Lecture 21: Light, shading, and color

Announcements

- Class presentations 12/11 and 12/12
- Project report due on 12/13
- Grading rubric will be released on Monday

What can we infer from intensity changes?

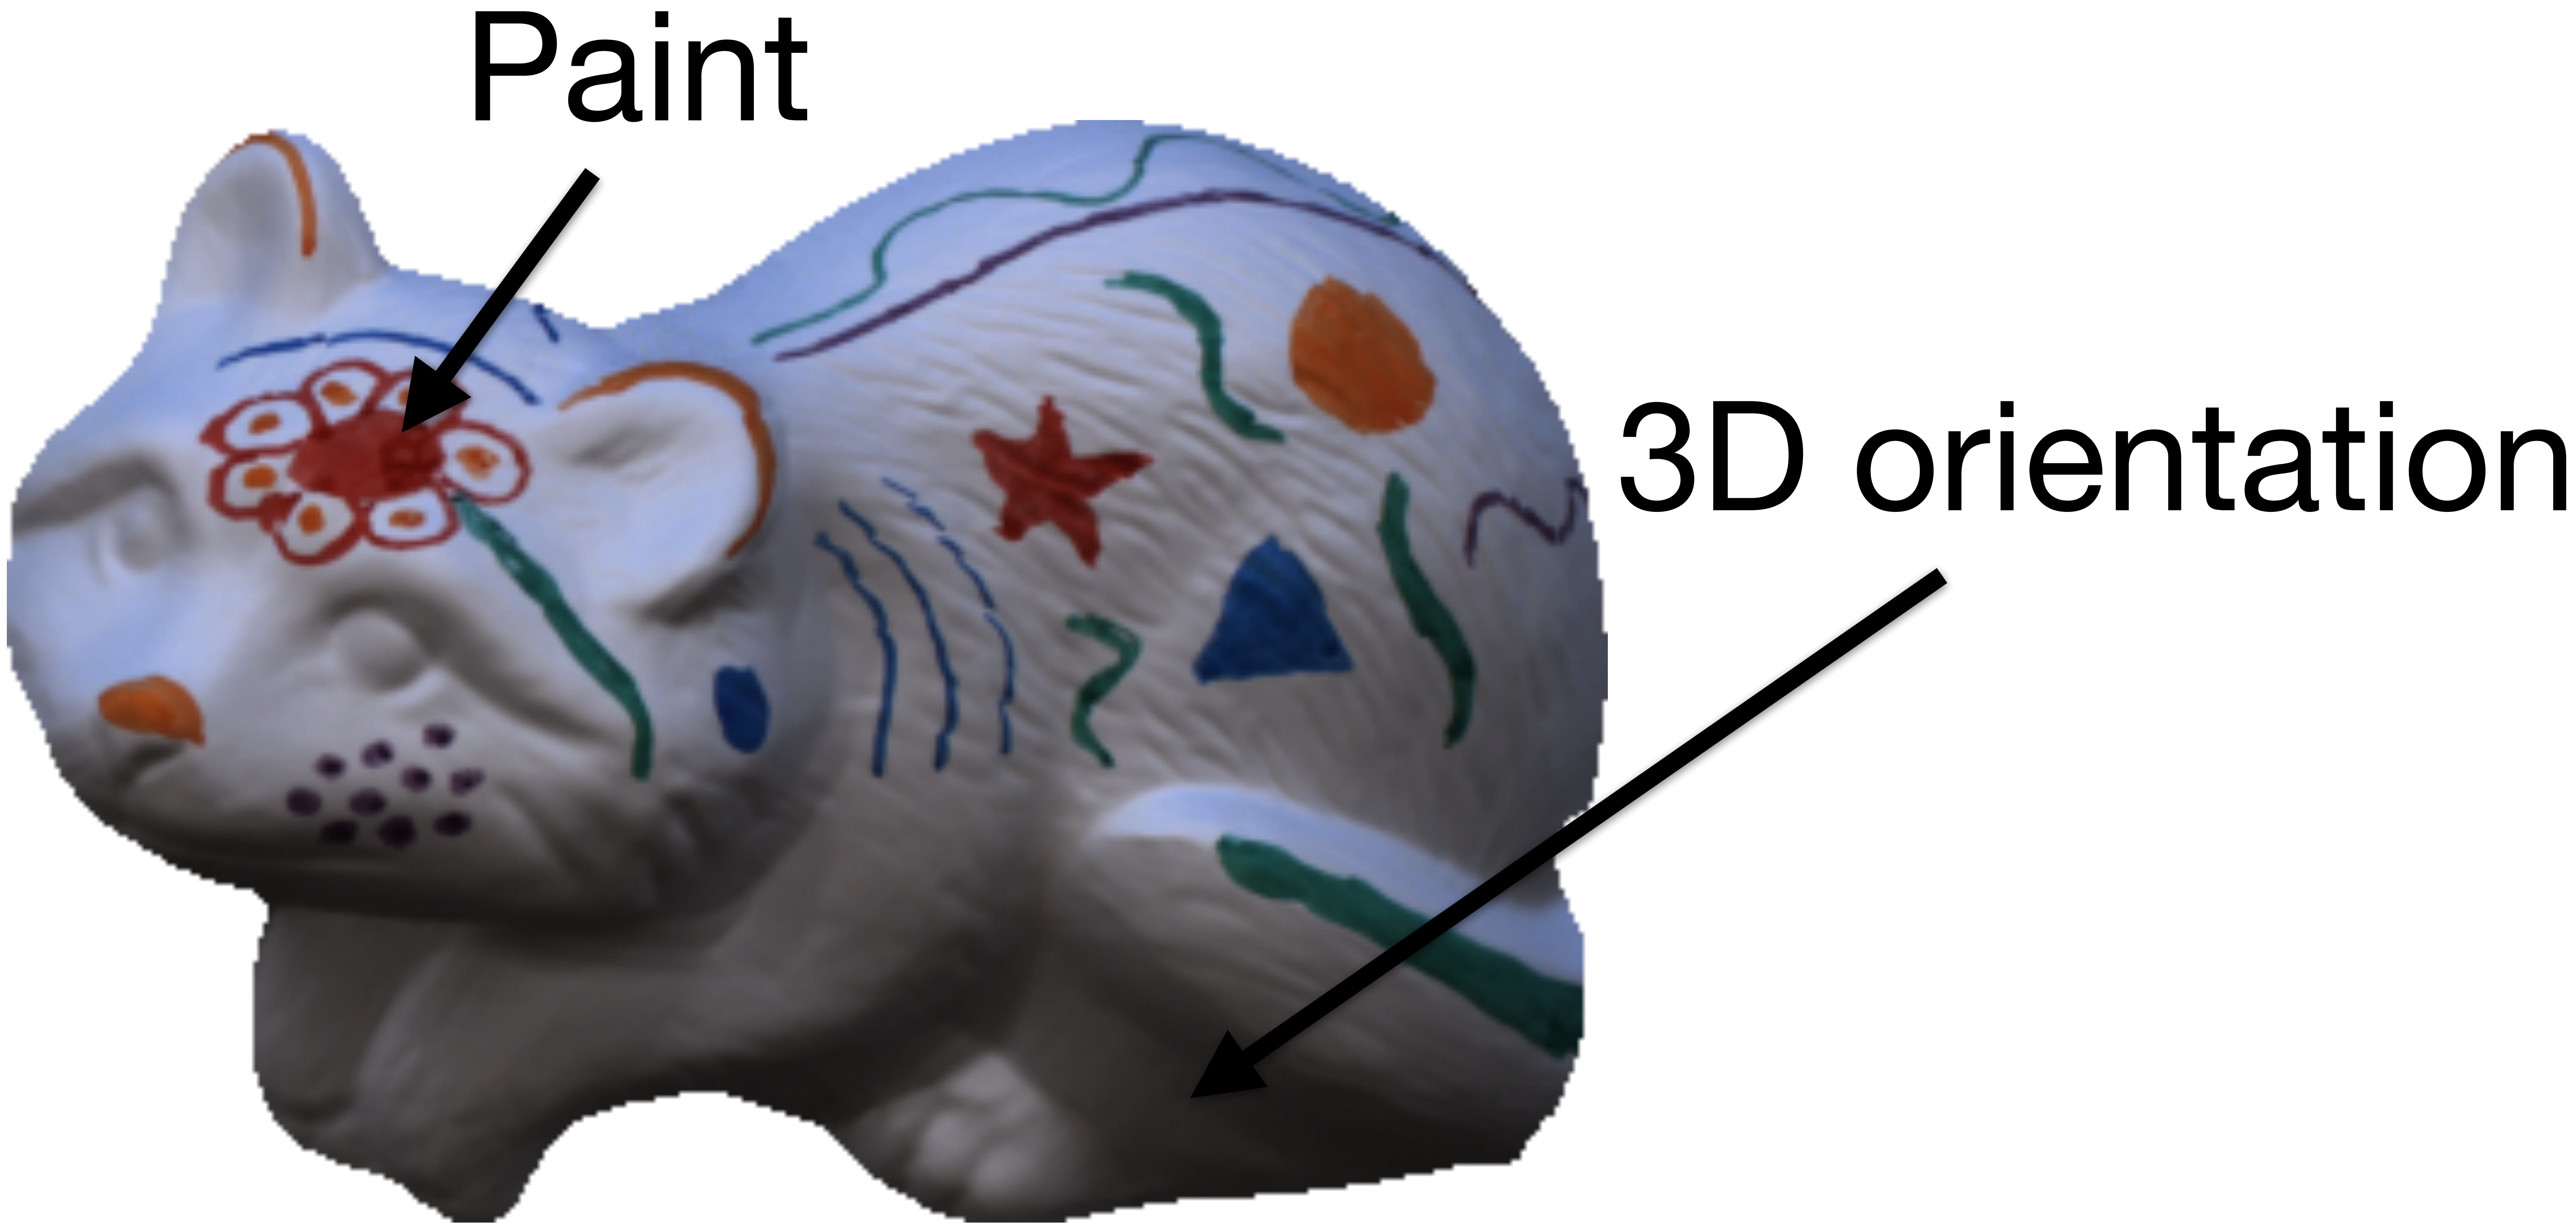

The Workshop Metaphor

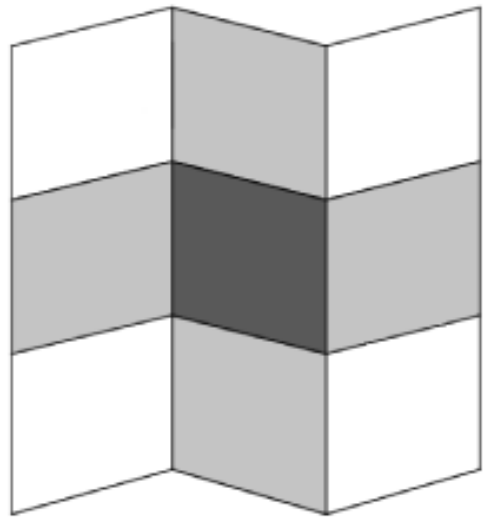

(a) an image

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Figure source: J. Barron

The Workshop Metaphor

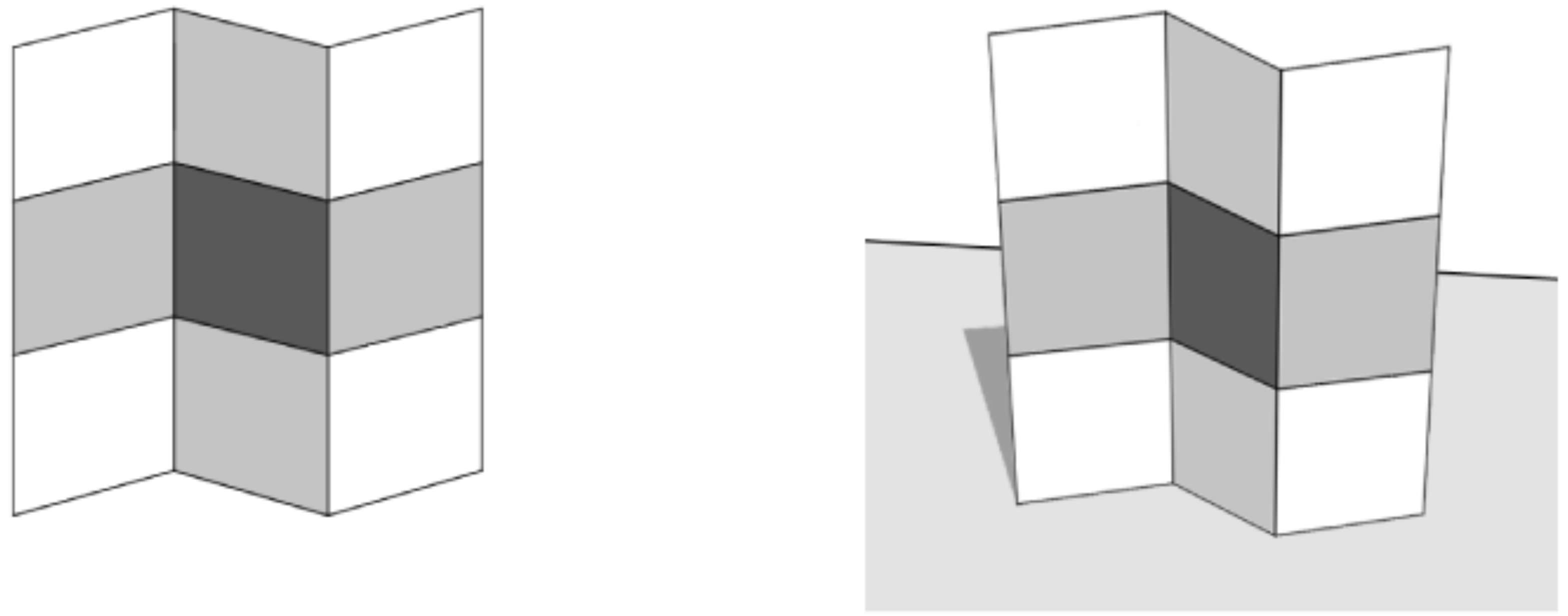

(a) an image (b) a likely explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Figure source: J. Barron

The Workshop Metaphor

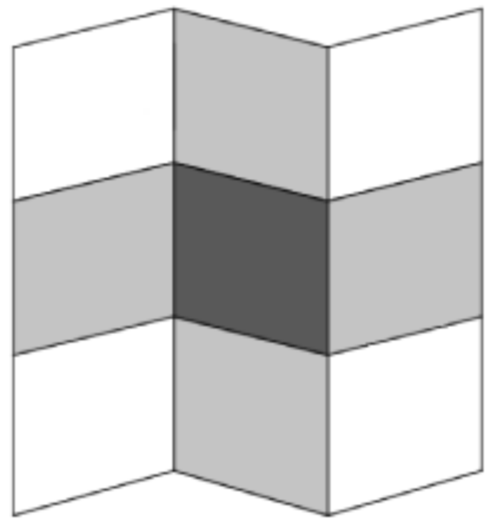

(a) an image

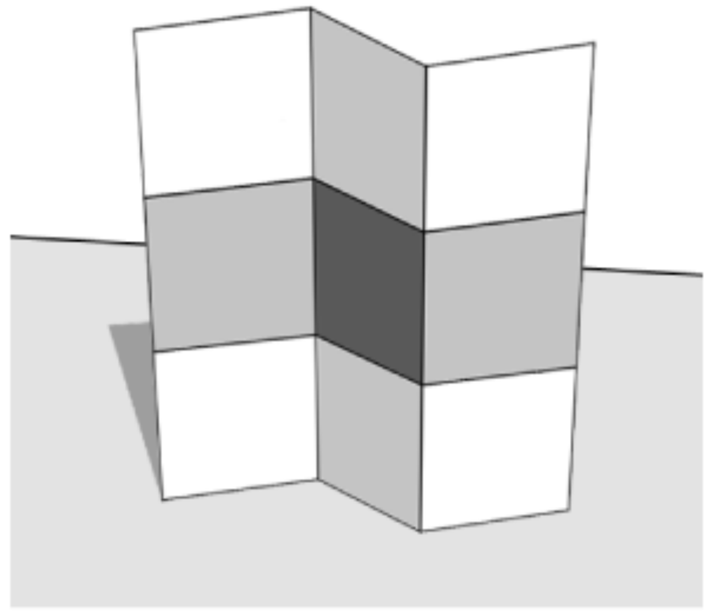

(b) a likely explanation

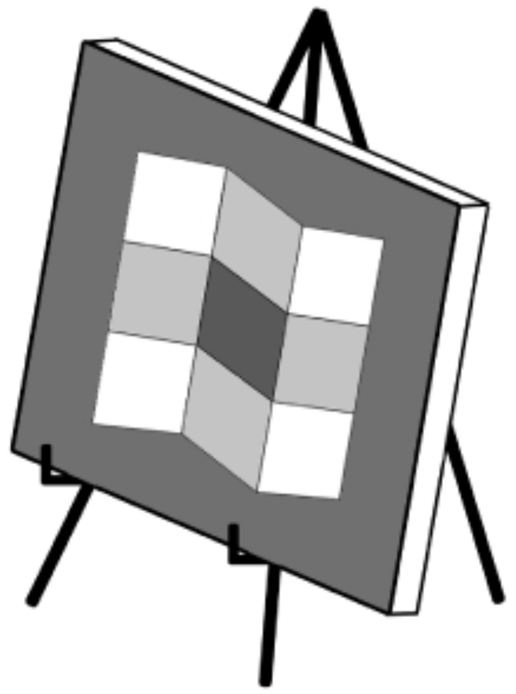

(c) painter's explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Figure source: J. Barron

The Workshop Metaphor

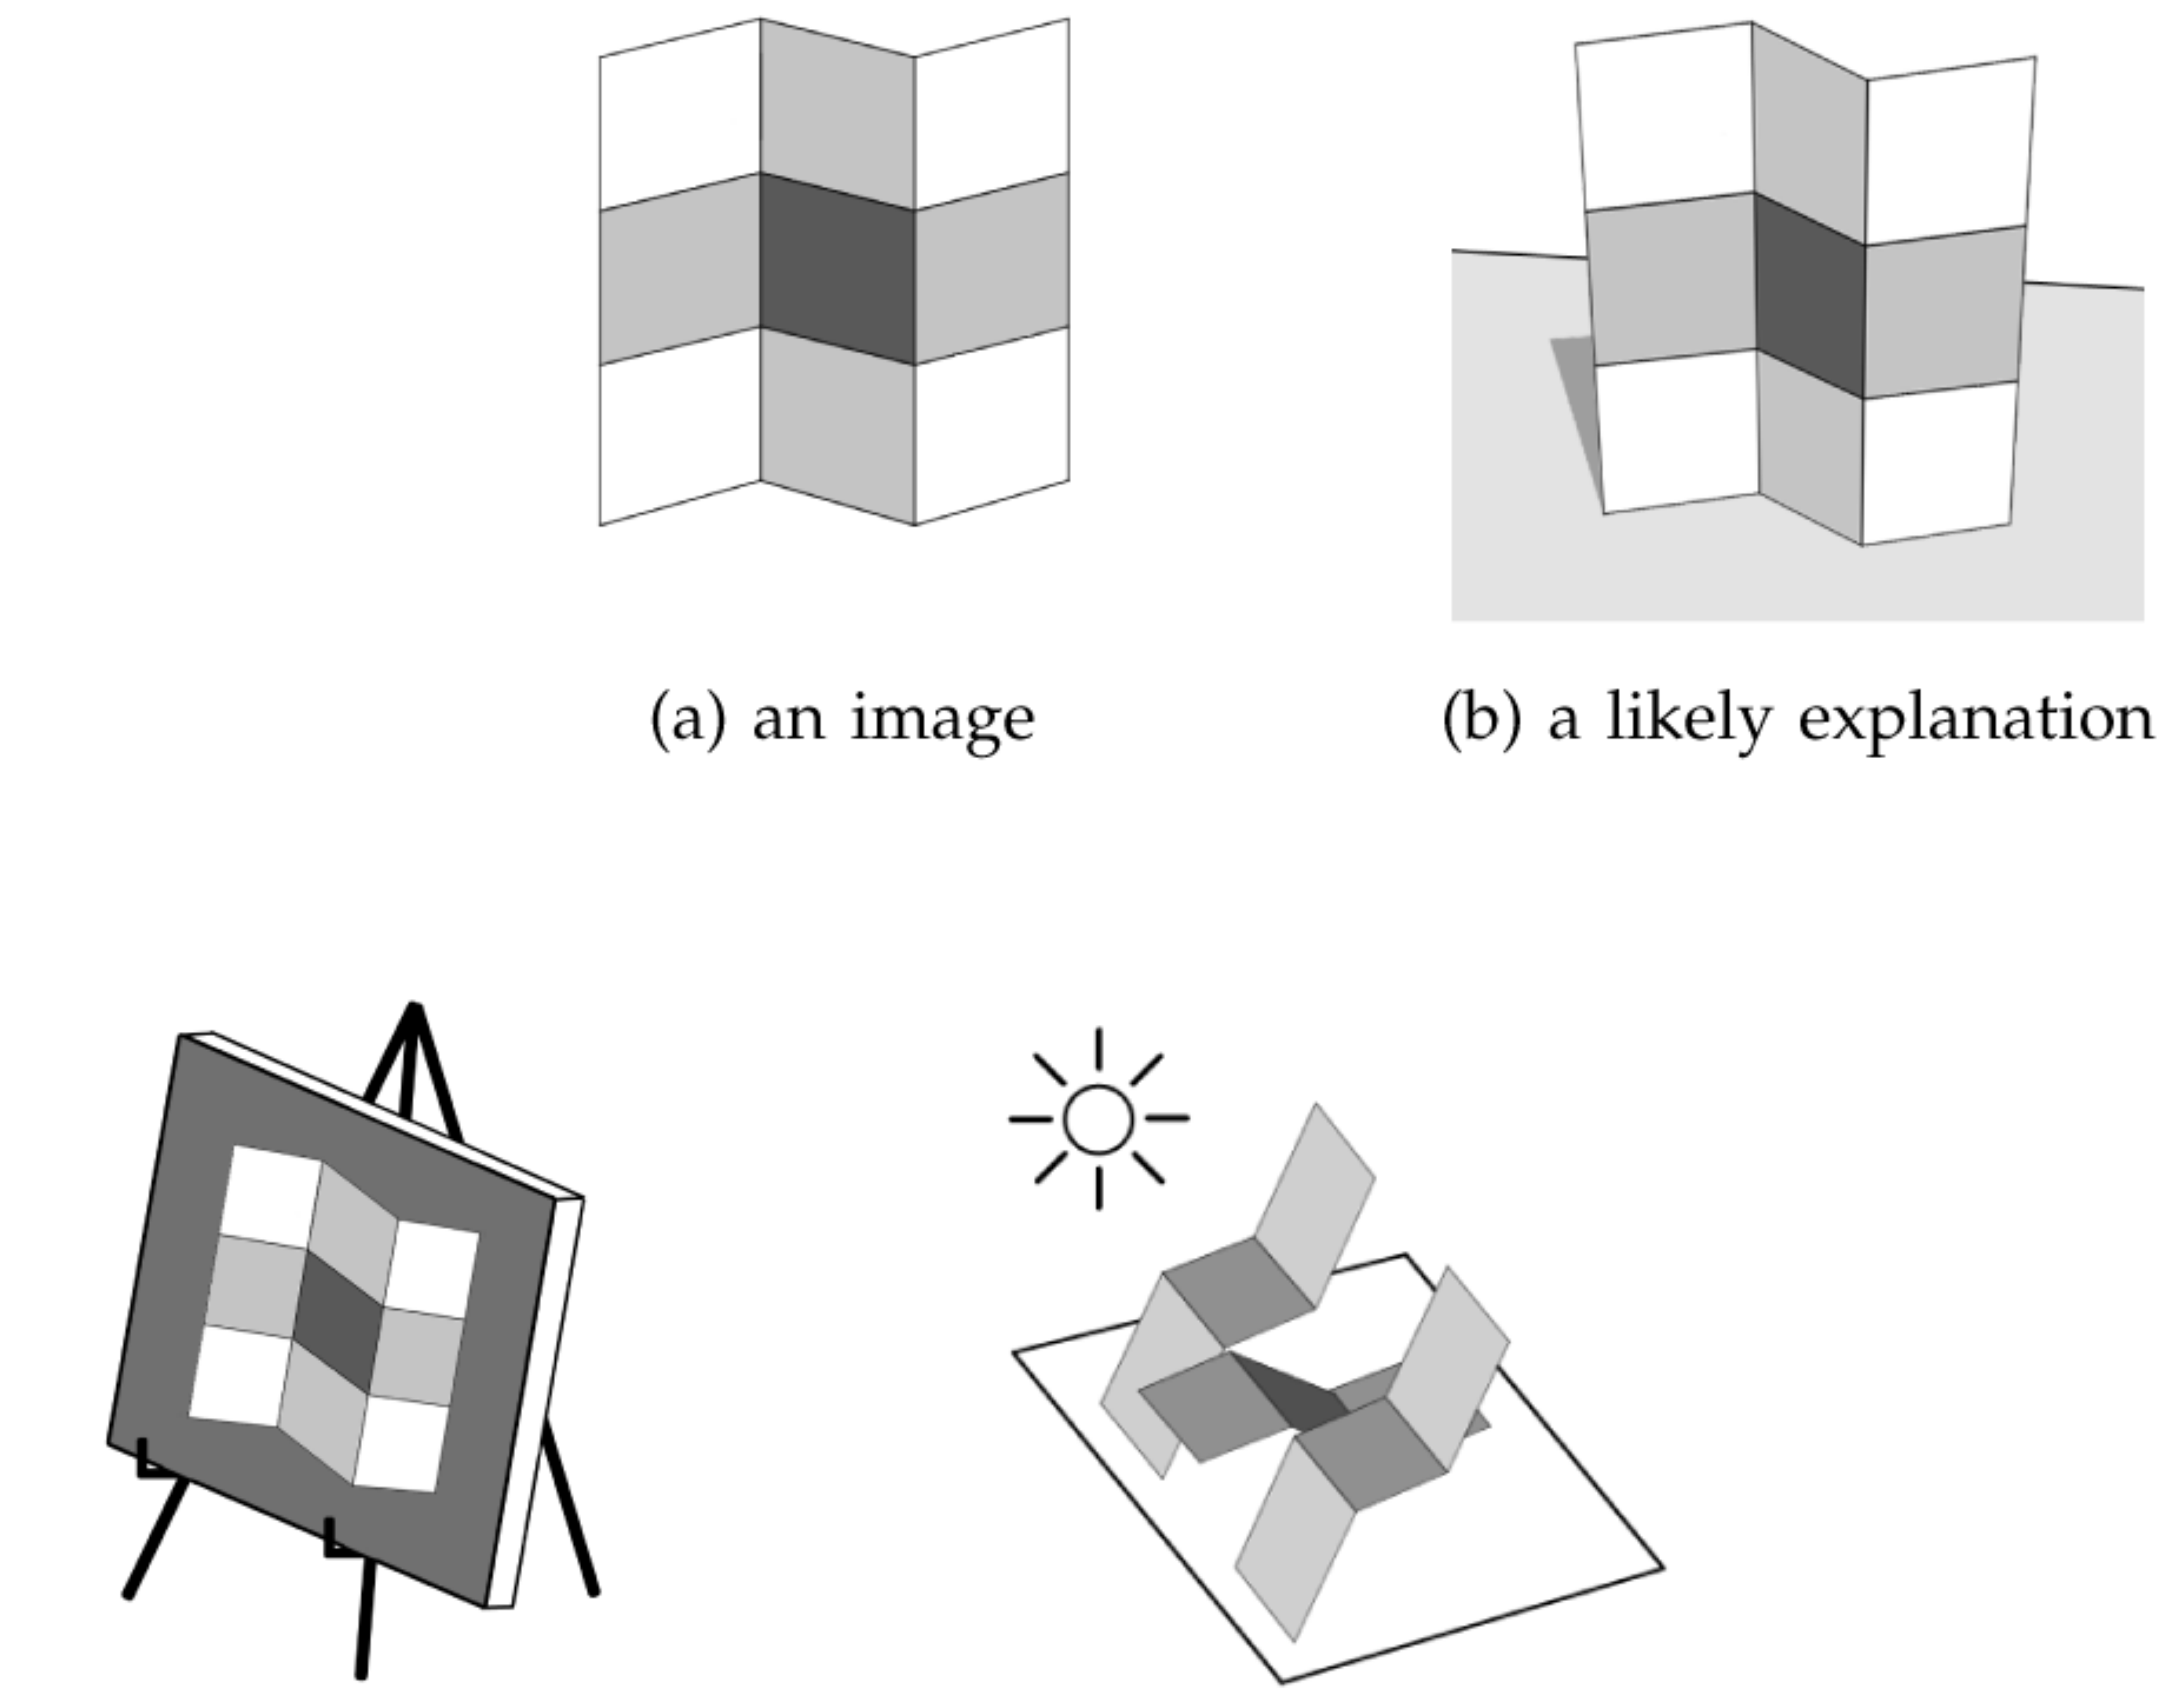

(a) an image (b) a likely explanation

(c) painter's explanation (d) sculptor's explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Figure source: J. Barron

The Workshop Metaphor

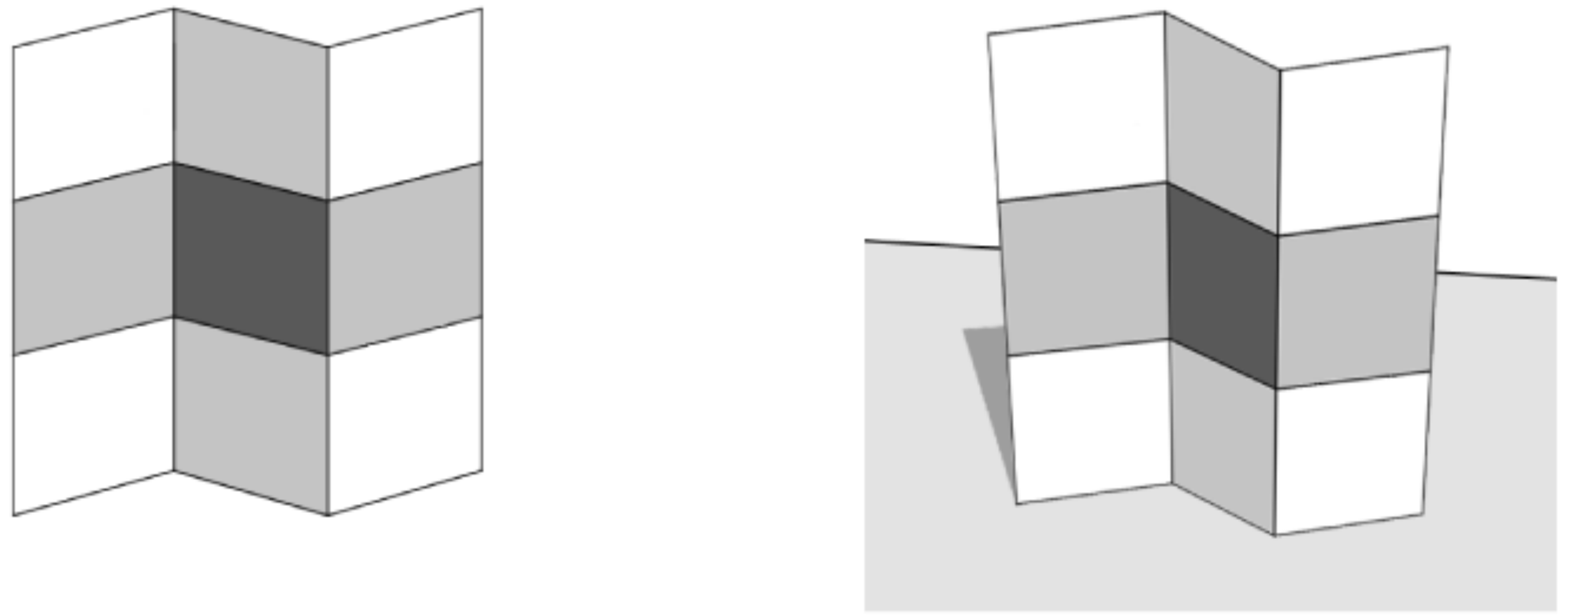

(a) an image (b) a likely explanation

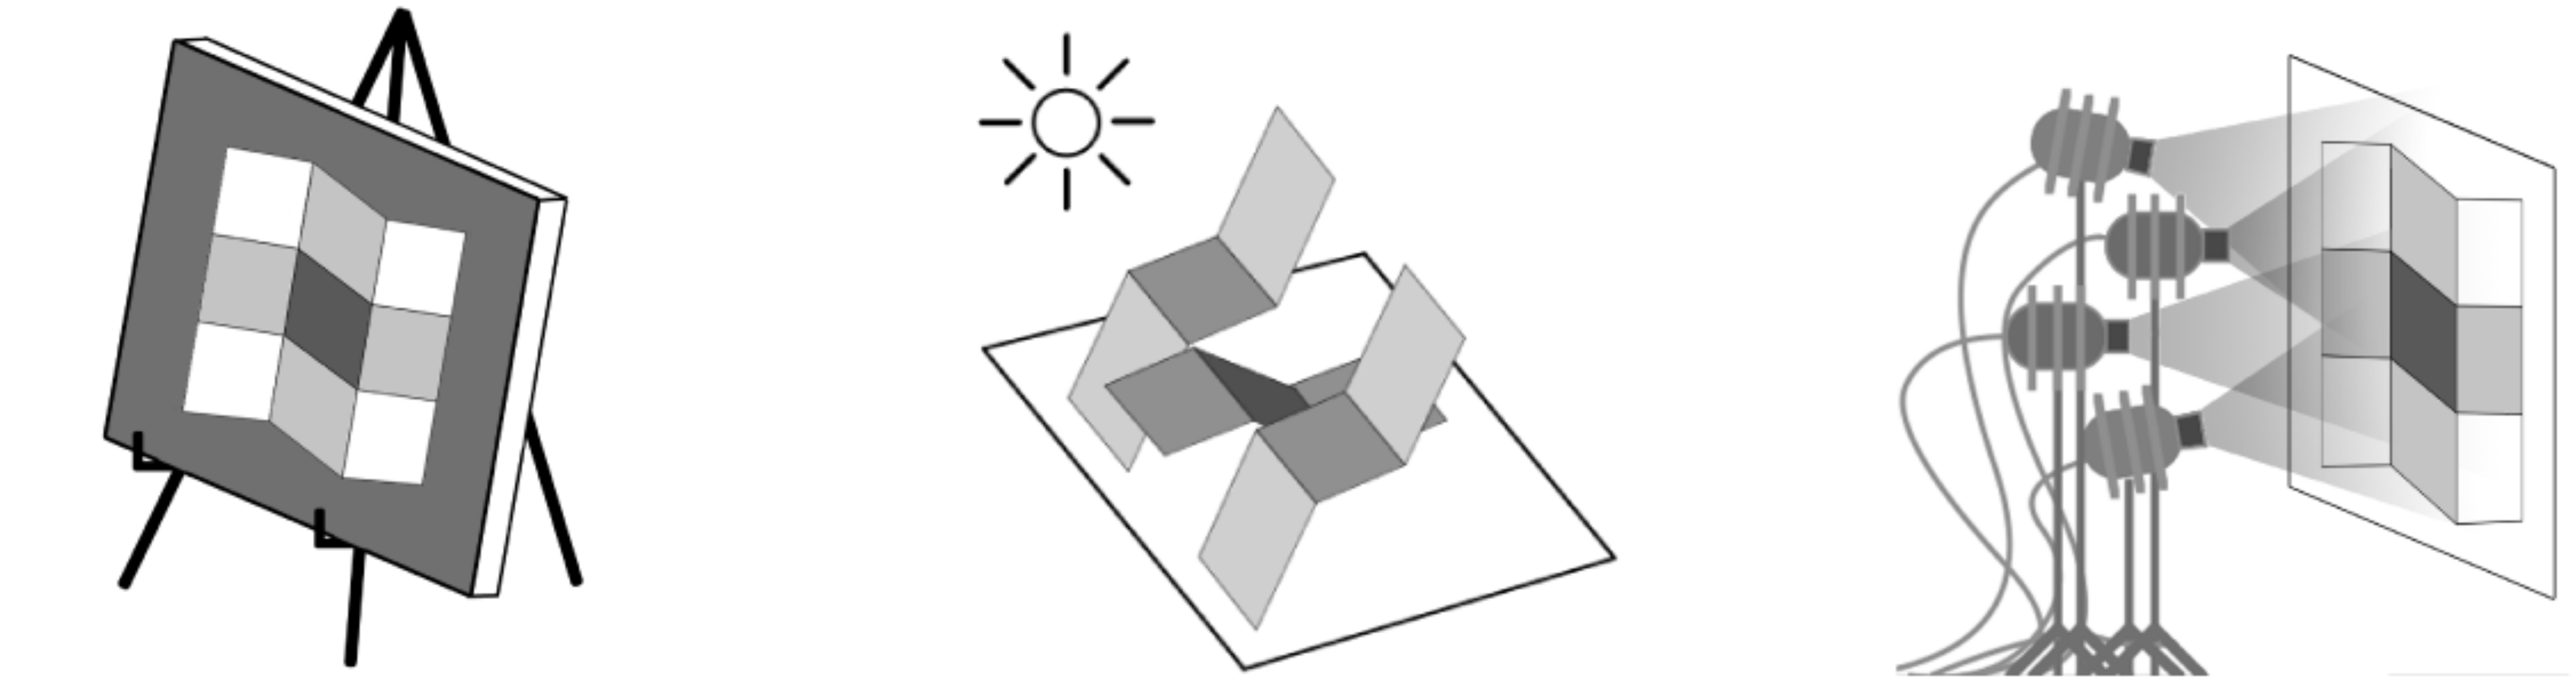

(c) painter's explanation (d) sculptor's explanation (e) gaffer's explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Figure source: J. Barron

Today

- **Shape from shading**
- Intrinsic image decomposition
- Color perception

Shape perception

Interaction of light and surfaces

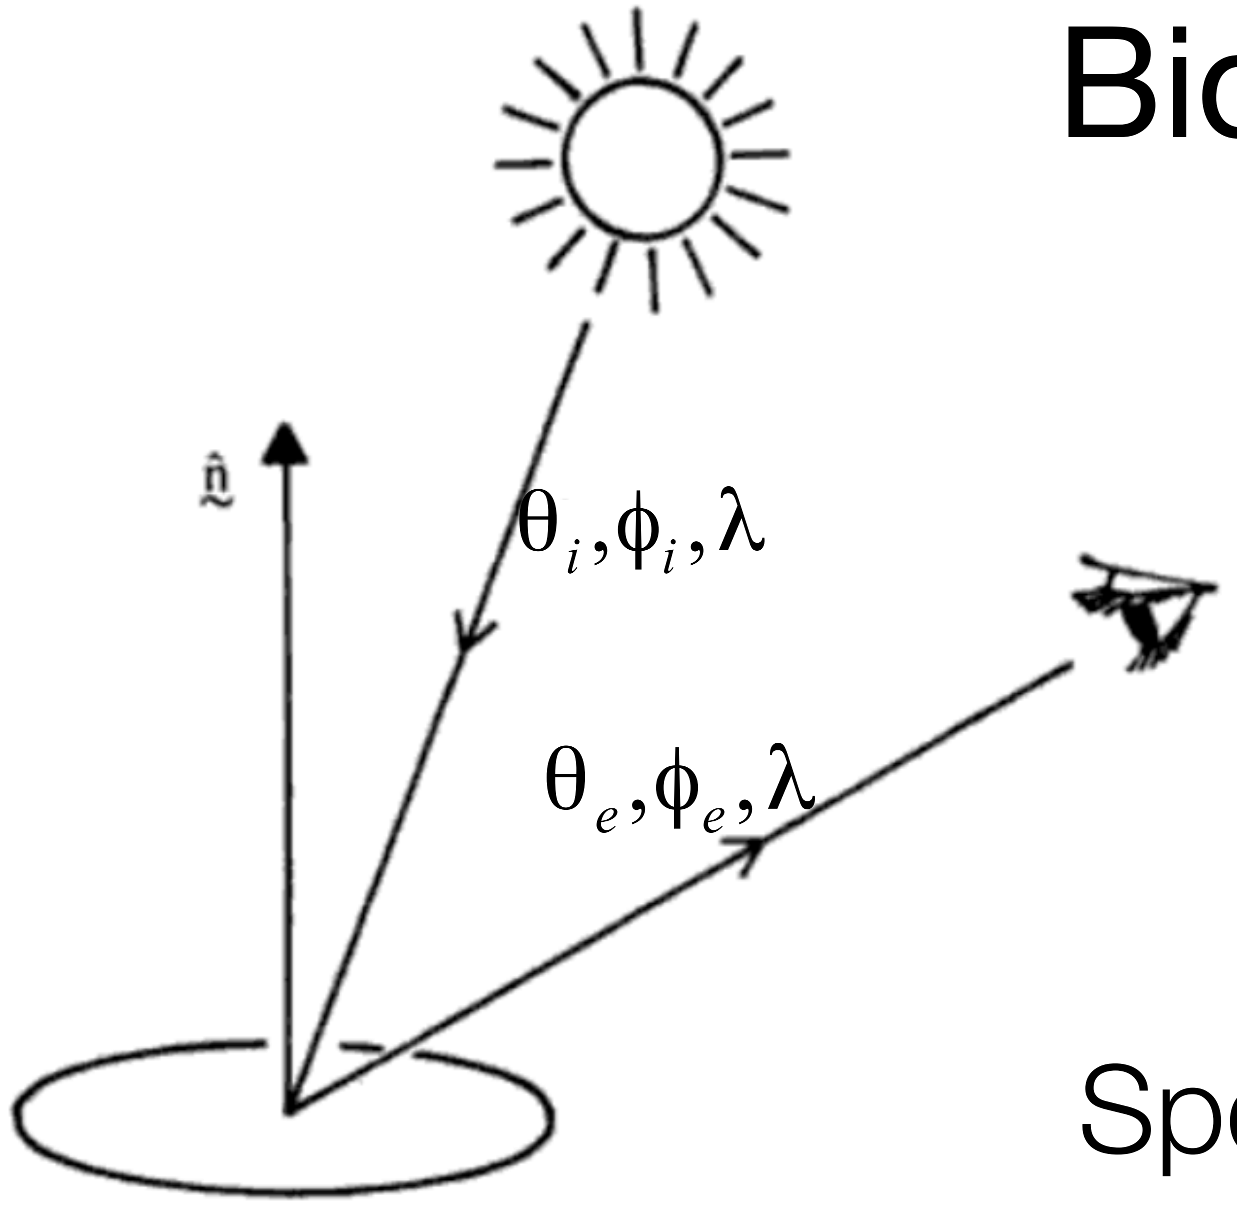

[Horn, 1986]

Bidirectional reflectance distribution function (BRDF)

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

Spectral irradiance: incident power per unit area, per unit wavelength.

Source: W. Freeman

Effect of BRDF on sphere rendering

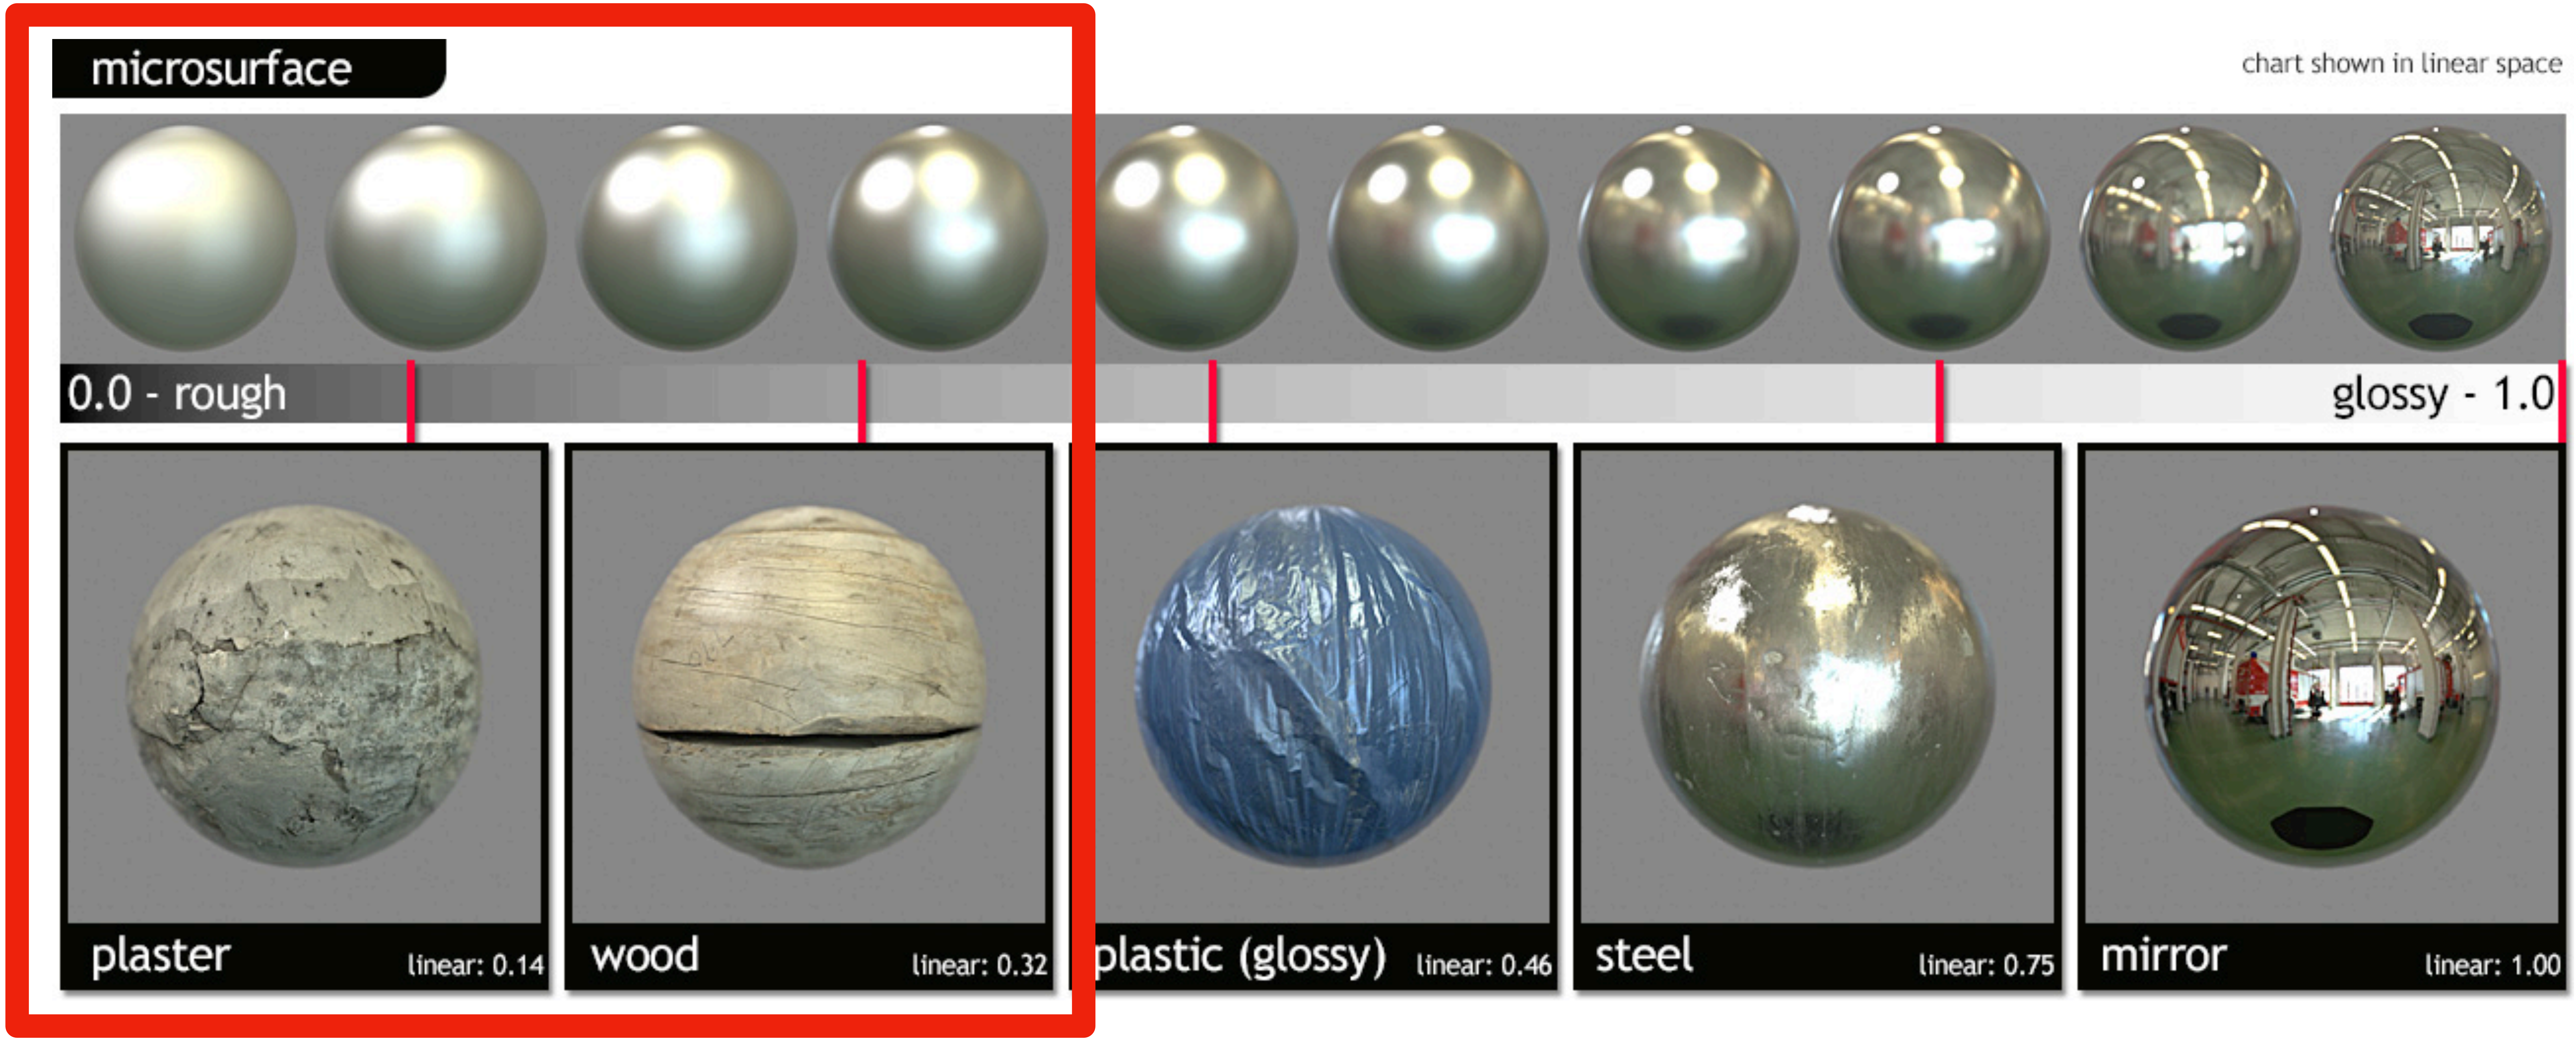

Diffuse reflection

https://marmoset.co/posts/physically-based-rendering-and-you-can-too/

Source: W. Freeman

For now, ignore specular reflection

Source: Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra. Changes by N. Snavely

And refraction…

Source: Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra. Changes by N. Snavely

And interreflections…

Source: Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra. Changes by N. Snavely

Diffuse reflection

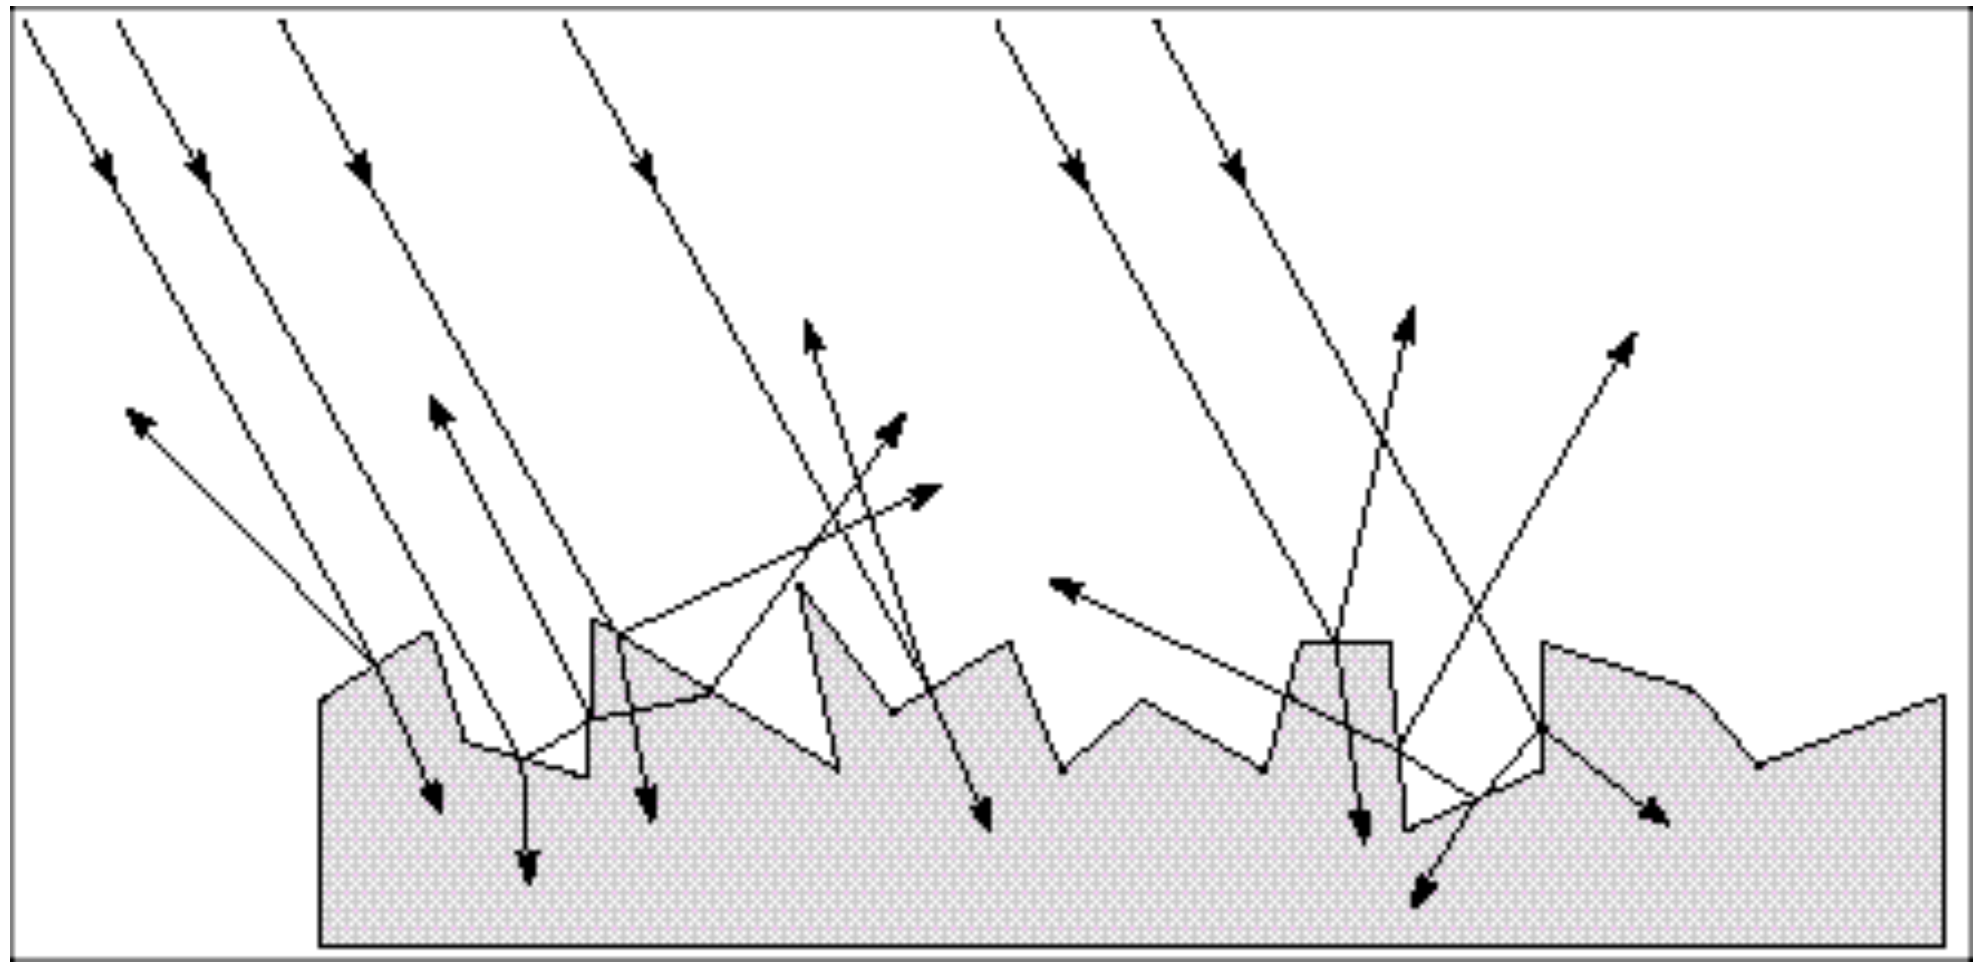

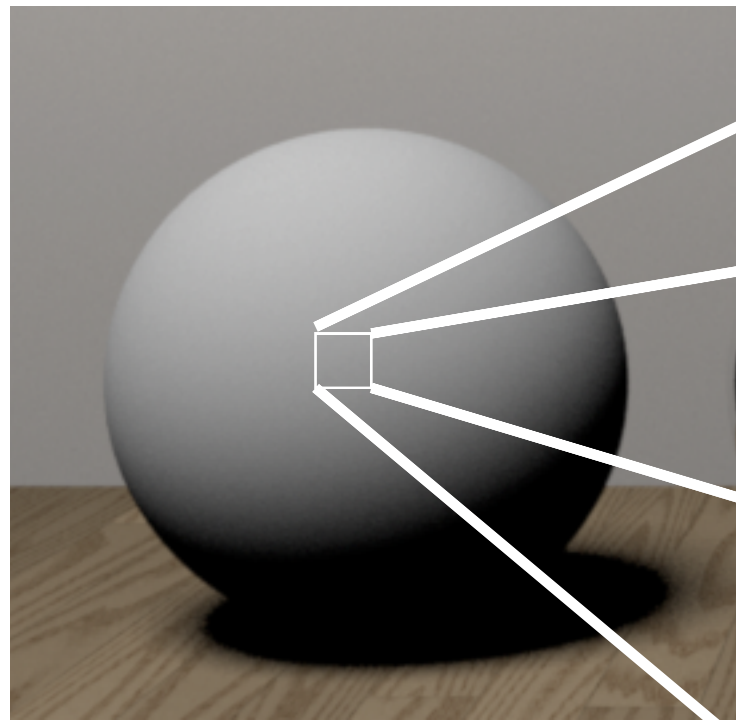

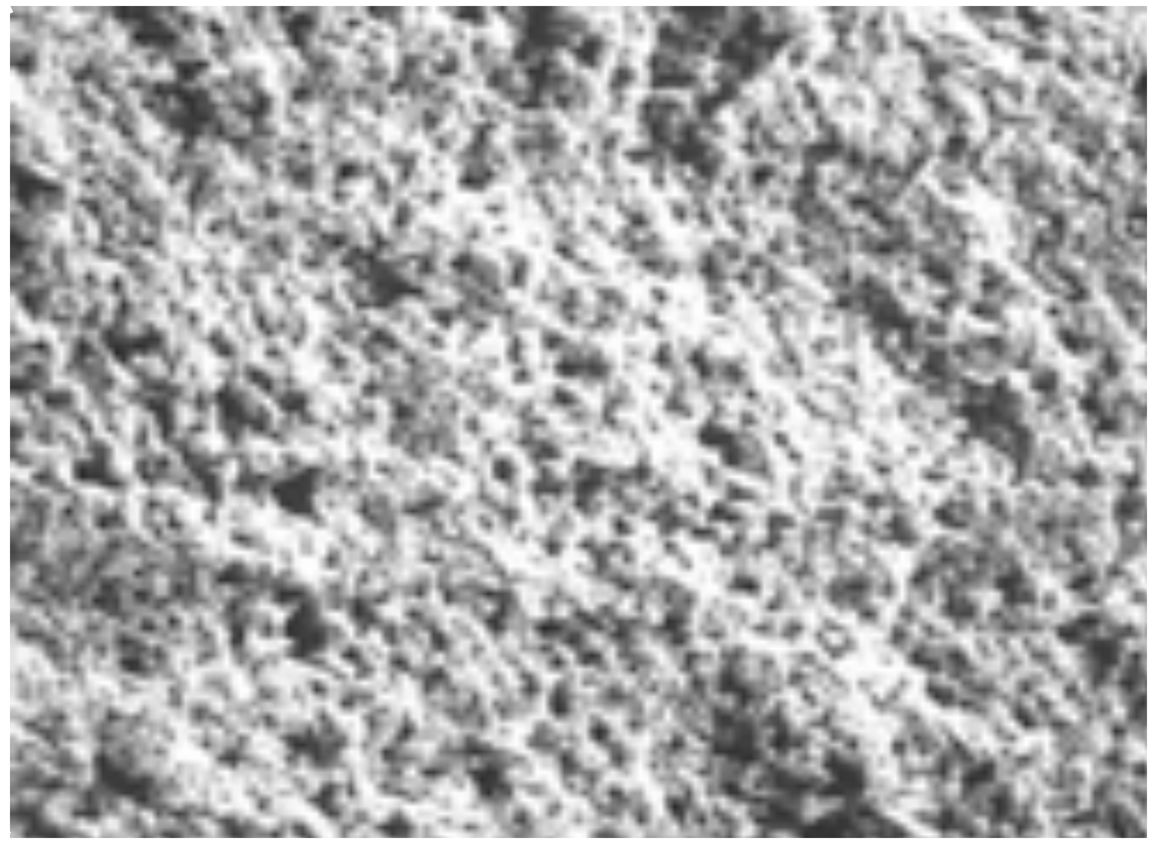

Diffuse reflection

- Dull, matte surfaces like chalk or latex paint
- Microfacets scatter incoming light randomly
- Effect is that light is reflected equally in all directions

Source: S. Lazebnik and K. Bala

Directional lighting

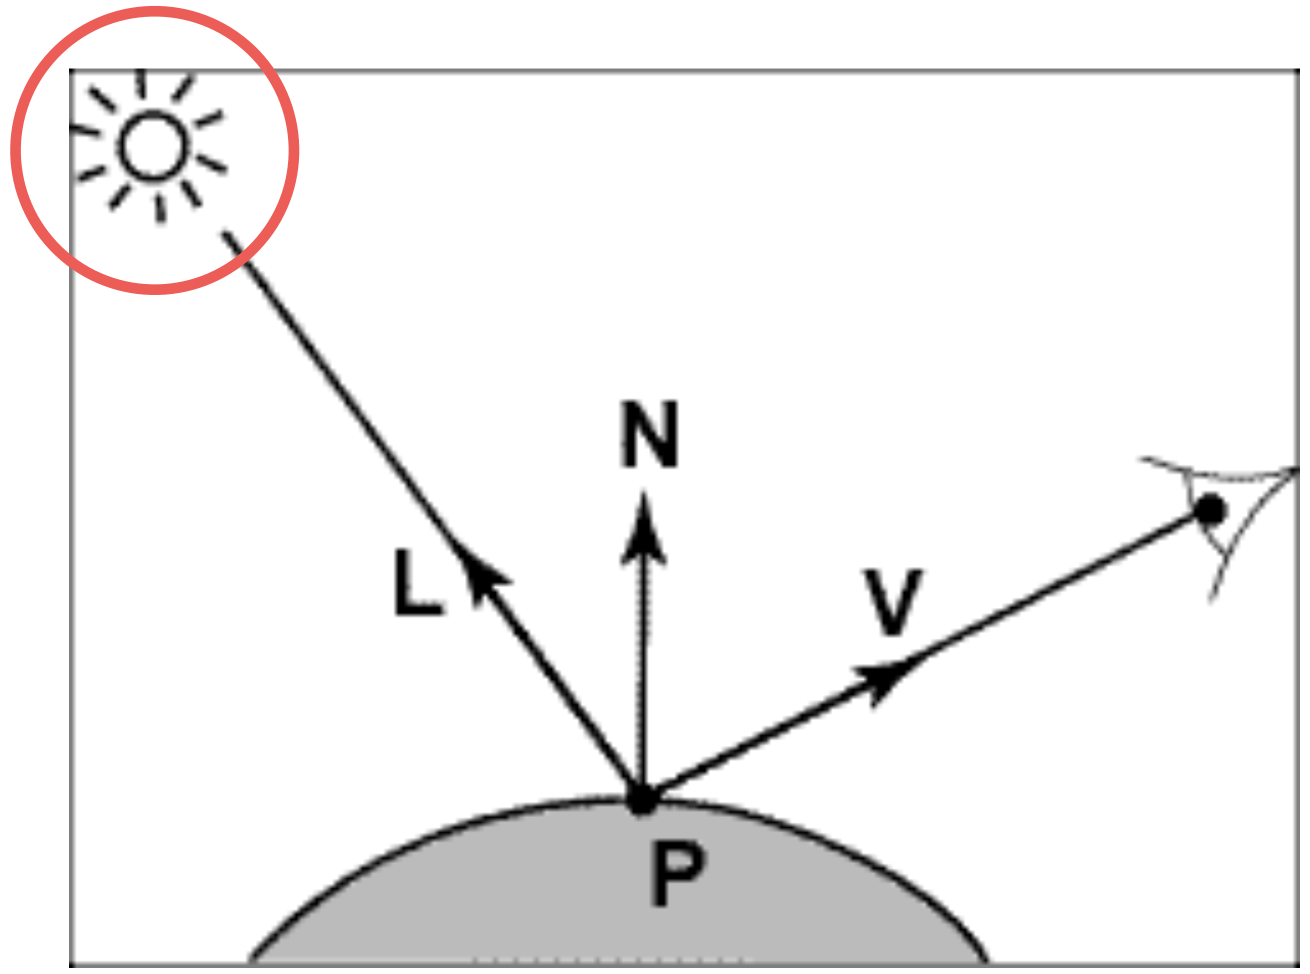

- All rays are parallel
- Equivalent to an infinitely distant point source

Source: N. Snavely

Diffuse reflection

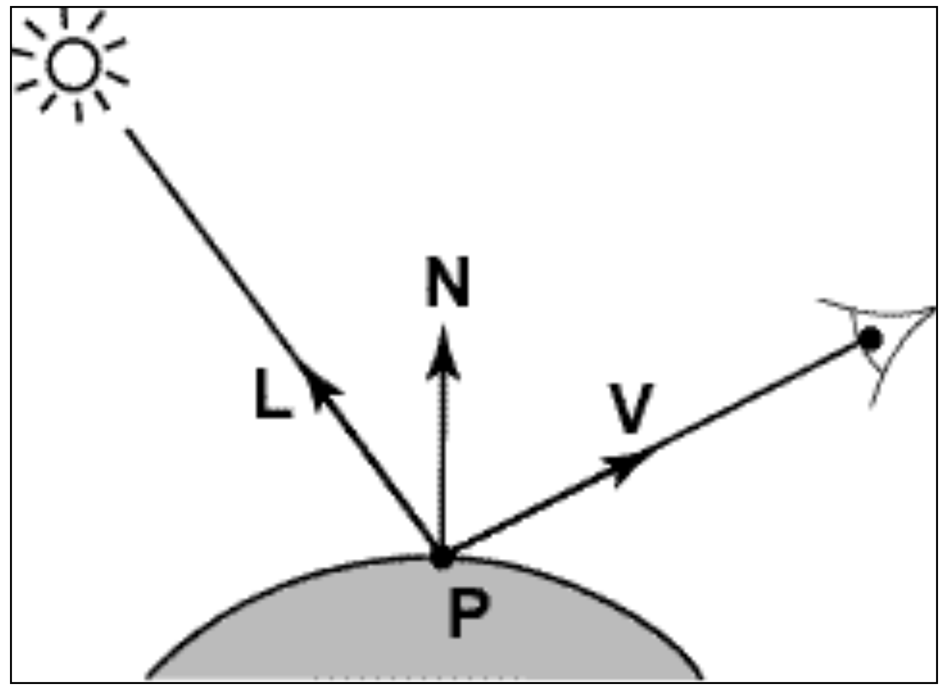

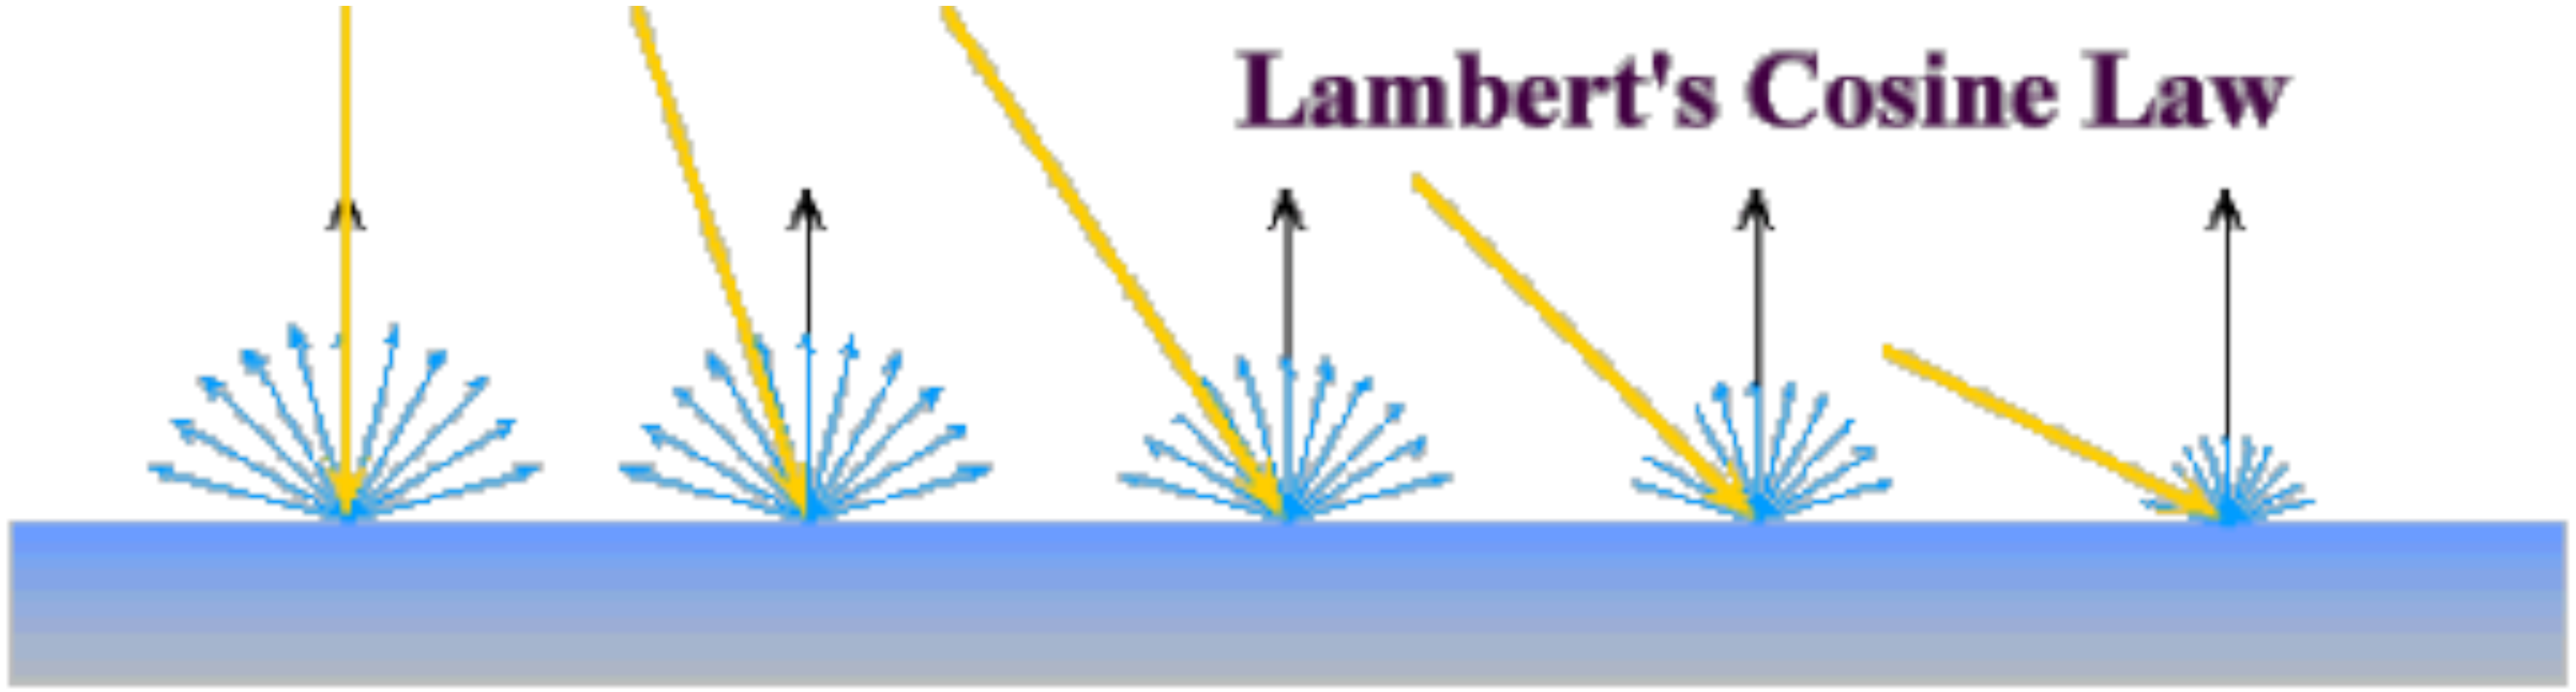

$$R_e = k_d \mathbf{N} \cdot \mathbf{L} R_i$$

image intensity of **P** → $I = k_d \mathbf{N} \cdot \mathbf{L}$

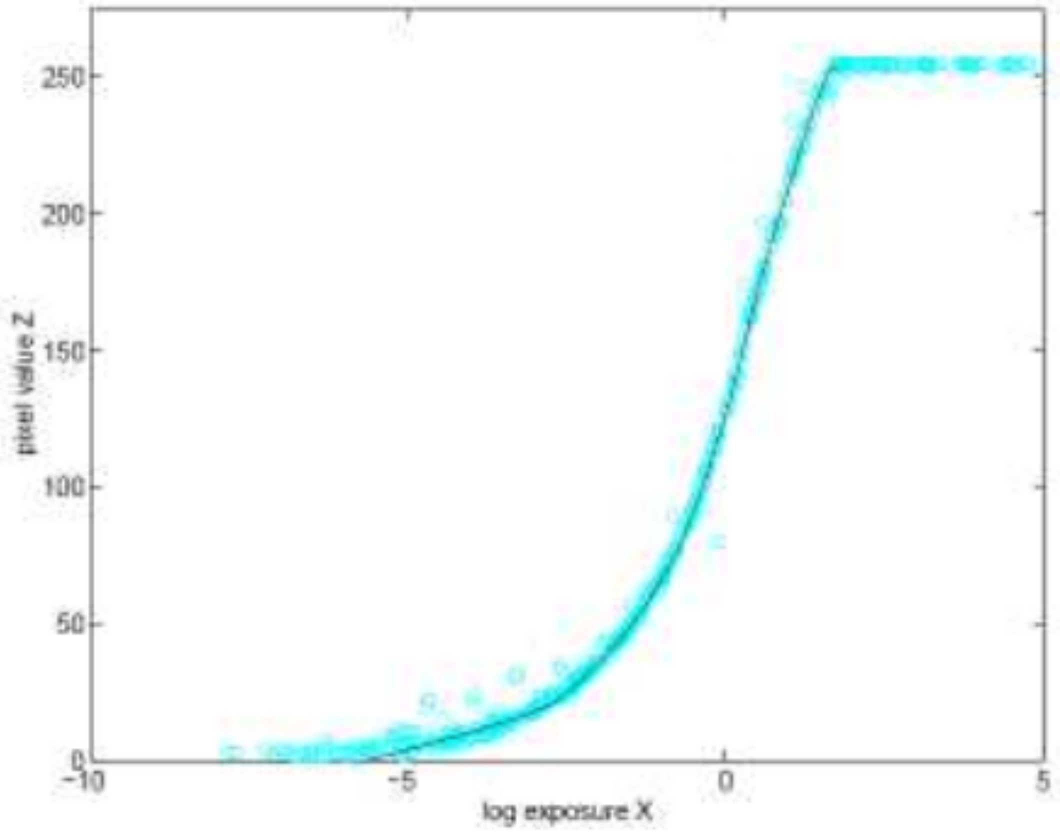

Source: [Debevec & Malik 1997]

Simplifying assumptions we'll often make:

- $I = R_e$: "camera response function" is the identity
  - In actual cameras it is a nonlinear function
  - Can always achieve this in practice by inverting it
- $R_i = 1$: light source intensity is 1
  - can achieve this by dividing each pixel in the image by R_i

Source: N. Snavely

Other BRDFs

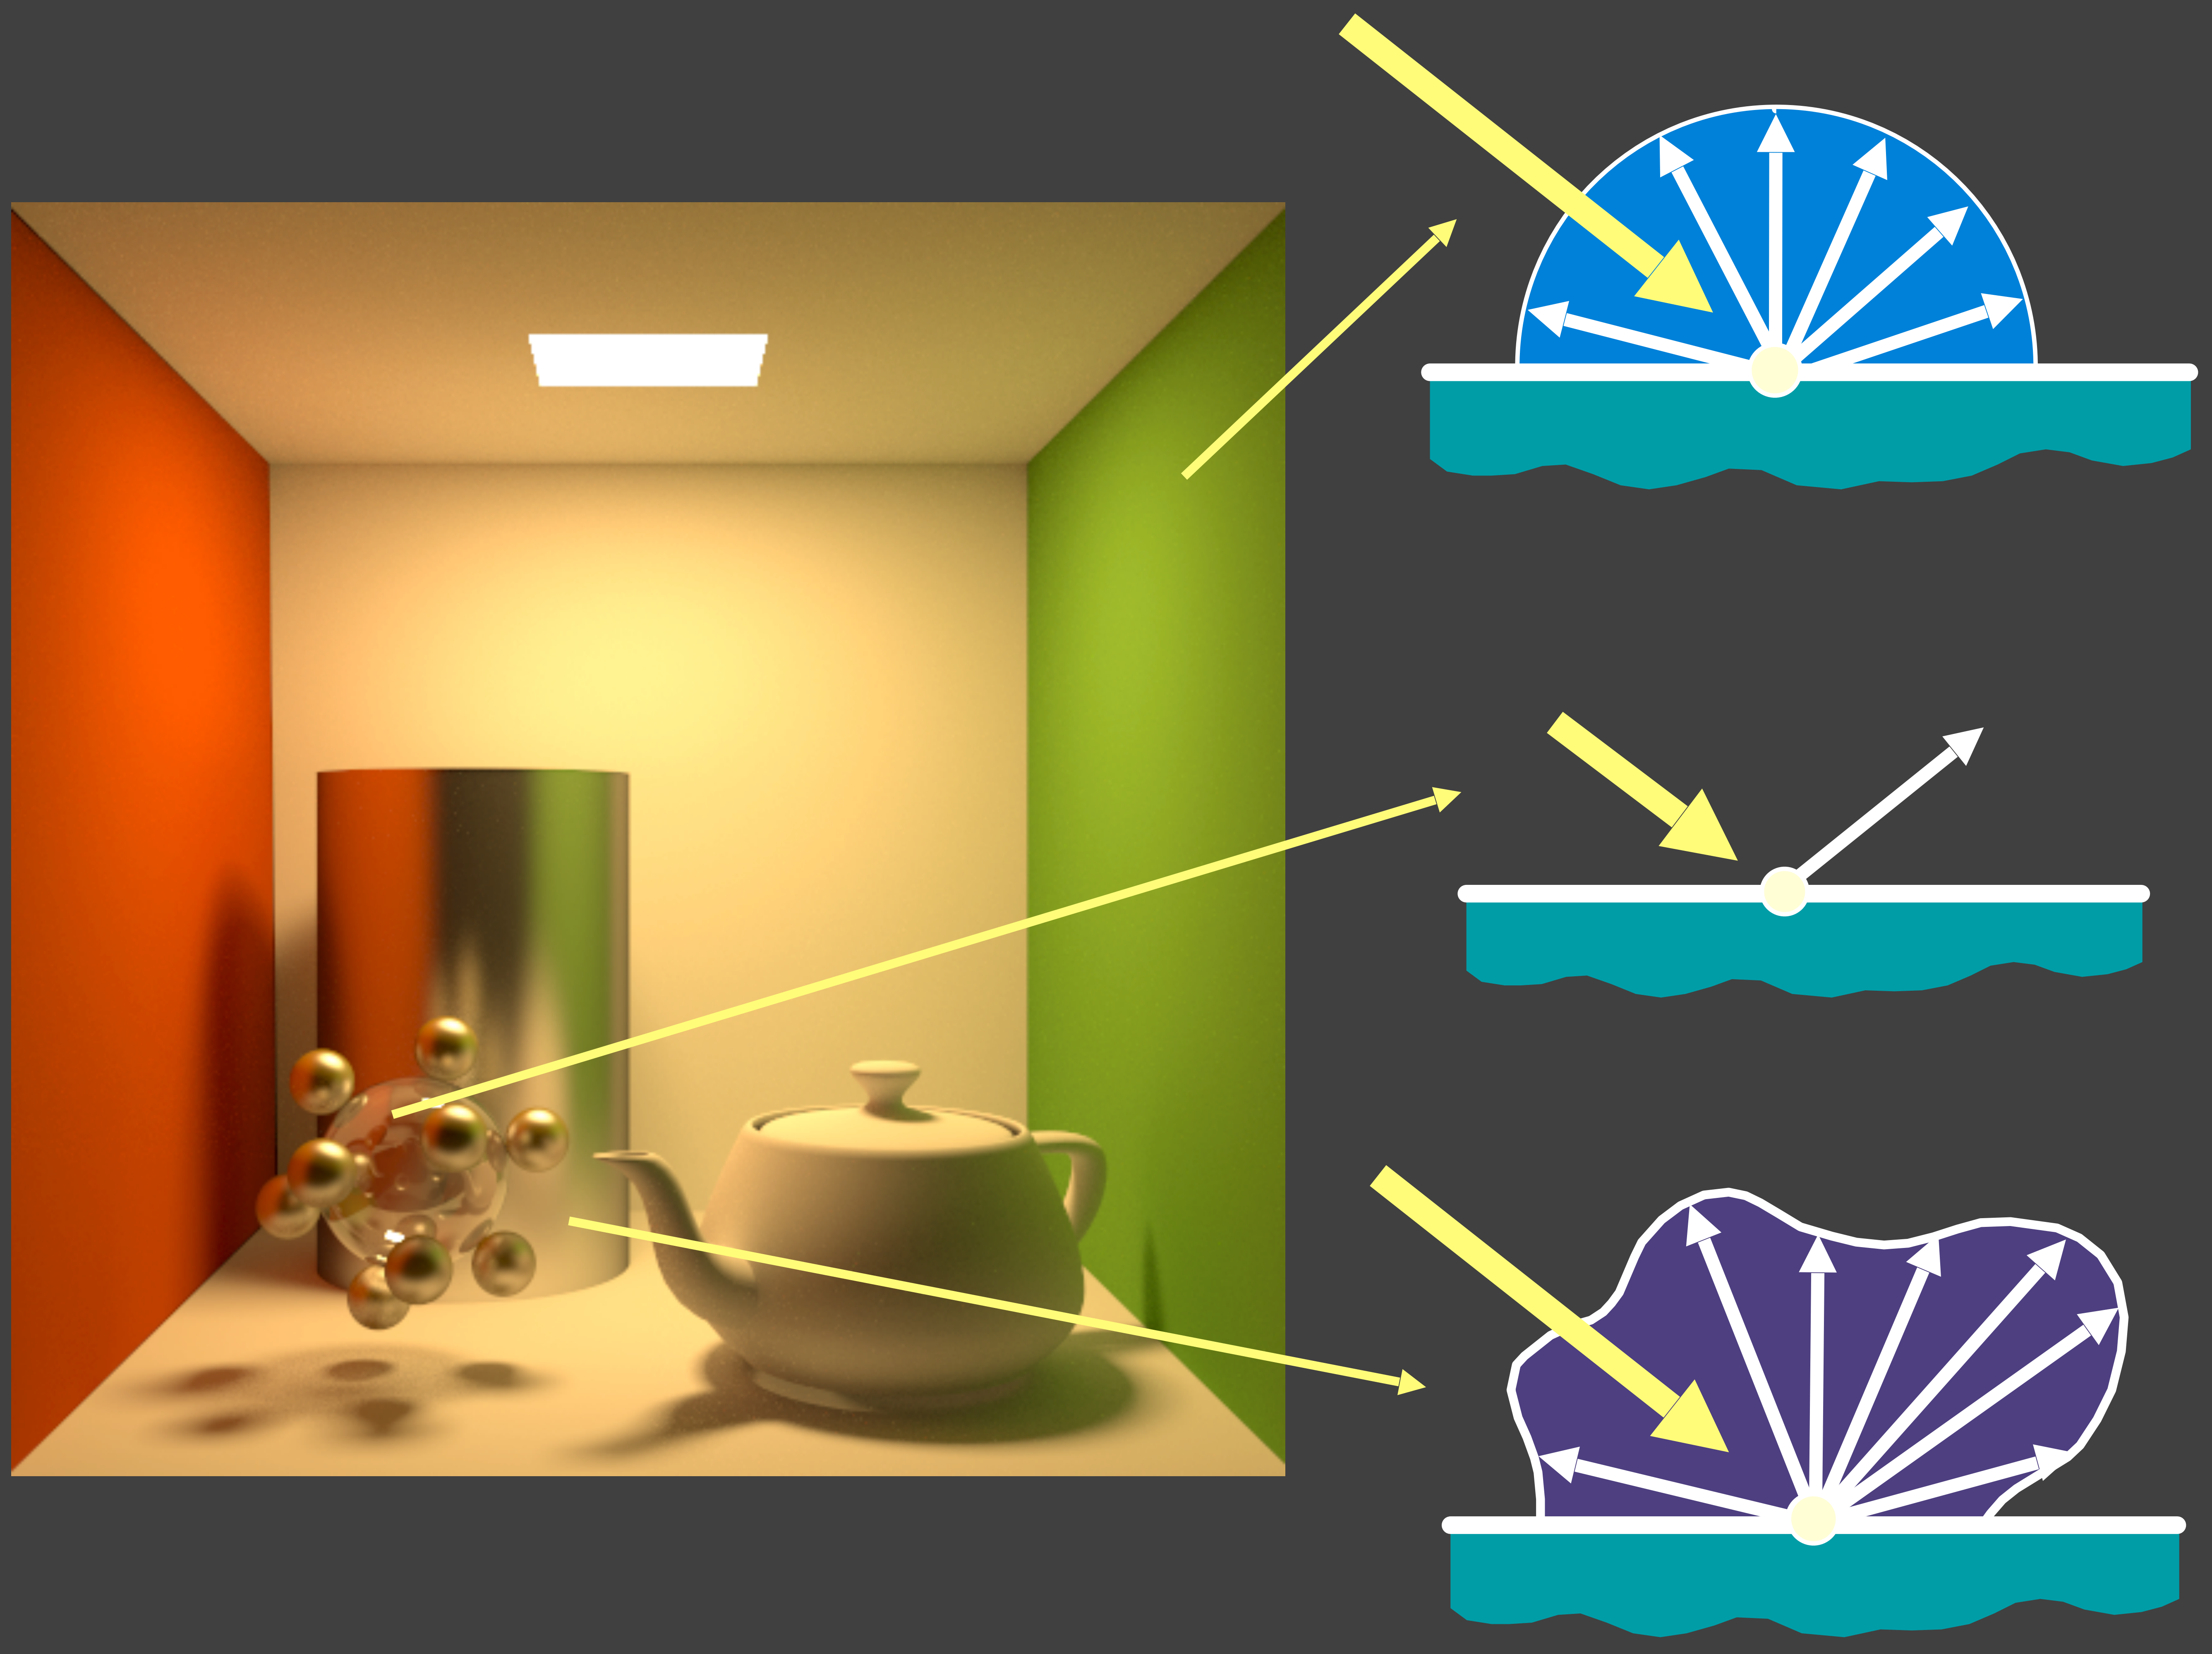

from Steve Marschner

Non-smooth-surfaced materials

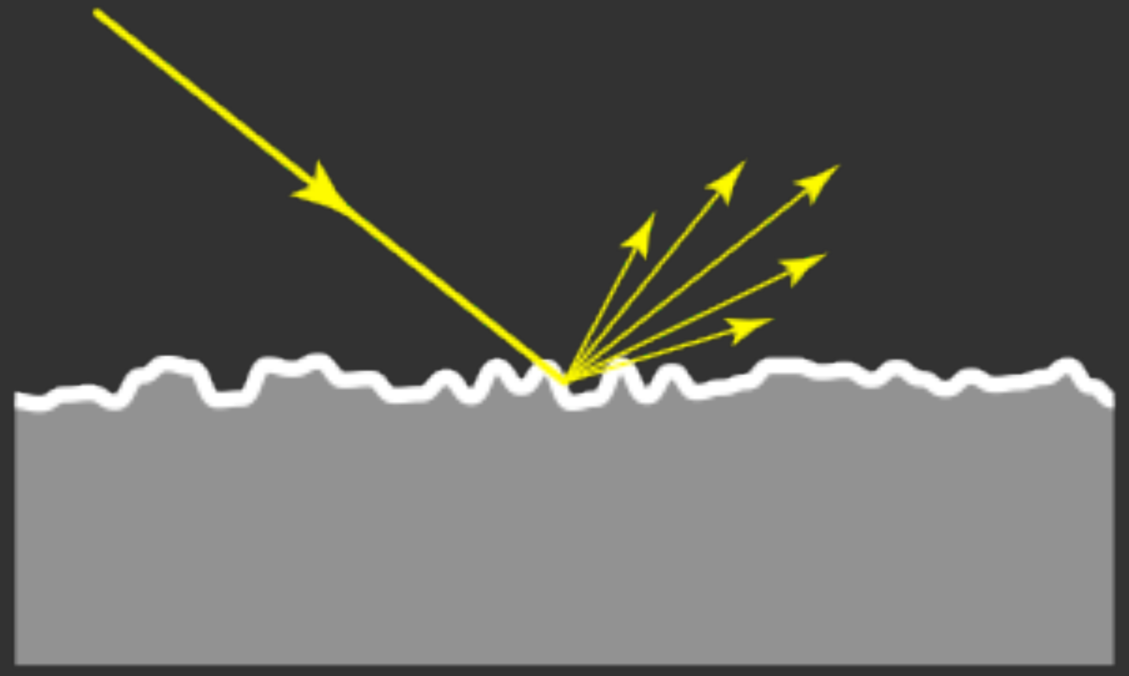

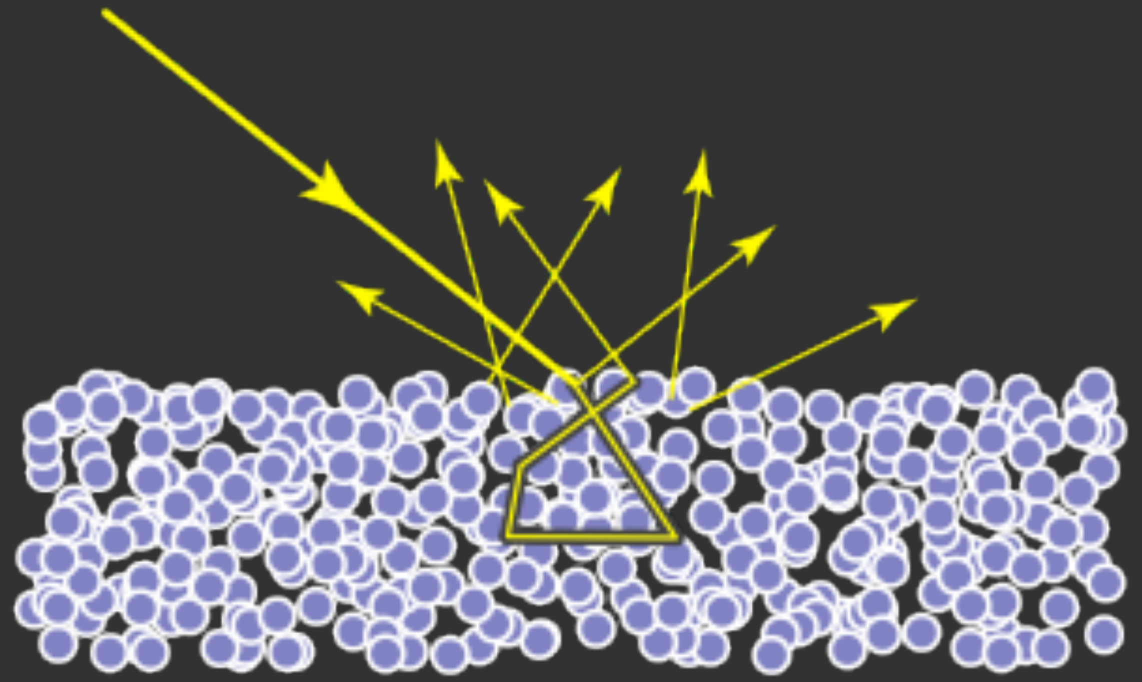

from Steve Marschner

Shape from shading

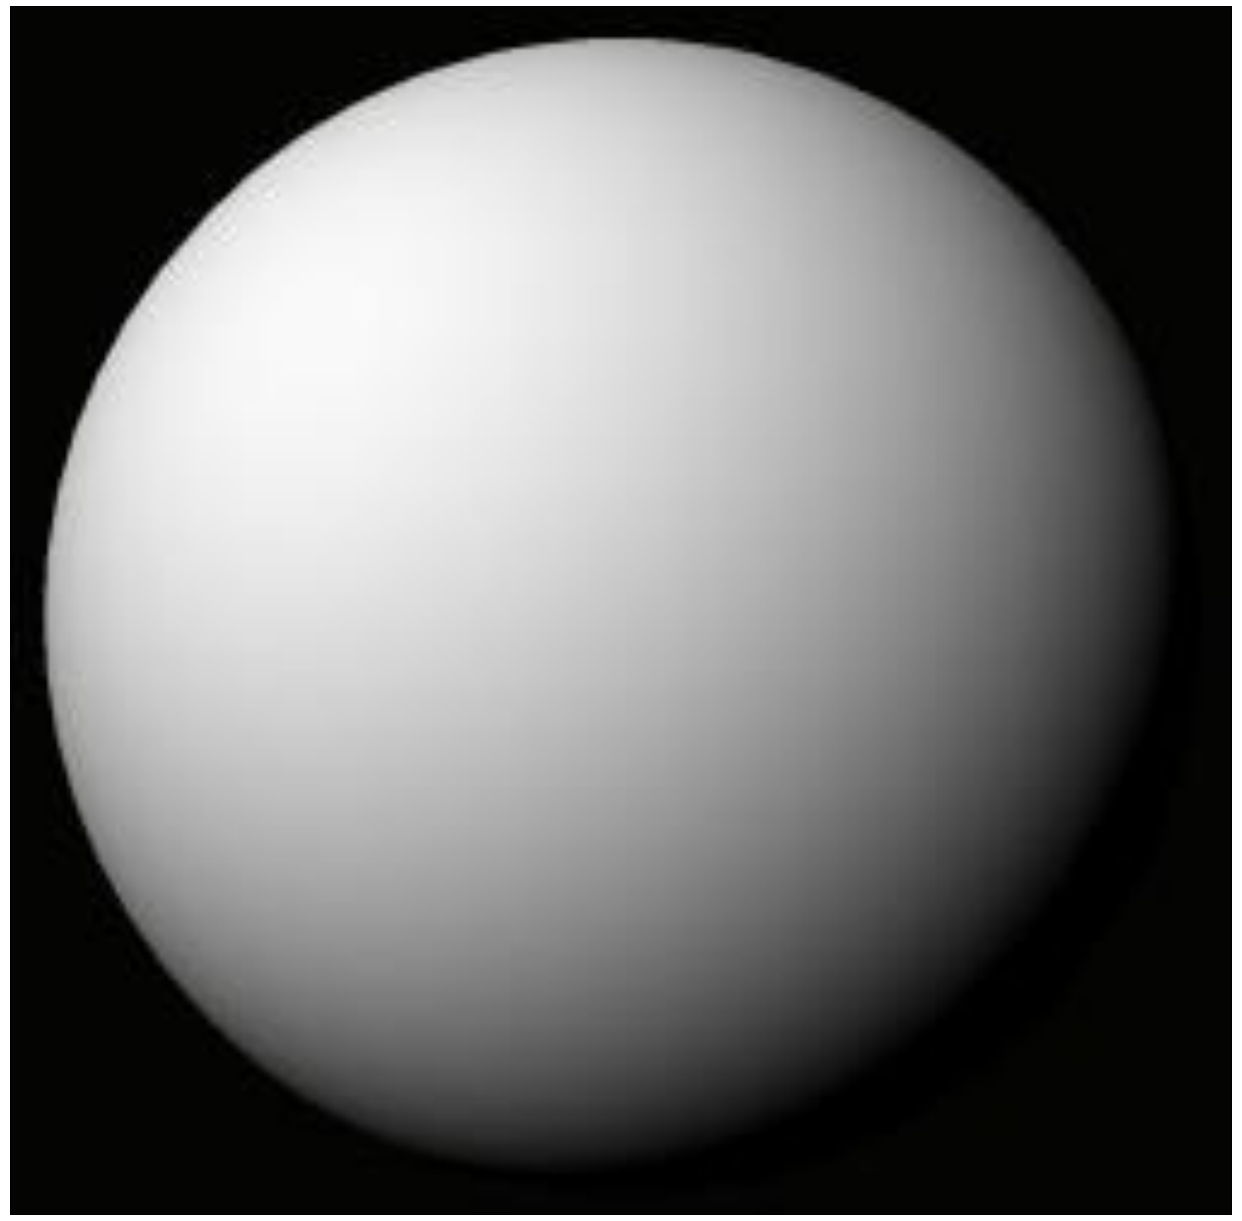

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

Assume k_d is 1 for now.
What can we measure from one image?

- $\cos^{-1}(I)$ is the angle between N and L
- Add assumptions:
  - Constant albedo
  - A few known normals (e.g. silhouettes)
  - Smoothness of normals

In practice, this doesn't work well: assumptions are too restrictive, too much ambiguity in nontrivial scenes.

Source: N. Snavely

An ambiguity that artists exploit!

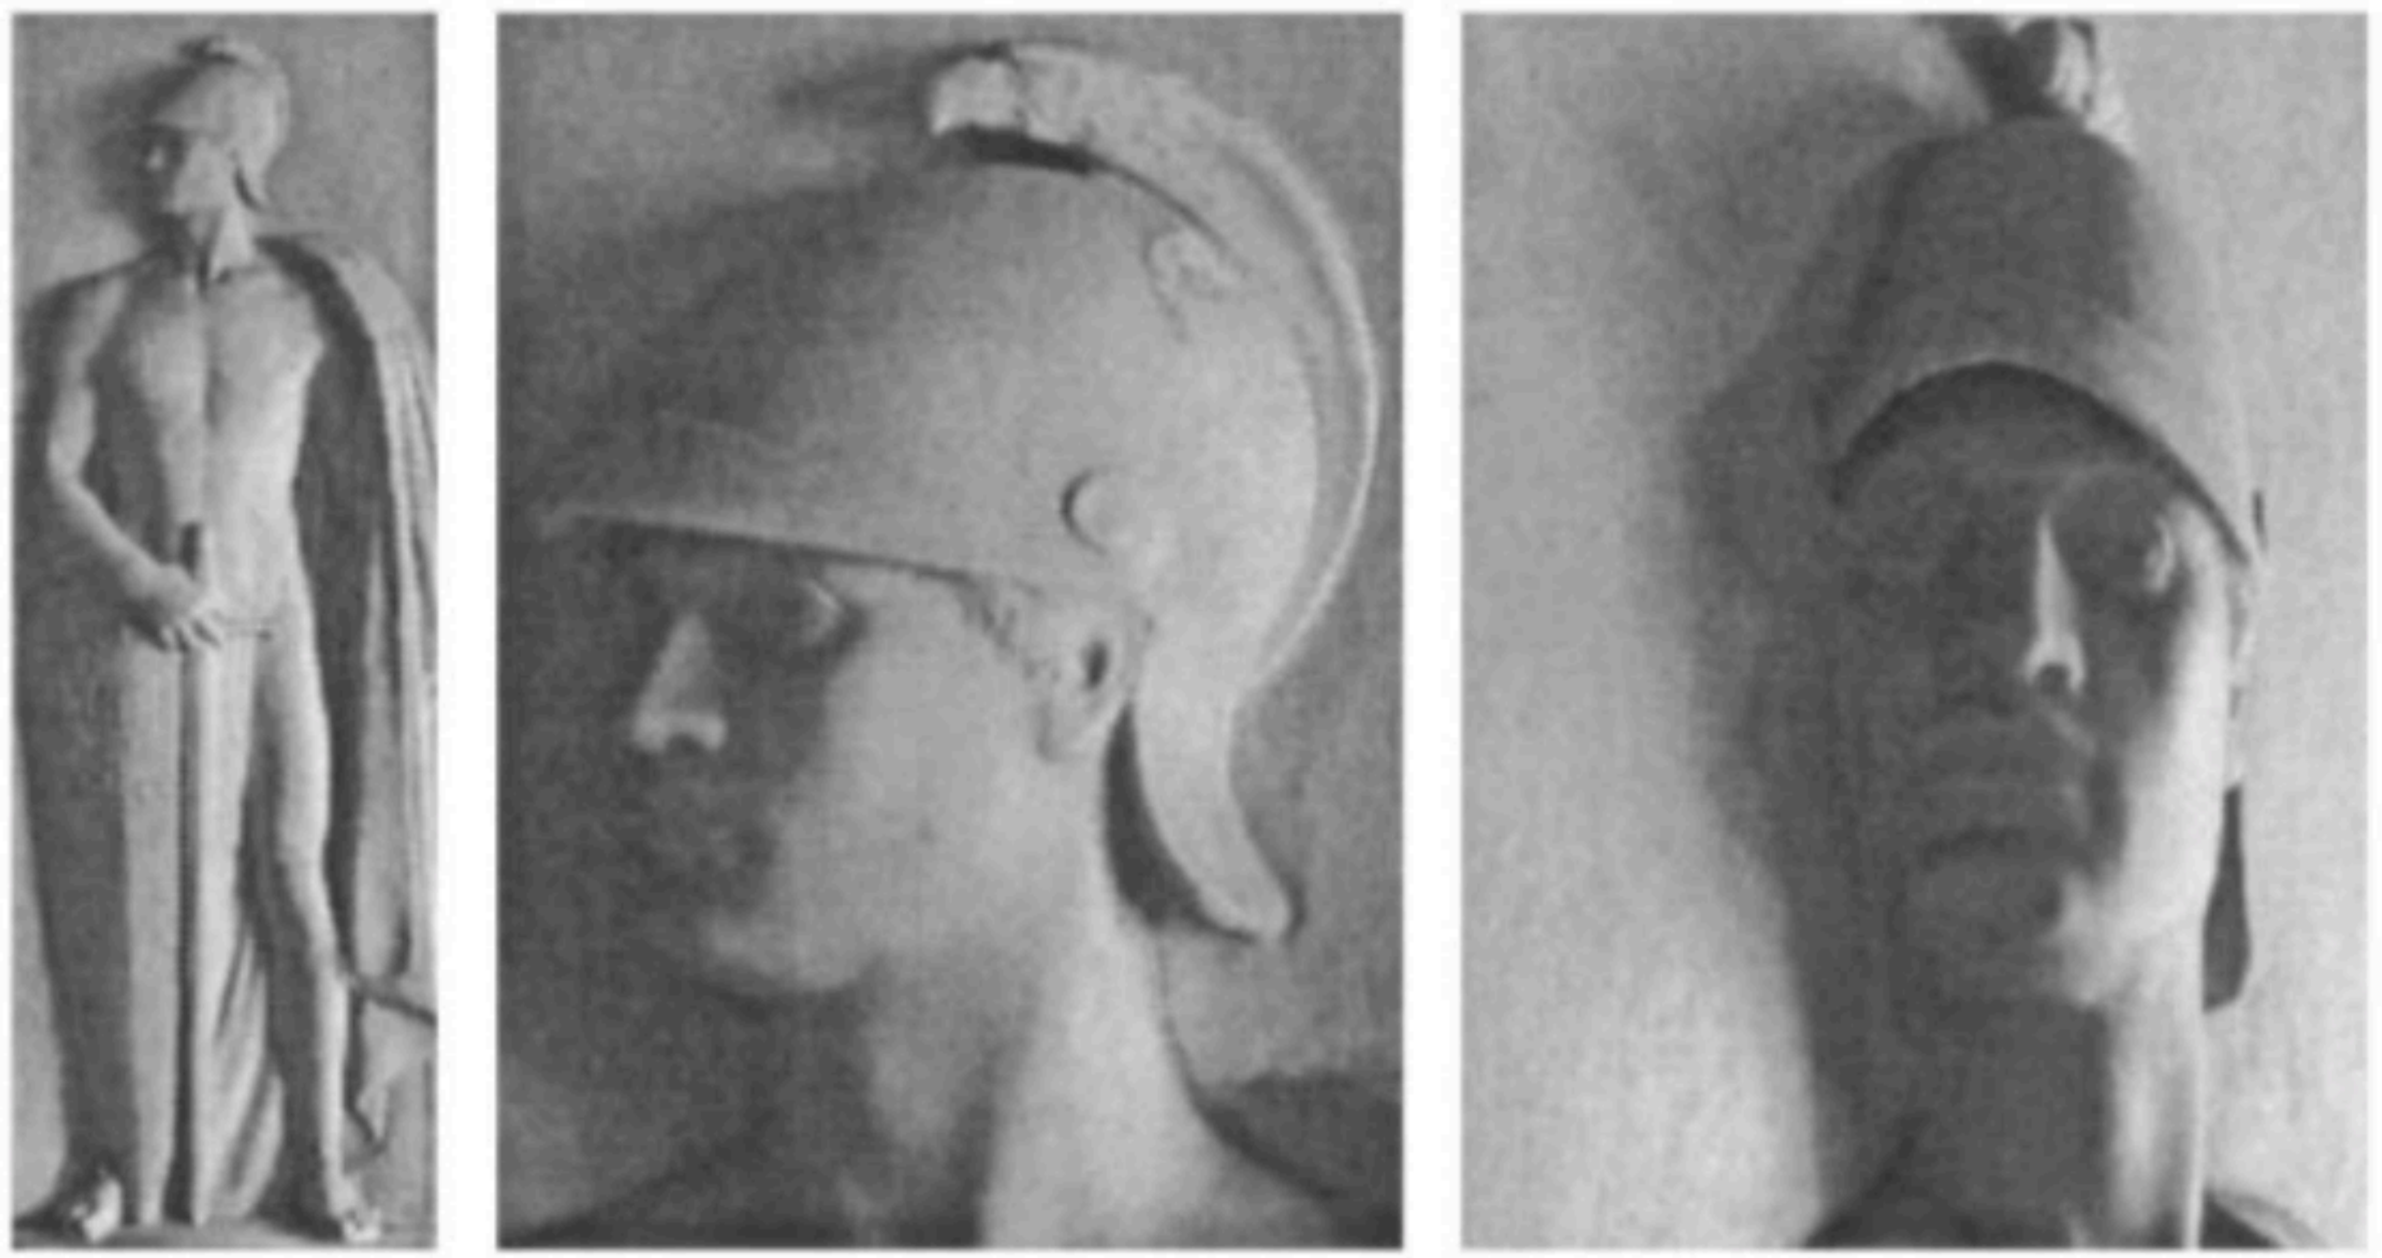

[Belhumeur et al. “The Bas-Relief Ambiguity”, 1999]

Contours provide extra shape information

Consider points on the *occluding contour*:

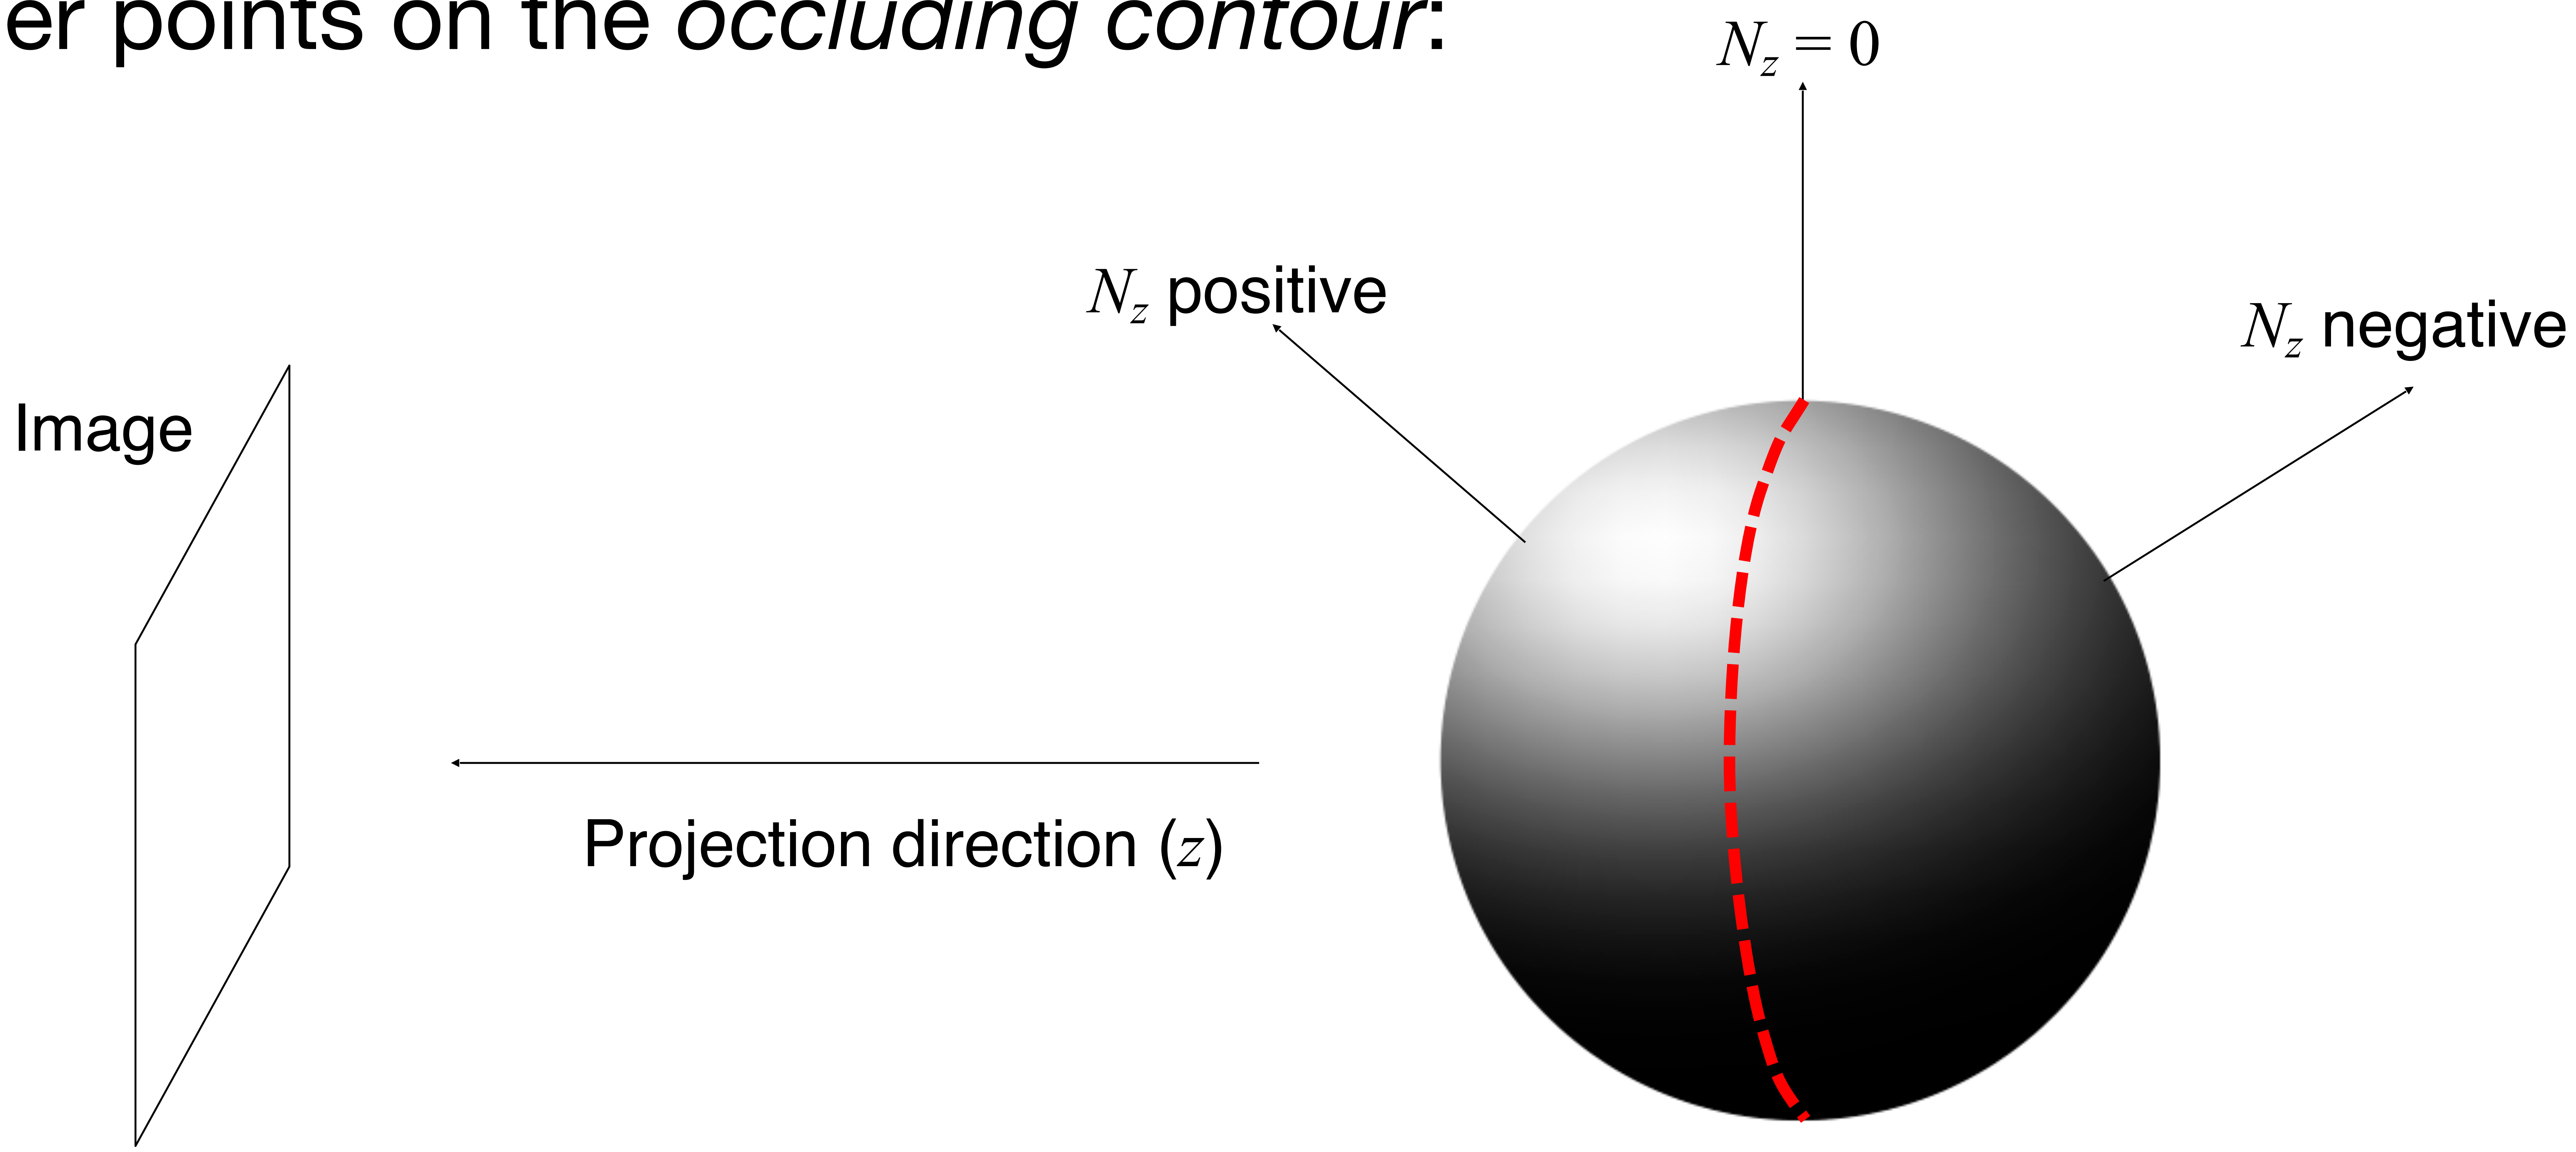

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Source: S. Lazebnik

Application: finding the direction of the light source

$$I(x,y) = \mathbf{N}(x,y) \cdot \mathbf{S}(x,y)$$

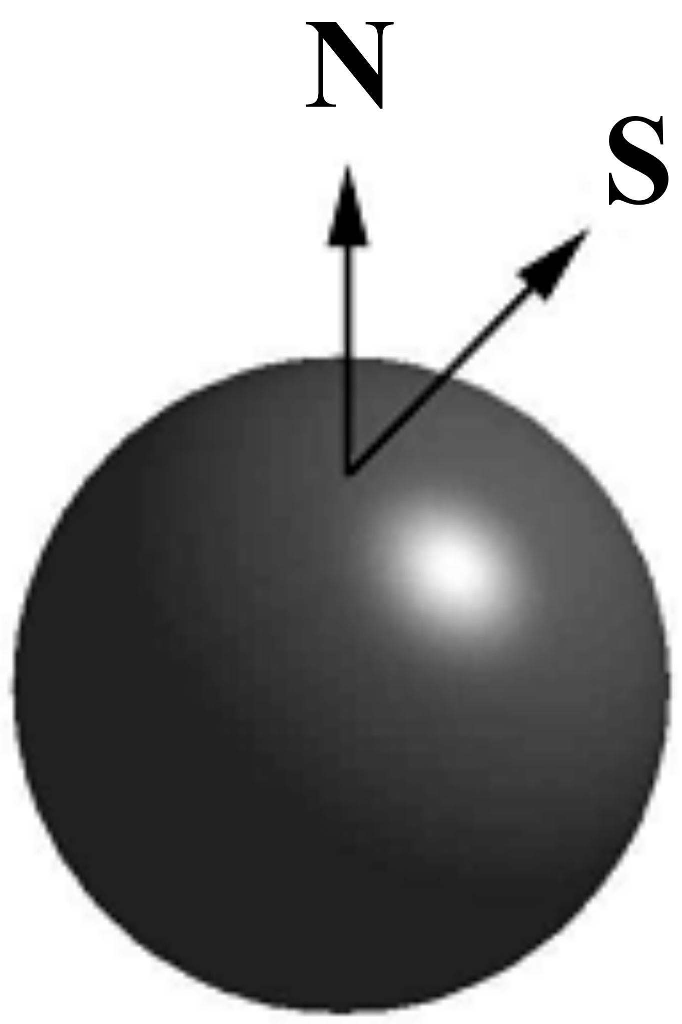

Full 3D case:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & N_z(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) & N_z(x_2,y_2) \\ \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & N_z(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ S_z \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

For points on the occluding contour, $N_z = 0$:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) \\ \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Source: S. Lazebnik

Finding the direction of the light source

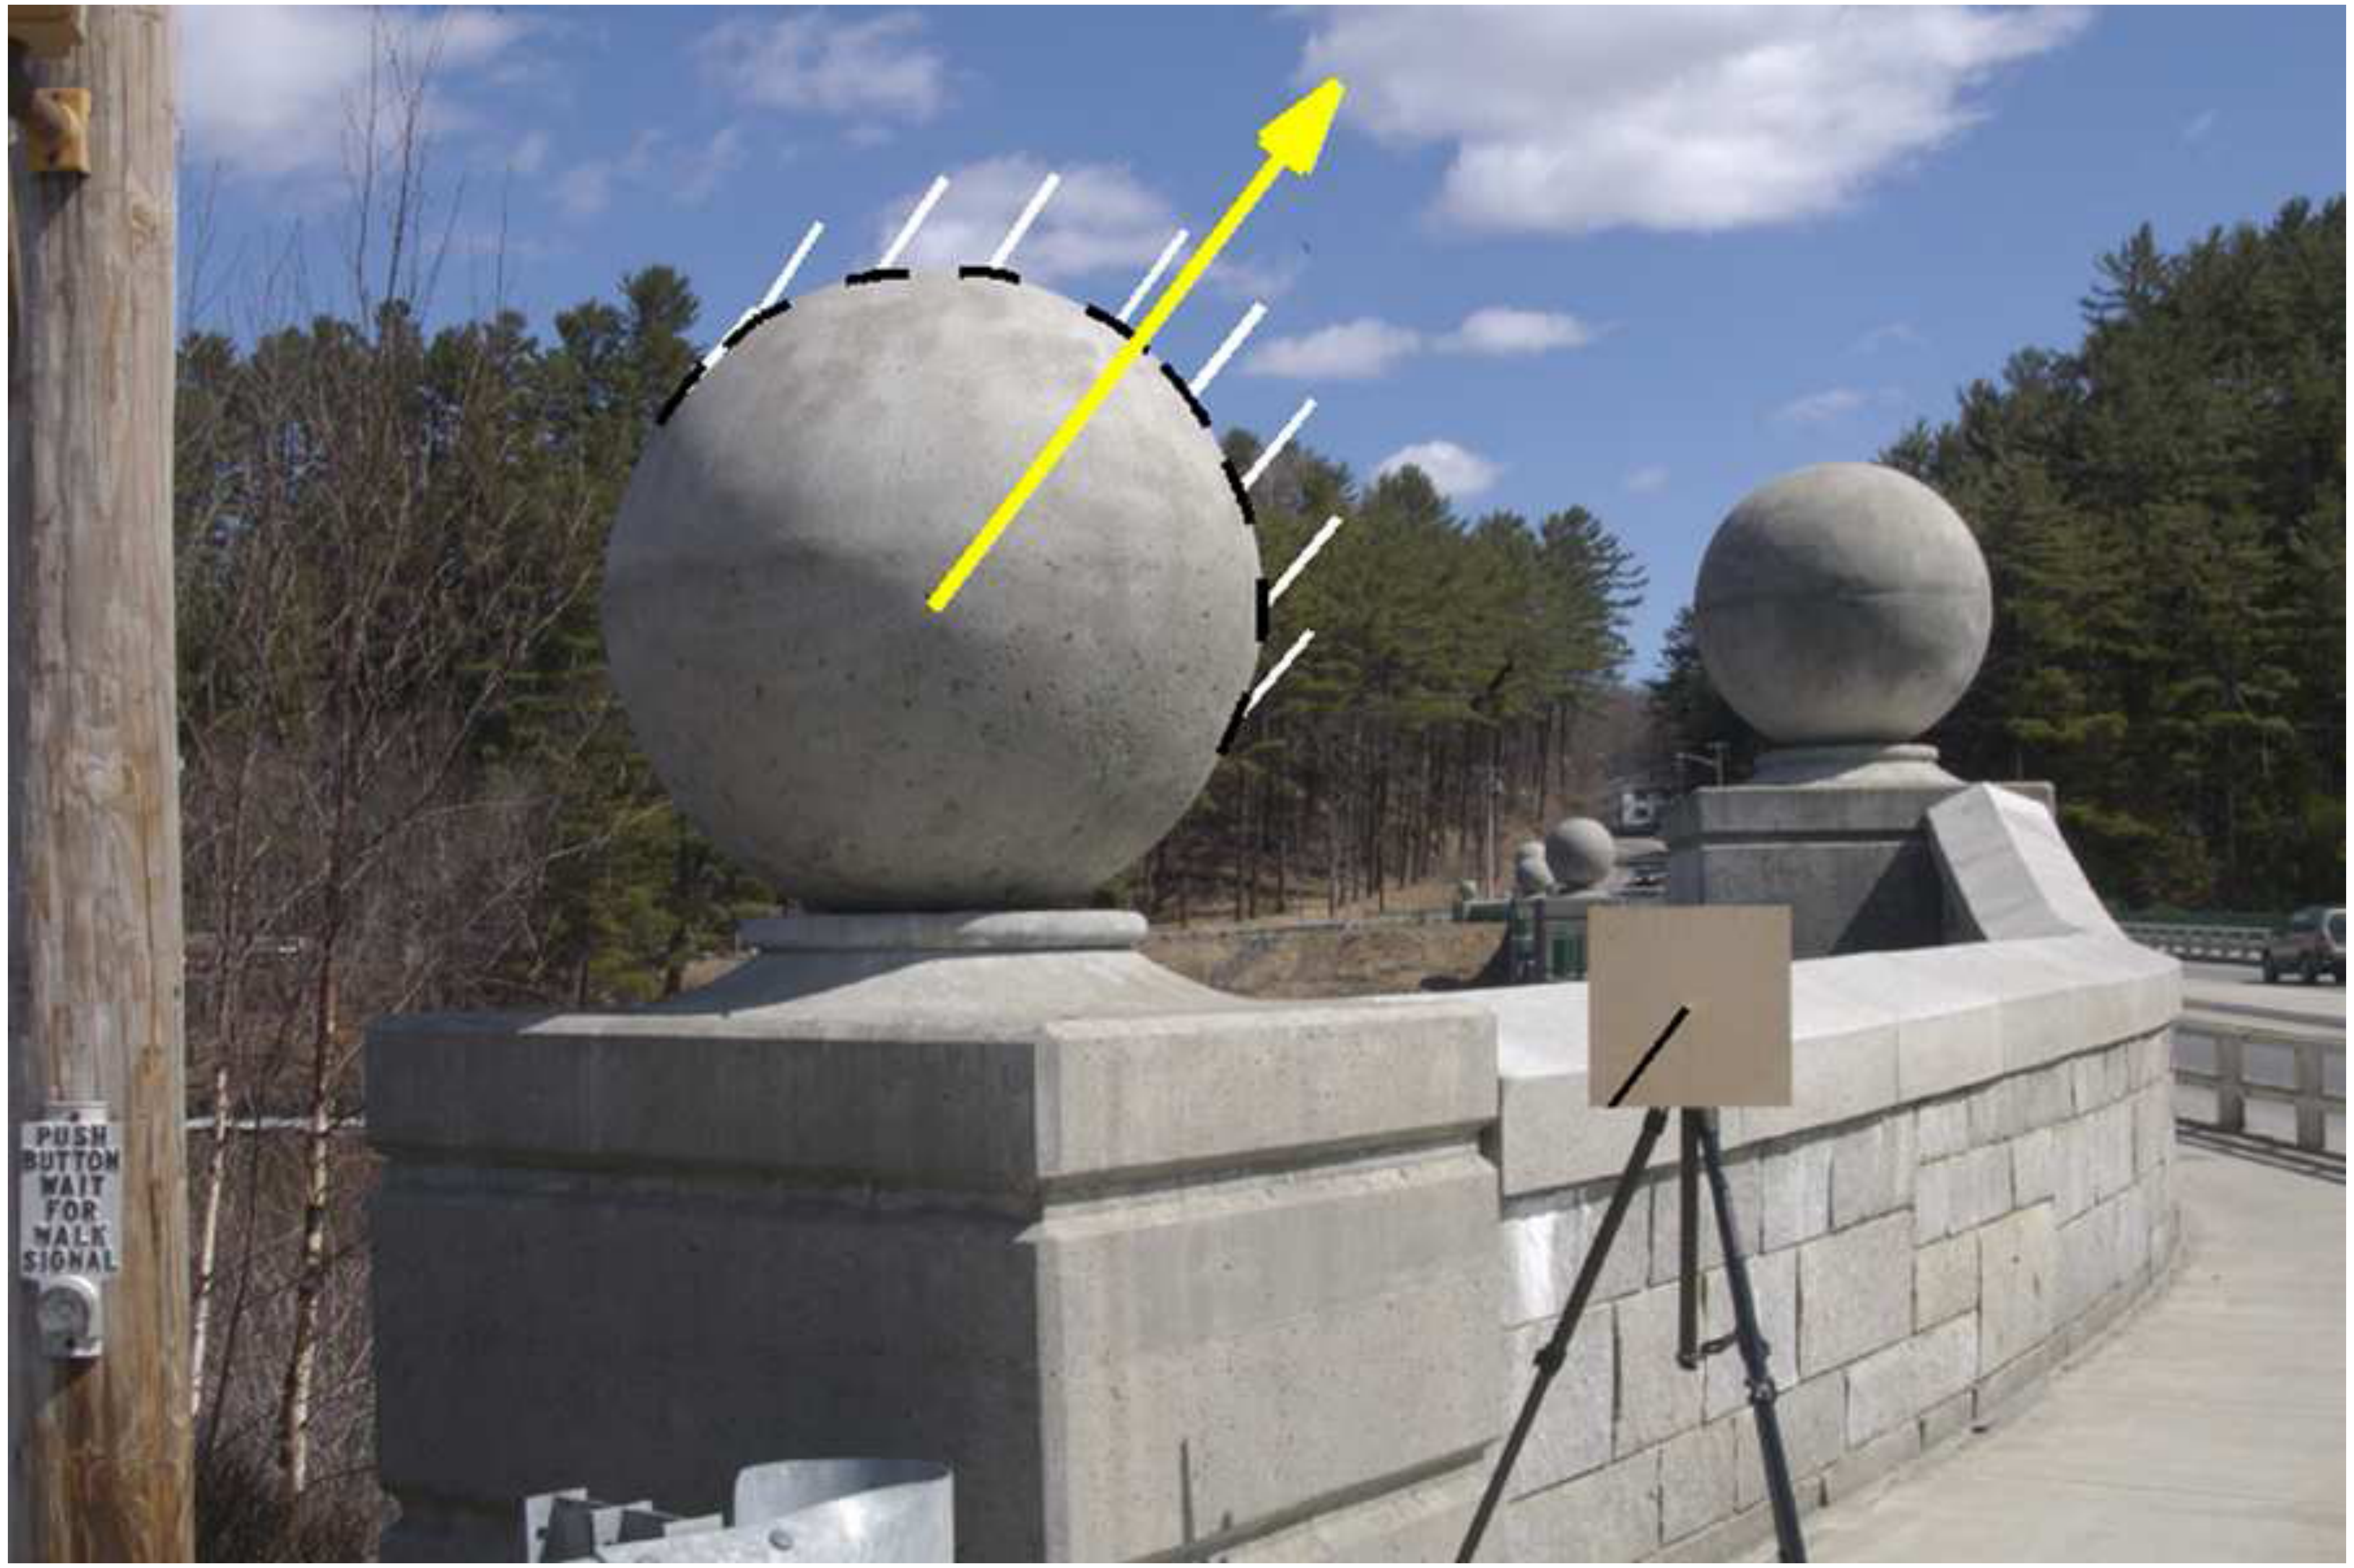

P. Nillius and J.-O. Eklundh, "Automatic estimation of the projected light source direction," CVPR 2001

Source: S. Lazebnik

Application: detecting image splices

Fake photo

Real photo

[Johnson and Farid, 2005]

Source: S. Lazebnik

Photometric stereo

Source: N. Snavely

Photometric stereo

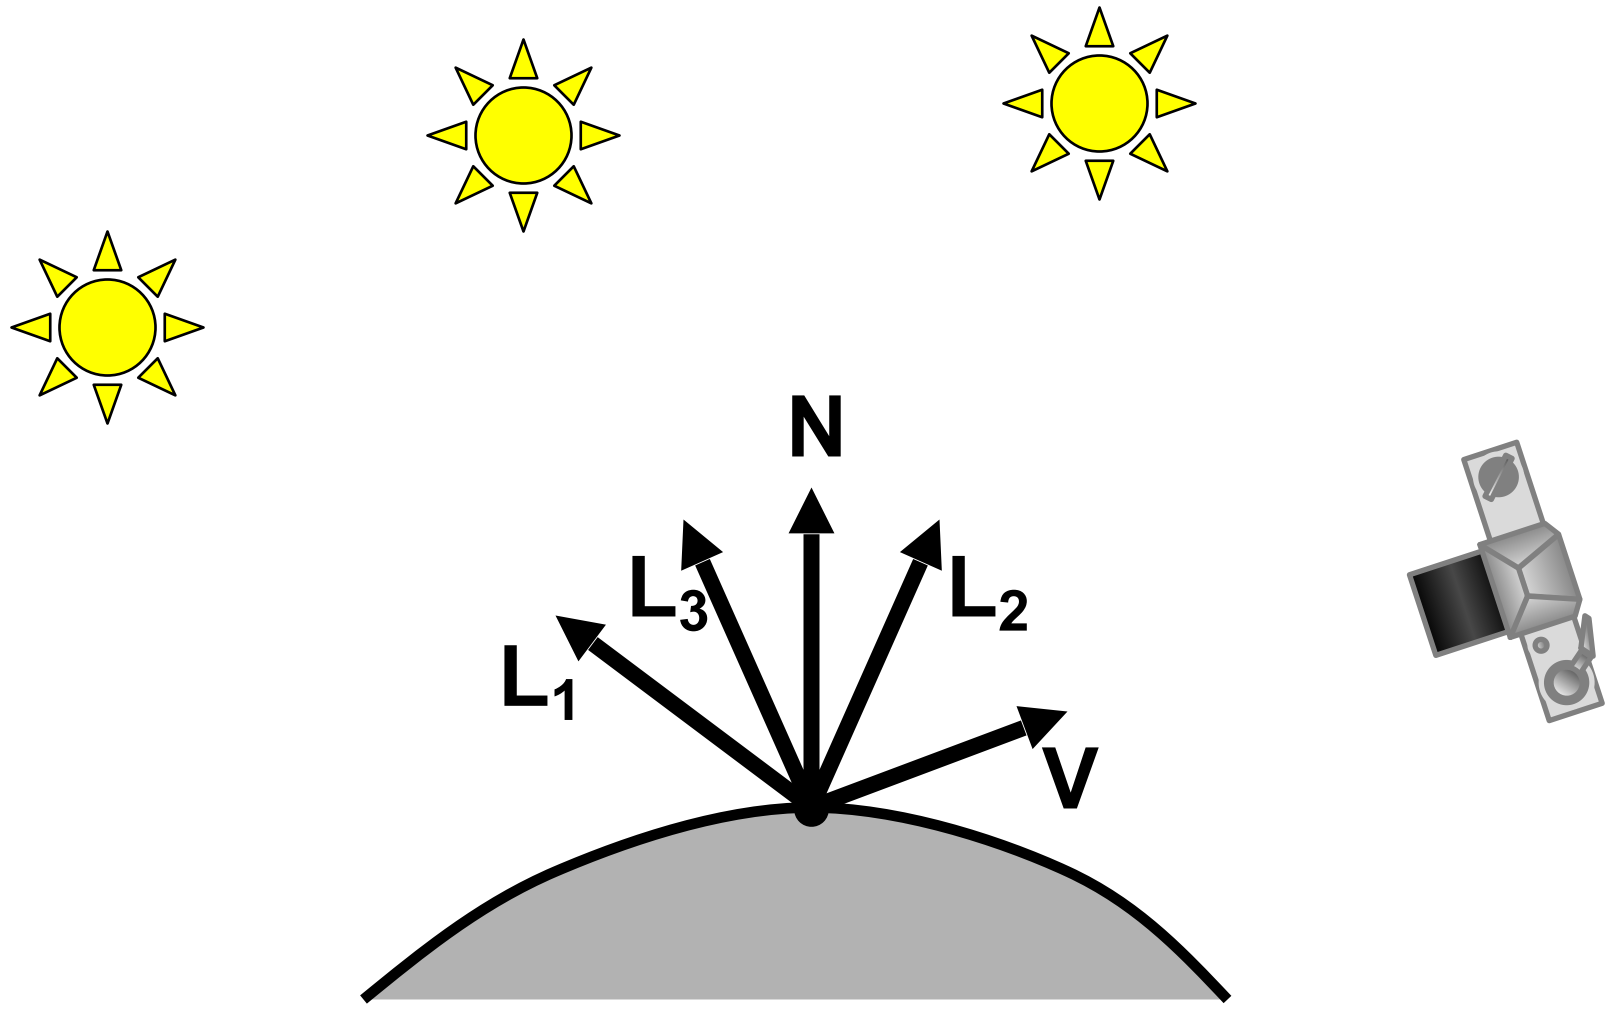

$$
\begin{aligned}
I_1 &= k_d \mathbf{N} \cdot \mathbf{L_1} \\
I_2 &= k_d \mathbf{N} \cdot \mathbf{L_2} \\
I_3 &= k_d \mathbf{N} \cdot \mathbf{L_3}
\end{aligned}
$$

Can write this as a linear system, and solve:

$$
\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = k_d \begin{bmatrix} \mathbf{L_1}^T \\ \mathbf{L_2}^T \\ \mathbf{L_3}^T \end{bmatrix} \mathbf{N}
$$

Source: N. Snavely

Photometric Stereo

Input

Recovered albedo

Recovered normal field

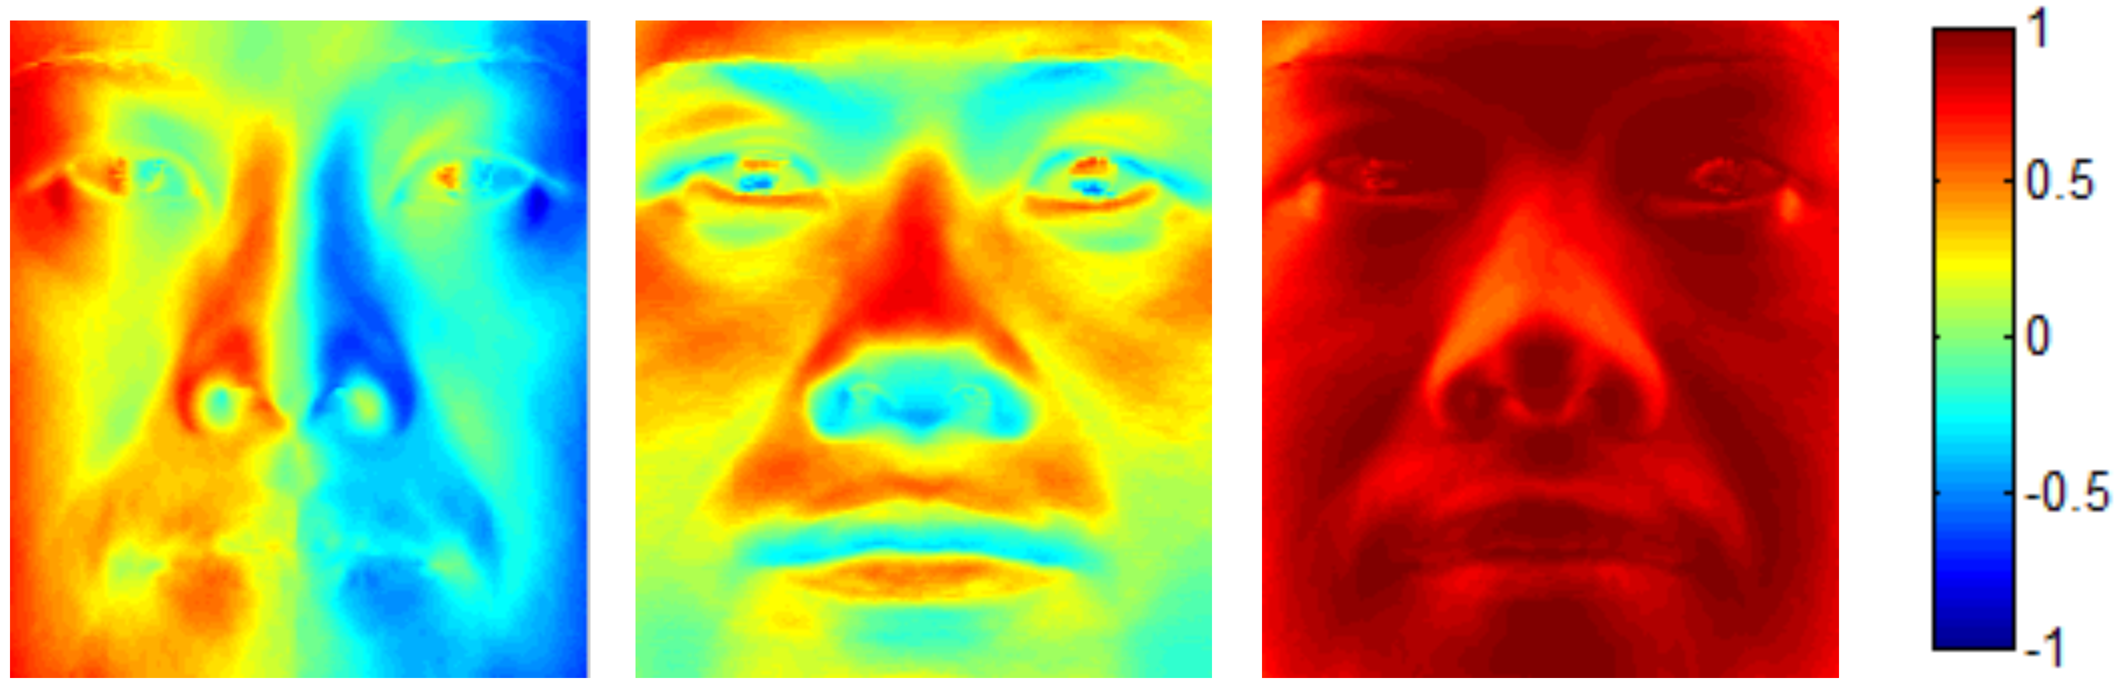

x y z

Recovered surface model

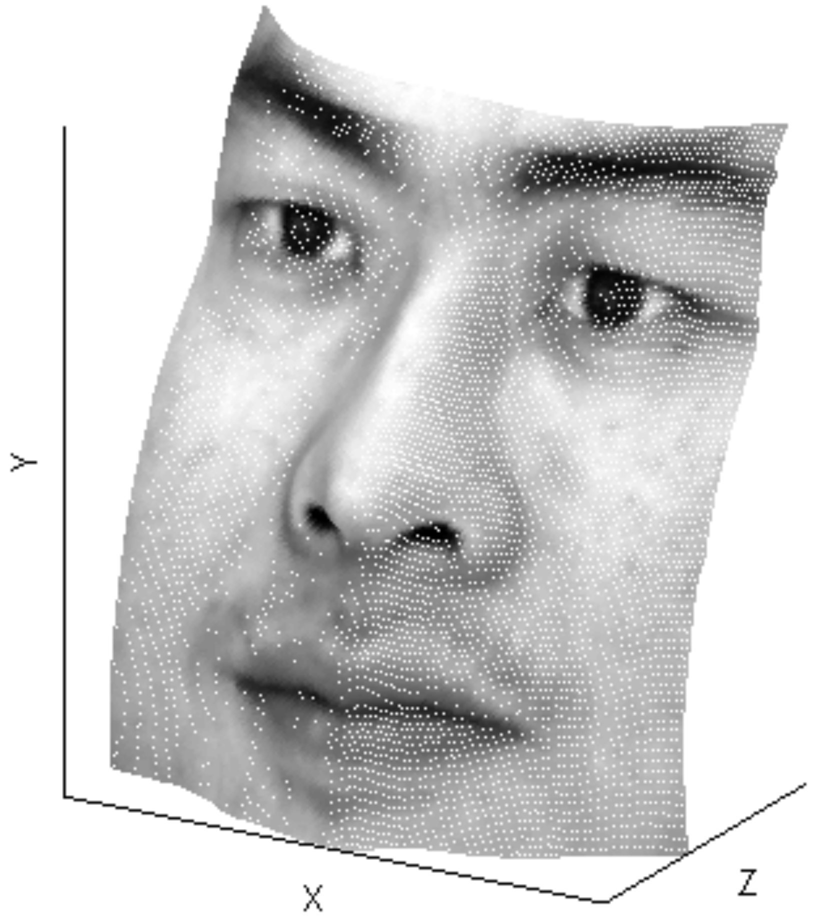

Source: Forsyth & Ponce, S. Lazebnik

Photometric Stereo

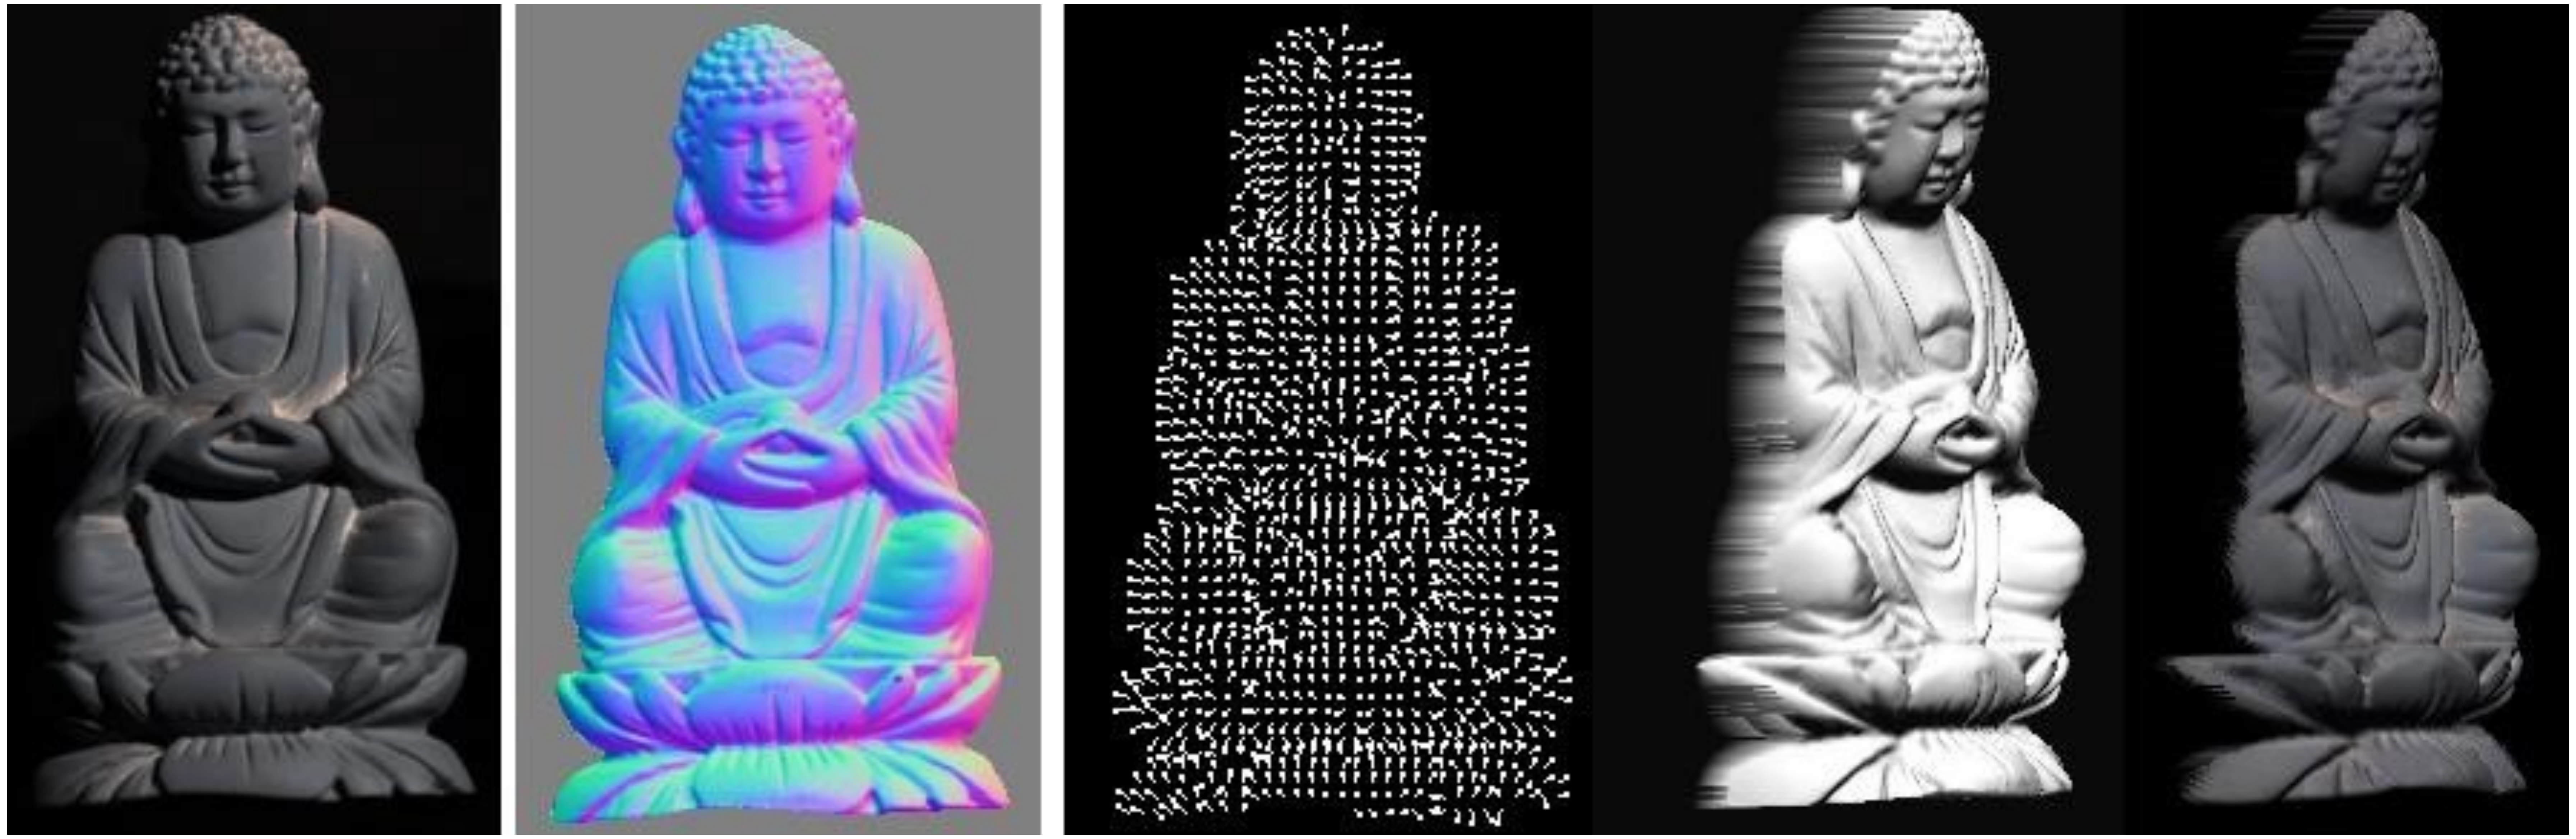

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Source: N. Snavely

Video photometric stereo

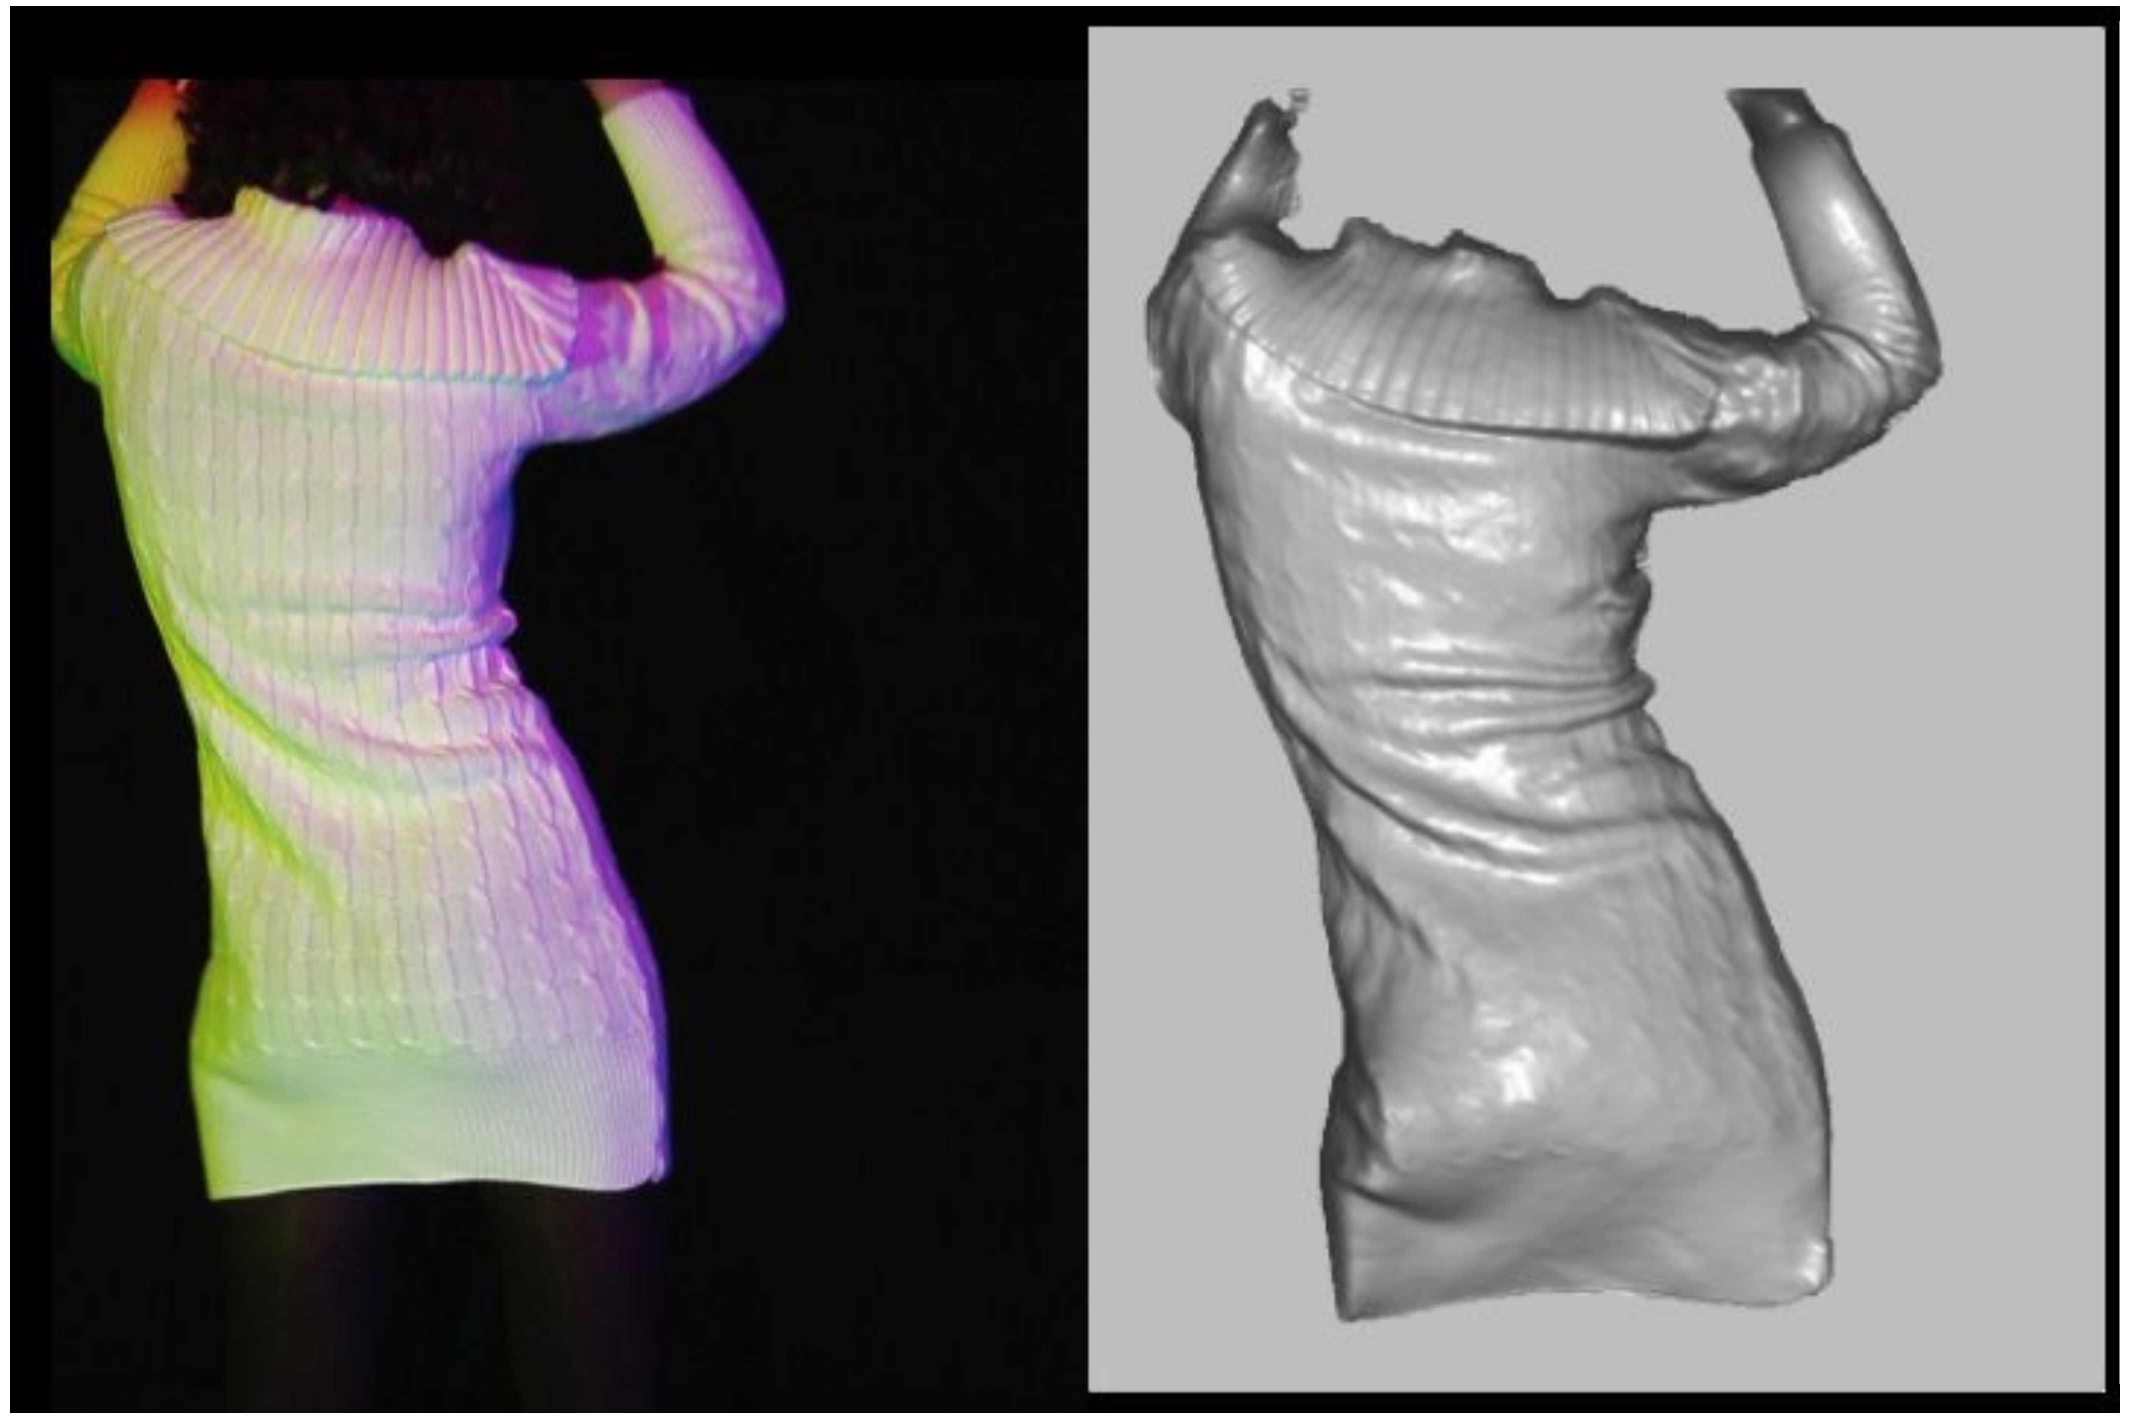

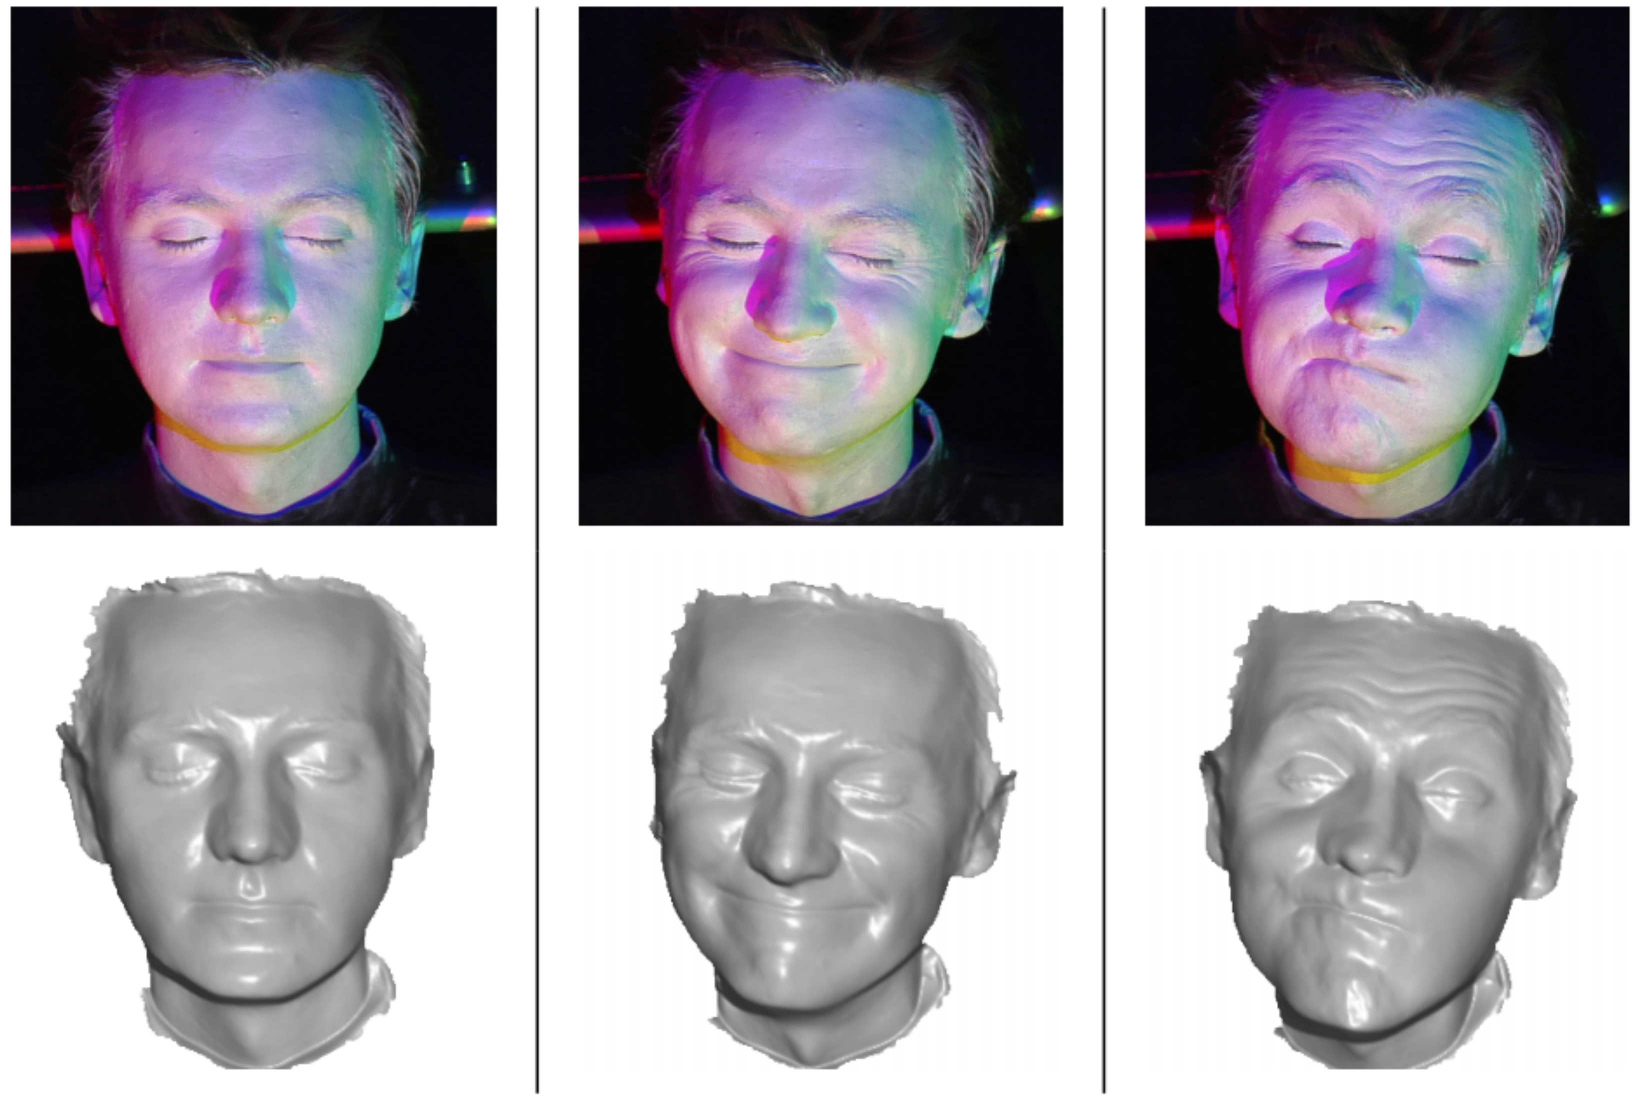

Fig. 2. **Applying the original algorithm to a face with white makeup.** Top: example input frames from video of an actor smiling and grimacing. Bottom: the resulting integrated surfaces.

Video Normals from Colored Lights
Gabriel J. Brostow, Carlos Hernández, George Vogiatzis, Björn Stenger, Roberto Cipolla
IEEE TPAMI, Vol. 33, No. 10, pages 2104-2114, October 2011.

Source: N. Snavely

But what if we don’t know the BRDF?

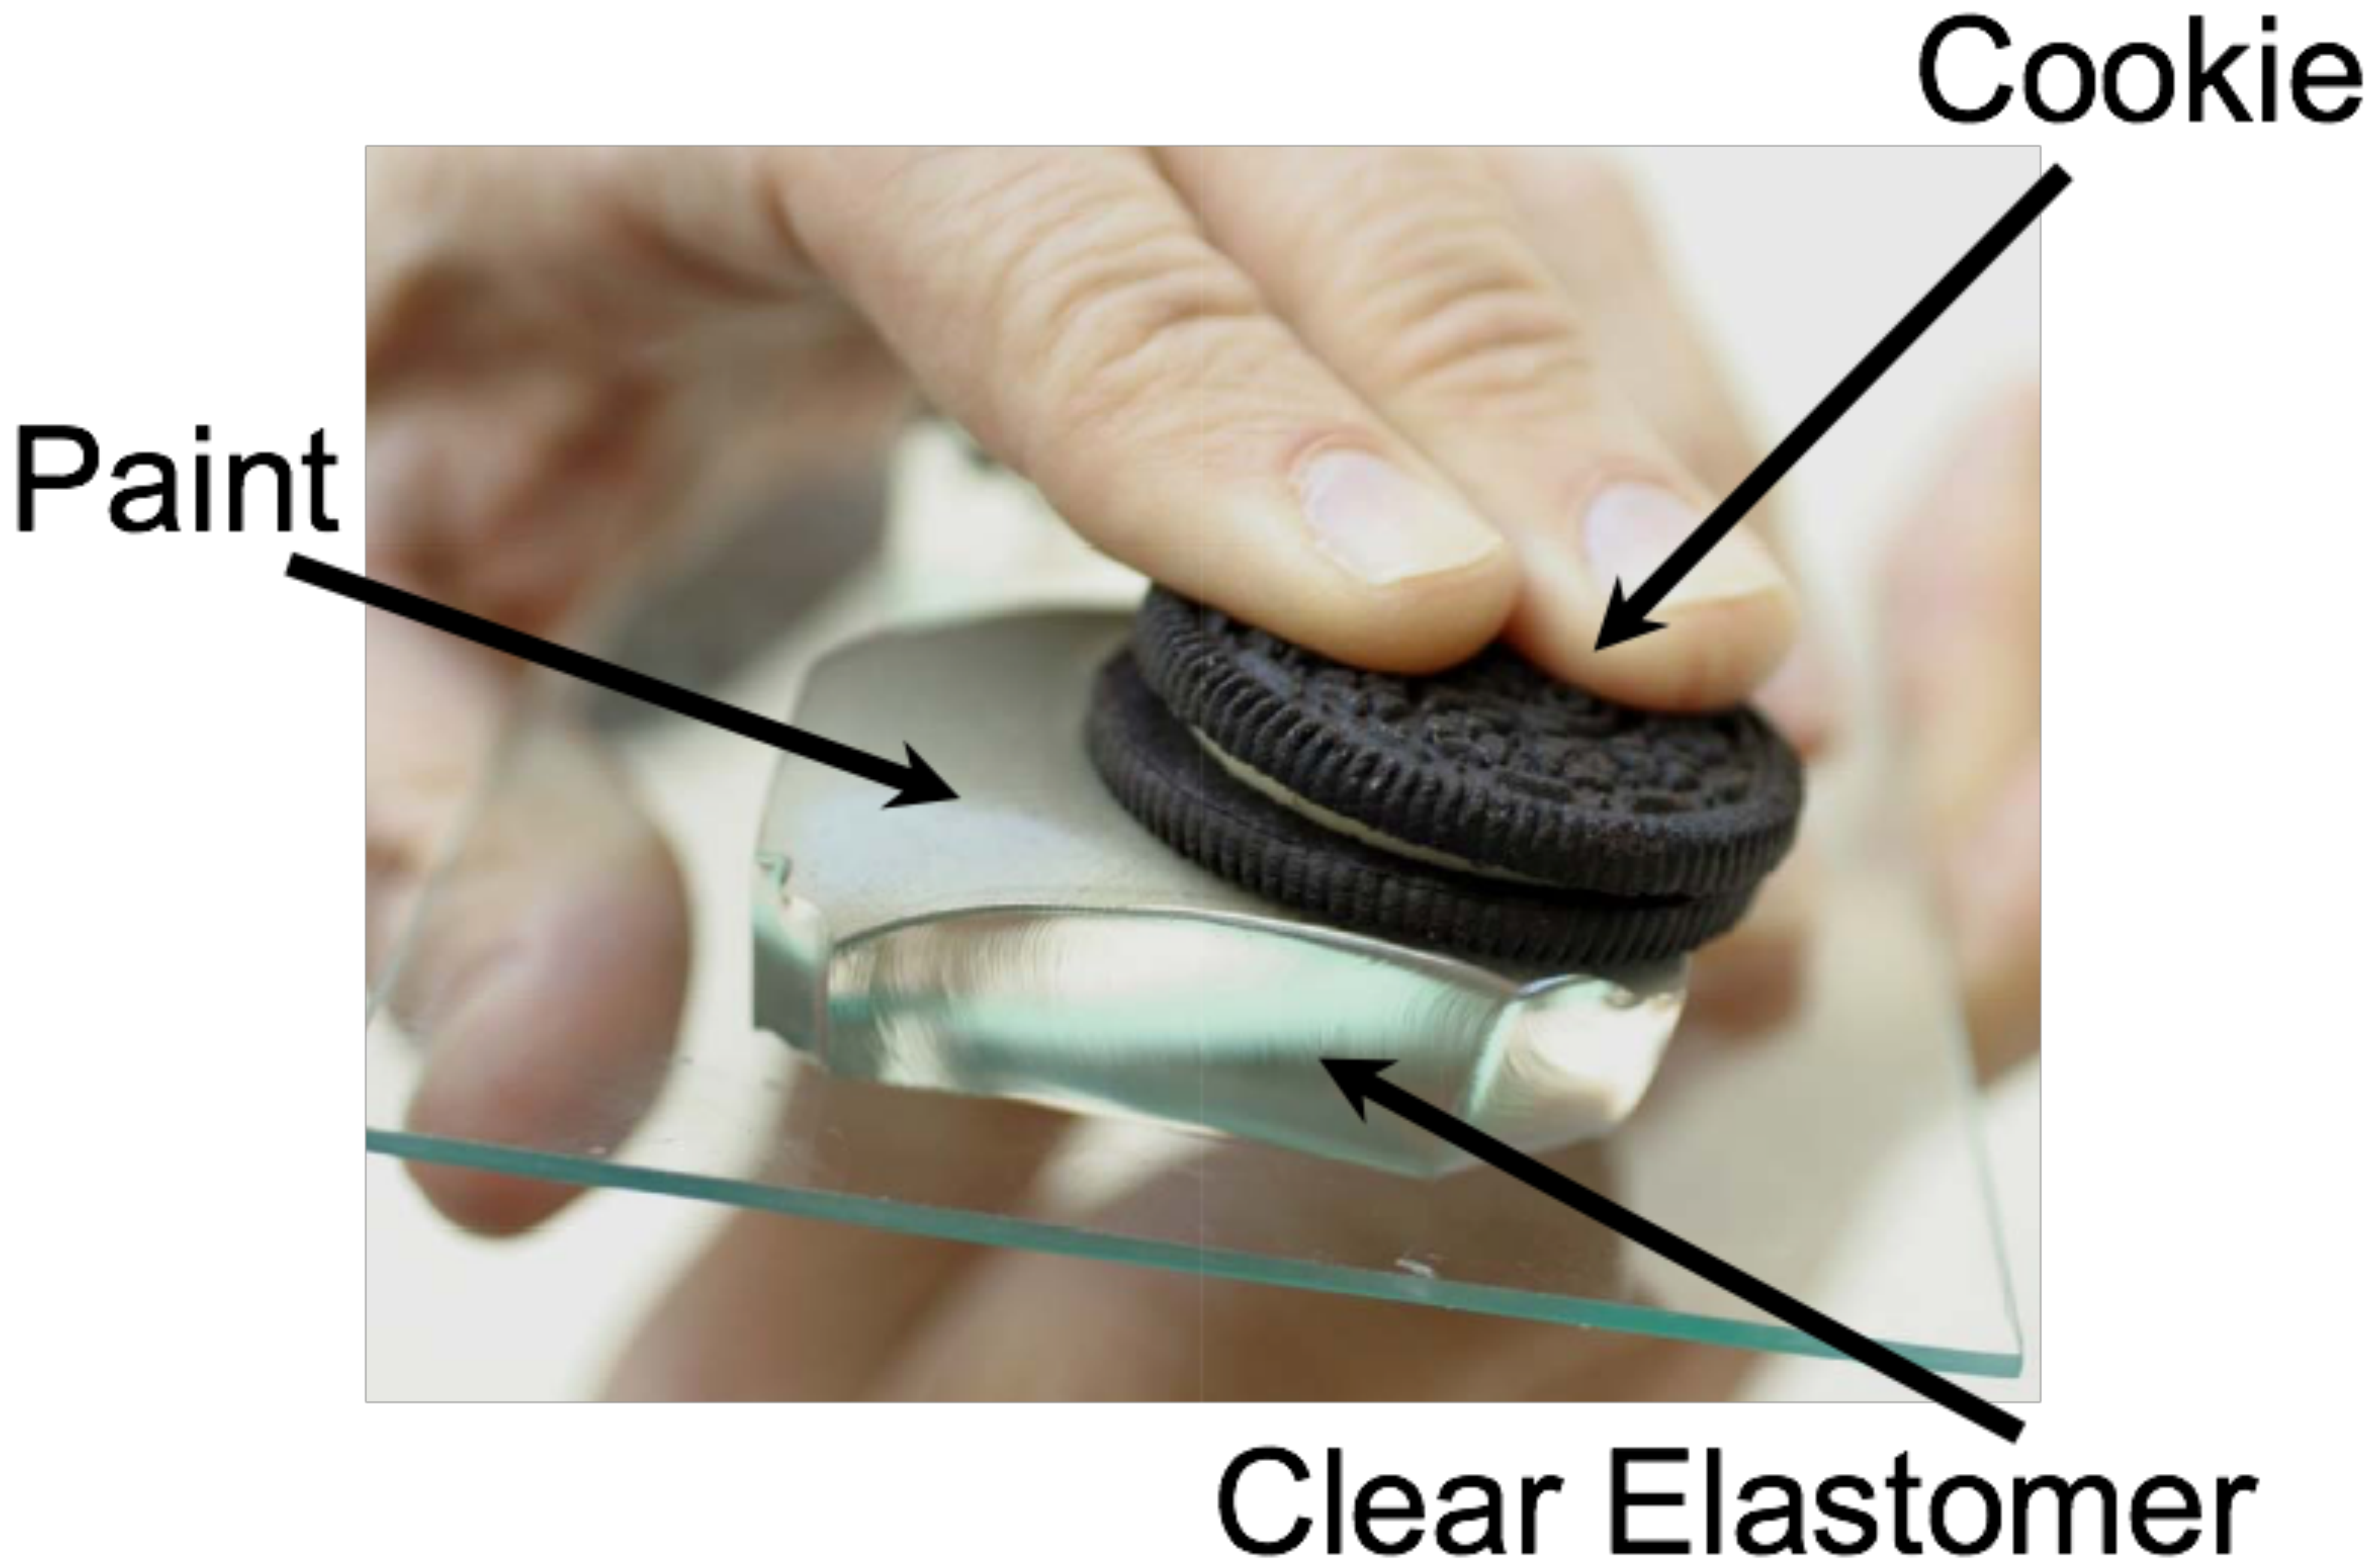

[Johnson and Adelson, 2009]

Source: N. Snavely

Source: N. Snavely

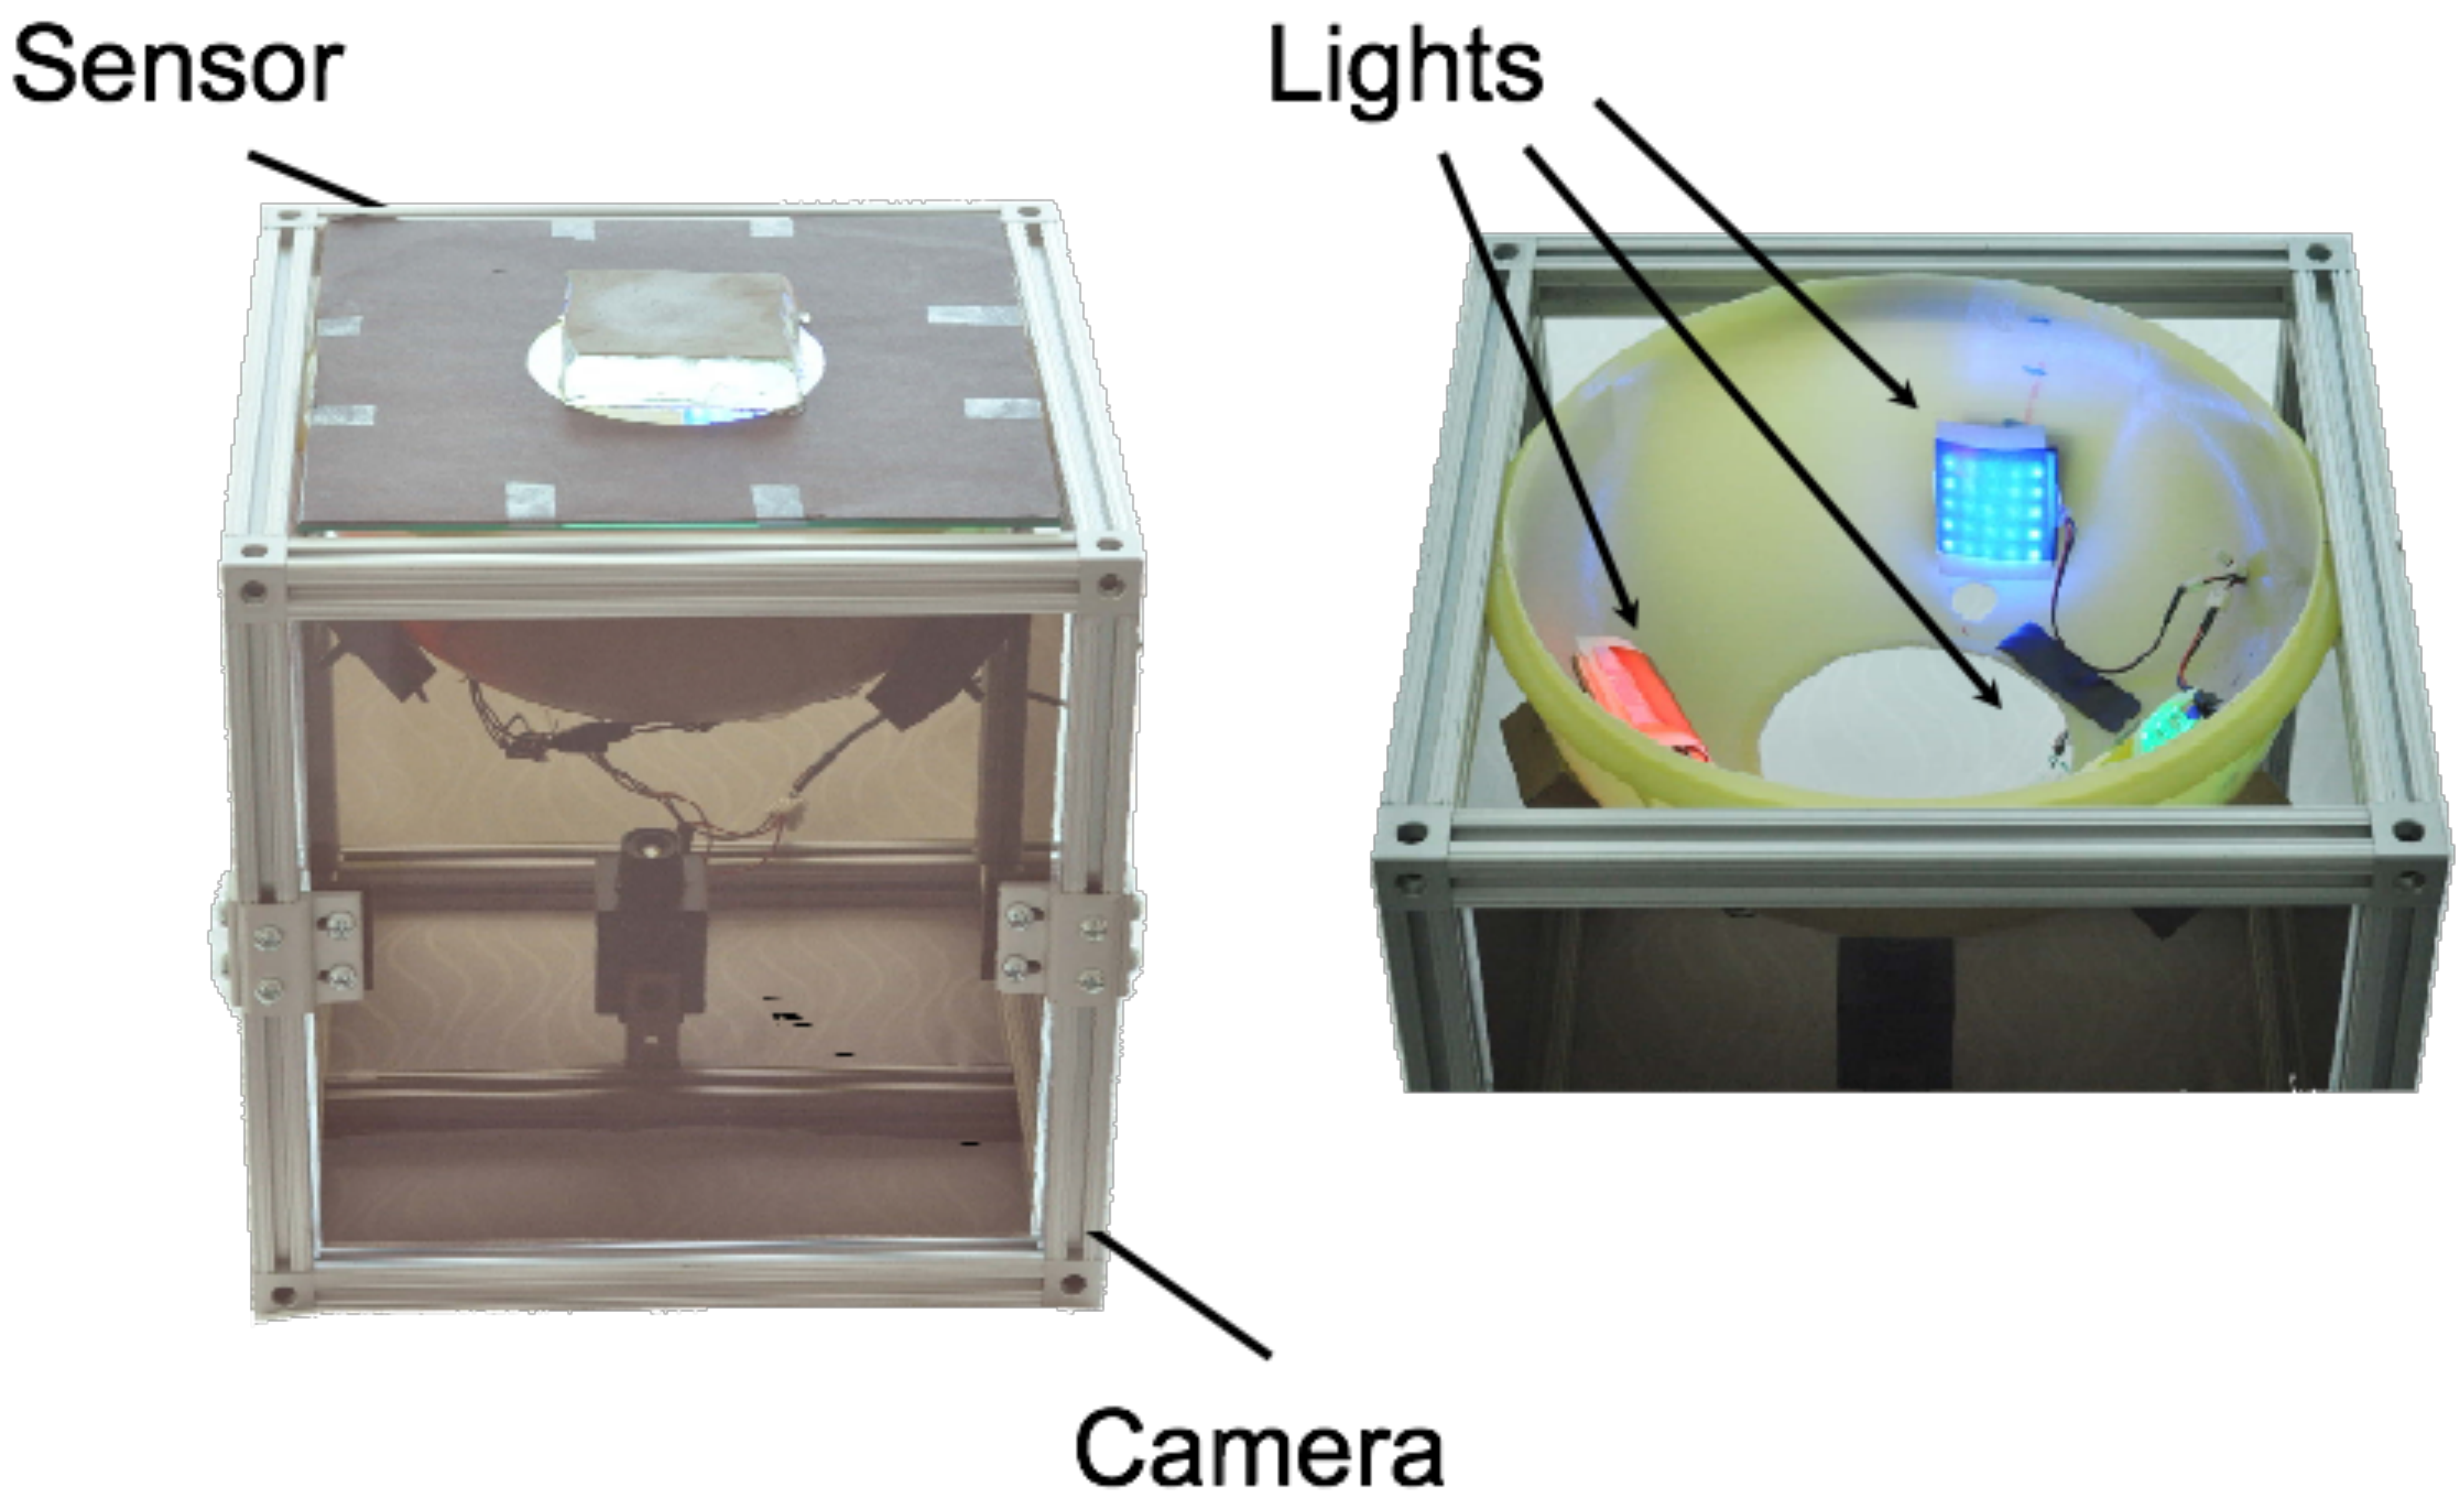

Source: N. Snavely

Source: N. Snavely

Today

- Shape from shading
- **Intrinsic image decomposition**
- Color perception

What about paint?

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

k_d is **reflectance** or **albedo**

Intrinsic image decomposition

Reflectance

Shading

Intrinsic image decomposition

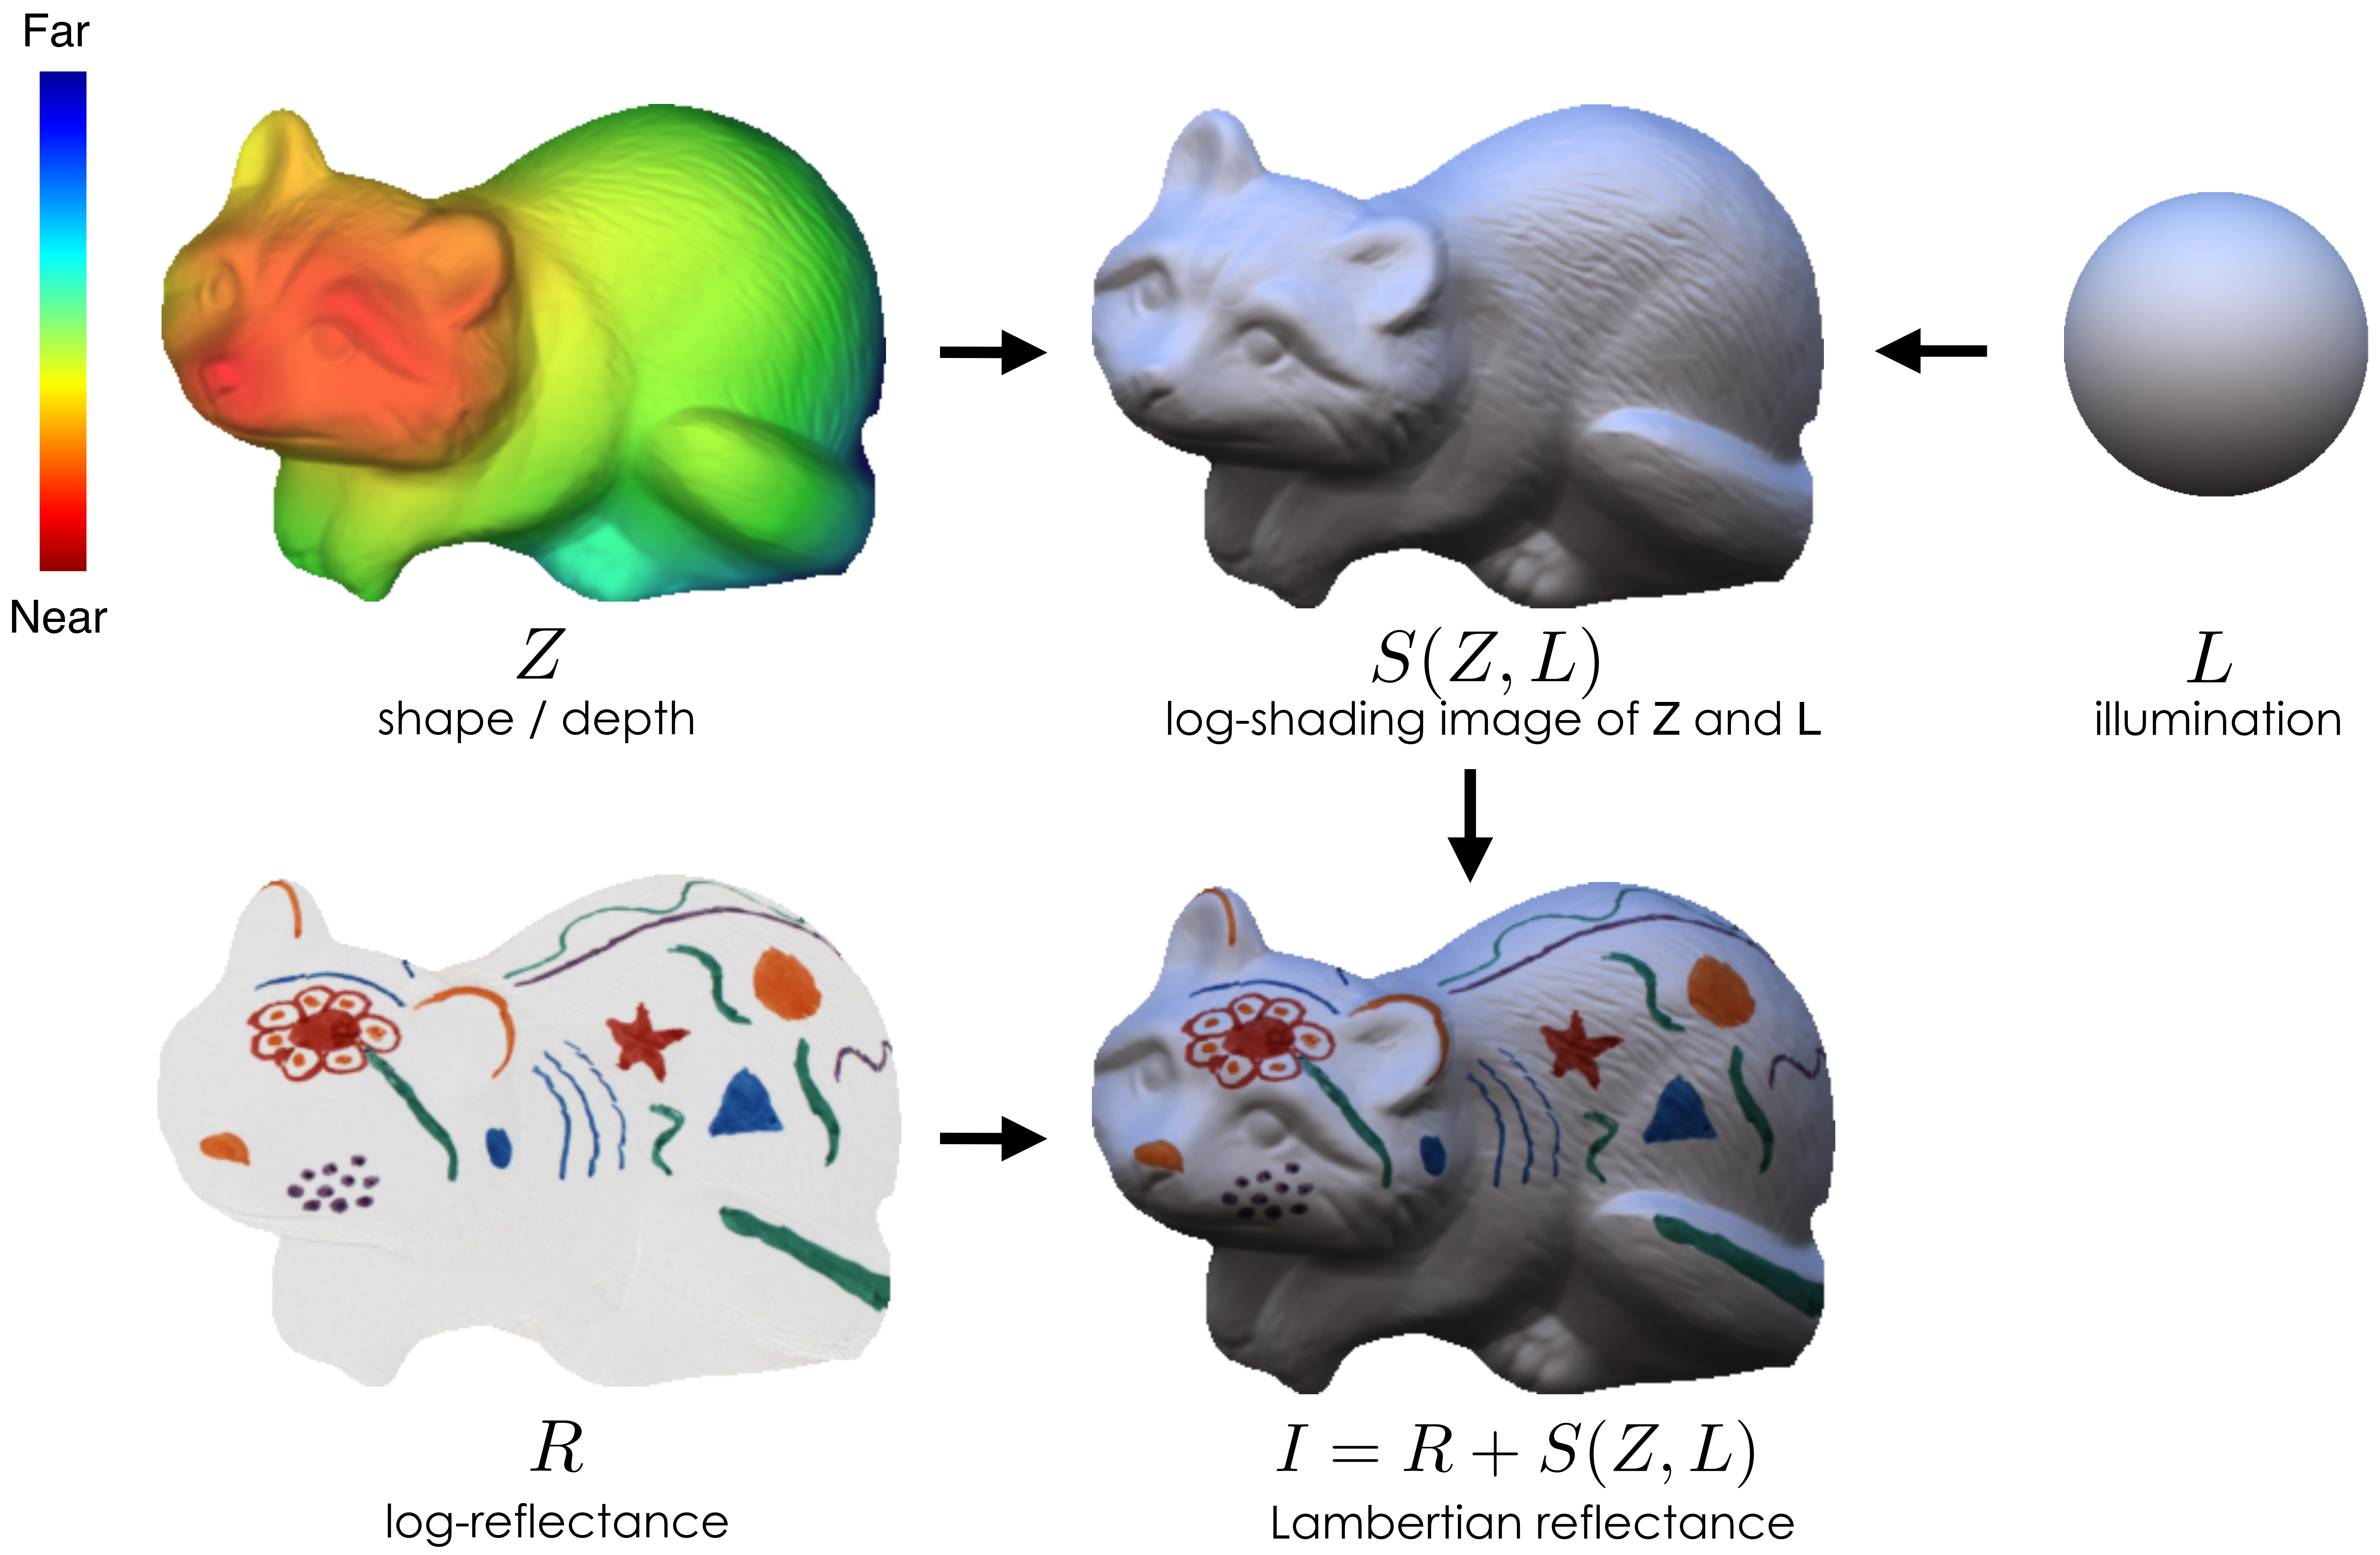

[Barrow and Tenenbaum 1978]

Source: J. Barron

Intrinsic image decomposition

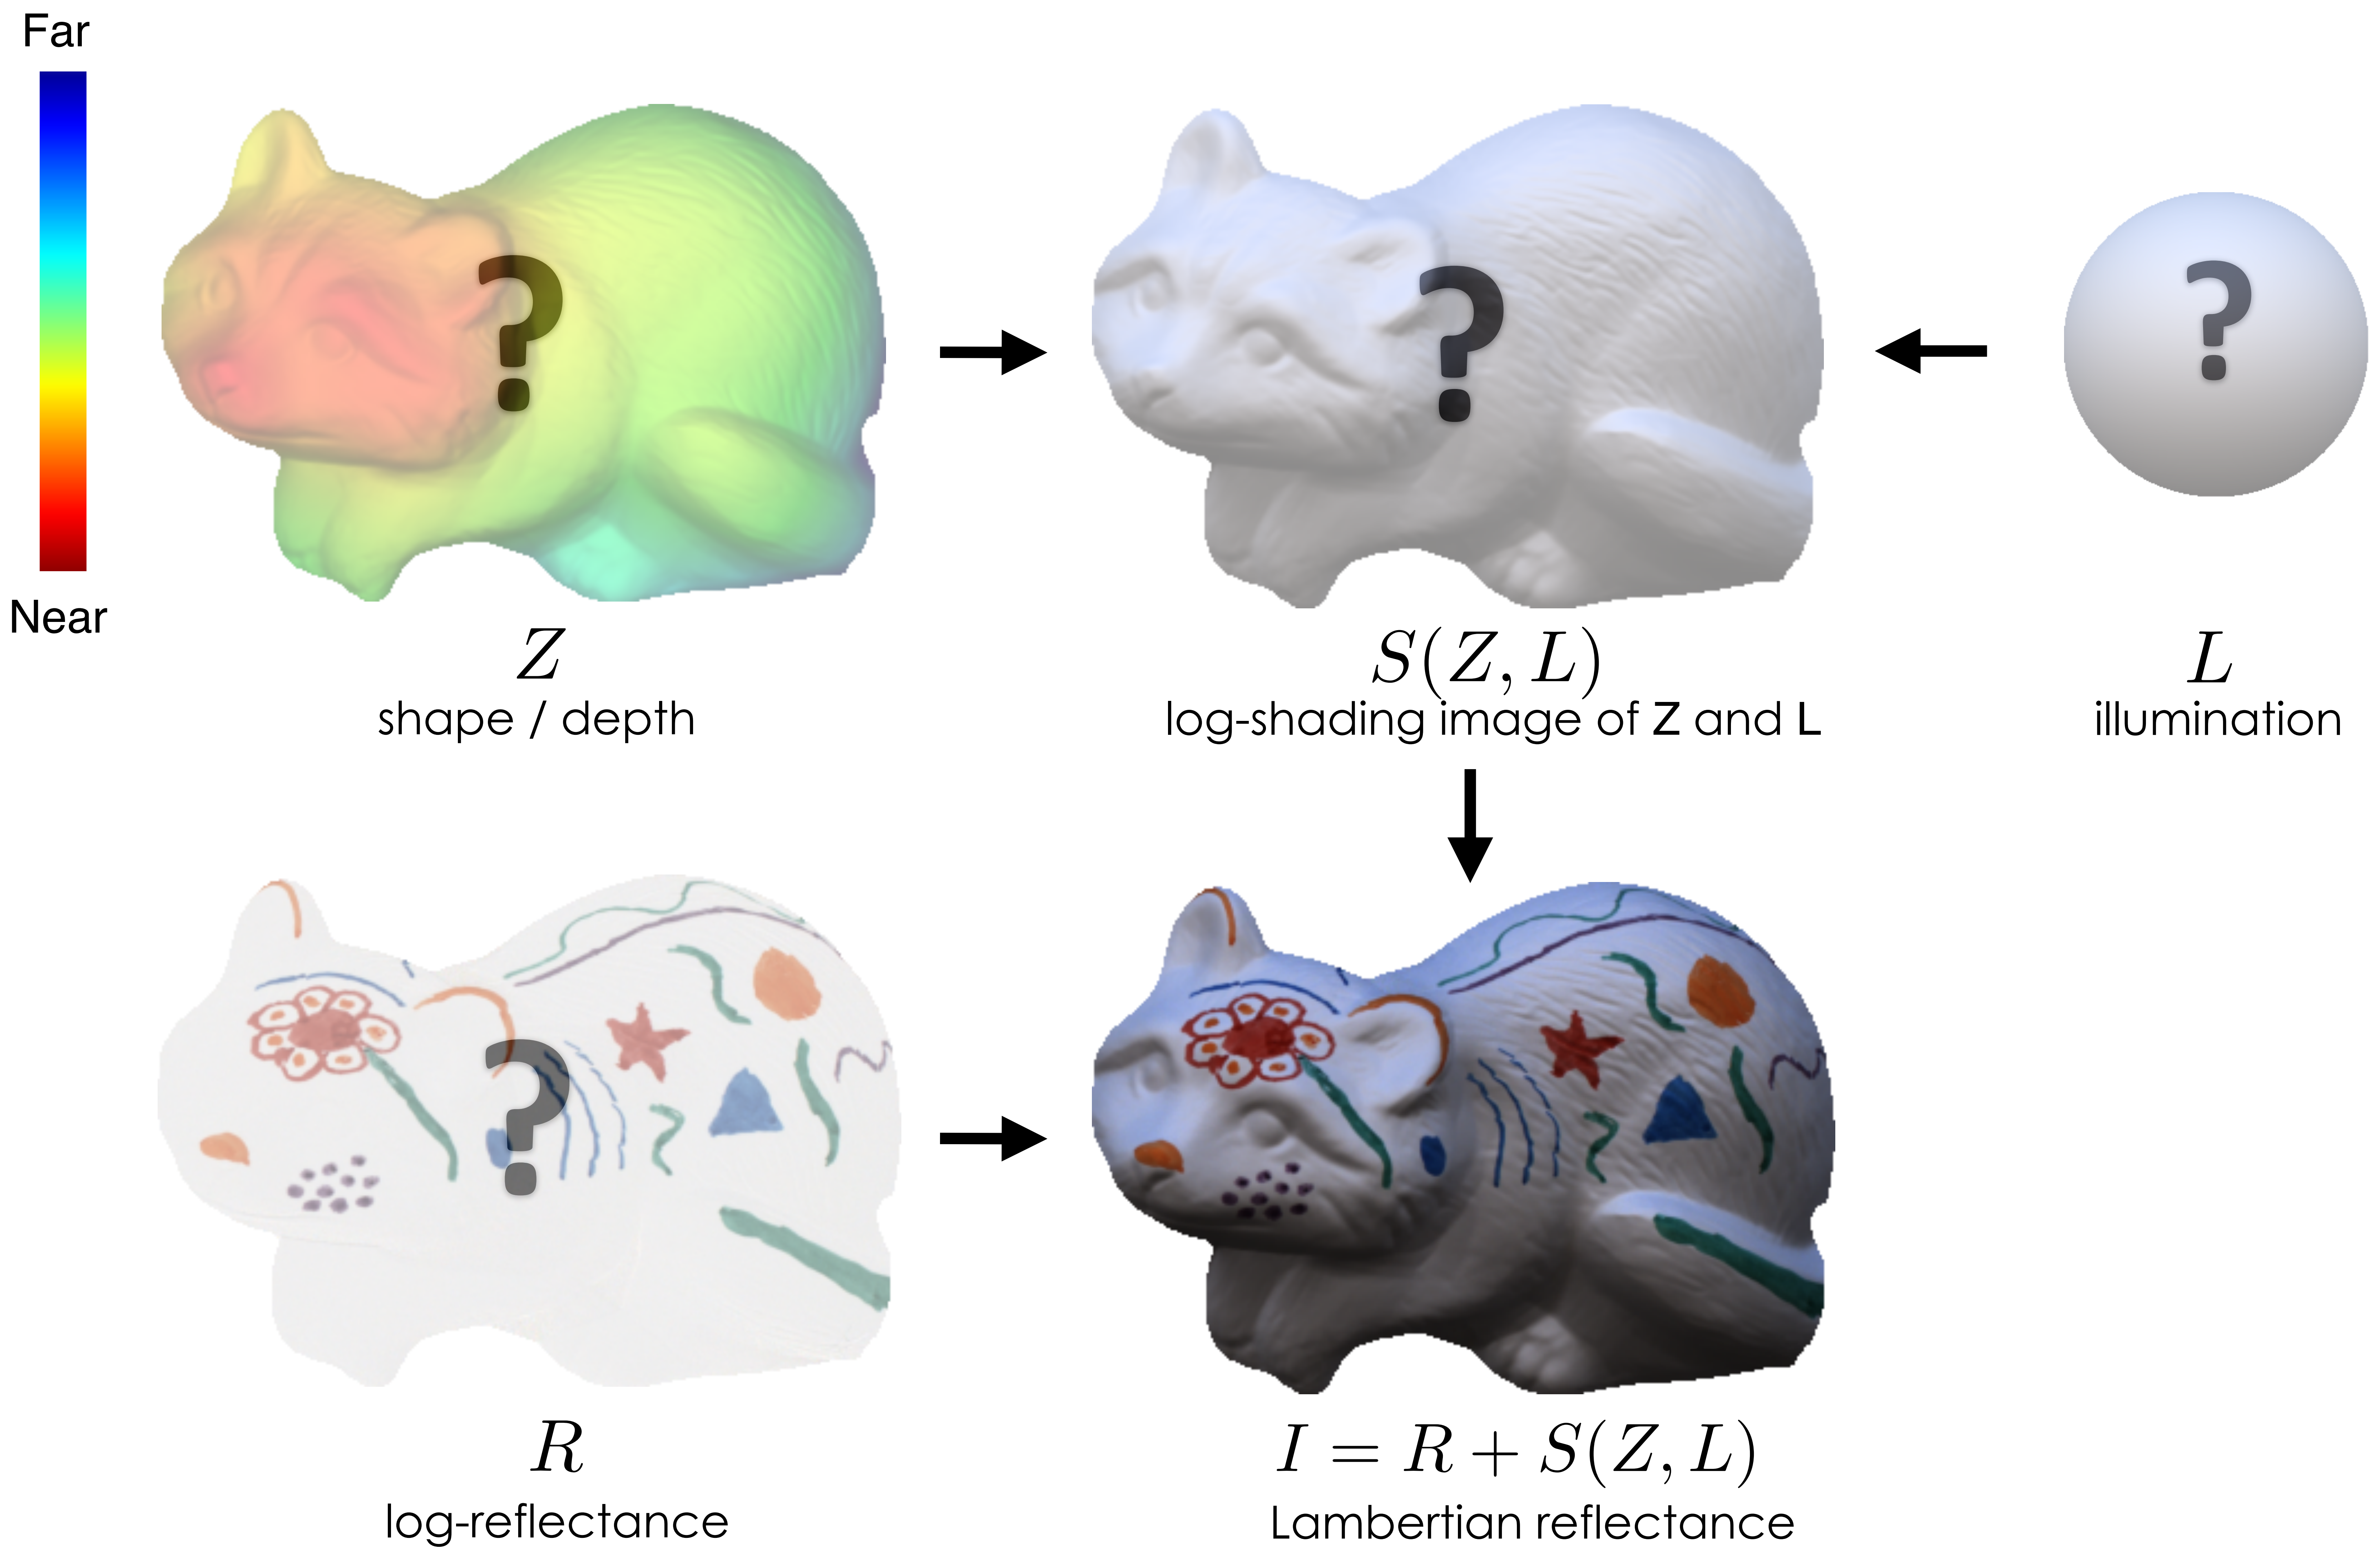

[Barrow and Tenenbaum 1978]

Source: J. Barron

Retinex: keep only large magnitude edges

Input

Estimated reflectance

[Land & McCann 1978], figure from [Grosse et al. 2019]

CNN-based reflectance estimation

[Bell et al., "Intrinsic images in the wild", 2014]

Applications of intrinsic image decomposition

[Barron and Malik “SIRFS”, 2012]

Today

- Shape from shading
- Intrinsic image decomposition
- **Color perception**

Color names for cartoon spectra

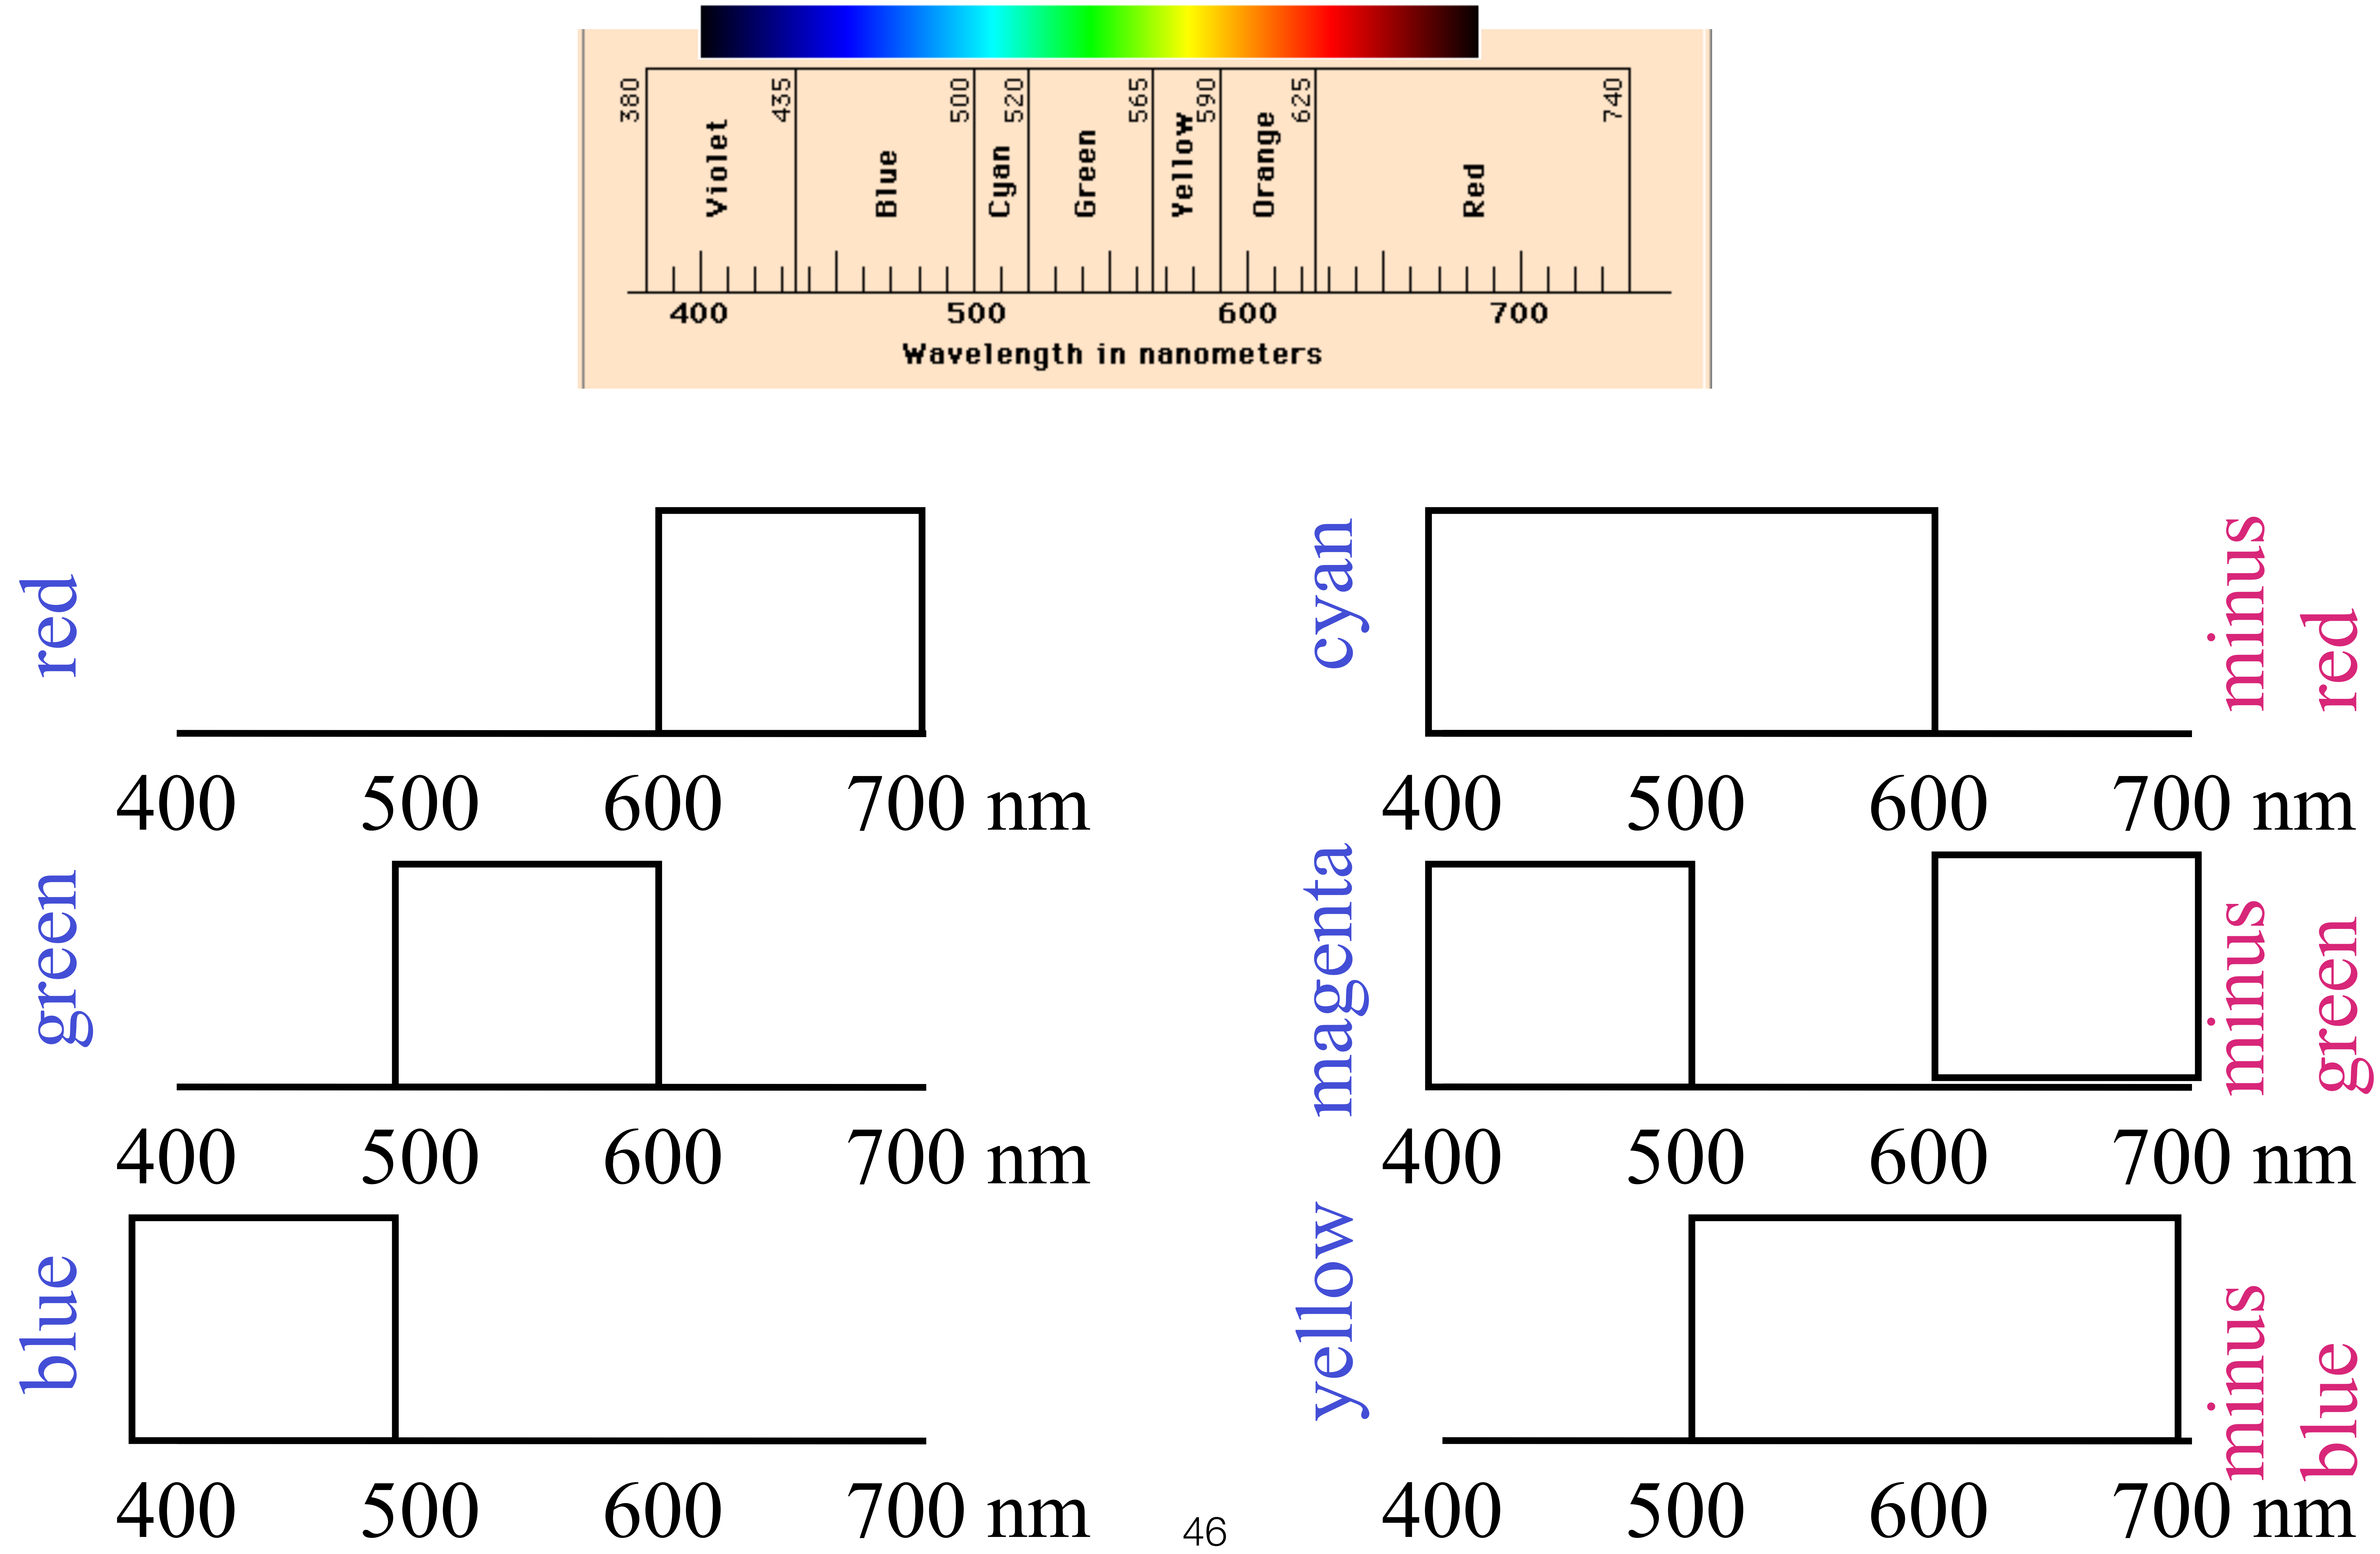

Source: W. Freeman

Additive color mixing

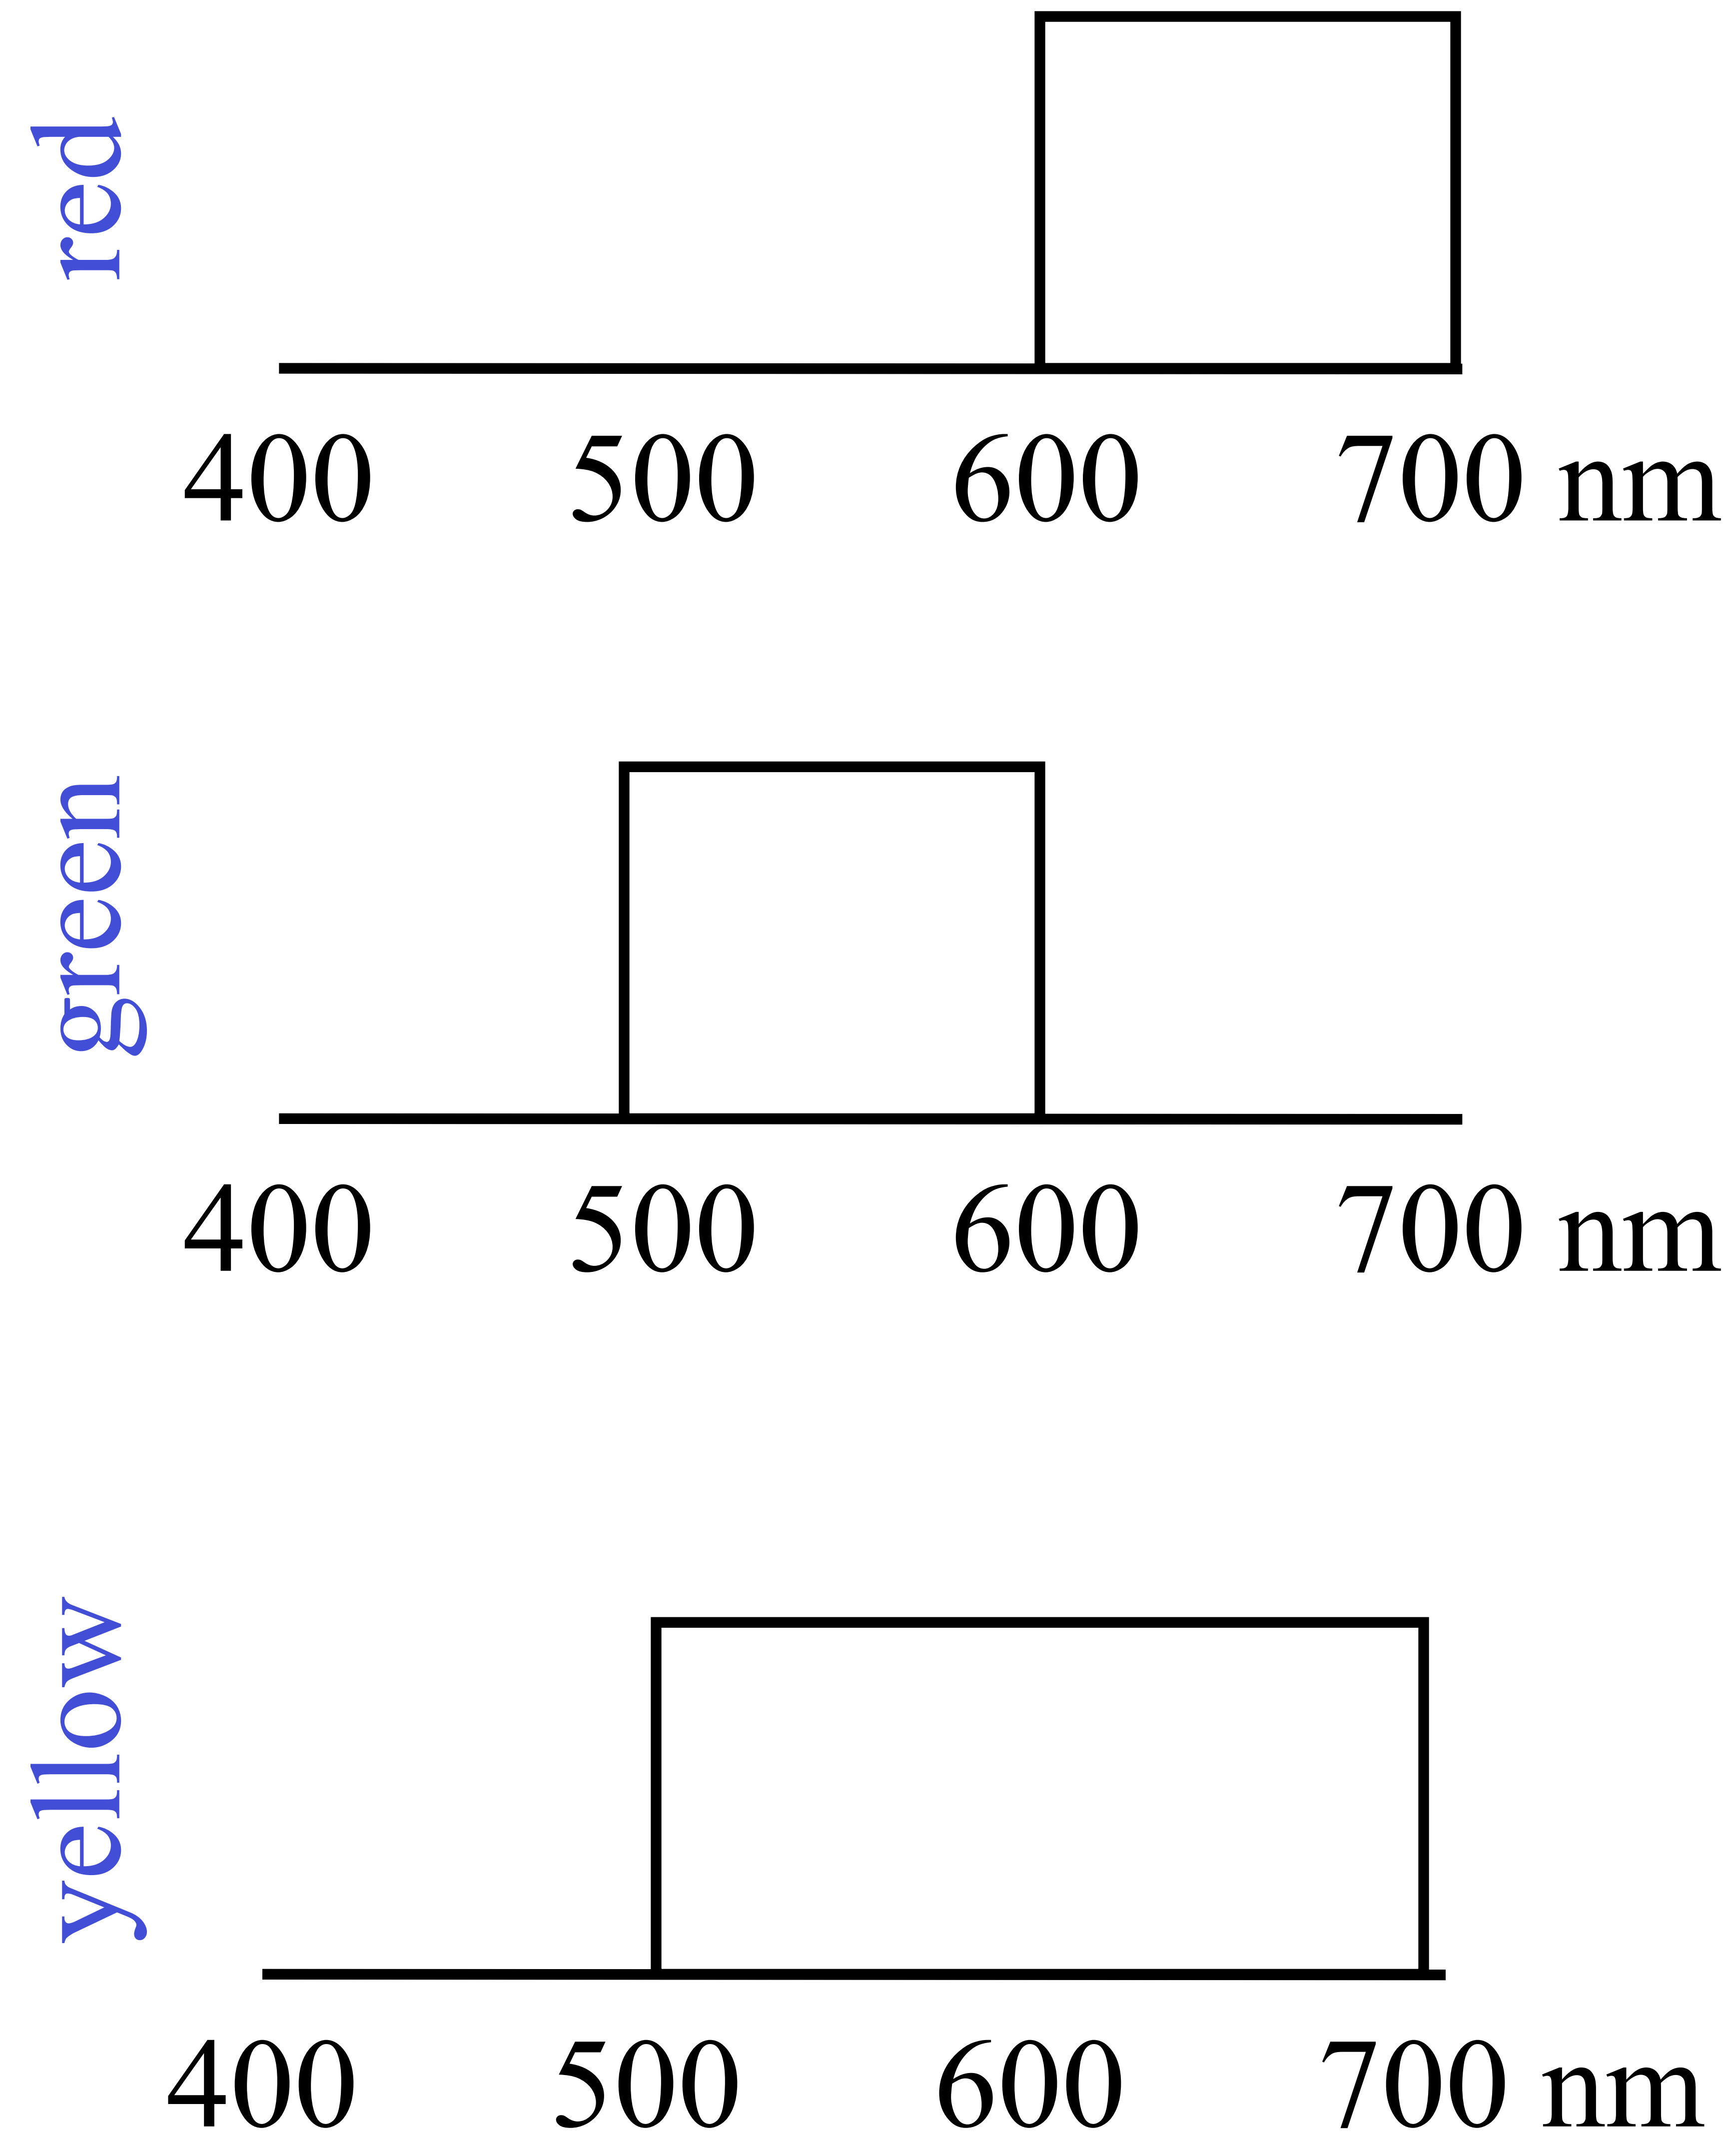

When colors combine by *adding* the color spectra. Example color displays that follow this mixing rule: CRT phosphors, multiple projectors aimed at a screen.

Red and green make…

Yellow!

Source: W. Freeman

How is color perceived in the eye?

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

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

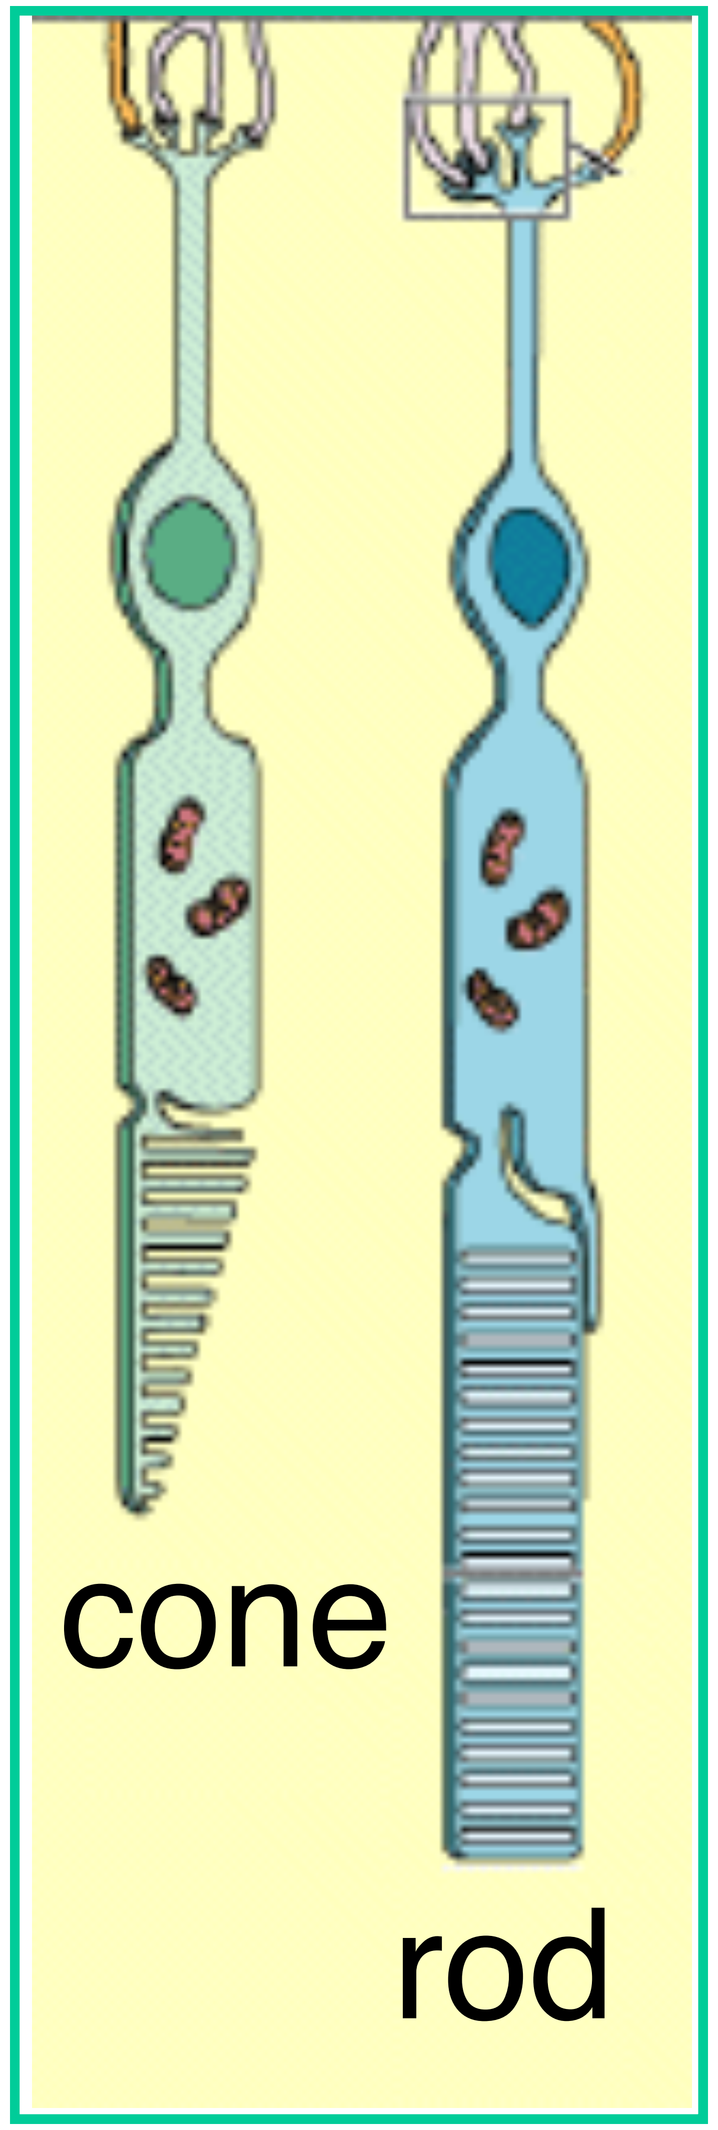

Source: A. Efros ?

Distribution of Rods and Cones

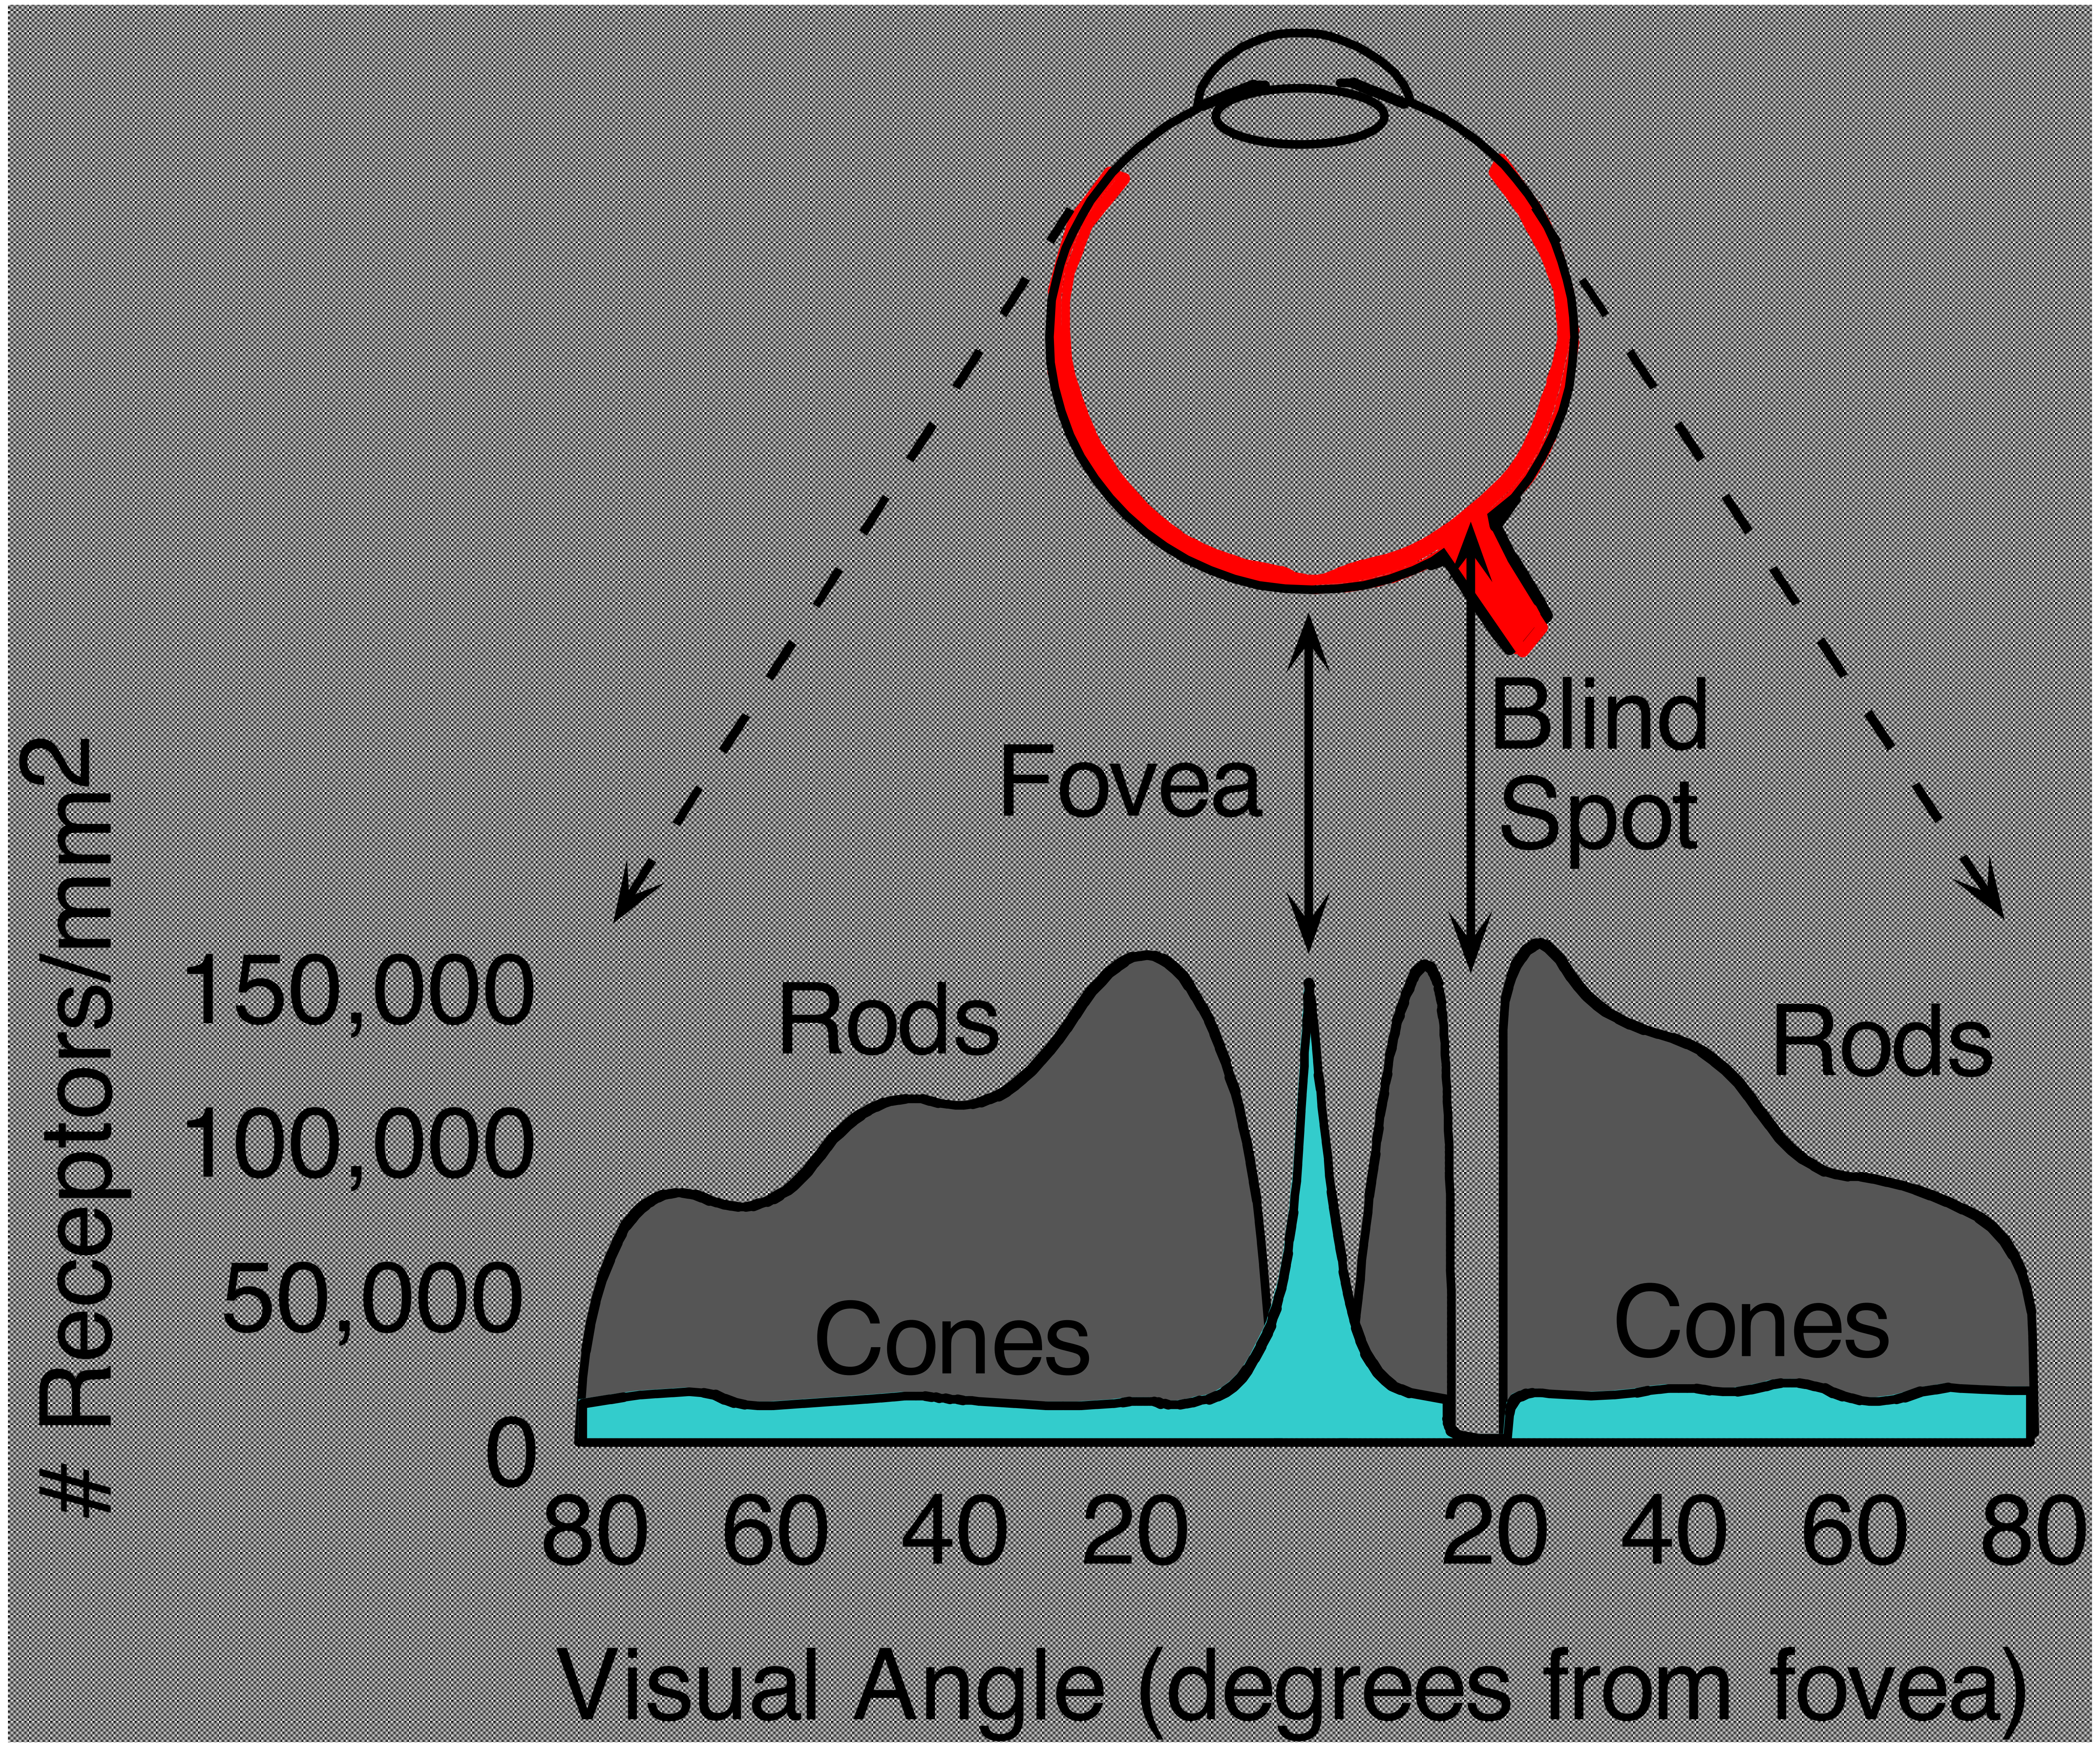

Source: A. Efros ?

Human eye photoreceptor spectral sensitivities

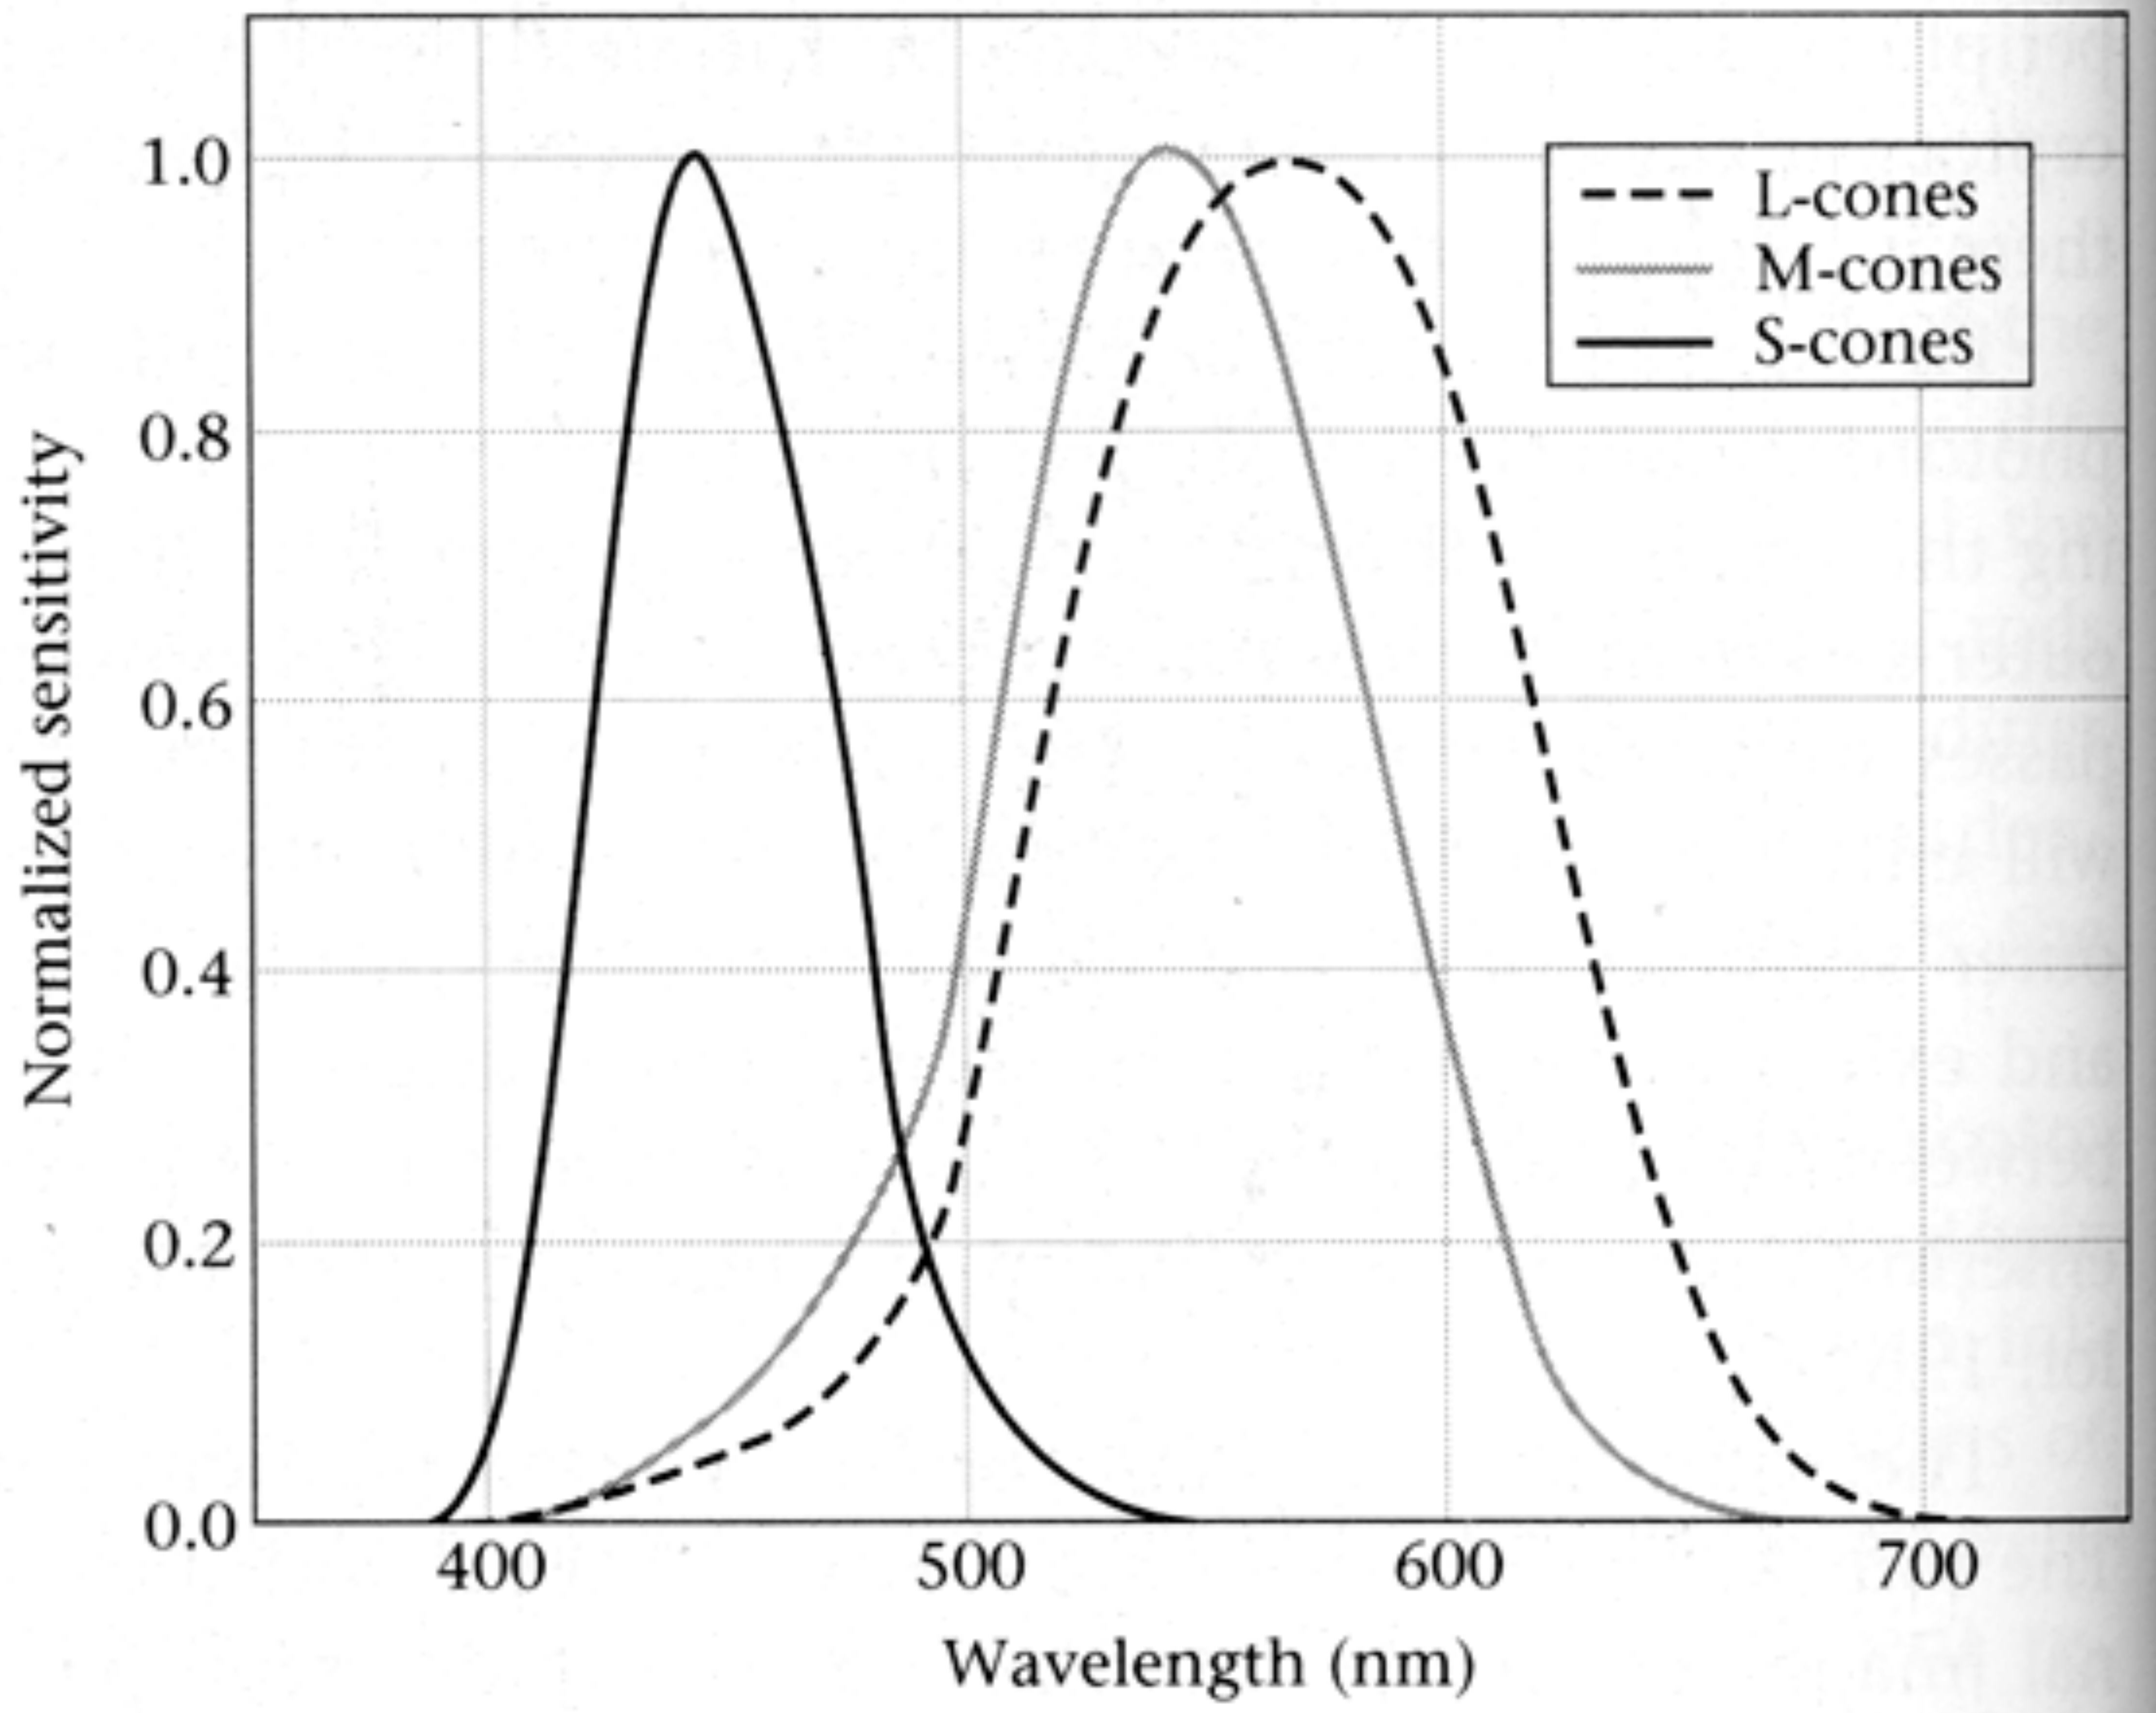

3.3 SPECTRAL SENSITIVITIES OF THE L-, M-, AND S-CONES in the human eye. The measurements are based on a light source at the cornea, so that the wavelength loss due to the cornea, lens, and other inert pigments of the eye plays a role in determining the sensitivity. Source: Stockman and MacLeod, 1993.

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

Source: W. Freeman

How we sense light spectra

Cone sensitivities

test light

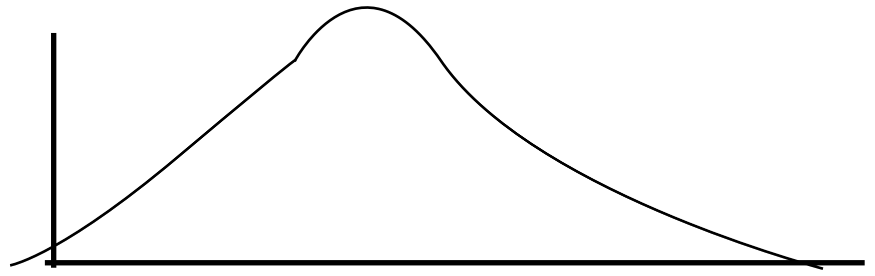

project

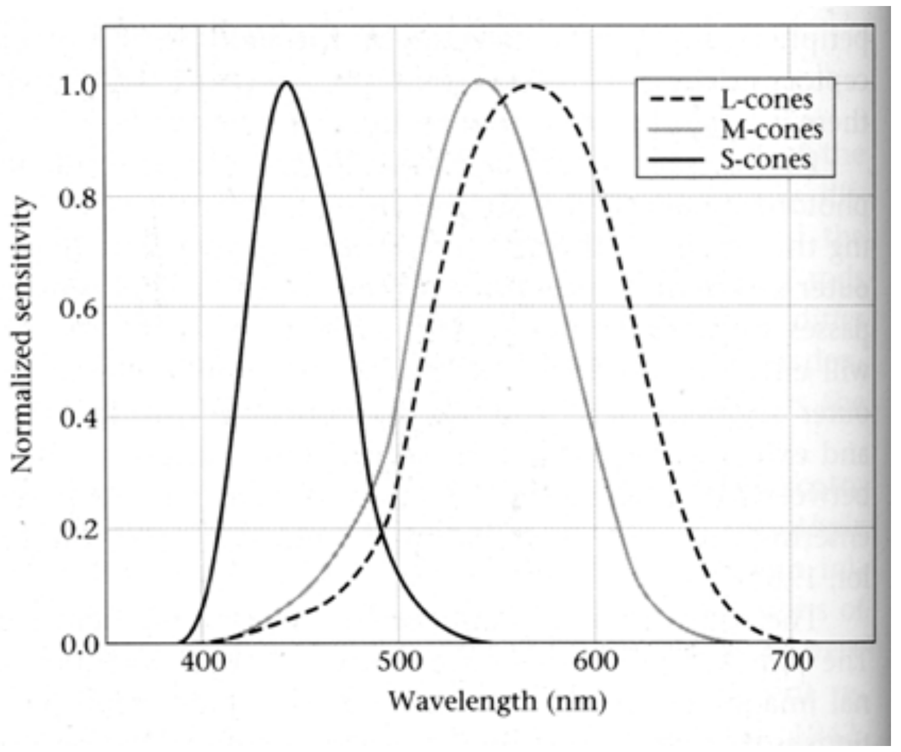

L, M, S responses

Biophysics: integrate the response over all wavelengths, weighted by the photosensor's sensitivity at each wavelength.

$$R = C\,t$$

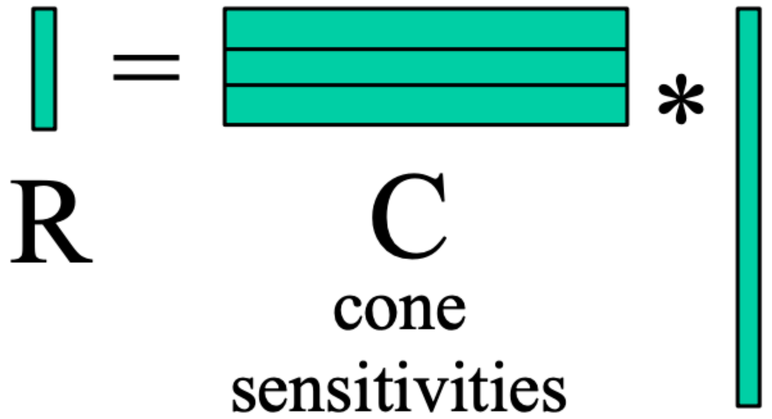

cone responses R

C cone sensitivities

t input spectrum

Mathematically: take dot product of input spectrum with the cone sensitivity basis vectors. Project the high-dimensional test light into a 3D space.

Source: W. Freeman

Cone response curves as basis vectors in a 3D subspace of light power spectra

3D depiction of the high-dimensional space of all possible power spectra

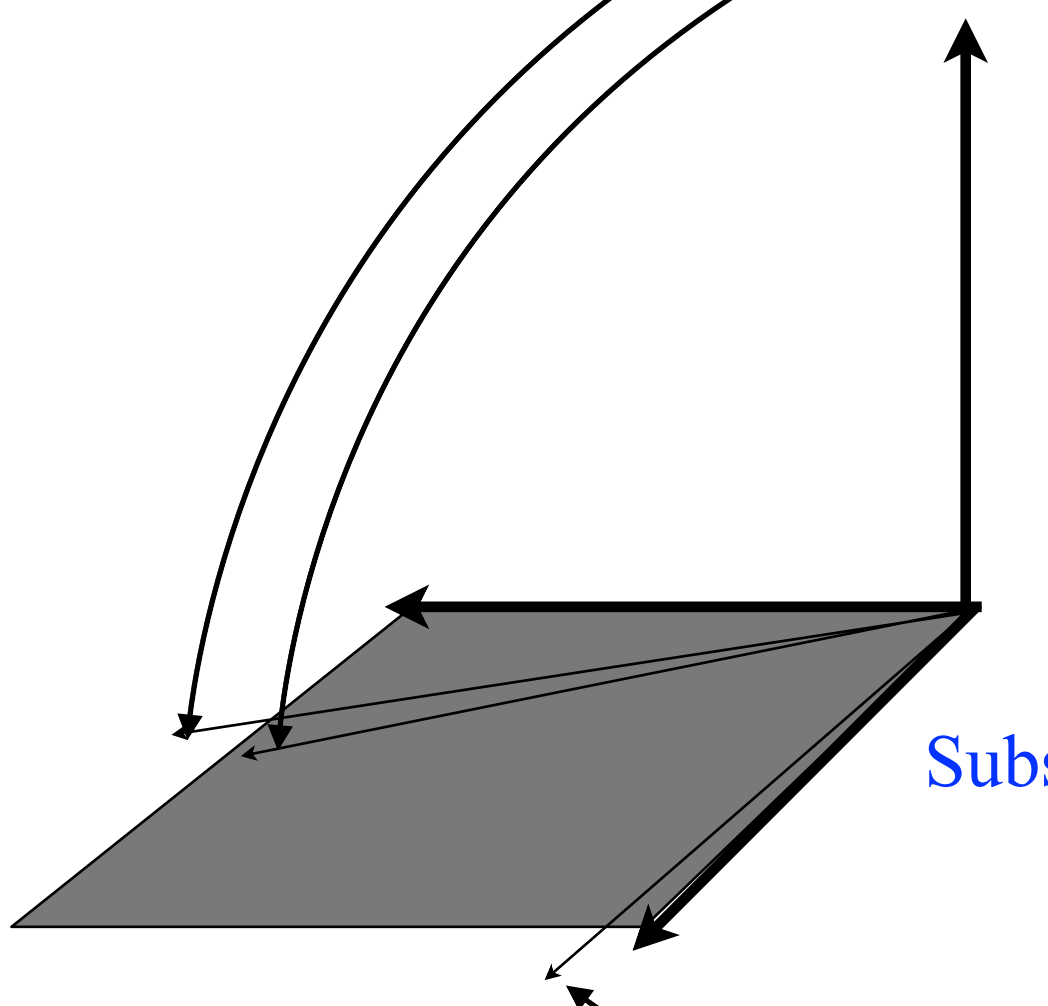

Subspace of sensor responses

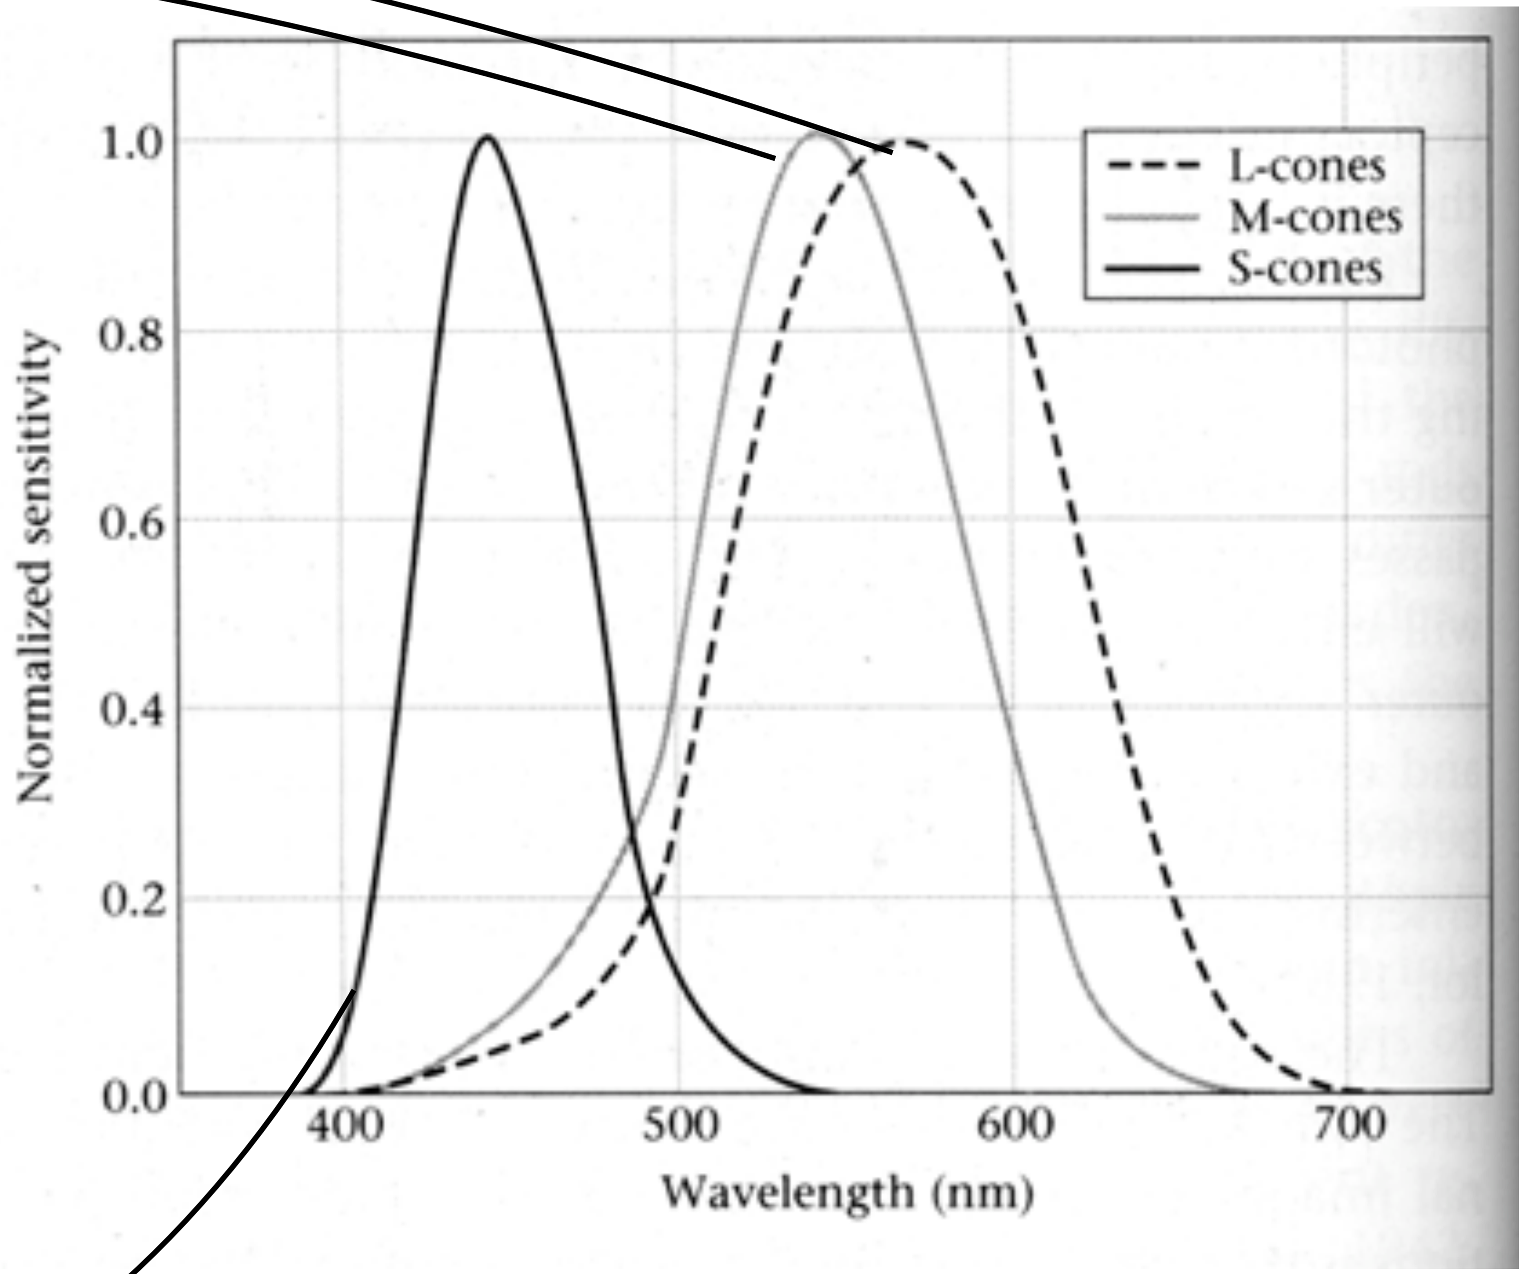

Spectral sensitivities of L, M, and S cones

Source: W. Freeman

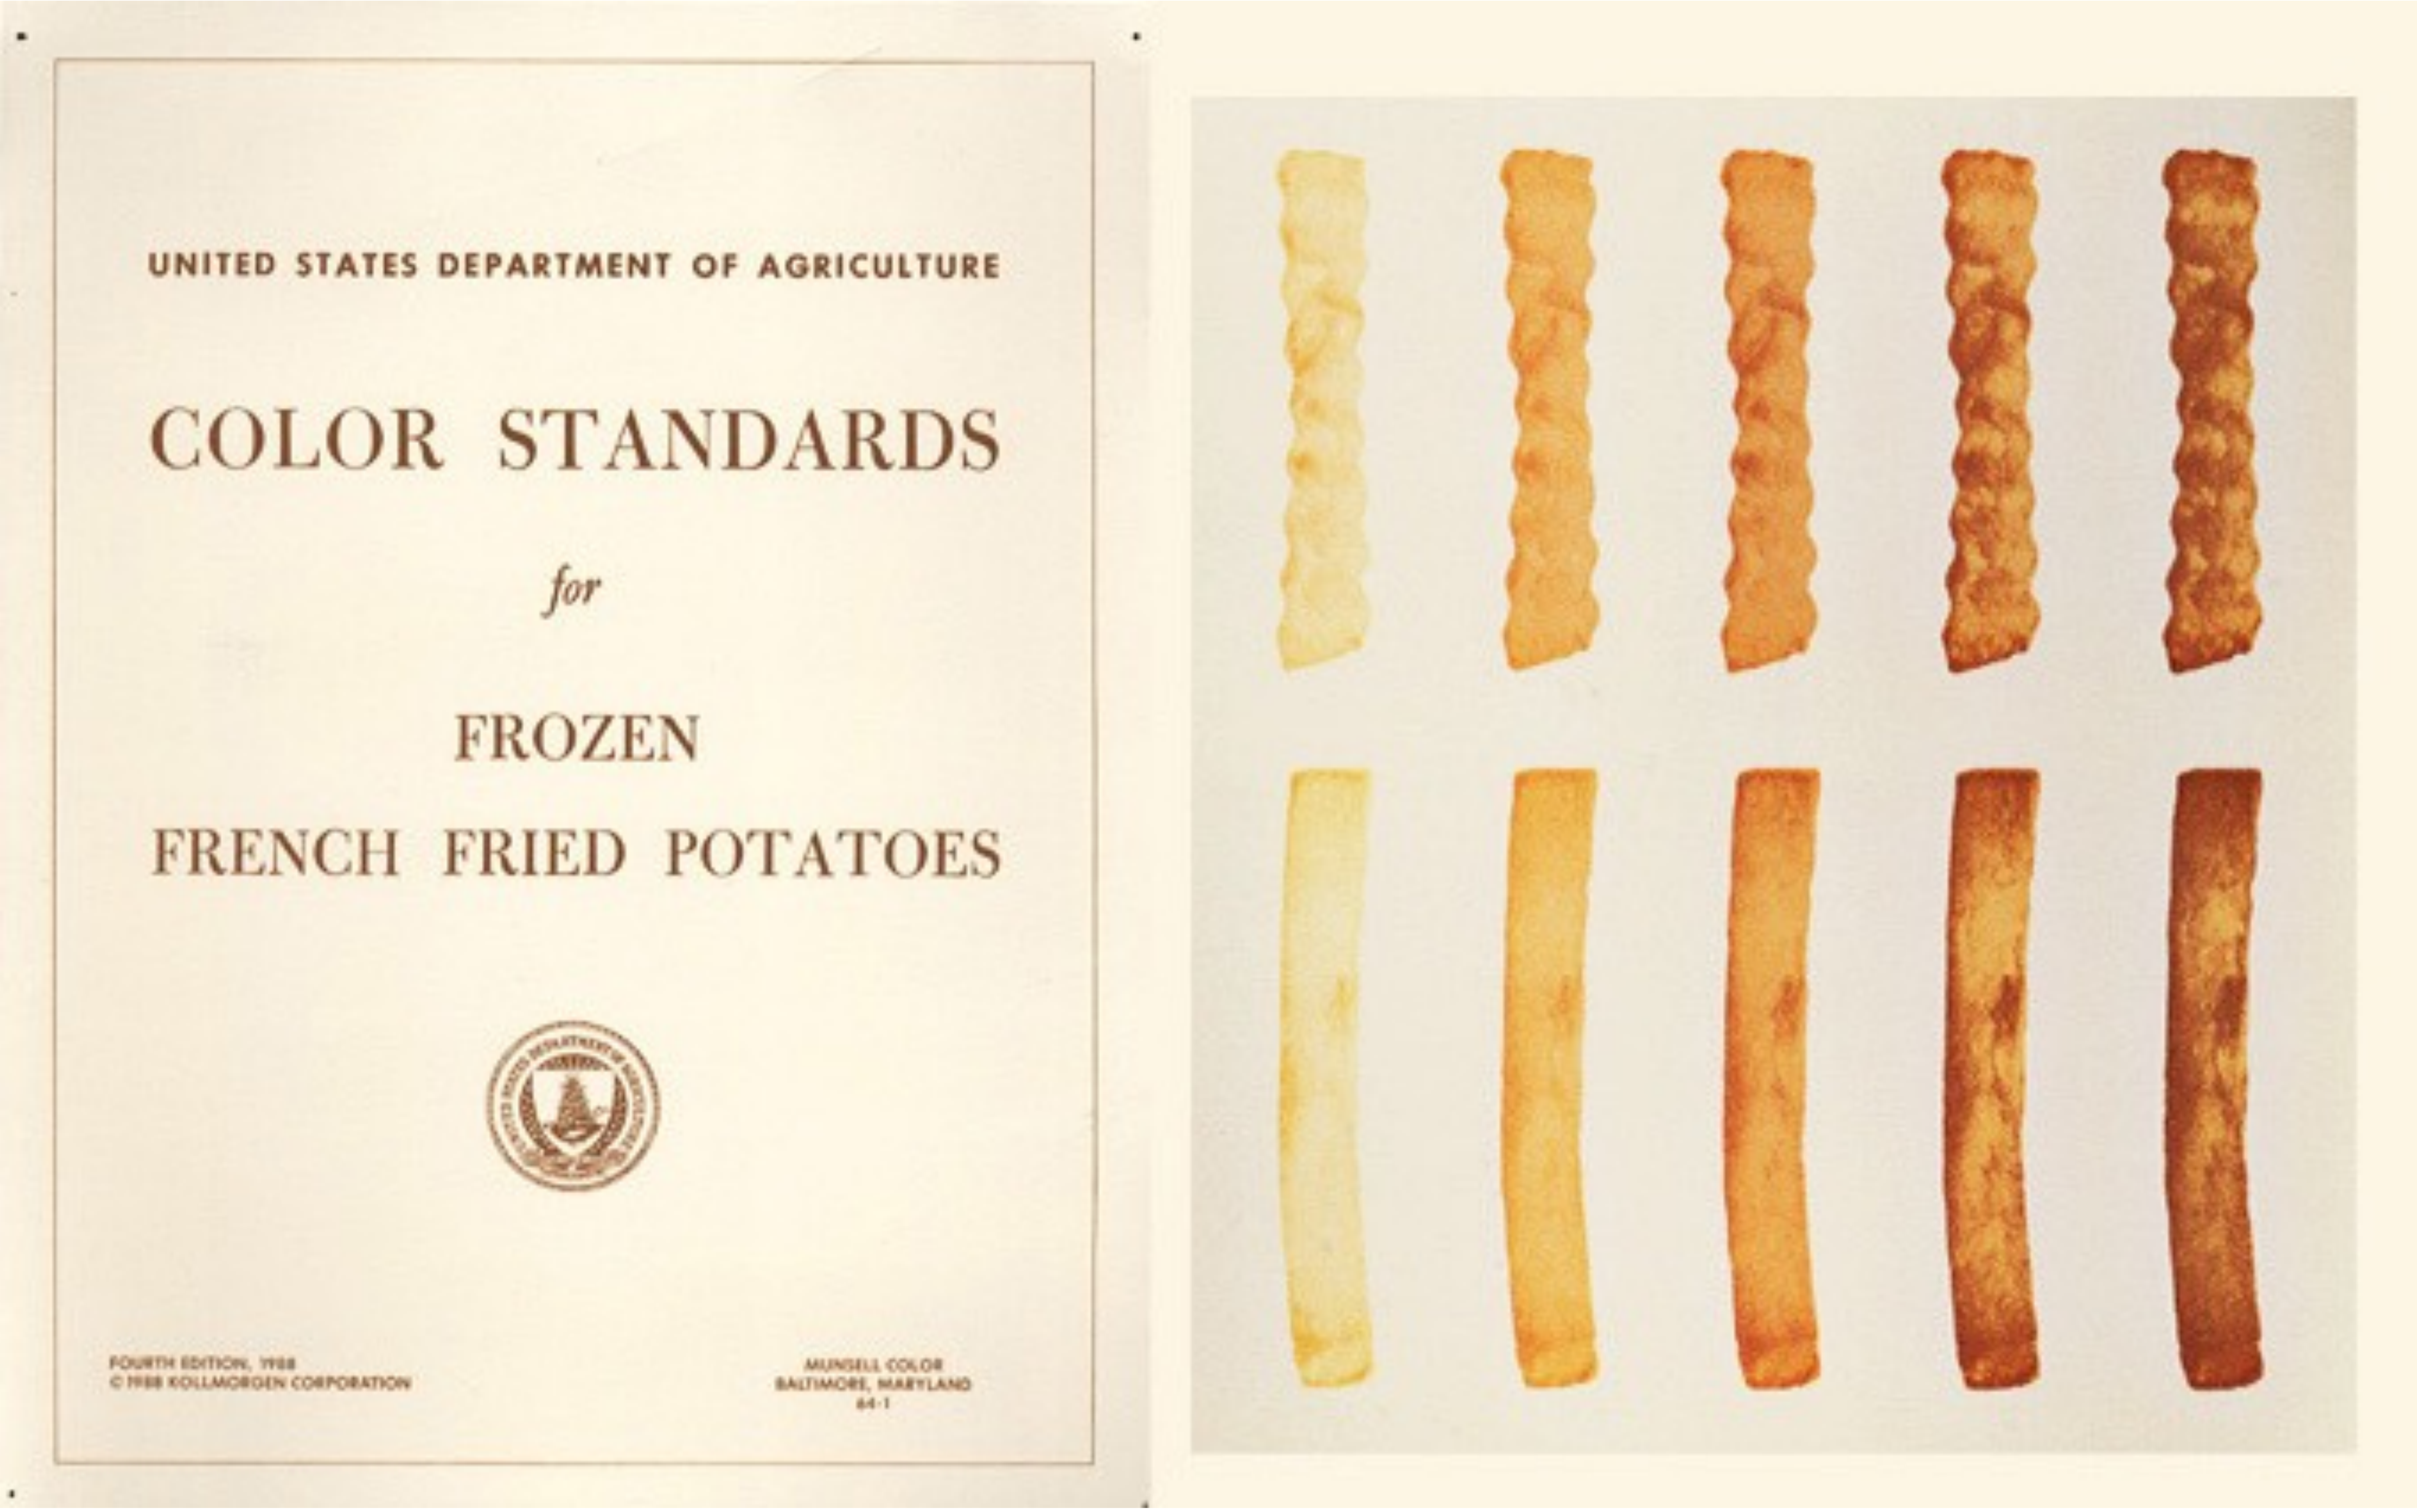

Source: W. Freeman

Color matching experiment

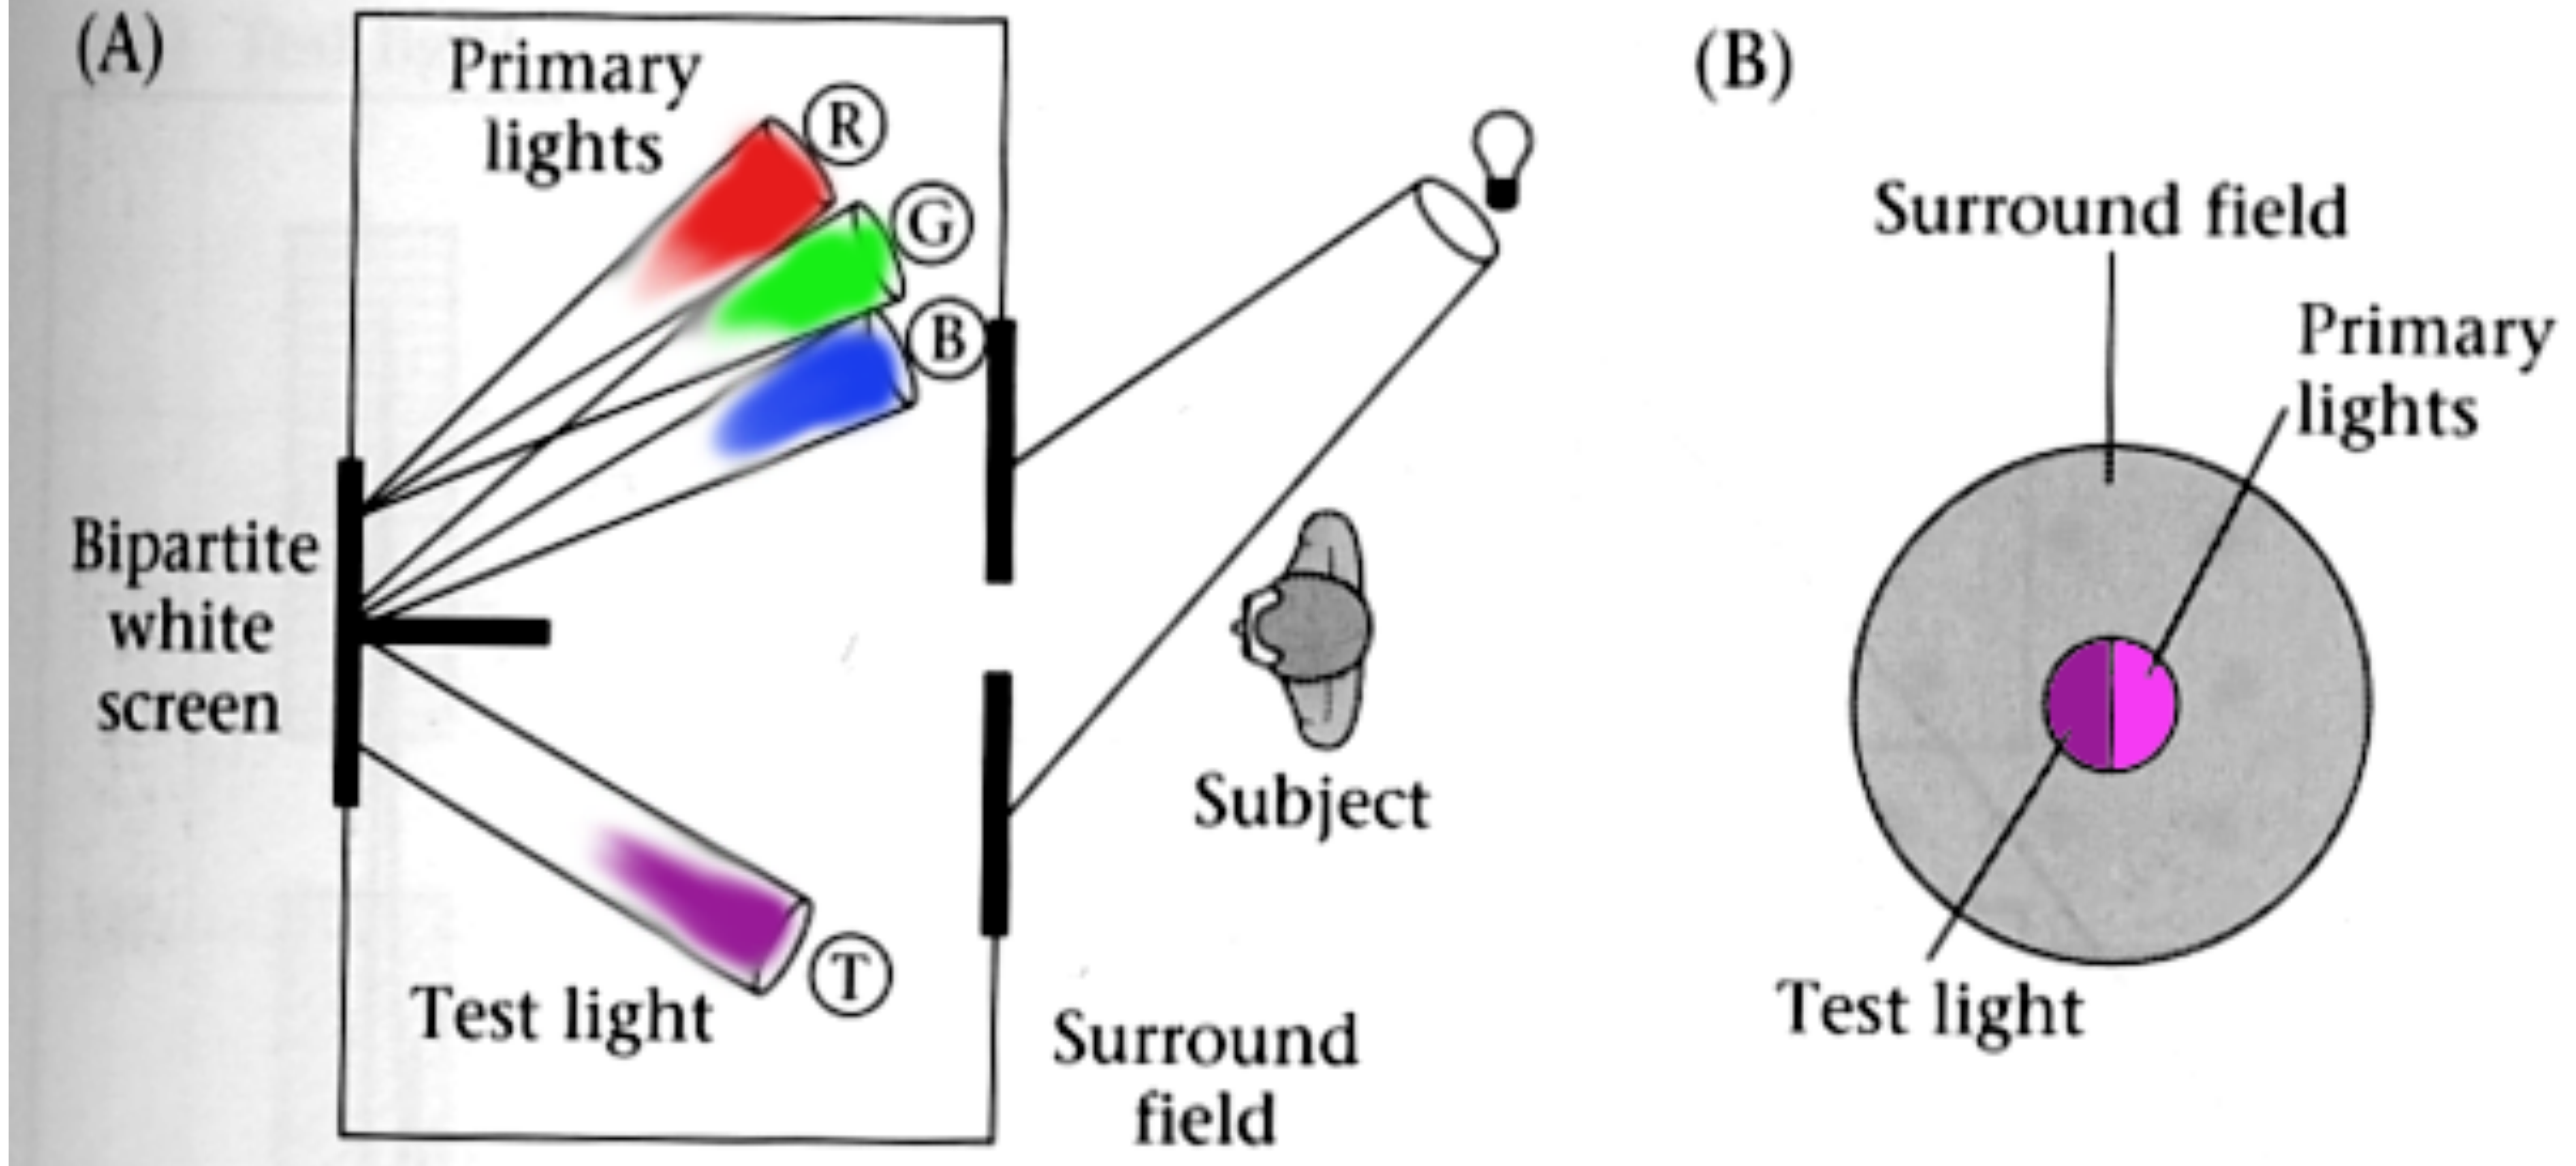

4.10 THE COLOR-MATCHING EXPERIMENT. The observer views a bipartite field and adjusts the intensities of the three primary lights to match the appearance of the test light. (A) A top view of the experimental apparatus. (B) The appearance of the stimuli to the observer. After Judd and Wyszecki, 1975.

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

Source: W. Freeman

Color matching experiment

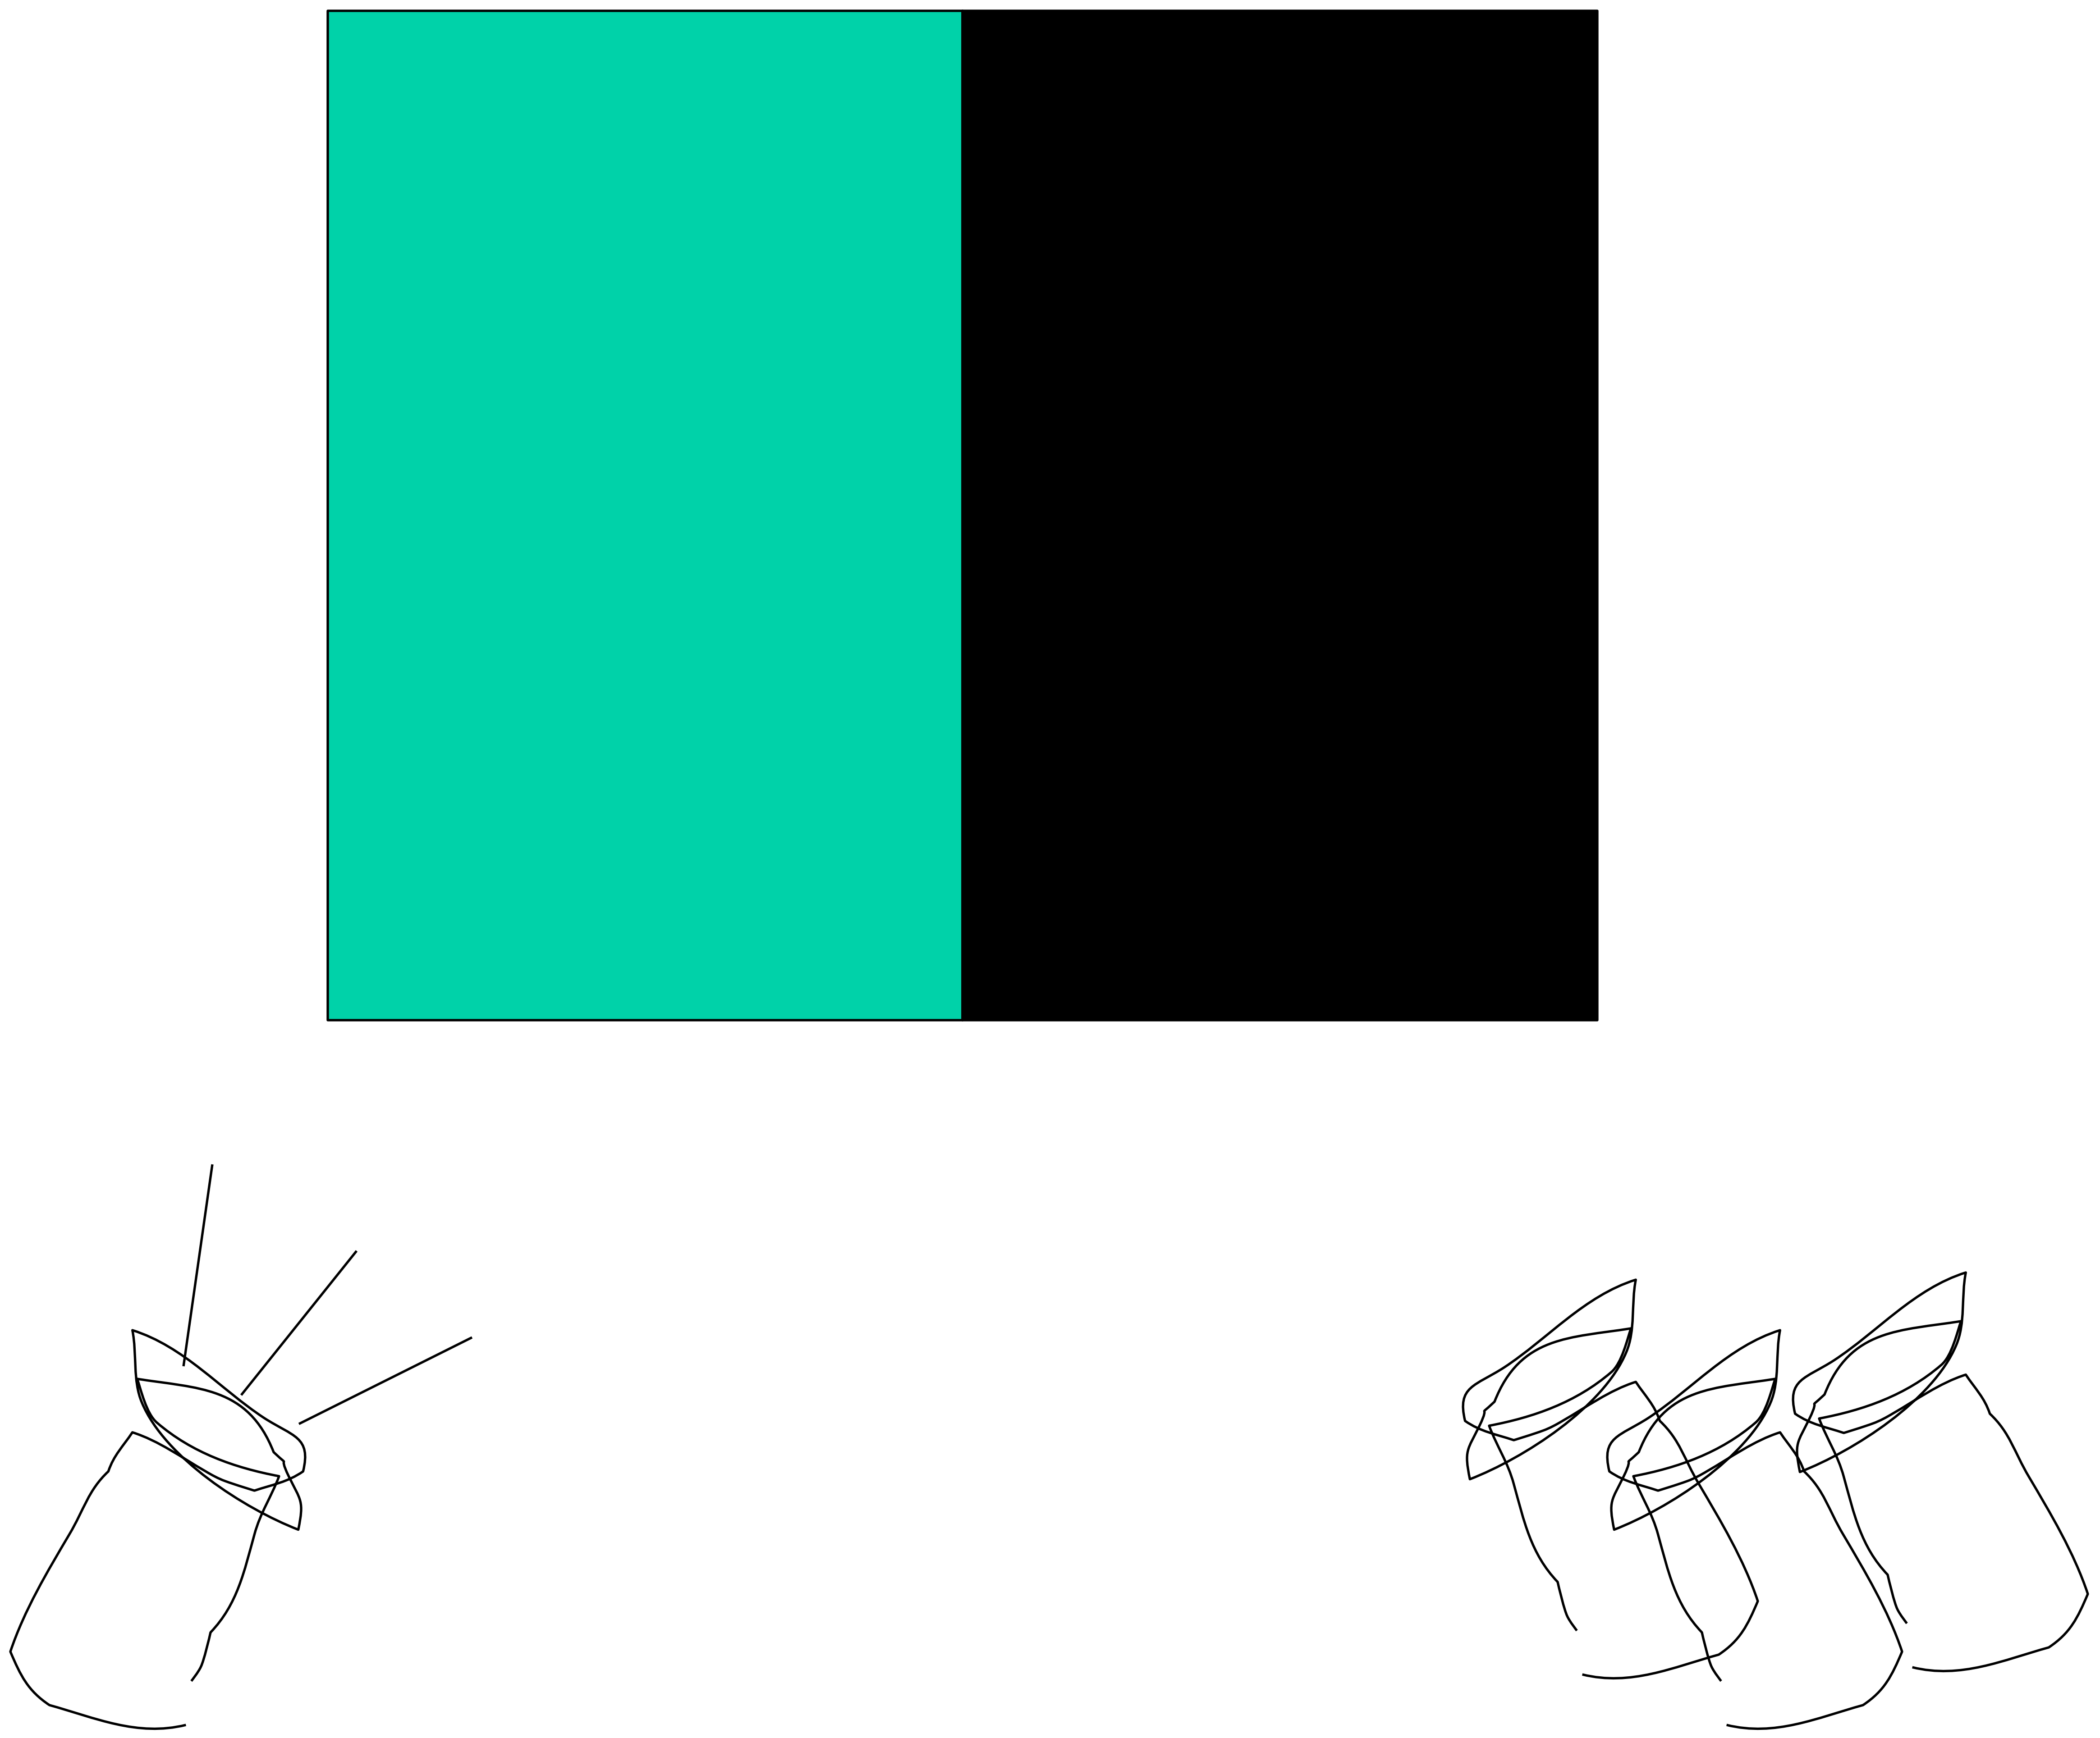

Source: W. Freeman

Color matching experiment

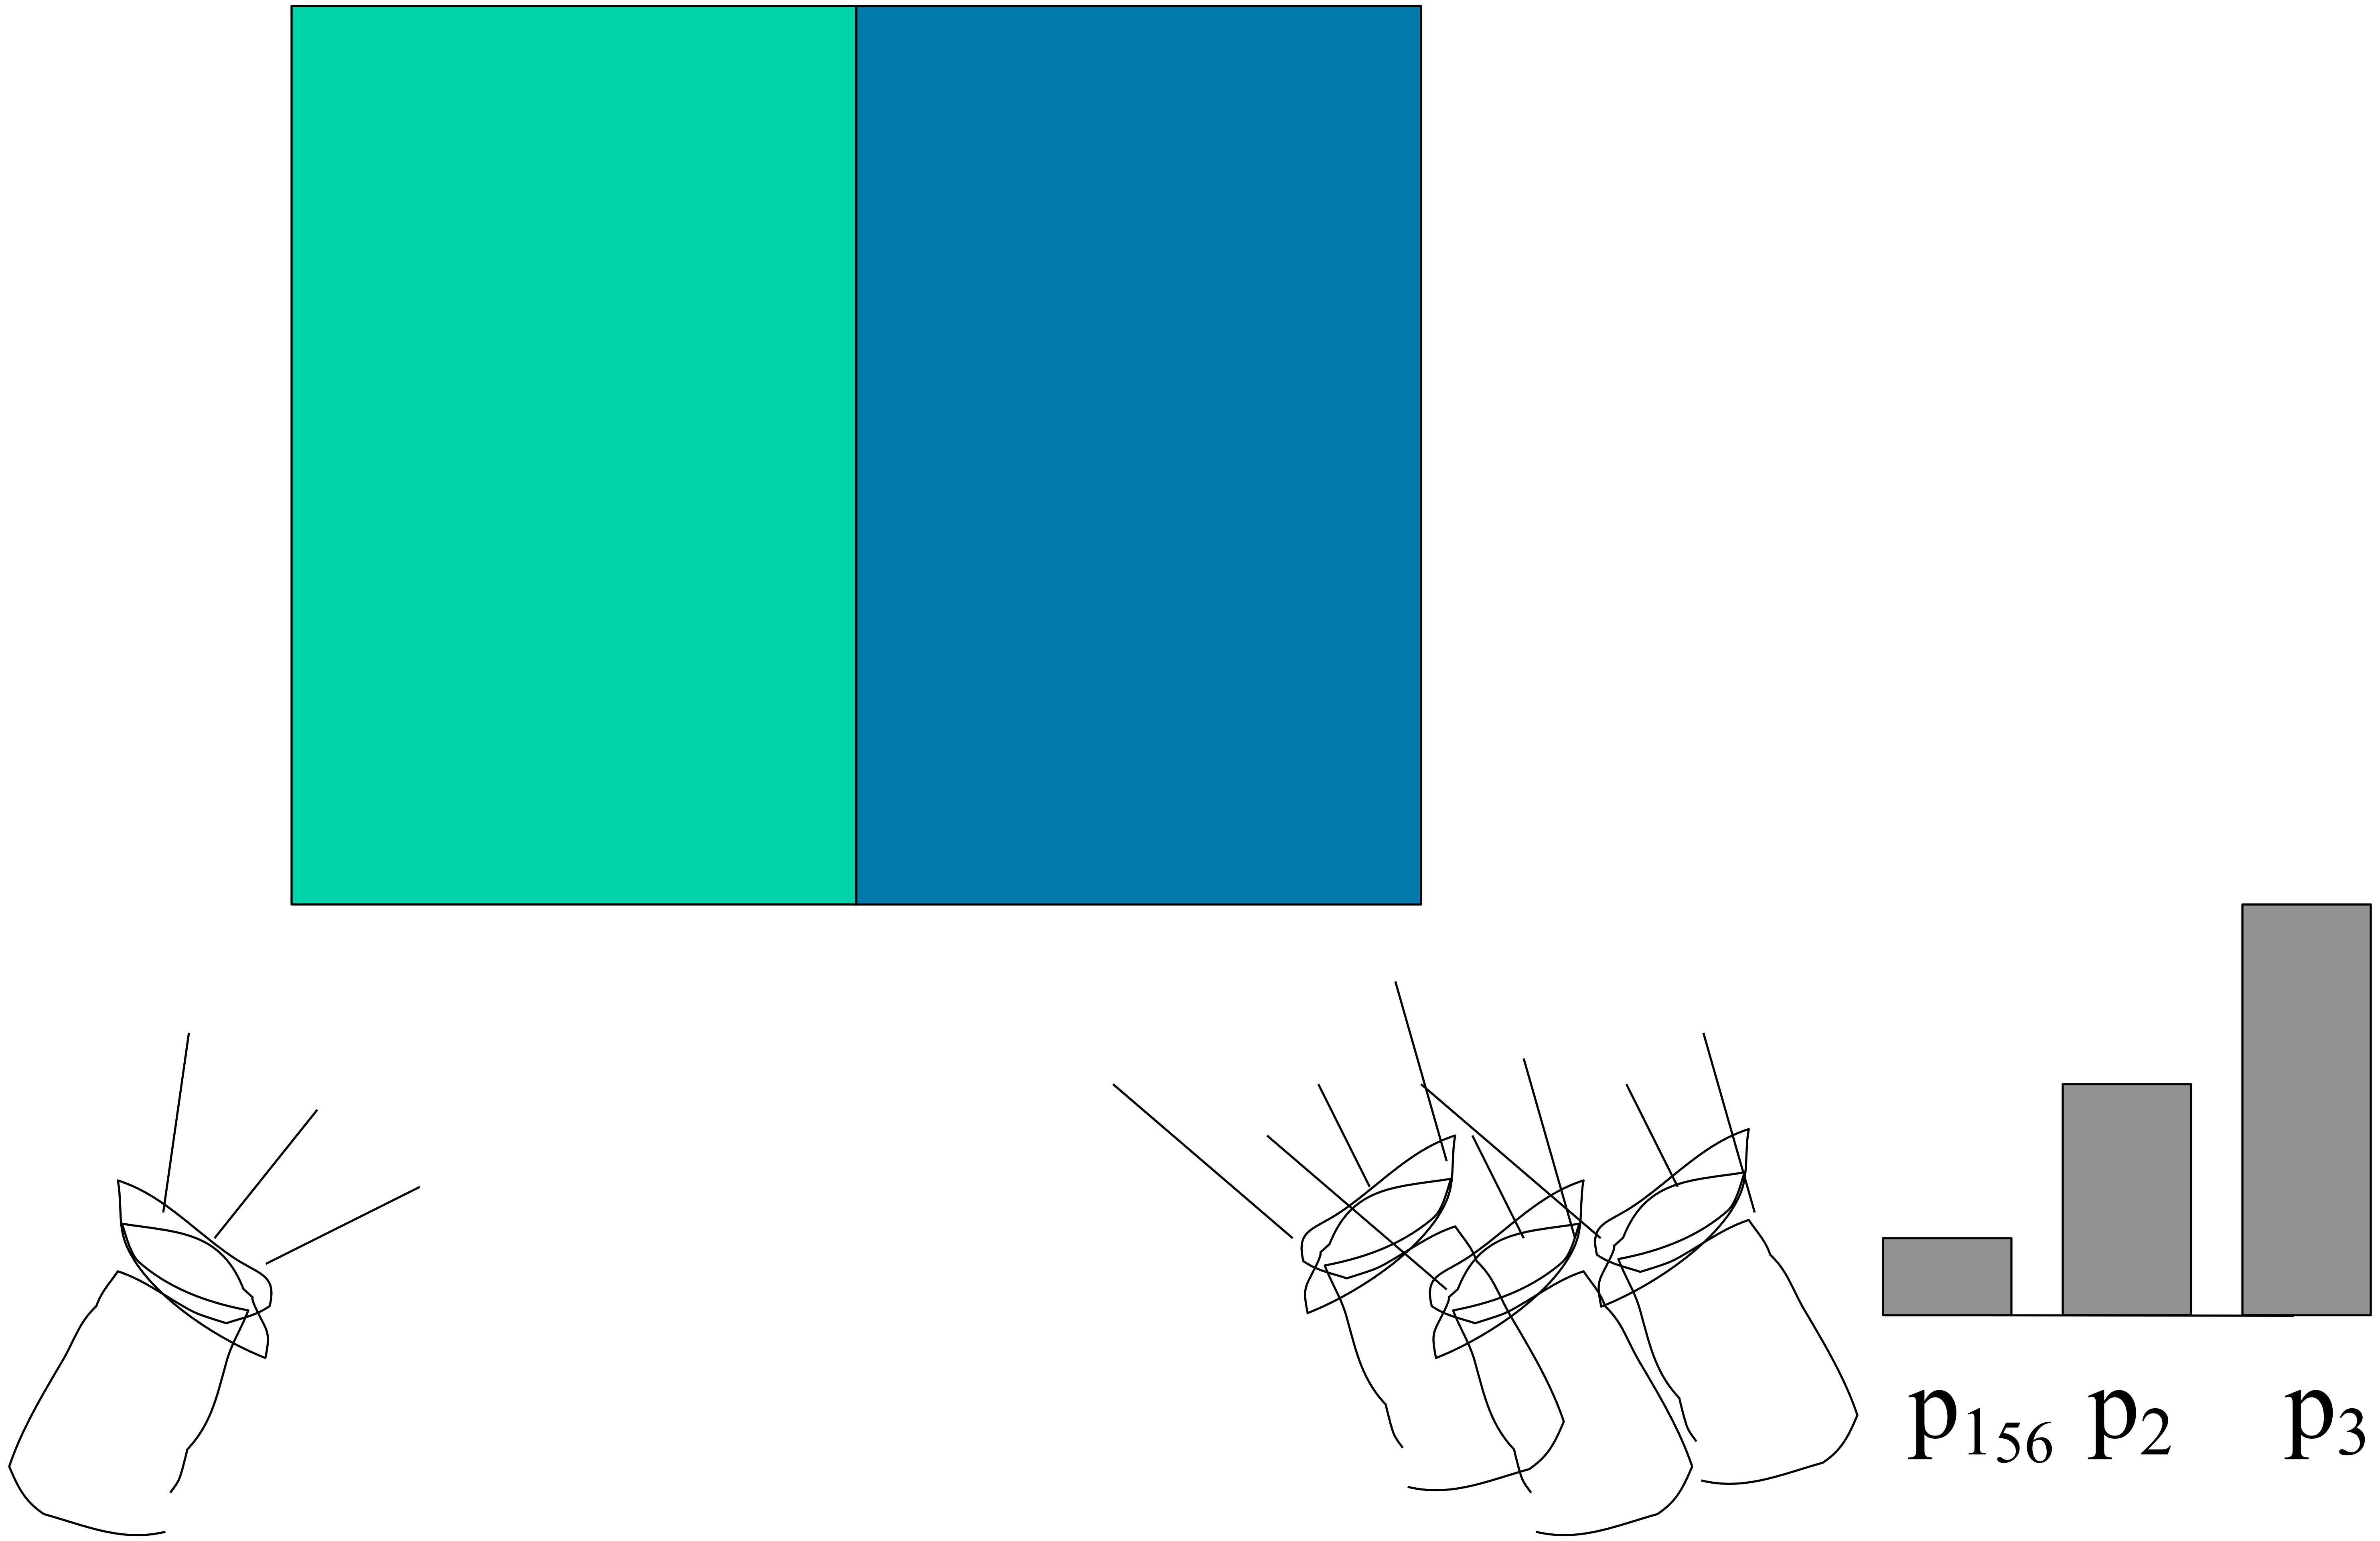

Source: W. Freeman

Color matching experiment

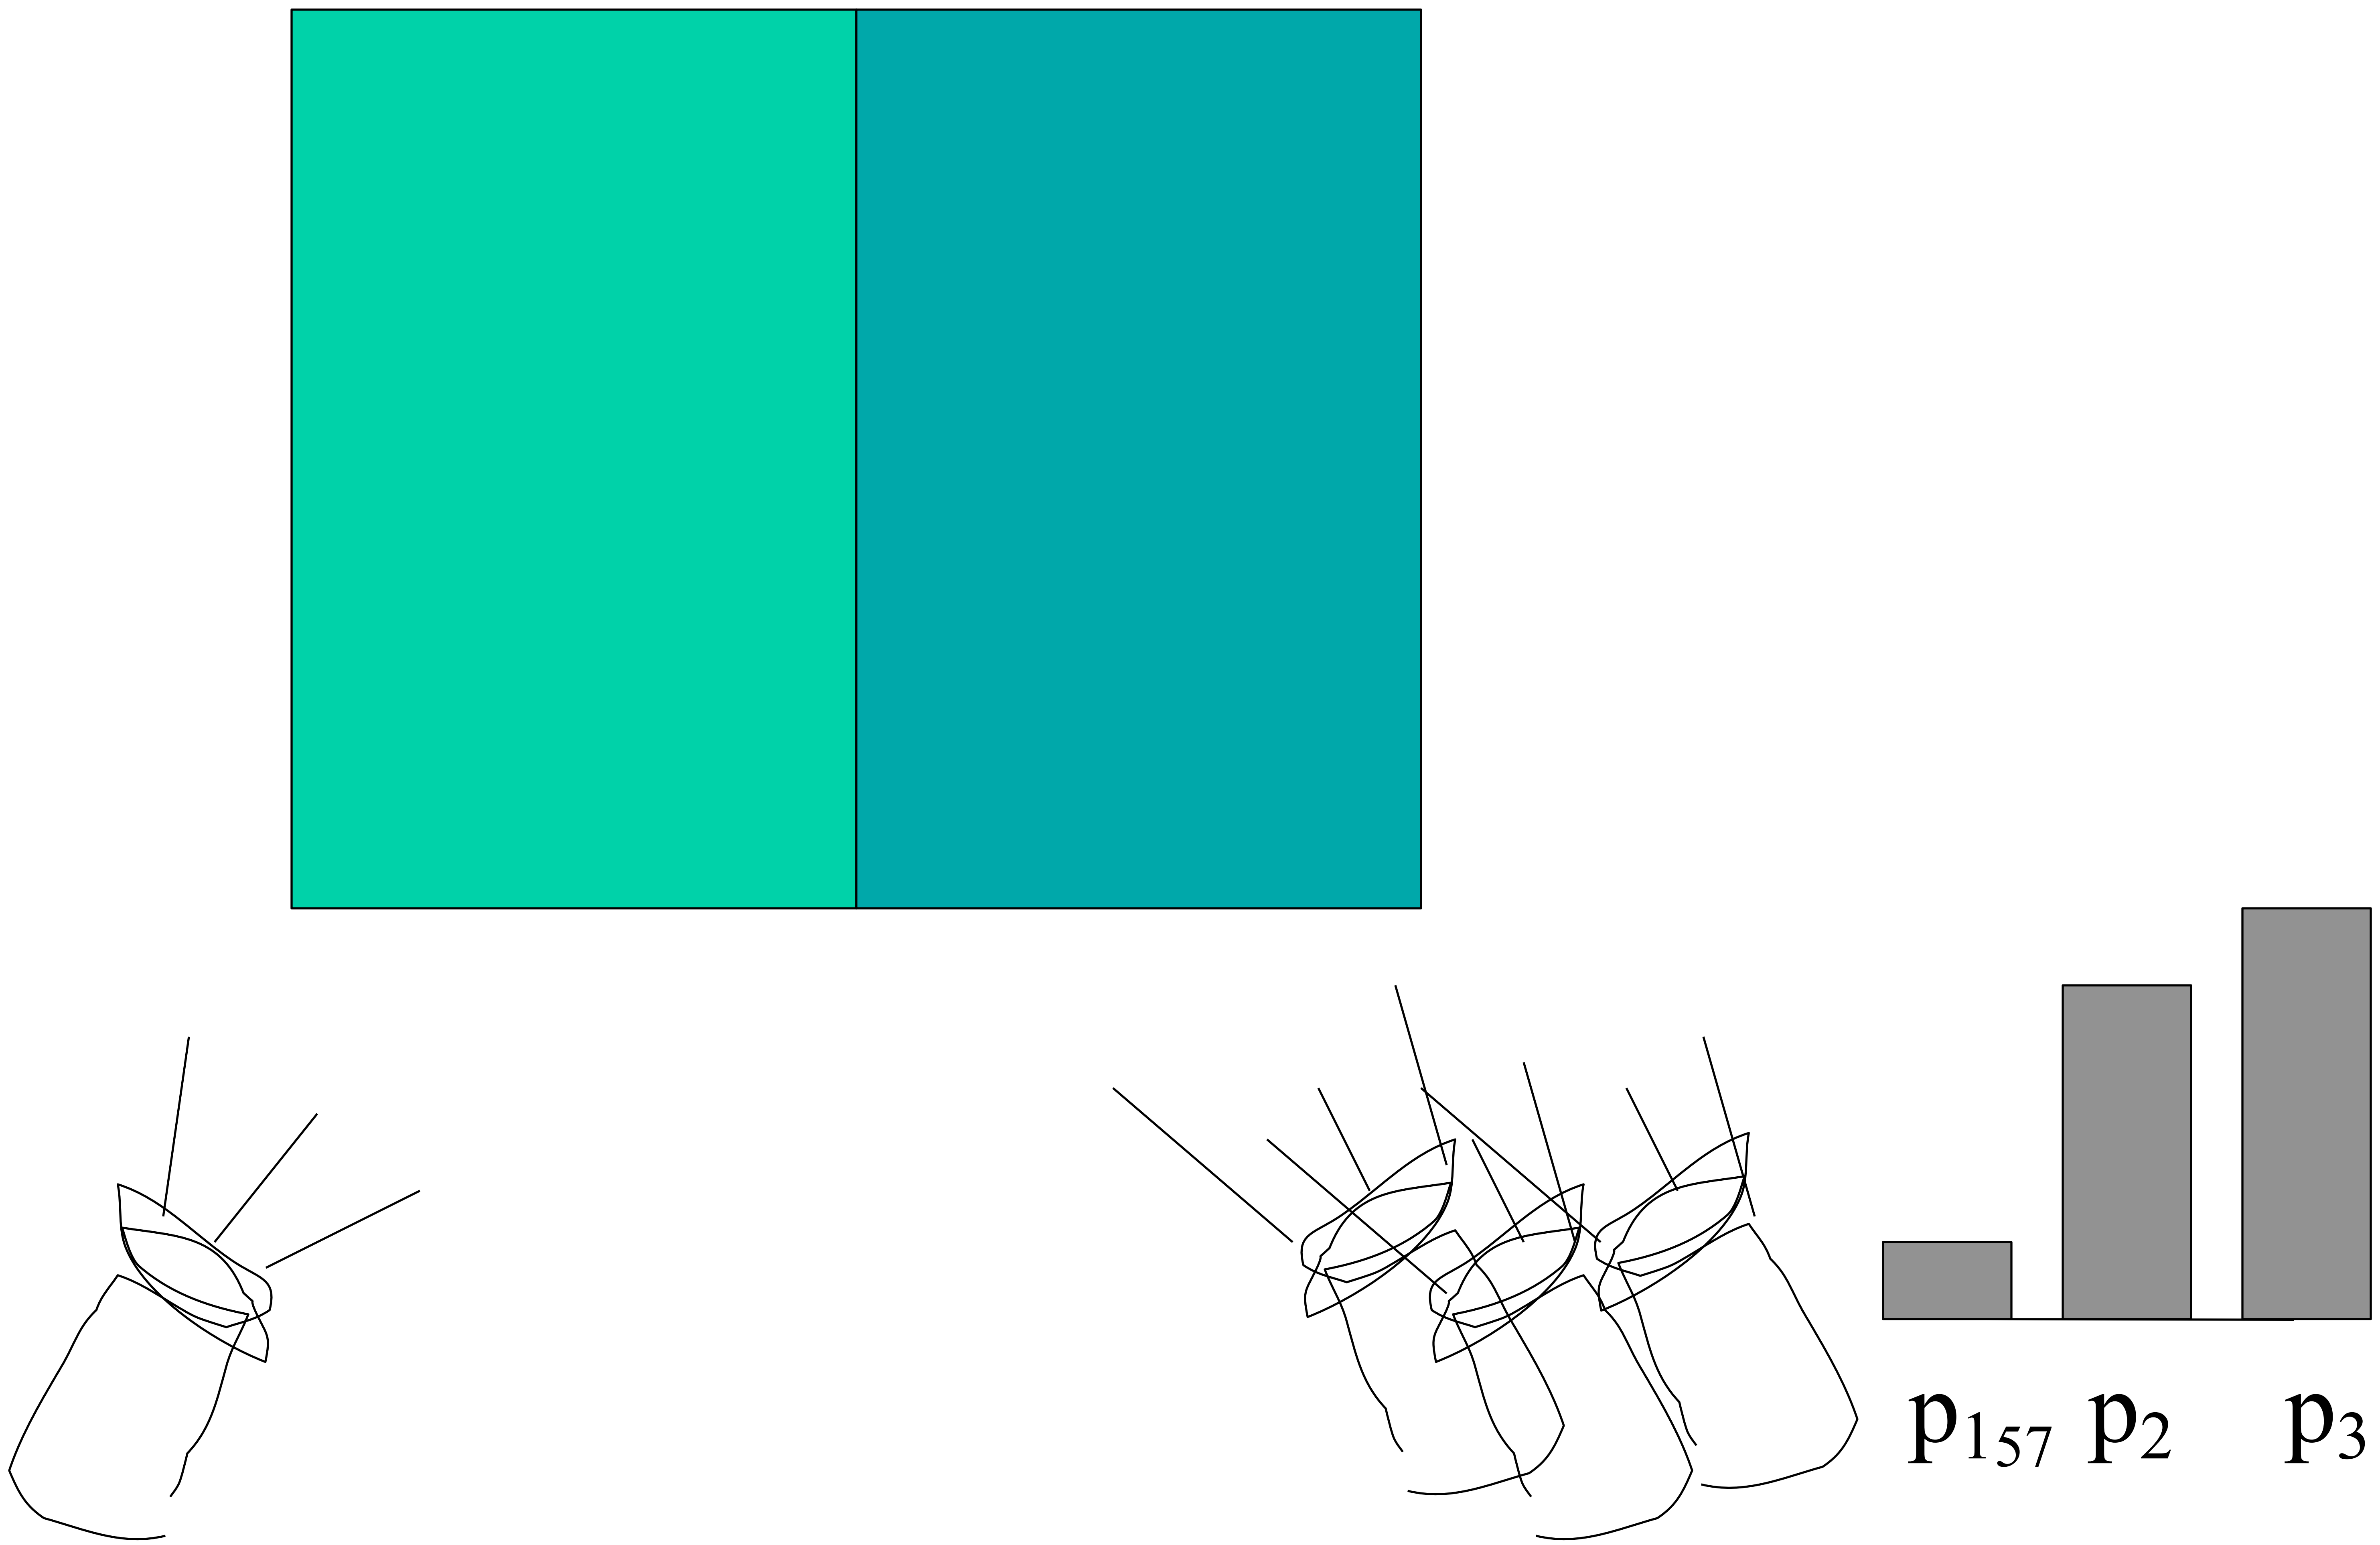

Source: W. Freeman

Color matching experiment

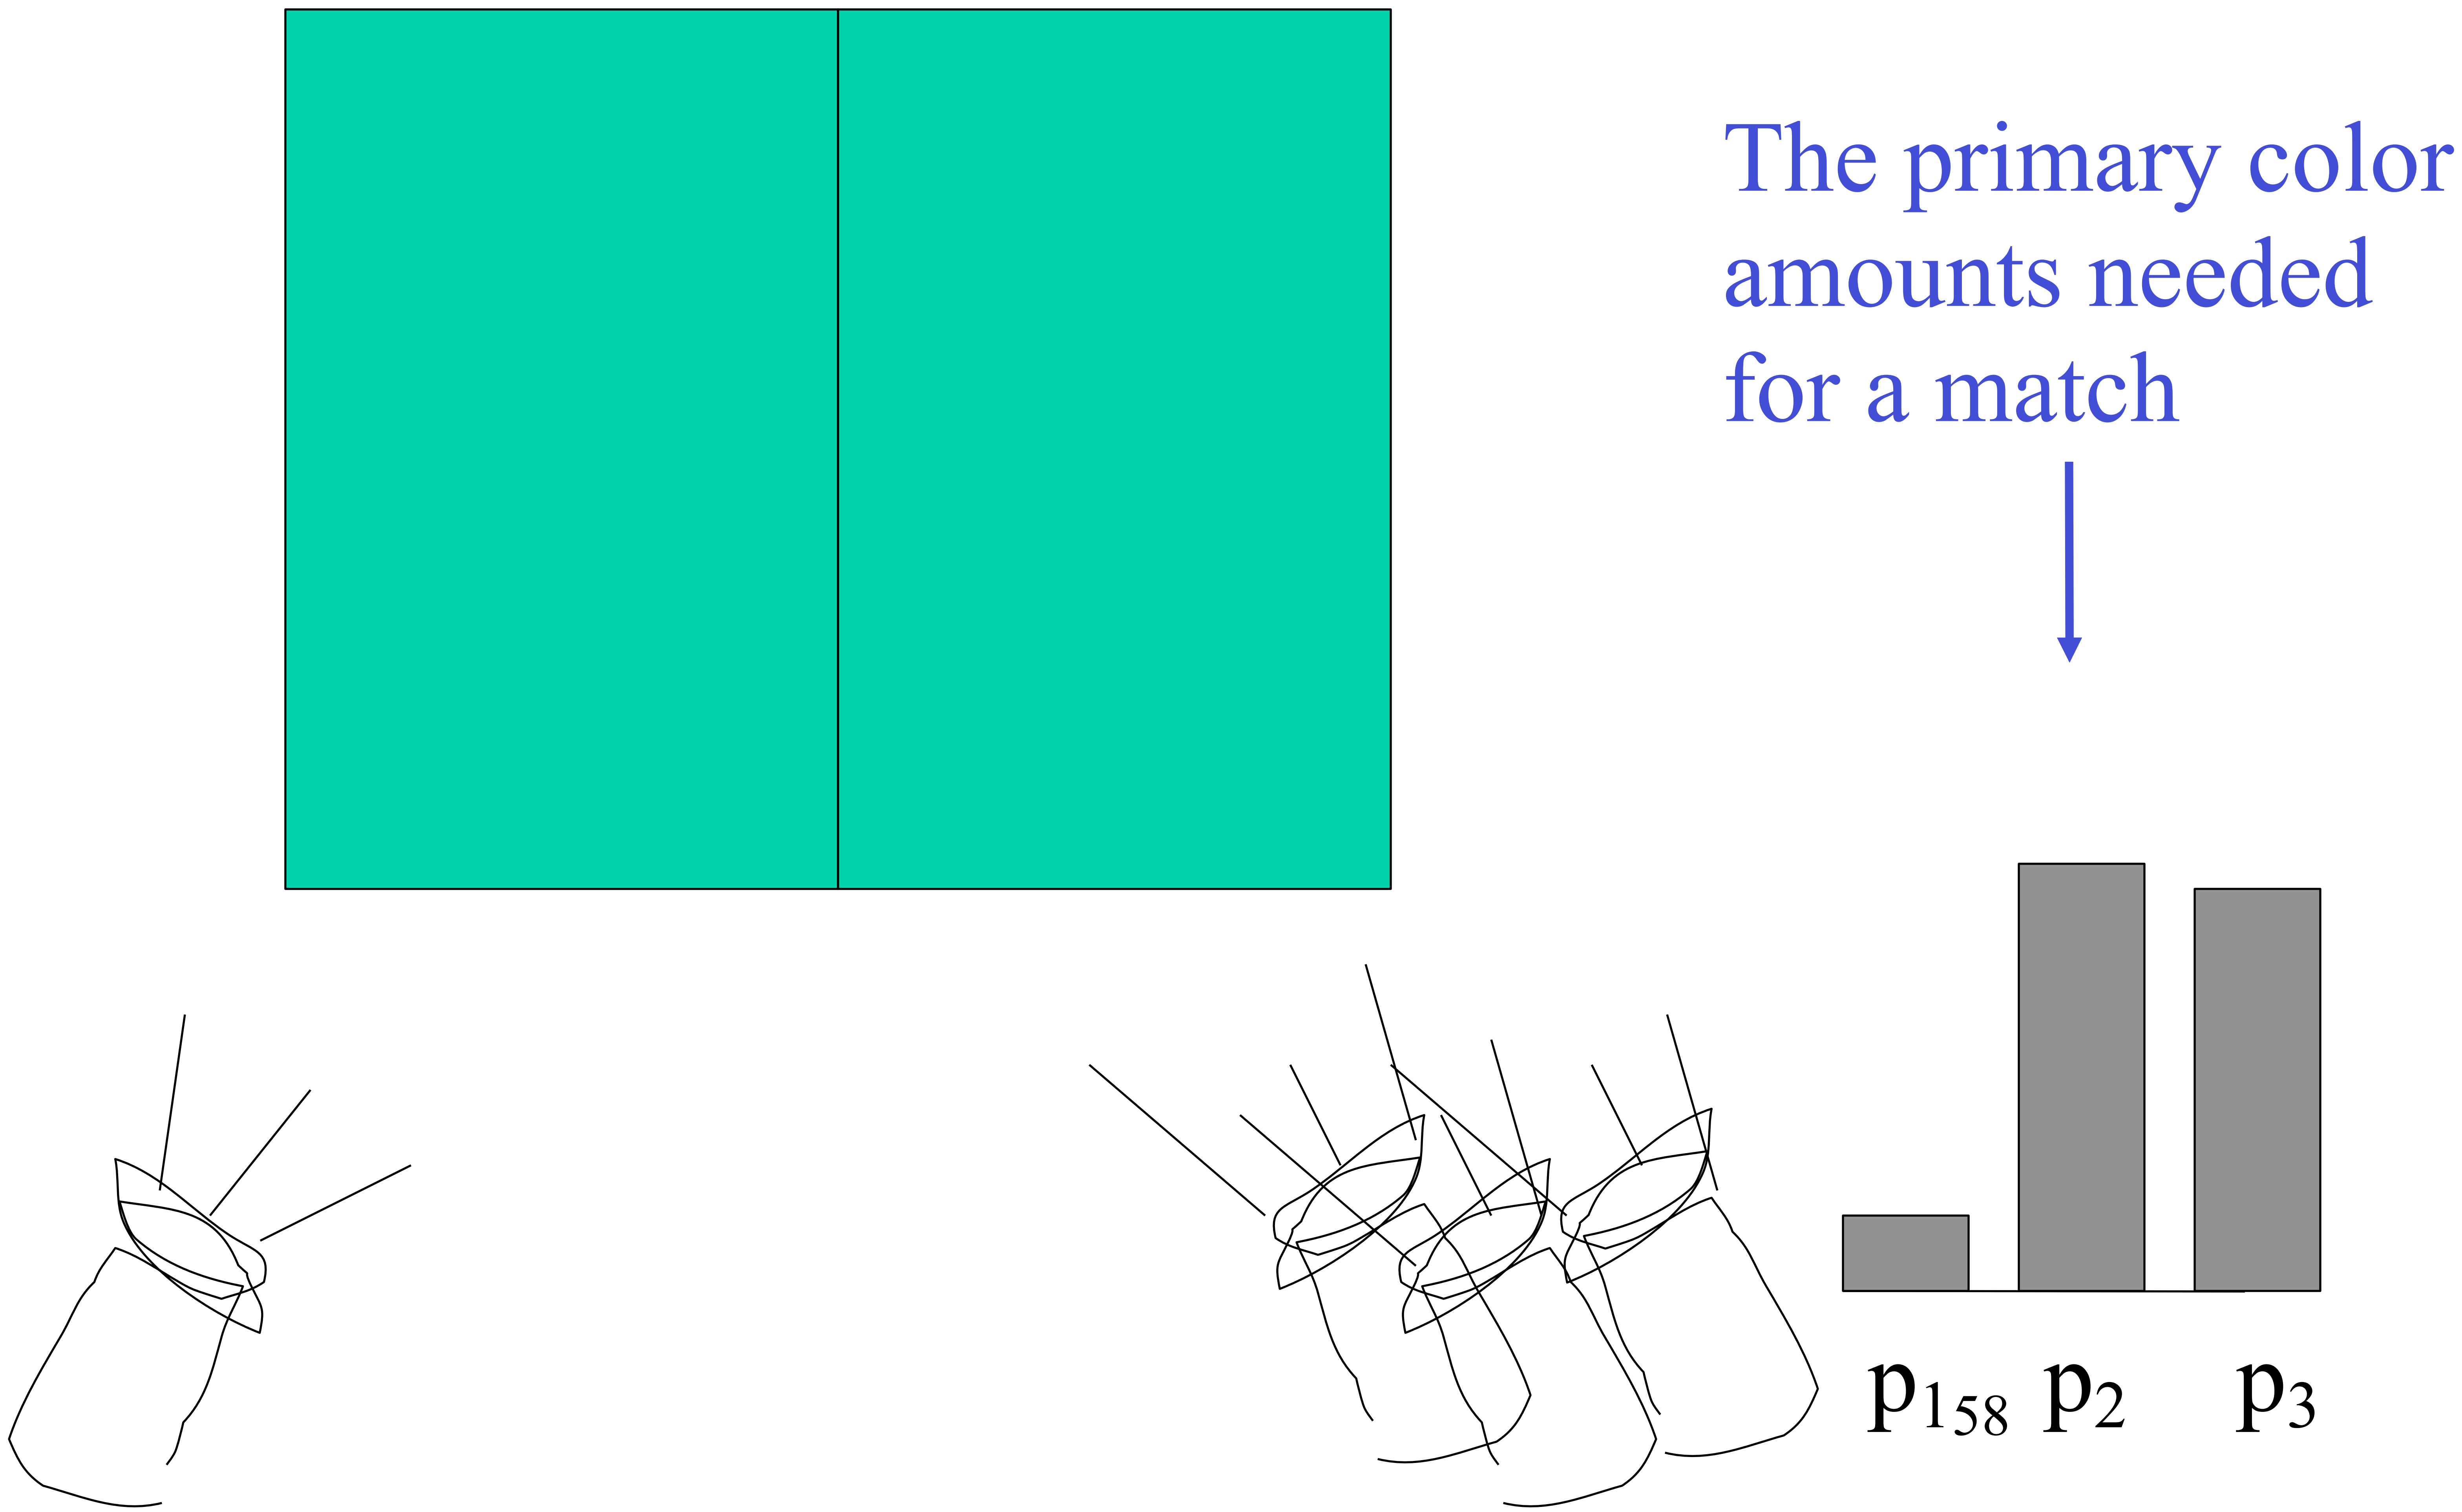

Source: W. Freeman

“Color matching functions” let us find other basis vectors for the eye response subspace of light power spectra

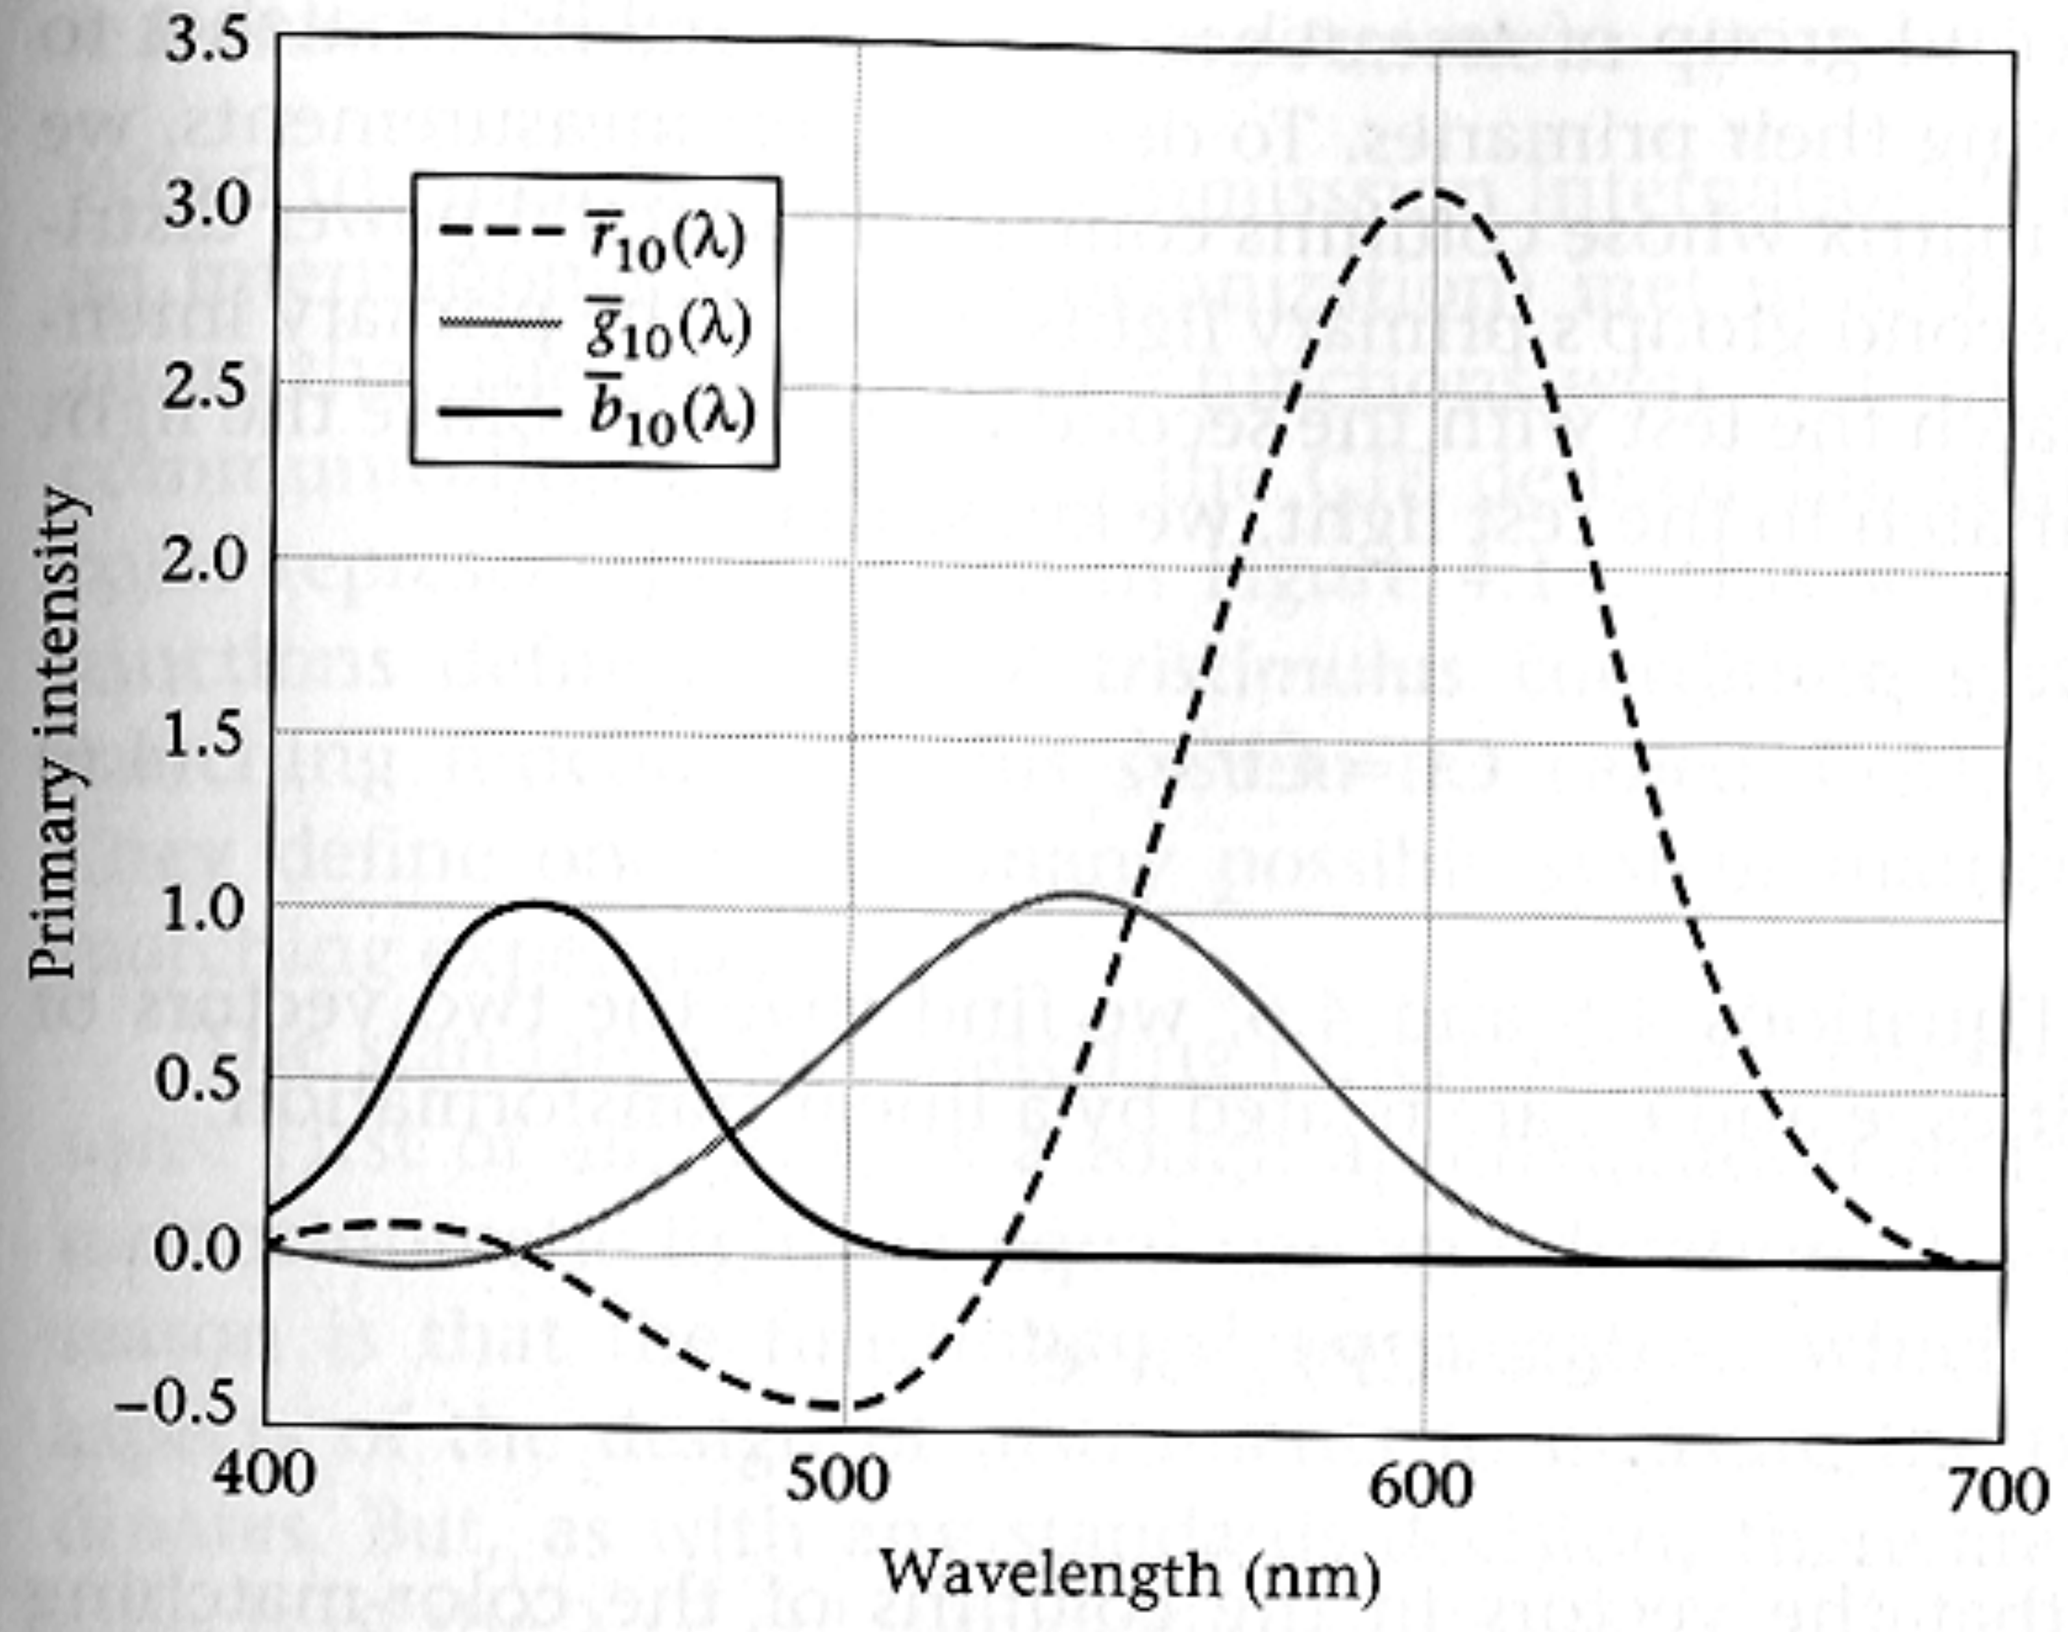

p_1 = 645.2 nm

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

Source: W. Freeman

Other color matching functions

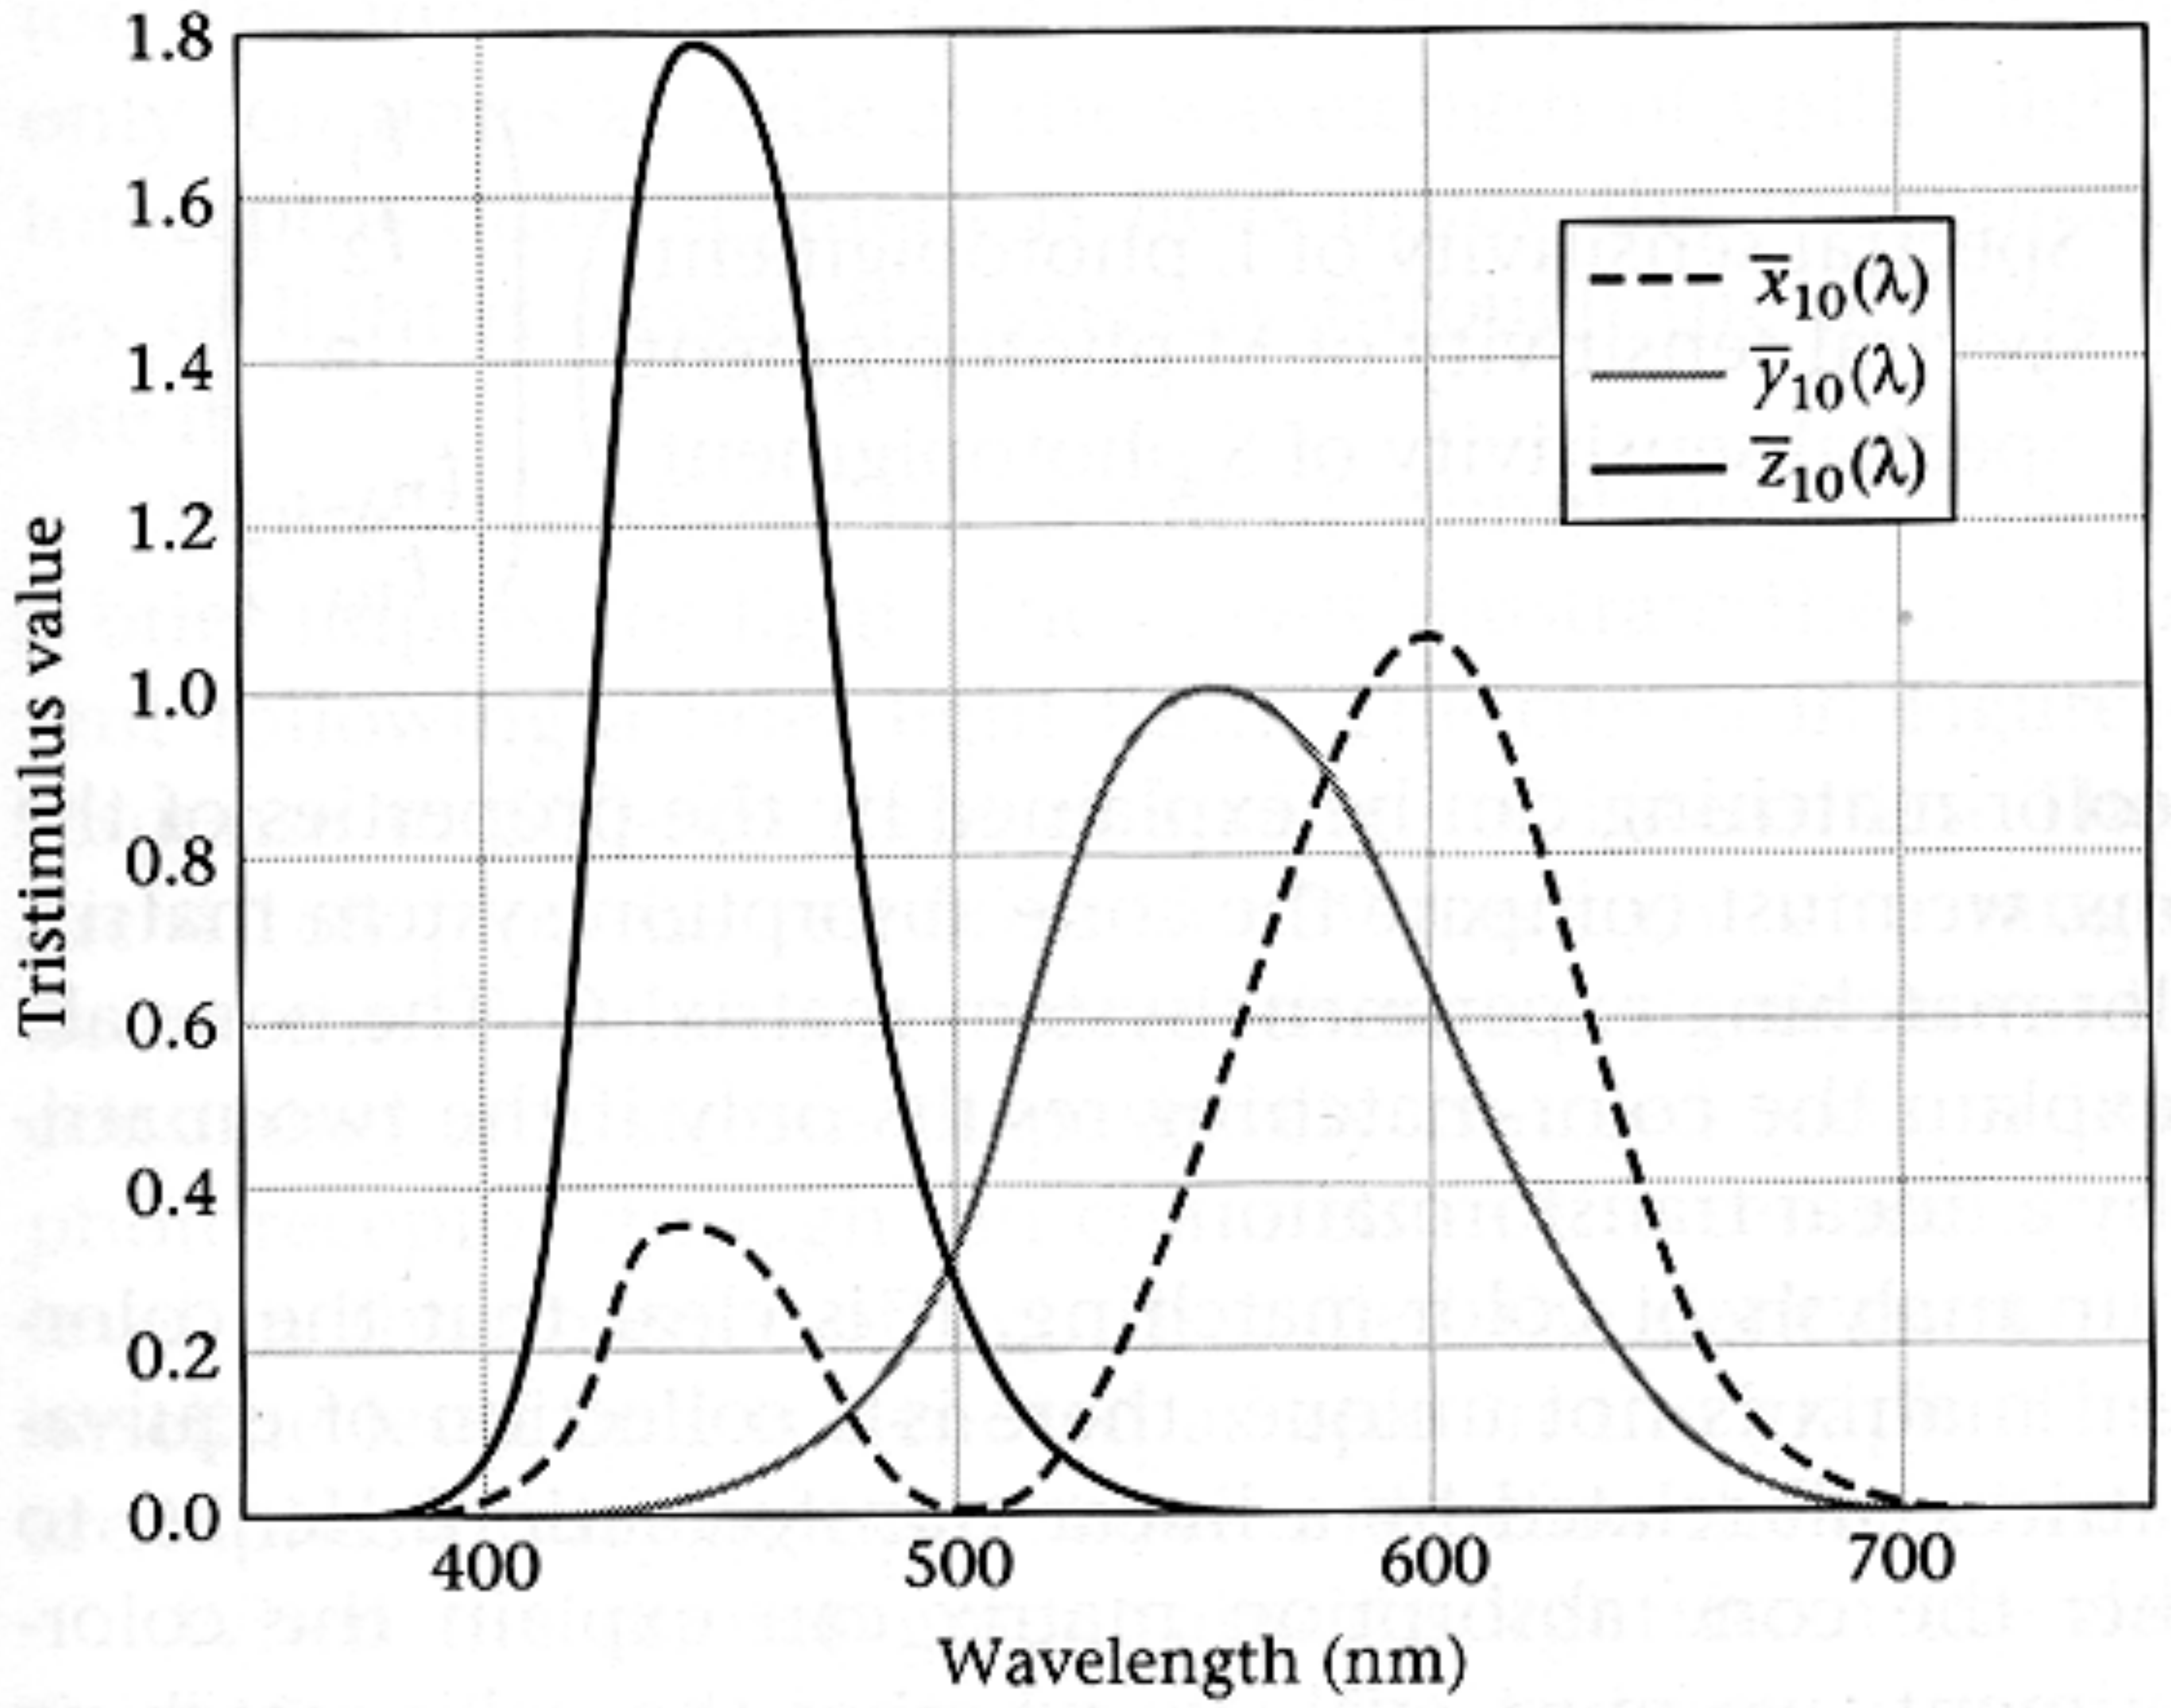

4.14 THE XYZ STANDARD COLOR-MATCHING FUNCTIONS. In 1931 the CIE standardized a set of color-matching functions for image interchange. These color-matching functions are called $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, and $\bar{z}(\lambda)$. Industrial applications commonly describe the color properties 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.

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

Source: W. Freeman

Metameric lights

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

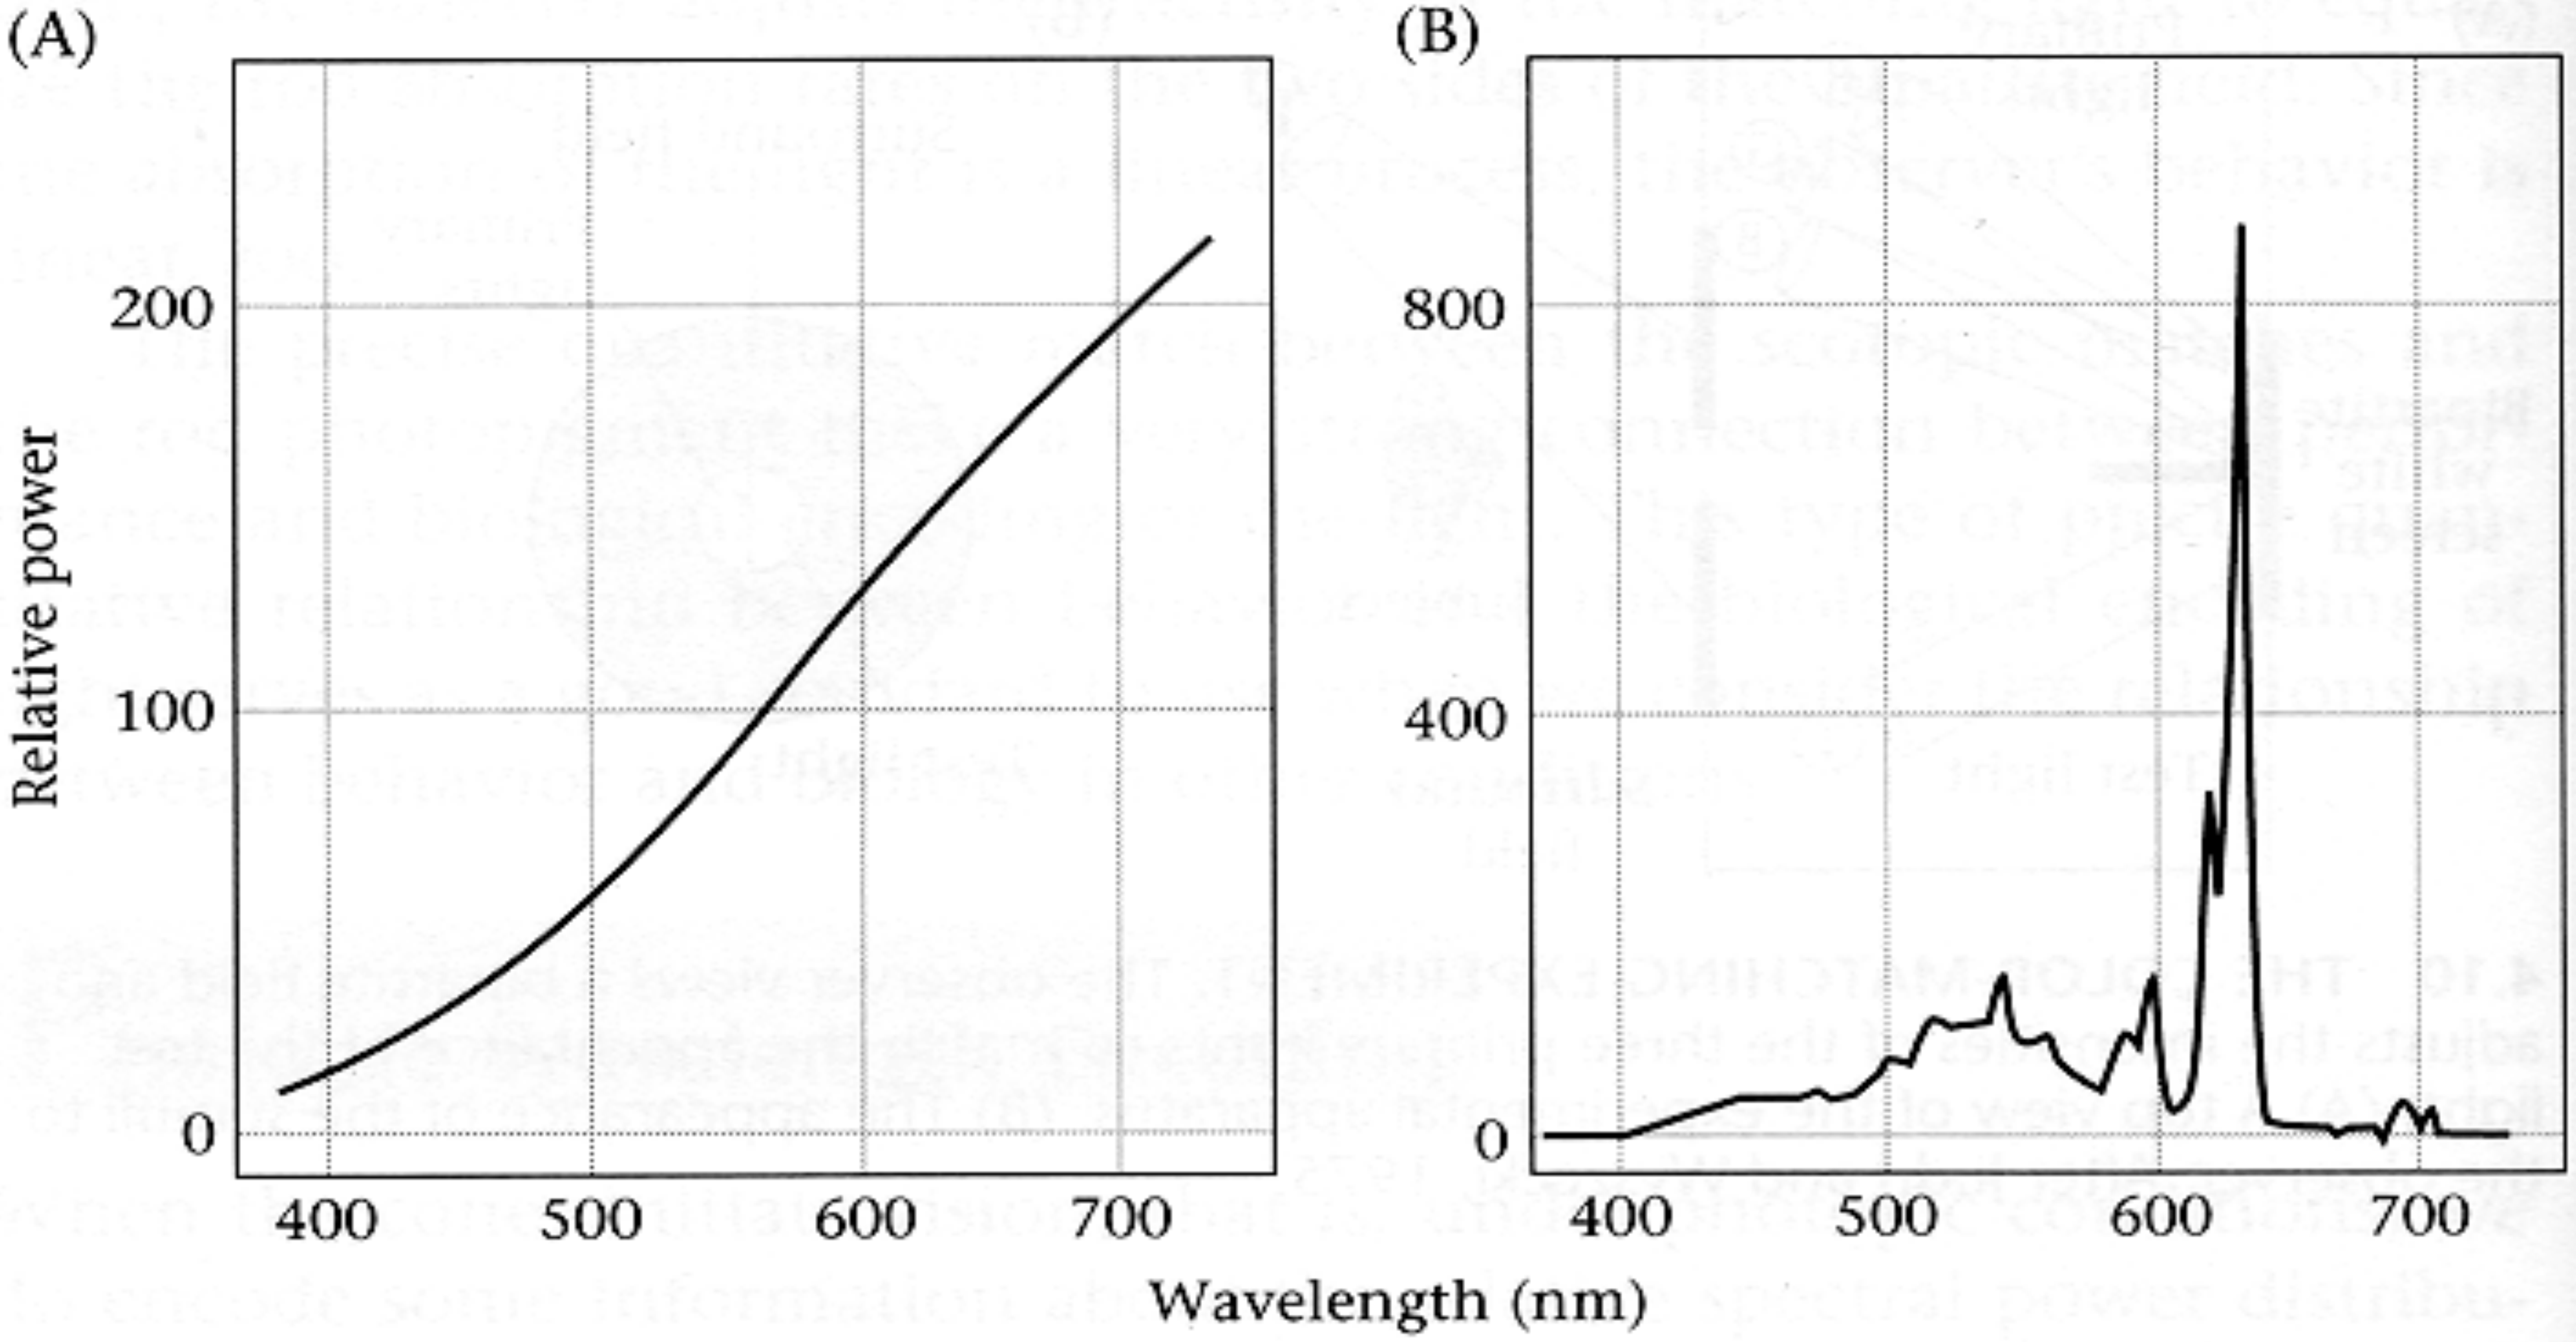

4.11 METAMERIC LIGHTS. Two lights with these spectral power distributions appear identical to most observers and are called metamers. (A) An approximation to the spectral power distribution of a tungsten bulb. (B) The spectral power distribution of light emitted from a conventional television monitor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

Source: W. Freeman

Color blindness

- Classical case: 1 type of cone is missing (e.g. red)
- Makes it impossible to distinguish some spectra
- There are also tetrachromats, who have 4 cones!

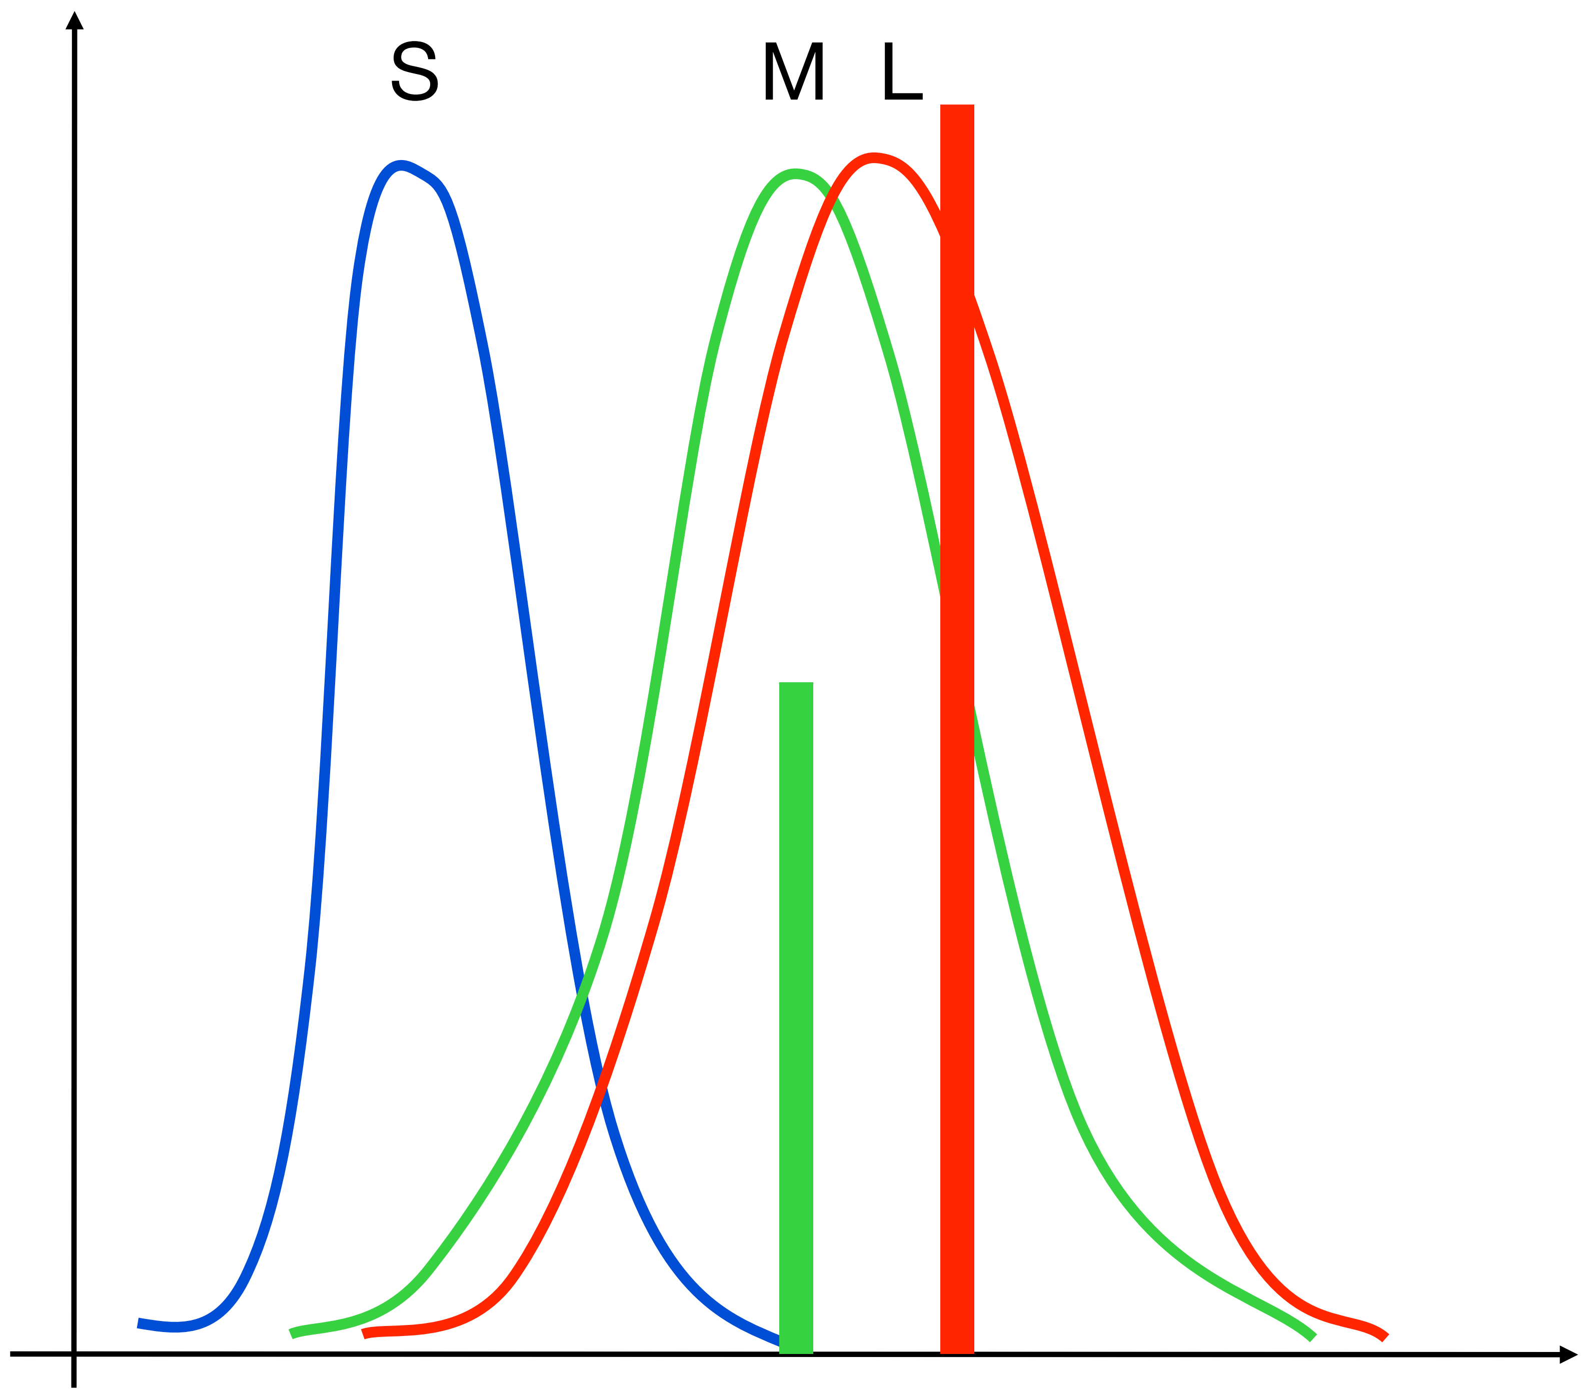

Differentiated

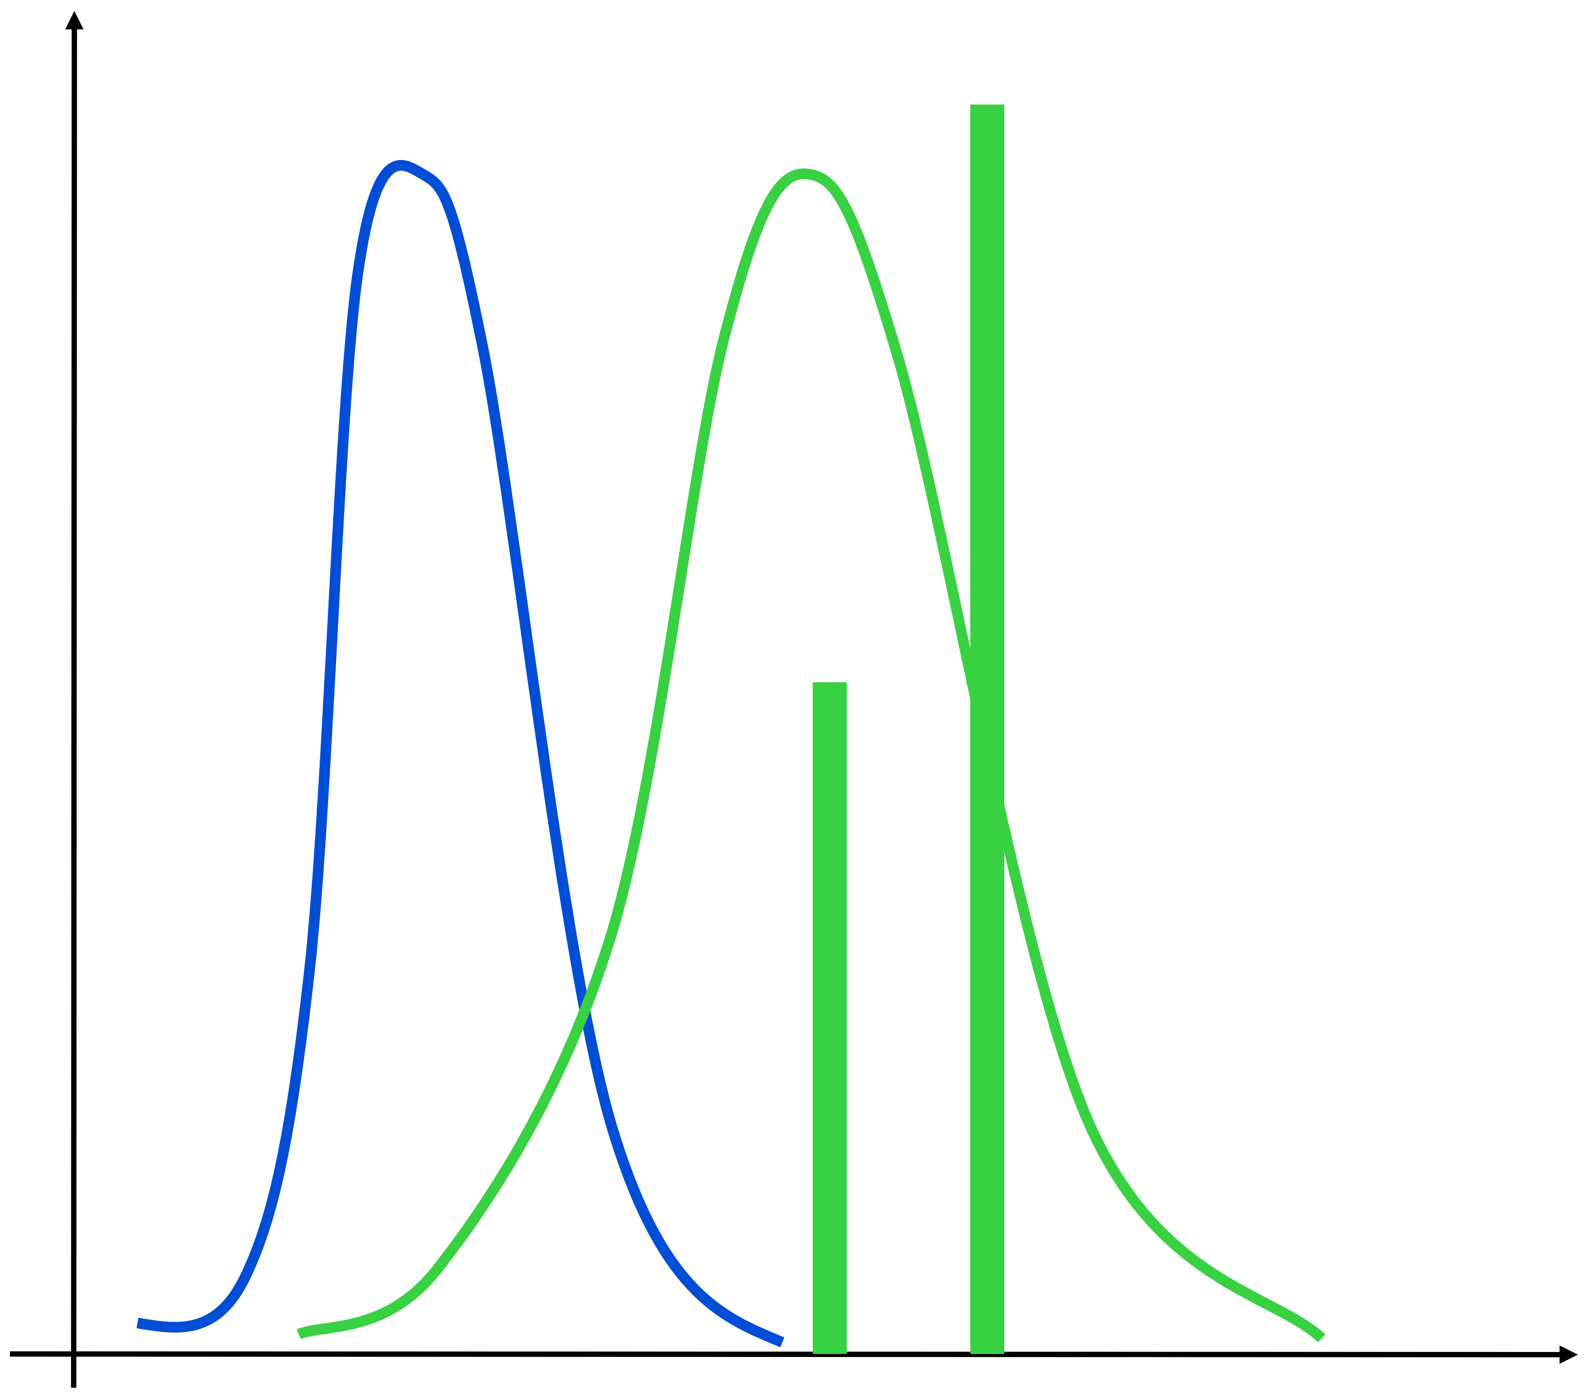

Same responses

Source: F. Durand

Light sources

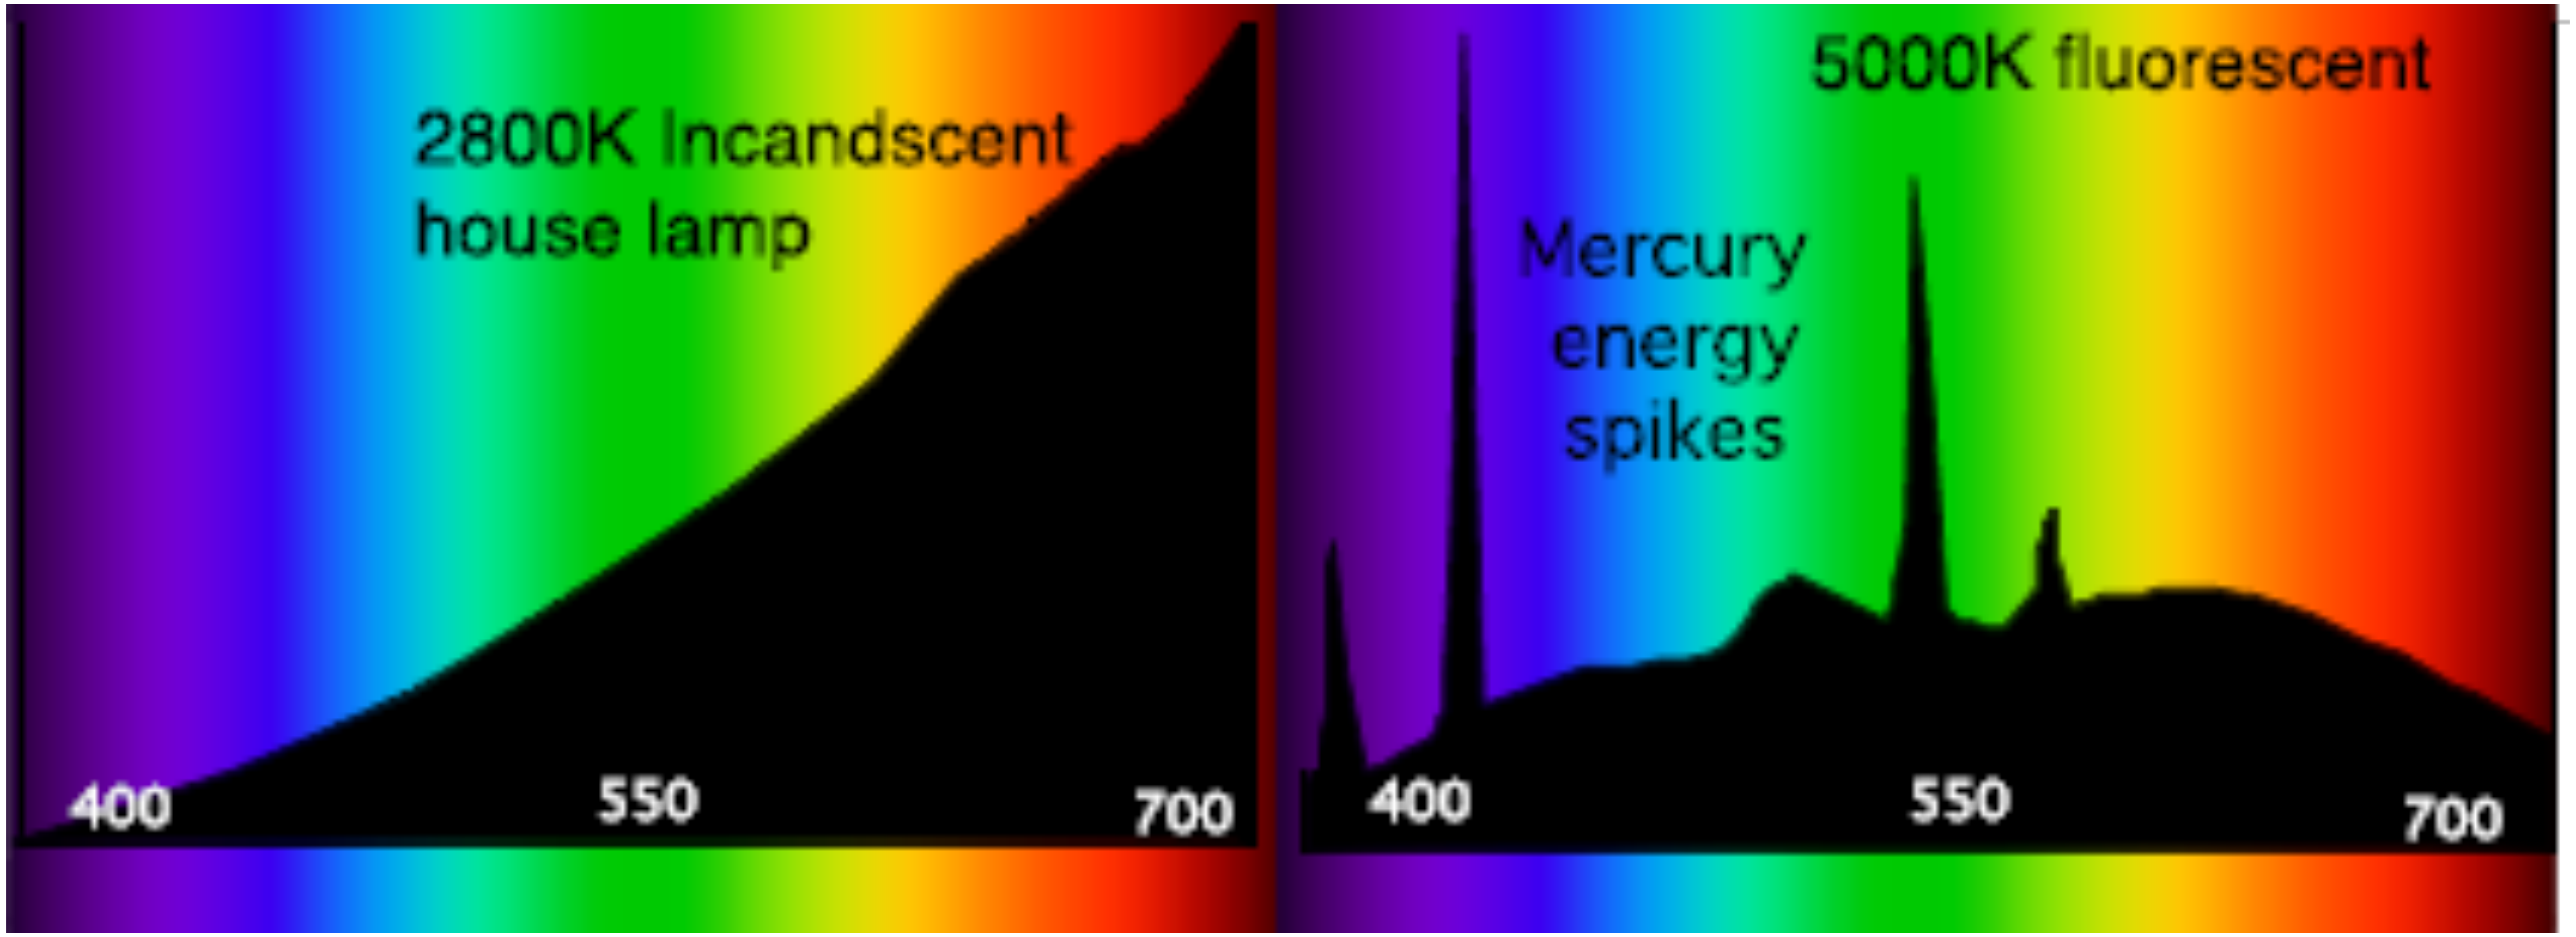

https://en.wikipedia.org/wiki/Color_temperature

Same scene under different illuminations

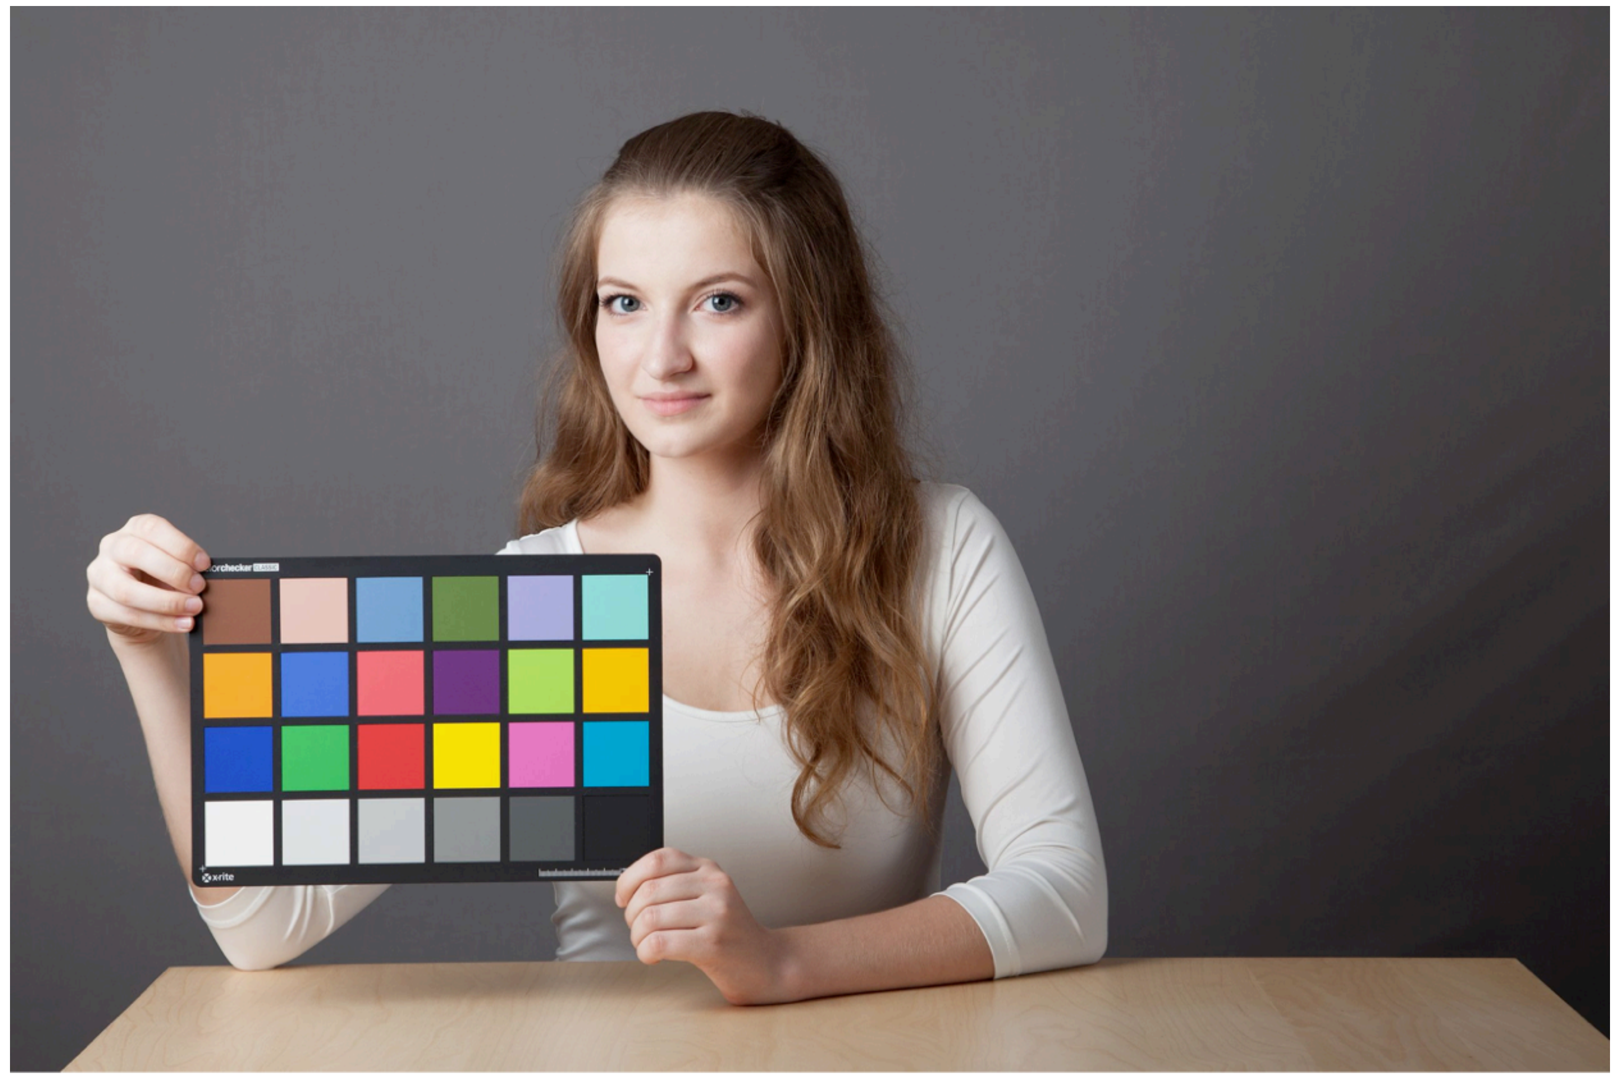

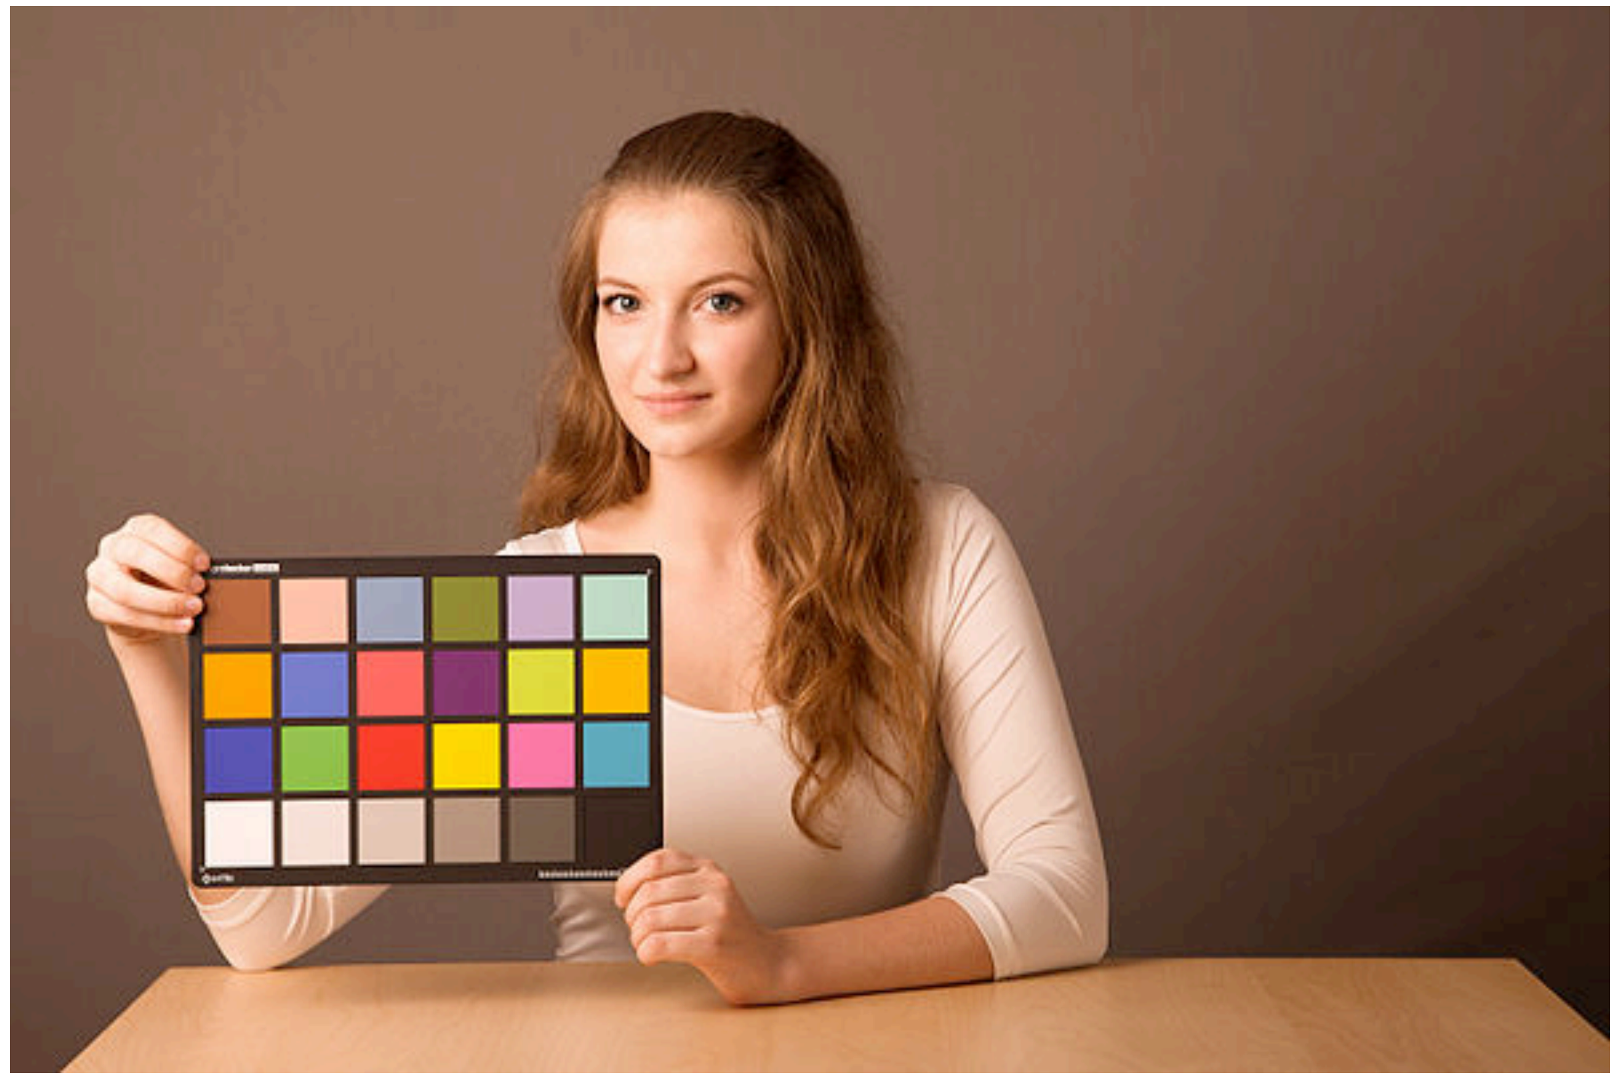

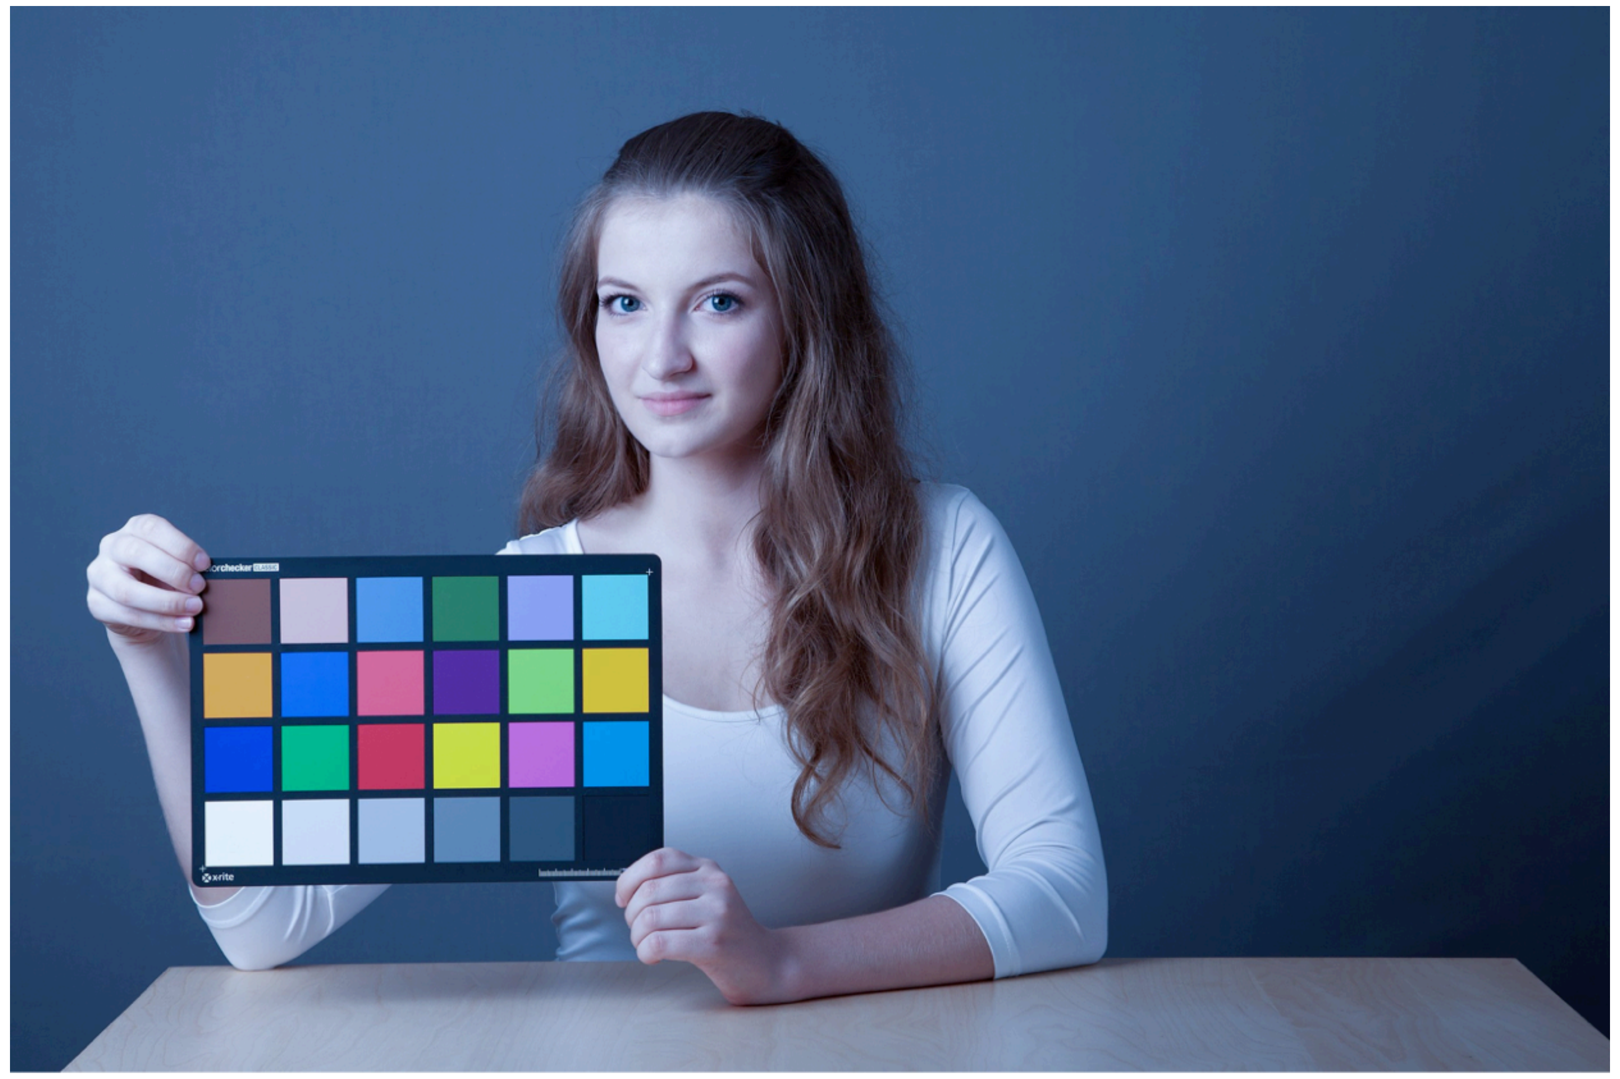

https://en.wikipedia.org/wiki/Color_balance

White balancing

Is this a green light and a white object, or a green object and a white light?

White balancing

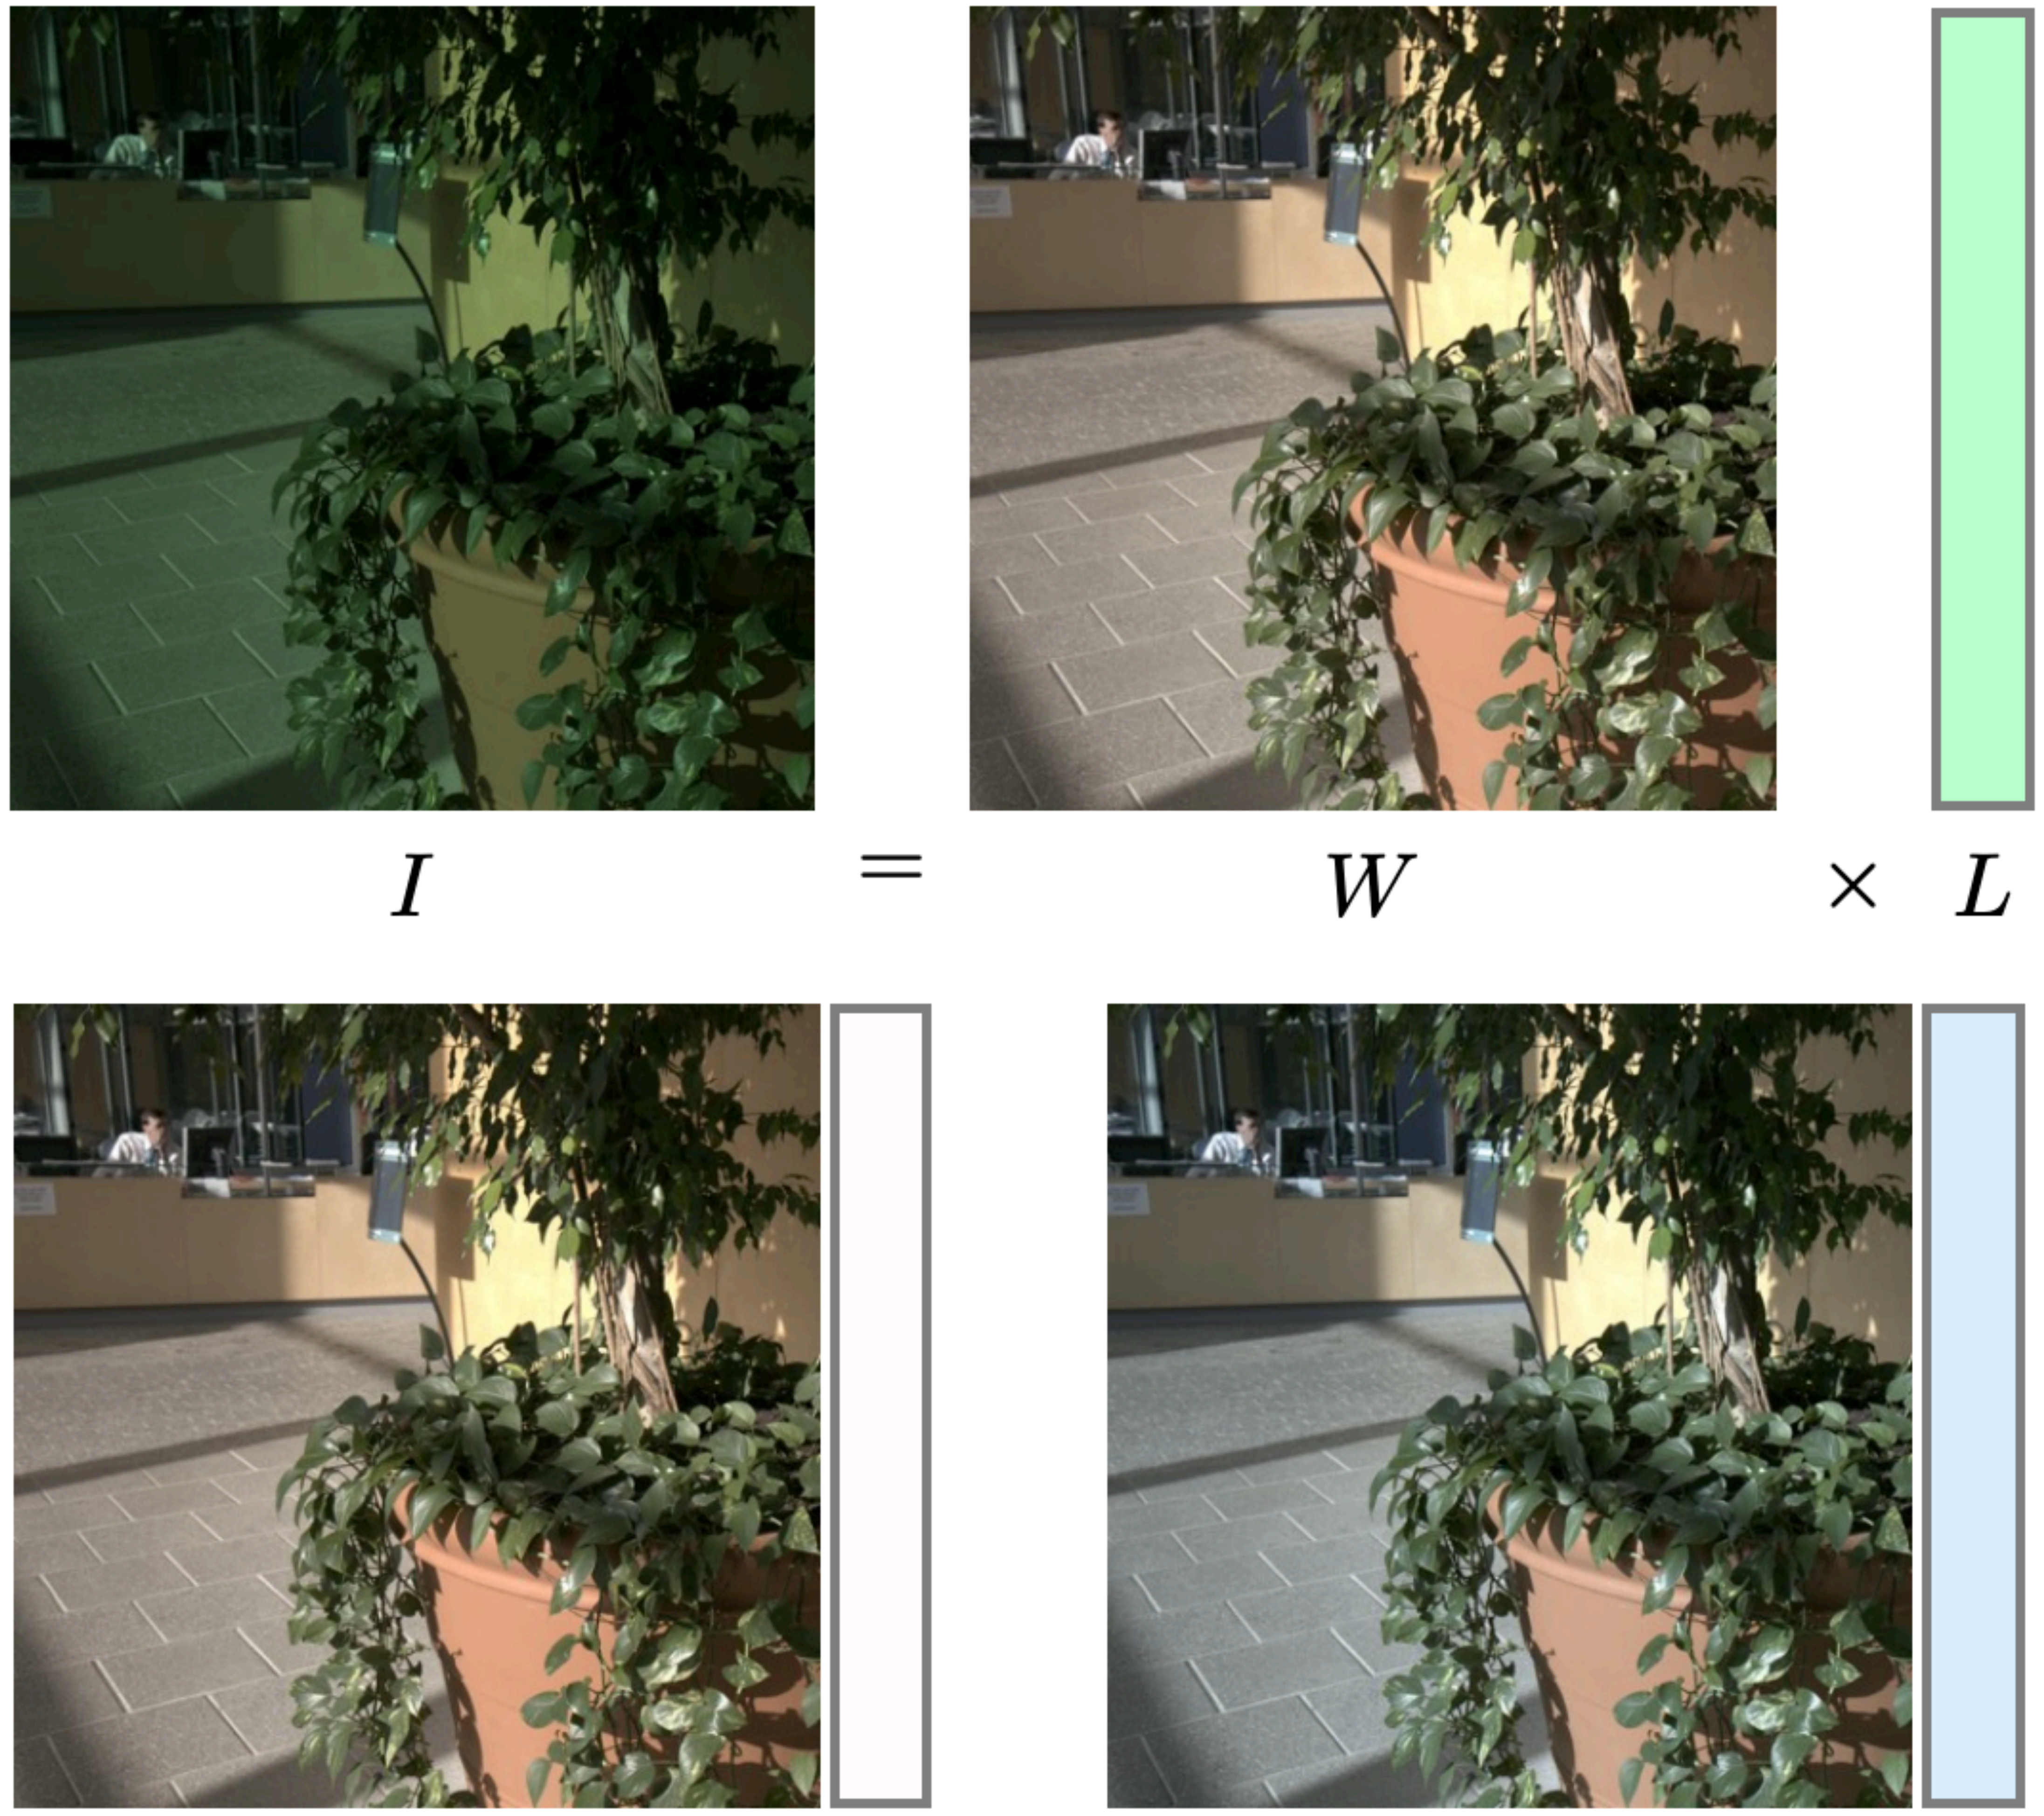

our $\hat{W}$, $\hat{L}$, err = $0.13°$ baseline $\hat{W}$, $\hat{L}$, err = $5.34°$

[Barron, “Convolutional Color Constancy”, 2015]

Color constancy

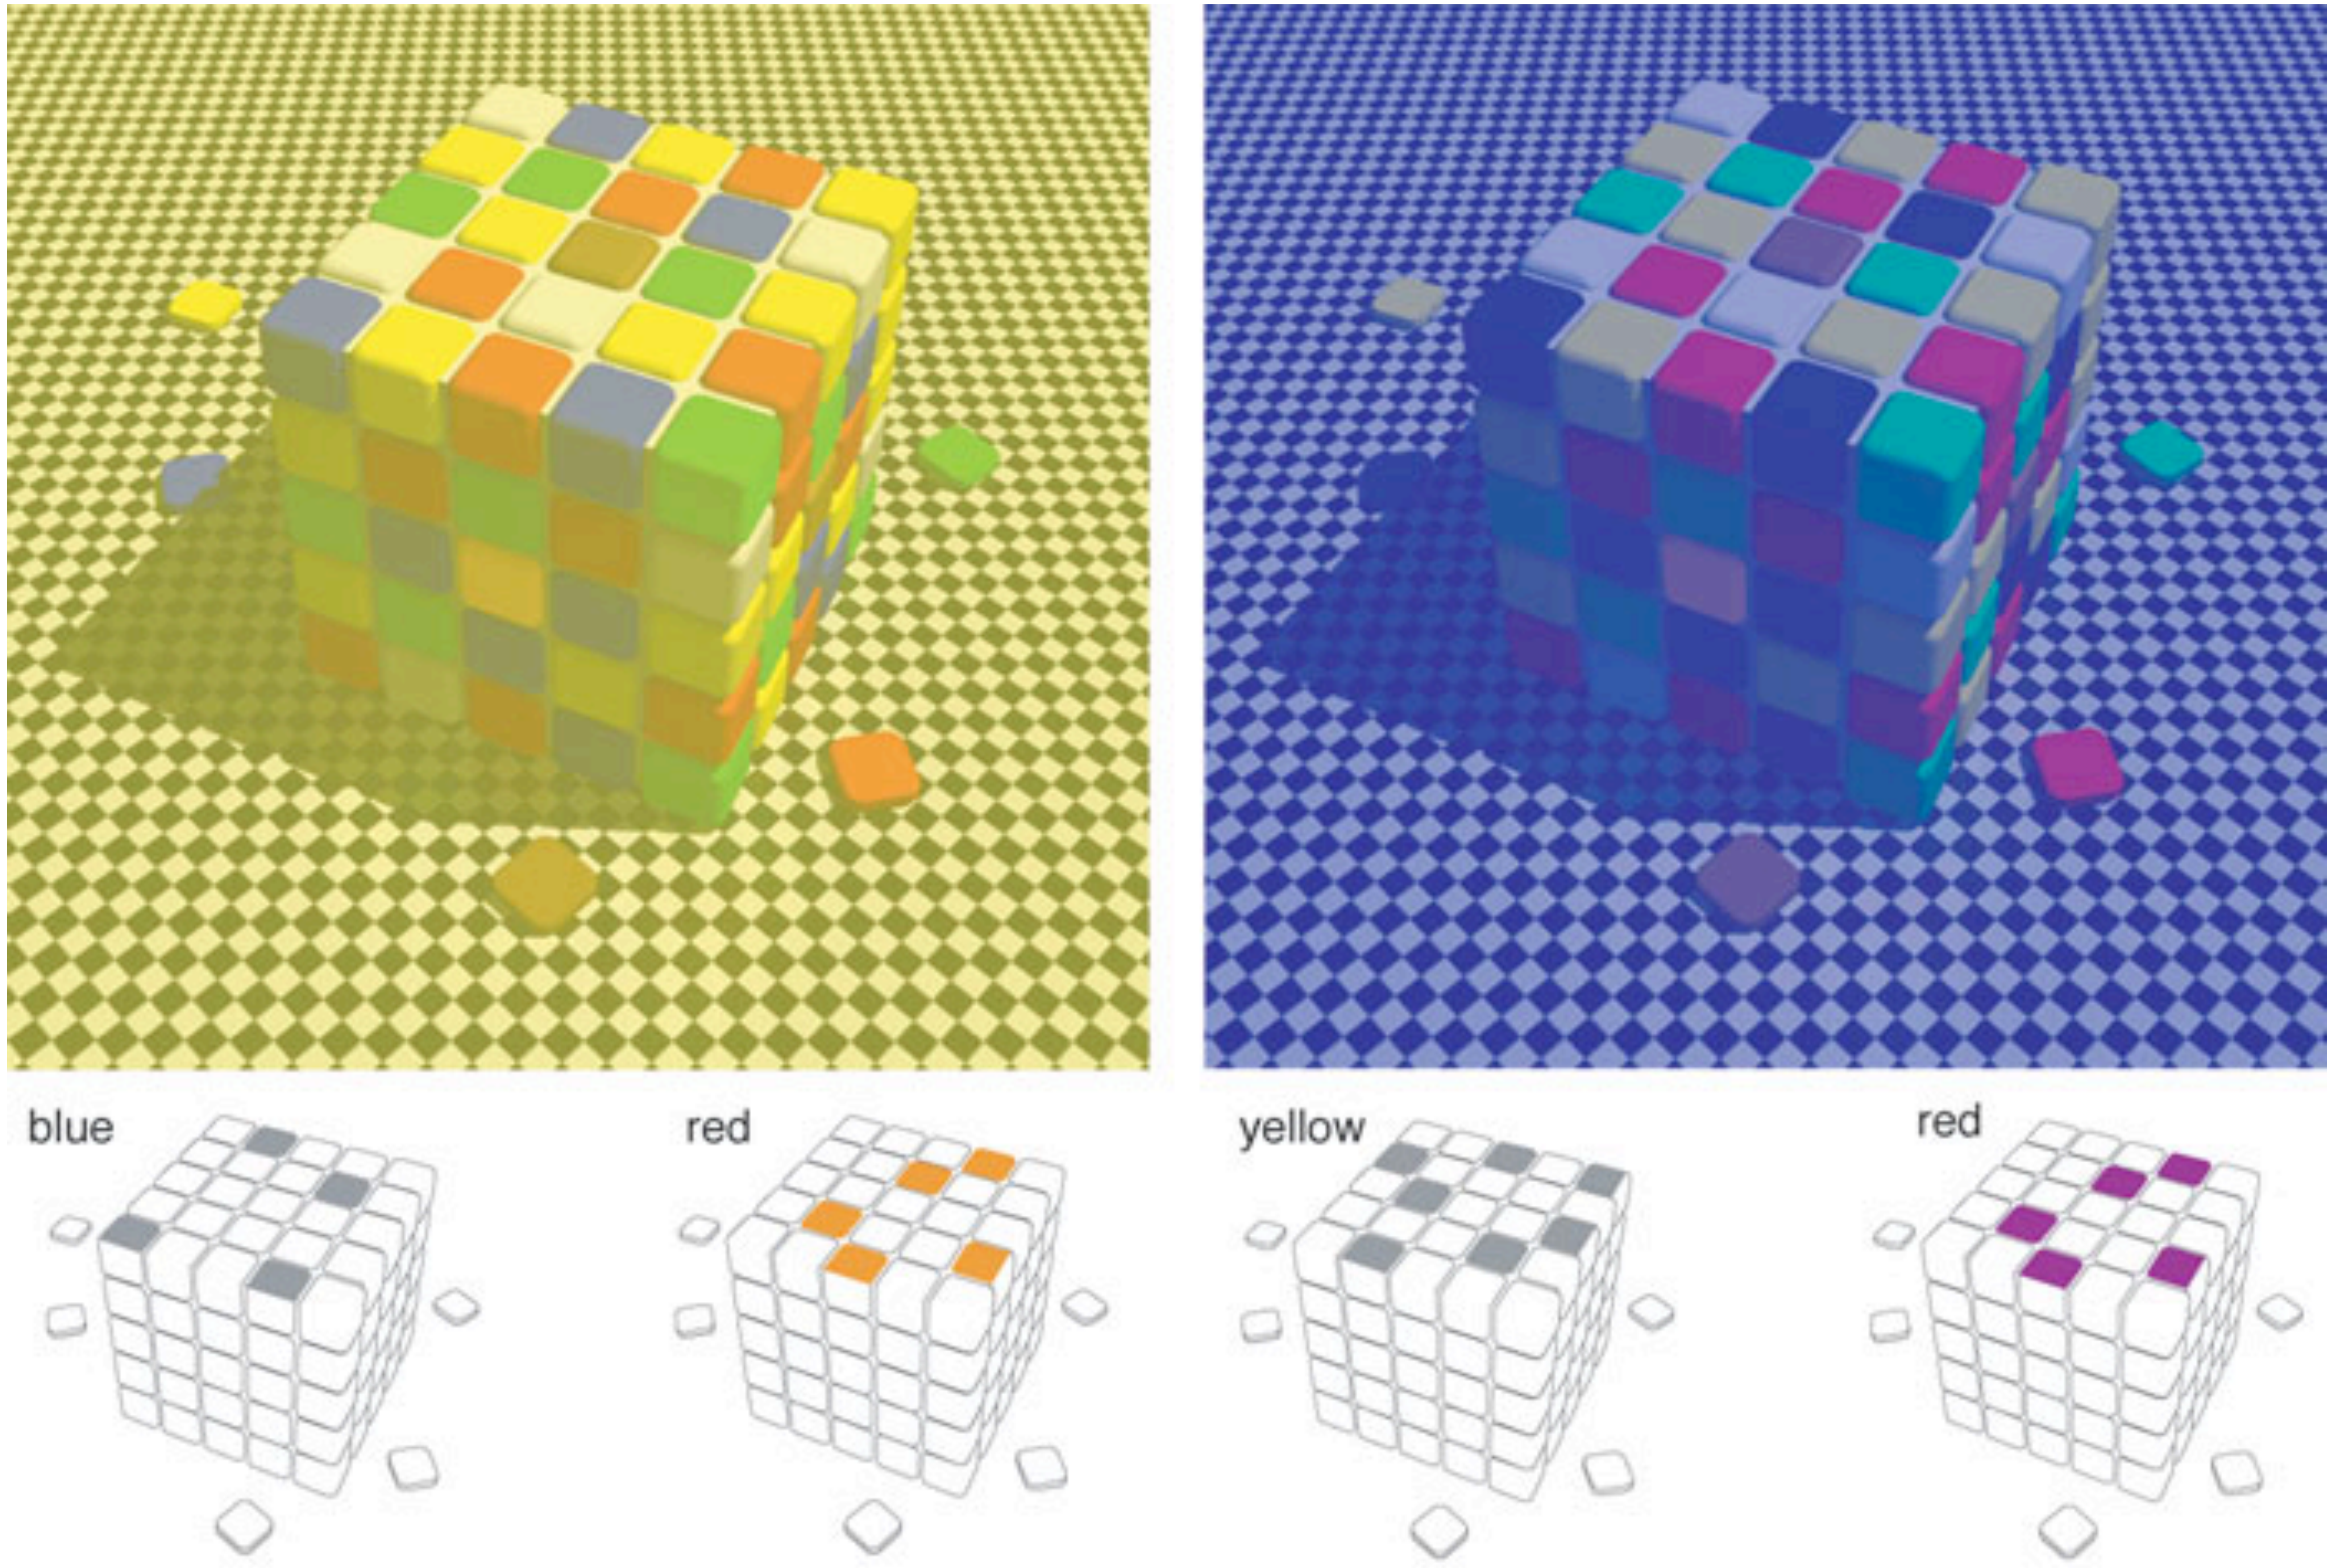

Dale Purves, R. Beau Lotto, Surajit Nundy, "Why We See What We Do."

The dress

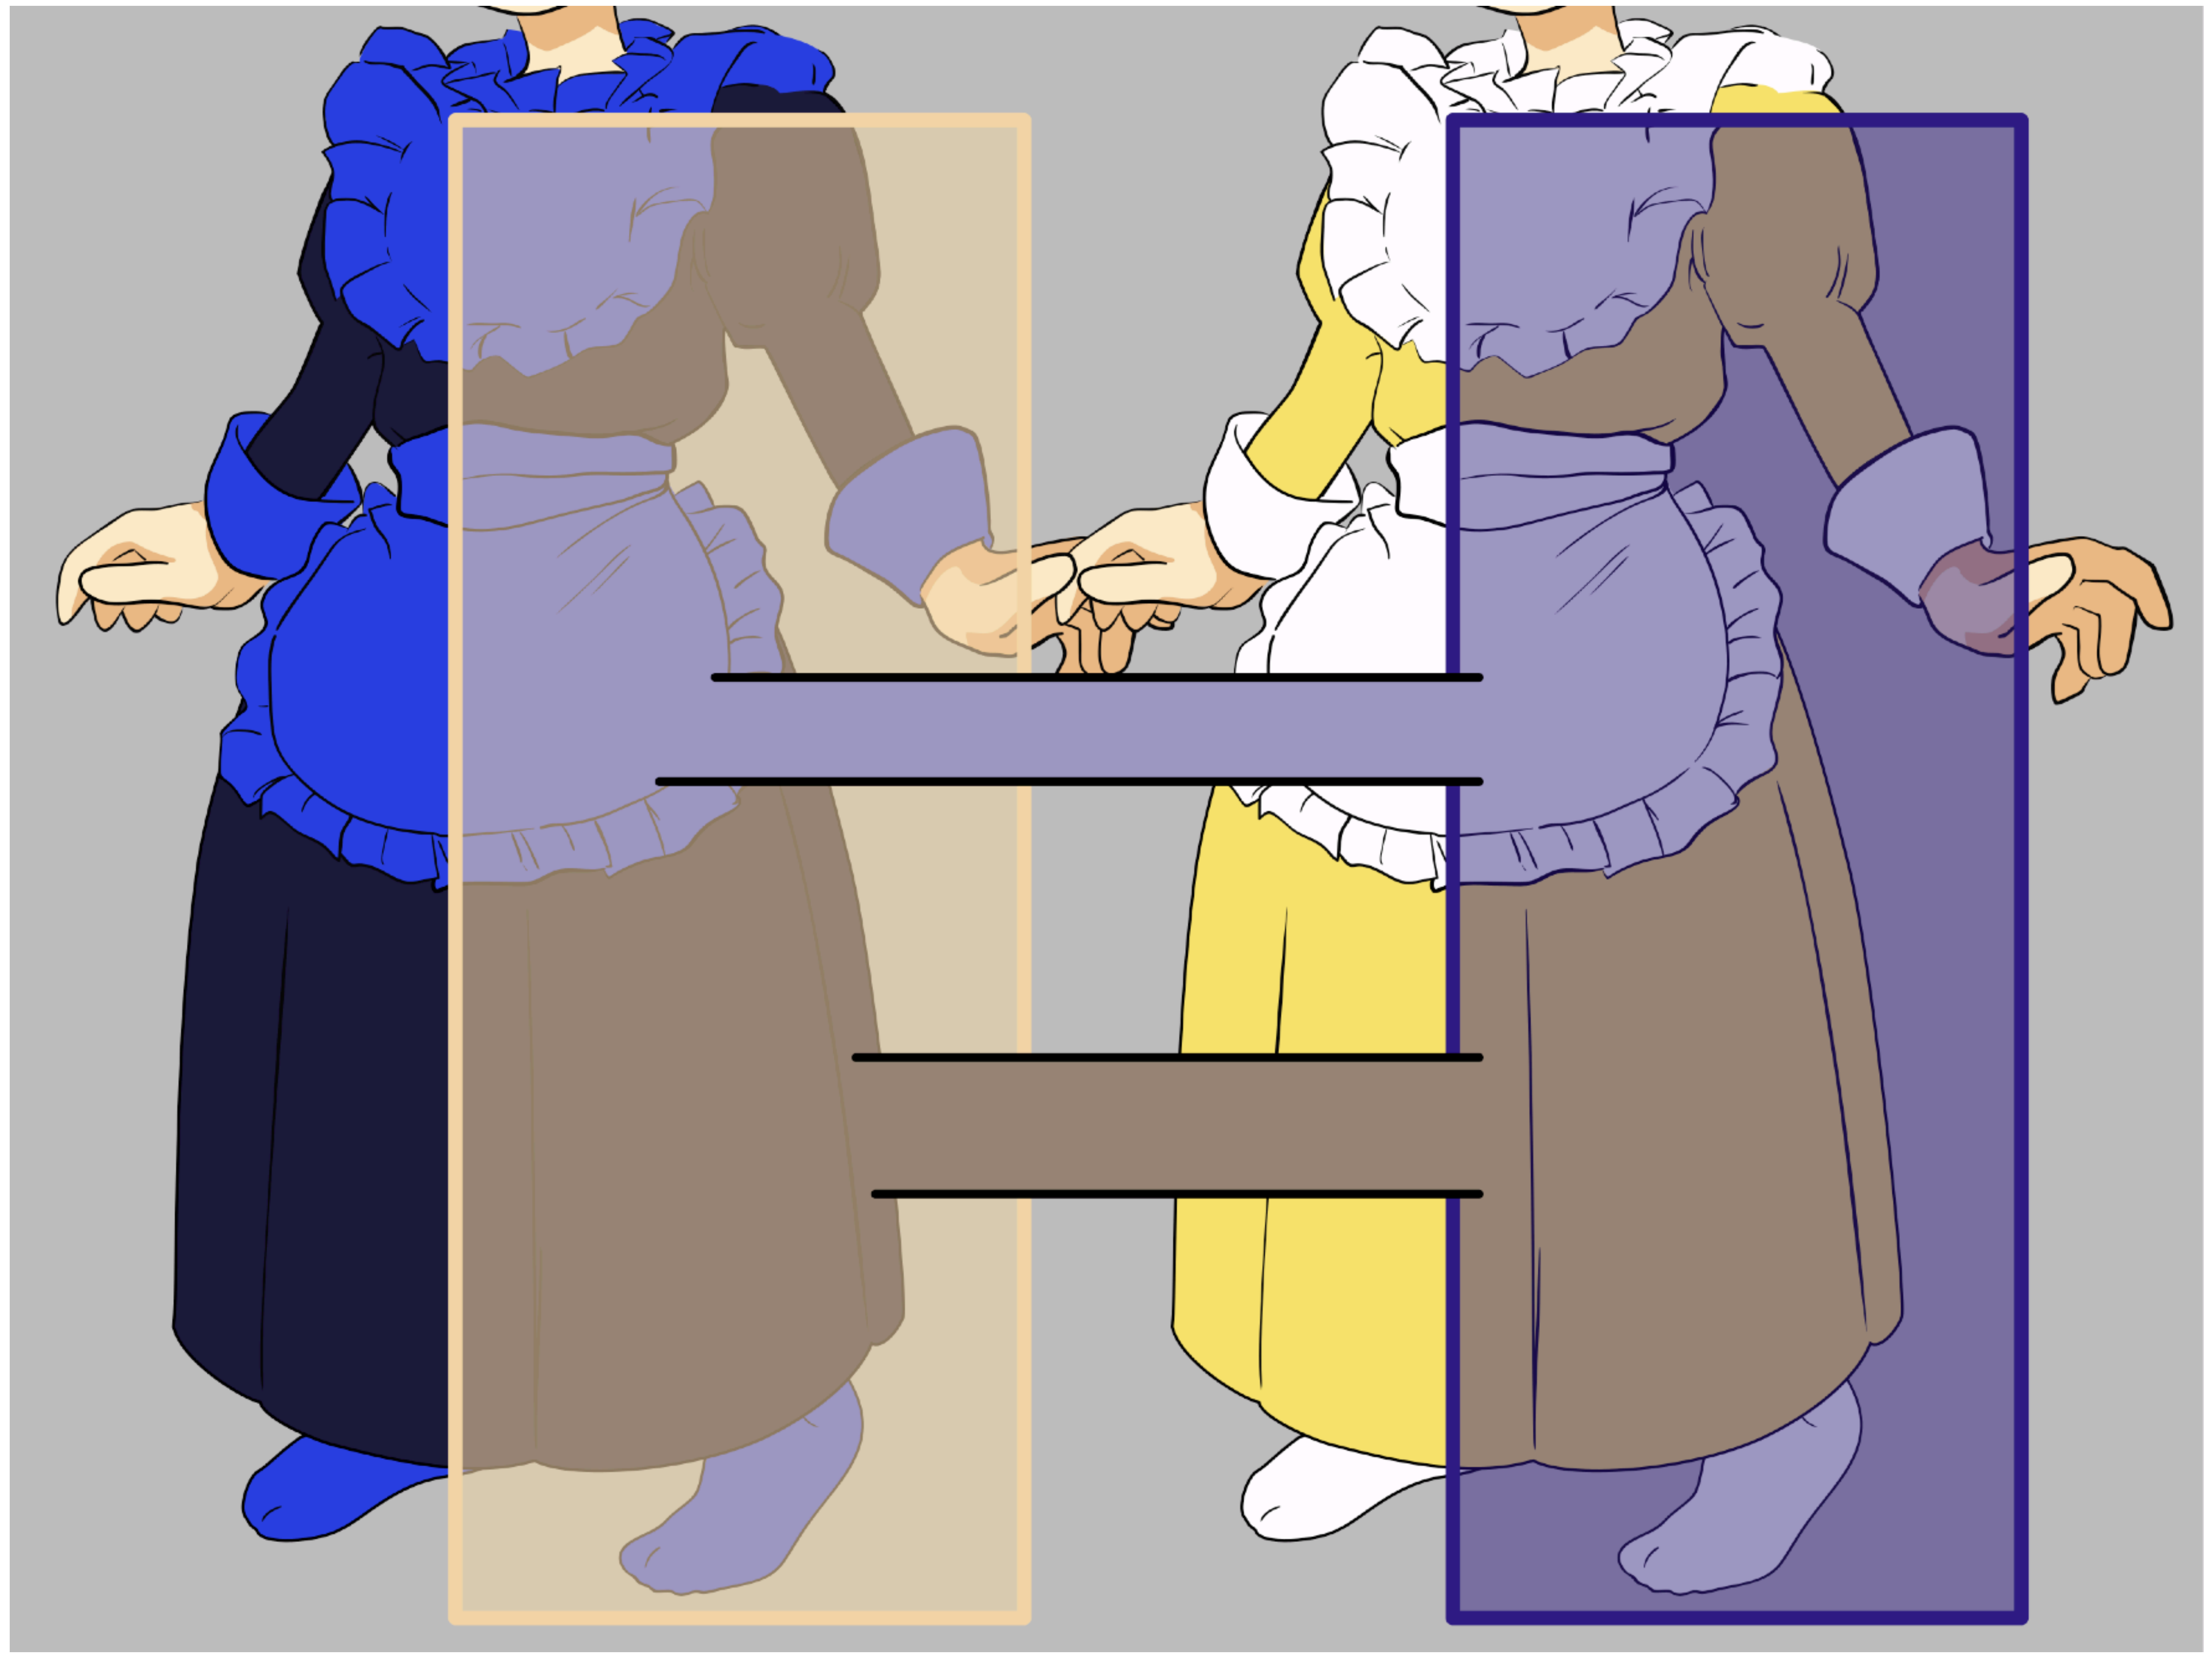

Two interpretations

Source: Wikipedia