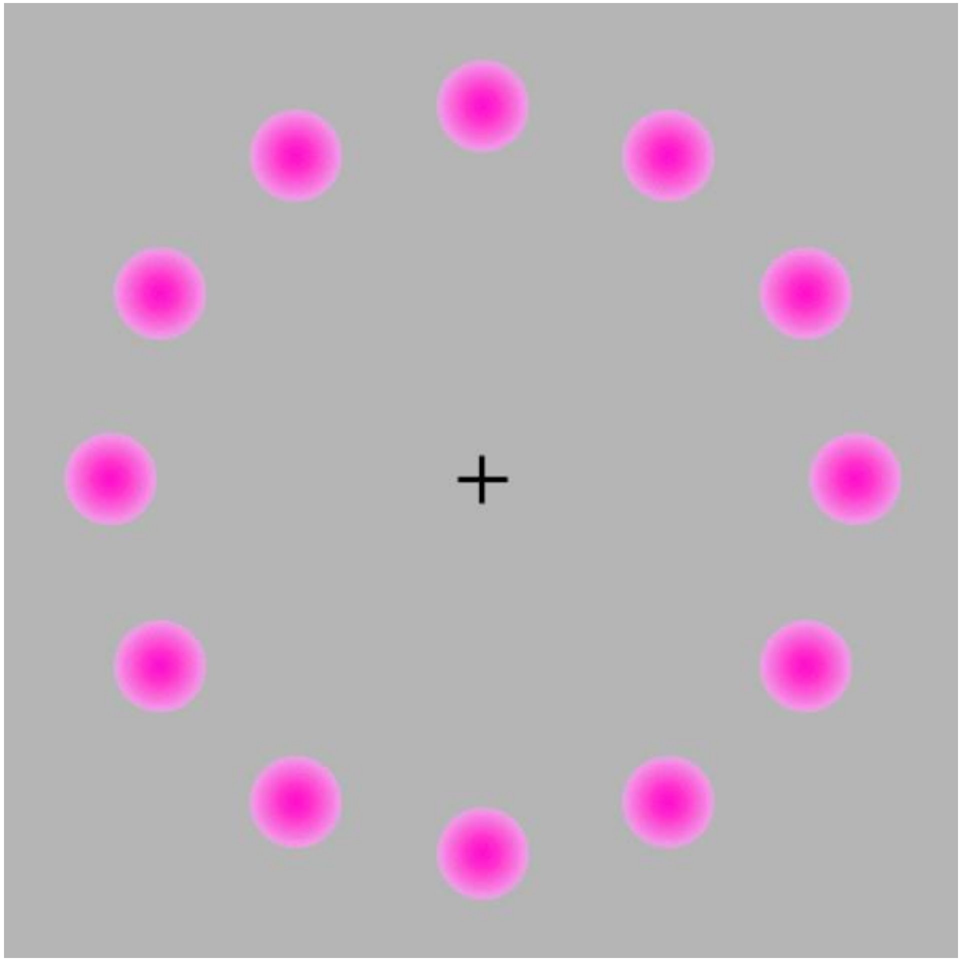

More Illusions

More Illusions

More Illusions

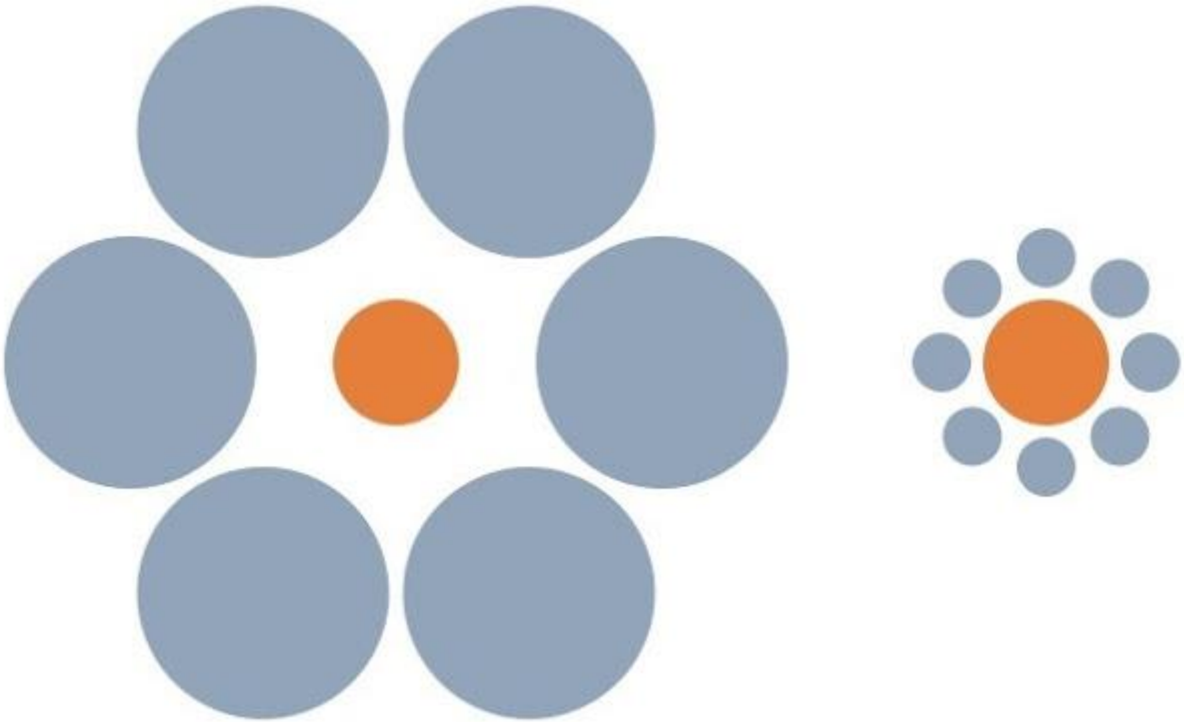

More Illusions

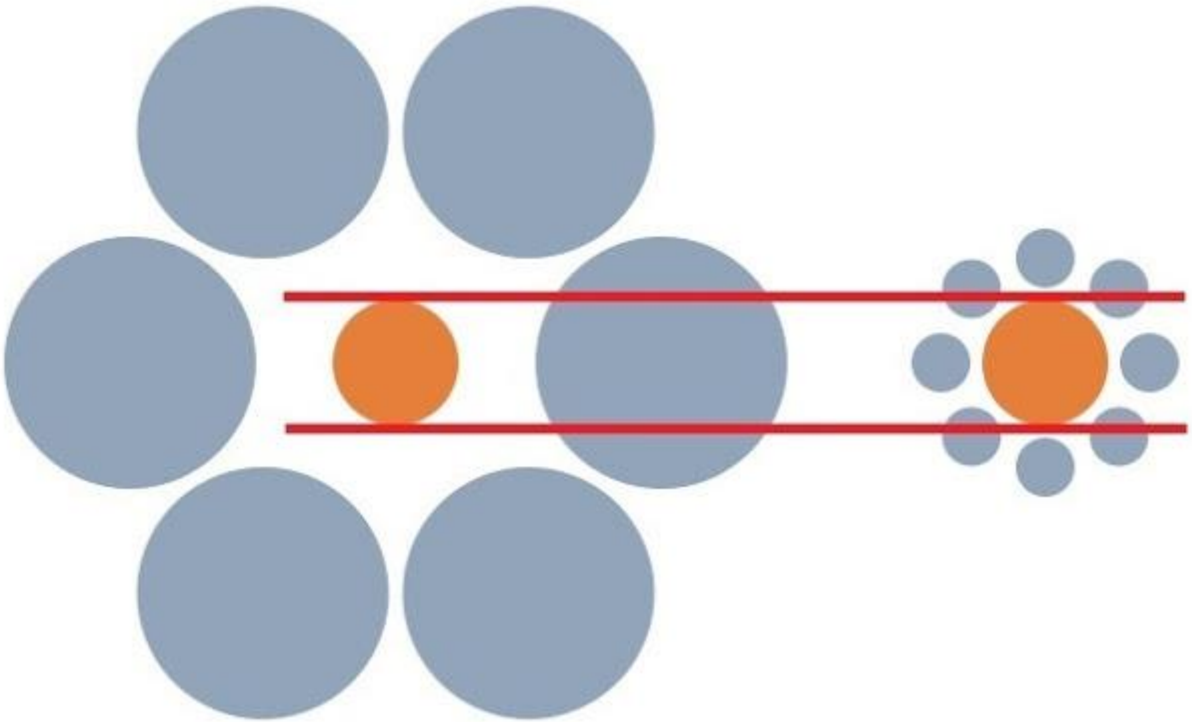

More Illusions

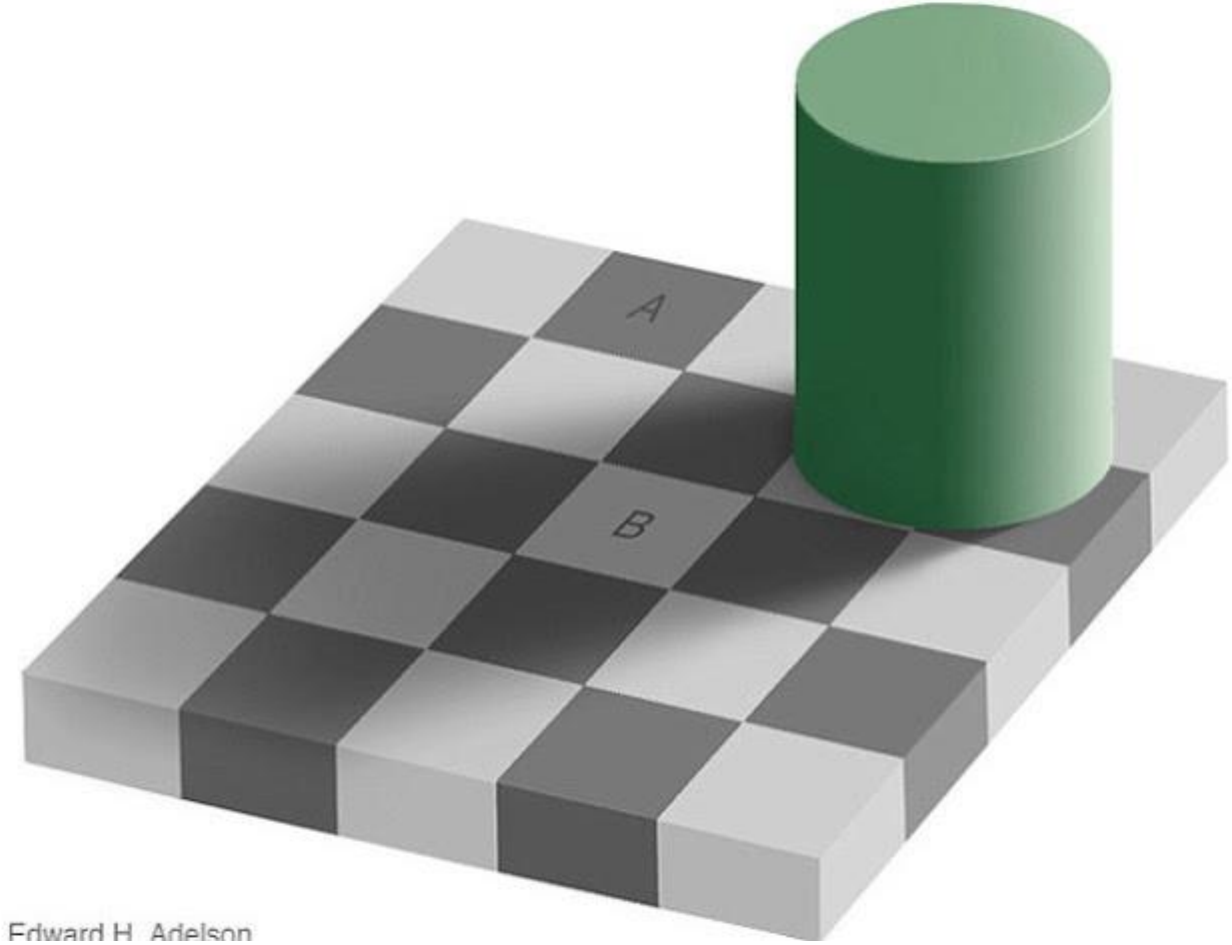

More Illusions

More Illusions

More Illusions

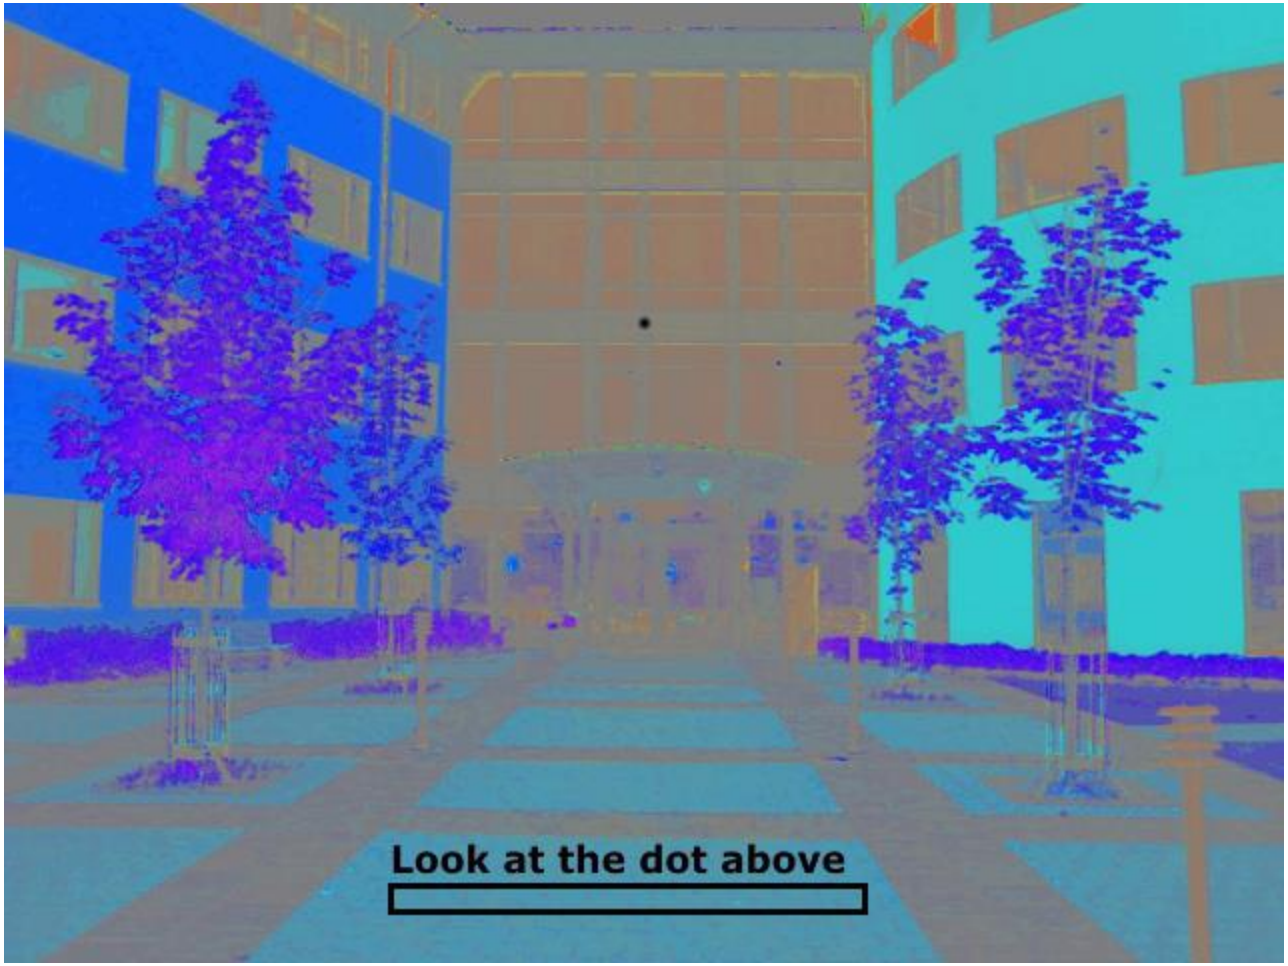

More Illusions

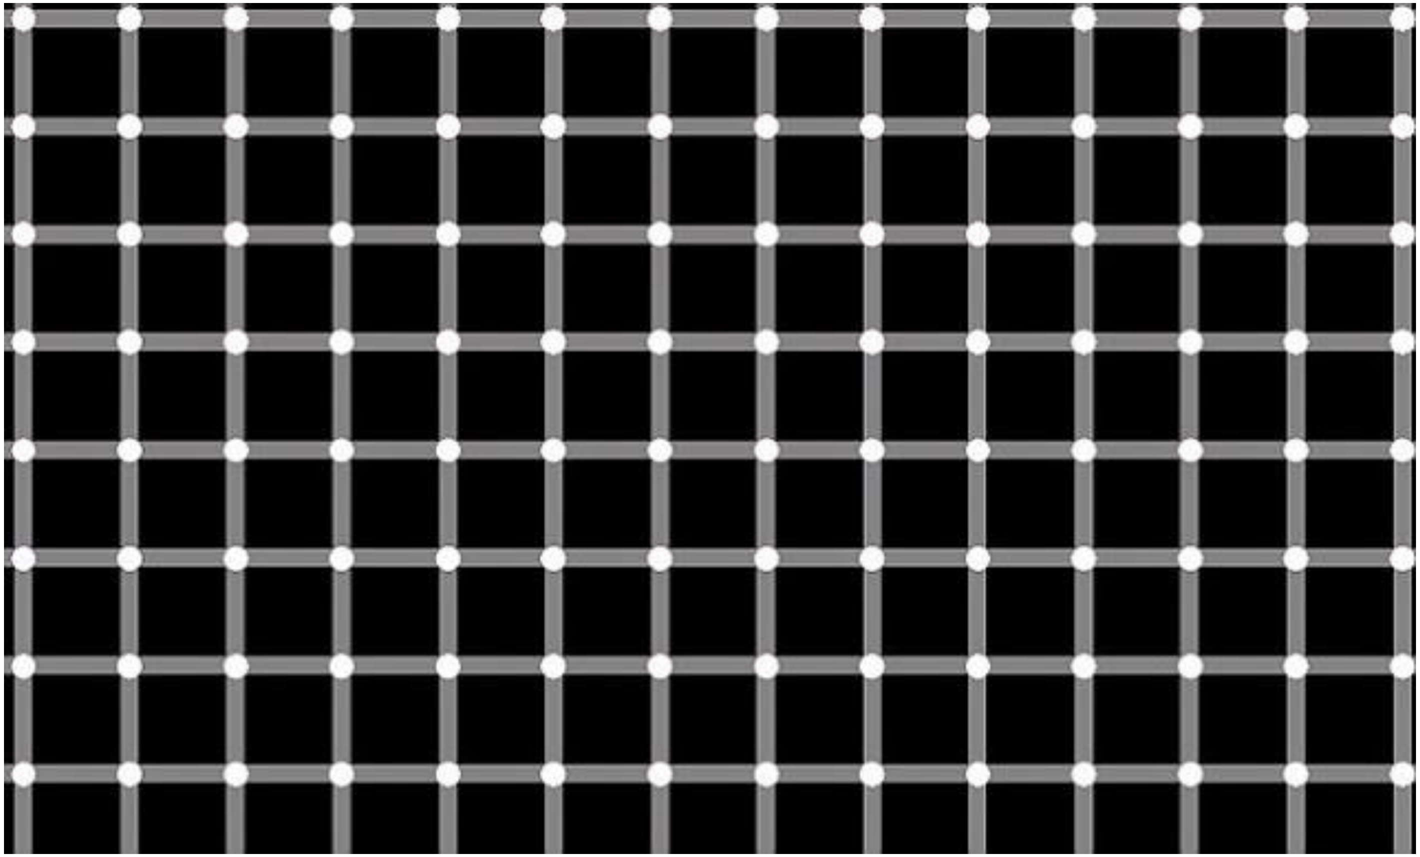

More Illusions

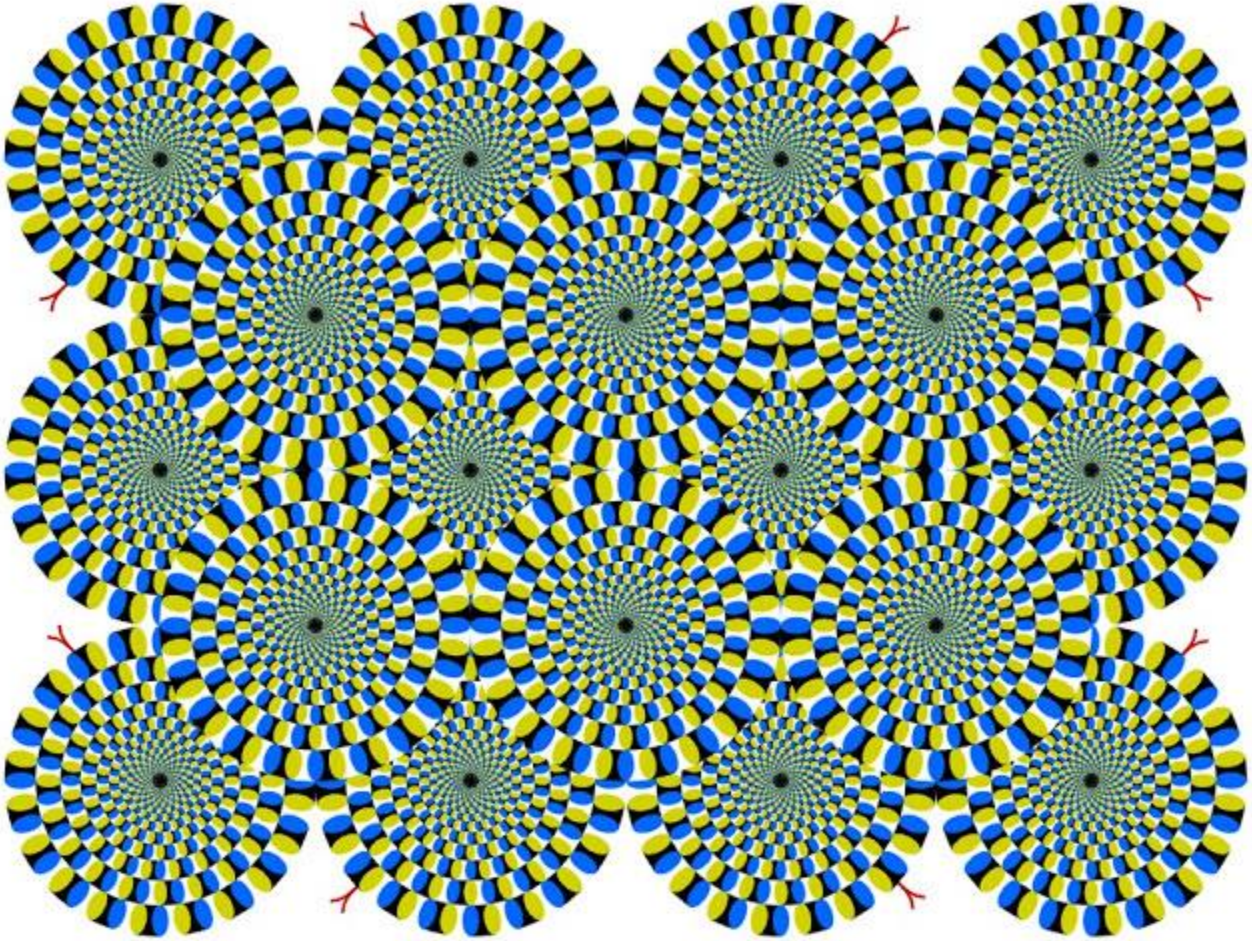

More Illusions

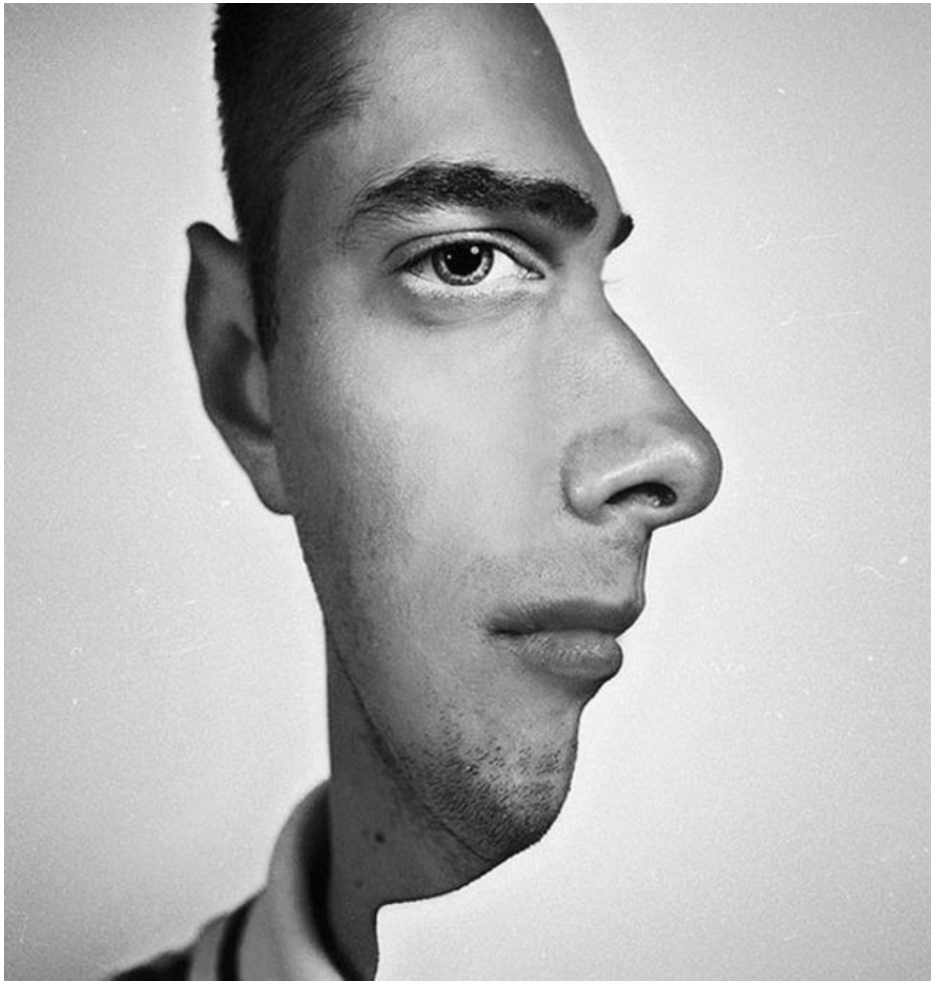

More Illusions

More Illusions

More Illusions

More Illusions

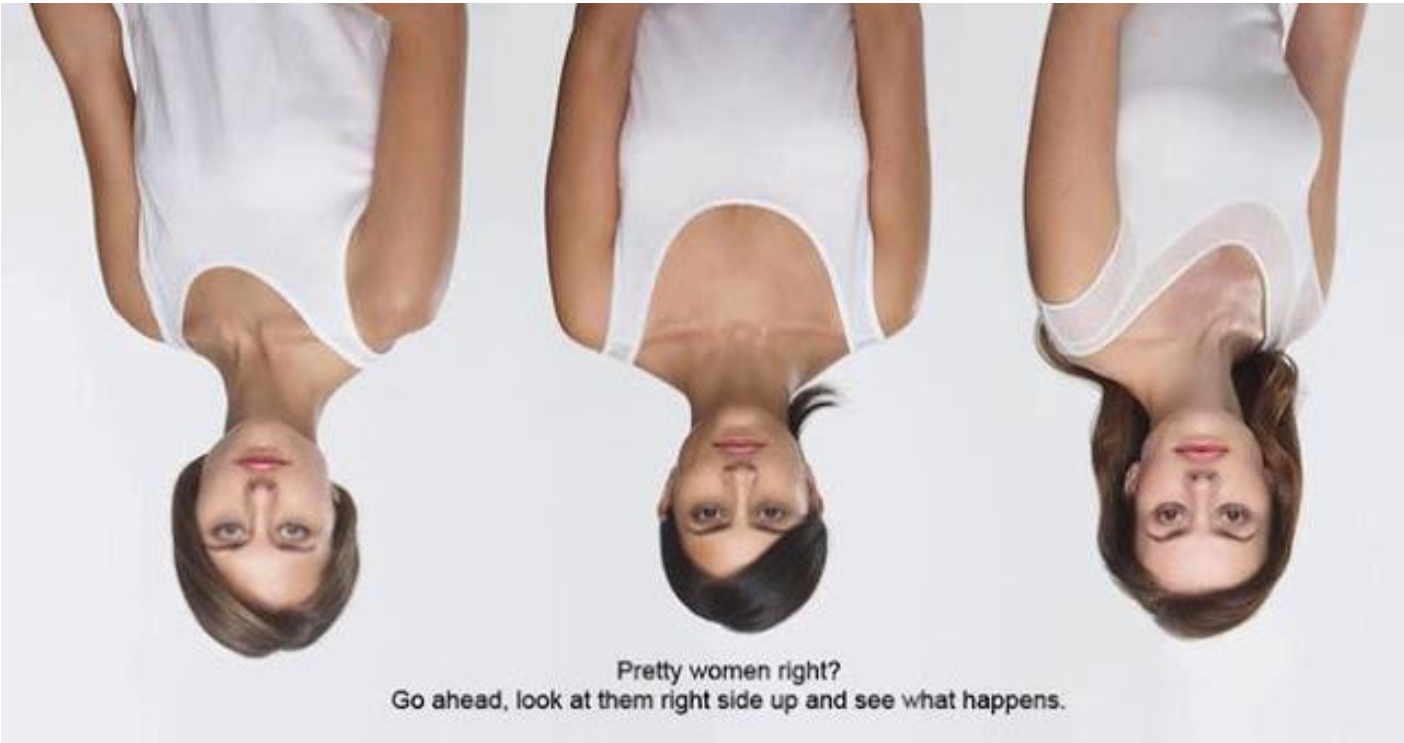

More Illusions

More Illusions

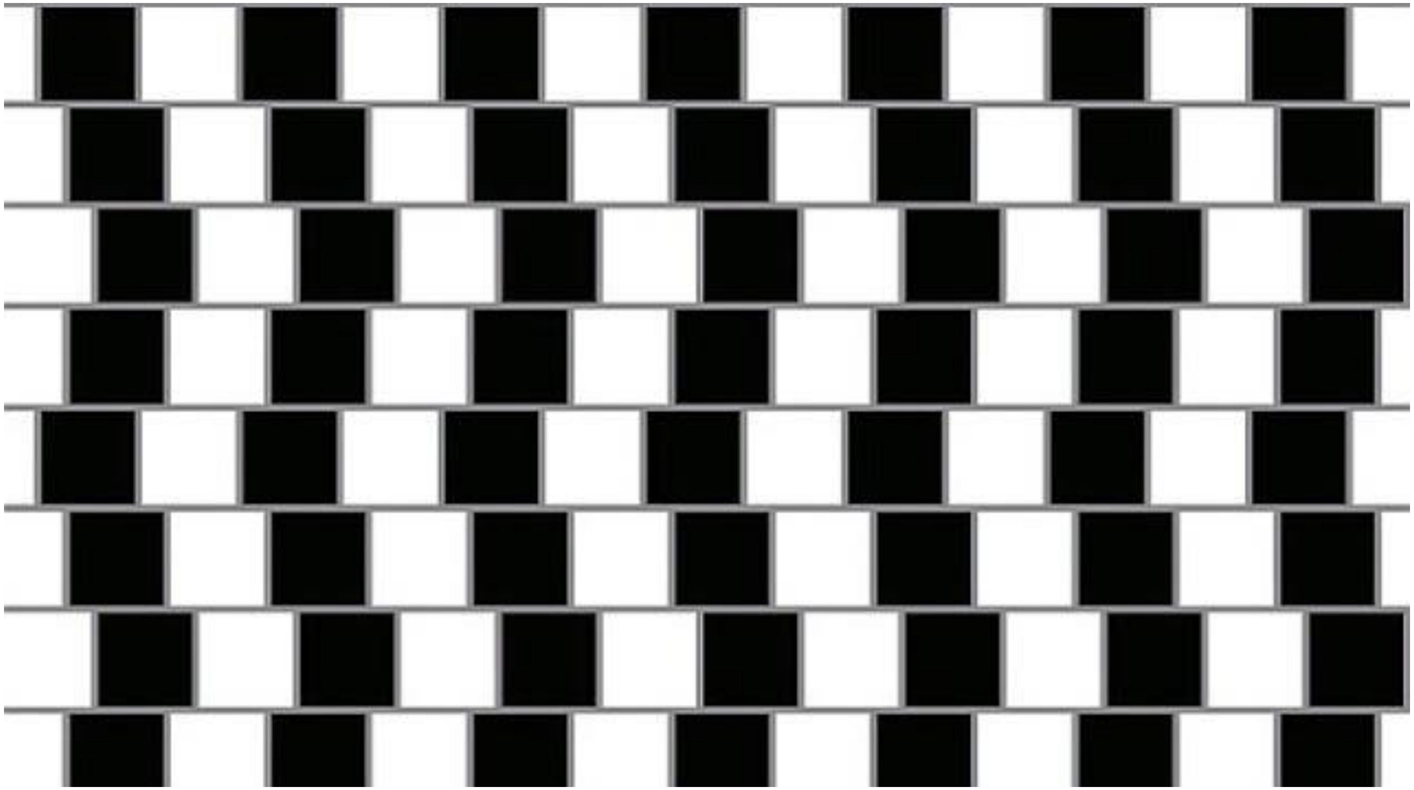

More Illusions

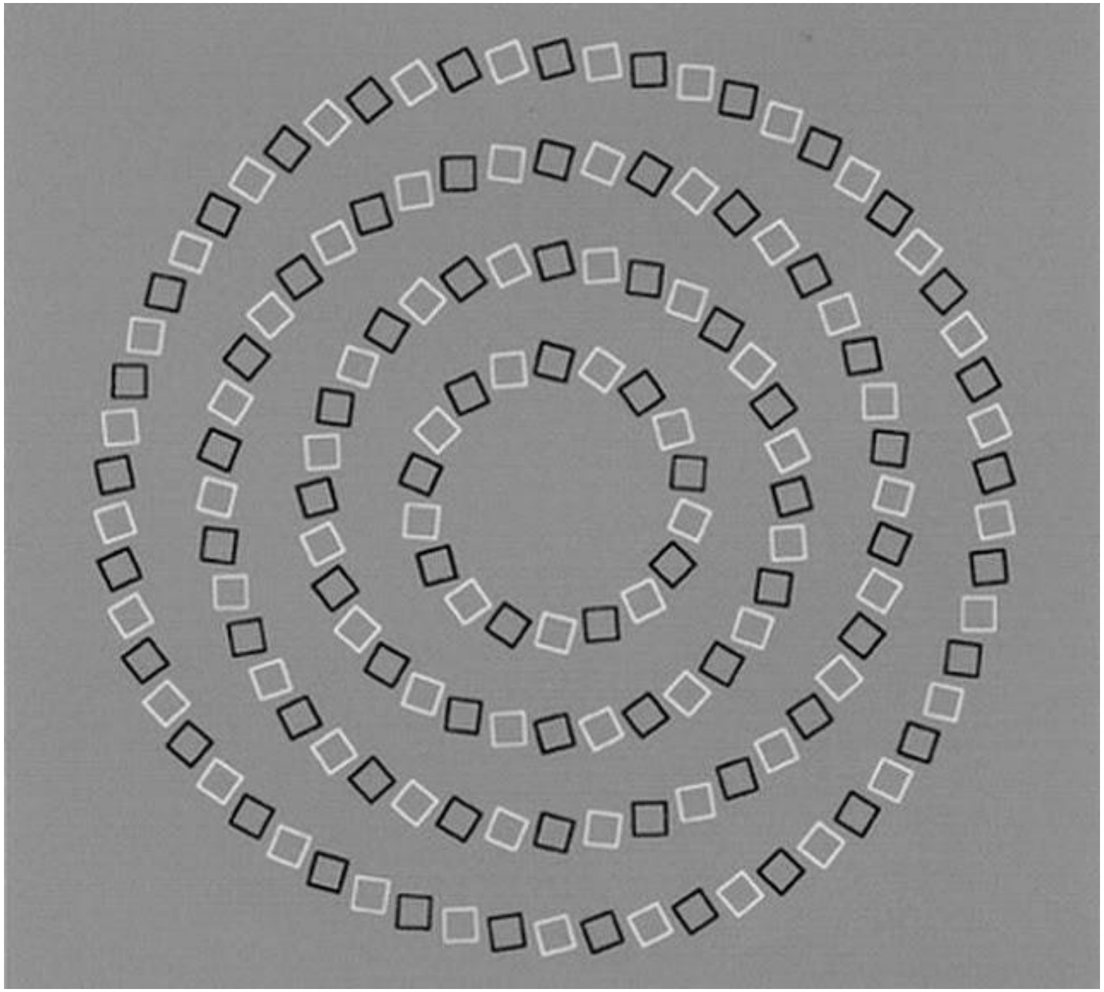

More Illusions

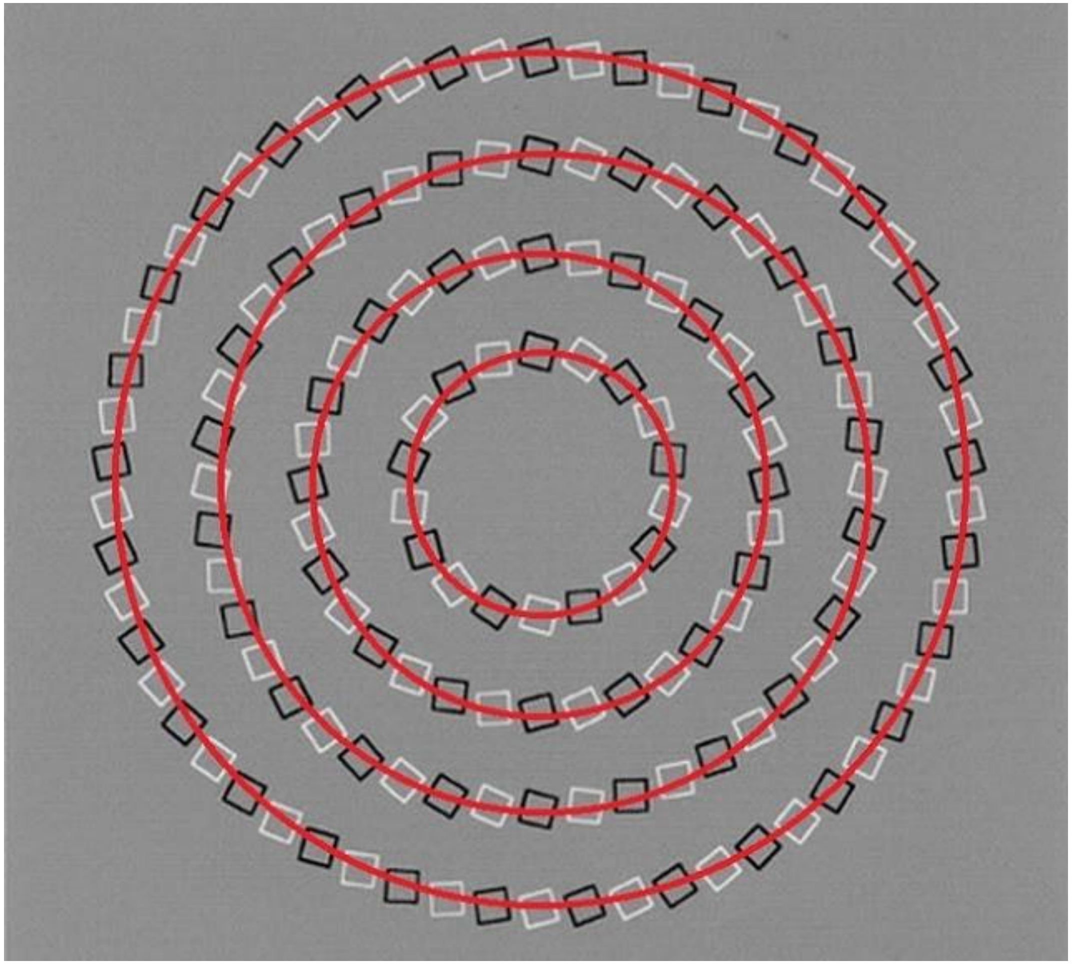

More Illusions

More Illusions

More Illusions

More Illusions

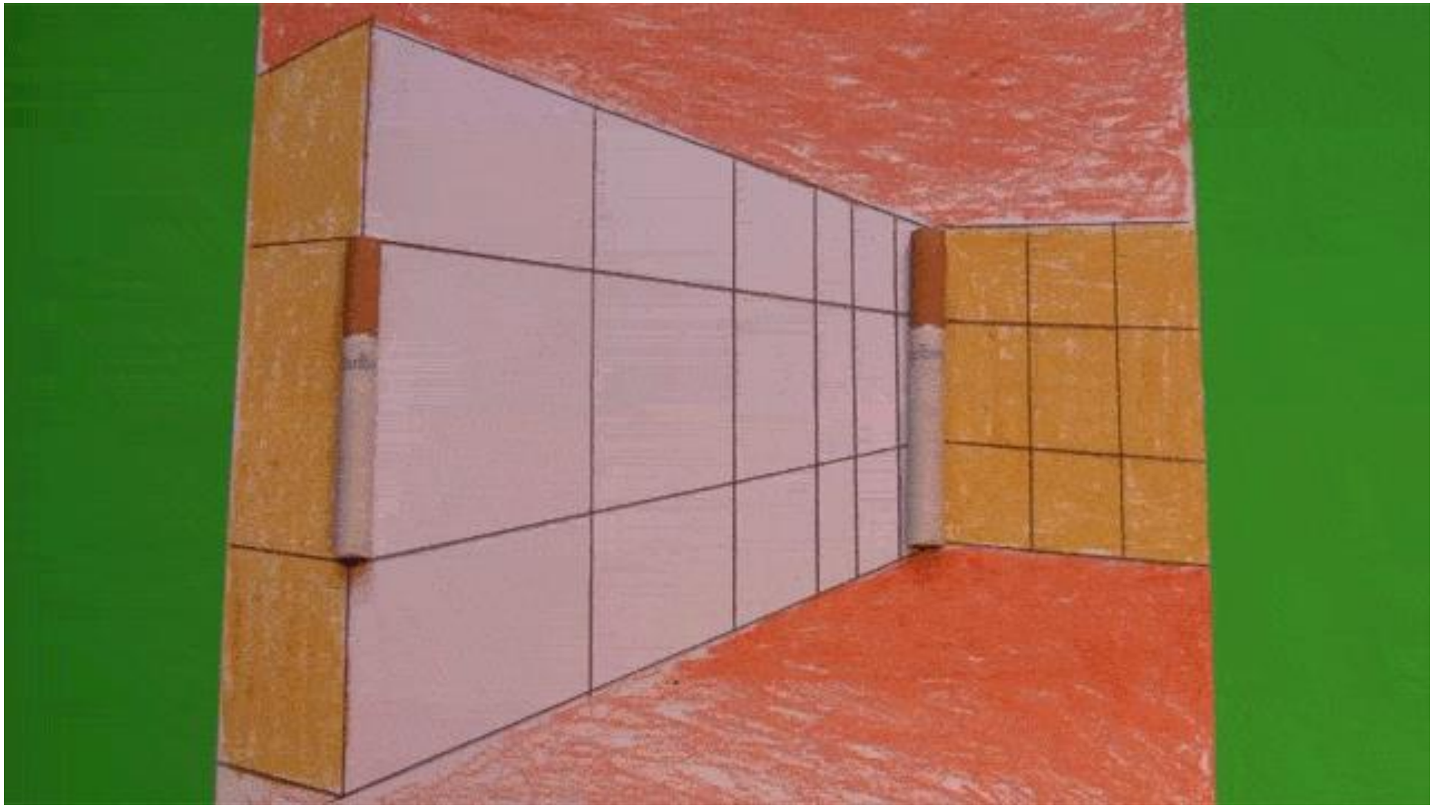

http://gizmodo.uol.com.br/20-ilusoes-opticas-para-confundir-a-sua-mente/

More Illusions

http://gizmodo.uol.com.br/20-ilusoes-opticas-para-confundir-a-sua-mente/

More Illusions

http://www.ilusoesopticas.com

More Illusions

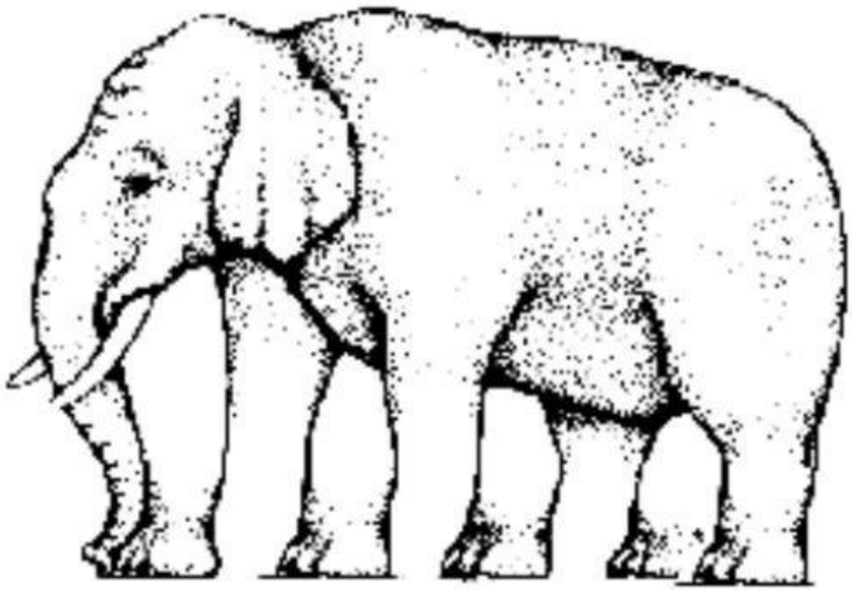

Quantas pernas tem o elefante?

Outline

- Image Formation
  - The Human Visual System
  - Image Capturing Systems
- Digital Images
- Colour and Noise

A Brief History of Images

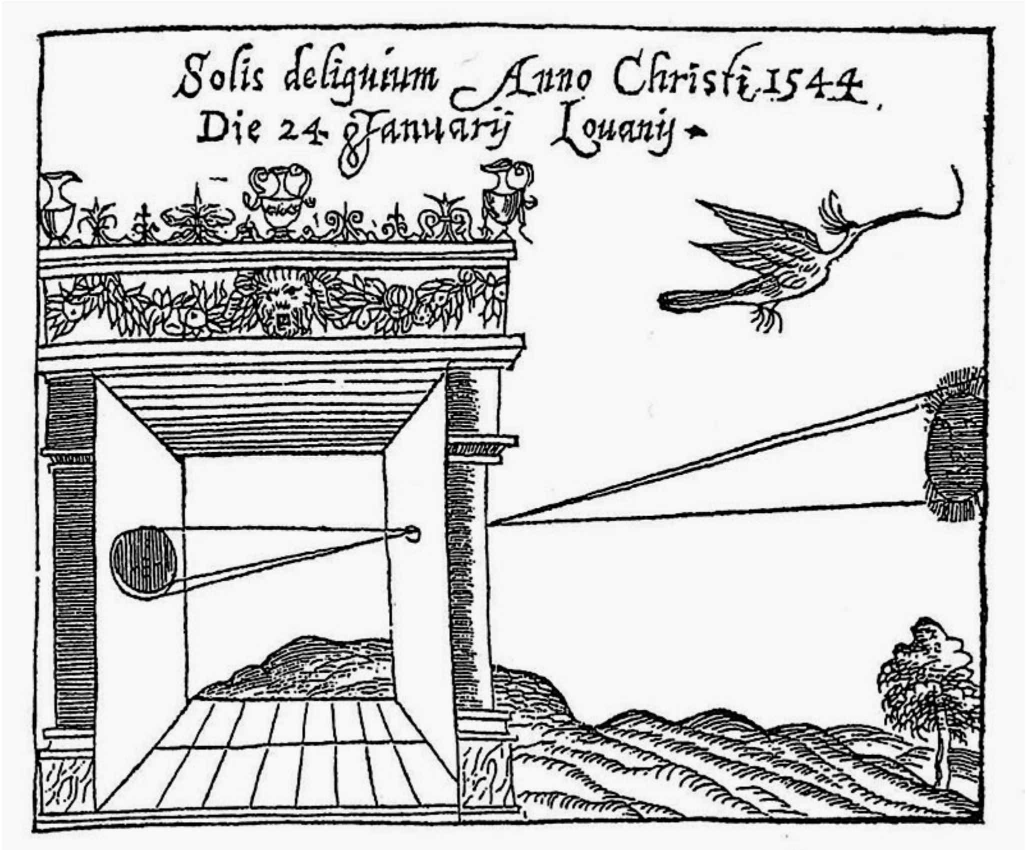

Camera Obscura, Gemma Frisius, 1544

A Brief History of Images

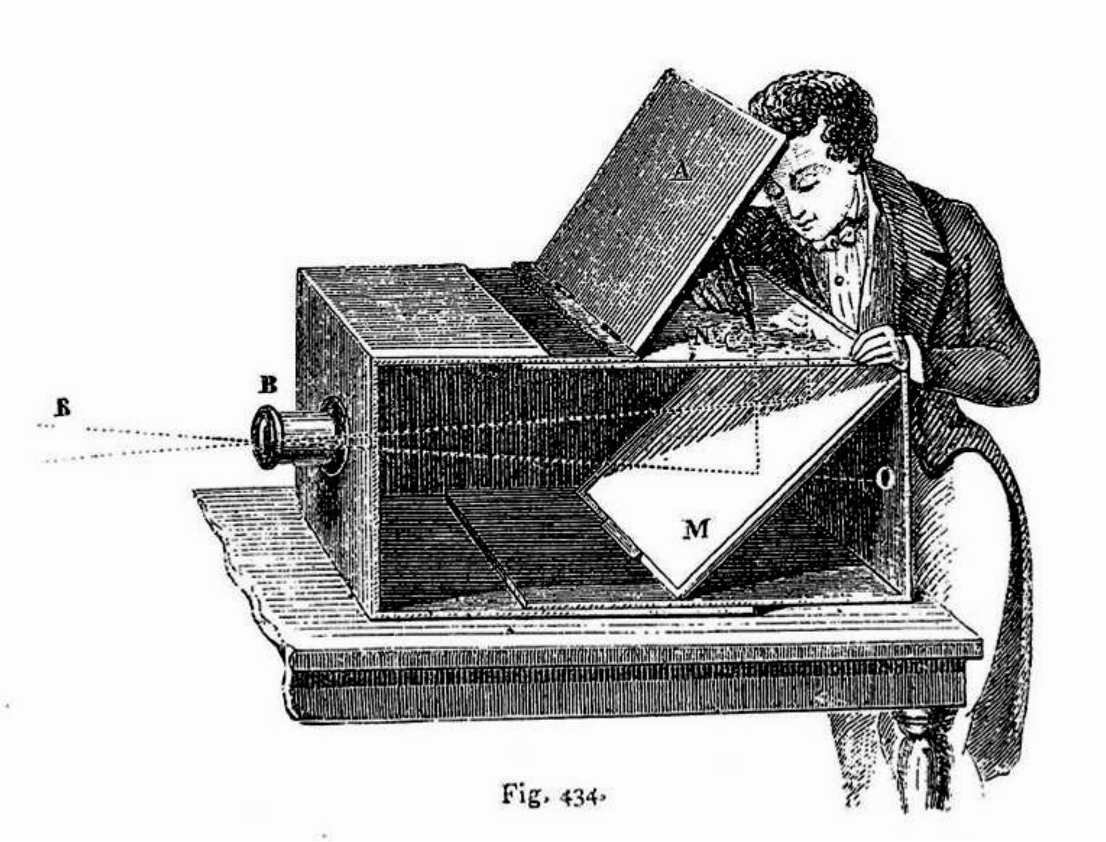

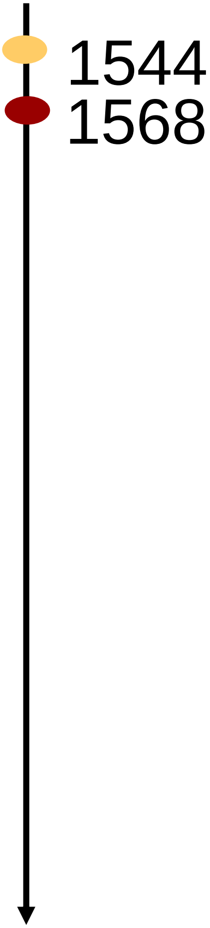

Lens Based Camera Obscura, 1568

A Brief History of Images

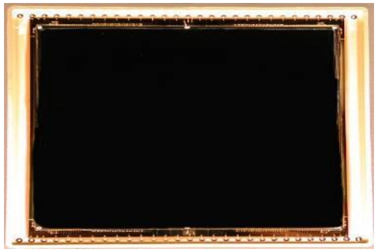

Silicon Image Detector, 1970

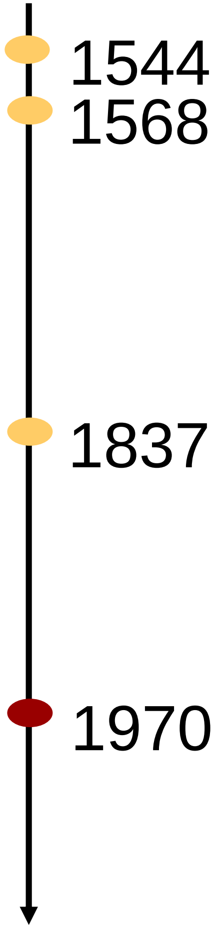

A Brief History of Images

Digital Cameras

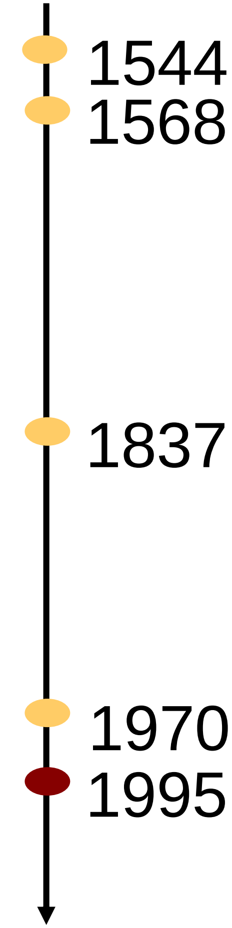

Components of a Computer Vision System

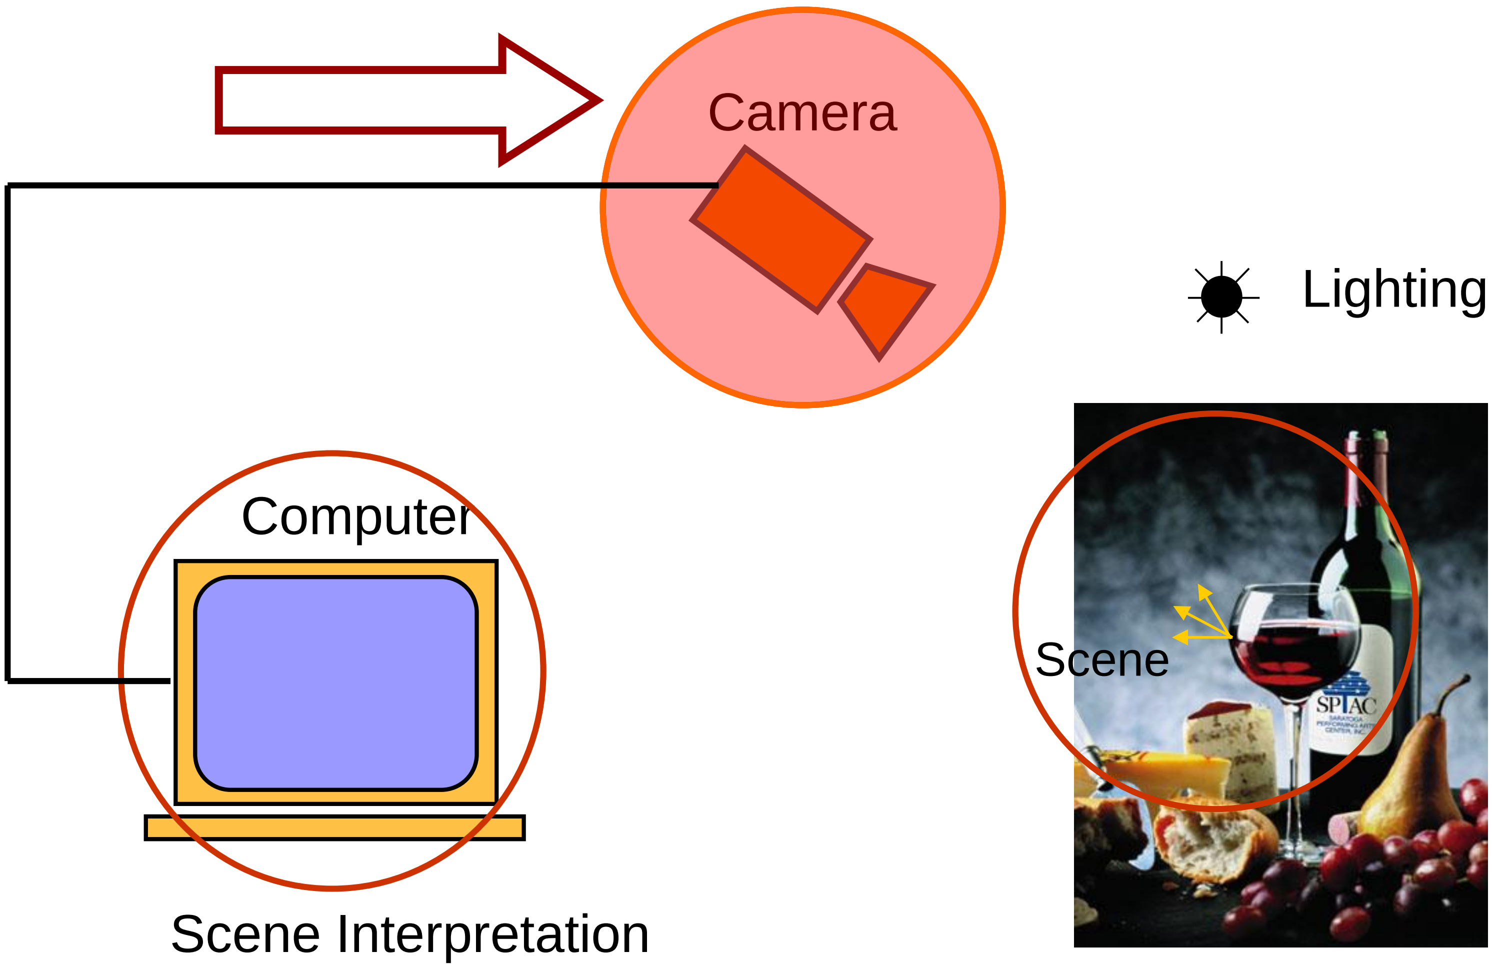

Pinhole and the Perspective Projection

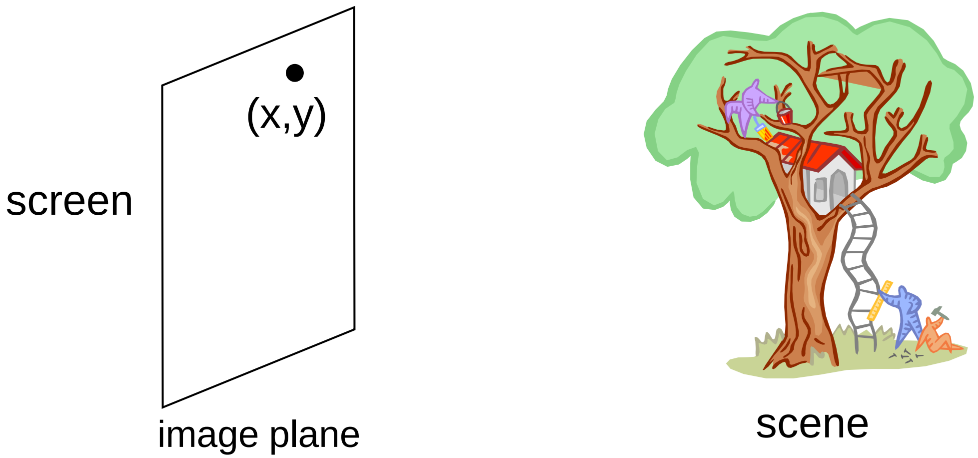

Is an image being formed on the screen?

YES! But, not a “clear” one.

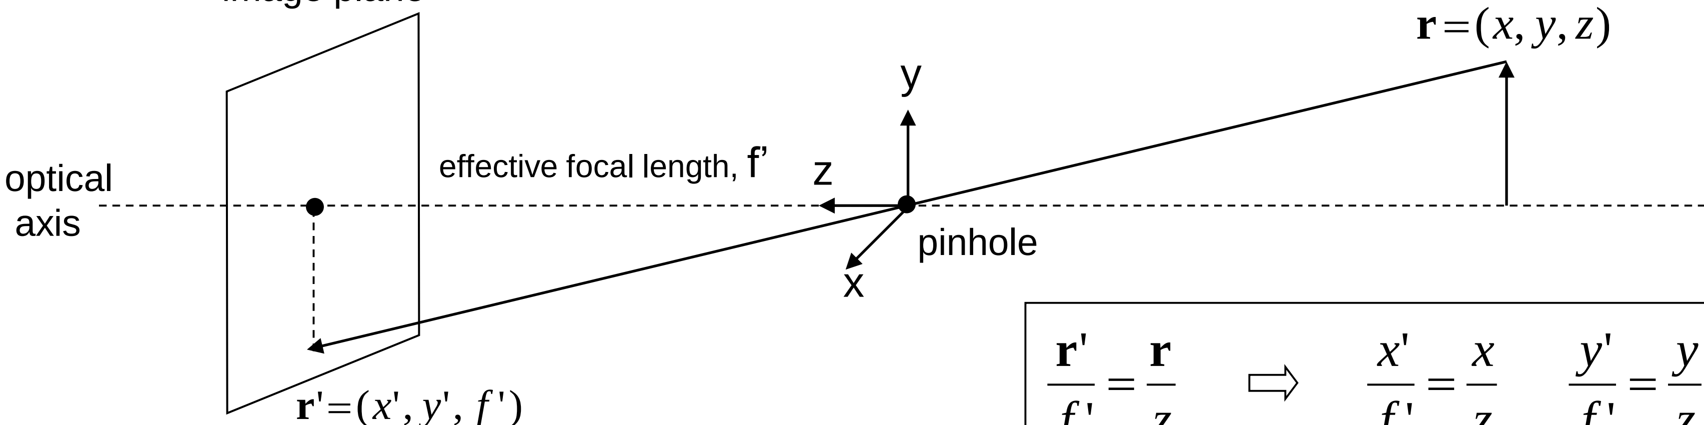

$$\frac{\mathbf{r}'}{f'} = \frac{\mathbf{r}}{z} \quad \Rightarrow \quad \frac{x'}{f'} = \frac{x}{z} \qquad \frac{y'}{f'} = \frac{y}{z}$$

Pinhole Camera

• Basically a pinhole camera is a box, with a tiny hole at one end and film or photographic paper at the other.

• Mathematically: out of all the light rays in the world, choose the set of light rays passing through a point and projecting onto a plane.

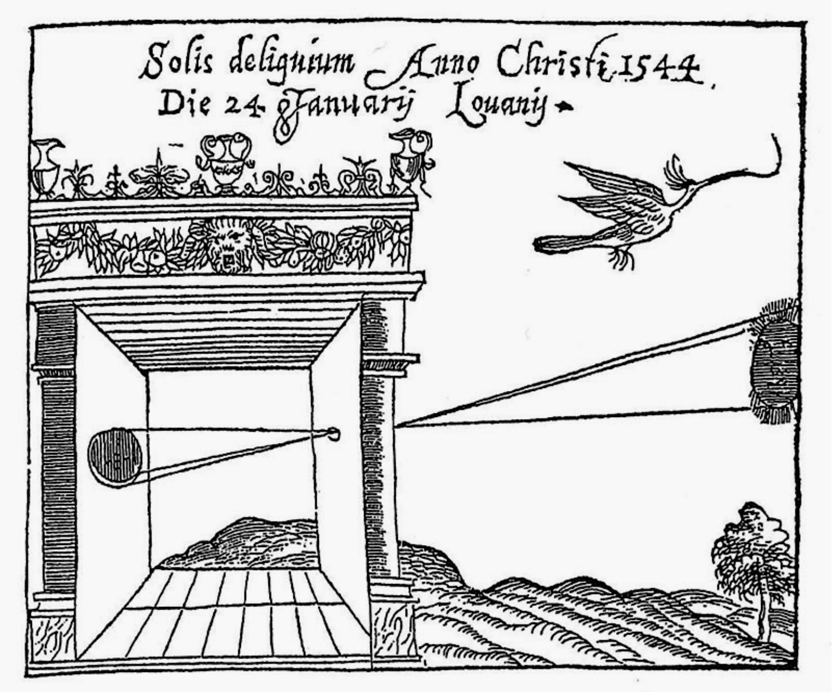

Do it by yourself!!!!
http://www.sbfisica.org.br/fne/Vol8/Num2/v08n02a05.pdf

Magnification

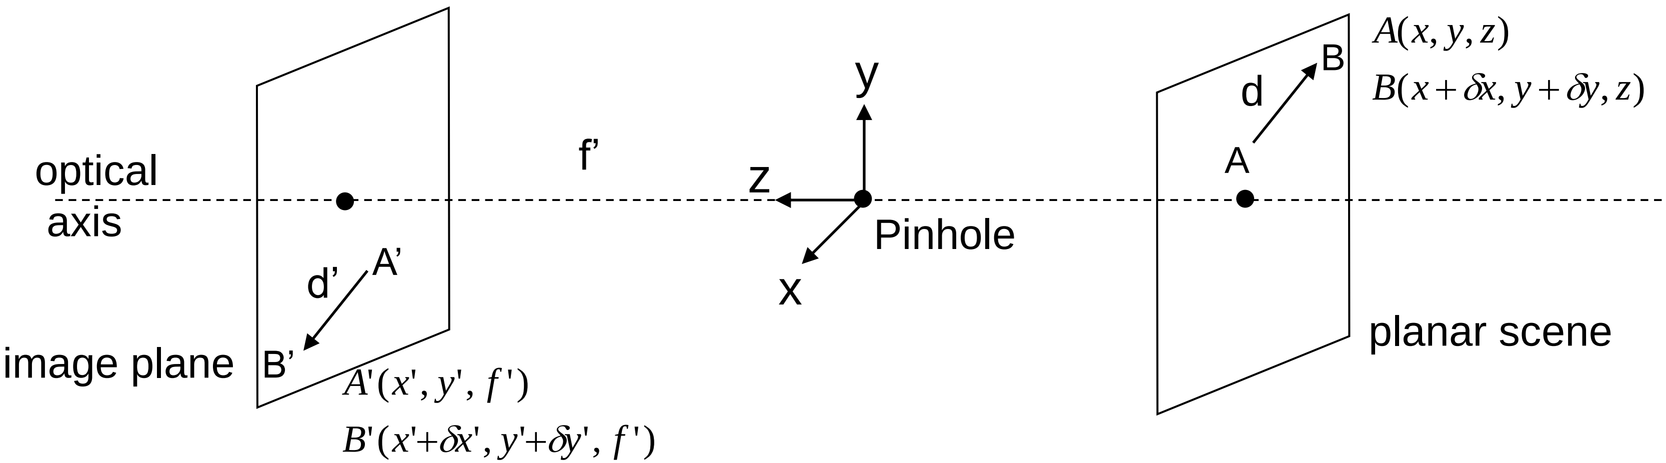

From perspective projection:

$$\frac{x'}{f'} = \frac{x}{z} \qquad \frac{y'}{f'} = \frac{y}{z}$$

$$\frac{x'+\delta x'}{f'} = \frac{x+\delta x}{z} \qquad \frac{y'+\delta y'}{f'} = \frac{y+\delta y}{z}$$

Magnification:

$$m = \frac{d'}{d} = \frac{\sqrt{(\delta x')^2 + (\delta y')^2}}{\sqrt{(\delta x)^2 + (\delta y)^2}} = \frac{f'}{z}$$

$$\frac{Area_{image}}{Area_{scene}} = m^2$$

Image Formation using Lenses

- Lenses are used to avoid problems with pinholes.
- Ideal Lens: Same projection as pinhole but gathers more light!

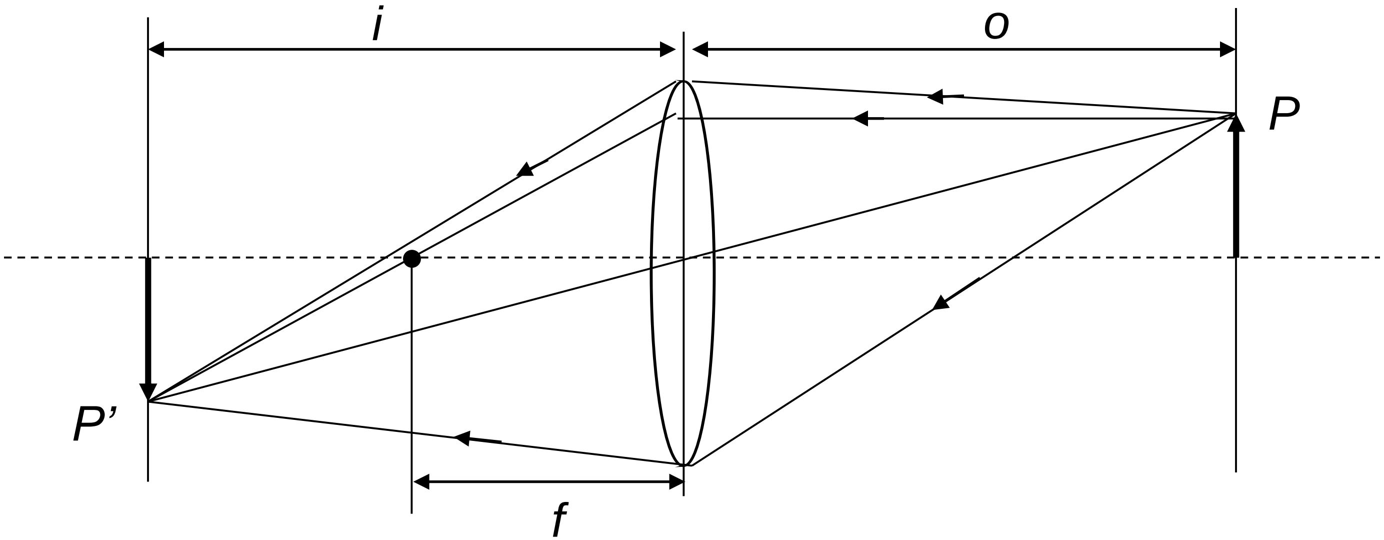

- Gaussian Thin Lens Formula: $$\frac{1}{i} + \frac{1}{o} = \frac{1}{f}$$

- f is the focal length of the lens – determines the lens's ability to refract light

Focus and Defocus

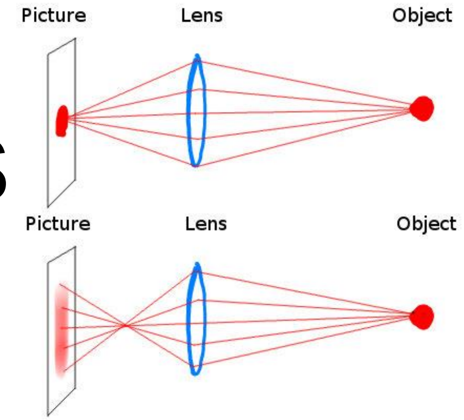

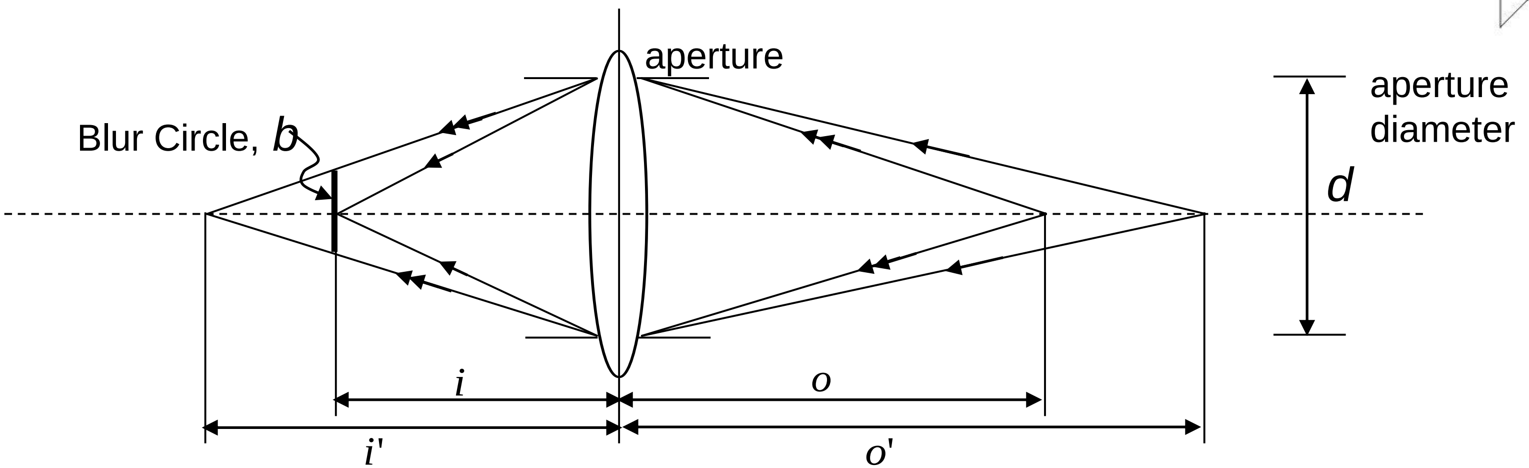

- Gaussian Law:

$$\frac{1}{i}+\frac{1}{o}=\frac{1}{f}$$

$$\frac{1}{i'}+\frac{1}{o'}=\frac{1}{f}$$

$$\Rightarrow \quad (i'-i)=\frac{f}{(o'-f)}\frac{f}{(o-f)}(o-o')$$

- In theory, only one scene plane is in focus.

Depth of Field

- Range of object distances over which image is <u>sufficiently well</u> focused.
- Range for which *blur circle* is less than the resolution of the sensor.

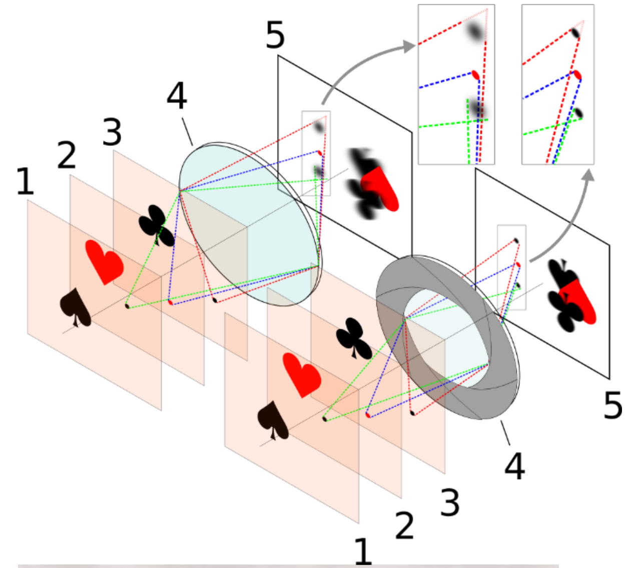

http://images.dpchallenge.com/images_portfolio/27920/print_preview/116336.jpg

Image Sensors

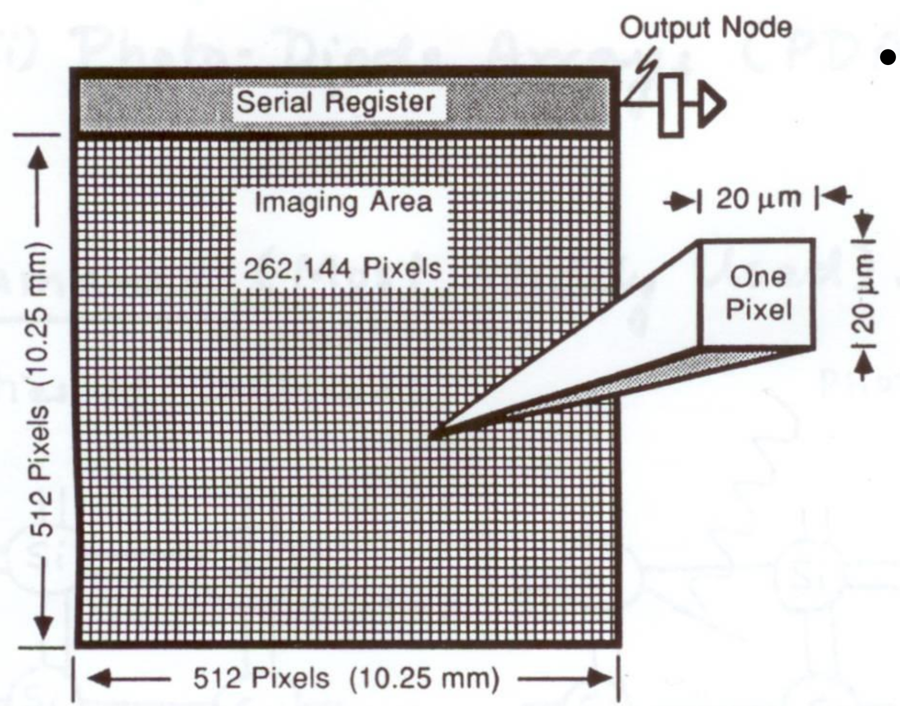

FIG. 4. Typical 512 × 512 CCD.

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

Image Sensors

- Convert light into an electric charge

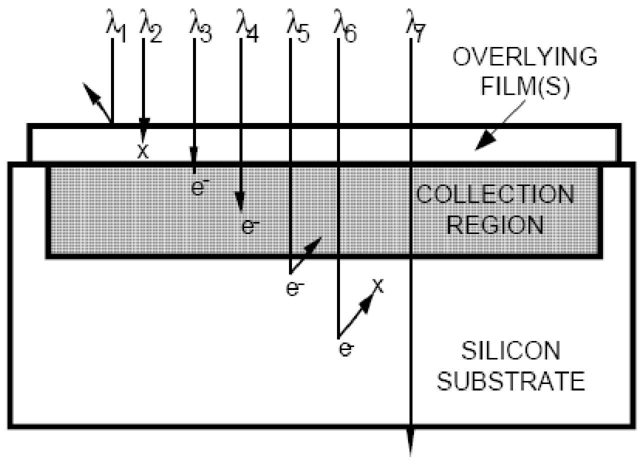

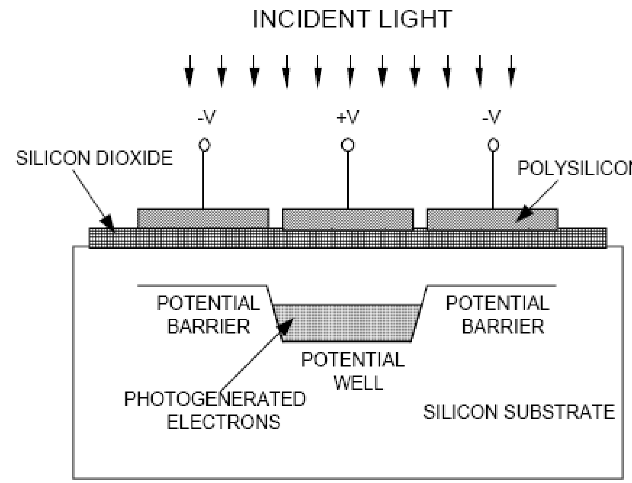

CCD (charge coupled device)

Higher dynamic range

High uniformity

Lower noise

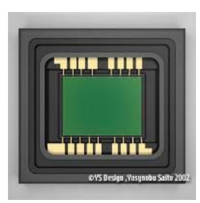

CMOS (complementary metal Oxide semiconductor)

Lower voltage

Higher speed

Lower system complexity

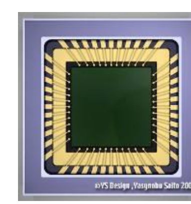

Sensing Brightness

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

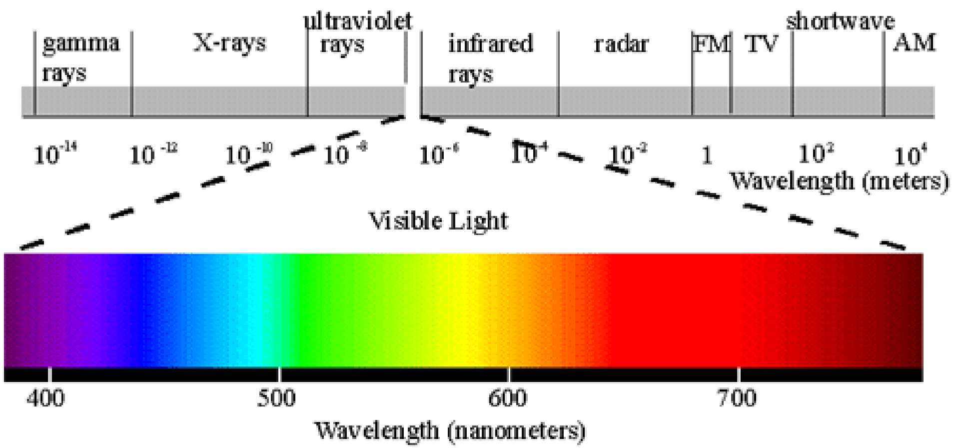

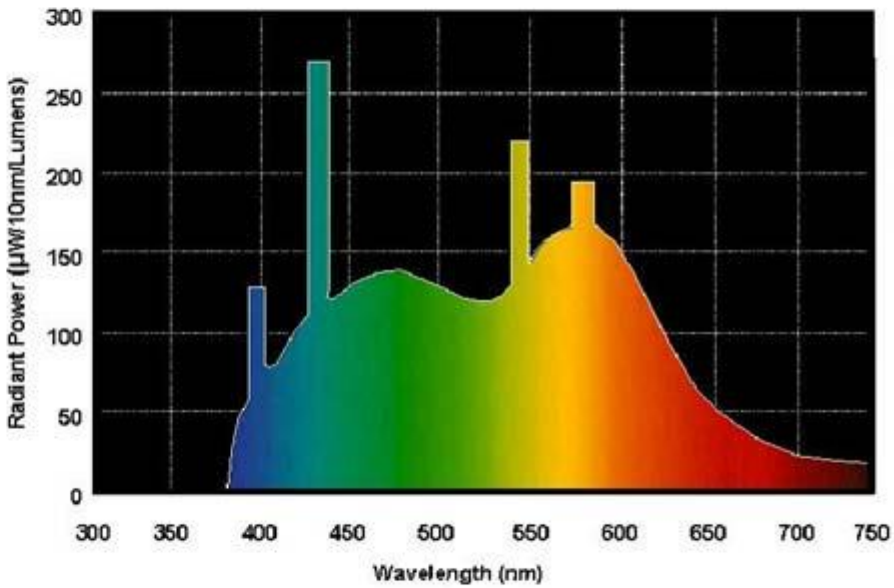

So the pixel intensity becomes

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

How do we sense colour?

- Do we have infinite number of filters?

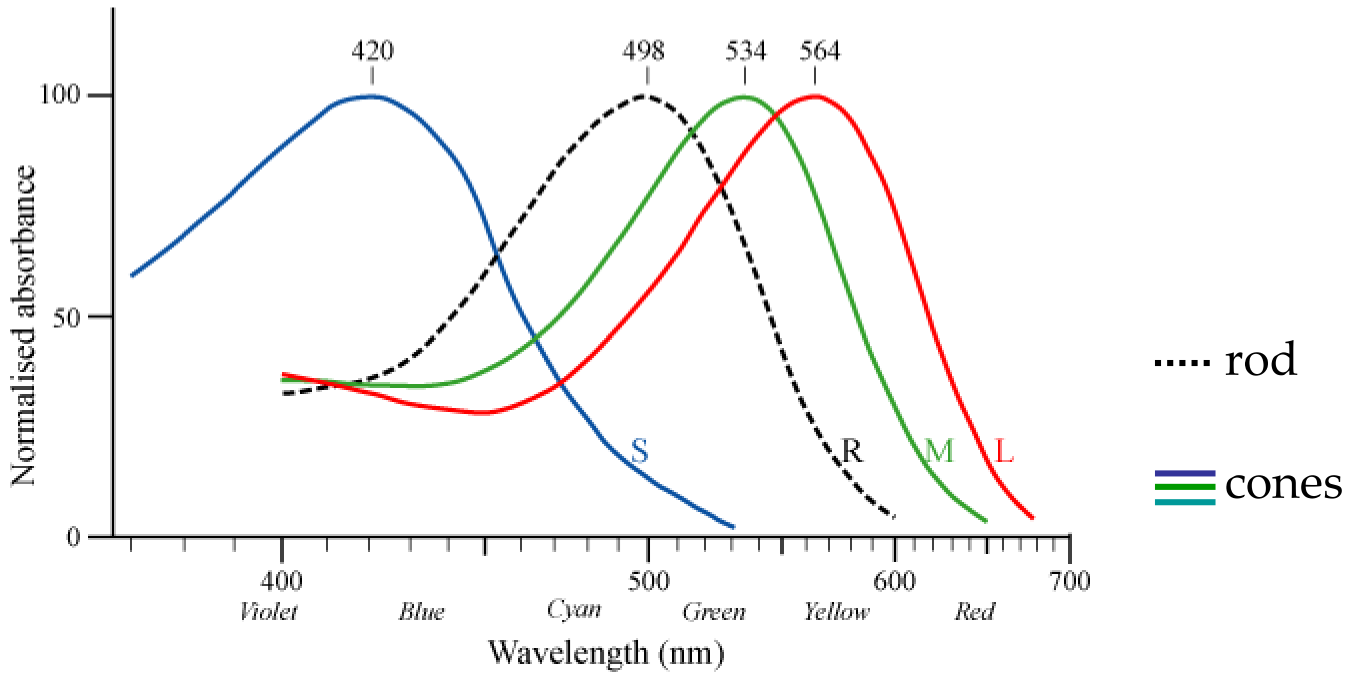

Three filters of different spectral responses

Sensing Colour

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

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

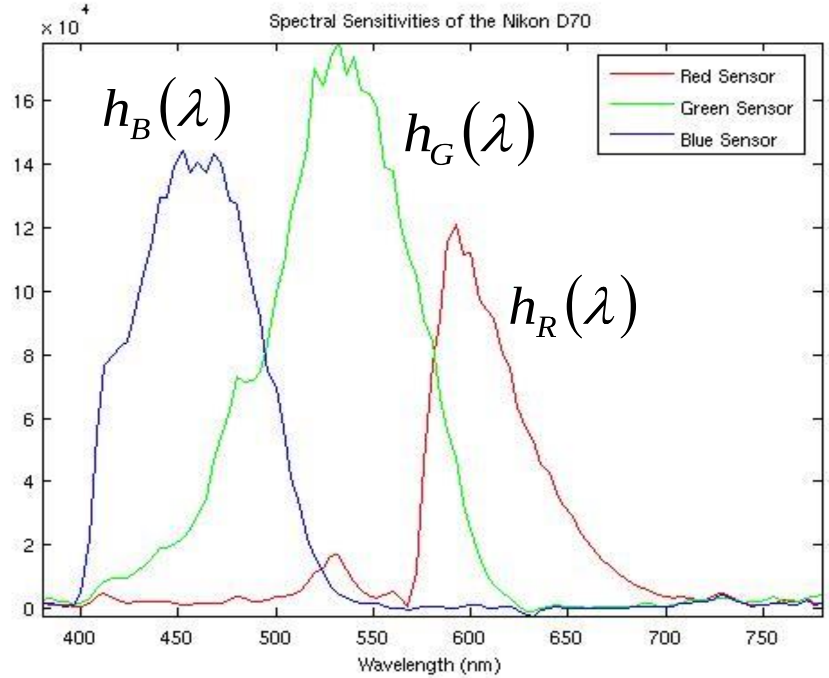

$$I_R = k \int_{-\infty}^{\infty} h_R(\lambda) p(\lambda) d\lambda$$

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

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

Sensing Colour

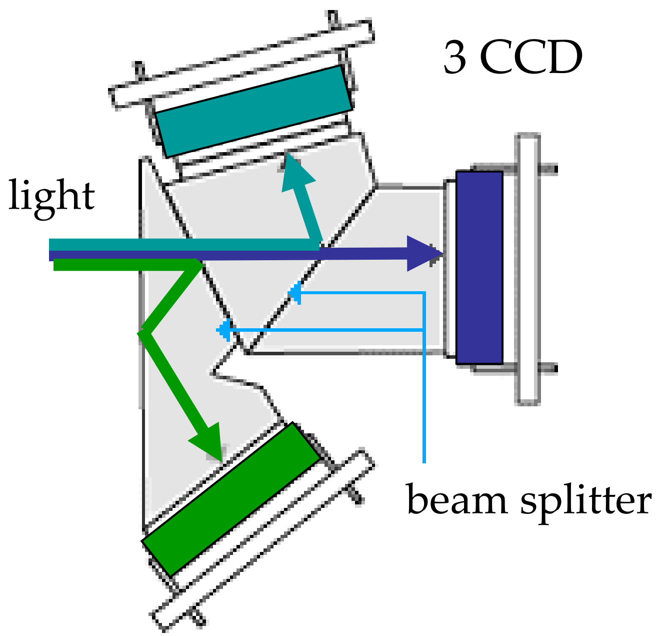

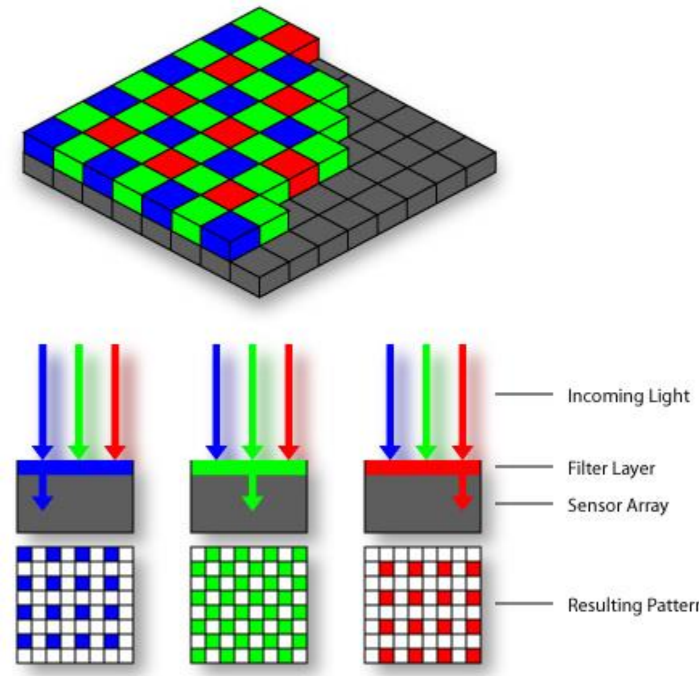

Bayer pattern

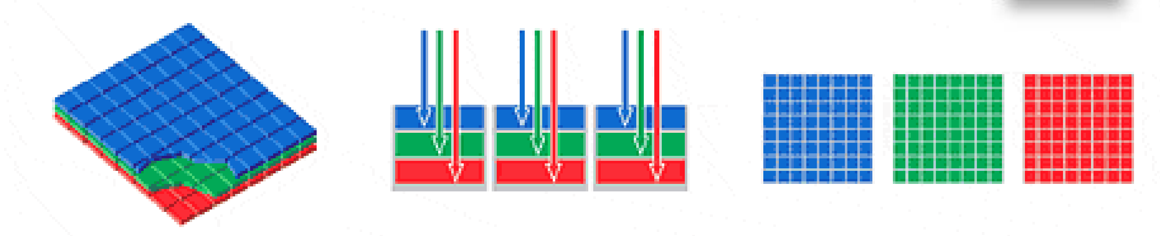

Foveon X3^{TM}

Outline

- Image Formation
- Digital Images
  - Sampling
  - Data Structures
  - Histograms
- Colour and Noise

Components of a Computer Vision System

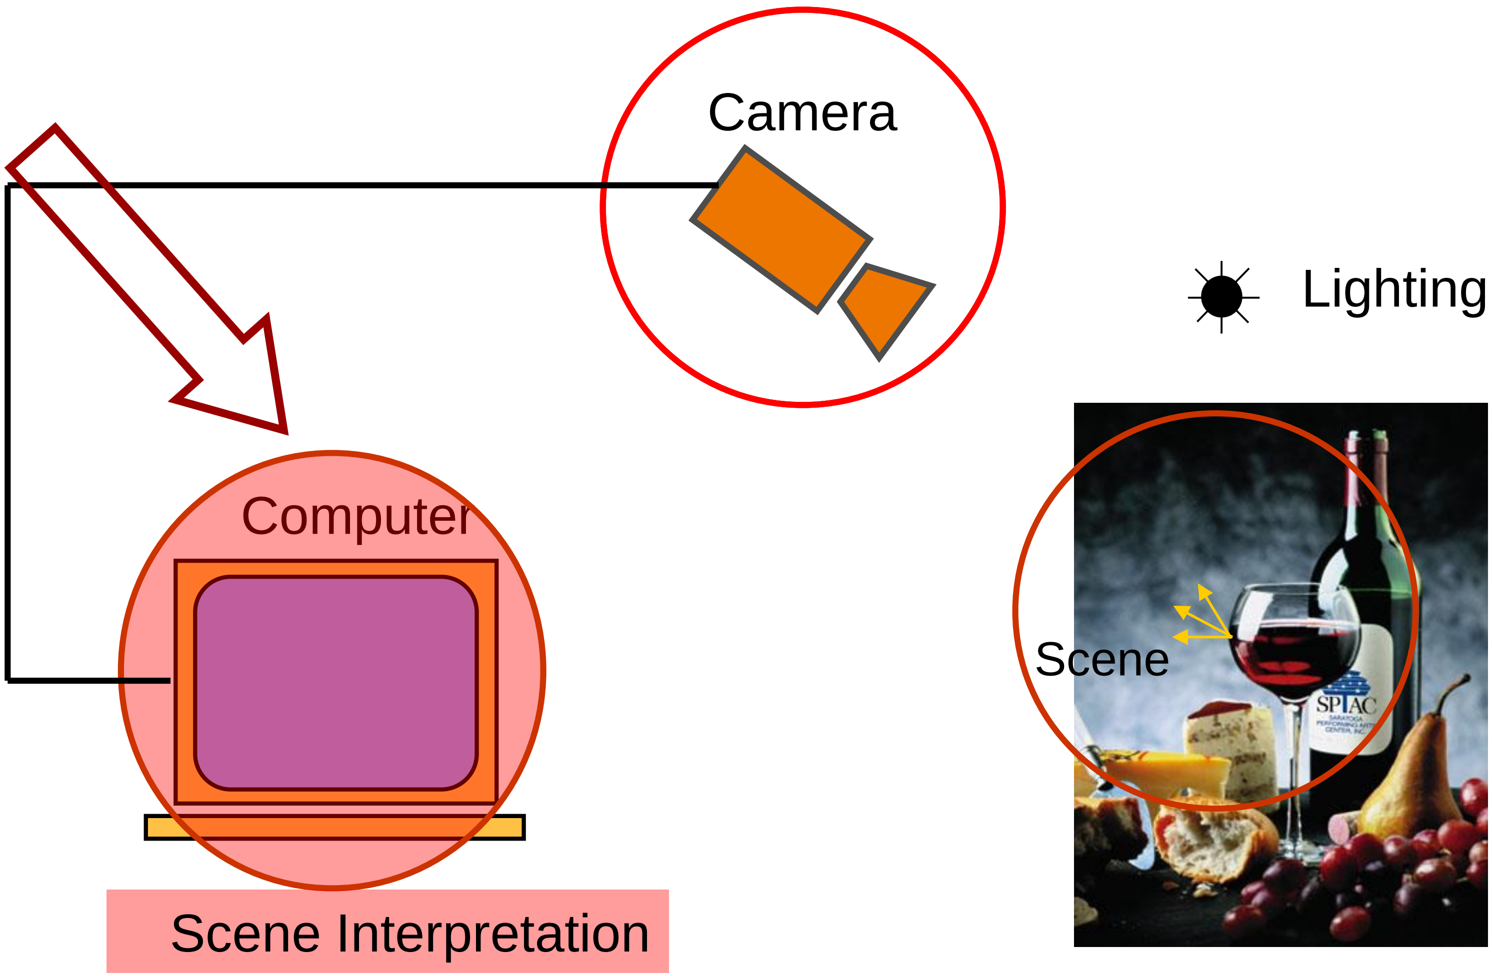

Digital Images

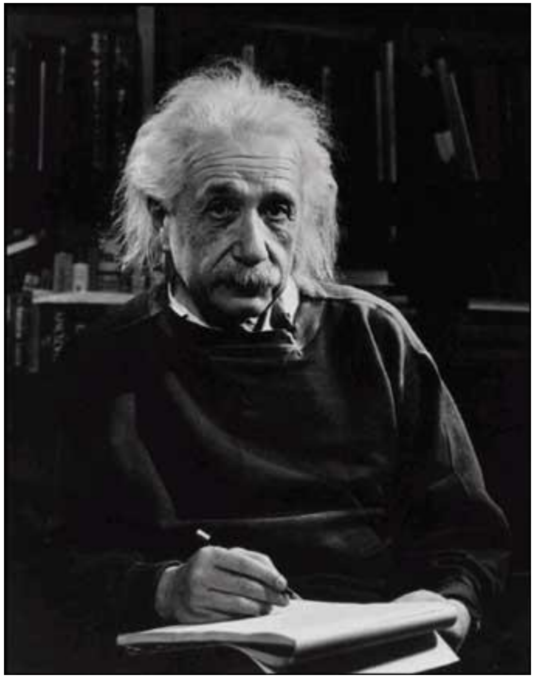

What we see

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

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

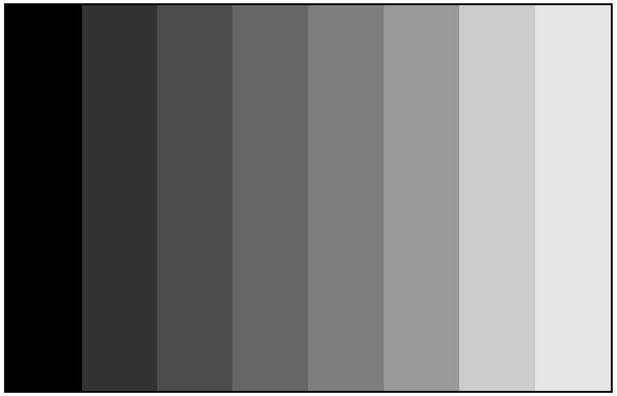

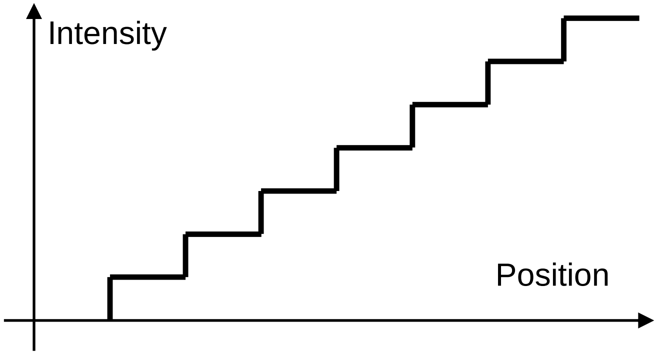

[Gonzalez & Woods]

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)

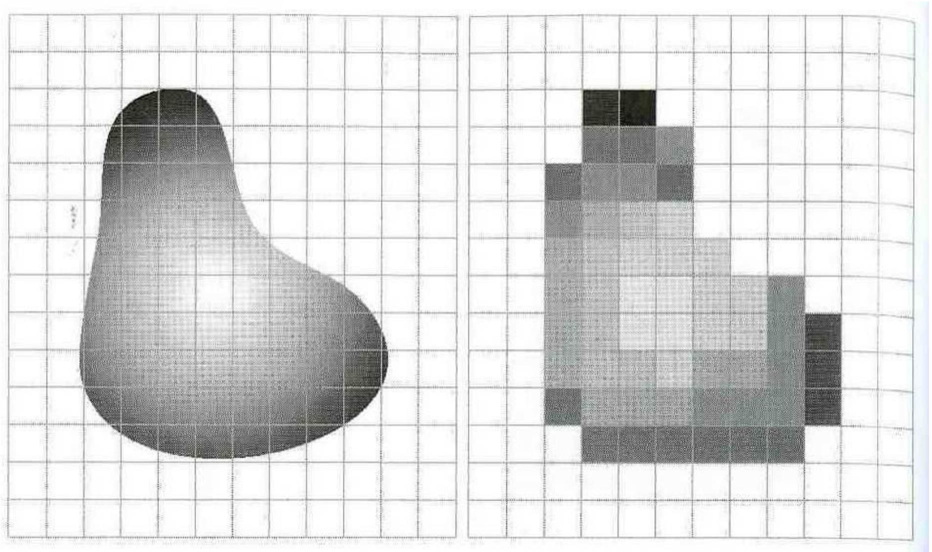

$$f(i, j) = \text{Quantize}\{f(i\Delta, j\Delta)\}$$

Images as Matrices

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

I = [(0,0) (0,1) …
[(1,0) (1,1) …
…

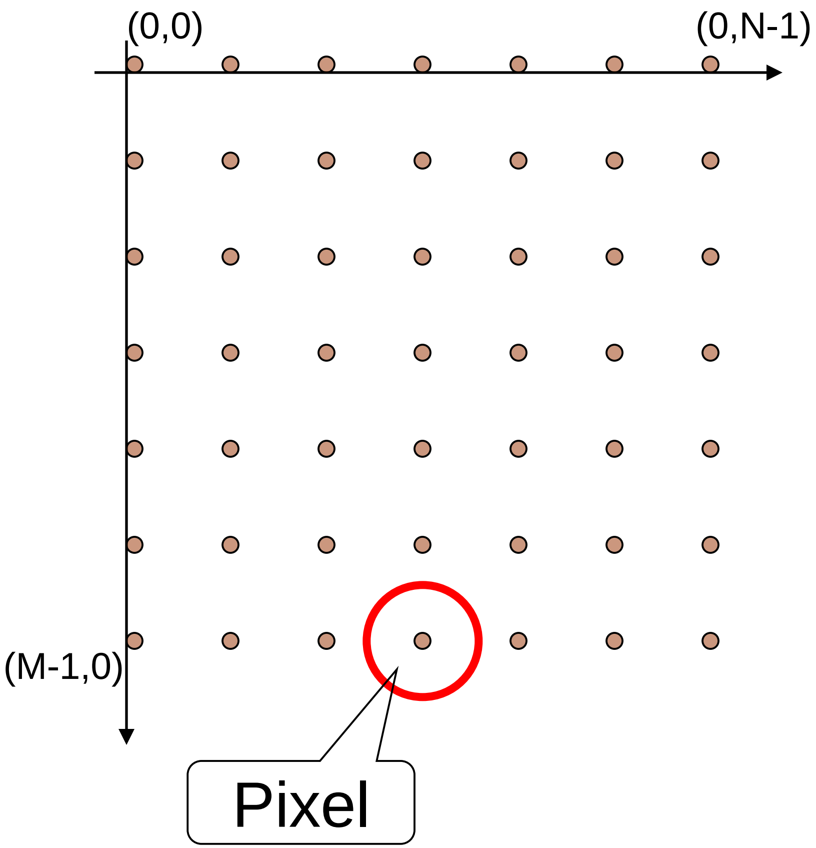

Sampling Theorem

Continuous signal:

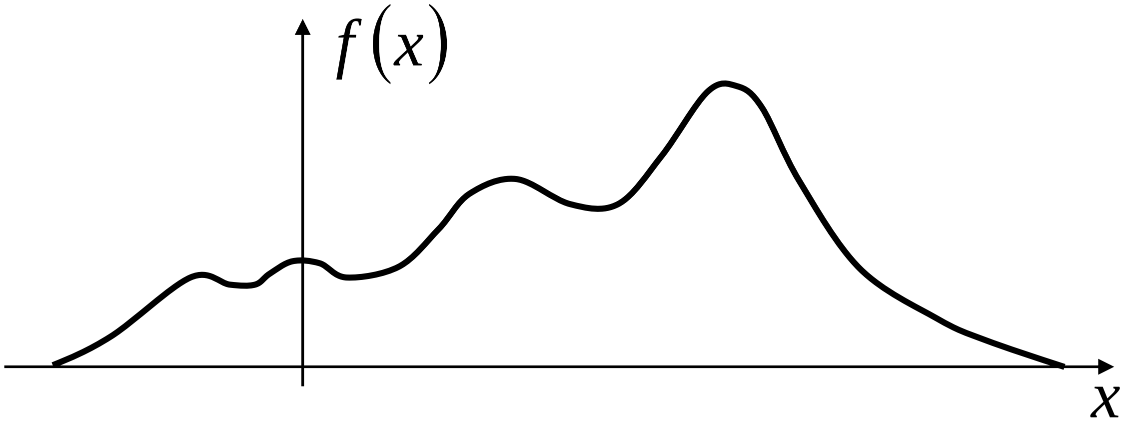

In the field of digital signal processing, the sampling theorem is a fundamental bridge between continuous-time signals (often called "analog signals") and discrete-time signals (often called "digital signals"). It establishes a sufficient condition for a sample rate that permits a discrete sequence of samples to capture all the information from a continuous-time signal of finite bandwidth.

Sampling Theorem

Continuous signal:

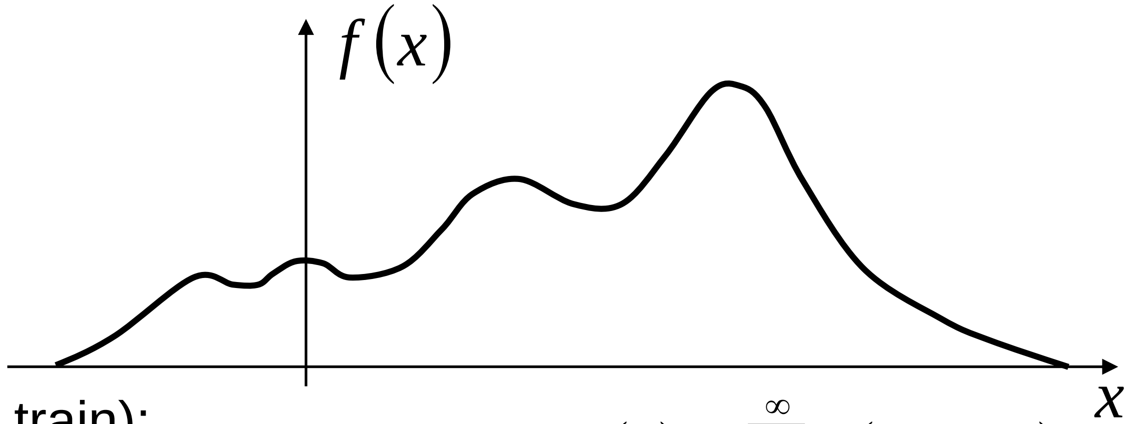

Shah function (Impulse train):

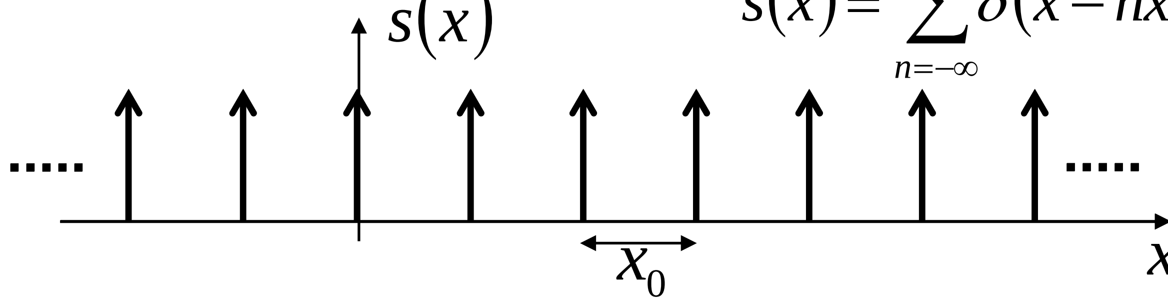

$$s(x) = \sum_{n=-\infty}^{\infty} \delta(x - nx_0)$$

Sampled function:

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

Sampling Theorem

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

Sampling frequency $\frac{1}{x_0}$

$$F_S(u) = F(u) * S(u) = F(u) * \frac{1}{x_0}\sum_{n=-\infty}^{\infty}\delta\left(u - \frac{n}{x_0}\right)$$

The Fourier transform decomposes a function of time (a signal) into the frequencies that make it up

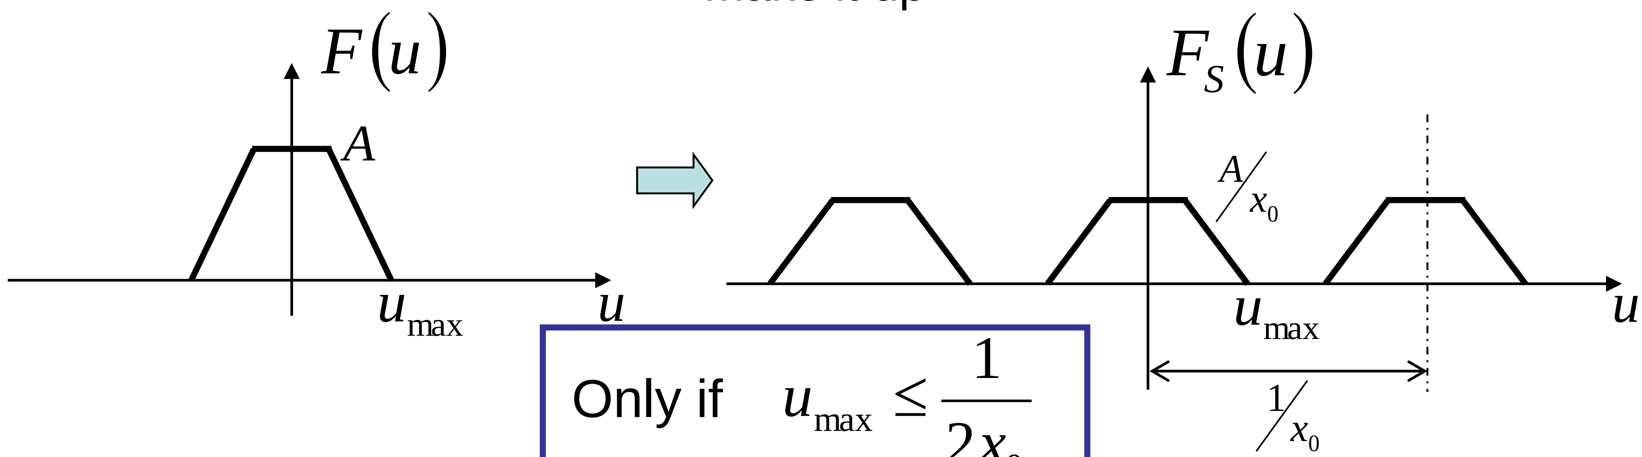

Only if $u_{max} \leq \frac{1}{2x_0}$

Nyquist Theorem

If $u_{\max} > \dfrac{1}{2x_0}$

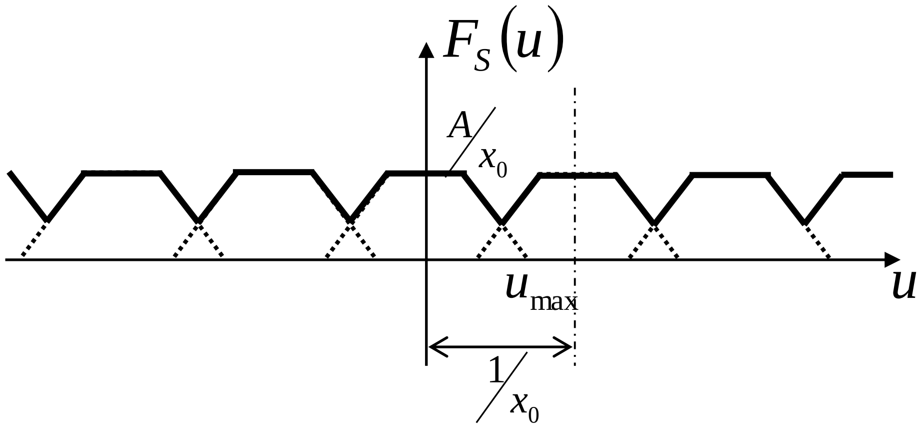

Aliasing

Sampling frequency must be greater than $2u_{\max}$

Quantization

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

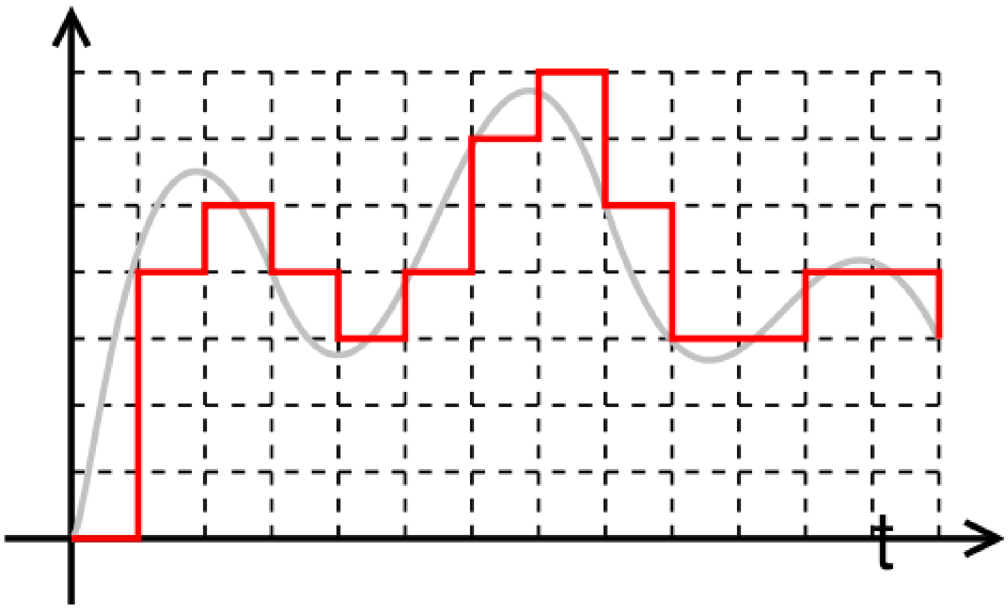

Quantization Levels

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

$$G = 2^m$$

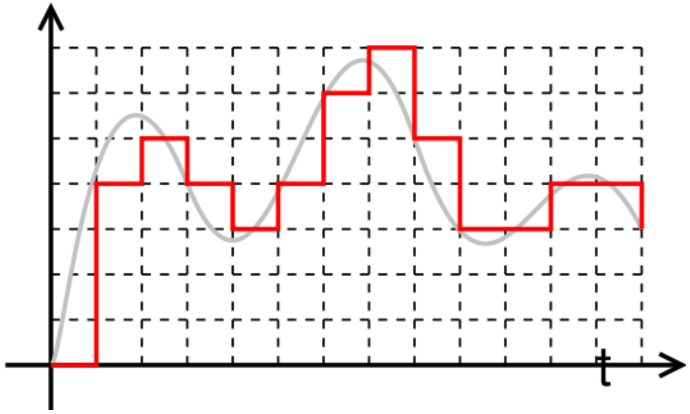

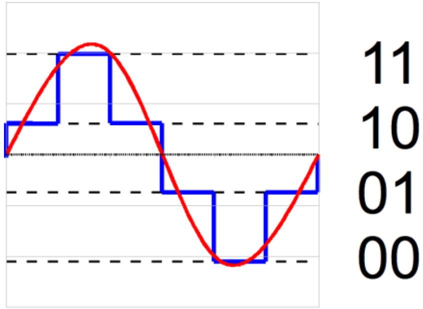

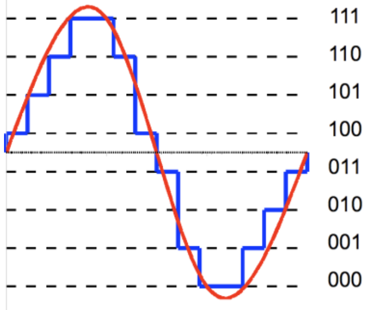

Effect of quantization

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?

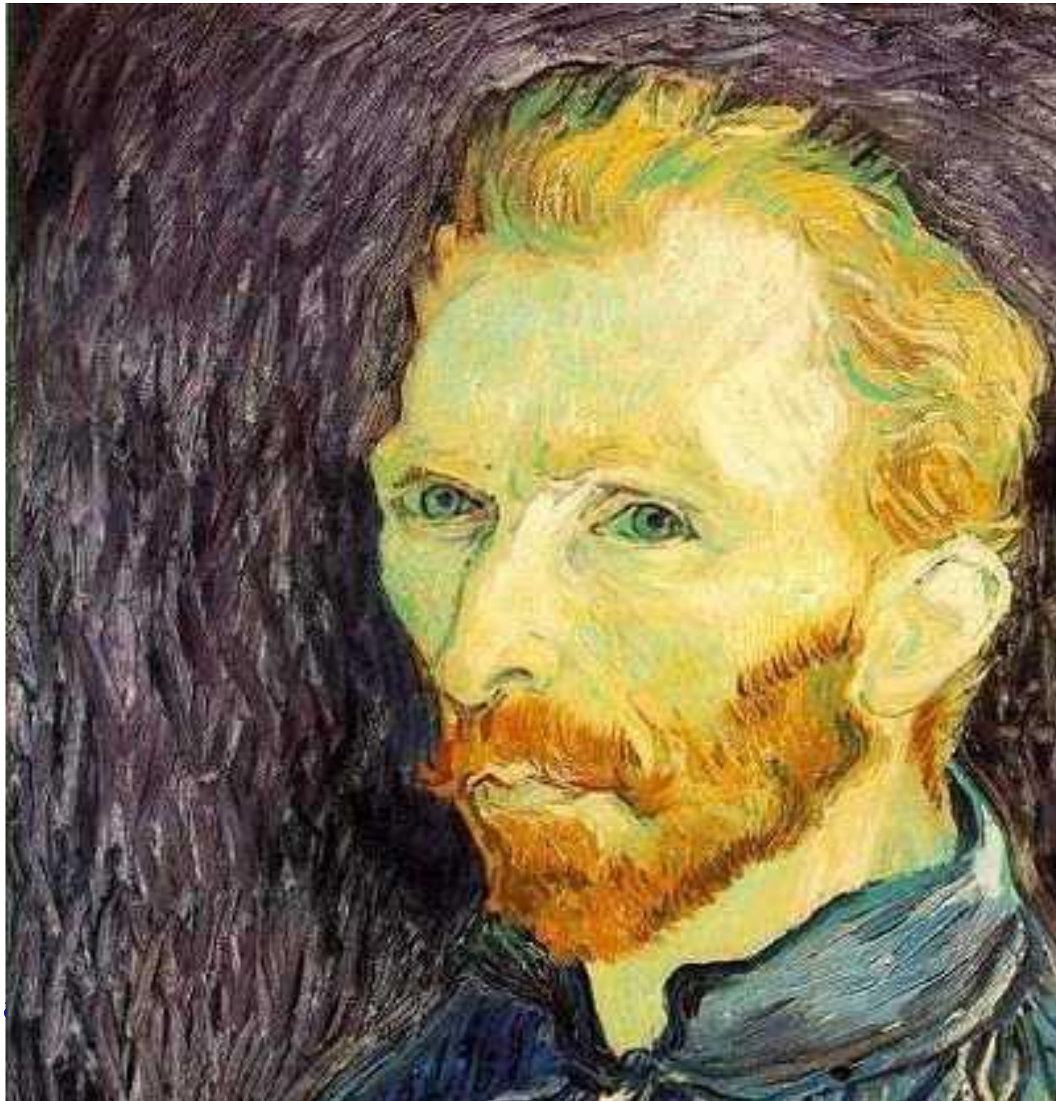

Ma

DOCTORAL PROGRAMME IN COMPUTER SCIENCE

Sub-sampling

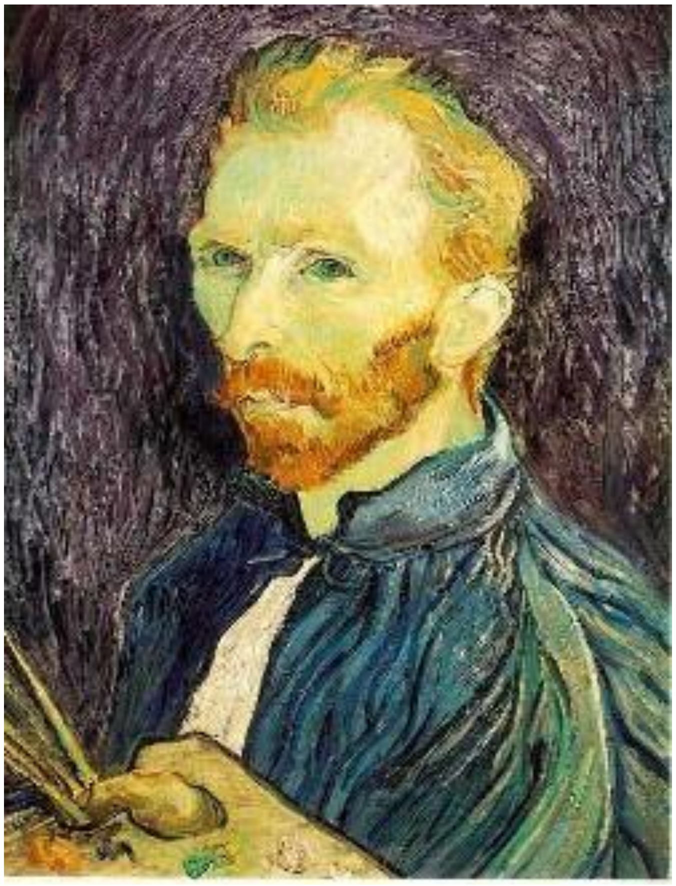

1/4

1/8

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

Sub-sampling

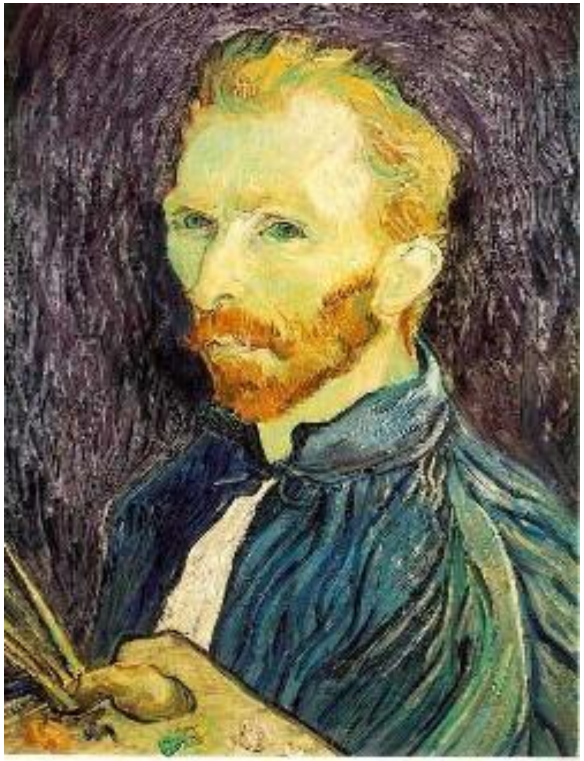

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Can we do better?
https://www.youtube.com/watch?v=6NcIJXTlugc

Outline

- Image Formation
- Digital Images
  - Sampling
  - Data Structures
  - Histograms
- Colour and Noise

Data Structures for Digital Images

- Are there other ways to represent digital images?

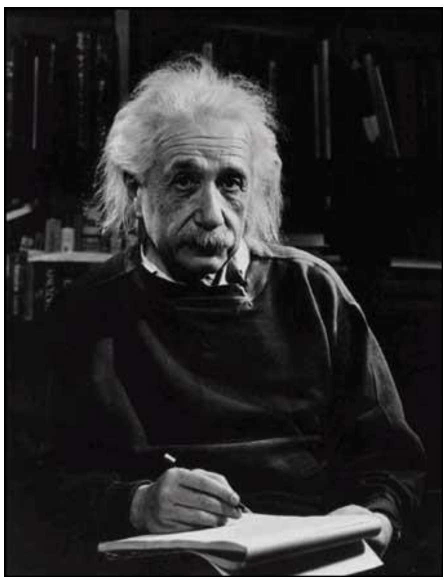

What we see

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!

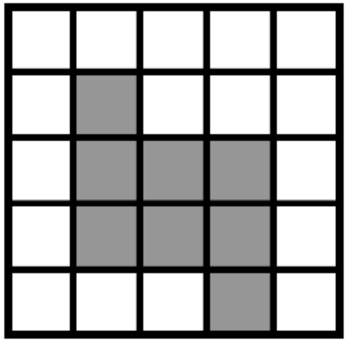

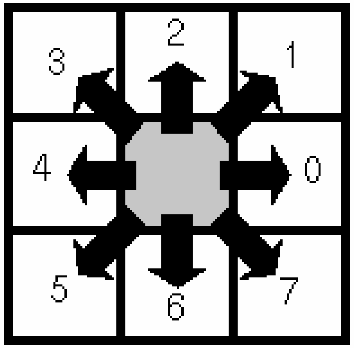

Freeman Chain Code

Using a Freeman Chain Code and considering the top-left pixel 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!

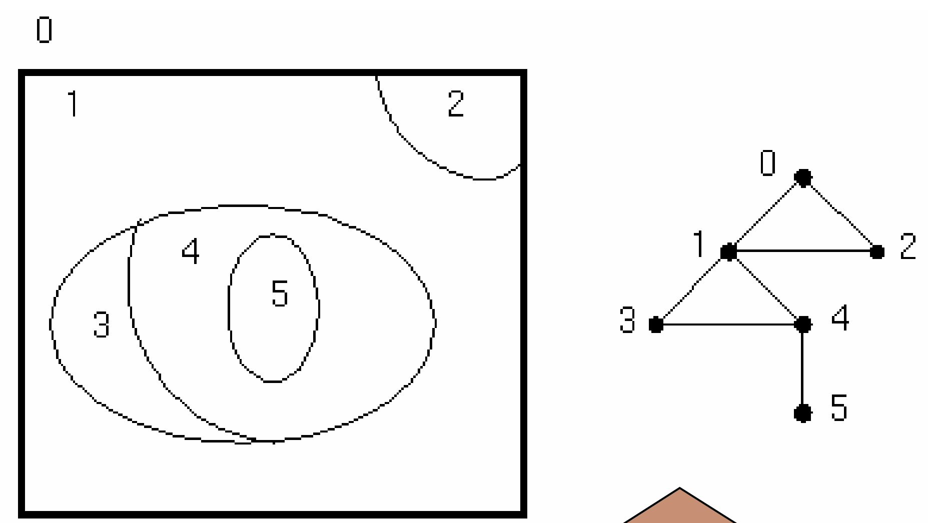

Region Adjacency Graph

Relational Structures

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

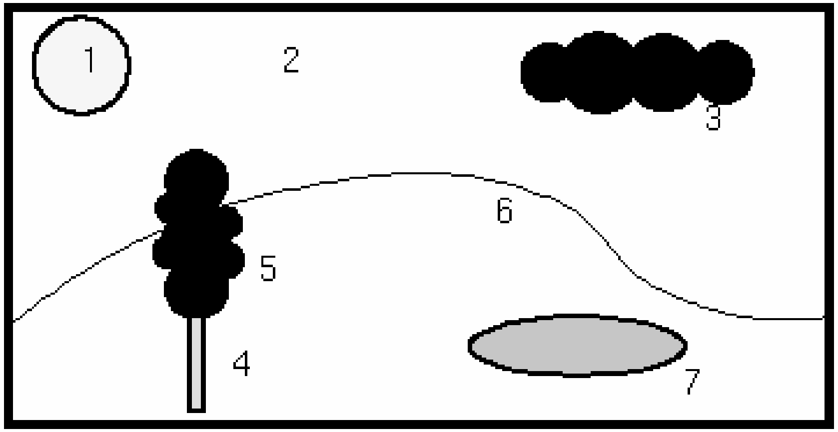

No.	Object name	Colour	Min. row	Min. col.	Inside
1	sun	white	5	40	2
2	sky	blue	0	0	-
3	cloud	grey	20	180	2
4	tree trunk	brown	95	75	6
5	tree crown	green	53	63	-
6	hill	light green	97	0	-
7	pond	blue	100	160	6

Relational Table

Outline

- Image Formation
- Digital Images
  - Sampling
  - Data Structures
  - Histograms
- Colour and Noise

Histograms

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

  [Wikipedia]

- Typically represented as a bar chart:

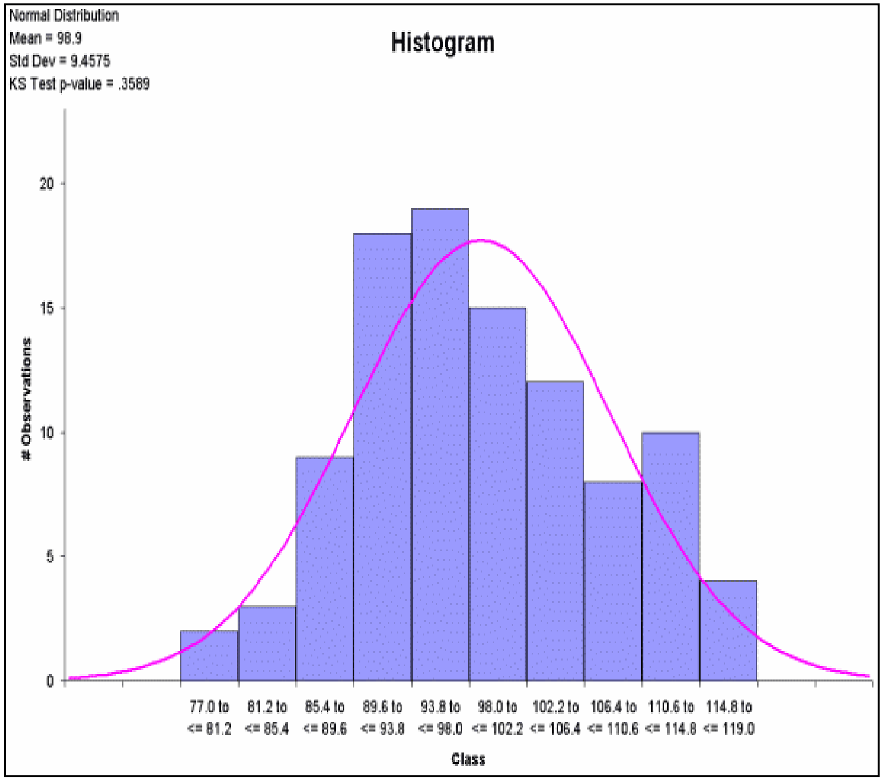

Image Histograms

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

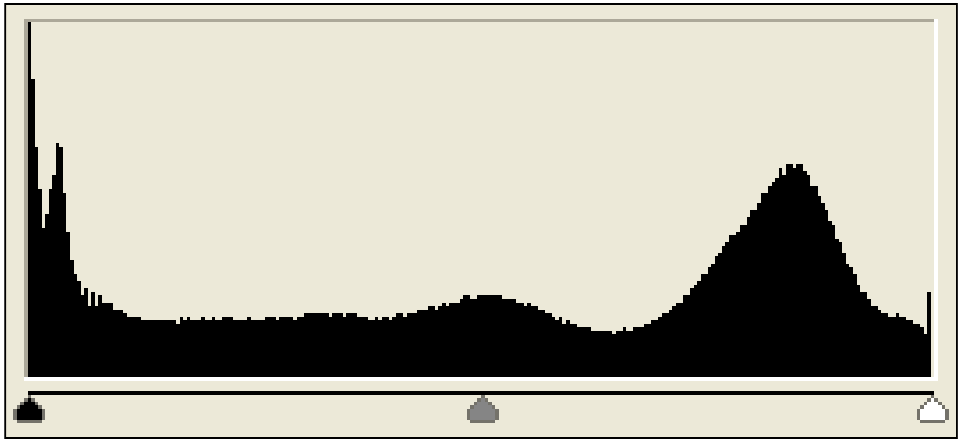

Colour Histogram

- As many histograms as axis of the colour space.

  Ex: RGB Colour space

  - Red Histogram
  - Green Histogram
  - Blue Histogram

- Combined histogram.

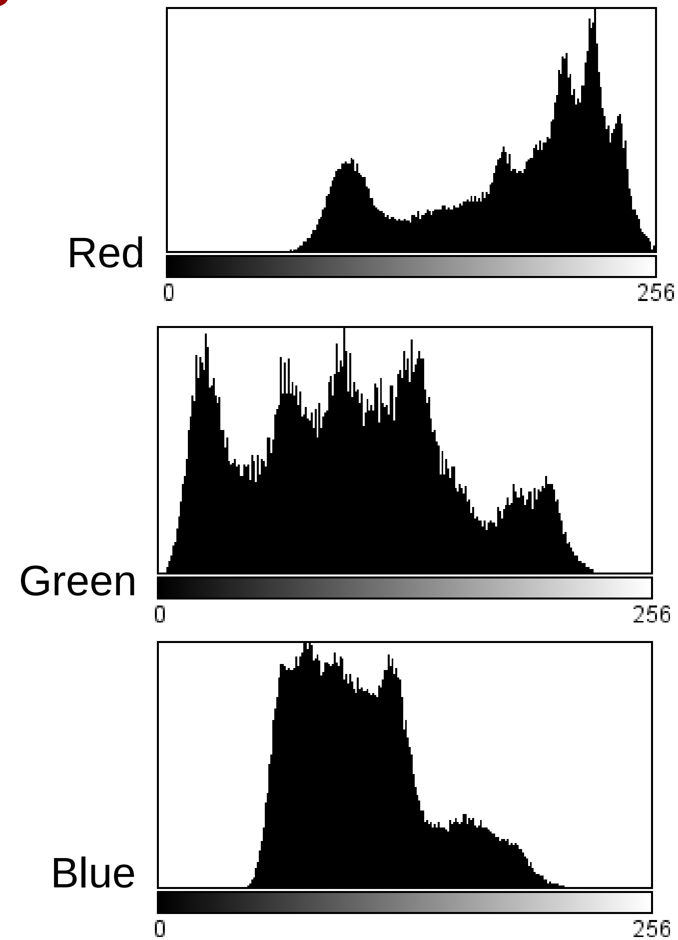

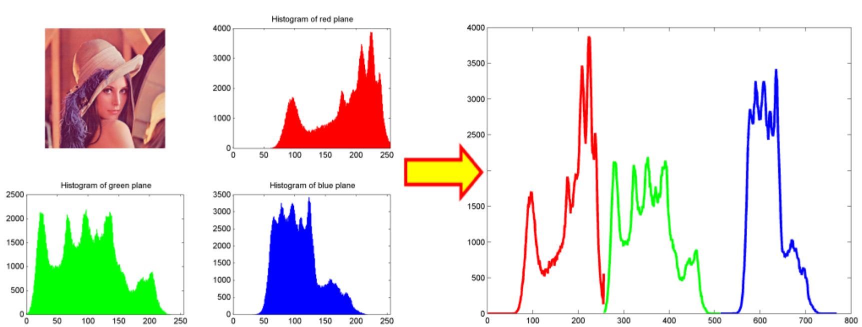

Outline

- Image Formation
- Digital Images
- Colour and Noise
  - Colour spaces
  - Colour processing
  - Noise

What is colour?

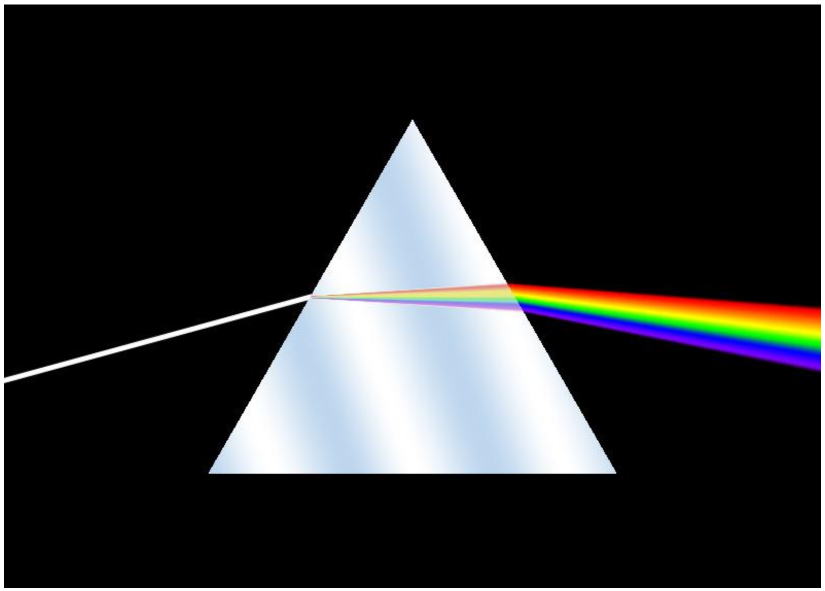

Optical Prism dispersing light

Visible colour spectrum

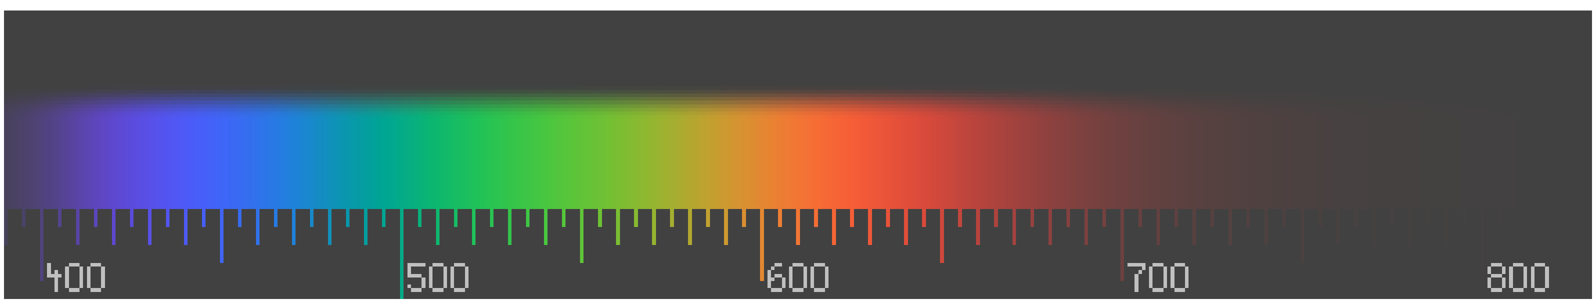

Visible Spectrum

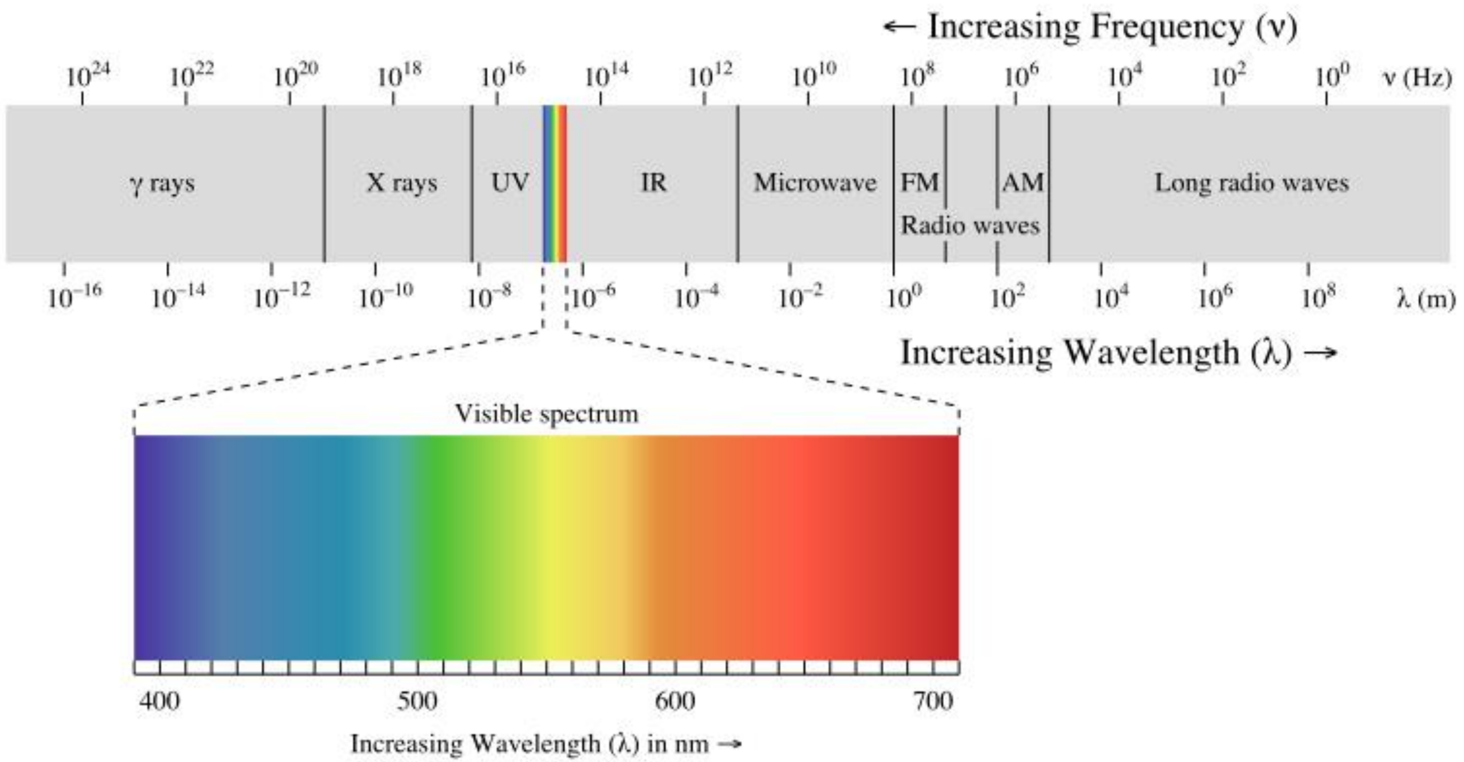

http://science.howstuffworks.com/light.htm

How do we see colour?

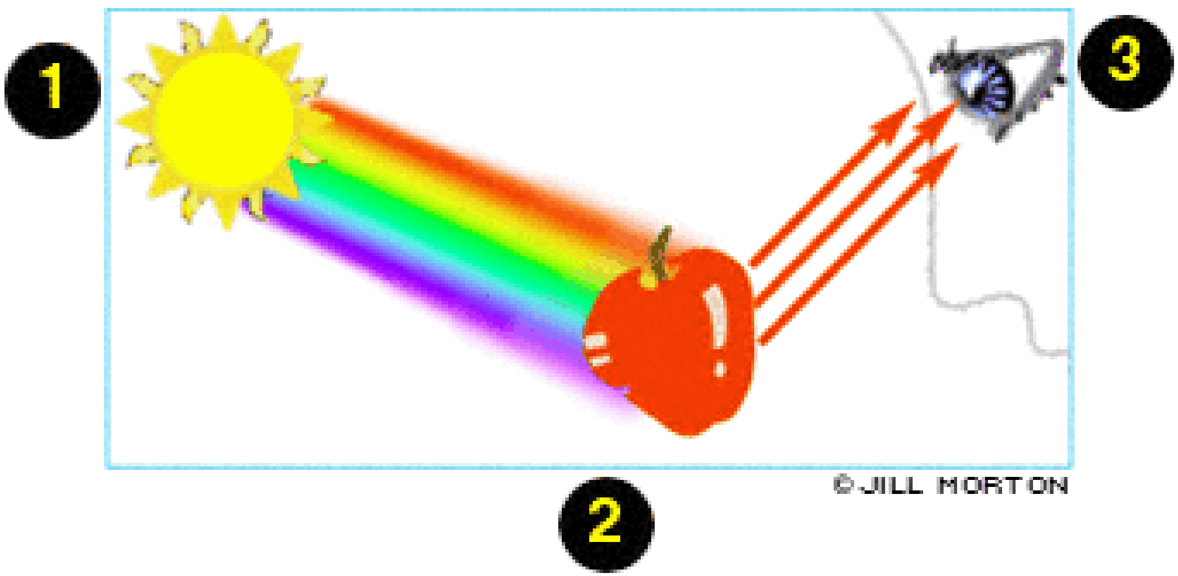

Primary Colours

- Not a fundamental property of light.
- Based on the physiological response of the human eye.
- Form an additive colour system.

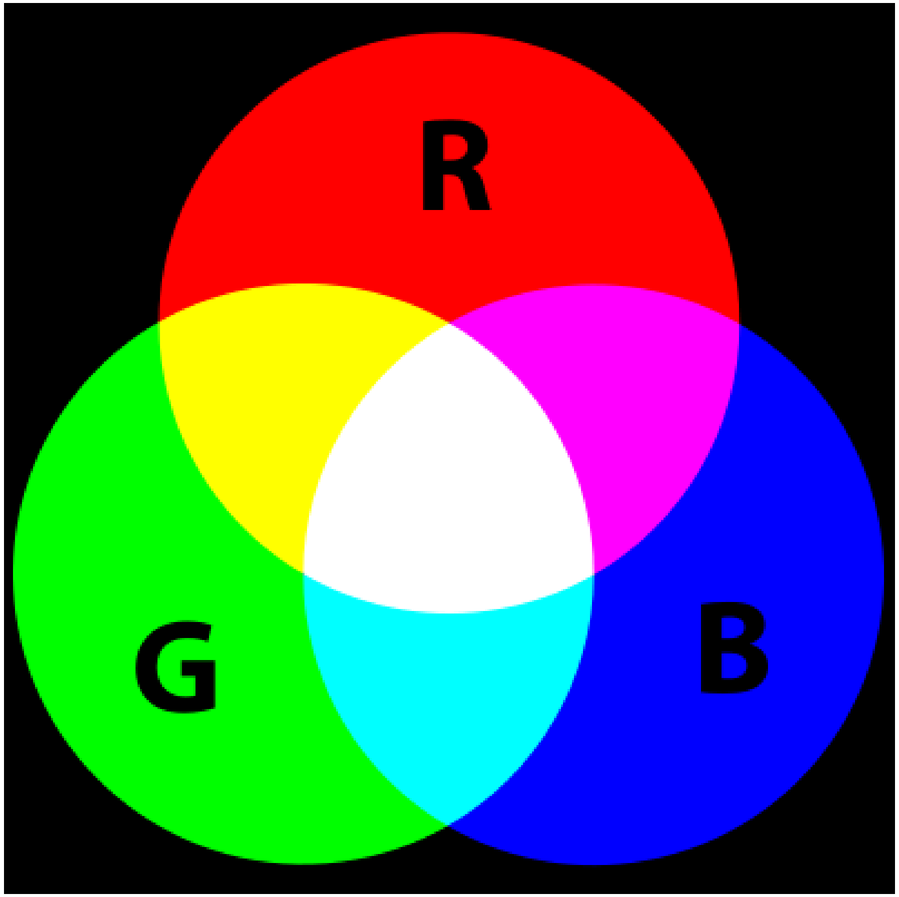

Colour Space

- “The purpose of a color model is to facilitate the specification of colours in some standard, generally accepted way”

  Gonzalez & Woods

- Colour space
  - Coordinate system
  - Subspace: One colour -> One point

RGB

- Red Green Blue
- Defines a colour cube.
- Additive components.
- Great for image capture.
- Great for image projection.
- Poor colour description.

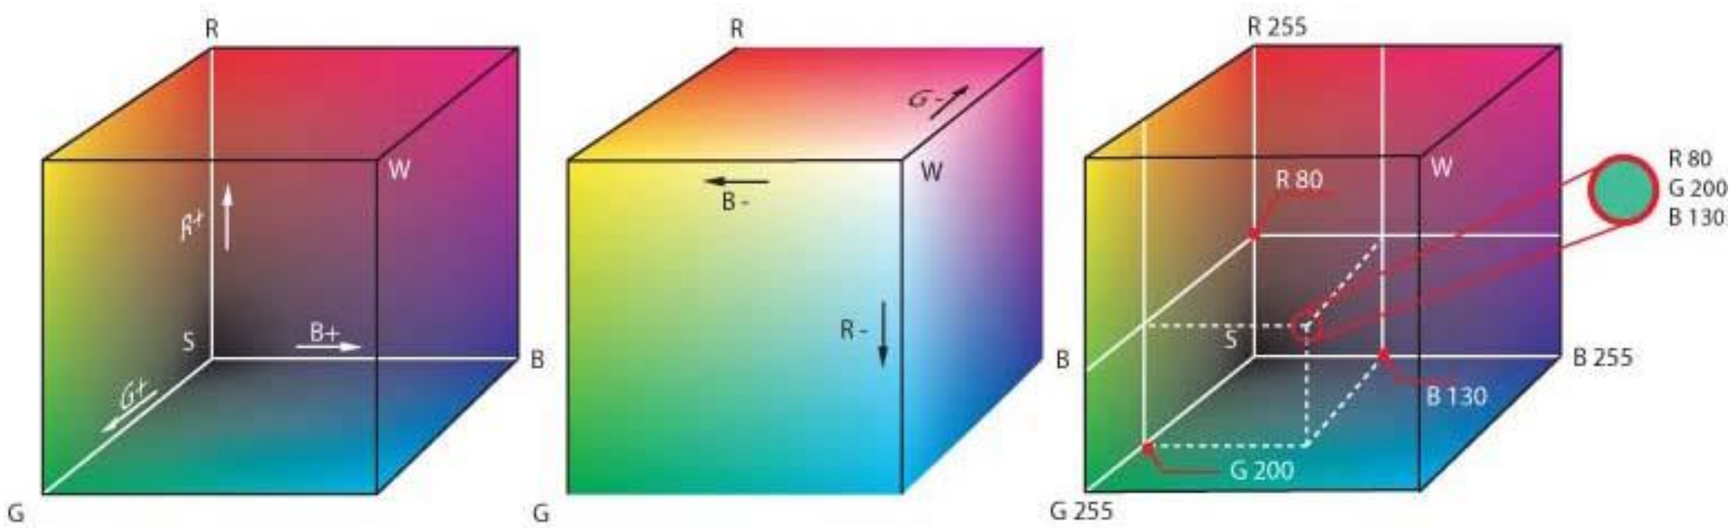

CMYK

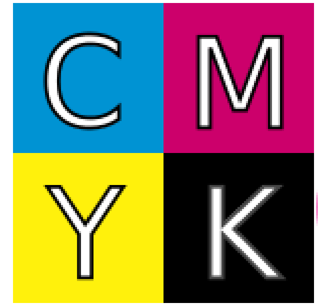

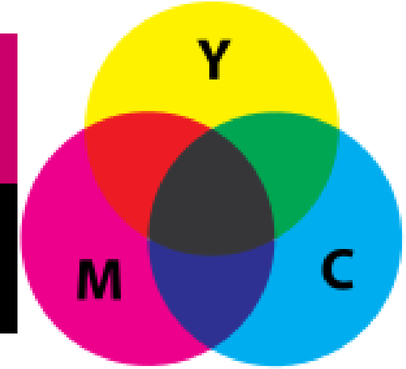

- Cyan Magenta Yellow Key.
- Variation of RGB.
- Technological reasons: great for printers.
- Subtractive model

CMYK

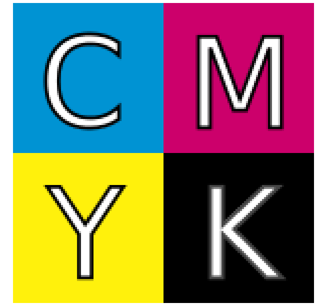

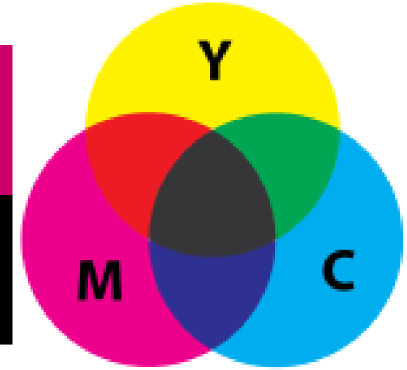

- In additive color models such as RGB, white is the "additive" combination of all primary colored lights, while black is the absence of light.
- In the CMYK model, it is the opposite: white is the natural color of the paper or other background, while black results from a full combination of colored inks.
- To save cost on ink, and to produce deeper black tones, unsaturated and dark colors are produced by using black ink instead of the combination of cyan, magenta and yellow.

HSI (HSV)

- Hue Saturation Intensity (Value)
- Defines a colour cone
- Great for colour description.
- cylindrical-coordinate representation of RGB
- Attempt to be more intuitive and perceptually relevant than cube

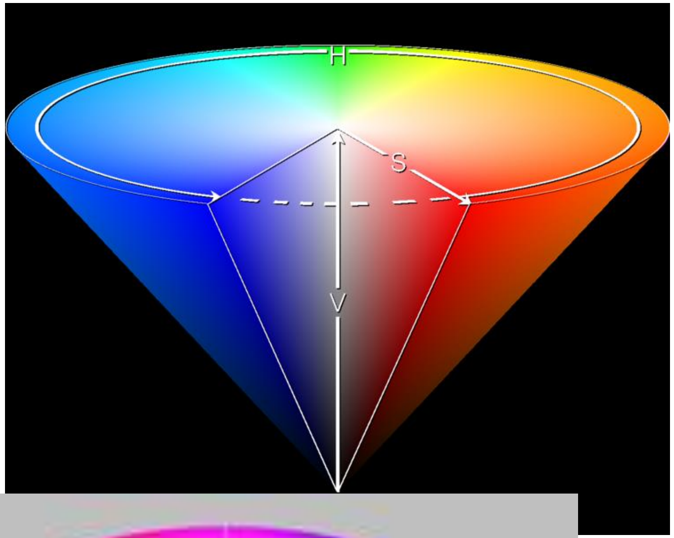

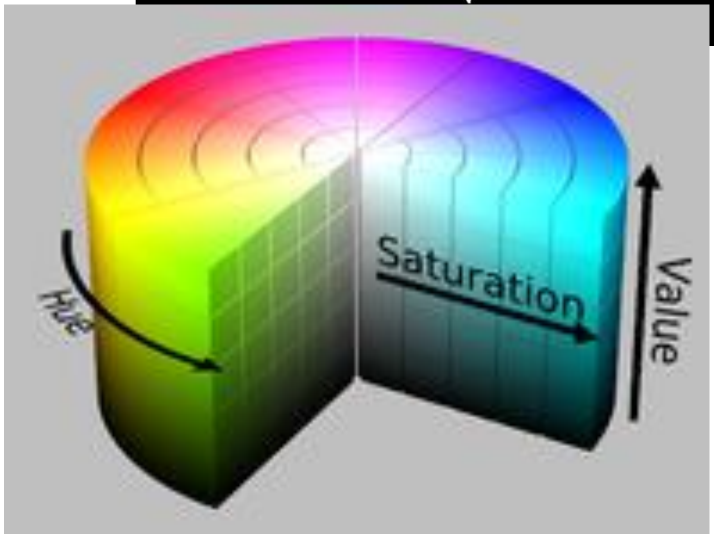

HSI (HSV)

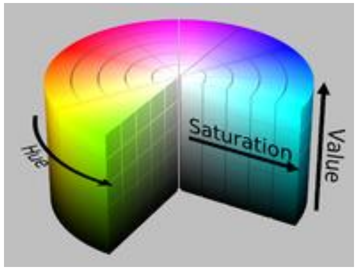

- The *hue* (H) of a color refers to which pure color it resembles. All tints, tones and shades of red have the same hue.
- Hues are described by a number that specifies the position of the corresponding pure color on the color wheel, as a fraction between 0 and 1. Value 0 refers to red; 1/6 is yellow; 1/3 is green; and so forth around the color wheel.

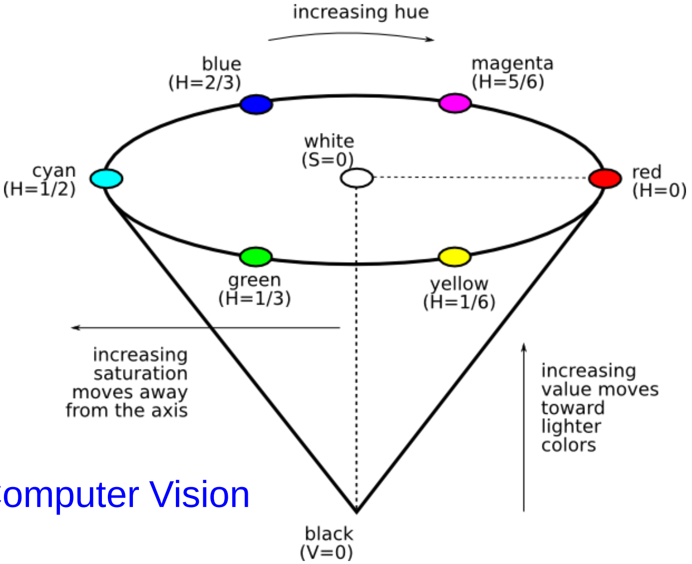

HSI (HSV)

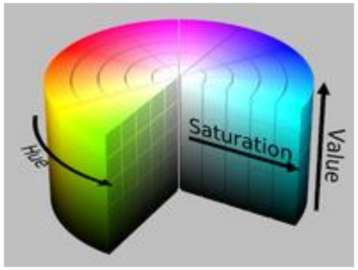

- The *saturation* (S) of a color describes how white the color is. A pure red is fully saturated, with a saturation of 1; tints of red have saturations less than 1; and white has a saturation of 0.

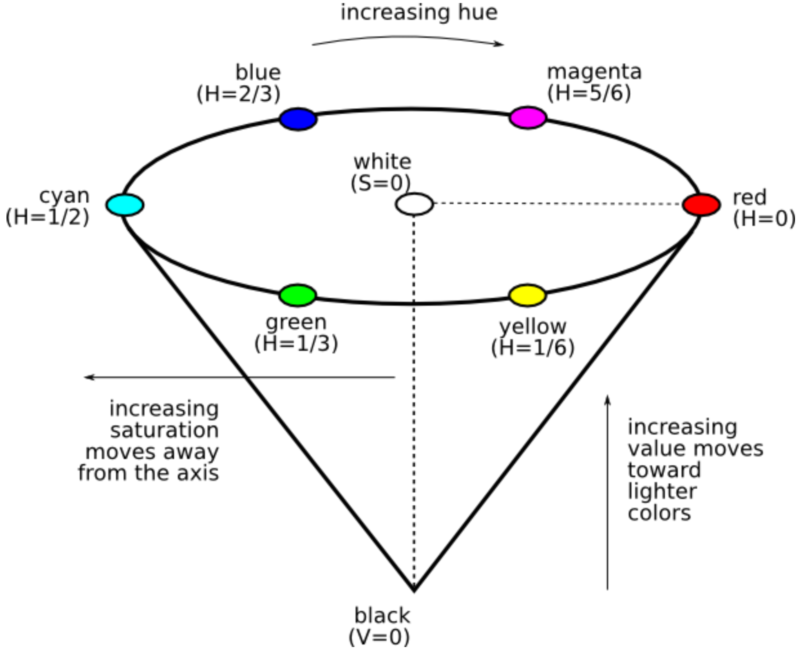

HSI (HSV)

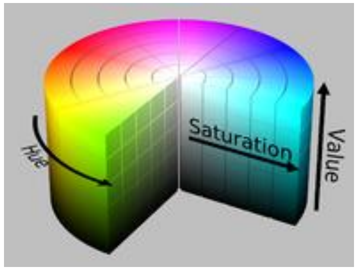

- The *value* (V) of a color, also called its *lightness*, describes how dark the color is. A value of 0 is black, with increasing lightness moving away from black.

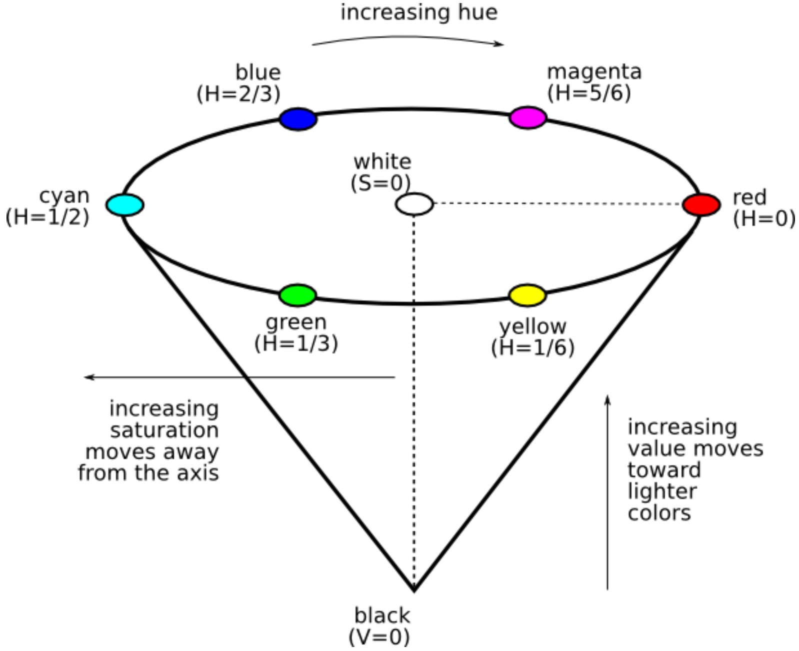

Chromaticity Diagram

- Axis:
  - Hue
  - Saturation
- Outer line represents our visible spectrum.

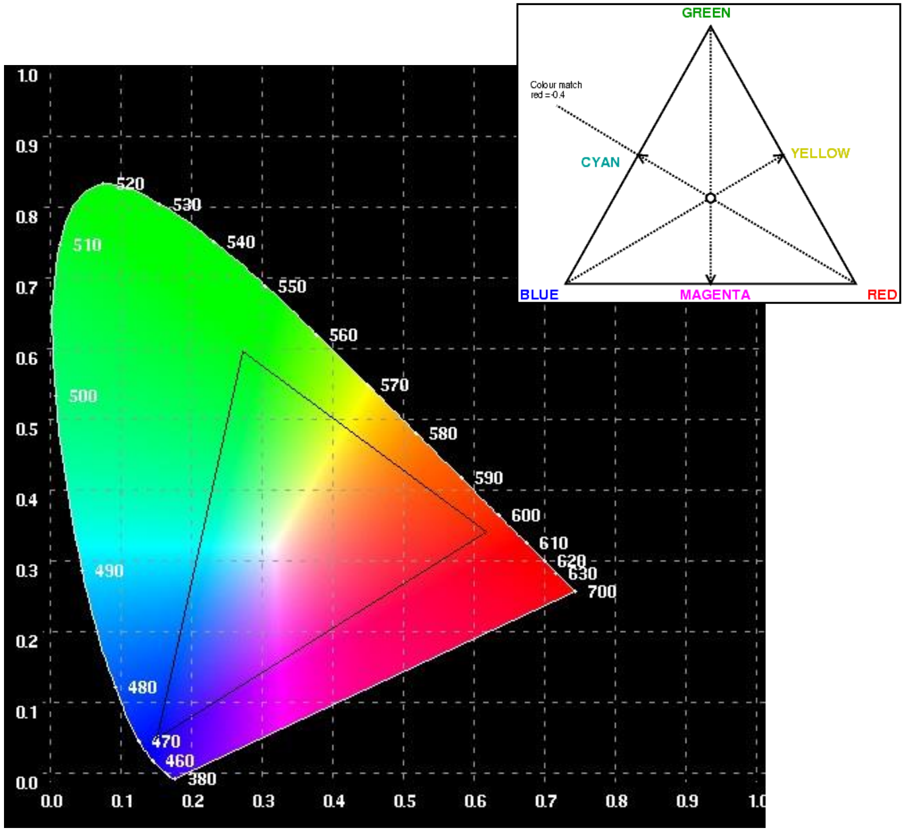

http://www.cs.rit.edu/~ncs/color/a_chroma.html

Chromaticity Diagram

- The **CIE 1931 XYZ color spaces** were the first defined quantitative links between physical pure colors (i.e. wavelengths) in the electromagnetic visible spectrum, and physiological perceived colors in human color vision. The mathematical relationships that define thesecolor spaces are essential tools for color management, important when dealing with color inks, illuminated displays, and recording devices such as digital cameras.

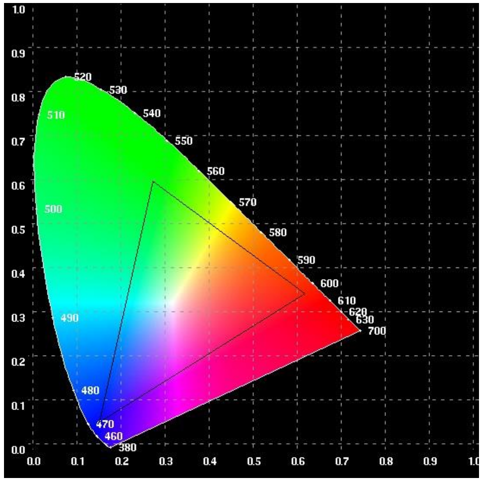

http://www.cs.rit.edu/~ncs/color/a_chroma.html

RGB to HSI

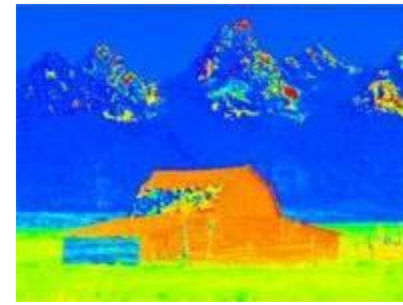

Hue:

$$H = \begin{cases} \theta & \Leftarrow B \le G \\ 360 - \theta & \Leftarrow B > G \end{cases}$$

$$\theta = \cos^{-1}\left\{ \frac{\frac{1}{2}\left[(R-G)+(R-B)\right]}{\left[(R-G)^2+(R-B)(G-B)\right]^{1/2}} \right\}$$

Saturation

$$S = 1 - \frac{3}{(R+G+B)}\left[\min(R,G,B)\right]$$

Intensity

$$I = \frac{1}{3}(R+G+B)$$

HSI to RGB

- Depends on the 'sector' of H

0 <= H < 120

$$B = I(1 - S)$$

$$R = I\left[1 + \frac{S\cos H}{\cos(60^{o} - H)}\right]$$

$$G = 3I - (R + B)$$

120 <= H < 240

$$H = H - 120^{o}$$

$$R = I(1 - S)$$

$$G = I\left[1 + \frac{S\cos H}{\cos(60^{o} - H)}\right]$$

$$B = 3I - (R + B)$$

240 <= H < 360

$$H = H - 240^{o}$$

$$G = I(1 - S)$$

$$B = I\left[1 + \frac{S\cos H}{\cos(60^{o} - H)}\right]$$

$$R = 3I - (R + B)$$

Outline

- Image Formation
- Digital Images
- Colour and Noise
  - Colour spaces
  - Colour processing
  - Noise

A WFPC2 image of a small region of the Tarantula Nebula in the Large Magellanic Cloud [NASA/ESA]

Pseudocolour

- Also called *False Colour*.
- Opposed to *True Colour* images.
- The colours of a pseudocolour image do not attempt to approximate the real colours of the subject.

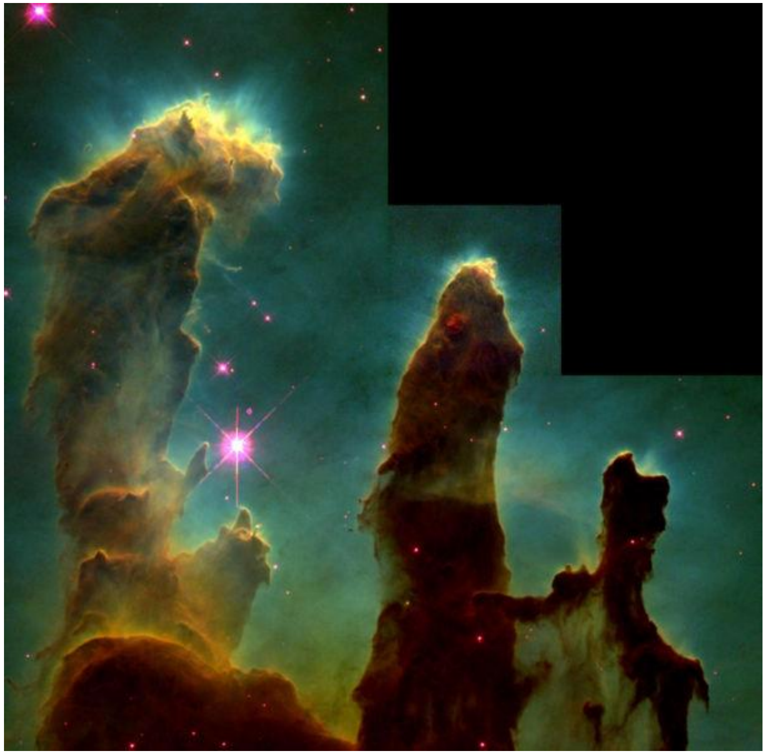

One of Hubble's most famous images: “pillars of creation” where stars are forming in the Eagle Nebula. [NASA/ESA]

Intensity Slicing

- Quantize pixel intensity to a specific number of values (*slices*).
- Map one colour to each *slice*.
- Loss of information.
- Enhanced human visibility.

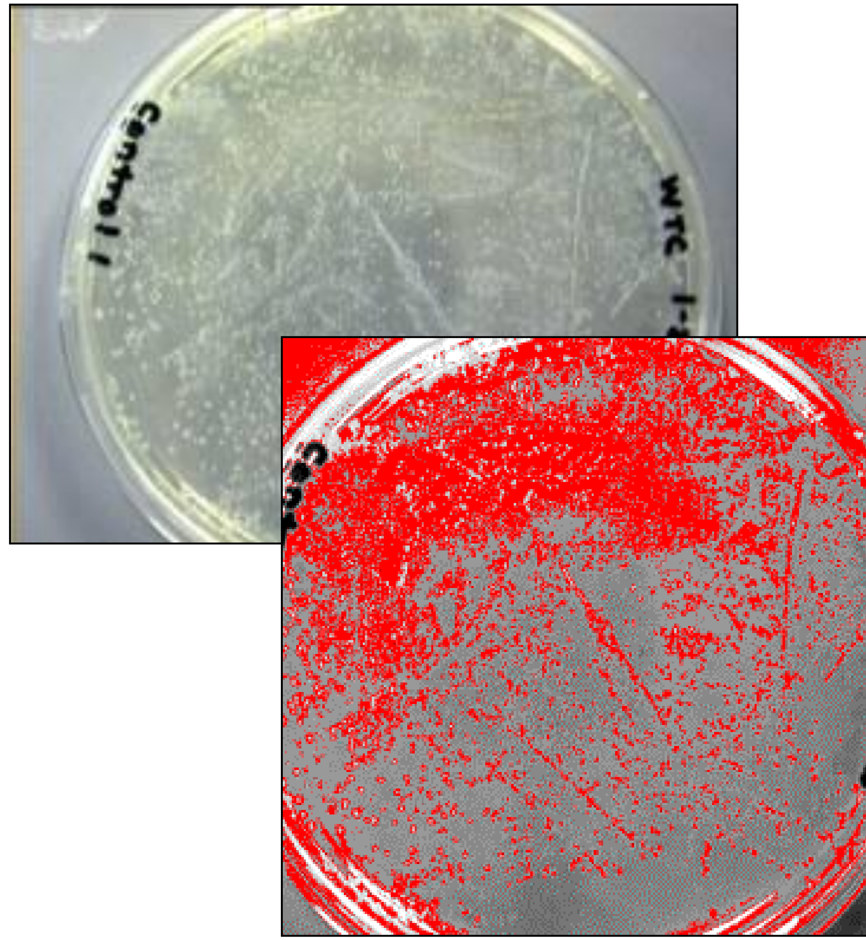

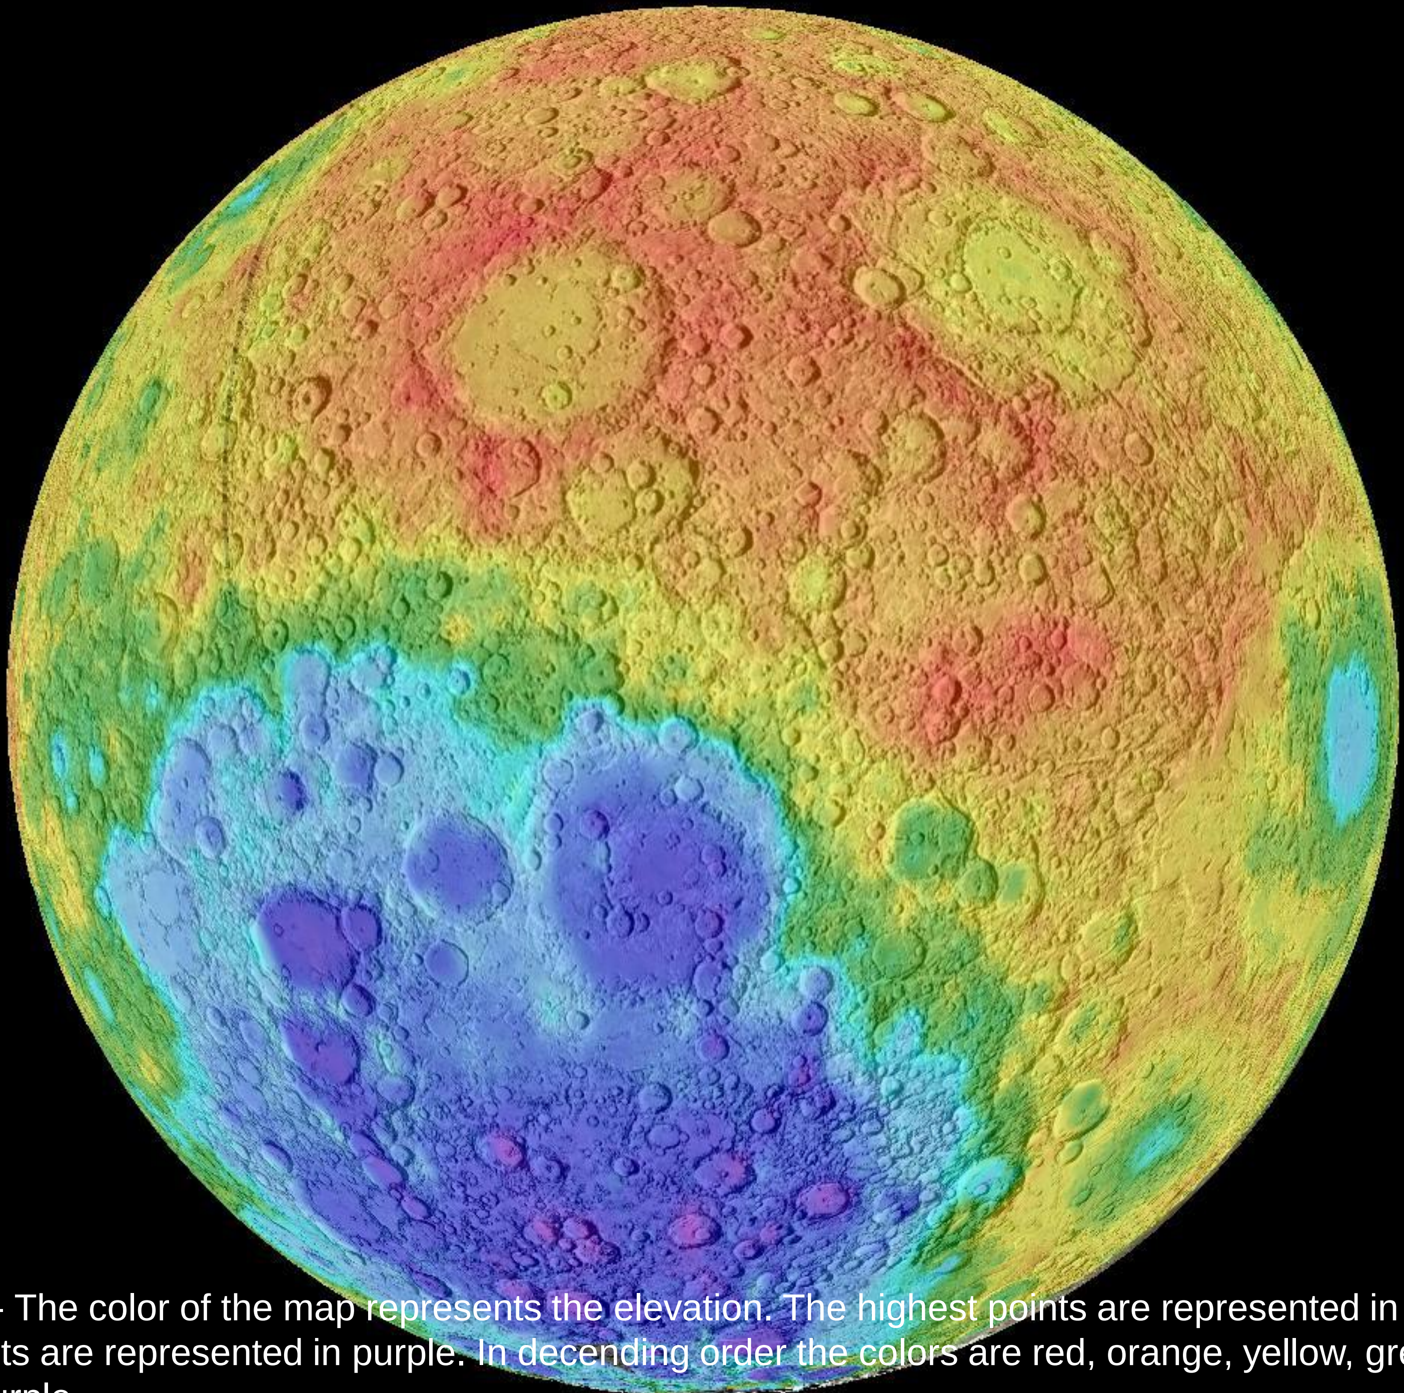

The Moon - The color of the map represents the elevation. The highest points are represented in red. The lowest points are represented in purple. In decending order the colors are red, orange, yellow, green, cyan, blue and purple.

Intensity to Colour Transformation

- Each colour component is calculated using a transformation function.
- Viewed as an Intensity to Colour map.
- Does not need to use RGB space!

$$f(x,y) \Rightarrow \begin{cases} f_R(x,y) = T_R\left[f(x,y)\right] \\ f_G(x,y) = T_G\left[f(x,y)\right] \\ f_B(x,y) = T_B\left[f(x,y)\right] \end{cases}$$

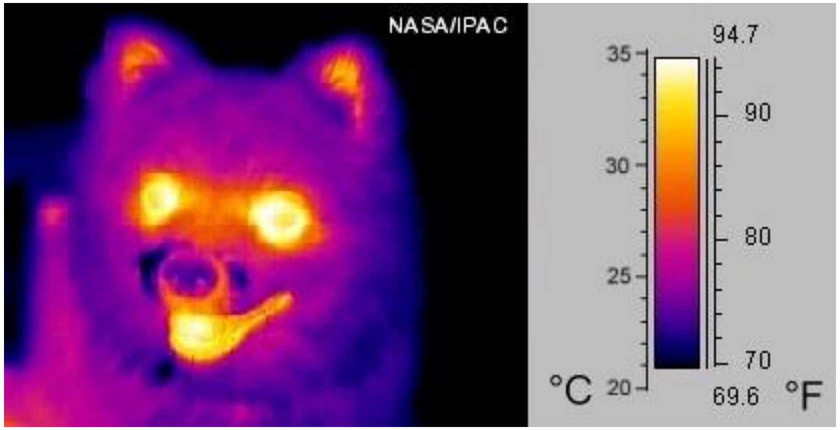

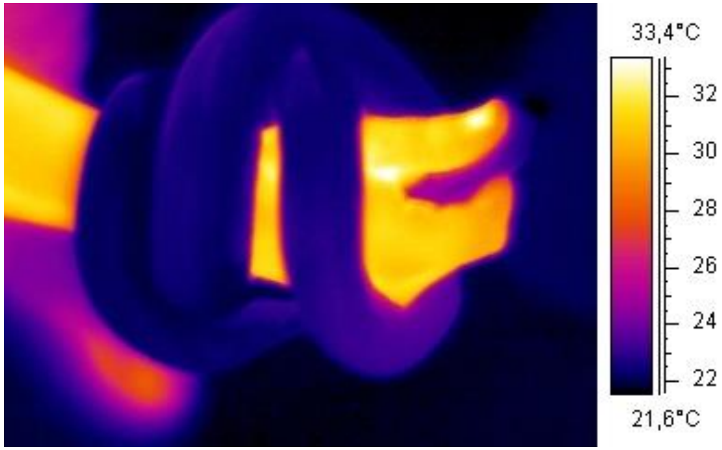

A supernova remnant created from the death of a massive star about 2,000 years ago.

http://chandra.harvard.edu/photo/false_color.html

http://landsat.gsfc.nasa.gov/education/compositor/

Colour Image Processing

- Grey-scale image
  - One value per position.

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

- Colour image
  - One vector per position.

$$f(x,y) = [R\ G\ B]^T$$

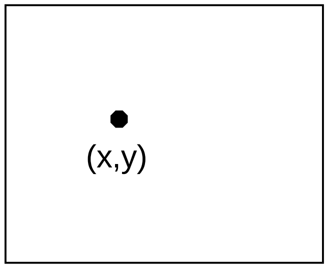

Grey-scale image

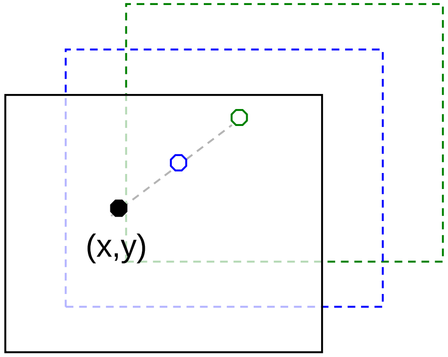

RGB Colour image

Colour Transformations

- Consider single-point operations:

  T_i: Transformation function for colour component *i*

  s_i, r_i: Components of *g* and *f*

$$g(x,y) = T\left[f(x,y)\right]$$

$$s_i = T_i(r_1, r_2, \ldots, r_n)$$

$$i = 1, 2, \ldots, n$$

- Simple example:
  - Increase Brightness of an RGB image

$$s_R = r_R + 20$$

$$s_G = r_G + 20$$

$$s_B = r_B + 20$$

What about an image negative?

Colour Complements

- Colour equivalent of an image negative.

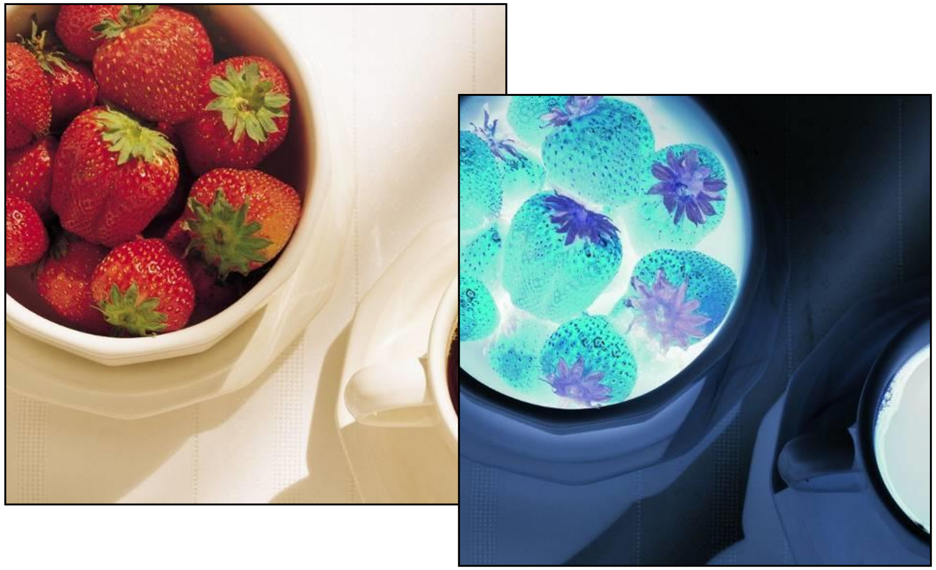

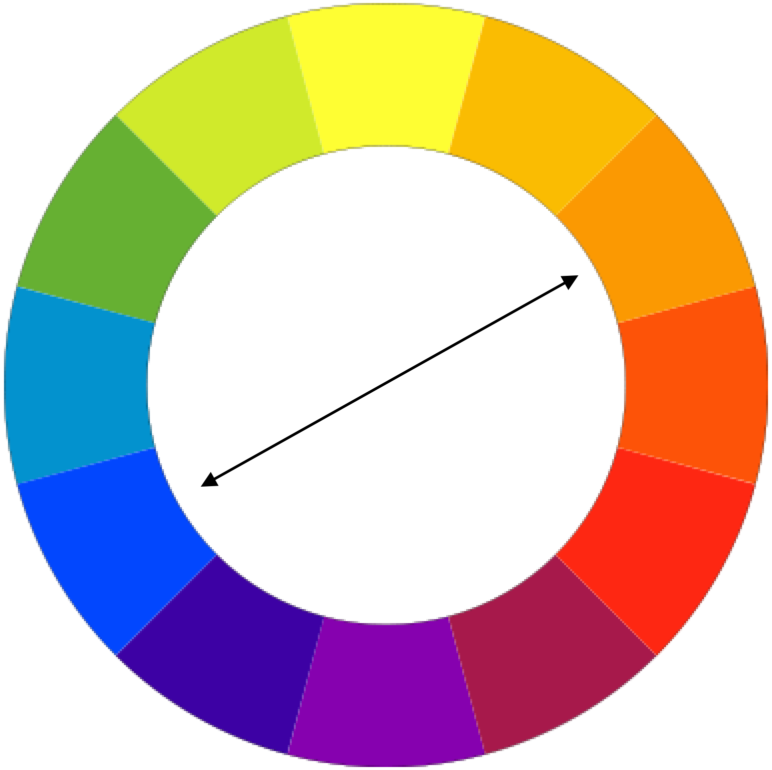

Complementary Colours

Colour Slicing

- Define a hyper-volume of interest inside my colour space.
- Keep colours if inside the hyper-volume.
- Change the others to a neutral colour.

Outline

- Image Formation
- Digital Images
- Colour and Noise
  - Colour spaces
  - Colour processing
  - Noise

Bring the Noise

- Noise is a distortion of the measured signal.
- Every physical system has noise.
- Images:
  - The importance of noise is affected by our human visual perception
  - Ex: Digital TV ‘block effect’ due to noise.

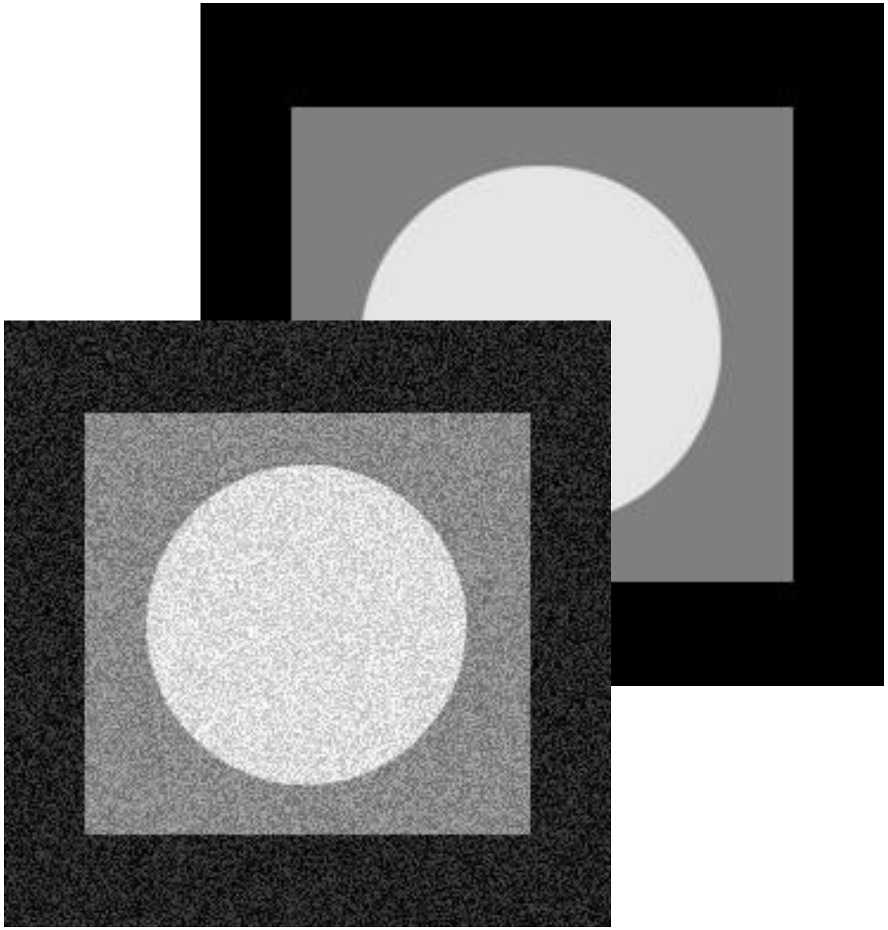

Where does it come from?

- ‘Universal’ noise sources:
  - – Thermal, sampling, quantization, measurement.
- Specific for digital images:
  - – The number of photons hitting each images sensor is governed by quantum physics: *Photon Noise.*
  - – Noise generated by electronic components of image sensors:
    - *On-Chip Noise*, *KTC Noise*, *Amplifier Noise*, etc.

Degradation / Restoration

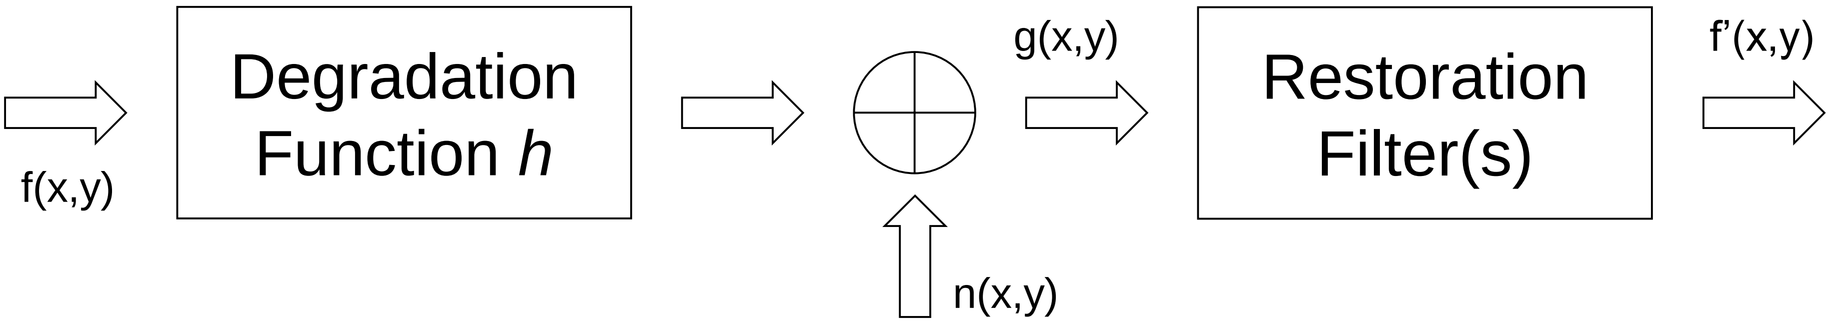

$$g(x,y) = h(x,y) * f(x,y) + n(x,y)$$

$$G(u,v) = H(u,v)F(u,v) + N(u,v)$$

Noise Models

- Noise models
  - We need to mathematically handle noise.
  - Spatial and frequency properties.
  - Probability theory helps!
- Advantages:
  - Easier to filter noise.
  - Easier to measure its importance.
  - More robust systems!

Model: Gaussian Noise

- Gaussian PDF (Probability Density Function).
- Great approximation of reality.
  - Models noise as a sum of various small noise sources, which is indeed what happens in reality.

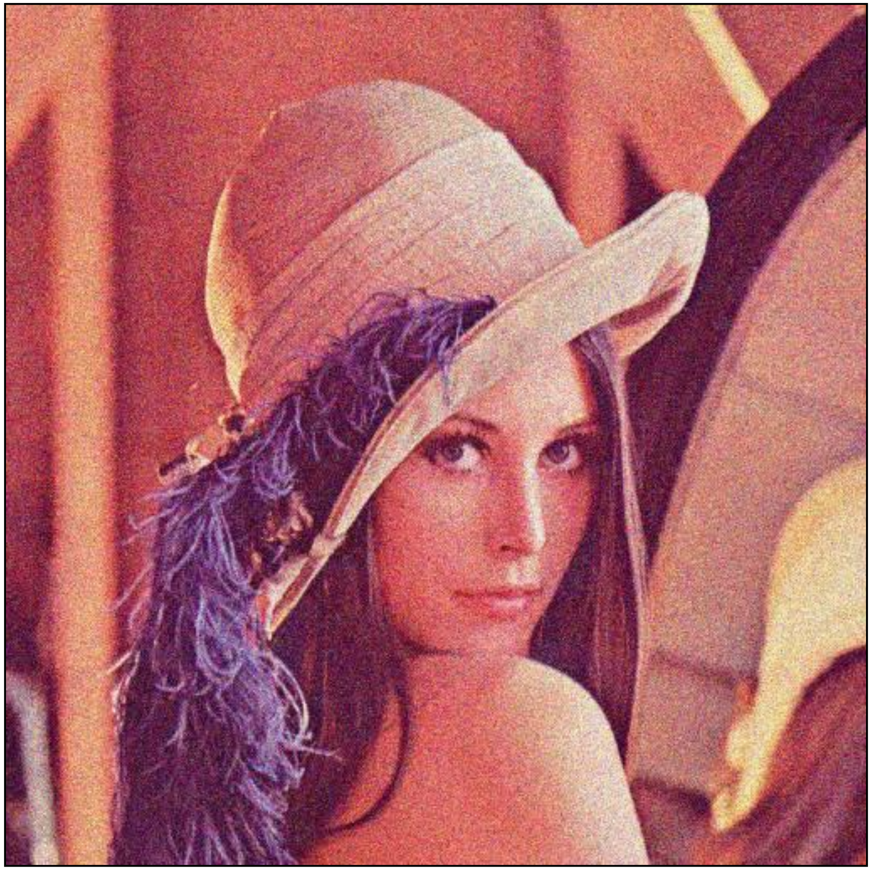

Model: Gaussian Noise

$$p(z) = \frac{1}{\sqrt{2\pi\sigma}} e^{-(z-\bar{z})^2/2\sigma^2}$$

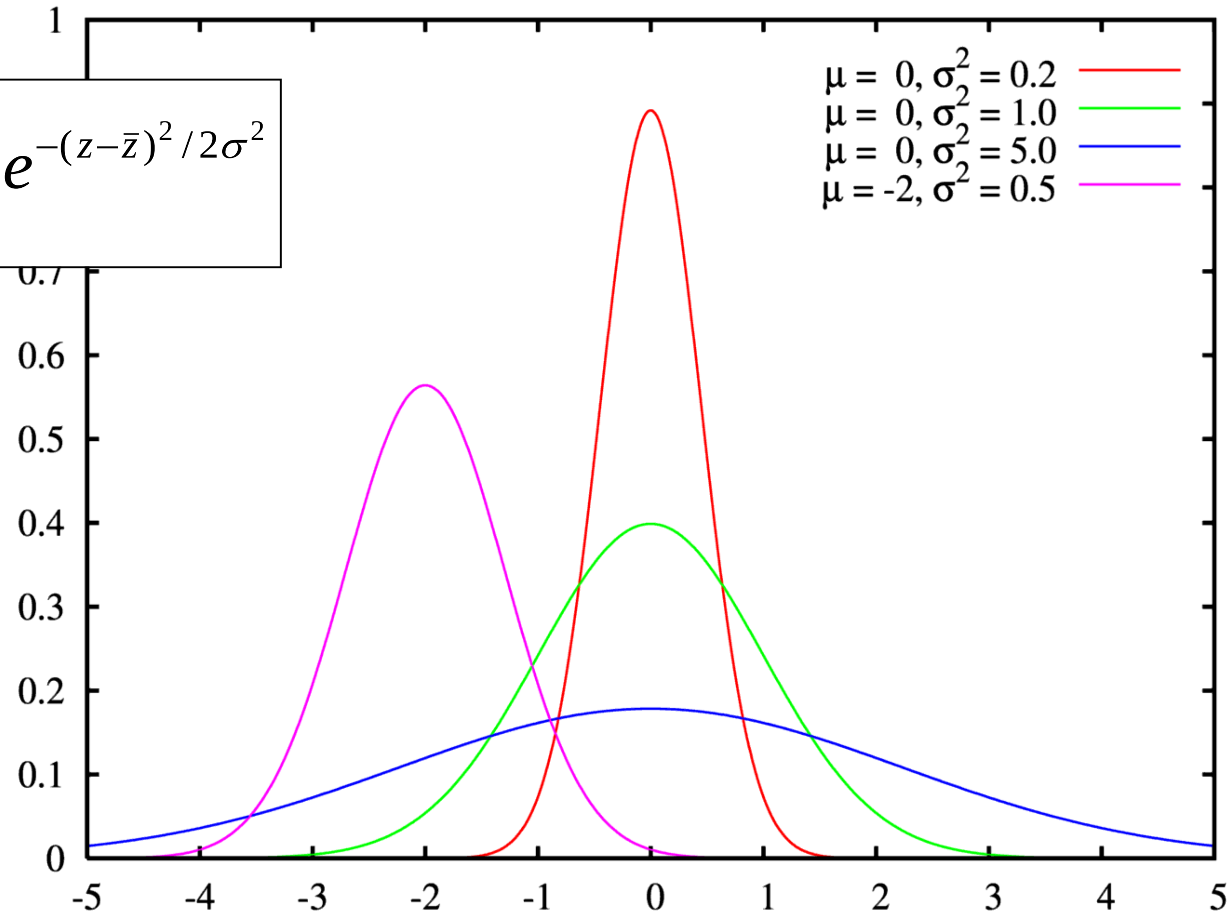

Model: Salt and Pepper Noise

- Considers that a value can randomly assume the MAX or MIN value for that sensor.
  - Happens in reality due to the malfunction of isolated image sensors.

How do we handle it?

- Not always trivial!
  - Frequency filters.
  - Estimate the degradation function.
  - Inverse filtering.
  - ...

One of the greatest challenges of signal processing!

Bibliography

Textbook

- Richard Szeliski, Computer Vision: Algorithms and Applications, Springer-Verlag London, 2011 (Available online: http://szeliski.org/Book/).
- **Other references**
- Making Things See, Greg Borenstein, O'Reilly 2012
- Learning OpenCV: Computer Vision in C++ with the OpenCV Library, Gary Bradski, Adrian Kaehler, O'Reilly 2012
- Machine vision: Theory, algorithms, practicalities, E. R. Davies, Morgan Kaufmann 2005.
- Digital Image Processing, Rafael C. Gonzalez, Richard E. Woods, Prentice Hall, 2007
- Image Processing: Analysis and Machine Vision, Milan Sonka et al., Chapman & Hall, 2007
- D. Forsyth, J. Ponce, Computer Vision: A Modern Approach, Prentice Hall, 2002.