Color

Lecture 2

Homework 1 is out

- Due on Mon 25th Feb
- Also start looking at ideas for projects
- Suggestions are welcome!

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

Why is color useful

- Find things to eat
- Spot dangerous things

Color

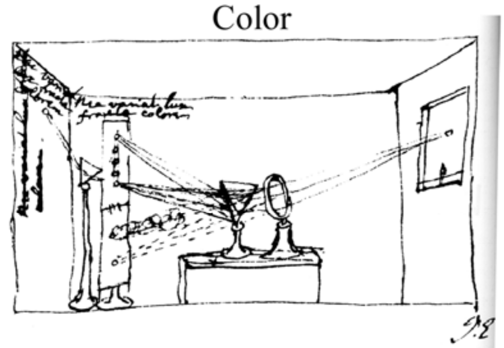

4.1 NEWTON'S SUMMARY DRAWING of his experiments with light. Using a point source of light and a prism, Newton separated sunlight into its fundamental components. By reconverging the rays, he also showed that the decomposition is reversible.

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

Electromagnetic Spectrum

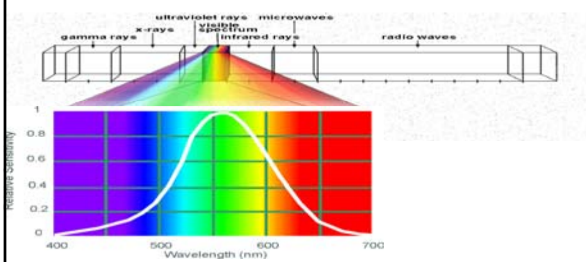

Human Luminance Sensitivity Function

http://www.yorku.ca/eye/photopik.htm

Visible Light

Plank's law for Blackbody radiation
Surface of the sun: ~5800K

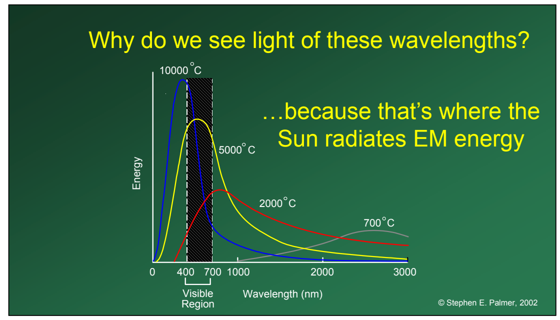

Spectral Image Formation

I(λ)

$S_c(λ)$

R(λ)

I(λ) R(λ)

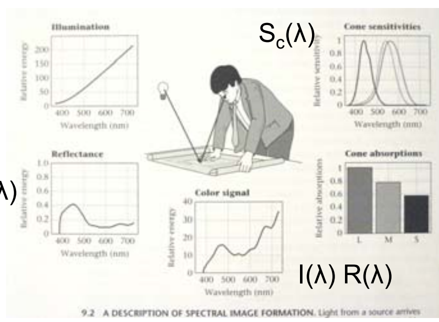

From Foundation of Vision by Brian Wandell, Sinauer Associates, 1995

Spectral Image Formation

$$p_c = \int I(\lambda) S_c(\lambda) R(\lambda)\ d\lambda$$

I(λ) – Illumination Spectrum

$S_c(λ)$ - Spectral sensitivity of channel c

R(λ) - Surface reflectance/transmission

Pixel value / Perceived color depends on all 3 terms!
→ Problem of color constancy

Simplified rendering models: reflectance

Often are more interested in relative spectral composition than in overall intensity, so the spectral BRDF computation simplifies a wavelength-by-wavelength multiplication of relative energies.

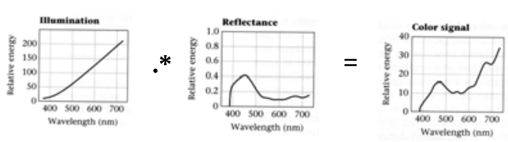

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Simplified rendering models: transmittance

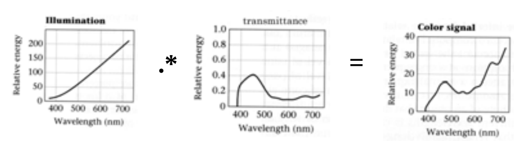

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

How measure those spectra: Spectrophotometer

(just like Newton's diagram…)

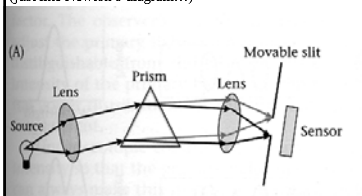

(B)

Wavelength (nm)	Color (viewed in isolation)
700	Red
610	Orange
580	Yellow
540	Green
480	Blue
400	Violet

4.2 A **SPECTRORADIOMETER** is used to measure the spectral power distribution of light. (A) A schematic design of a spectroradiometer includes a means for separating the input light into its different wavelengths and a detector for measuring the energy at each of the separate wavelengths. (B) The color names associated with the appearance of lights at a variety of wavelengths are shown. After Wyszecki and Stiles, 1982.

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

Spectrometer Demo

- Lights
- Bulb
- LEDs

Some reflectance spectra

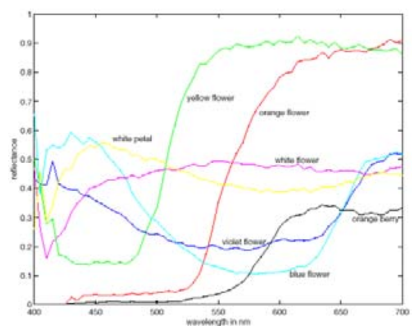

Spectral albedoes for several different leaves, with color names attached. Notice that different colours typically have different spectral albedo, but that different spectral albedoes may result in the same perceived color (compare the two whites). Spectral albedoes are typically quite smooth functions. Measurements by E.Koivisto.

Forsyth, 2002

The Physics of Light

Some examples of the reflectance spectra of surfaces

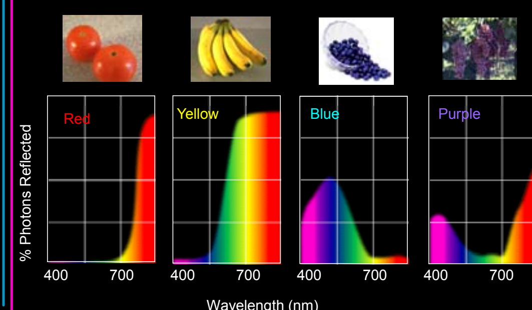

© Stephen E. Palmer, 2002

Transmission Demo

- Gels
- Glasses

Spectra are smooth

- Physical process involved with transmission/reflection means that spectra are similar at nearby wavelengths
- Can model spectra in low dimensional space
  - Principal Components Analysis (PCA)

Low-dimensional models for color spectra - PCA

$$\begin{pmatrix} \vdots \\ e(\lambda) \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots & \vdots & \vdots \\ E_1(\lambda) & E_2(\lambda) & E_3(\lambda) \\ \vdots & \vdots & \vdots \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix}$$

How to find a linear model for color spectra:

- --form a matrix, D, of measured spectra, 1 spectrum per column.
- --[u, s, v] = svd(D) satisfies D = u*s*v‘
- --the first n columns of u give the best (least-squares optimal) n-dimensional linear bases for the data, D:

$$D \approx u(:,1:n) * s(1:n,1:n) * v(1:n,:)'$$

Slide credit:
W. Freeman

Basis functions for Macbeth color checker

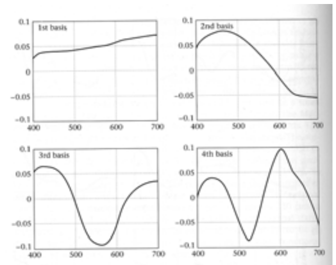

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

n-dimensional linear models for color spectra

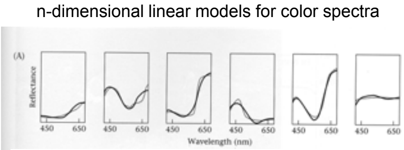

n = 3

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

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

n-dimensional linear models for color spectra

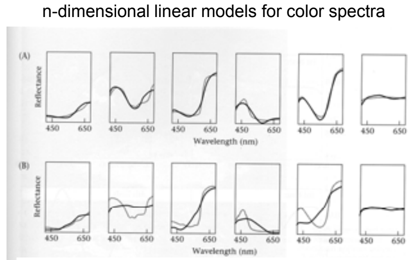

n = 3

n = 2

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

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

n-dimensional linear models for color spectra

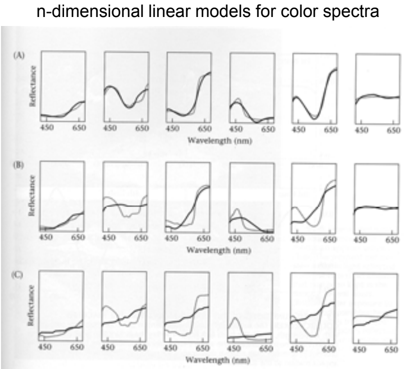

n = 3

n = 2

n = 1

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

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

Reconstruction of MacBeth color chart

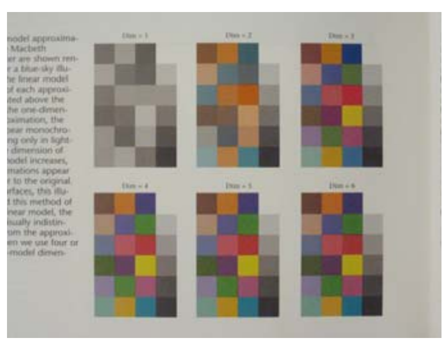

From Foundation of Vision by Brian Wandell, Sinauer Associates, 1995

Color names for cartoon spectra

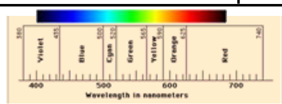

Slide credit: W. Freeman

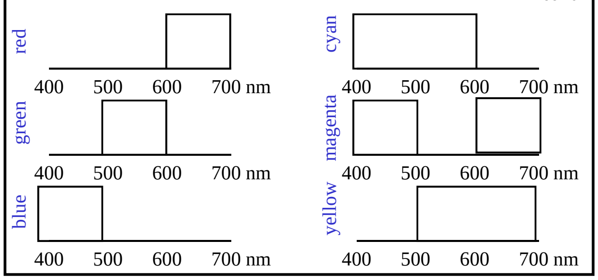

Additive color mixing

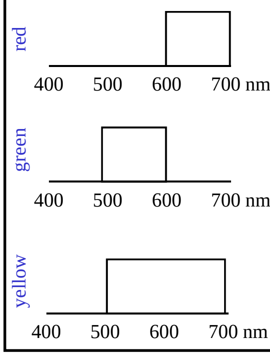

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen, Polachrome slide film.

Red and green make…

Yellow!

Slide credit: W. Freeman

Additive Color mixing

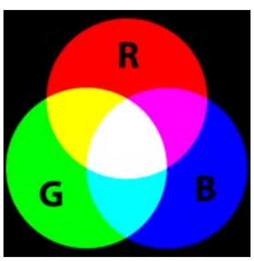

Wikipedia

Subtractive color mixing

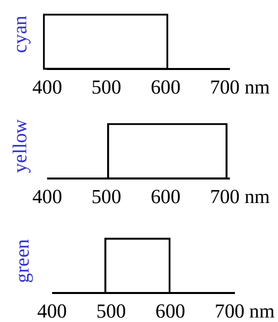

When colors combine by *multiplying* the color spectra. Examples that follow this mixing rule: most photographic films, paint, cascaded optical filters, crayons.

Cyan and yellow (in crayons, called “blue” and yellow) make…

Green!

Slide credit: W. Freeman

Subtractive color mixing

- Light reflecting off colored object
- E.g. printing inks

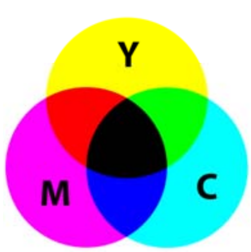

Wikipedia

Subtractive mixing demo

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

The Eye

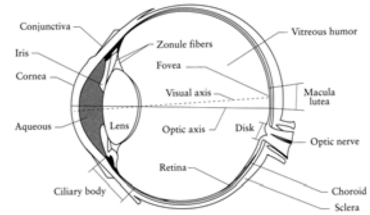

The human eye is a camera!

- **Iris** - colored annulus with radial muscles
- **Pupil** - the hole (aperture) whose size is controlled by the iris
- What's the "film"?
  - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

The Retina

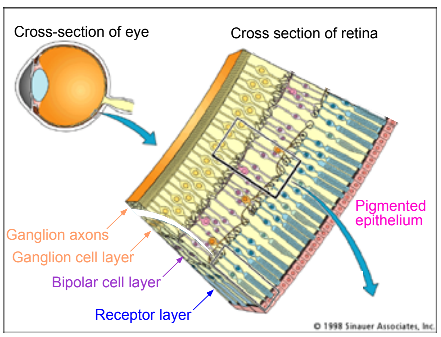

Retina up-close

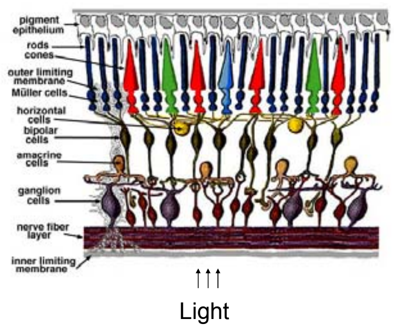

Two types of light-sensitive receptors

Cones
cone-shaped
less sensitive
operate in high light
color vision

Rods
rod-shaped
highly sensitive
operate at night
gray-scale vision

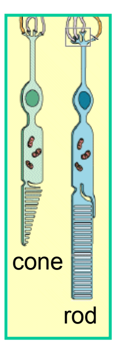

Human Photoreceptors

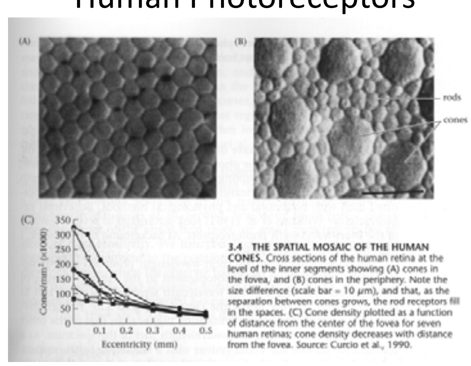

(From Foundations of Vision, by Brian Wandell, Sinauer Assoc.)

Distribution of Rods and Cones

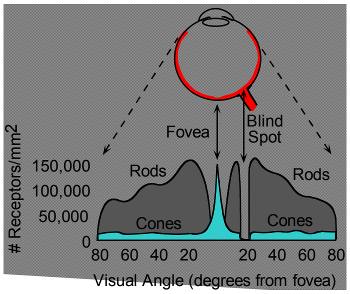

Night Sky: why are there more stars off-center?

Rod / Cone sensitivity

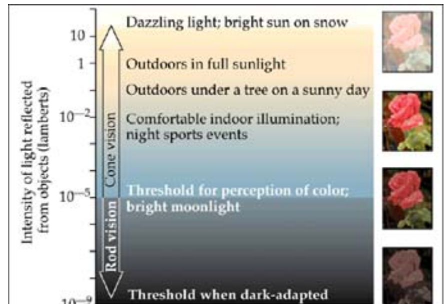

The famous sock-matching problem…

Human eye photoreceptor spectral sensitivities

Measurement of cone photocurrents

3.3 SPECTRAL SENSITIVITIES OF THE L-, M-, AND S-CONES in the human eye. The measurements are based on a light source at the cornea, so that the wavelength loss due to the cornea, lens, and other inert pigments of the eye plays a role in determining the sensitivity. Source: Stockman and MacLeod, 1993.

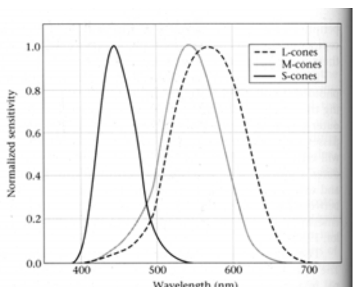

What colors would these look like?

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

Physiology of Color Vision

Three kinds of cones:

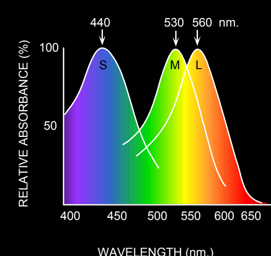

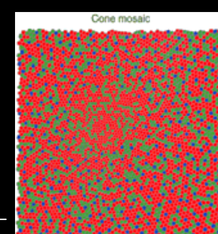

- Why are M and L cones so close?
- Are are there 3?

Why specify color numerically?

- Accurate color reproduction is commercially valuable
  - Many products are identified by color ("golden" arches);
- Few color names are widely recognized by English speakers -
  - About 10; other languages have fewer/more, but not many more.
  - It's common to disagree on appropriate color names.
- Color reproduction problems increased by prevalence of digital imaging - eg. digital libraries of art.
  - How do we ensure that everyone sees the same color?

Forsyth & Ponce

An assumption that sneaks in here

- For now we will assume that the spectrum of the light arriving at your eye completely determines the perceived color.
- But we know color appearance really depends on:
  - The illumination
  - Your eye's adaptation level
  - The colors and scene interpretation surrounding the observed color.

Slide credit: W. Freeman

Color matching experiment

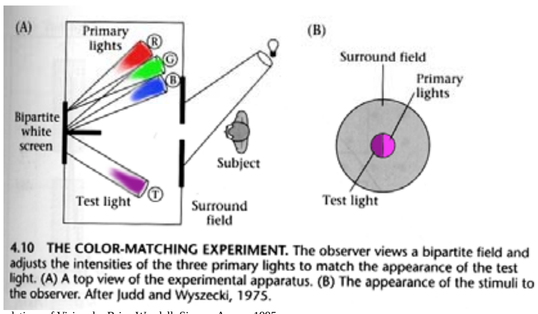

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

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

Color matching experiment 1

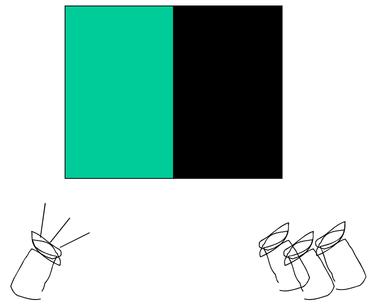

Color matching experiment 1

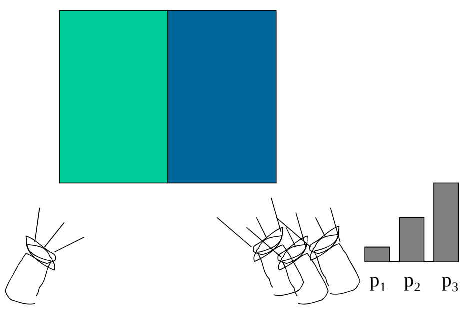

Color matching experiment 1

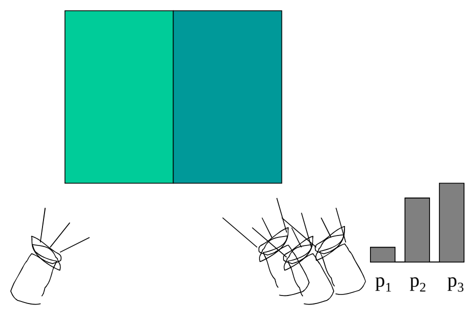

Color matching experiment 1

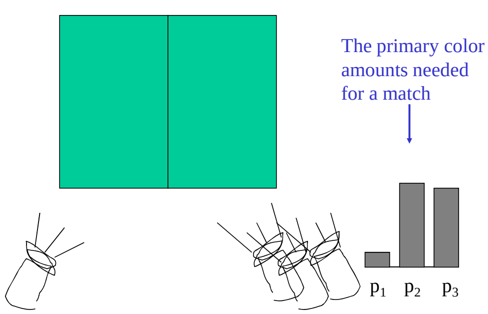

Color matching experiment 2

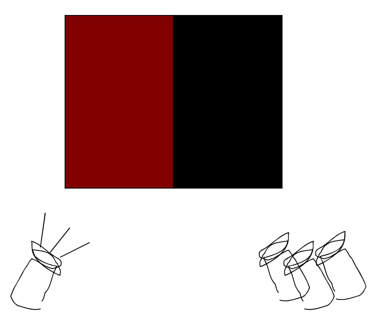

Color matching experiment 2

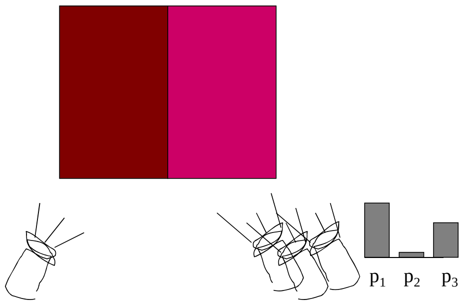

Color matching experiment 2

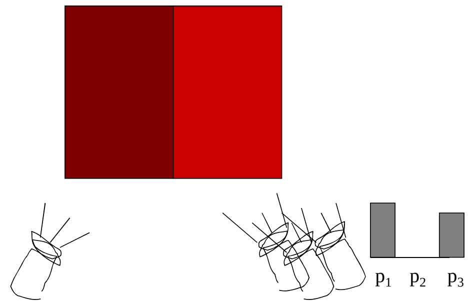

Color matching experiment 2

We say a "negative" amount of p_2 was needed to make the match, because we added it to the test color's side.

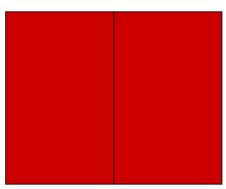

The primary color amounts needed for a match:

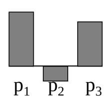

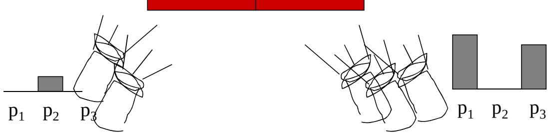

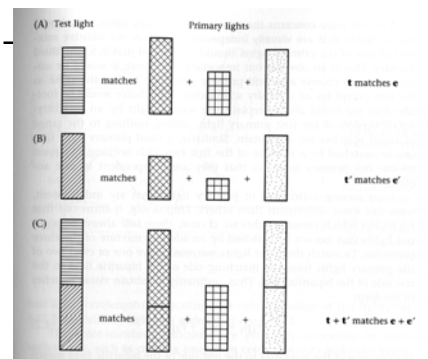

4.12 THE COLOR-MATCHING EXPERIMENT SATISFIES THE PRINCIPLE OF SUPERPOSITION. In parts (A) and (B), test lights are matched by a mixture of three primary lights. In part (C) the sum of the test lights is matched by the additive mixture of the primaries, demonstrating superposition.

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

Grassman's Laws

- For color matches:
  - symmetry: U=V <=>V=U
  - transitivity: U=V and V=W => U=W
  - proportionality: U=V <=> tU=tV
  - additivity: if any two (or more) of the statements
    - U=V,
    - W=X,
    - (U+W)=(V+X) are true, then so is the third
- These statements are as true as any biological law. They mean that additive color matching is linear.

Forsyth & Ponce

Measure color by color-matching paradigm

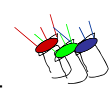

- Pick a set of 3 primary color lights.
- Find the amounts of each primary, e_1, e_2, e_3, needed to match some spectral signal, t.
- Those amounts, e_1, e_2, e_3, describe the color of t. If you have some other spectral signal, s, and s matches t perceptually, then e_1, e_2, e_3 will also match s, by Grassman's laws.
- Why this is useful—it lets us:
  - Predict the color of a new spectral signal
  - Translate to representations using other primary lights.

Goal: compute the color match for any color signal for any set of primary colors

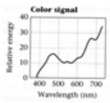

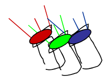

- Examples of why you'd want to do that:
  - Want to paint a carton of Kodak film with the Kodak yellow color.
  - Want to match skin color of a person in a photograph printed on an ink jet printer to their true skin color.
  - Want the colors in the world, on a monitor, and in a print format to all look the same.

How to compute the color match for any color signal for any set of primary colors

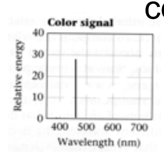

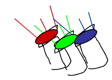

- Pick a set of primaries, $p_1(\lambda), p_2(\lambda), p_3(\lambda)$
- Measure the amount of each primary, $c_1(\lambda), c_2(\lambda), c_3(\lambda)$ needed to match a monochromatic light, $t(\lambda)$ at each spectral wavelength λ (pick some spectral step size). These are called the color matching functions.

Color matching functions for a particular set of monochromatic primaries

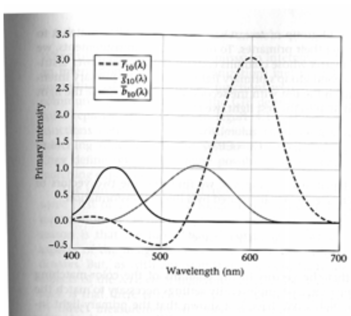

p_1 = 645.2 nm
p_2 = 525.3 nm
p_3 = 444.4 nm

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

Using the color matching functions to predict the primary match to a new spectral signal

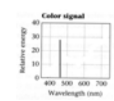

We know that a monochromatic light of wavelength λ_i will be matched by the amounts $c_1(\lambda_i), c_2(\lambda_i), c_3(\lambda_i)$ of each primary.

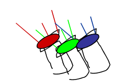

And any spectral signal can be thought of as a linear combination of very many monochromatic lights, with the linear coefficient given by the spectral power at each wavelength.

$$\vec{t} = \begin{pmatrix} t(\lambda_1) \\ \vdots \\ t(\lambda_N) \end{pmatrix}$$

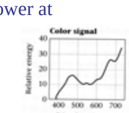

Using the color matching functions to predict the primary match to a new spectral signal

Store the color matching functions in the rows of the matrix, C

$$C = \begin{pmatrix} c_1(\lambda_1) & \cdots & c_1(\lambda_N) \\ c_2(\lambda_1) & \cdots & c_2(\lambda_N) \\ c_3(\lambda_1) & \cdots & c_3(\lambda_N) \end{pmatrix}$$

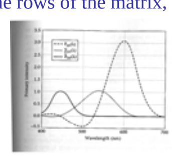

Let the new spectral signal be described by the vector t.

$$\vec{t} = \begin{pmatrix} t(\lambda_1) \\ \vdots \\ t(\lambda_N) \end{pmatrix}$$

Then the amounts of each primary needed to match t are: $C\vec{t}$

Switching between matching functions

- Color matching functions are not unique
- Depend on set of primaries!
- Map between coordinates in two bases using 3x3 matrix

How do you translate colors between different systems of primaries?

$p_1 = (0\ 0\ 0\ 0\ 0 \ldots\ 0\ 1\ 0)^T$
$p_2 = (0\ 0 \ldots 0\ 1\ 0 \ldots 0\ 0)^T$
$p_3 = (0\ 1\ 0\ 0 \ldots 0\ 0\ 0\ 0)^T$
Primary spectra, P
Color matching functions, C

$p'_1 = (0\ 0.2\ 0.3\ 4.5\ 7 \ldots 2.1)^T$
$p'_2 = (0.1\ 0.44\ 2.1 \ldots\ 0.3\ 0)^T$
$p'_3 = (1.2\ 1.7\ 1.6 \ldots\ 0\ 0)^T$
Primary spectra, P'
Color matching functions, C'

$$C\vec{t} = CP'C'\vec{t}$$

Any input spectrum, t
The amount of each primary in P needed to match the color with spectrum t.
The amount of each P' primary needed to match t
The spectrum of a perceptual match to t, made using the primaries P'
The color of that match to t, described by the primaries, P.

So, how to translate from the color in one set of primaries to that in another:

The values of the 3 primaries, in the unprimed system
The values of the 3 primaries, in the primed system

$$e = CP'e'$$

a 3x3 matrix

C
P'

P' are the old primaries
C are the new primaries' color matching functions

Comparison of color matching functions with best 3x3 transformation of cone responses

4.20 COMPARISON OF CONE PHOTOCURRENT RESPONSES AND THE COLOR-MATCHING FUNCTIONS. The cone photocurrent spectral responsivities are within a linear transformation of the color-matching functions, after a correction has been made for the optics and inert pigments in the eye. The smooth curves show the Stiles and Burch (1959) color-matching functions. The symbols show the matches predicted from the photocurrents of the three types of macaque cones. The predictions included a correction for absorption by the lens and other inert pigments in the eye. Source: Baylor, 1987.

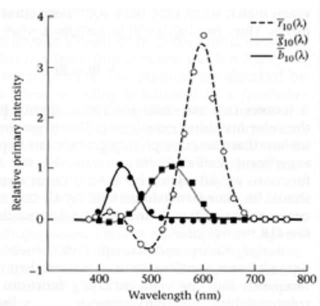

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Summary so far

- Set of non-unique color matching functions that can be used to describe color in 3-coordinate system
- Can describe physiology experiments that directly electrical signals from cone cell
- Mapping between matching functions takes form of 3x3 matrix
- Let's now look at the various color spaces used to describe color

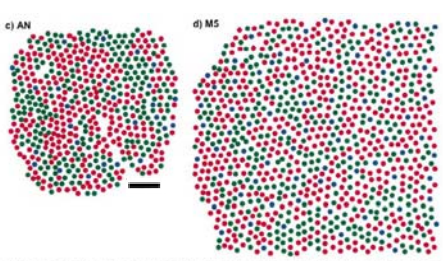

Fig. 2. L, M and S cone mosaics for two humans: JW (a nasal and a temporal location is shown, labeled JWN and JWT, respectively); and AN, and one macaque, M5. L, M and S cones are shown as red, green and blue dots respectively. For JWN, a patch of central cones was not identified due to a capillary that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

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

Basic issues

- How to represent a continuous spectrum in a three coordinate system
- Cover wide range of colors
- Limited dynamic range (8bits/channel)
- Preferably be intuitive

CIE Standard Observer Color matching functions

- 3-D coordinate space
- Difficult to visualize in 2-D

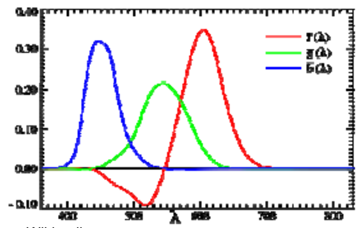

Wikipedia

Define:

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

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

Constant brightness

CIE rg Chromaticity space

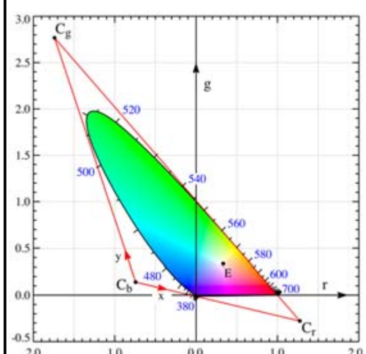

$$r = \frac{R}{R+G+B},$$

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

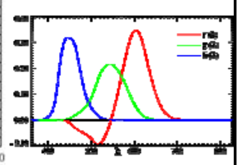

Wikipedia

CIE XYZ color-matching functions (1931)

- No negative light
- Green matches luminosity (overall brightness)
- Primaries not physically realizable
- Not an intuitive representation

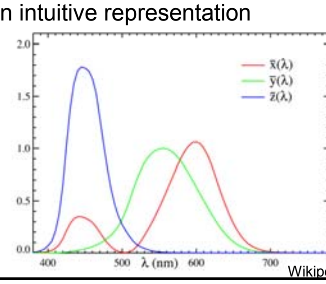

Wikipedia

Mapping to CIE XYZ (1931)

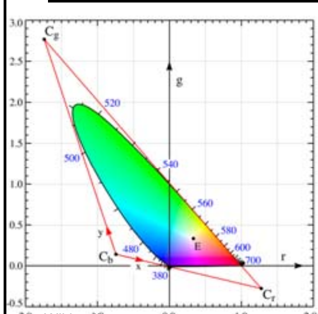

Wikipedia

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \frac{1}{b_{21}} \begin{bmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} = \frac{1}{0.17697} \begin{bmatrix} 0.49 & 0.31 & 0.20 \\ 0.17697 & 0.81240 & 0.01063 \\ 0.00 & 0.01 & 0.99 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

X,Y,Z – tristimulus values

3x3 mapping matrix

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

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

$$z = \frac{Z}{X+Y+Z} = 1 - x - y$$

CIE XYZ Chromaticity space

- 2-D slice of inherently 3-D space
- Many other spaces but are derivatives of this
- Used to show range of colors obtainable, known as the *gamut*

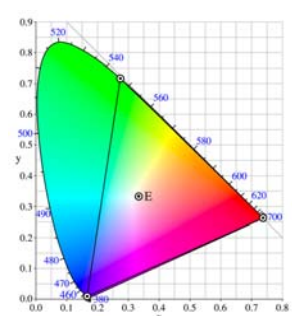

Wikipedia

sRGB color space

- Standard space used for 8-bit/channel color display
- Representation implicit for most images/application.
- Limited gamut (other spaces exist with larger gamut)
- Incorporates non-linearities inherent in display devices (1/
- Mapping from CIE XYZ to sRGB has two parts:

1. $$\begin{bmatrix} R_{linear} \\ G_{linear} \\ B_{linear} \end{bmatrix} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

2. $$C_{srgb} = \begin{cases} 12.92 C_{linear}, & C_{linear} \le 0.0031308 \\ (1+a) C_{linear}^{1/2.4} - a, & C_{linear} > 0.0031308 \end{cases}$$

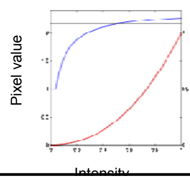

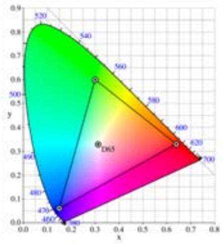

Gamma correction

- Tone response curve

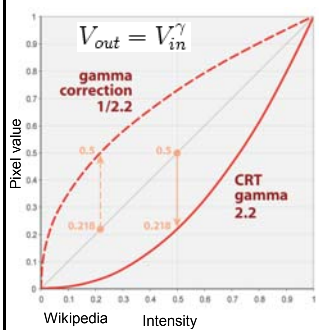

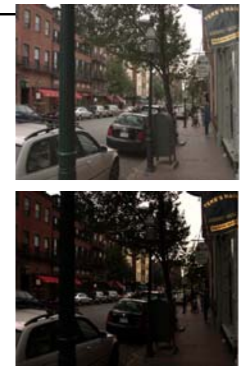

Typical CMYK color space

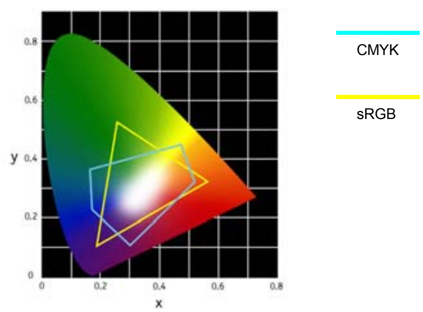

http://www.tasi.ac.uk/images/cie_gamut.jpg

Intuitive color space

- Hue, Saturation, Brightness
- Very similar to Hue, Satuation, Value

The Psychophysical Correspondence

There is no simple functional description for the perceived color of all lights under all viewing conditions, but ……

A helpful constraint:
Consider only physical spectra with normal distributions

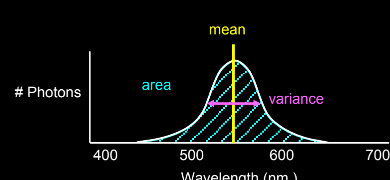

The Psychophysical Correspondence

Mean ⟷ Hue

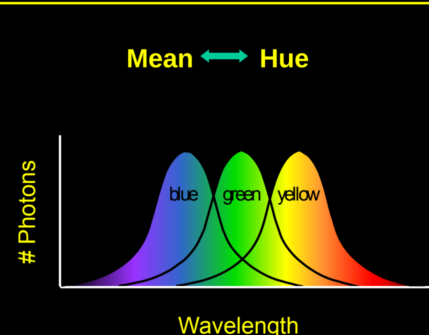

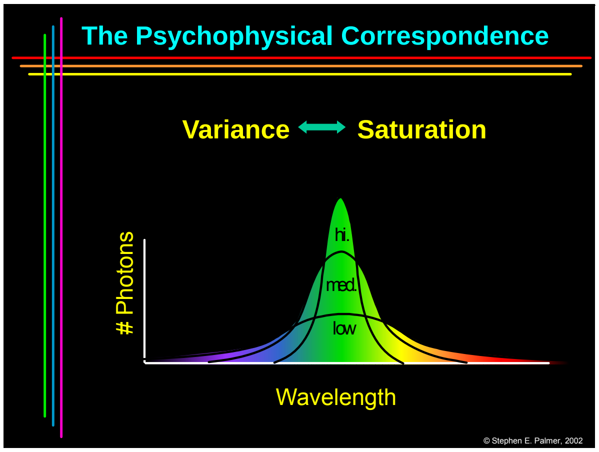

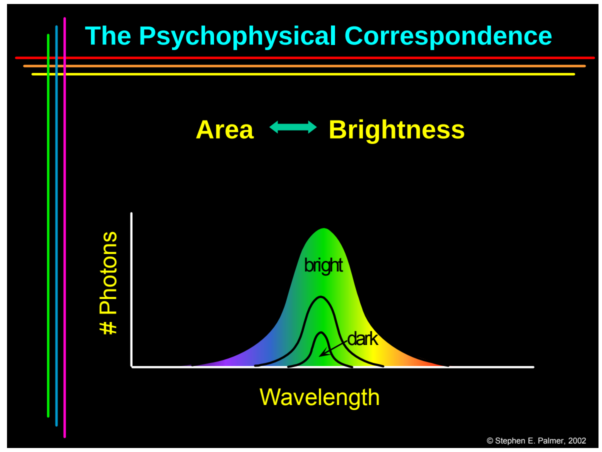

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

DSLR Camera sensor

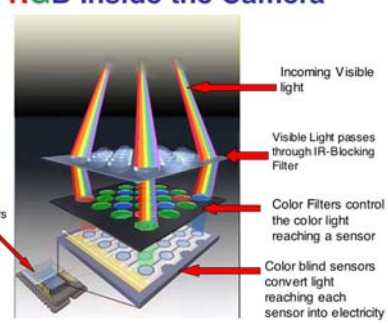

http://www.photoaxe.com/wp-content/uploads/2007/04/camera-sensor.jpg

Image Acquisition Pipeline

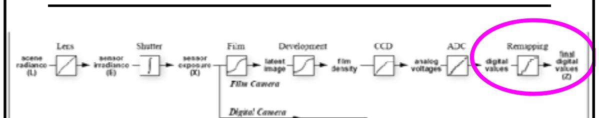

Debevec & Malik '97

Camera primaries given by dies on sensor

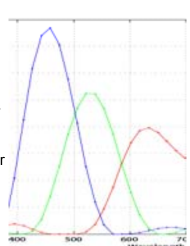

3x3 matrix (dependent on camera model) maps to XYZ color space

Then convert to sRGB or Adobe RGB etc.

RAW file is raw output of the sensor And allow manual control of these conversion steps.

Sensitivity (ISO)

- Third variable for exposure
- Linear effect (200 ISO needs half the light as 100 ISO)
-

-

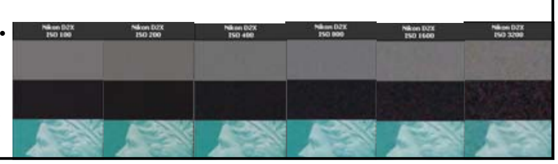

From dpreview.com

CCD color sampling

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

What are some approaches to sensing color images?

- Scan 3 times (temporal multiplexing)
- Use 3 detectors (3-ccd camera)
- Use offset color samples (spatial multiplexing)
- Multiplex in the depth of the sensor (Foveon)

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
- Use spatially offset color samples (spatial multiplexing)
  - Single-chip CCD color cameras
  - Human eye
- Multiplex in the depth of the sensor
  - Foveon

Bayer RGB mosaic

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

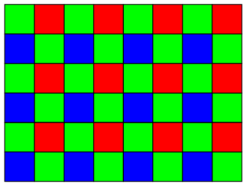

Foveon image sensor

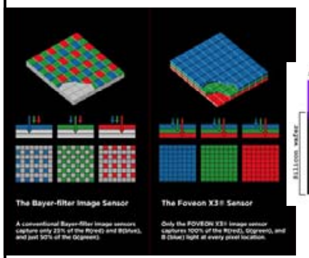

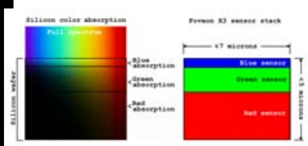

DSLR Camera sensor

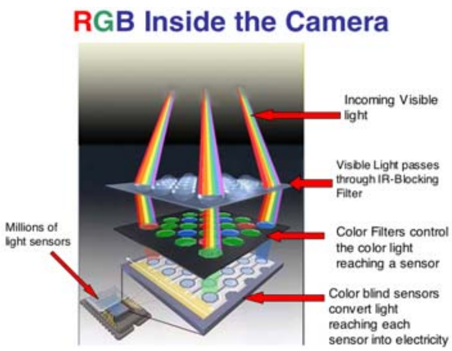

http://www.photoaxe.com/wp-content/uploads/2007/04/camera-sensor.jpg

Camera response curves with IR block filter removed

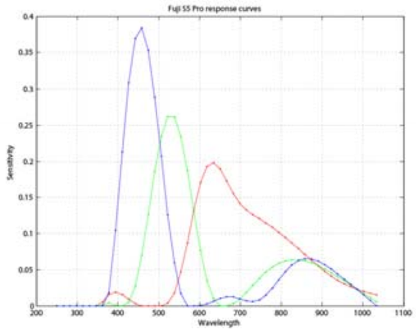

Ultra-Violet photography

http://www.naturfotograf.com/uvstart.html

***Potentilla anserina* © Bjørn Rørslett/NN 2001**

Infra-Red Photography

Normal Visible light

Infra Red (only >715nm)

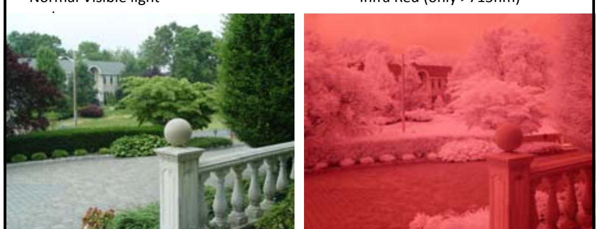

Photos from MaxMax.com

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

Demosaicing

- Interpolate missing values

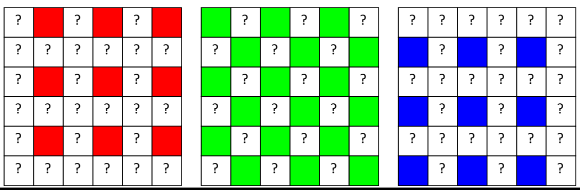

Demosaicing

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

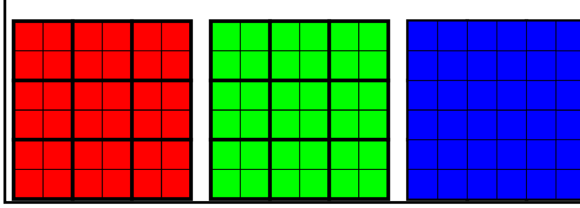

Linear interpolation

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

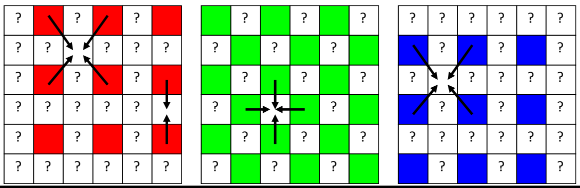

Color Moire patterns

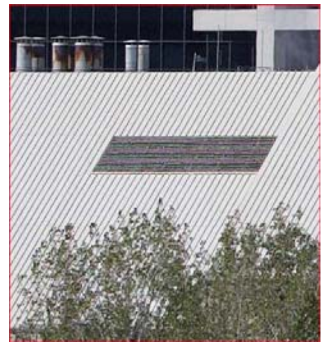

http://forums.dpreview.com/forums/read.asp?forum=1034&message=8115658

CCD color filter pattern

detector

(simplified for simpler visualization)

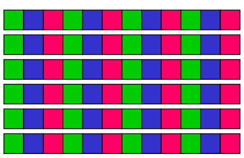

Slide credit: F. Durand

Typical color moire patterns

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

Slide credit: F. Durand

The cause of color moire

detector

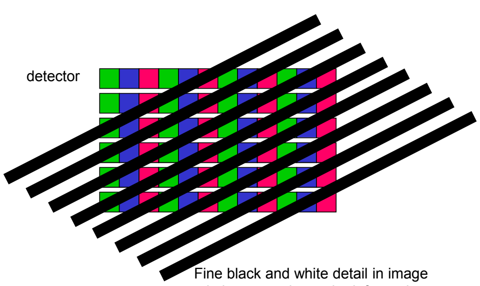

Fine black and white detail in image mis-interpreted as color information.

Slide credit: F. Durand

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…

Slide credit: F. Durand

Median filter

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

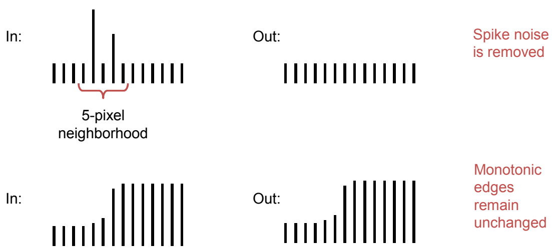

Slide credit: F. Durand

Recombining the median filtered colors

Linear interpolation

Median filter interpolation

Slide credit: F. Durand

Overview of today

- Physics of color
- Human encoding of color
- Color spaces
- Camera sensor & color
- Demosaicing
- White balancing

White balance problem

- Spectrum of light source affects color of objects in scene
- We need to discount the color of the light source
- We often don't notice since we adapt to the illuminant of the room, not that of the scene in the picture

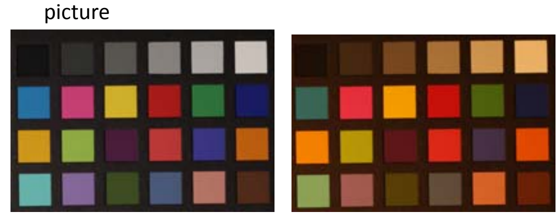

Von Kries adaptation

- Multiply each channel by a gain factor
- Note that the light source could have a more complex effect
  - Arbitrary 3x3 matrix
  - More complex spetrum transformation

Slide credit: F. Durand

Best way to do white balance

- Grey card:
- Take a picture of a neutral object (white or gray)
- Deduce the weight of each channel
- If the object is recoded as r_w, g_w, b_w use weights $1/r_w$, $1/g_w$, $1/b_w$

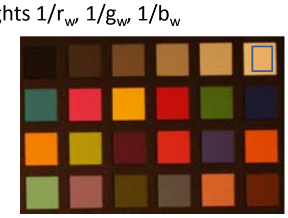

Slide credit: F. Durand

Grey world assumption

- The average color in the image is grey
- Use weights

$$\frac{1}{\int_{image} r}, \frac{1}{\int_{image} g}, \frac{1}{\int_{image} b}$$

- Note that this also sets the exposure/brightness
- Usually assumes 18% grey

Slide credit: F. Durand

Brightest pixel assumption

- Highlights usually have the color of the light source
  - At least for dielectric materials
- Do white balance by using the brightest pixels
  - Plus potentially a bunch of heuristics
  - In particular use a pixel that is not saturated/clipped

Slide credit: F. Durand

End

Sensitivity to blur