# Color receptors

- **Principle of univariance**: cones give the same kind of response, in different *amounts*, to different wavelengths. The output of the cone is obtained by summing over wavelengths. Responses are measured in a variety of ways (comparing behaviour of color normal and color deficient subjects).
- All experimental evidence suggests that the response of the k'th type of cone can be written as

$$\int \sigma_k(\lambda) E(\lambda) d\lambda$$

where $\sigma_k(\lambda)$ is the sensitivity of the receptor and spectral energy density of the incoming light.

# Color receptors

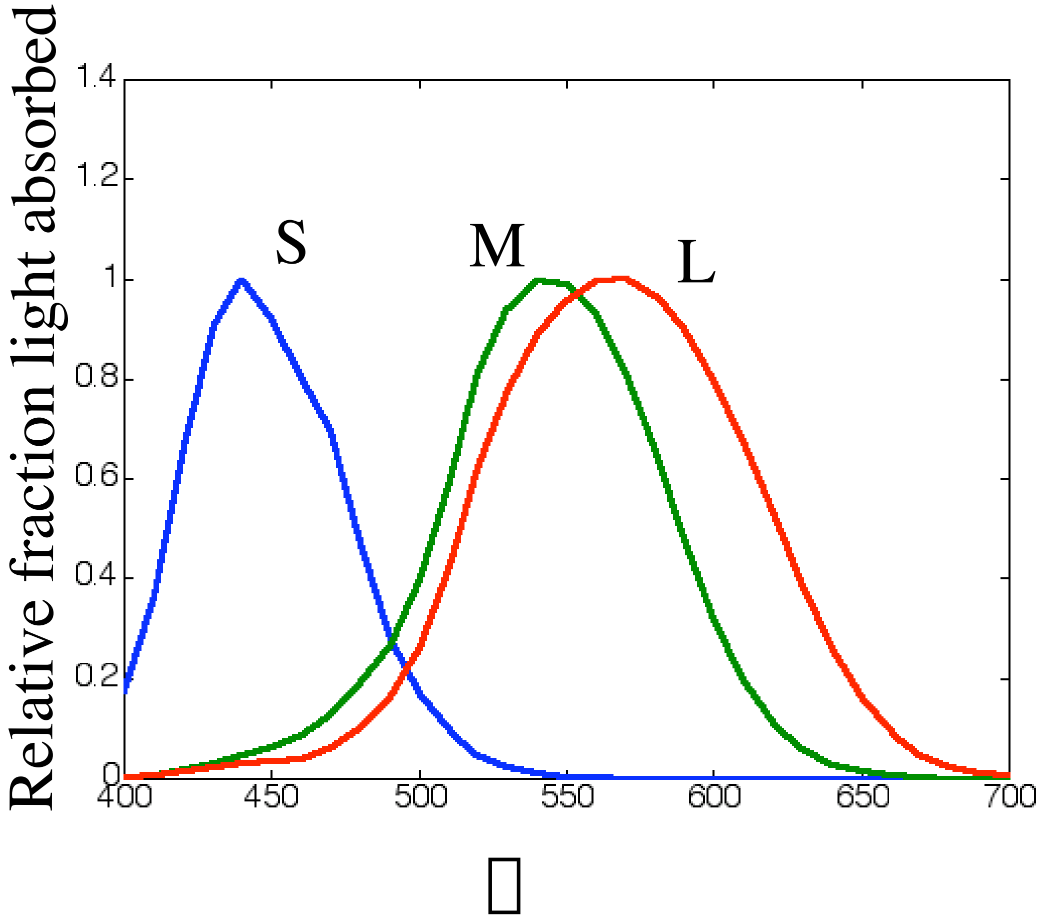


Plot shows relative sensitivity as a function of wavelength, for the three cones. The S (for short) cone responds most strongly at short wavelengths; the M (for medium) at medium wavelengths and the L (for long) at long wavelengths.

These are occasionally called B, G and R cones respectively, but that's misleading - you don't see red because your R cone is activated.

# Simpler if we discretize frequency

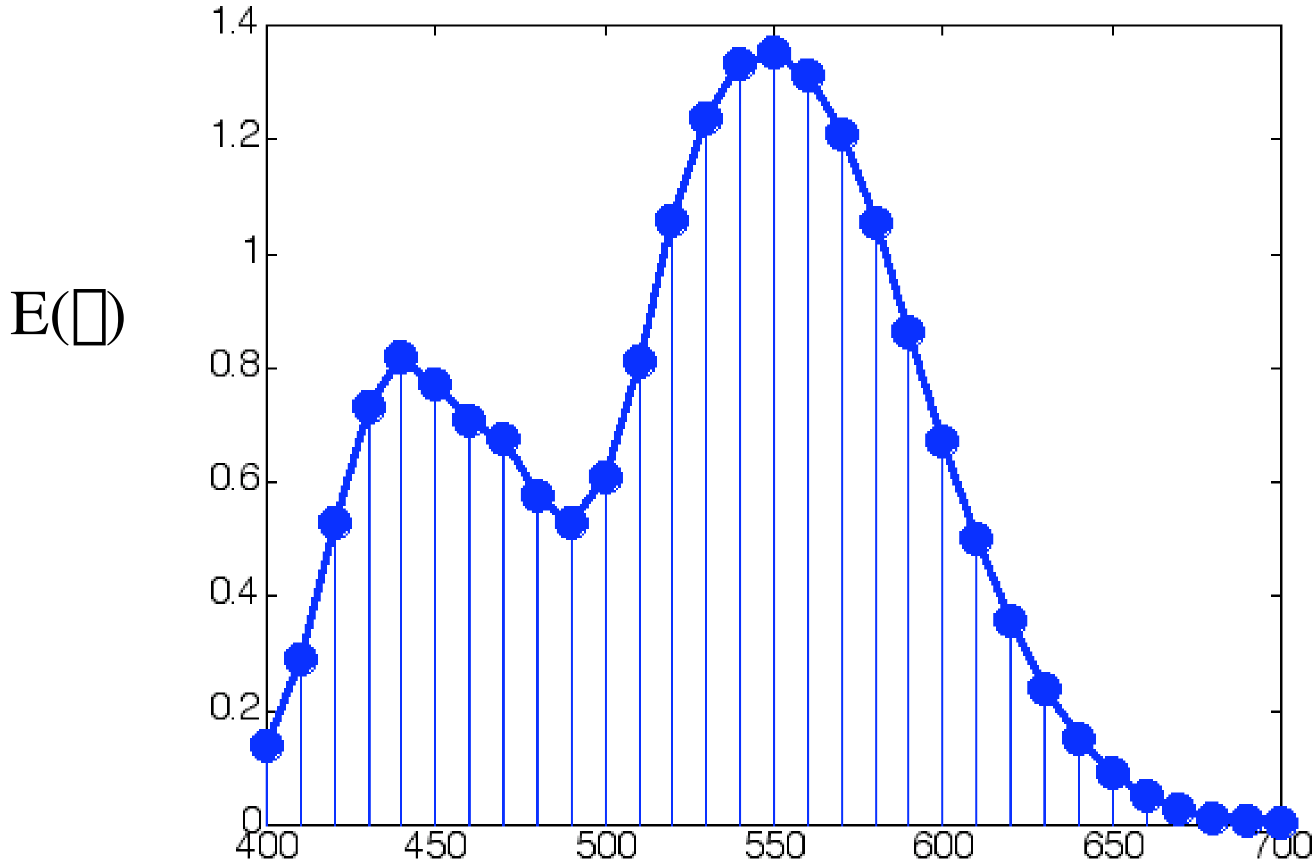

# What does color matching do?

Choose
$$\begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} = \begin{bmatrix} \vec{a}^T \\ \vec{b}^T \\ \vec{c}^T \end{bmatrix} \vec{E}$$

Such that when:
$$\begin{bmatrix} I_r \\ I_g \\ I_b \end{bmatrix} = \begin{bmatrix} \vec{\phi}_r^T \\ \vec{\phi}_g^T \\ \vec{\phi}_b^T \end{bmatrix} \vec{E}$$

Then the combined primaries yields:
$$\begin{bmatrix} I_r \\ I_g \\ I_b \end{bmatrix} = \begin{bmatrix} \vec{\phi}_r^T \\ \vec{\phi}_g^T \\ \vec{\phi}_b^T \end{bmatrix} \begin{bmatrix} \vec{A} & \vec{B} & \vec{C} \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix}$$

# Color matching functions

Choose
$$\begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} = \begin{bmatrix} \vec{a}^T \\ \vec{b}^T \\ \vec{c}^T \end{bmatrix} \vec{E}$$

Then the combined primaries yields:
$$\begin{bmatrix} \vec{\square}_r^T \\ \vec{\square}_g^T \\ \vec{\square}_b^T \end{bmatrix} \vec{E} = \begin{bmatrix} \vec{\square}_r^T \\ \vec{\square}_g^T \\ \vec{\square}_b^T \end{bmatrix} \left[ \vec{A} \quad \vec{B} \quad \vec{C} \right] \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} \Rightarrow \begin{bmatrix} \vec{\square}_r^T \\ \vec{\square}_g^T \\ \vec{\square}_b^T \end{bmatrix} \vec{E} = M \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix}$$

$$M^{-1} \begin{bmatrix} \vec{\square}_r^T \\ \vec{\square}_g^T \\ \vec{\square}_b^T \end{bmatrix} \vec{E} = \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} \Rightarrow M^{-1} \begin{bmatrix} \vec{\square}_r^T \\ \vec{\square}_g^T \\ \vec{\square}_b^T \end{bmatrix} = \begin{bmatrix} \vec{a}^T \\ \vec{b}^T \\ \vec{c}^T \end{bmatrix}$$

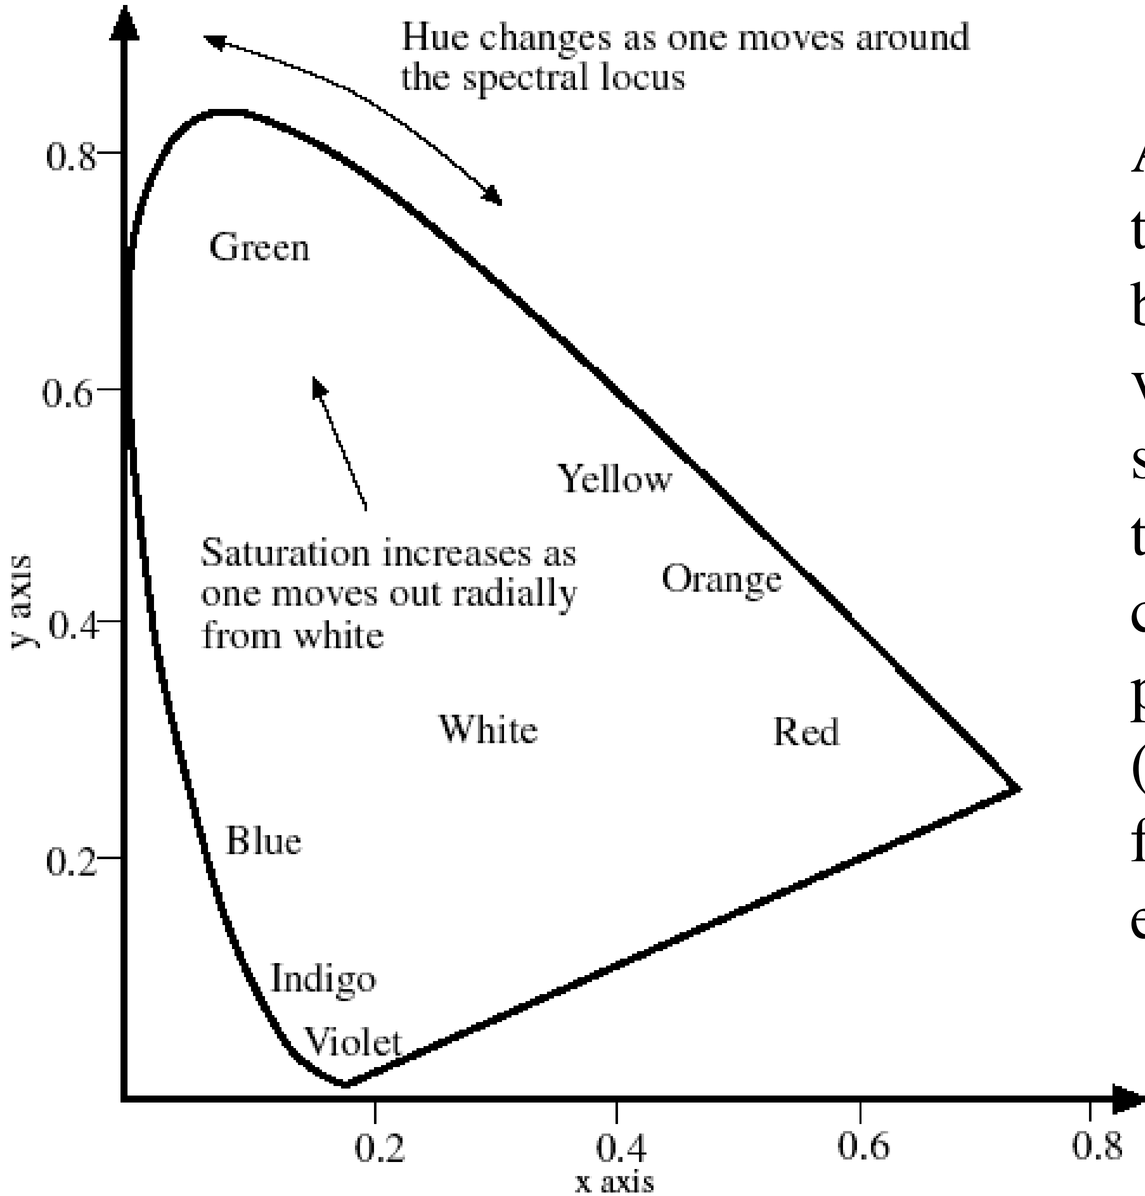


A qualitative rendering of the CIE (x,y) space. The blobby region represents visible colors. There are sets of (x, y) coordinates that don’t represent real colors, because the primaries are not real lights (so that the color matching functions could be positive everywhere).

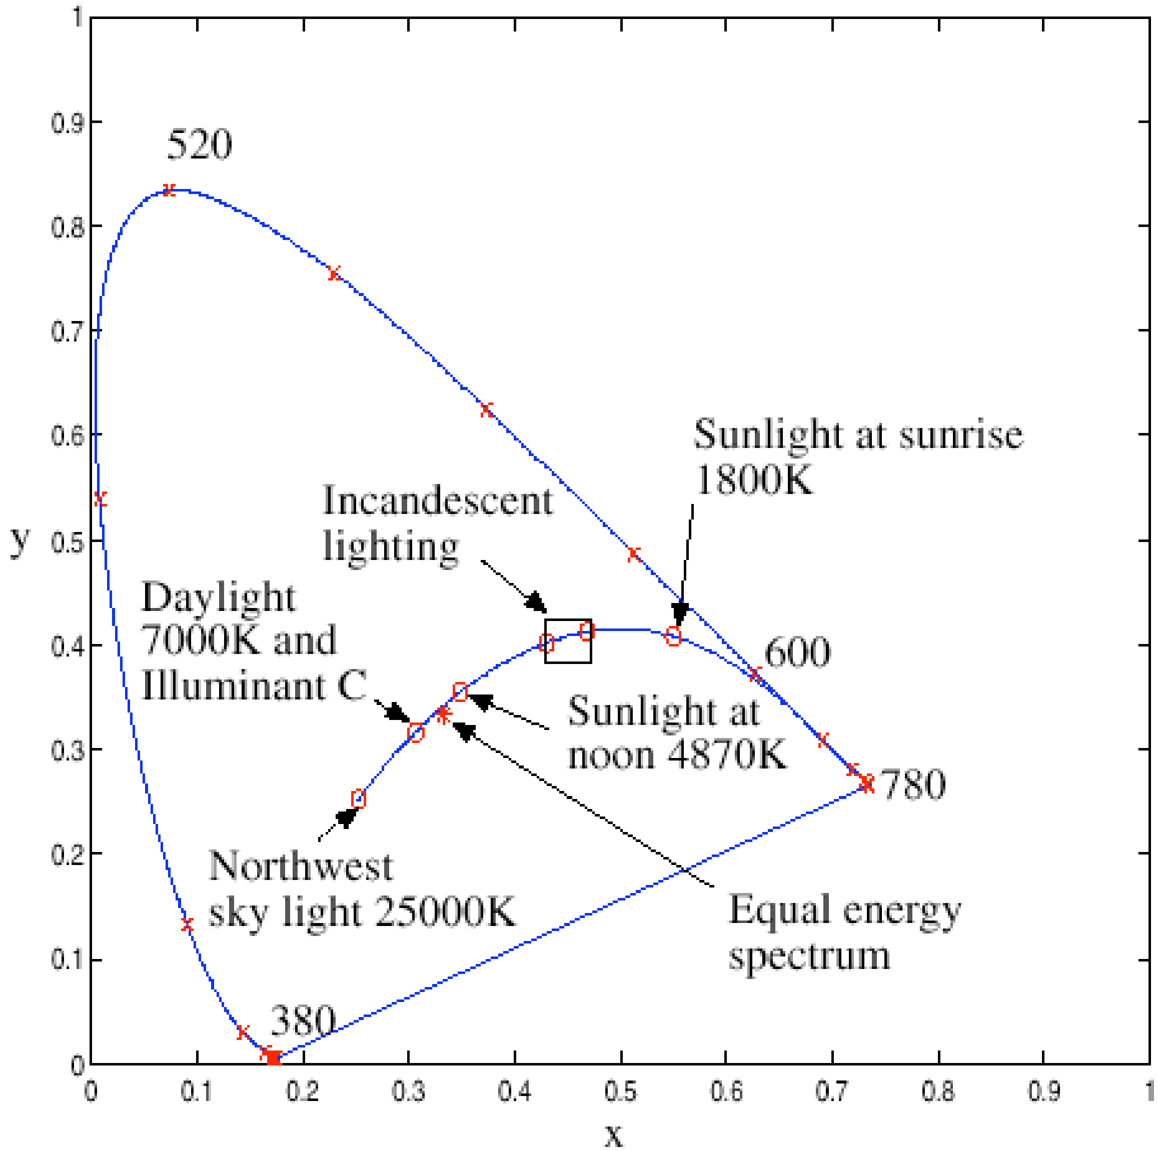


A plot of the CIE (x,y) space. We show the spectral locus (the colors of monochromatic lights) and the black-body locus (the colors of heated black-bodies). I have also plotted the range of typical incandescent lighting.

# Non-linear colour spaces

- HSV: Hue, Saturation, Value are non-linear functions of XYZ.
  - because hue relations are naturally expressed in a circle

- Uniform: equal (small!) steps give the same perceived color changes.

- Munsell: describes surfaces, rather than lights - less relevant for graphics. Surfaces must be viewed under fixed comparison light

# HSV hexcone

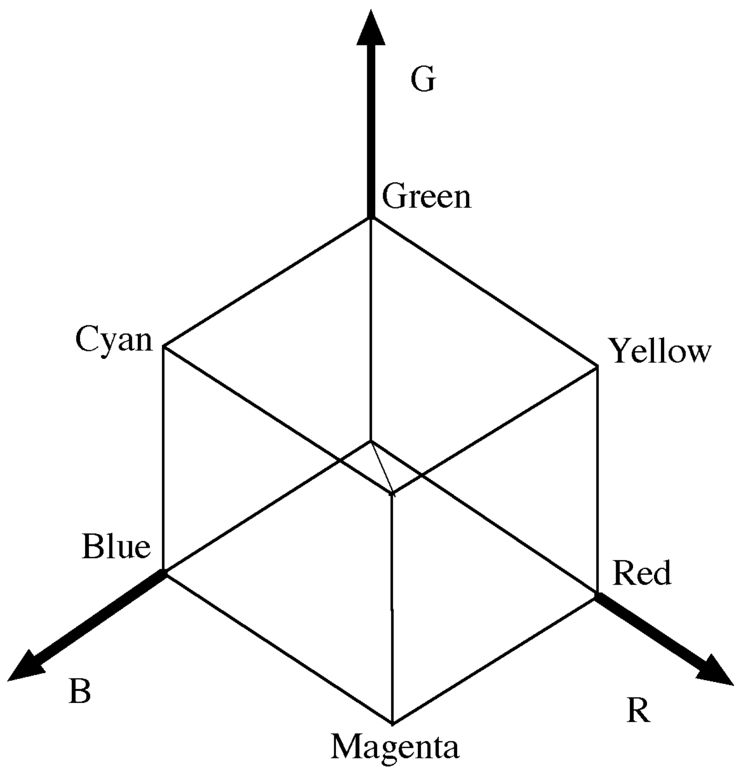
G
Green
Cyan
Yellow
Blue
Red
B
R
Magenta

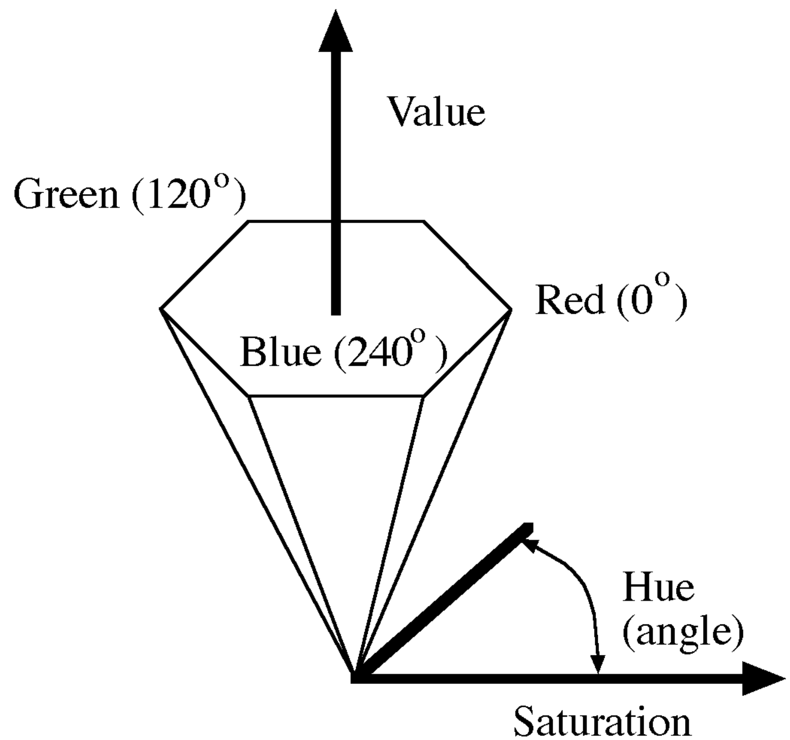
Value
Green (120°)
Red (0°)
Blue (240° )
Hue
(angle)
Saturation

# Uniform color spaces

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

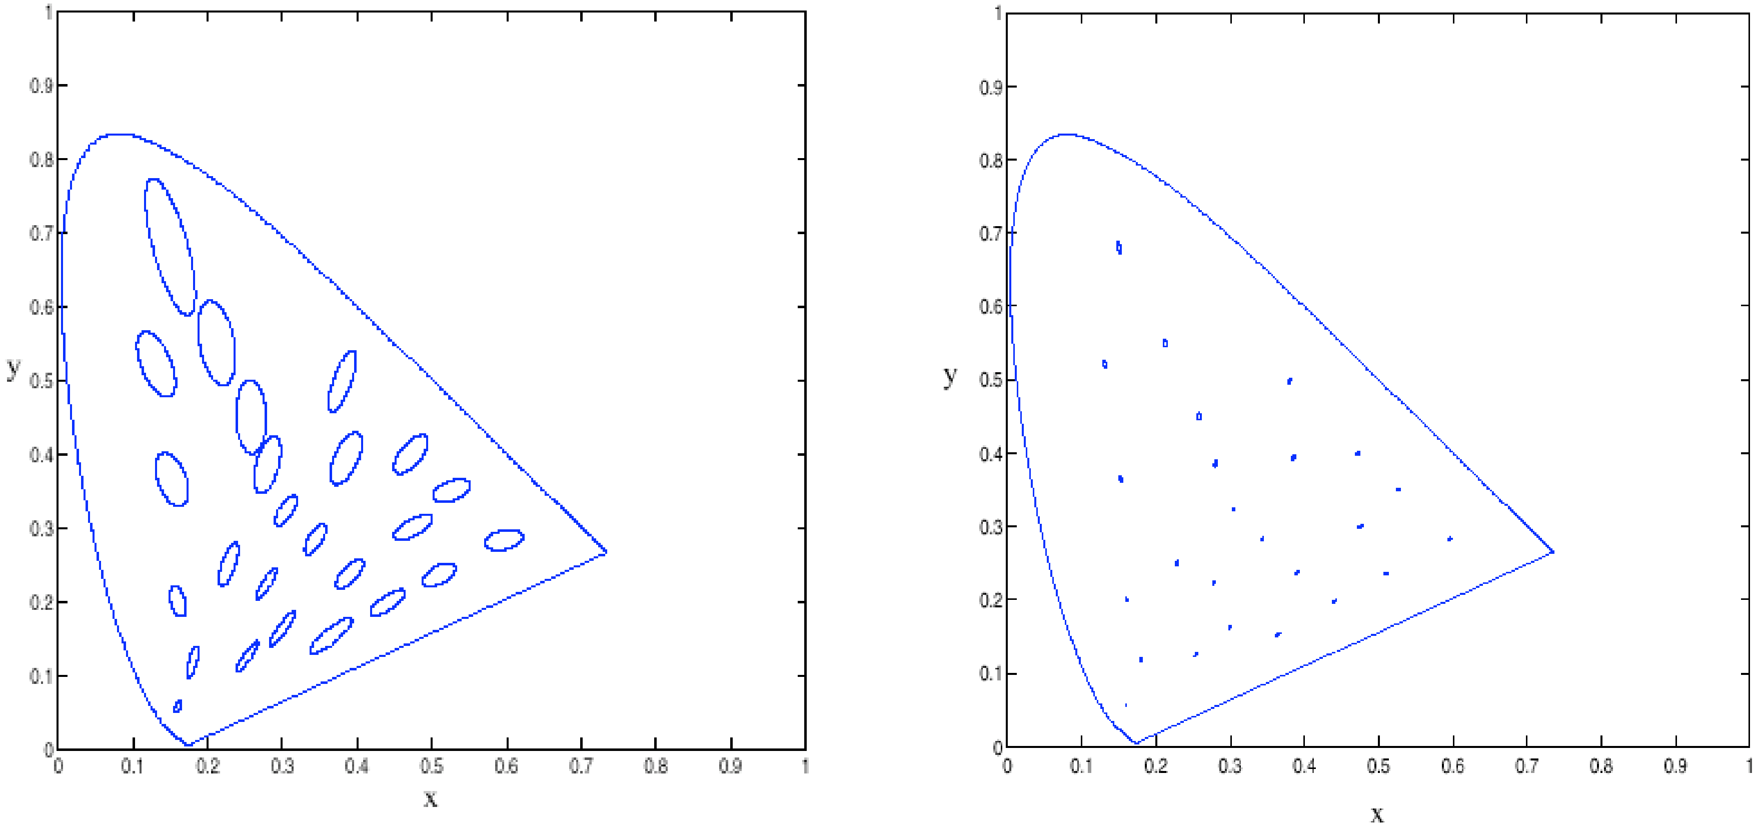


Variations in color matches on a CIE x, y space. At the center of the ellipse is the color of a test light; the size of the ellipse represents the scatter of lights that the human observers tested would match to the test color; the boundary shows where the just noticeable difference is. The ellipses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

CIE u'v' which is a projective transform of x, y. We transform x,y so that ellipses are most like one another. Figure shows the transformed ellipses.

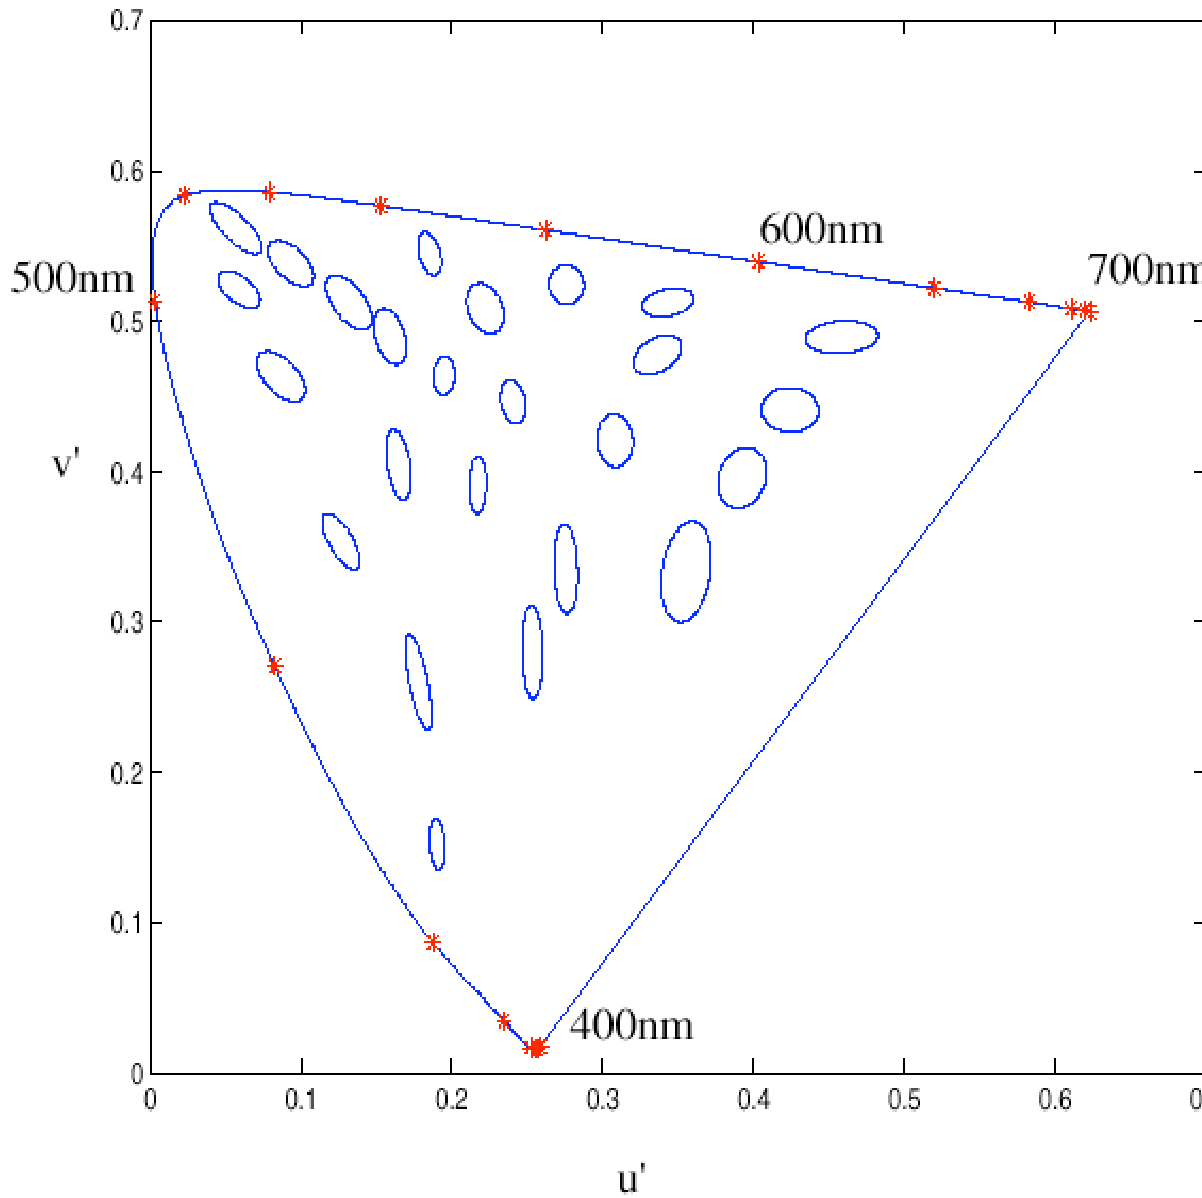

# Color constancy

- Assume we've identified and removed specularities
- The spectral radiance at the camera depends on two things
  - surface albedo
  - illuminant spectral radiance
  - the effect is much more pronounced than most people think (see following slides)
- We would like an illuminant invariant description of the surface
  - e.g. some measurements of surface albedo
  - need a model of the interactions

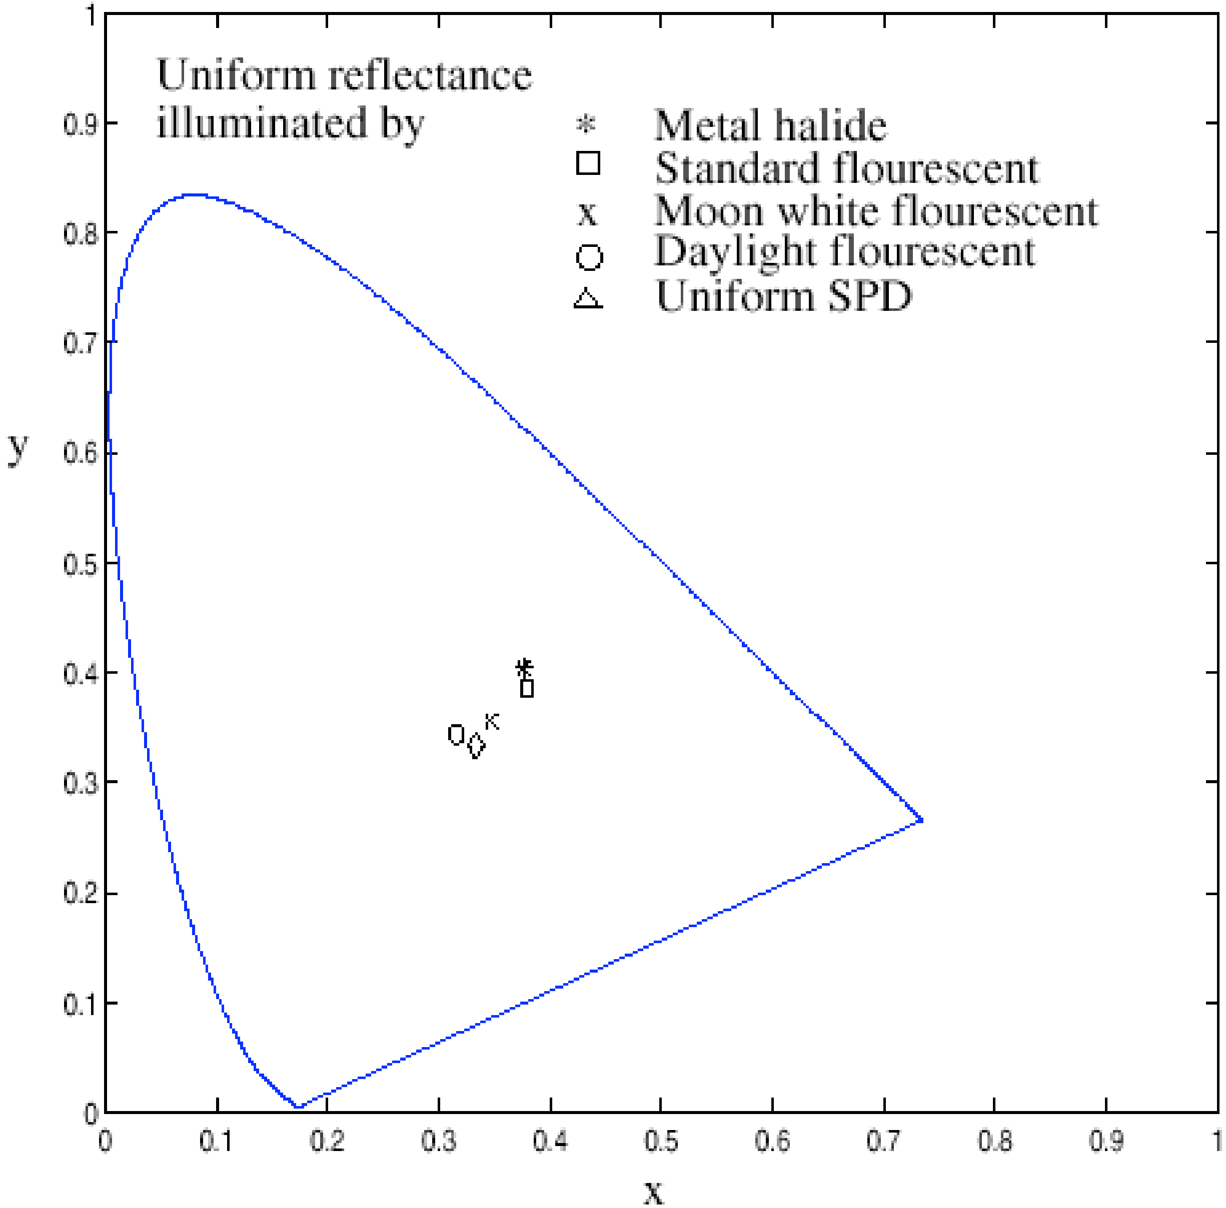


Notice how the color of light at the camera varies with the illuminant color; here we have a uniform reflectance illuminated by five different lights, and the result plotted on CIE x,y

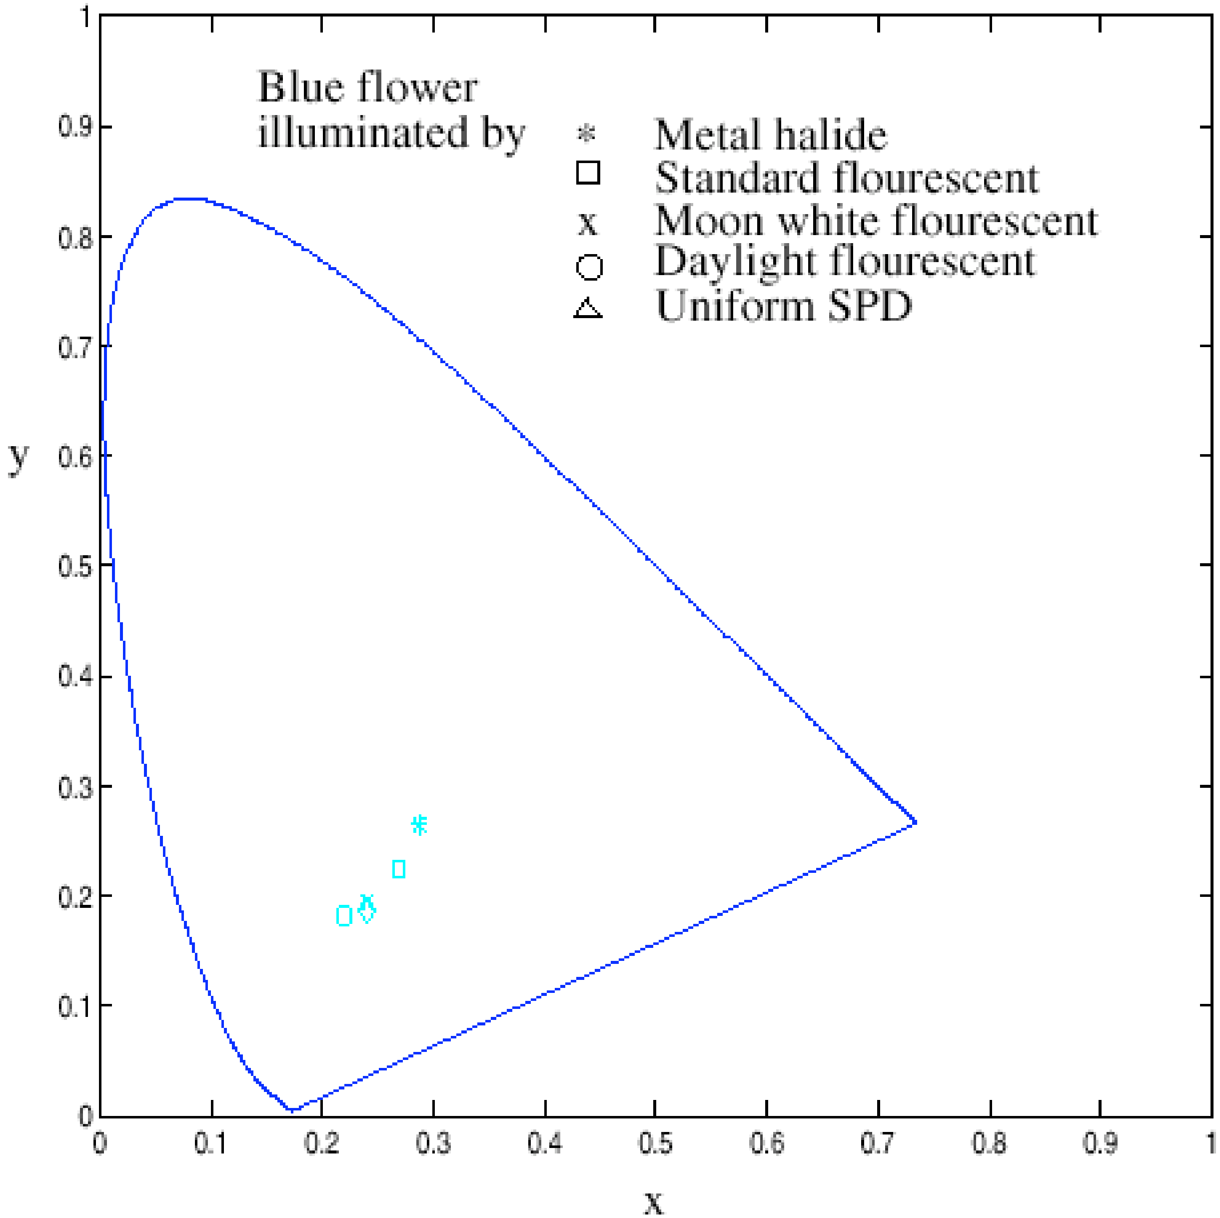


Notice how the color of light at the camera varies with the illuminant color; here we have the blue flower illuminated by five different lights, and the result plotted on CIE x,y. Notice how it looks significantly more saturated under some lights.

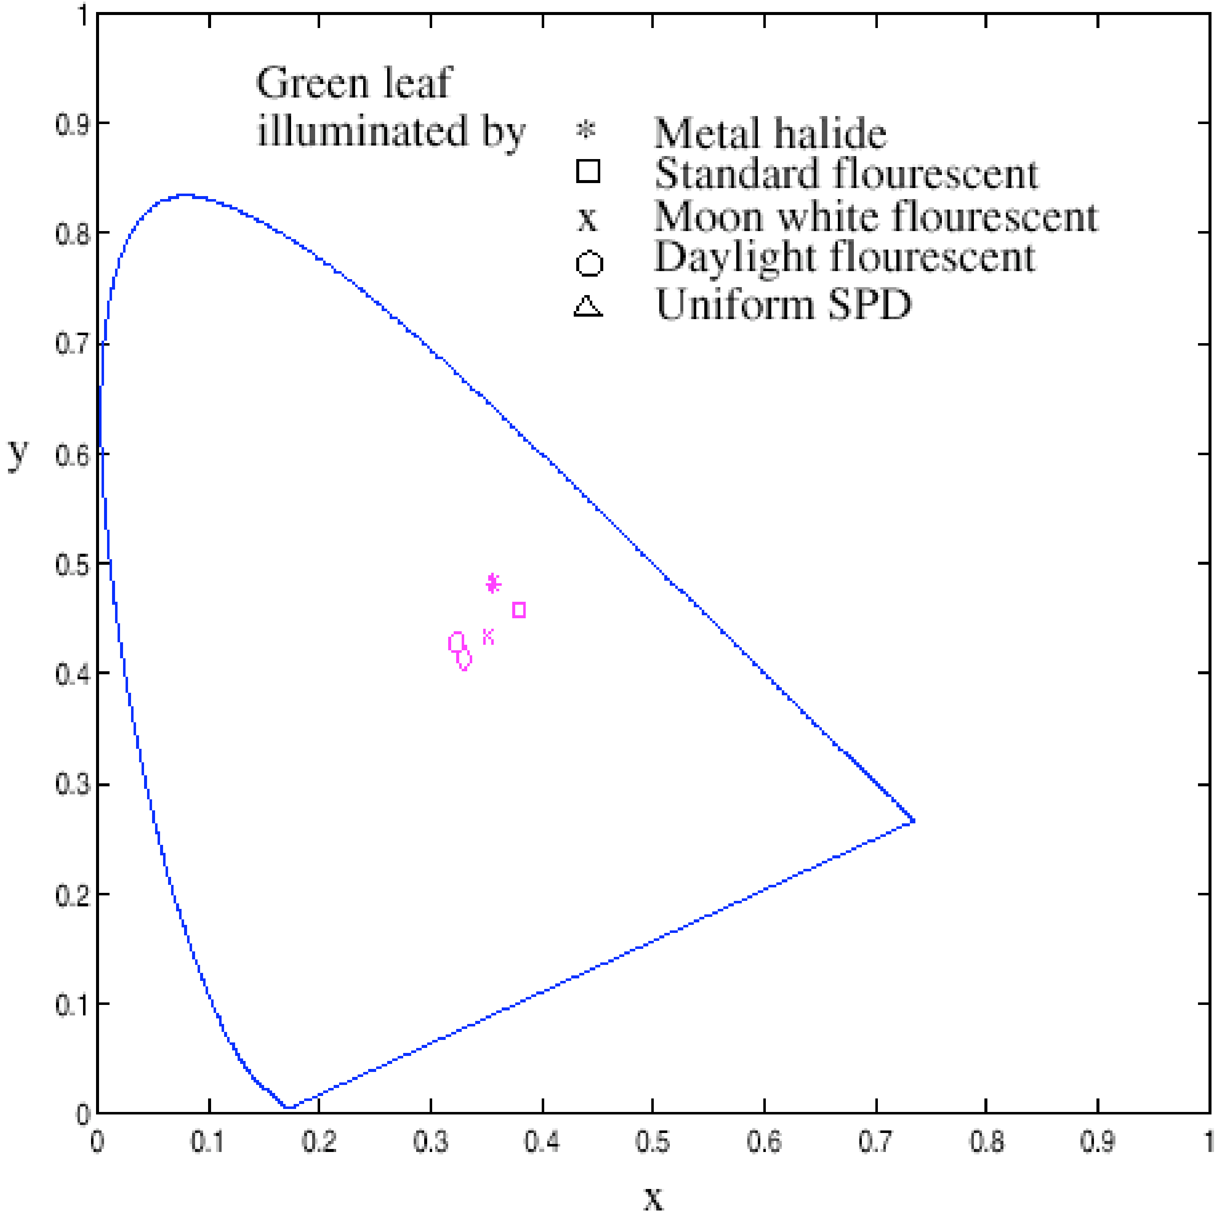


Notice how the color of light at the camera varies with the illuminant color; here we have a green leaf illuminated by five different lights, and the result plotted on CIE x,y

# Viewing coloured objects

- Assume diffuse+specular model
- Specular
  - specularities on dielectric objects take the colour of the light
  - specularities on metals can be coloured
- Diffuse
  - colour of reflected light depends on both illuminant and surface
  - people are surprisingly good at disentangling these effects in practice (colour constancy)
  - this is probably where some of the spatial phenomena in colour perception come from

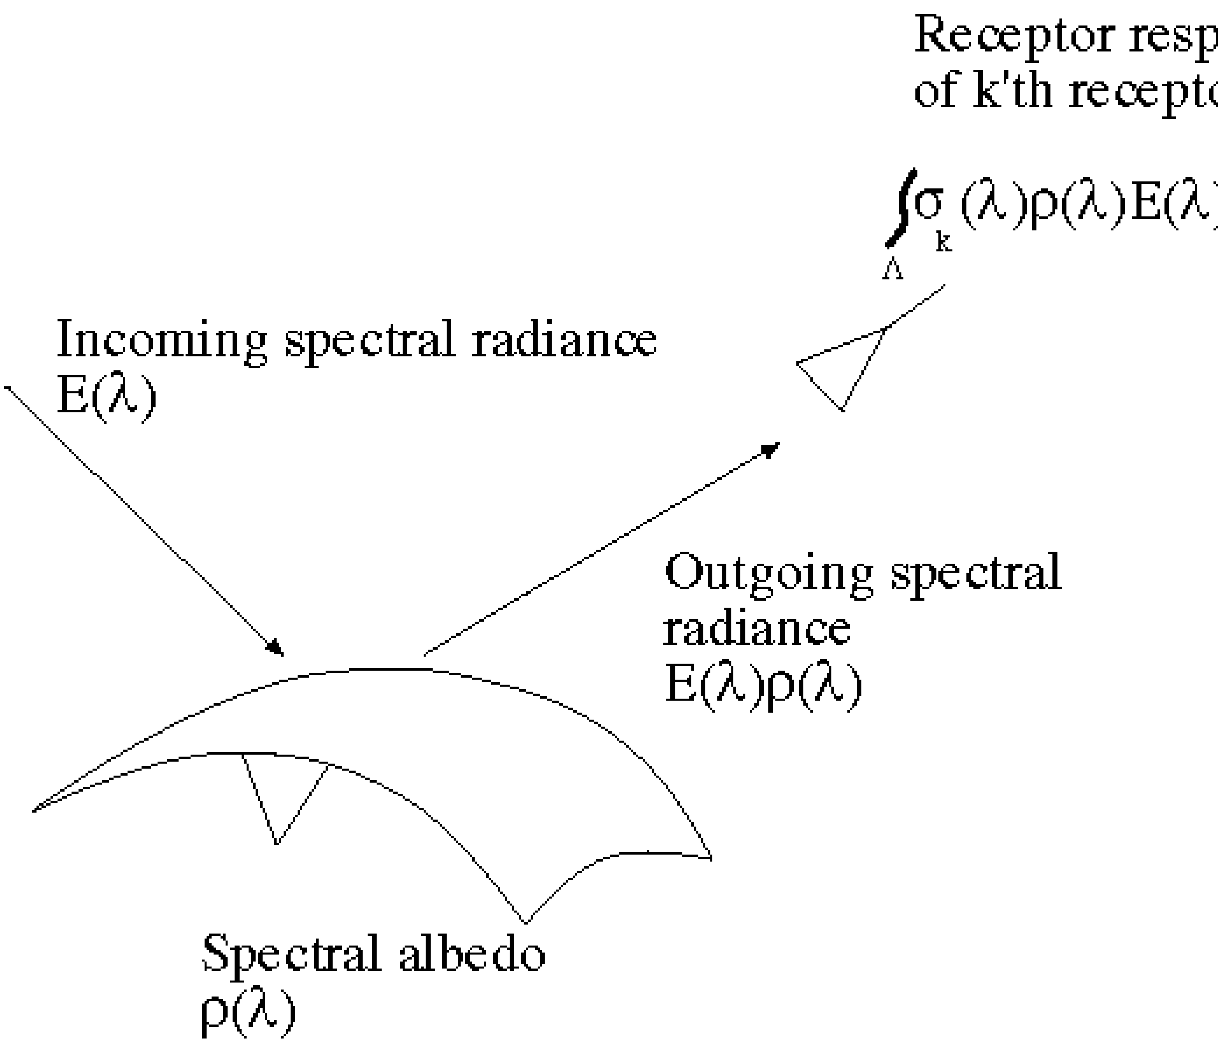


When one views a colored surface, the spectral radiance of the light reaching the eye depends on both the spectral radiance of the illuminant, and on the spectral albedo of the surface. We're assuming that camera receptors are linear, like the receptors in the eye. This is usually the case.

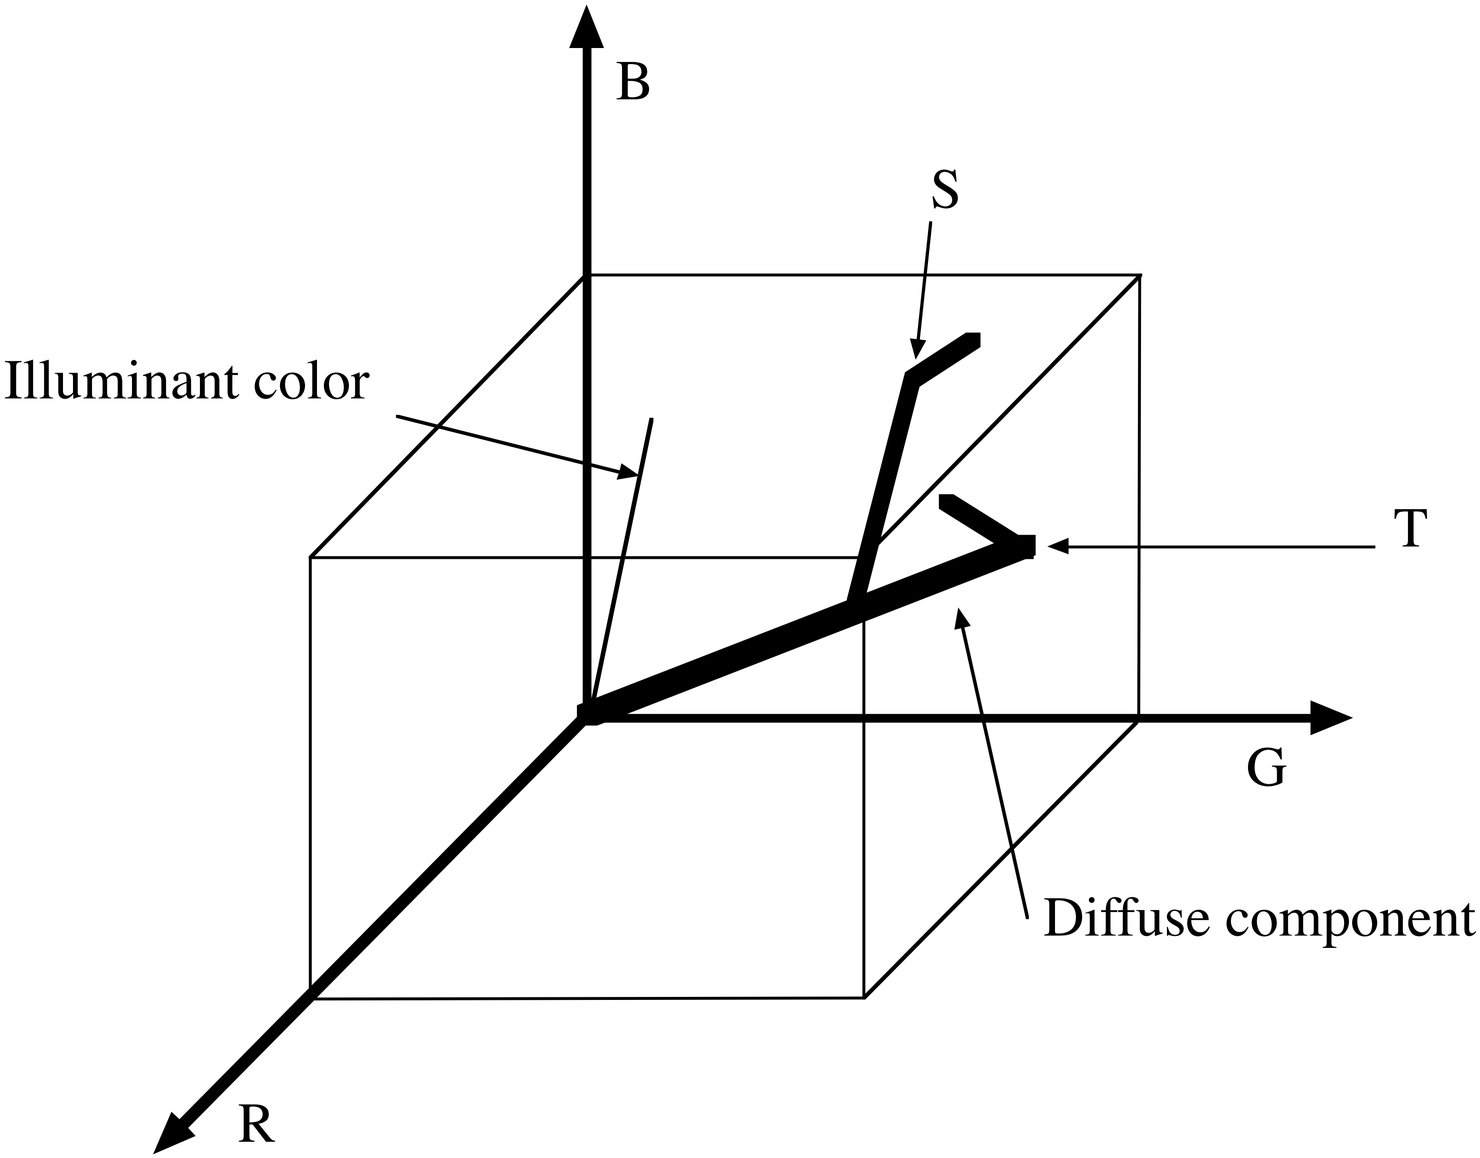
B
S
Illuminant color
T
G
Diffuse component
R

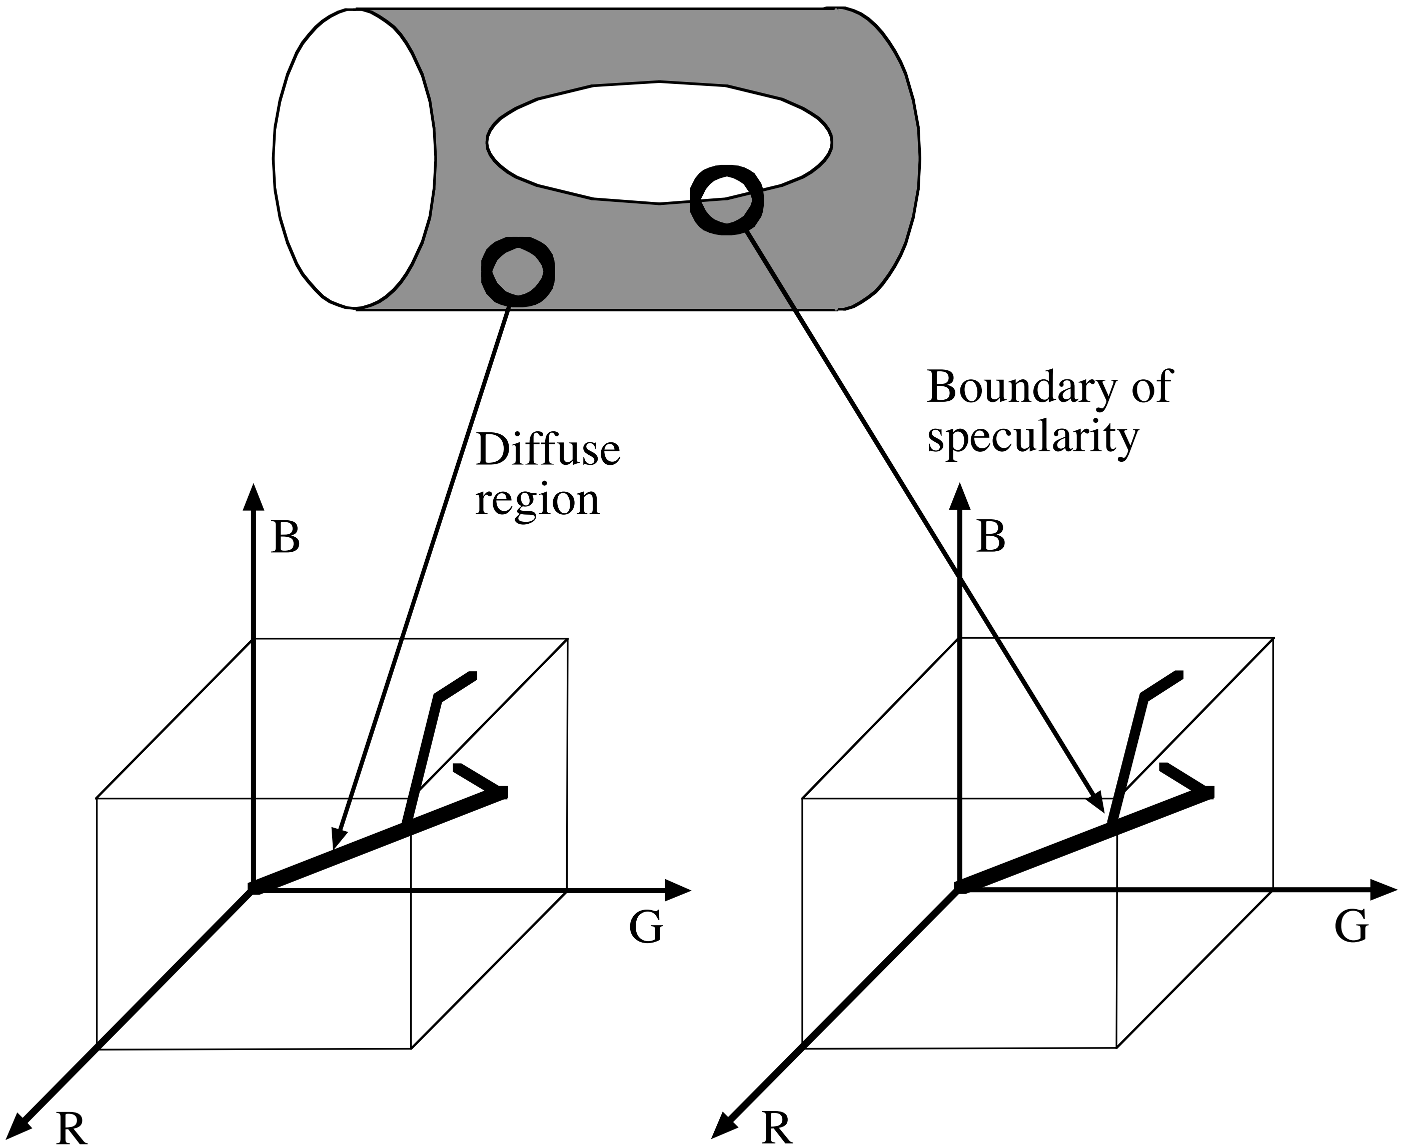

Diffuse
region
Boundary of
specularity
B
G
R
B
G
R

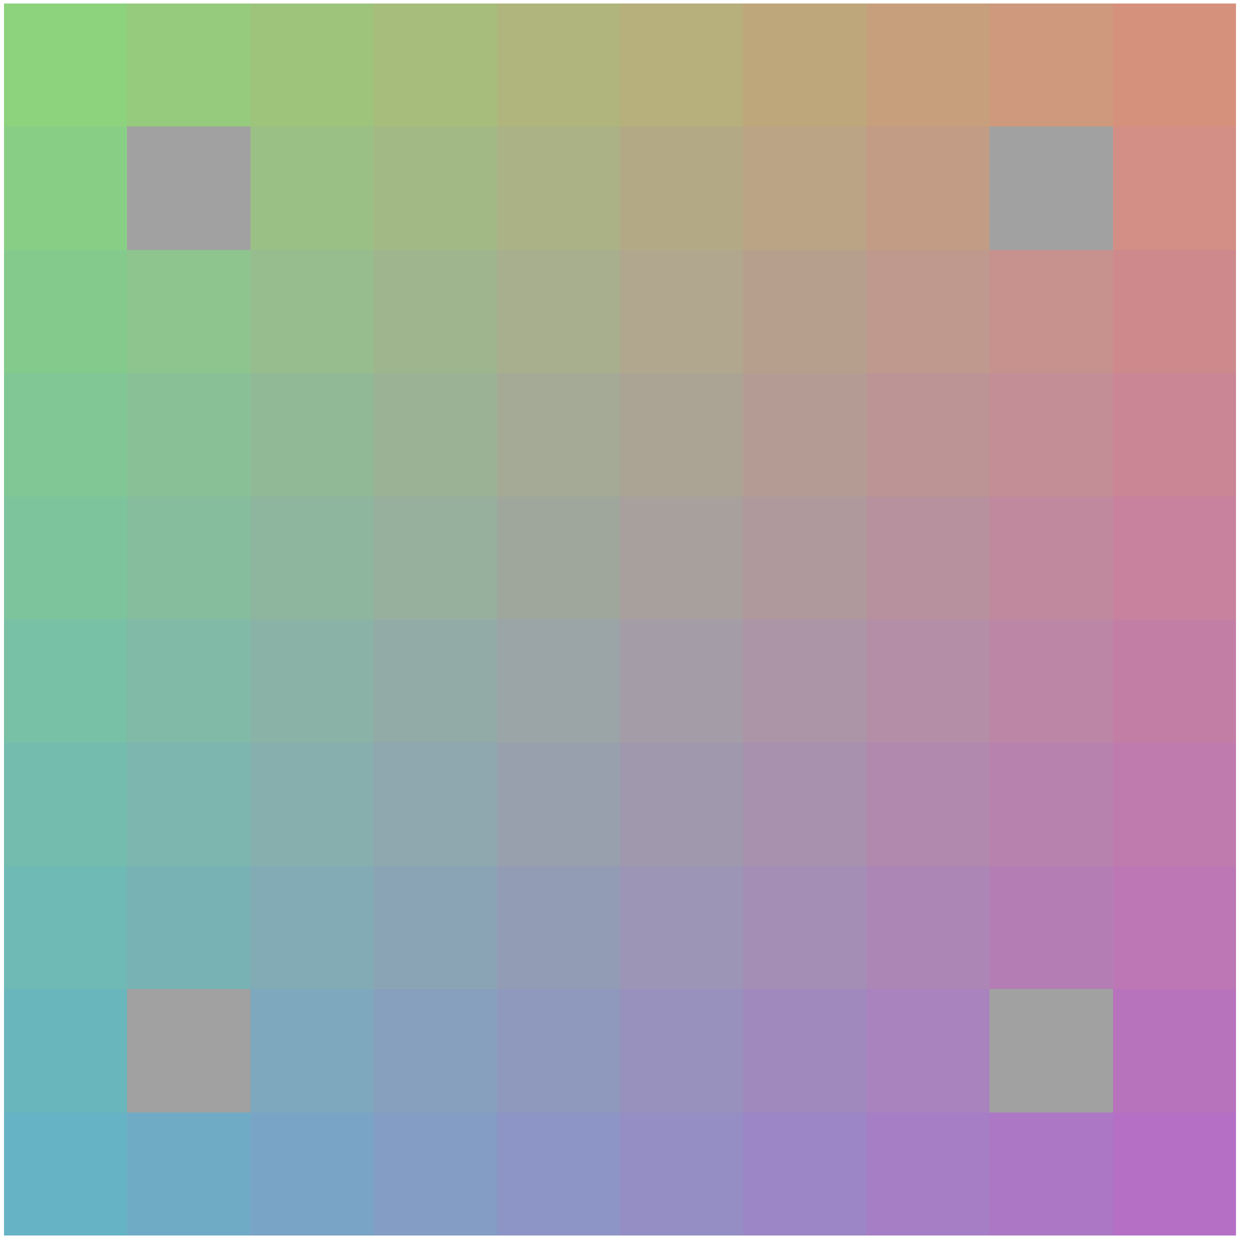

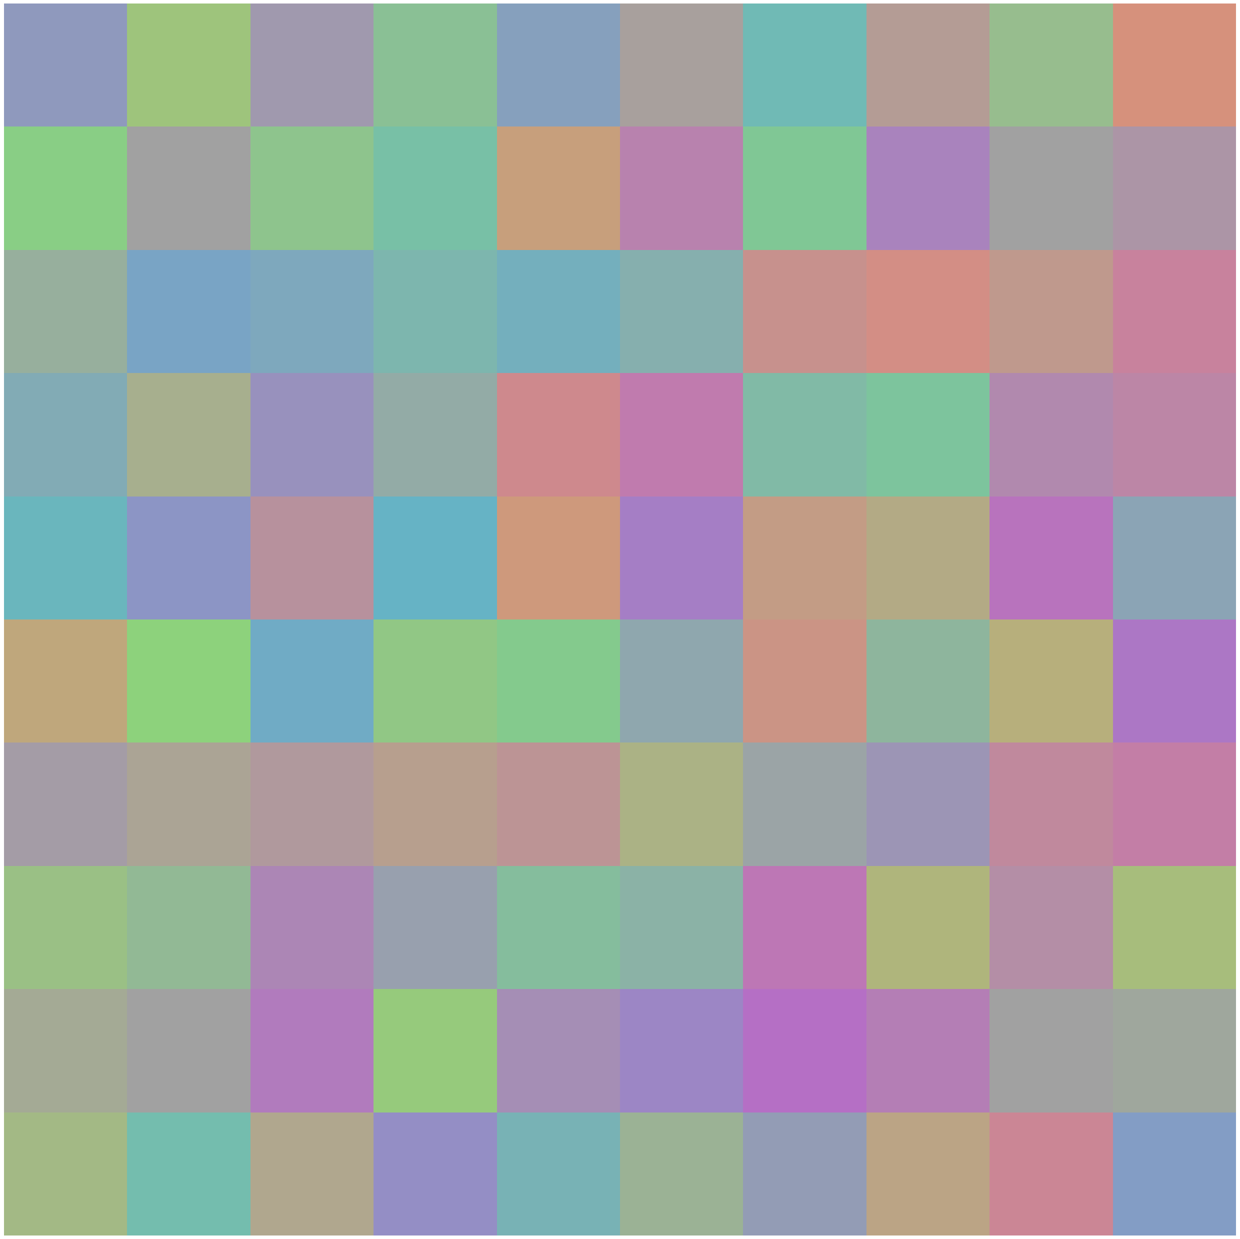

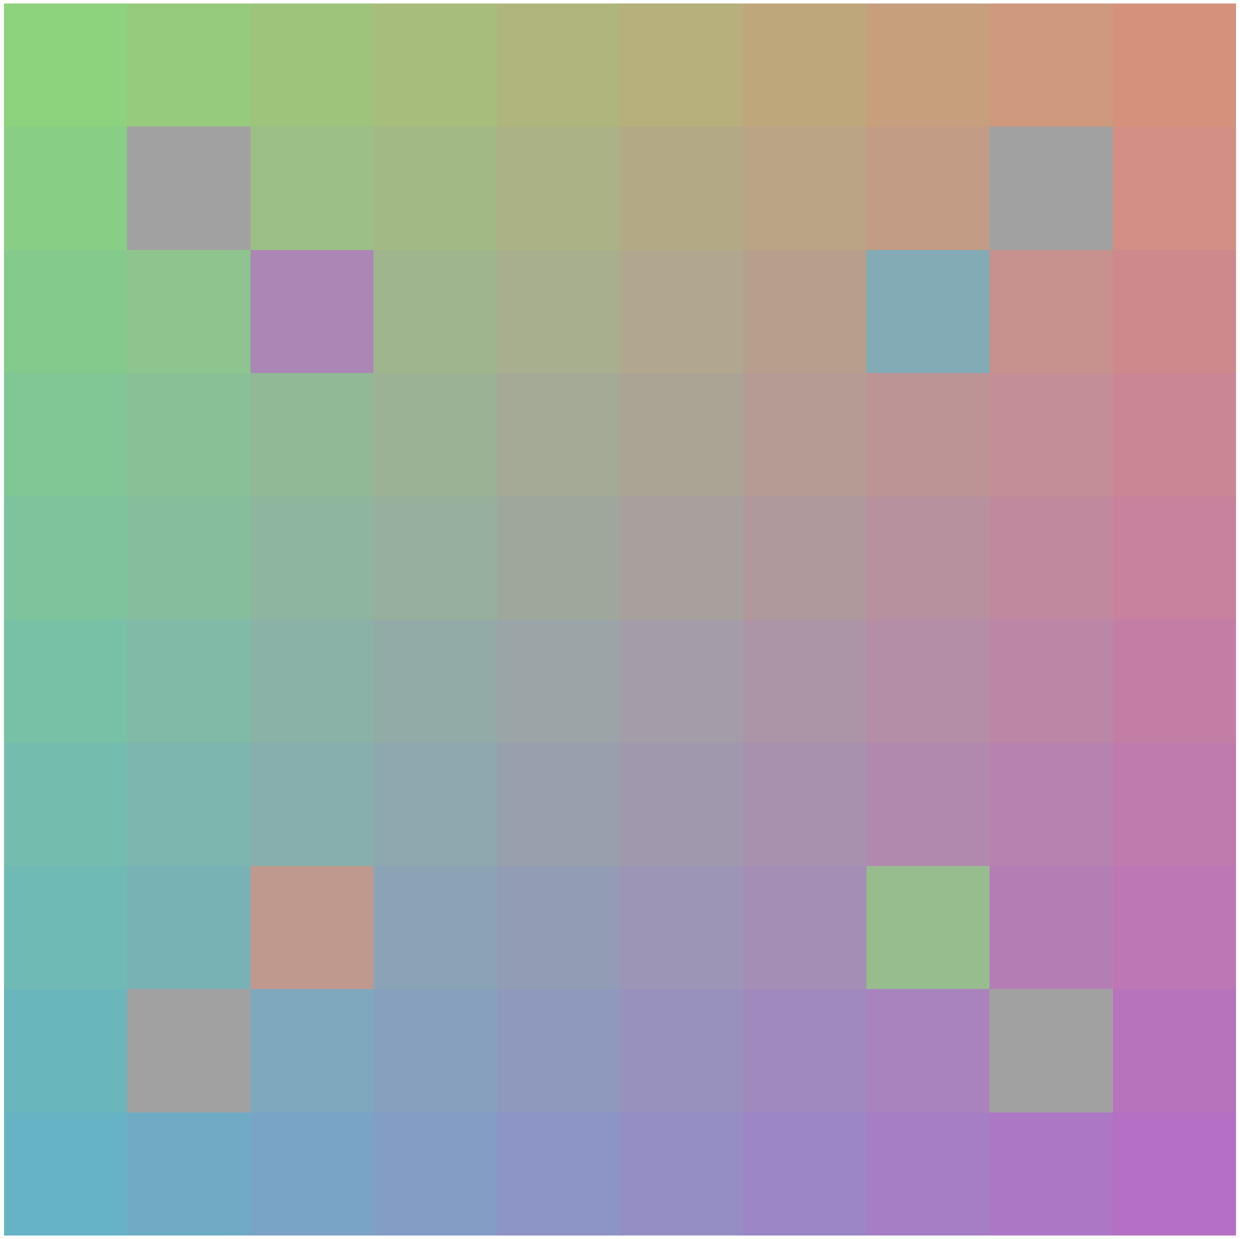

**Figure 1. Same objects imaged under two natural illuminants.** Top) The patches show a rectangular region extracted from images of the same object under different outdoor illuminants. Bottom) The images from which the patches were taken. Images were acquired by the author in Merion Station, PA using a Nikon CoolPix 995 digital camera. The automatic white balancing calculation that is a normal part of the camera's operation was disabled during image acquisition.

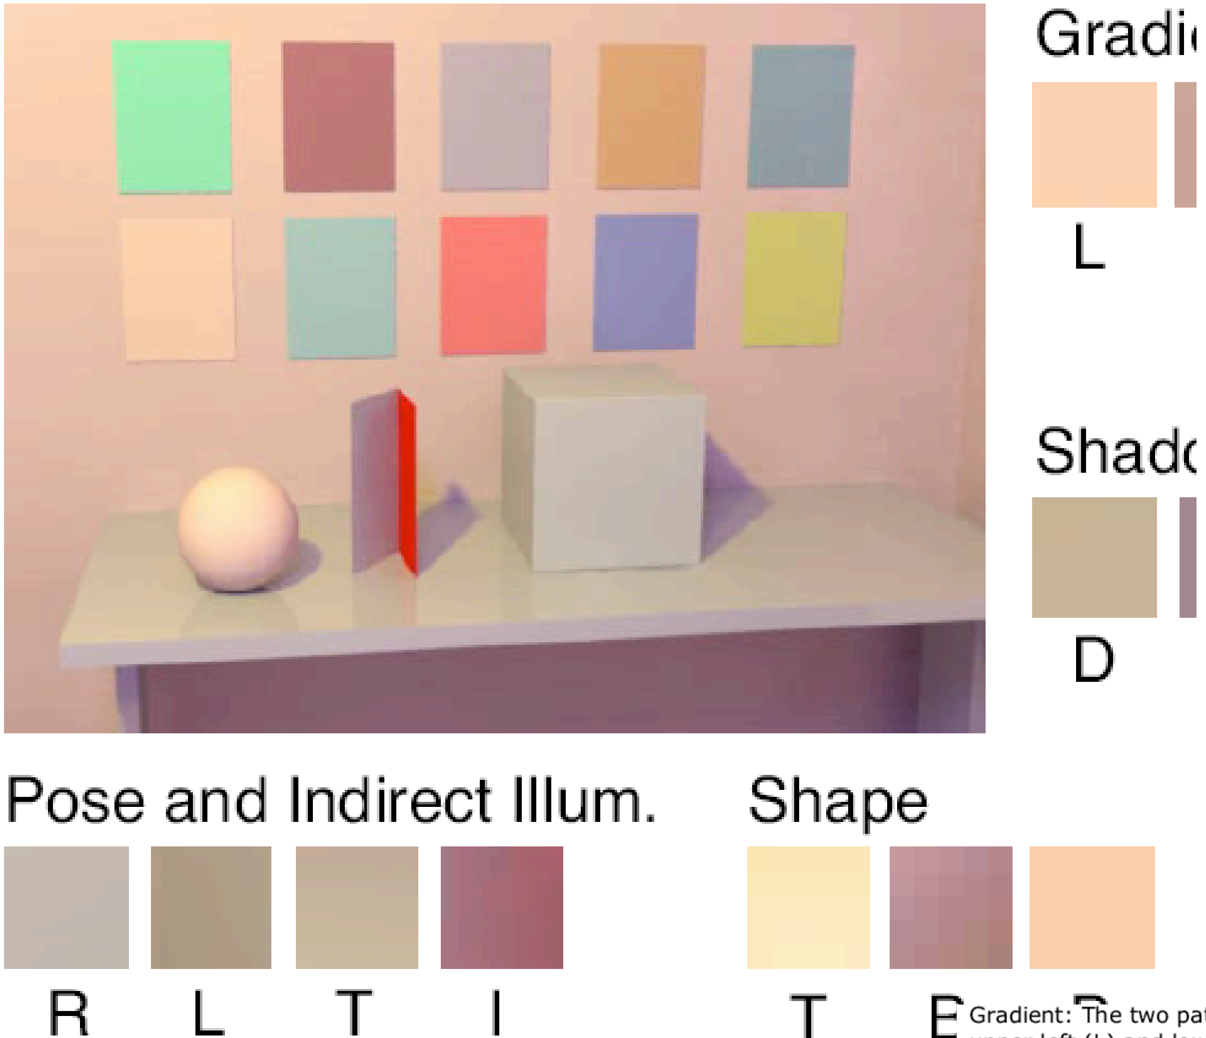


Gradient: The two patches shown were extracted from the upper left (L) and lower right (R; above table) of the back wall of the scene. Shadow. The two patches were extracted from the tabletop in direct illumination (D) and shadow (S). Shape: The three patches shown were extracted from two regions of the sphere (T and B; center top and right bottom respectively) and from the colored panel directly above the sphere (P; the panel is the leftmost of the four in the bottom row). Both the sphere and panel have the same simulated surface reflectance function. Pose and Indirect Illum. The four patches were extracted from the three visible sides of the cube (R, L, and T; right, left and top visible sides respectively) and from the left side of the folded paper located between the cube and sphere (I). The simulated surface reflectance of all sides of the cube and of the left side of the folded paper are identical. The image was

# Land's Demonstration

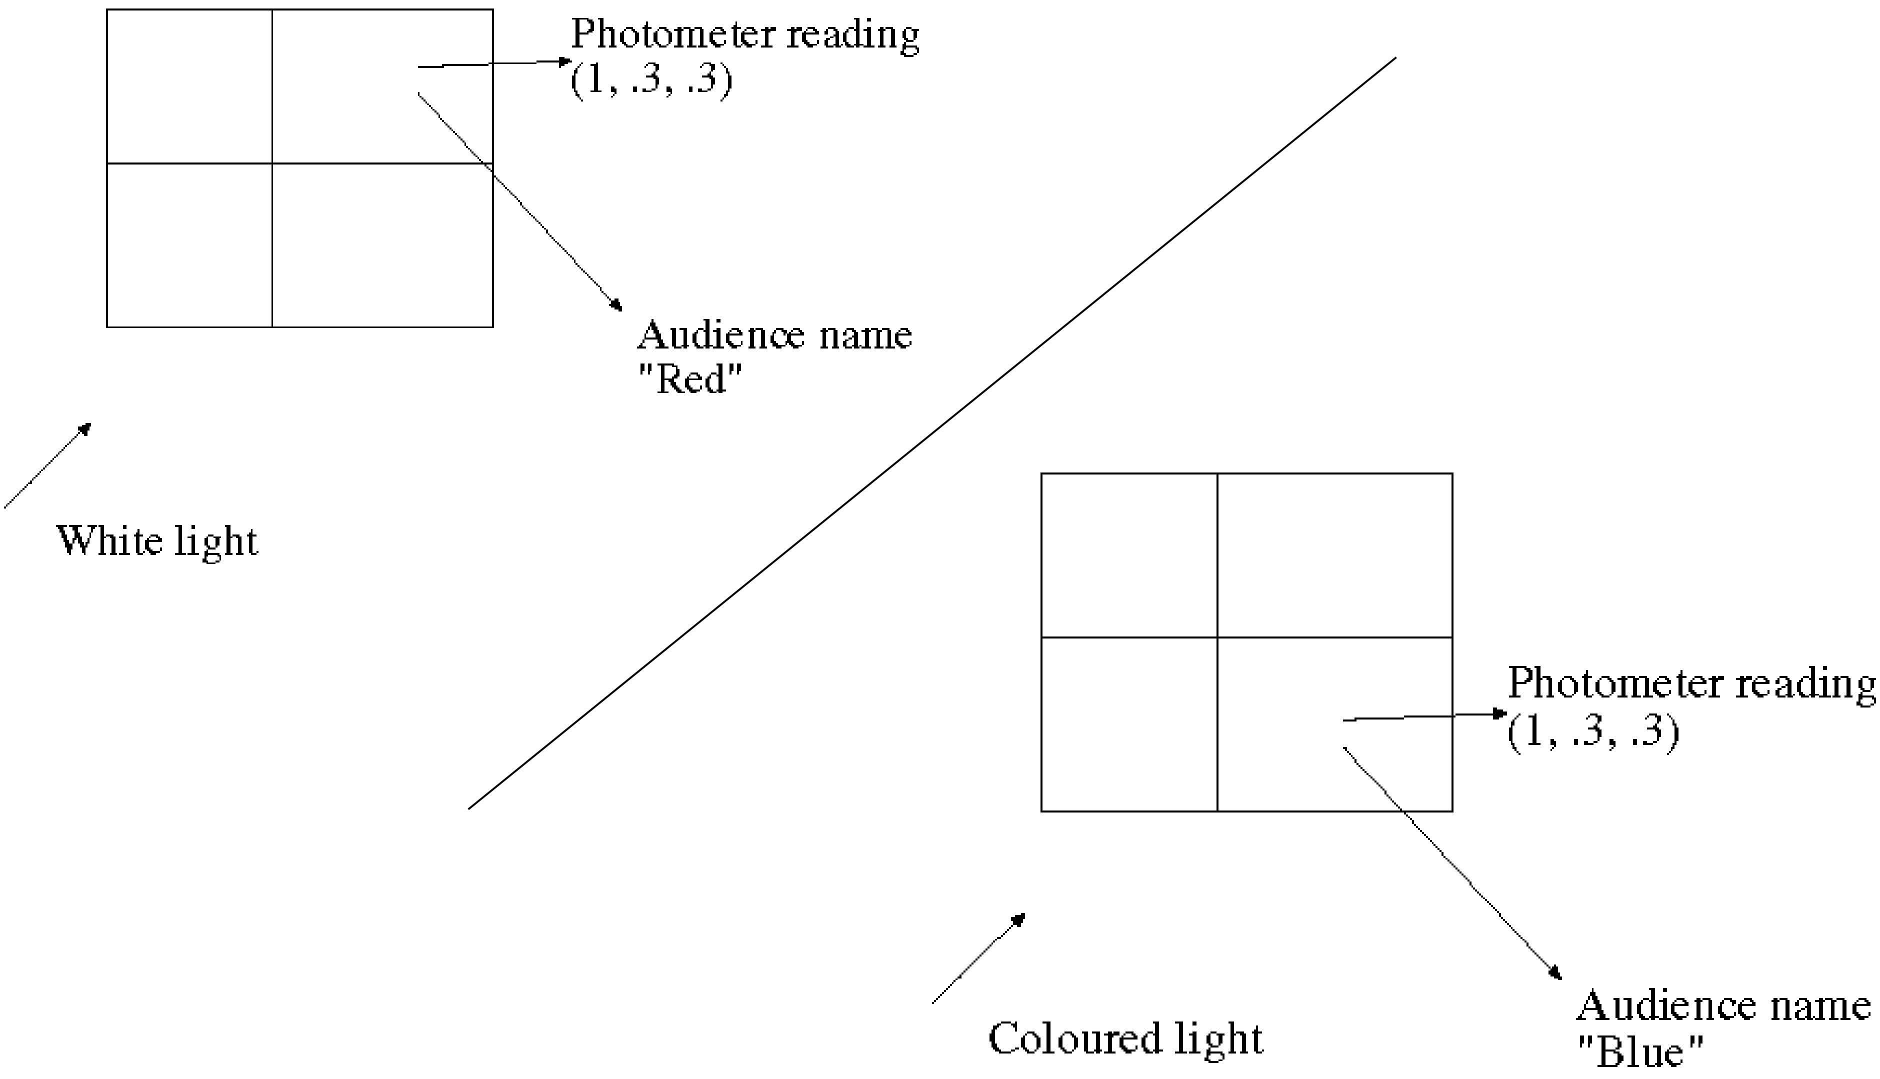

# Lightness Constancy

- Lightness constancy
  - how light is the surface, independent of the brightness of the illuminant
  - issues
    - spatial variation in illumination
    - absolute standard
  - Human lightness constancy is very good
- Assume
  - frontal 1D “Surface”
  - slowly varying illumination
  - quickly varying surface reflectance

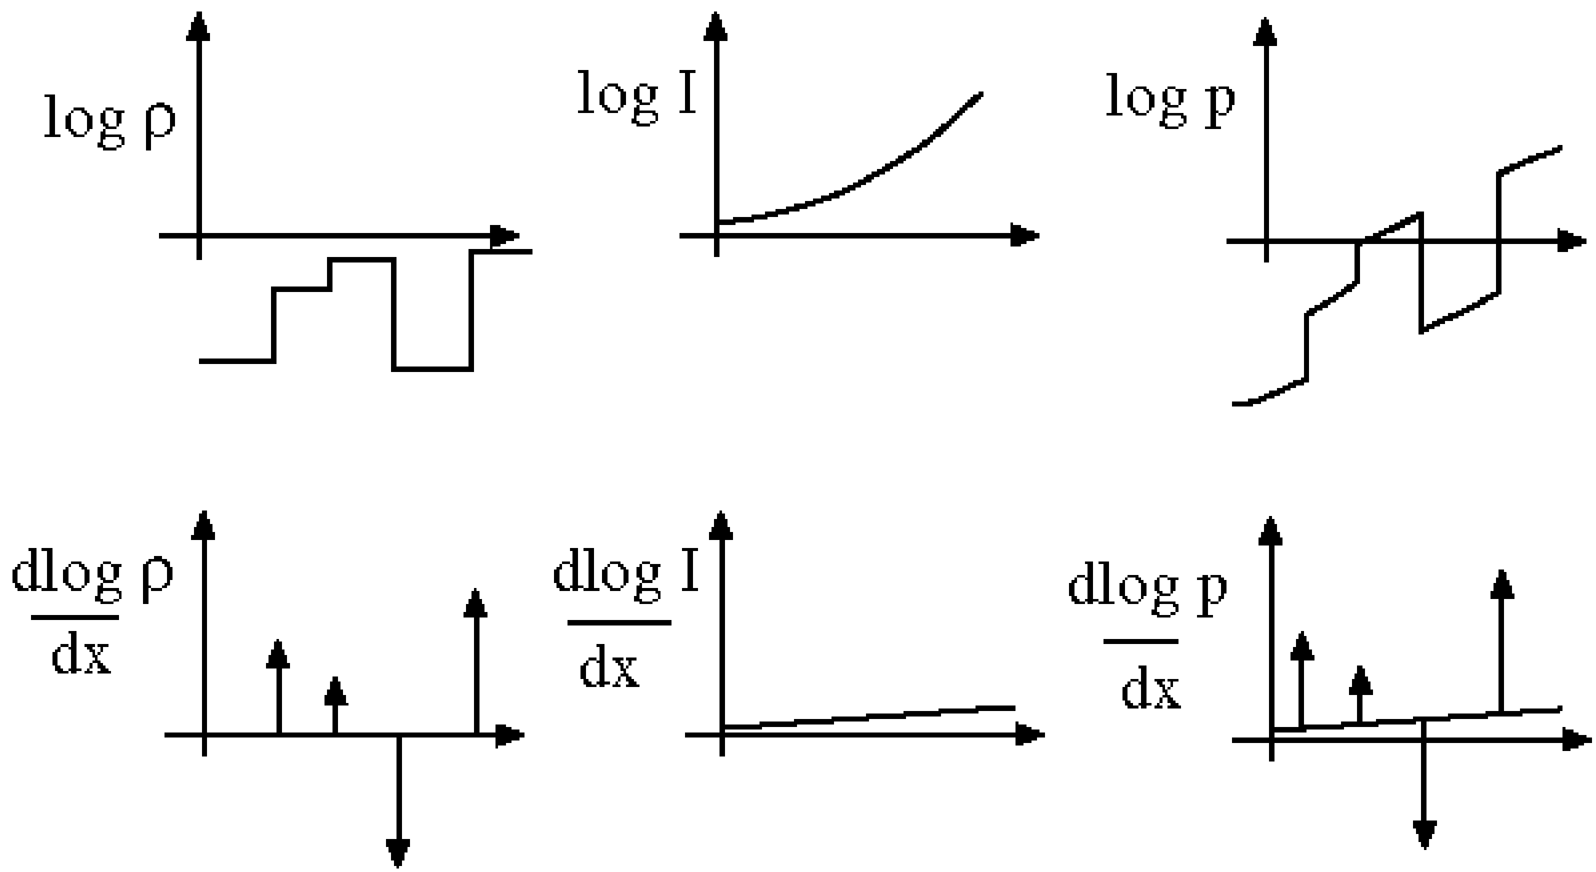
log ρ
log I
log p
dlog ρ
dx
dlog I
dx
dlog p
dx

Thresholded $\frac{d\log p}{dx}$

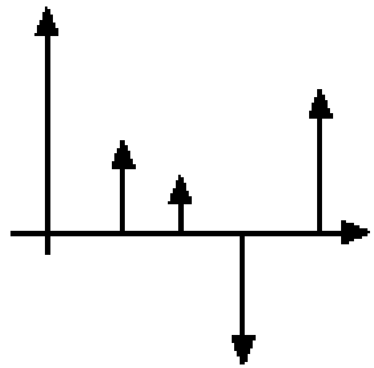

Integrate
This to get

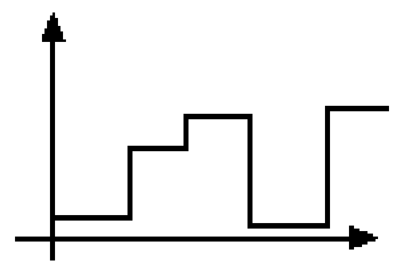

# Lightness Constancy in 2D

- Differentiation, thresholding are easy
  - integration isn't
  - problem - gradient field may no longer be a gradient field
- One solution
  - Choose the function whose gradient is "most like" thresholded gradient
- This yields a minimization problem
- How do we choose the constant of integration?
  - average lightness is grey
  - lightest object is white
  - ?

# Simplest colour constancy

- Adjust three receptor channels independently
  - Von Kries
  - Where does the constant come from?
    - White patch
    - Averages
    - Some other known reference (faces, nose)

# Colour Constancy - I

- We need a model of interaction between illumination and surface colour
  - finite dimensional linear model seems OK
- Finite Dimensional Linear Model (or FDLM)
  - surface spectral albedo is a weighted sum of basis functions
  - illuminant spectral exitance is a weighted sum of basis functions
  - This gives a quite simple form to interaction between the two

# Finite Dimensional Linear Models

$$E(\lambda) = \sum_{i=1}^{m} \varepsilon_i \psi_i(\lambda)$$

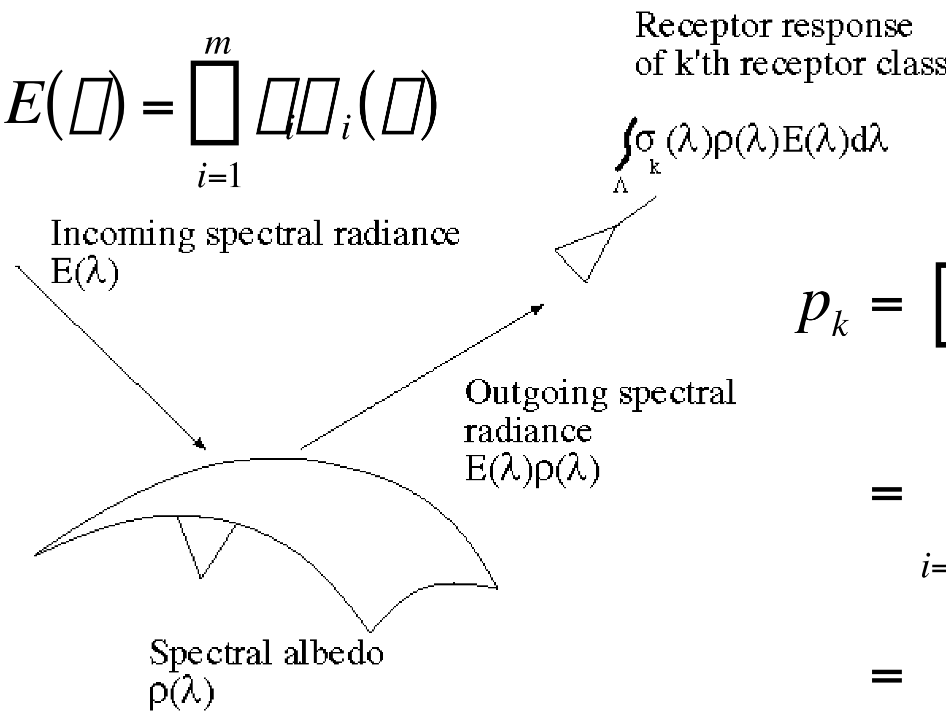


$$\rho(\lambda) = \sum_{j=1}^{n} r_j \varphi_j(\lambda)$$

$$p_k = \int \sigma_k(\lambda) \left\{ \sum_{i=1}^{m} \varepsilon_i \psi_i(\lambda) \right\} \left\{ \sum_{j=1}^{n} r_j \varphi_j(\lambda) \right\} d\lambda$$

$$= \sum_{i=1,\, j=1}^{m,n} \varepsilon_i r_j \int \sigma_k(\lambda) \psi_i(\lambda) \varphi_j(\lambda) d\lambda$$

$$= \sum_{i=1,\, j=1}^{m,n} \varepsilon_i r_j g_{ijk}$$

# General strategies

- Determine what image would look like under white light
- 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

# Obtaining the illuminant from specularities

- Assume that a specularity has been identified, and material is dielectric.
- Then in the specularity, we have

$$p_k = \int \rho_k(\lambda) E(\lambda) d\lambda$$

$$= \sum_{i=1}^{m} \varepsilon_i \int \rho_k(\lambda) \psi_i(\lambda) d\lambda$$

- Assuming
  - we know the sensitivities and the illuminant basis functions
  - there are no more illuminant basis functions than receptors
- This linear system yields the illuminant coefficients.

(unfortunately, not true)

# Obtaining the illuminant from average color assumptions

- Assume the spatial average reflectance is known

$$\bar{\rho}(\lambda) = \sum_{j=1}^{n} \bar{r}_j \psi_j(\lambda)$$

- We can measure the spatial average of the receptor response to get

$$\bar{p}_k = \sum_{i=1,j=1}^{m,n} \varepsilon_i \bar{r}_j g_{ijk}$$

- Assuming
  - g_ijk are known
  - average reflectance is known
  - there are not more receptor types than illuminant basis functions
- We can recover the illuminant coefficients from this linear system

# Computing surface properties

- Two strategies
  - compute reflectance coefficients
  - compute appearance under white light.
- These are essentially equivalent.
- Once illuminant coefficients are known, to get reflectance coefficients we solve the linear system

$$p_k = \sum_{i=1,j=1}^{m,n} \epsilon_i r_j g_{ijk}$$

- to get appearance under white light, plug in reflectance coefficients and compute

$$p_k = \sum_{i=1,j=1}^{m,n} \epsilon_i^{white} r_j g_{ijk}$$

# Bayesian Color Constancy

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} = Cp(\mathbf{y}|\mathbf{x})p(\mathbf{x}).$$

$$\bar{L}(\tilde{\mathbf{x}}|\mathbf{y}) = \int_{\mathbf{x}} L(\tilde{\mathbf{x}},\mathbf{x})p(\mathbf{x}|\mathbf{y})\mathrm{d}\mathbf{x}.$$

$$s_k = \mathop{\square}_{i=1,\, j=1}^{m,n} \square_i r_j g_{ijk}$$

$$y_k = f(x), \qquad x = \square_i r_j$$

Simplest example:

$$y = ab$$

$$p(\mathbf{y}|\mathbf{x}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{|\mathbf{y} - f(\mathbf{x})|^2}{2\sigma^2}\right].$$

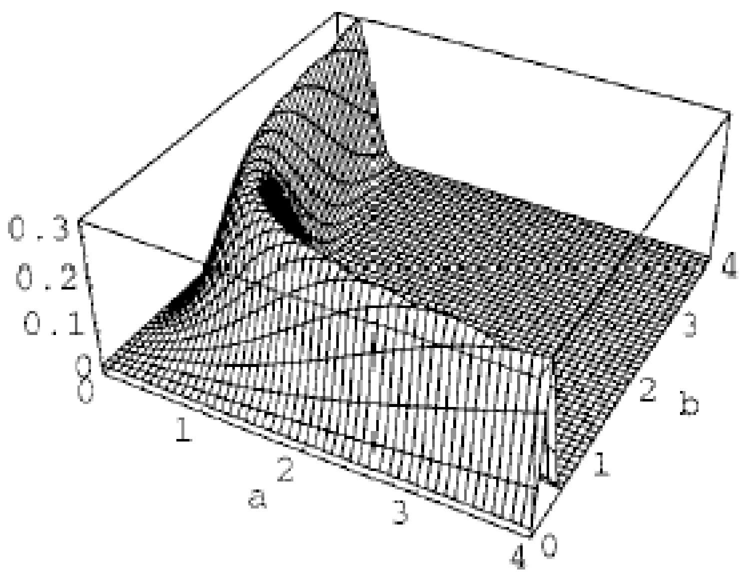


(a)

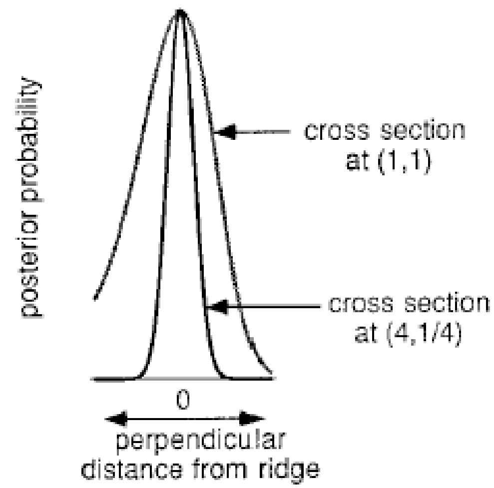


(b)

Fig. 2. Bayesian analysis of the product example: (a) posterior probability for the observed data $ab = 1$ for Gaussian observation noise of variance $\sigma^2 = 0.18$ and uniform prior probabilities over the plotted region, (b) cross section through the posterior at two different locations. Note the different thicknesses of the ridge; some local regions have more probability mass than others, even though the entire ridge has a constant maximum height.

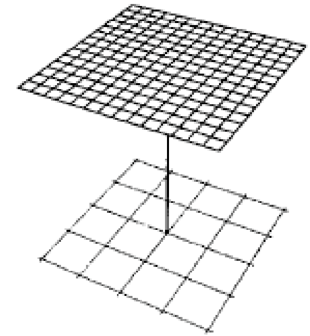

(a) MAP loss function

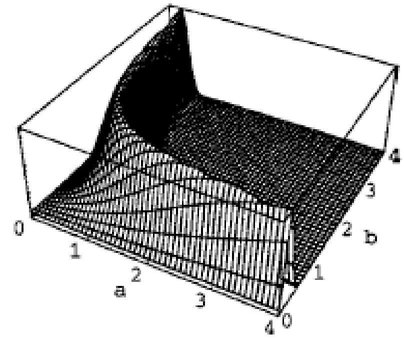


(d) (minus) MAP expected loss

$$L(\tilde{\mathbf{x}}, \mathbf{x}) = -\delta(\tilde{\mathbf{x}} - \mathbf{x}).$$

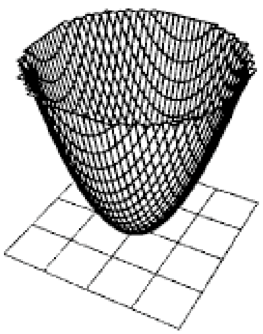

(b) MMSE loss function

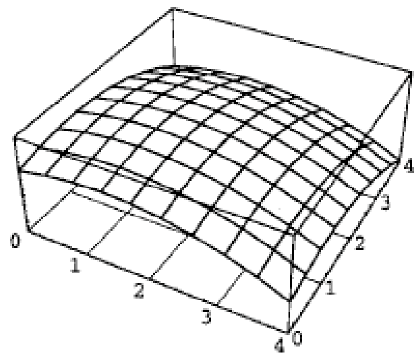


(e) (minus) MMSE expected loss

$$L(\tilde{\mathbf{x}}, \mathbf{x}) = |\tilde{\mathbf{x}} - \mathbf{x}|^2.$$

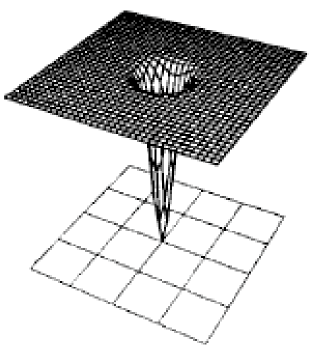

(c) MLM loss function

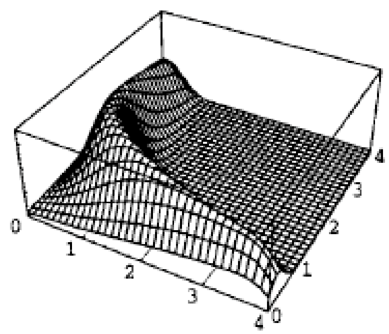


(f) (minus) MLM expected loss

$$L(\tilde{\mathbf{x}}, \mathbf{x}) = -\exp[-|\mathbf{K}_L{}^{-1/2}(\tilde{\mathbf{x}} - \mathbf{x})|^2],$$

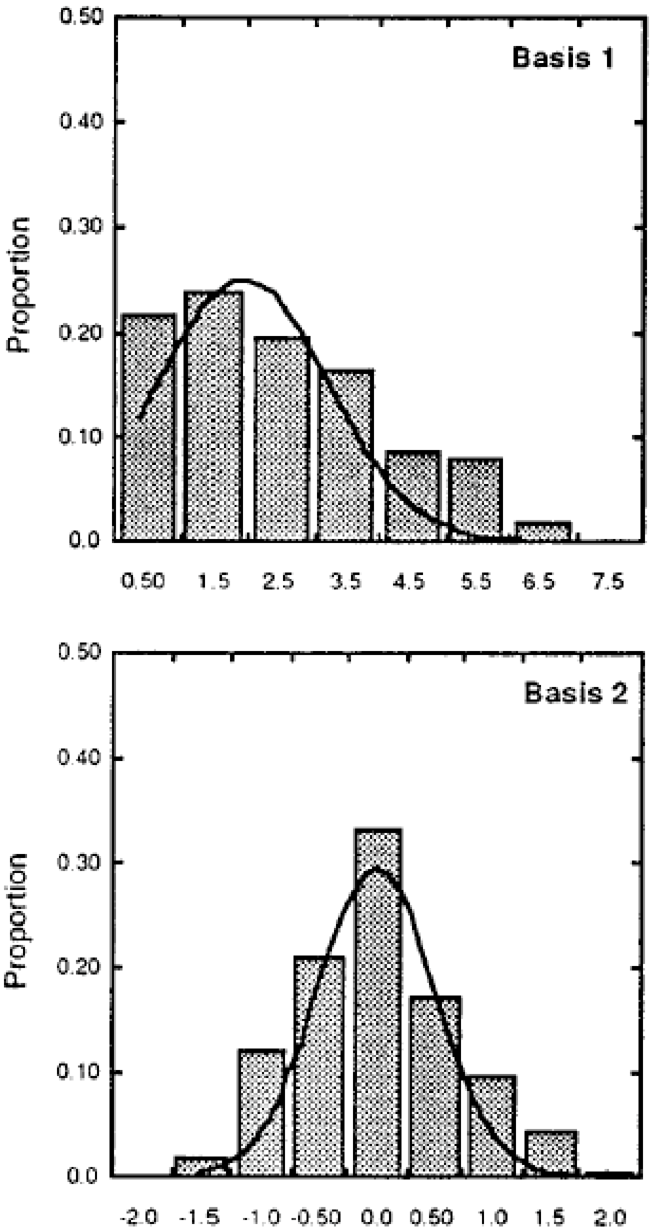
Basis 1
Proportion
0.50
0.40
0.30
0.20
0.10
0.0
0.50
1.5
2.5
3.5
4.5
5.5
6.5
7.5
Basis 2
Proportion
0.50
0.40
0.30
0.20
0.10
0.0
-2.0
-1.5
-1.0
-0.50
0.0
0.50
1.0
1.5
2.0

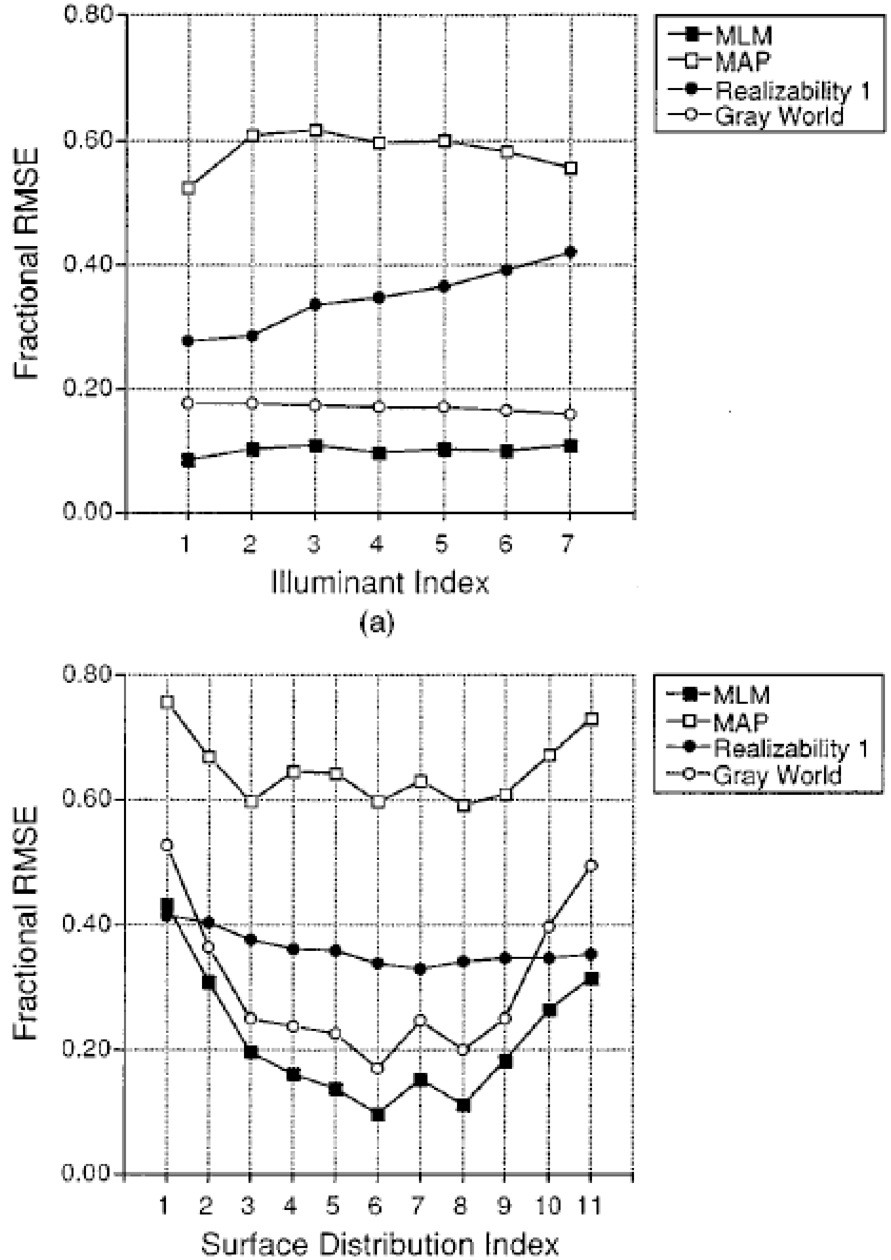
MLM
MAP
Realizability 1
Gray World
Fractional RMSE
0.80
0.60
0.40
0.20
0.00
1
2
3
4
5
6
7
Illuminant Index
(a)
MLM
MAP
Realizability 1
Gray World
Fractional RMSE
0.80
0.60
0.40
0.20
0.00
1
2
3
4
5
6
7
8
9
10
11
Surface Distribution Index

(b)