Color

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

Oct. 1, 2002
MIT 6.801/6.866
Profs. Freeman and Darrell

Shape-from-shading

- Reading:
  - Chapter 11 (esp. 11.1, 11.7), Horn

linear filters (Thursday & next Tuesday)

- Reading:
  - Chapter 7, F&P
- Recommended Reading:
  - Chapter 7, 8 Horn

Oct. 1, 2002
MIT 6.801/6.866
Profs. Freeman and Darrell

Color matching experiment

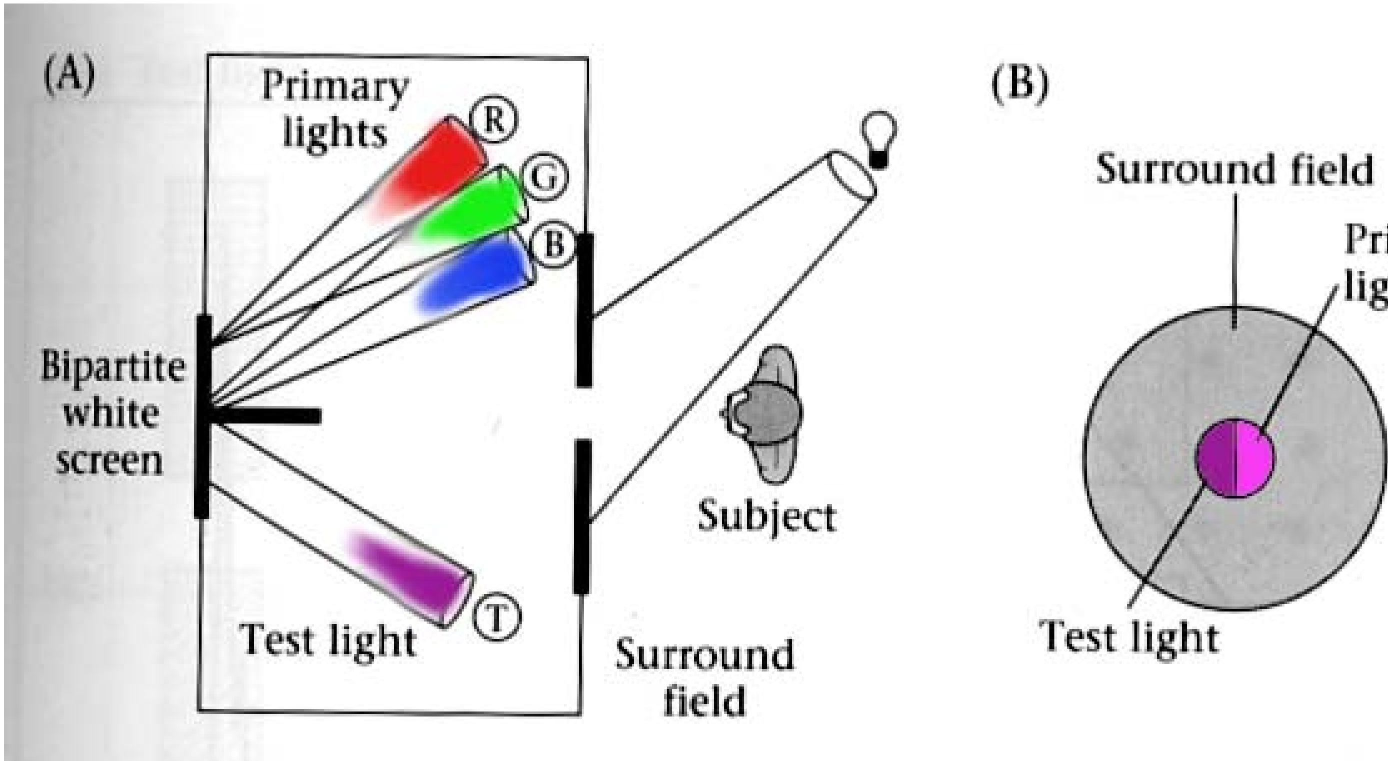

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

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

Color matching functions for a particular set of monochromatic primaries

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

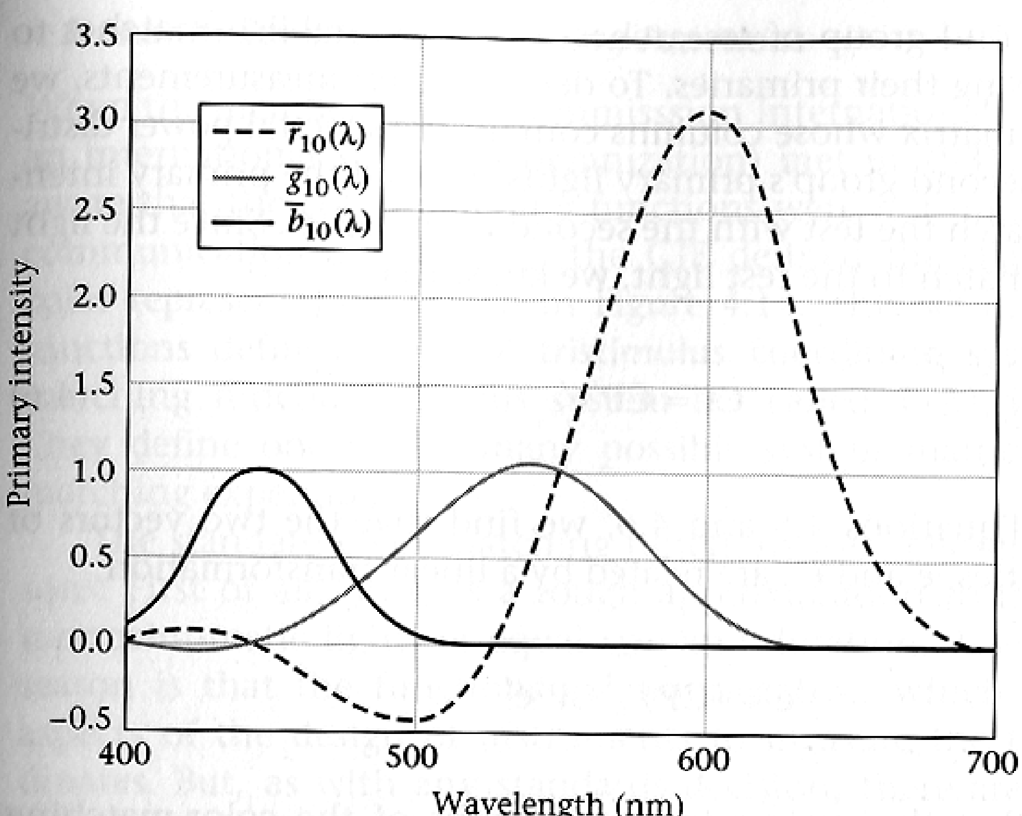

4.13 THE COLOR-MATCHING FUNCTIONS ARE THE ROWS OF THE COLOR-MATCHING SYSTEM MATRIX. The functions measured by Stiles and Burch (1959) using a 10-degree bipartite field and primary lights at the wavelengths 645.2 nm, 525.3 nm, and 444.4 nm with unit radiant power are shown. The three functions in this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

test light wavelengths

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

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

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

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

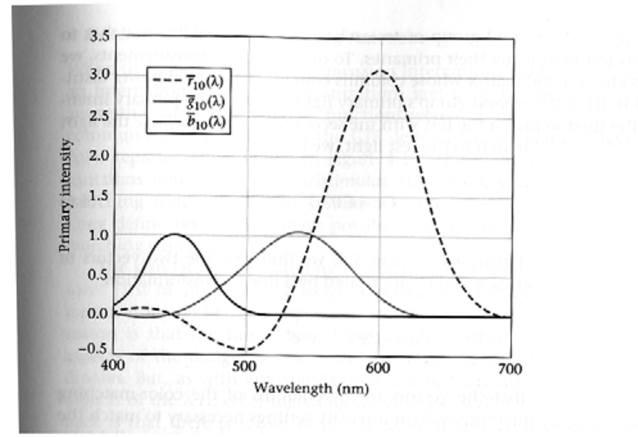

Let the new spectral signal to be characterized be the vector t.

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

Then the amounts of each primary needed to match t are:

$$C\vec{t}$$

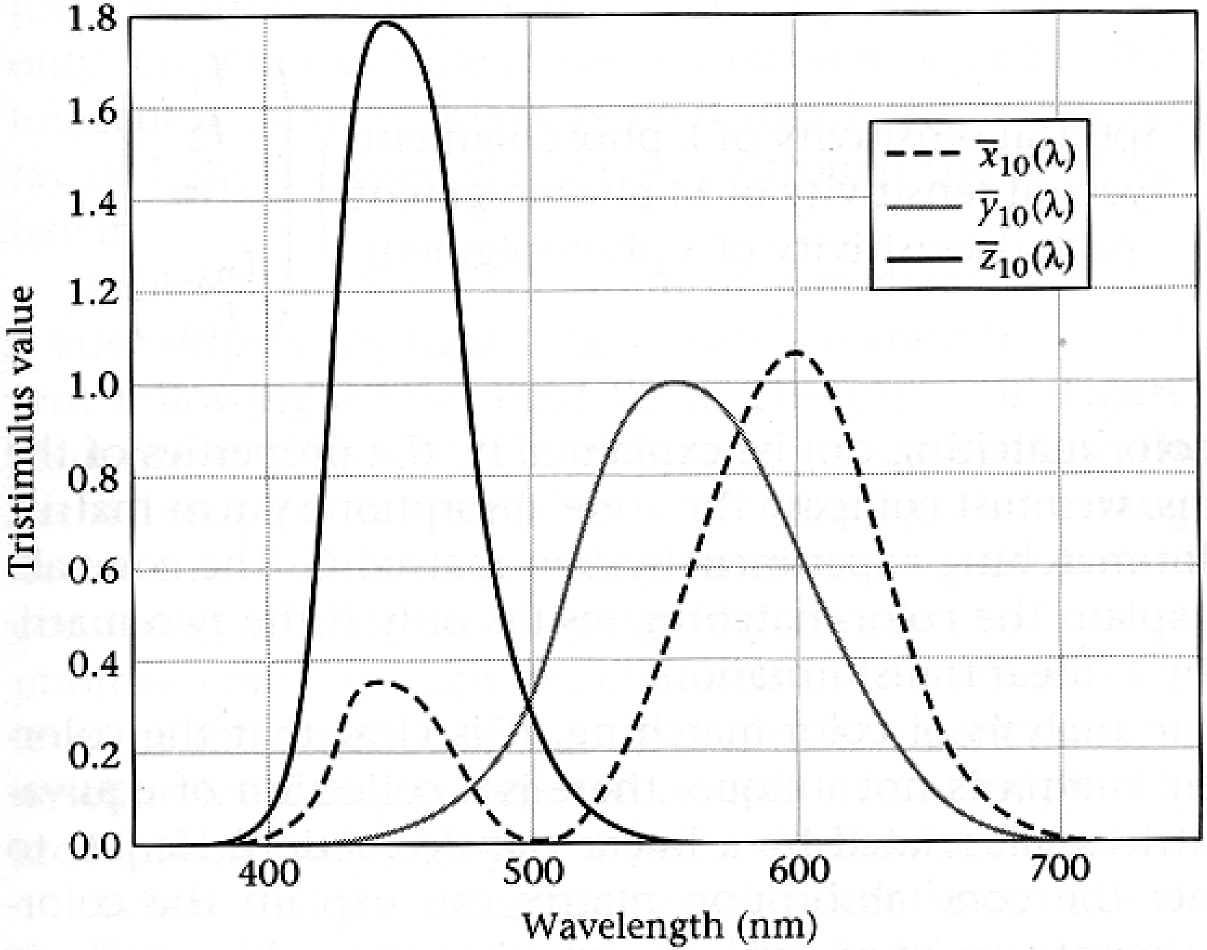

4.14 THE XYZ STANDARD COLOR-MATCHING FUNCTIONS. In 1931 the CIE standardized a set of color-matching functions for image interchange. These color-matching functions are called $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$. Industrial applications commonly describe the color properties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

CIE XYZ: Color matching functions are positive everywhere, but primaries are imaginary. Usually draw x, y, where $x=X/(X+Y+Z)$

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

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

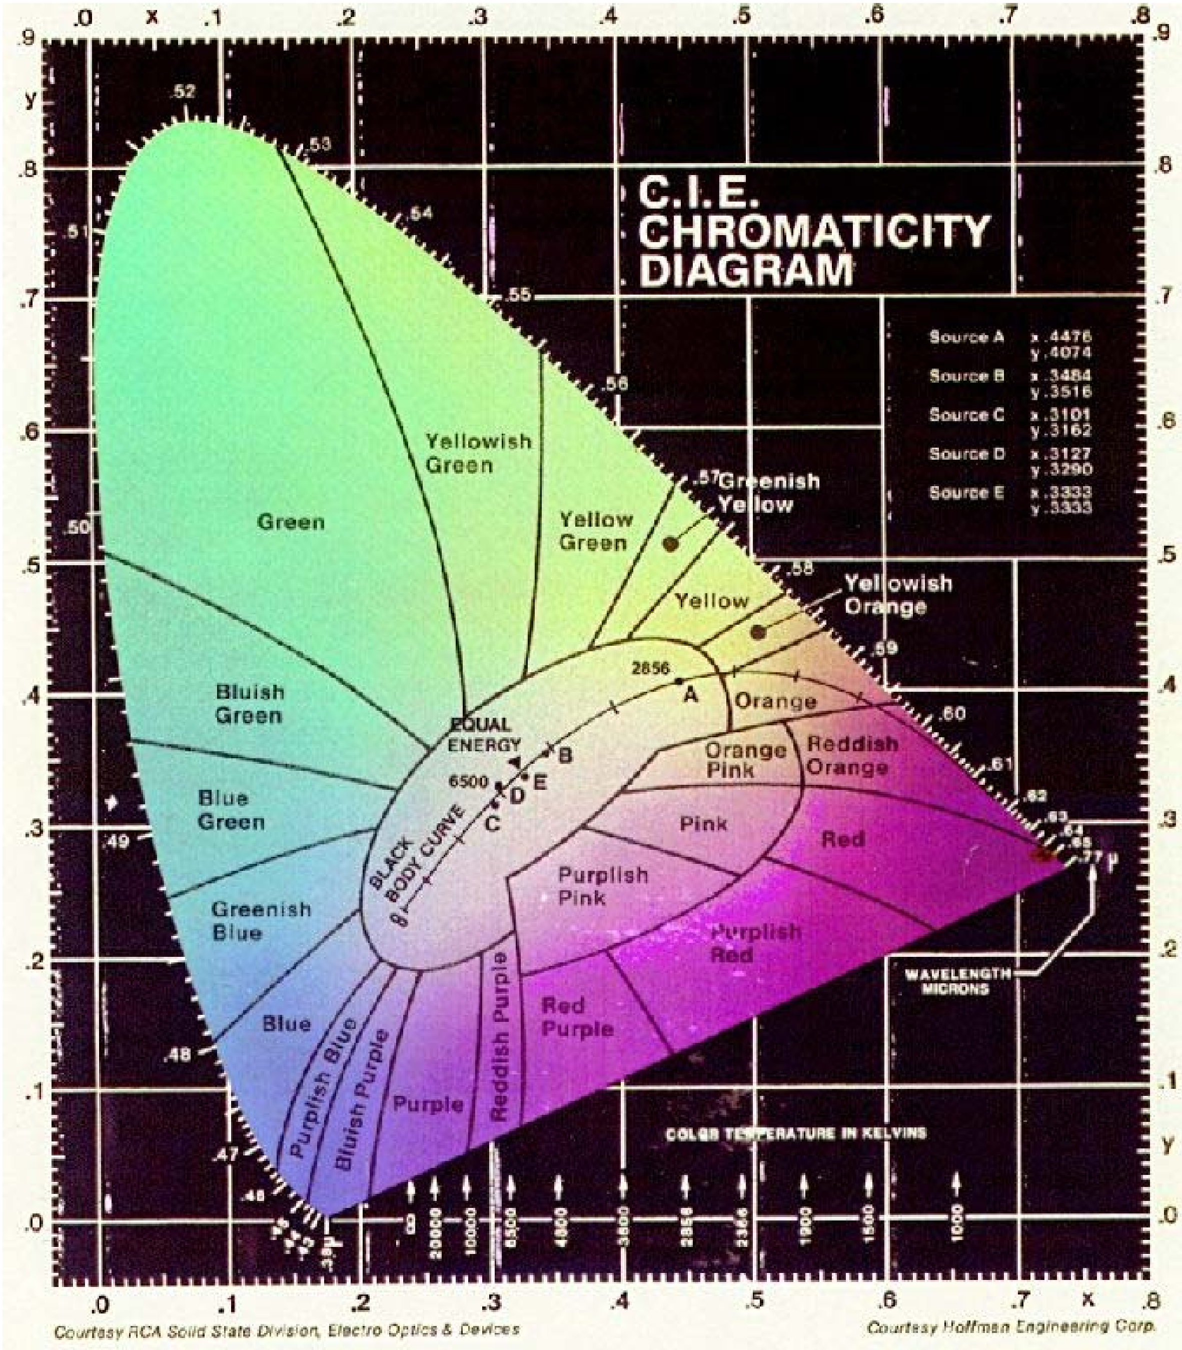
C.I.E. CHROMATICITY DIAGRAM
Source A x .4476 y .4074
Source B x .3484 y .3516
Source C x .3101 y .3162
Source D x .3127 y .3290
Source E x .3333 y .3333
Green
Yellowish Green
Yellow Green
Greenish Yellow
Yellow
Yellowish Orange
Orange
Reddish Orange
Orange Pink
Pink
Red
Purplish Pink
Purplish Red
Red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

Color metamerism

Two spectra, t and s, perceptually match when

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

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

Graphically,

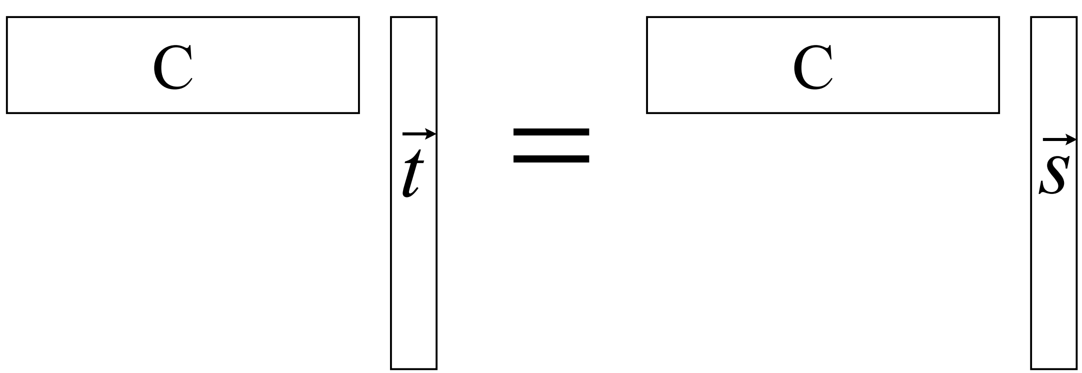

Metameric lights

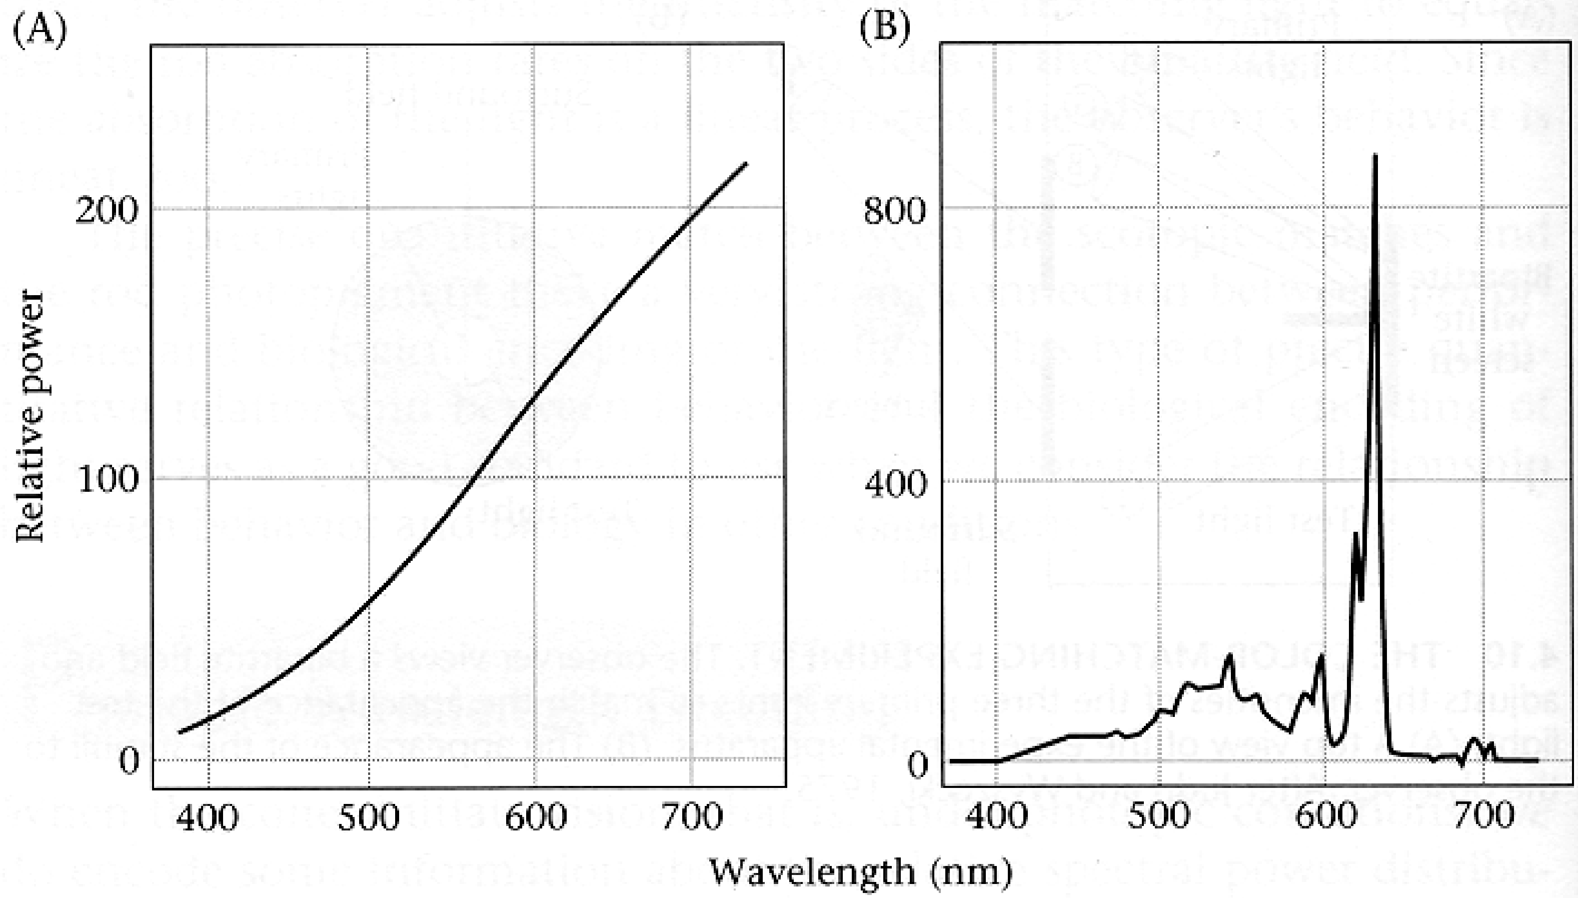

4.11 METAMERIC LIGHTS. Two lights with these spectral power distributions appear identical to most observers and are called metamers. (A) An approximation to the spectral power distribution of a tungsten bulb. (B) The spectral power distribution of light emitted from a conventional television monitor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

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

Metamerism applications

The word “VOID” wasn’t visible on the check, when viewed by eye, only after xerox copying did it appear.

Color appearance

- Or, color matching outside of the controlled experimental setup.

Spatial effects on color appearance

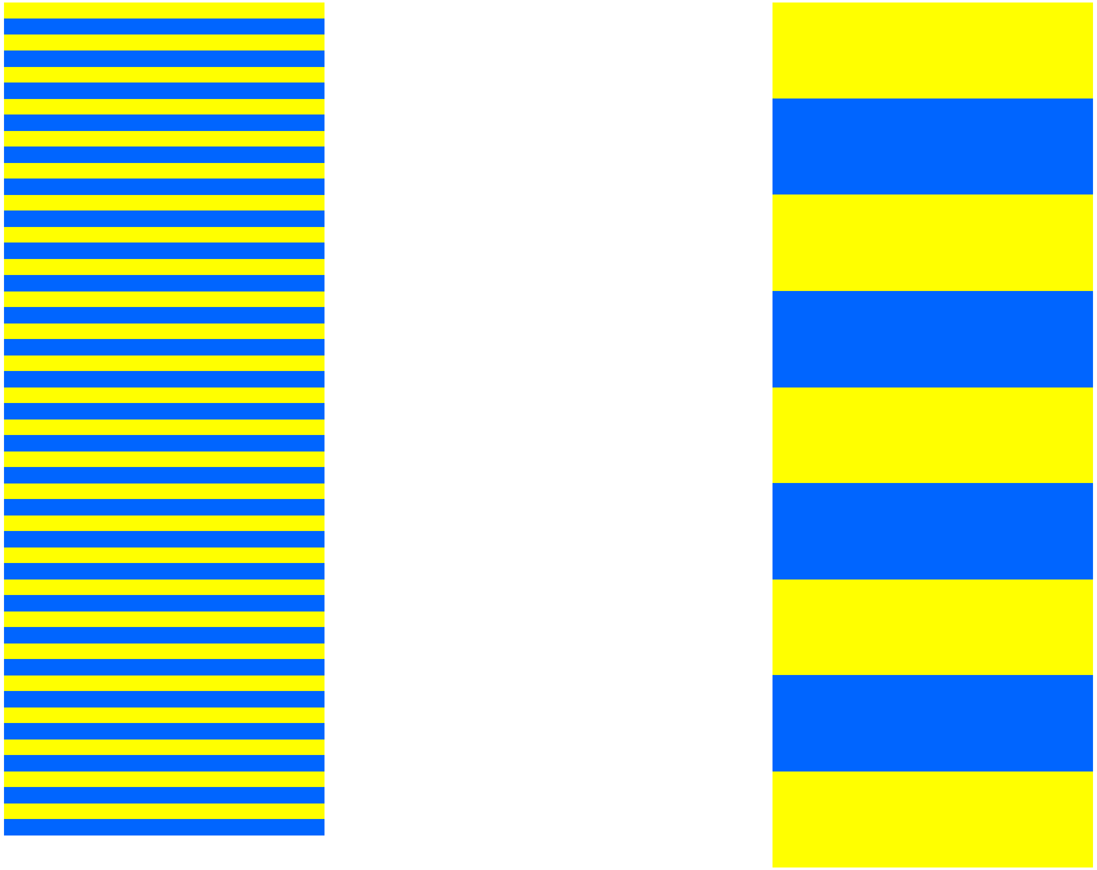

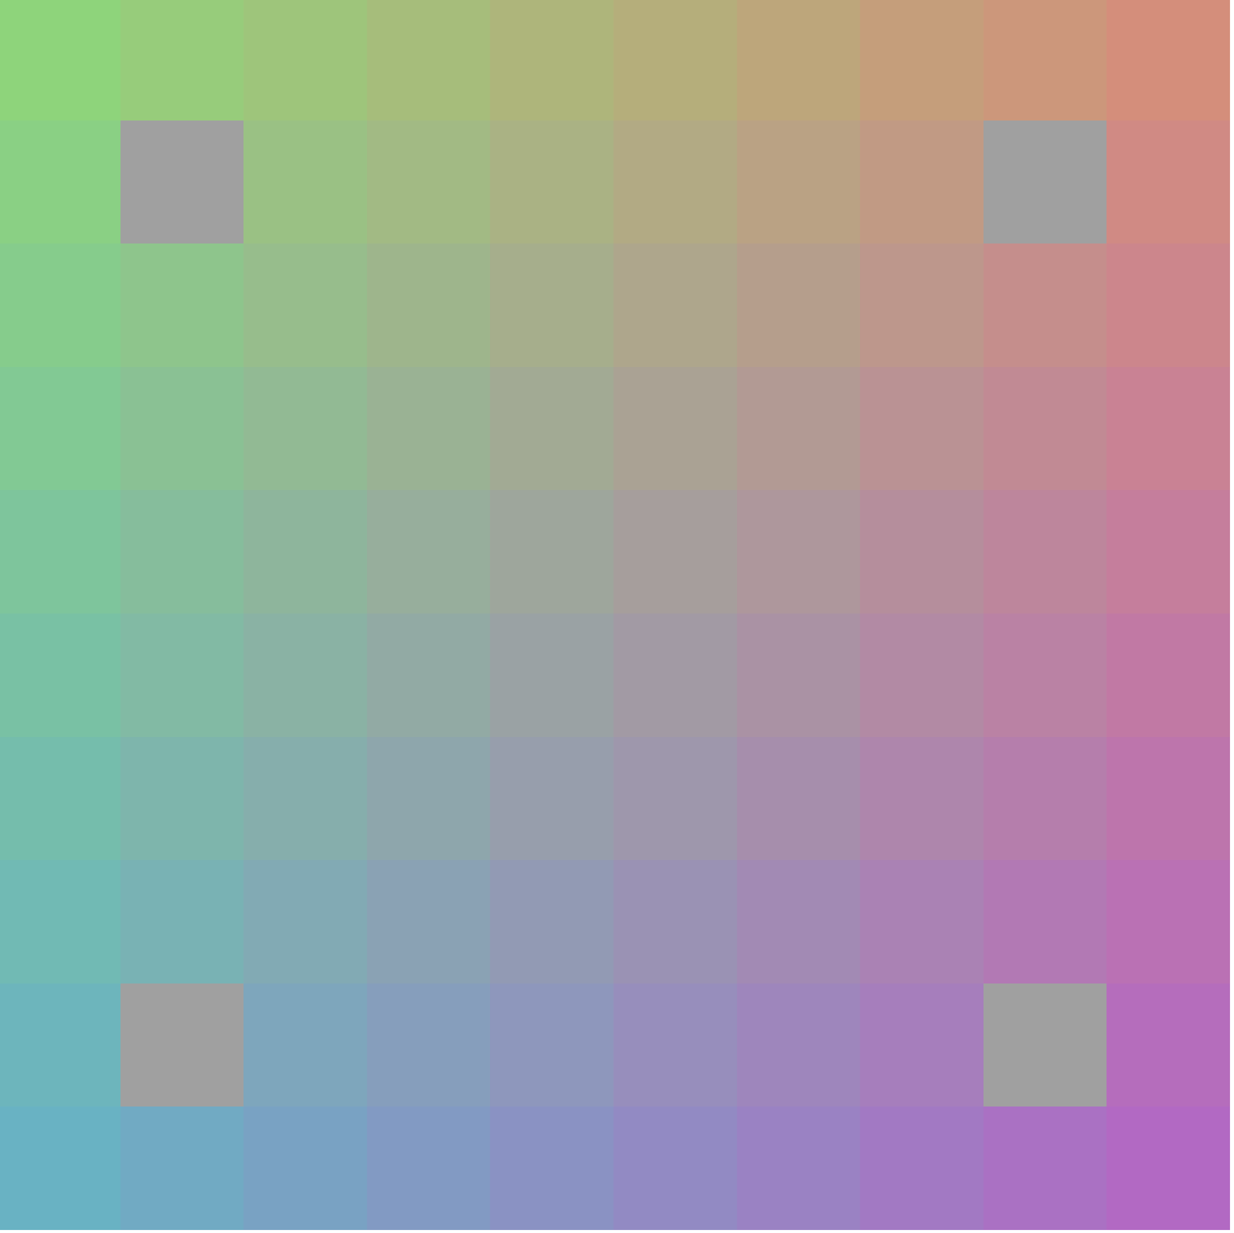

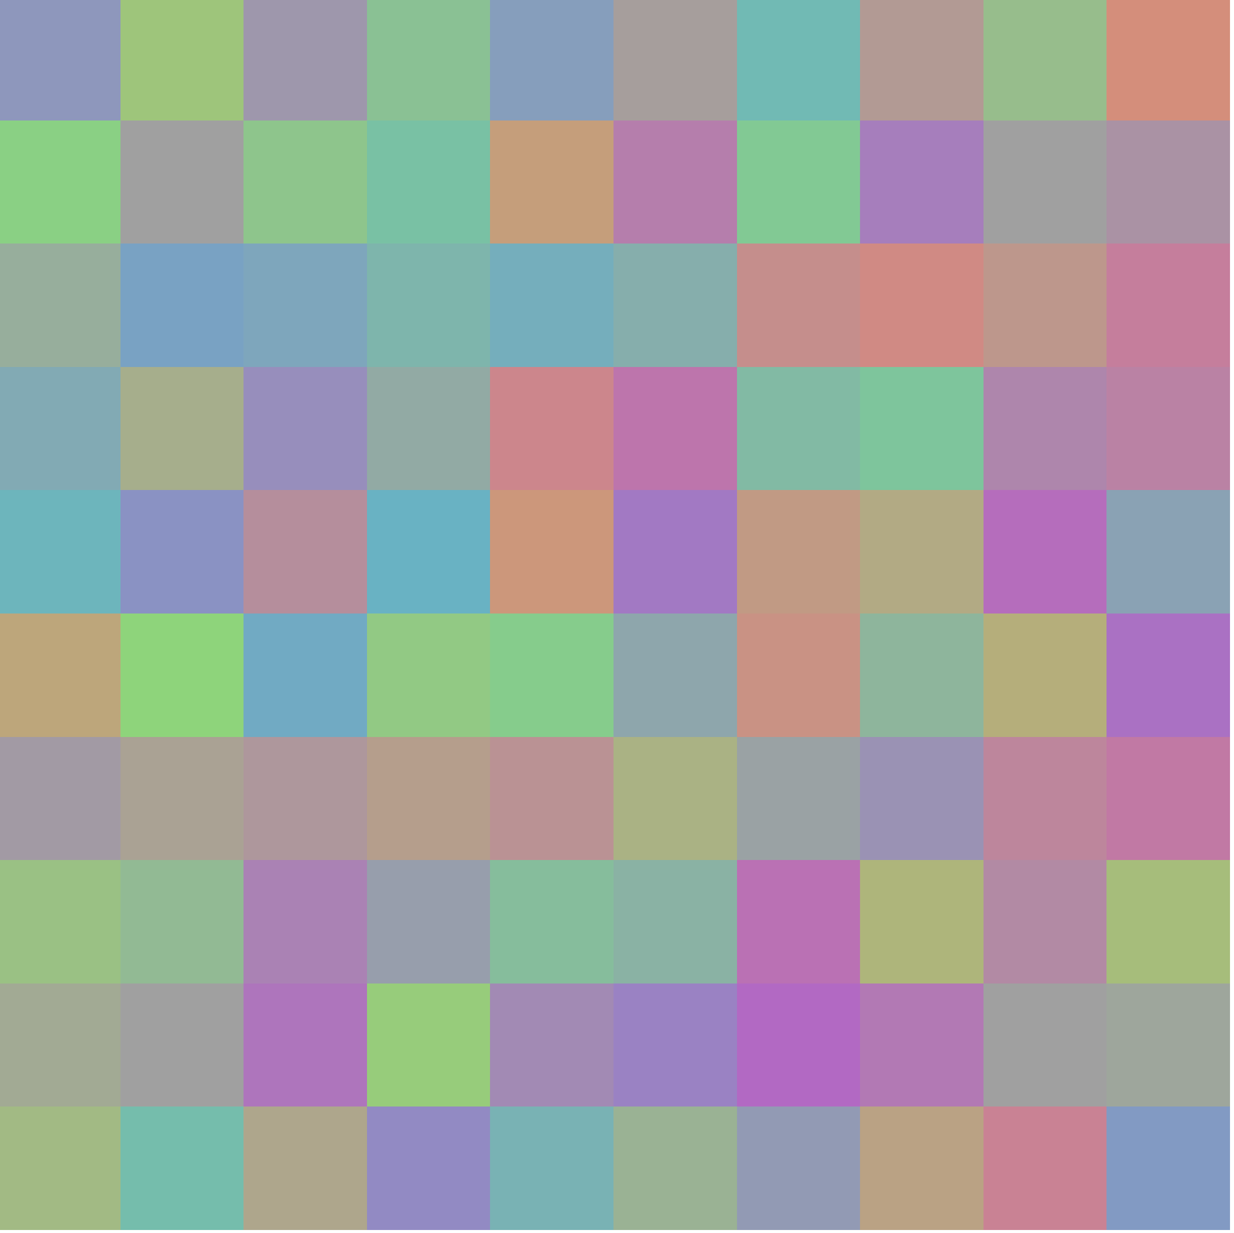

demos

lightness

color

Outline

- Color physics.
- Color perception and color matching.
- Inference about the world from color observations.

What does a color reaching your eye tell you about the world?

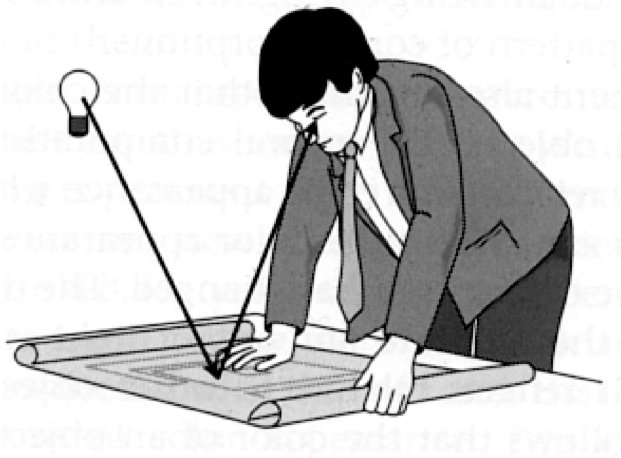

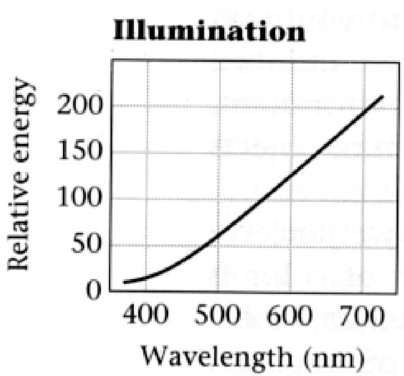

.*

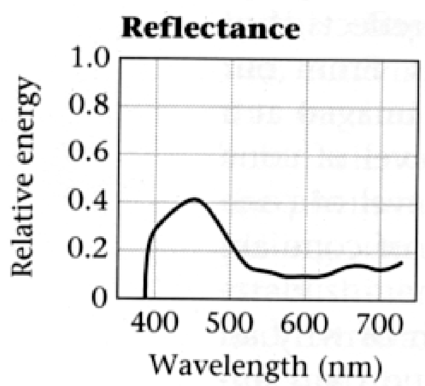

=

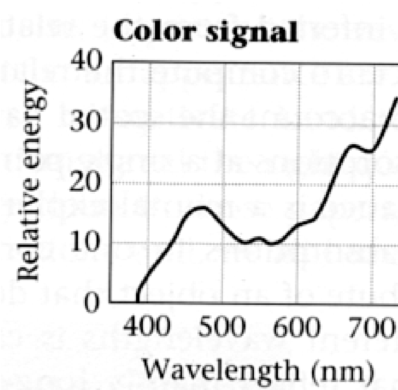

Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

General strategies

- (a) Determine what image would look like under white light, or (b) surface reflectances
- Assume
  - that we are dealing with flat frontal surfaces
  - We've identified and removed specularities
  - no variation in illumination
- We need some form of reference
  - brightest patch is white
  - spatial average is known
  - gamut is known
  - specularities
  - prior probabilities for lights and surfaces

Low-dimensional models for color spectra

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

Basis functions for Macbeth color checker

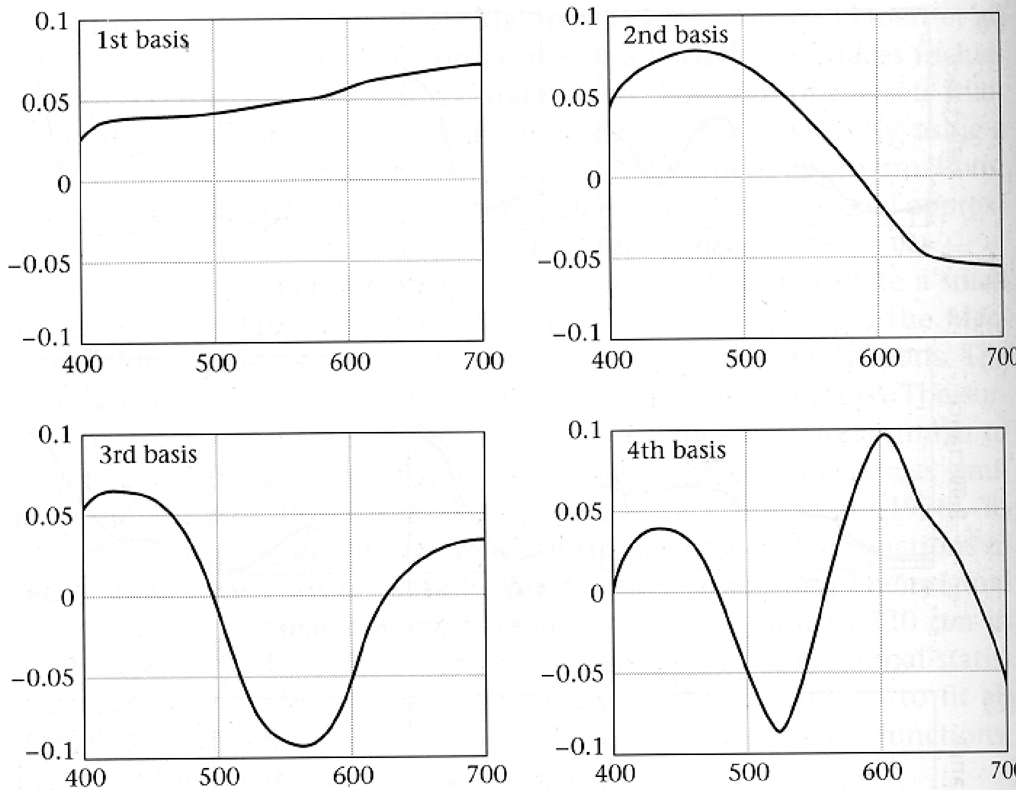

9.9 BASIS FUNCTIONS OF THE LINEAR MODEL FOR THE MACBETH COLORCHECKER. The surface-reflectance functions in the collection vary smoothly with wavelength, as do the basis functions. The first basis function is all positive and explains the most variance in the surface-reflectance functions. The basis functions are ordered in terms of their relative significance for reducing the error in the linear-model approximation to the surfaces.

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

Low-dimensional models for color spectra

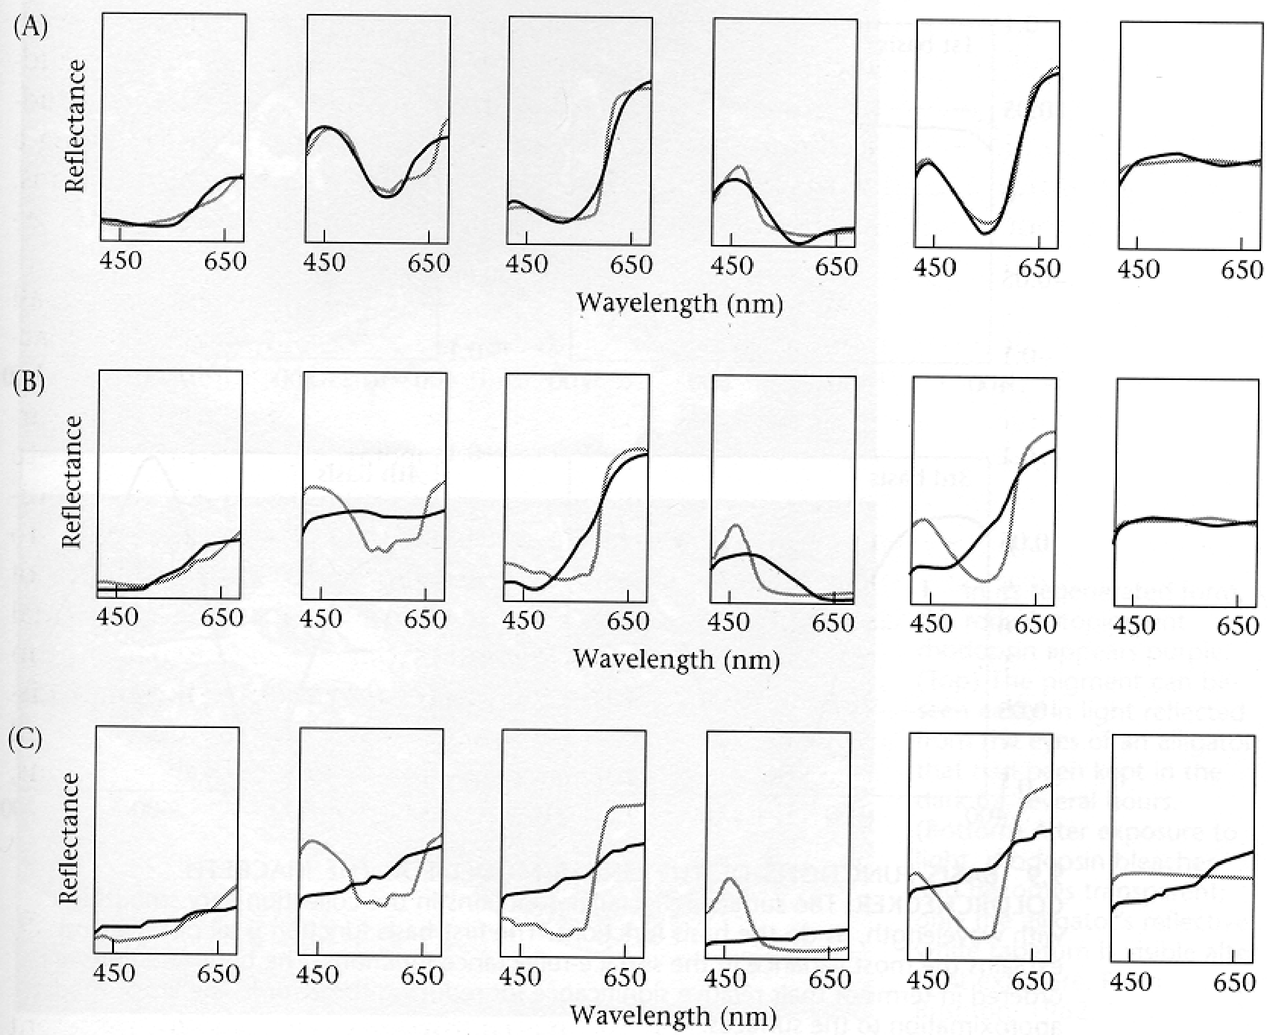

9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using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

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

The rendering model

(using linear combinations of spectral basis functions to represent both surfaces and lights)

$$\boxed{p_k^s} = \boxed{c_k(\lambda)} \left(\boxed{\phi_j(\lambda)} \; \boxed{r_j^s} \right) .* \left(\boxed{\psi_i(\lambda)} \; \boxed{e_i} \right)$$

Response of k[th] sensor to the j[th] surface patch

Spectral response of sensors

Surface spectral basis functions

Surface basis coefficients

Illuminant spectral basis functions

illuminant basis coefficients

$$p_k^s = \sum_{\lambda} c_k(\lambda) \left(\sum_j r_j^s \phi_j(\lambda) \right) \left(\sum_i e_i \psi_i(\lambda) \right)$$

Count unknowns and equations

- Suppose you use a 3-dimensional linear model for both the illuminant and surfaces (two different linear models of the same dimensionality). Suppose have N surfaces.
- 3 + 3N unknowns
- 3N measurements

A possible assumption: know the mean surface value

Sensor responses, as function of surfaces and lights, from before:

$$p_k^s = \sum_{\lambda} c_k(\lambda) \left(\sum_j r_j^s \phi_j(\lambda) \right) \left(\sum_i e_i \psi_i(\lambda) \right)$$

Average over all the N surfaces:

$$\bar{p}_k^s = \sum_{\lambda} c_k(\lambda) \left(\sum_j \bar{r}_j^s \phi_j(\lambda) \right) \left(\sum_i e_i \psi_i(\lambda) \right)$$

Sum over λ and j to write

$$\bar{\mathrm{p}} = \mathrm{A}\mathrm{e}$$

In general, A will be invertible and the estimated illuminant, e, under the "gray world" assumption, is:

$$\mathrm{A}^{-1}\bar{\mathrm{p}} = \mathrm{e}$$

Shape from shading

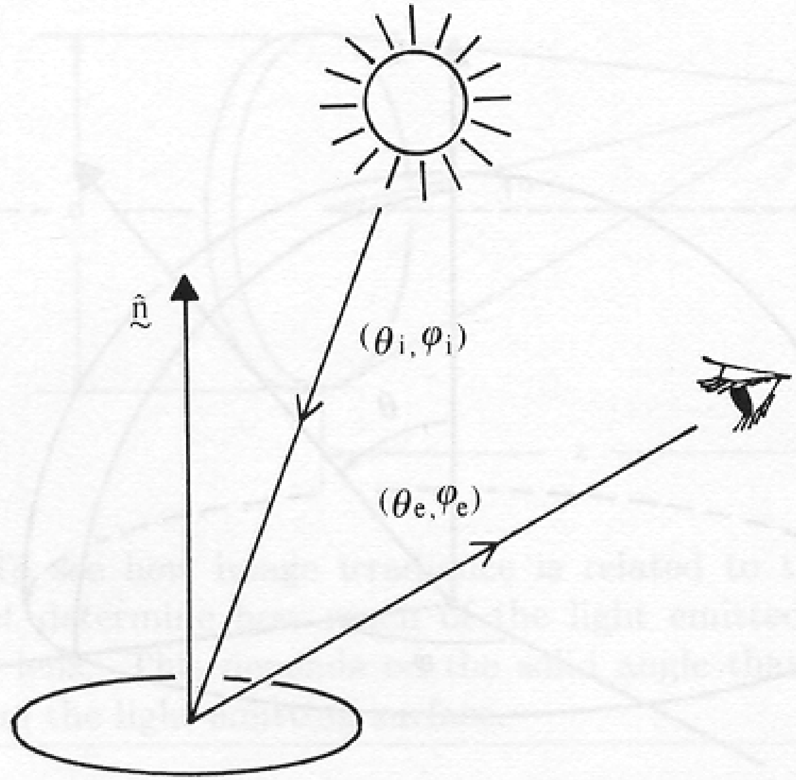

Horn, 1986

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{L(\theta_e, \phi_e)}{E(\theta_i, \phi_i)}$$

Generic reflectance map

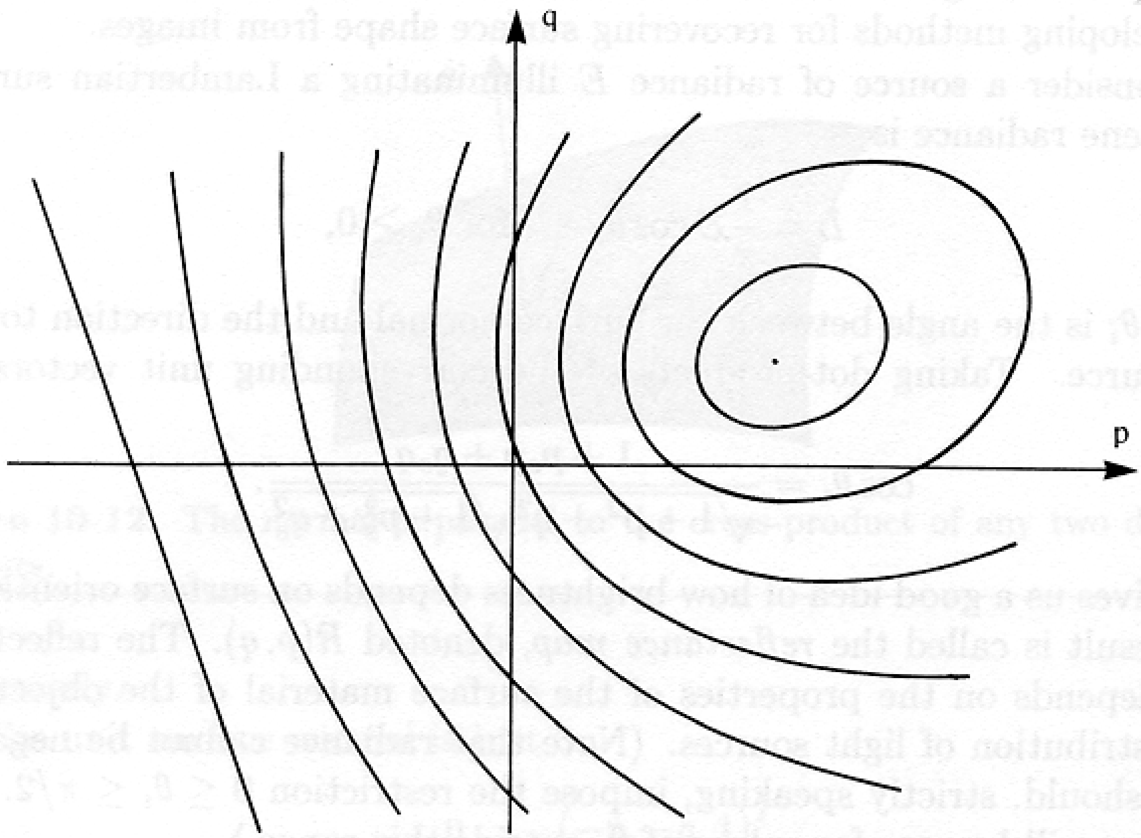

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p,q)$ occurs at the point $(p,q) = (p_s, q_s)$, found inside the nested conic sections, while $R(p,q) = 0$ all along the line on the left side of the contour map.

Horn, 1986

Remember, this model doesn't handle properly:

- Occluding edges
- Albedo changes
- Perspective effects (small)
- Interreflections
- Material changes across surfaces in the image

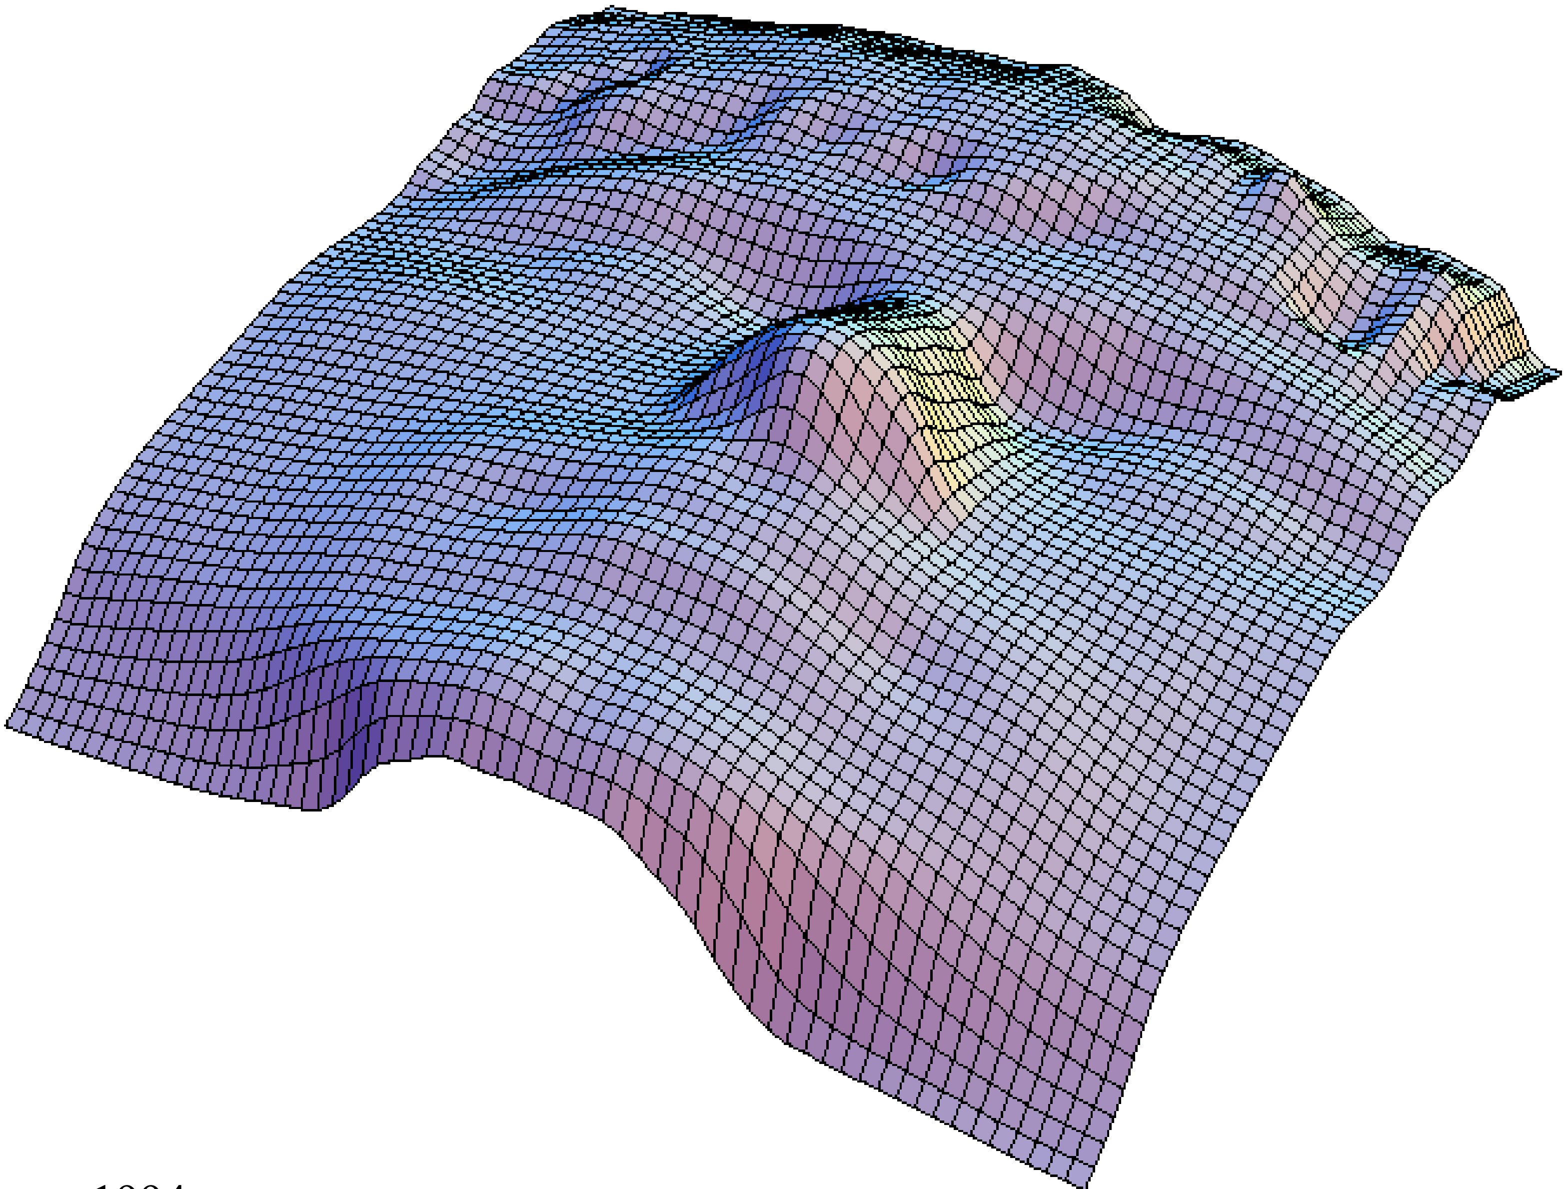

Freeman, 1994

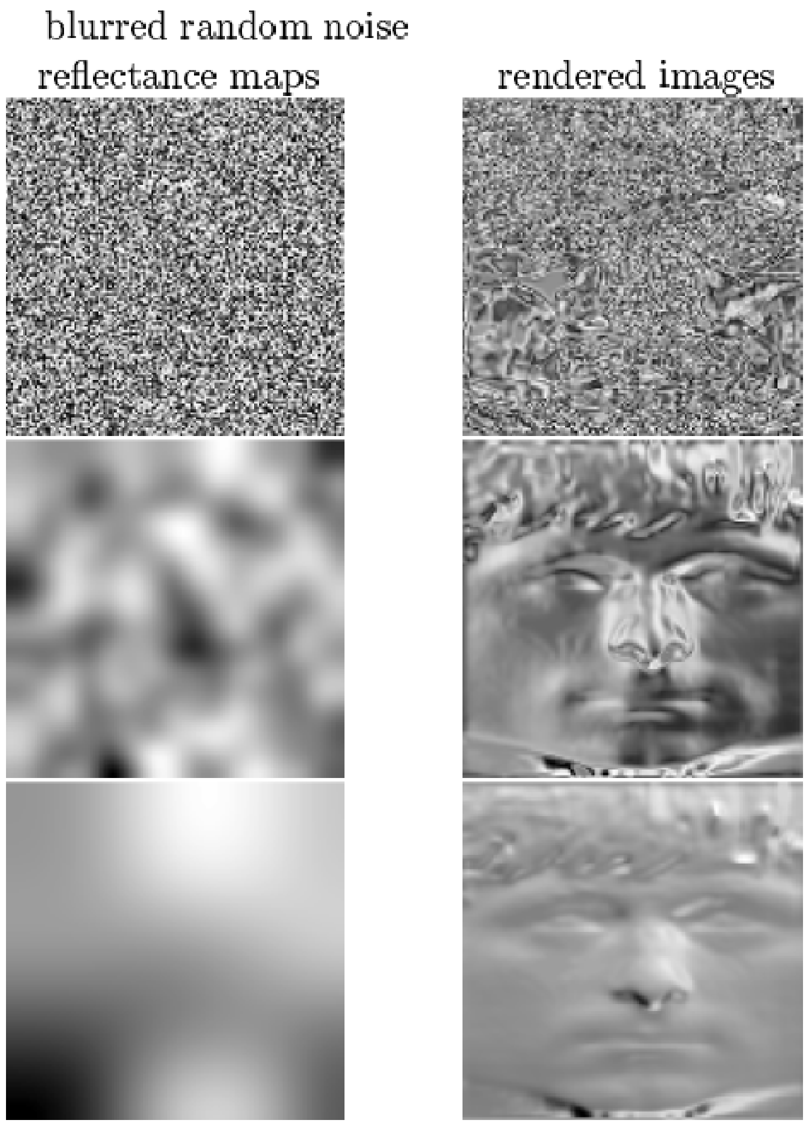

Freeman, 1994

Figure 6: Showing that blurred reflectance maps lead to shapes which are easier to interpret.

Linear shading map

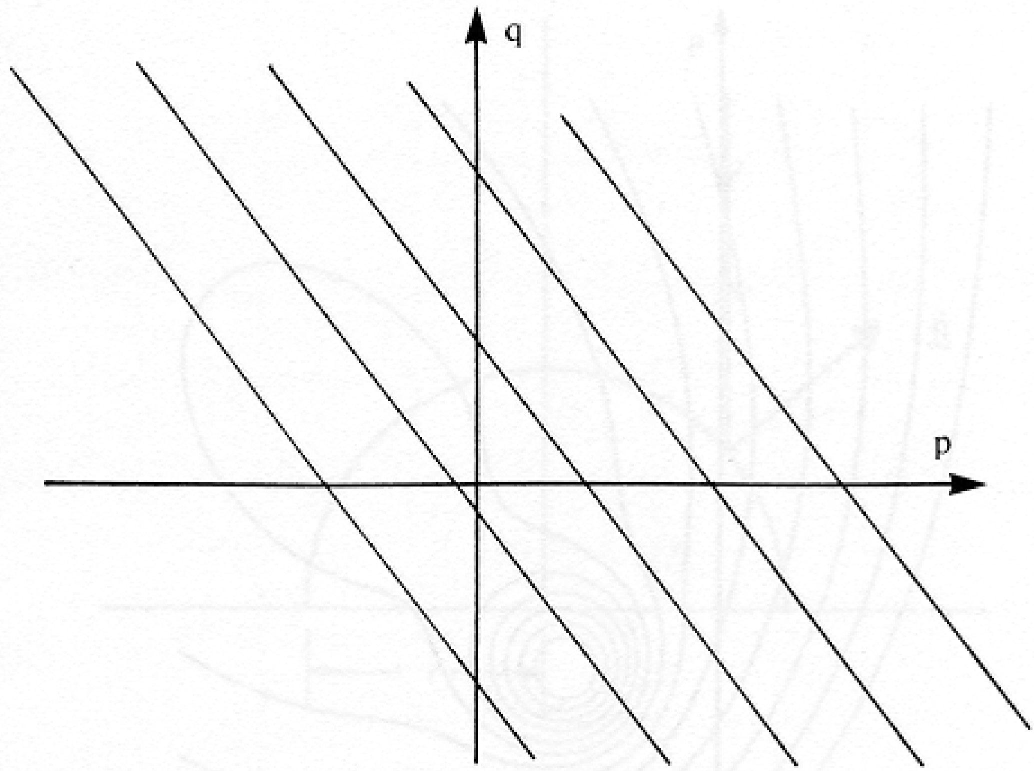

Figure 10-14. In the case of the material in the maria of the moon, the reflectance map can be closely approximated by a function of a linear combination of the components of the gradient. The contours of constant brightness are parallel straight lines in gradient space.

Horn, 1986

Linear shading: 1st order terms of Lambertian shading

Lambertian point source

$$R(p,q) = k\frac{1 + p_s p + q_s q}{\sqrt{1 + p^2 + q^2}\sqrt{1 + p_s^2 + q_s^2}}$$

1st order Taylor series about p=q=0

$$\approx k_2 + \left.\frac{\partial R(p,q)}{\partial p}\right|_{p=0,q=0} p + \left.\frac{\partial R(p,q)}{\partial q}\right|_{p=0,q=0} q$$

$$= k_2(1 + p_s p + q_s q)$$

See Pentland, IJCV vol. 1 no. 4, 1990.

Linear shading approximation

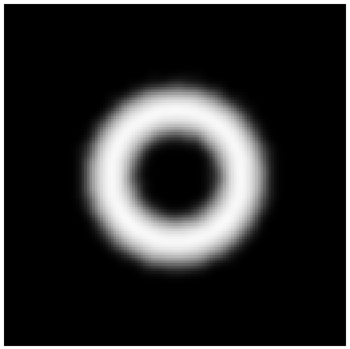

range image

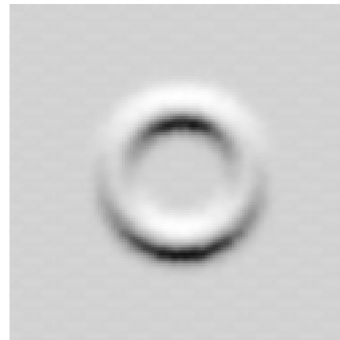

Lambertian shading

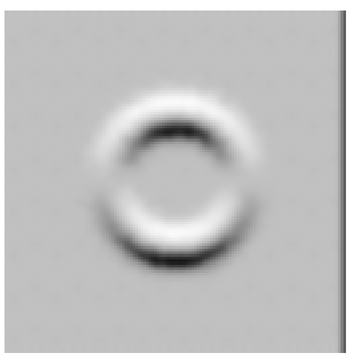

linear shading

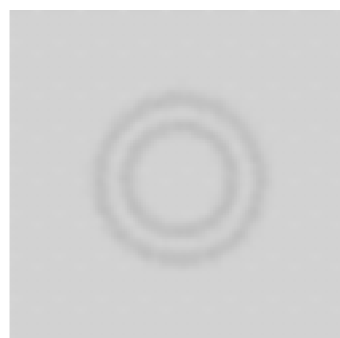

quadratic terms

higher-order terms

Pentland 1990, Adelson&Freeman, 1991

Simplified lighting assumption

- For simplicity, and without loss of generality, let’s consider light source from the left.
- Then rendered image is proportional to the derivative of the range image along a row

Simplified linear shape from shading computation, for $I(n) = S(n+1)-S(N)$

- Image, one row:

  $I(n) = [17\ \ 14\ \ 12\ \ 4\ \ 4\ \ 9\ \ 25\ \ldots]$

- Shape, same row:

  $S(n) = [0\ \ \ 17\ \ \ 31\ \ 43\ \ 47\ \ 51\ \ 60\ \ 85\ldots]$

or $S(n) = [-17\ \ \ 0\ \ \ 14\ \ 26\ \ 30\ \ 34\ \ 41\ \ 66\ldots]$

or etc.

Illustrates a general problem

- Each row can have a different constant of integration.
- Typical of shape-from-shading problems: extra knowledge has to be put in, or extra assumptions have to be made:
  – Boundary conditions
  – Regularization terms

Shape-from-shading as prototypical vision computation

Want to estimate scene parameters (surface slopes p and q at every image position, i, j).

Have a rendering function that takes you from some given set of scene parameters to observation data (reflectance map R(p,q) gives image intensity for any p, q).

Want to find the parameters ($p_{i,j}$, $q_{i,j}$) that minimize the difference from the observations ($E_{i,j}$).

But the problem is "ill-posed", or underspecified from that constraint alone. So add-in additional requirements that the scene parameters must satisfy (the surface slopes p and q must be smooth at every point).

Shape-from-shading as prototypical vision computation

smoothness at i,j [Bayesian: prior]

$$s_{i,j} = \frac{1}{4}\left(\left(p_{i+1,j} - p_{i,j}\right)^2 + \left(p_{i,j+1} - p_{i,j}\right)^2 + \left(q_{i+1,j} - q_{i,j}\right)^2 + \left(q_{i,j+1} - q_{i,j}\right)^2\right)$$

error in image irradiance at i,j [Bayesian: likelihood]

$$r_{ij} = \left(E_{ij} - R_s\left(p_{ij}, q_{ij}\right)\right)^2$$

Shape-from-shading as prototypical vision computation

$$s_{i,j} = \frac{1}{4}\left(\left(p_{i+1,j} - p_{i,j}\right)^2 + \left(p_{i,j+1} - p_{i,j}\right)^2 + \left(q_{i+1,j} - q_{i,j}\right)^2 + \left(q_{i,j+1} - q_{i,j}\right)^2\right)$$

$$r_{ij} = \left(E_{ij} - R_s\left(p_{ij}, q_{ij}\right)\right)^2$$

Minimize regularized function [Bayesian: maximize posterior]

$$e = \sum_i \sum_j \left(s_{ij} + \lambda r_{ij}\right)$$

This approach is used in many different vision problems. See poggio nature paper, horn work, szeliski paper, etc. Relation to Bayesian methods will be discussed later.

Look for stationary points

$$s_{i,j} = \frac{1}{4}\left(\left(p_{i+1,j} - p_{i,j}\right)^2 + \left(p_{i,j+1} - p_{i,j}\right)^2 + \left(q_{i+1,j} - q_{i,j}\right)^2 + \left(q_{i,j+1} - q_{i,j}\right)^2\right)$$

$$r_{ij} = \left(E_{ij} - R_s\left(p_{ij}, q_{ij}\right)\right)^2$$

$$e = \sum_i \sum_j \left(s_{ij} + \lambda r_{ij}\right)$$

Differentiate error term w.r.t. shape parameters

$$\frac{\partial e}{\partial p_{kl}} = 2\left(p_{kl} - \overline{p}_{kl}\right) - 2\lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial p}$$

$$\frac{\partial e}{\partial q_{kl}} = 2\left(q_{kl} - \overline{q}_{kl}\right) - 2\lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial q}$$

where $$\overline{p}_{kl} = \frac{1}{4}\left(p_{i+1,j} + p_{i,j+1} + p_{i-1,j} + p_{i,j-1}\right)$$

$$\overline{q}_{kl} = \frac{1}{4}\left(q_{i+1,j} + q_{i,j+1} + q_{i-1,j} + q_{i,j-1}\right)$$

Set derivatives = 0

$$s_{i,j} = \frac{1}{4}\left(\left(p_{i+1,j} - p_{i,j}\right)^2 + \left(p_{i,j+1} - p_{i,j}\right)^2 + \left(q_{i+1,j} - q_{i,j}\right)^2 + \left(q_{i,j+1} - q_{i,j}\right)^2\right)$$

$$r_{ij} = \left(E_{ij} - R_s\left(p_{ij}, q_{ij}\right)\right)^2$$

$$e = \sum_i \sum_j \left(s_{ij} + \lambda r_{ij}\right)$$

$$\frac{\partial e}{\partial p_{kl}} = 2\left(p_{kl} - \overline{p}_{kl}\right) - 2\lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial p}$$

$$\frac{\partial e}{\partial q_{kl}} = 2\left(q_{kl} - \overline{q}_{kl}\right) - 2\lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial q}$$

Rearranging terms, find a necessary condition for local max of function to be optimized, e:

$$p_{kl} = \overline{p}_{kl} + \lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial p}$$

$$q_{kl} = \overline{q}_{kl} + \lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial q}$$

Set derivatives = 0

$$s_{i,j} = \frac{1}{4}\left(\left(p_{i+1,j} - p_{i,j}\right)^2 + \left(p_{i,j+1} - p_{i,j}\right)^2 + \left(q_{i+1,j} - q_{i,j}\right)^2 + \left(q_{i,j+1} - q_{i,j}\right)^2\right)$$

$$r_{ij} = \left(E_{ij} - R_s\left(p_{ij}, q_{ij}\right)\right)^2$$

$$e = \sum_i \sum_j \left(s_{ij} + \lambda r_{ij}\right)$$

$$\frac{\partial e}{\partial p_{kl}} = 2\left(p_{kl} - \overline{p}_{kl}\right) - 2\lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial p}$$

$$\frac{\partial e}{\partial q_{kl}} = 2\left(q_{kl} - \overline{q}_{kl}\right) - 2\lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial q}$$

$$p_{kl} = \overline{p}_{kl} + \lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial p}$$

$$q_{kl} = \overline{q}_{kl} + \lambda\left(E_{kl} - R_s\left(p_{kl}, q_{kl}\right)\right)\frac{\partial R_s}{\partial q}$$

This suggests an iterative algorithm (keep insisting that the condition be met, until it is):

$$p_{kl}^{n+1} = \overline{p}_{kl}^{n} + \lambda\left(E_{kl} - R_s\left(p_{kl}^{n}, q_{kl}^{n}\right)\right)\frac{\partial R_s}{\partial p}$$

$$q_{kl}^{n+1} = \overline{q}_{kl}^{n} + \lambda\left(E_{kl} - R_s\left(p_{kl}^{n}, q_{kl}^{n}\right)\right)\frac{\partial R_s}{\partial q}$$

Results using this approach

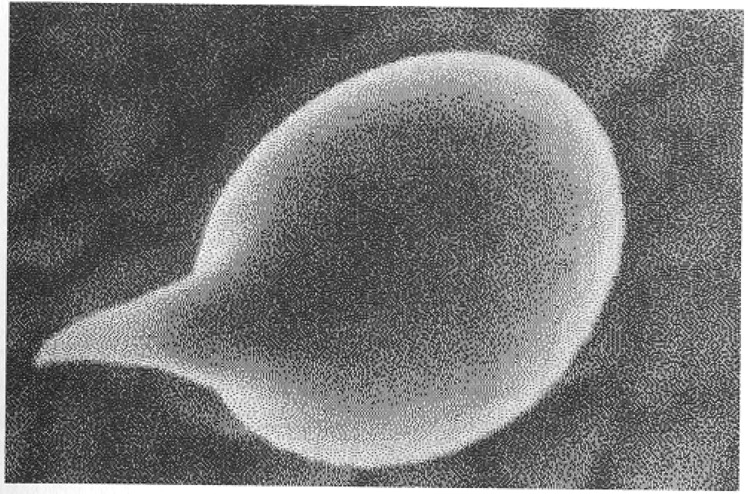

Figure 11-10. Display of the image of a small resin droplet on a flower of a *Cannabis sativa* plant. (Reproduced by permission from the book *Magnifications—Photography with the Scanning Electron Microscope* by David Scharf, Schocken Books, New York, 1977.)

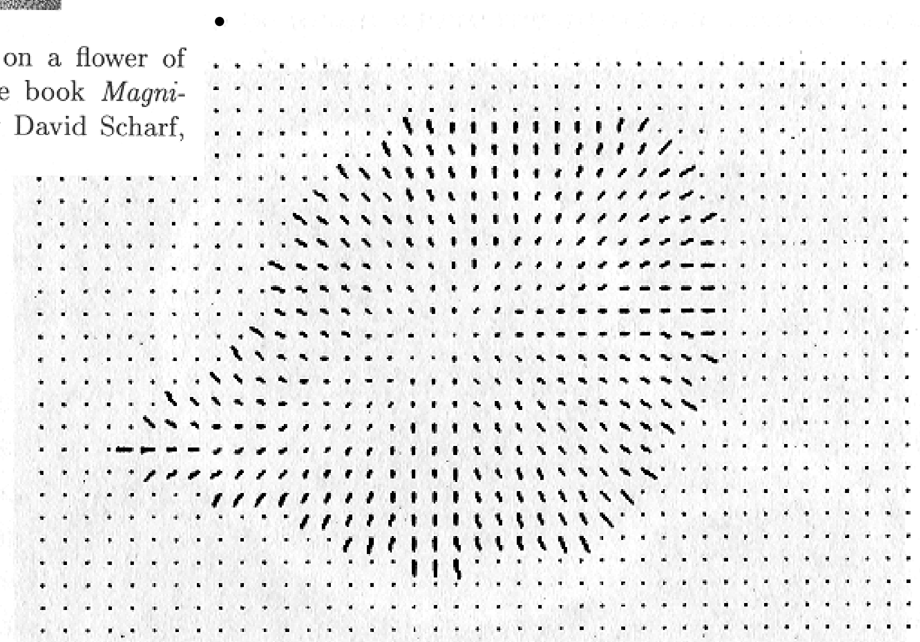

Figure 11-11. Needle diagram calculated by the iterative scheme under the assumption that the reflectance map is $\sec\theta_i$. (The surface orientation data are actually available on a finer grid; they are sampled coarsely here for display purposes.) The needle diagram is the estimate of the shape of the surface of the resin droplet shown in the previous figure. (Figure kindly provided by Katsushi Ikeuchi.)

Horn, 1986