NEURL-GA 3042.005 -- Spring, 2014
(cross-listed as: MATH-GA 2855, CSCI-GA 2715, PSYCH-GA 3405.005)

Representation and Analysis of Visual Images

Eero Simoncelli

01 - Light, reflectance, imaging, color

From M. Landy and J. A. Movshon (eds), *Computational Models of Visual Processing* (pp. 3-20). Cambridge, MA: MIT Press (1991).

1

The Plenoptic Function and the Elements of Early Vision

Edward H. Adelson and
James R. Bergen

What are the elements of early vision? This question might be taken to mean, What are the fundamental atoms of vision?—and might be variously answered in terms of such candidate structures as edges, peaks, corners, and so on. In this chapter we adopt a rather different point of view and ask the question, What are the fundamental *substances* of vision? This distinction is important because we wish to focus on the first steps in extraction of visual information. At this level it is premature to talk about discrete objects, even such simple ones as edges and corners.

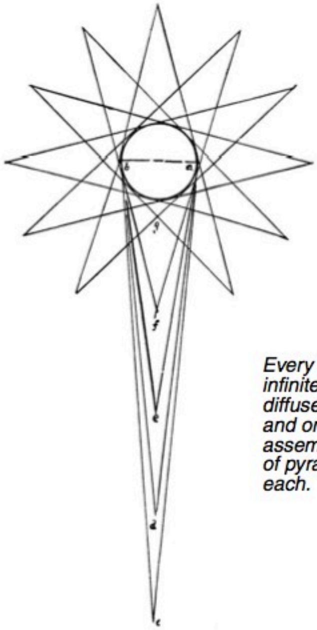

Every body in light and shade fills the surrounding air with infinite images of itself; and these, by infinite pyramids diffused in the air, represent this body throughout space and on every side. Each pyramid that is composed of a long assemblage of rays includes within itself an infinite number of pyramids and each has the same power as all, and all as each.

– The Notebooks of Leonardo da Vinci

(late 1400's)

Spectral nature of light

[Newton, 1665]

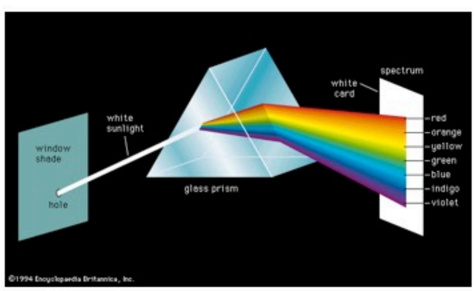

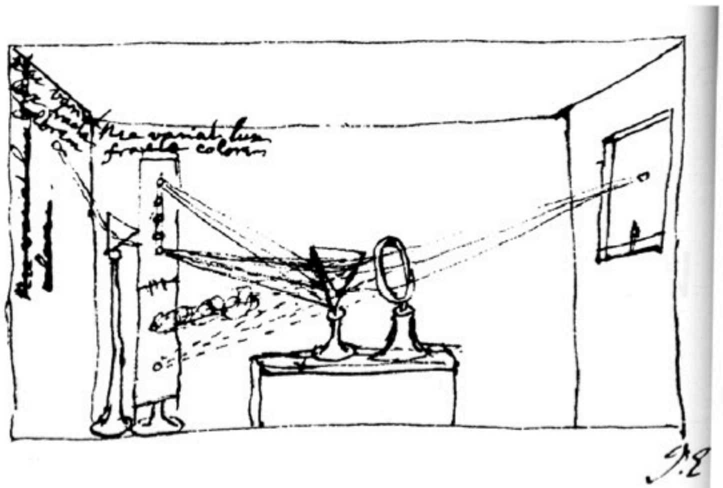

4.1 **NEWTON'S SUMMARY DRAWING** of his experiments with light. Using a point source of light and a prism, Newton separated sunlight into its fundamental components. By reconverging the rays, he also showed that the decomposition is reversible.

[from Wandell: Foundations of Vision, 1995]

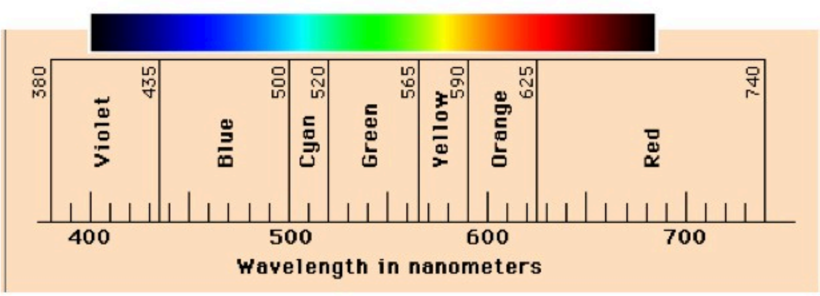

The "Plenoptic" function

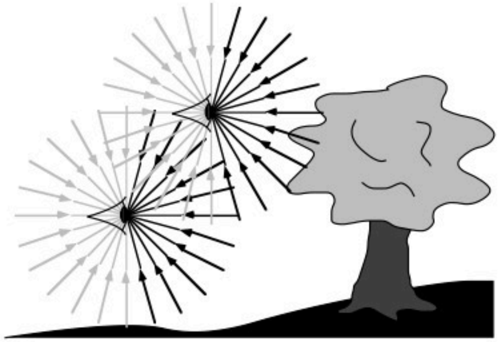

$$I(x, y, \lambda, t, V_x, V_y, V_z)$$

(everything there is to see)

[Adelson & Bergen 91]

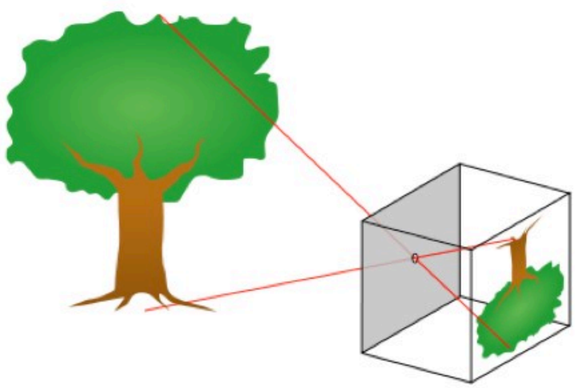

Imaging: capture of the plenoptic function on a 2D sensor surface

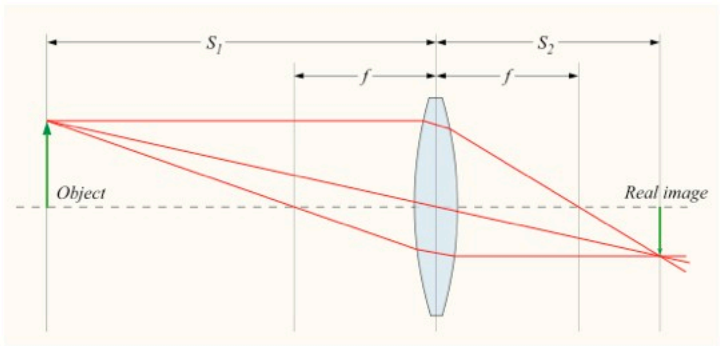

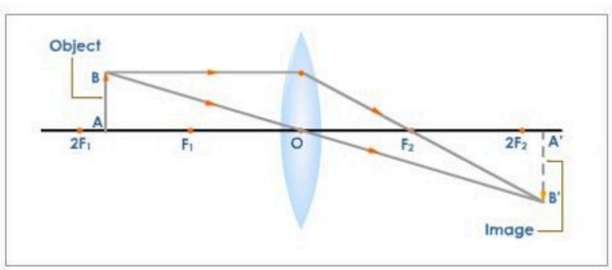

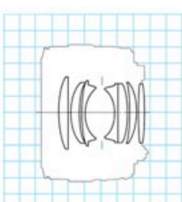

Canon, 50mm, f1.4

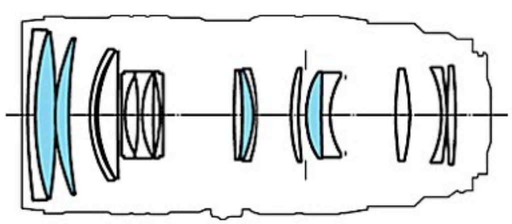

Canon, 70-200, f2.8

First production camera ("Daguerrotype"), 1839

14-bit CMOS sensor

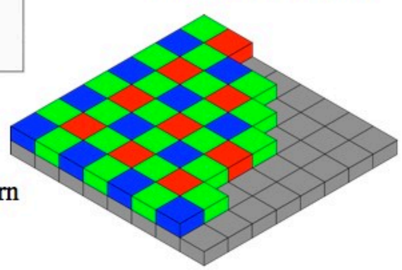

"Bayer" pattern

$I(x, y, \lambda, t, V_x, V_y, V_z)$? Integrated over viewpoint and time, sampled in (x,y) and wavelength

Perceptual color matching experiment

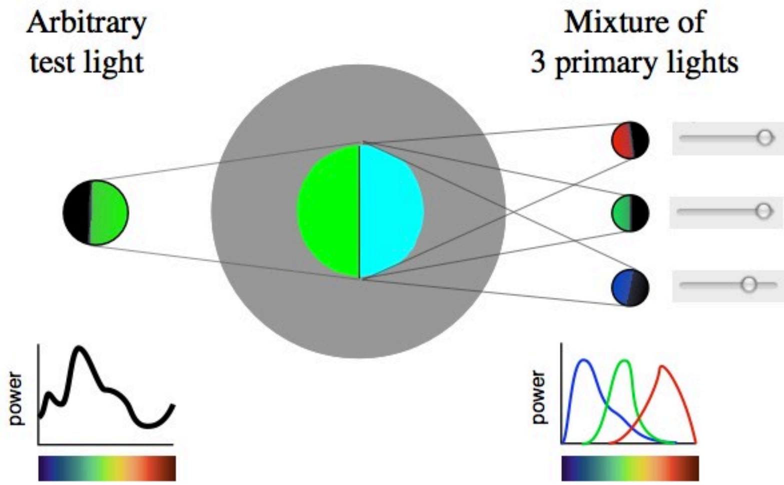

[Young, Helmholtz, Grassman, etc, 1800's; slide c/o D. Brainard]

Perceptual color matching experiment

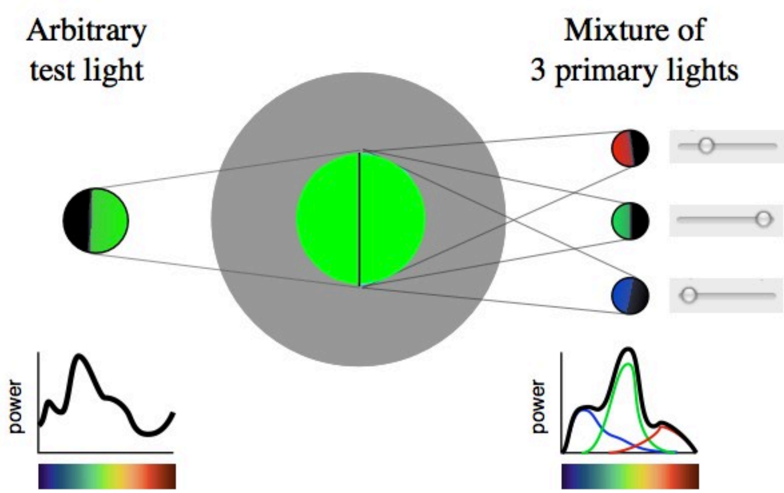

[Young, Helmholtz, Grassman, etc, 1800's; slide c/o D. Brainard]

Grassmann's Laws, 1853

1) Any light can be matched with a mixture of 3 primaries
2) Sum of 2 lights results in a sum of the corr. mixtures
3) Rescaling the light results in a rescaled mixture

➡ Matching can be described by an Nx3 linear system!

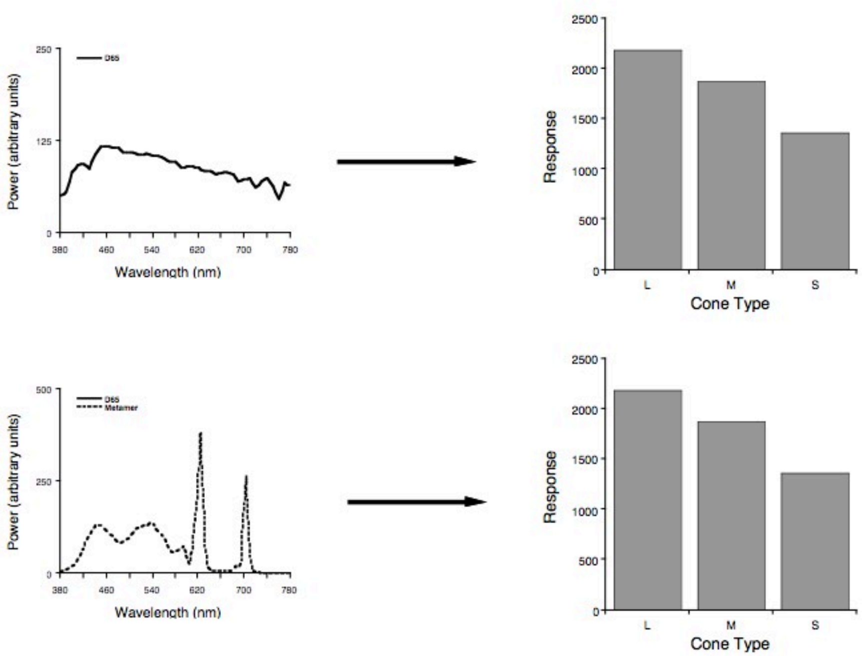

Theory: the visual system projects the wavelength spectra of light onto a 3-dimensional space

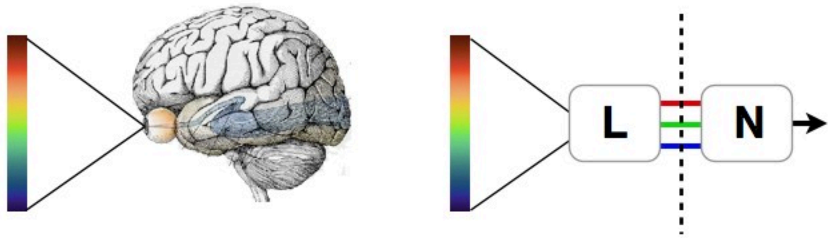

- Accurately predicts perceptual limitations
- Basis for color technology standards (CIE, 1931)
- Underlying mechanism (cones) verified 100+ years later

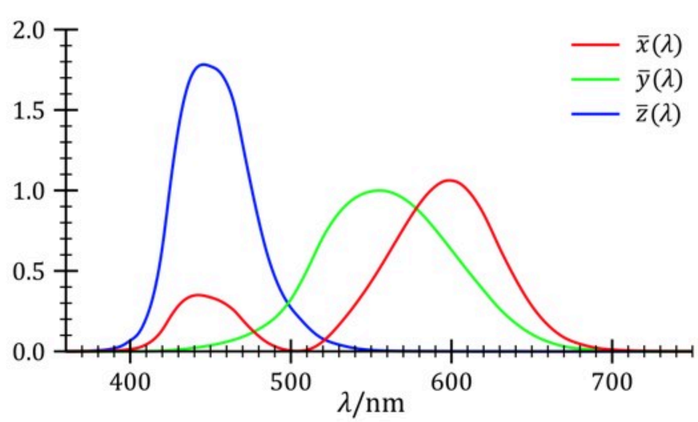

CIE “standard observer” color matching functions (1931)

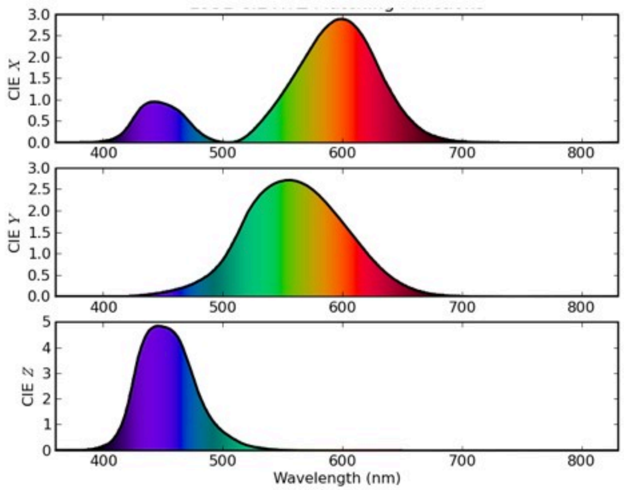

CIE "standard observer" color matching functions (1931) used to define the standard "XYZ" color coordinate system

CIE-LUV color space (1976)

One (of many) attempts to create a perceptually-uniform color space.

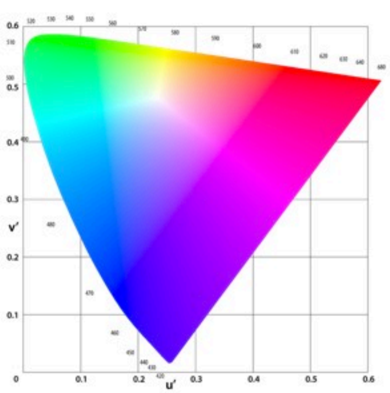

$$L^* = \begin{cases} \left(\frac{29}{3}\right)^3 Y/Y_n, & Y/Y_n \le \left(\frac{6}{29}\right)^3 \\ 116\,(Y/Y_n)^{1/3} - 16, & Y/Y_n > \left(\frac{6}{29}\right)^3 \end{cases}$$

$$u^* = 13L^* \cdot (u' - u'_n)$$

$$v^* = 13L^* \cdot (v' - v'_n)$$

$$u' = \frac{4X}{X + 15Y + 3Z}$$

$$v' = \frac{9Y}{X + 15Y + 3Z}$$

ICC profiles

- Standard for specifying color coordinate system for displays and sensors (often embedded in digital photo files)
- 3x3 matrix [relative to XYZ coords]
- point nonlinearities [power, or lookup table]

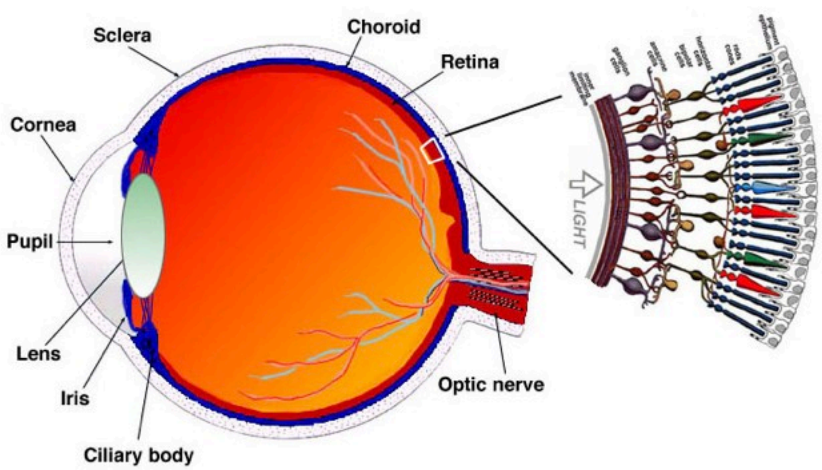

[figure: John Moran Eye Center, U. Utah]

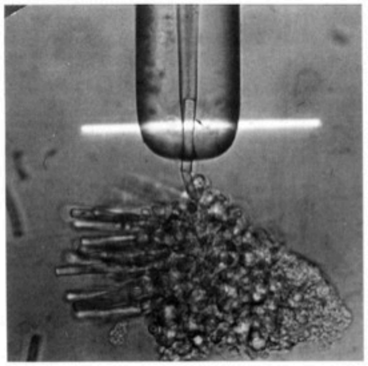

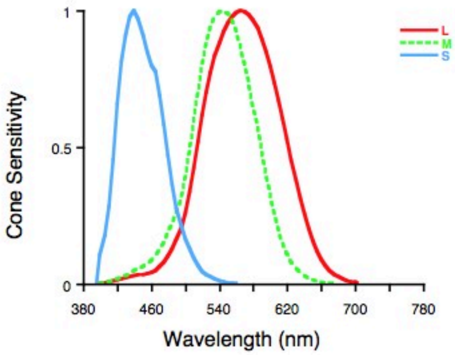

[Baylor, Nunn & Schnapf, 1987]

[figure: David Brainard]

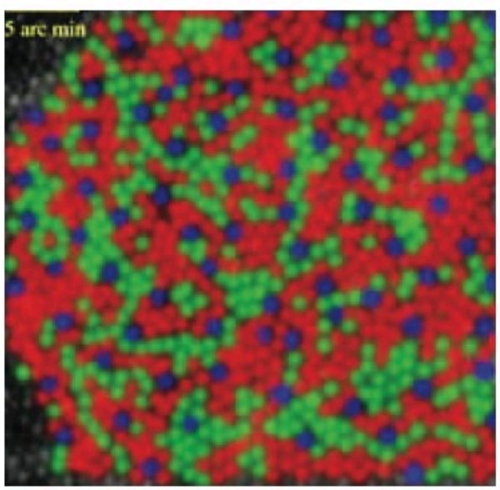

Cone photoreceptor mosaic near fovea

- Roorda and Williams (1999)

$I(x, y, \lambda, t, V_x, V_y, V_z)$? Integrated over viewpoint and time, sampled in (x,y) and wavelength!

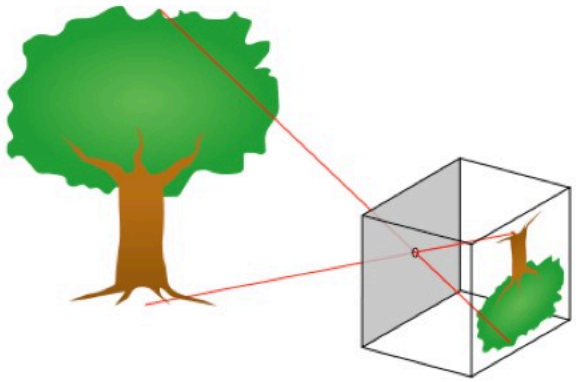

Light sources (illuminants)

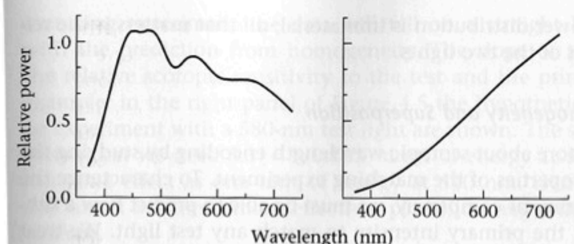

4.4 THE SPECTRAL POWER DISTRIBUTION of two important light sources are shown: (left) blue skylight and (right) a tungsten bulb.

[Wandell: Foundations of Vision, 1995]

Reflectance

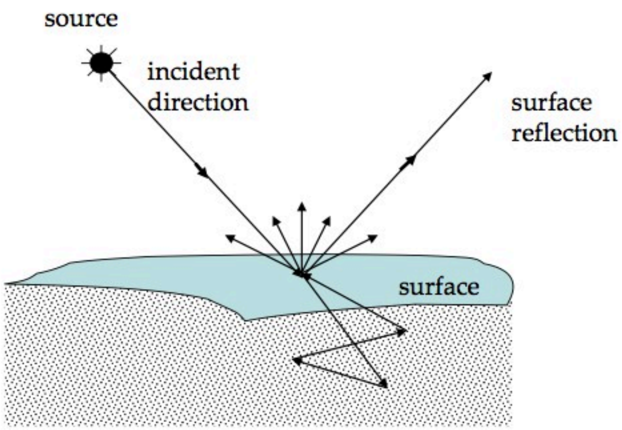

Reflectance functions

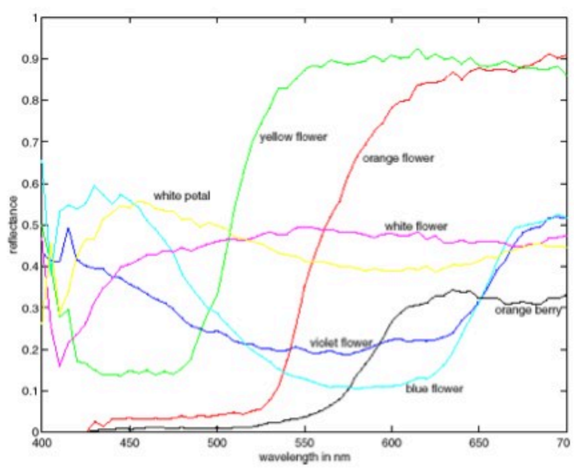

[Forsyth, 2002]

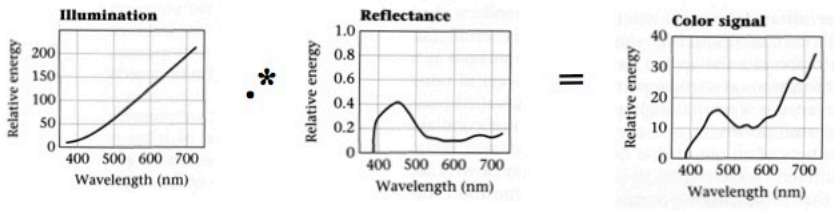

1) Graphics is easy, vision is hard/impossible :)

2) In fact, humans are quite good (but not perfect) at separating illumination from reflectance.

[Wandell: Foundations of Vision, 1995]

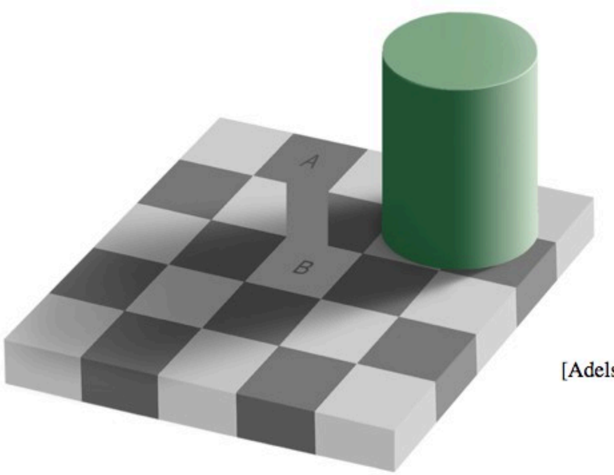

[Adelson, 1995]

Even in grayscale, humans are pretty good at estimating **reflectance**, and local intensity changes, but are very poor at estimating absolute intensities

What's color perception for?

Perhaps humans try to estimate reflectance, regardless of lighting?

[c/o David Brainard]

Change of lighting, or change of paint?

[photo: Brainard 2003]

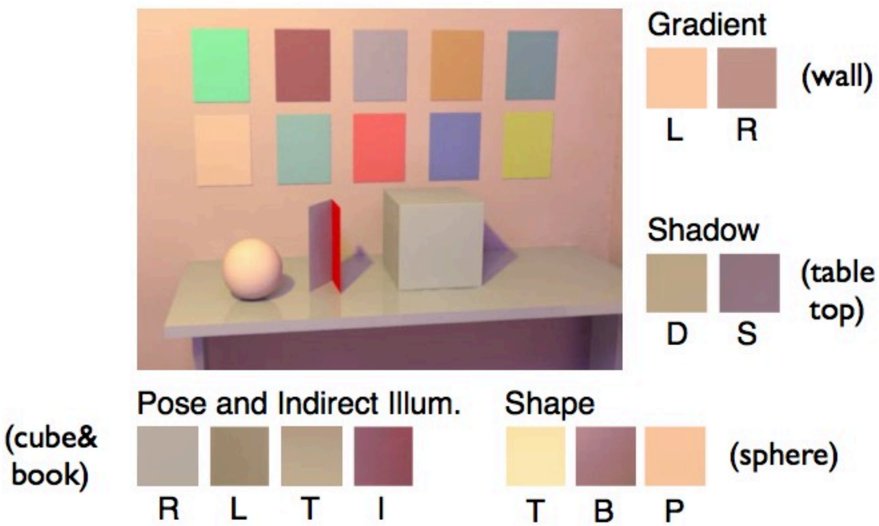

[Brainard 2003]

Why don't colors look right in photographs? Sources of failure:

- viewing conditions
- multiple light sources
- sensor not matched to cone absorption space

Auto white balance? Sources of information:

- distribution of pixel colors in scene
- brightest point in scene (specularities)
- prior assumptions about illuminants, reflectances
- inter-reflections

red
rose
or
red
light?

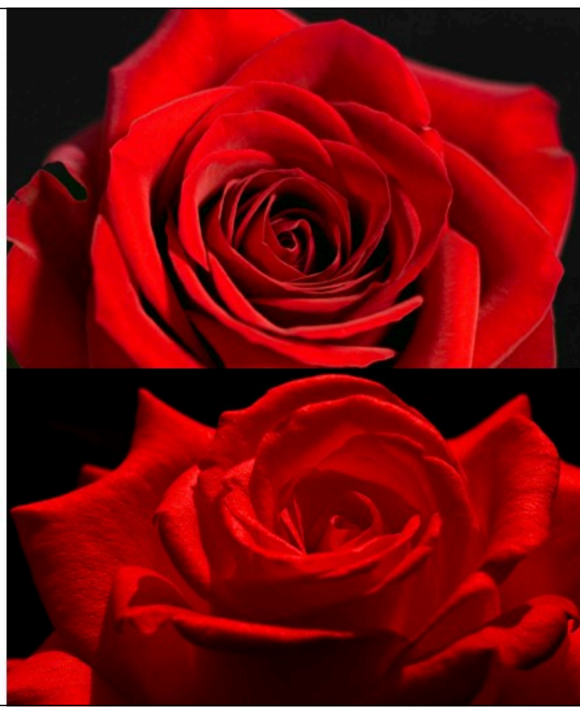

red
rose
or
red
light?

$$\text{red rose, white light :} \quad \int_\lambda c_i(\lambda)\; w(\lambda) \sum_{n=1} \alpha_n\; r(\lambda)^n$$

$$\text{red light, white rose :} \quad \int_\lambda c_i(\lambda)\; r(\lambda) \sum_{n=1} \alpha_n\; w(\lambda)^n$$

$c_i(\lambda)$: sensitivity of ith cone type
$w(\lambda)$: white (either reflectance or spectral power rescaled to $[0, 1]$
$r(\lambda)$: red(...)
α_n : proportion of light with n bounces

Inter-reflections...

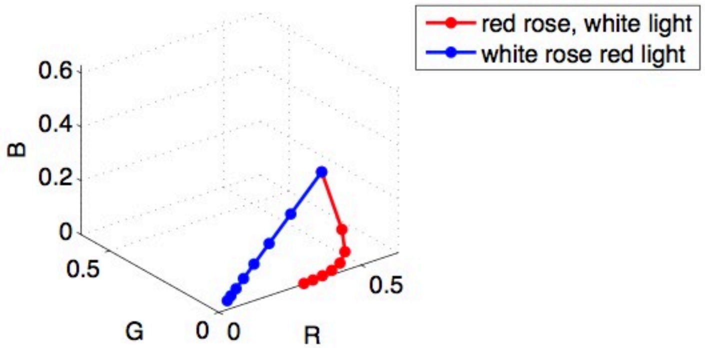

Bayesian color constancy

David H. Brainard

Department of Psychology, University of California, Santa Barbara, California 93106

William T. Freeman

MERL, a Mitsubishi Electric Research Laboratory, Cambridge, Massachusetts 02139

Received August 30, 1996; revised manuscript received December 23, 1996; accepted January 7, 1997

The problem of color constancy may be solved if we can recover the physical properties of illuminants and surfaces from photosensor responses. We consider this problem within the framework of Bayesian decision theory. First, we model the relation among illuminants, surfaces, and photosensor responses. Second, we construct prior distributions that describe the probability that particular illuminants and surfaces exist in the world. Given a set of photosensor responses, we can then use Bayes's rule to compute the posterior distribution for the illuminants and the surfaces in the scene. There are two widely used methods for obtaining a single best estimate from a posterior distribution. These are maximum *a posteriori* (MAP) and minimum mean-squared-error (MMSE) estimation. We argue that neither is appropriate for perception problems. We describe a new estimator, which we call the maximum local mass (MLM) estimate, that integrates local probability density. The new method uses an optimality criterion that is appropriate for perception tasks: It finds the most probable approximately correct answer. For the case of low observation noise, we provide an efficient approximation. We develop the MLM estimator for the color-constancy problem in which flat matte surfaces are uniformly illuminated. In simulations we show that the MLM method performs better than the MAP estimator and better than a number of standard color-constancy algorithms. We note conditions under which even the optimal estimator produces poor estimates: when the spectral properties of the surfaces in the scene are biased.

Reflectance Prior

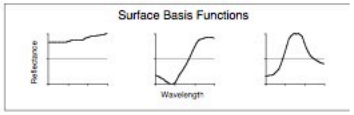

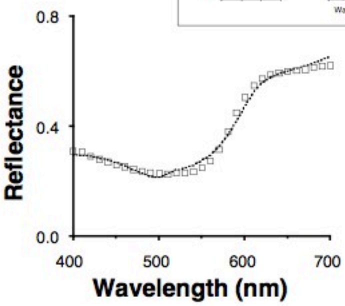

$$S(\lambda)=w_1S_1(\lambda)+w_2S_2(\lambda)+w_3S_3(\lambda)$$

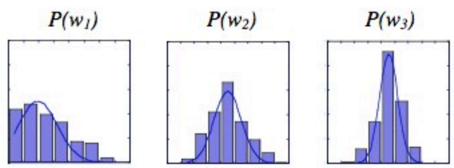

[Brainard & Freeman]

Illuminant Prior

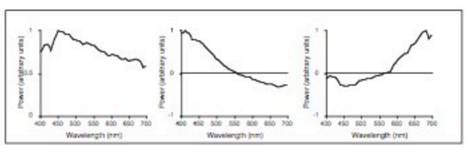

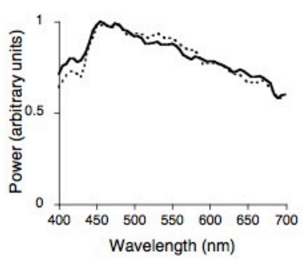

$$E(\lambda)=w[E_1^{'}(\lambda)+uE_2^{'}(\lambda)+v\,E_3^{'}(\lambda)]$$

Daylight measurements

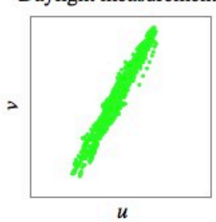

[Brainard & Freeman]

Algorithm Performance

scene 1 scene 2 scene 3

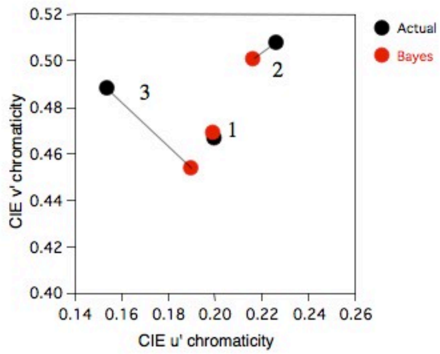

[Brainard & Freeman]

Bayes Predictions: Broad Prior

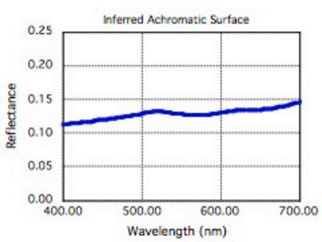

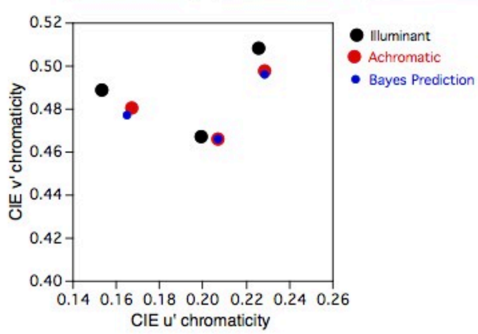

[Brainard & Freeman]