Color II: applications in photography

CS 178, Spring 2010

Begun 5/13/10, finished 5/18, and recap slides added.

Marc Levoy
Computer Science Department
Stanford University

Outline

- spectral power distributions
- color response in animals and humans
- 3D colorspace of the human visual system
  - and color filter arrays in cameras
- 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

© Marc Levoy

Newton's color circle

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

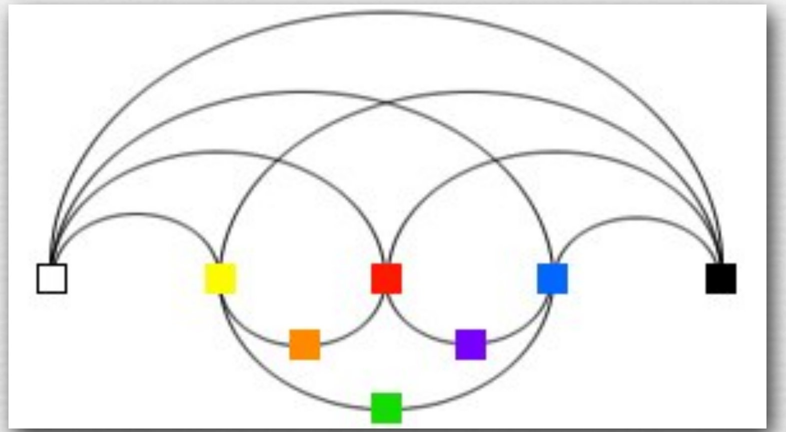

Peter Paul Rubens and
François d'Aguilon (1613)

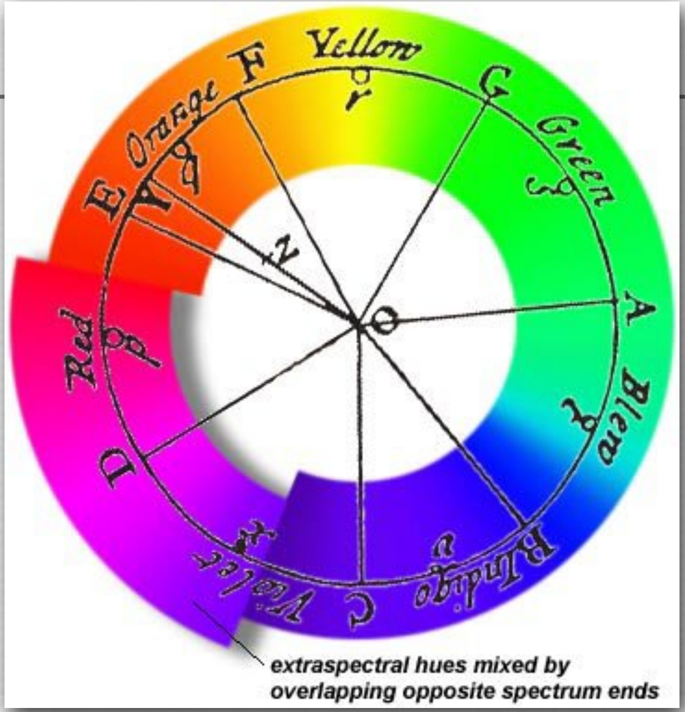

Isaac Newton (1708)

- 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

© Marc Levoy

Cylindrical color spaces (contents of whiteboard)

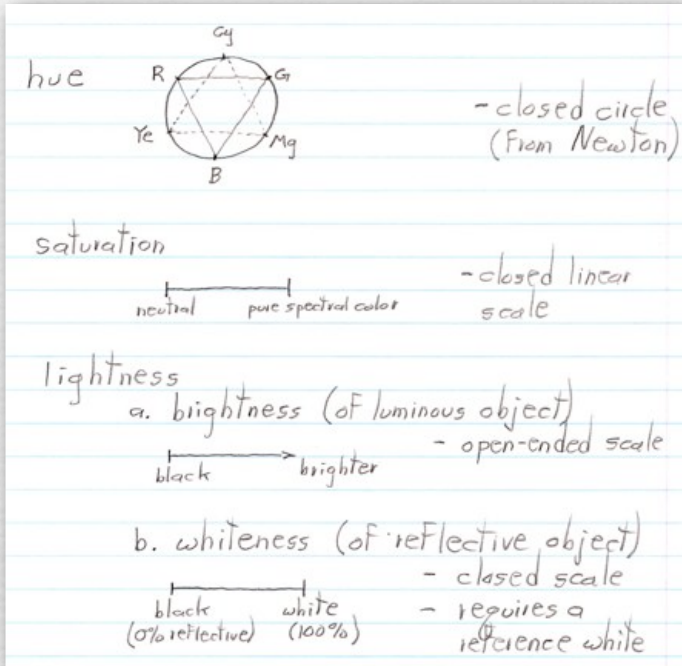

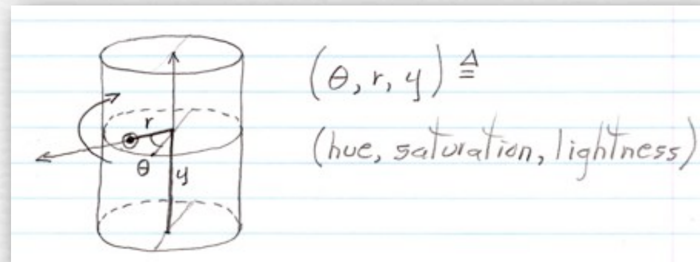

- given one circular scale and two linear scales, i.e. one angle and two lengths, the logical coordinate system is a cylindrical one
- selection of colors within such a system is easily done using 1D scales for H, S, and L, or 2D surfaces of constant H, S, or L

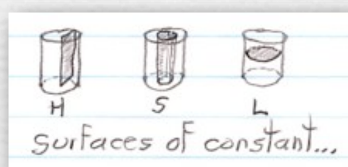

© Marc Levoy

Cylindrical color spaces

(wikipedia)

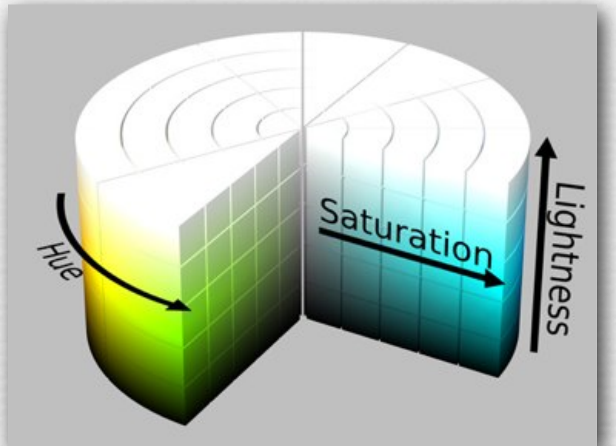

HSL cylinder

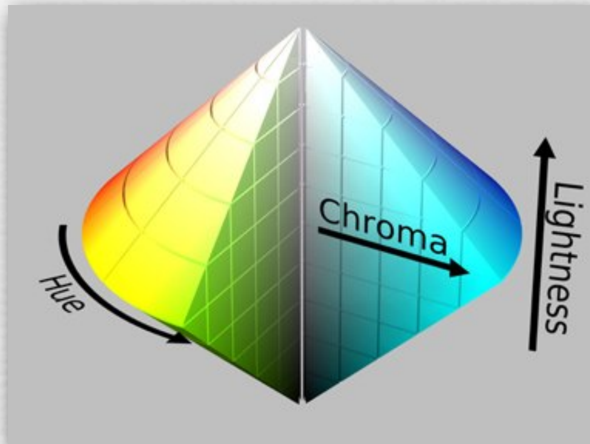

HSL double cone

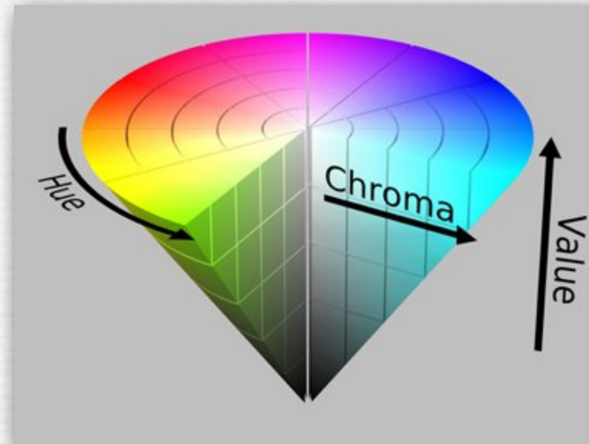

HSV single cone

- a cylinder is easy to understand, but colors near the top and bottom are indistinguishable
  - single cone solves this by compressing top & bottom to a point
- when artists mix "complementary" lights, they expect to get white, but halfway from red to cyan in HSL space is gray
  - HSV pushes the white point down onto the max-S plane
  - painters might prefer an inverted cone, with black on this plane

© Marc Levoy

Munsell color system

Albert Munsell
(1858-1918)

(wikipedia)

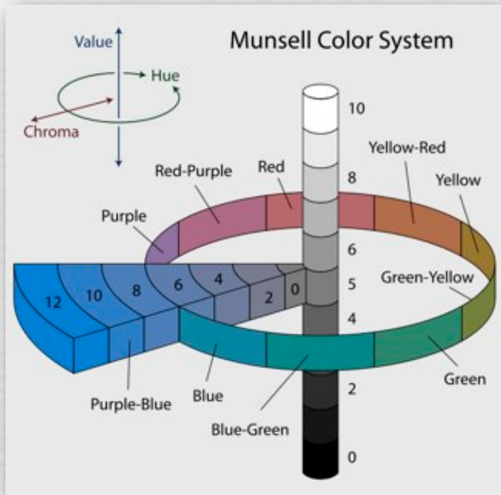

3-axis colorspace

1905 book

CG rendering of
1929 measurements

- spacing of colors is perceptually uniform (by experiment)
- outer envelope of solid determined by available inks

© Marc Levoy

A menagerie of color selectors

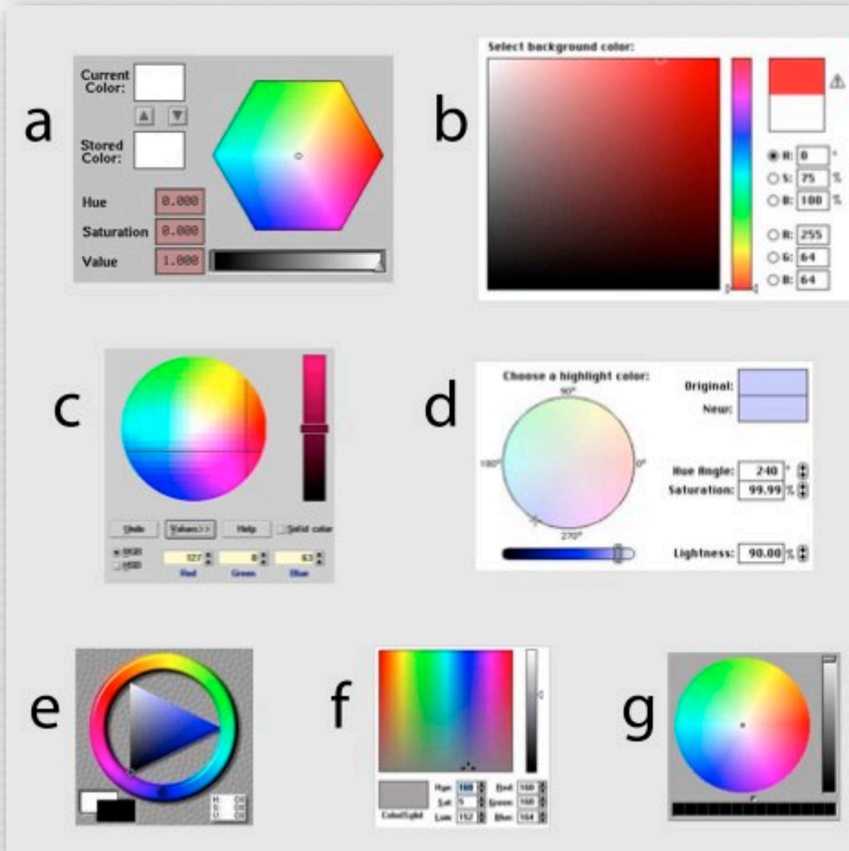

© Marc Levoy

Color selection in Photoshop

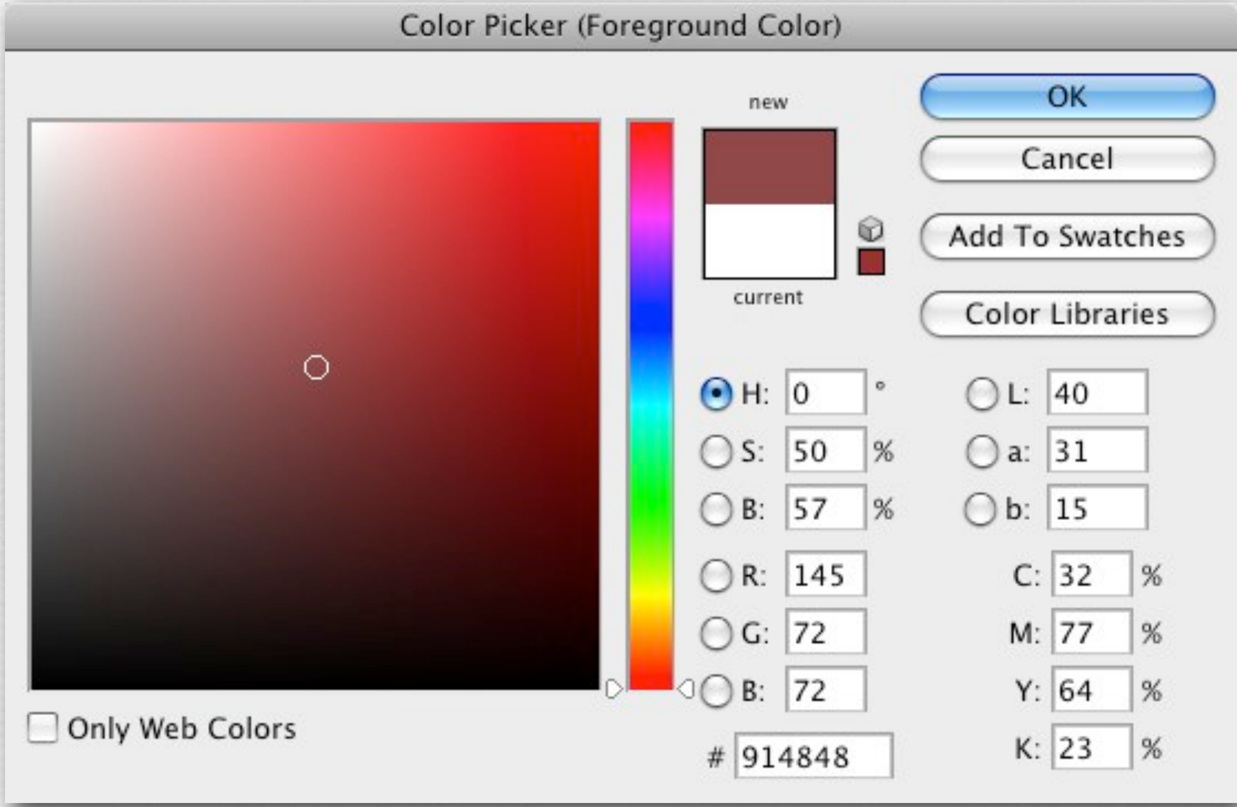

© Marc Levoy

Photoshop's color selector in HSL space (contents of whiteboard)

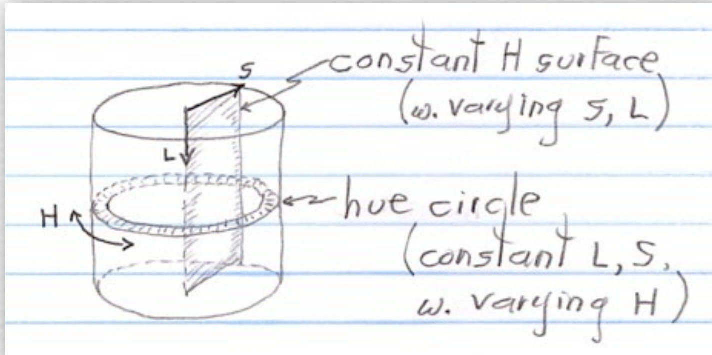

- the main rectangle in Photoshop's color selector is a 2D surface of constant hue in cylindrical color space, hence varying saturation and lightness
- the vertical rainbow to its right (in the dialog box) is a circumference along the outside surface of the cylinder, hence a 1D scale of varying hue and constant lightness and saturation

© Marc Levoy

Color selection in Photoshop

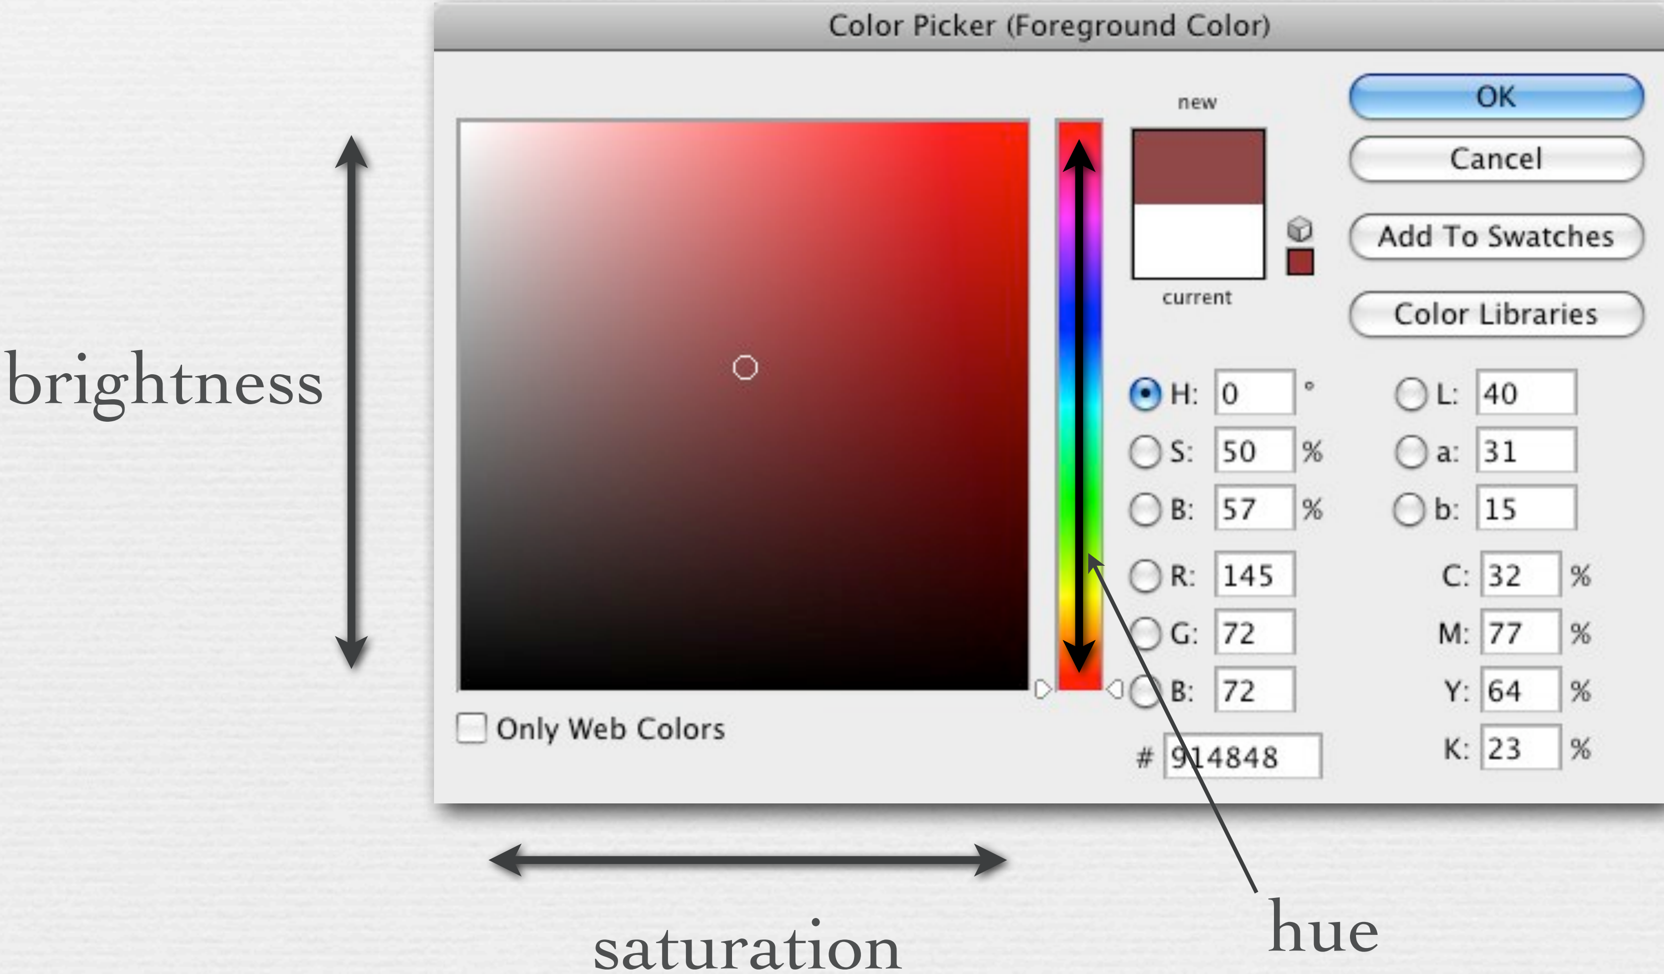

© Marc Levoy

Color selection in Photoshop

Cartesian to cylindrical coordinate conversion

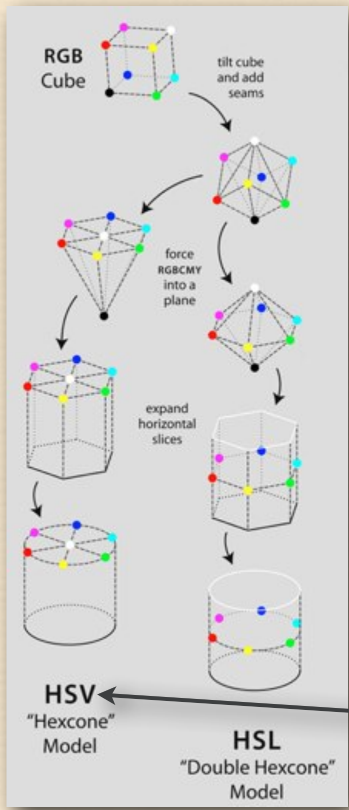

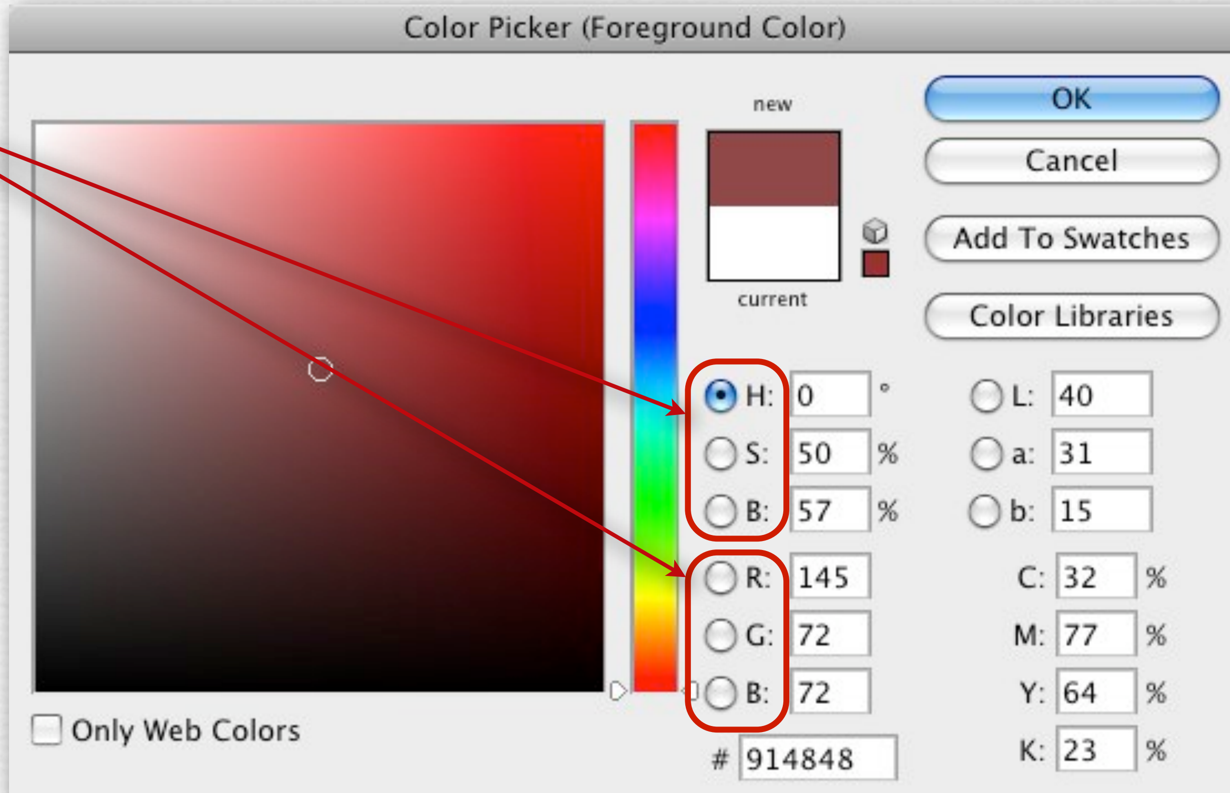

HSV ≈ HSB

© Marc Levoy

Color selection in Photoshop

3 x 3 matrix conversion

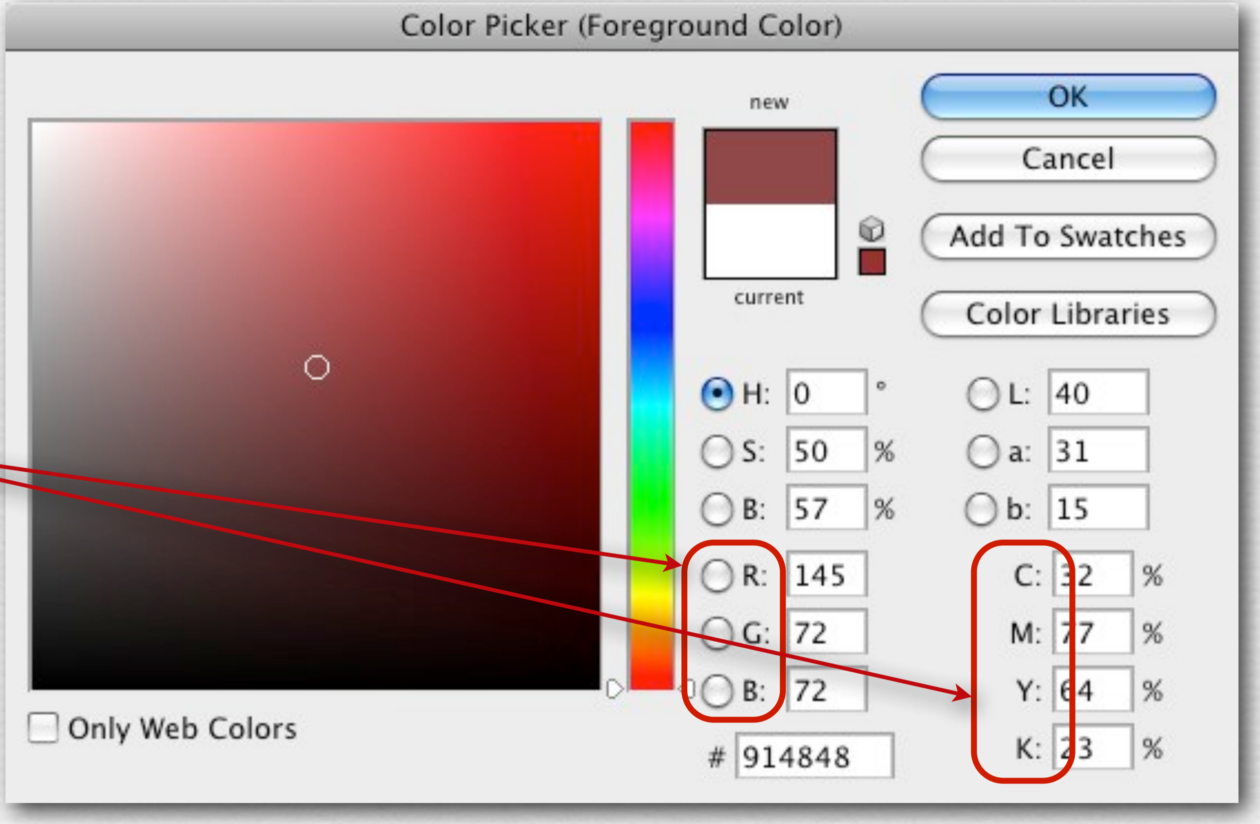

© Marc Levoy

Color selection in Photoshop

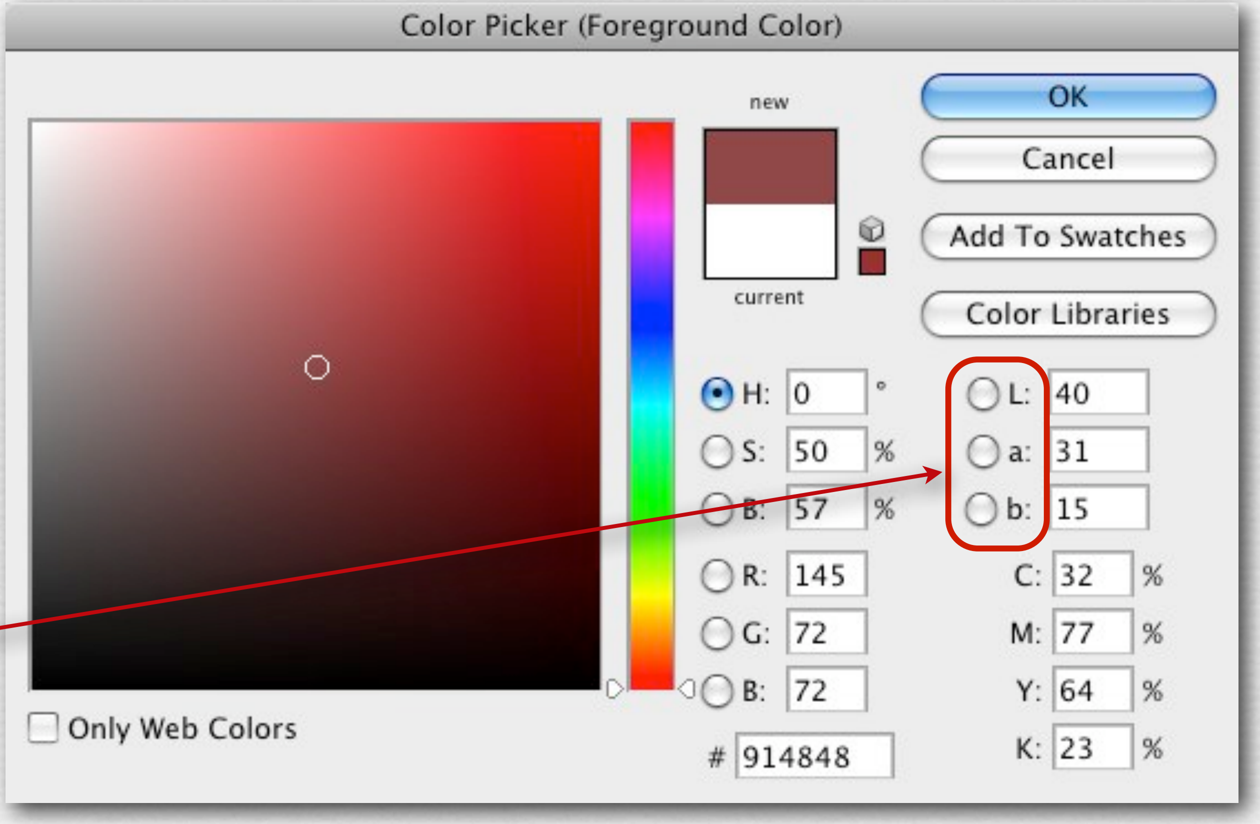

we'll cover this later in the lecture

© Marc Levoy

Recap

- *hue* is well represented by a color circle, formed from the rainbow plus mixtures of the two ends to form purples
- *saturation* is well represented by a linear scale, from neutral (black, gray, or white) to fully saturated (single wavelength)
- *lightness* is well represented by a linear scale, either open-ended if representing the brightness of luminous objects or closed-ended if representing the whiteness of reflective objects
- given one circular scale and two linear scales, the logical coordinate system is cylindrical where (H, S, L) = (θ, r, y)
- selection of colors within such a system is easily done using 1D scales for each of H, S, and L, or that in combination with 2D surfaces of constant H, S, or L

Questions?

© Marc Levoy

Outline

- ✦ spectral power distributions
- ✦ color response in animals and humans
- ✦ 3D colorspace of the human visual system
  - and color filter arrays in cameras
- ✦ 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

© Marc Levoy

Chromaticity diagrams

- 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*

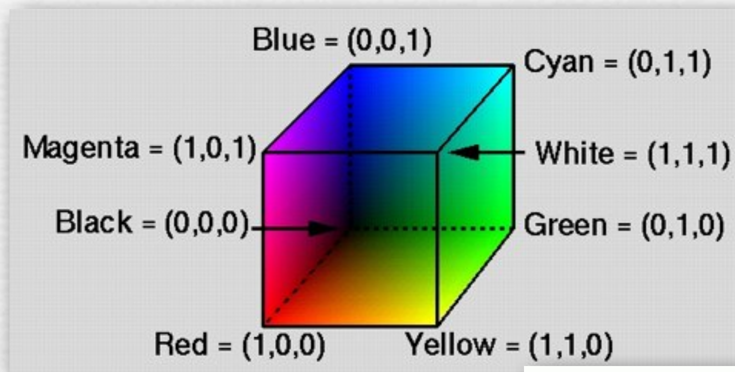

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/threedgamut.html

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

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

- 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*

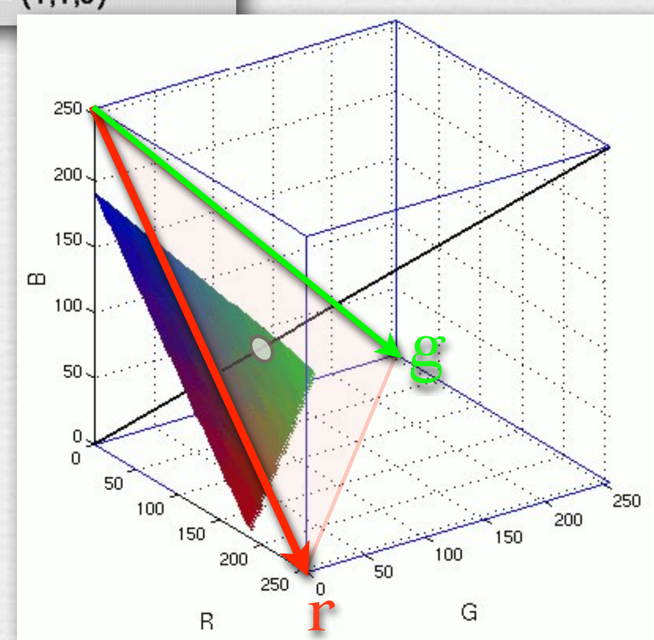

© Marc Levoy

Chromaticity diagrams

- this triangle is called the *rgb chromaticity diagram* for the chosen RGB primaries
  - mixtures of colors lie along straight lines
  - neutral (black to white) lies at (⅓, ⅓)
  - r>0, g>0 does not enclose spectral locus
- the same construction can be performed using <u>any</u> set of 3 vectors as primaries, even impossible ones ($\rho < 0$ or $\gamma < 0$ or $\beta < 0$)
- the CIE has defined a set of primaries XYZ, and the associated *xyz chromaticity diagram*
  - x>0, y>0 <u>does</u> enclose spectral locus
  - one can connect red and green on the locus with a *line of extra-spectral purples*
  - x,y is a standardized way to denote colors

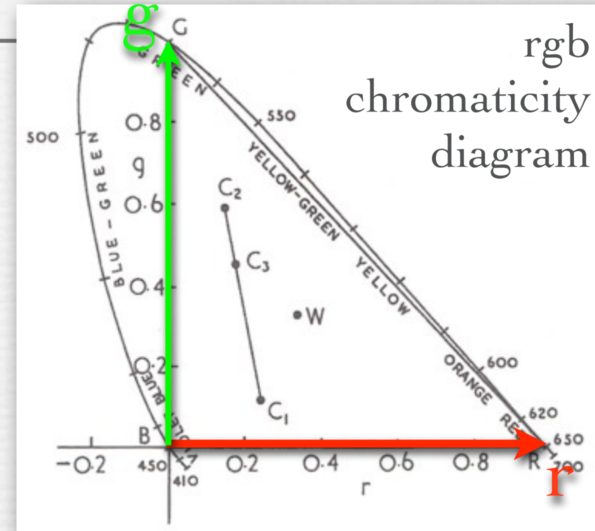

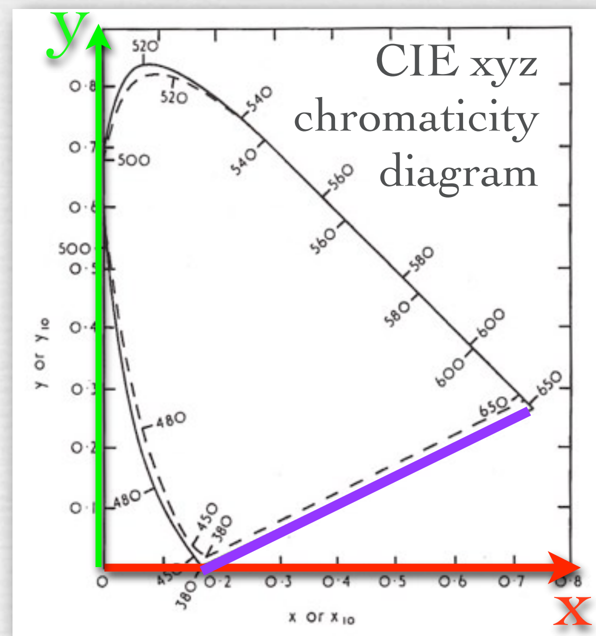

(Hunt)

© Marc Levoy

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

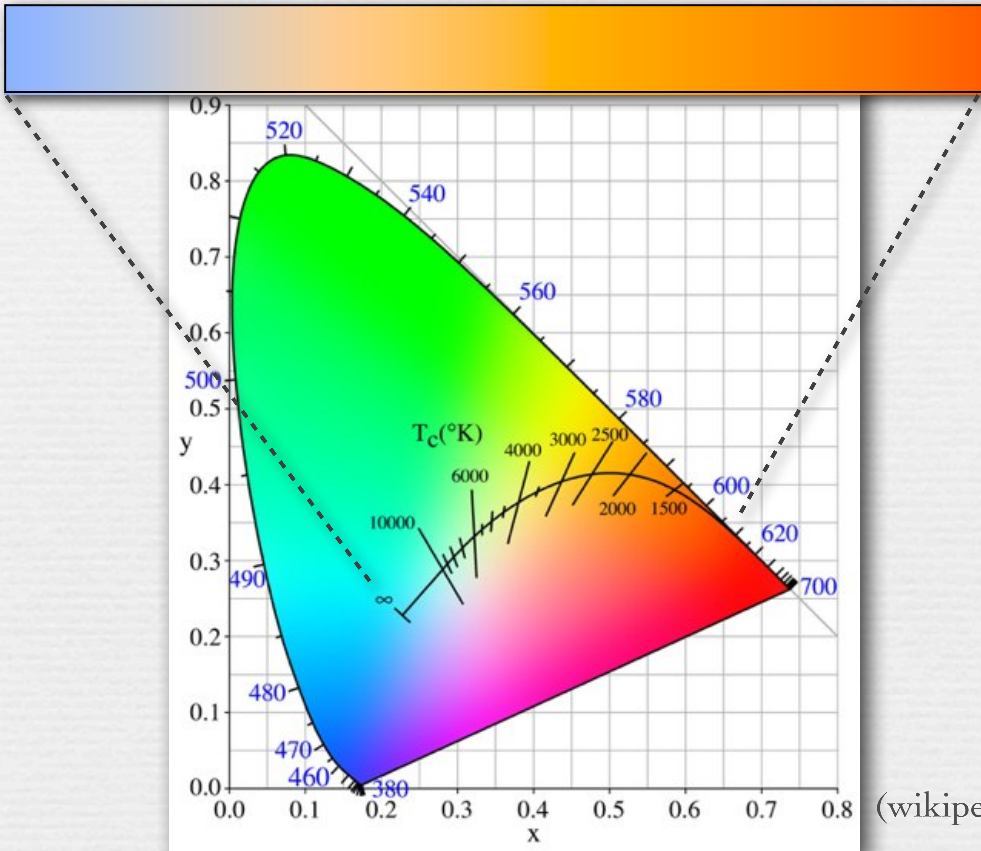

(wikipedia)

<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*

© Marc Levoy

White balancing in digital photography

- 1. choose an object in the photograph you think is neutral (somewhere between black and white) in the real world
- 2. compute scale factors (S_R,S_G,S_B) that map the object's (R,G,B) to neutral (R=G=B), i.e. $S_R = \frac{1}{3}(R+G+B) / R$, etc.
- 3. apply the same scaling to all pixels in the sensed image
- the eventual appearance of R=G=B, hence of your chosen object, depends on the color space of the camera
  - the color space of most digital cameras is sRGB
  - the reference white for sRGB is D65 (6500°K)
- thus, white balancing on an sRGB camera forces your chosen object to appear 6500°K (blueish white)
- if you trust your object to be neutral, this procedure is equivalent to finding the color temperature of the illumination

© Marc Levoy

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

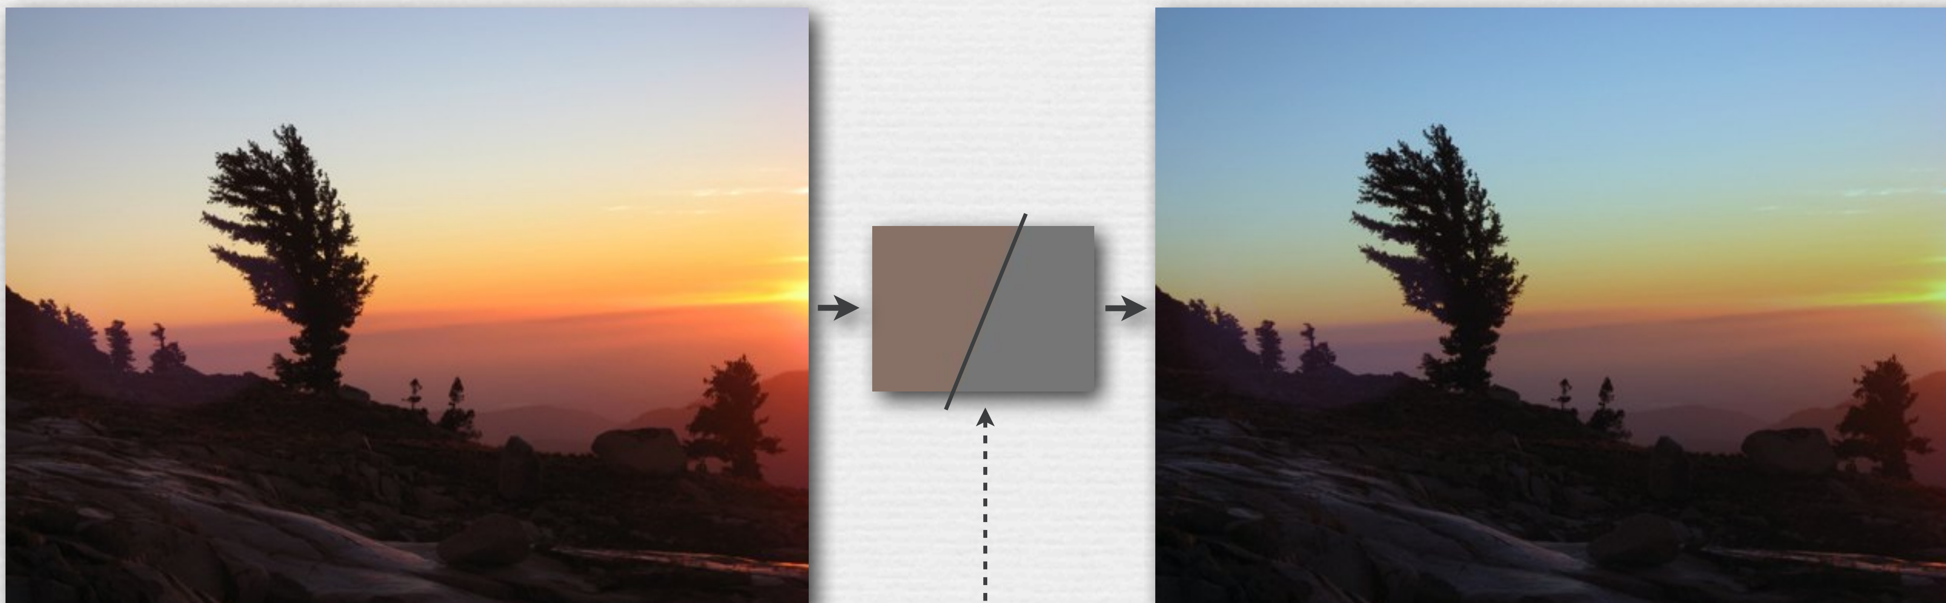

(Marc Levoy)

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

(S_R, S_G, S_B) = (0.84, 1.04, 1.15)

© Marc Levoy

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 should be white
    - fails if pixel is saturated
    - fails if object is metallic - gold has gold-colored highlights
  - find a neutral-colored object in the scene
    - but how??

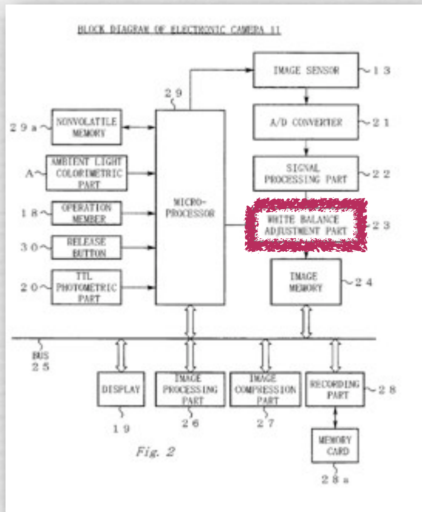

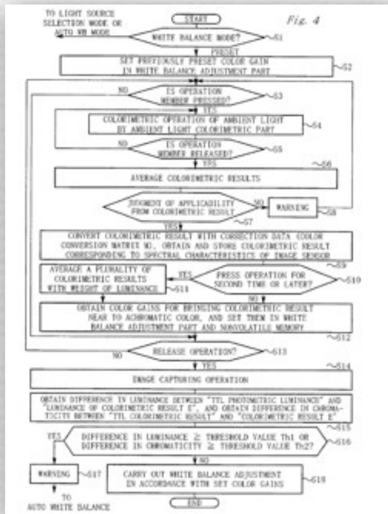

(Nikon patent)

© Marc Levoy

Finding the color temperature of the illumination

- Auto White Balance (AWB)
- manually specify the illumination's color temperature
  - each color temperature corresponds to a unique (x,y)
  - for a given camera, one can measure the (R,G,B) values recorded when a neutral object is illuminated by this (x,y)
  - compute scale factors (S_R,S_G,S_B) that map this (R,G,B) to neutral (R=G=B); apply this scaling to all pixels as before

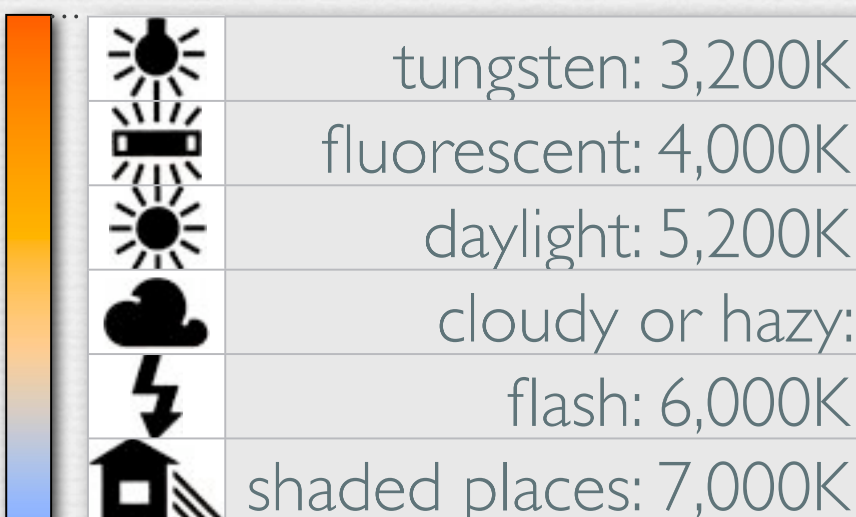

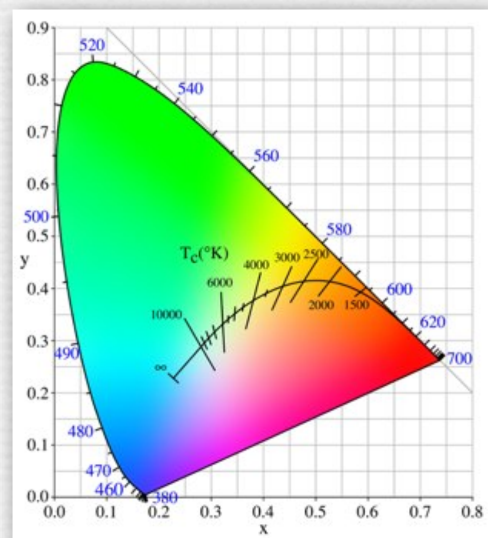

© Marc Levoy

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 very reddish, so it boosted blues
  - same thing would have happened if originally shot with tungsten WB

© Marc Levoy

Recap

- by choosing three primaries (defined by three matching functions) and a *reference white* (defined by three "hidden scales"), one defines an *RGB cube*, with black at one corner and your reference white at the opposite corner
- by projecting points in an RGB cube towards the origin (black point) and intersecting them with the R+G+B=1 plane, one factors out brightness, yielding the 2D *rgb chromaticity diagram*
- repeating this for a standard but non-physical set of primaries called XYZ, one obtains the *xyz chromaticity diagram*; in this diagram the spectral locus falls into the all-positive octant
- by identifying a feature you believe is neutral (it reflects all wavelengths equally), to the extent its RGB values are not the same, you are identifying the color of the illumination; by rescaling all pixel values until that feature is neutral, you correct for the illumination, a process called *white balancing*
- a common scale for illumination color is *correlated color temperature*, which forms a curve in the xyz chromaticity diagram

Questions?

© Marc Levoy

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

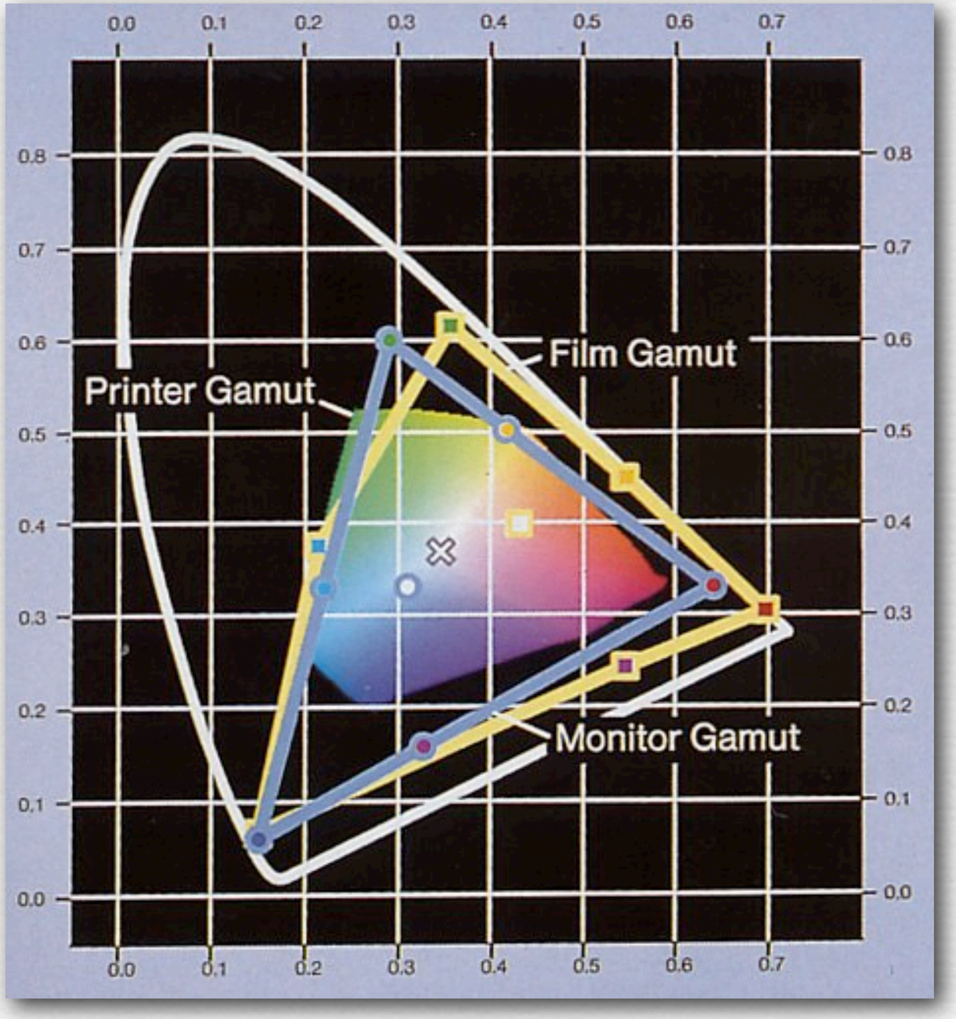

(Foley)

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

© Marc Levoy

Digitizing the paint colors at Hanna-Barbera Productions

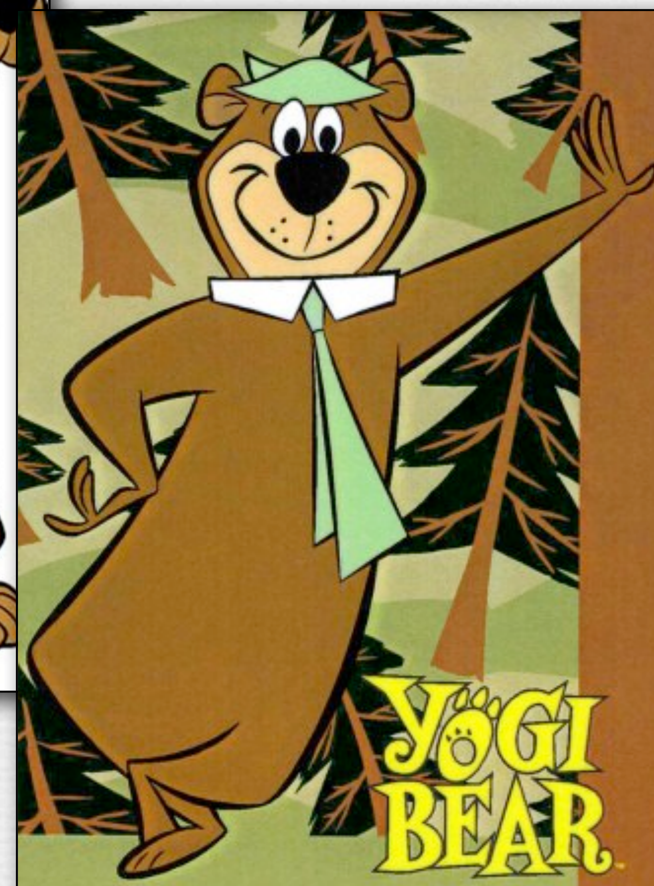

© Marc Levoy

Digitizing the paint colors at Hanna-Barbera Productions

physical color samples

spectroreflectometer

spectrum for each color

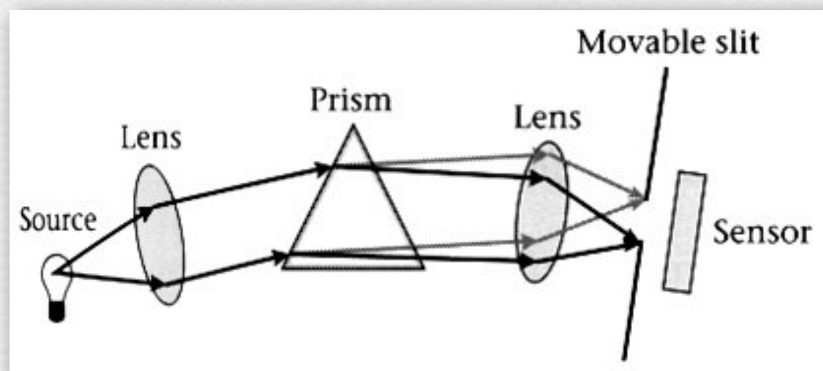

© Marc Levoy

Digitizing the paint colors at Hanna-Barbera Productions

physical color samples

↓ spectroreflectometer

spectrum for each color

↓ CIE matching functions

XYZ coordinates

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

© Marc Levoy

Digitizing the paint colors at Hanna-Barbera Productions

physical color samples

↓ spectroreflectometer

spectrum for each color

↓ CIE matching functions

XYZ coordinates

↓ projection onto X=Y=Z=1 plane

xy chromaticity coordinates

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

© Marc Levoy

Digitizing the paint colors at Hanna-Barbera Productions

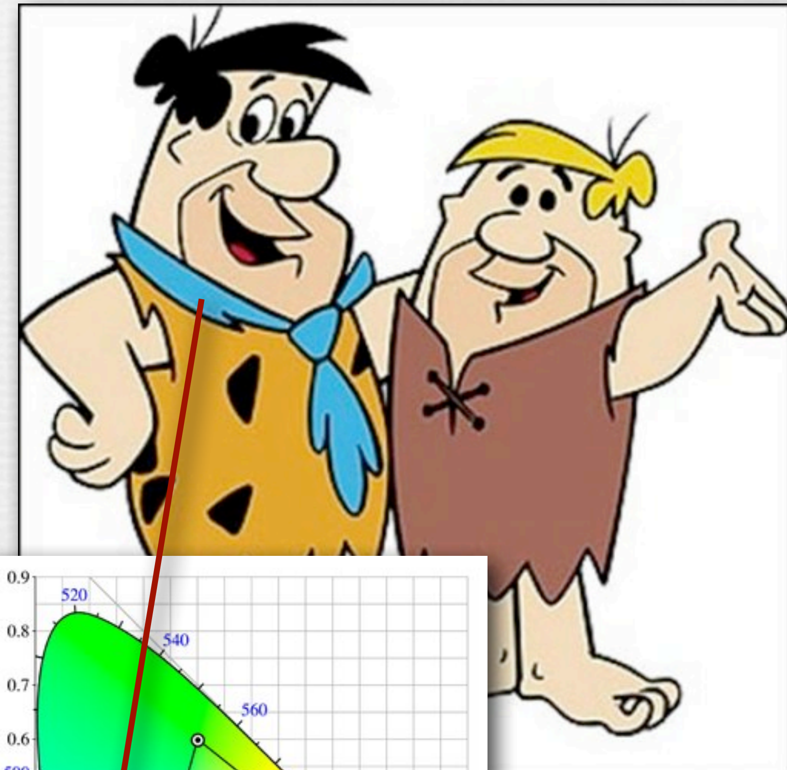

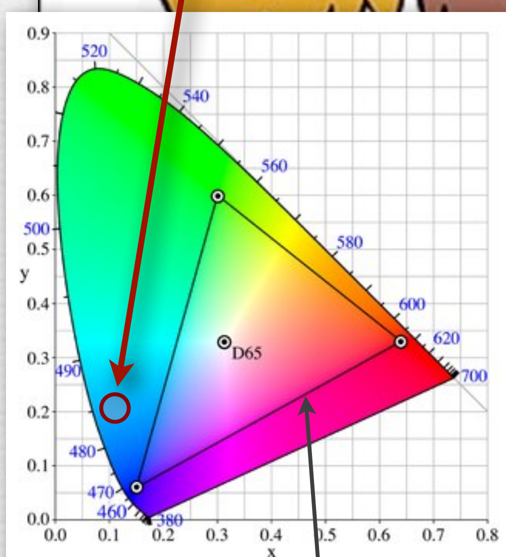

NTSC gamut

physical color samples

spectroreflectometer

spectrum for each color

CIE matching functions

XYZ coordinates

projection onto X=Y=Z=1 plane

xy chromaticity coordinates

DANGER: NECKTIE OUT OF GAMUT !!

© Marc Levoy

Uniform perceptual color spaces

equally perceivable
“MacAdam ellipses”

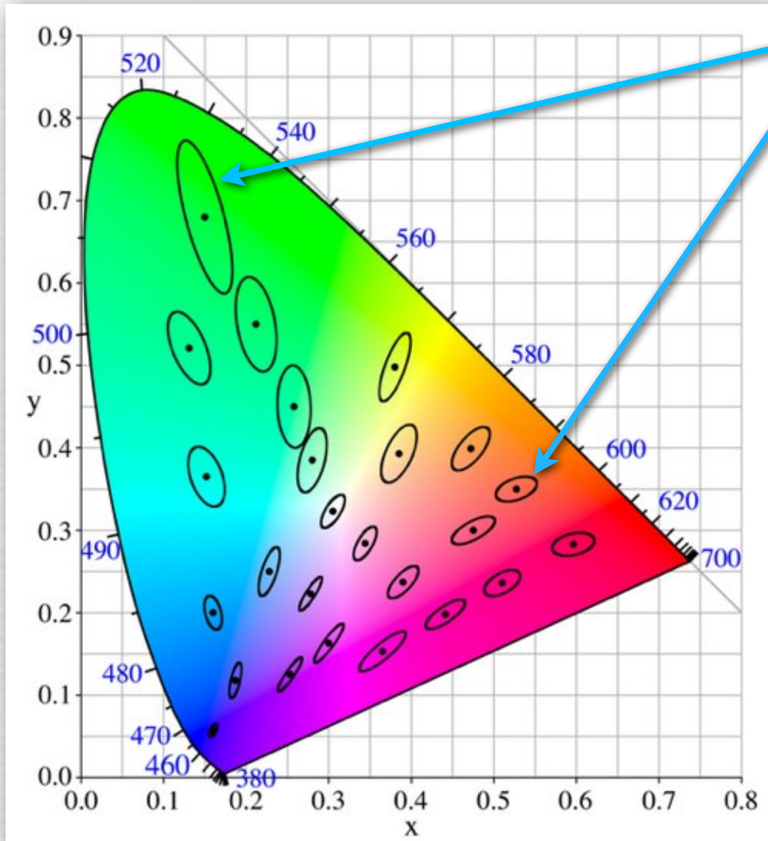

(wikipedia)

(Wyszecki and Stiles)

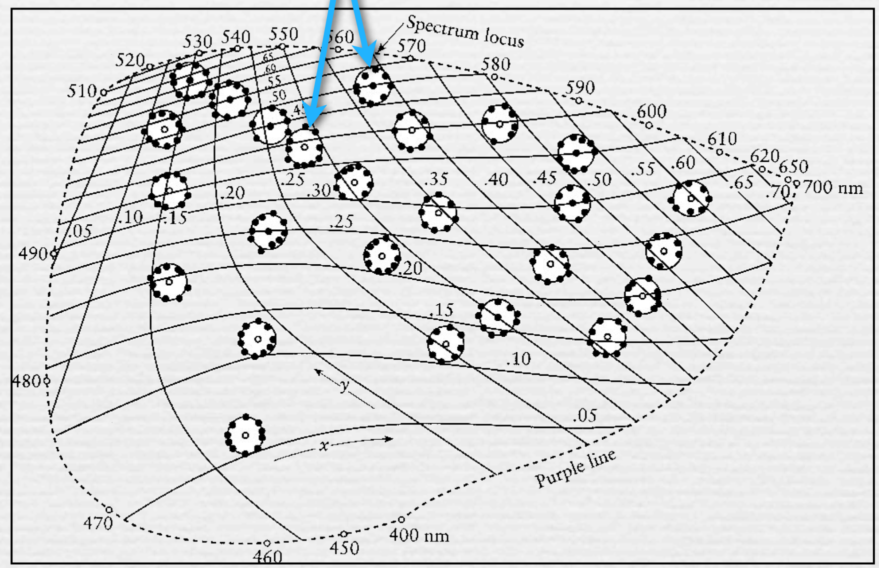

a non-linear mapping

- 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 xy diagram) non-linearly

© Marc Levoy

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

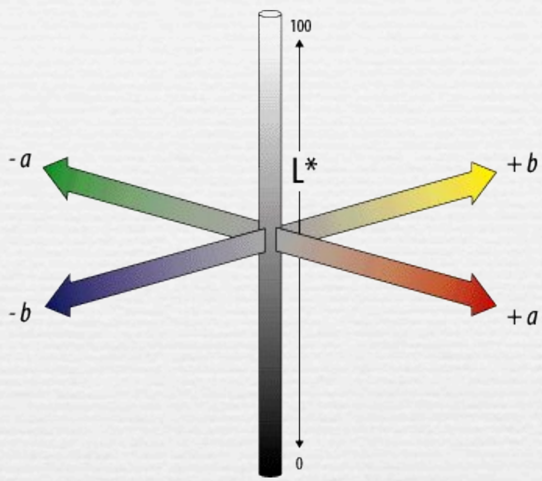

$$L = 25\left(100\frac{Y}{Y_0}\right)^{1/3} - 16$$

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

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

non-linear mapping (a gamma transform)

- L* is lightness
- a* and b* are color-opponent pairs
  - a* is red-green, and b* is blue-yellow
- gamma transform is because for humans, perceived brightness ∝ scene intensity$^{\gamma}$, where $\gamma \approx ⅓$
  - similar to Weber-Fechner Law: dB = k dI/I, so B = k ln(I/I_0)

© Marc Levoy

Complementary colors

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

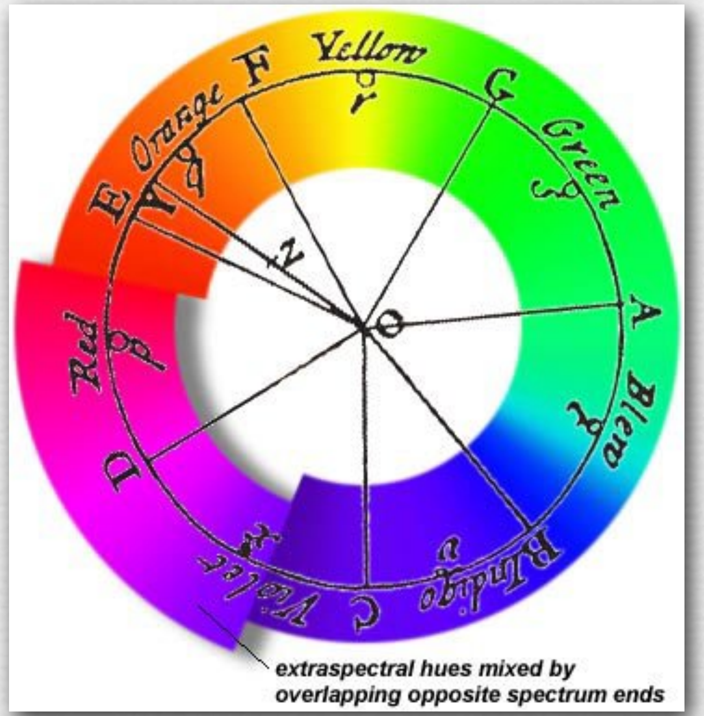

- Leonardo described complementarity of certain pairs of colors
- Newton arranged them opposite one another across his circle
- Comte de Buffon (1707-1788) observed that afterimage colors were exactly the complementary colors

© Marc Levoy

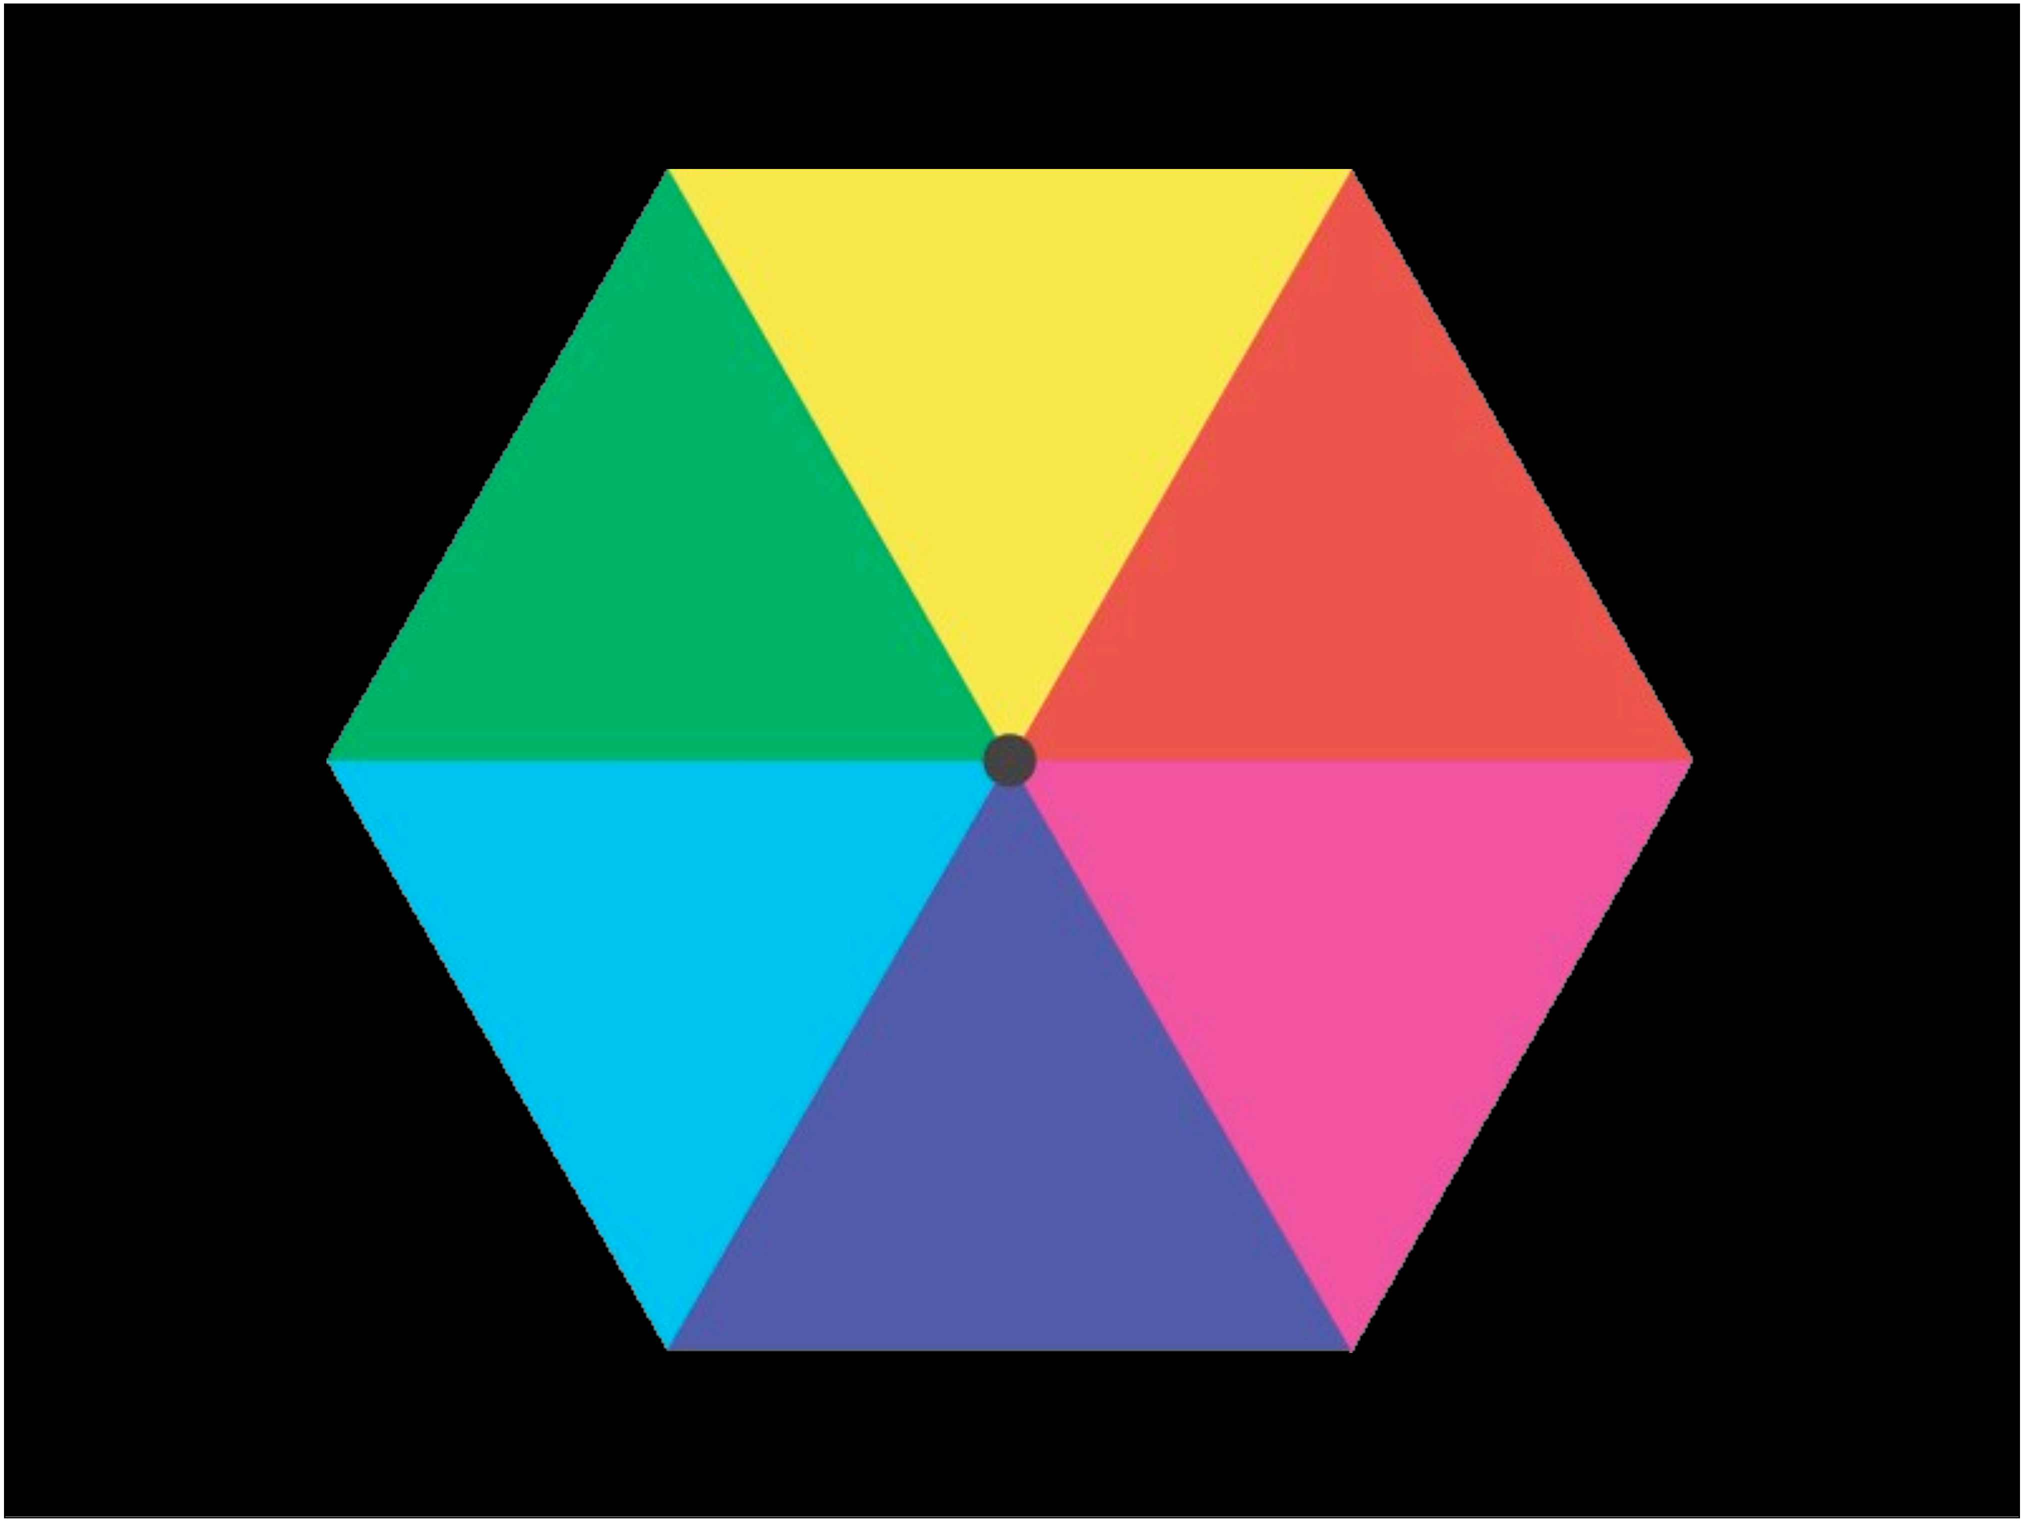

To get the effect, stare at the N-2 slide for 30 seconds, fixating on the gray dot in the middle of the pattern, then without looking at anything else, advance to the N-1 slide. What do you see? You should see the afterimage shown at right below. Each color is the compliment (opponent) of the corresponding color on the left below.

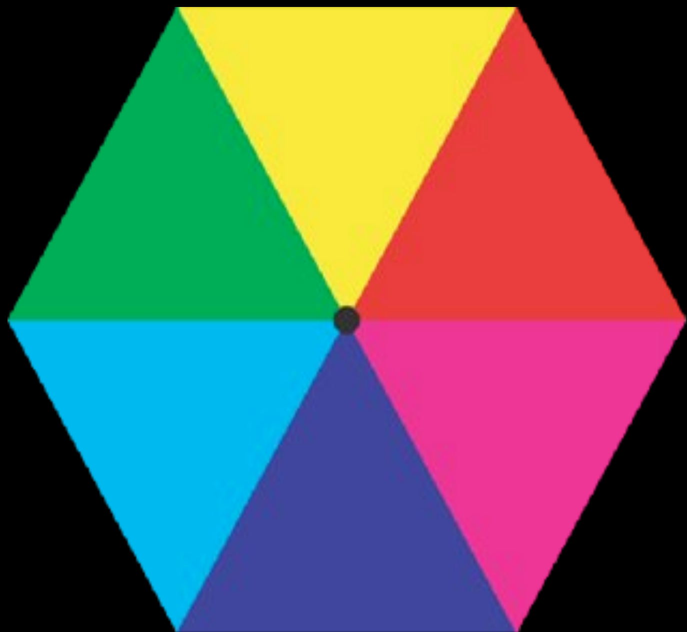

image

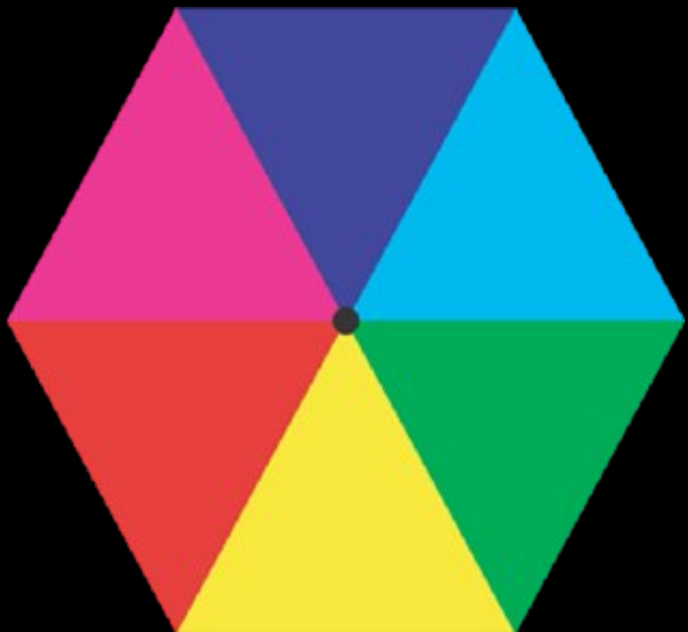

afterimage

Opponent colors

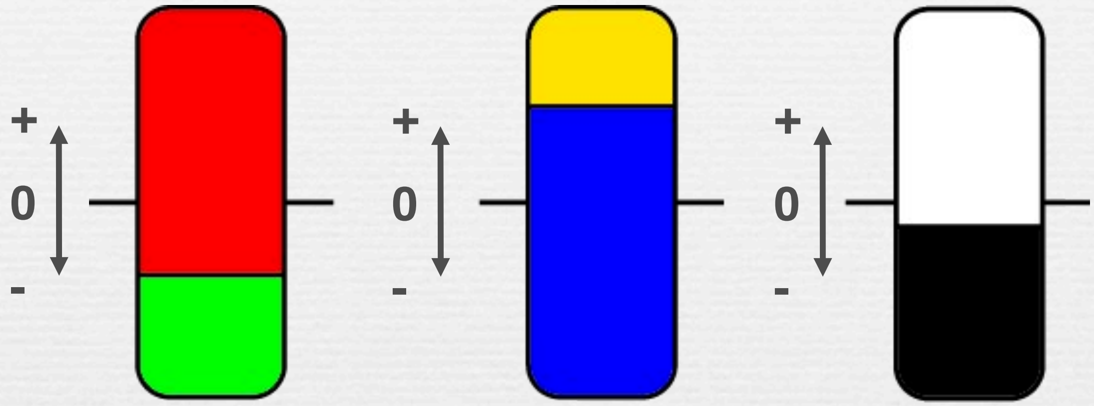

red/green receptors

blue/yellow receptors

black/white receptors

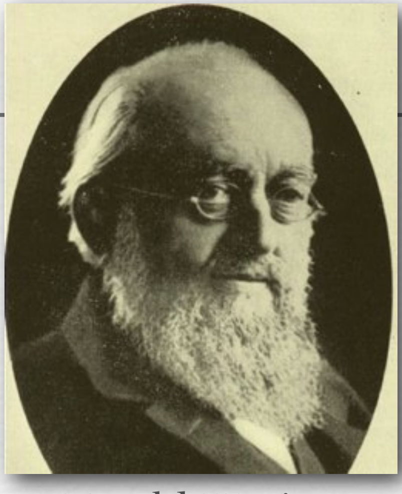

Ewald Hering
(1834-1918)

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

© Marc Levoy

Opponent colors wiring

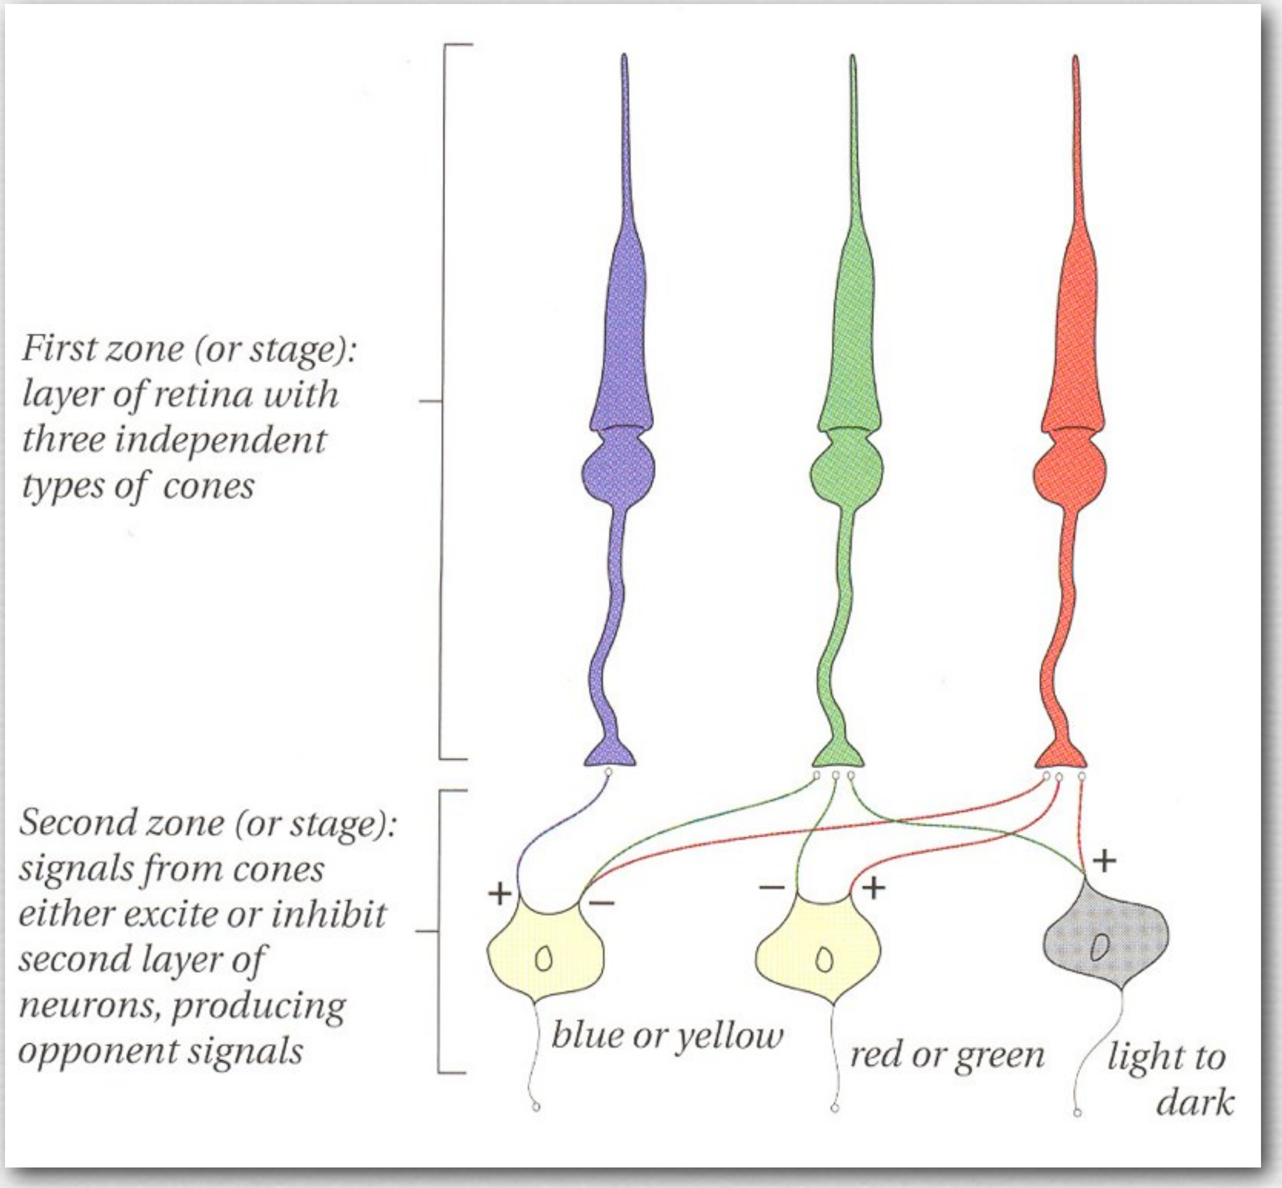

© Marc Levoy

Practical use of opponent colors: NTSC color television

(wikipedia)

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

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

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

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

RGB & YIQ are axes in (ρ, γ, β) space, hence these transforms are 3×3 matrix multiplications

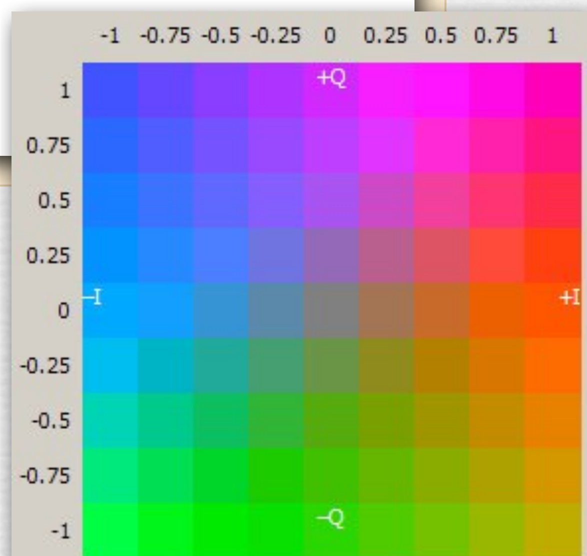

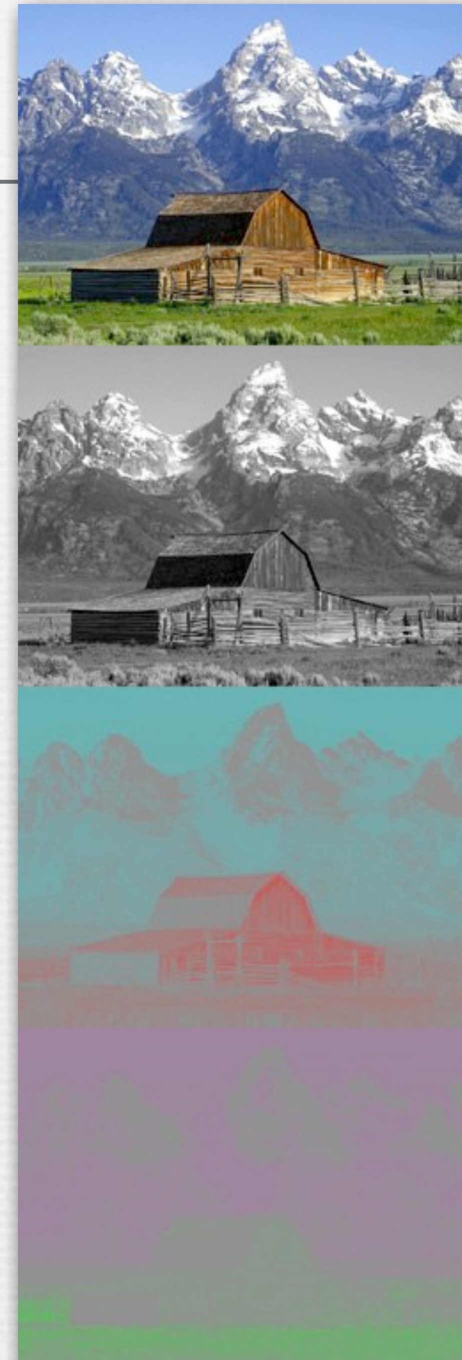

© Marc Levoy

Practical use of opponent colors: JPEG image compression

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

$$
\begin{aligned}
Y' &= 16 + (65.481 \cdot R' + 128.553 \cdot G' + 24.966 \cdot B') \\
C_B &= 128 + (-37.797 \cdot R' - 74.203 \cdot G' + 112.0 \cdot B') \\
C_R &= 128 + (112.0 \cdot R' - 93.786 \cdot G' - 18.214 \cdot B')
\end{aligned}
$$

I replaced the above set of equations after class, to keep the notation consistent.

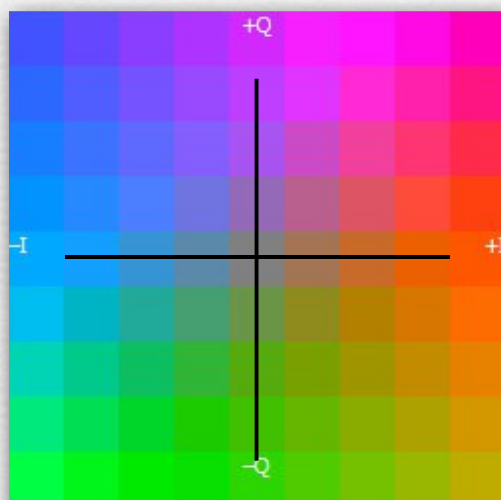

(wikipedia)

Y'

Cb

Cr

© Marc Levoy

Practical use of opponent colors: JPEG compression

(wikipedia)

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

we are more sensitive to high frequencies in Y than CbCr, so use more bits for Y (~2×)

$$
\begin{aligned}
Y' &= 16 + (65.481 \cdot R' + 128.553 \cdot G' + 24.966 \cdot B') \\
C_B &= 128 + (-37.797 \cdot R' - 74.203 \cdot G' + 112.0 \cdot B') \\
C_R &= 128 + (112.0 \cdot R' - 93.786 \cdot G' - 18.214 \cdot B')
\end{aligned}
$$

inputs R’, G’, B’ are R^γ, G^γ, B^γ for some gamma $\gamma < 1$

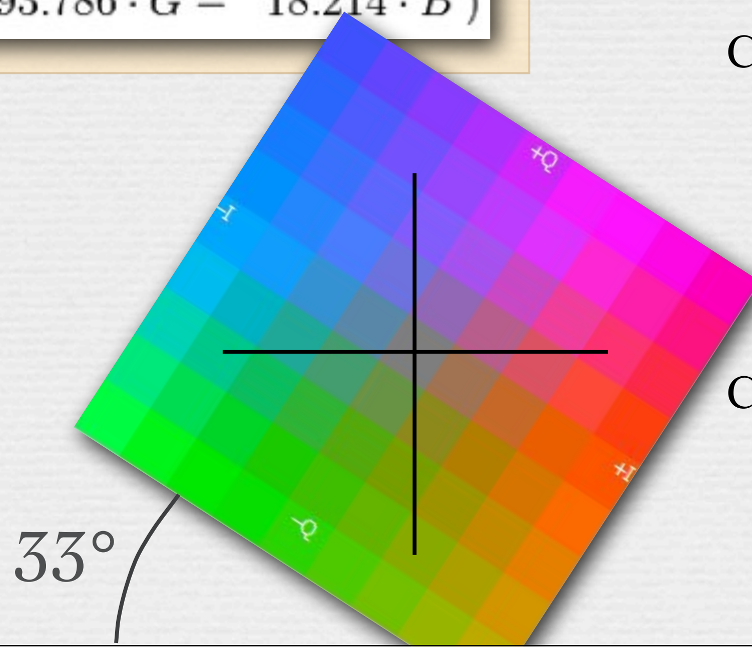

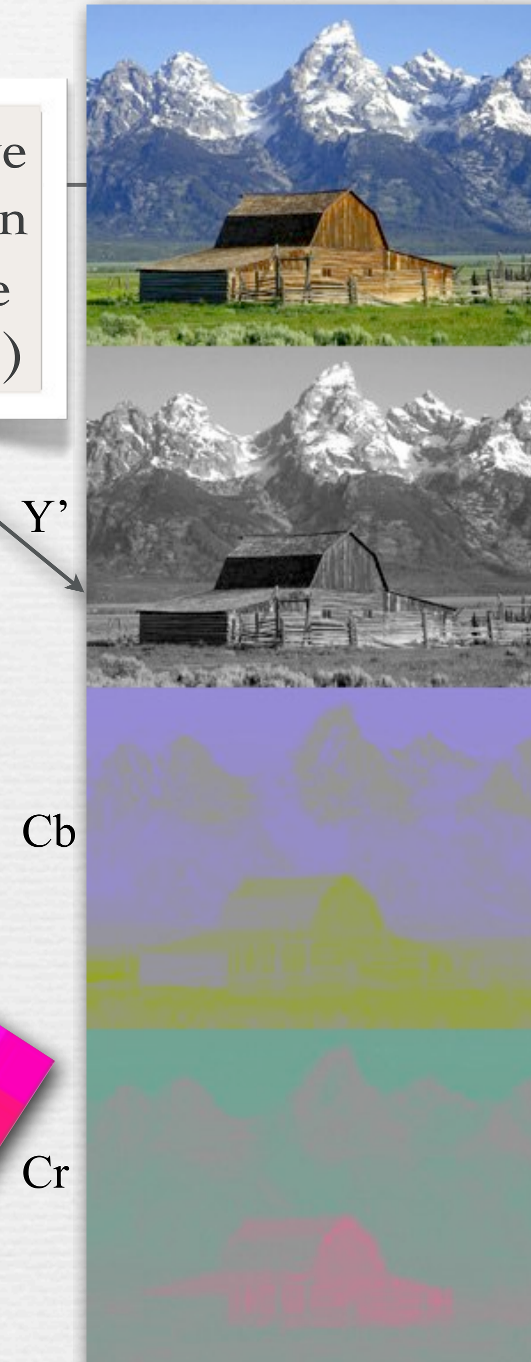

© Marc Levoy

Apparent spatial sharpness depends mainly on luminance, not chrominance

original image

(Wandell)

© Marc Levoy

Apparent spatial sharpness depends mainly on luminance, not chrominance

red-green
channel (Cr)
blurred

(Wandell)

© Marc Levoy

Apparent spatial sharpness depends mainly on luminance, not chrominance

original image

(Wandell)

© Marc Levoy

Apparent spatial sharpness depends mainly on luminance, not chrominance

blue-yellow
channel (Cb)
blurred

(Wandell)

© Marc Levoy

Apparent spatial sharpness depends mainly on luminance, not chrominance

original image

(Wandell)

Y’

Cb

Cr

© Marc Levoy

Apparent spatial sharpness depends mainly on luminance, not chrominance

luminance channel (Y') blurred

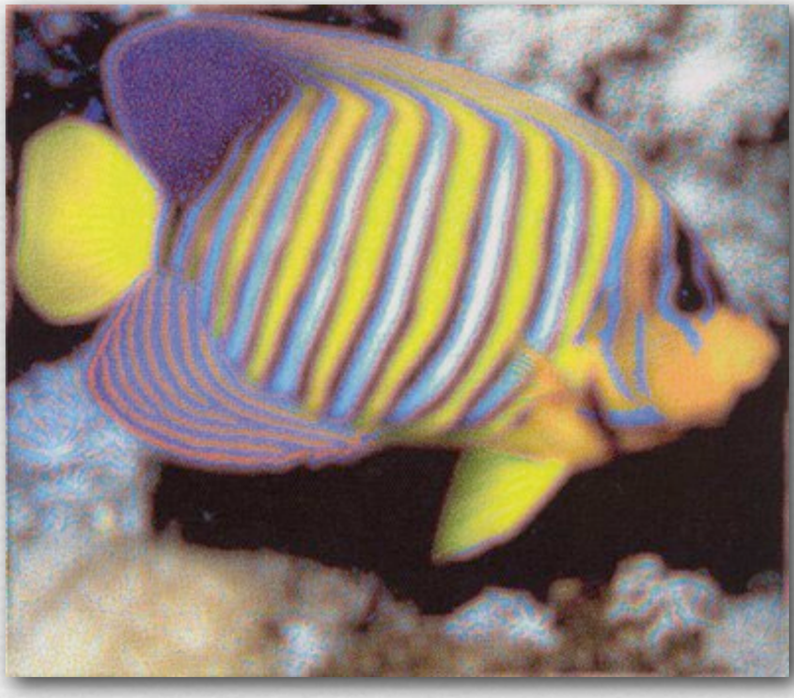

(Wandell)

© Marc Levoy

The color spaces used in cameras

- to define an RGB color space, one needs
  - the location of the R,G,B axes in (ρ, γ, β) space, i.e. what color are the 3 primaries?
  - the location of the R=G=B=1 point in (ρ, γ, β) space, i.e. what is the reference white?
  - these locations can be given in X,Y,Z coordinates, or x,y and max luminance

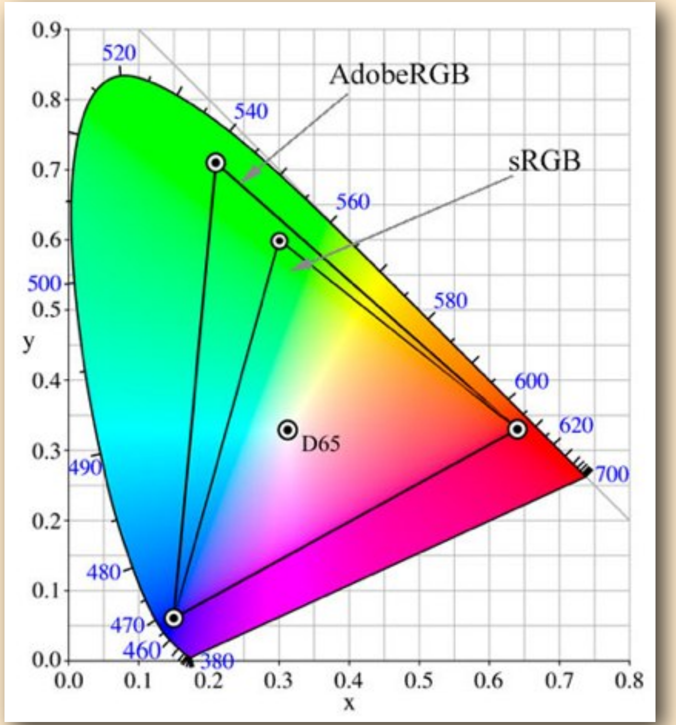

- the mapping from the RGB space to (ρ, γ, β) may be
  - a linear transformation (i.e. 3 x 3 matrix) or a non-linear mapping (like L*a*b*)
  - sRGB and Adobe RGB use a non-linear mapping, but are not perceptually uniform

Not responsible on exams for orange-tinted material

© Marc Levoy

Back to gamut mapping

(now in a perceptually uniform space)

input color space
(like sRGB)

non-linear mapping ↓

perceptually
uniform space
(like L*A*b*)

gamut mapping ↓

reduced gamut

non-linear mapping ↓

output color space
(like CMYK)

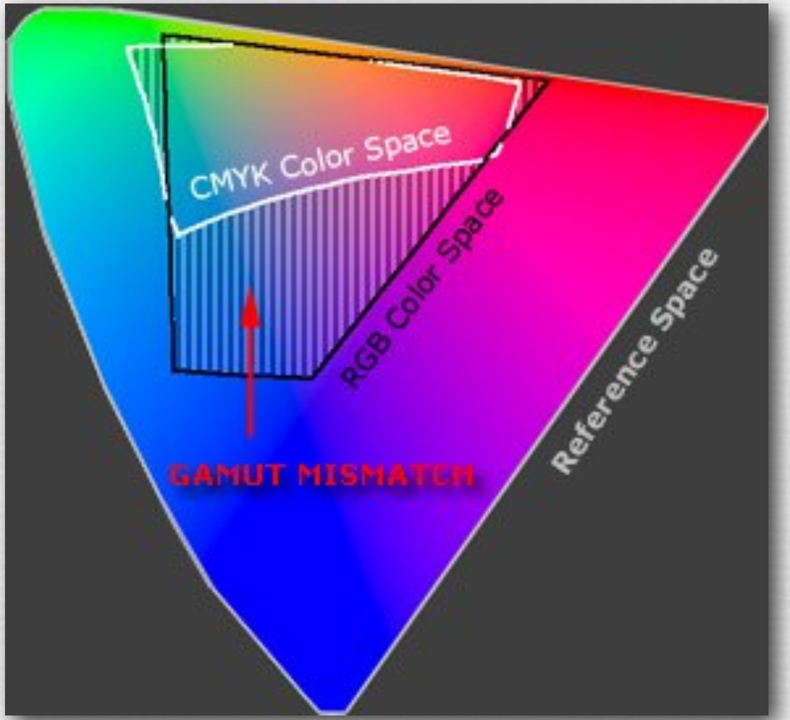

(cambridgeincolour.com)

© Marc Levoy

Rendering intents

you can do this explicitly in Photoshop, or you can let the printer do it for you

- called "color space conversion options" in Photoshop
  - *relative colorimetric* - shrinks only out-of-gamut colors, towards N
  - *absolute colorimetric* - same but shrinks to nearest point on gamut
  - *perceptual* - smoothly shrinks all colors to fit in target gamut
  - *saturated* - sacrifices smoothness to maintain saturated colors

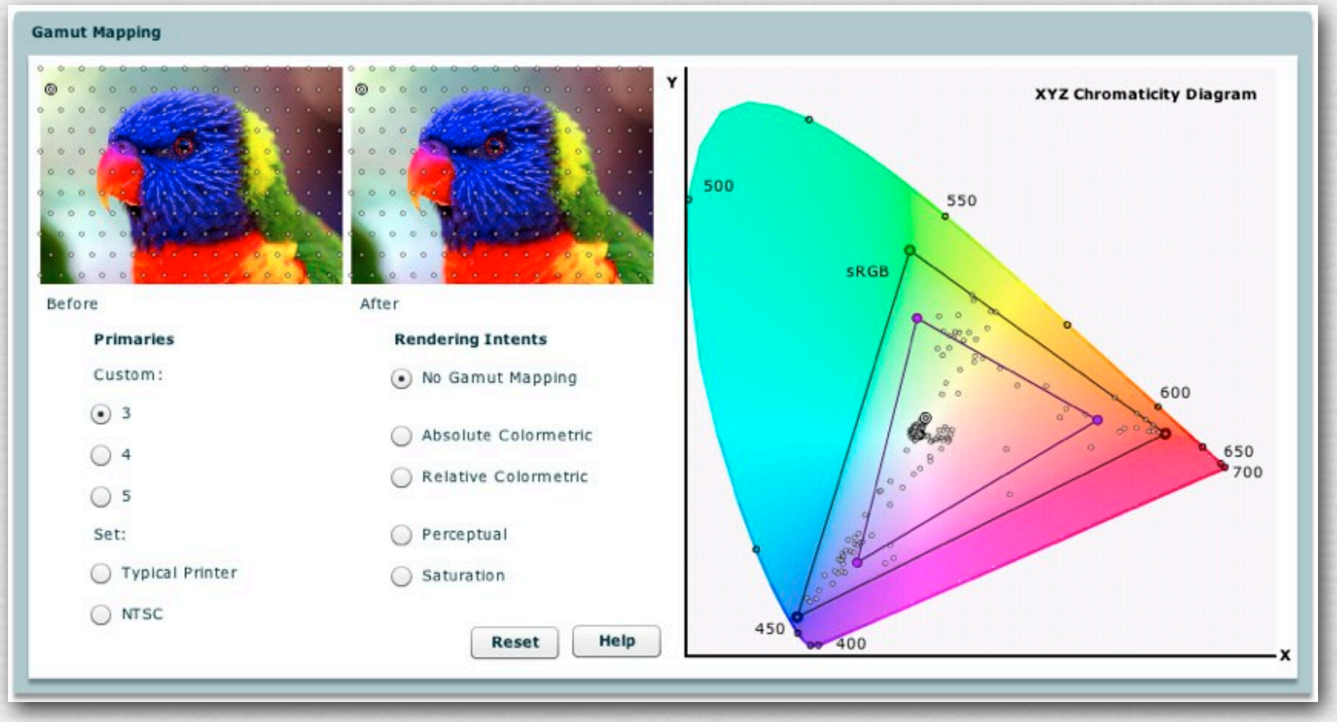

(FLASH DEMO)

http://graphics.stanford.edu/courses/cs178/applets/gamutmapping.html

© Marc Levoy

Color spaces and color management

- 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!

© Marc Levoy

Recap

In class I forgot to explain the first point below adequately. You can think of a gamut as the triangle through the middle of an RGB cube, i.e. the lightly shaded triangle in the bottom-right figure of slide 16, now drawn on the chromaticity diagram.

- the R+G+B=1 surface of a practical reproduction system (e.g. a display or printer) forms a triangle in the xyz chromaticity diagram, or more complicated figure if more than 3 primaries; the boundaries of this figure is the *gamut* for this system
- if a color to be reproduced falls outside the gamut of a target system, it must be replaced by a color lying inside the gamut, perhaps replacing other colors in the image at the same time to maintain color relationships; this is called *gamut mapping*
- gamut mapping can be performed manually (e.g. in Photoshop) or automatically by display or printer software, typically in a *perceptually uniform colorspace* like L*a*b*; how you perform the mapping is governed by a *rendering intent*, four of which are conventionally defined

The four rendering intents are defined in prose only. How each one is translated to a mathematical mapping is left up to the implementers of color management systems,. In other words, Photoshop may do "relative colorimetric" gamut mapping differently than your printer does.

Questions?

© Marc Levoy

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.

© Marc Levoy

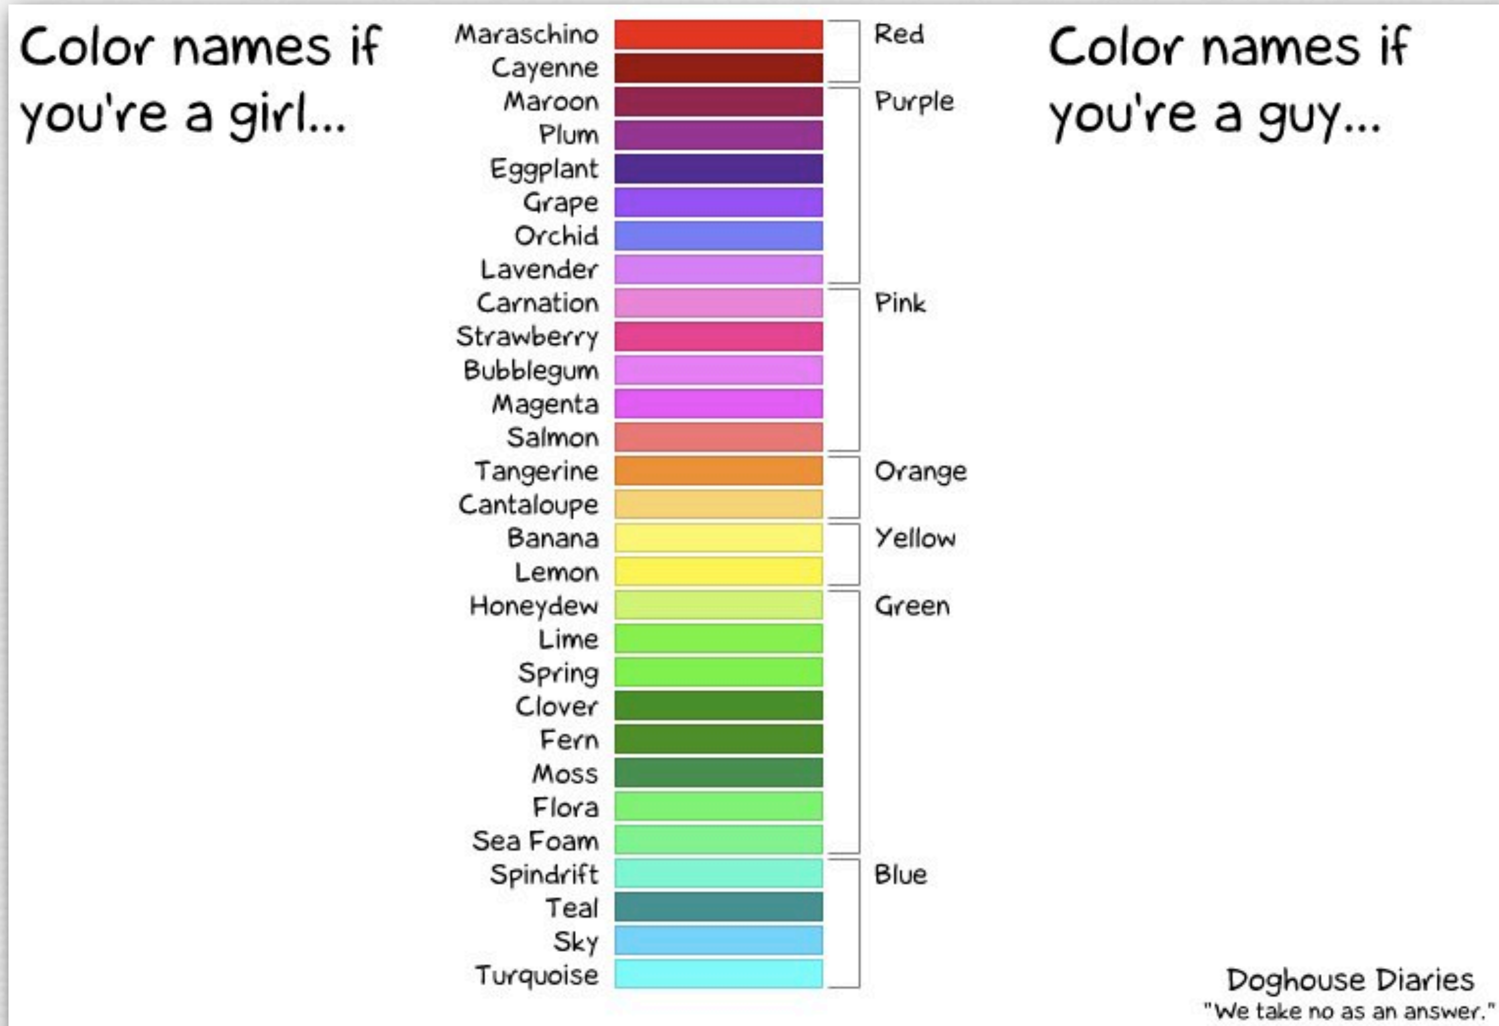
Color names if you're a girl...
Maraschino
Cayenne
Maroon
Plum
Eggplant
Grape
Orchid
Lavender
Carnation
Strawberry
Bubblegum
Magenta
Salmon
Tangerine
Cantaloupe
Banana
Lemon
Honeydew
Lime
Spring
Clover
Fern
Moss
Flora
Sea Foam
Spindrift
Teal
Sky
Turquoise
Red
Purple
Pink
Orange
Yellow
Green
Blue
Color names if you're a guy...
Doghouse Diaries
"We take no as an answer."


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.