Color II, and CFA interpolation

Bill Freeman and Fredo Durand
MIT EECS 6.098/6.882

Feb. 16, 2006

“Could you remind people that I'll be conducting an SLR intro tomorrow during my office hours (2:30) “ --Fredo

Internal summary

- What are colors?
  - Arise from power spectrum of light.
- How represent colors:
  - Pick primaries
  - Measure color matching functions (CMF's)
  - Matrix mult power spectrum by CMF's to find color as the 3 primary color values.
- How share color descriptions between people?
  - Translate colors between systems of primaries
  - Standardize on a few sets of primaries.

Color matching experiment

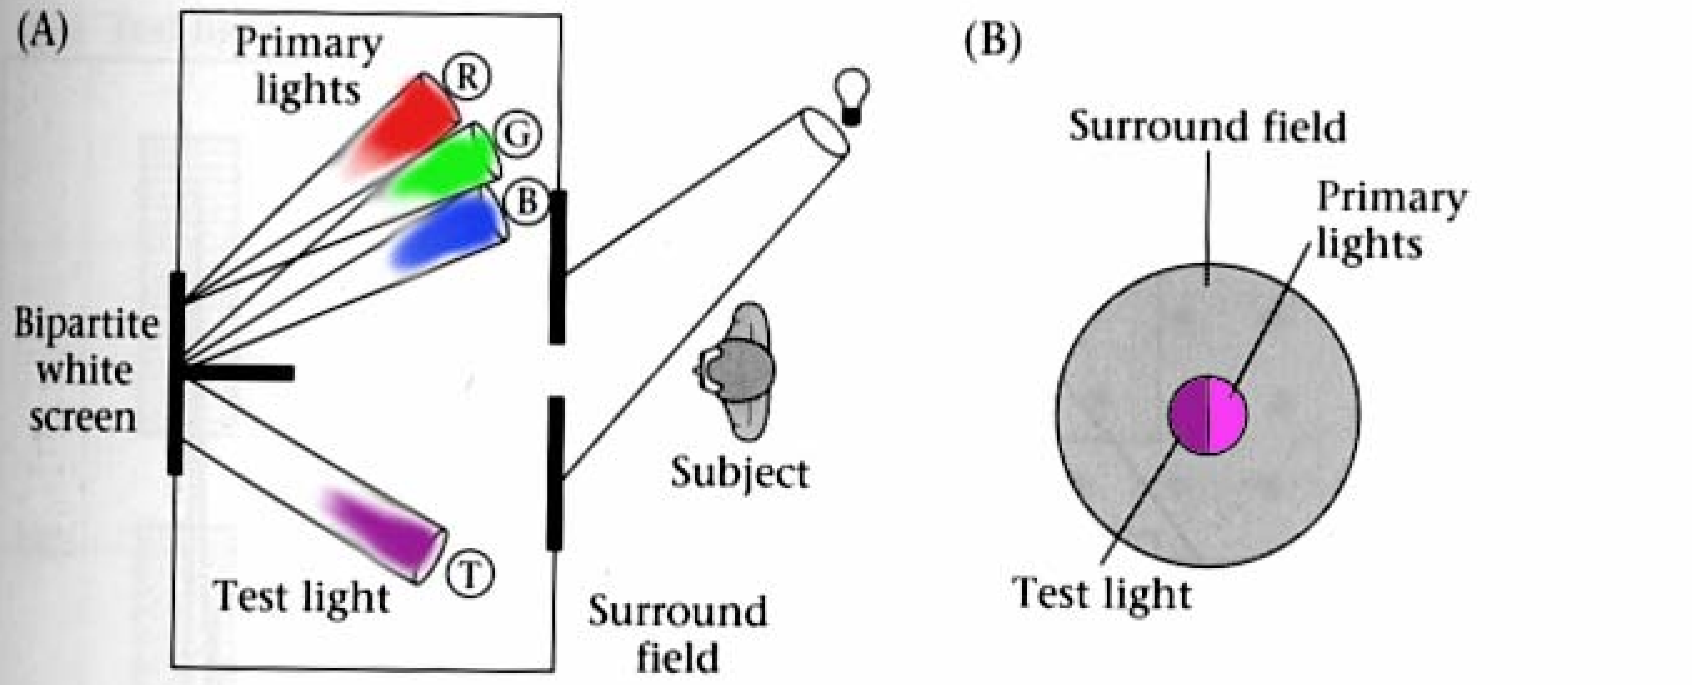

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Color matching functions for a particular set of monochromatic primaries

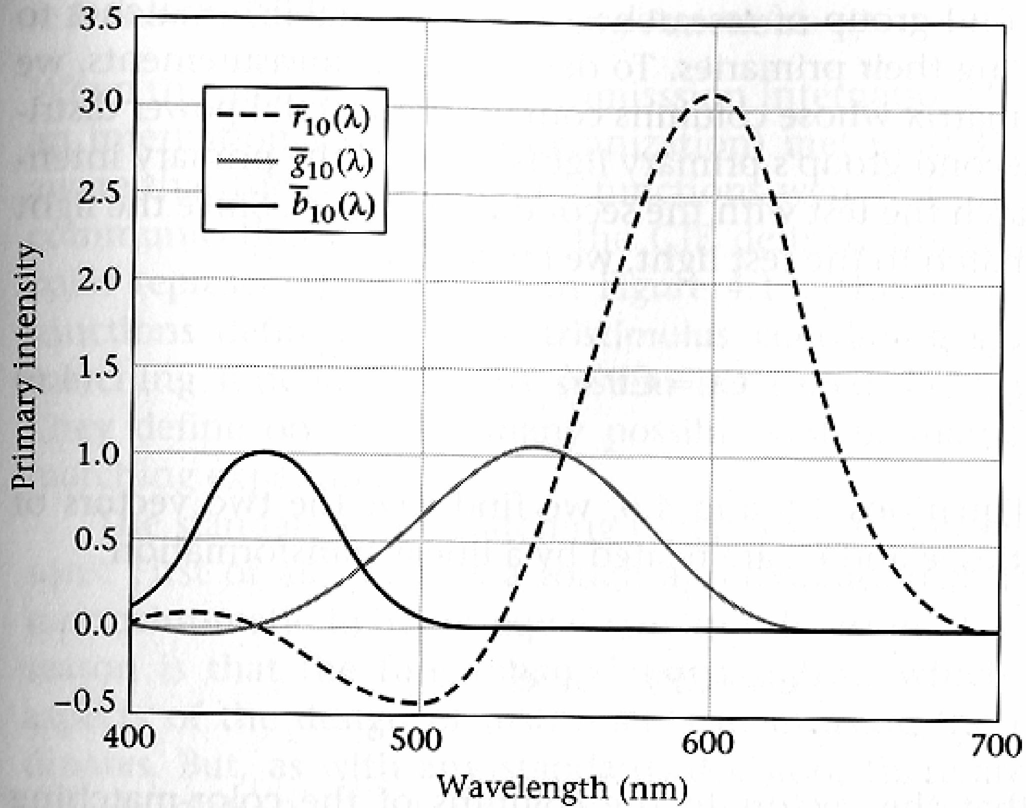

p_1 = 645.2 nm

p_2 = 525.3 nm

p_3 = 444.4 nm

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this figure are called $\bar{r}_{10}(\lambda)$, $\bar{g}_{10}(\lambda)$, and $\bar{b}_{10}(\lambda)$.

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Suppose you invent a new color display…

Given C and P

And some new set of primaries, P’

How do you find C’ ?

3 useful facts

(1) Translate color values from the primed set to the unprimed set by the matrix CP'

(2) Color matching functions, C, translate to the primed system by some 3x3 matrix R.

(3) CP = **1**, the identity matrix.

How to find the color matching functions for new primaries, P’

1 = C’ P’

1 = (R C) P’

so

R = (C P’)$^{-1}$

and

C’ = (C P’)$^{-1}$ C

This also tells you conditions on P’: CP’ must be of full rank in order to be invertible.

Color metamerism: different spectra looking the same color

Two spectra, t and s, perceptually match when

$$C\vec{t} = C\vec{s}$$

where C are the color matching functions for some set of primaries.

Graphically,

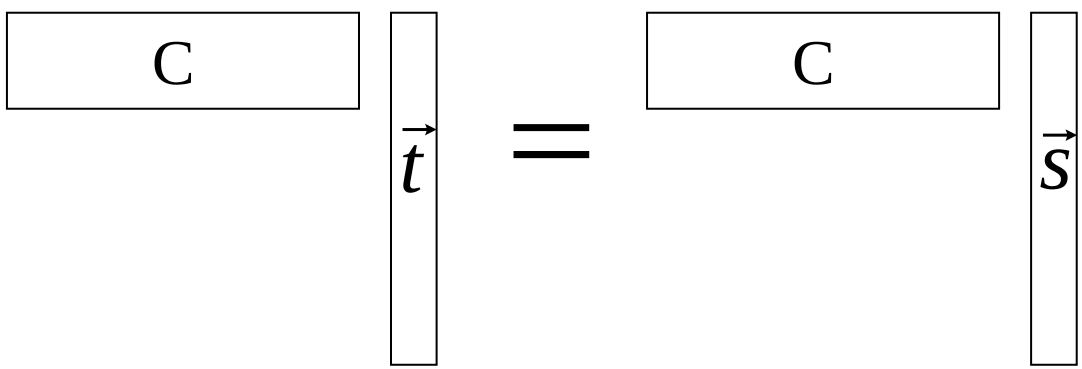

Metameric lights

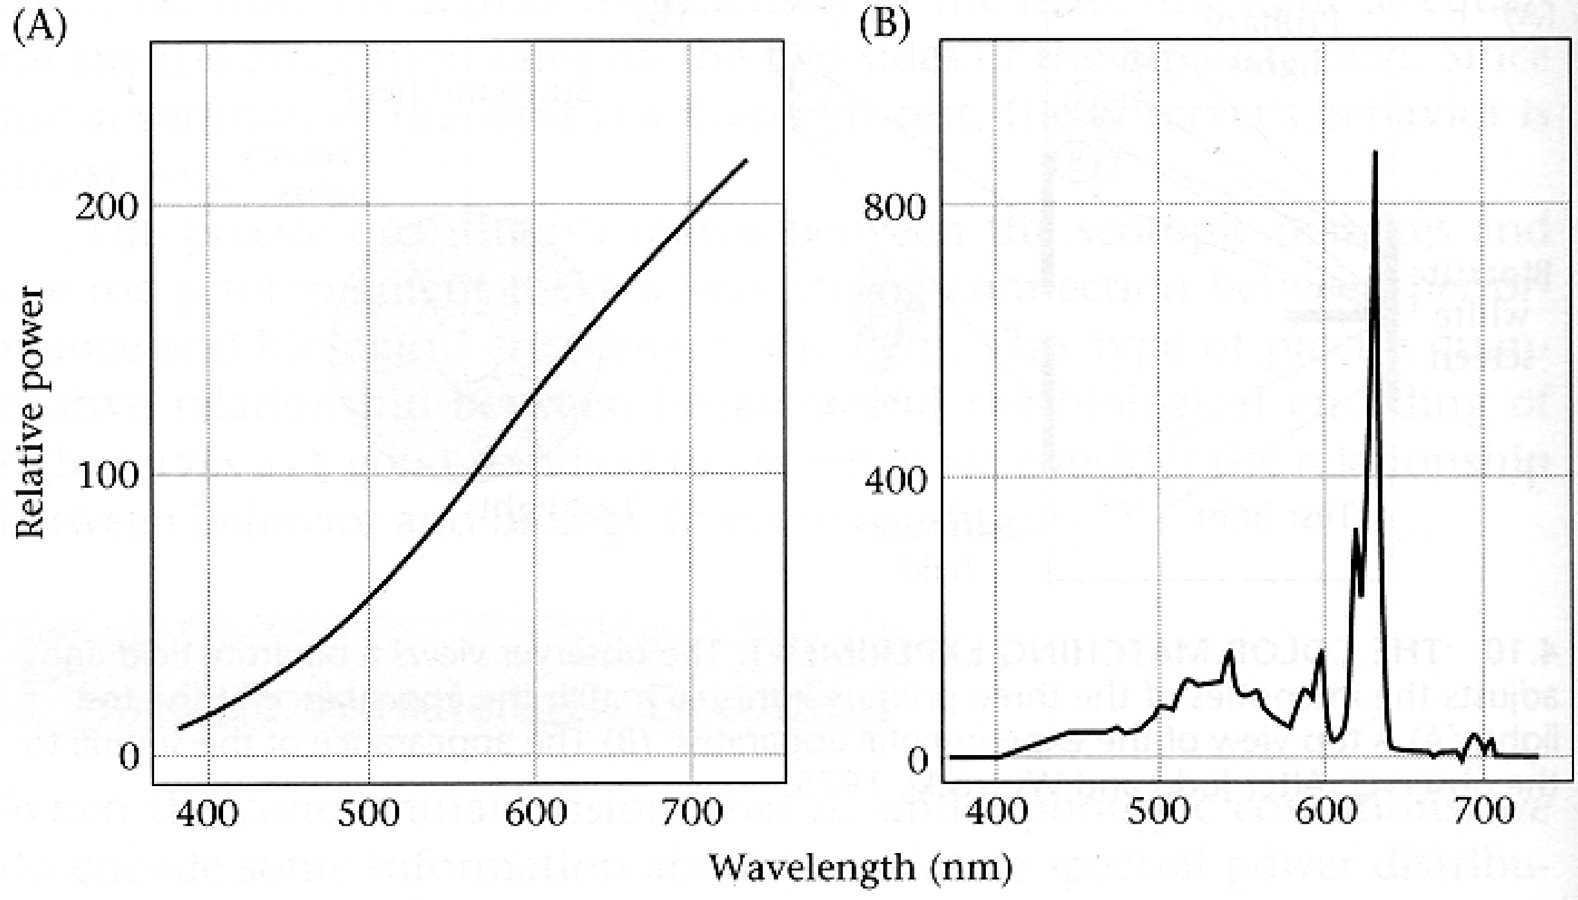

4.11 METAMERIC LIGHTS. Two lights with these spectral power distributions appear identical to most observers and are called metamers. (A) An approximation to the spectral power distribution of a tungsten bulb. (B) The spectral power distribution of light emitted from a conventional television monitor whose three phosphor intensities were set to match the light in panel A in appearance.

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Internal summary

- What are colors?
  - Arise from power spectrum of light.
- How represent colors:
  - Pick primaries
  - Measure color matching functions (CMF's)
  - Matrix mult power spectrum by CMF's to find color as the 3 primary color values.
- How share color descriptions between people?
  - Translate colors between systems of primaries
  - Standardize on a few sets of primaries.

Since we can define colors using almost any set of primary colors, let’s agree on a set of primaries and color matching functions for the world to use…

CIE XYZ color space

- Commission Internationale d'Eclairage, 1931
- "…as with any standards decision, there are some irratating aspects of the XYZ color-matching functions as well…no set of physically realizable primary lights that by direct measurement will yield the color matching functions."
- "Although they have served quite well as a technical standard, and are understood by the mandarins of vision science, they have served quite poorly as tools for explaining the discipline to new students and colleagues outside the field."

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

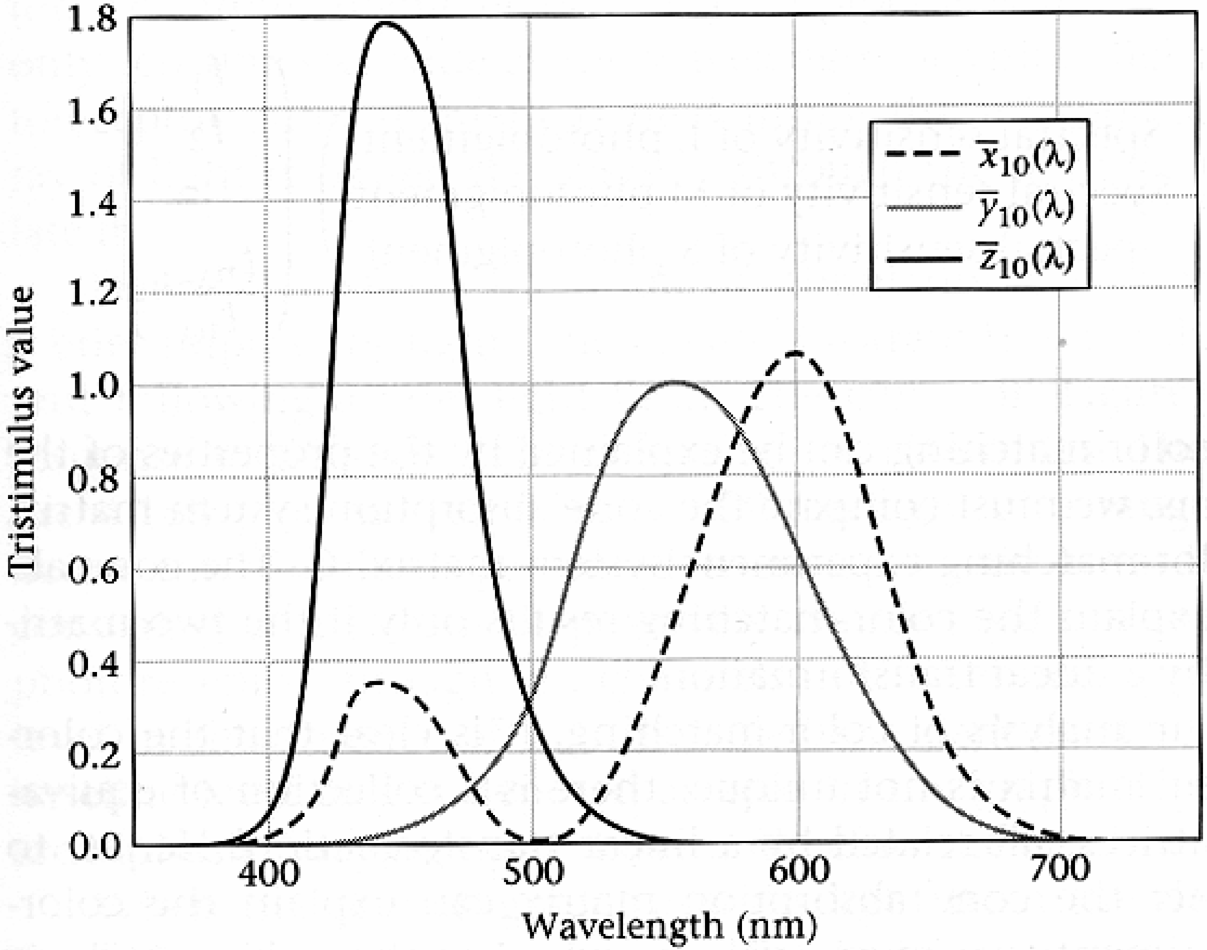

4.14 THE XYZ STANDARD COLOR-MATCHING FUNCTIONS. In 1931 the CIE standardized a set of color-matching functions for image interchange. These color-matching functions are called $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$. Industrial applications commonly describe the color properties of a light source using the three primary intensities needed to match the light source that can be computed from the XYZ color-matching functions.

CIE XYZ: Color matching functions are positive everywhere, but primaries are "imaginary" (require adding light to the test color's side in a color matching experiment). Usually compute x, y, where

$$x = X/(X+Y+Z)$$
$$y = Y/(X+Y+Z)$$

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

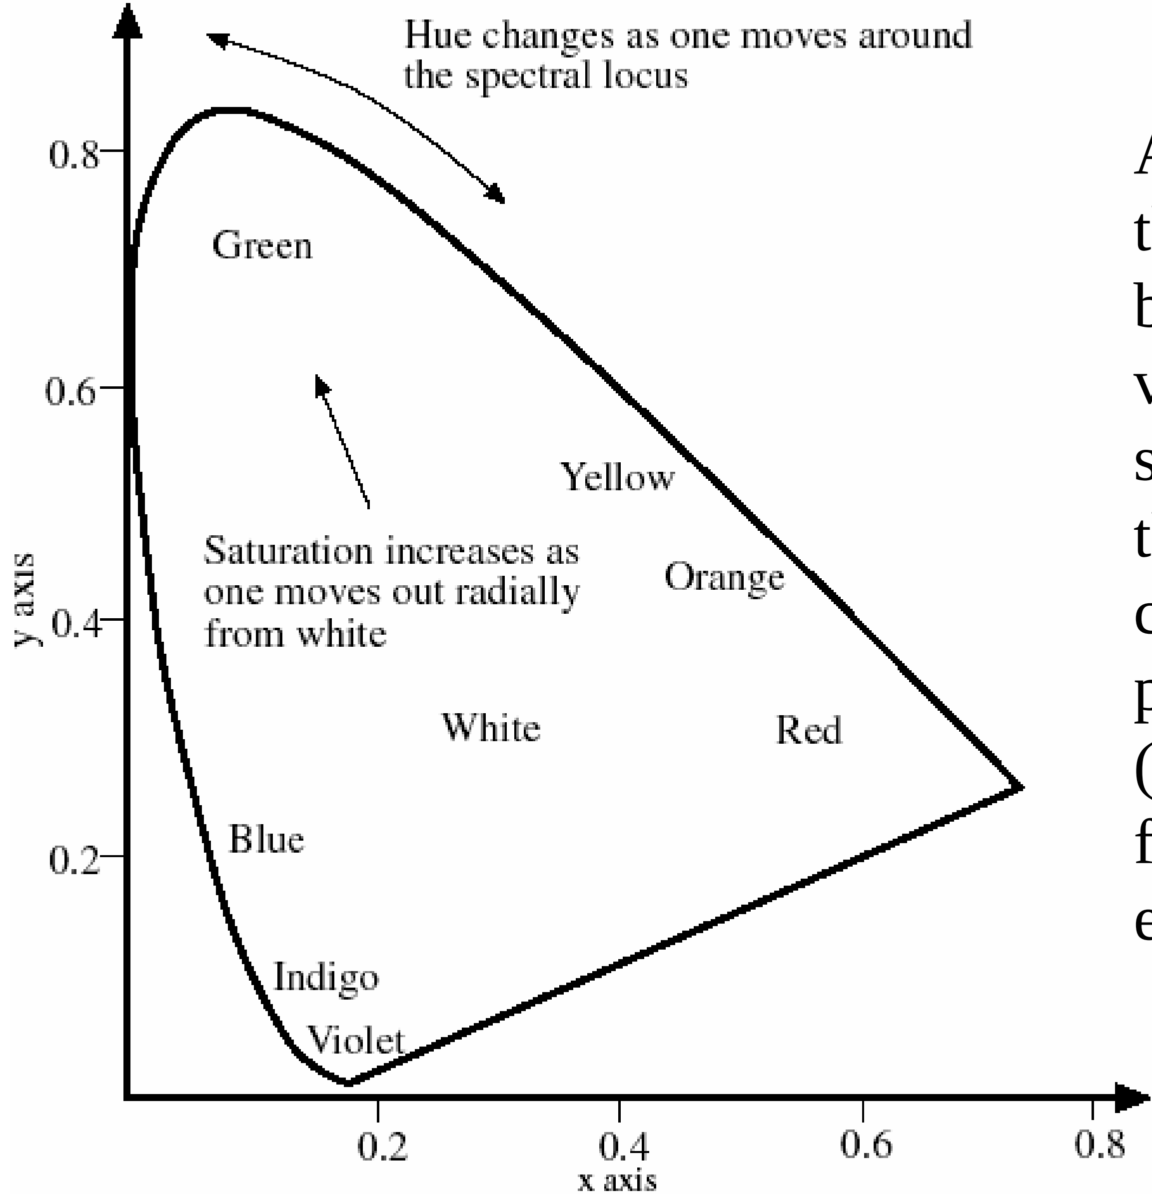

A qualitative rendering of the CIE (x,y) space. The blobby region represents visible colors. There are sets of (x, y) coordinates that don't represent real colors, because the primaries are not real lights (so that the color matching functions could be positive everywhere).

Forsyth & Ponce

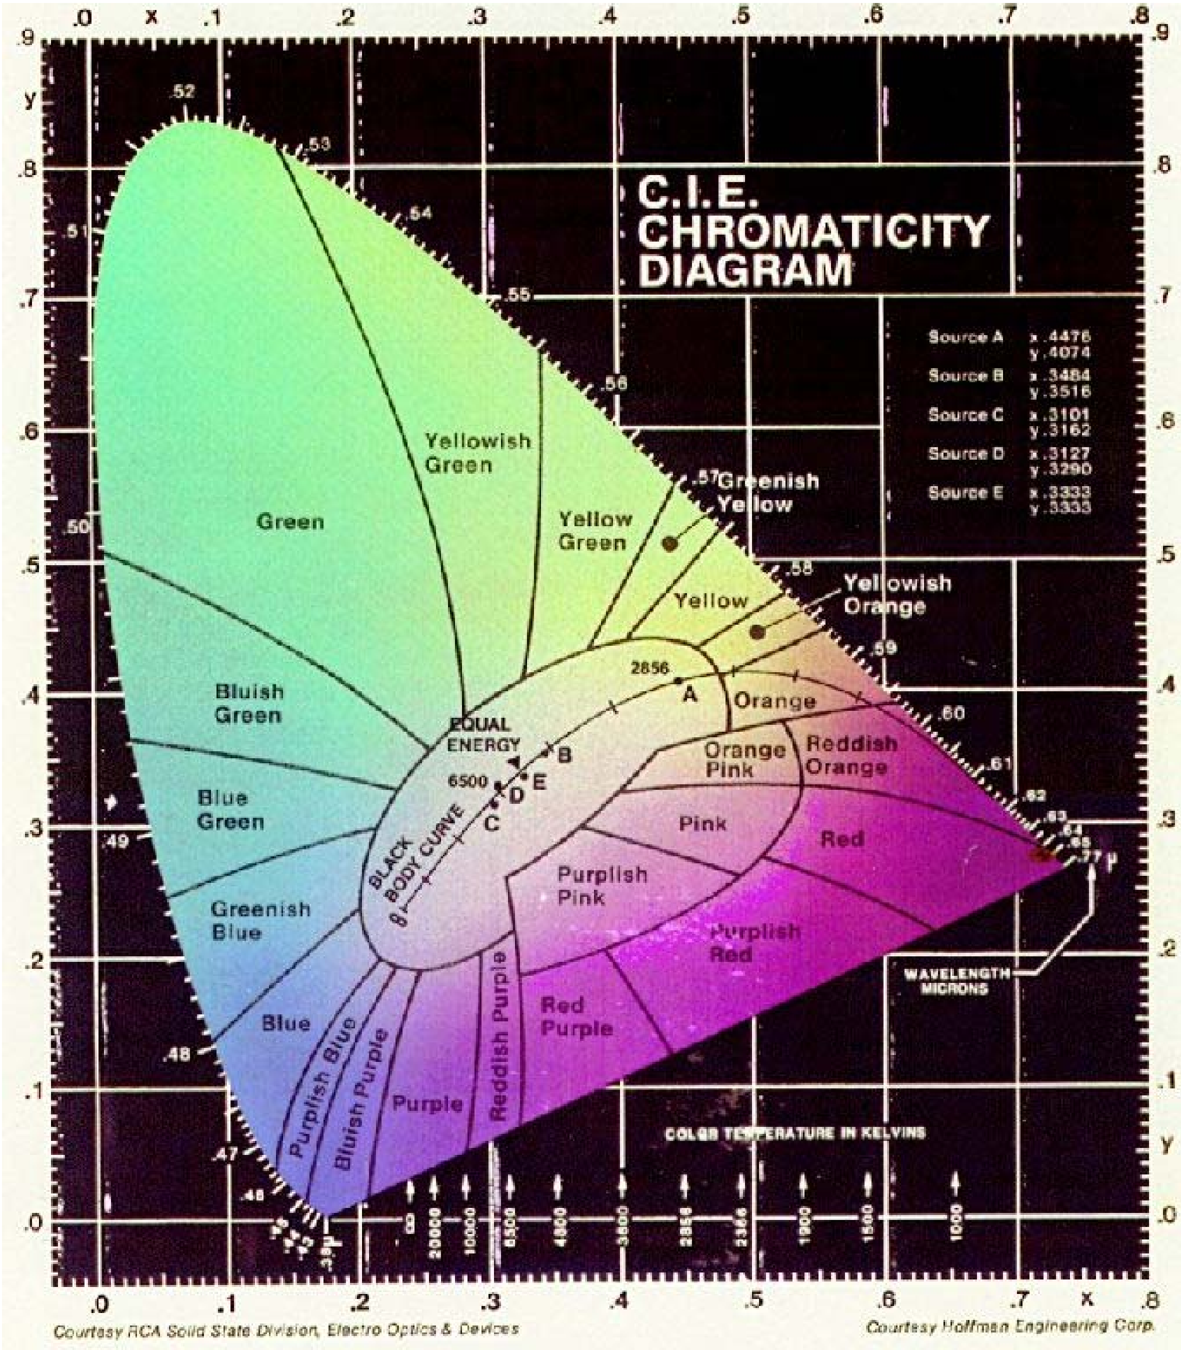
C.I.E.
CHROMATICITY
DIAGRAM
Source A x .4476 y .4074
Source B x .3484 y .3516
Source C x .3101 y .3162
Source D x .3127 y .3290
Source E x .3333 y .3333
Yellowish Green
Green
Yellow Green
Greenish Yellow
Yellow
Yellowish Orange
Orange
Bluish Green
EQUAL ENERGY
Orange Pink
Reddish Orange
Blue Green
Pink
Red
BLACK BODY CURVE
Purplish Pink
Greenish Blue
Purplish Red
WAVELENGTH MICRONS
Blue
Purplish Blue
Bluish Purple
Purple
Reddish Purple
Red Purple
COLOR TEMPERATURE IN KELVINS
Courtesy RCA Solid State Division, Electro Optics & Devices
Courtesy Hoffman Engineering Corp.

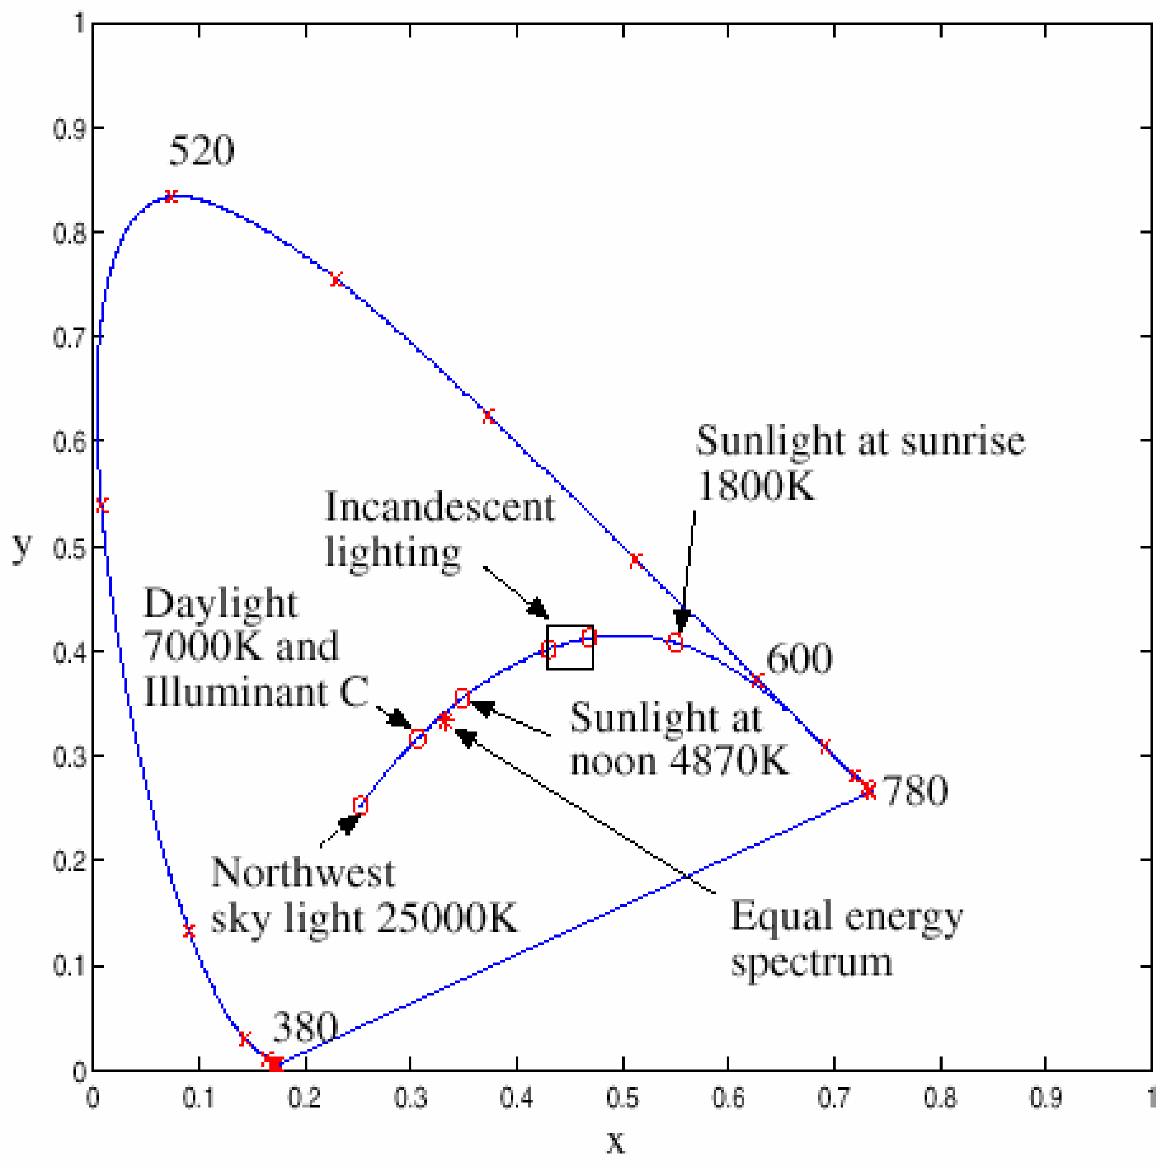

A plot of the CIE (x,y) space. We show the spectral locus (the colors of monochromatic lights) and the black-body locus (the colors of heated black-bodies). I have also plotted the range of typical incandescent lighting.

Pure wavelength in chromaticity diagram

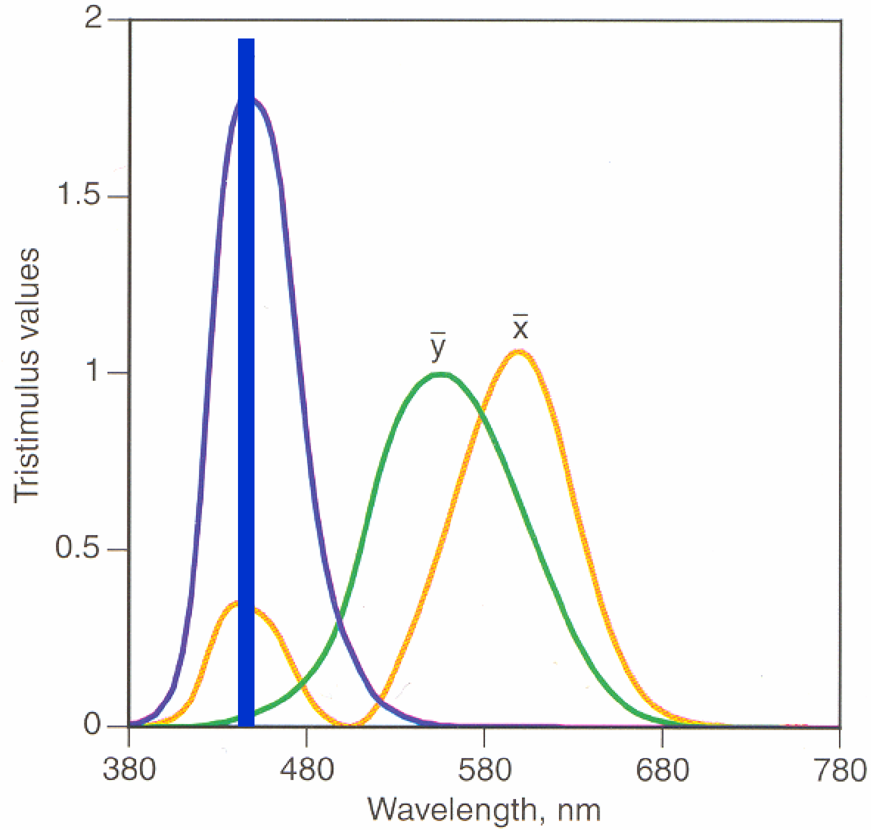

The 1931 standard observer, as it is usually shown.

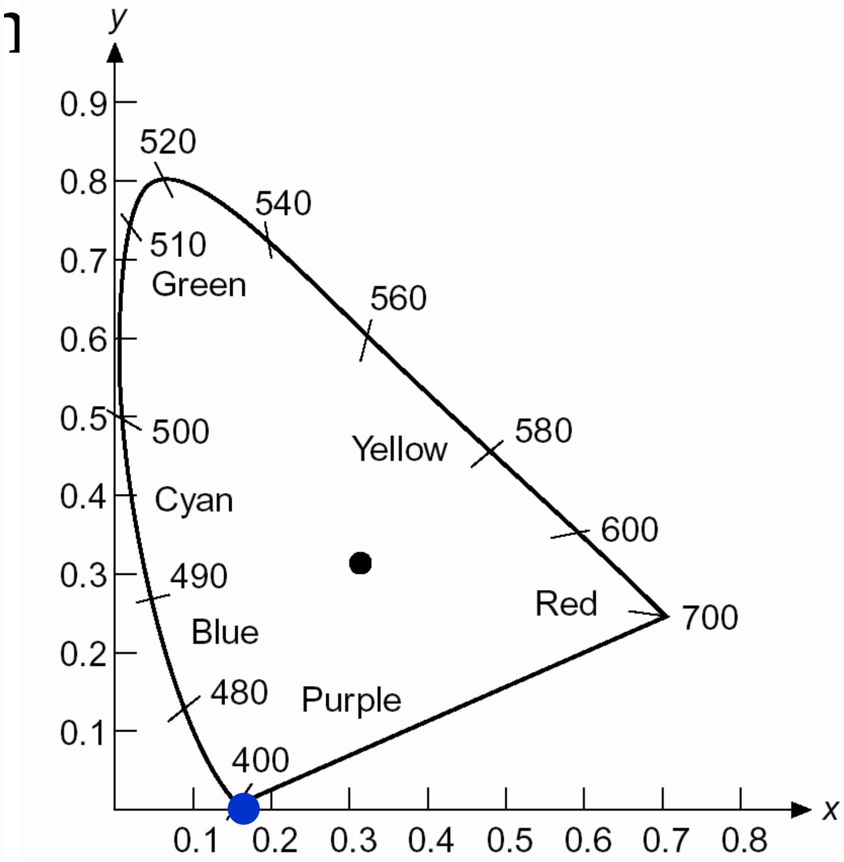

Pure wavelength in chromaticity diagram

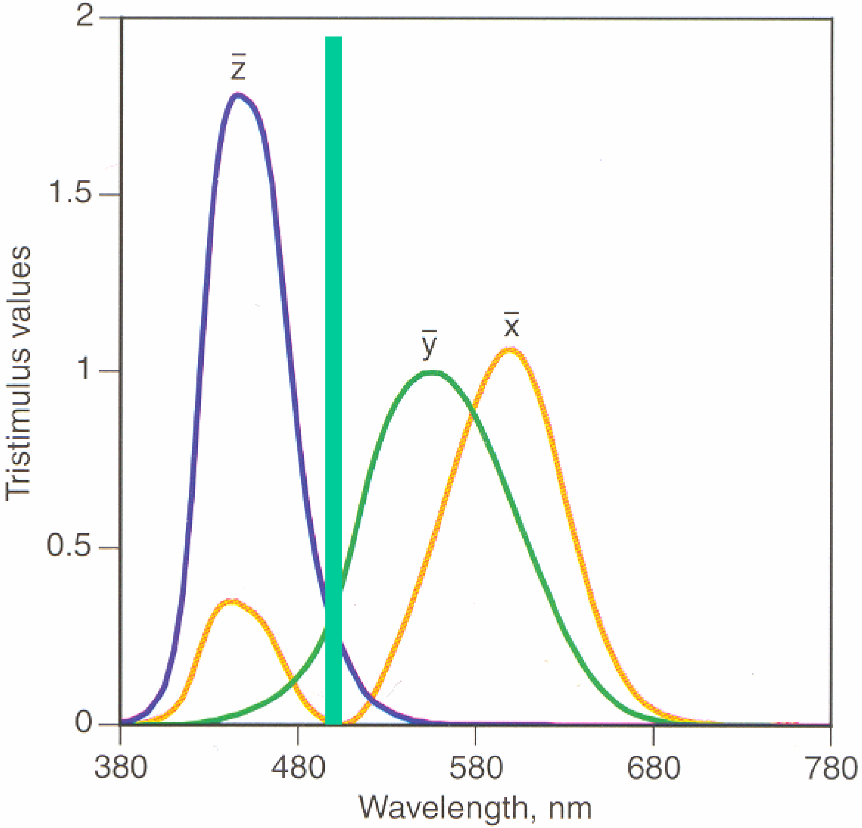

The 1931 standard observer, as it is usually shown.

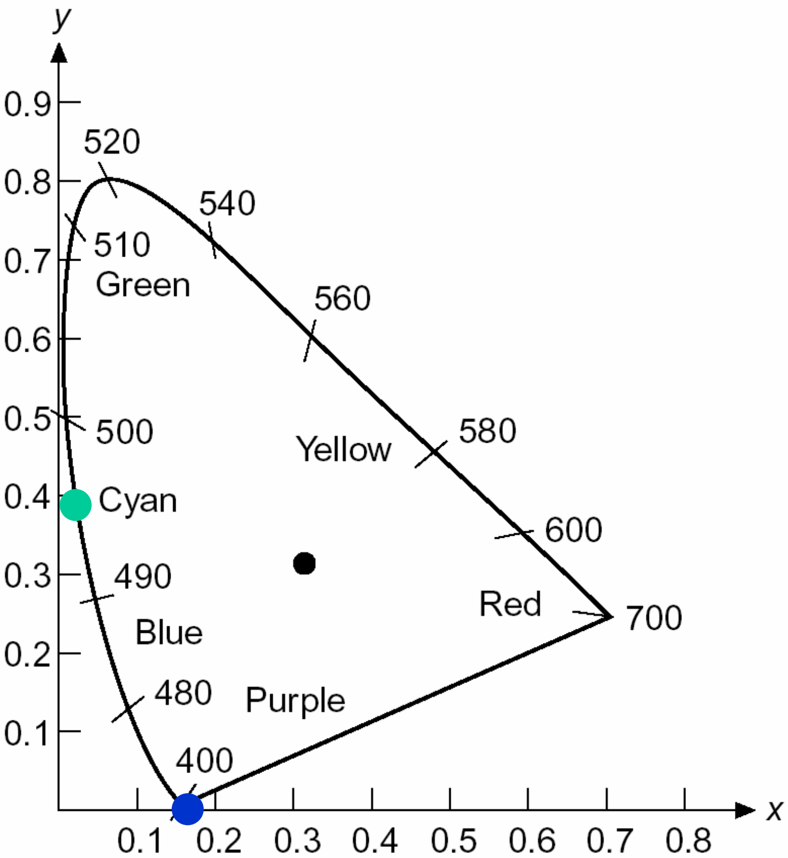

Pure wavelength in chromaticity diagram

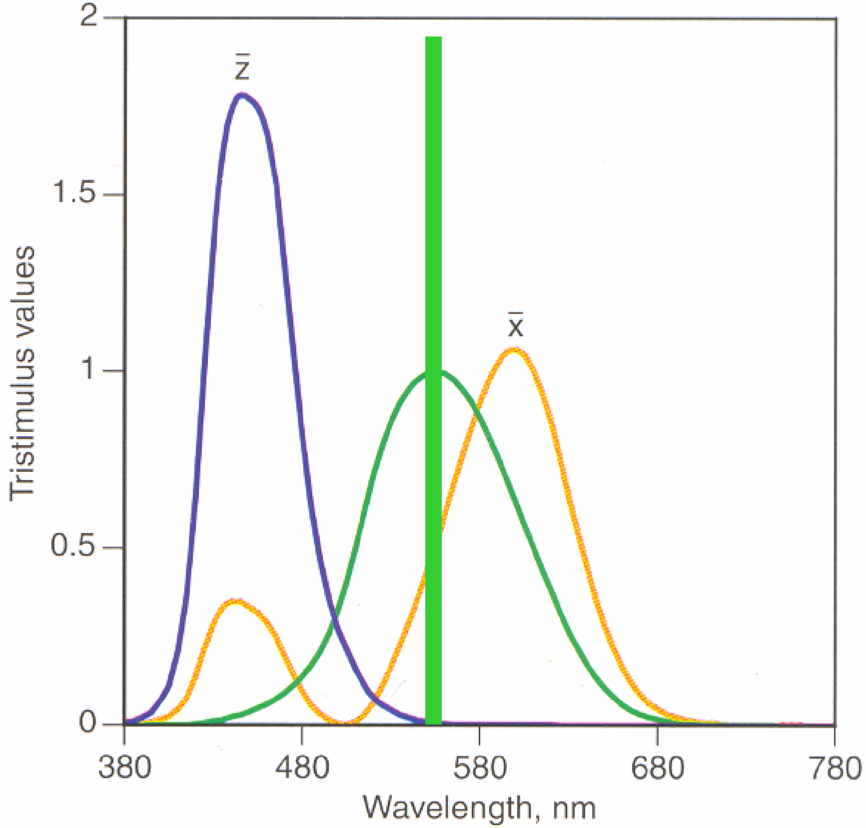

The 1931 standard observer, as it is usually shown.

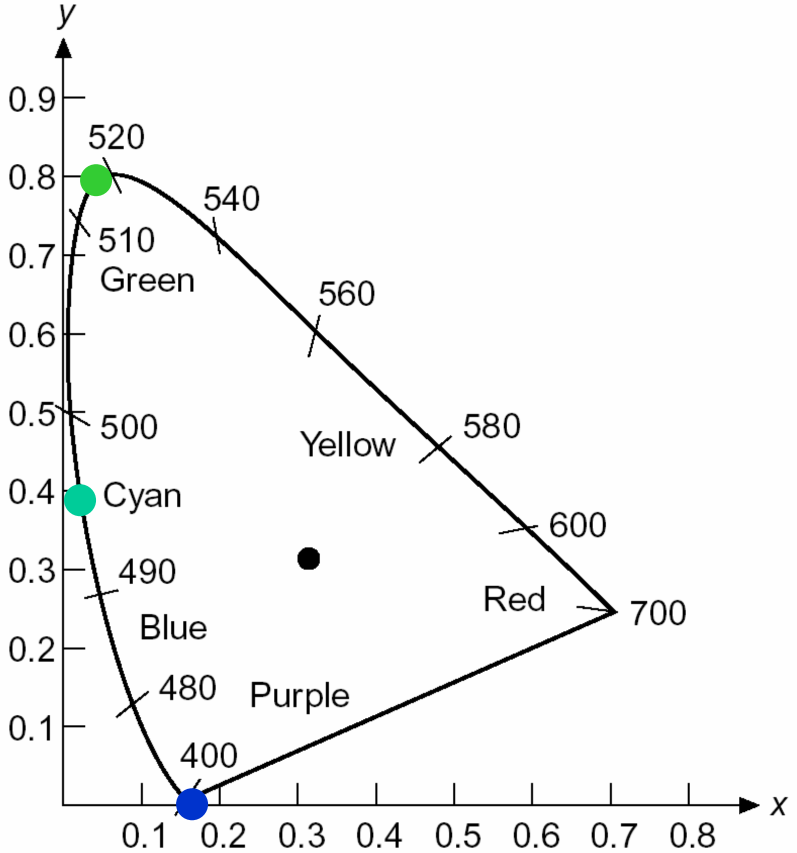

Pure wavelength in chromaticity diagram

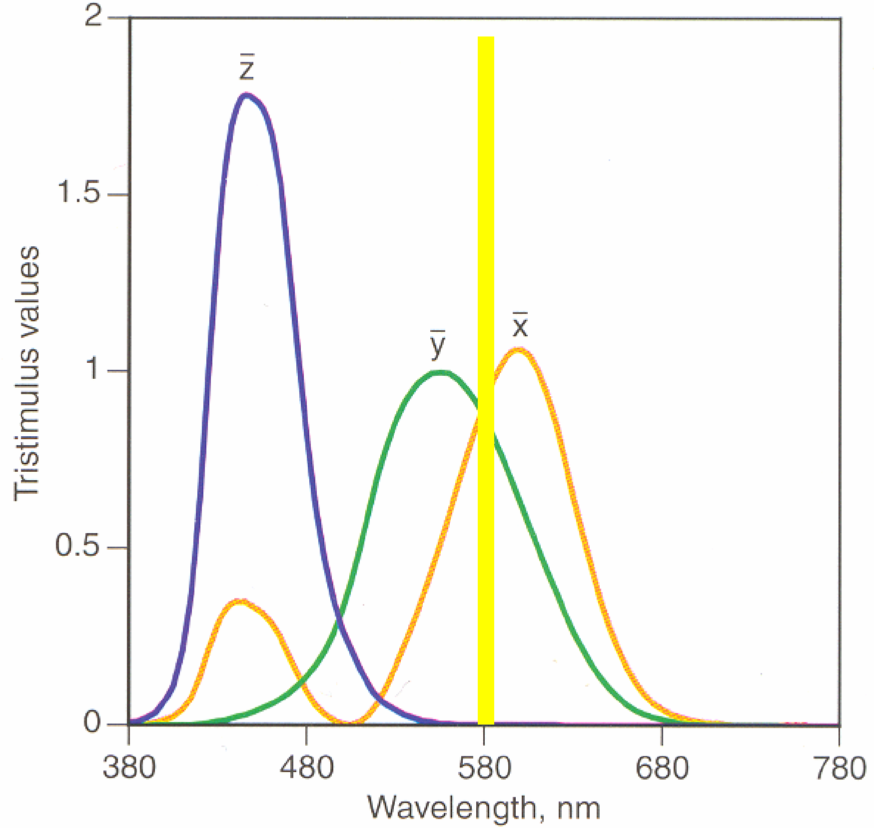

The 1931 standard observer, as it is usually shown.

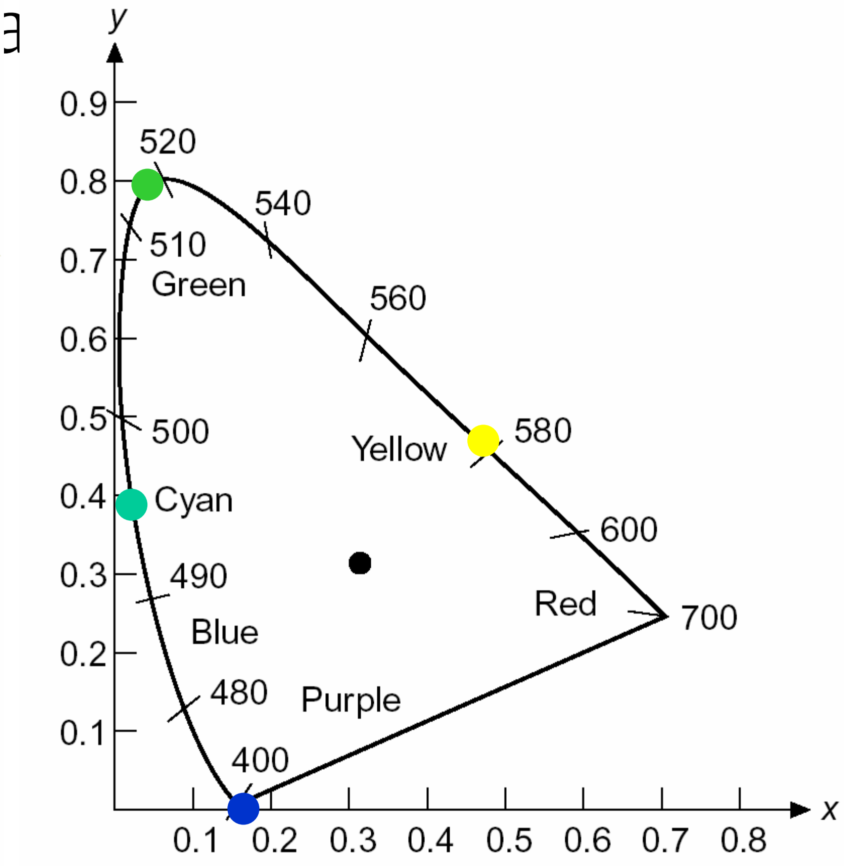

Pure wavelength in chromaticity diagram

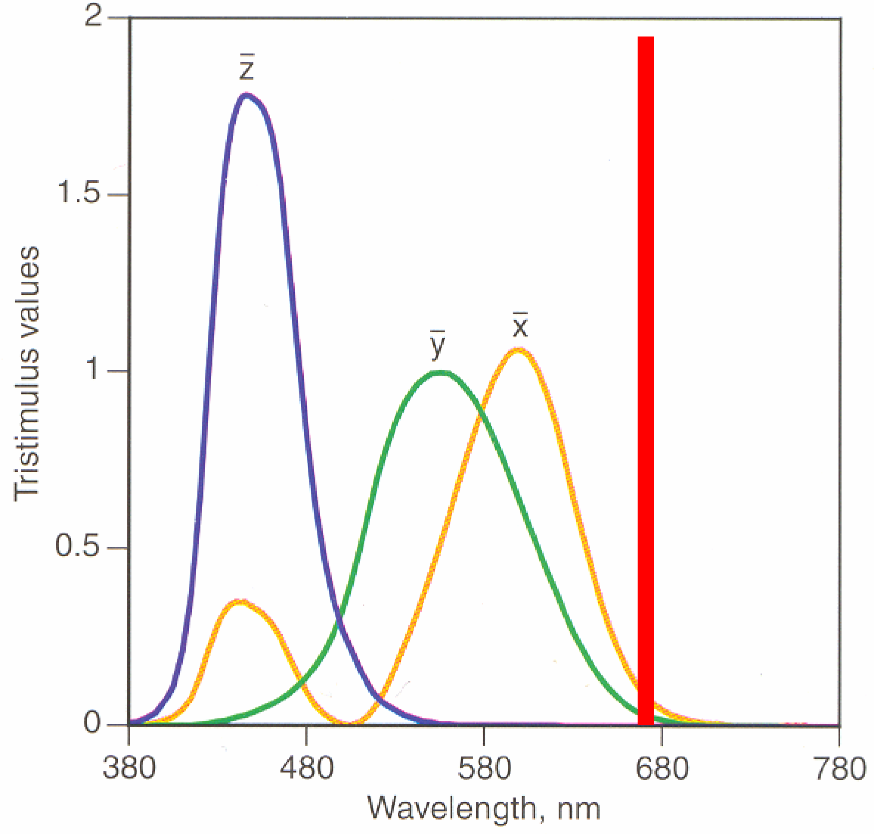

The 1931 standard observer, as it is usually shown.

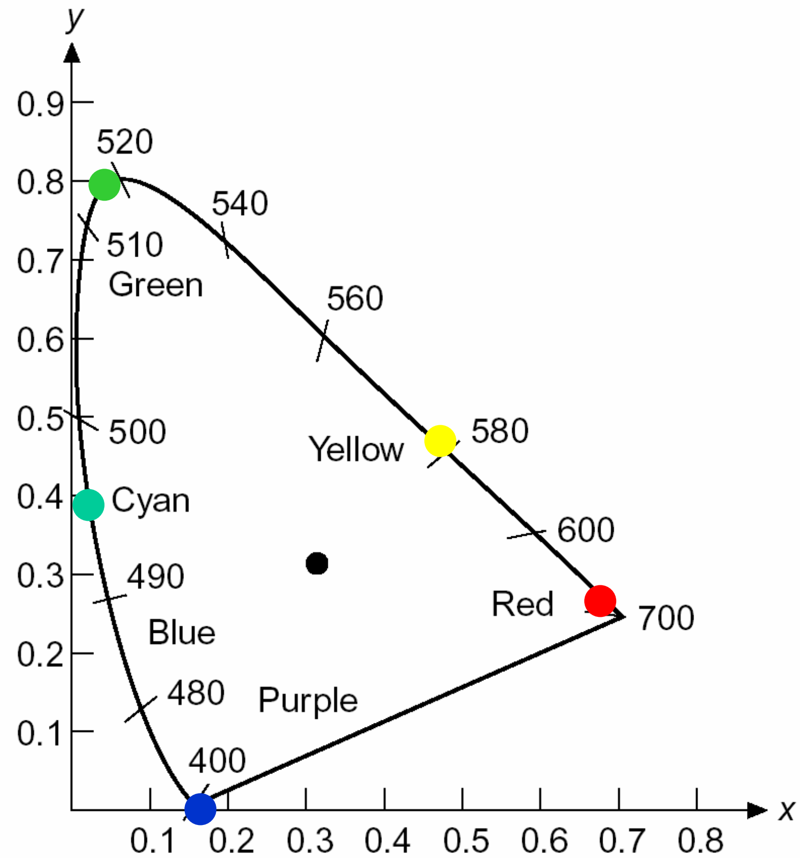

CIE chromaticity diagram

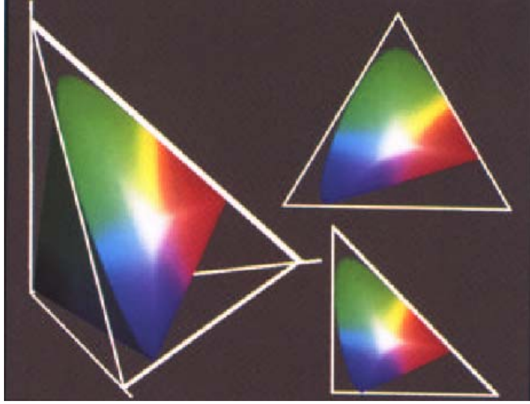

- Spectrally pure colors lie along boundary
- Weird shape comes from shape of matching curves and restriction to positive stimuli
- Note that some hues do not correspond to a pure spectrum (purple-violet)
- Standard white light (approximates sunlight) at C

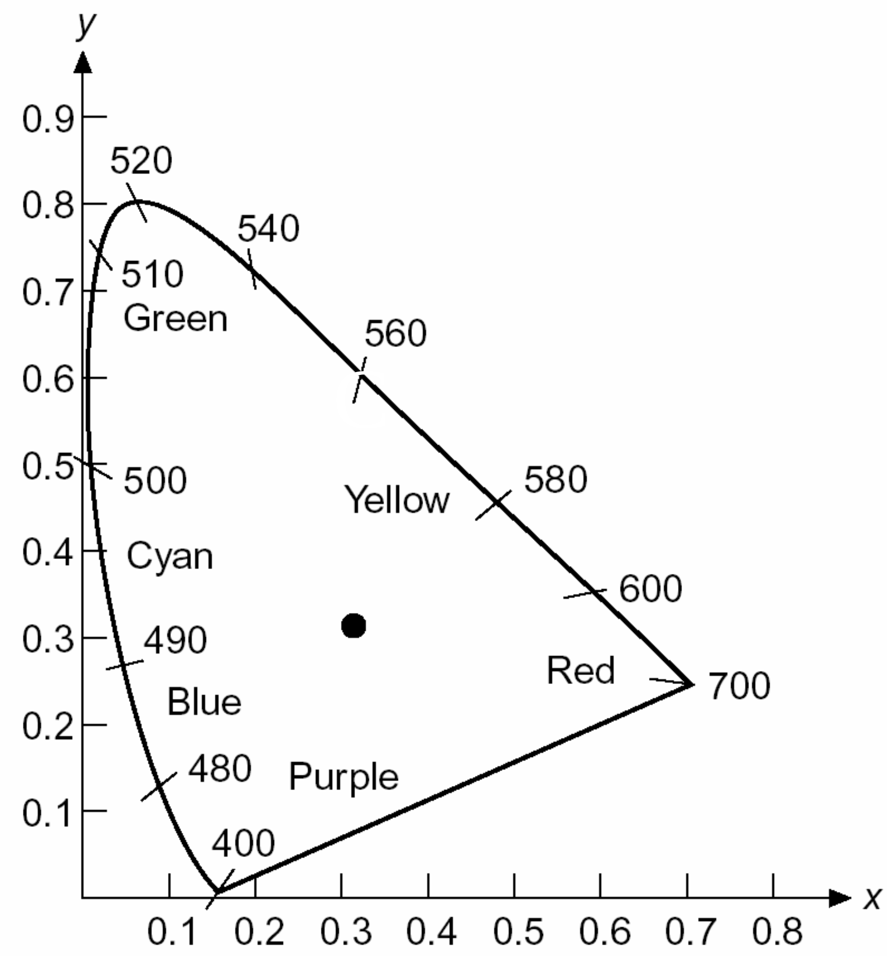

CIE color s

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 3.24 & -1.54 & -0.50 \\ -0.97 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

- Can think of X, Y , Z as coordinates
- Linear transform from typical RGB or LMS
- Always positive (because physical spectrum is positive and matching curves are positives)
- Note that many points in XYZ do not correspond to visible colors!

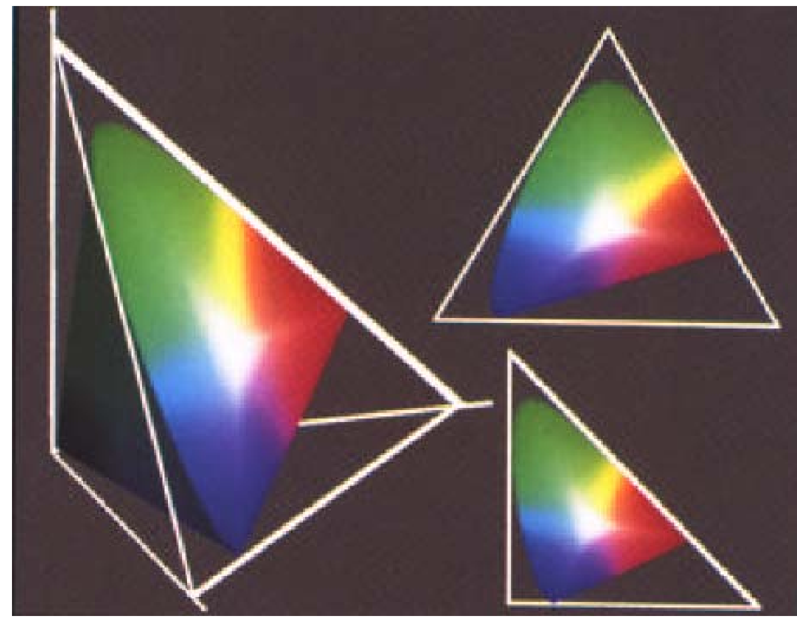

XYZ vs. RGB

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 3.24 & -1.54 & -0.50 \\ -0.97 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

- Linear transform
- XYZ is rarely used for storage
- There are tons of flavors of RGB
  - sRGB, Adobe RGB
  - Different matrices!
- XYZ is more standardized
- XYZ can reproduce all colors with positive values
- XYZ is not realizable physically !!
  - What happens if you go “off” the diagram
  - In fact, the orthogonal (synthesis) basis of XYZ requires negative values.

Questions?

Some other color spaces…

NTSC color components: Y, I, Q

$$\begin{pmatrix} Y \\ I \\ Q \end{pmatrix} = \begin{pmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.274 & -0.322 \\ 0.211 & -0.523 & 0.312 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

NTSC - RGB

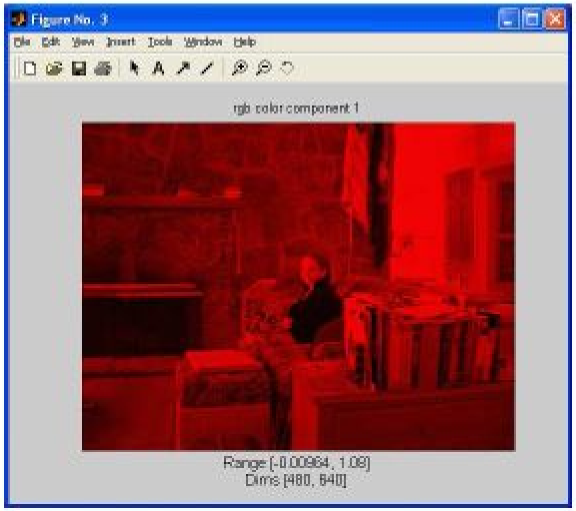
Figure No. 3
rgb color component 1
Range [-0.00964, 1.08]
Dims [480, 640]

R

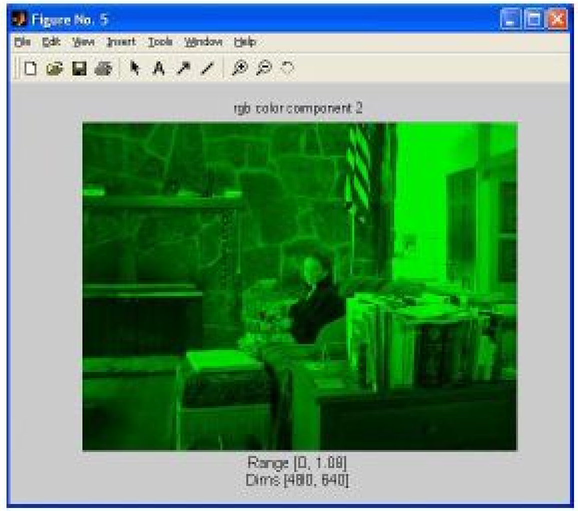
Figure No. 5
rgb color component 2
Range [0, 1.08]
Dims [480, 640]

G

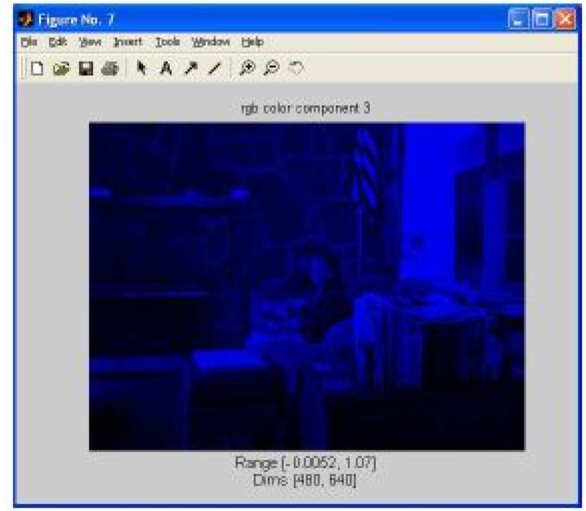
Figure No. 7
rgb color component 3
Range [-0.0052, 1.07]
Dims [480, 640]

B

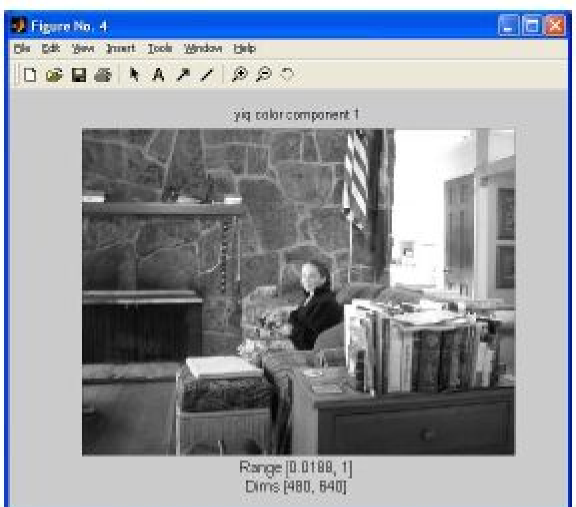
Figure No. 4
yiq color component 1
Range [0.0188, 1]
Dims [480, 640]

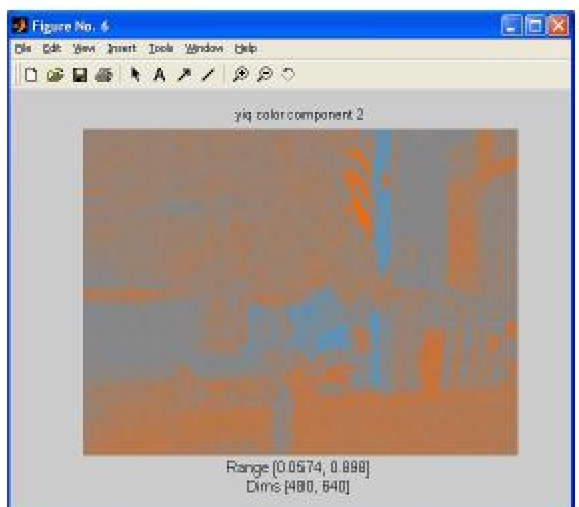
Figure No. 6
yiq color component 2
Range [0.0574, 0.898]
Dims [480, 640]

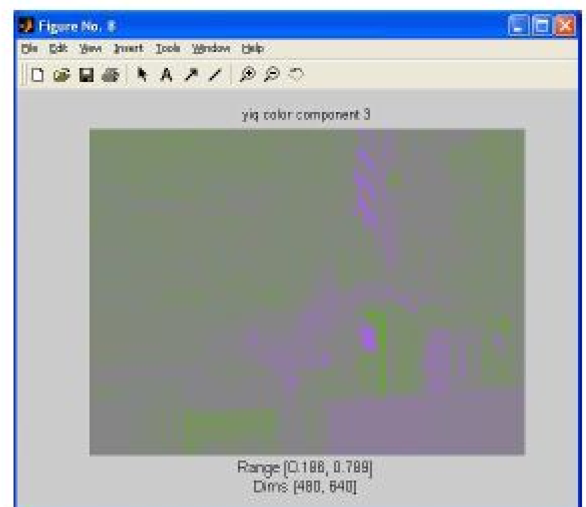
Figure No. 8
yiq color component 3
Range [0.186, 0.789]
Dims [480, 640]

HSV hexcone

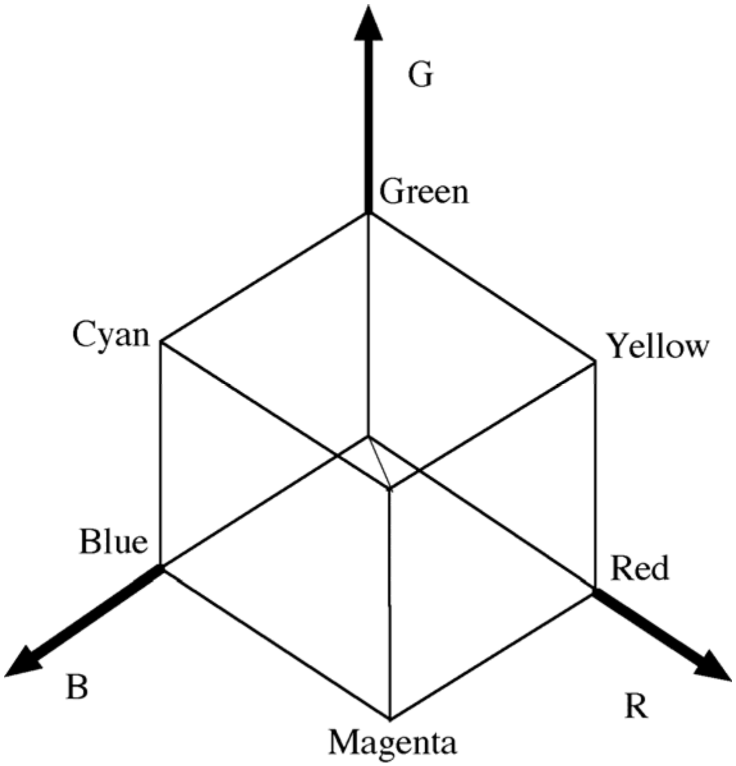

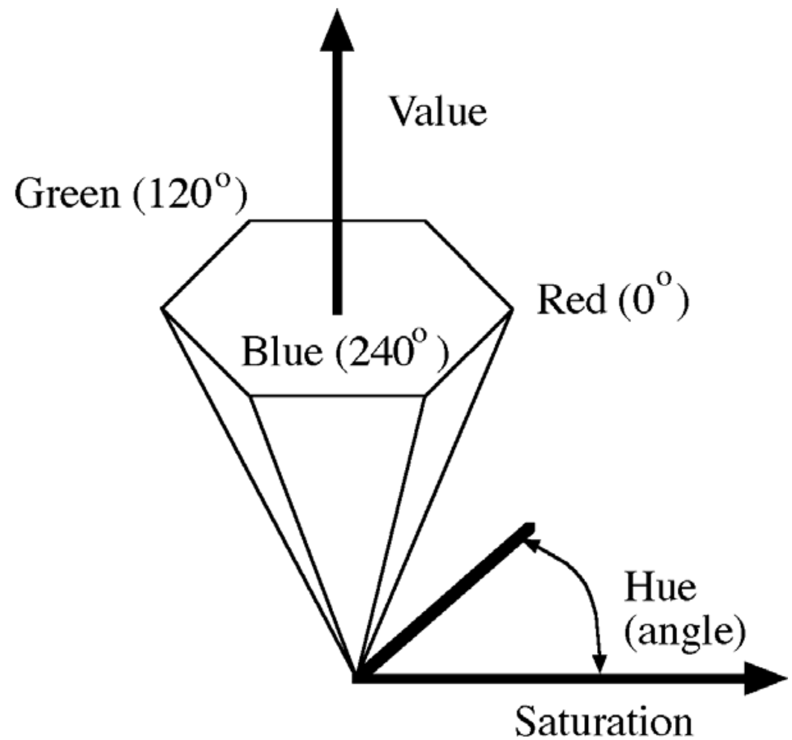

Forsyth & Ponce

Hue Saturation Value

- Value: from black to white
- Hue: dominant color (red, orange, etc)
- Saturation: from gray to vivid color
- HSV double cone

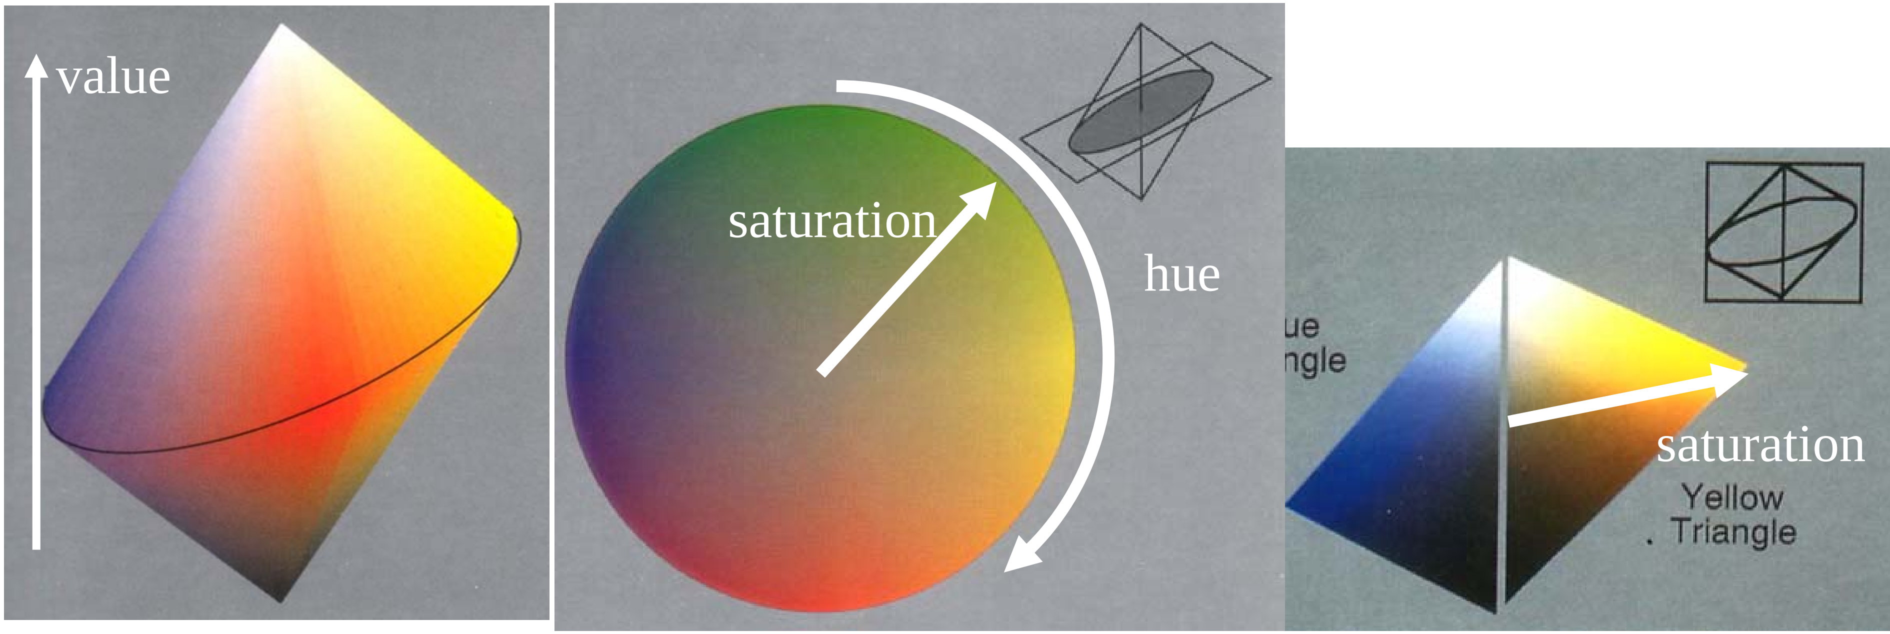

Uniform color spaces

- McAdam ellipses (next slide) demonstrate that differences in x,y are a poor guide to differences in color
- Construct color spaces so that differences in coordinates are a good guide to differences in color.

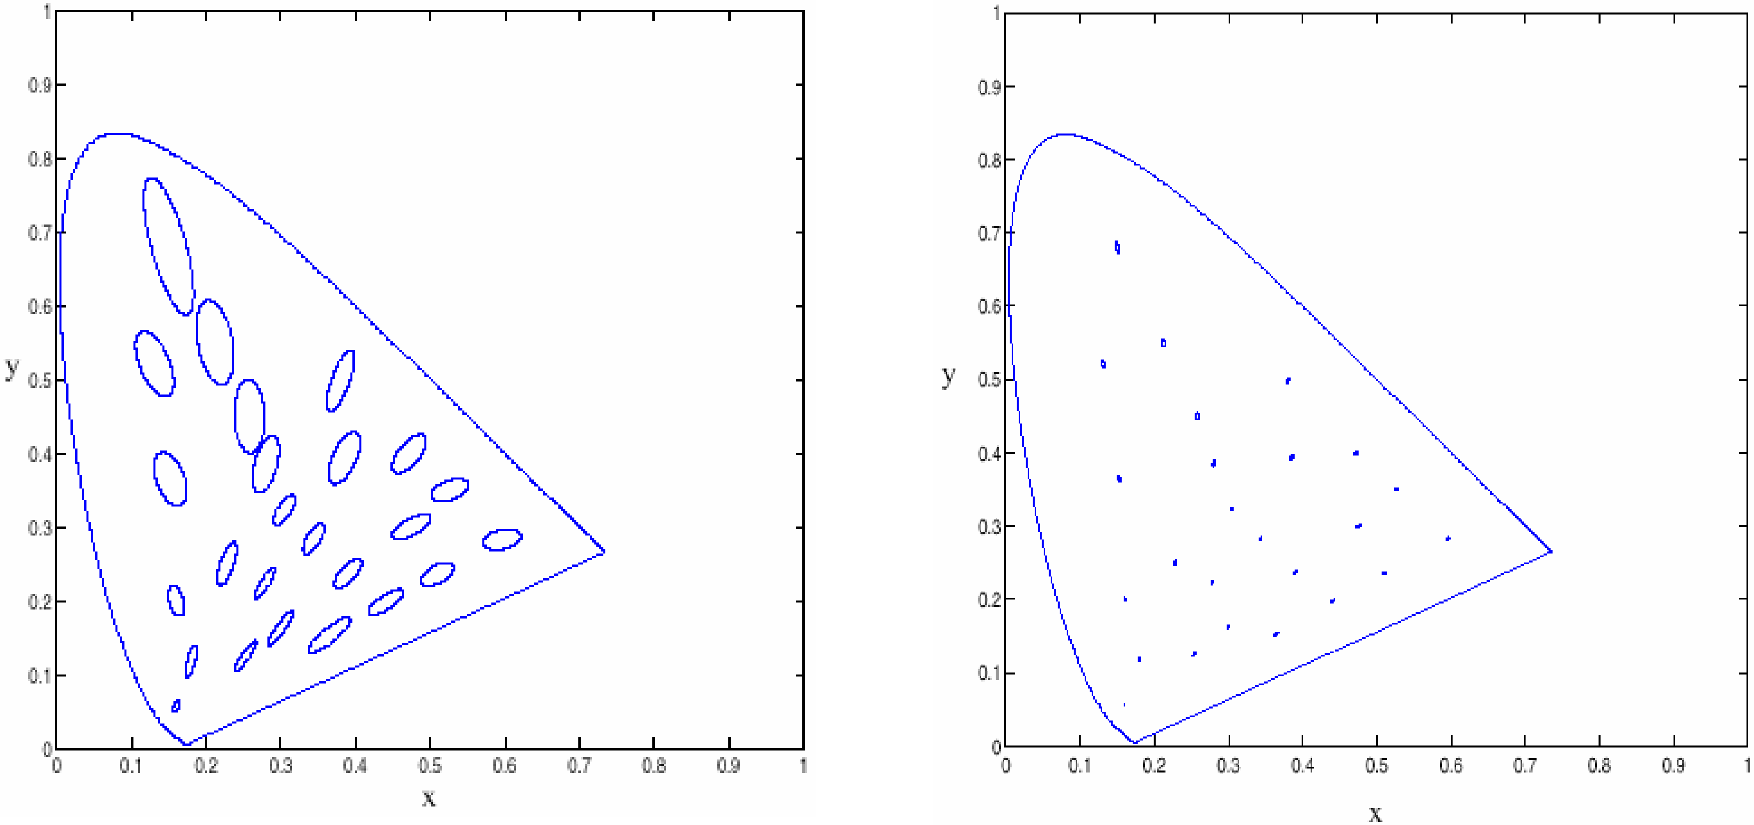

Variations in color matches on a CIE x, y space. At the center of the ellipse is the color of a test light; the size of the ellipse represents the scatter of lights that the human observers tested would match to the test color; the boundary shows where the just noticeable difference is. The ellipses on the left have been magnified 10x for clarity; on the right they are plotted to scale. The ellipses are known as MacAdam ellipses after their inventor. The ellipses at the top are larger than those at the bottom of the figure, and that they rotate as they move up. This means that the magnitude of the difference in x, y coordinates is a poor guide to the difference in color.

Forsyth & Ponce

Perceptually Uniform Space: MacAdam

- In perceptually uniform color space, Euclidean distances reflect perceived differences between colors
- MacAdam ellipses (areas of unperceivable differences) become circles
- Non-linear mapping, many solutions have been proposed

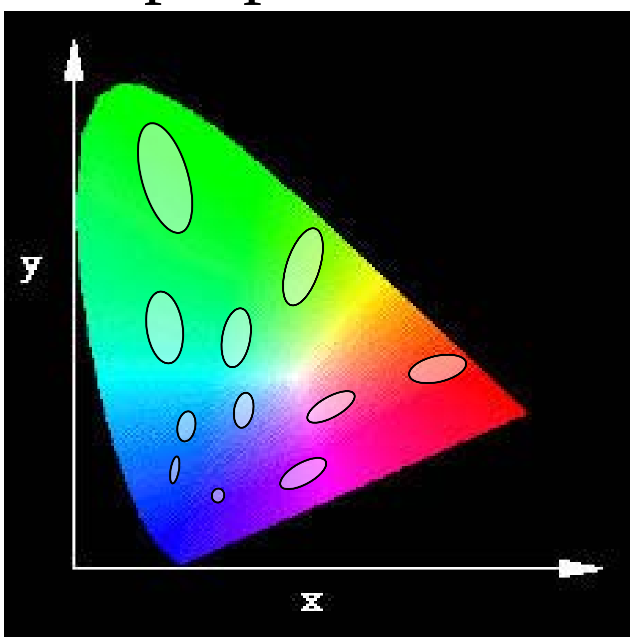

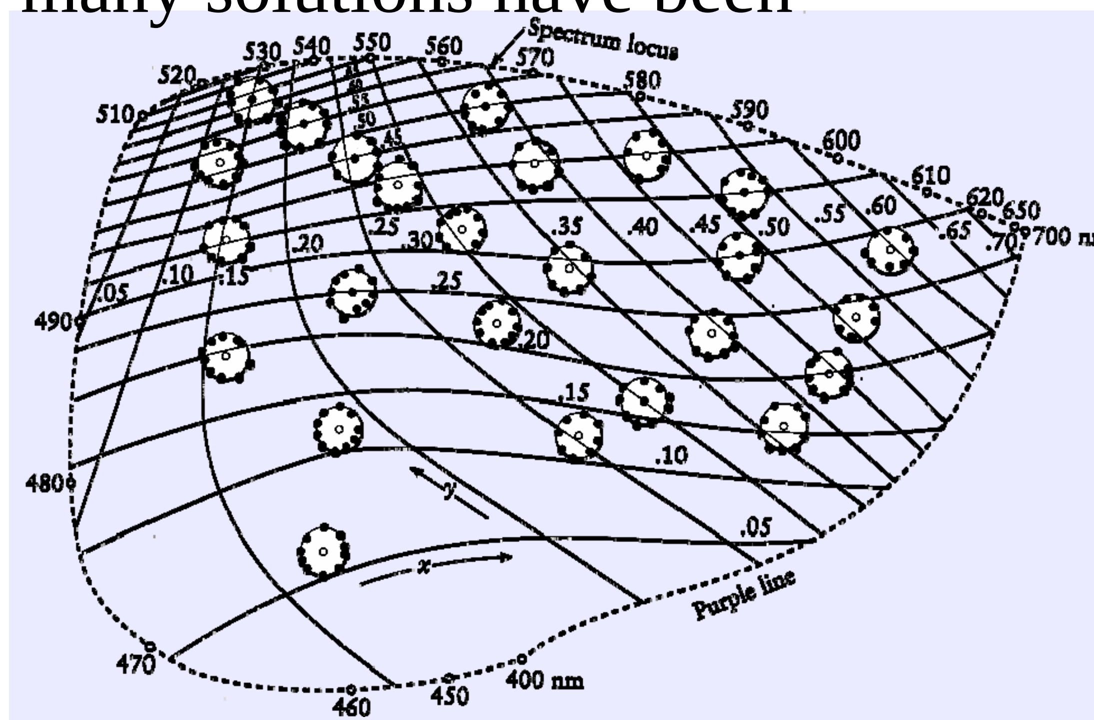

Source: [Wyszecki and Stiles '82]

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

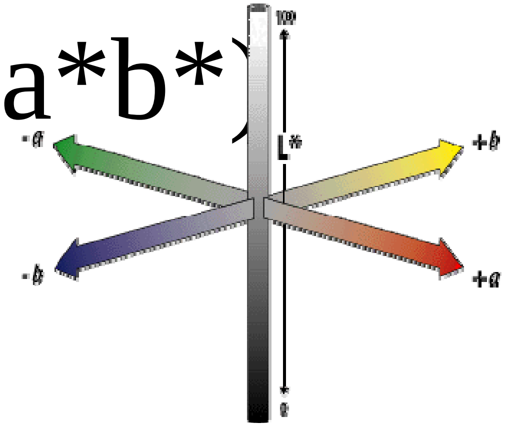

$$L = 25\left(100\frac{Y}{Y_0}\right)^{1/3} - 16$$

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

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

- The reference perceptually uniform color space
- L: lightness
- a and b: color opponents

- X_0, Y_0, and Z_0 are used to color-balance: they're the color of the reference white

Source: [Wyszecki and Stiles '82]

Some class project ideas using the lecture material on color

Class project idea 1

- How best convert from hyperspectral image to rgb image?
- A related paper: http://www.ee.washington.edu/research/guptalab/publications/grspaperJacobsonGupta.pdf
- But the focus is on display of satellite data.

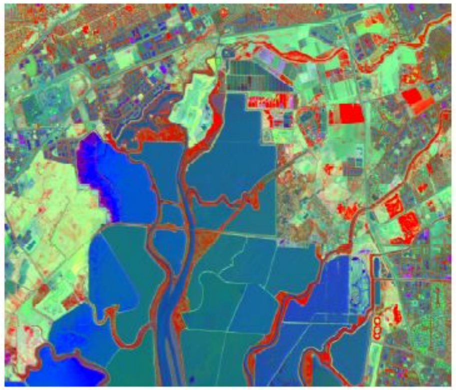

Fig. 4. PCA visualization of Moffett Field, with display mapping $(P_1, P_2, P_3) \rightarrow (R, G, B)$

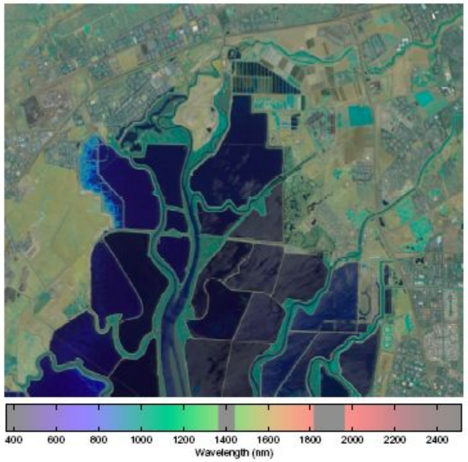

Class project idea 1

(I'll be happy to help you with this project)

- Start from a hyperspectral photograph.
- Re-render the image into RGB to try to meet these two criteria:
  - Having the perceptual distance between colors correspond to the distance between their power spectra, and
  - Having the colors relate somewhat to their true colors.
- Why is there any hope? Because you do this optimization on a per-image basis, and any given image has lots of unused colors you can exploit.
- This optimization would reveal the invisible metameric color changes, while maintaining a natural looking image.
- Or a simpler problem: render to make perceptual distances correspond to hyperspectral distances, without requiring that the colors look right.

Class project idea 2: time-lapse photography temporal color filtering

- Some colors change slowly over time and we can’t easily perceive those long-term changes.
- Take photographs over time of imagery you want to analyze, and include a color calibration card in the scene.
- From the measurements over the card, you can pull out the illumination spectrum for each photo, and show each image as if they were all taken under the same illumination.
- Then color differences between images should correspond to true surface color changes. Temporally filter the color-balanced time-lapse imagery to accentuate the color changes of your subject over time. This will give you a color magnifying glass to exaggerate color changes over time.

Class project idea 3,
the hair-brained one: revealing hidden colors

- Color magnification II: fit hyperspectral photographic measurements as an illuminant spectrum times a surface spectrum that is a product of two or three fundamental dye spectra. Redisplay the image to show the small variations in concentration of the invisible spectra. This might allow you to see color changes that would otherwise be masked.

Color constancy demo

- We assumed that the spectrum impinging on your eye determines the object color. That’s often true, but not always.   Here’s a counter-example…

Selected Bibliography

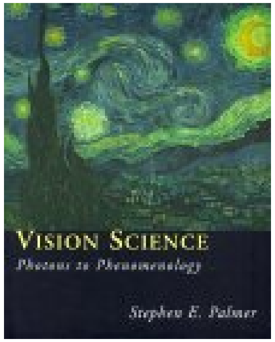

Vision Science
by Stephen E. Palmer
MIT Press; ISBN: 0262161834
760 pages (May 7, 1999)

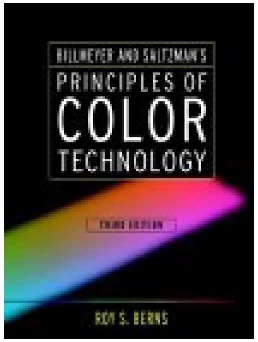

Billmeyer and Saltzman's Principles of Color Technology, 3rd Edition
by Roy S. Berns, Fred W. Billmeyer, Max Saltzman
Wiley-Interscience; ISBN: 047119459X
304 pages 3 edition (March 31, 2000)

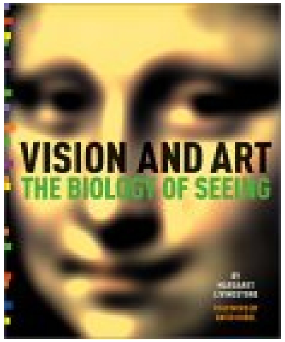

Vision and Art : The Biology of Seeing
by Margaret Livingstone, David H. Hubel
Harry N Abrams; ISBN: 0810904063
208 pages (May 2002)

Selected Bibliography

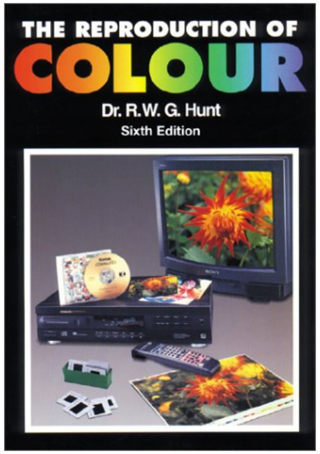

The Reproduction of Color
by R. W. G. Hunt
Fountain Press, 1995

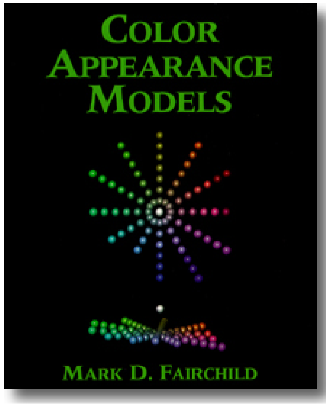

Color Appearance Models
by Mark Fairchild
Addison Wesley, 1998

Other color references

- Reading:
  - Chapter 6, Forsyth & Ponce
  - Chapter 4 of Wandell, Foundations of Vision, Sinauer, 1995 has a good treatment of this.

Class photos

CCD color sampling

What are some approaches to sensing color images?

- Scan 3 times (temporal multiplexing)
- Use 3 detectors (3-ccd camera, and color film)
- Use offset color samples (spatial multiplexing)

Some approaches to color sensing

- Scan 3 times (temporal multiplexing)
  - Drum scanners
  - Flat-bed scanners
  - Russian photographs from 1800’s
- Use 3 detectors
  - High-end 3-tube or 3-ccd video cameras
  - Photographic film
- Use spatially offset color samples (spatial multiplexing)
  - Single-chip CCD color cameras
  - Human eye

Typical errors in spatial multiplexing approach.

- Color fringes.

CCD color filter pattern

detector

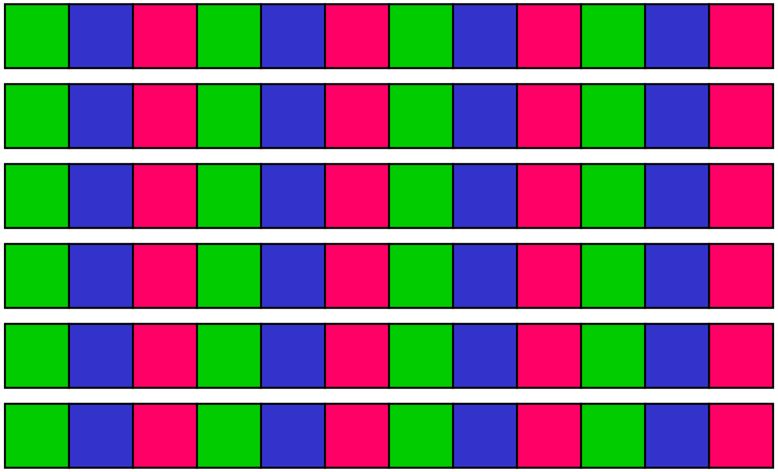

The cause of color moire

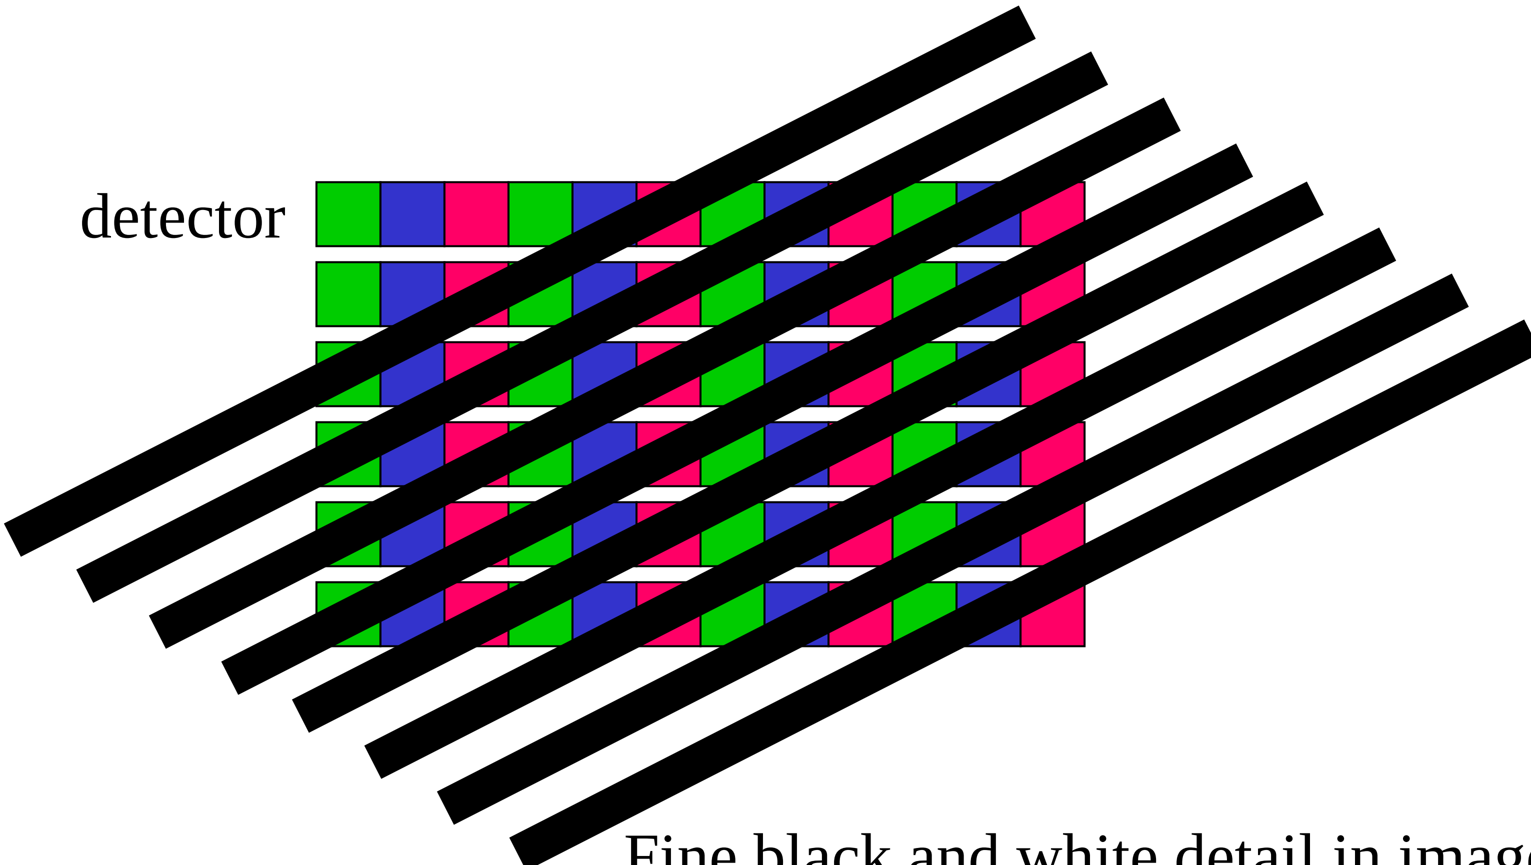

Fine black and white detail in image mis-interpreted as color information.

The Fourier transform of a sampled signal

$$F\left(\mathrm{Sample}_{2\mathrm{D}}(f(x,y))\right) = F\left(f(x,y) \sum_{i=-\infty}^{\infty} \sum_{i=-\infty}^{\infty} \delta(x-i, y-j) \right)$$

$$= F\left(f(x,y)\right) ** F\left(\sum_{i=-\infty}^{\infty} \sum_{i=-\infty}^{\infty} \delta(x-i, y-j) \right)$$

$$= \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} F(u-i, v-j)$$

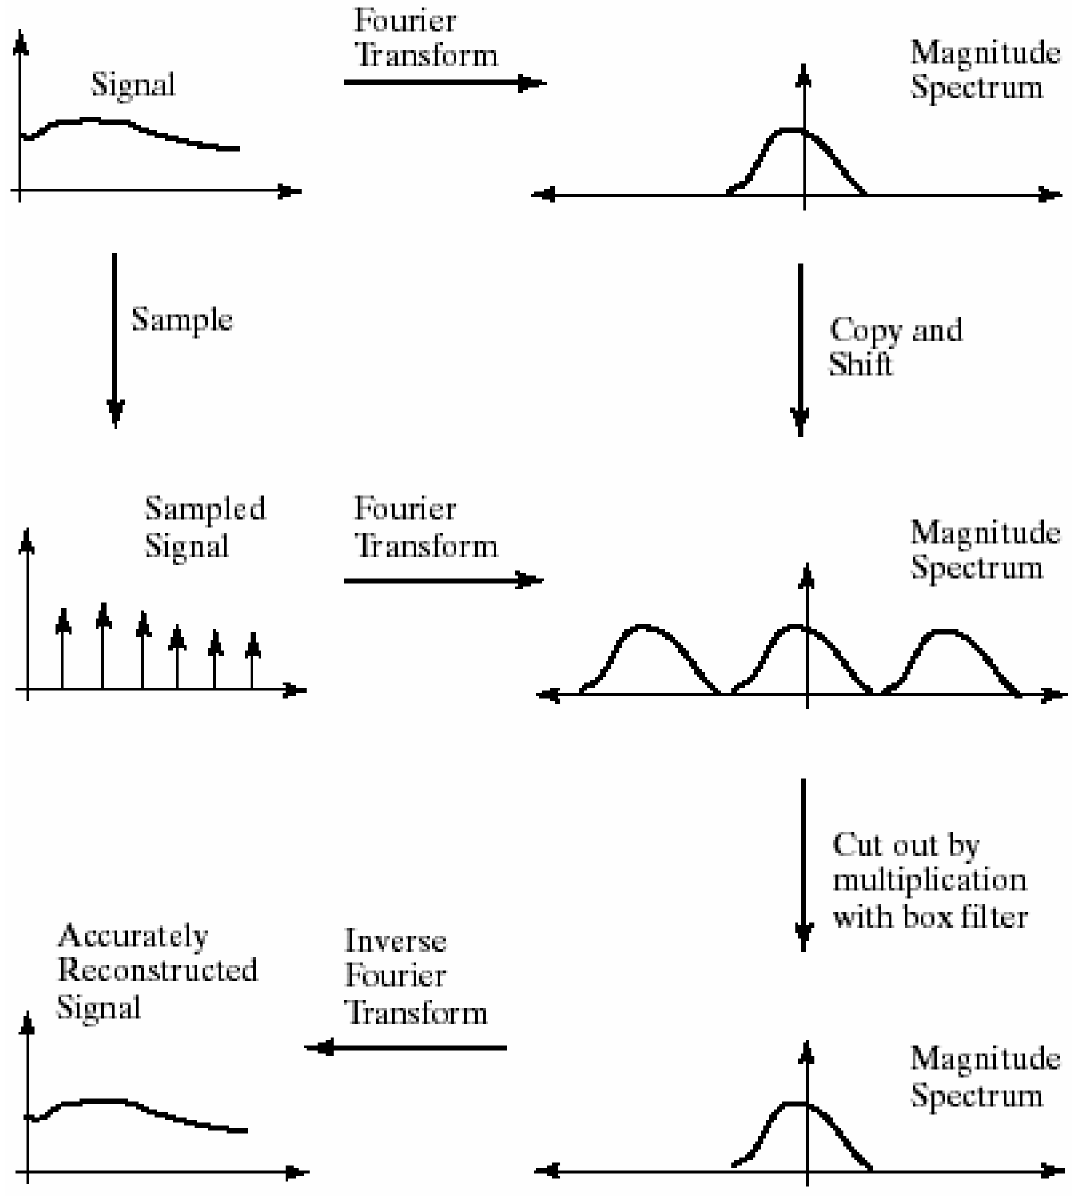
Signal
Fourier Transform
Magnitude Spectrum
Sample
Copy and Shift
Sampled Signal
Fourier Transform
Magnitude Spectrum
Cut out by multiplication with box filter
Accurately Reconstructed Signal
Inverse Fourier Transform
Magnitude Spectrum

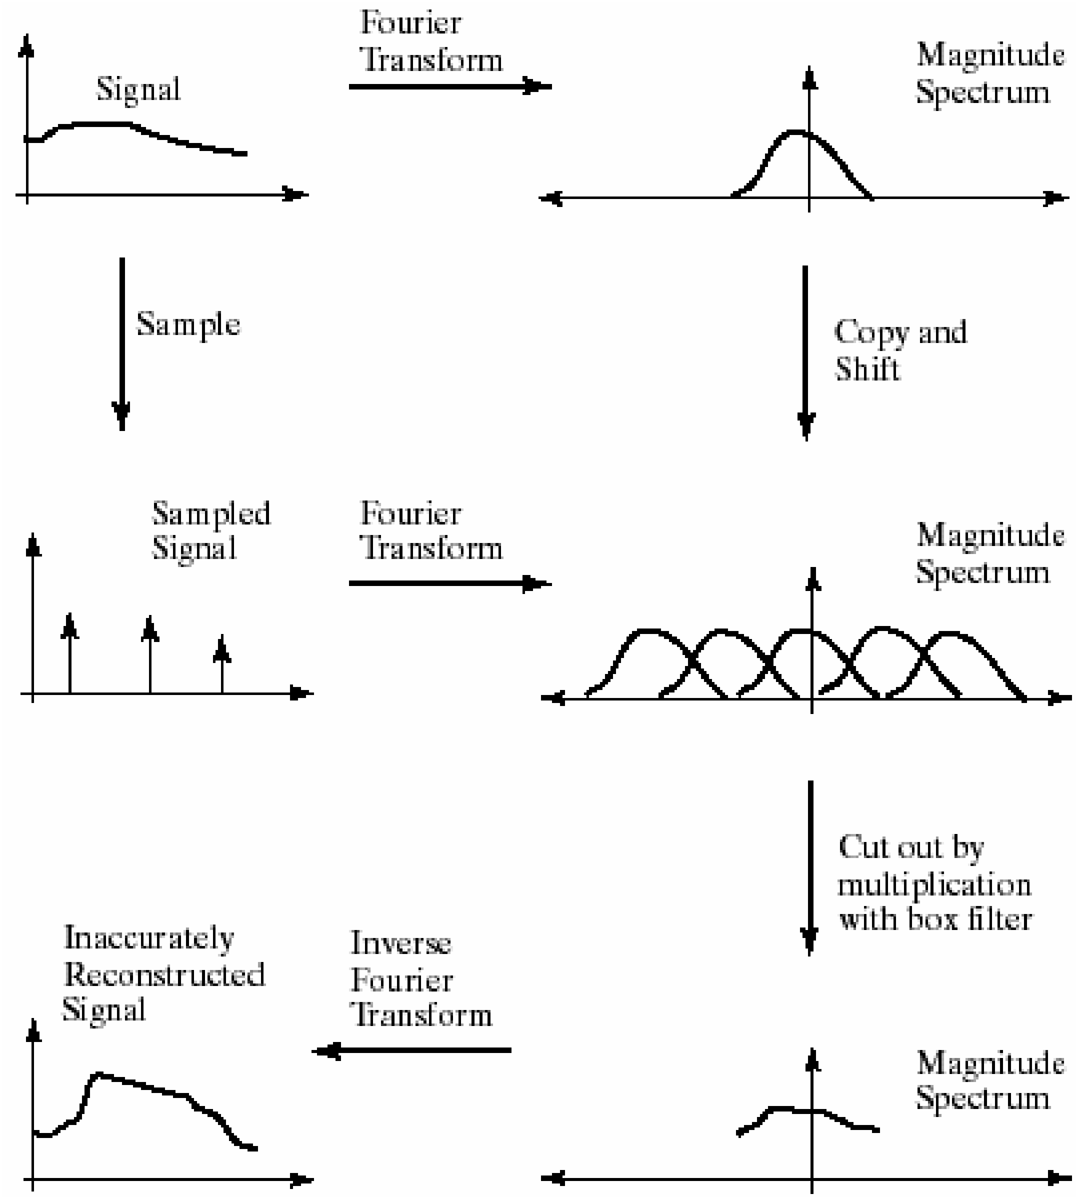
Signal
Fourier
Transform
Magnitude
Spectrum
Sample
Copy and
Shift
Sampled
Signal
Fourier
Transform
Magnitude
Spectrum
Cut out by
multiplication
with box filter
Inaccurately
Reconstructed
Signal
Inverse
Fourier
Transform
Magnitude
Spectrum

Black and white edge falling on color CCD detector

Black and white image (edge)

Detector pixel colors

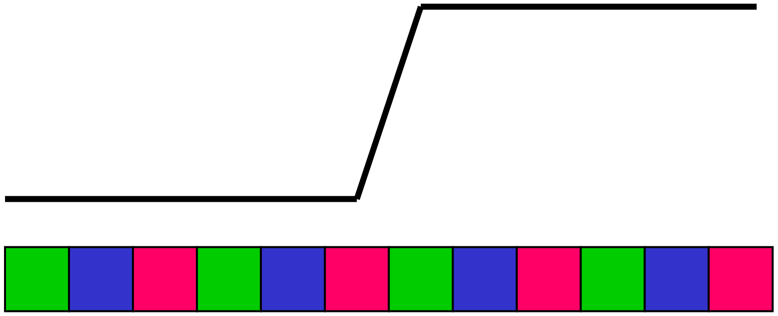

(previous slides were the freq domain interpretation of aliasing. Here’s the spatial domain interpretation.)

Color sampling artifact

Interpolated pixel colors,
for grey edge falling on colored
detectors (linear interpolation).

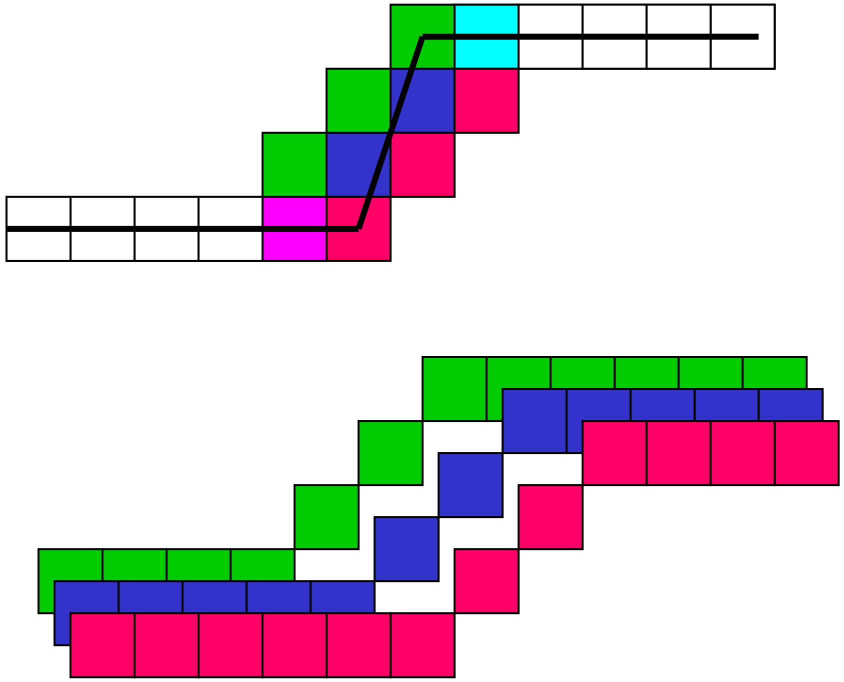

Typical color moire patterns

Blow-up of electronic camera image. Notice spurious colors in the regions of fine detail in the plants.

Color sampling artifacts

How many of you have seen color fringe artifacts from the camera sensor mosaics of cameras you own?

Human Photoreceptors

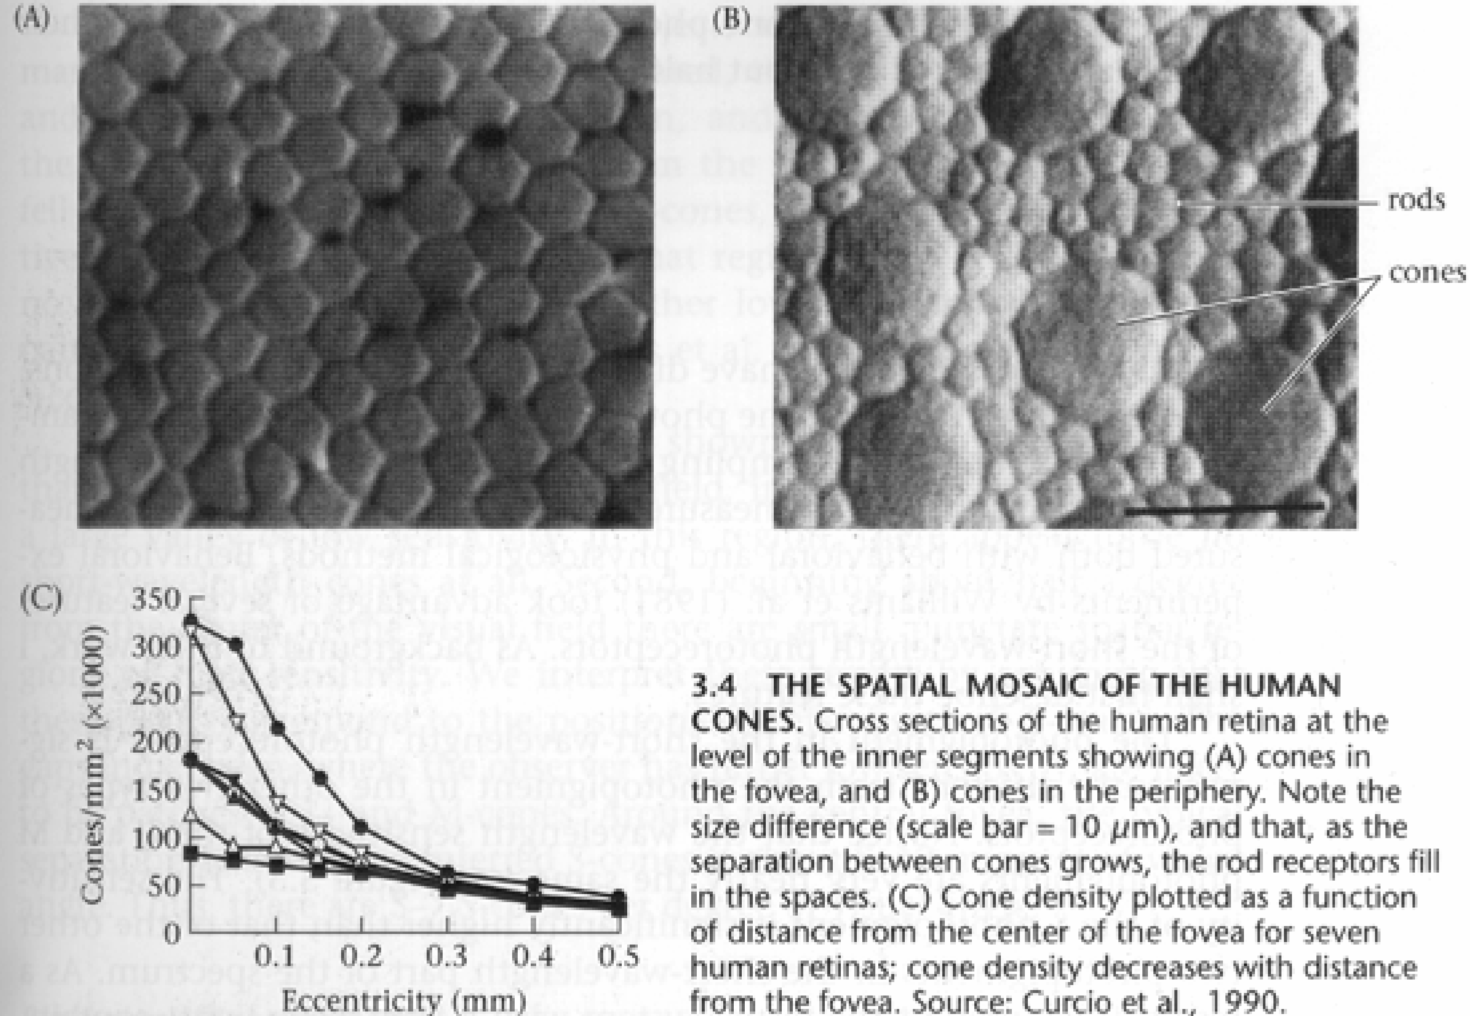

3.4 THE SPATIAL MOSAIC OF THE HUMAN CONES. Cross sections of the human retina at the level of the inner segments showing (A) cones in the fovea, and (B) cones in the periphery. Note the size difference (scale bar = 10 μm), and that, as the separation between cones grows, the rod receptors fill in the spaces. (C) Cone density plotted as a function of distance from the center of the fovea for seven human retinas; cone density decreases with distance from the fovea. Source: Curcio et al., 1990.

(From Foundations of Vision, by Brian Wandell, Sinauer Assoc.)

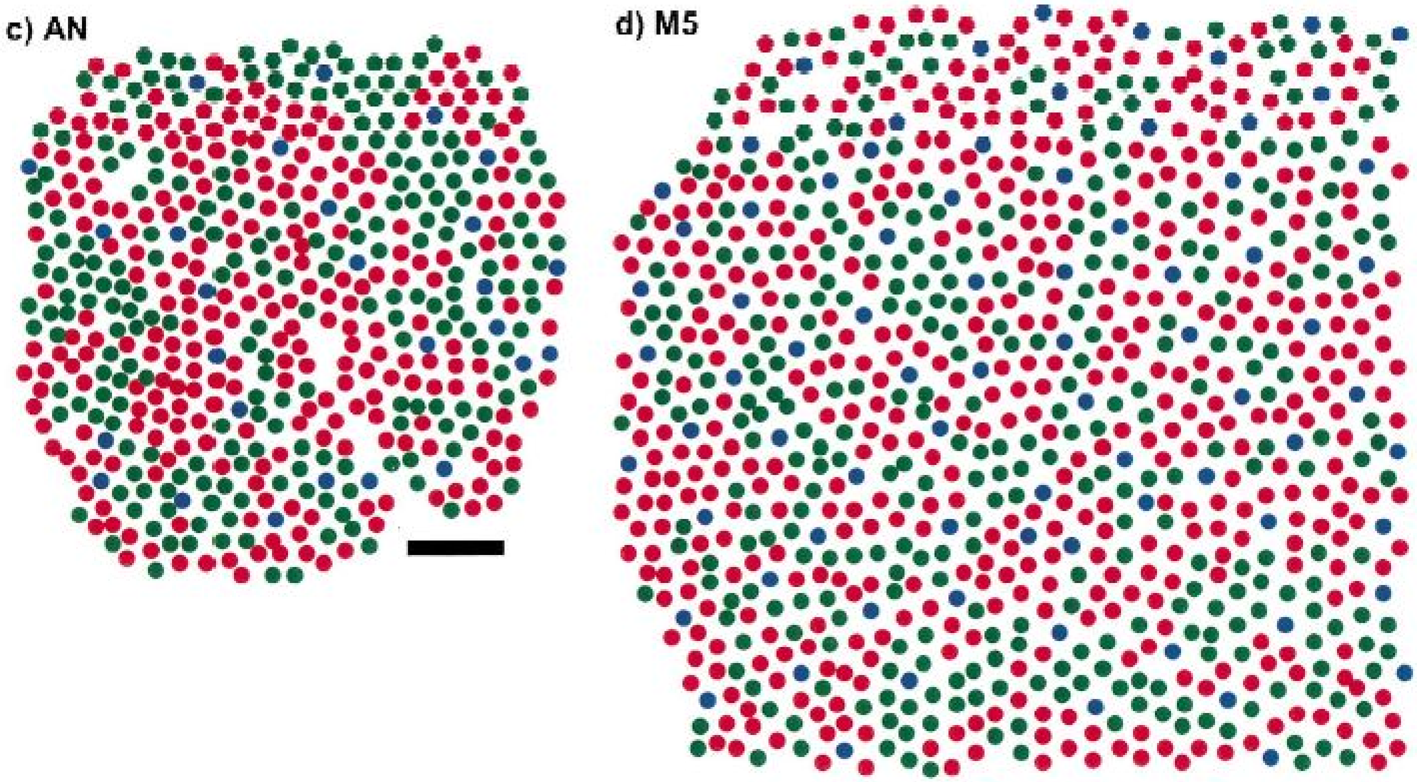

Fig. 2. L, M and S cone mosaics for two humans: JW (a nasal and a temporal location is shown, labeled JWN and JWT, respectively); and AN, and one macaque, M5. L, M and S cones are shown as red, green and blue dots respectively. For JWN, a patch of central cones was not identified due to a capillary that obscured those cones. All mosaics are shown to the same scale. Scale bar = 5 μm.

http://www.cns.nyu.edu/~pl/pubs/Roorda_et_al01.pdf

Have any of you seen color sampling artifacts from the spatially offset color sampling in your own visual systems?

Where I've seen color fringe reconstruction artifacts in my ordinary world

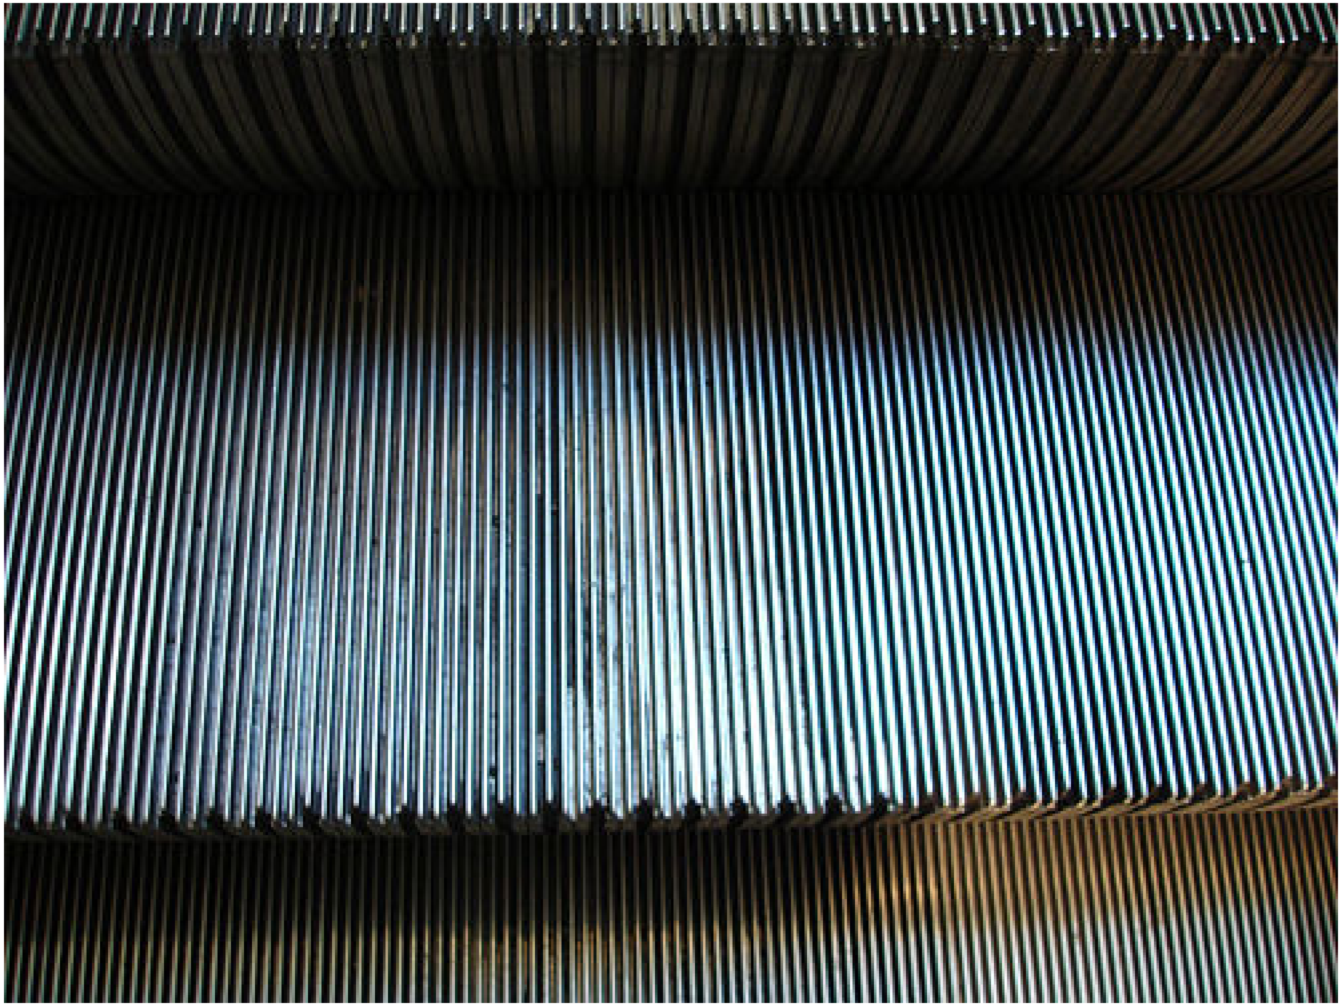

http://static.flickr.com/21/31393422_23013da003.jpg

Brewster's colors—evidence of interpolation from spatially offset color samples

Scale relative to human photoreceptor size: each line covers about 7 photoreceptors.

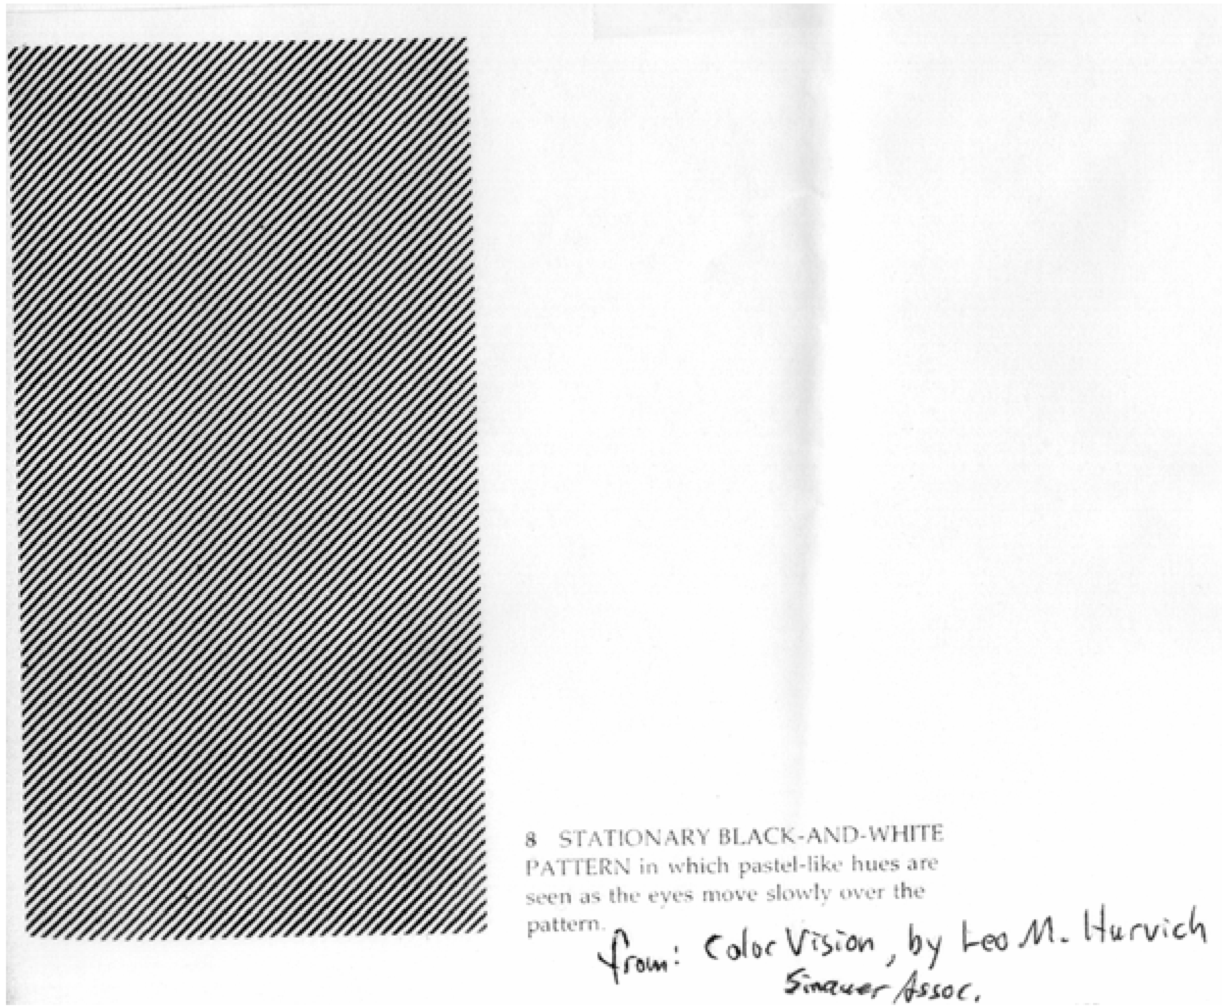

Motivation for median filter interpolation

The color fringe artifacts are obvious; we can point to them. Goal: can we characterize the color fringe artifacts mathematically? Perhaps that would lead to a way to remove them…

R-G, after linear interpolation

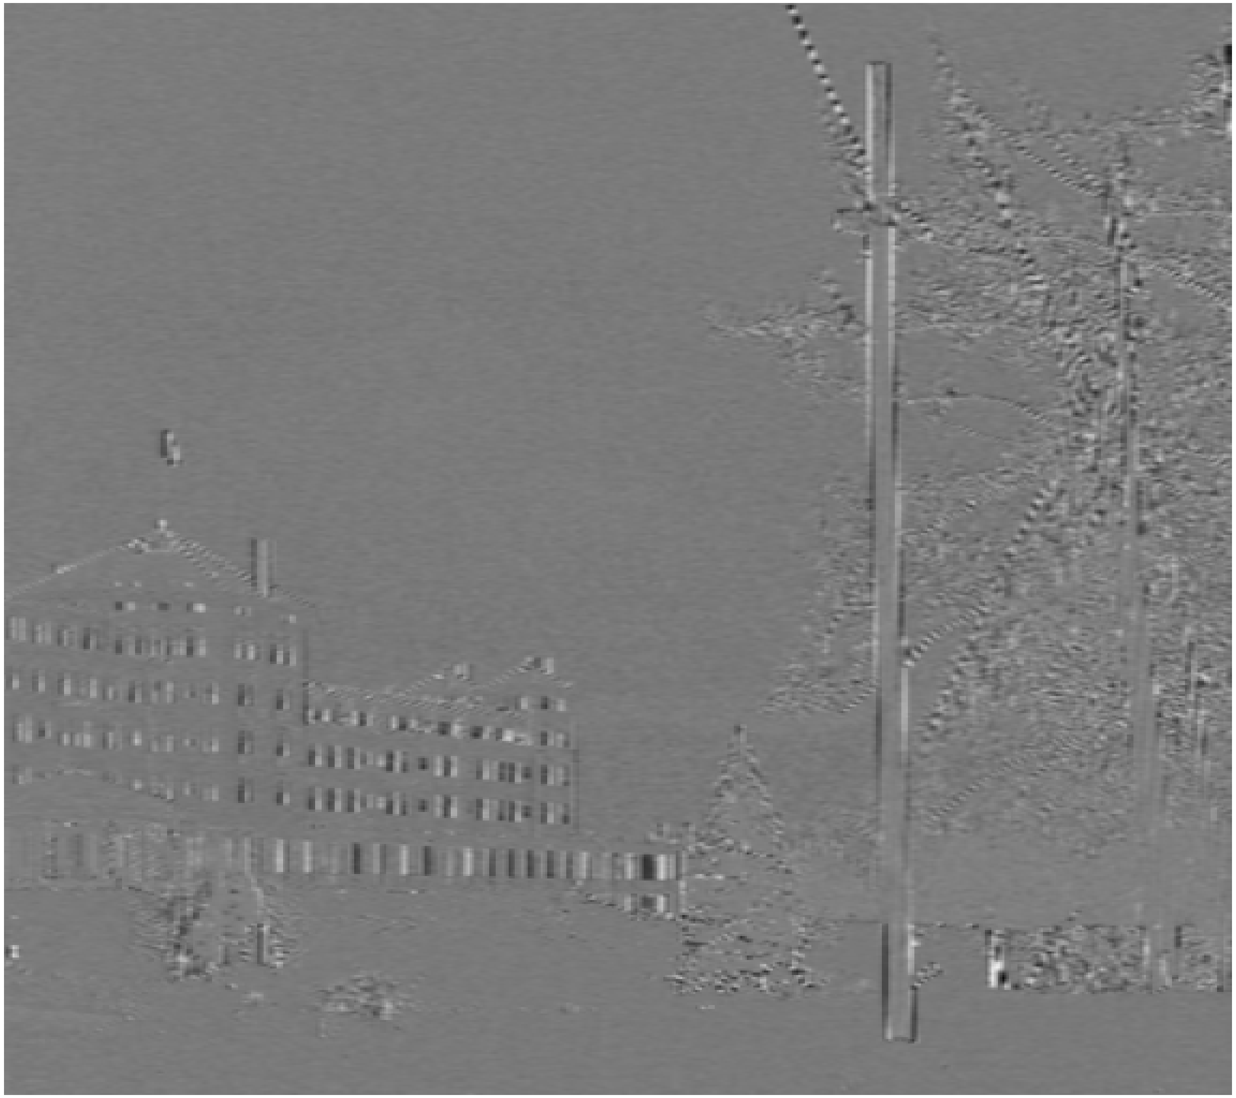

Median filter

Replace each pixel by the median over N pixels (5 pixels, for these examples). Generalizes to “rank order” filters.

In:

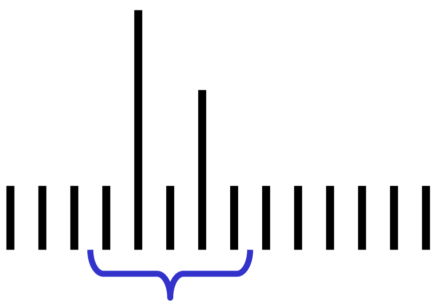

5-pixel neighborhood

Out:

Spike noise is removed

In:

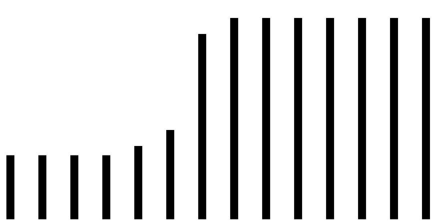

Out:

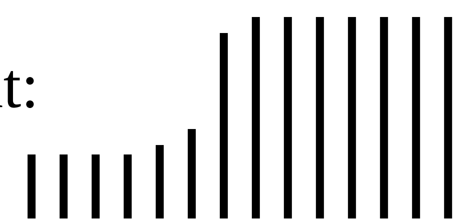

Monotonic edges remain unchanged

Degraded image

Radius 1 median filter

Radius 2 median filter

R – G, median filtered (5x5)

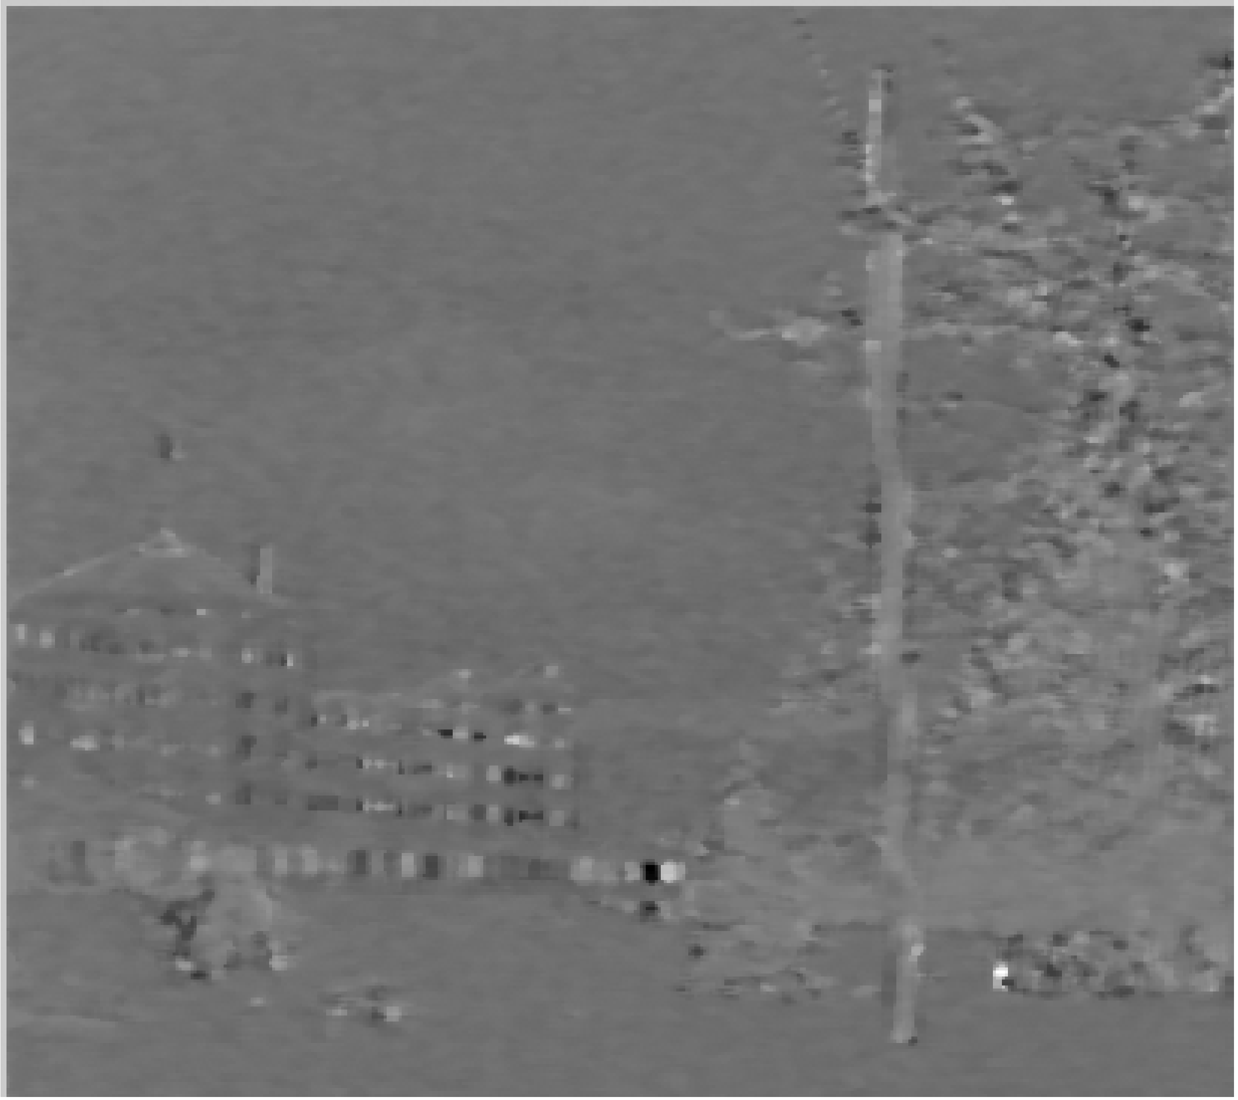

R – G

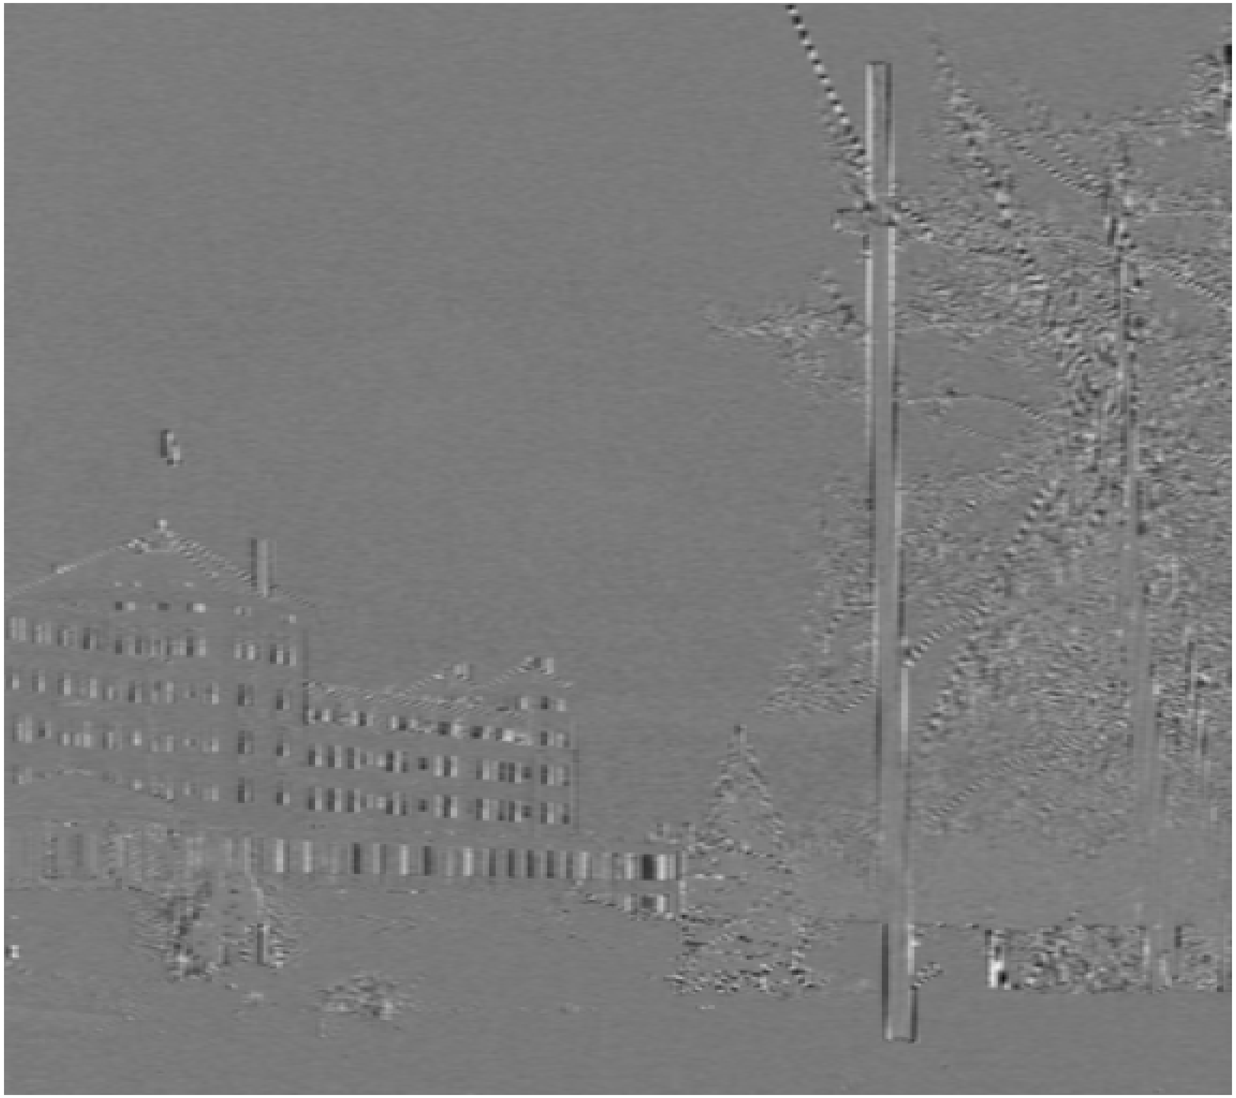

Median Filter Interpolation

- Perform first interpolation on isolated color channels.
- Compute color difference signals.
- Median filter the color difference signal.
- Reconstruct the 3-color image.

Two-color sampling of BW edge

Luminance profile

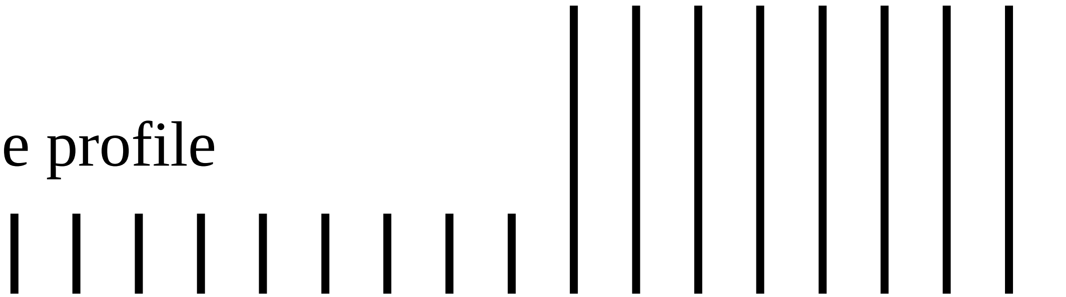

Two-color sampling of BW edge

Luminance profile

True full-color image

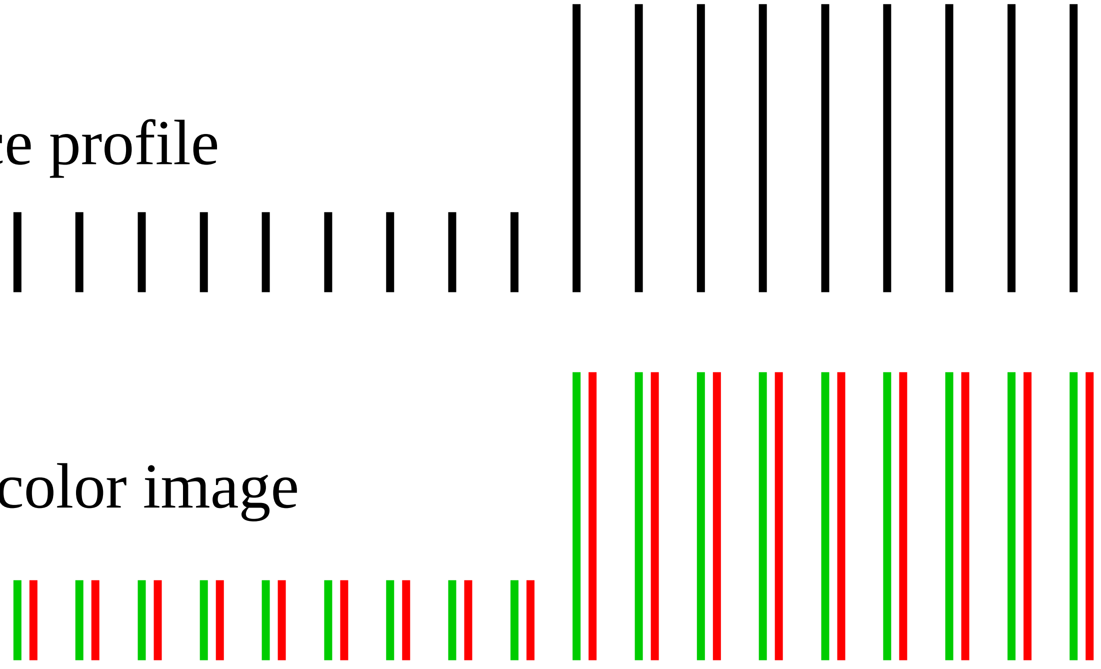

Two-color sampling of BW edge

Sampled data

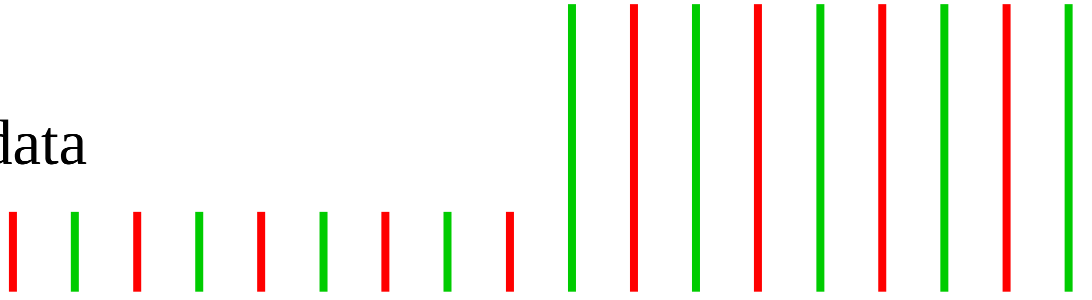

Two-color sampling of BW edge

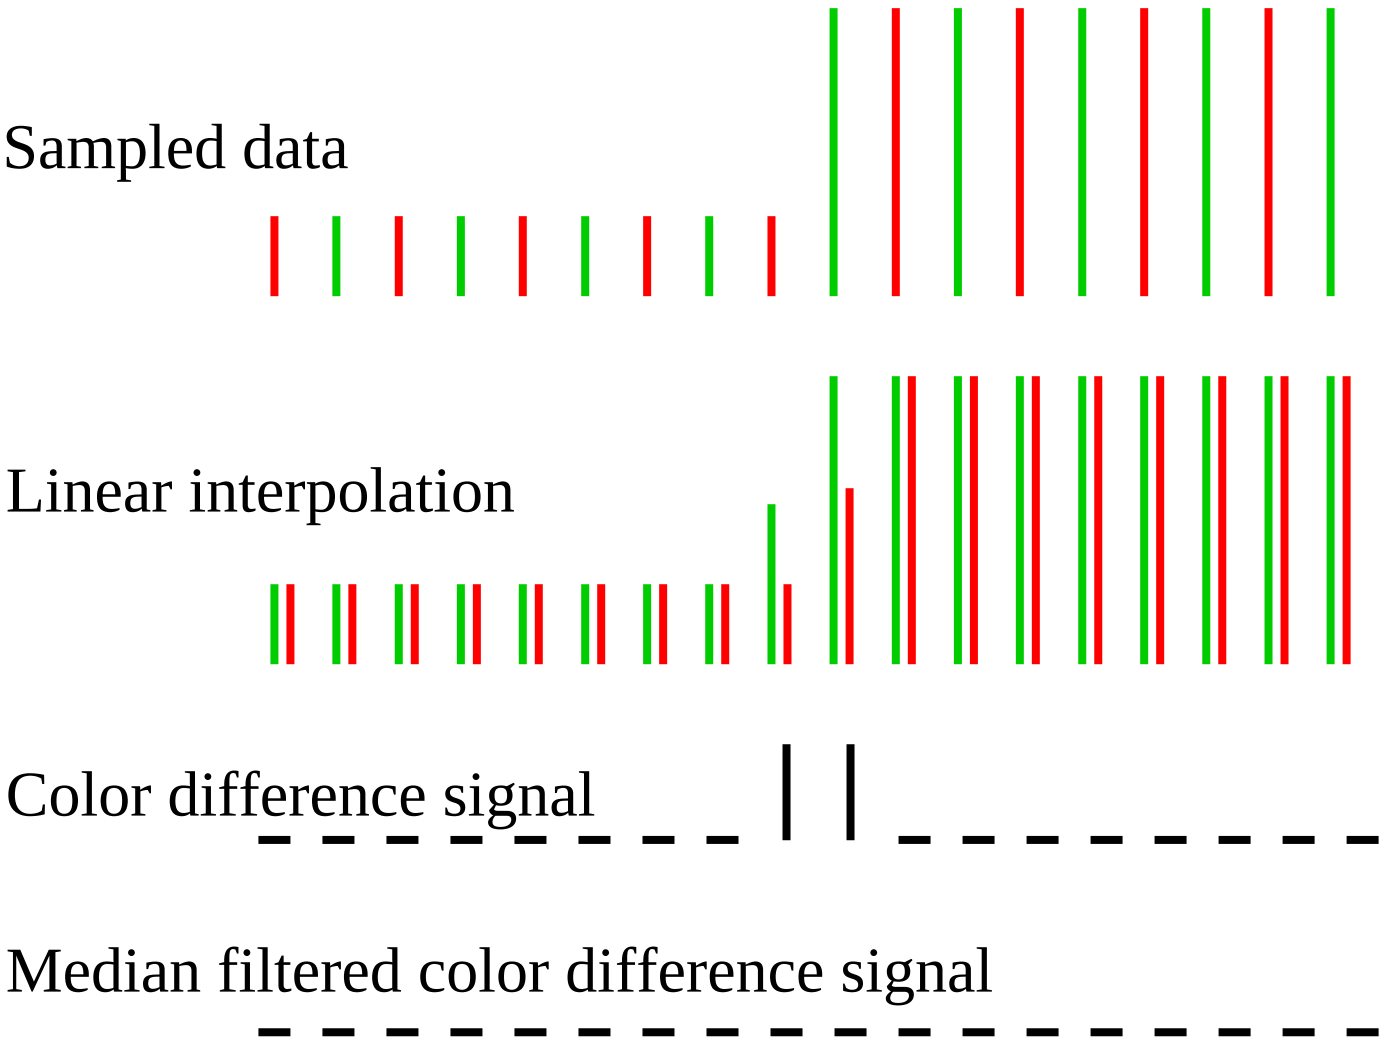

Recombining the median filtered colors

Linear interpolation

Median filter interpolation

Beyond linear interpolation between samples of the same color

- Luminance highs
- Median filter interpolation
- Regression
- Gaussian method
- Regression, including non-linear terms
- Multiple linear regressors

Project ideas

(1) Develop a new color interpolation algorithm

(2) Study the tradeoffs in sensor color choice for image reconstruction:

human vision uses randomly placed, very unsaturated color sensors;

cameras typically use regularly spaced, saturated color sensors.

end