Color II: applications in photography

CS 178, Spring 2009

Begun 5/14/09, finished 5/19/09.

Marc Levoy
Computer Science Department
Stanford University

Outline

- ✦ spectral power distributions
- ✦ color response in animals and humans
- ✦ 3D colorspace of the human visual system
- ✦ reproducing colors using three primaries
- ✦ additive versus subtractive color mixing

- ☞ ✦ cylindrical color systems used by artists (and Photoshop)
- ✦ chromaticity diagrams
  - color temperature and white balancing
  - standardized color spaces and gamut mapping
  - uniform perceptual spaces and opponent colors

Newton's color circle

(http://www.handprint.com/HP/WCL/color6.html)

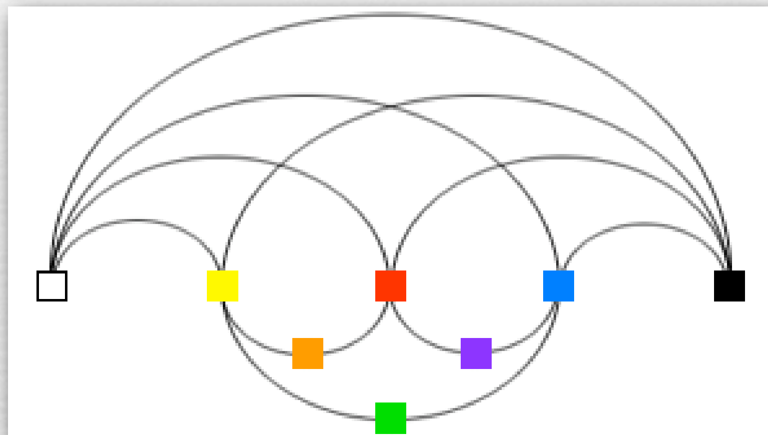

Peter Paul Rubens and
François d'Aguilon (1613)

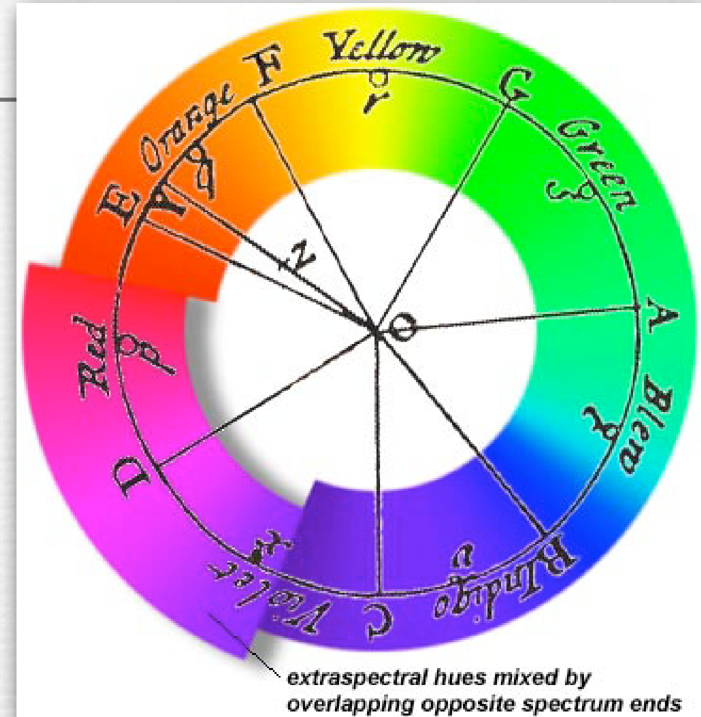

Isaac Newton (1708)

- previous authors could not move beyond linear scales, because they felt compelled to include black and white as endpoints
- Newton closed the circle by removing black and white, then adding extra-spectral purples not found in the rainbow
    - by mixing red at one end with violet at the other end

Scales for cylindrical color spaces

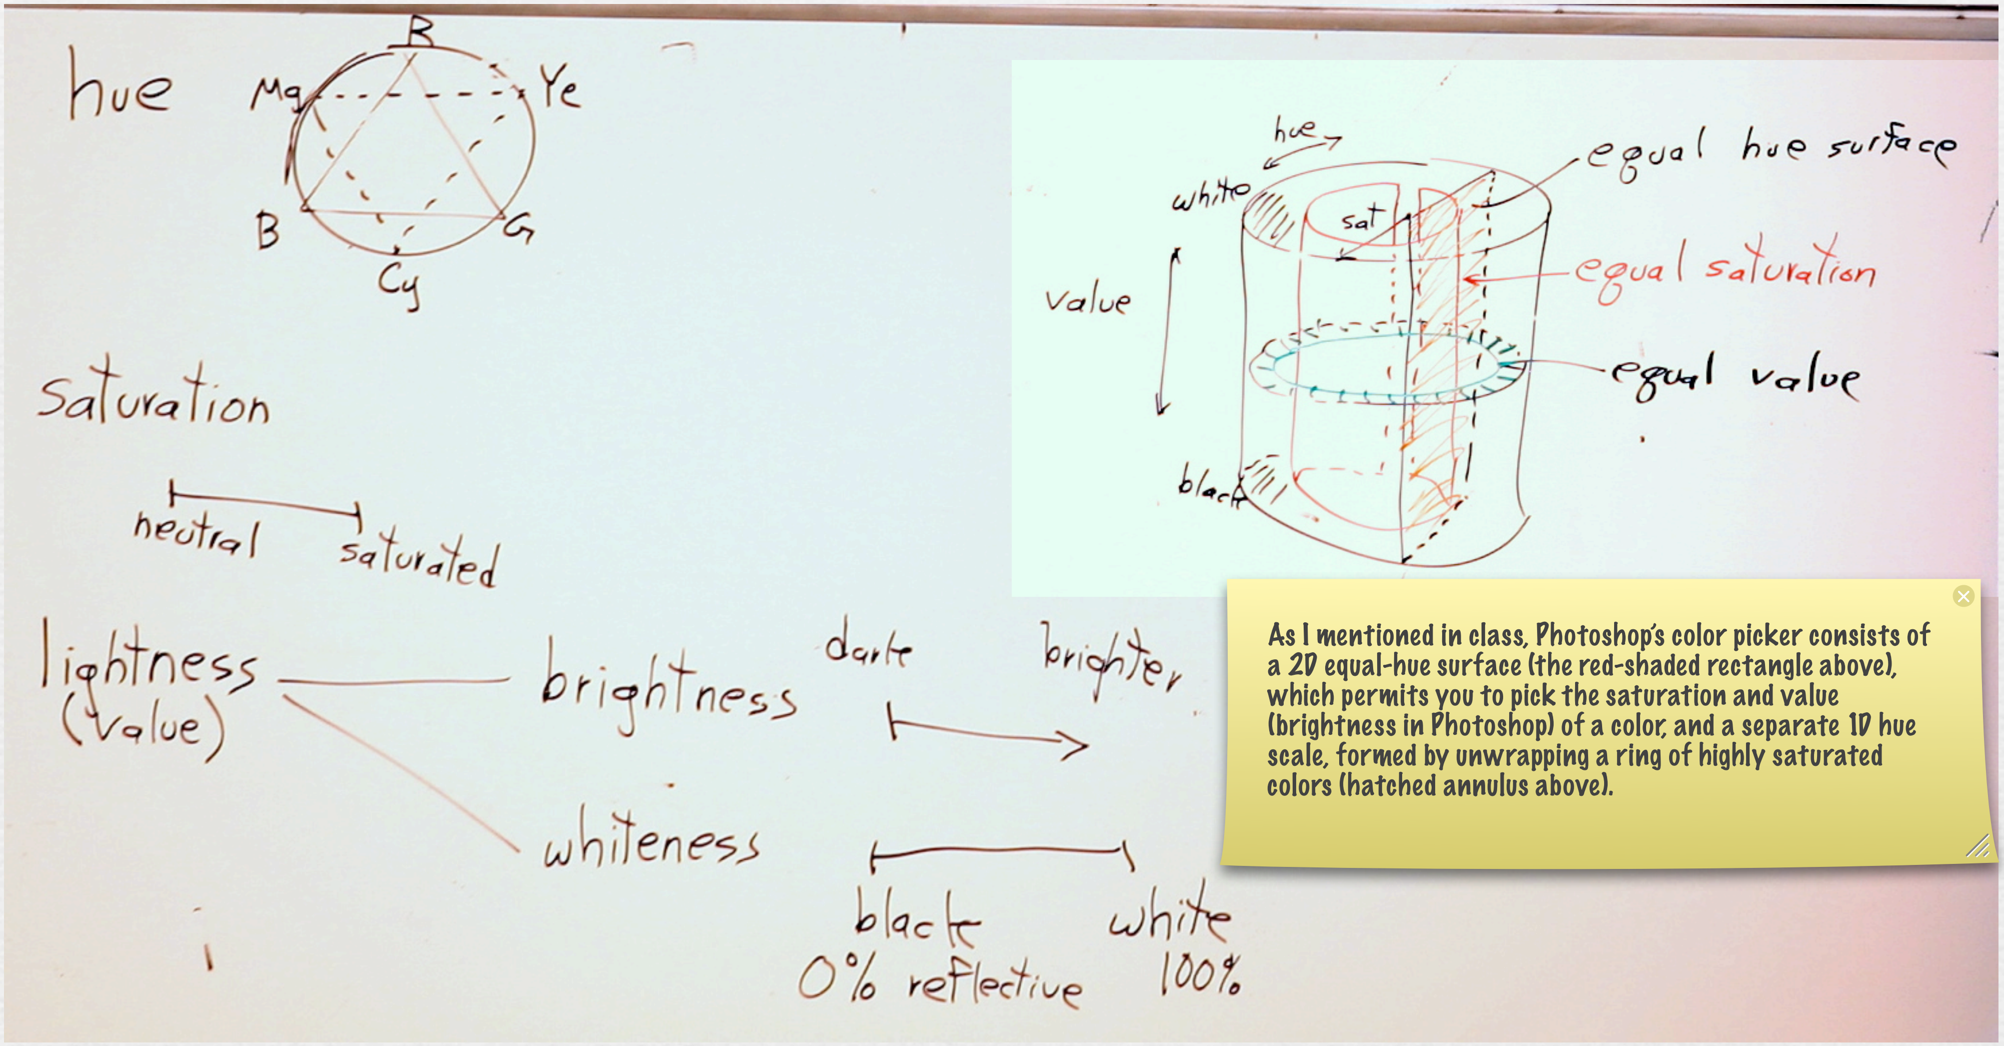

As I mentioned in class, Photoshop's color picker consists of a 2D equal-hue surface (the red-shaded rectangle above), which permits you to pick the saturation and value (brightness in Photoshop) of a color, and a separate 1D hue scale, formed by unwrapping a ring of highly saturated colors (hatched annulus above).

Cylindrical color spaces

As I defined cylindrical space on the previous slide, the top surface (a circular disk) should be all white, and the bottom surface should be all black. In the version at right below, grabbed from the Internet, the programmer has squeezed together the useless black values at the bottom of the cylinder. Imagine that he (or she) also squeezed together the useless white values at the top of the cylinder. The result would be two opposing cones, like a child's spinning toy. Now imagine that he shoved the top of the double cone downward until it became level with the midpoint (vertically). This would produce the figure at right below, where white is in the middle of a disk of highly saturated colors. If he reopened the squeezed-together bottom of the cone, but left the top shoved downward, he would produce the figure at left below. There are many other variants one could create - single or double cones, single or double hexcones, even spheres or ellipsoids. All are 3D shapes, and all are nonlinear distortions of the same basic cylindrical coordinate system.

(wikipedia)

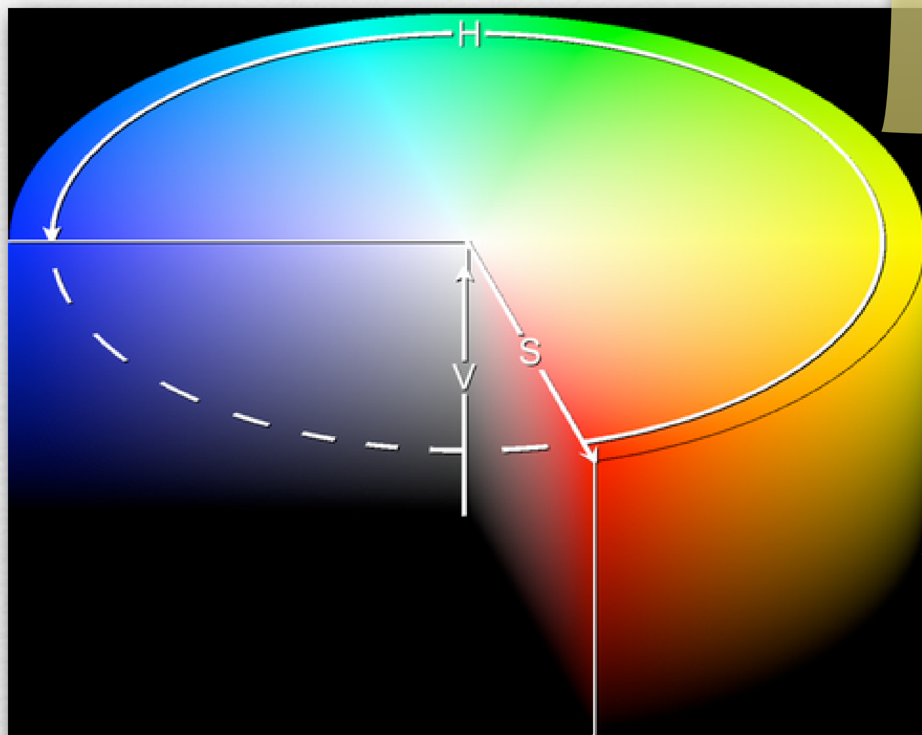

cylinder

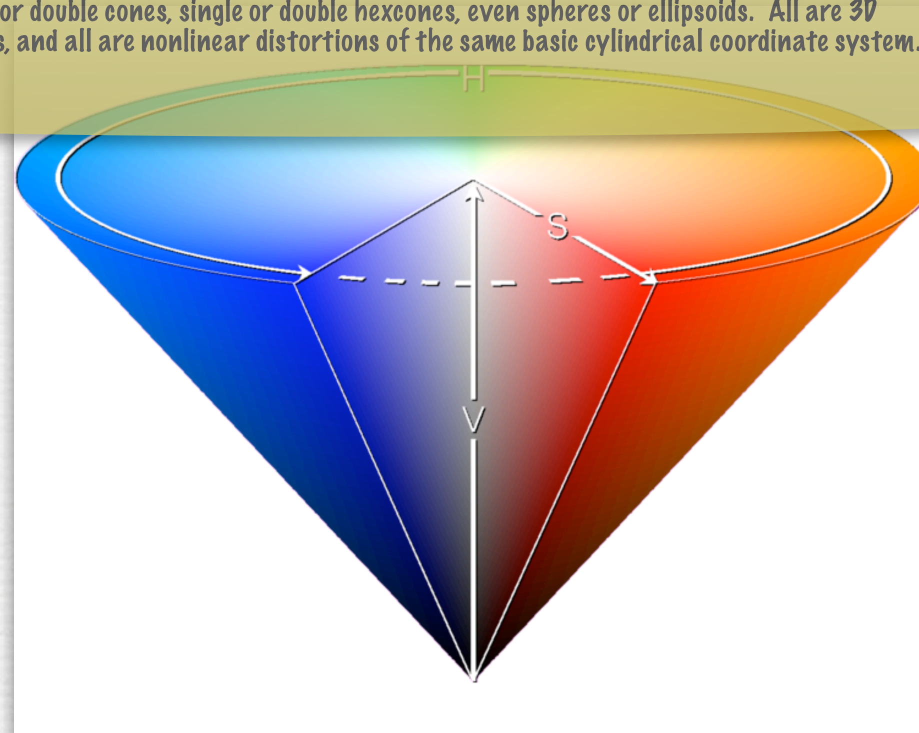

single cone

- a cylinder is easy to understand, but colors near the bottom are all dark, so they are barely distinguishable
  - the single cone solves this by compressing the bottom to a point

Color selection in Photoshop

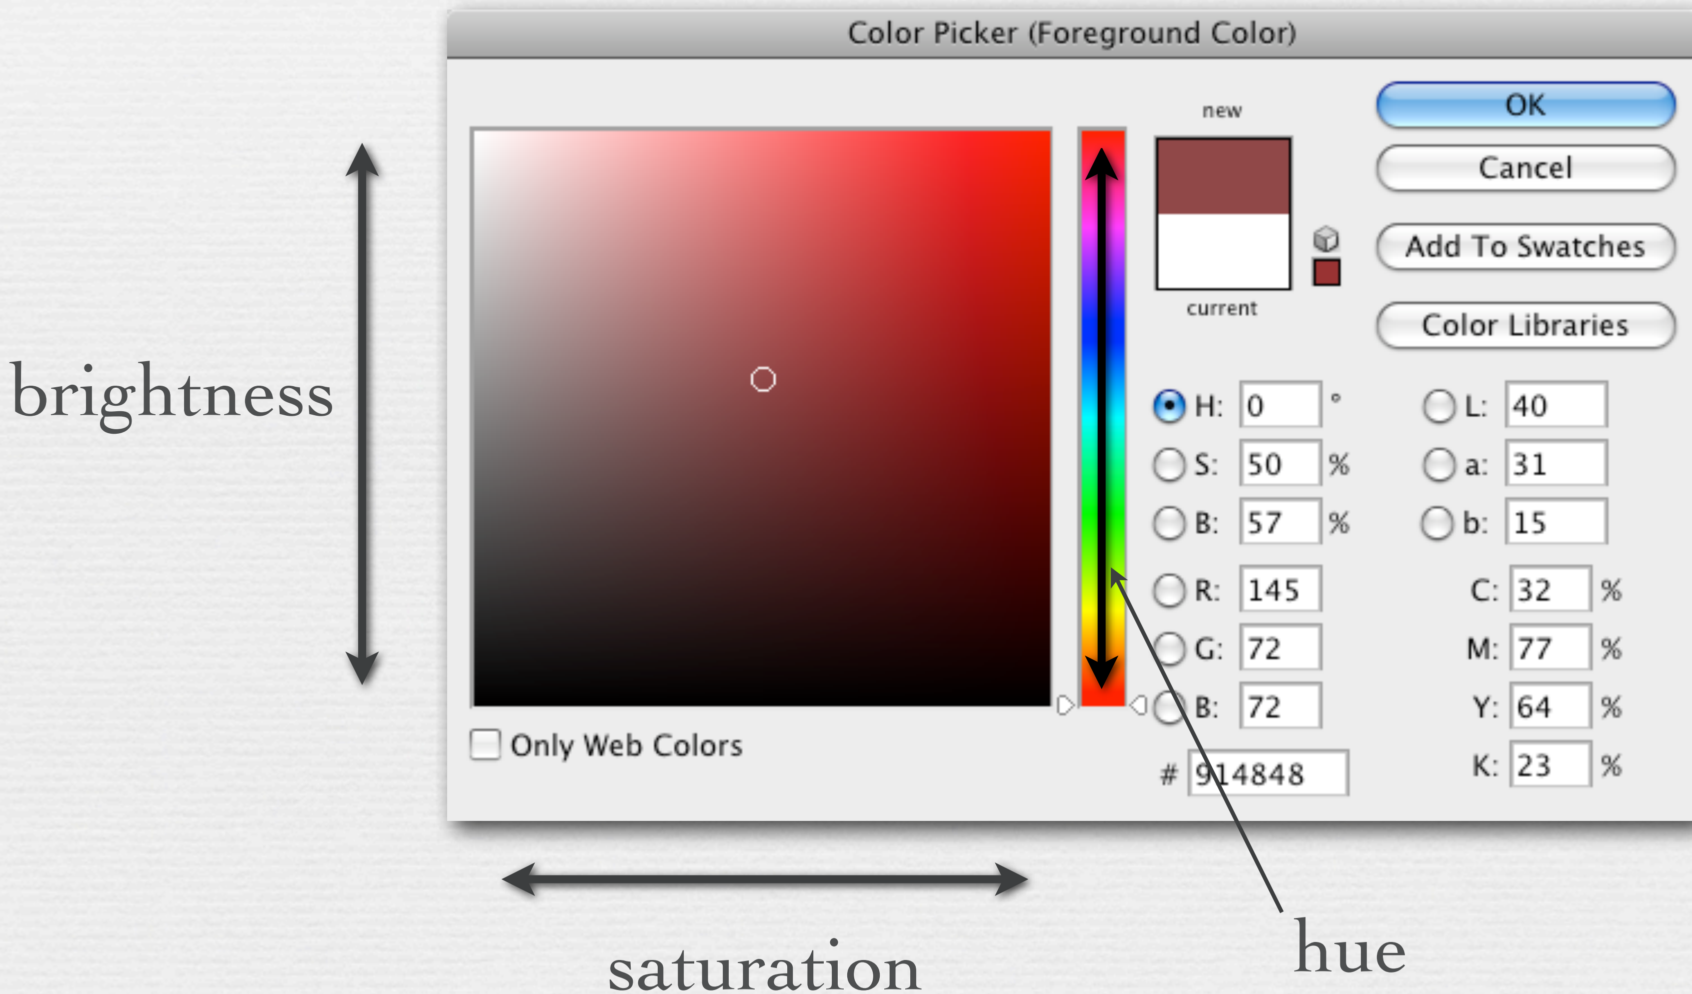

Color selection in Photoshop

Cartesian to cylindrical coordinate conversion

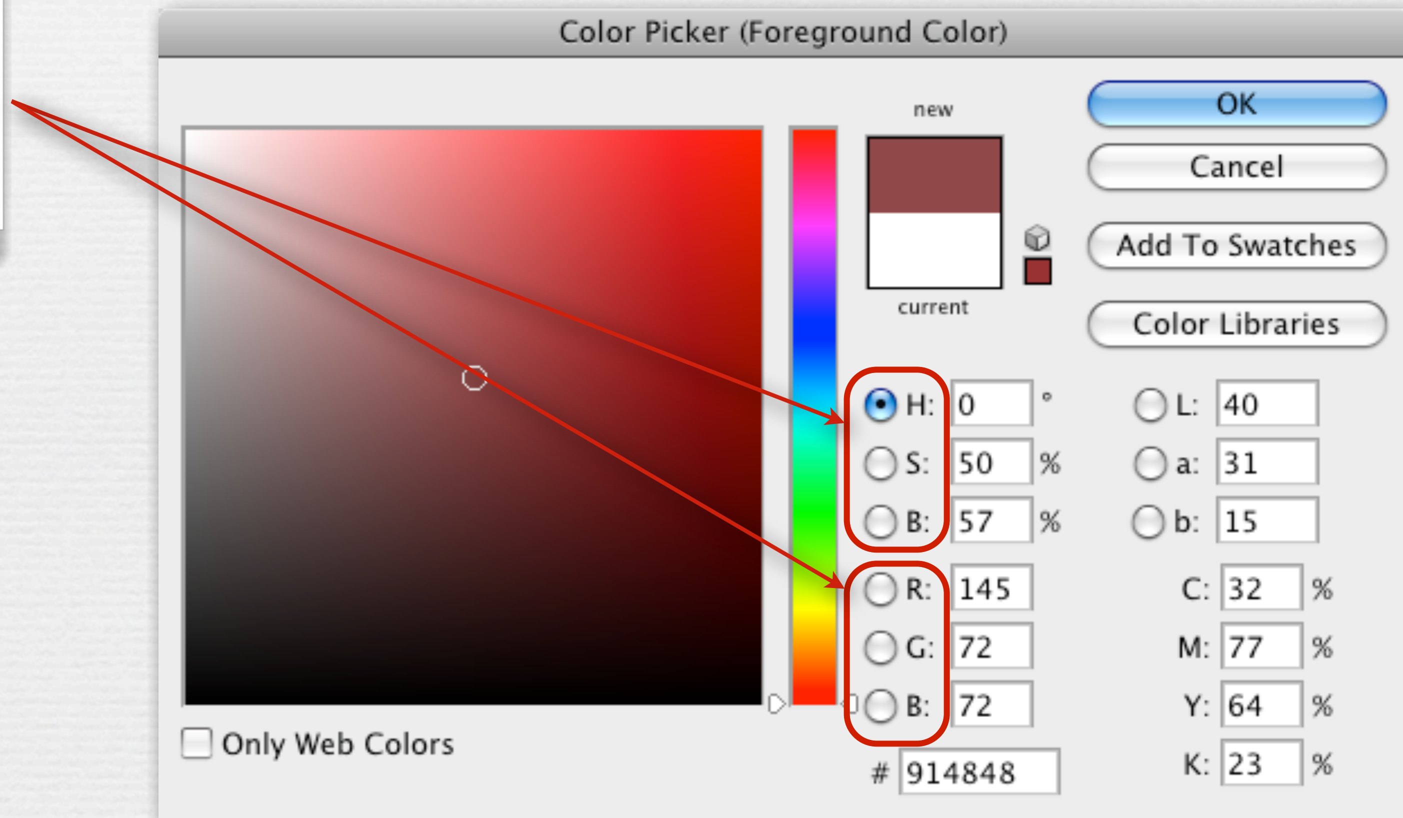

Color selection in Photoshop

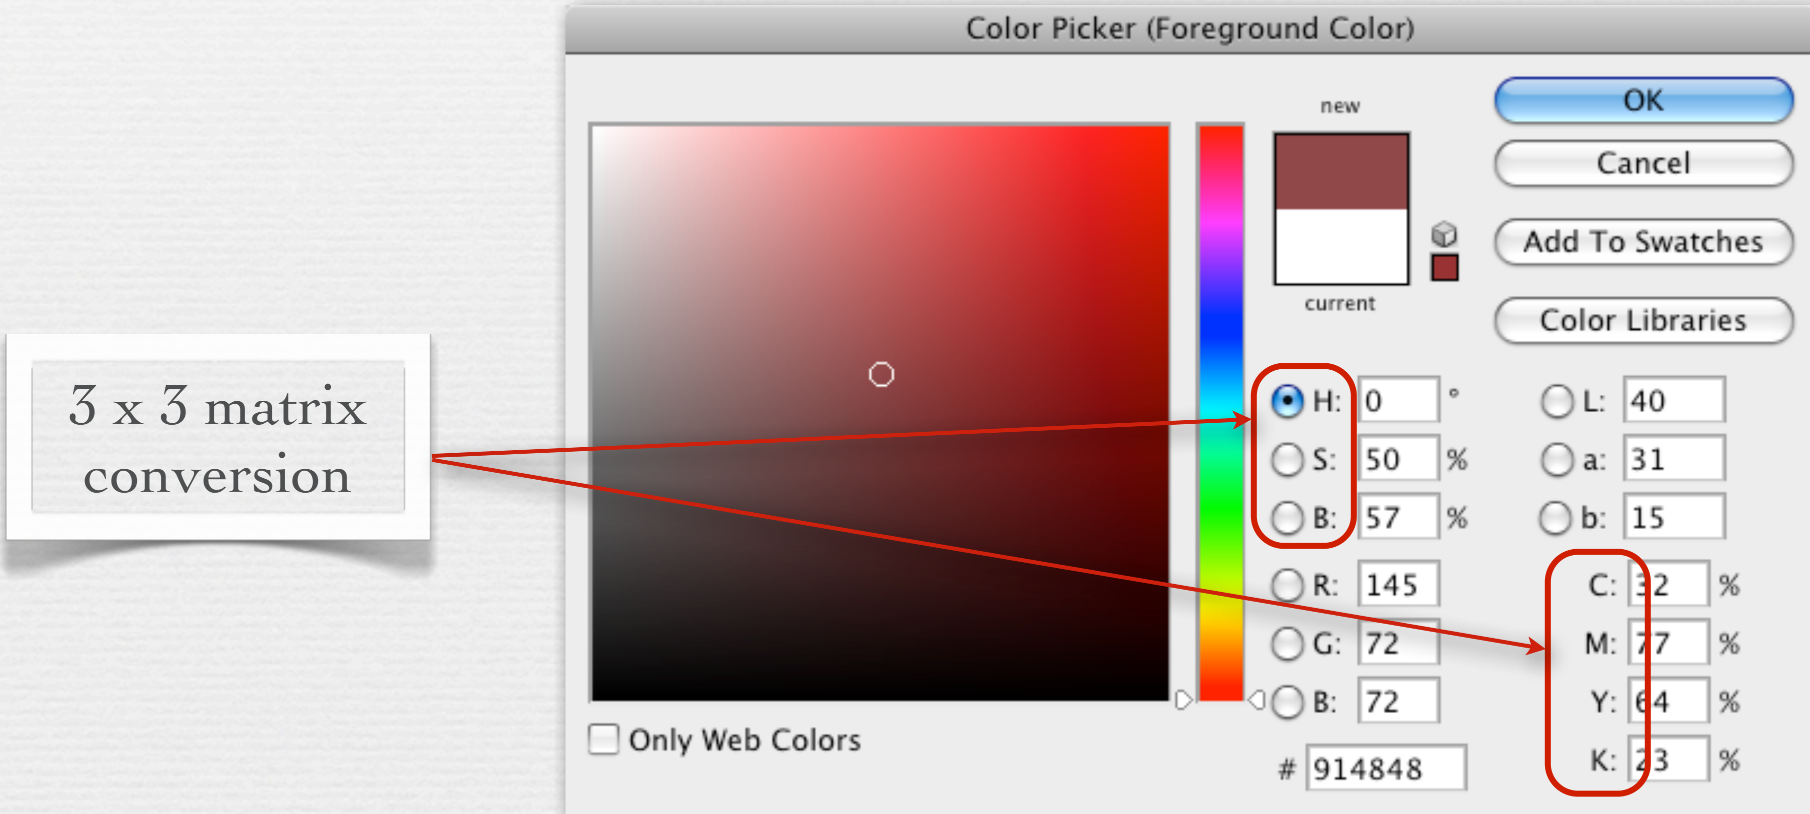

Color selection in Photoshop

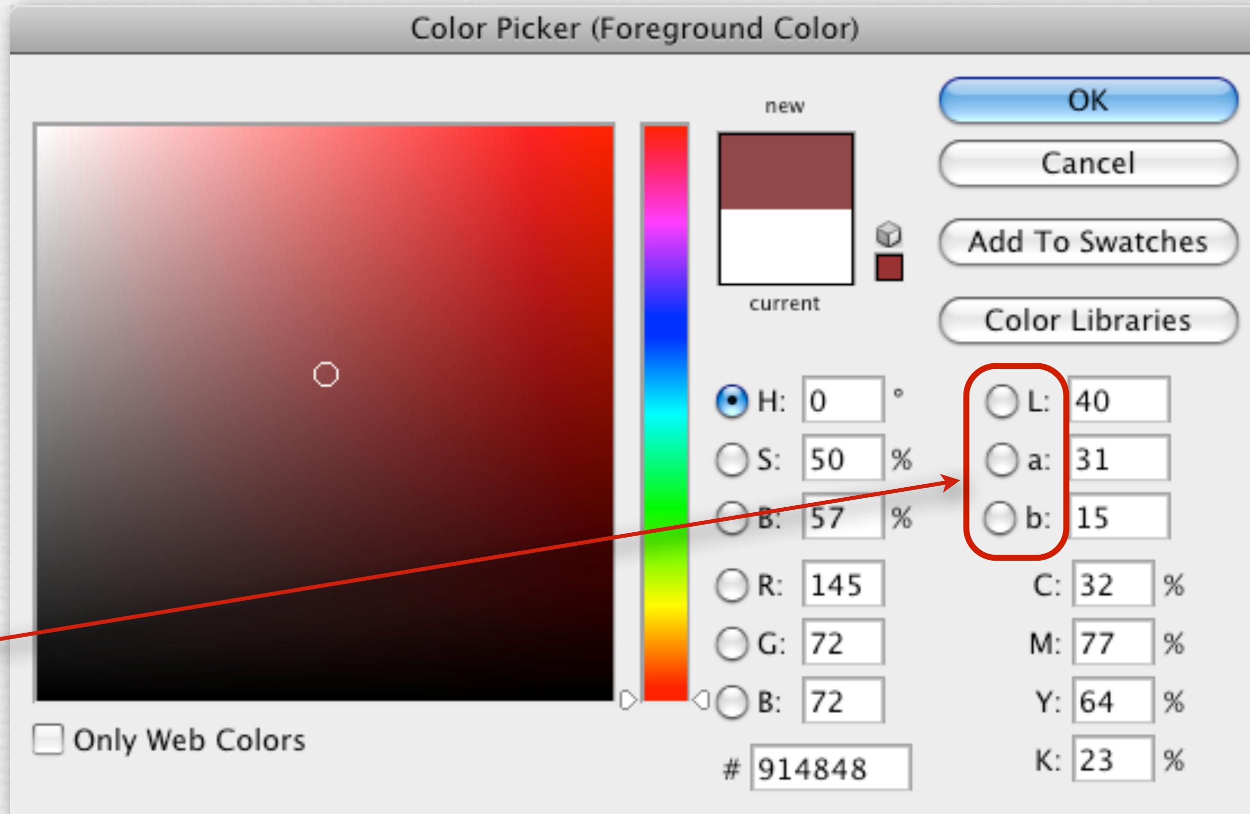

we'll cover this later in the lecture

Outline

- ✦ spectral power distributions
- ✦ color response in animals and humans
- ✦ 3D colorspace of the human visual system
- ✦ reproducing colors using three primaries
- ✦ additive versus subtractive color mixing
- ✦ cylindrical color systems used by artists (and Photoshop)
- ☞ ✦ chromaticity diagrams
  - color temperature and white balancing
  - standardized color spaces and gamut mapping
  - uniform perceptual spaces and opponent colors

Constructing the rgb chromaticity diagram

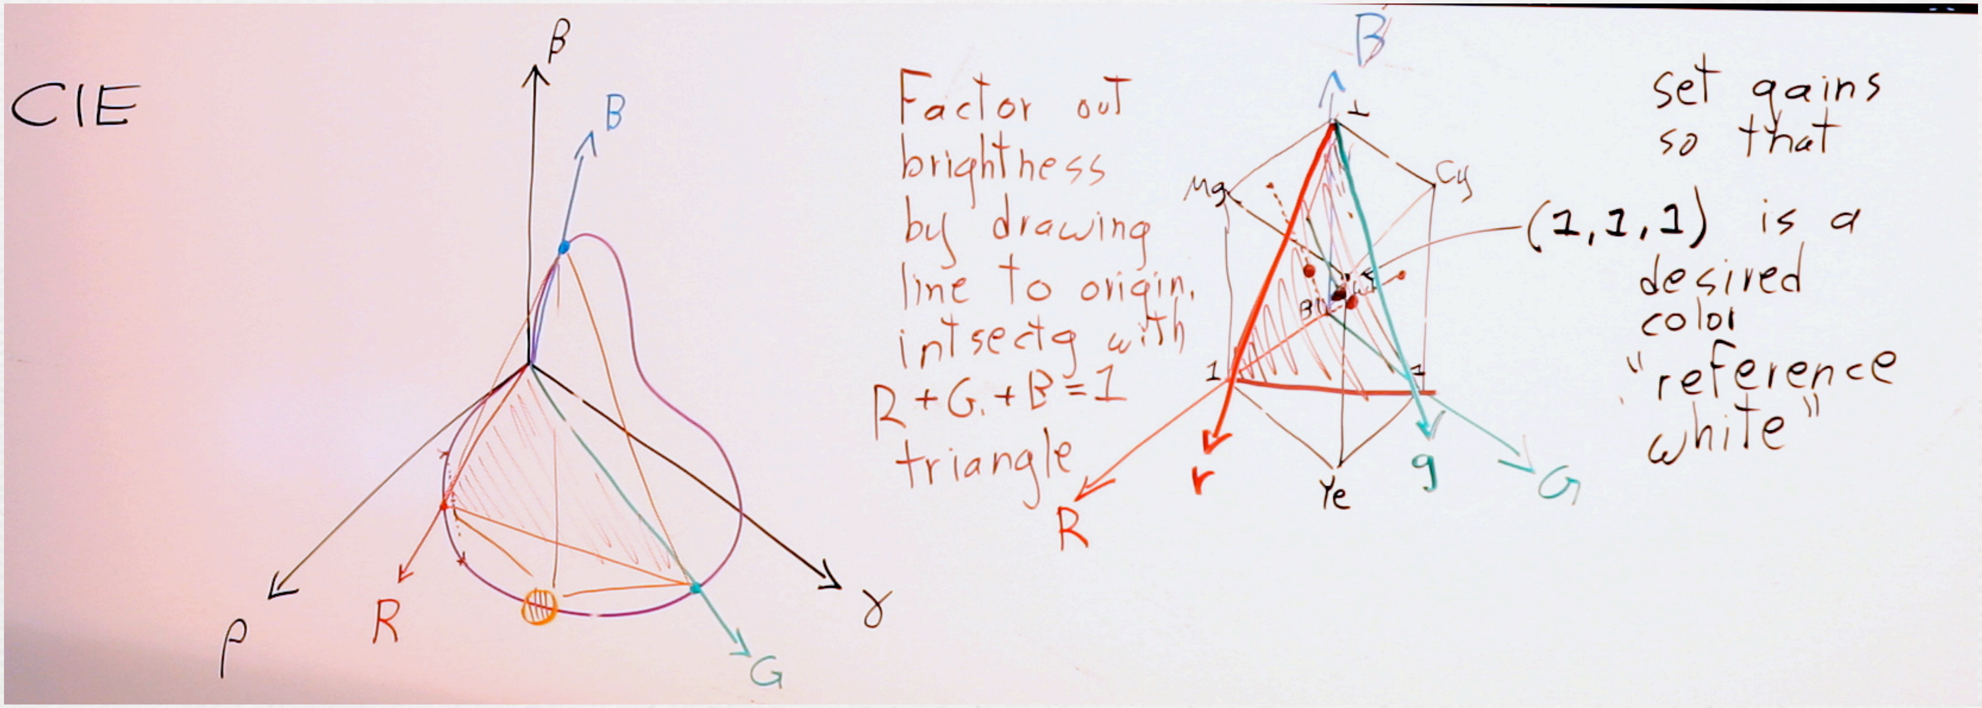

- for a point (R,G,B) in the RGB cube, the formulae that project it onto the brown triangle (i.e. draw a line to the origin and intersect that line with the triangle) are

$$r = \frac{R}{R+G+B} \qquad g = \frac{G}{R+G+B}$$

I didn't give these formulae in class. I should have. I gave them later for the analogous case of xy coordinates and the XYZ cube.

Summary of chromaticity diagrams (1 of 2)

- choose three primaries R,G,B, pure wavelengths or not
- adjust R=1,G=1,B=1 to obtain a desired *reference white*
- this yields an *RGB cube*

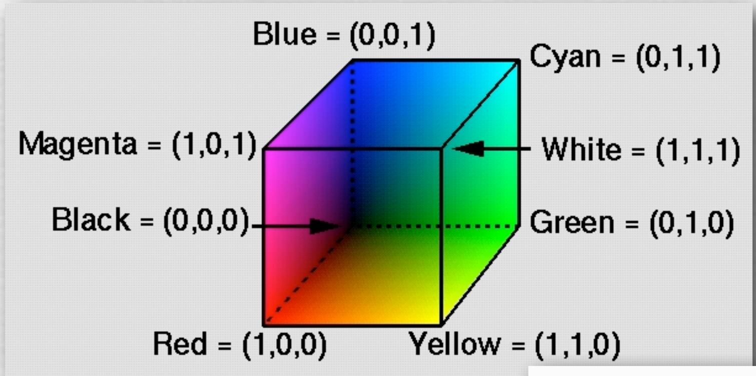

- one may factor the brightness out of any point in the cube by drawing a line to the origin and intersecting this line with the triangle made by corners Red, Green, Blue
- all points on this triangle, which are addressable by two coordinates, have the same brightness but differing *chromaticity*

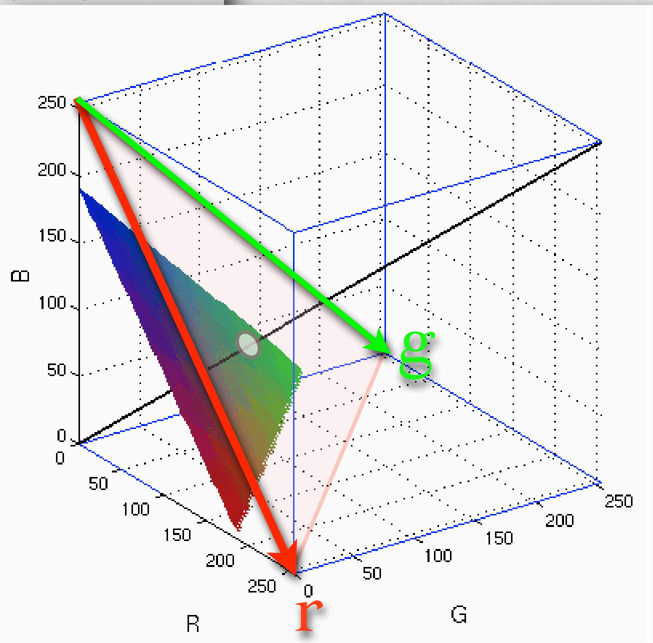

Summary of chromaticity diagrams (2 of 2)

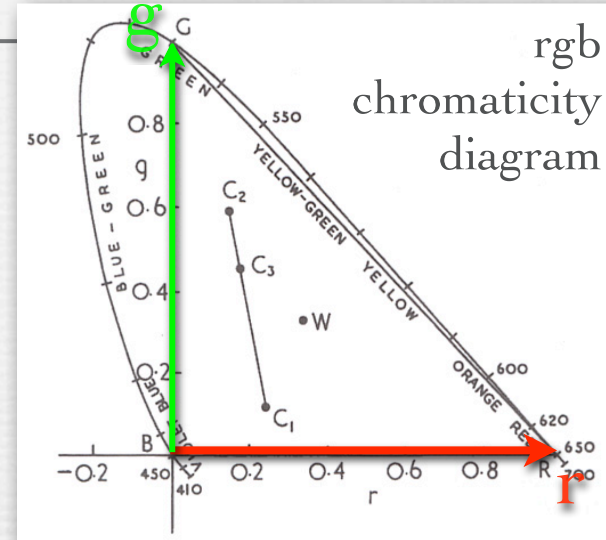

- this triangle is called the *rgb chromaticity diagram* for the chosen RGB primaries; note that it does not enclose the spectral locus
- the same construction can be performed using <u>any</u> set of 3 vectors in (ρ, γ, β) space as primaries, even physically impossible ones
- choosing a set XYZ that encloses the locus yields new matching functions $\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)$ and the CIE *xyz chromaticity diagram*

no negative values!

$\bar{r}(\lambda), \bar{g}(\lambda), \bar{b}(\lambda)$ matching functions for some primaries

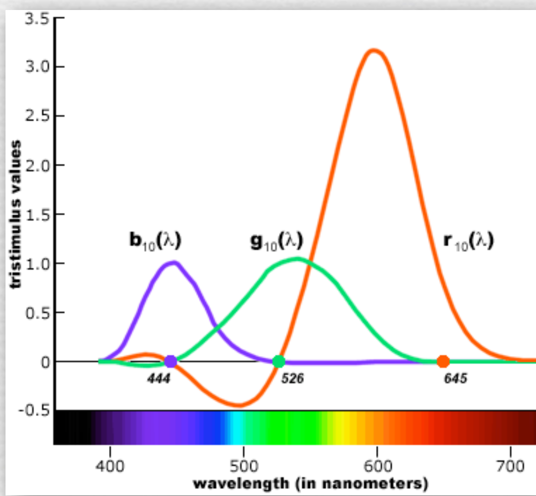

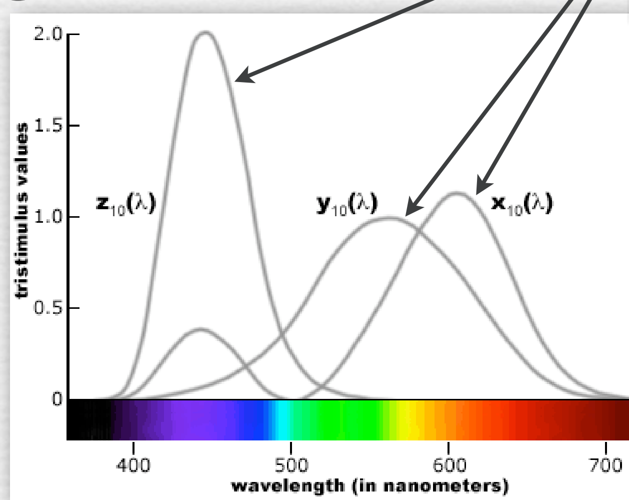

$\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)$ standardized matching functions

Summary of chromaticity diagrams (2 of 2)

- this triangle is called the *rgb chromaticity diagram* for the chosen RGB primaries; note that it does not enclose the spectral locus
- the same construction can be performed using <u>any</u> set of 3 vectors in (ρ, γ, β) space as primaries, even physically impossible ones
- choosing a set XYZ that encloses the locus yields new matching functions $\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)$ and the CIE *xyz chromaticity diagram*

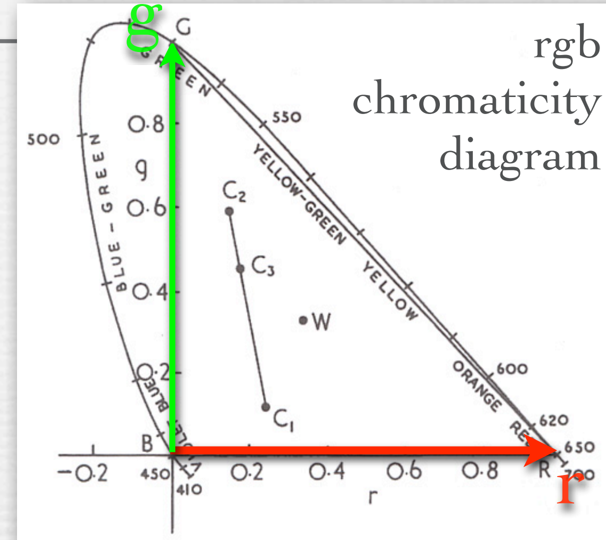

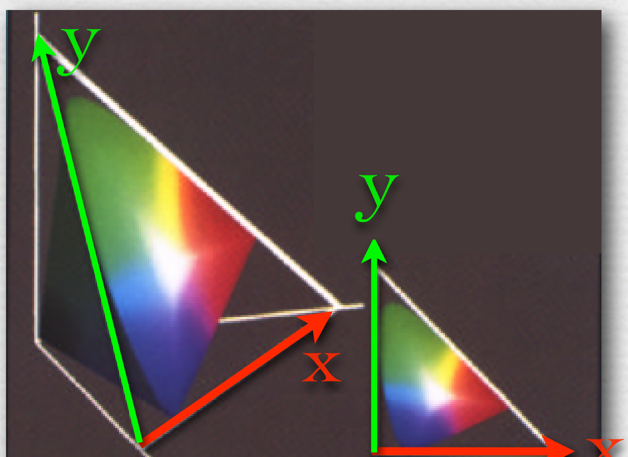

(Foley)

line of extra-spectral purples

(Hunt)

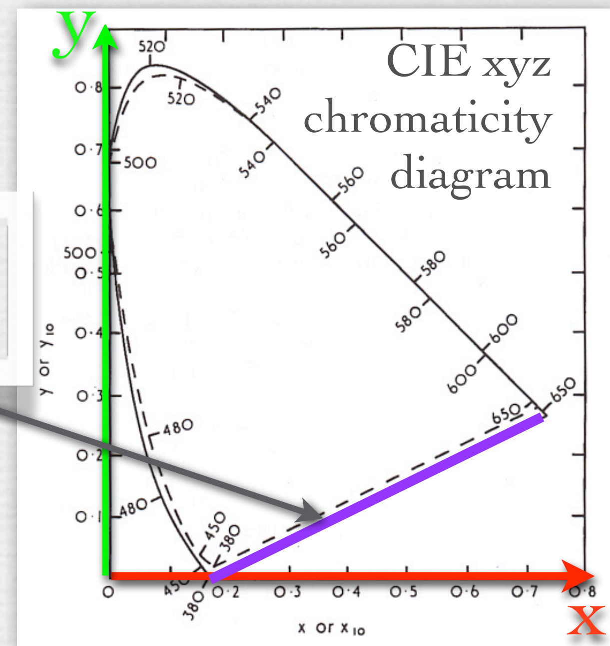

Constructing the xyz chromaticity diagram

- the X, Y, and Z axes are placed in (ρ, γ, β) space so as to enclose the locus of spectral colors in their all-positive octant

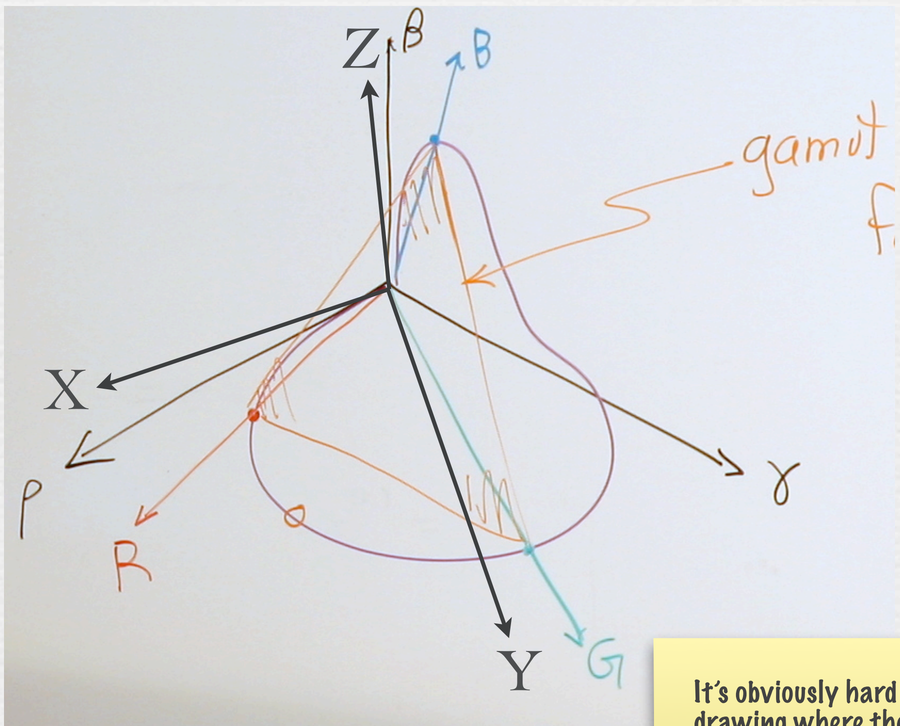

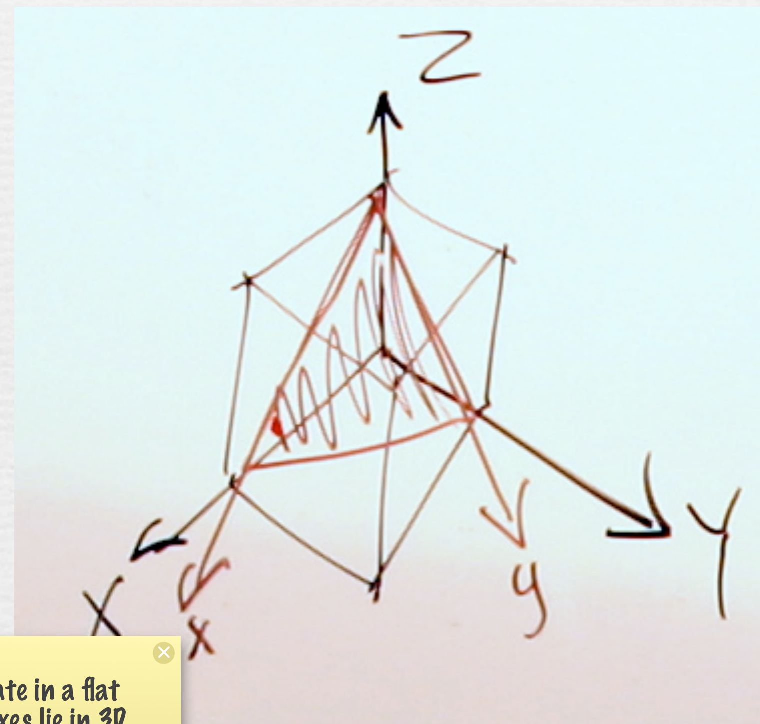

It's obviously hard to indicate in a flat drawing where the X,Y,Z axes lie in 3D. Just think of them as the indicated R,G,B axes, but pushed away from the spectral locus until the locus lies on the positive side of each axis, i.e. to the right of X, above Y, and to the right of Z is it's drawn here.

Application of chromaticity diagrams #1: color temperature and white balancing

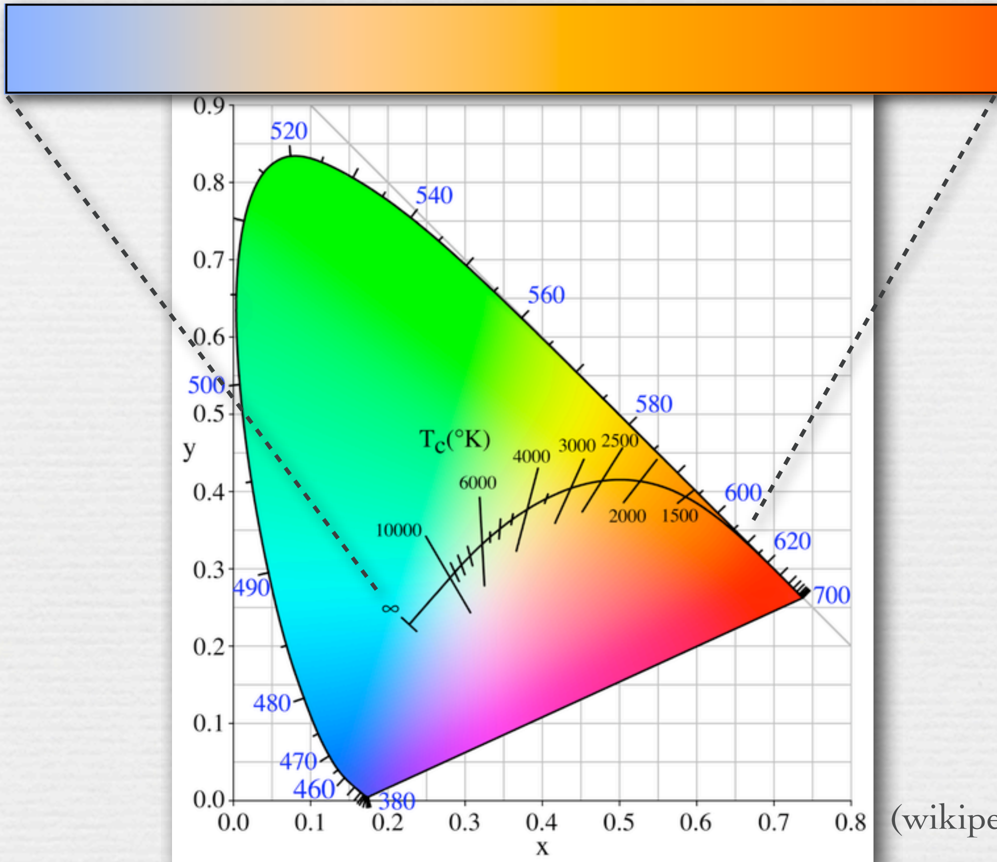

(wikipedia)

<u>correlated color temperatures</u>

3200°K incandescent light
4000°K cool white fluorescent
5000°K equal energy white (D50, E)
6000°K midday sun, photo flash
6500°K overcast, television (D65)
7500°K northern blue sky

- the apparent colors emitted by a *black-body radiator* heated to different temperatures fall on a curve in the chromaticity diagram
- for non-blackbody sources, the nearest point on the curve is called the *correlated color temperature*

White balancing in digital photography

- 1. find the color temperature of the illumination
  - 1a. note the (R,G,B) of an object in the photograph you think was white in the real world, or equivalently,...
  - 1b. choose an (R,G,B) you think represents the color that white objects in the scene ended up as in the photograph, then
- 2. scale the RGB values of all pixels in the photograph proportionately up or down so that the chosen (R,G,B) becomes (1,1,1) ($0 \leq RGB \leq 1$)
- the eventual appearance of (1,1,1) depends on the reference white of the camera or file format (assuming no later fiddling with white balance)
  - the color space of most digital cameras is sRGB
  - the reference white for sRGB is D65 (6500°K)
- thus, white balancing forces chosen (R,G,B) to appear 6500°K

Finding the color temperature of the illumination

- Auto White Balance (AWB)
  - gray world: assume the average color of a scene is gray, so force the average color to be gray - often inappropriate

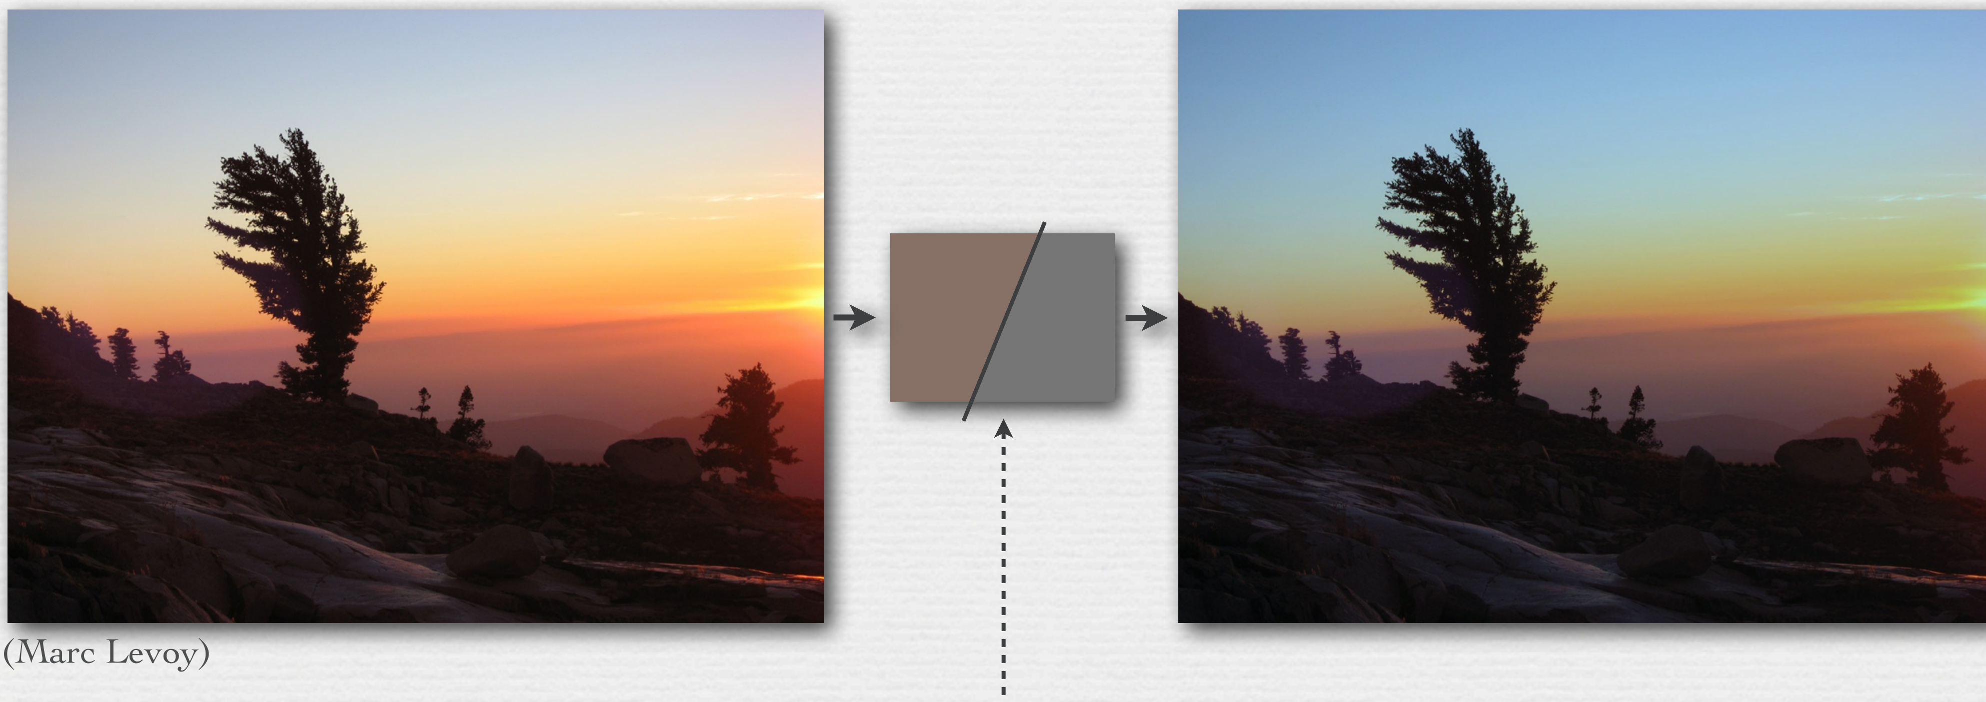

(Marc Levoy)

(R, G, B) = (100%, 81%, 73%) ➔ (100%, 100% 100%)

Finding the color temperature of the illumination

- Auto White Balance (AWB)
  - gray world: assume the average color of a scene is gray, so force the average color to be gray - often inappropriate
  - assume the brightest pixel (after demosaicing) is a specular highlight and therefore white - fails if that pixel is saturated

As I mentioned in class, the specular highlight method also fails if the specular highlight occurs on a colored metal, like gold or brass. In that case, the specular highlight will be gold-colored even if the illumination is perfectly white. Using the color of that highlight to judge the color of the illumination will yield incorrect results.

Finding the color temperature of the illumination

- Auto White Balance (AWB)
  - gray world: assume the average color of a scene is gray, so <u>force</u> the average color to be gray - often inappropriate
  - assume the brightest pixel (after demosaicing) is a specular highlight and therefore white - fails if that pixel is saturated
  - find a white object in the scene - but how??

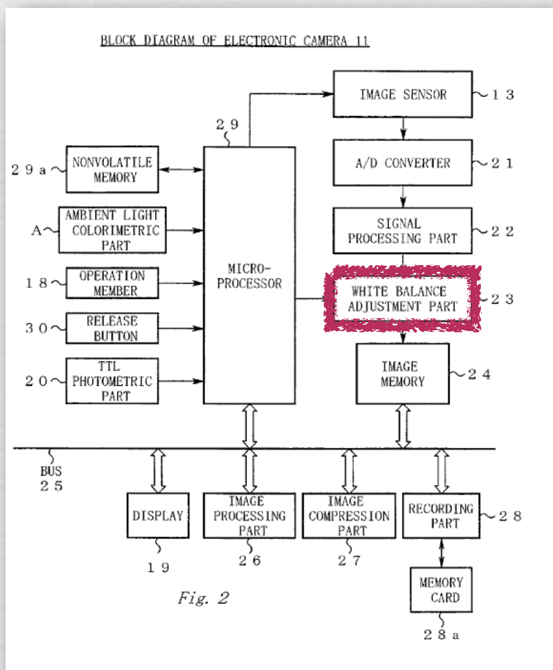

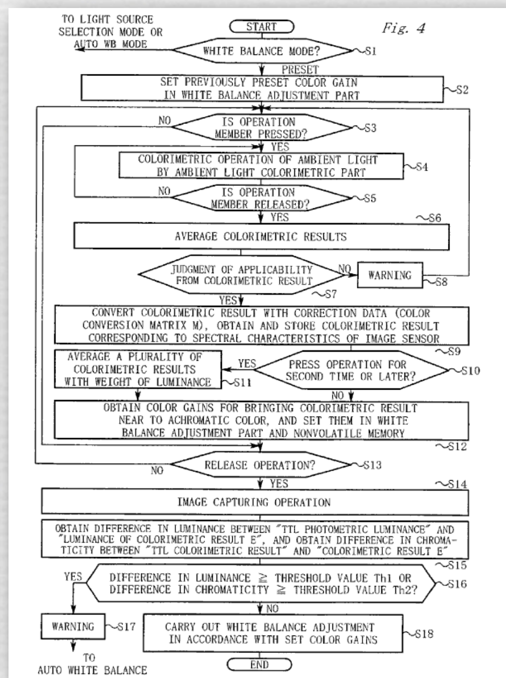

(Nikon patent)

Finding the color temperature of the illumination

- Auto White Balance (AWB)
  - gray world: assume the average color of a scene is gray, so force the average color to be gray - often inappropriate
  - assume the brightest pixel (after demosaicing) is a specular highlight and therefore white - fails if that pixel is saturated
  - find a white object in the scene - but how??

- manually specify the color temperature of the illumination
  - each color temperature maps to a unique (R,G,B)
  - scale all pixels so that (R,G,B) becomes (1,1,1)

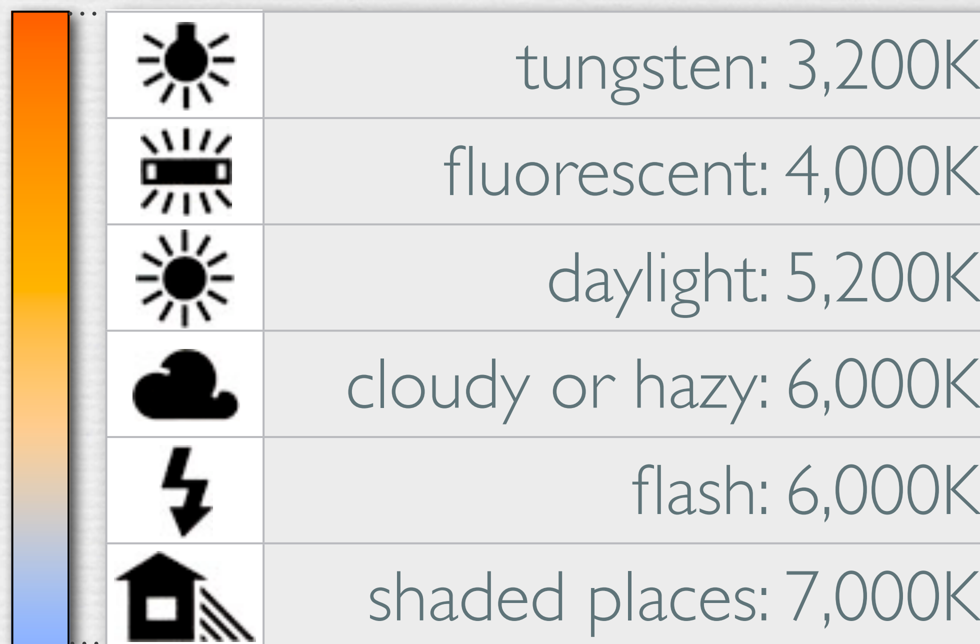

Incorrectly chosen white balance

(Eddy Talvala)

- scene was photographed in sunlight, then re-balanced as if it had been photographed under something warmer, like tungsten
  - re-balancer assumed illumination was reddish, so it boosted blues
  - same thing would have happened if originally shot with tungsten WB

Application of chromaticity diagrams #2: standardized color spaces and gamut mapping

- the chromaticities reproducible using 3 primaries fill a triangle in the xyz chromaticity diagram, a different triangle for each choice of primaries; this is called the *device gamut* for those primaries

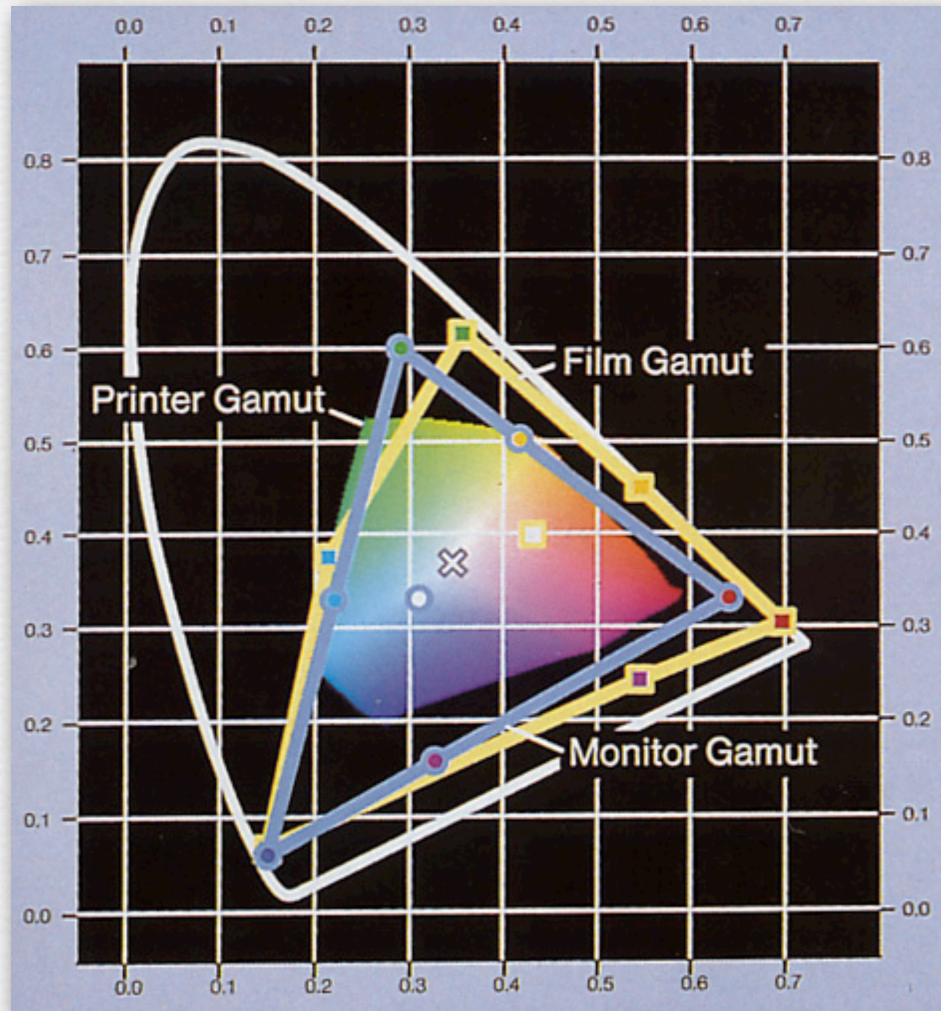

Q. Why is the diagram, scanned from a book, black outside the printer gamut?

(Foley)

Digitizing the paint colors at Hanna-Barbera Productions

physical color samples

spectrophotometer ↓

spectrum for each color

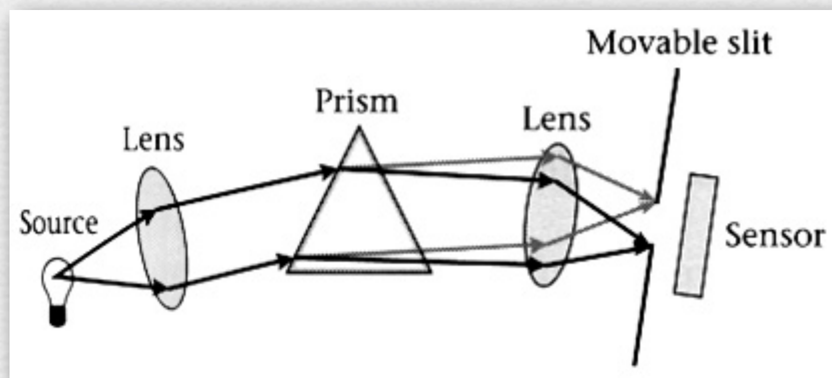

Digitizing the paint colors at Hanna-Barbera Productions

physical color samples

↓ spectrophotometer

spectrum for each color

↓ CIE matching functions

XYZ coordinates

$$(X,Y,Z) = \left(\int_{400nm}^{700nm} L_e(\lambda)\, \bar{x}(\lambda)\, d\lambda,\ \int_{400nm}^{700nm} L_e(\lambda)\, \bar{y}(\lambda)\, d\lambda,\ \int_{400nm}^{700nm} L_e(\lambda)\, \bar{z}(\lambda)\, d\lambda \right)$$

Digitizing the paint colors at Hanna-Barbera Productions

physical color samples

↓ spectrophotometer

spectrum for each color

↓ CIE matching functions

XYZ coordinates

↓ projection onto X=Y=Z=1 plane

xy chromaticity coordinates

$$x = \frac{X}{X+Y+Z} \qquad y = \frac{Y}{X+Y+Z}$$

Digitizing the paint colors at Hanna-Barbera Productions

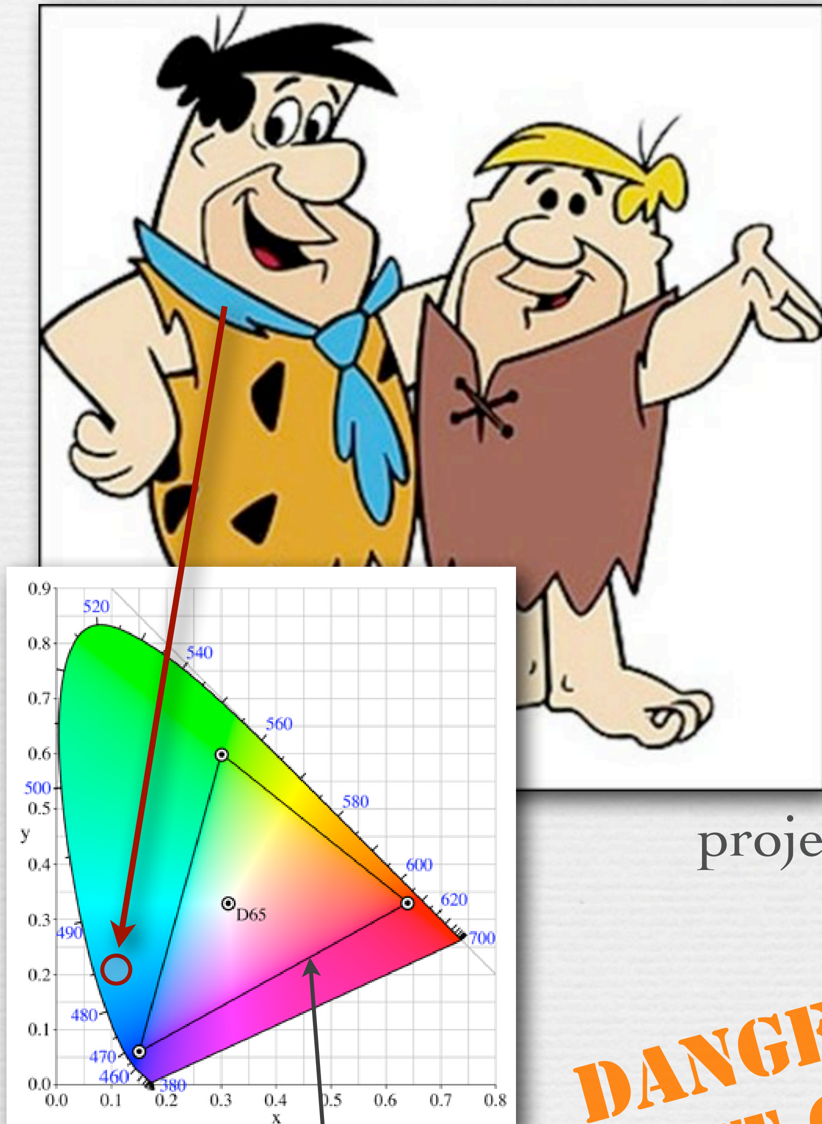

NTSC gamut

physical color samples

↓ spectrophotometer

spectrum for each color

↓ CIE matching functions

XYZ coordinates

↓ projection onto X=Y=Z=1 plane

xy chromaticity coordinates

DANGER: NECKTIE OUT OF GAMUT !!

Fixing out-of-gamut colors

- in the drawing below, the right answer is to project the errant color back towards the reference white until it is back inside the gamut triangle
  - the effect of this projection is to desaturate the color
- the follow-on question becomes, does one also compress all other colors towards the reference white?
  - so that color relationships stay unchanged?

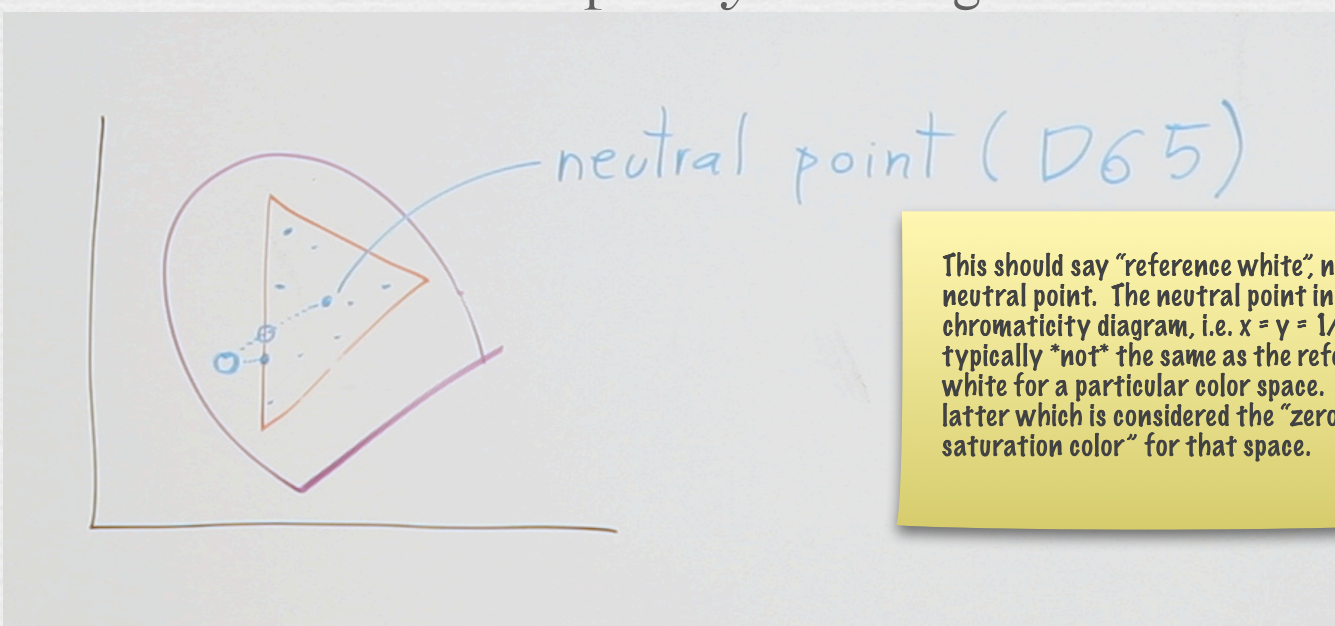

This should say "reference white", not neutral point. The neutral point in the xyz chromaticity diagram, i.e. x = y = 1/3, is typically *not* the same as the reference white for a particular color space. It is the latter which is considered the "zero saturation color" for that space.

Uniform perceptual color spaces

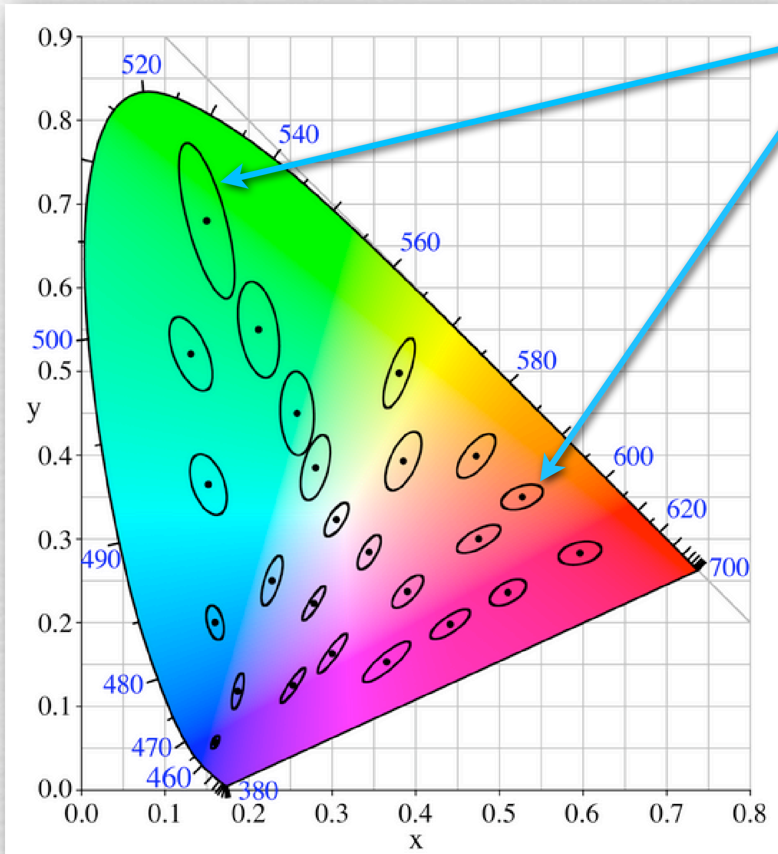

equally perceivable
"MacAdam ellipses"

(wikipedia)

(Wyszecki and Stiles)

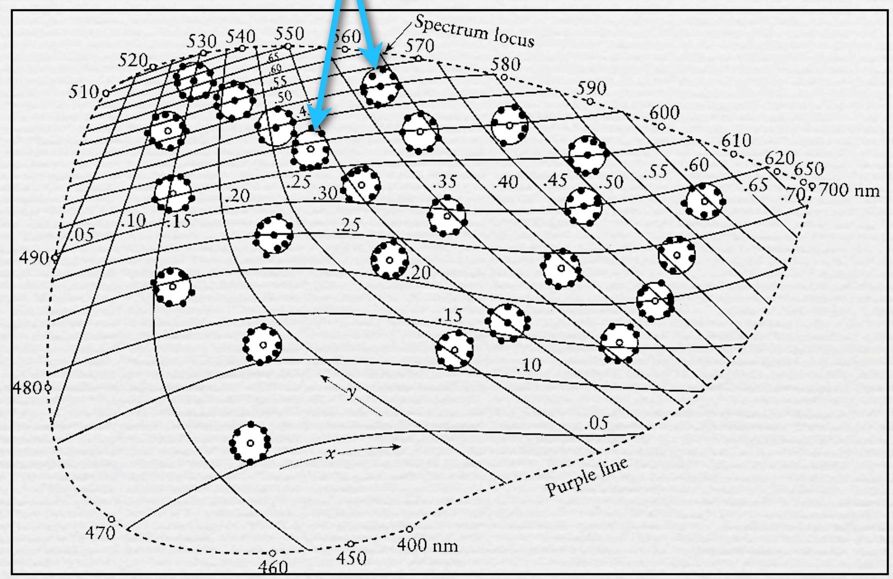

a non-linear mapping

- in the xyz chromaticity diagram, equal distances on the diagram are not equally perceivable to humans
- to create a space where they are equally perceivable, one must distort XYZ space (and the xyz diagram) non-linearly

CIELAB space (a.k.a. L*a*b*)

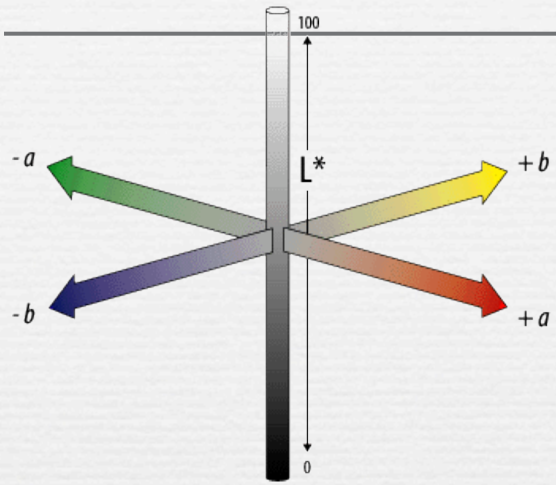

$$L = 25\left(100\frac{Y}{Y_0}\right)^{1/3} - 16$$

$$a = 500\left[\left(\frac{X}{X_0}\right)^{1/3} - \left(\frac{Y}{Y_0}\right)^{1/3}\right]$$

$$b = 200\left[\left(\frac{Y}{Y_0}\right)^{1/3} - \left(\frac{Z}{Z_0}\right)^{1/3}\right]$$

non-linear mapping (a gamma transform)

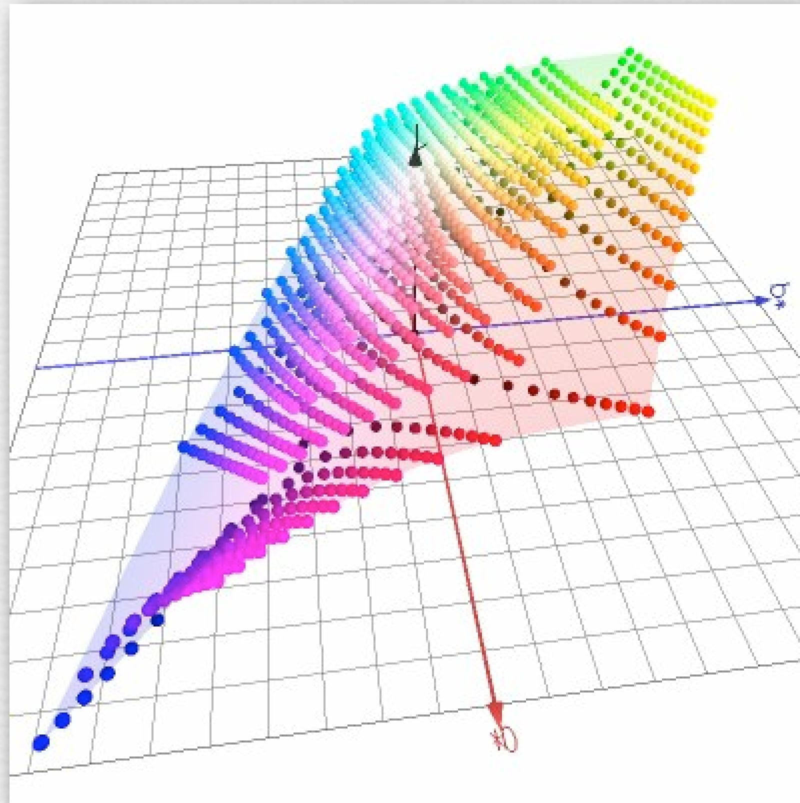

the XYZ cube in L*a*b* space

- L* is lightness
- a* and b* are color-opponent pairs
  - a* is red-green, and b* is blue-yellow

Complementary colors

(http://www.handprint.com/HP/WCL/color6.html)

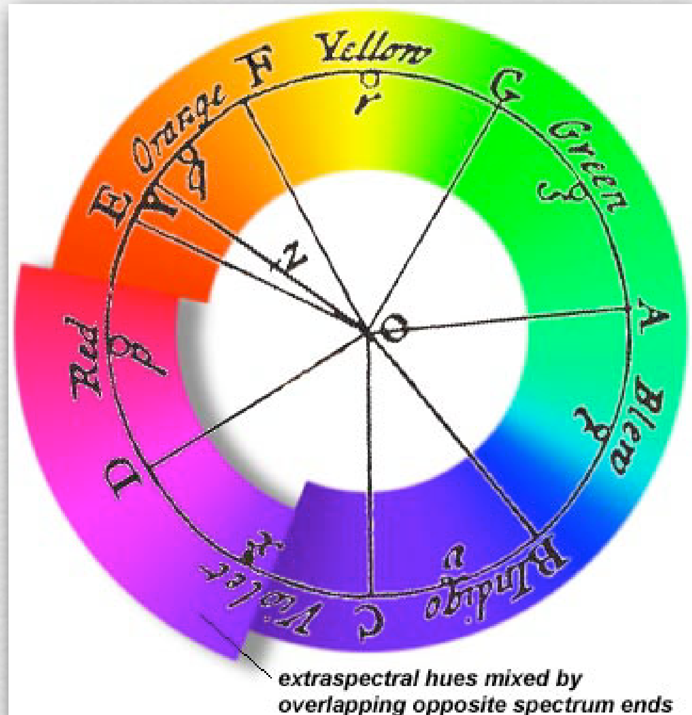

- Leonardo described complementarity of certain pairs of colors
- Newton arranged them opposite one another across his circle
- Comte de Buffon (1707-1788) observed that afterimage colors were exactly the complementary colors

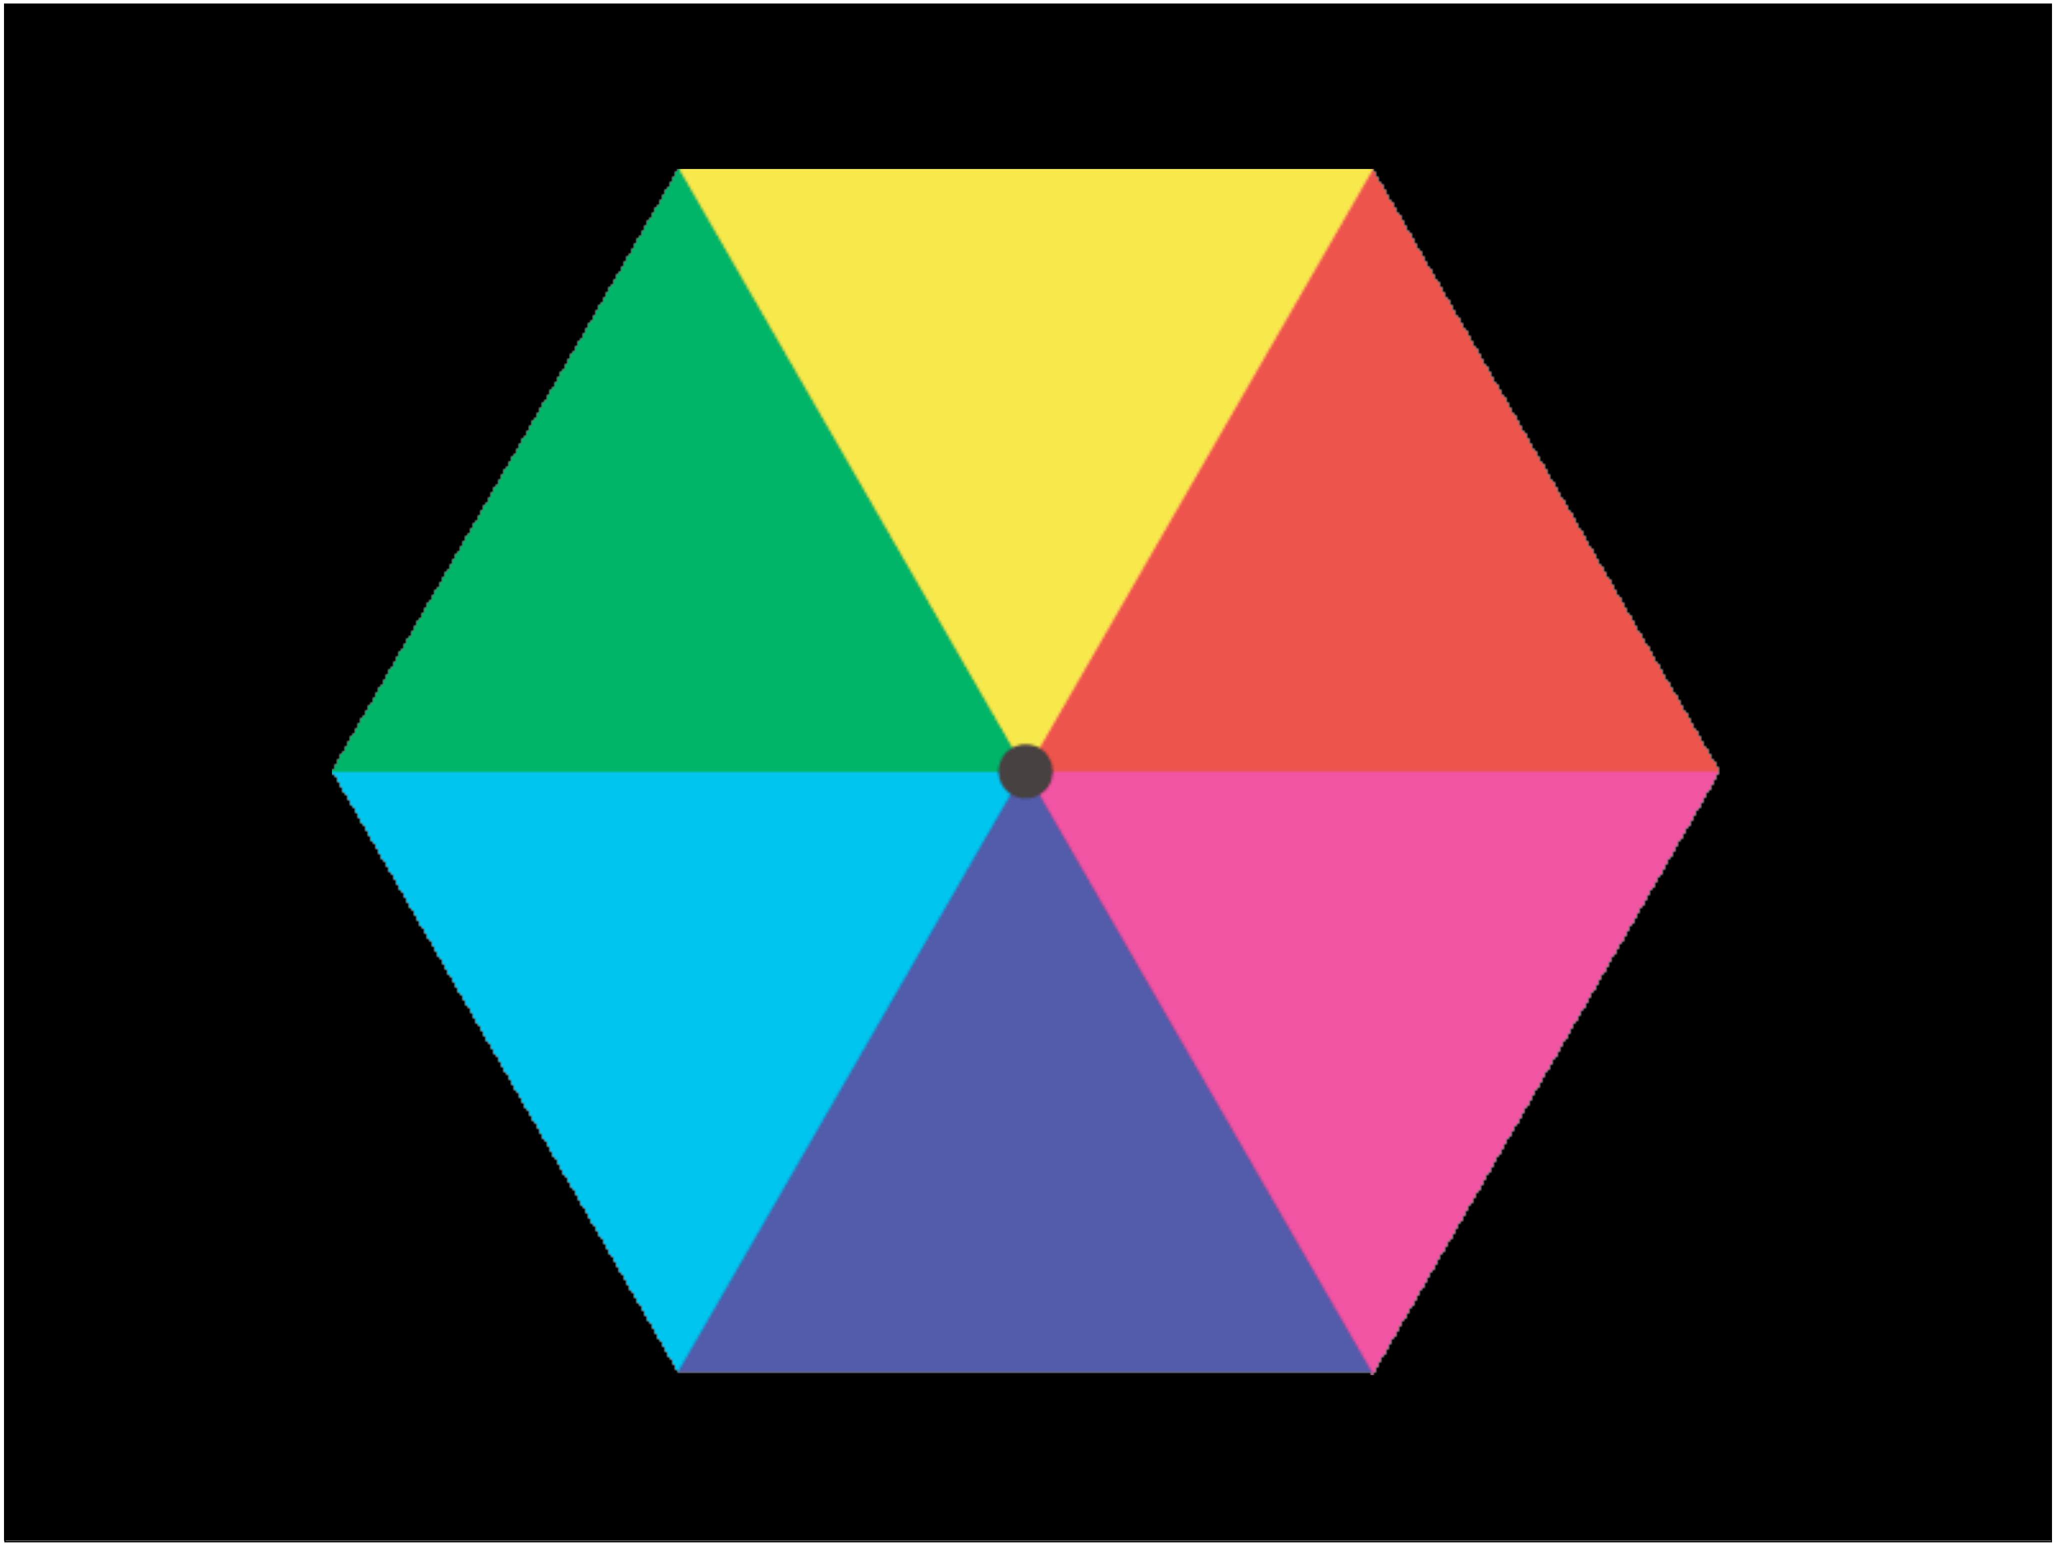

Stare at slide #31 (2 slides back) for 30 seconds, fixating your gaze as best as possible on the black dot in the middle, then without averting your gaze, switch to the white slide that follows it. You should see an afterimage floating in your vision that looks like the example below.

image

afterimage

Opponent colors

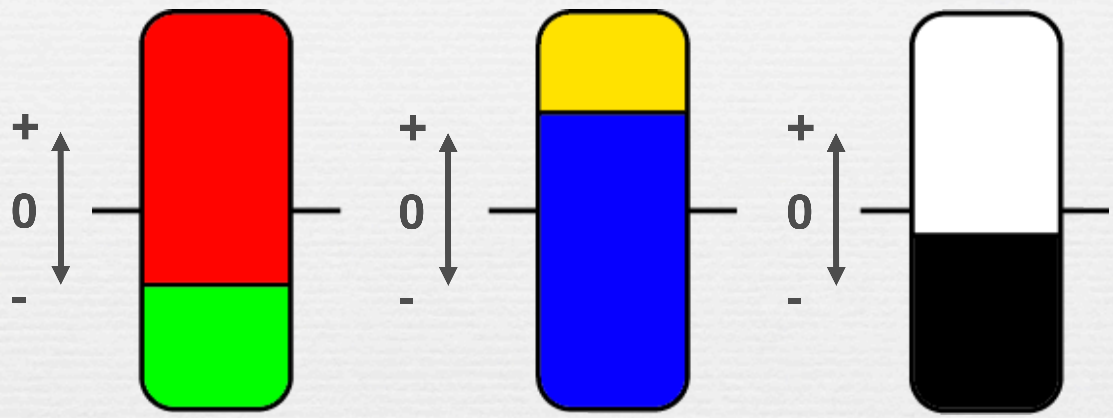

red/green receptors

blue/yellow receptors

black/white receptors

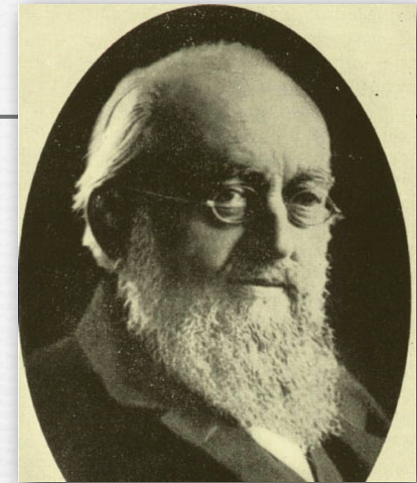

Ewald Hering
(1834-1918)

- observed that humans don't see reddish-green colors or blueish-yellow colors
- hypothesized three receptors, as shown above

Opponent colors wiring

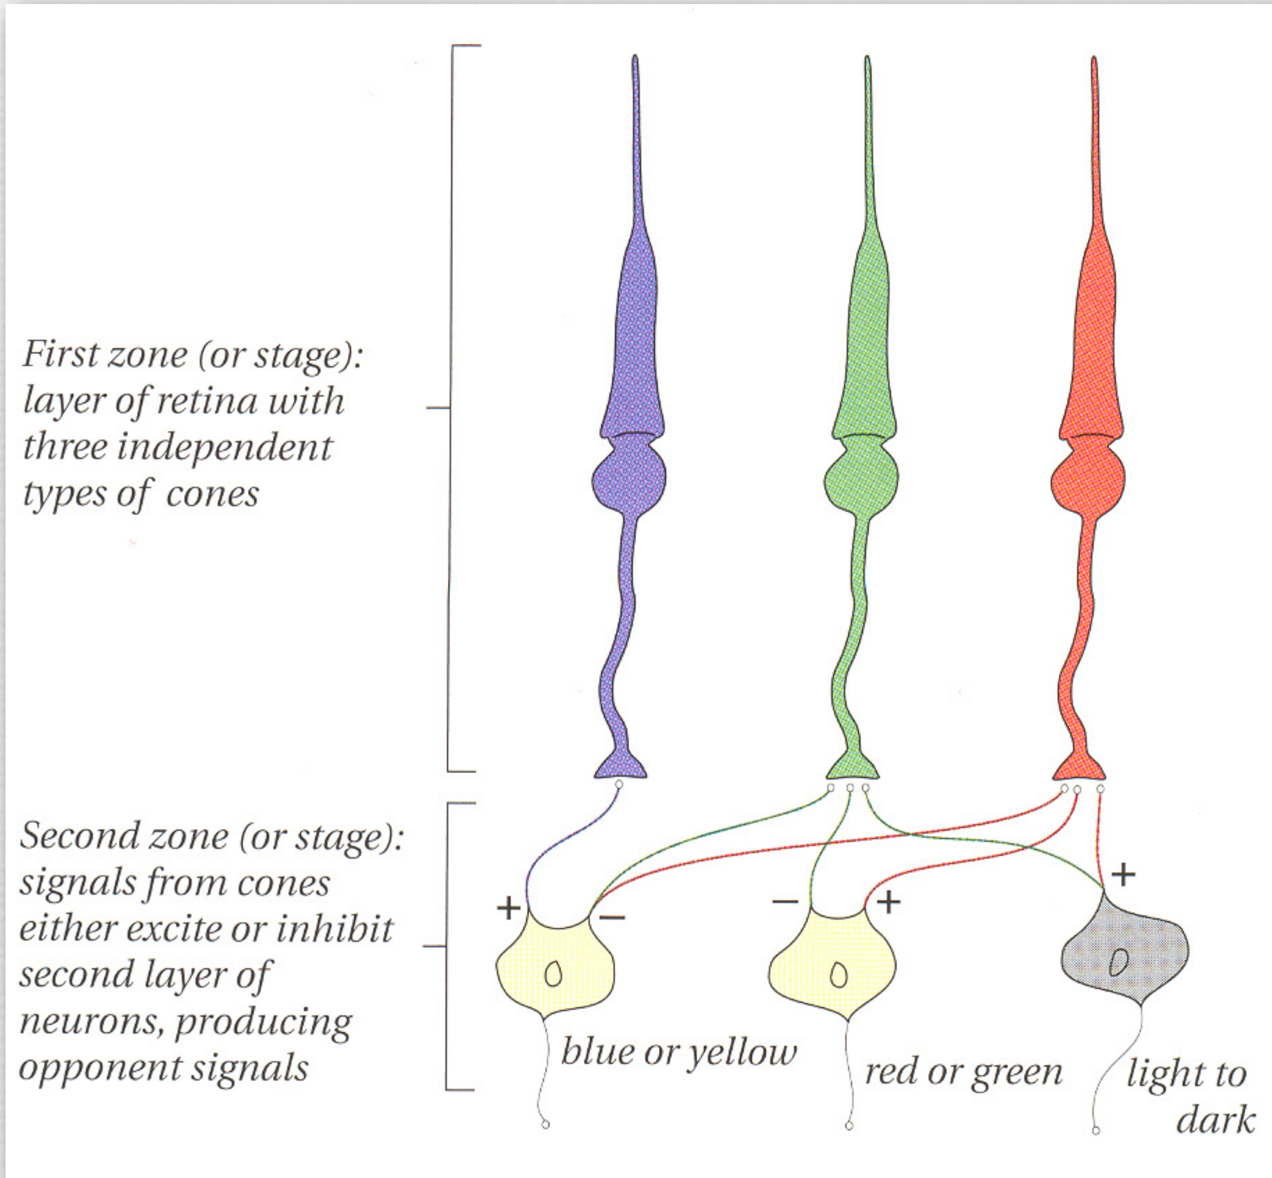

(wikipedia)

Practical use of opponent colors: NTSC color television

- color space is YIQ
  - Y = luminance
  - I = orange-blue axis
  - Q = purple-green axis

$$R, G, B, Y \in [0,1], \quad I \in [-0.5957, 0.5957], \quad Q \in [-0.5226, 0.5226]$$

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.595716 & -0.274453 & -0.321263 \\ 0.211456 & -0.522591 & 0.311135 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

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

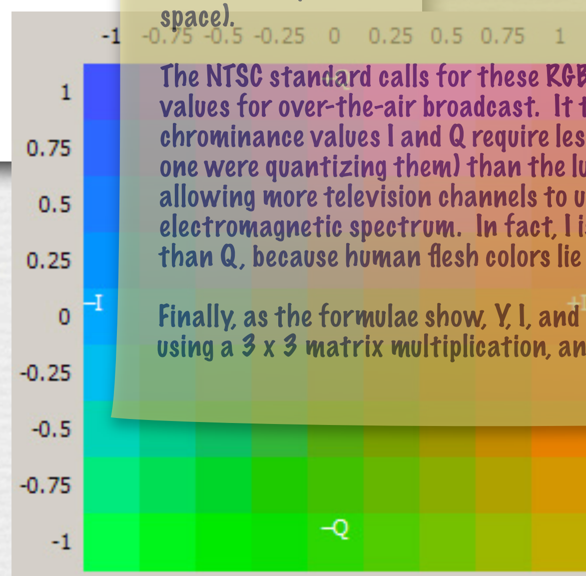

The NTSC standard assumes that American color television sets employ three particular primaries, denoted here as R, G, and B. These primaries originally represented three particular phosphorescent materials. Now they merely refer to three particular pairs of x,y chromaticity coordinates (see slide #38 on how to define an RGB color space).

The NTSC standard calls for these RGB values to be converted to YIQ values for over-the-air broadcast. It then turns out that the chrominance values I and Q require less radio bandwidth (fewer bits if one were quantizing them) than the luminance value Y, thereby allowing more television channels to use the same chunk of electromagnetic spectrum. In fact, I is allocated more bandwidth than Q, because human flesh colors lie in the +I direction.

Finally, as the formulae show, Y, I, and Q can be computed from R, G, B using a 3 x 3 matrix multiplication, and vice versa. You are not

Practical use of opponent colors: JPEG compression

(wikipedia)

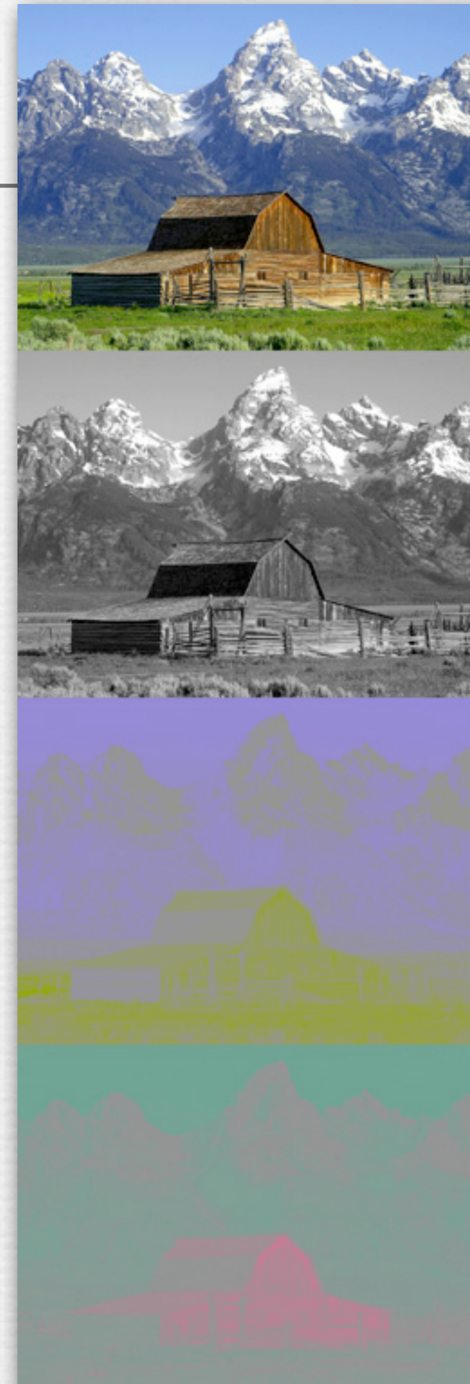

- color space is YCbCr
  - Y = luminance
  - Cb = yellow-blue axis
  - Cr = red-green axis

```
JFIF-Y'CbCr (601) from "digital 8-bit R'G'B'   "
========================================================================
Y' =         + 0.299     * R'd + 0.587     * G'd + 0.114     * B'd
Cb = 128     - 0.168736 * R'd - 0.331264 * G'd + 0.5       * B'd
Cr = 128     + 0.5      * R'd - 0.418688 * G'd - 0.081312 * B'd
........................................................................
R'd, G'd, B'd   in {0, 1, 2, ..., 255}
Y', Cb, Cr      in {0, 1, 2, ..., 255}
```

As on the previous slide, this is a color space in which one coordinate (Y) is luminance and the other two (Cb and Cr) are chrominance. Cb and Cr can be compressed more heavily than Y without creating objectionable visual artifacts. That these axes are aligned with human opponent color pairs is one of the reasons this works well.

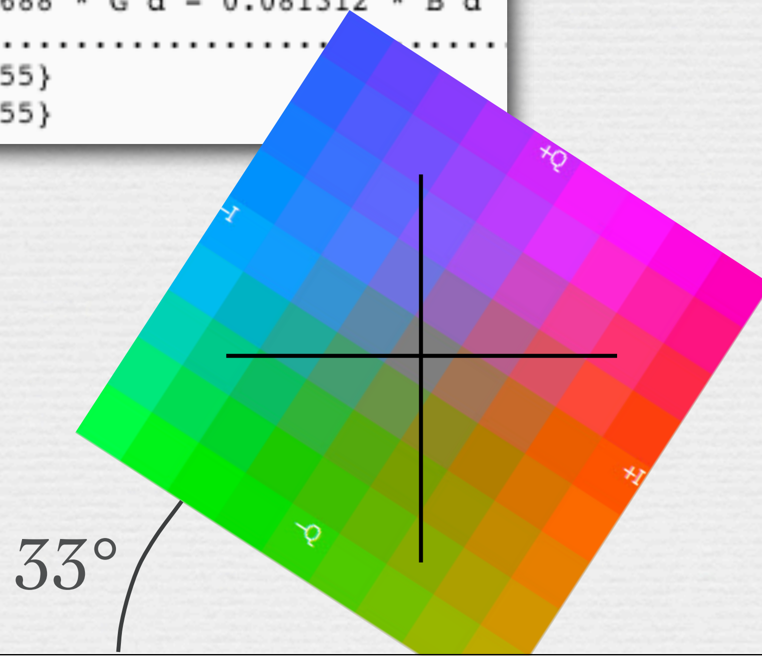

Back to applications of chromaticity diagrams: standardized RGB color spaces

- to define a new 3D color space, one needs
  - chromaticity coordinates (x,y) of each of 3 primaries
  - chromaticity coordinates of the reference white
  - maximum allowable luminance

- the mapping of a color space to xyz may be linear (3 x 3 matrix) or non-linear (like CIELAB)
  - sRGB and Adobe RGB employ non-linear mappings
  - but are not perceptually uniform

(Jeffrey Friedl)

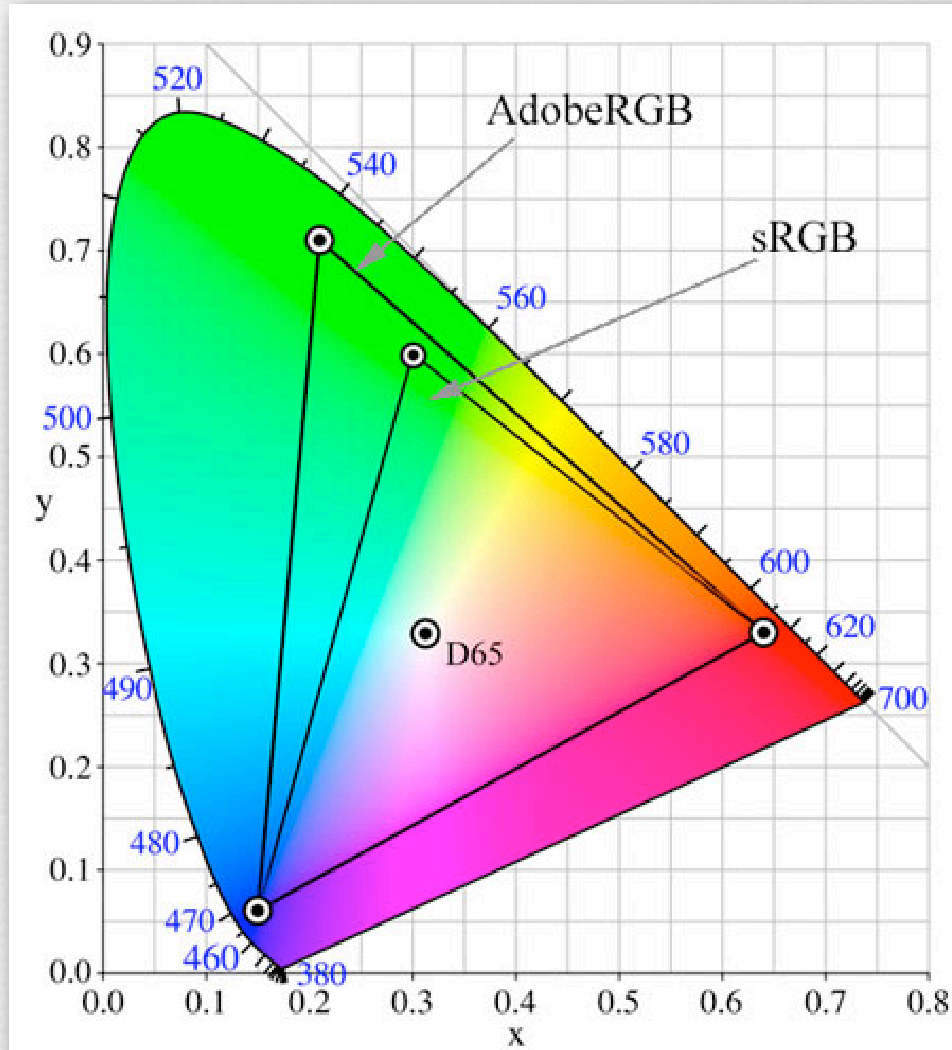

Gamut mapping, 2nd try

diagram upside-down relative to previous slides

input color space (like sRGB)

non-linear mapping ↓

perceptually uniform space (like CIELAB)

gamut mapping ↓

reduced gamut

non-linear mapping ↓

output color space (like CMYK)

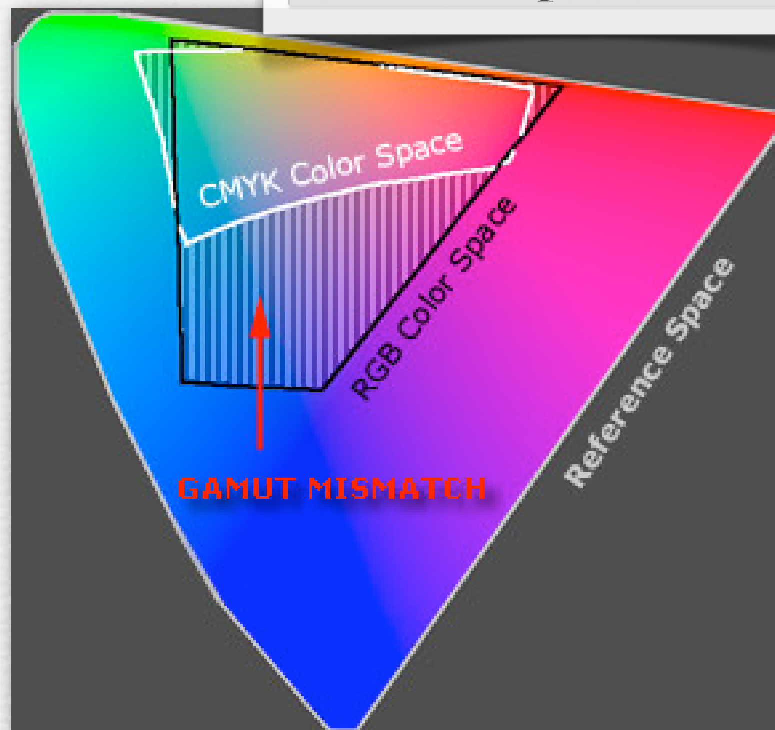

(cambridgeincolour.com)

- you can do this mapping in Photoshop, or let the printer (or other output device) do it for you

Gamut mapping options

- a.k.a. *rendering intents* (“color space conversion options”) in Photoshop
  - perceptual - smoothly shrinks all colors to fit in target gamut
  - saturated - sacrifices smoothness to maintain saturated colors
  - relative coloimetric - shrinks only out-of-gamut colors
  - absolute colorimetric - handles the white point differently

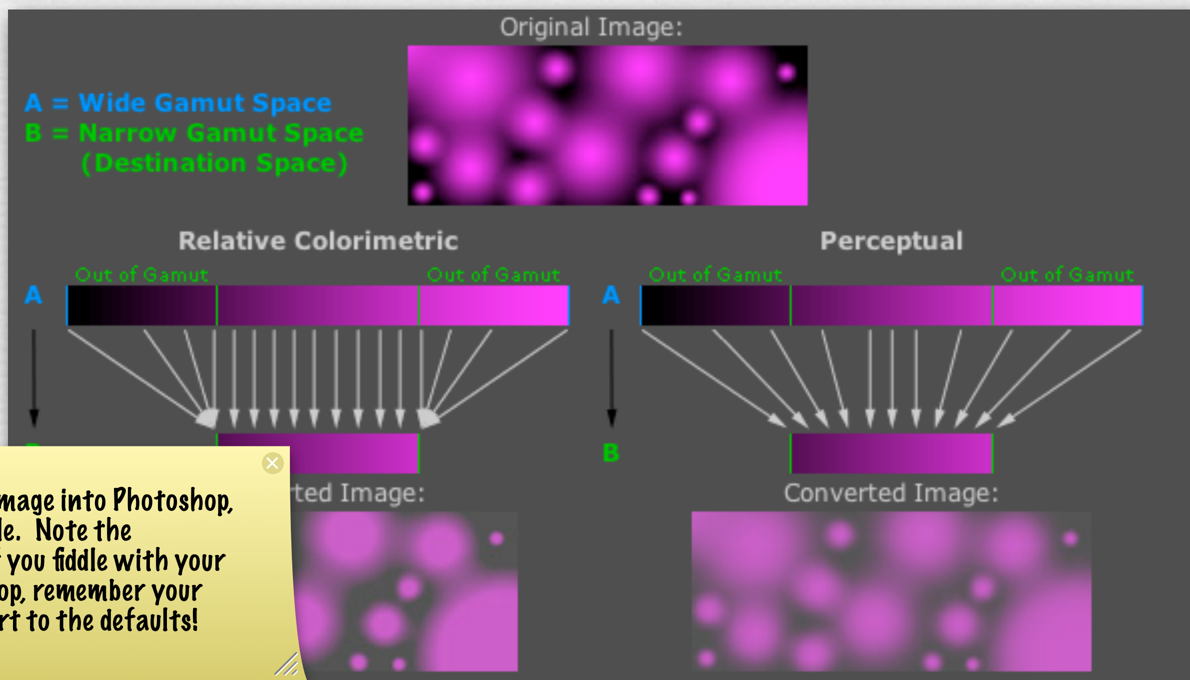

(cambridgeincolour.com)

To try this yourself, load a color image into Photoshop, then invoke Edit/Convert To Profile. Note the “rendering intents”. Warning - if you fiddle with your default color settings in Photoshop, remember your original settings, so you can revert to the defaults!

Common RGB color spaces

- Canon cameras
  - sRGB or Adobe RGB
- Nikon cameras
  - same, with additional options
- HP printers
  - ColorSmart/sRGB, ColorSync, Grayscale, Application Managed Color, Adobe RGB
- Canon desktop scanners
  - no color management (as of two years ago)
- operating systems' color management infrastructure
  - Apple ColorSync and Microsoft ICM
  - not used by all apps, disabled by default when printing

What a mess!

Slide credits

- Fredo Durand
- Bill Freeman
- Jennifer Dolson

- Robin, H., *The Scientific Image*, W.H. Freeman, 1993.
- Wandell, B., *Foundations of Vision*, Sinauer Associates, 1995.
- Hunt, R.W.G., *The Reproduction of Color* (6th ed.), John Wiley & Sons, 2004.
- Wyszecki, G. and Stiles, W.S., *Color Science* (2nd ed.), John Wiley & Sons, 1982.
- Foley, van Dam, et al., *Computer Graphics* (2nd ed.), Addison-Wesley, 1990.
- Berns, R.S., *Billmeyer and Saltzman's Principles of Color Technology* (3rd ed.), John Wiley, 2000.

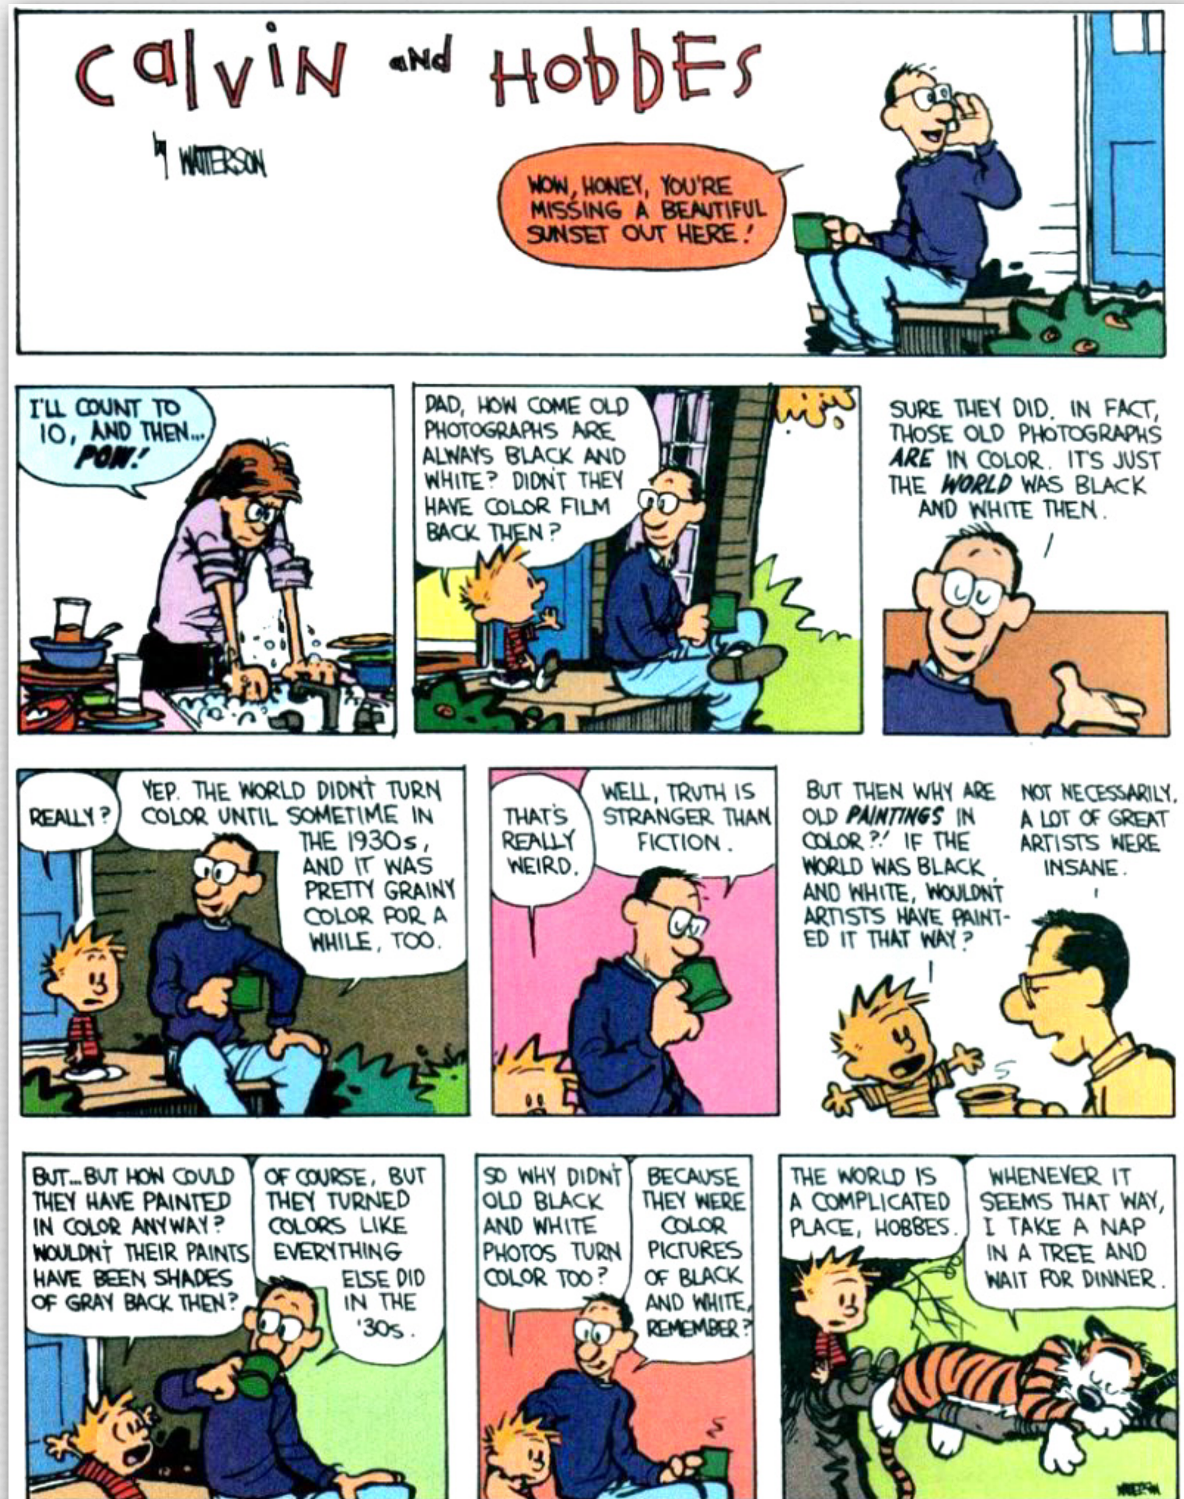
Calvin and Hobbes
by WATTERSON
WOW, HONEY, YOU'RE MISSING A BEAUTIFUL SUNSET OUT HERE!
I'LL COUNT TO 10, AND THEN... POW!
DAD, HOW COME OLD PHOTOGRAPHS ARE ALWAYS BLACK AND WHITE? DIDN'T THEY HAVE COLOR FILM BACK THEN?
SURE THEY DID. IN FACT, THOSE OLD PHOTOGRAPHS ARE IN COLOR. IT'S JUST THE WORLD WAS BLACK AND WHITE THEN.
REALLY?
YEP. THE WORLD DIDN'T TURN COLOR UNTIL SOMETIME IN THE 1930s, AND IT WAS PRETTY GRAINY COLOR FOR A WHILE, TOO.
THAT'S REALLY WEIRD.
WELL, TRUTH IS STRANGER THAN FICTION.
BUT THEN WHY ARE OLD PAINTINGS IN COLOR?! IF THE WORLD WAS BLACK AND WHITE, WOULDN'T ARTISTS HAVE PAINTED IT THAT WAY?
NOT NECESSARILY. A LOT OF GREAT ARTISTS WERE INSANE.
BUT... BUT HOW COULD THEY HAVE PAINTED IN COLOR ANYWAY? WOULDN'T THEIR PAINTS HAVE BEEN SHADES OF GRAY BACK THEN?
OF COURSE, BUT THEY TURNED COLORS LIKE EVERYTHING ELSE DID IN THE '30s.
SO WHY DIDN'T OLD BLACK AND WHITE PHOTOS TURN COLOR TOO?
BECAUSE THEY WERE COLOR PICTURES OF BLACK AND WHITE, REMEMBER?
THE WORLD IS A COMPLICATED PLACE, HOBBES.
WHENEVER IT SEEMS THAT WAY, I TAKE A NAP IN A TREE AND WAIT FOR DINNER.