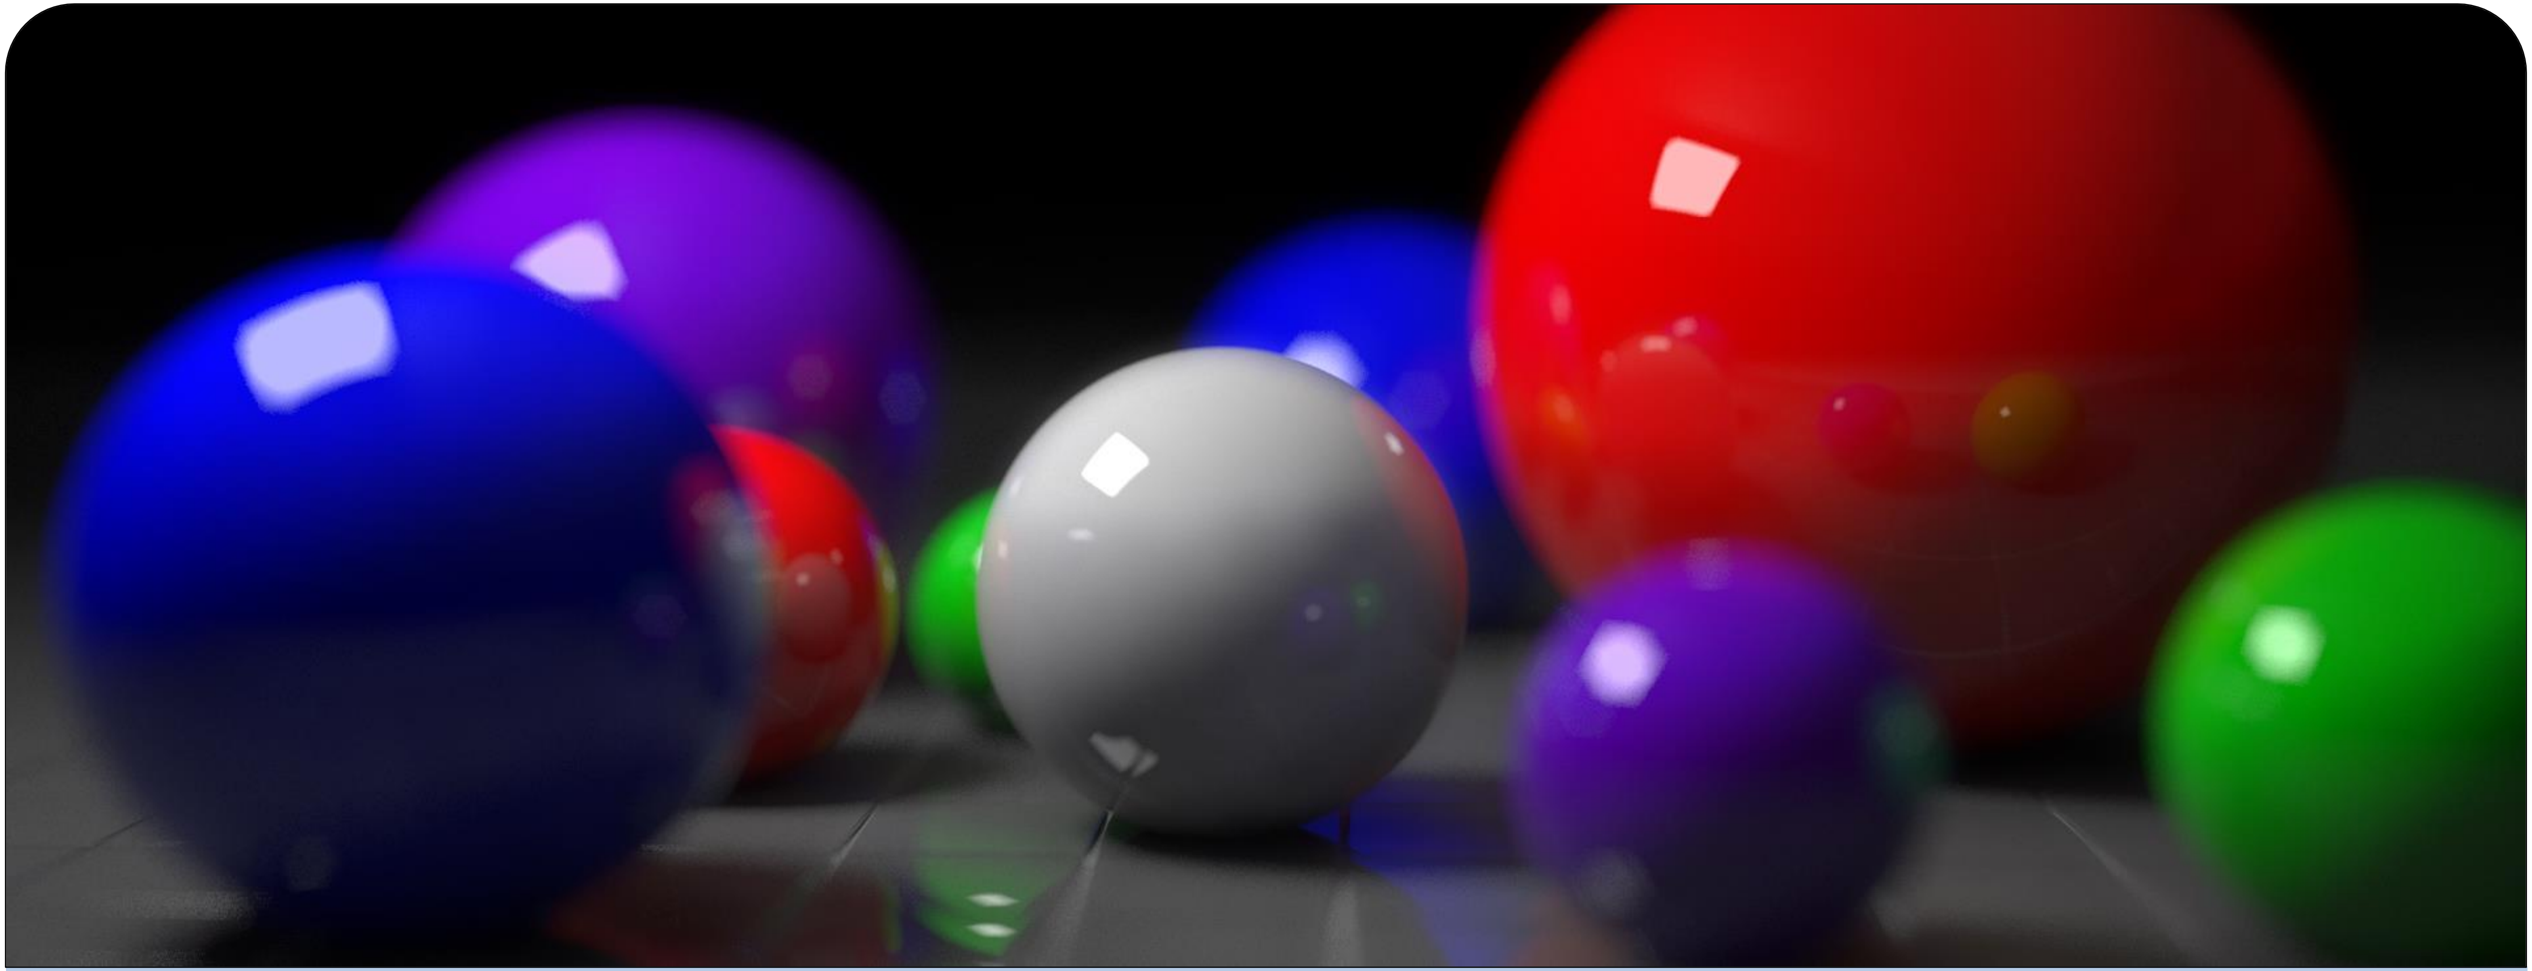

Computer Graphics

Ray tracing - Colors

Andreas Aristidou
andarist@ucy.ac.cy
http://www.andreasaristidou.com

Colors

Lecture Roadmap

- Motivation
- Achromatic light
- Chromatic color
- Color models for raster graphics
- Reproducing color
- Using color in computer graphics

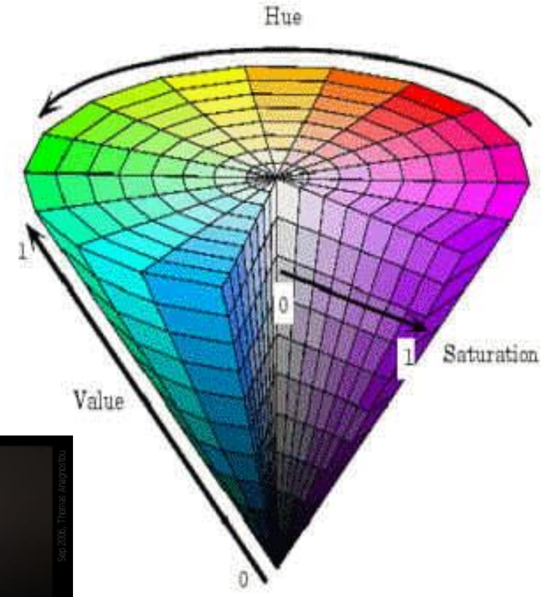

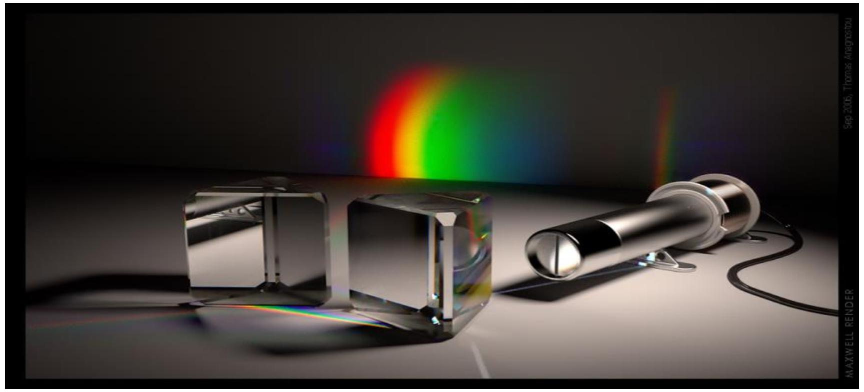

(rendered with Maxwell)

Why Study Color in Computer Graphics?

- Measurement for realism - what does coding an RGB triple mean?
- Aesthetics for selecting appropriate user interface colors
  - How to put on matching pants and shirt in the morning
- What are the perceptual/physical forces driving one's "taste" in color?
- Understanding color models for providing users with easy color selection
  - systems for naming and describing colors
- Color models, measurement, and color gamuts for converting colors between media
  - why colors on your screen may not be printable, and vice-versa
  - managing color in systems that interface computers, screens, scanners, and printers together

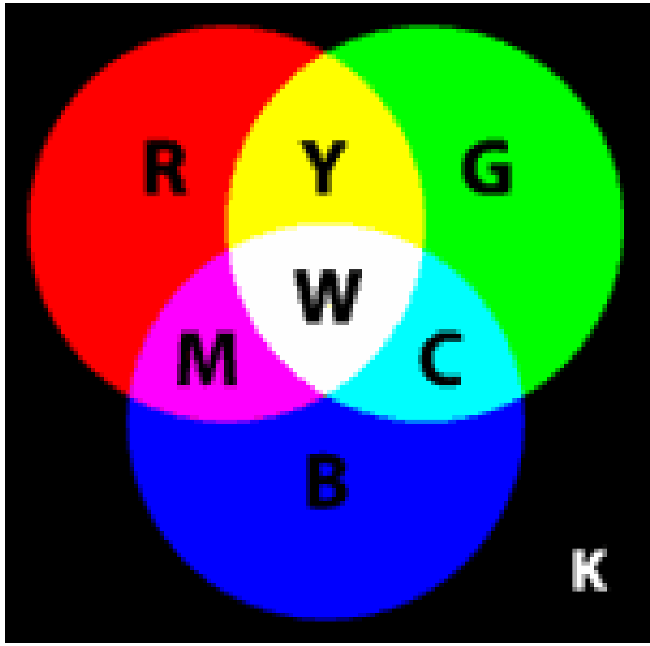

Vision

- **Eye:** is a sensor that can detect electromagnetic waves coming from an object
- More than 70% of the information we collect comes from vision
- Vision is the most important sense (compared to others: hearing, taste, smell, touch), of human
- The main part of a multimedia system is visual (images, video, animation, 3D models, text)

→ It is important to understand the **human visual system (HVS)**

→ allows us to use multimedia technology efficiently

What is color?

- Color is a property of objects that our minds create – an interpretation of the world around us
- Ability to differentiate between colors varies greatly among different animal species
- Color perception stems from two main components:
  - **Physical** properties of the world around us
    - electromagnetic waves interact with materials in world and eventually reach eyes
    - visible light comprises the portion of the electromagnetic spectrum that our eyes can detect (380nm/violet – 740nm/red)
    - photoreceptors in the eye (rods and cones) convert light (photons) into electro-chemical signals, then processed by the brain
  - **Physiological** interpretation of signals (“raw data”) output by receptors – light has no intrinsic “color”(Newton)
    - less well-understood and incredibly complex higher level processing – we understand the eye, but the brain much less so!
    - very dependent on past experience and object associations
- Both are important to understanding our sensation/perception of color
  - Newton (“Optics”, 1704) concentrated on physics/objective measurement, Goethe (“Zur Farbenlehe”, 1810) disputed this reductionist view and added perceptual and artistic elements to his theory: “Colour is a degree of darkness”

Introduction to the color

- Electromagnetic waves (EM) may have different wavelengths
  - Long wavelengths (π.χ. dc/power band)
  - Short wavelengths (π.χ. gamma rays)

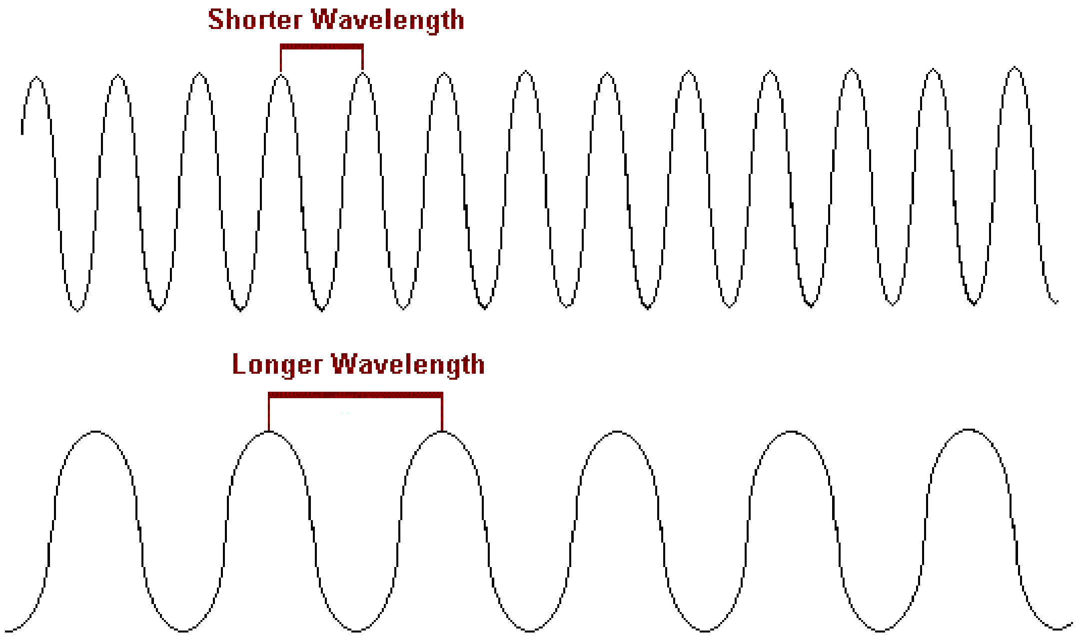

Introduction

- The eye can detect waves in a narrow spectrum
  - 400nm–700nm (nm:nanometer 1nm=1×10^{-9}m)
  - The electromagnetic spectrum is the region with all possible electromagnetic radiation frequencies
  - Different colors have different wavelengths
    - E.g. red = 700nm, blue = 400nm

- Sunlight:
  - Its spectrum is within this wavelength
  - It is the main source of electromagnetic waves of the earth
  - It is believed that the evolution of HVS has been influenced by the sun's radiation

Introduction

- How do we see objects

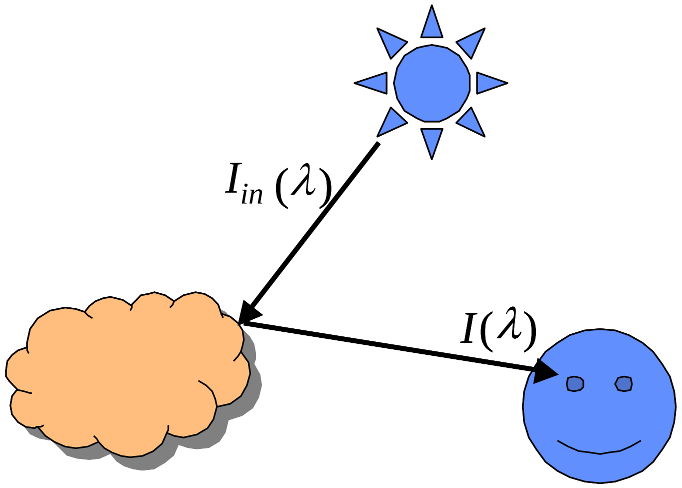

$$I(\lambda) = \rho(\lambda) \times I_{in}(\lambda)$$

$I_{in}(\lambda)$: the distribution of energy coming from the light source

$I(\lambda)$: the light received by the observer;

$\rho(\lambda)$: Reflectivity or refractiveity of the object

λ: wavelength

Human Visual System

- The human visual system can be considered as an optical system
-

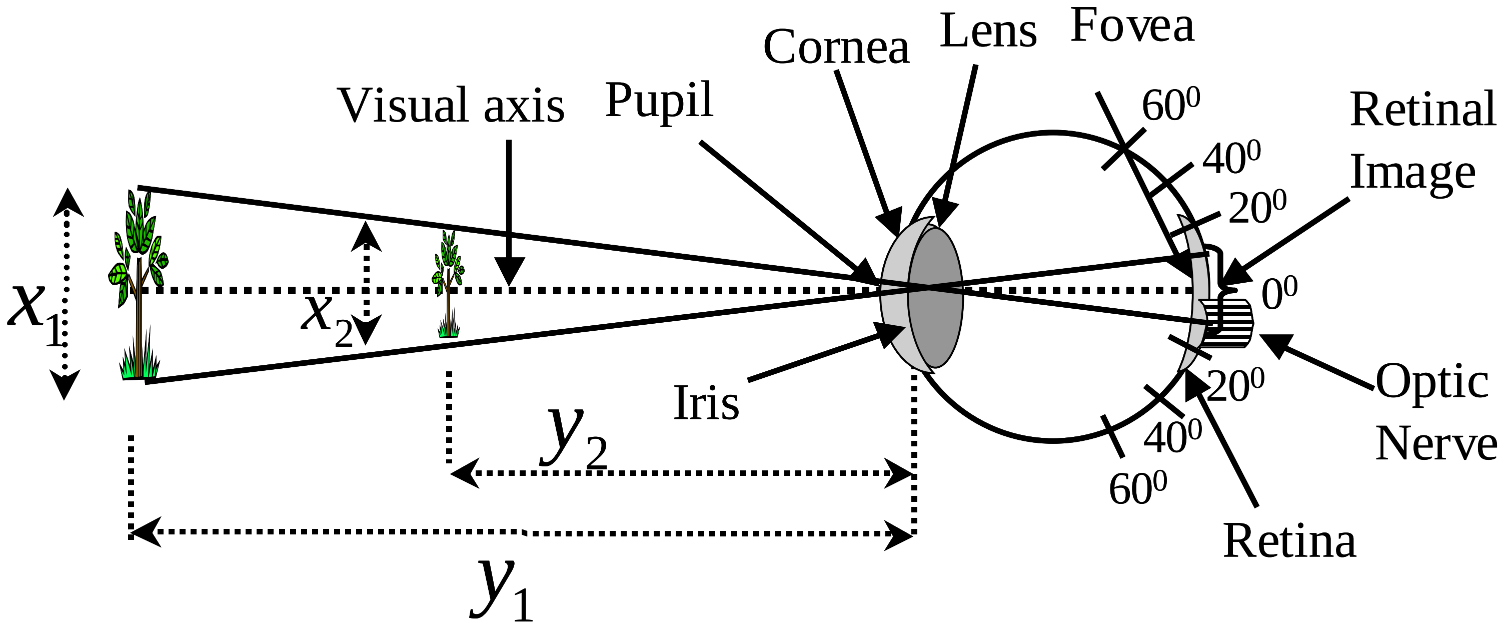

Human Visual System

- Optical system
  - The light of an object falls on the eye
  - The pupil of the eye acts as an opening (aparture)
  - An image is created on the retina
  - The observer sees the picture

- The perceived size of an object depends on the angle created in the retina

- The retina can see the details better when the image it creates is larger and has more light

Human Visual System

- The pupil of the eye controls the size of the light entering the eye
  - Usual lighting: the pupil is about 2 mm
  - Low light: size grows allowing more light
  - High lighting: size diminishes by limiting the entry of lightlight
- The Retina contains two types of light signal receivers (photoreceptors):
  - rods
  - cones

Human Visual System

- Mesopic vision (the lighting is low, but not dark enough): average lighting
- Rods and cones are active

- **Color vision:** Given by cones
  - Cones are not active in low light → the color of an object may not be correctly detected in low light

Relative Luminous Efficiency

- the luminance or intensity of an object with an incoming energy i(λ) can be calculated by using the equation below:

$$L = I(\lambda) \times V(\lambda), \qquad where: \quad I(\lambda) = \rho(\lambda)\ x\ I_{in}(\lambda)$$

- V(λ): relative luminous efficiency

- The brightness is measured on cd/m2 where cd (candela) is the unit of the reflected light intensity
  - e.g. usually computers have a brightness of 50 - 300 cd/m2
  - Sunlight: 100 000 cd/m^2

- L brightness is more important than I(λ) in the process of forming the image
  - Why? Because it is what the observer will see

Relative Luminous Efficiency

- Light with high L brightness always appears bright to a human observer
- Light with low L brightness always appears dark to a human observer
- High I(λ) (Radiation energy) may appear dark (if the energy is close to two ends of the shape of the curve)

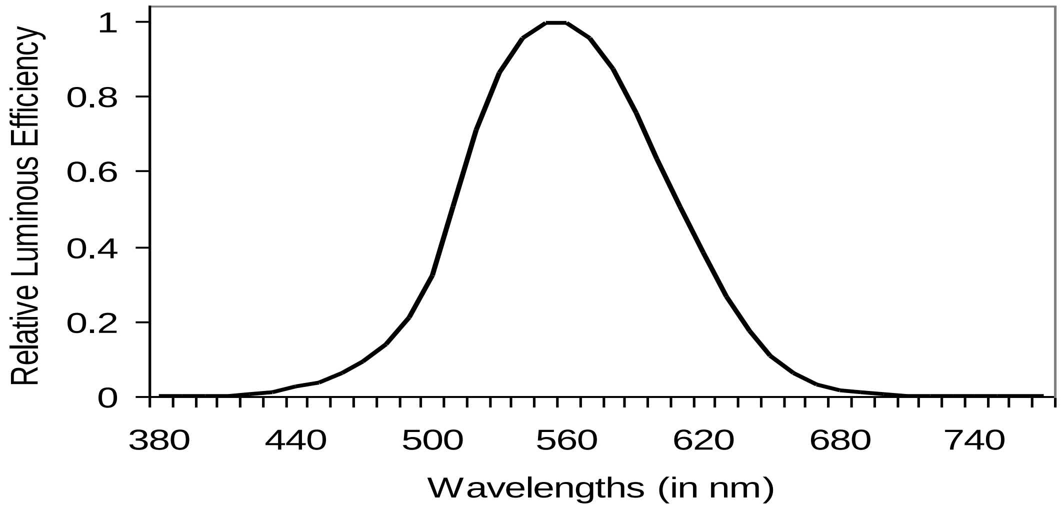

Relative Luminous Efficiency

- Brightness is an important factor in determining whether an area of the image will be dark or bright
- The **Brightness** of surrounding areas also plays an important role in displaying a particular image
- 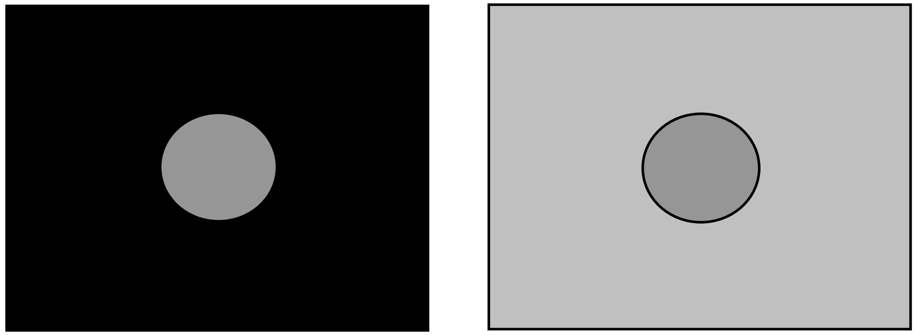

The circles in the center have the same brightness in both cases

- The circle on the dark background seems to be brighter

Luminous – Brightness - Contrast

- The perceived brilliance of an object is called **brightness**
- **Brightness** depends on the environment

Why?

- Our perception is sensitive to contrasts of brightness
- **Contrast** is a term that is often used to describe differences in the brightness values of two adjacent image areas

Our Visual System Constructs our Reality

- Visual processing of perception (such as reflex actions) is generally faster than higher-level cognitive processing
  - Roads use symbols instead of words for signs (but "double encoding" is always better)
  - It can have a more emotional impact
  - Often "a picture is worth a thousand words..."
  - Our visual system constructs our reality
  - Perceptual invariants are crucial for sense-making
    - Optical illusions

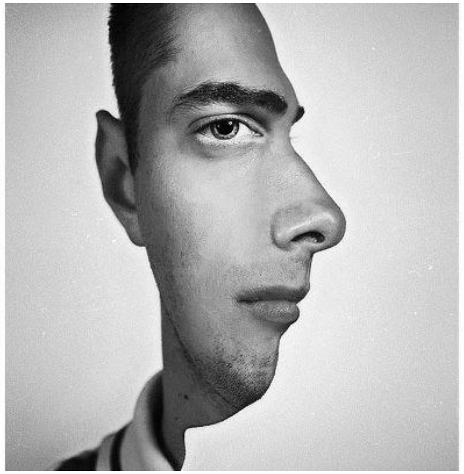

Optical illusions

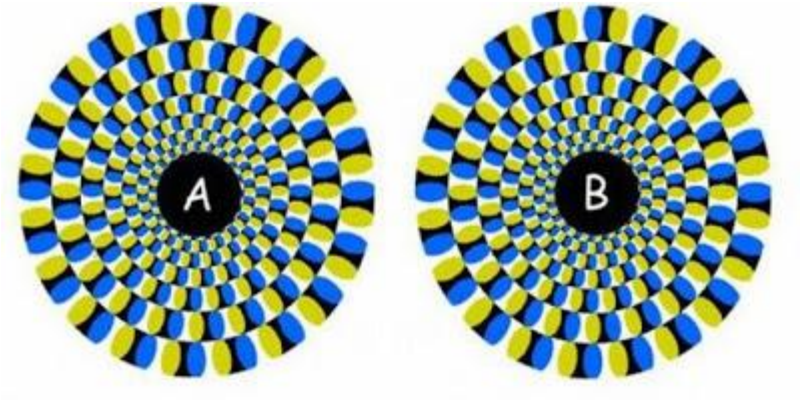

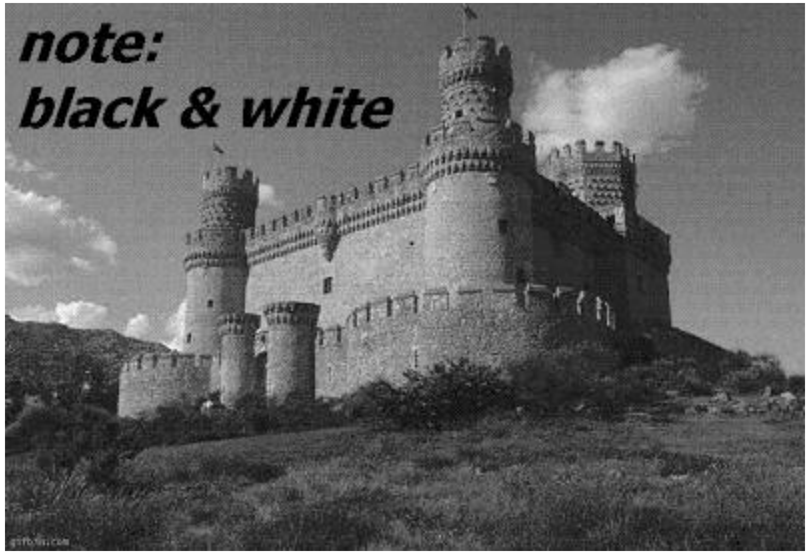

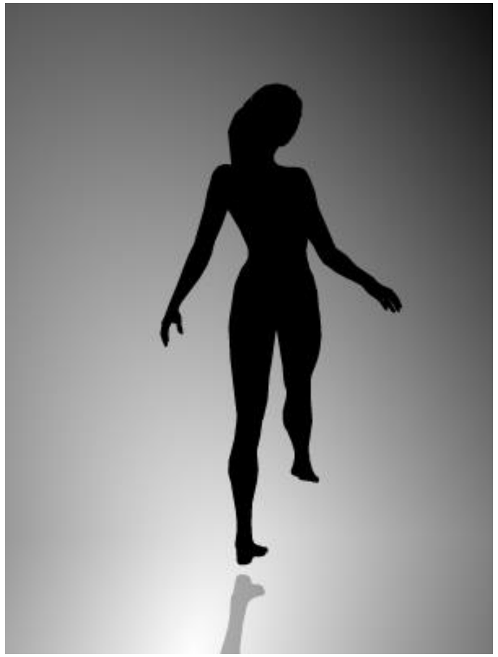

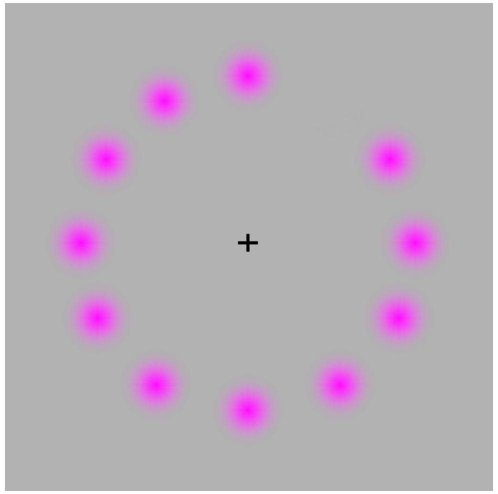

More

- Vection
- Forced Perspective

Our Visual System Constructs our Reality

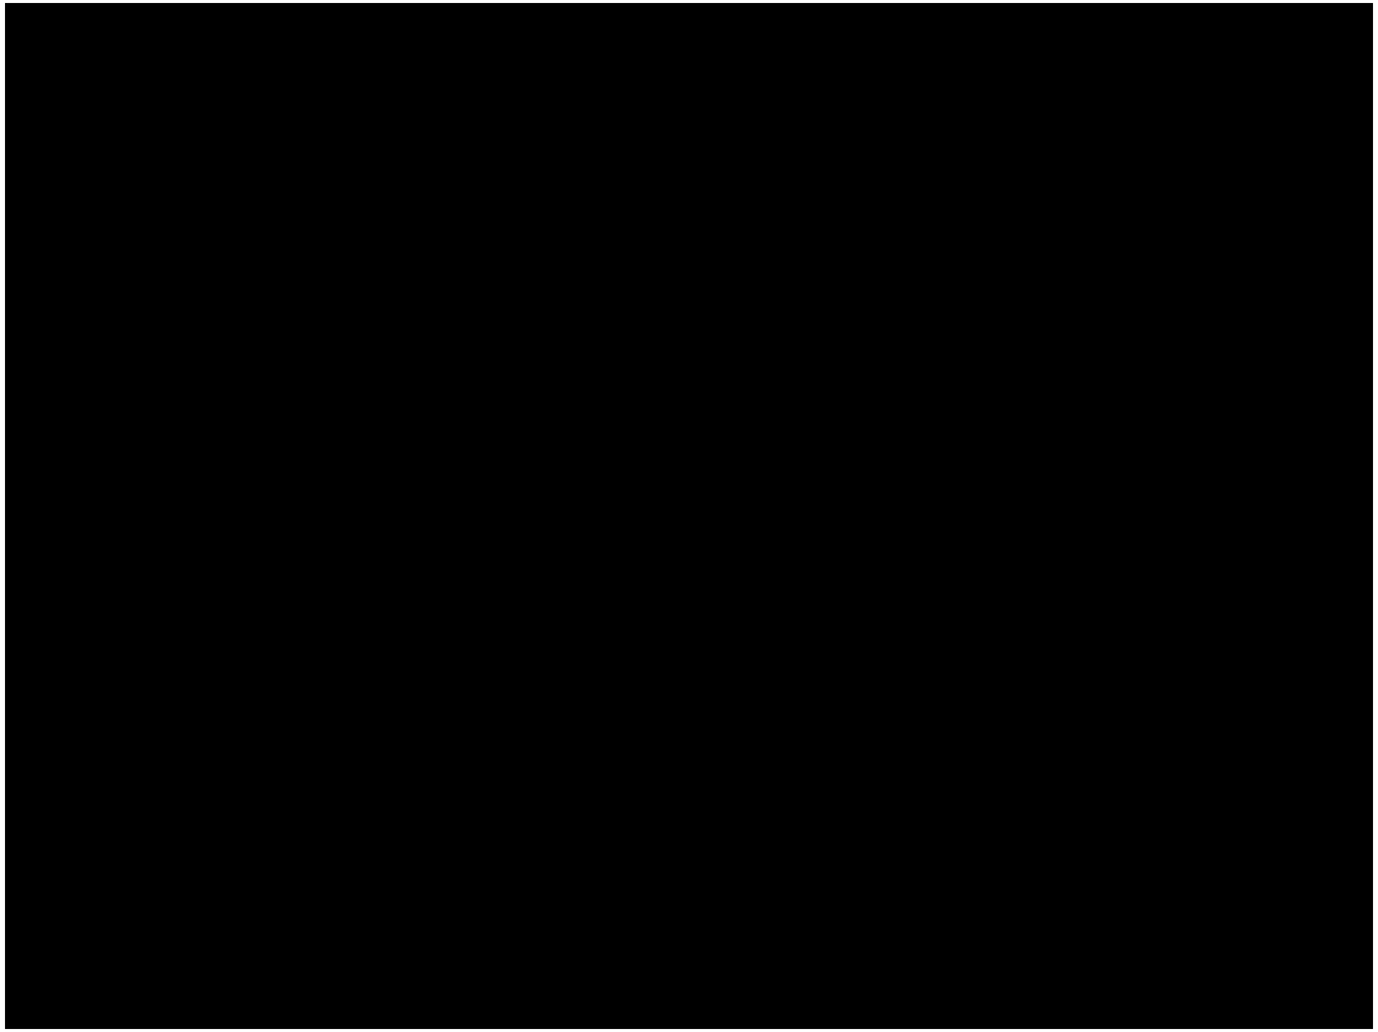

“The Dress”

- Photo that went viral because people could not agree if the dress was black and blue or white and gold
- Aroused interest in how the brain perceives color for individuals
- A similar situation occurs in the sound
  - Yanny or Laurel?

Achromatic Light

- Achromatic Light: “without color,” quantity of light only
  - Called intensity, luminance, or measure of light’s energy/brightness
    - The psychophysical sense of perceived intensity
  - Gray levels (e.g., from 0.0 to 1.0)
    - We can distinguish approximately 128 gray levels
  - Seen on black and white displays
  - Note Mach banding/edge enhancement – stay tuned

Achromatic Light

- Eye is much more sensitive to slight changes in luminance (intensity) of light than slight changes in color (hue)
  - "Colors are only symbols. Reality is to be found in luminance alone… When I run out of blue, I use red." (Pablo Picasso)
  - “Picasso's Poor People on the Seashore uses various shades of blue that differ from each other in luminance but hardly at all in color (hue). The melancholy blue color serves an emotional role, but does not affect our recognition of the scene.” Also, consider how realistic black and white images/movies look – suspension of disbelief! In short, luminance for recognition, color for mood
  - “The biological basis for the fact that color and luminance can play distinct roles in our perception of art, or of real life, is that color and luminance are analyzed by different subdivisions of our visual system, and these two subdivisions are responsible for different aspects of visual perception. The parts of our brain that process information about color are located several inches away from the parts that analyze luminance -- as anatomically distinct as vision is from hearing. ” (source below)
- *Source:* Includes in-depth explanations of many interesting natural phenomena relating to color (including several interactive applications) http://www.webexhibits.org/causesofcolor/

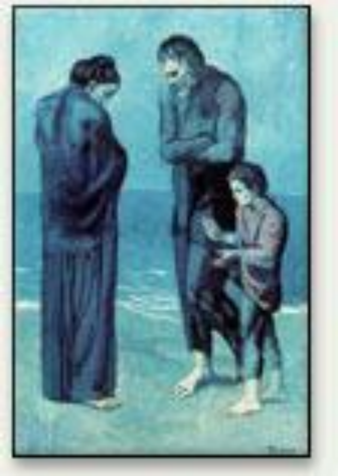

Color representation

- The study of colors is very important in computer graphics

- Color:
  - It gives a pleasant feeling
  - It allows us to quickly capture more information

- Scientific experiments have shown that human:
  - Can distinguish small differences in color
  - It is estimated that the number of colors we can perceive is almost 10 million

Chromatic Light

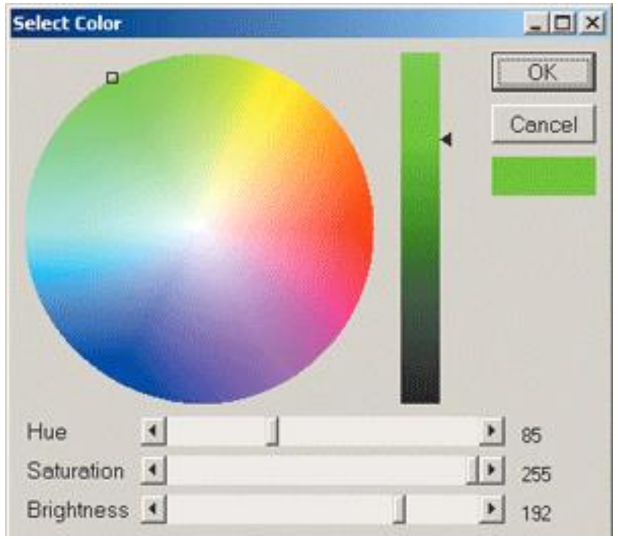

Example of an HSV color picker

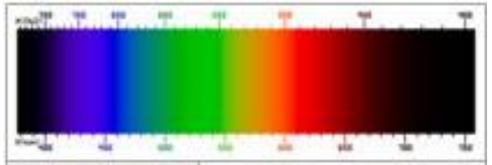

Color	Wavelength
violet	380–450 nm
blue	450–495 nm
green	495–570 nm
yellow	570–590 nm
orange	590–620 nm
red	620–750 nm

Ingredients of a Rainbow

- Factors of color sensation – a 3-space of largely independent perceptual parameters (hmmm…coincidence?!?)
  - Brightness / intensity (this circular color picker shows single brightness)
  - Chromaticity / color:
    - Hue / position in spectrum(red, green, …) - angle in polar coordinates (circular color picker)
    - Saturation / vividness – radius in polar coordinates (circular color picker)
- Interesting factoid: early civilizations didn't identify blue as a unique color with its own name, e.g. the *Odyssey*
  - https://www.sciencealert.com/humans-didn-t-see-the-colour-blue-until-modern-times-evidence-science

Color representation

- The theory of trichromatic representation means that any color can be produced by mixing appropriate quantities of three additive main colors.

3-receptor model

- It has been discovered that there are three different types of cones in the human retina
  - "Red" cones
  - "Green" cones
  - "Blue" cones

- The excitation of the different types of cones determines the color that the observer will see

3-receptor model

The absorption spectra of the three types of con

Colors on websites, traffic lights, danger, etc.

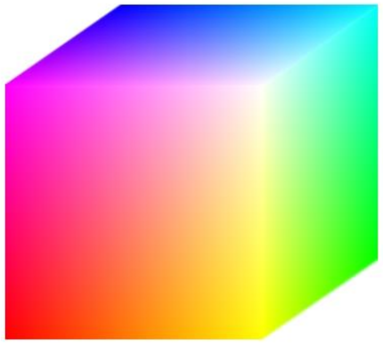

$p_1 = 645.2$ nm
$p_2 = 525.3$ nm
$p_3 = 444.4$ nm

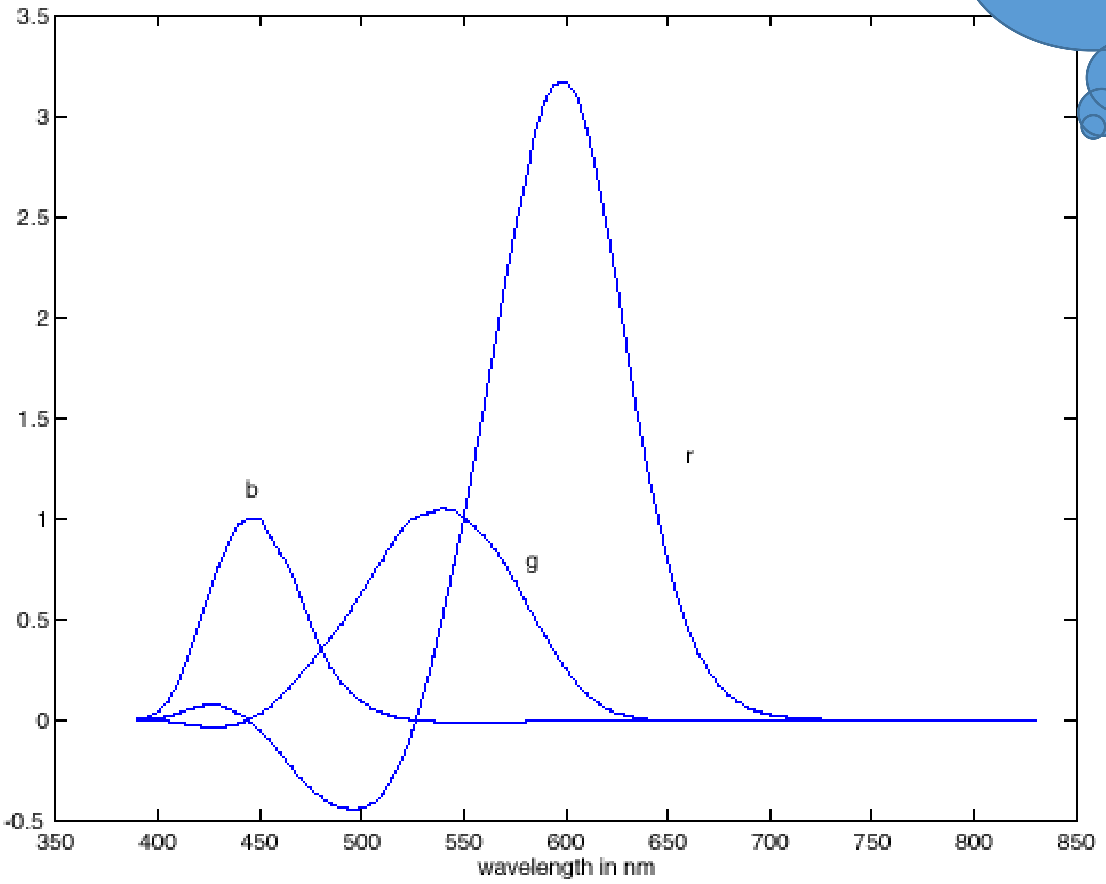

3-receptor model

- Chromatic blindness
- Monochromatic people have only one rod and only one type of cone
- Discolored humans have only two types of cones (one type of cone is absent)
  - They can see some colors perfectly
  - Ignore other colors

3-receptor model

- Several thousand natural colors can be distinguished from humans
- → difficult to design a system that will be able to display such a large number of colors
- Special properties of HVS make it possible to use mechanisms to design a simplified color display system
- Any color can be produced by appropriately blending the three main colors

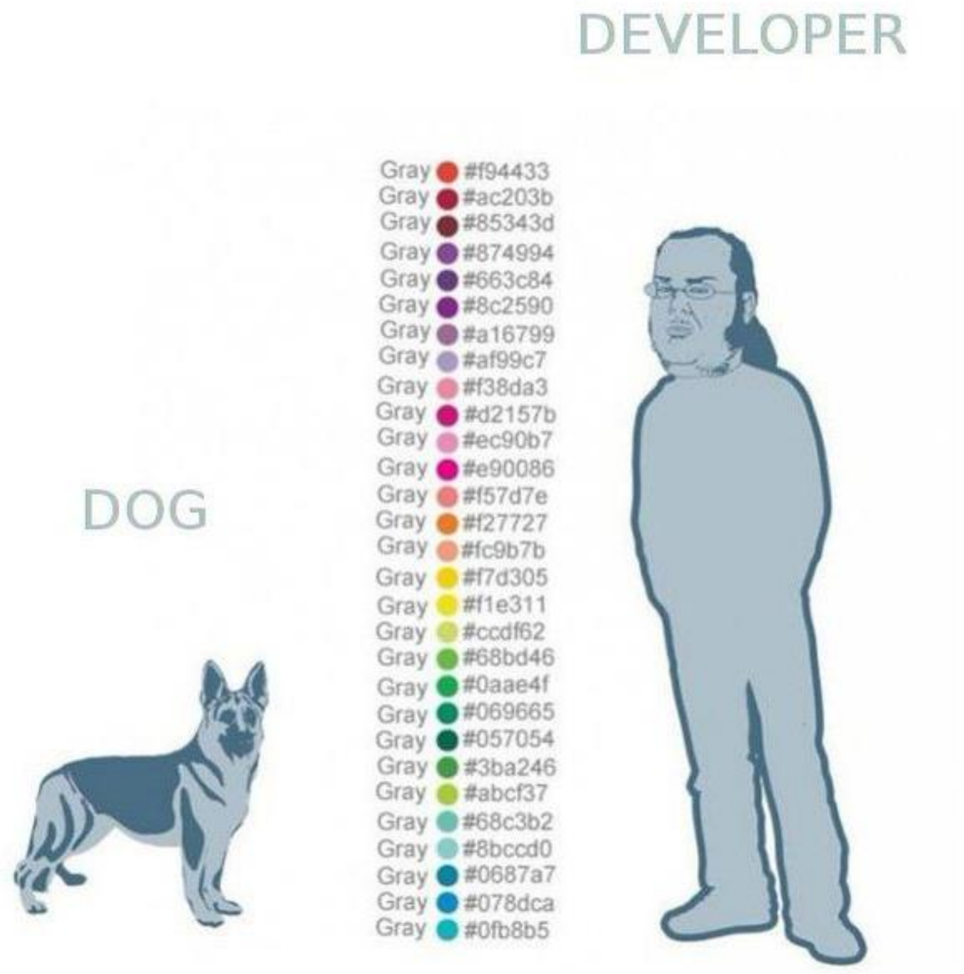

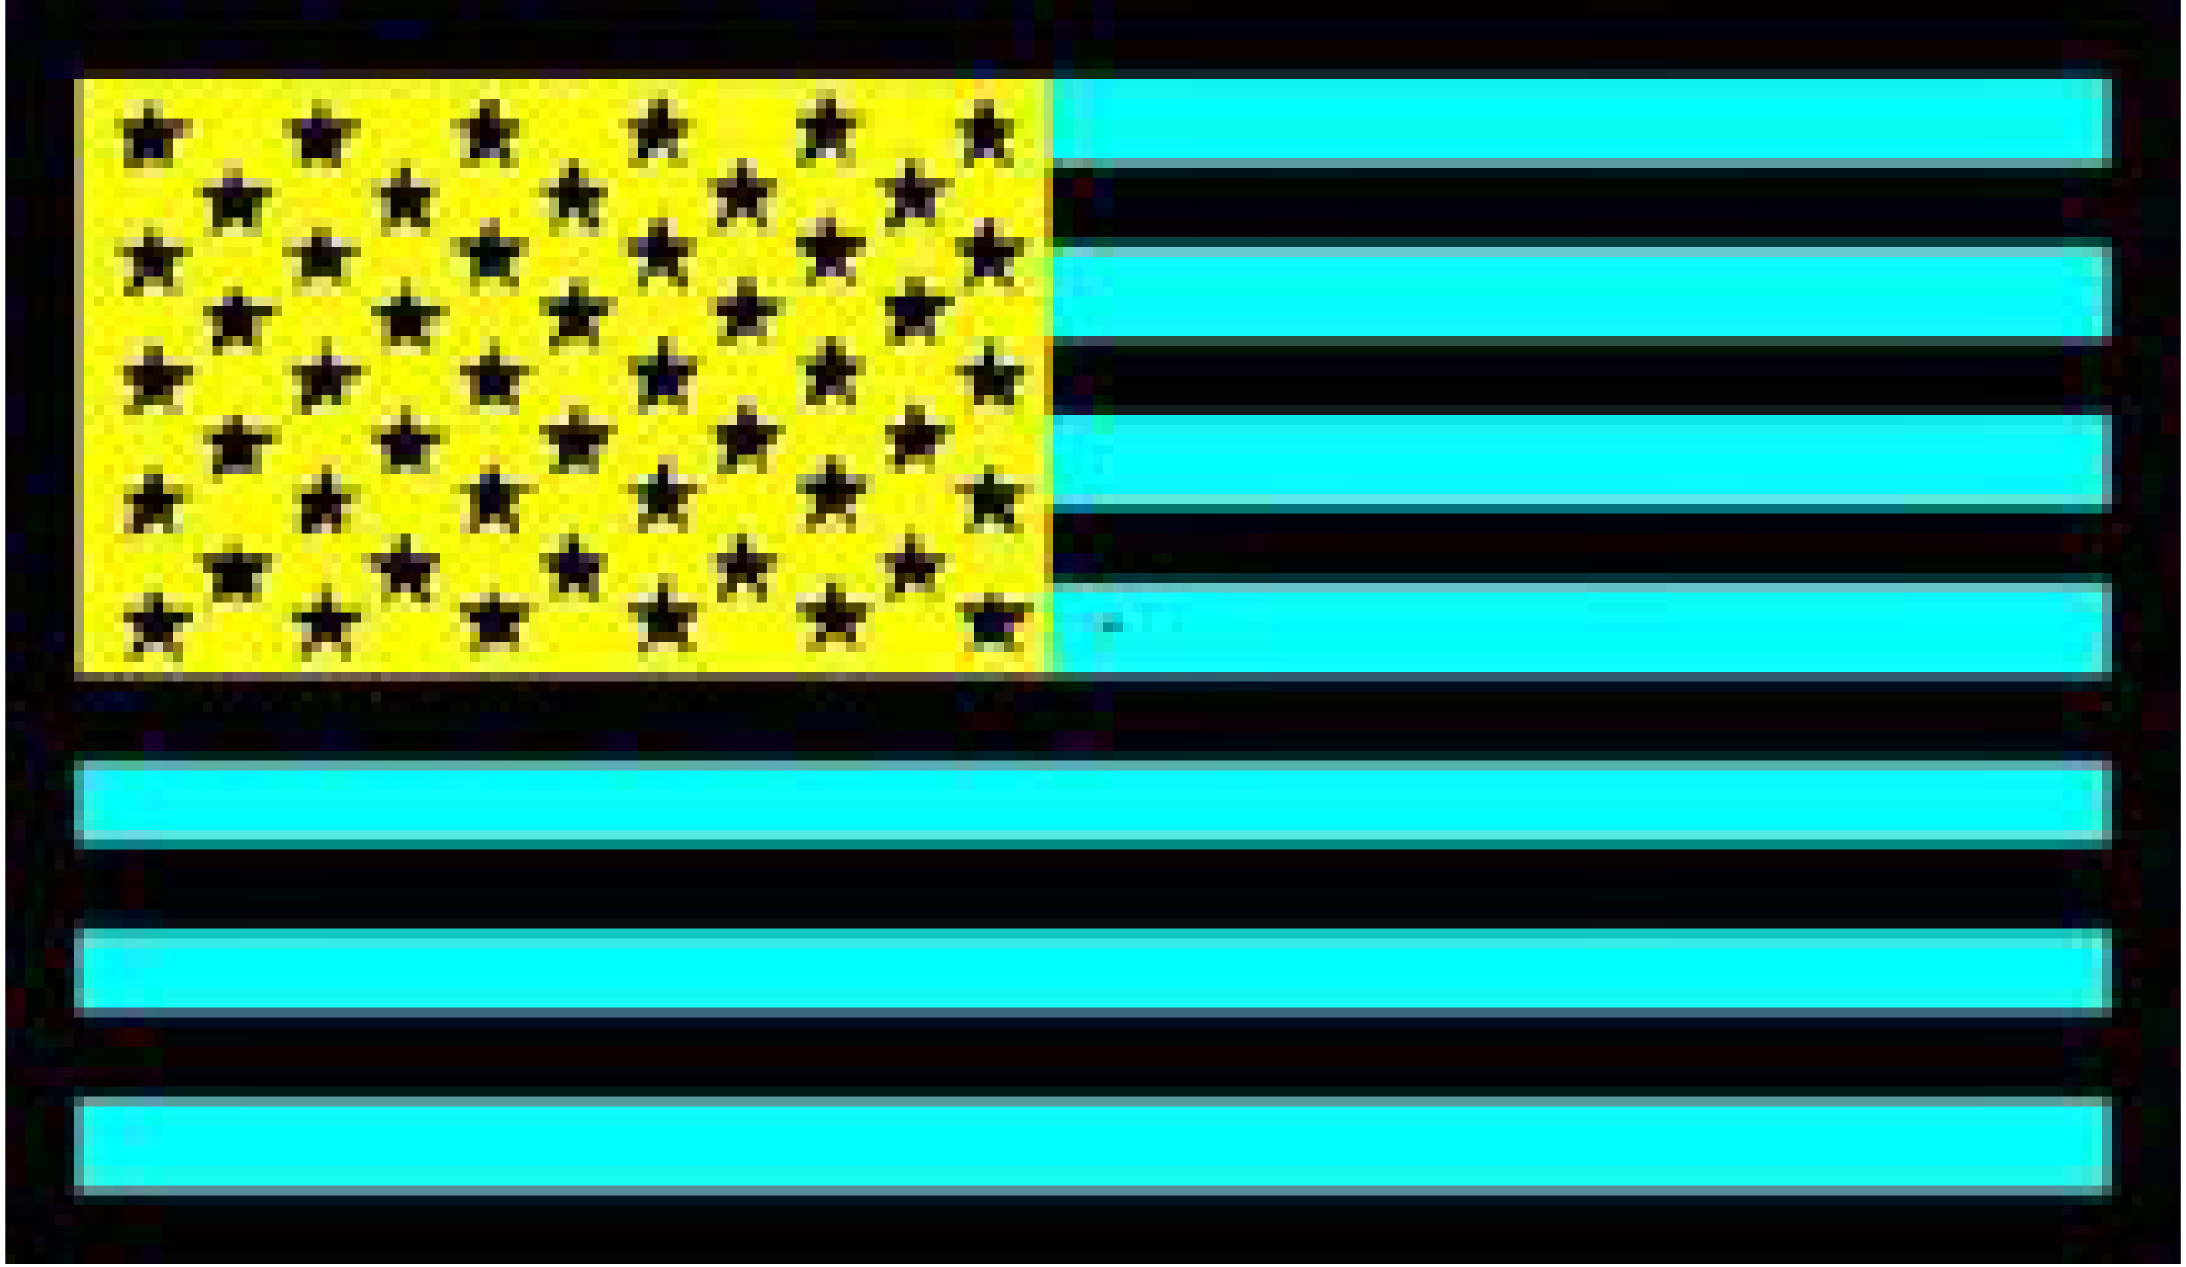

Color matching

- Most color reproduction systems use the 3-receptors model
- To implement such a system we need to know the percentage of the main colors that need to be mixed to reproduce a particular color
- The science of measuring color is known as **colourimetry**

Color matching

- Laws to be applied during color matching:
  1. Any color can be represented by blending a maximum of three colors
  2. The brightness of a color is equal to the sum of the illuminance of its elements
  3. Color addition
  4. Color subtraction

Color representation

The RGB color space

The RGB color space

- In the RGB color space, the three primary colors are red, green, and blue.
- It is also important to know that – regardless of the the amount of red, green or blue – it is not possible to create each color as a combination of the red, green and blue elements.
- Only colors in **rgb gamut** can be represented in this way

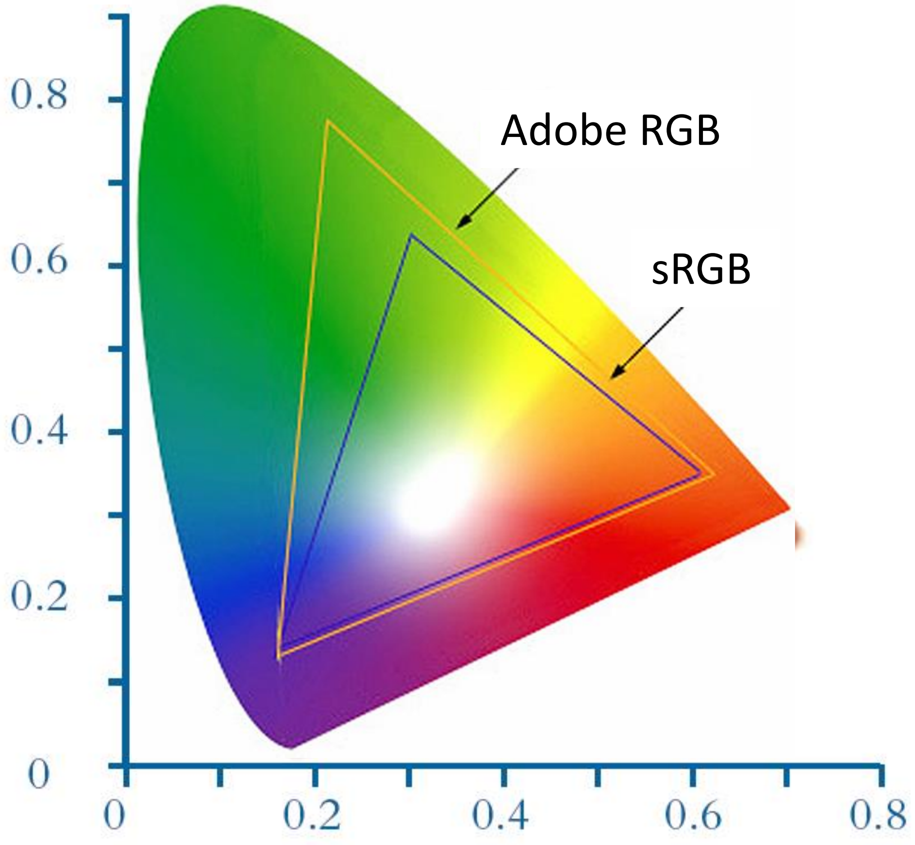

The RGB colour gamut

The RGB color space

The three headlights of the **red**, **blue**, and **green** light

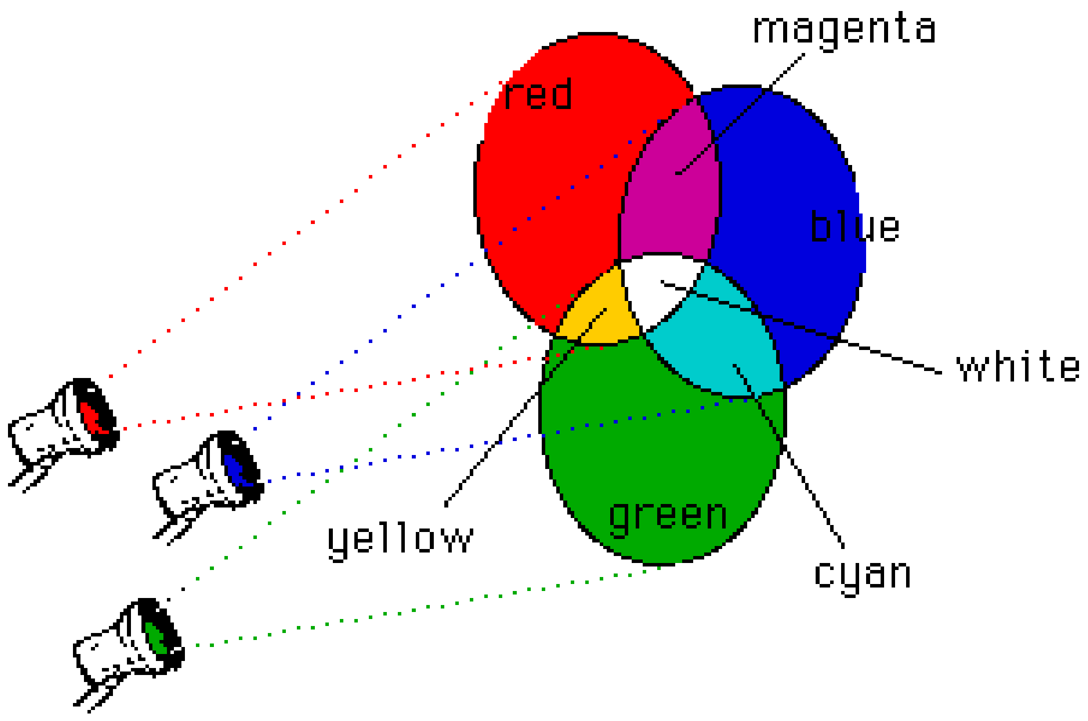

RGB: Adding colors

Red + Blue = Magenta

Blue + Green = Cyan

Red + Green = Yellow

Red + Green + Blue = White

The RGB model

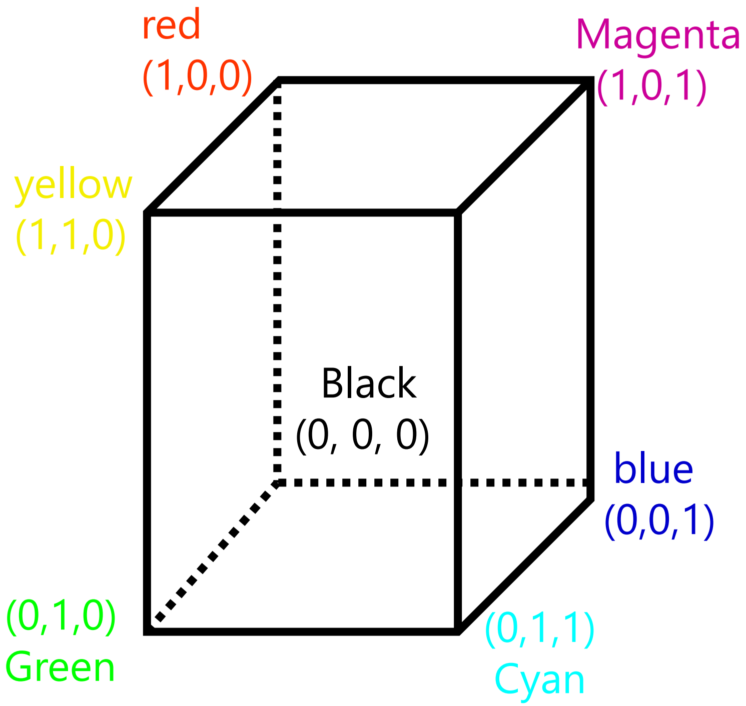

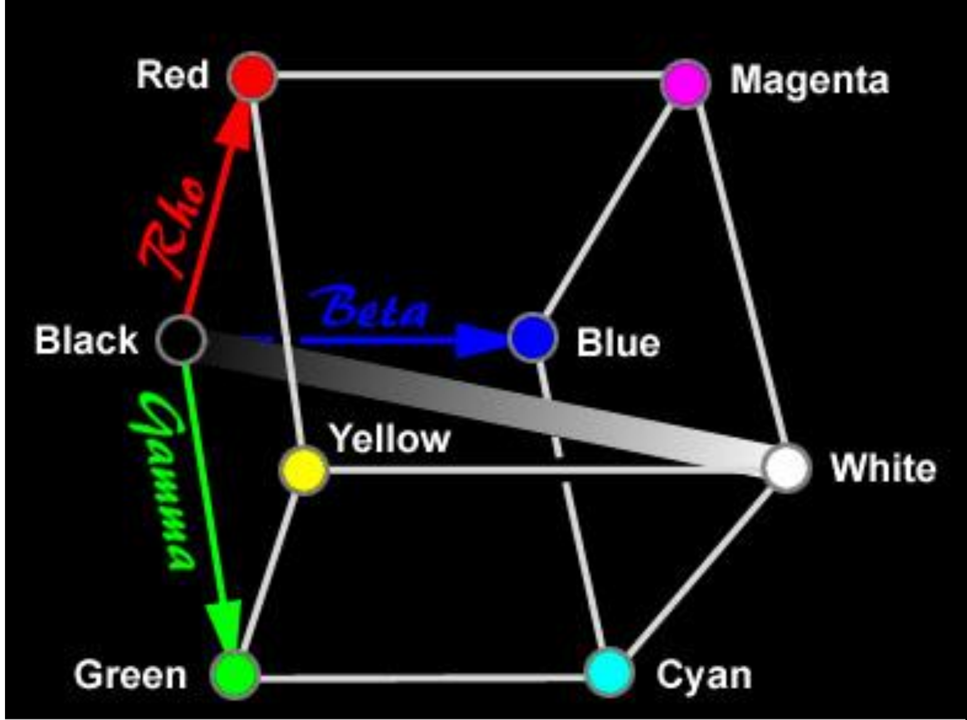

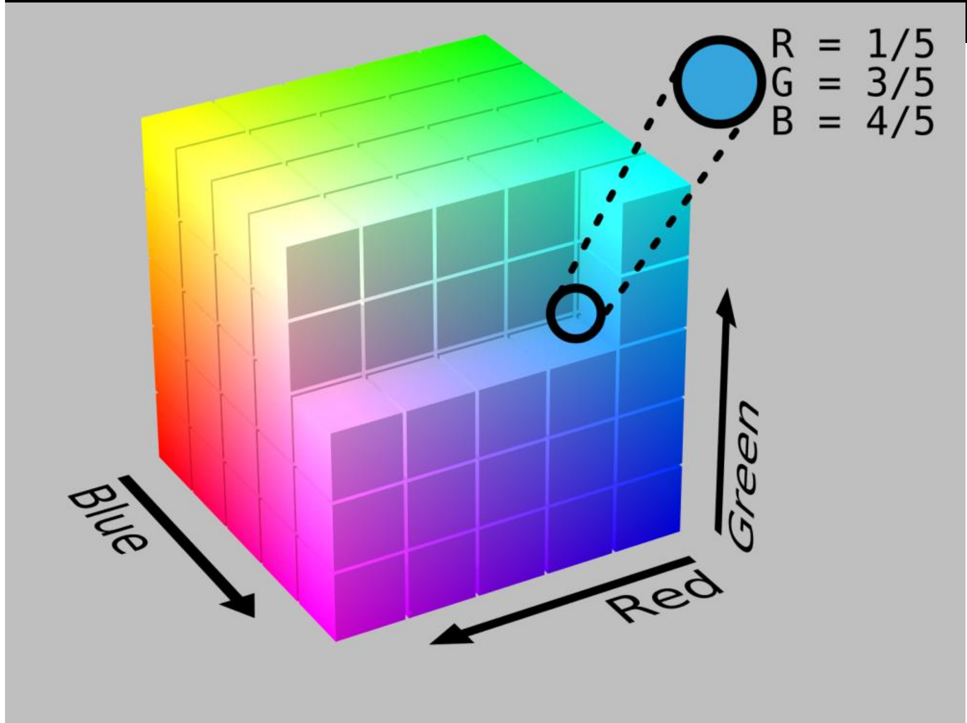

RGB: Adding colors

- This feature of human vision is very useful in photography, printing, art and television.
- The color of a TV screen or a color computer monitor consists of thousands of small dots—called phosphors—which are illuminated in red, green, or blue by an electron beam.
- All the colors we see on a tv screen are created by different intensities and combinations of red, green and blue

The RGB color space

- Each color is determined by three values (R, G, B) giving the relative proportions of the three basic colors.
- we usually have one byte (8-bits) for each of the 3 color elements
  - So we have 256 possible values for each item
  - That gives us 2^{24} possible colors *(24-bit display, true-color, real color)*
  - This is often written as a 6-digit hexadecimal number, with R, G and B being between 0 and 255 (FF),
  - A color value usually occupies 24 bit.

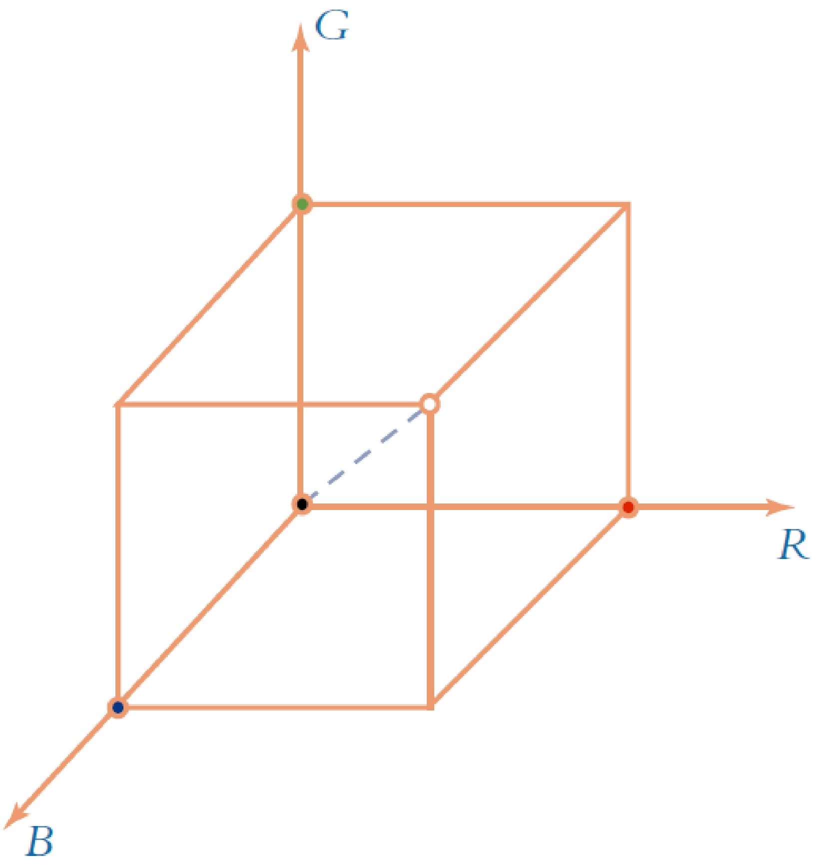

The RGB colour space

The RGB color space

- White light is produced by mixing with equal light proportions from all primary colors, so that the RGB color value of white is (100%, 100%, 100%).
- When the amounts of red, green, and blue light are equal, the result is a shade of gray, so that each value in the form (x, x, x) represents a gray color:
  - the higher the value of x, the lighter the gray will be.

The RGB color space

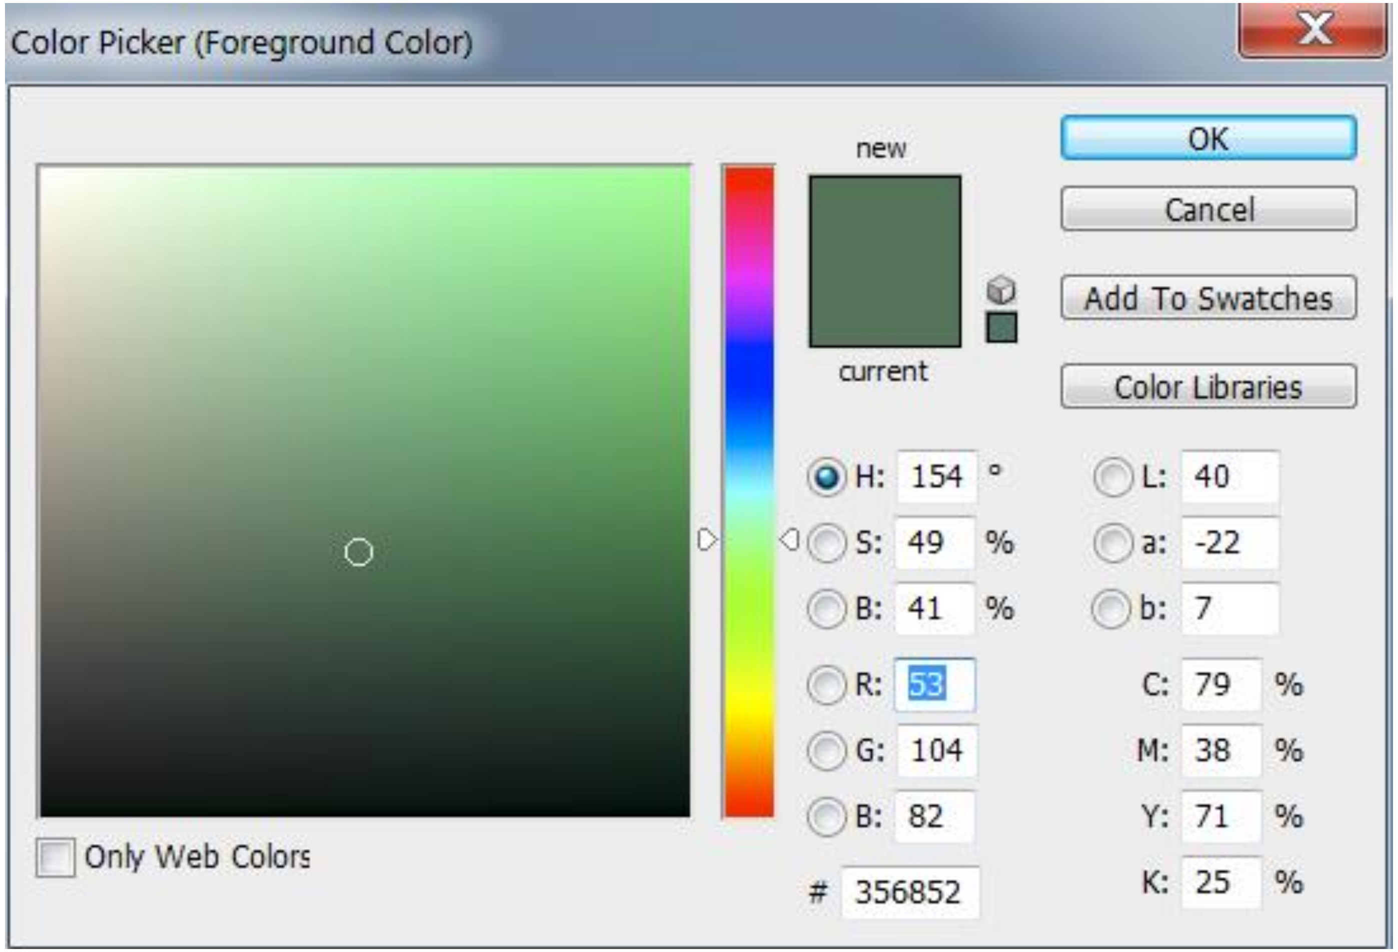

The RGB color space

- The number of bits used to store a color value – the color depth – determines how many different colors can be represented. The use of low color depth leads to posterization and loss of image detail, but reduces file size.
- Most computers store and reproduce 24-bit colors,
  - Mobile phones used to use less (8,12,16 bits), but now use 24 or 32 bits.
  - What does it mean when we say 32 bit display? (alpha)

The RGB color space

A photograph in 24, 8 (top), 4 and 1 (bottom) bit colour

The RGB color space

- Indexed color models, instead of storing 24-bit colors for each pixel, use 8-bit values that are essentially markers in a color table. A color table is a color palette used in the image.
- Using the index, the color value does not give the color, but a marker with a palette.
-

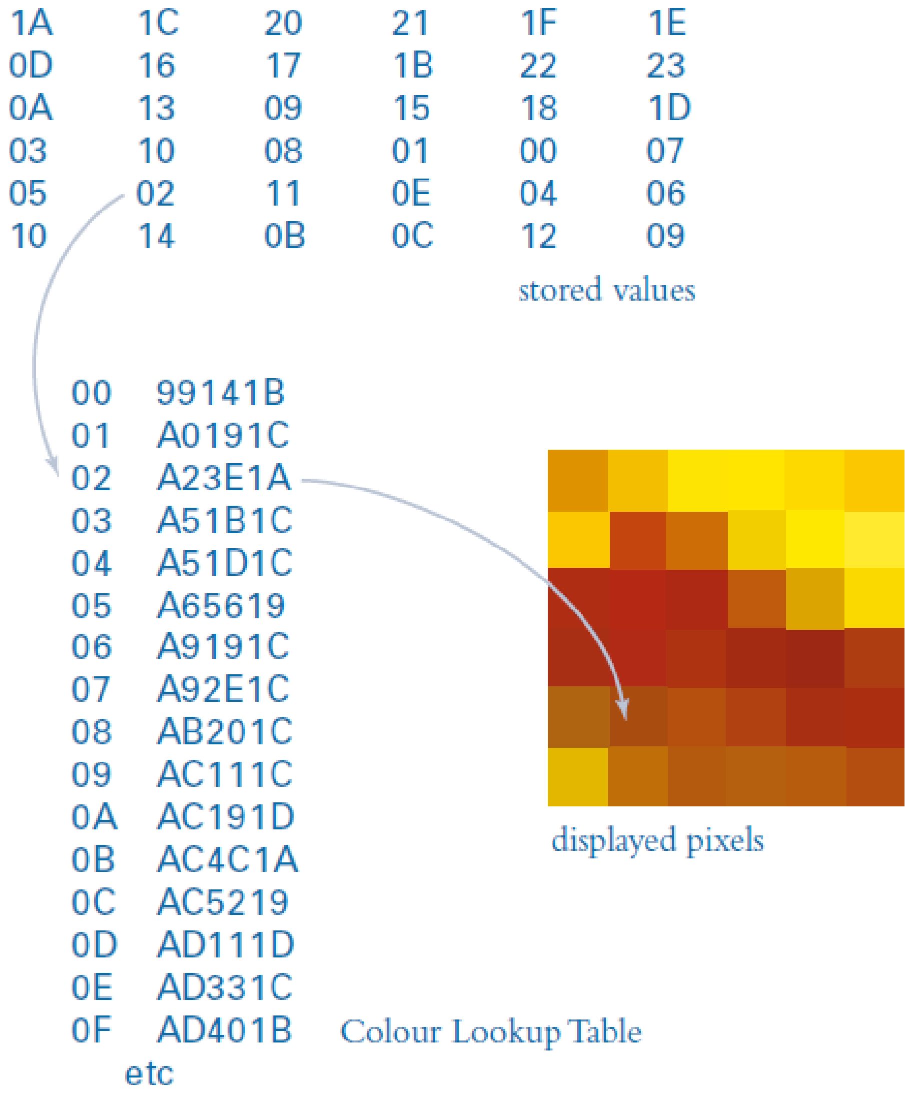

Using a colour table

The RGB color space

Images and their palettes

The RGB color space

- What if we have an 8 bits set (R, G, B)?
- And what if we have 16 bit for the whole color?
- Graphics memory (on graphics card)
- The screen memory is used to freshen up the screen.
  - We do not need to calculate the 76Hz image (76 times per second)

Color representation

The CMYK color space

CMYK: Colour subtraction

- When red, green and blue lights are mixed they produce white light.
- But you already know already that mixing the red, green and blue colors does not produce white--produces a dark gray.
- What's the difference?
  We can see what is not absorbed by the pigments, i.e. what is reflected and comes back to us
- White paint reflects all the pigments

Color subtraction

- To create red:
  - small particles reflect red light
  - **Reject** other colors

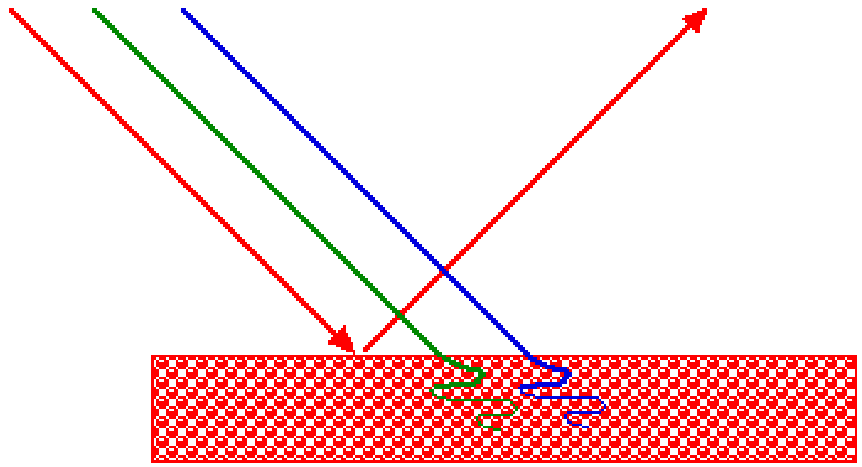

Color subtraction

- The union of red and green pigments to create white:
  - The red pigment will reject the green and blue light
  - The green pigment will reject the red and blue light.
  - In the end no color will be reflected

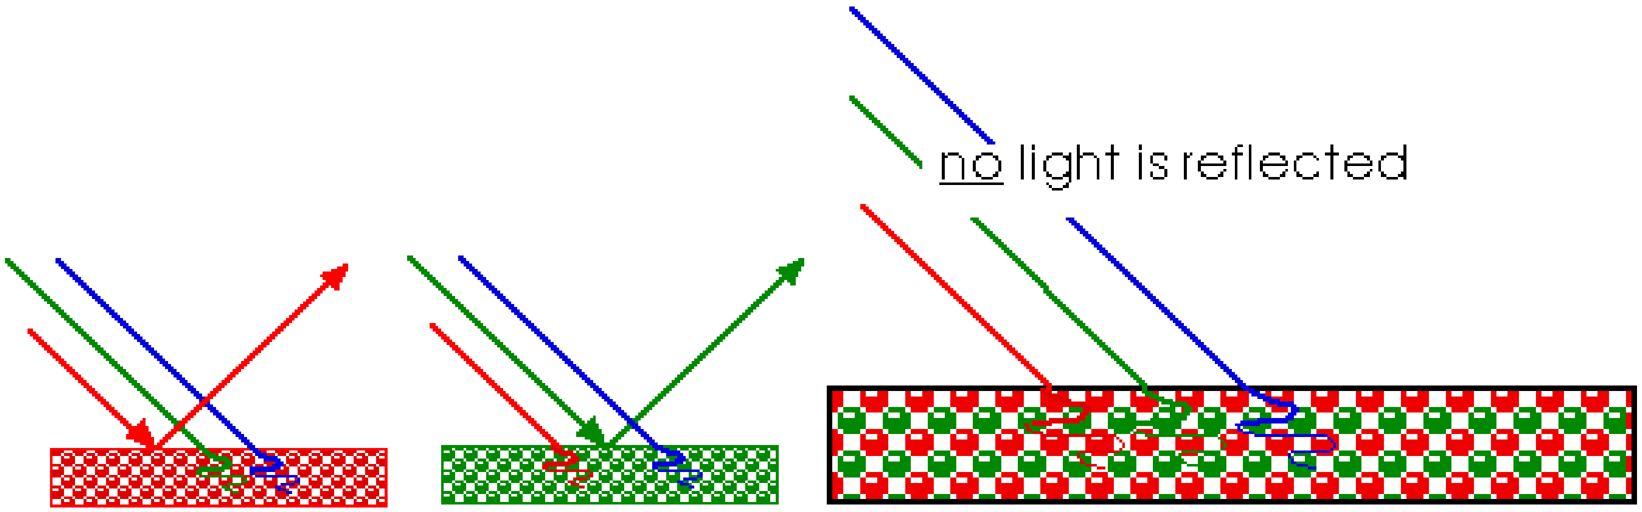

Color subtraction

- Practically speaking, most pigments do not discard all the color, so the end result will be a dark gray, instead of black
- The main role of Pigment is to reflect what you are not absorbing.
- When the paints come together the result will be the set of discarded characteristics.
- This is why it is known as **colour subtraction**

The CMYK color space

- Cyan, magenta and yellow are the main colors of the CMYK subtractive method. These are the complementary colors of red, green and blue, respectively.
- Thin layers of ink absorb certain elements of the light incident, and are used as a printing method.

Color subtraction

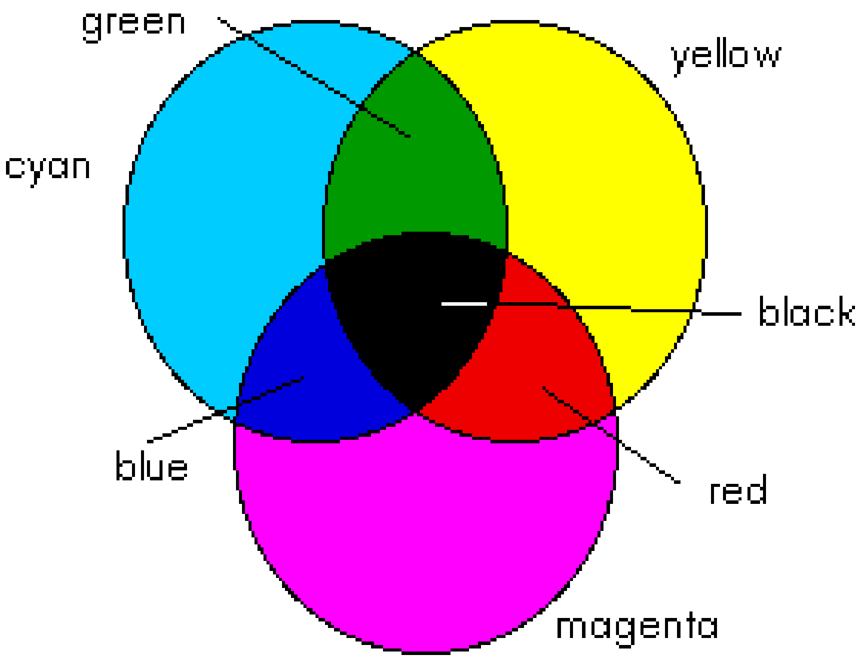

- **Color subtraction** is the basis for painting in printing. Why;

Color subtraction

- The pigments of cyan, yellow, and magenta are used to create infinite colors.
- **Black** is used to make shades more intense

The CMYK color space

- CMYK colour gamut, represents easily printable colors, is smaller than RGB gamut, but some CMYK colors are outside of RGB gamut.

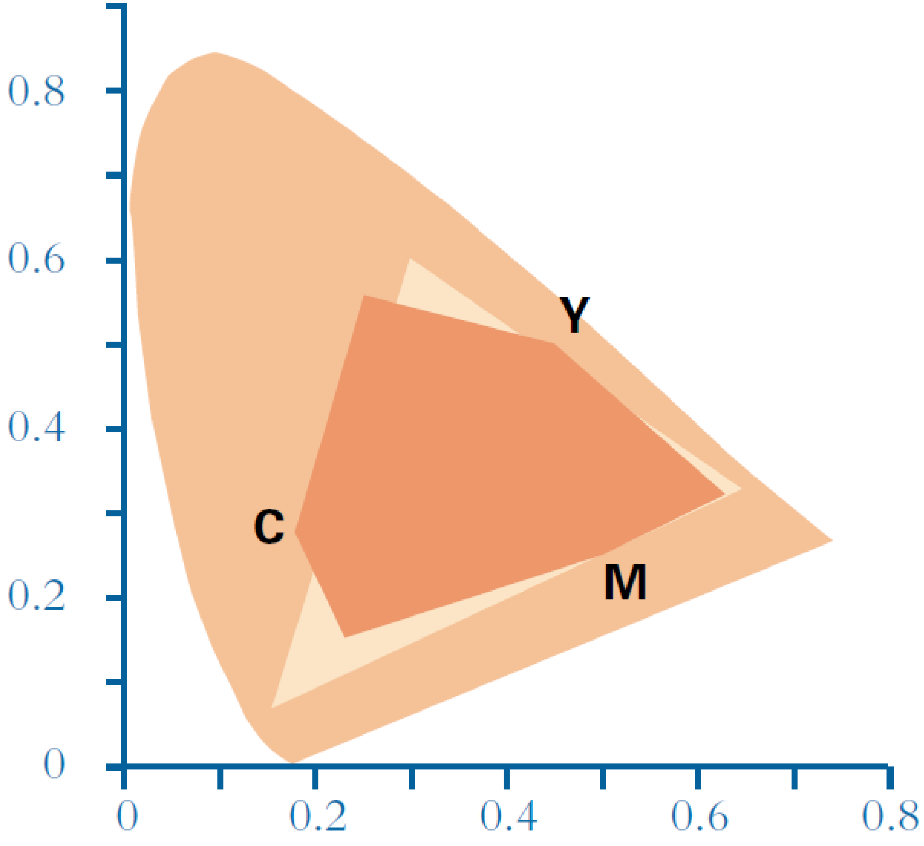

The CMYK gamut

Colour Matching

- Color addition

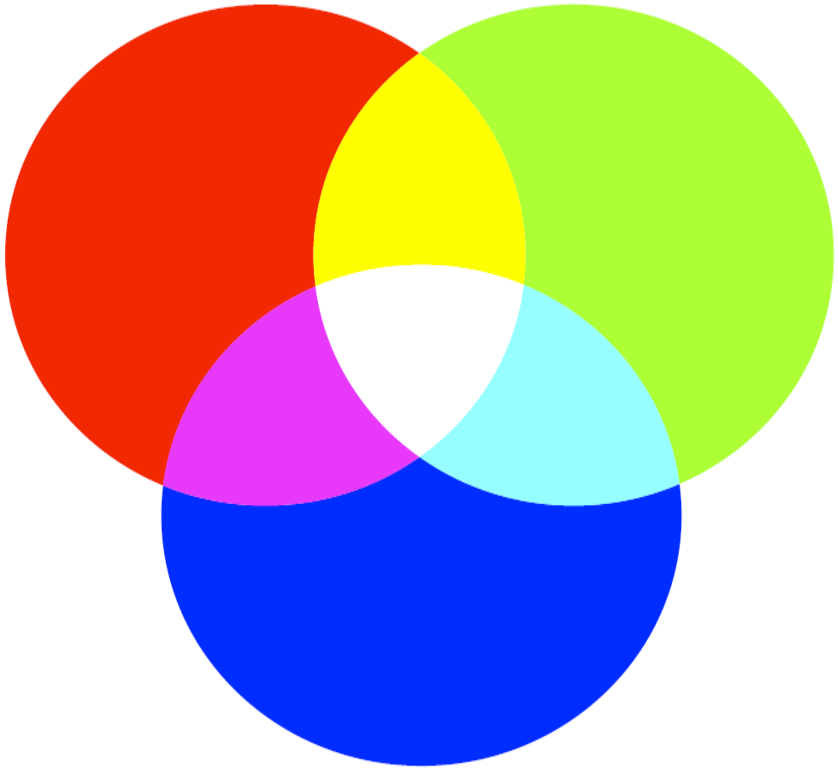

- Color subtraction

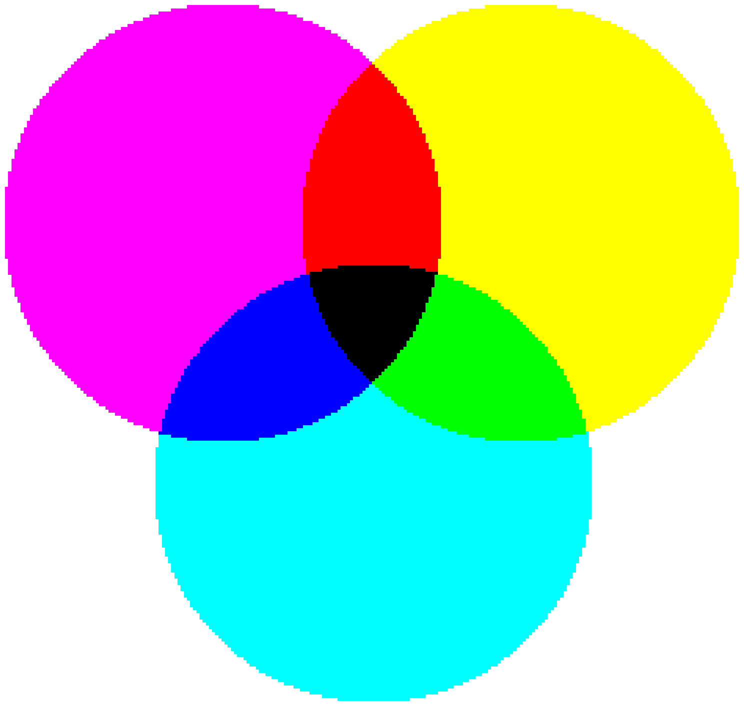

Color representation

The HSL color space

The HSL color space

- Hue, saturation and lightness can be combined into a three-dimensional double-cone. Any color can be determined by its H, S and L elements.

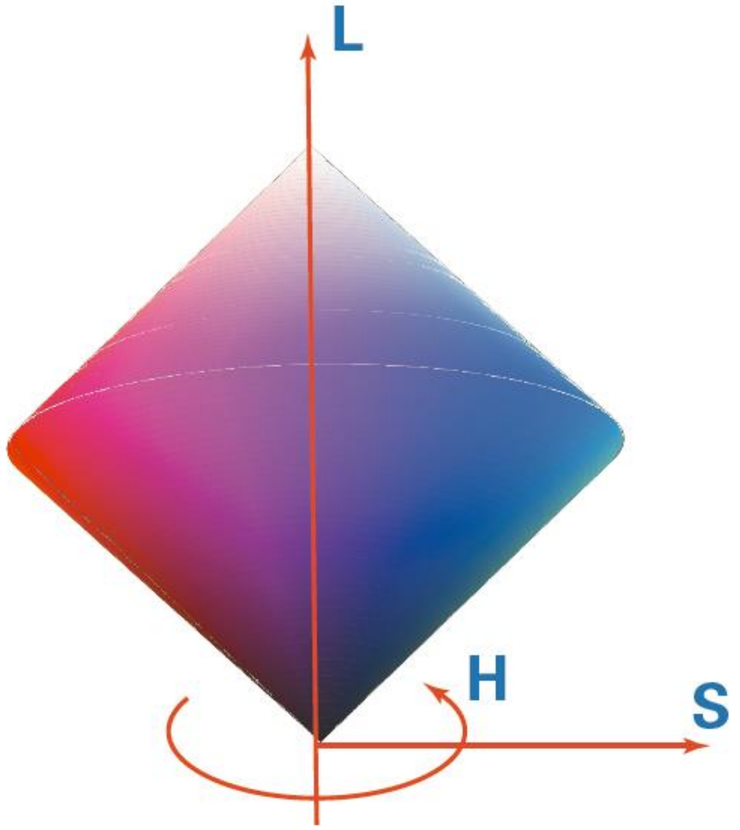

The HSL colour solid

The HSL color space

- 3 main color characteristics
  - **H**ue
  - **S**aturation
  - Brightness
    - **L**ightness
  - **HSL**

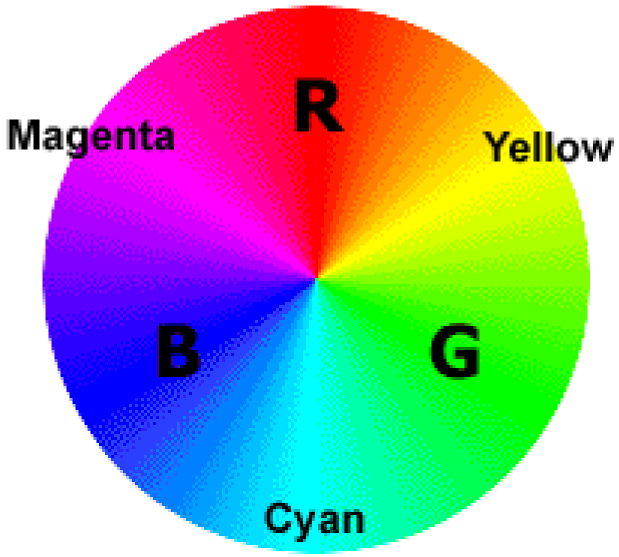

- Wheel of colors
  - The main colors R, G and B
    are at equal distances inside the wheel
  - the other colors of the spectrum can be created from mixtures between the main colors

The HSL color space

- Hue
  - *each specific point inside the wheel, from 0 to 360 degrees, is referred to as* **hue**
  - Specifies the specific style of the color.
  - "Hue" differs from "color" because color is also determined by saturation or brightness, along with its hue

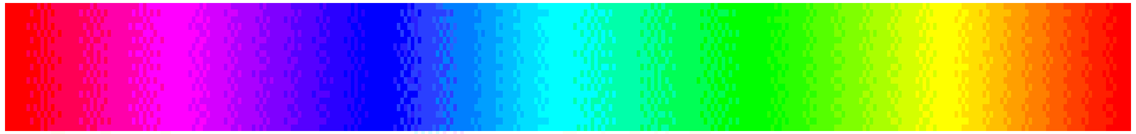

The HSL color space

- Saturation
  - the intensity of the hue:
  - from gray tone (no saturation)
  - In deep color (high saturation).

The HSL color space

- Brightness:
  - is the relative brightness or darkness of a color
  - From black (no brightness)
  - to white (full brightness)

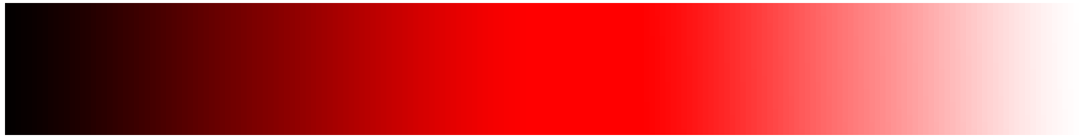

no brightness **full brightness**

Color Mixing/Color Gamuts

- Small size of the print gamut

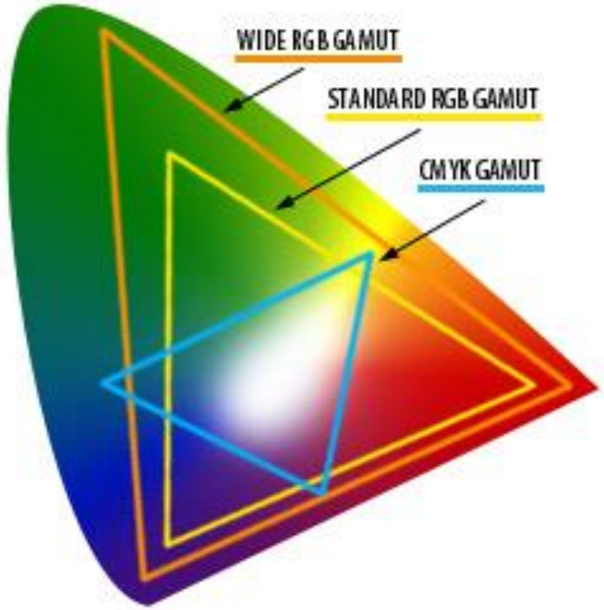

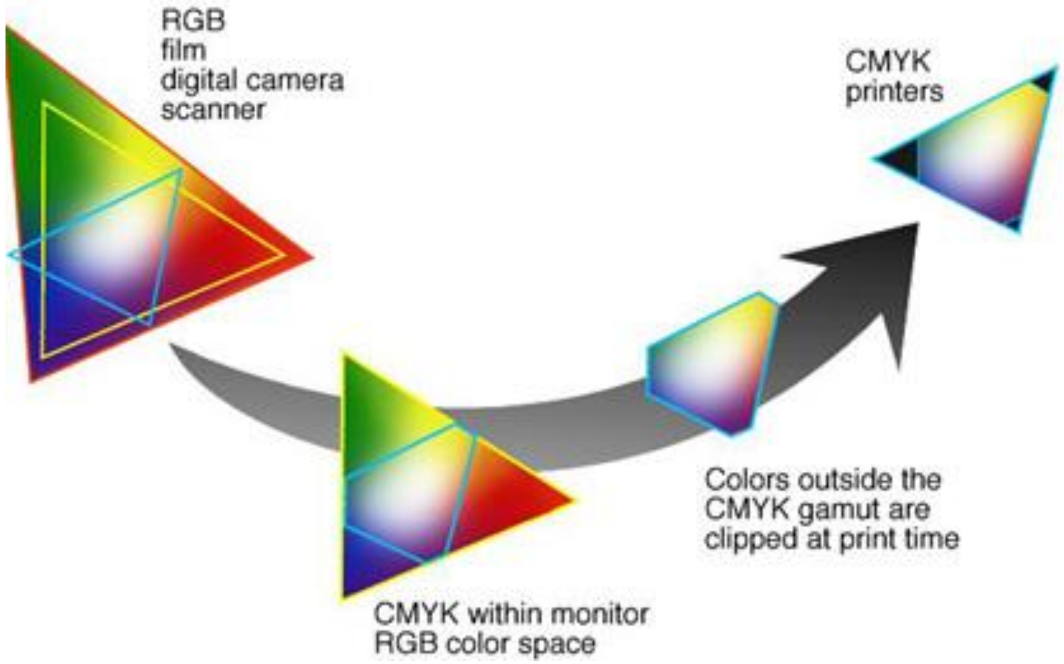

Why perceptual uniformity?

- Imagine encoding a gradient from green to blue in 10 steps
- Change hue to equal spacing increments in x-y space

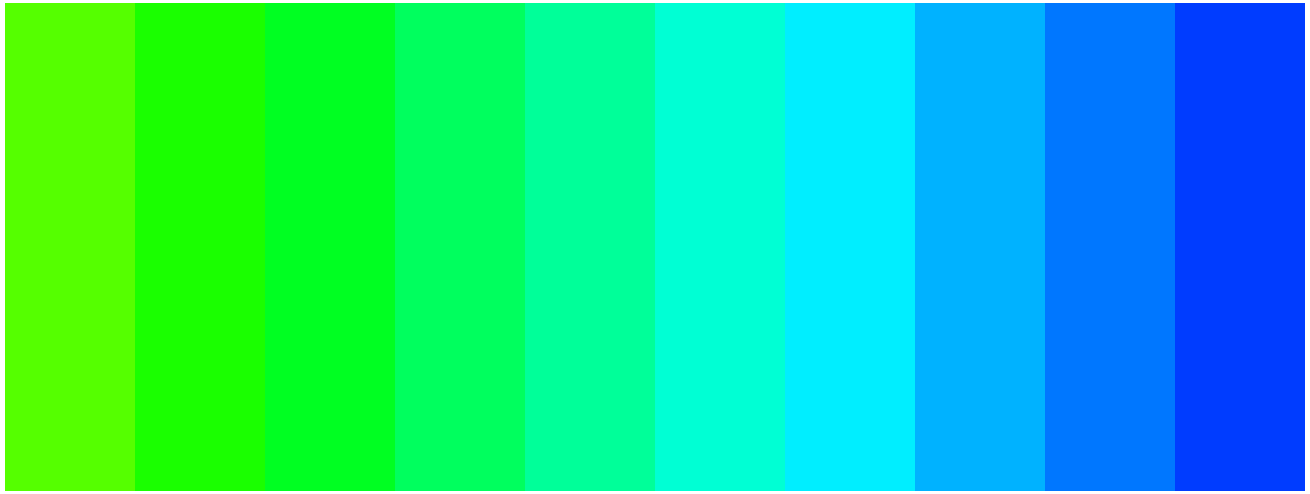

- What went wrong?

Why perceptual uniformity?

- Look back at the CIE chromaticity diagram
  - 520nm at 560nm: everything looks green
  - 460nm to 500nm: much wider range of colors
  - Explains why the gradient looked mostly green with a sudden turn to blue at the end
- Use a perceptually uniform color space

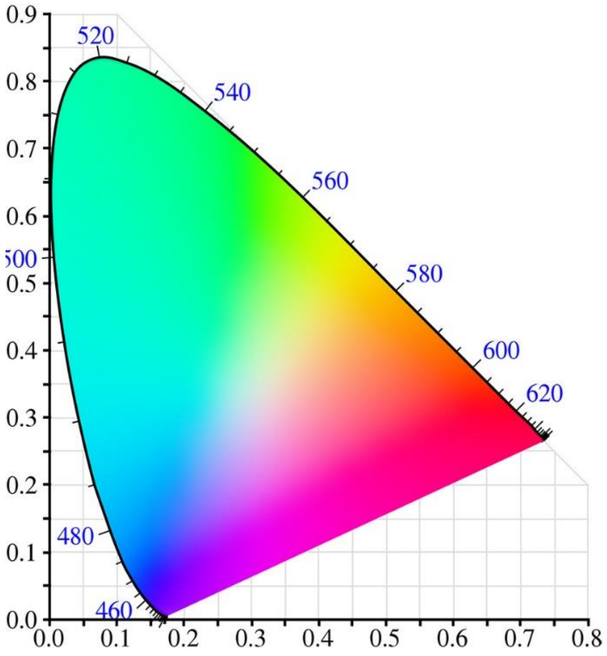

- https://programmingdesignsystems.com/color/perceptually-uniform-color-spaces/

Midterm Exam

- Date: 08/03/2022
- Time: 10:30 – 12:00
- Room: 147

- **Note: It is strictly forbidden to carry food and drink within the classrooms**.
- You are not allowed to carry notes, cassette, mobile, smart watch, tablet with you. You can have ONLY pen/pencil (you are encouraged to have 4 different colors of pen), ruler, simple calculator (cos/sin/tan).

Ray Shooting / Casting
Βολή Ακτίνας

The first humbling steps

Why Raytracing?

- Ray tracing is the basis of many algorithms that aim to have a realistic visualization!
-

$$\boldsymbol{Physically\ Based\ Rendering\ Algorithms\ = f\ (Ray\ Tracing)}$$


GENESIS
PINK FLOYD

https://en.wikipedia.org/wiki/Ray_tracing_(graphics) https://benedikt-bitterli.me/resources/

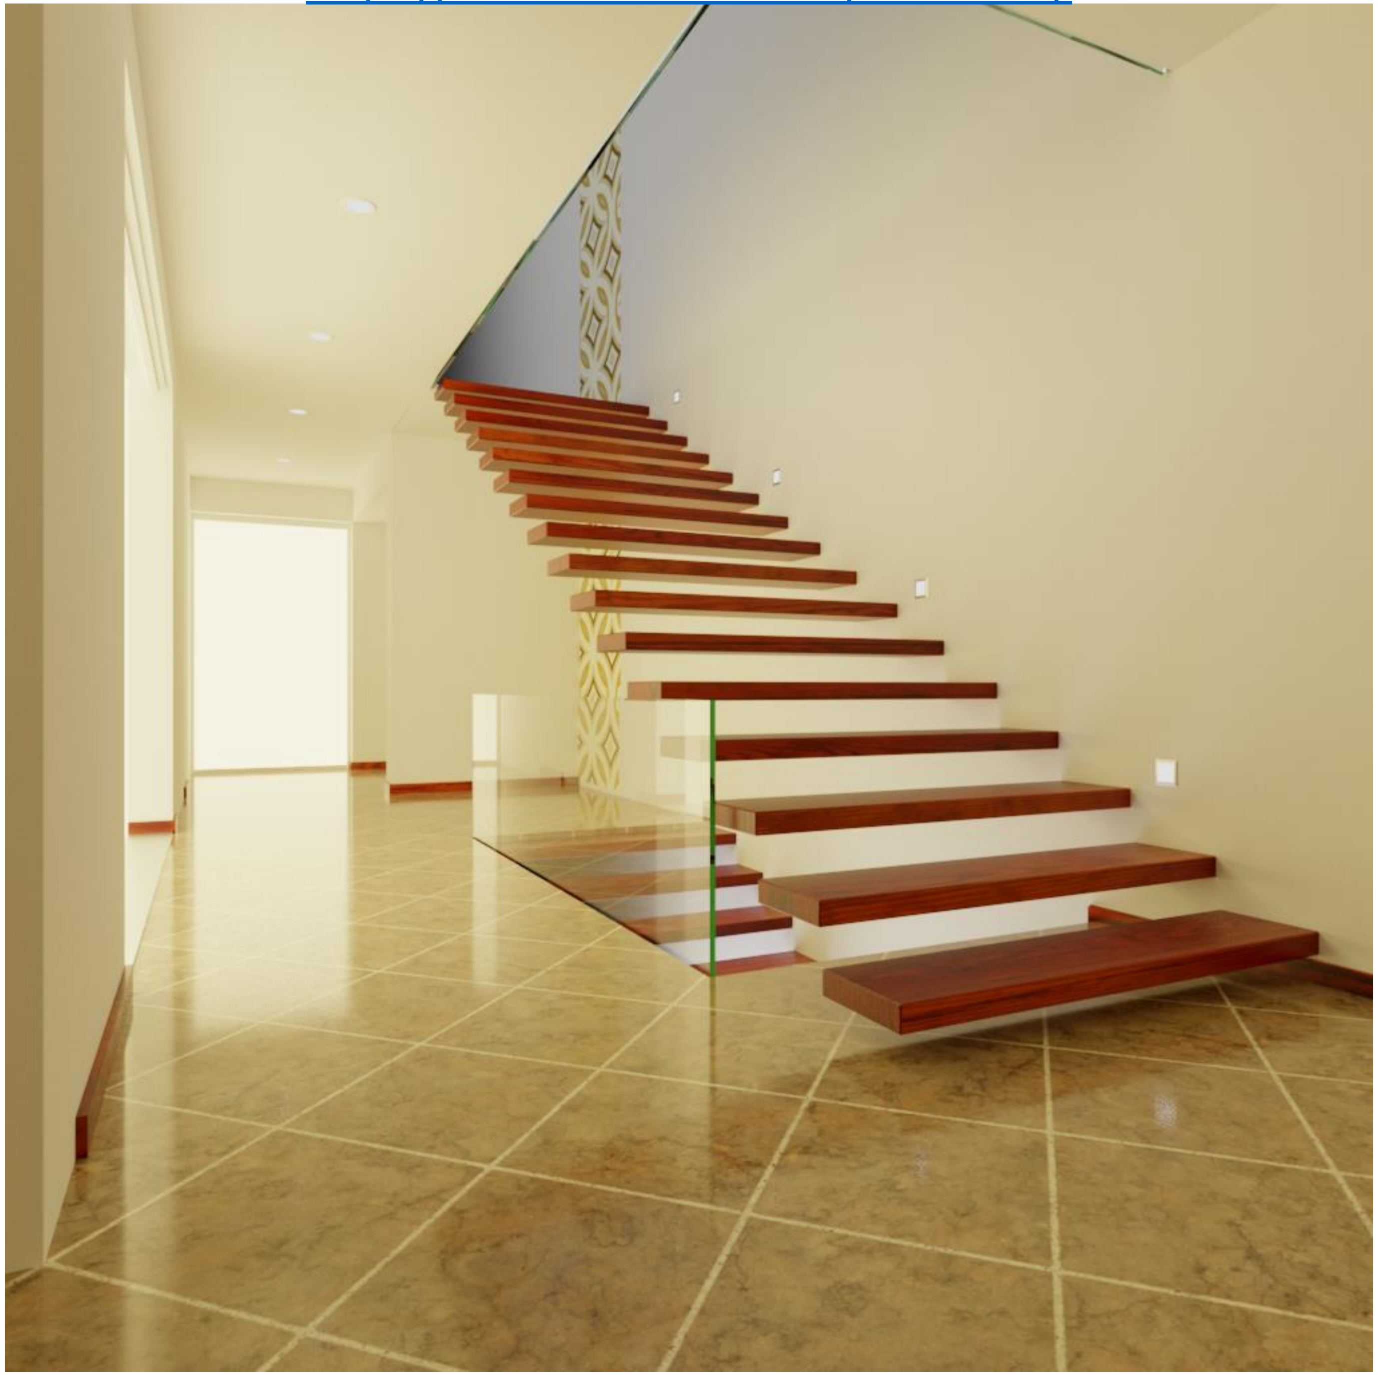

https://www.mitsuba-renderer.org/

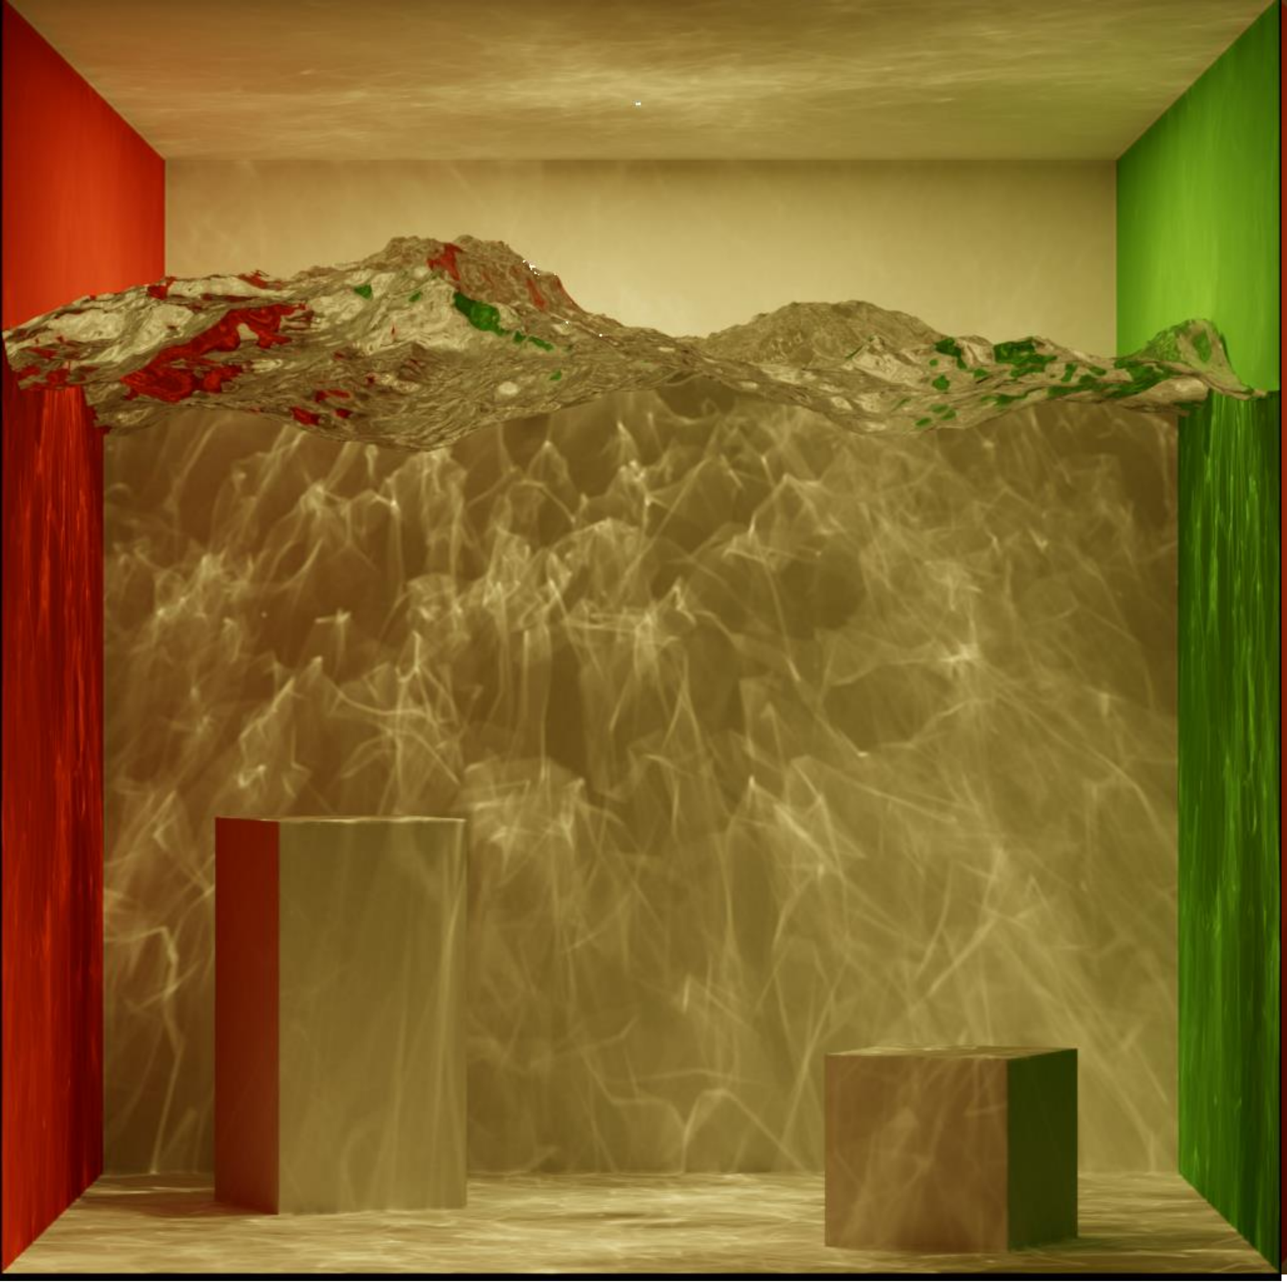
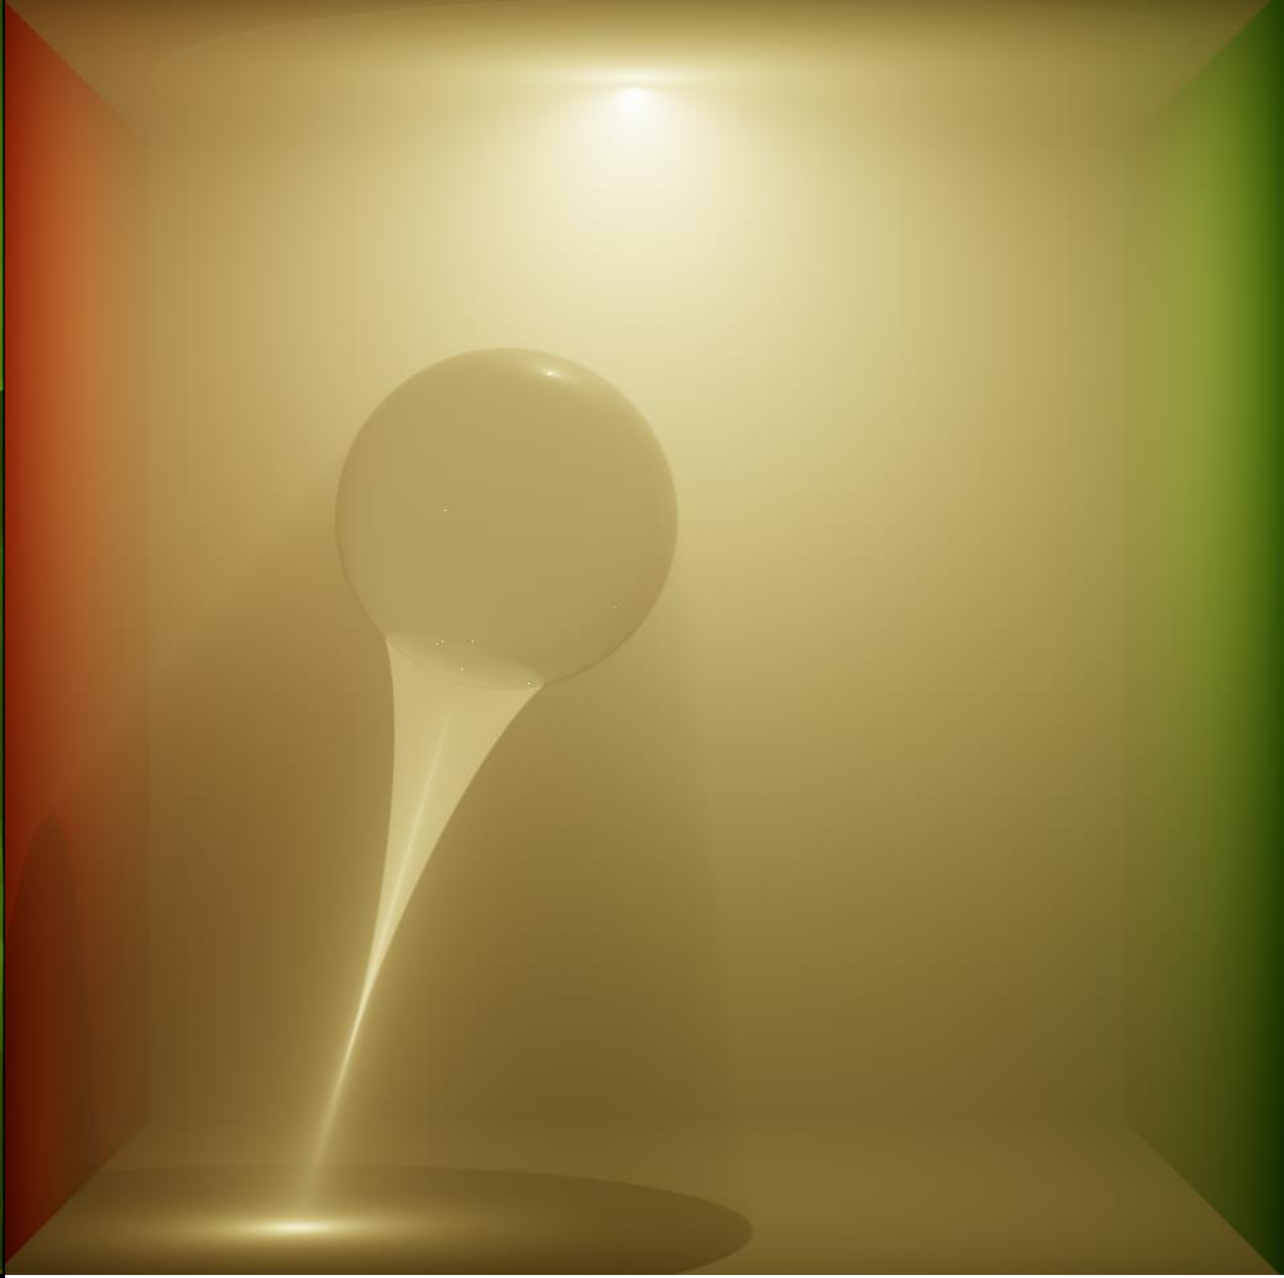

https://benedikt-bitterli.me/resources/

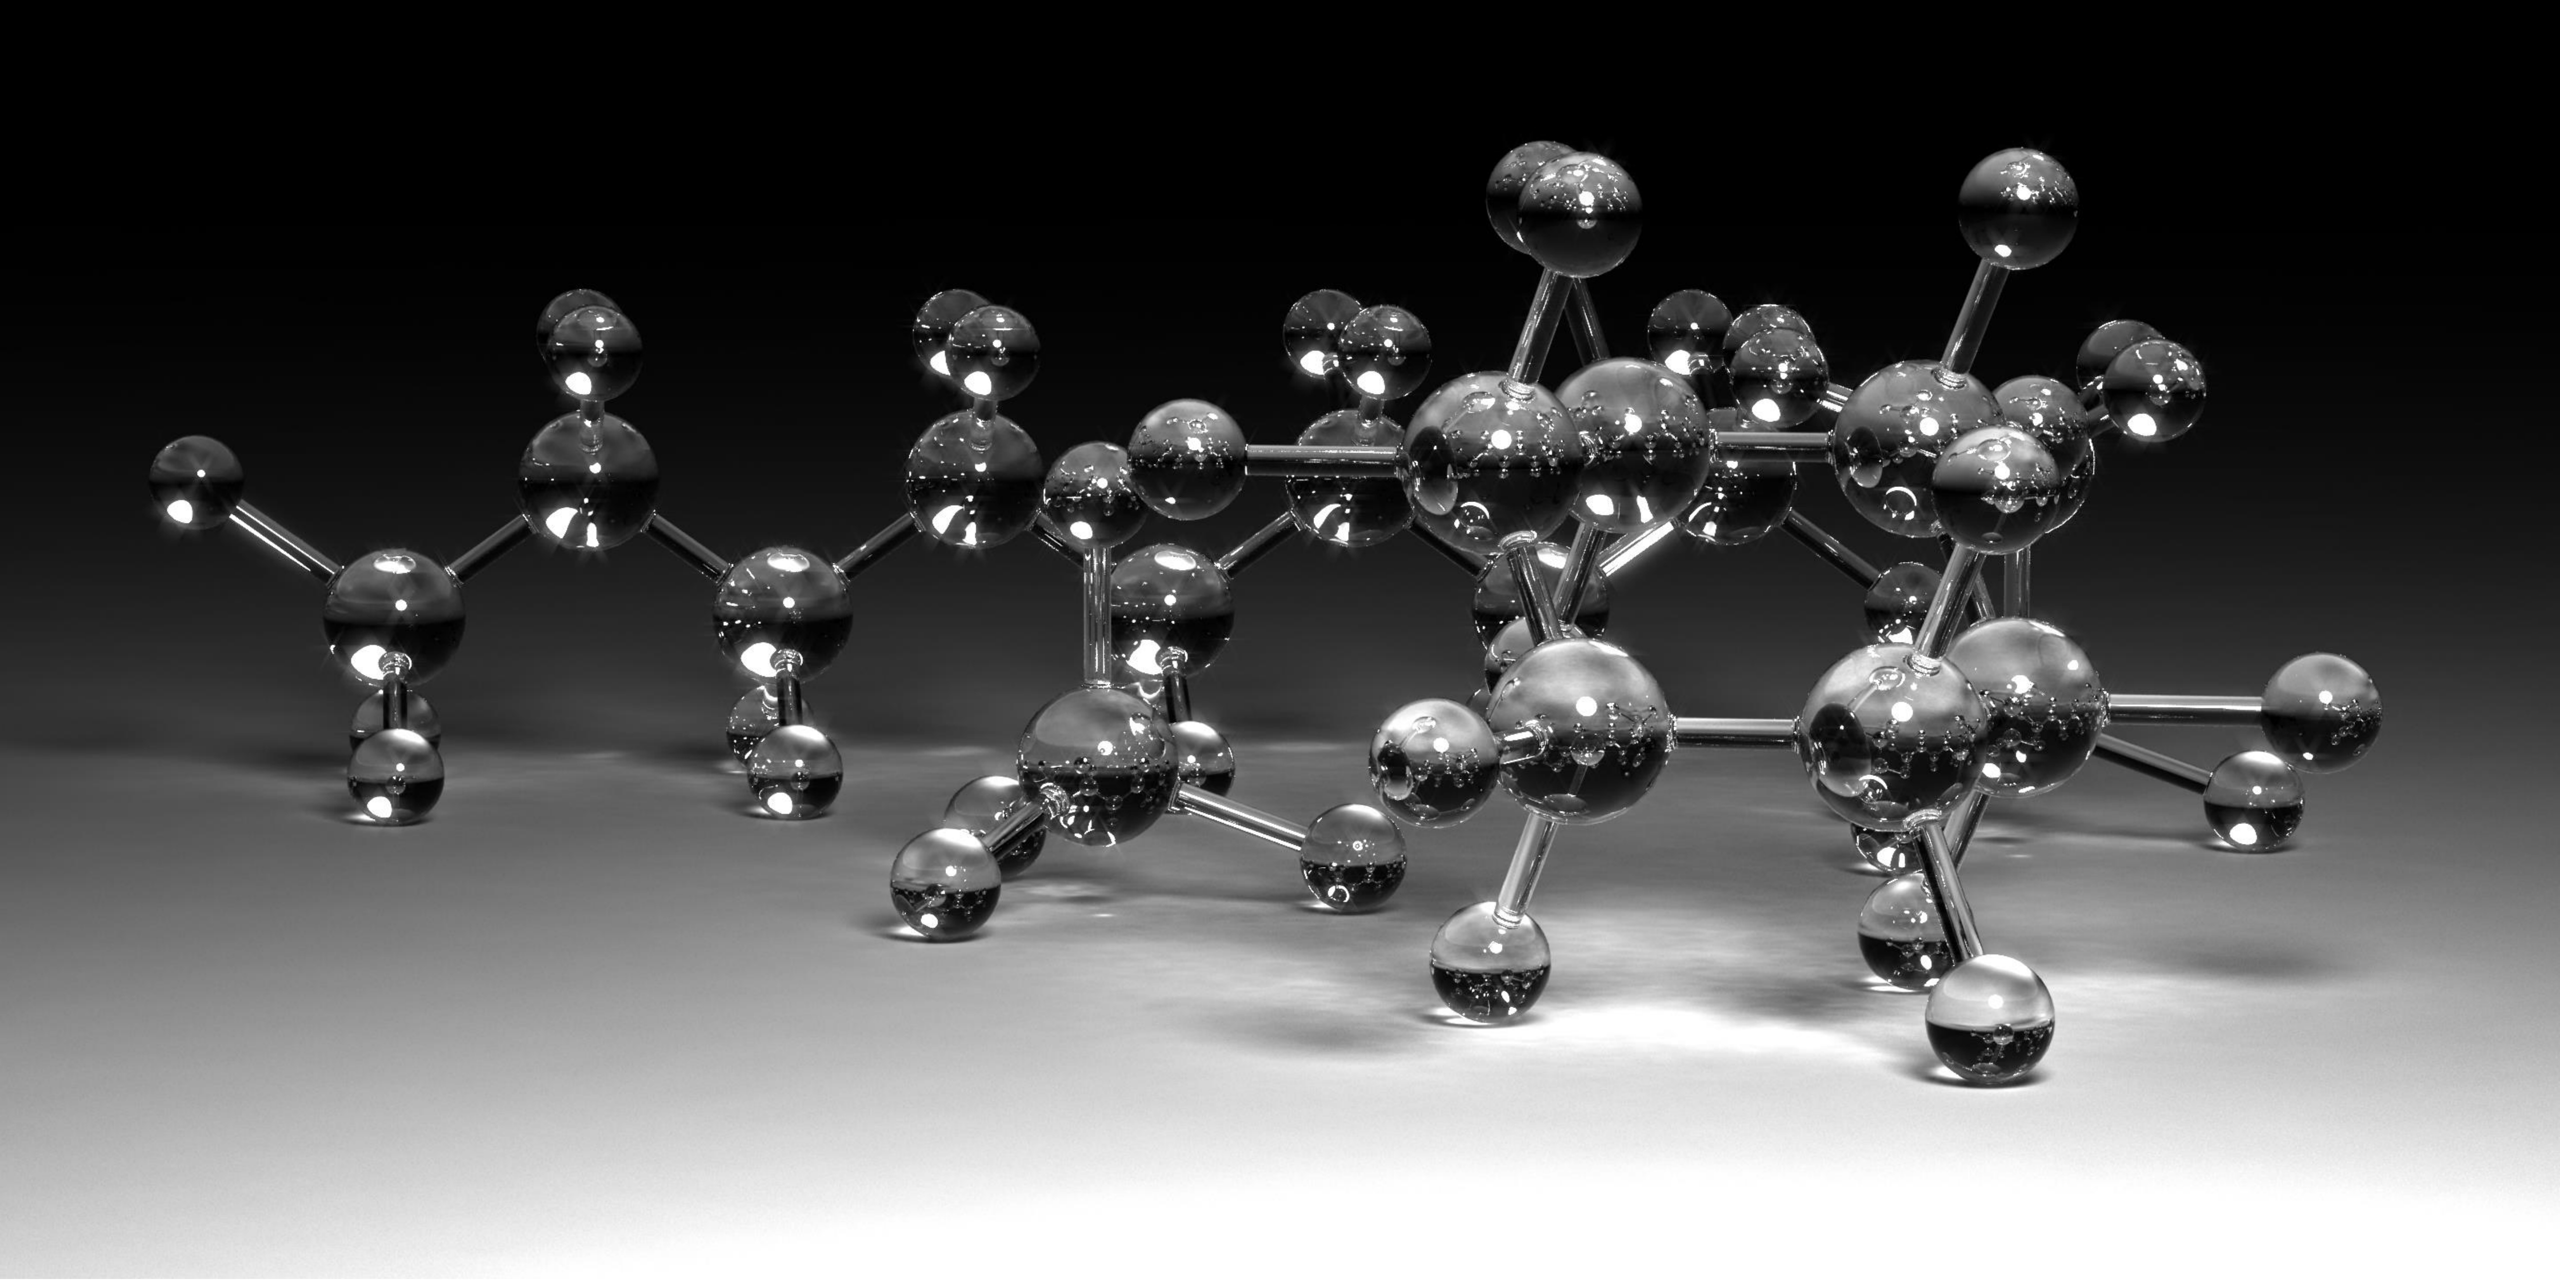

https://www.pbrt.org

The Art of Justin Holt
http://justinmichaelholt.com/tank-girl/

The Art of David Moratilla
http://www.davidmoratilla.com/
(With Post-Processing)

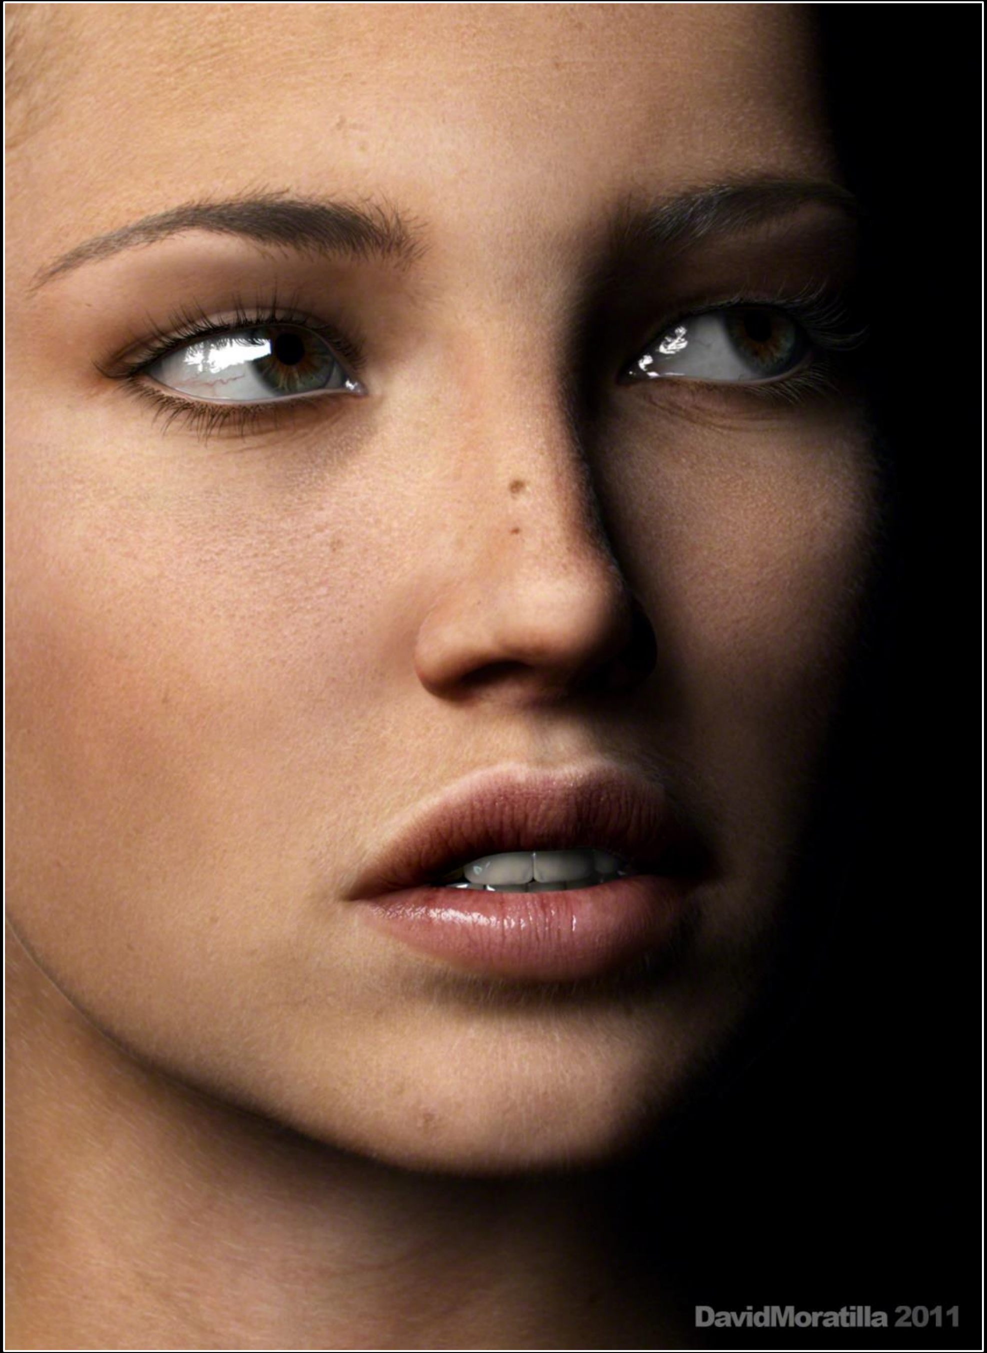

The Art of David Moratilla
http://www.davidmoratilla.com/
(With Post-Processing)

Two general approaches to creating images

A. Rendering Pipeline

- For each triangle
  - For each projected pixel

Project scene to the pixels

B. Ray Casting

- For each pixel
  - For each object

Send pixels to the scene

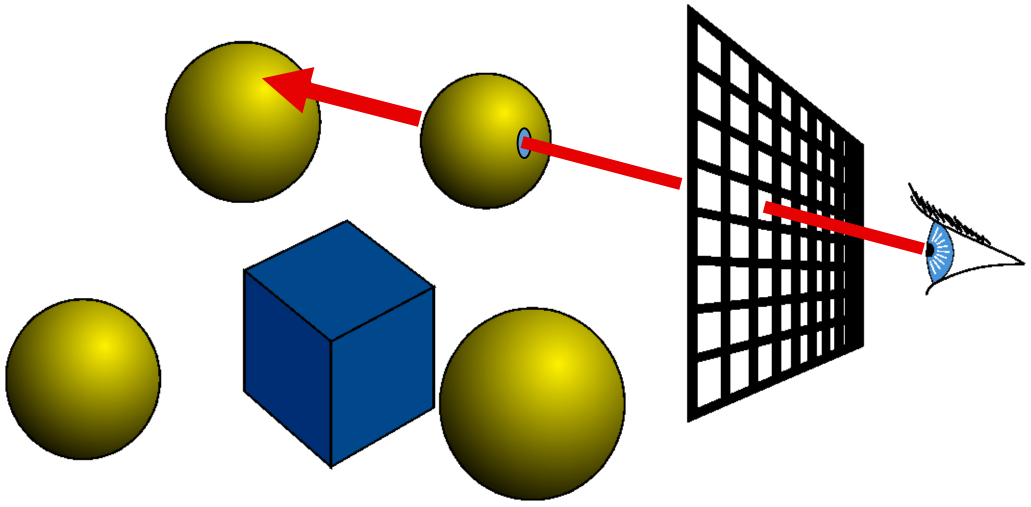

An allegory with painting

Albrecht Durer 1525

Image from Hitachi's Viewseum at http://www.viewseum.com/

- We have an artist who looks through a frame and carries what he sees on the canvas
- The frame consists of a grid with squares; similarly the canvas
  - equal in number but not necessarily in size
  - (*scene pixels, screen pixels*)
- Conditions/problems
  - Quality, movement of the head, process is instantaneous

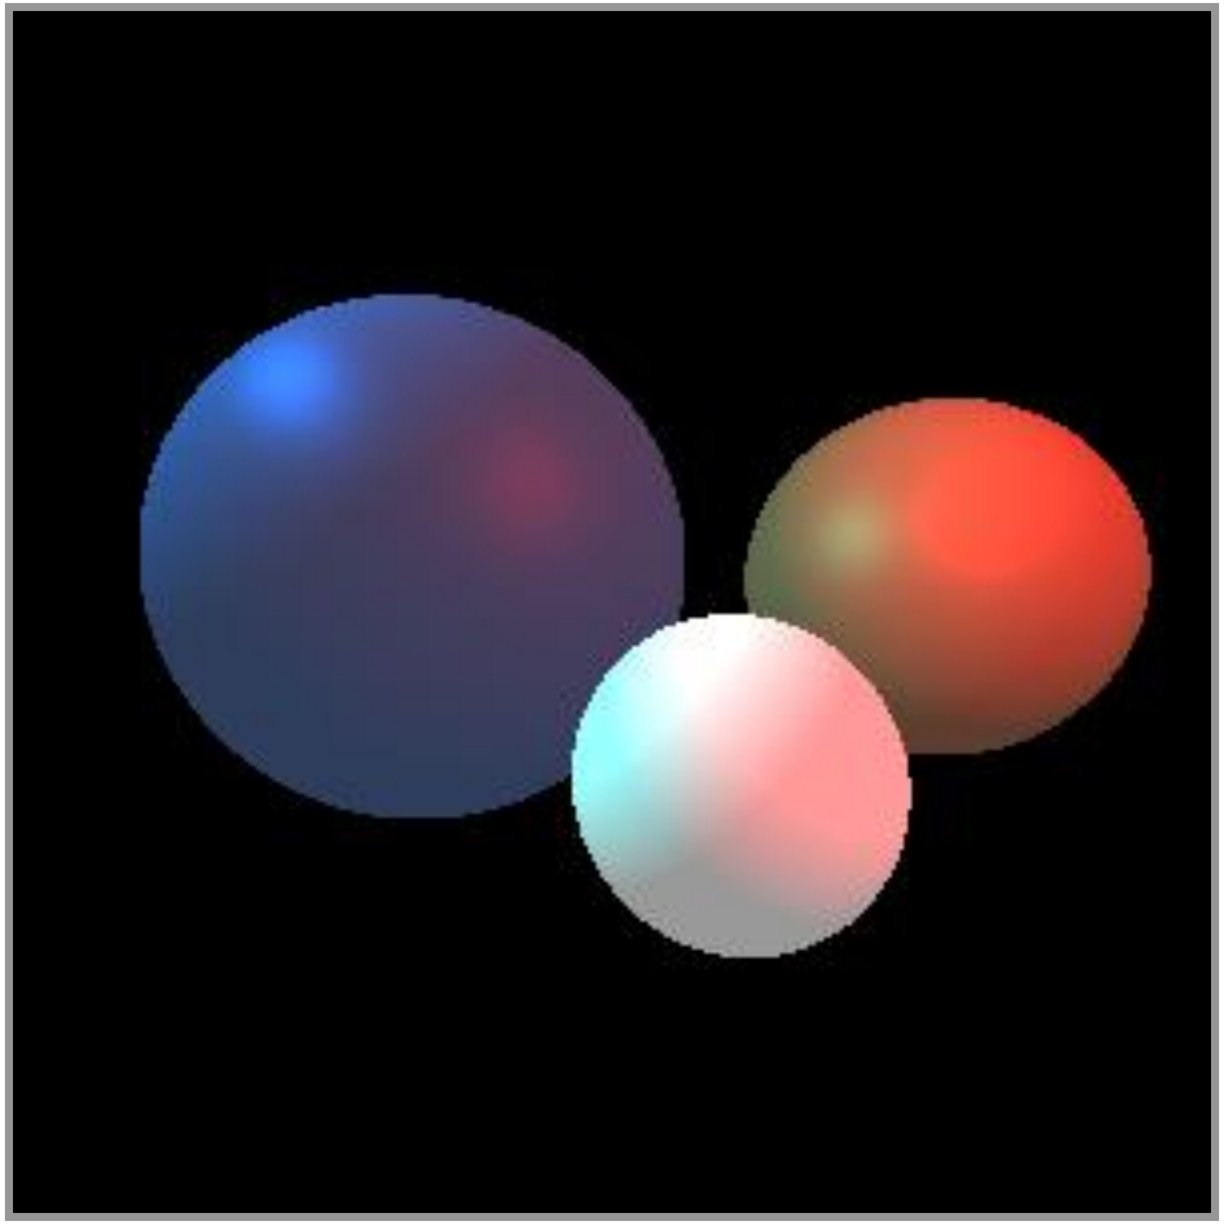

The 3 elements of the allegory

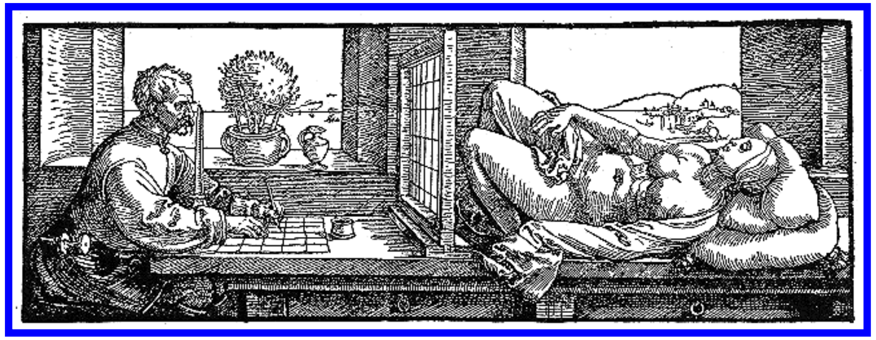

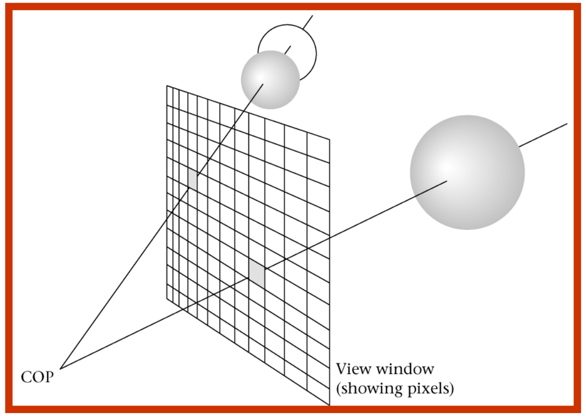

A. The scene
   the lady — the spheres
B. the view
   the eye & the tilar — a simple camera
C. the rendering process
   painting — ray casting
   i. specify the color of the pixel
   ii. painting the pixel on the screen

A. The Scene

- A group of objects
- Each object has:
  - geometric properties => the shape
  - material properties = > how light reflects, how and if it emits light
- World coordinates
  - right-hand system

B. The view - Simplified Camera

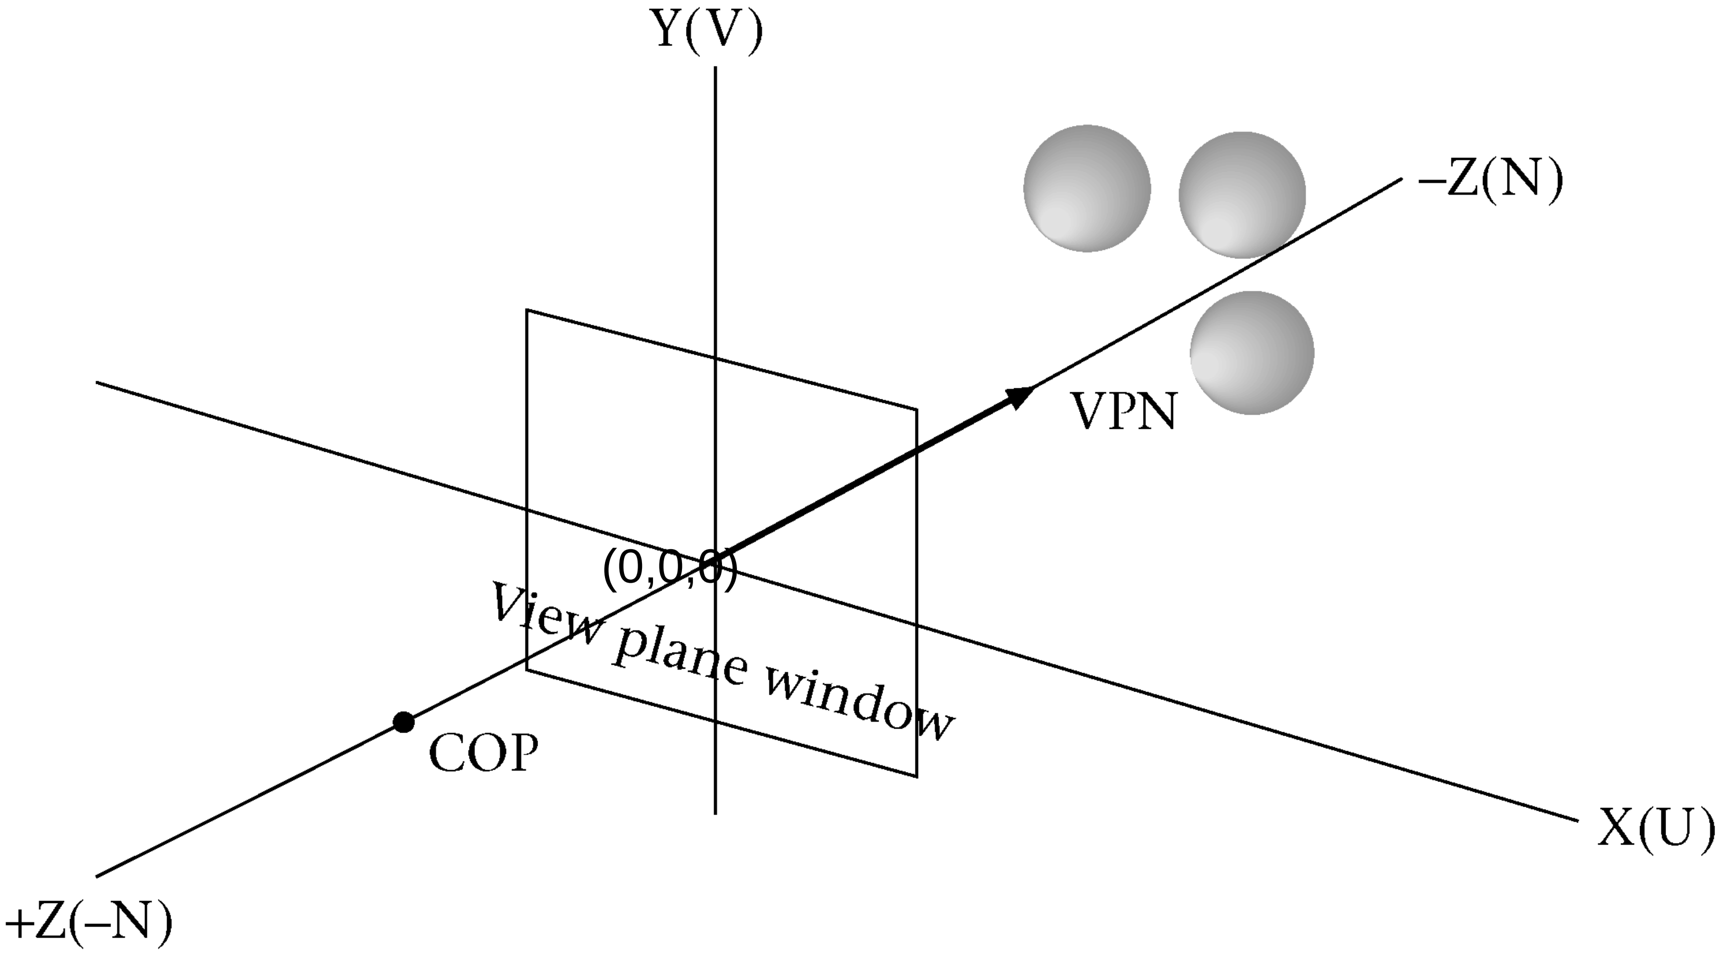

B. Simplified camera

- center of projection, *COP* or viewpoint, *vp*
- view plane distance, *d*
- view plane
- view plane normal, *VPN*
- view plane window

- Viewing coordinate system **uvn**
  - We put the u-axis on the positive x-axis, the v-axis on the positive y-axis and the n-axis on the negative z- axis
  - It's left-hand system

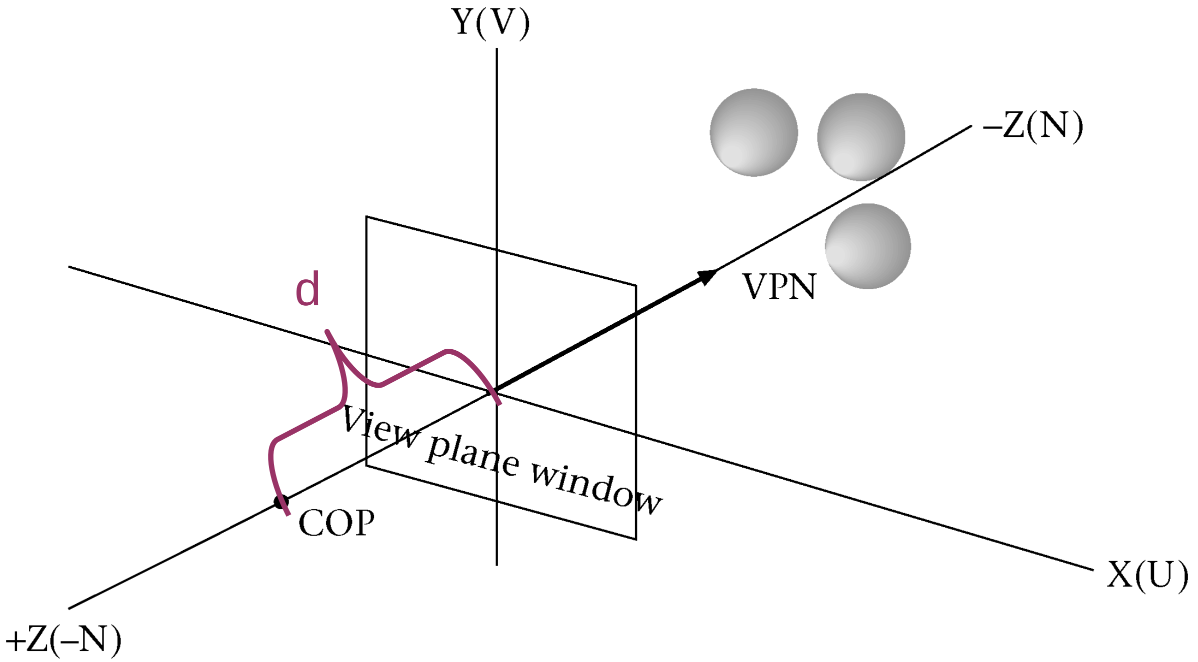

View volume

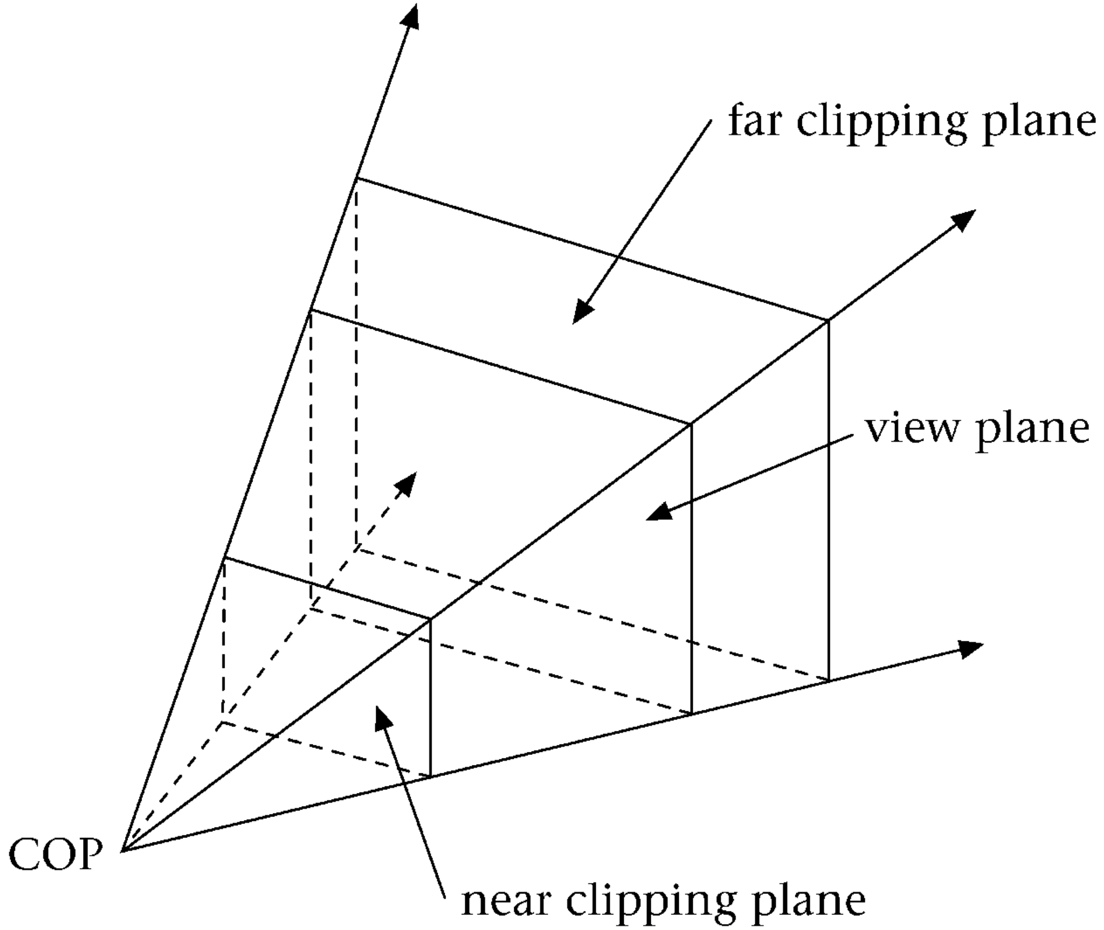

C. Rendering Process

- We shoot a ray from the center of each pixel
  - ‘Backwards’ ray-casting
- To define each ray we need 2 points: the COP, and the center of the pixel
- We look for the closest intersection of the ray with the scene and take the color of the object at this point
  - Repeat
- If there is no intersection, return black (or the color we set as background)

C. The definition of ray

- $M \times N$ pixels on view plane window
- Suppose we have a pixel i, j
- Corresponds to the rectangle
  - *width* $W = \frac{x_{max} - xmin}{M}$
  - *height* $H = (y_{max} - ymin)/N$
- The ray passes through the center of pixel
- Therefore the ray passes through the point

$$(x_{min} + W \times (i + 0.5), y_{min} + H \times (j + 0.5), 0.0)$$

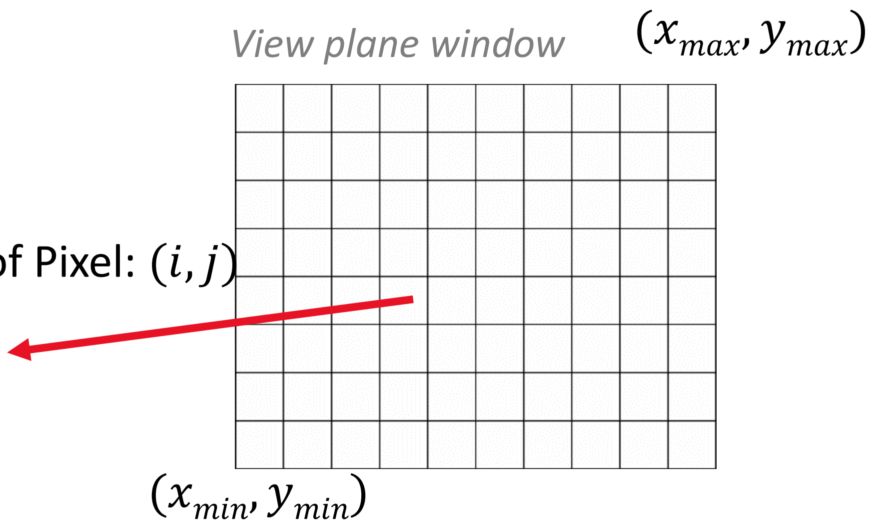

C. The definition of ray

- So a ray that starts at COP and passes through the pixel (i,j) is defined as:

$$p(t) = \big(x(t), y(t), z(t)\big) =$$
$$= (t \times (x_{min} + W \times (i + 0.5)),\ t \times (y_{min} + H \times (j + 0.5)),\ d - t \times d)$$

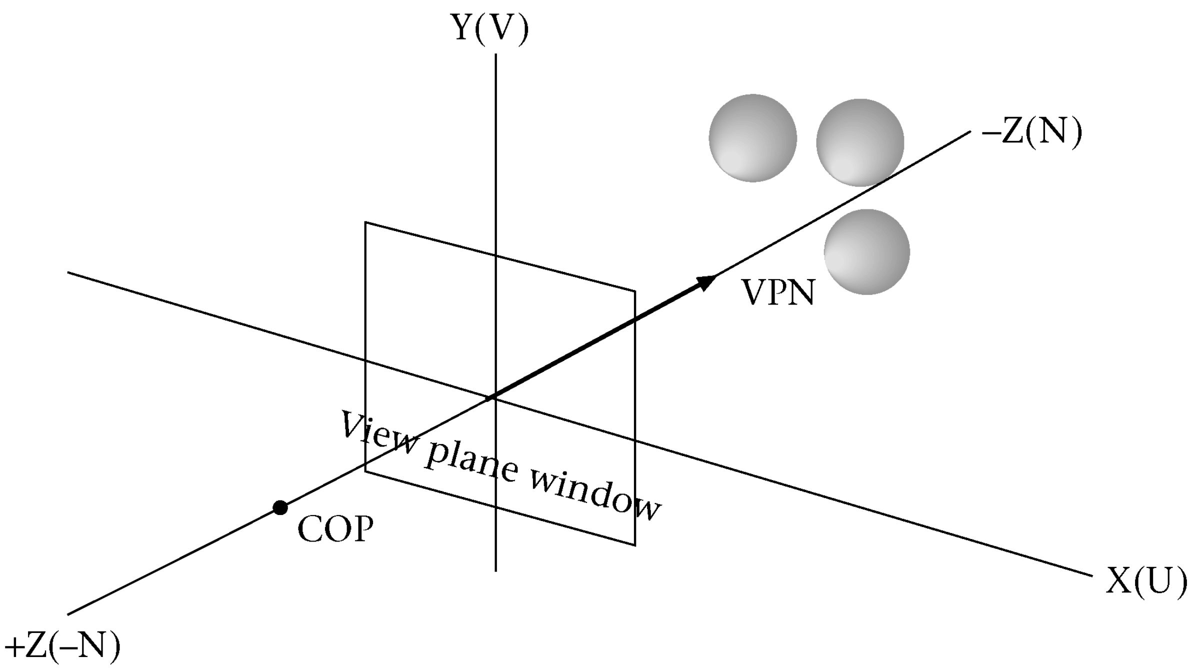

C. Rendering process, creation of the final image

- Mapping from the viewplane window to the viewport (e.g. screen)
  - window, viewport
  - linear mapping
- Map screen pixels (M by N window) to points in camera view plane

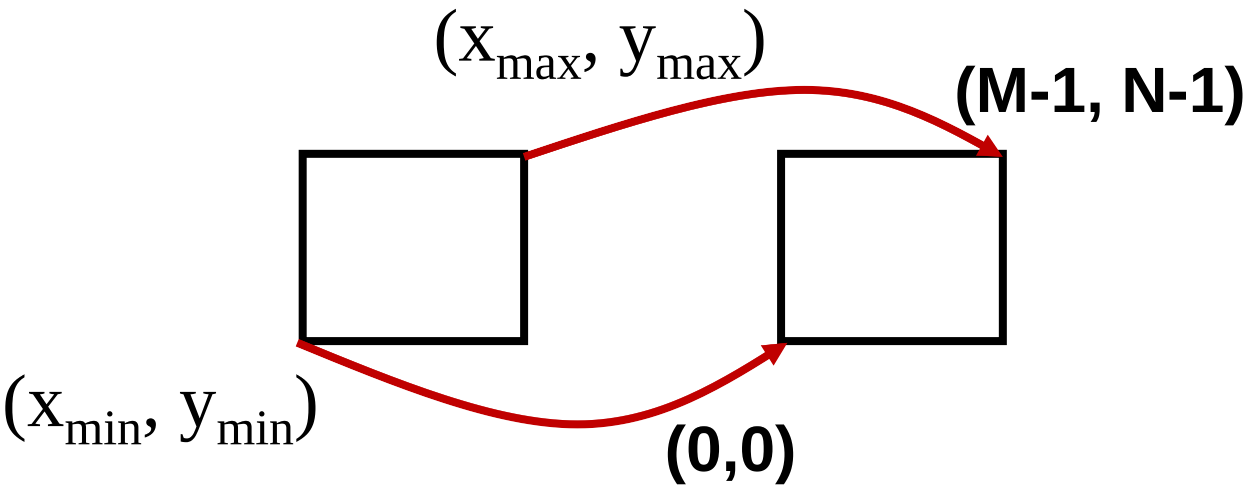

Painting metaphor

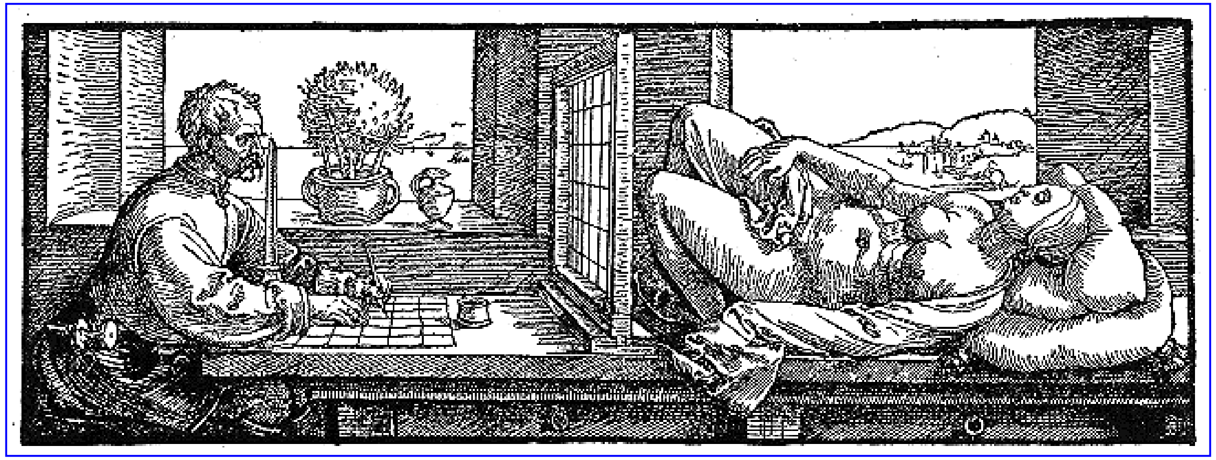

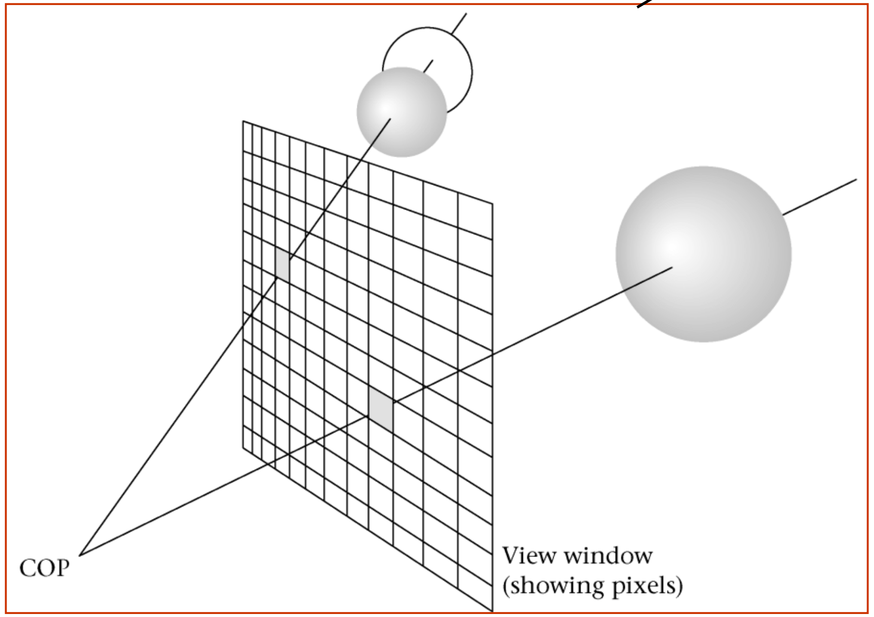

C. Mapping

from the viewplane window to the viewport

In the general case the correspondence between the two windows has:

$$x_v = vx\min + \frac{dv_x}{dw_x}(x - wx\min)$$

$$y_v = vy\min + \frac{dv_y}{dw_y}(y - wy\min)$$

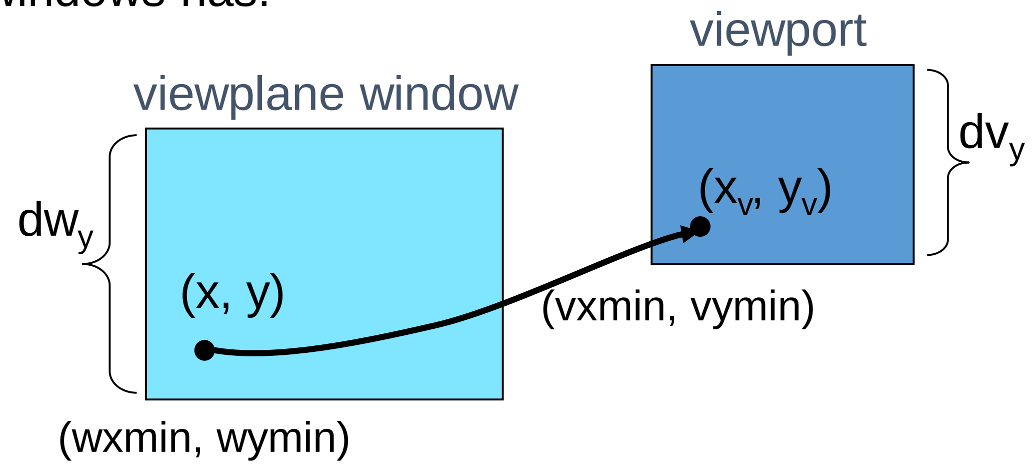

In our case we go from the center of the pixel to (integer) pixels on the screen

$$(wx\min + \frac{width}{2}, wy\min + \frac{height}{2}) \rightarrow (0,0)$$

(left bottom corner)

Ray-Casting

- **Ray-casting** is typically used to describe csg objects when objects are described with certain limits
- To ray casting is applied by specifying the objects intersected by a set of parallel lines coming from the *xy* layer along the *z* axis
- The *xy* level is known as the **firing plane**

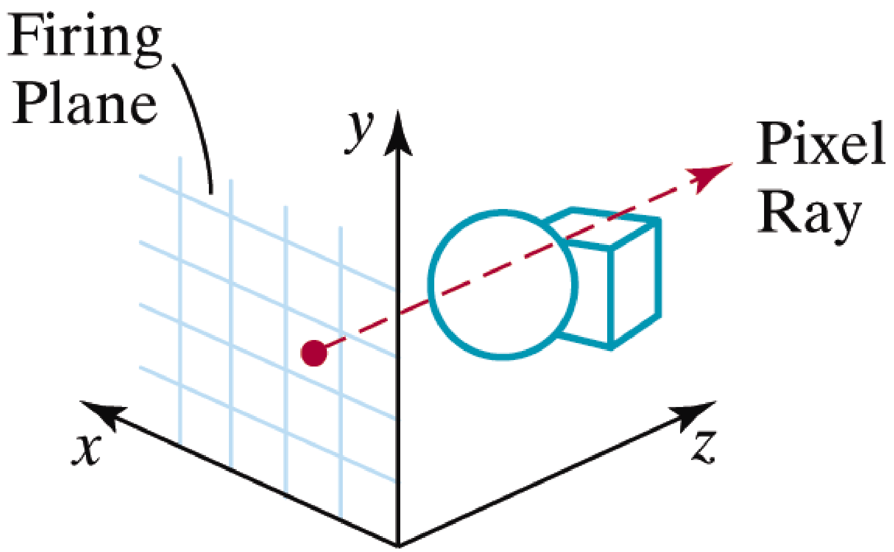

Ray-Casting

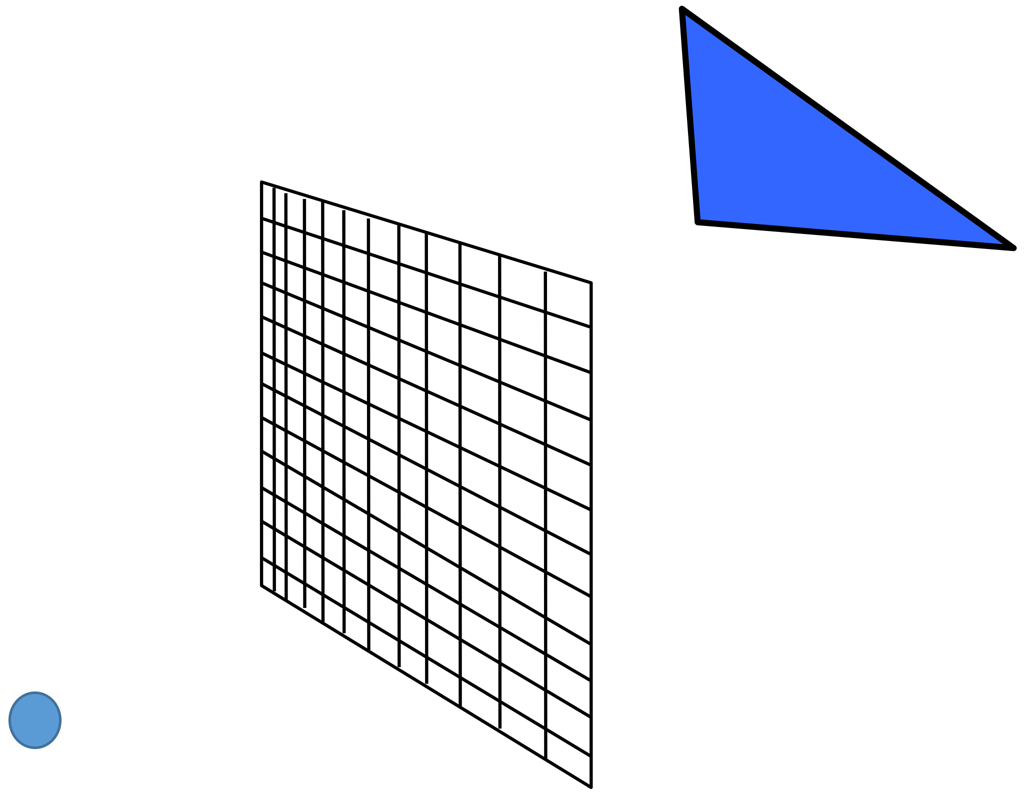

Ray-Casting

- Find out which rays will hit the edges of the triangle
- these rays are defined by a 2d triangle on the image layer
- Scan this 2D triangle

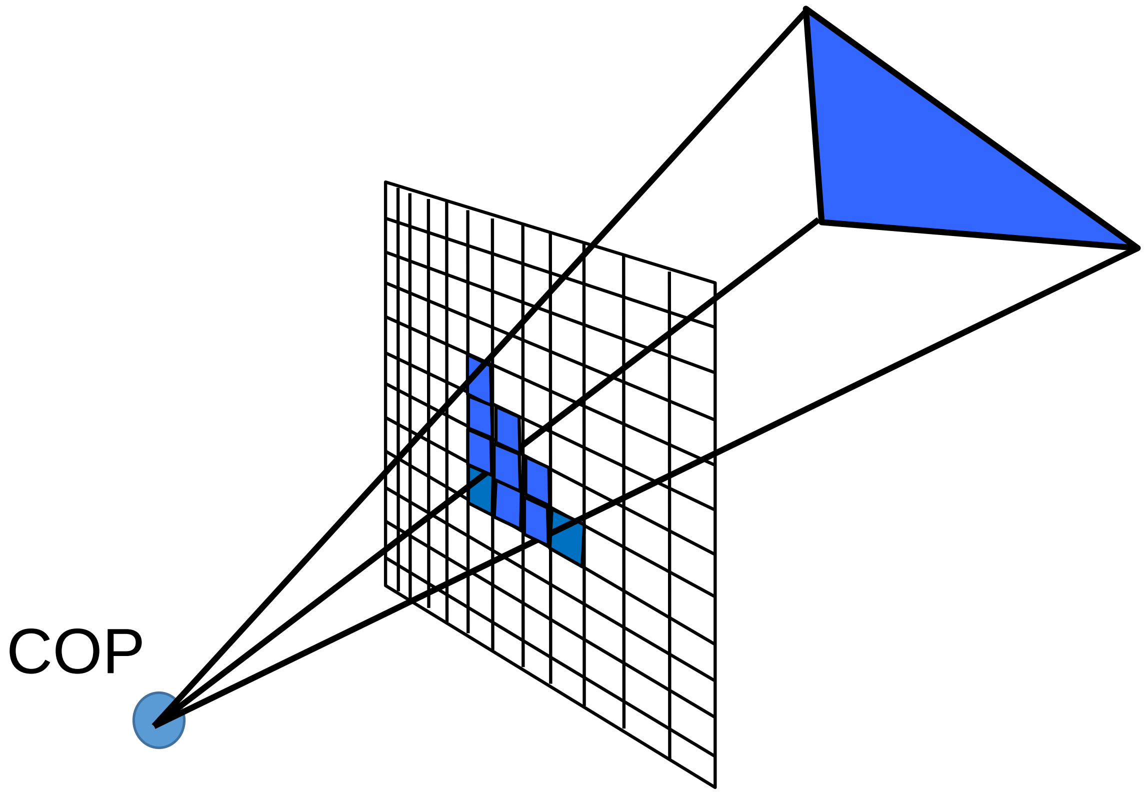

- **Instead of detecting the vertices, view them**

Rasterization

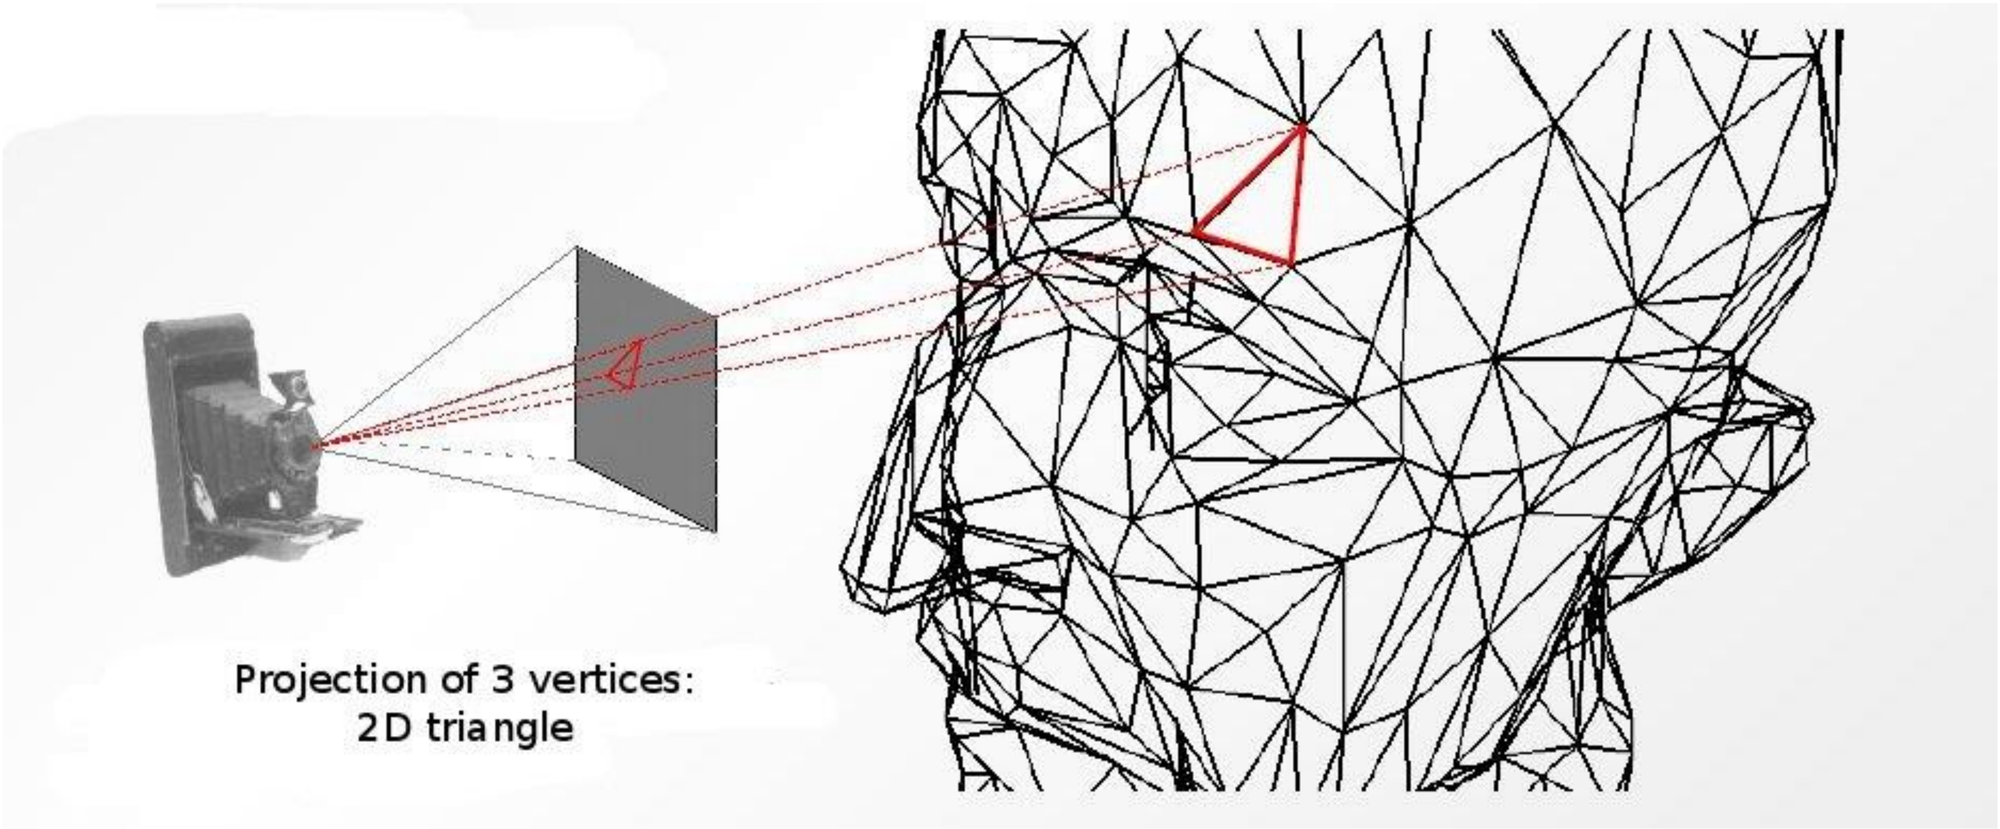

Projecting the vertices

- Start at the vertices and view it to find the pixel they cover
-

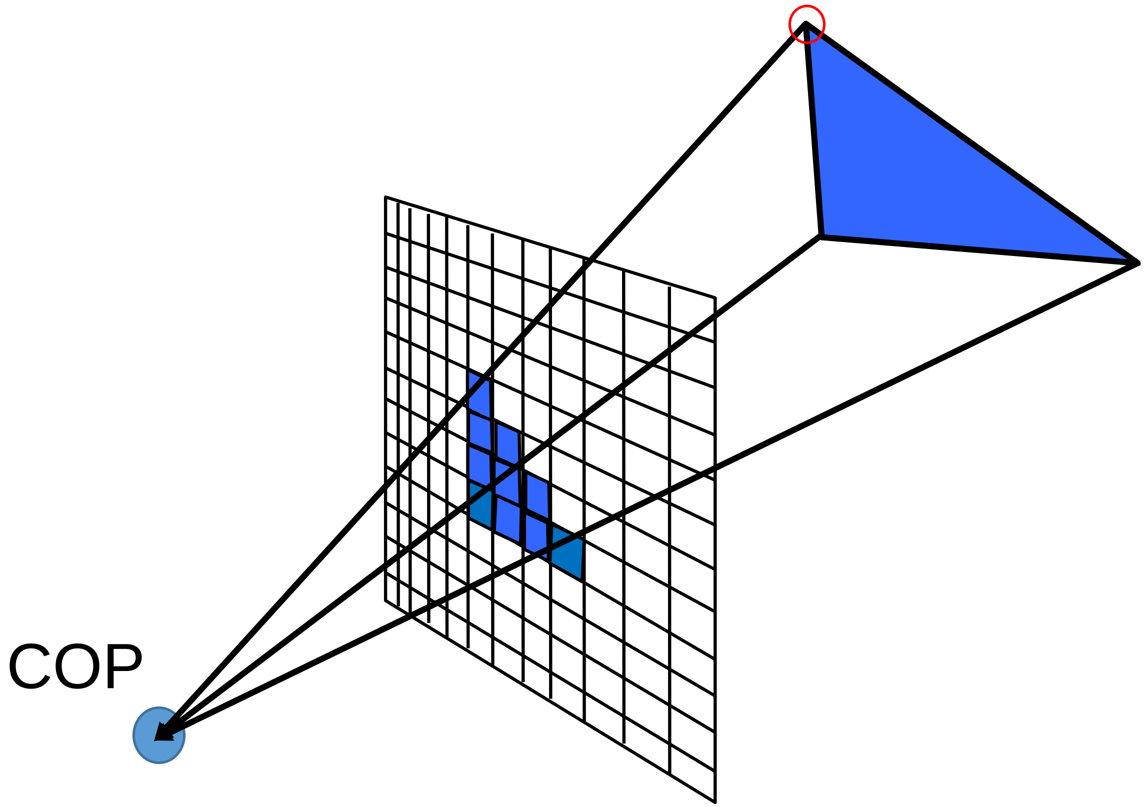

Ray Tracing

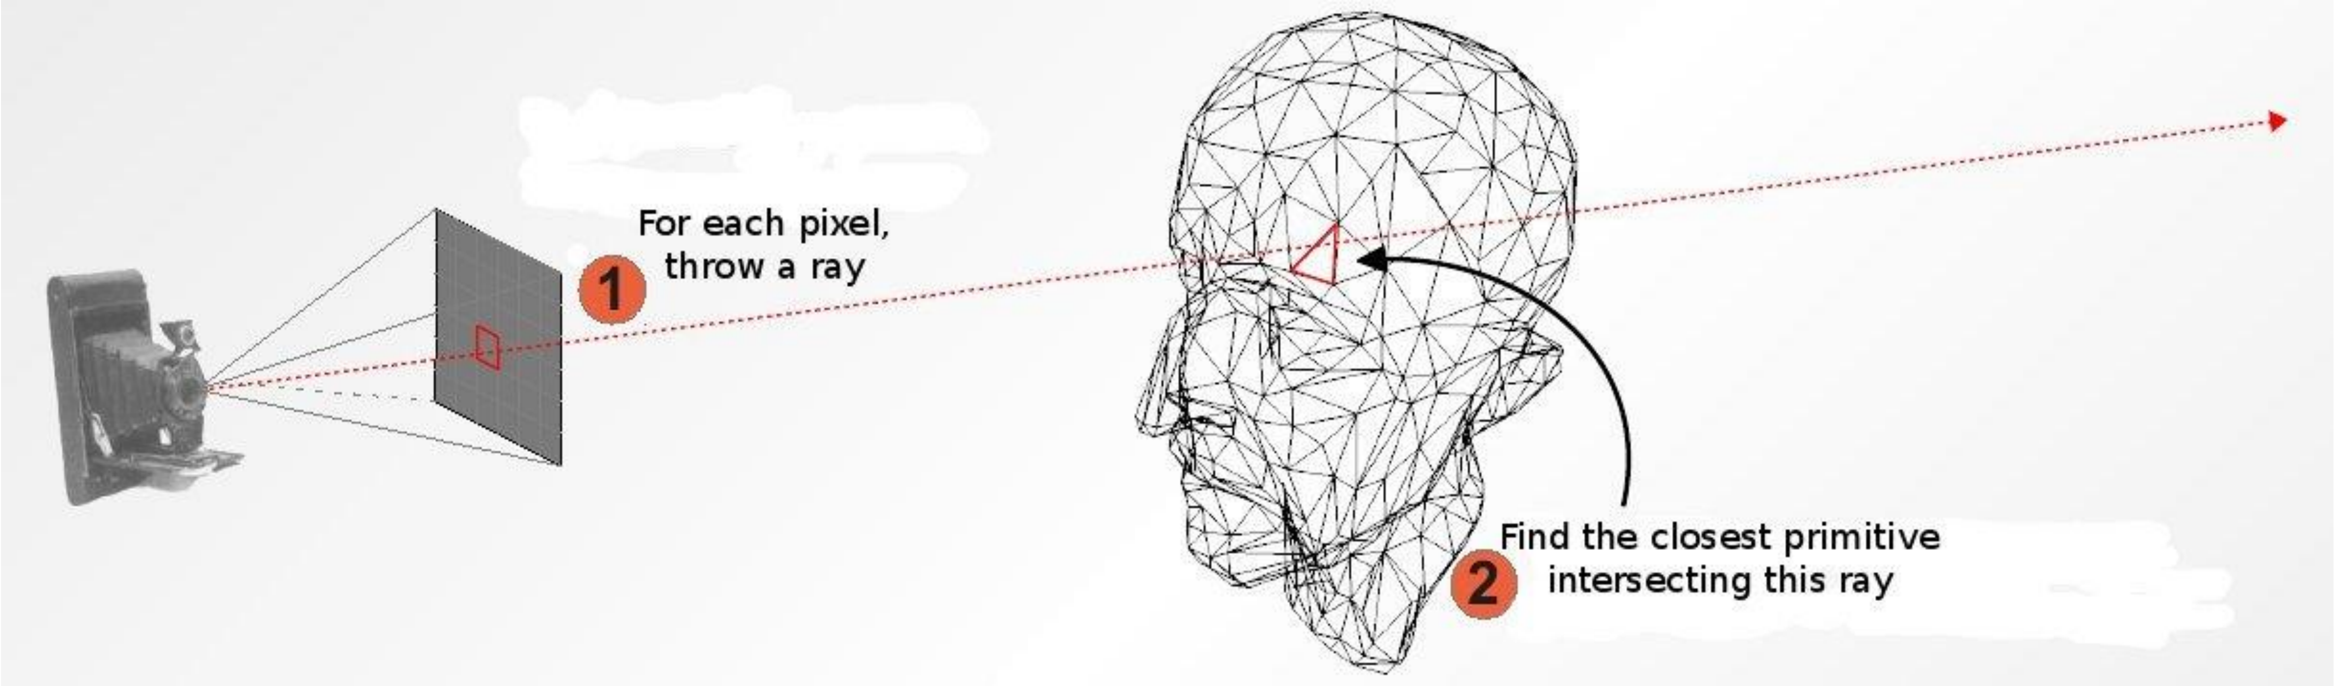

For each pixel, throw a ray
1
Find the closest primitive intersecting this ray
2

Μπορούμε να επιταχύνουμε απλοποιώντας

- Ray tracing is a very slow process
- E.g. for a cow we have 10,000 triangles, but to create an image of 1000 x 1000 pixels we need to shoot 1,000,000 primary rays, each one having to intersect the 10,000 triangles (and then add reflections, shadows, etc.)
- Reduce global lighting (i.e. no recursion)
  - Reduction of lighting, left only lighting from the environment
  - We believe that there are only polygons
  - Instead of detecting rays in each pixel, we simply locate them on the vertices and fill the interval in between
    - **More in upcoming lectures**