EECS 487: Interactive Computer Graphics

Lecture 22:
- Light and Color
- Color Spaces

Radiometry: Measuring Light

Radiant power or radiant flux or just flux: $\Phi = dQ/dt$
energy per time unit, in joules per sec ($=$ watt (W))

Spectral Energy Density/
Spectral Power Distribution
(or just Energy):
- radiant power per unit spectrum interval
- amount of light present at each wavelength
- $\mathrm{W/nm}$

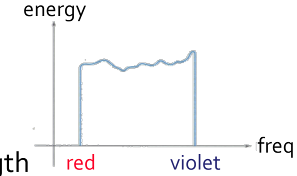

Intensity (I) is radiant flux per unit solid angle ($\mathrm{W/sr}$)
- solid angle: a $3\mathrm{D}$ angle, in steradians (sr)
  4π sr covers the whole area of a unit sphere

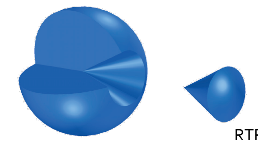

RTR3

Photometry

Photometry: quantifying the sensitivity of the average human eye in perceiving the energy of various wavelengths (not colors, which can be affected by many other factors)

- human perception is a non-linear function
- deals only with the visible spectrum, wavelength: 380 to 780 nm

Radiometry	Photometry	measures
radiant energy (in joules)	luminous energy (in talbots)	energy
radiant flux (watt)	lumen (lm)	power (energy/time)
radiant intensity (I, $\mathrm{W/sr}$)	candela (cd, lm/sr)	power/solid angle
radiant flux density (watt/m^2)	lux (lumens/m^2) (foot-candle, fc)	power/area
radiance	luminance (candela/m^2, nit)	power/(area*solid angle)
irradiance	illuminance	integrated radiance

Purpose of Computer Graphics?

Communication

Human perception is the context
- CG techniques leverage visual perception abilities

Fidelity is a tool, not (necessarily) the goal
- no apology is required for "approximations"
  - especially for interactive graphics
  - but not for scientific visualization and engineering drawing!

Three levels of realism:

1. Physical realism:
   - fidelity to visual stimulation of scene
   - highly computationally expensive
2. Photo realism:
   - fidelity to visual response to scene
   - takes observer's visual system into account
3. Functional realism:
   - fidelity to visual information of scene
   - task-oriented info presentation
   - takes advantage of observer's visual attention to task

Akeley07
Ferweda03

Two Types of Light-Sensitive Cell

Rods:
- highly sensitive, $1000\times$ more sensitive than cones
- spread over the retina, but predominate in peripheral vision
- effective in low light, night-vision
- monochrome vision: brightness perception only, no color

Cones:
- three types, each sensitive to a different frequency distribution
- concentrated in the fovea (center of the retina)
- less sensitive
- effective in bright light
- give color vision

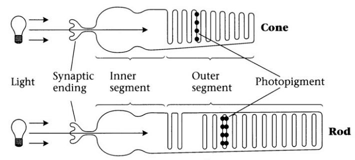

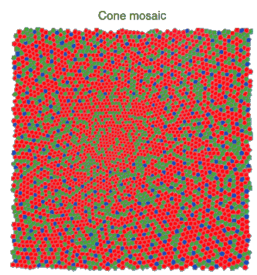

Zhang08

Light Absorbance

The rods and each of the three types of cone absorb light differently

Given a monochromatic spectrum (λ), we denote the response (light absorbance) of the three cones to λ as $s(\lambda)$, $m(\lambda)$, and $l(\lambda)$

The response curves $S = s(\lambda)$, $M = m(\lambda)$, and $L = l(\lambda)$ were experimentally determined in the 1980s

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.3816 & 0.5785 & 0.0399 \\ 0.1969 & 0.7246 & 0.0785 \\ 0.0248 & 0.1248 & 0.8504 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

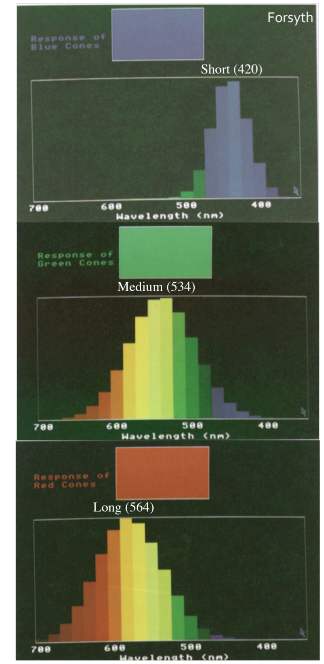

Cones and Rod Sensitivity

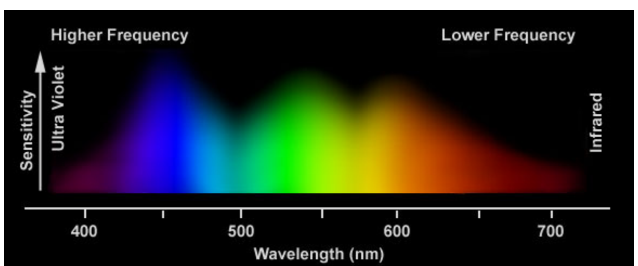

Rods and cones can be thought of as filters
- rods detect average intensity across spectrum

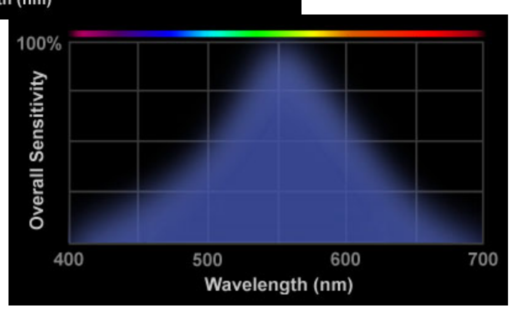

Tristimulus Response Functions

The eye's response to a unit amount of light with wavelength λ_1, $r(\lambda_1)$, is the linear combination $s(\lambda_1) + m(\lambda_1) + l(\lambda_1)$

$r(\lambda_2) = s(\lambda_2) + m(\lambda_2) + l(\lambda_2)$

$r(\lambda_1 + \lambda_2) = (s(\lambda_1) + s(\lambda_2)) + (m(\lambda_1) + m(\lambda_2)) + (l(\lambda_1) + l(\lambda_2))$

Let an arbitrary spectrum be the sum of its constituent "monochromatic spectra," each with radiance $L_i(\lambda_i) \Rightarrow \sum_i L_i(\lambda_i)$

Assuming linearity, the response of a cone (e.g., s) to an arbitrary spectrum is given by a convolution integral: $r(L, s) = \int s(\lambda) L(\lambda)\, d\lambda$

And the eye's response to an arbitrary light is a linear combination of the responses of the three cones:
$$r(L) = \int (s(\lambda) + m(\lambda) + l(\lambda)) L(\lambda)\, d\lambda = \int r(\lambda)\, L(\lambda)\, d\lambda$$

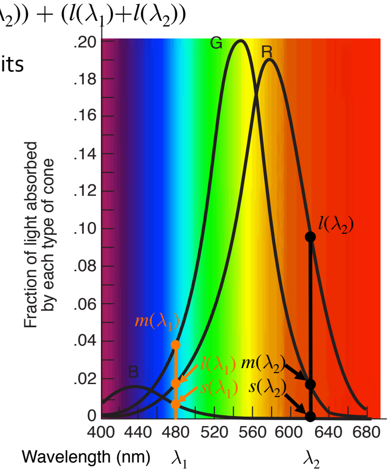

Luminous Efficiency of the Eye

The eye's total response to a multitude of wavelengths, all having the same intensity:

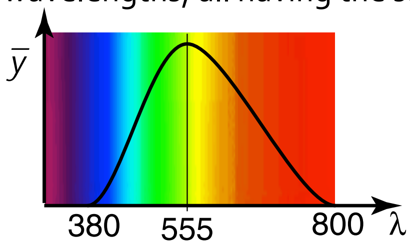

a.k.a. the CIE photopic spectral luminous efficiency curve, centered around 555 nm

Luminance (Y or V):
- is the overall magnitude of visual response to a spectrum (independent of color, $\approx$ "brightness/intensity")
- from radiance to luminance: $Y = 683$ lm/W $\int \bar{y} L(\lambda) d\lambda$
  - 1 Watt of radiant energy at $\lambda = 555$ nm equals to 683 lumens

Luminous Efficiency of the Eye

photopic: lighting conditions brighter than 3.4 cd/m^2 (twilight or brighter, cones effective)

scotopic: moonless night or darker, below 0.034 cd/m^2, rods effective (response curve centered around 507 nm)

- spectral sensitivity changes to bluish monochrome

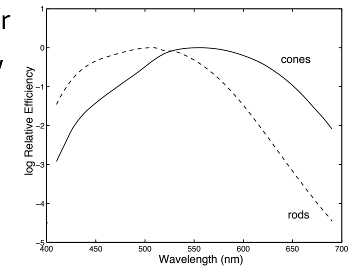

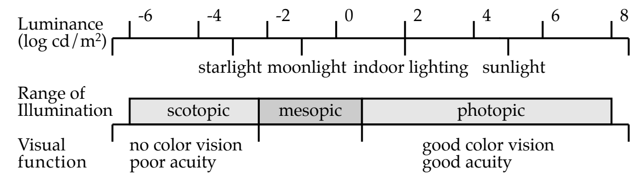

Colorimetry

Science of color measurement

Color of an object determine by reflected (i.e., not absorbed) wavelengths
- white = all wavelengths or frequencies
- blue is 450 nm, green 540 nm, red 650 nm

Color is not intrinsic in wavelength, but related to how a collection of photons with a spectral distribution is interpreted by our eyes

> "Light is not a thing that can be reproduced, but something that has to be represented with something else, with colors."
> – *Paul Cézanne*

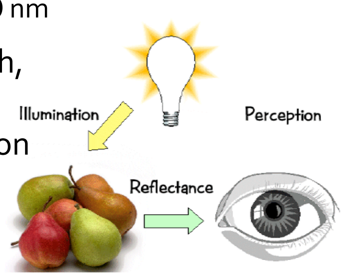

What the Eyes See

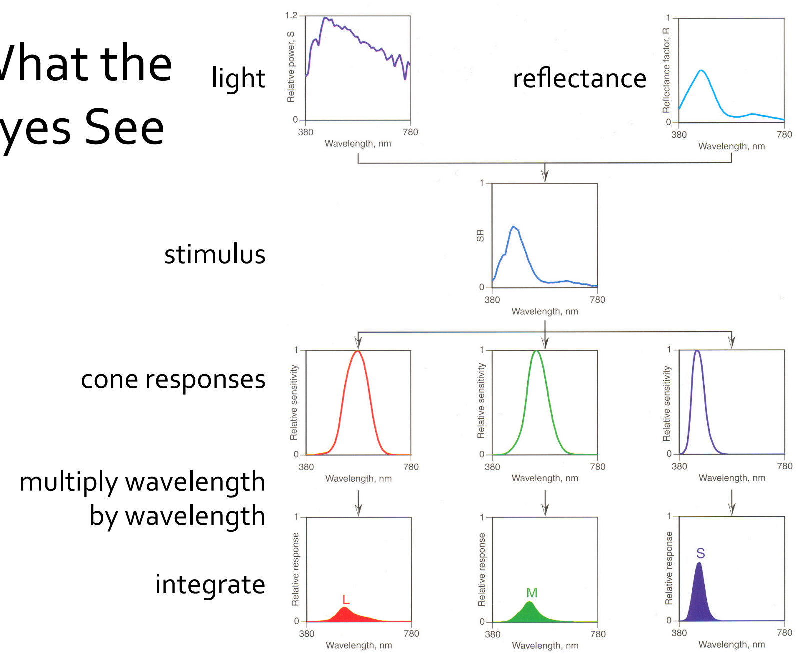

Durando6

Color Blindness

Classical case: 1 type of cone is missing (e.g., red)

Now project onto a lower-dimension space (2D)

Makes it impossible to distinguish some spectra

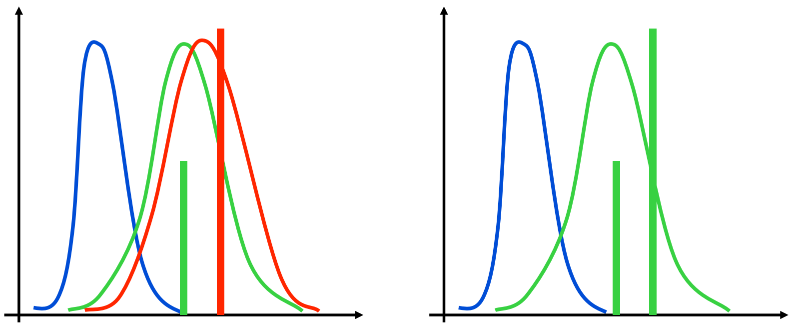

differentiated   same responses

Color Dimensions

Cones do not "see" colors

- they just respond to intensities of different wavelengths

A physical spectrum is a complex function of wavelengths

An arbitrary spectrum is infinitely dimensional

- but our eye-response has only three dimensions

How can we encode an infinite dimensional spectrum with just three dimensions?

- we can't: information is lost ⇒ we're all color blind!

Color Representation

Our ability to respond to only three primary colors is good news for computers

- only need three values to reproduce the full color spectrum visible to humans!
- imagine how much more expensive a display would be if each pixel must encode 6 values instead of 3, how much more storage, bandwidth, etc.

von Helmholtz 1859

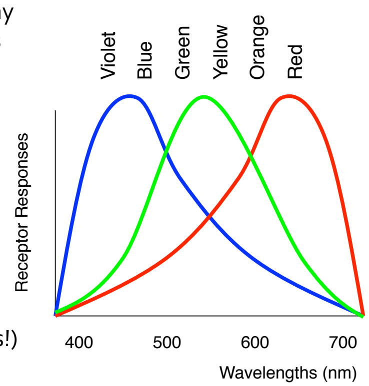

Trichromatic Theory:

- claims that any color can be represented as a weighted sum of three primary colors
- proposed red, green, blue as primaries
- developed in the 18th, 19th century (before the discovery of photoreceptor cells!)

Metamers

Metamers: when cones give the same response to different spectra

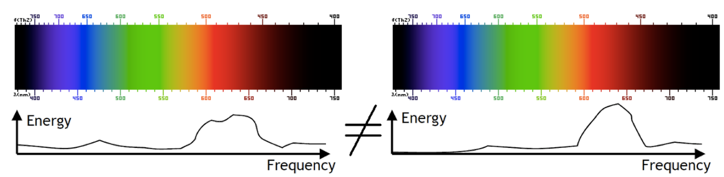

Perceived color: red = Perceived color: red

Color Matching

Metamers allows for color matching experiment

- reproduce the color of any reference light with the addition of three given primary lights
- how strong must each of the three primaries be for a match? (known as tristimulus weight or value)
- if no match can be obtained, add a primary to reference (negative weight)

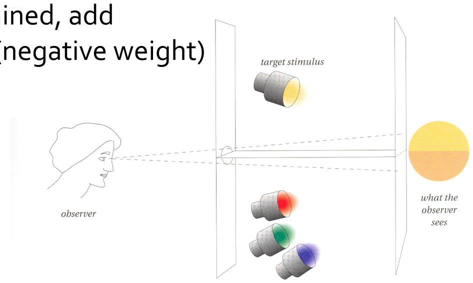

Color Matching Experiment

Result: RGB amounts needed to match all wavelengths of the visible spectrum:

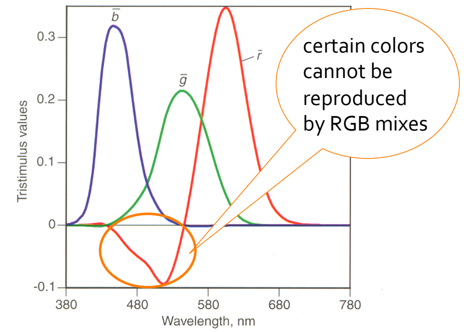

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

The tristimulus values were determined experimentally c. 1920 by the CIE (Commission Internationale de l'Eclairage (Illumination))

Durando6

Grassman's Laws

For color matches, where $\mathbf{u}$, $\mathbf{v}$, $\mathbf{w}$ are colors

- symmetry: $\mathbf{u} = \mathbf{v} \Leftrightarrow \mathbf{v} = \mathbf{u}$
- transitivity: $\mathbf{u} = \mathbf{v}$ and $\mathbf{v} = \mathbf{w} \Rightarrow \mathbf{u} = \mathbf{w}$
- proportionality: $\mathbf{u} = \mathbf{v} \Leftrightarrow t\mathbf{u} = t\mathbf{v}$
- additivity: if any two of the statements: $\mathbf{u} = \mathbf{v}$, $\mathbf{w} = \mathbf{x}$, $(\mathbf{u}+\mathbf{w}) = (\mathbf{v}+\mathbf{x})$ are true, so is the third
- these are natural laws and they mean additive color matching is linear

Forsyth&Ponce

Color Spaces

Color vision is linear and 3D, any color space based on color matching can be described by a "coordinate system" with 3 basis

Lots of different color spaces — related by matrix transforms!

- full spectrum
  - allows any radiation (visible or invisible) to be described
  - usually unnecessary and impractical
- RGB
  - convenient for display (CRT uses red, green, and blue phosphors)
  - not very intuitive
- HSV
  - an intuitive color space
  - Hue is cyclic so HSV is a non-linear transformation of RGB
- CIE XYZ
  - a linear transform of RGB used by color scientists
- Why not use the L, M, S cone responses as the basis?
  - not discovered until the 1980s …

Additive Primaries: RGB

http://www.jgiesen.de/ColorTheory/RGBColorApplet/rgbcolorapplet.html

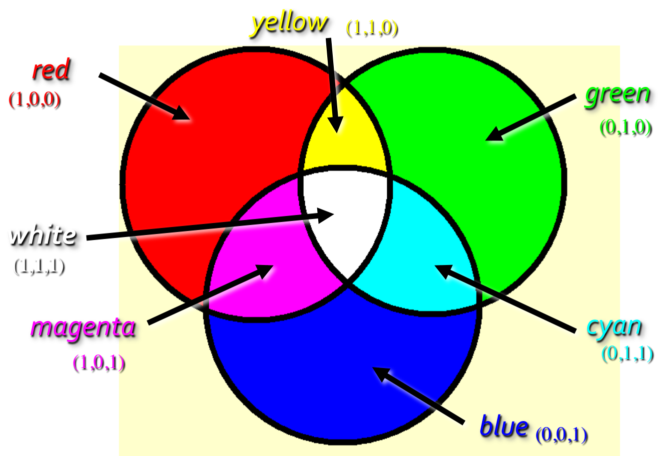

Riesenfeld 2006

Subtractive Primaries: CMY

http://www.jgiesen.de/ColorTheory/CMYColorApplet/cmycolorapplet.html

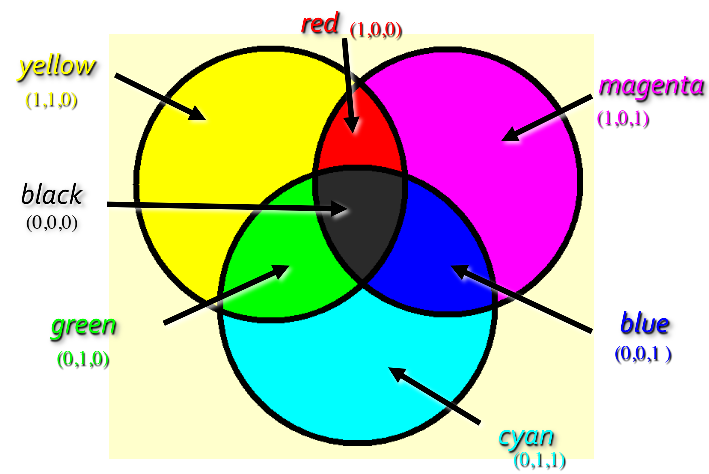

Riesenfeld 2006

CMY(K) Color Model

Typical inks: Cyan, Magenta, Yellow, (blacK)

- the pigments remove parts of the spectrum

pigment	absorbs	reflects
cyan	red	blue and green
magenta	green	blue and red
yellow	blue	red and green
black	all	none

black ink used to ensure high quality printed black

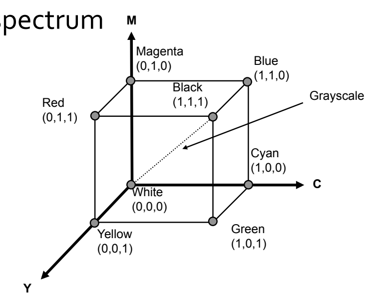

- complements of RGB

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

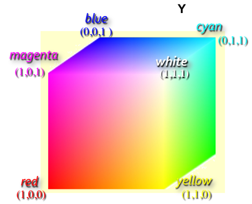

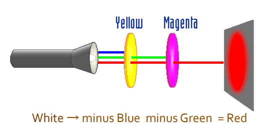

Hodgins

Problems with RGB

Non-intuitive: how much R, G, and B is there in "brown"? (answer: .64, .16, .16)

Represent only a small range of human perceptible colors (not perceptually based)

Perceptually non-linear

- the perceived difference between two colors is not consistently proportional to their separation in the color space

Chenney

HSV Color Model

What do we perceive?

- Hue (or chromaticity): what color is it?
  - red vs. green vs. blue
- Saturation (or purity or chroma): how non-gray is it?
  - vivid red vs. pastel grayish pink
- Value (or luminance or intensity): how bright is it?
  - perceived intensity reflected by object
  - the term "brightness" is used only for emitters e.g., light-bulbs

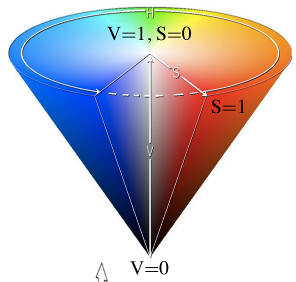

HLS: represented as a double cone with black at the bottom and white at the top

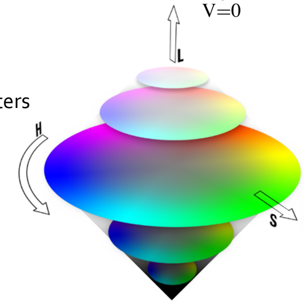

Rossignac,Hodgins,Gillies,wikipedia

HSV: Non-linear Distortion of the RGB Color Cube

Top of HSV hexcone corresponds to the projection of the RGB color cube along the principal diagonal

- the main diagonal of RGB space becomes the **V** axis of the HSV space

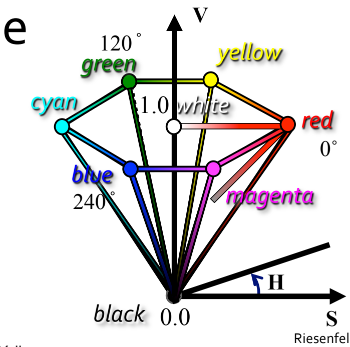

The RGB cube has subcubes

Each plane of constant **V** in HSV space corresponds to a view of a subcube of RGB space

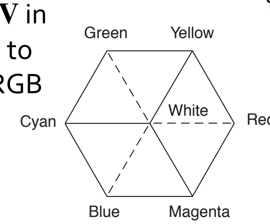

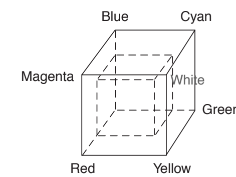

Foley, van Dam 90

Artist's Color Model

Artists discuss color (hue) in terms of:

- Tint: strength of color
  - the amount of white added to pure pigment to decrease saturation
- Shade: brightness of color
  - the amount of black added to decrease lightness (value)
- Tone:
  - the amount of black and white added to a pure pigment (given hue)

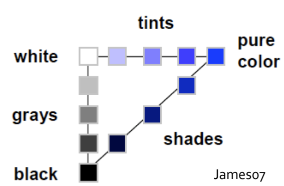

Color Spaces

Traditional, Artistic:
RGB
CMY(K)
HSV
HLS

Any set of real lights would need negative weighting to cover the whole visible light spectrum

Perceptually Based:
XYZ (Tristimulus)
xyY
Hunter-Lab
CIE-L*u'v'
CIE-L*a*b*
CIE-L*CH°

Riesenfeld 2006

CIE XYZ Color Space

Standardized by the International Commission on Illumination (Commission Internationale de l'Éclairage)

- defines three "imaginary" primary lights ($\mathbf{X}, \mathbf{Y}, \mathbf{Z}$) representing an imaginary basis that does not correspond to human perception of color
  - color $\mathbf{c} = X\,\mathbf{X} + Y\,\mathbf{Y} + Z\,\mathbf{Z}$
- all visible colors can be matched with positive X, Y, and Z $\Rightarrow$ they are in the positive octant of the $\mathbf{XYZ}$ space
- provides a standard for sharing color information between disciplines, e.g., computer graphics and fabric design

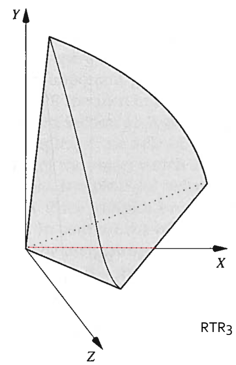

RTR3

Recall: Color Matching Experiment

Certain visible colors cannot be reproduced by RGB mixes

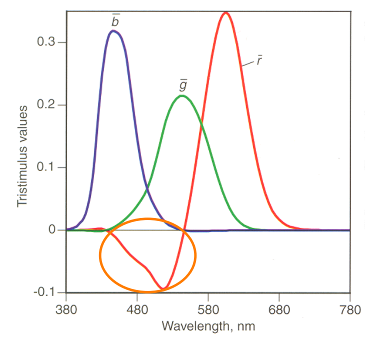

These curves are the color-matching functions for the 1931 standard observer, The average results of 17 color-normal observers having matched each wavelength of the equal-energy spectrum with primaries of 435.8 nm, 546.1 nm, and 700 nm.

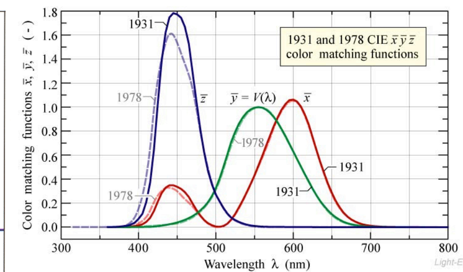

All visible colors can be matched with positive X, Y, and Z

Durando6

2D Visualization of CIE Color Space

Given a color $\mathbf{c} = X\,\mathbf{X} + Y\,\mathbf{Y} + Z\,\mathbf{Z}$

It's cumbersome to visualize it in 3D

Want a 2D representation:

First map it to a triangle in 3D:

$$x = \frac{X}{X+Y+Z};\quad y = \frac{Y}{X+Y+Z};\quad z = \frac{Z}{X+Y+Z}$$

$x + y + z = 1$ and $x, y, z \geq 0$,
i.e., x, y, and z are coefficients of a convex combination and the point is in the $(X + Y + Z = 1)$ plane

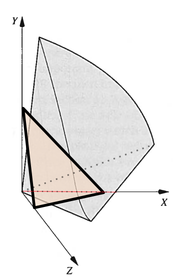

2D Visualization of CIE Color Space

Then map it to a 2D chromaticity diagram by dropping the z coordinate

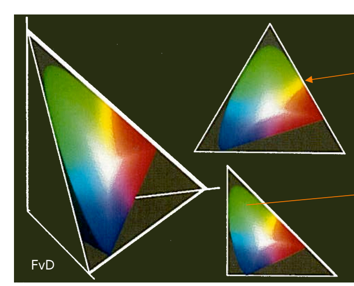

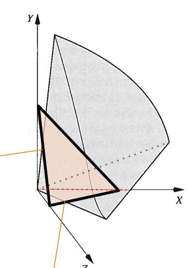

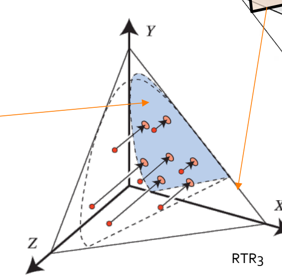

RTR3

CIE Chromaticity Diagram

The horseshoe region represents all visible chromaticity values

- all perceivable colors with the same chromaticity but different luminance map into the same point within this region
- shaded area: color found in nature

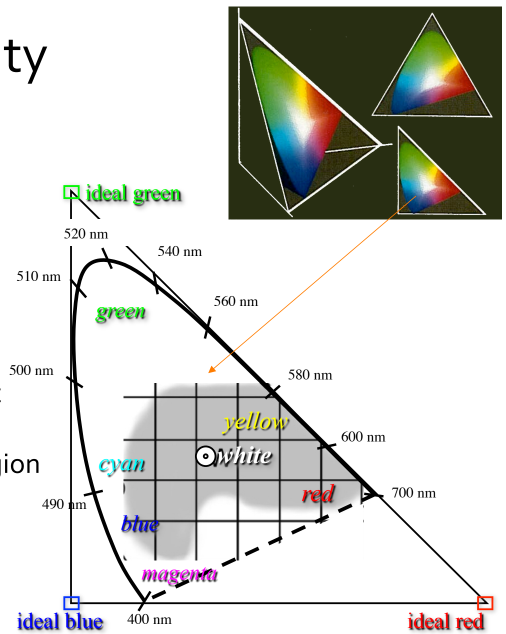

Riesenfeld,TP3

Color Gamut

The color gamut of a device is the convex hull defined by the convex combinations of its primary colors

- a device can only produce colors within its color gamut
  - different devices have different color gamuts
  - to match gamuts between devices can be difficult
- points outside a gamut correspond to negative weights of the primaries

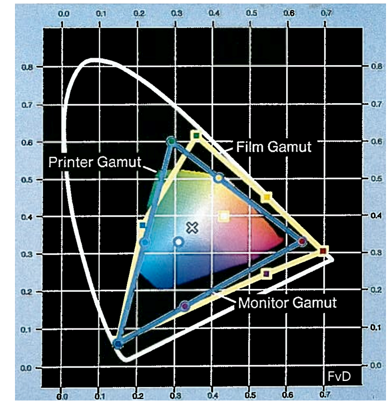

HSV ↔ CIE XYZ

Hue (chromaticity)

- the dominant wavelength/frequency, identified by drawing a line from $\mathbf{c}$ (white) to spectral color
  - $\mathbf{c}_6 \Rightarrow \mathbf{c}_s$
  - $\mathbf{c}_7$ is nonspectral, cannot be identified with a dominant wavelength, instead the dominant wavelength is the complement of $\mathbf{c}_c \Rightarrow \mathbf{c}_c^c$

Saturation (chroma)

- "purity" of the color, distance from $\mathbf{c}$ (white) to spectral color
  - $\mathbf{c}_4$ very pure
  - $\mathbf{c}_6$ not so pure

Value (brightness/lightness)

- perceived luminance (intensity)
  - lightness: reflecting objects
  - brightness: emitting objects
- not represented (not representable) in 2D

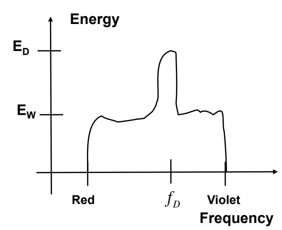

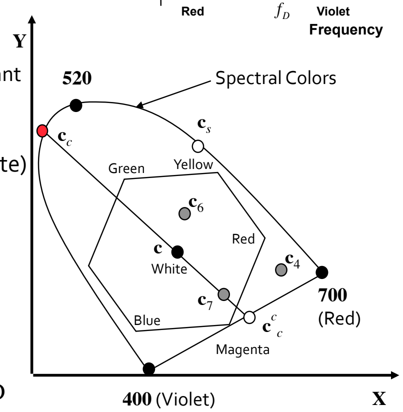

RGB ↔ CIE XYZ

RGB color model

- additive primaries $\mathbf{c}_\lambda = R\,\mathbf{R} + G\,\mathbf{G} + B\,\mathbf{B}$
- is contained within the CIE XYZ color space

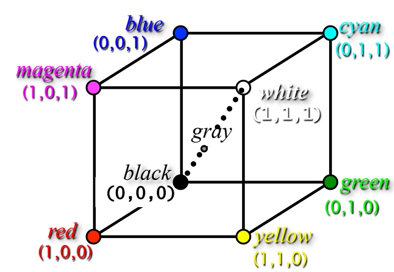

Riesenfeld 2006

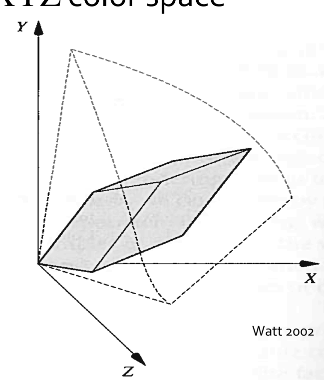

Watt 2002

RGB Color Gamut

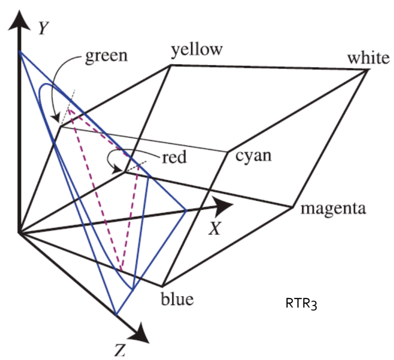

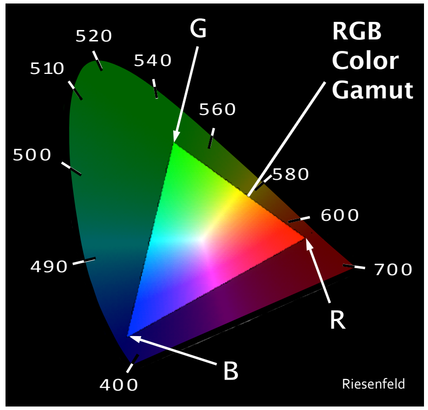

RGB chromaticity coordinates:
R = (0.735, 0.265)
G = (0.274, 0.717)
B = (0.167, 0.009)

sRGB ↔ CIE XYZ Transforms

sRGB: RGB with standard conversion to CIE XYZ

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7151 & 0.0721 \\ 0.0193 & 0.1192 & 0.9505 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

- each matrix is the inverse of the other
- Y encodes luminance; to go from color to gray: $Y = 0.213\,R + 0.712\,G + 0.072\,B$
- grayscale intensity (Y) is not equal parts RGB because to the eye, the intensity of red and green contribute more than that of blue to overall perceived brightness

Chenney

Pyschophysics

Humans are
- not so sensitivity to low frequencies
- more sensitive to medium to high frequencies

Separate intensity from the color of a pixel:
- reduce the contrast of low frequencies
- but keep the color

Luminous Efficiency Curve

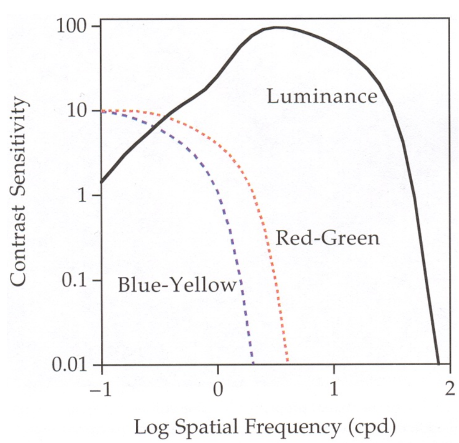

Freeman&Durand06

YIQ Color Model

For NTSC color TV

Separate luminance from chrominance
- we perceive brightness ranges better than color ranges, so give it more bandwidth (samples): Y: 4.5 MHz, I: 1.5 MHz, Q: 0.6 MHz
- backward compatibility: B&W TV displays only the Y component
- two colors on B&W with different luminance show up as different grades of gray

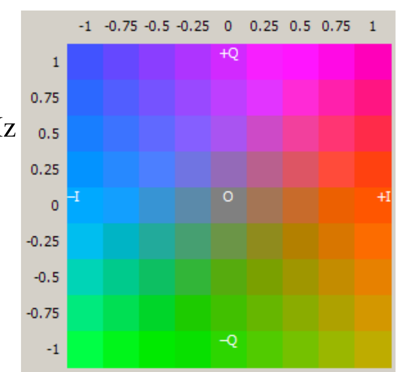

Merrello8

RGB ↔ YIQ

$$\begin{array}{r} \text{brightness} \\ \text{color} \left\{ \begin{array}{c} \\ \\ \end{array} \right. \end{array} \begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.144 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.528 & 0.311 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1 & 0.9563 & 0.6210 \\ 1 & -0.2721 & -0.6474 \\ 1 & -1.1070 & 1.7046 \end{bmatrix} \begin{bmatrix} Y \\ I \\ Q \end{bmatrix}$$

- for modern CRT and HDTV:
  $Y = 0.2125\ R + 0.7154\ G + 0.0721\ B$
  comparable to sRGB:
  $Y = 0.213\ R + 0.712\ G + 0.072\ B$

xyY Coordinate

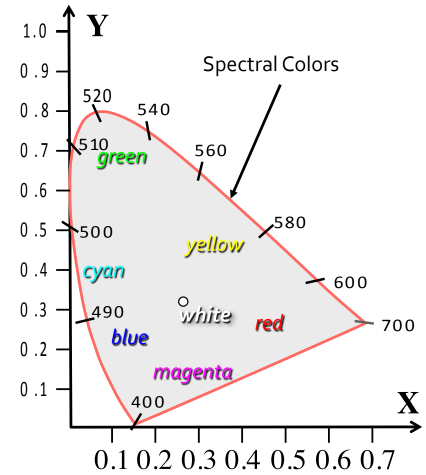

Riesenfeld 2006

The chromaticity diagram does not represent a complete color palette:

- doesn't account for color changes due to luminance
  - brown, which is orange-red with low luminance, is not represented
- to recover full palette add back luminance: (x,y,Y)

$$X = \frac{x}{y}Y,\ \ Y = Y,\ \ Z = \frac{(1-x-y)}{y}Y$$

Characteristics of Color Spaces

Color Space	Intuitive?	Perceptually based?	Perceptually linear?	Device independent?
HSV	✔	✗	✗	✗
RGB	✗	✗	✗	✗
CMY	✗	✗	✗	✗
CIE XYZ	✗	✔	✗	✔
CIE L*u′v′	✗	✔	✔	✔
CIE L*a*b*	✗	✔	✔	✔

TP3

Perceptual Non-linearity

MacAdam ellipses: regions in CIE xy color space that are perceived as the same color (based on a 1942 experiment)

- in perceptually uniform color space, the ellipses should be circles

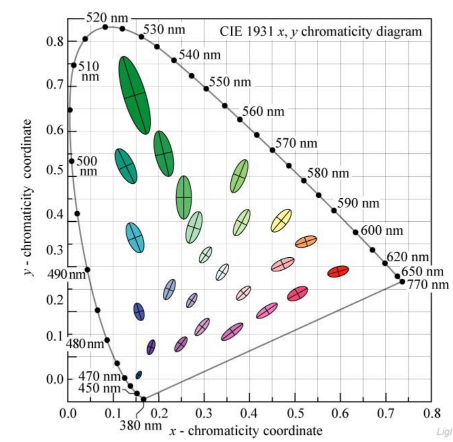

Schulze,Schubert

Perceptually Uniform Color Spaces

Two major color spaces standardized by the CIE

- designed so that equal steps in coordinates produce equally visible differences in color
- L*u'v': non-linear color space that gives perceptual linearity: MacAdam ellipses now look more like circles

$$\begin{bmatrix} u' \\ v' \end{bmatrix} = \frac{1}{X + 15Y + 3Z} \begin{bmatrix} 4X \\ 9Y \end{bmatrix}$$

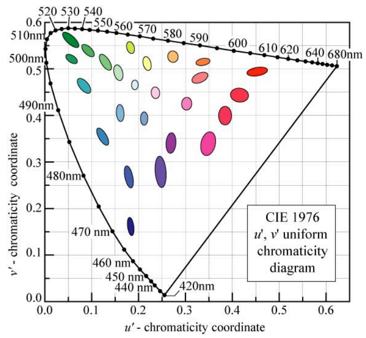

- L*a*b*: more complex but more uniform
- both separate luminance (L*) from chromaticity

James,Chenney,Schubert

CIE L*a*b*

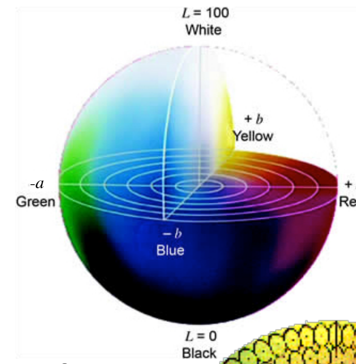

Perceptually uniform color space:

- L*: luminance
- a*: red-green
- b*: blue-yellow
- the '*'s are to differentiate it from Hunter-Lab color space . . .

Non-linear conversion between CIE XYZ and CIE L*a*b*

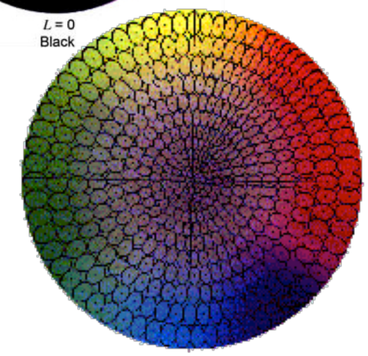

Opponent Color Theory

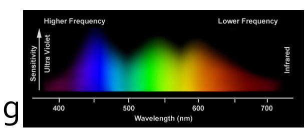

Since the cones respond to overlapping wavelengths, the visual system can more efficiently record the differences between total cone responses, instead of individual cone's responses

The 0^{th}, 1^{st}, 2^{nd} derivatives of a Gaussian weighting function of the wavelengths are similar to the luminance, blue-yellow, and red-green weighting functions found in the human visual system

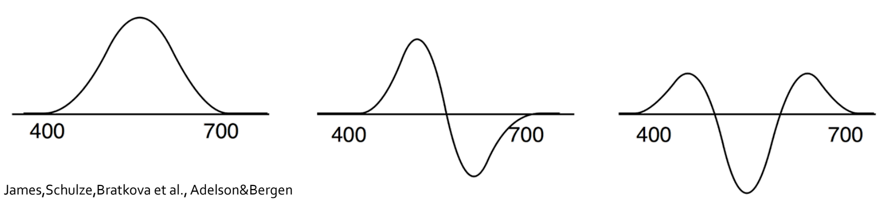

James,Schulze,Bratkova et al., Adelson&Bergen

Opponent Color Theory

The brain seems to encode color using three axes: white – black, red – green, yellow – blue

- the white – black axis determines luminance (V or Y), the others determine chromaticity
- first proposed in the 19^{th} century, physiological evidence in the 1950s

You can have light green, dark green, yellow-green (chartreuse), or a blue-green (teal), but not reddish green
⇒ red is the opponent to green

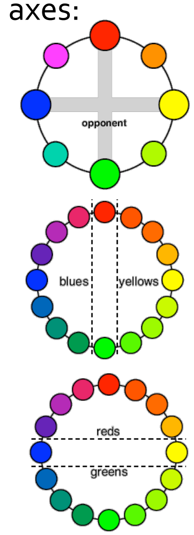

James,Schulze,Bratkova et al., Adelson&Bergen