Color

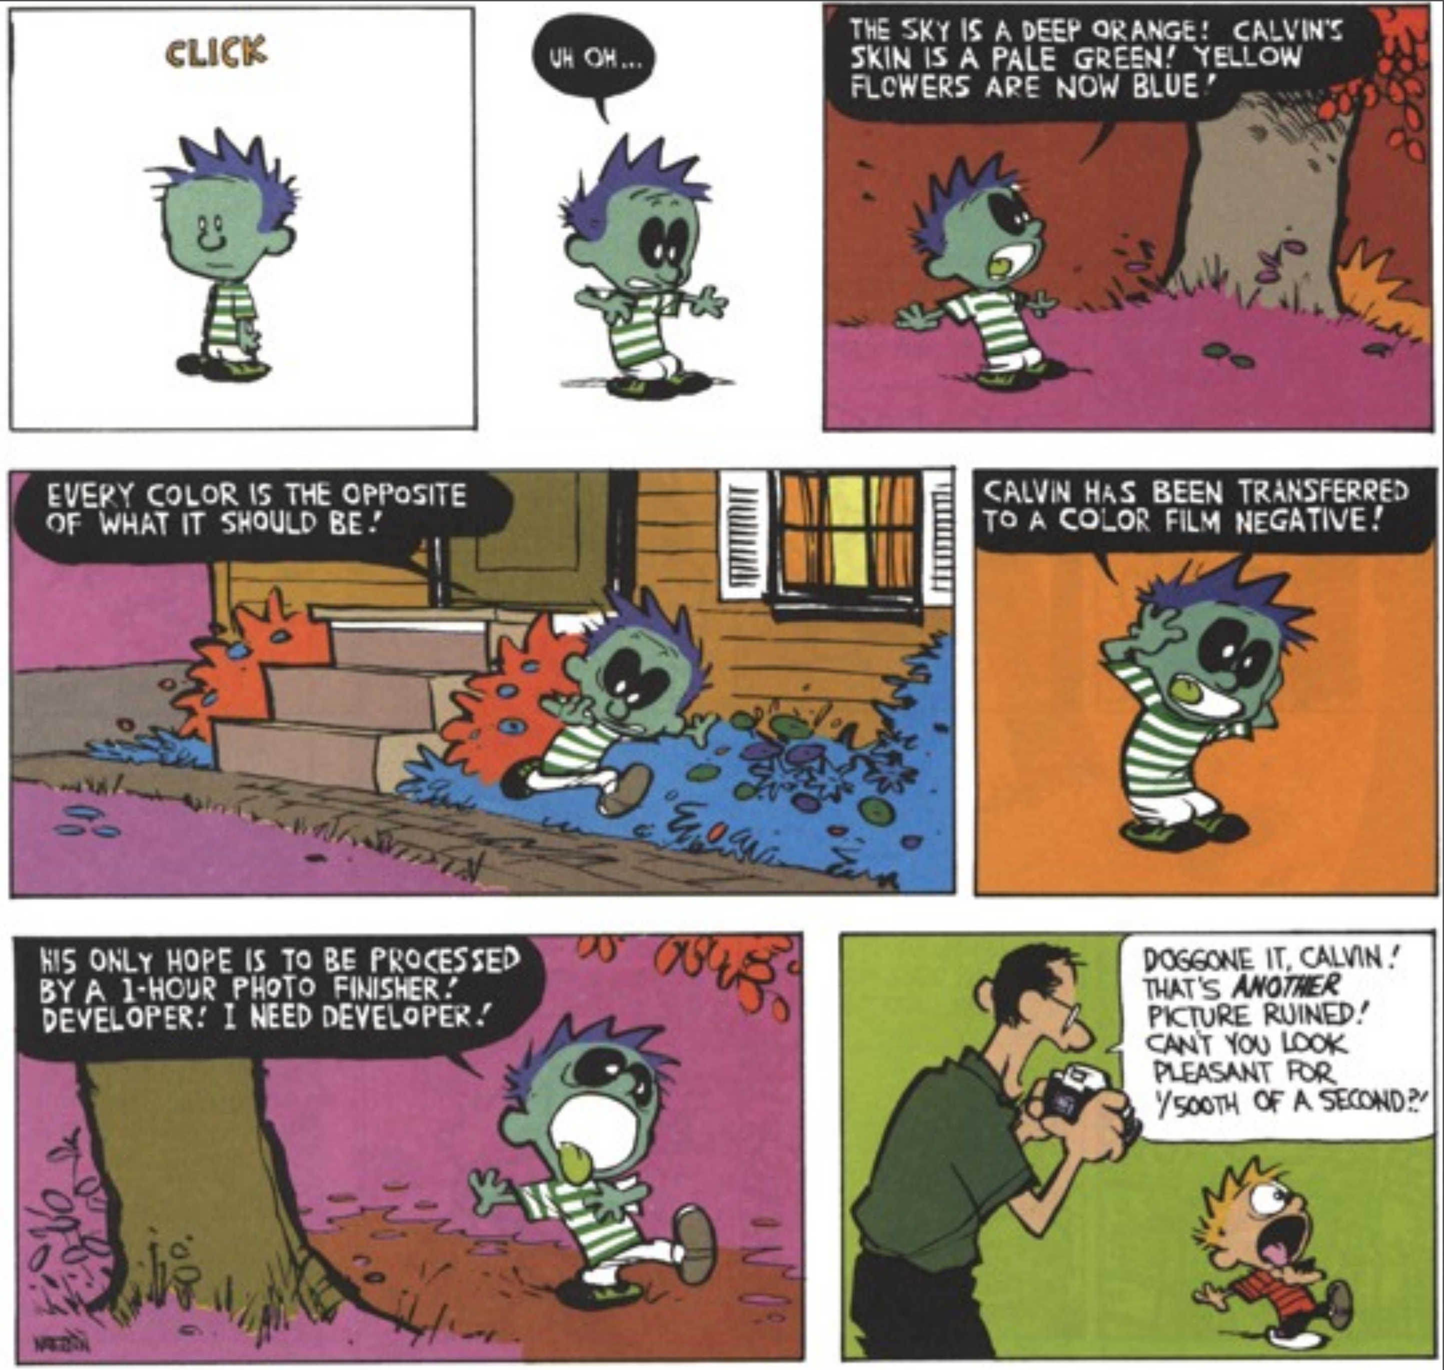
CLICK
UH OH...
THE SKY IS A DEEP ORANGE! CALVIN'S SKIN IS A PALE GREEN! YELLOW FLOWERS ARE NOW BLUE!
EVERY COLOR IS THE OPPOSITE OF WHAT IT SHOULD BE!
CALVIN HAS BEEN TRANSFERRED TO A COLOR FILM NEGATIVE!
HIS ONLY HOPE IS TO BE PROCESSED BY A 1-HOUR PHOTO FINISHER! DEVELOPER! I NEED DEVELOPER!
DOGGONE IT, CALVIN! THAT'S *ANOTHER* PICTURE RUINED! CAN'T YOU LOOK PLEASANT FOR 1/500TH OF A SECOND?!

Color & Demosaicing

Frédo Durand
MIT- EECS

Many slides courtesy of Victor Ostromoukhov, Leonard McMillan, Bill Freeman

Oliver Sacks

- Great books about people with neurological disorders (typically missing one part of the brain)
- http://www.oliversacks.com/
- http://www.oliversacks.com/island.htm
- http://www.oliversacks.com/mars.htm
- http://www.oliversacks.com/hat.htm

Plan

- Wrap up CIE XYZ
- More human color perception
- Extra quick overview of JPEG
- Demosaicing
- Other color imaging strategies

CIE XYZ : recap

- The most widely recognized color space
- Linear transform of the original tristimuls curves
- Y corresponds to brightness (1924 CIE standard photometric observer)
- No negative value of matching curve
- But no physically-realizable primary
(negative values in primary rather than in matching curve)

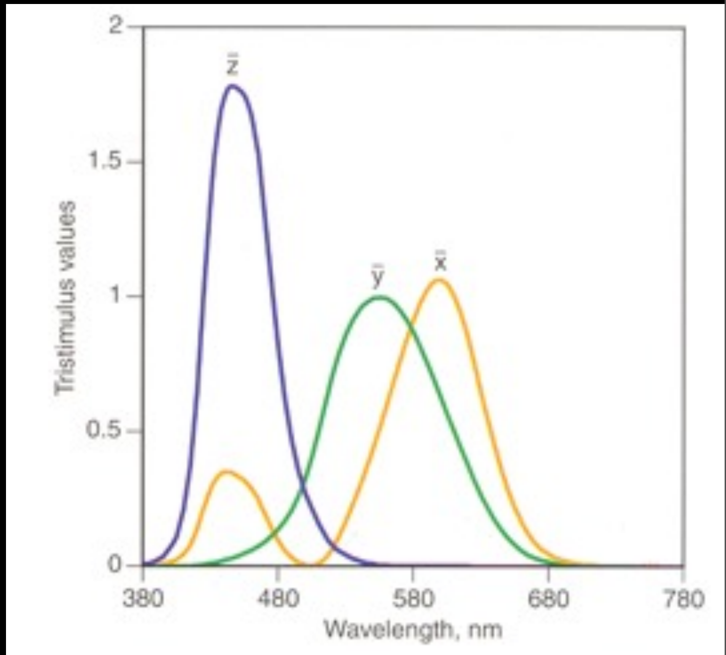

The 1931 standard observer, as it is usually shown.

CIE color space

- 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)

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

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

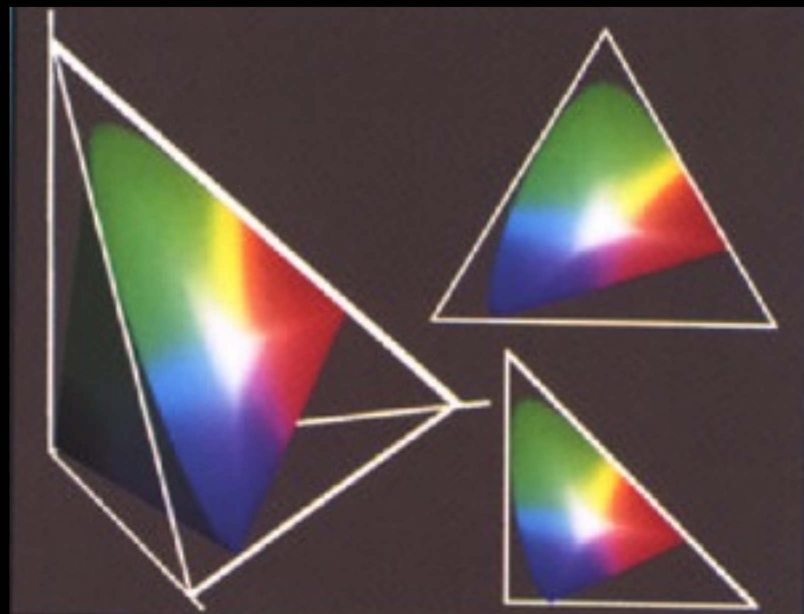

XYZ to RGB & back

- e.g. http://www.brucelindbloom.com/index.html?Eqn_RGB_XYZ_Matrix.html
- sRGB to XYZ

0.412424	0.212656	0.0193324
0.357579	0.715158	0.119193
0.180464	0.0721856	0.950444

XYZ to sRGB

3.24071	-0.969258	0.0556352
-1.53726	1.87599	-0.203996
-0.498571	0.0415557	1.05707

- Adobe RGB to XYZ

0.576700	0.297361	0.0270328
0.185556	0.627355	0.0706879
0.188212	0.0752847	0.991248

XYZ to Adobe RGB

2.04148	-0.969258	0.0134455
-0.564977	1.87599	-0.118373
-0.344713	0.0415557	1.01527

- NTSC RGB to XYZ

0.606734	0.298839	0.000000
0.173564	0.586811	0.0661196
0.200112	0.114350	1.11491

XYZ to NTSC RGB

1.91049	-0.984310	0.0583744
-0.532592	1.99845	-0.118518
-0.288284	-0.0282980	0.898611

- http://dba.med.sc.edu/price/irf/Adobe_tg/manage/images/gamuts.jpg
- http://www.petrvodna
ColorSpace.htm

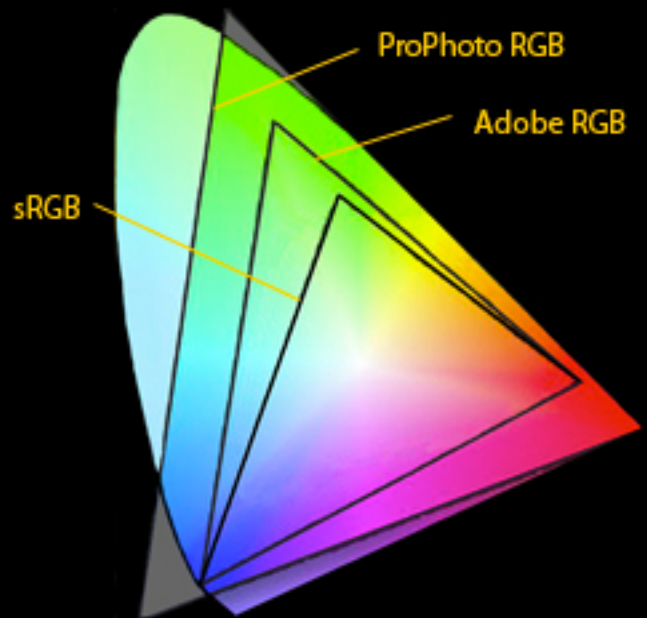

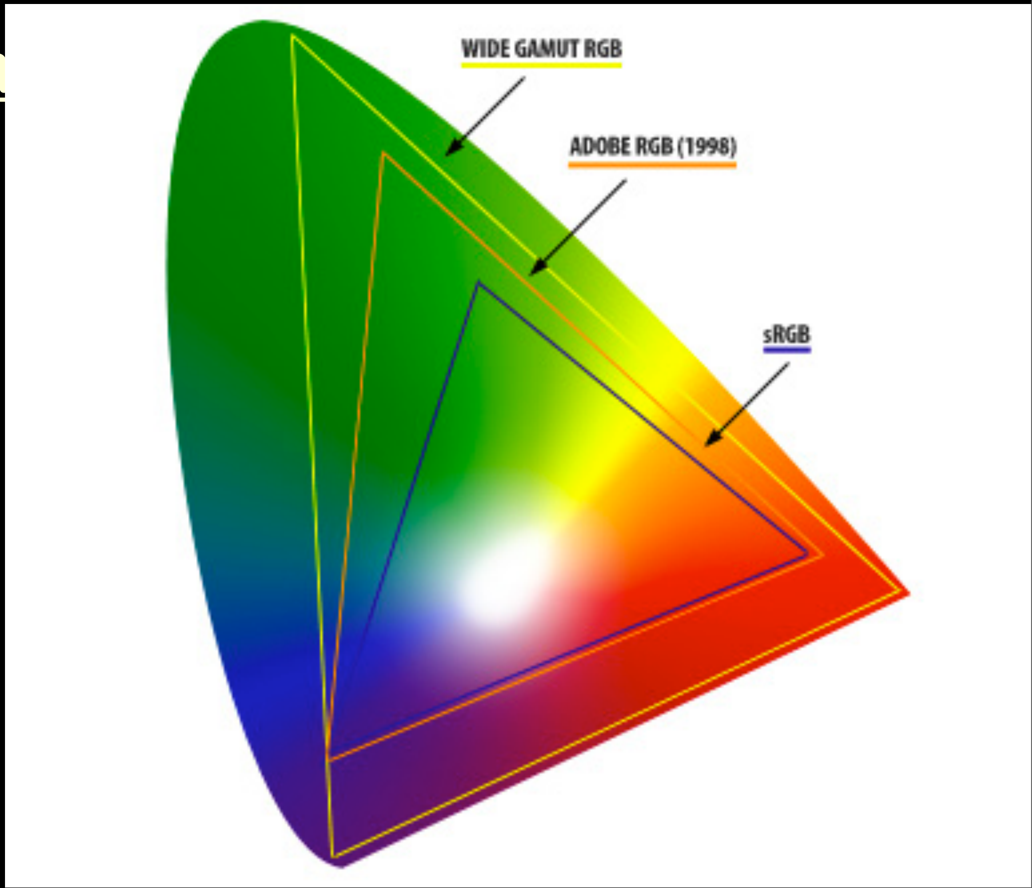

- http://www.brucelindbloom.com/index.html?

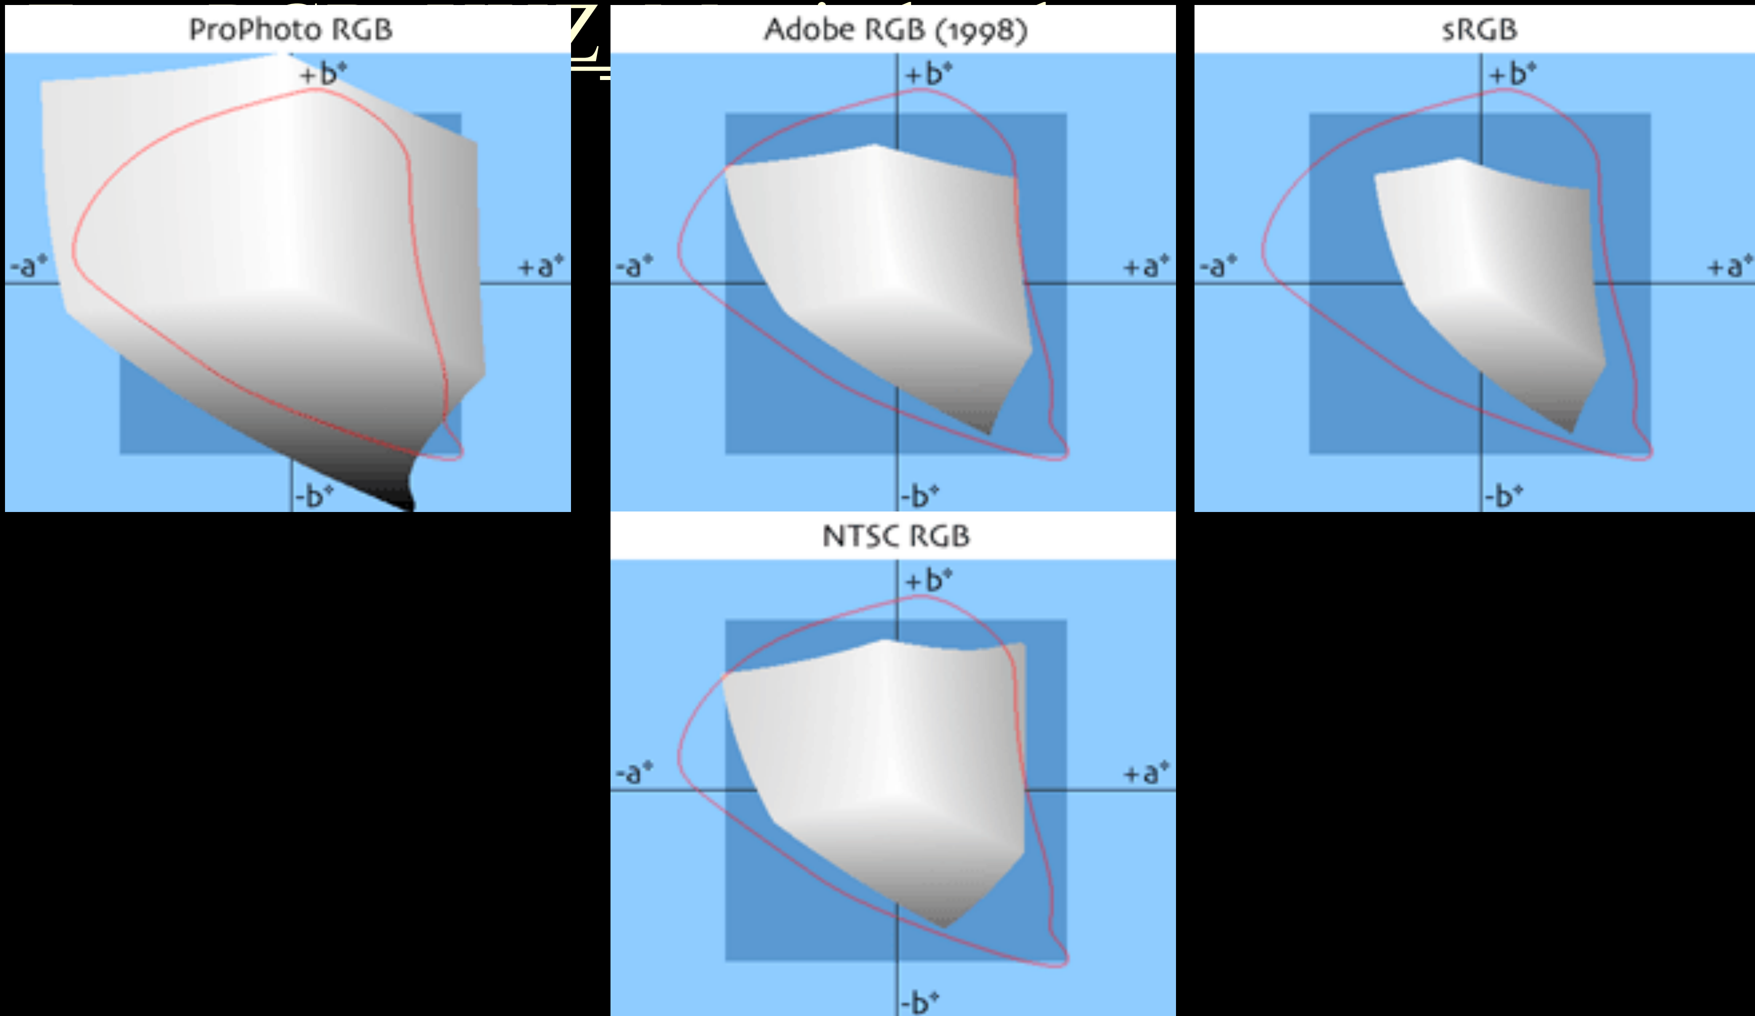

Chromaticity diagrams

- 3D space are tough to visualize
- Usually project to 2D for clarity
- Chromaticity diagram:
  - normalize against X + Y + Z:

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

  - Perspective projection to plane X+Y+Z=1

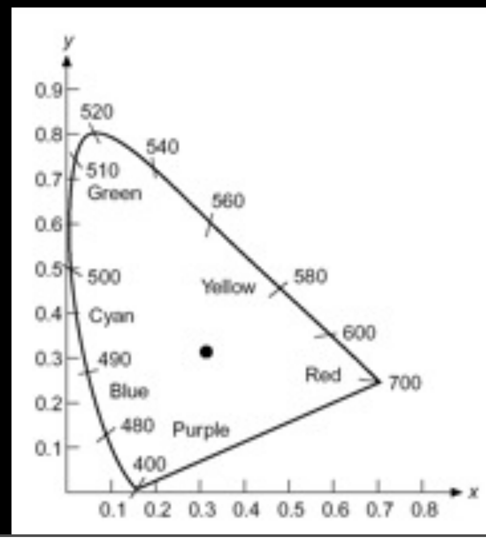

Color gamut

- Given 3 primaries
- The realisable chromaticities lay in the triangle in xy chromaticity diagram

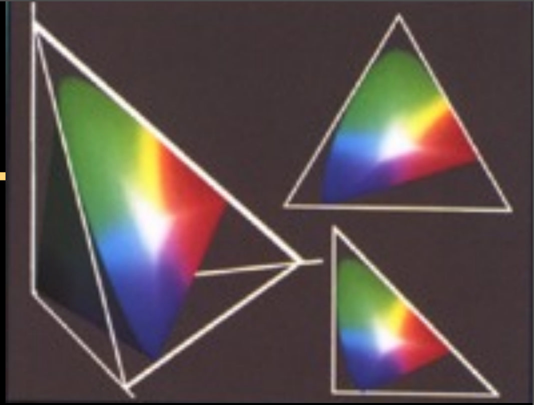

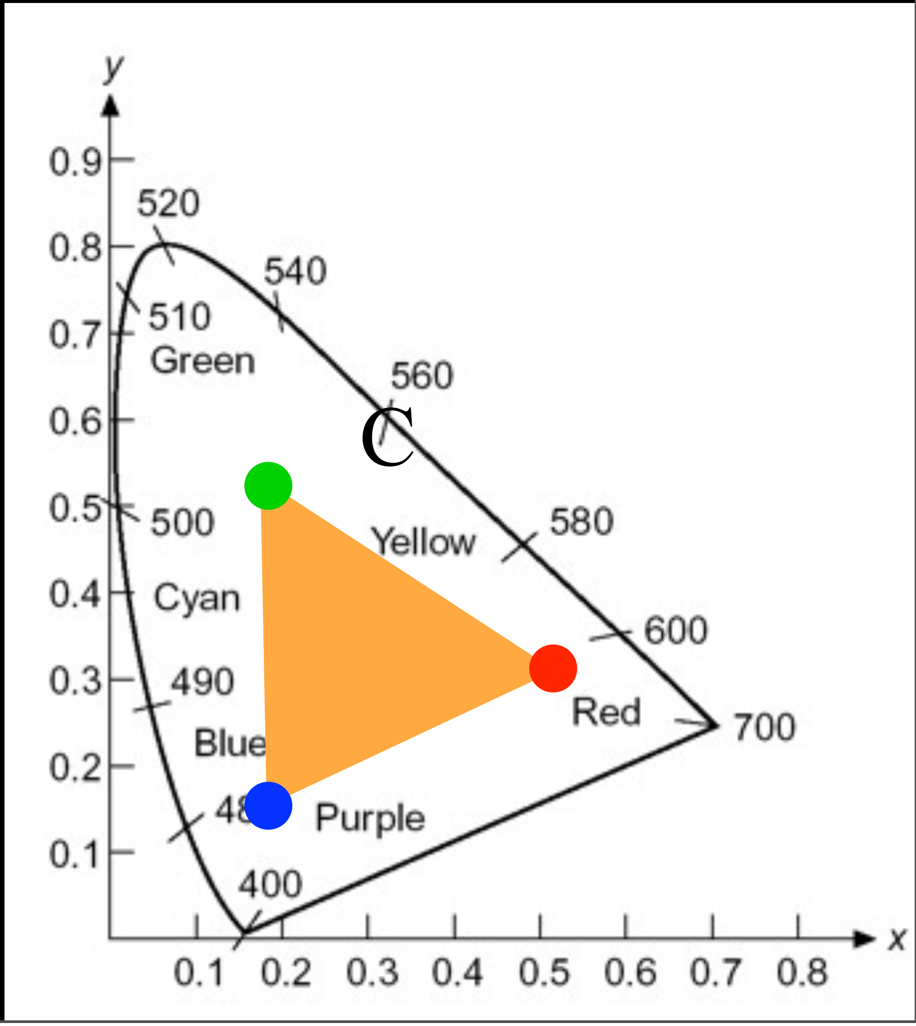

Questions?

VIII. Philipp Otto Runge, *Colour Sphere*, 1809, Hamburg Kunsthalle.

Plan

- Wrap up CIE XYZ
- **More human color perception**
- Extra quick overview of JPEG
- Demosaicing
- Other color imaging strategies

Remember von Helmholtz

- Colors as relative responses (ratios)

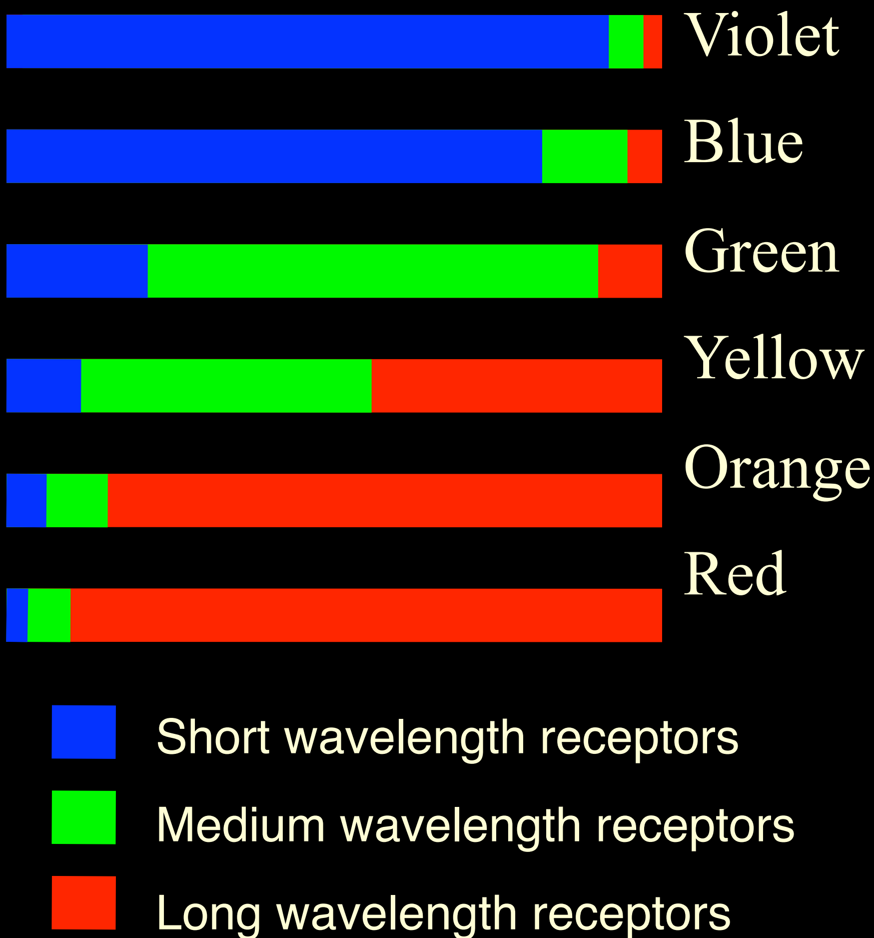

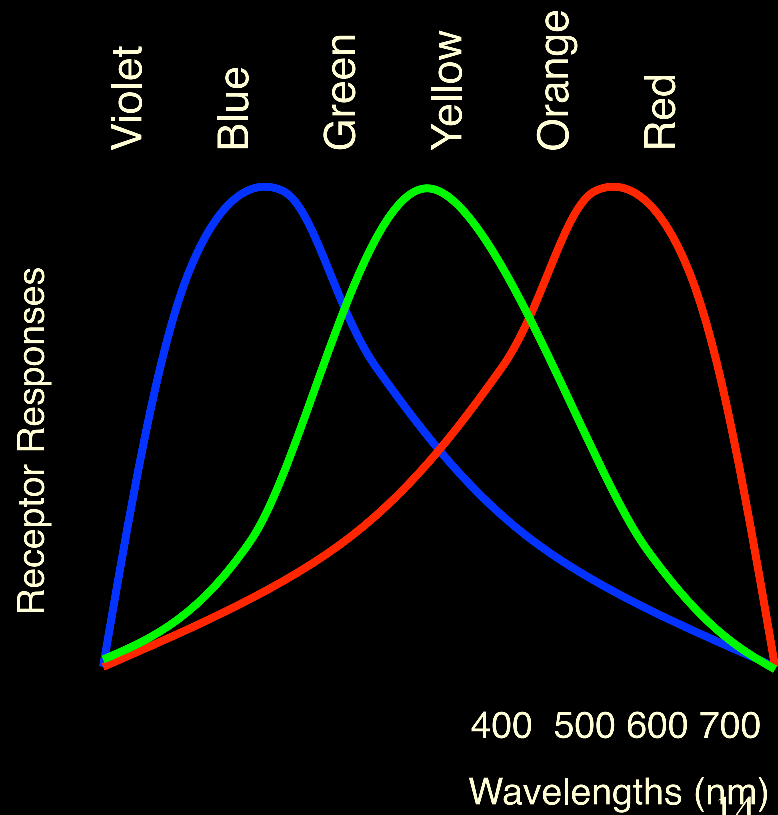

Hering 1874: Opponent Colors

- Hypothesis of 3 types of receptors: Red/Green, Blue/Yellow, Black/White
- Explains well several visual phenomena

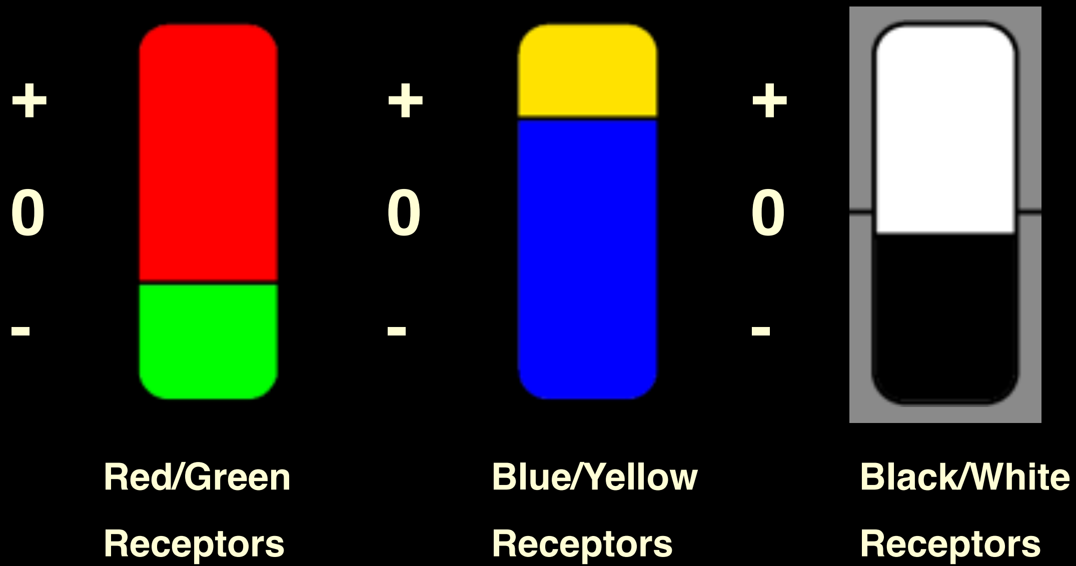

e.g. Spanish Castle illusion (pset 0)

- Duplicate layer

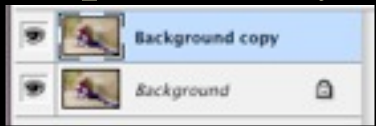

- Desaturate bottom layer (ctrl shift U)

- Switch to Lab color mode (don't flatten)

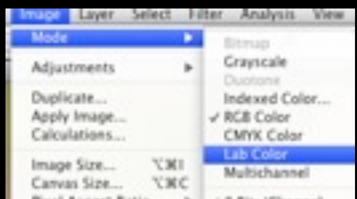

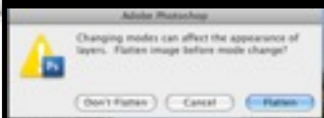

- Select Lightness channel of 2nd layer

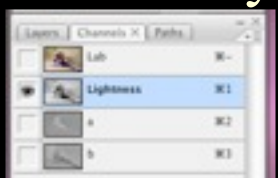

- Fill with 60% grey

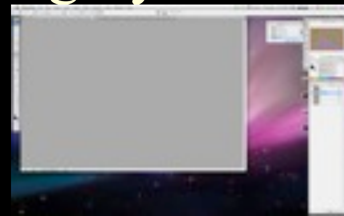

- Select all channels and invert (ctrl I)

- Add a black dot in the center

- Potential plus: blur color version, increase saturation

Dual Process Theory

- The input is LMS
- The output has a different parameterization:
  - Light-dark
  - Blue-yellow
  - Red-green

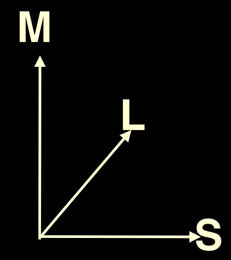

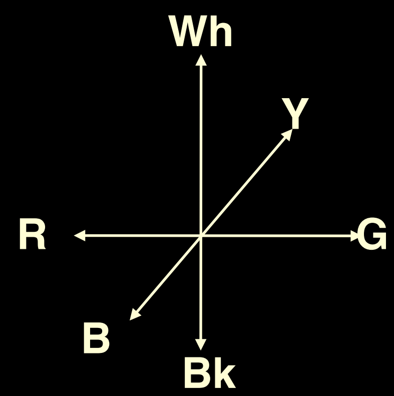

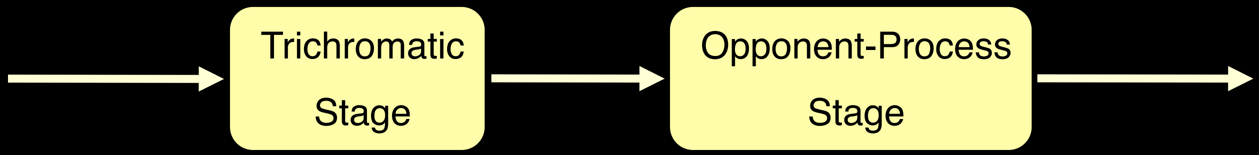

Color opponents wiring

- Sums for brightness
- Differences for color opponents
- At the end, it's just a 3x3 matrix compared to LMS

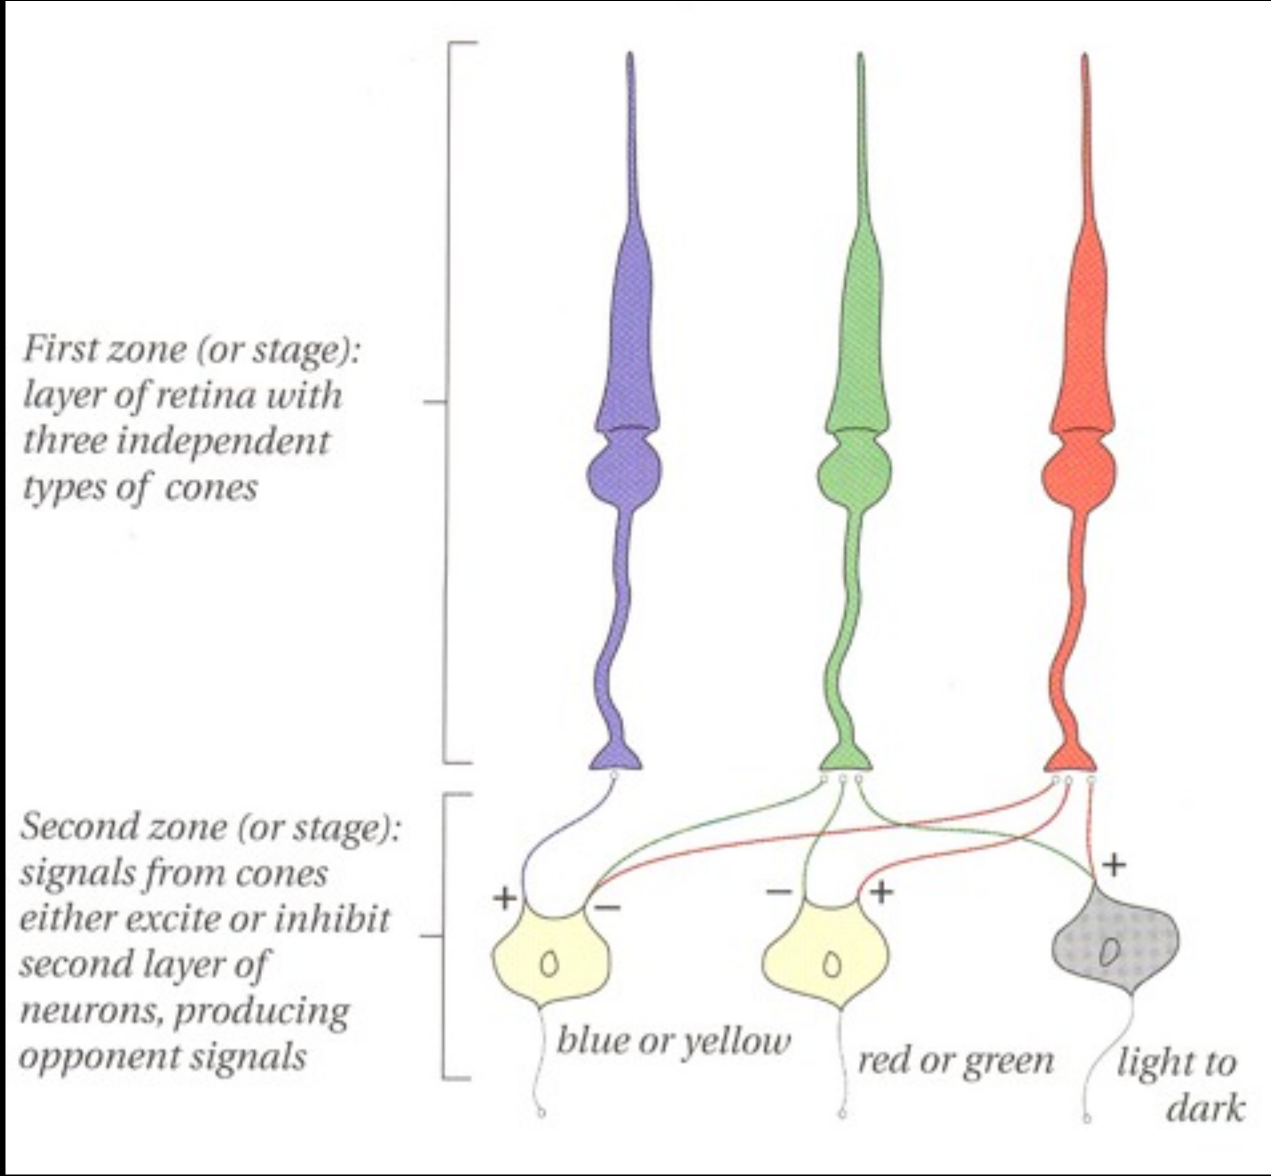

Color reparameterization

- The input is LMS
- The output has a different parameterization:
  - Light-dark
  - Blue-yellow
  - Red-green
- A later stage may reparameterize:
  - Brightness or Luminance or Value
  - Hue
  - Saturation

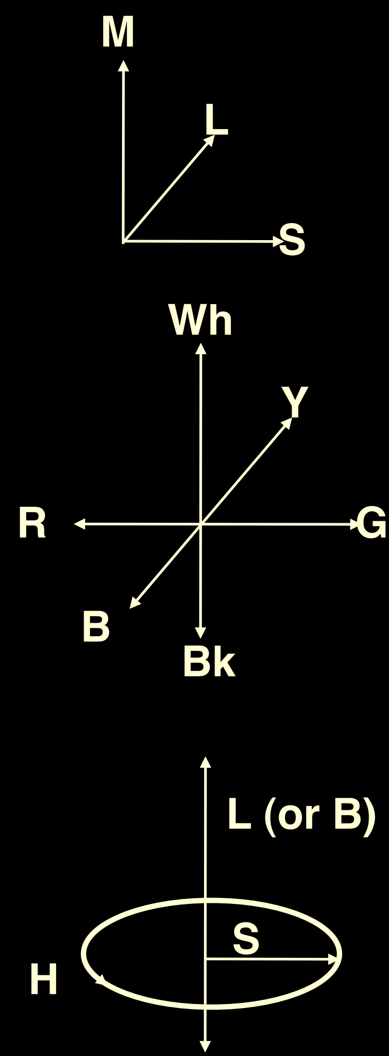

<table>
<tr><th>PHYSICS</th><th>PSYCHOPHYSICS</th><th colspan="2">PSYCHOLOGY</th></tr>
<tr><td>Light Source</td><td>Source x Eye</td><td colspan="2">Color Perception</td></tr>
<tr><td>Radiant energy
Spectral composition
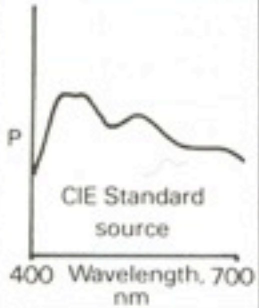
</td><td>Luminous energy
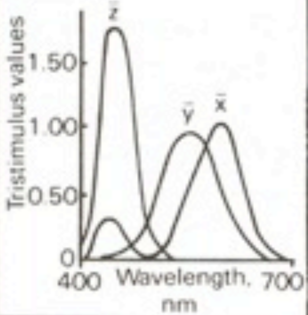
</td><td>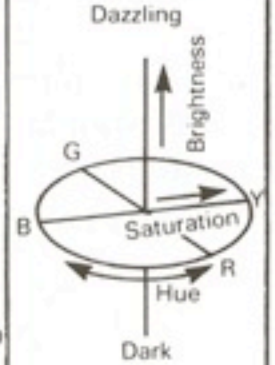
</td><td></td></tr>
<tr><td>Characteristics of Radiant Energy</td><td>Characteristics of Luminous Energy</td><td>Attributes of Color</td><td>Corresponding Modes of Appearance</td></tr>
<tr><td>Radiant flux
Radiance
Irradiance
Radiant reflectance
Radiant transmittance</td><td>Luminous flux
Luminance
Illuminance
Luminous reflectance
Luminous transmittance</td><td>Brightness</td><td>Aperture
Illuminant
Illumination
Object modes:
Surface
Volume</td></tr>
<tr><td>Spectral distribution
(Relative spectral composition, quality)
Radiant purity</td><td>Chromaticity
Dominant wavelength (or complementary)
Purity
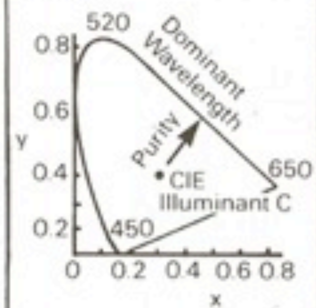
</td><td>Chromaticness
Hue
Saturation</td><td>Attributes of modes of appearance:
1. Brightness (or lightness)
2. Hue
3. Saturation
4. Size
5. Shape
6. Location
7. Flicker
8. Sparkle
9. Transparency
10. Glossiness
11. Luster</td></tr>
</table>

Hue Saturation Value

- Value: from black to white
- Hue: dominant color (red, orange, etc)
- Saturation: from gray to vivid color
- HSV double cone
  - careful, not a cleanly defined color space

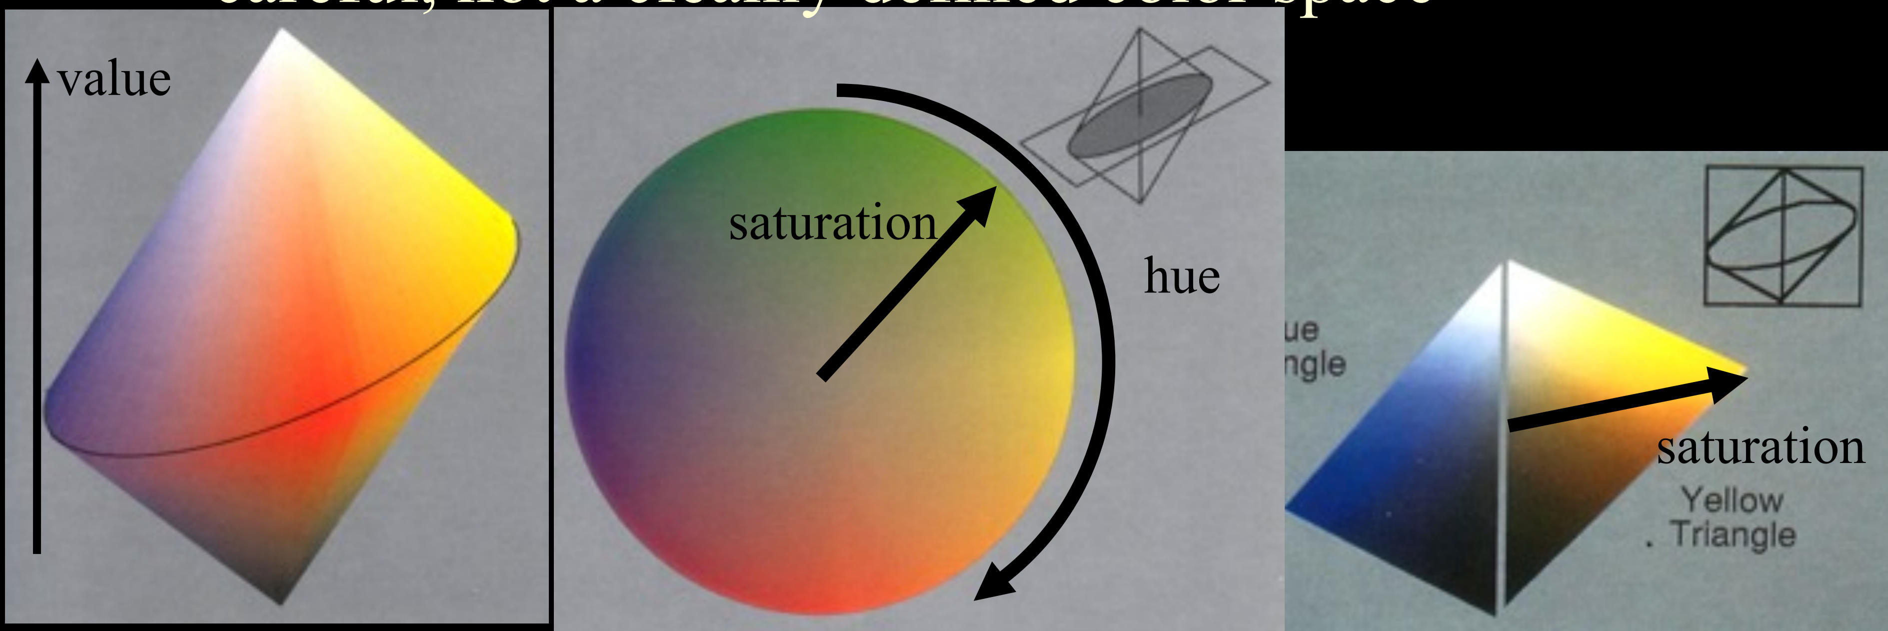

Questions?

Perceptual color difference

- In color space CIE-XYZ, the perceived distance between colors is not equal everywhere
- Can be represented by ellipses of perceived differences (set of colors that look no more different than a given threshold)
- Measured by MacAdam
- Same for all linear color spaces (RBG, LMS, etc.)

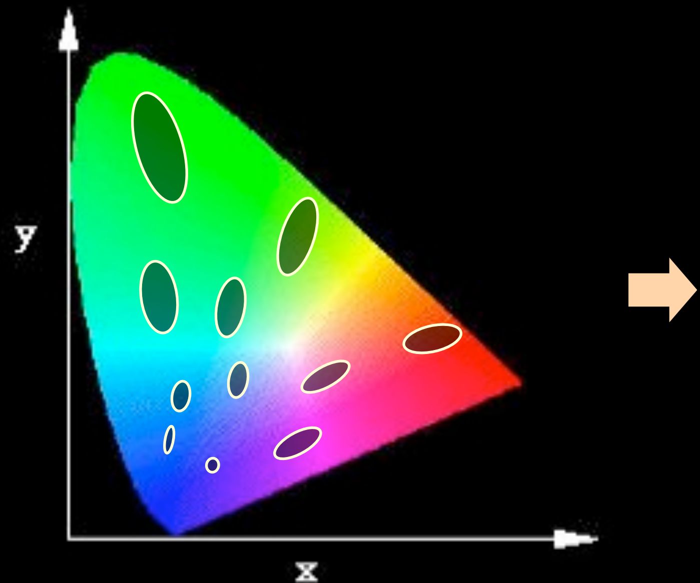

Source: [Wyszecki and Stiles '82]

Perceptually Uniform Space: MacAdam

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

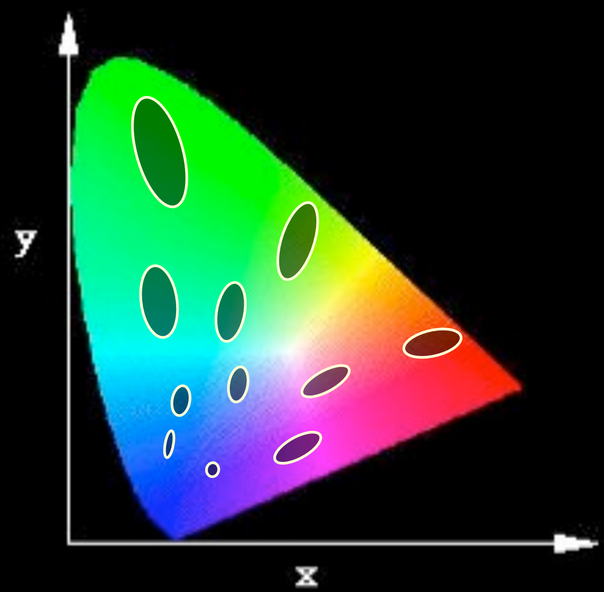

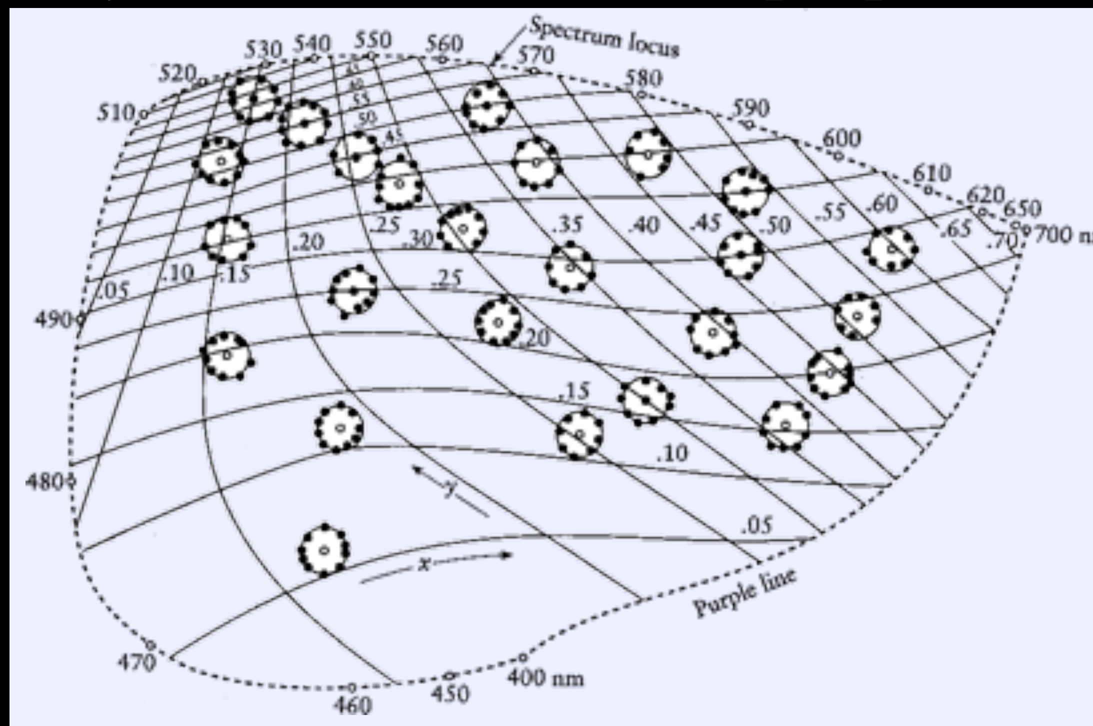

Source: [Wyszecki and Stiles '82]

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

- The reference perceptually uniform color space
- L: lightness
- a and b: color opponents
- Note the cubic roots

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

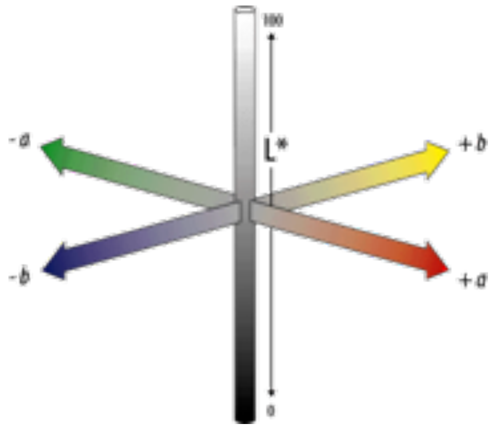

$$L = 25\left(100\frac{Y}{Y_0}\right)^{1/3} - 16$$

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

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

Source: [Wyszecki and Stiles '82]

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

- Euclidean distance between two color

$$\sqrt{(L_1 - L_2)^2 + (a_1 - a_2)^2 + (b_1 - b_2)^2}$$

- correlates well with ability to discriminate these two colors
  - a human can tell the difference between two colors if their CIE L*a*b* distance is greater than 2
  - called “just noticeable difference”

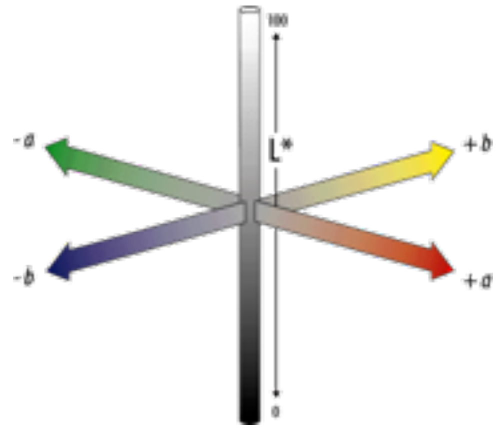

$$L = 25\left(100\frac{Y}{Y_0}\right)^{1/3} - 16$$

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

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

Source: [Wyszecki and Stiles ’82]

Question?

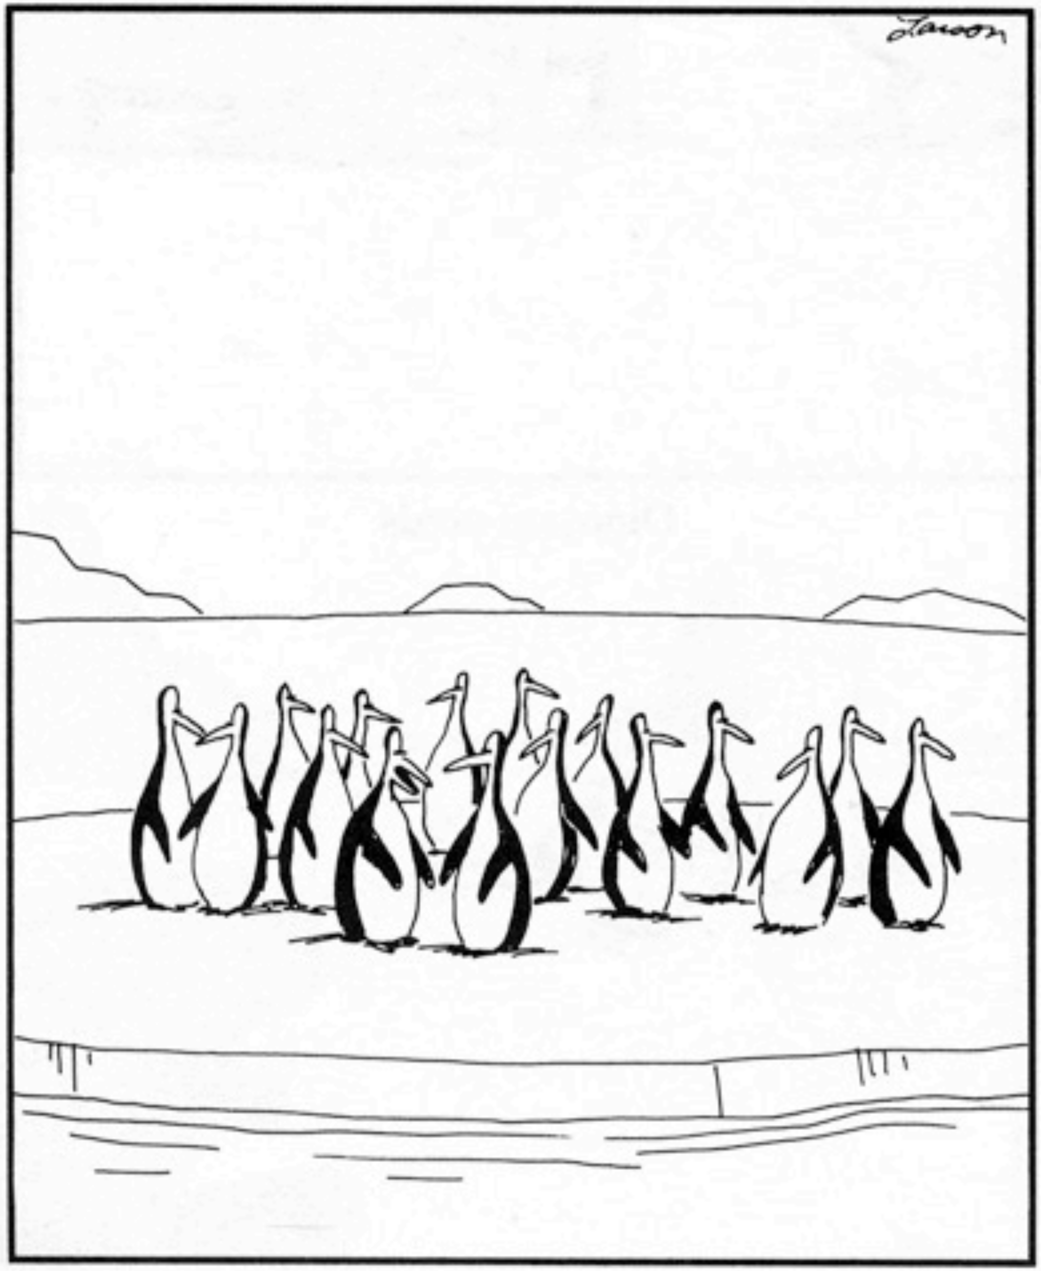

"Well, that's an interesting bit of trivia — I guess I *do* only dream in black and white."

Contrast Sensitivity

- Sine Wave grating
- What contrast is necessary to make the grating visible?
- How does it change with spatial frequency

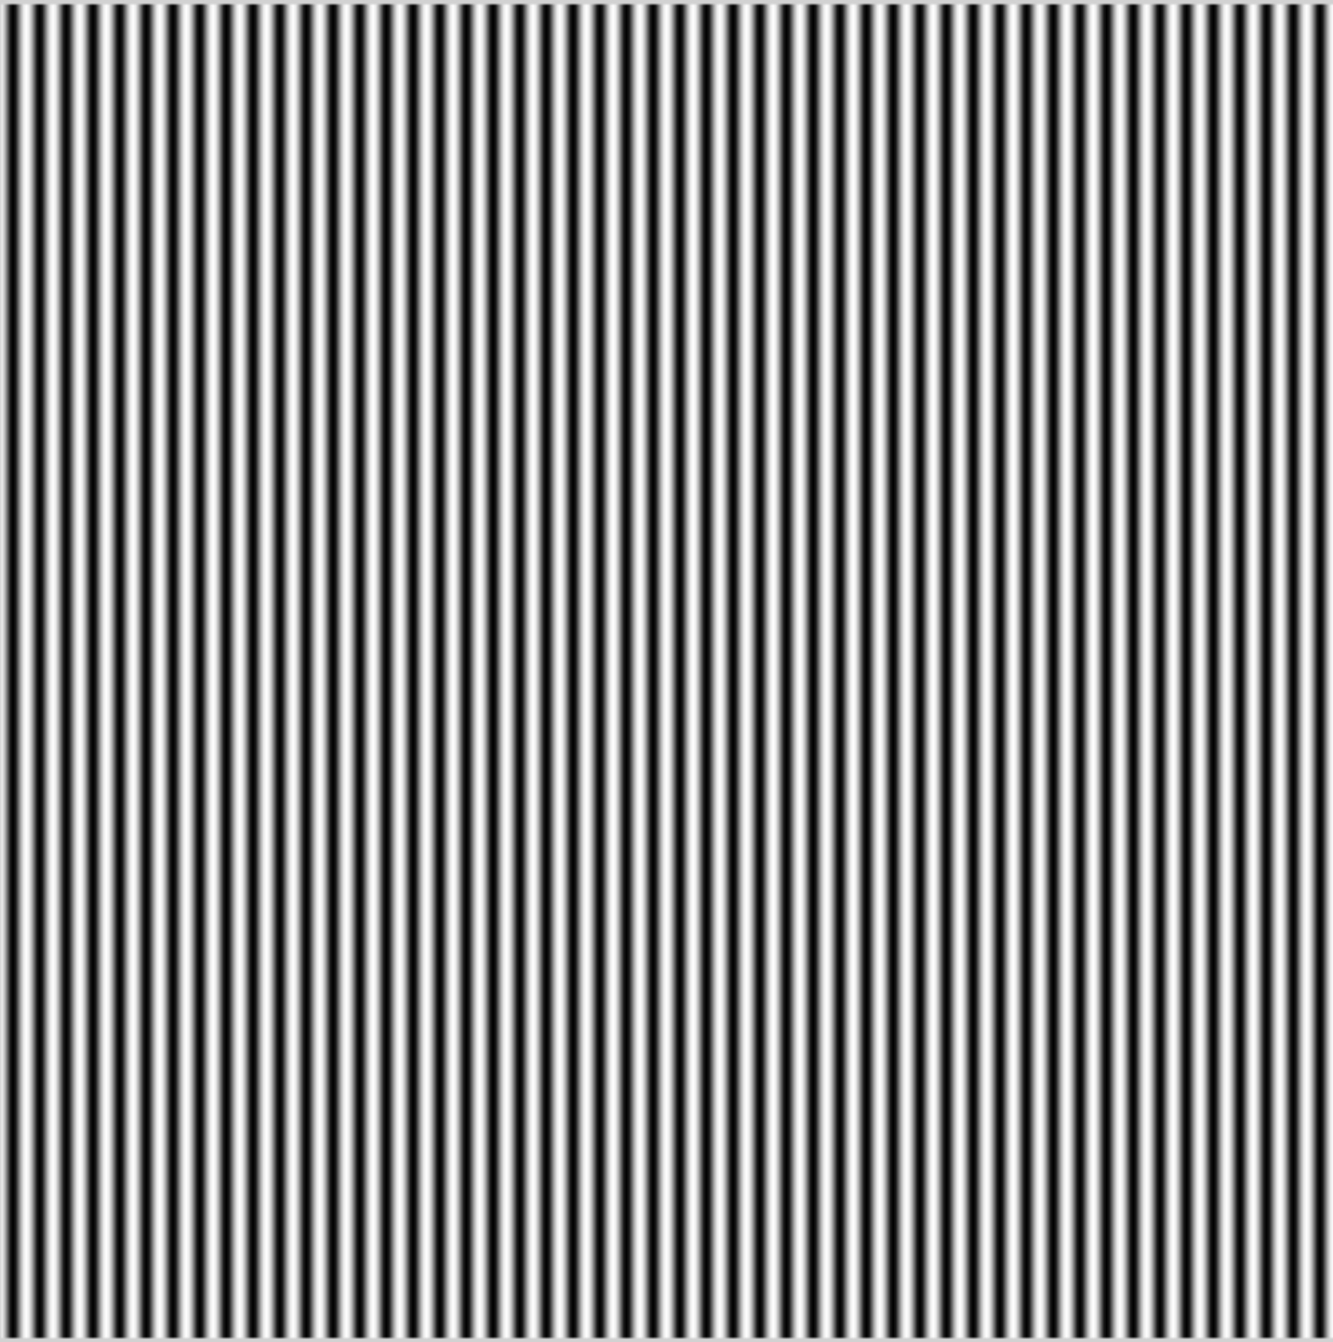

Contrast Sensitivity Function (CSF)

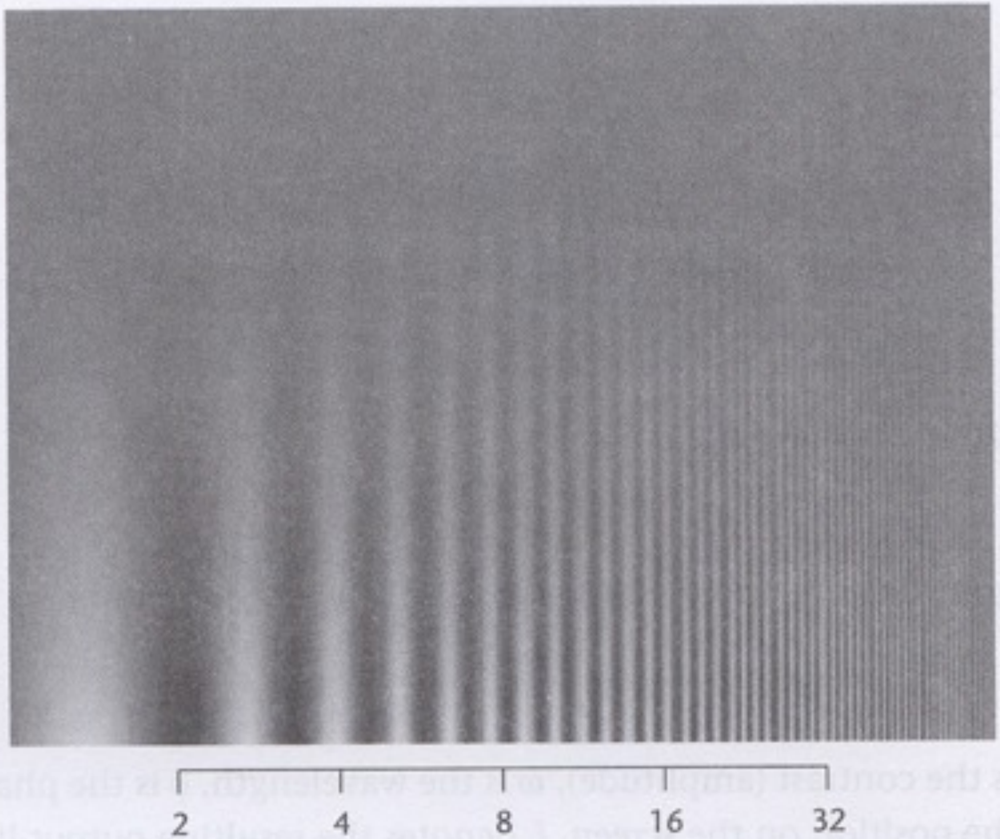

Figure 2.21 This grating pattern changes frequency exponentially from left to right and varies in contrast in a vertical direction. The highest frequency you can resolve depends on the distance from which you view the pattern. The scale gives the spatial frequency if it is viewed from 2.3 m.

Contrast Sensitivity Function (CSF)

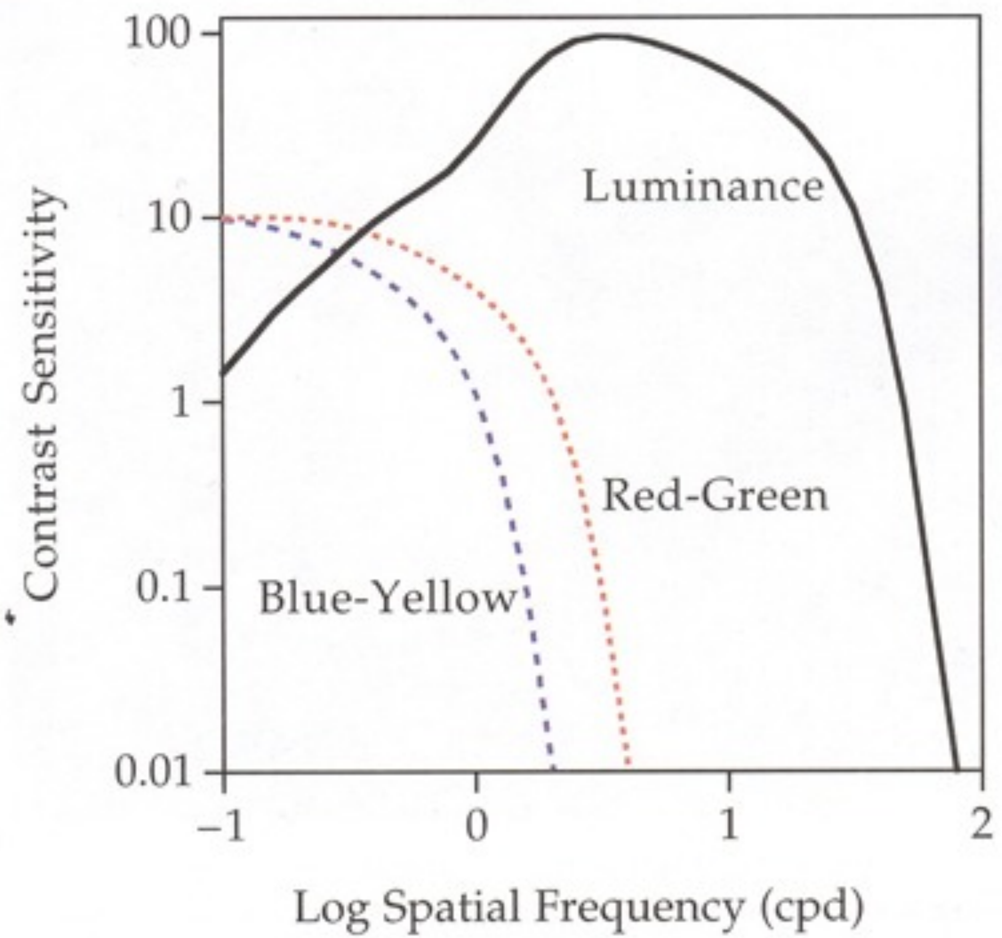

Figure 1-18. Spatial contrast sensitivity functions for luminance and chromatic contrast.

Photoshop demo

- Image > Mode > Lab color
- Go to channel panel, select Lightness

- Filter > Blur > Gaussian Blur , e.g. 4 pixel radius
  - very noticeable

- Undo, then select a & b channels
- Filter > Blur > Gaussian Blur , same radius
  - hardly visible effect

Original

Blur Lightness

Blur a & b

Questions?

A BALANCED COLOR SPHERE

Plan

- Wrap up CIE XYZ
- More human color perception
- **Extra quick overview of JPEG**
- Demosaicing
- Other color imaging strategies

Opponents and image compression

- JPG, MPG, television
- Color opponents instead of RGB
- Compress color more than luminance
  - downsample by factor of two for jpeg
  - less bandwidth for TV
- Exploit contrast sensitivity function
  - Compress high frequencies more

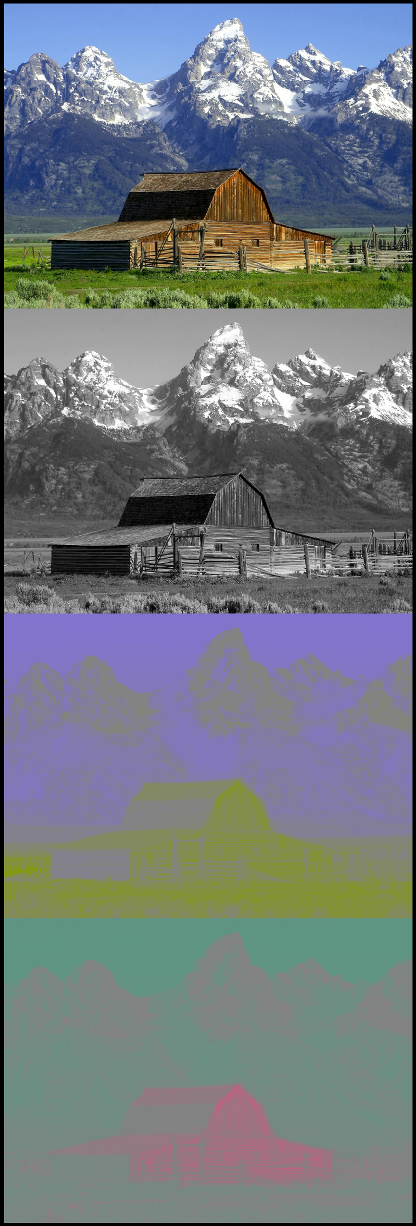

JPEG Compression

- convert to YCbCr
  - half the resolution for Cr & Cb

- Perform Discrete Cosine Transform to work in frequency space
  - Local DCT, 8x8 blocks

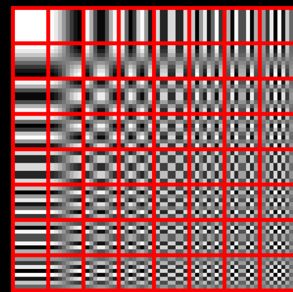

- Use CSF for quantization
  - more bits for frequencies with more sensitivity (medium)

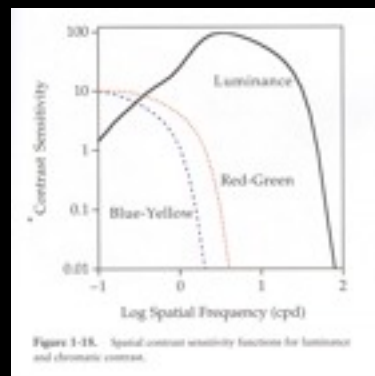

- Other usual coding tricks
  - entropy coding, smart order of blocks

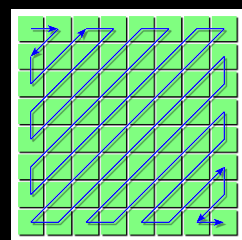

Example

- 800 x 533 image

	RGB	Lightness	a
Low quality JPEG (0 in Photoshop, 172 KB)			
High quality JPEG (12 in Photoshop, 460 KB)			

Questions?

Plan

- Wrap up CIE XYZ
- More human color perception
- Extra quick overview of JPEG
- **Demosaicing**
- Other color imaging strategies

Color sampling

- Problem: a photosite can record only one number
- We need 3 numbers for color
- What can we do?

What are some approaches to sensing color images?

- Scan 3 times (temporal multiplexing)
- Use 3 detectors (3-ccd camera)
- Use offset color samp (spatial multiplexing)
- Multiplex in depth (Tripack film, Foveon)
- Interferences (Lipmann)

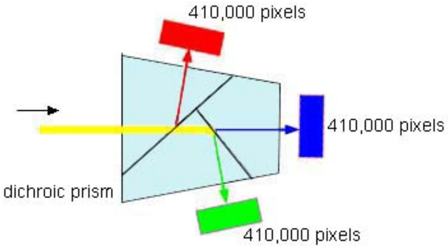

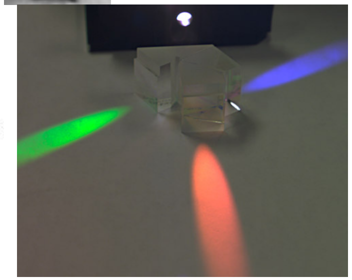

Dicklyon at the English Wikipedia project.

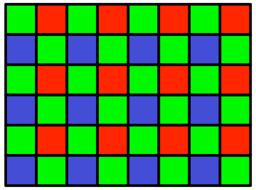

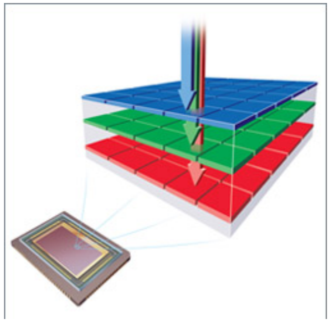

Bayer RGB mosaic

- Each photosite has a different color filter

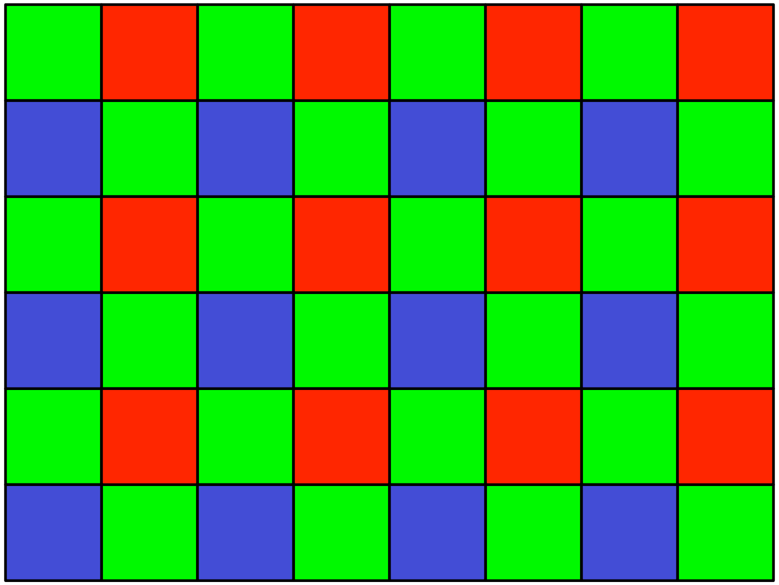

Note that which one is the upper left color is arbitrary and depends on the camera

Bayer RGB mosaic

- Why more green?
  - We have 3 channels and square lattices don't like odd numbers
  - It's the spectrum "in the middle"
  - More important to human perception of luminance

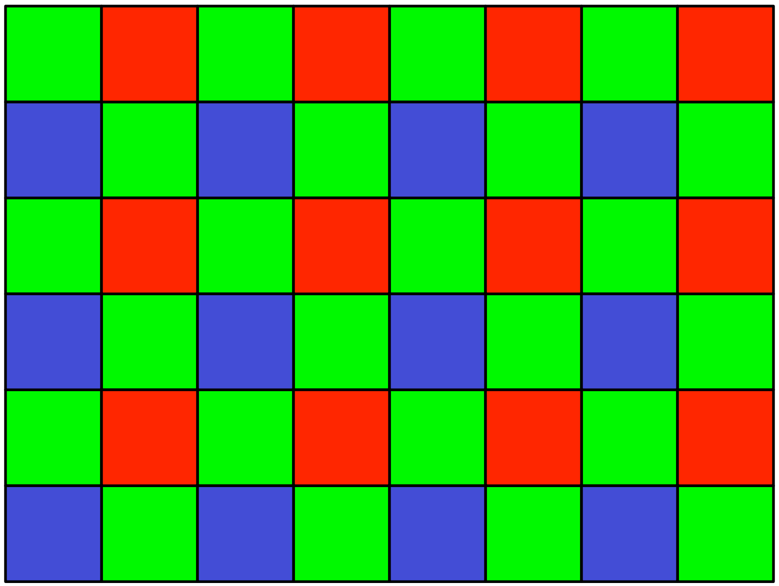

Demosaicing

- Interpolate missing values

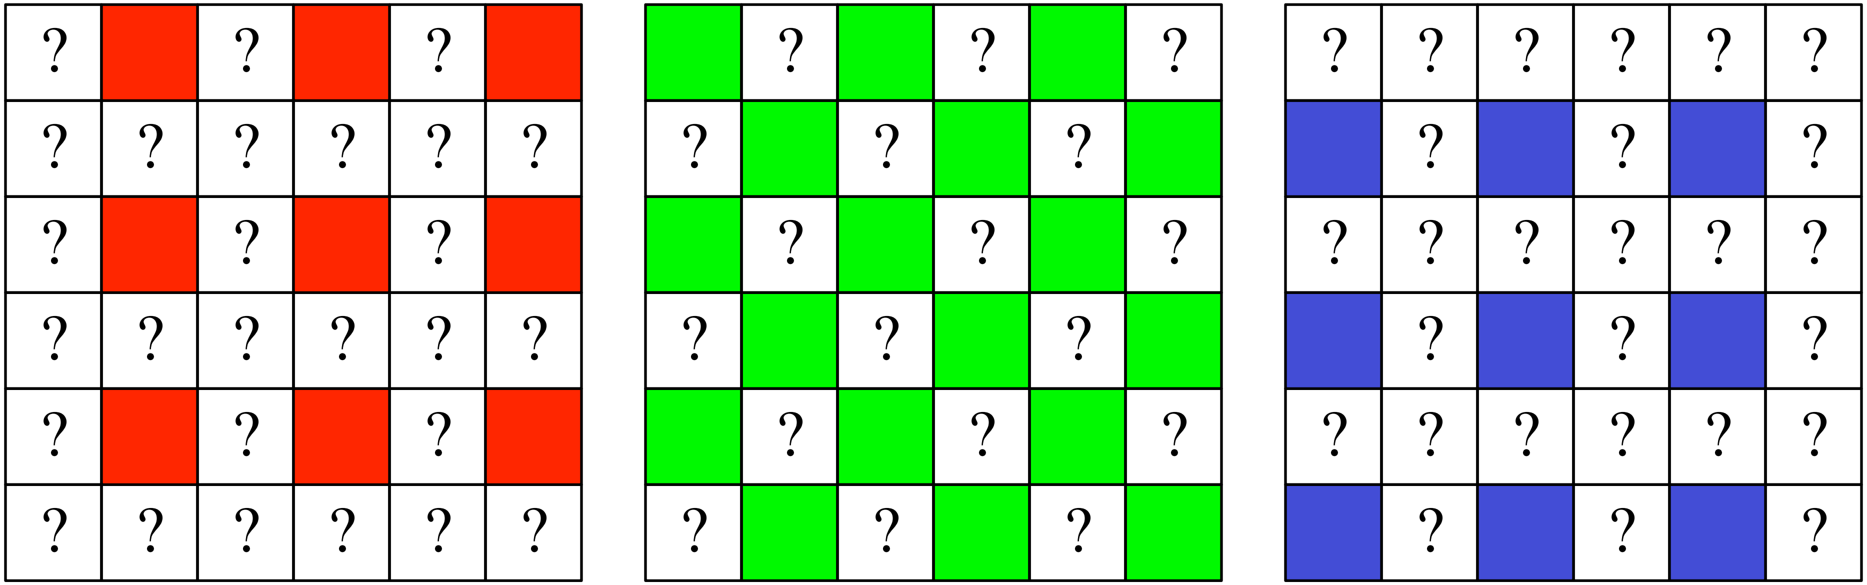

Demosaicing

- Simplest solution: downsample!
  - Nearest-neighbor reconstruction
- Problem: resolution loss (and megapixels are so important for marketing!)

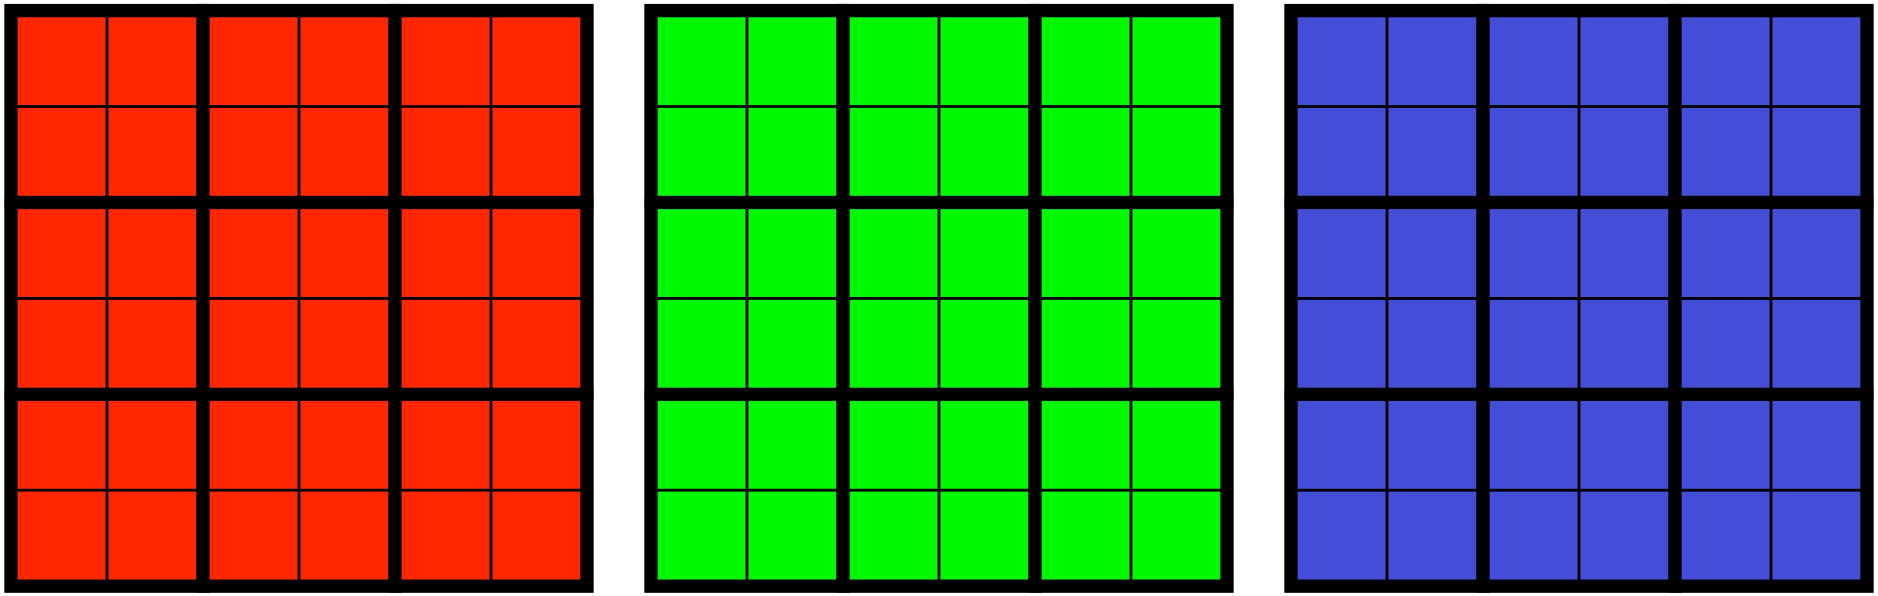

Linear interpolation

- Average of the 4 or 2 nearest neighbors
  - Linear (tent) kernel
- Smoother kernels can also be used (e.g. bicubic) but need wider support

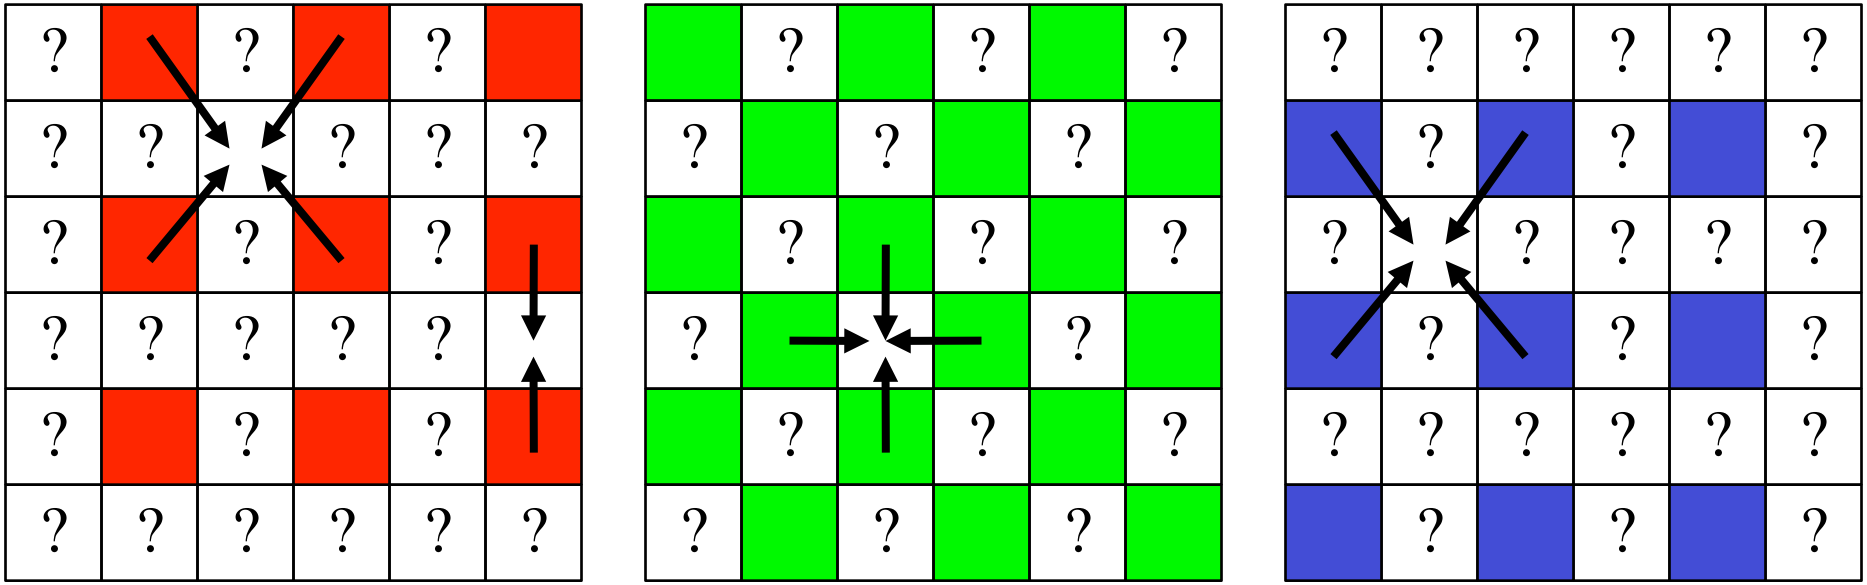

Questions?

Typical errors in spatial multiplexing approach.

- Color fringes or jaggies

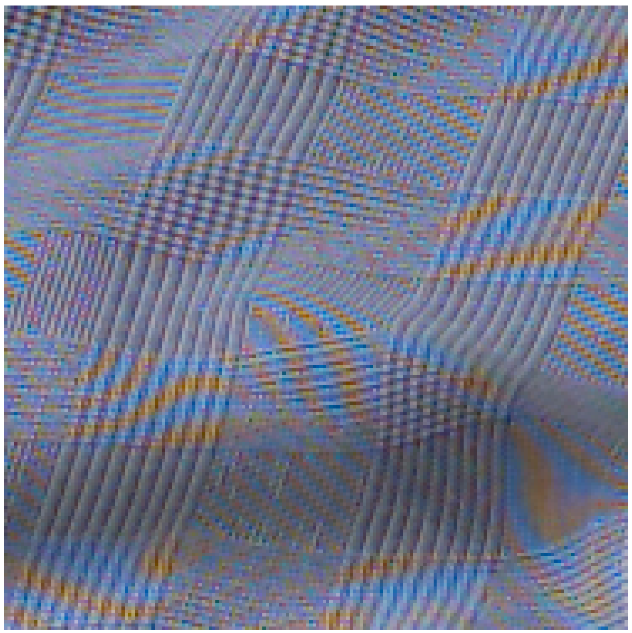

The cause of color moire

detector

(simplified for simpler visualization)

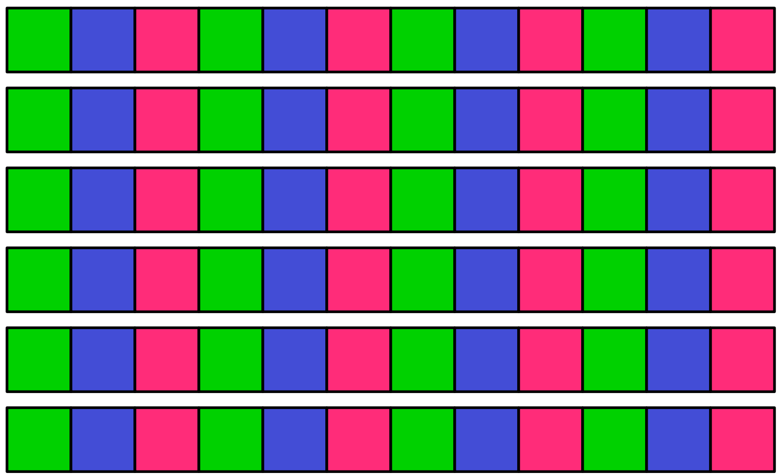

The cause of color moire

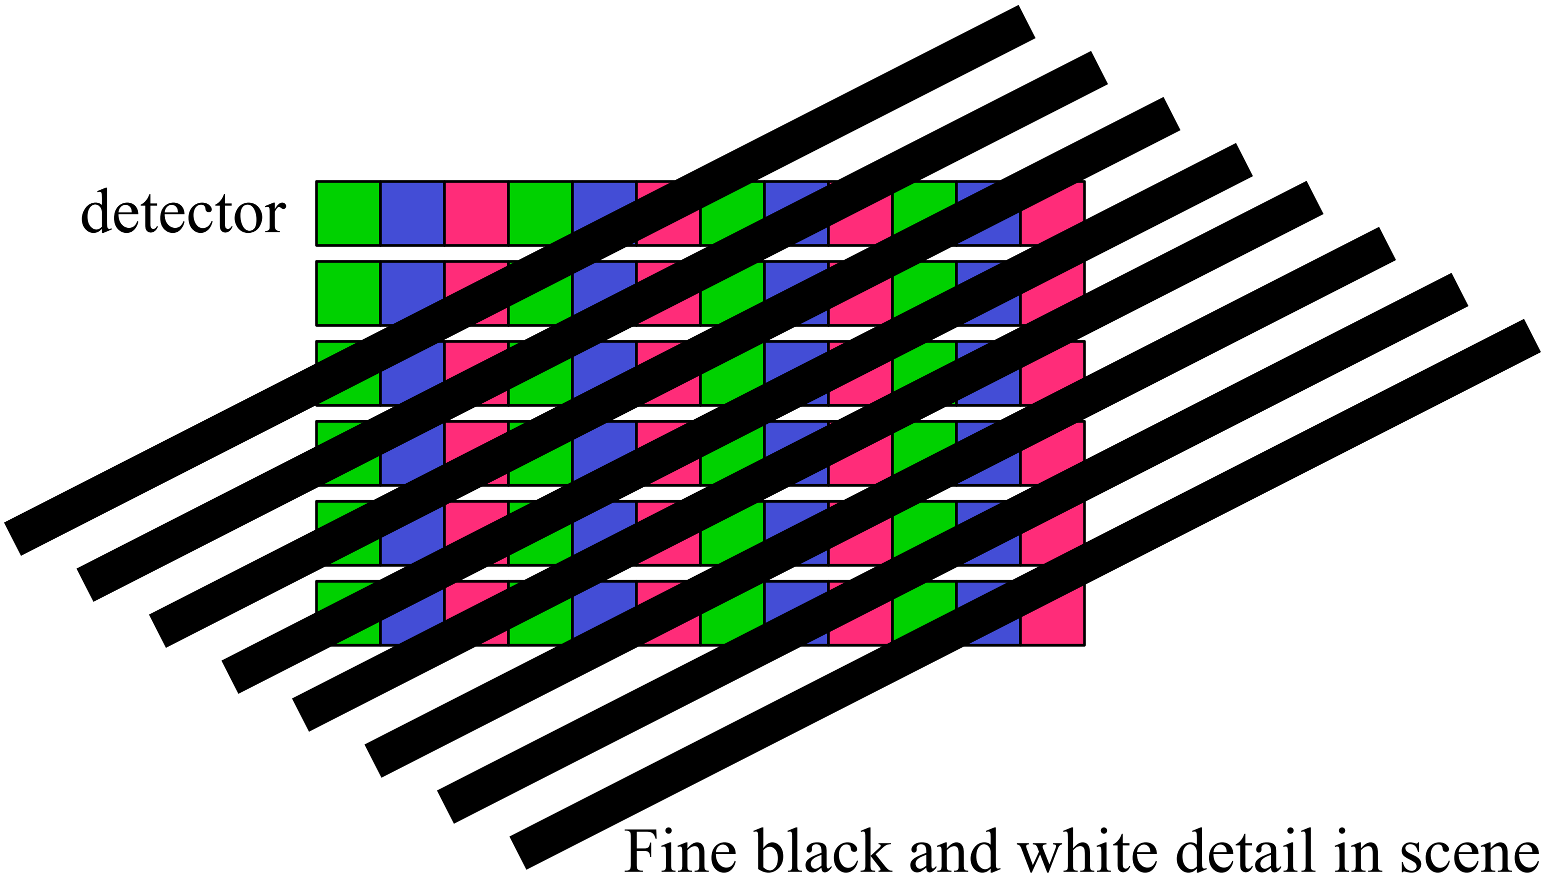

Fine black and white detail in scene mis-interpreted as color information.

Even simpler: Black and white edge falling on color sensor

Black and white image (edge)

Detector pixel colors

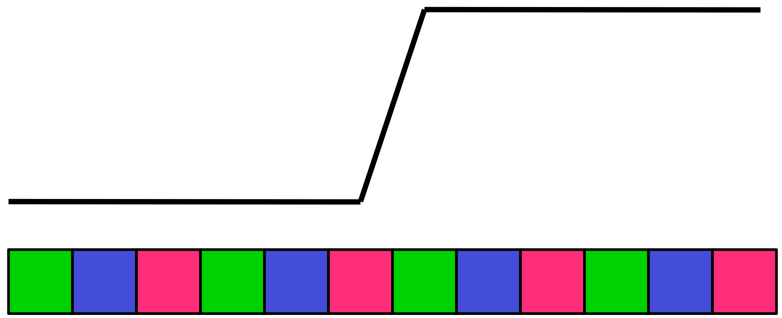

Color sampling artifact

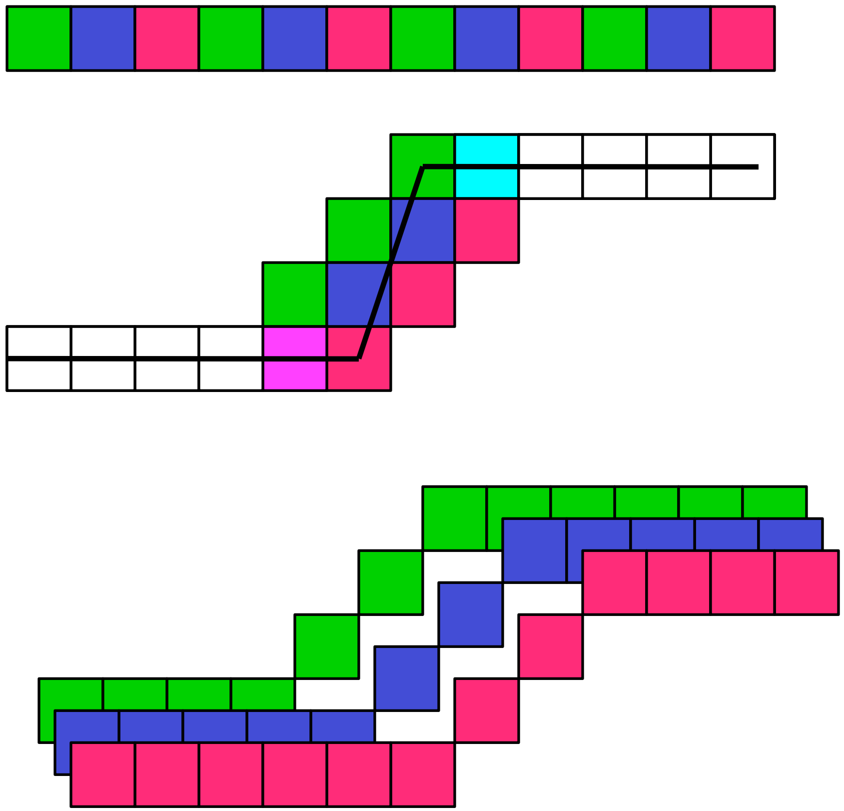

Interpolated pixel colors,
for grey edge falling on colored
detectors (linear interpolation).

Color sampling artifacts

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…

Motivation for median filter interpolation

The color fringe artifacts are places where **chrominance** changes abruptly and only temporarily.

Color opponent R-G, after linear interpolation

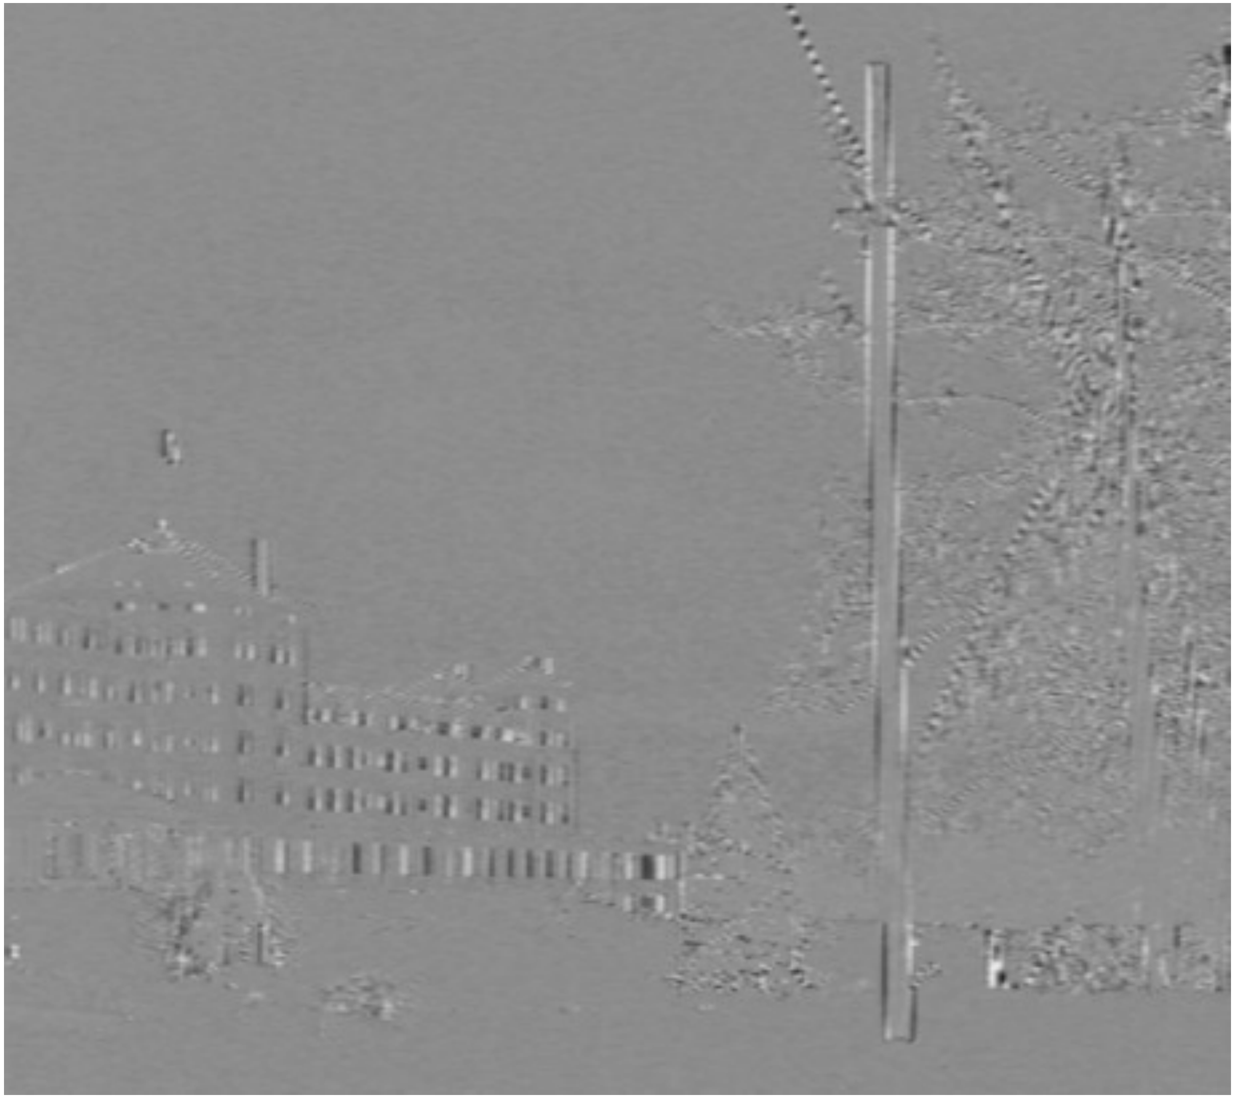

Median Demosaicking

- W. T. Freeman, “Median filter for reconstructing missing color samples,” US Patent, 4724395, 1988.
- Idea: Use median filter on chrominance to remove outlier transient chrominance changes

Median filter

Replace each pixel by the median over N pixels (5 pixels, for these examples).

In:

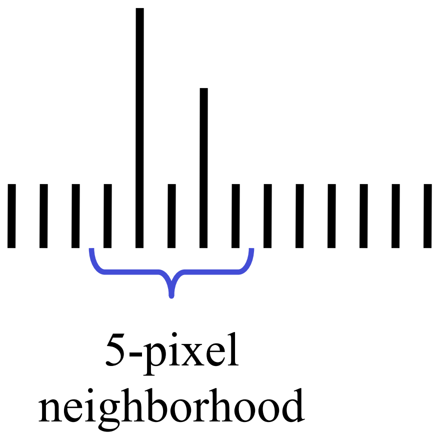

Out:

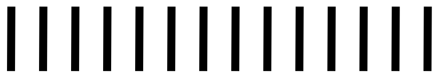

Spike noise is removed

In:

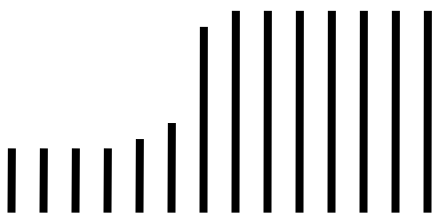

Out:

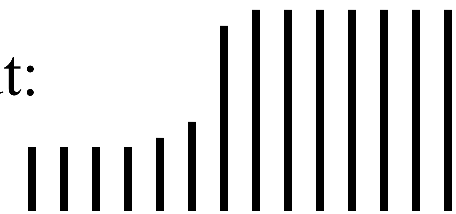

Monotonic edges remain unchanged

Degraded image

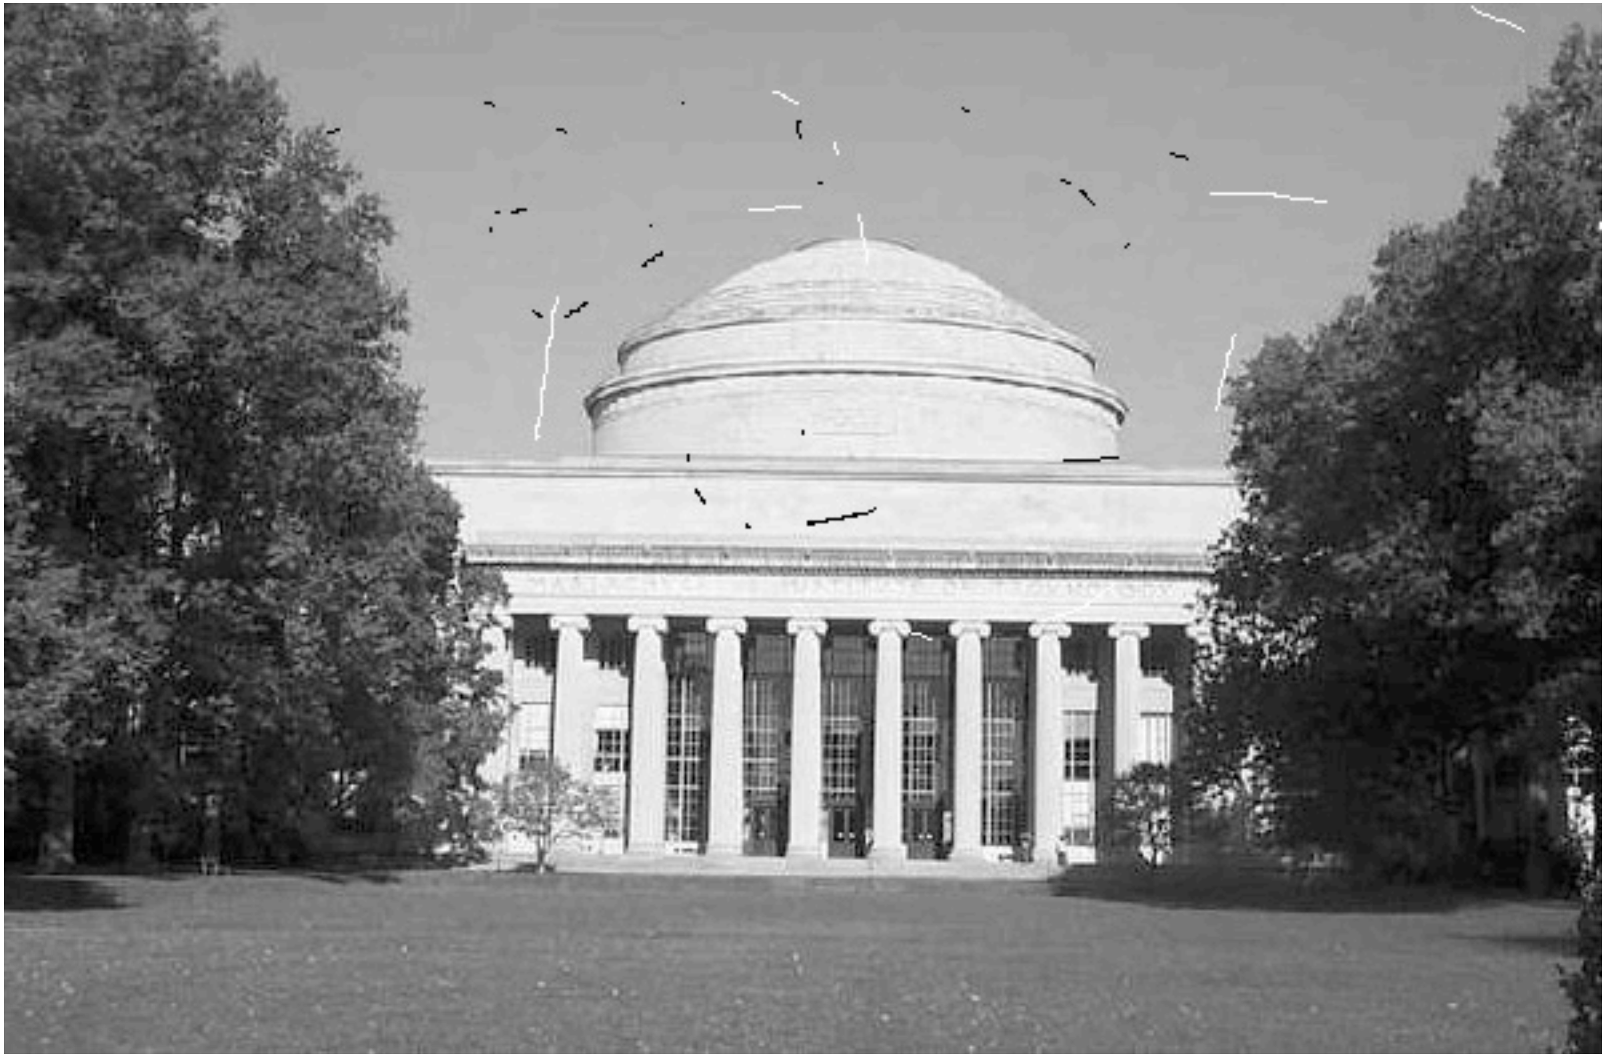

Radius 1 median filter

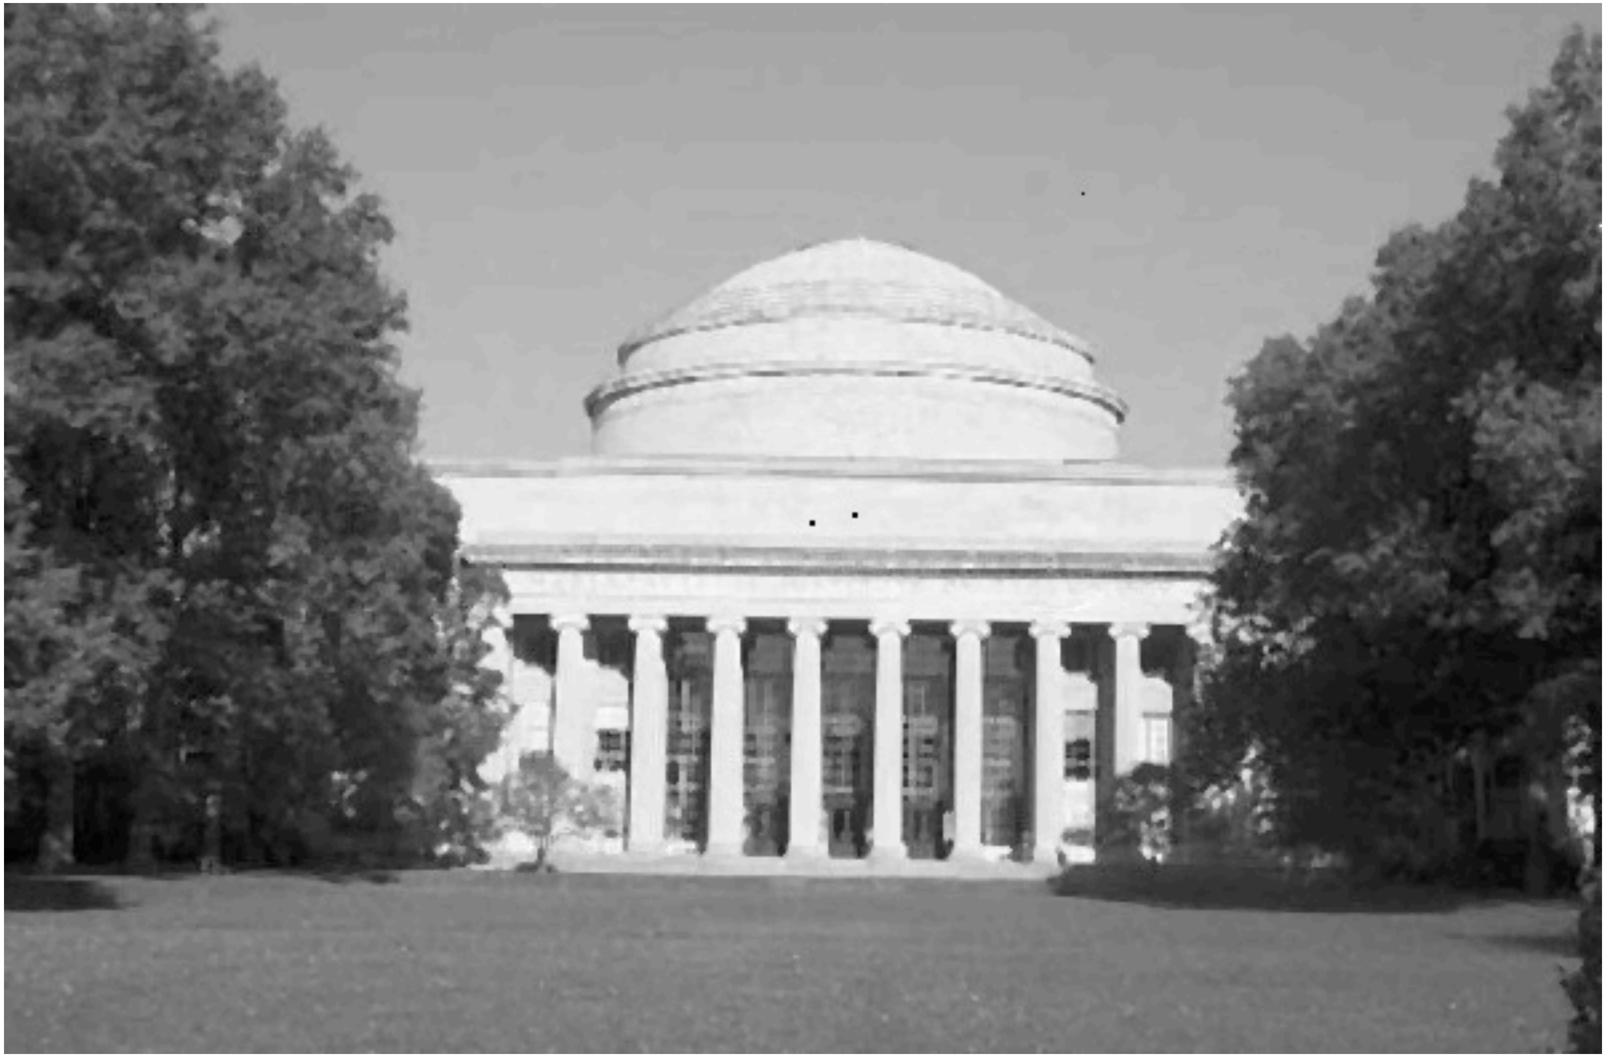

Radius 2 median filter

R – G

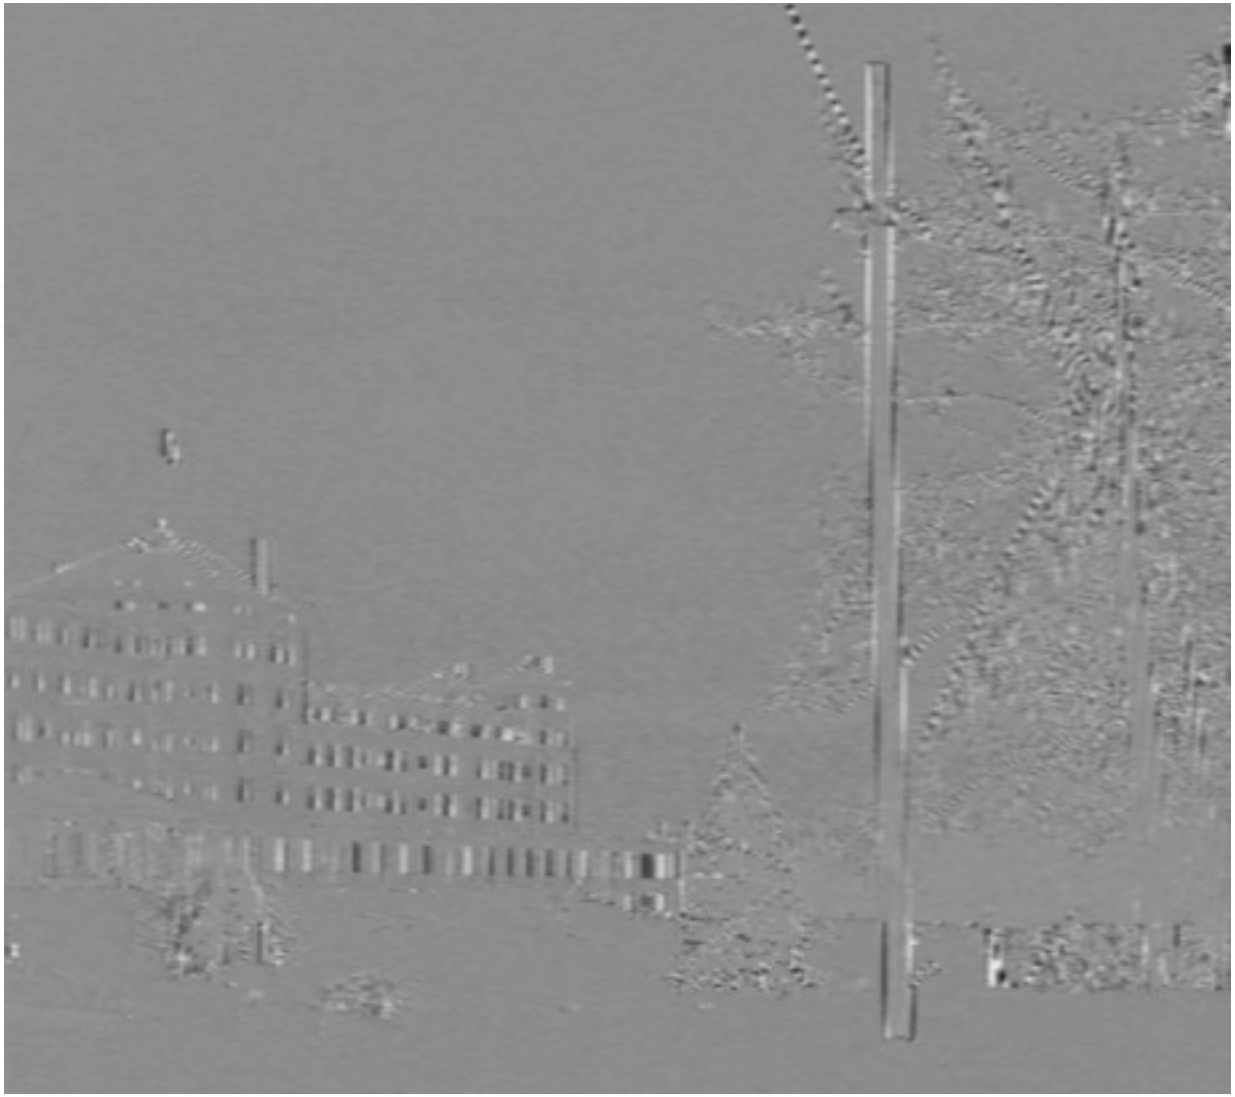

R – G, median filtered (5x5)

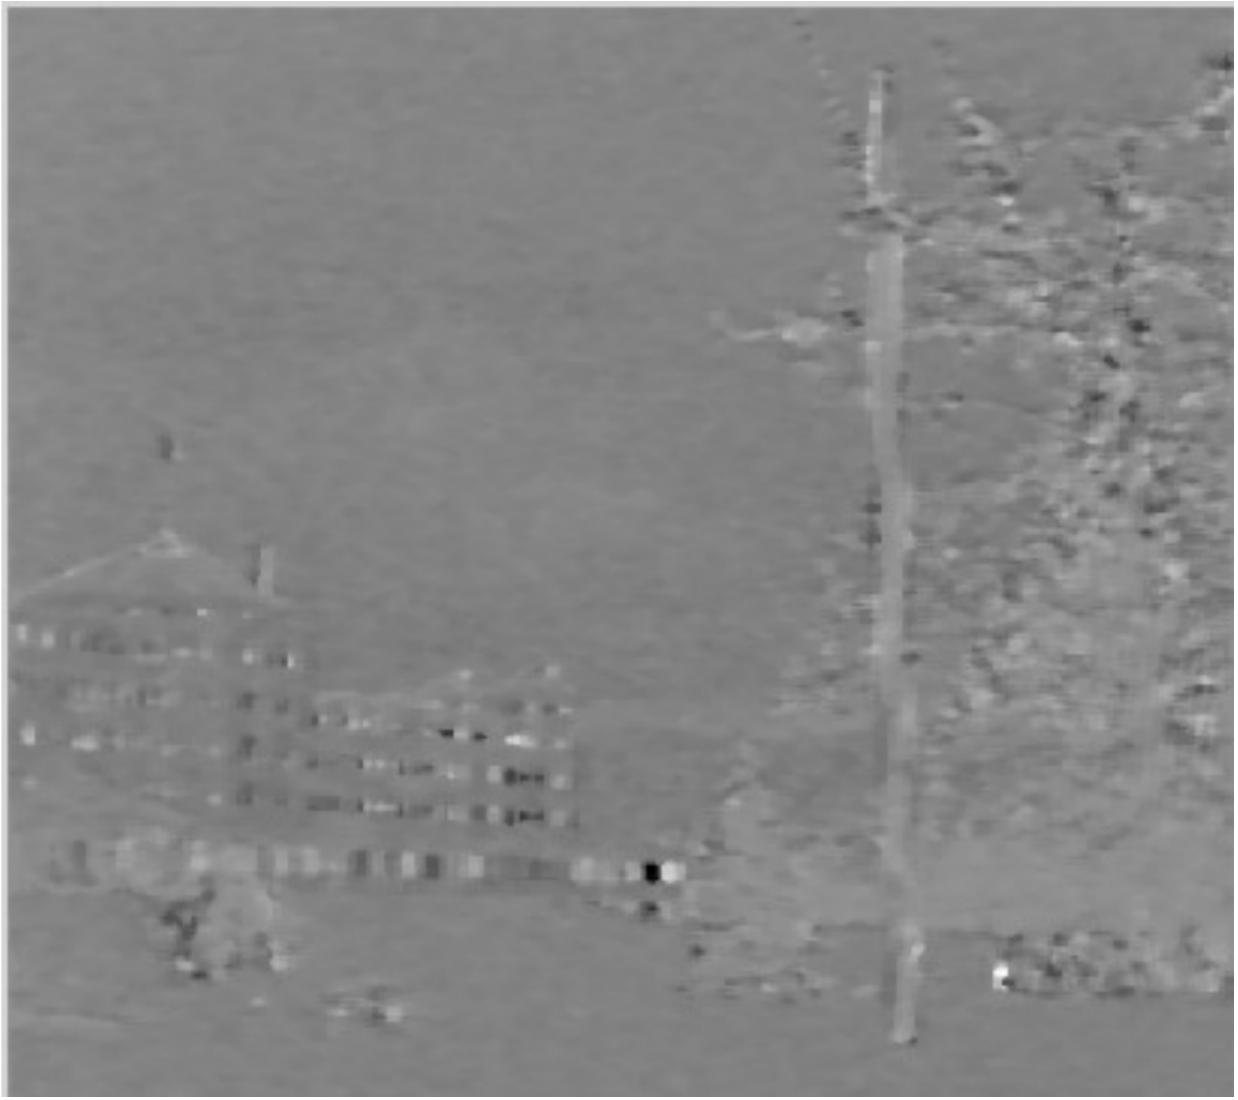

Median Filter Demosaicing

- Perform first linear interpolation on 3 color channels.
- Compute color difference signals (2 channels).
- Median filter the color difference signal.
- Reconstruct the 3-color image based on
  - median filtered color difference (2 channels)
  - original CCD data (1 value per pixel)

Two-color illustration, BW edge

Luminance profile

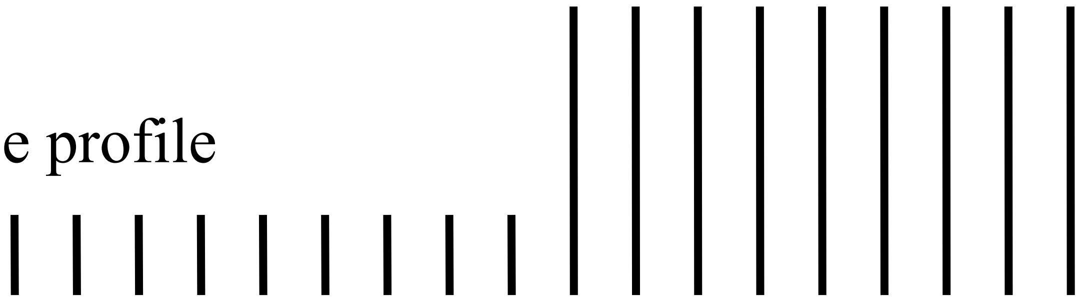

Two-color illustration, BW edge

Luminance profile

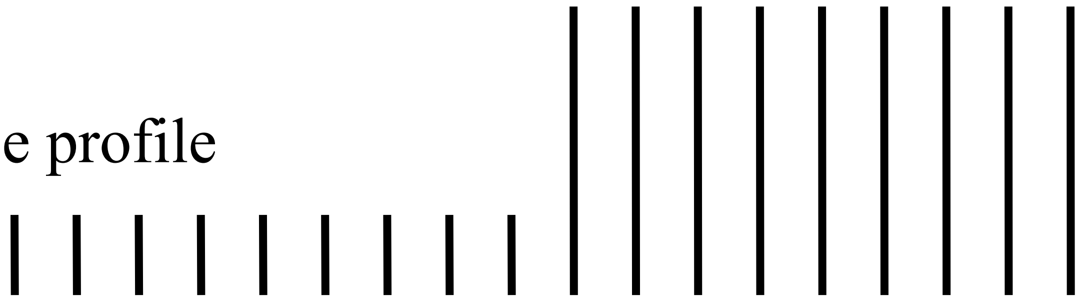

True full-color image

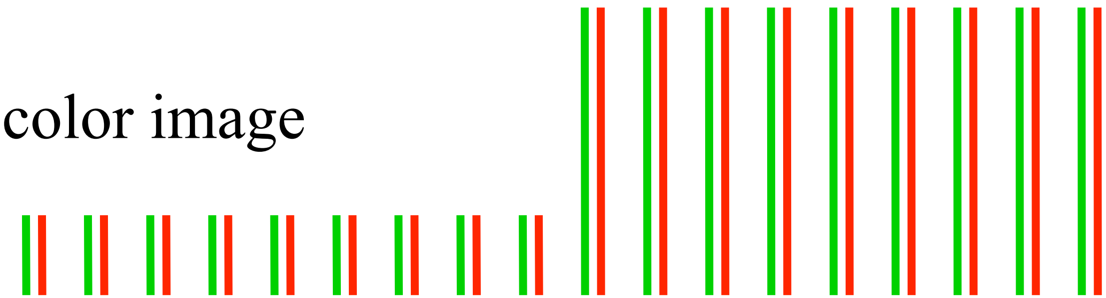

Two-color sampling of BW edge

Input (sensor mosaiced data)

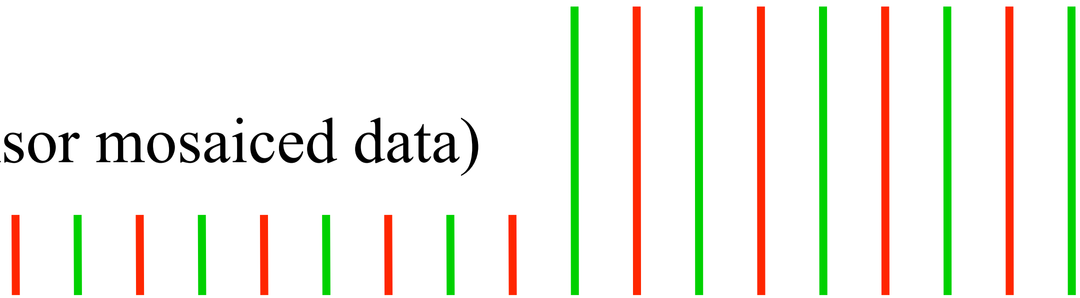

step 1: Linear interpolation

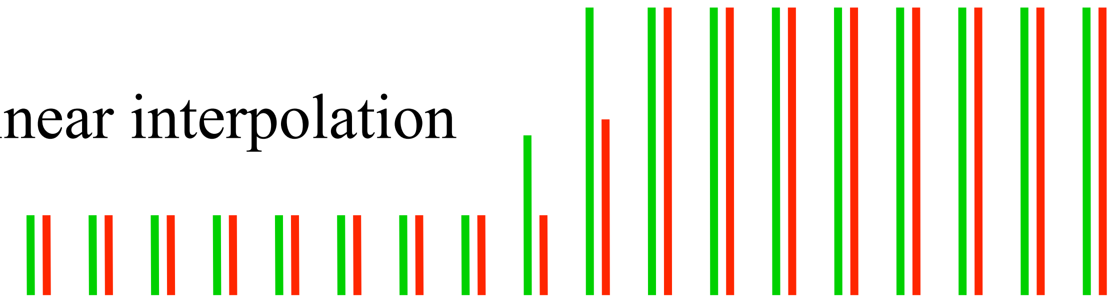

step 2:
Color difference signal G-R

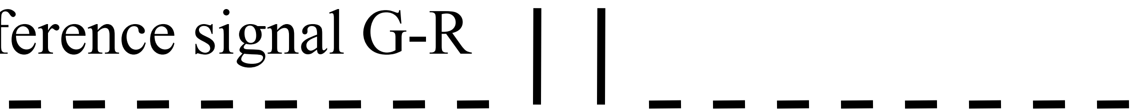

step 3:
Median filtered color difference signal

Final reconstruction (simplified version)

Input (sensor mosaiced data)

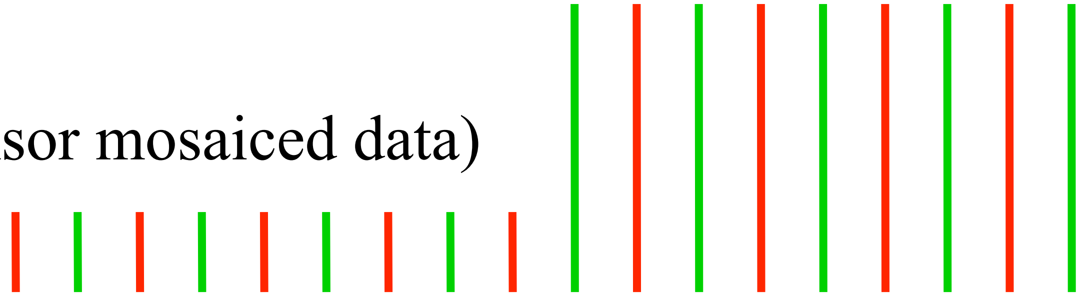

Median filtered color difference signal (from step 3)

step 4: For each pixel, reconstruct value that satisfy both:

- sampled data (1 number per pixel)
- median-filtered color difference (here 1 number per pixel, 2 for real)

e.g. for a green pixel: green=recorded value; red=green + difference

Output full-color image

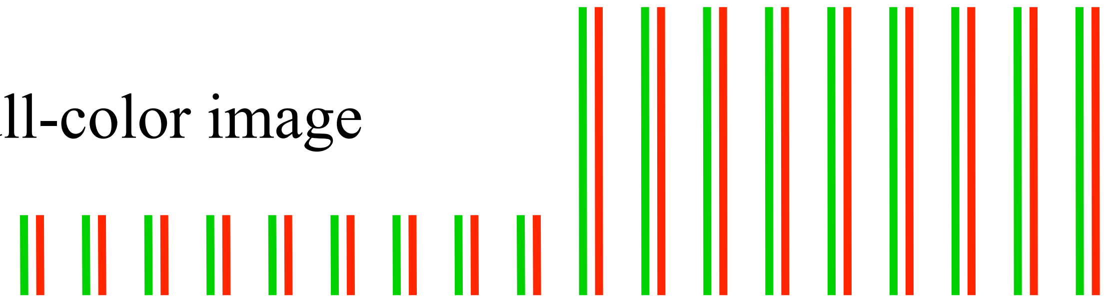

Questions?

- Recap:
  - linear interpolation
  - compute color opponent
  - median filter color opponent
  - reconstruct from color opponent and sensor data
- Note: what we saw was just for a simplified version with only R & G, and a simplistic color opponent.

The real case

- How many numbers per pixel in linear interpolated image? 3
- How many numbers for color difference/ chrominance? 2
- How many recorded number per pixel? 1
- How many numbers do we want? 3

YCbCr color space

- One of many “decorrelated” or opponent color spaces
  - used in JPEG, TV
  - a little too focused on 8-bit but oh well
- Y: luminance
- Cb: blue-yellow difference
- Cr: green-red difference
- e.g.
- $Y = 16 + 1/256 * (65.738 * R + 129.057 * G + 25.064 * B)$
- $Cb = 128 + 1/256 * (- 37.945 * R - 74.494 * G + 112.439 * B)$
- $Cr = 128 + 1/256 * (112.439 * R - 94.154 * G - 18.285 * B)$
- Just remember it’s linear

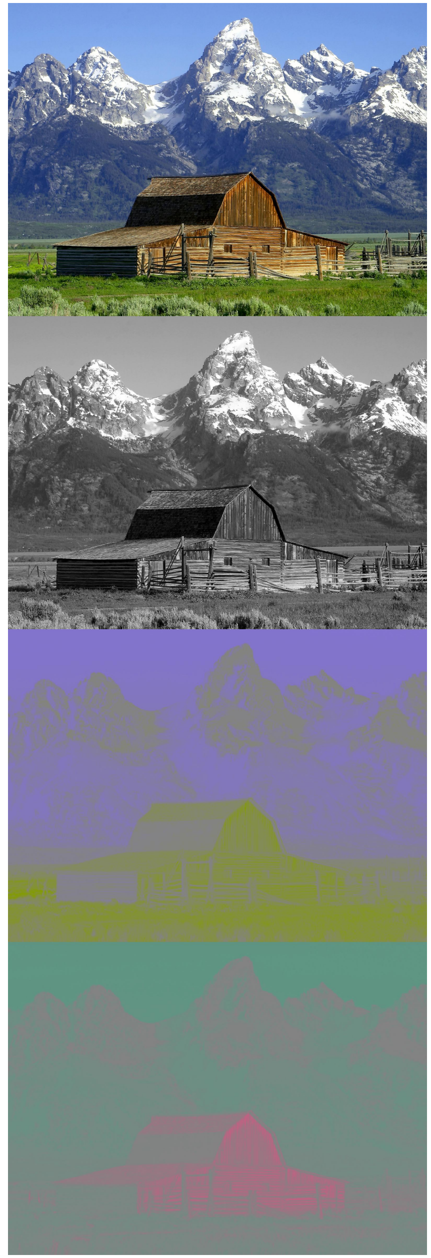

Y

Cb

Cr

Median color interpolation

- Apply naive interpolation
- Convert to YCbCr
  - throw away luminance Y
- Median filter Cr & Cb
  - independently for simplicity
  - vector-median if you're hard core
- At each pixel, reconstruct R, G, and B from sensor value + Cr & Cb

Per-pixel reconstruction

- For each pixel we have
  - One of either R, G, or B (sensor data)
  - 2 chrominance numbers Cr & Cb (after median)
- We want three numbers R, G, and B
- We know
  - $Cb = 128 + 1/256 * (- 37.945 * R - 74.494 * G + 112.439 * B)$
  - $Cr = 128 + 1/256 * (112.439 * R - 94.154 * G - 18.285 * B)$
- Think you can do it?

Recombining the median filtered colors

Linear interpolation

Median filter interpolation

Questions?

Typical bells and whistles

- Focus first on green channel
- Oriented interpolation (detect gradient or grating orientation)
- Use green to help interpolate blue and red
- Post-process step such as the median to remove color artifacts

Other possibilities

- CMY mosaic
  - Pro: gather more light per photosite
  - Con: not directly what we want, potential loss of color sensitivity

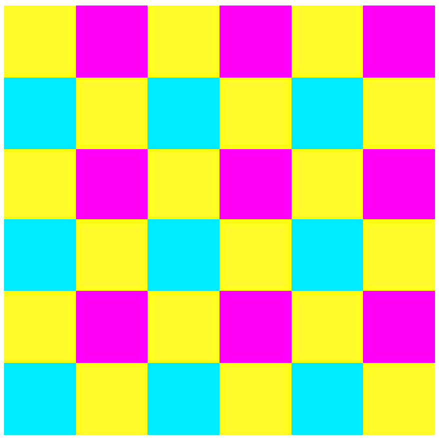

CMY Bayer Pattern
DCS 620x

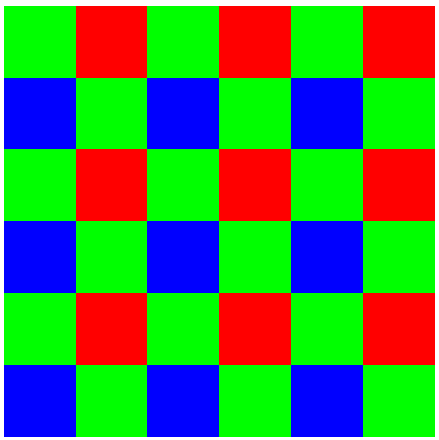

RGB Bayer Pattern
DCS 620

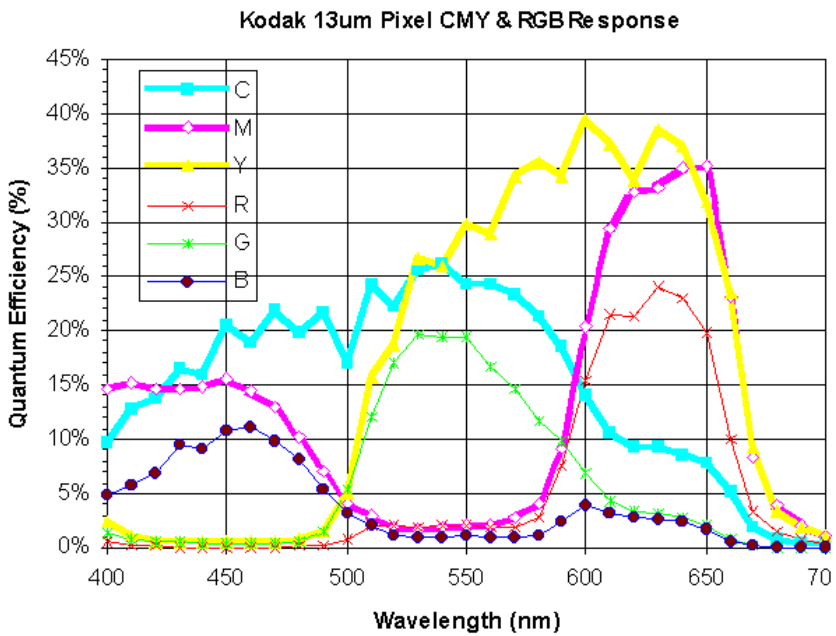

Extension

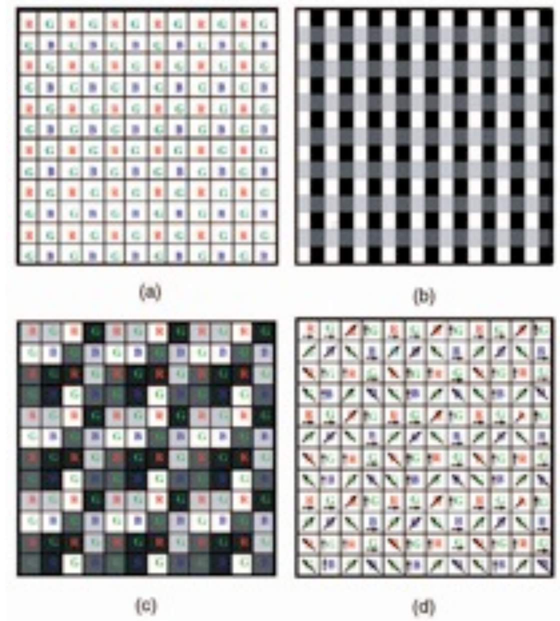

- Mosaicing can be used to gather more information
  - Use neutral density filters to get more dynamic range
  - Polarizers
  - Etc.

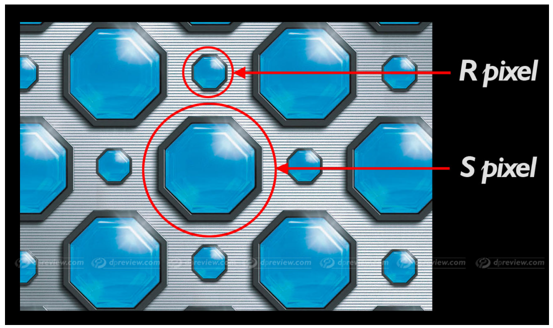

- Nayar et al.'s work
  - http://www1.cs.columbia.edu/CAVE/publications/pdfs/Narasimhan_PAMI05.p

- Fuji's super CCD

Questions?

Plan

- Wrap up CIE XYZ
- More human color perception
- Extra quick overview of JPEG
- Demosaicing
- **Other color imaging strategies**

Discussion: Approaches to color imaging

- Scan 3 times (temporal multiplexing)
- Use 3 detectors (3-ccd camera)
- Use offset color samp (spatial multiplexing)
- Multiplex in depth (Tripack film, Foveon)
- Interferences (Lipmann)

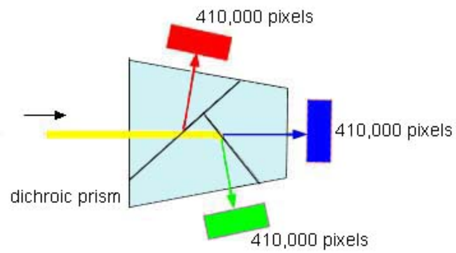

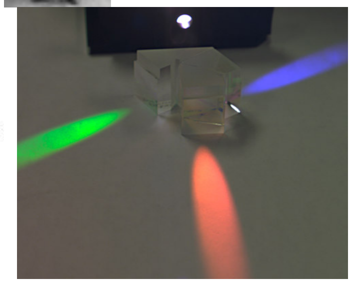

Dicklyon at the English Wikipedia project.

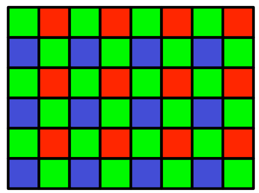

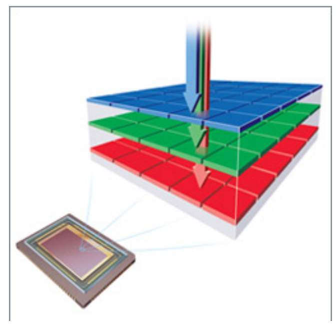

Temporal multiplexing

- Examples:
  - Drum scanners
  - Flat-bed scanners
  - Russian photographs from 1800’s
- Pros:
  - 3 real values per pixel
  - Can use a single sensor
- Cons
  - Only for static scenes, slow

3-chip

- High-end 3-ccd video cameras
- Use separation prisms
  - prisms that split wavelengths
- Pros
  - 3 real values per pixel
  - Little photon loss
- Cons
  - costly (needs 3 sensors)
  - space

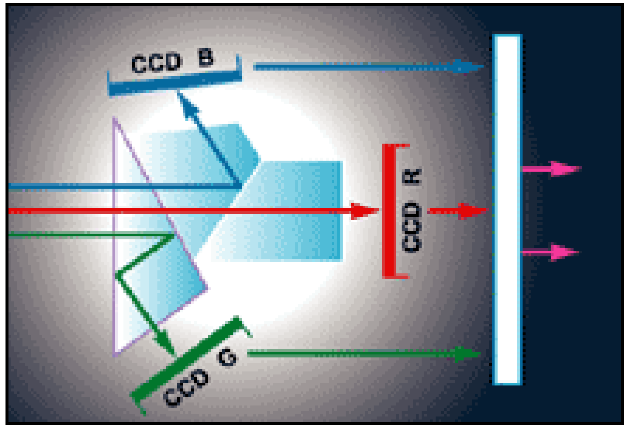

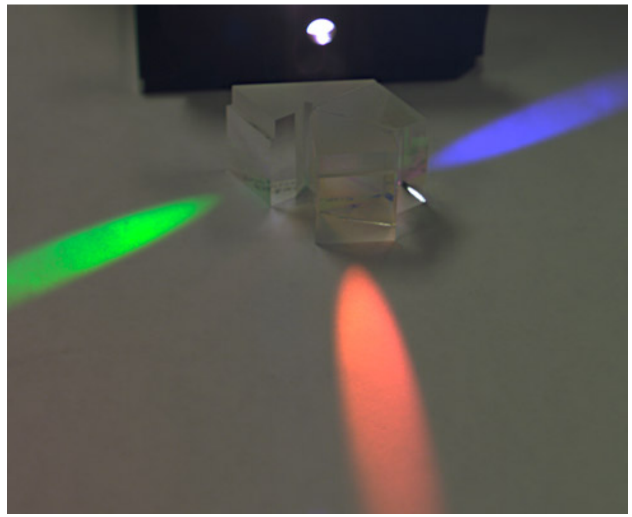

Dicklyon at the English Wikipedia project.

Spatial multiplexing

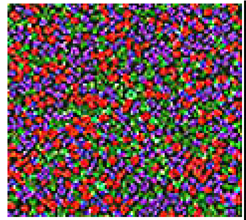

- Human eye (cone mosaic)
- Older color film (Autochrome)
- Bayer mosaic or CFA (color filter array)
- Most still cameras, most cheap camcorder, some high-end video cameras (e.g. RED)

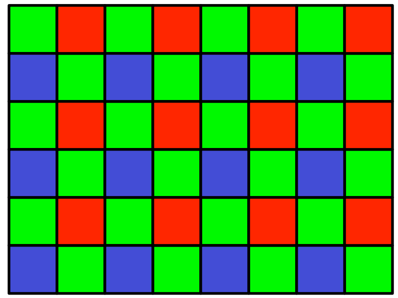

- Pros
  - single sensor
  - well mastered technology, high resolution
- Cons
  - needs interpolation, color jaggies
  - requires antialiasing filter (reduces sharpness)
  - loss of light

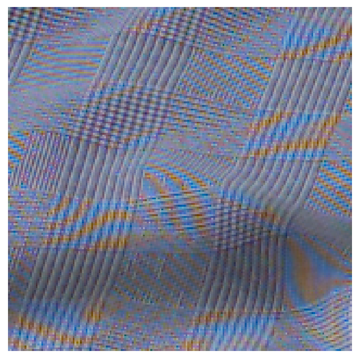

Combination: pixel shift

- 3-ccd with prisms + spatial multiplexing
- The 3 ccds are shifted by 1/2 pixel to provided resolution increase
  – usually selectable (not shifted for lower-res, shifted to get HD)
  – Often horizontal only

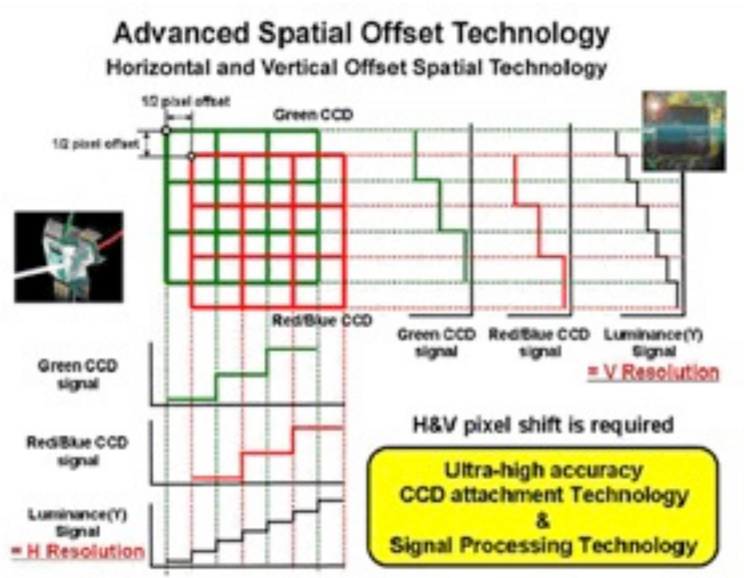

From Panasonic

Depth multiplexing (Foveon X3 sensor)

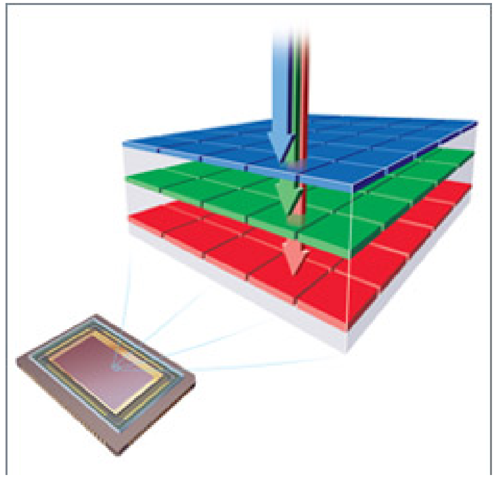

- Leverage difference in absorption per wavelength

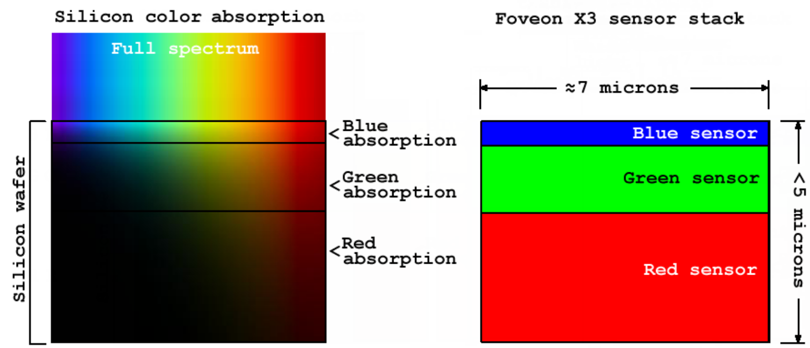

- Pros
  - 3 real numbers per pixel
  - Less light loss
- Cons
  - Requires more color processing (3 numbers must be multiplied by matrix to get RGB)
  - Tends to be noisier (because color processing and because shallow blue layer)

Depth multiplexing

- Good old color film (tripack)

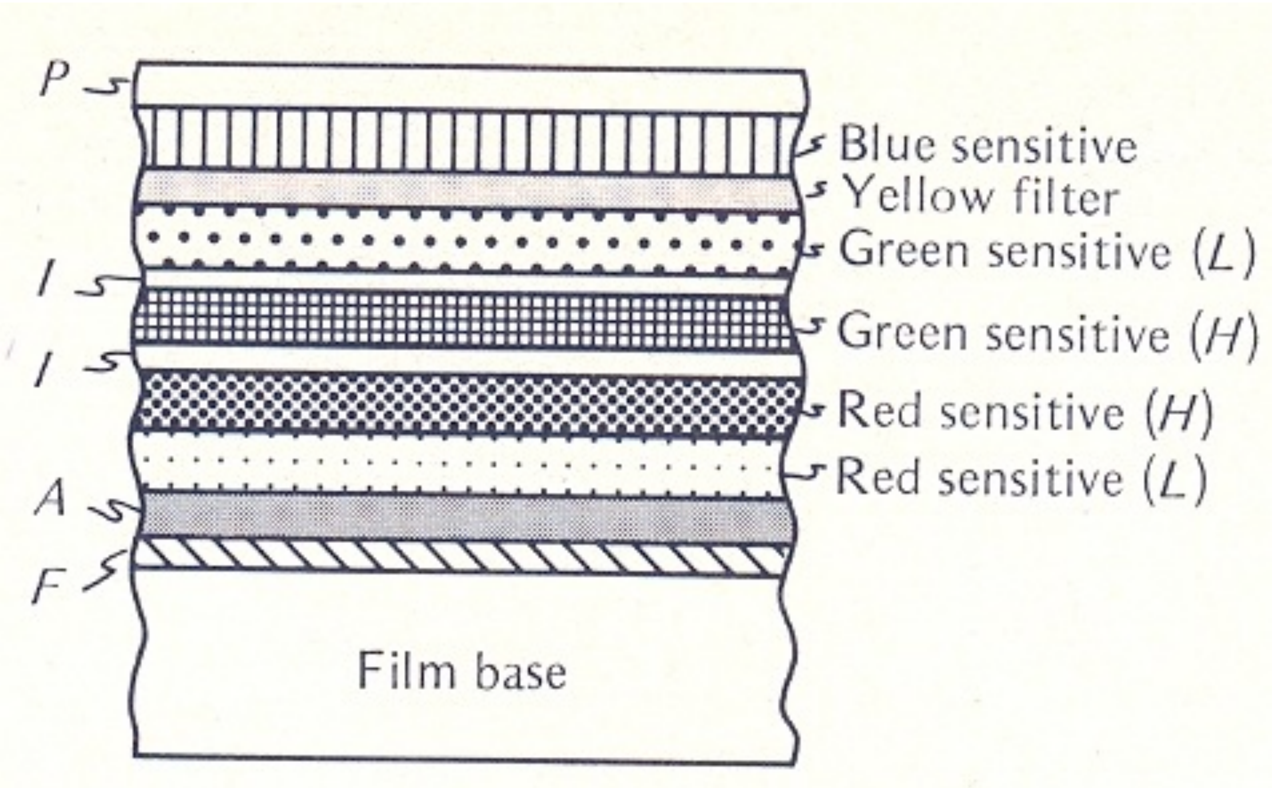

Interferences (Lippmann process)

- Metal mirror to create interferences
  - ancestor of holography
  - similar to colors in thin oil film

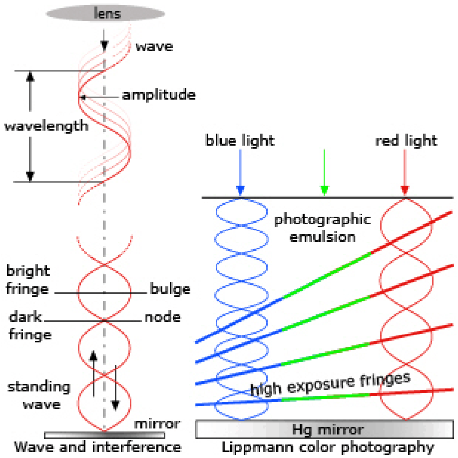

http://nobelprize.org/nobel_prizes/physics/articles/biedermann/index.html

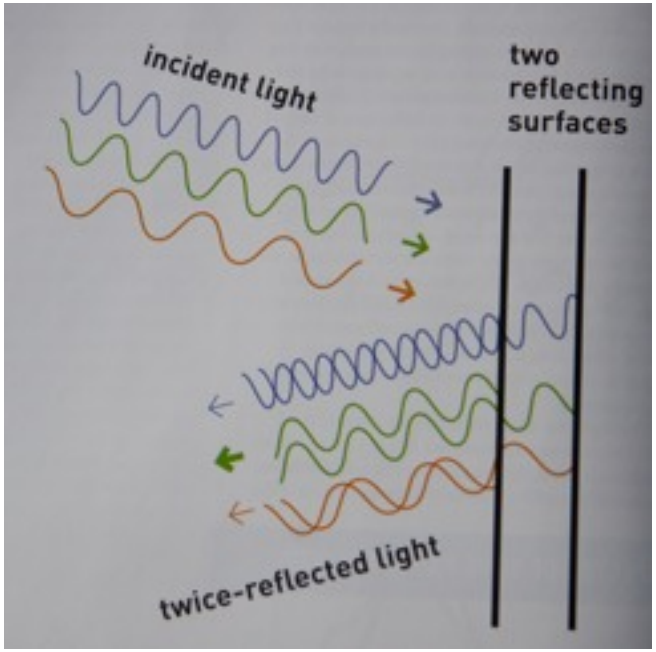

Interferences (Lippmann process)

- Metal mirror to create interferences
  - ancestor of holography
  - similar to colors in thin oil film
- Pros
  - Full spectrum!!!!!
  - Gets you the Nobel if you invent it ;-)
- Cons
  - Needs high-resolution sensor/film
  - limited field of view

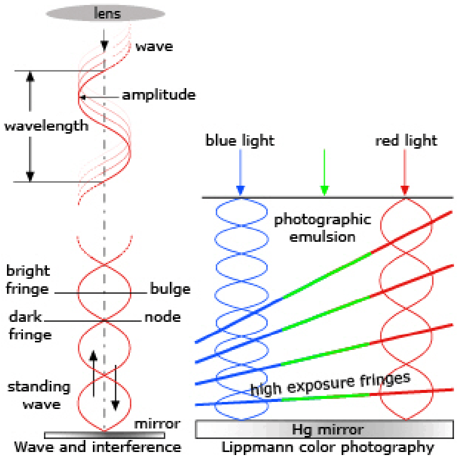

http://nobelprize.org/nobel_prizes/physics/articles/biedermann/index.html

'Saint-Maxime", 1891-1899
Photographer: Gabriel Lippmann (1845-1929)

Refs

- http://www.ists.dartmouth.edu/library/edf0205.pdf
- http://ieeexplore.ieee.org/iel5/79/30519/01407714.pdf?isnumber=30519&prod=JNL&arnumber=1407714&arSt=+44&ared=+54&arAuthor=Gunturk%2C+B.K.%3B+Glotzbach%2C+J.%3B+Altunbasak%2C+Y.%3B+Schafer%2C+R.W.%3B+Mersereau%2C+R.M.

end