Light and Color

Computer Vision

Jia-Bin Huang, Virginia Tech

Empire of Light, 1950 by Rene Magritte

Administrative stuffs

- Signed up Piazza discussion board?
- Sample final project ideas posted
- Installed MATLAB?
  - Akrit (TA) will hold a tutorial session next Friday
- Reviewed Linear Algebra?
- Questions about the course logistics?

Search for Teammates!

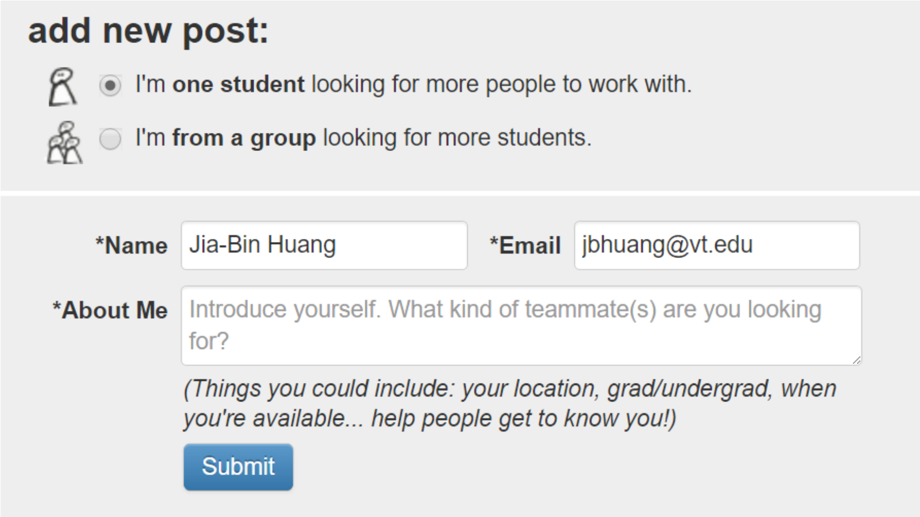

Previous class: Introduction

- Overview of computer vision

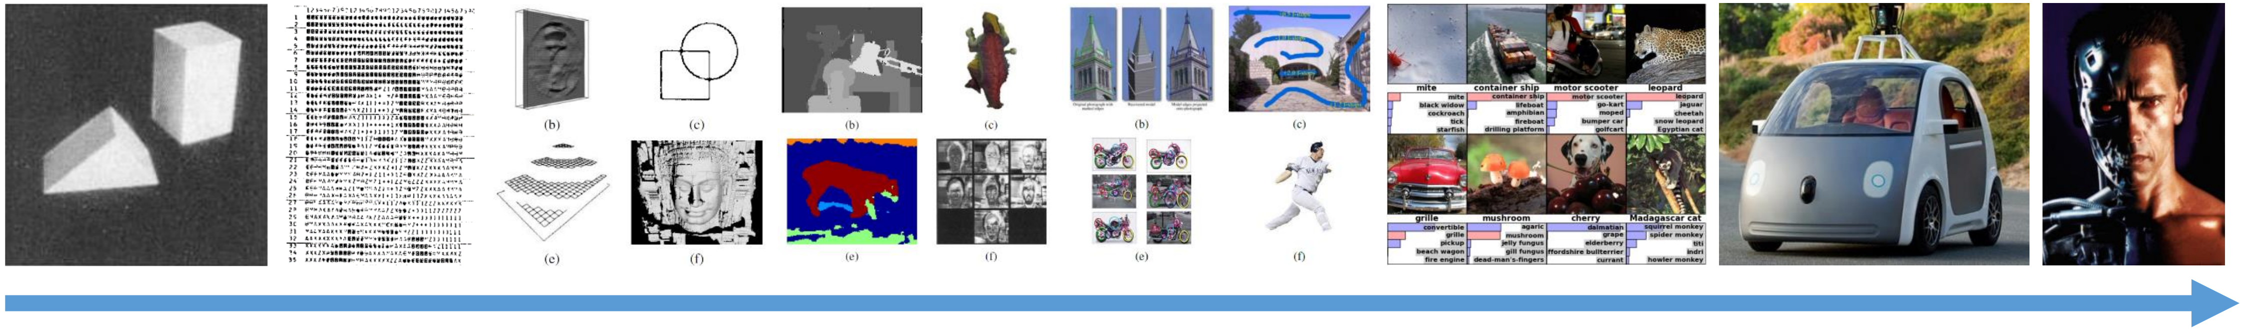

- Examples of computer vision applications

Safety

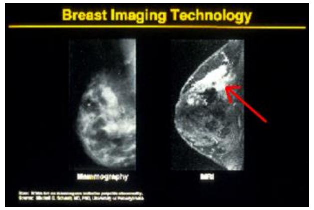

Health

Security

Comfort

Fun

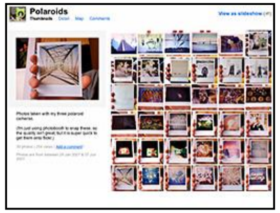

Access

Today’s class

- What determines pixels’ **brightness**?
- What determines pixels’ **color**?
- What can we infer about the scene from pixel intensities?

Why should we care?

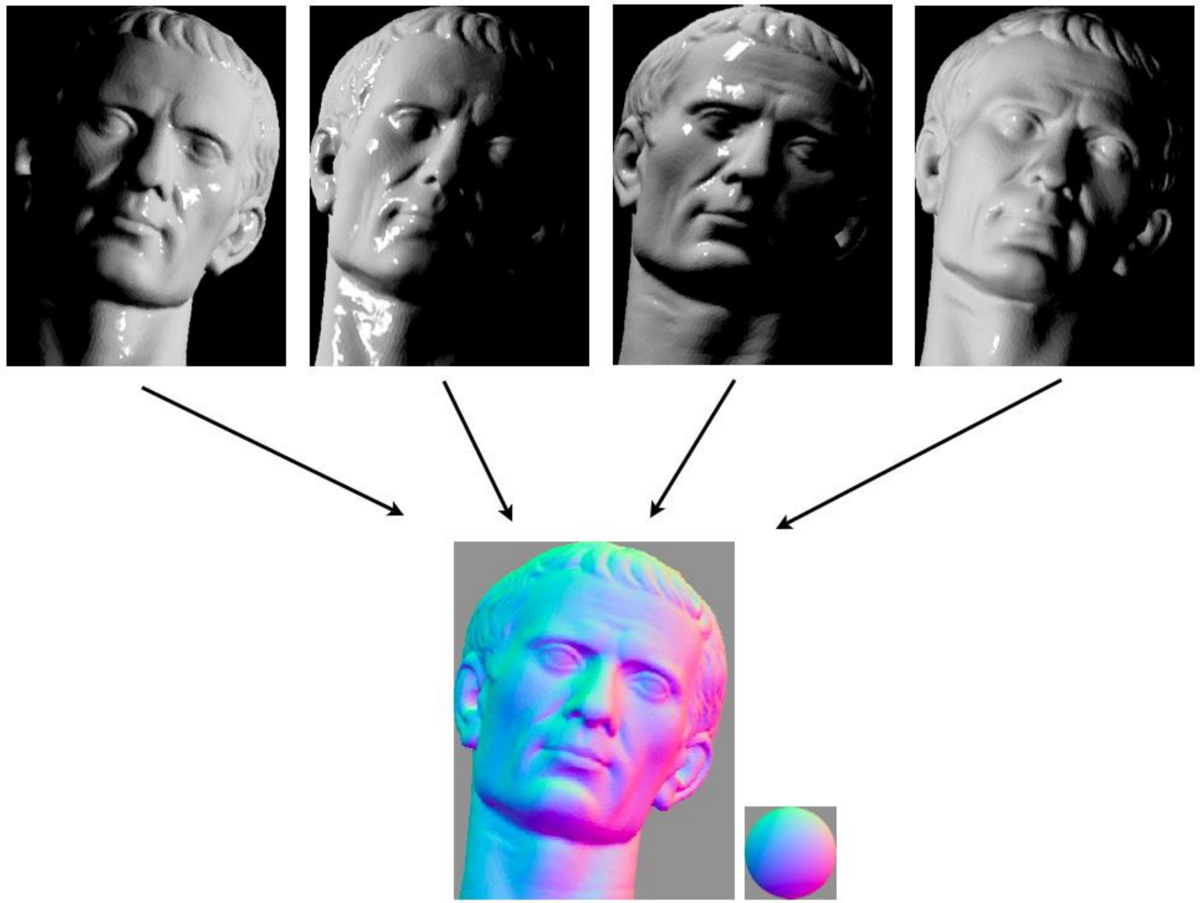

[Photometric Stereo]

Why should we care?

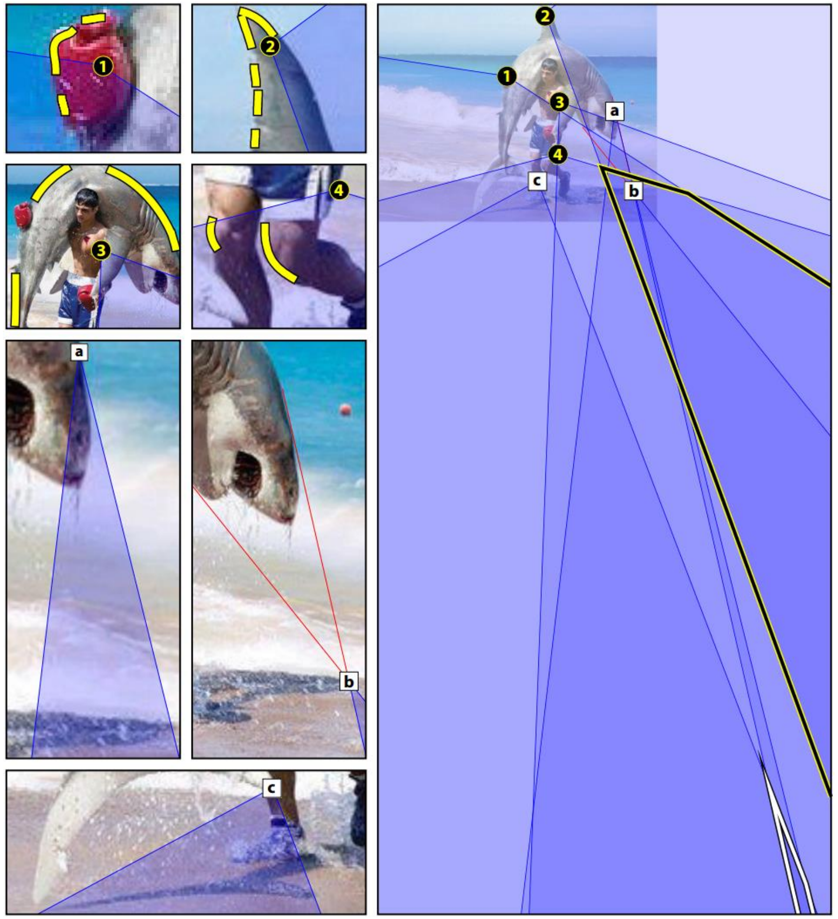

Exposing Photo Manipulation from Shading and Shadows [Kee et al. TOG 14]

Why should we care?

White and gold?

Or

Black and blue?

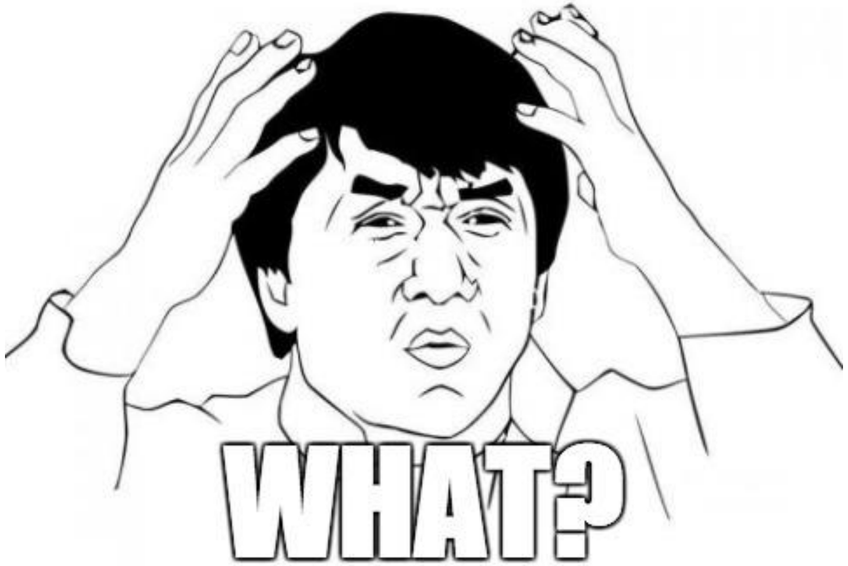

Why should we care?

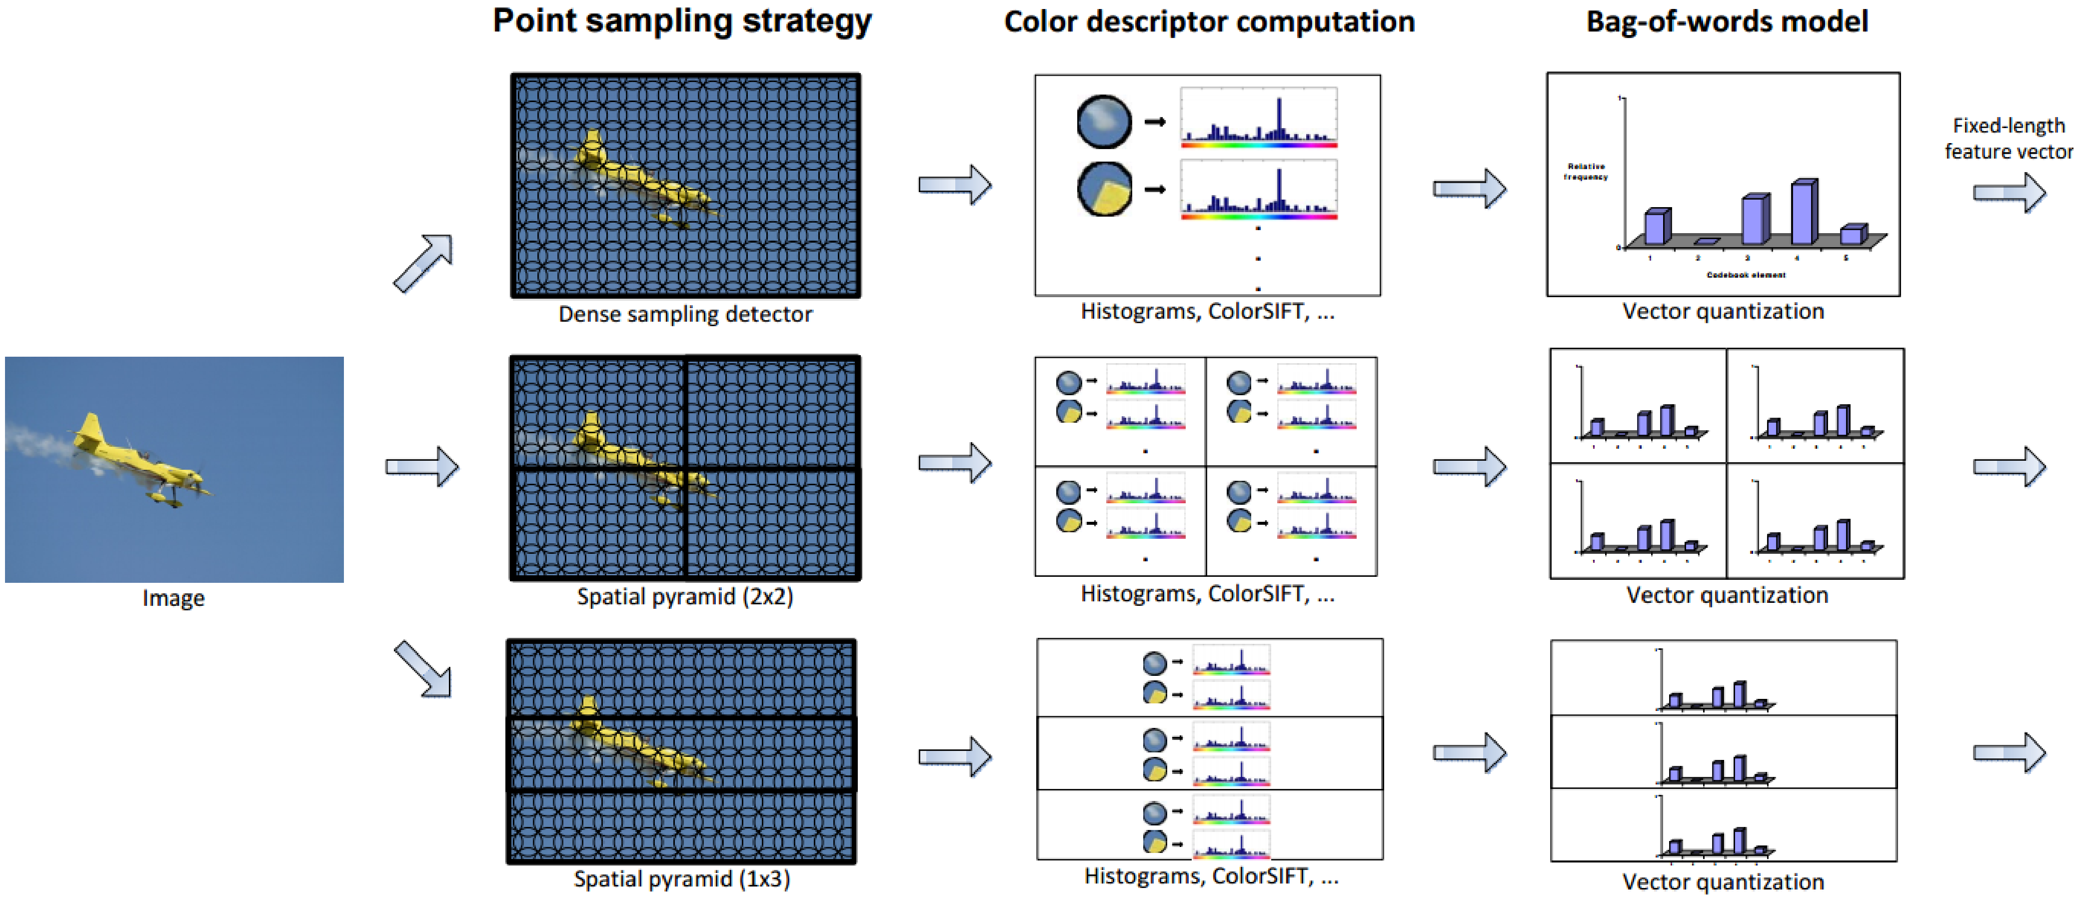

Object and scene categorization [Sande et al. PAMI 2010]

What determines pixels’ brightness?

Image Formation

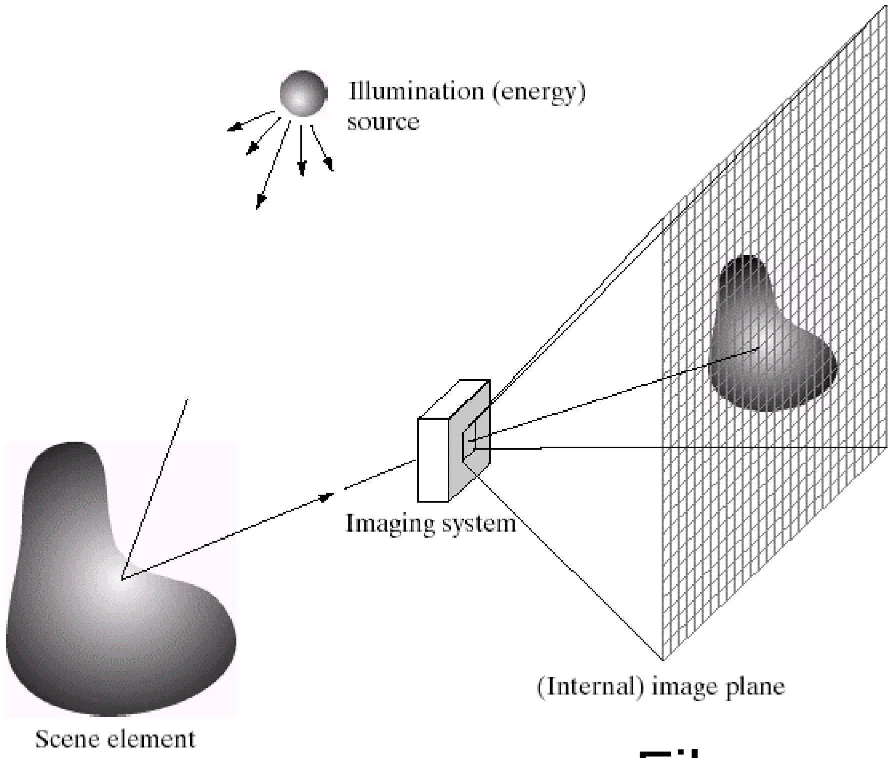

Film

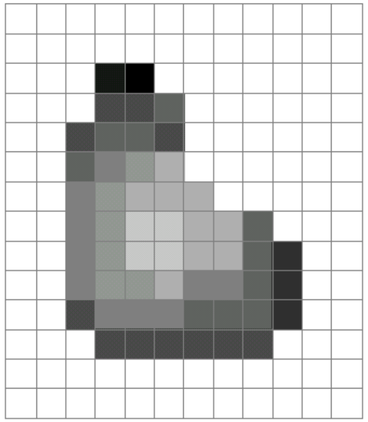

Digital Camera

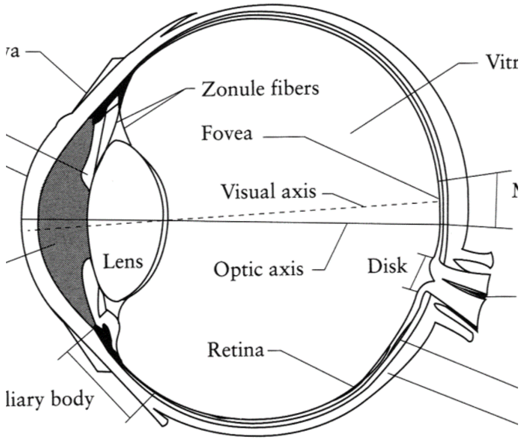

The Eye

Sensor Array

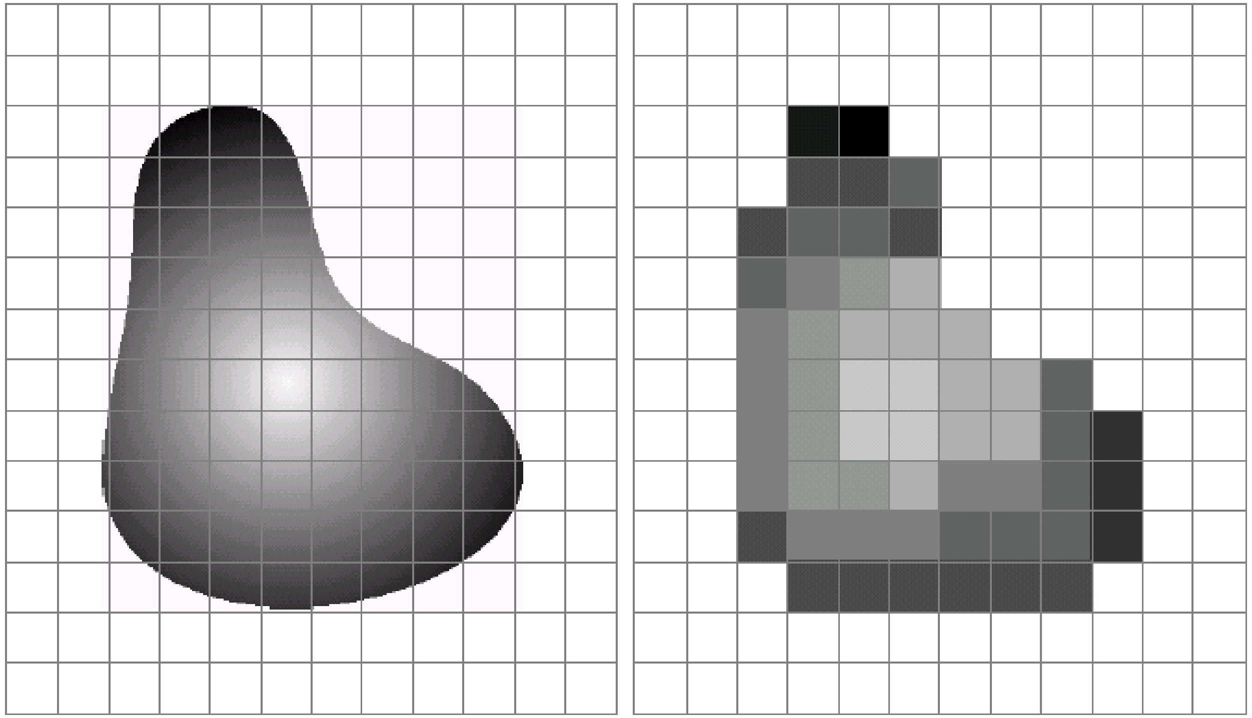

a b

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

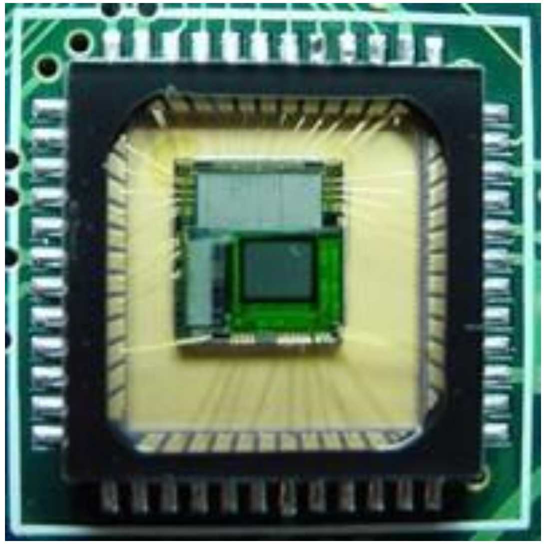

CMOS sensor

What humans see

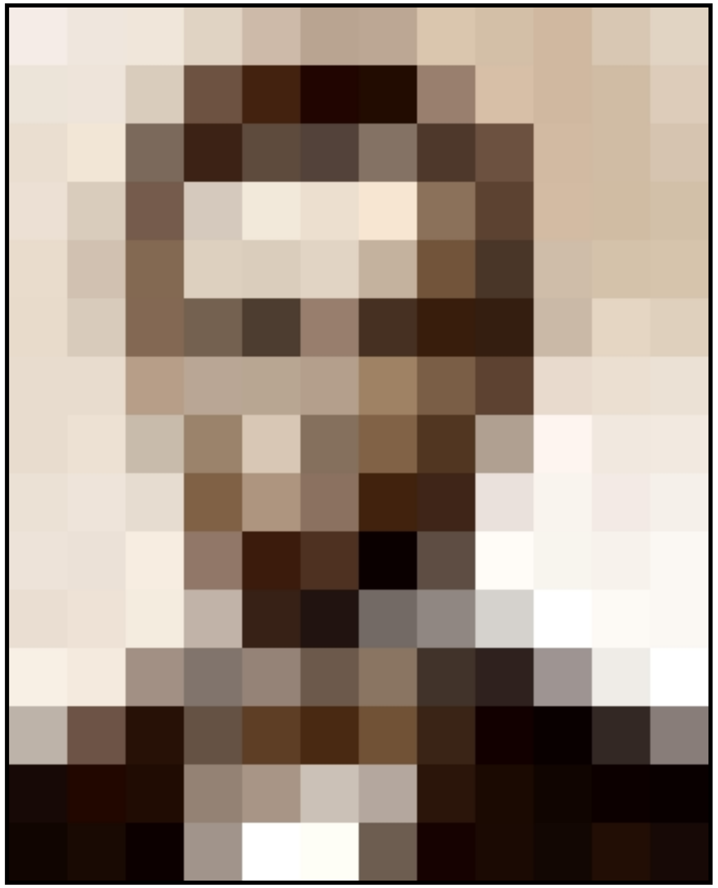

Slide credit: Larry Zitnick

What computers see

243	239	240	225	206	185	188	218	211	206	216	225
242	239	218	110	67	31	34	152	213	206	208	221
243	242	123	58	94	82	132	77	108	208	208	215
235	217	115	212	243	236	247	139	91	209	208	211
233	208	131	222	219	226	196	114	74	208	213	214
232	217	131	116	77	150	69	56	52	201	228	223
232	232	182	186	184	179	159	123	93	232	235	235
232	236	201	154	216	133	129	81	175	252	241	240
235	238	230	128	172	138	65	63	234	249	241	245
237	236	247	143	59	78	10	94	255	248	247	251
234	237	245	193	55	33	115	144	213	255	253	251
248	245	161	128	149	109	138	65	47	156	239	255
190	107	39	102	94	73	114	58	17	7	51	137
23	32	33	148	168	203	179	43	27	17	12	8
17	26	12	160	255	255	109	22	26	19	35	24

Slide credit: Larry Zitnick

How does a pixel get its value?

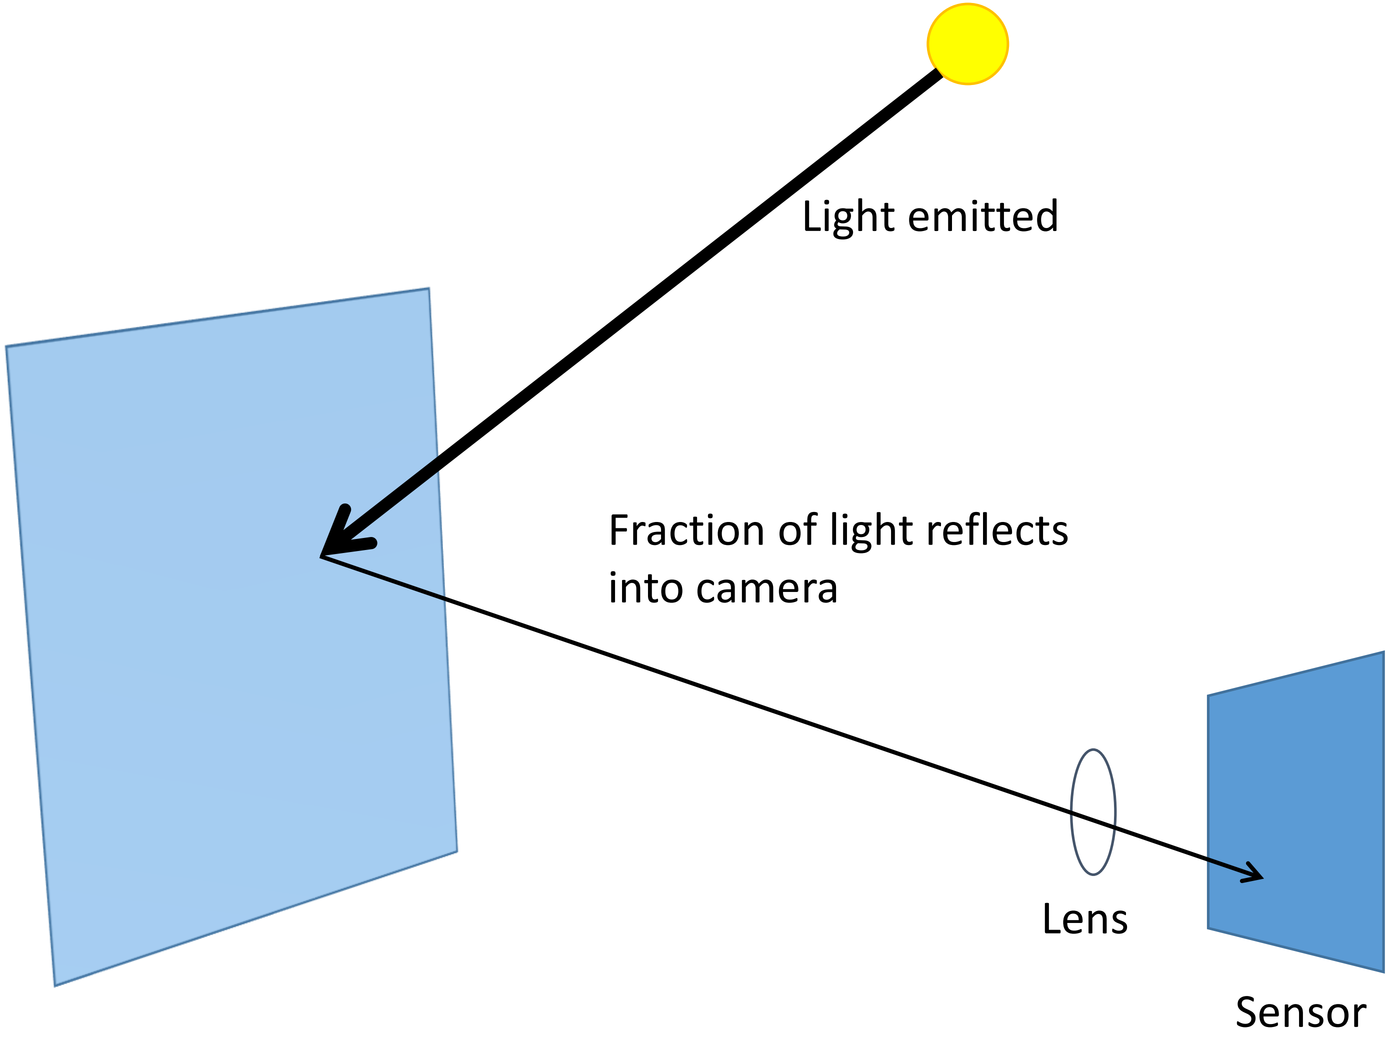

Slide credit: Derek Hoiem

How does a pixel get its value?

- Major factors
  - Illumination strength and direction
  - Surface geometry
  - Surface material
  - Nearby surfaces
  - Camera gain/exposure

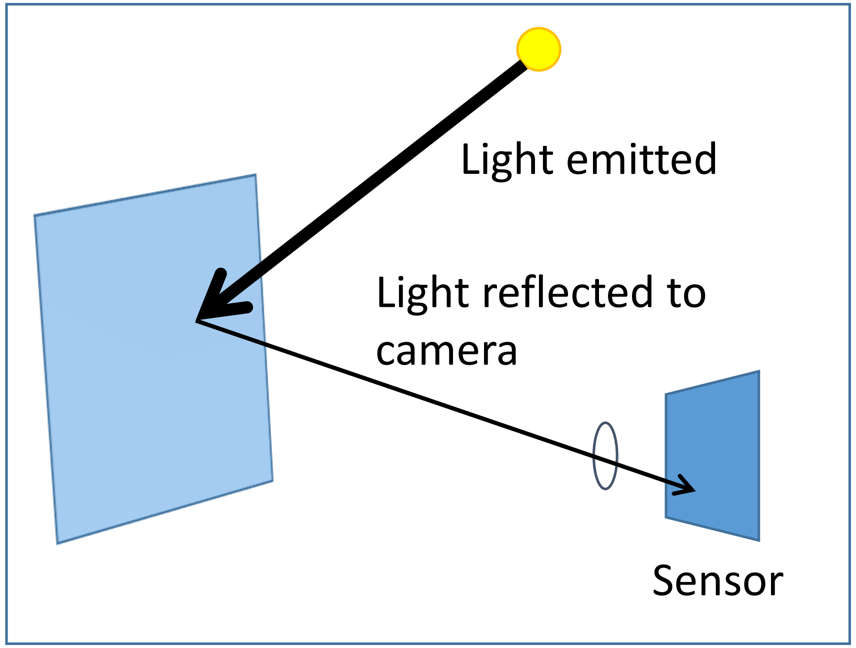

Slide credit: Derek Hoiem

Basic models of reflection

- Specular: light bounces off at the incident angle
  - E.g., mirror

- Diffuse: light scatters in all directions
  - E.g., brick, cloth, rough wood

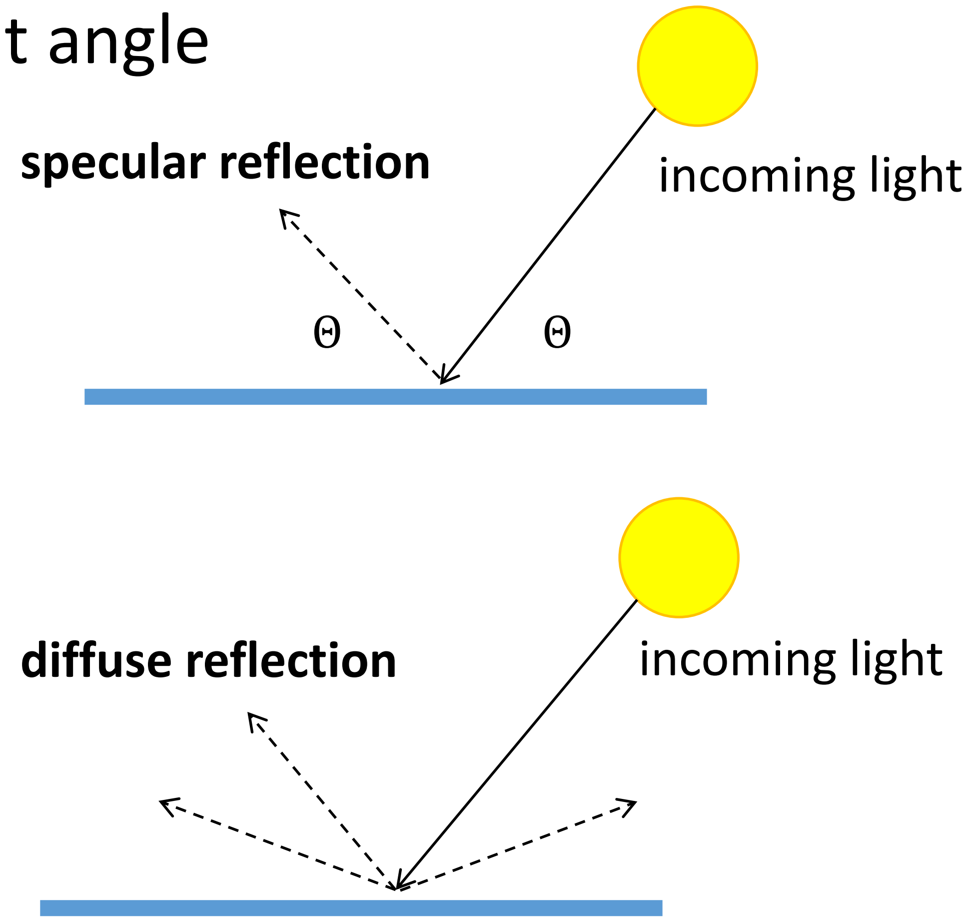

Lambertian reflectance model

- Some light is absorbed (function of albedo ρ)
- Remaining light is scattered (diffuse reflection)
- Examples: soft cloth, concrete, matte paints

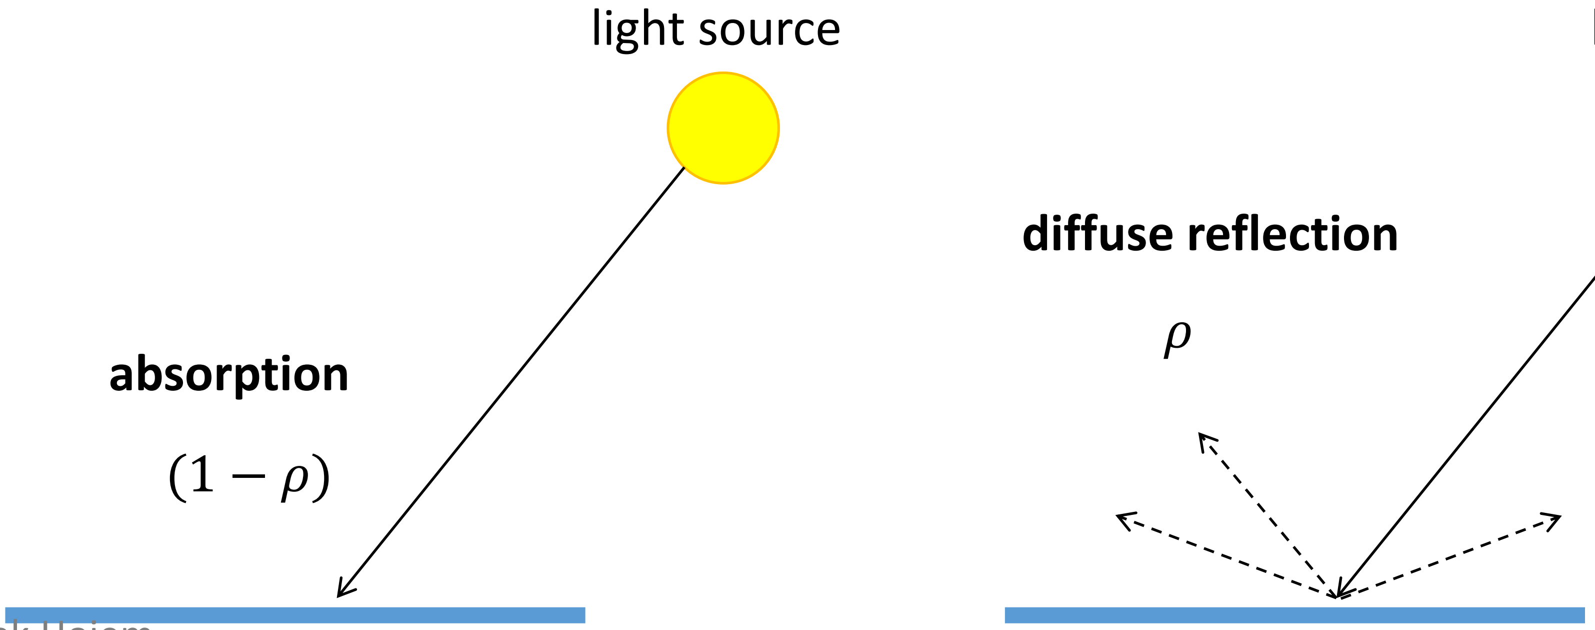

Lambertian reflectance model

- Some light is absorbed (function of albedo ρ)
- Remaining light is scattered (diffuse reflection)
- Examples: soft cloth, concrete, matte paints

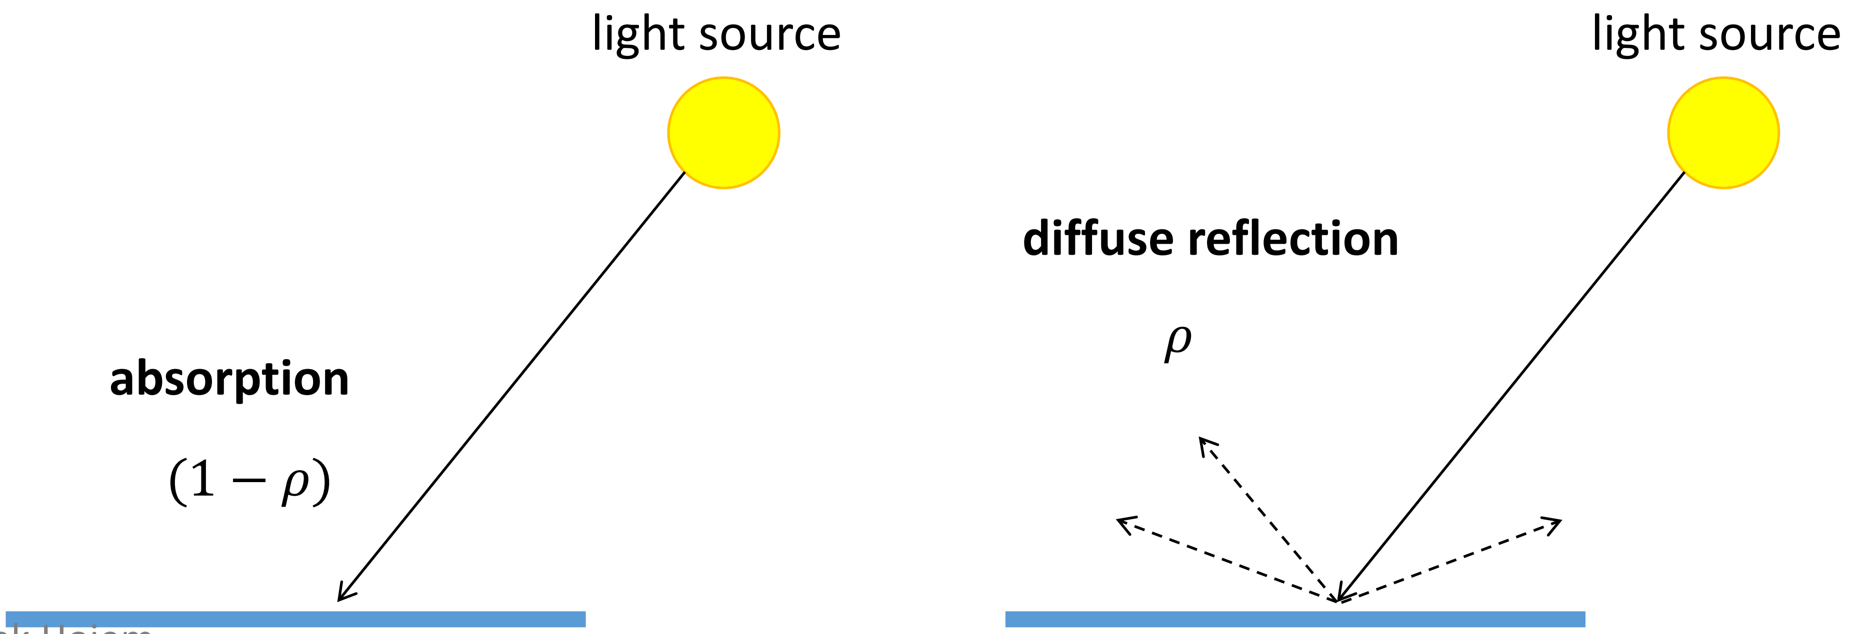

Slide credit: Derek Hoiem

Diffuse reflection: Lambert's cosine law

Intensity does *not* depend on viewer angle.

- Amount of reflected light proportional to $\cos(\theta)$
- Visible solid angle also proportional to $\cos(\theta)$

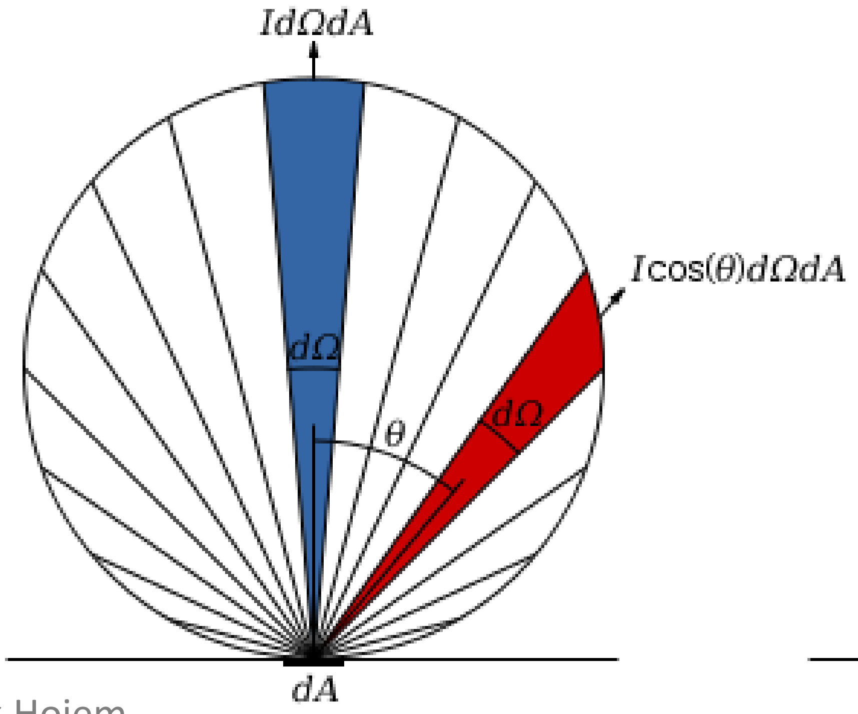

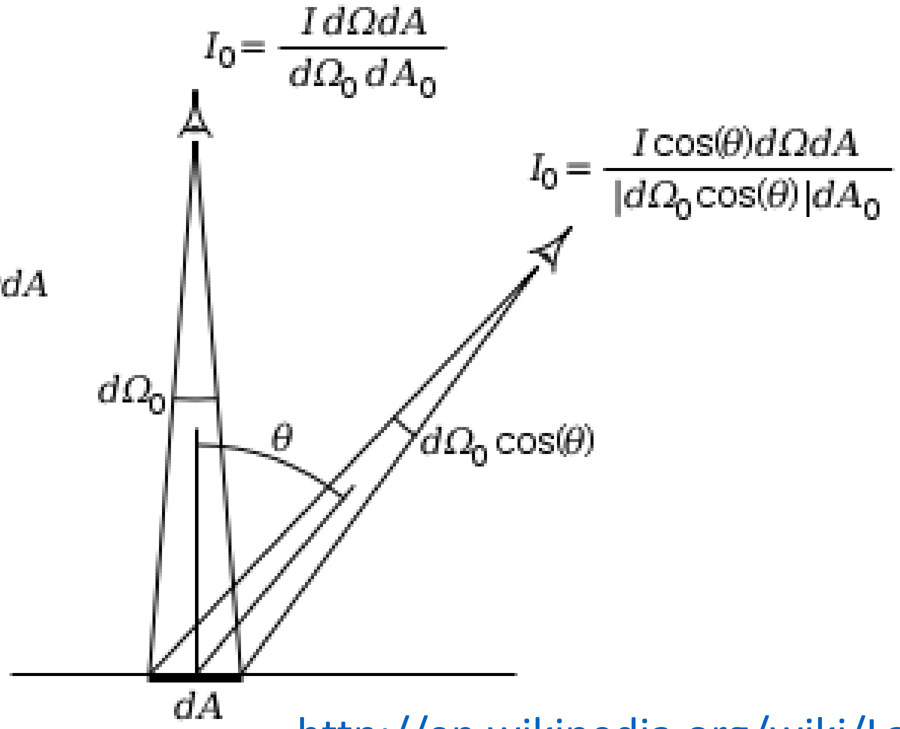

Slide credit: Derek Hoiem

http://en.wikipedia.org/wiki/Lambert%27s_cosine_law

Specular Reflection

- Reflected direction depends on light orientation and surface normal
  - E.g., mirrors are fully specular
  - Most surfaces can be modeled with a mixture of diffuse and specular components

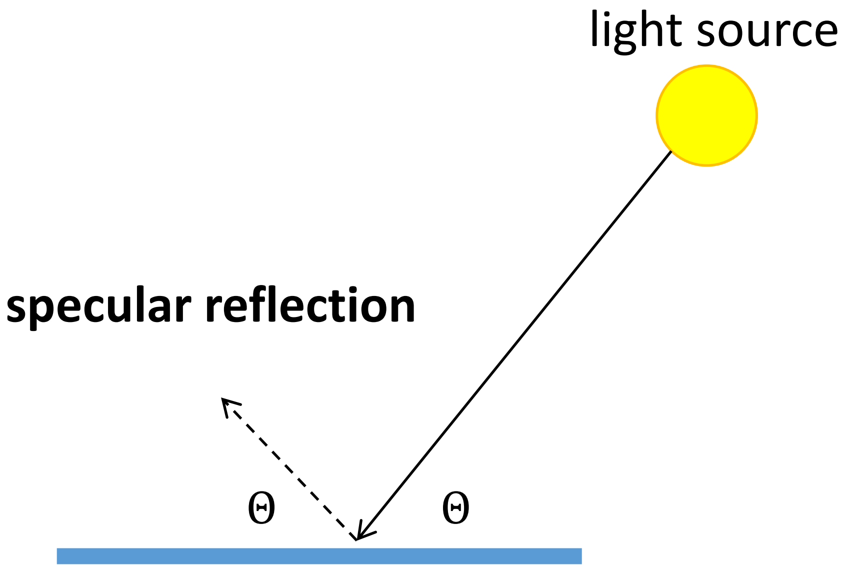

Flickr, by suzysputnik

Flickr, by piratejohnny

Slide credit: Derek Hoiem

Most surfaces have both specular and diffuse components

- Specularity = spot where specular reflection dominates (typically reflects light source)

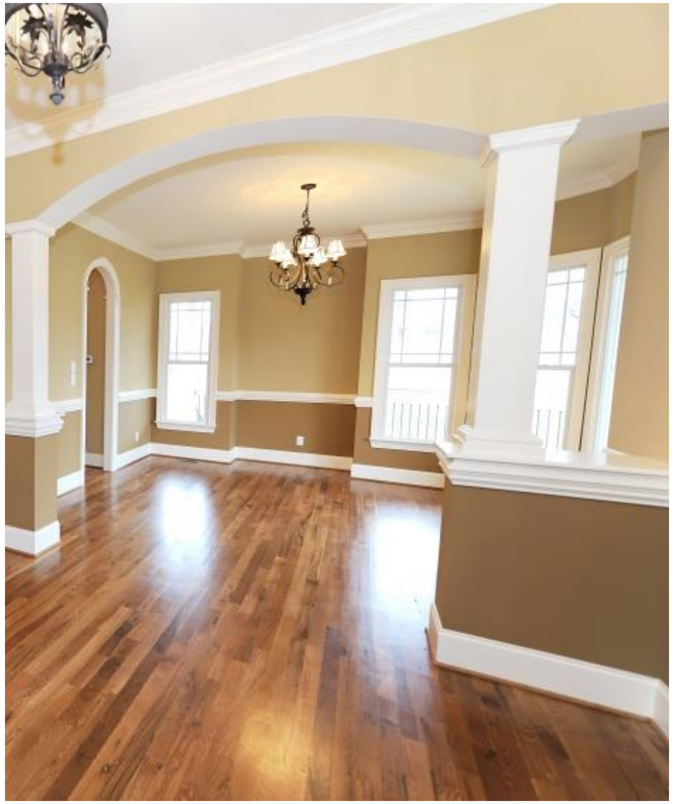

Photo: northcountryhardwoodfloors.com

Typically, specular component is small

Slide credit: Derek Hoiem

Intensity and Surface Orientation

Intensity depends on illumination angle because less light comes in at oblique angles.

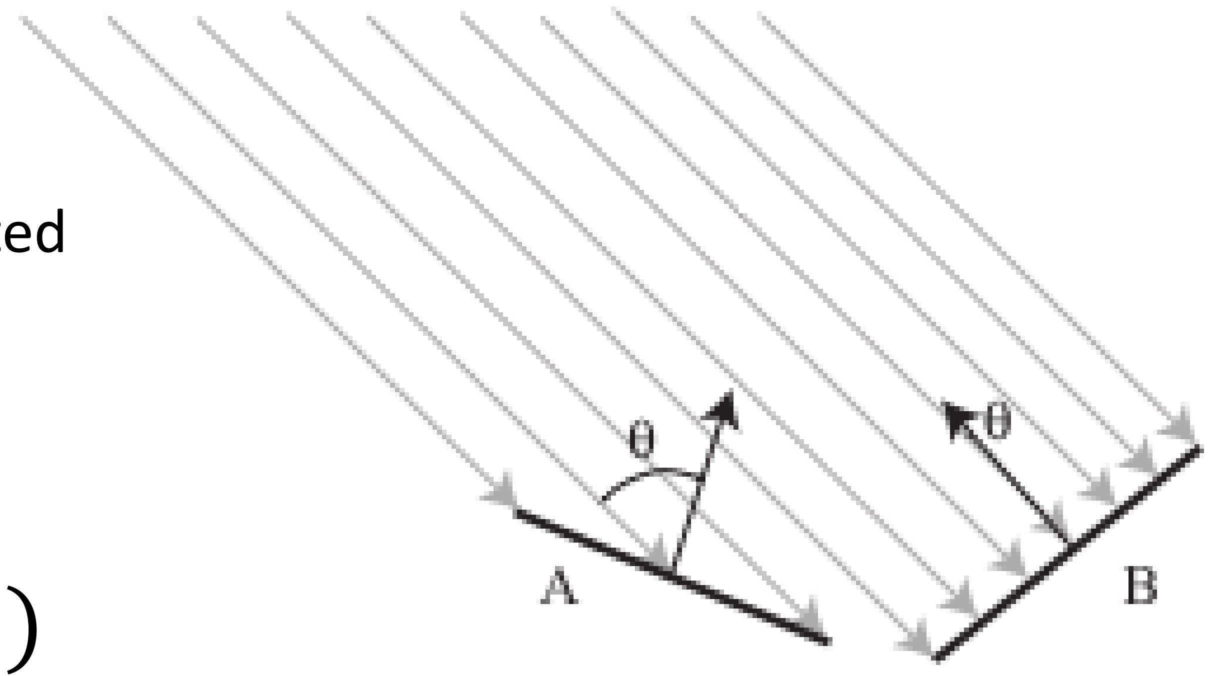

ρ = Albedo: fraction of light that is reflected

$\boldsymbol{S}$ = directional source

$\boldsymbol{N}$ = surface normal

I = reflected intensity

$$I(x) = \rho(x)(\boldsymbol{S} \cdot \boldsymbol{N}(x))$$


1
2

Recap

- When light hits a typical surface
  - Some light is absorbed (1-ρ)
    - More absorbed for low albedos
  - Some light is reflected diffusely
    - Independent of viewing direction
  - Some light is reflected specularly
    - Light bounces off (like a mirror), depends on viewing direction

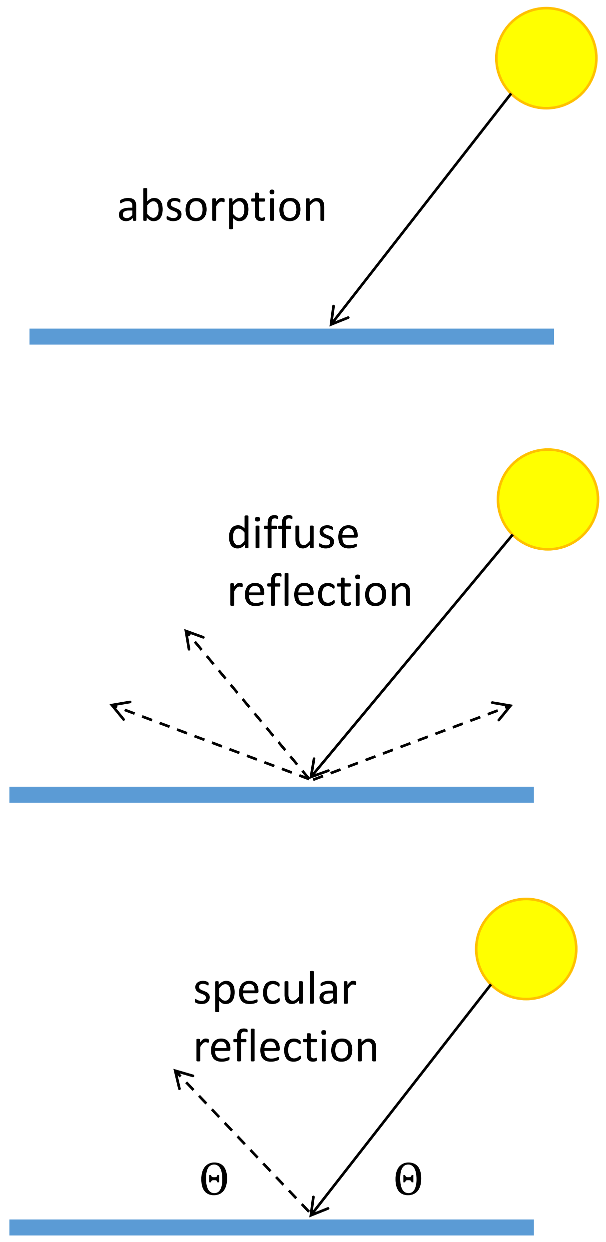

Other possible effects

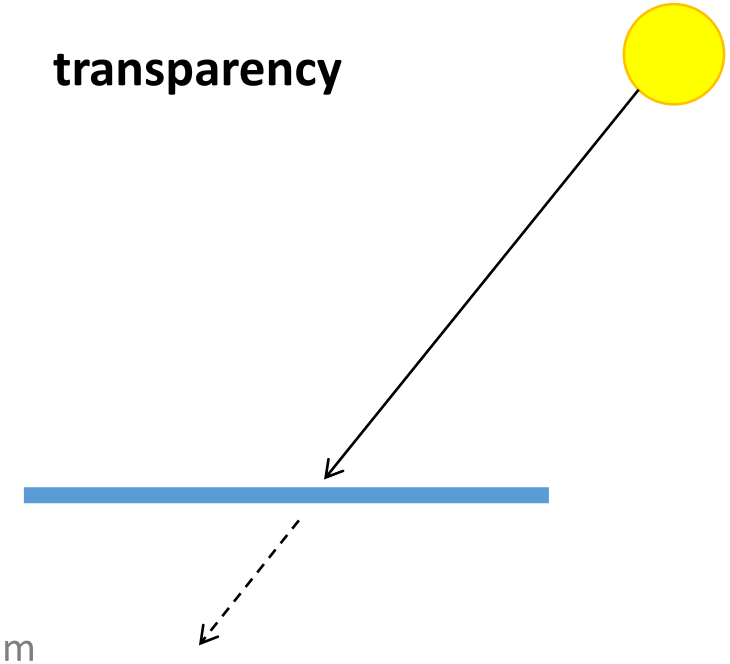

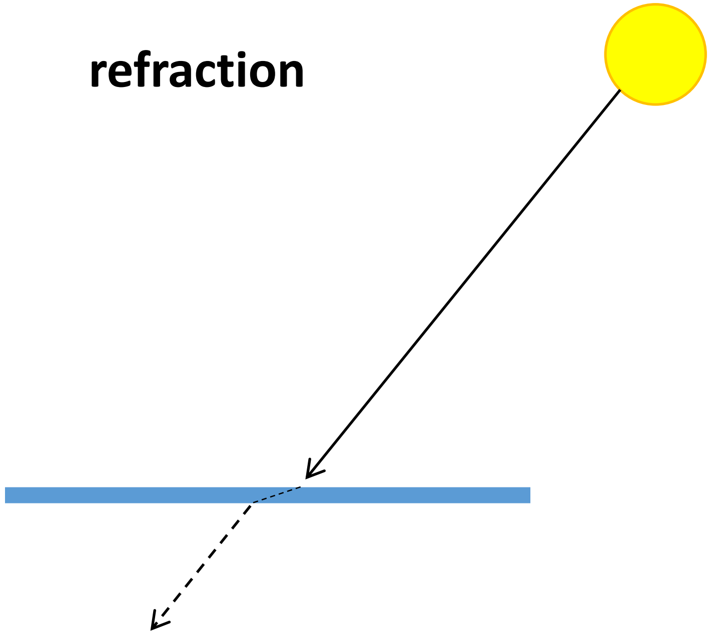

fluorescence

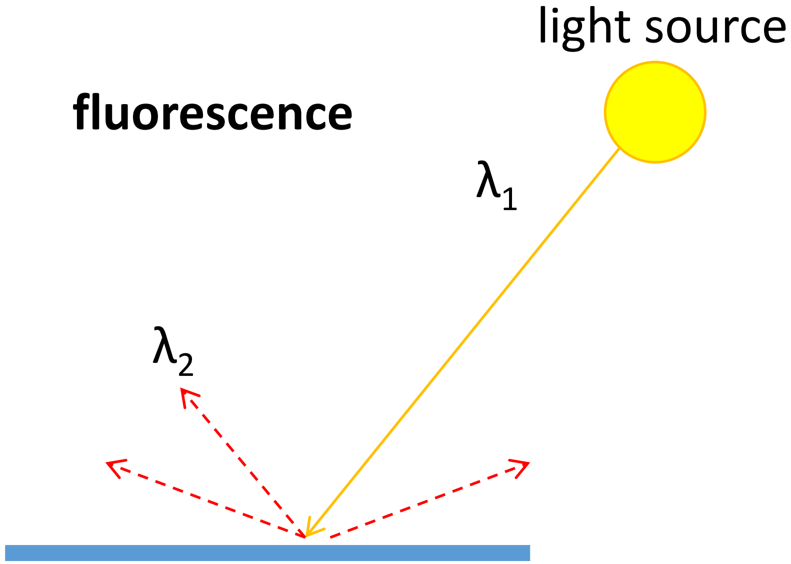

phosphorescence

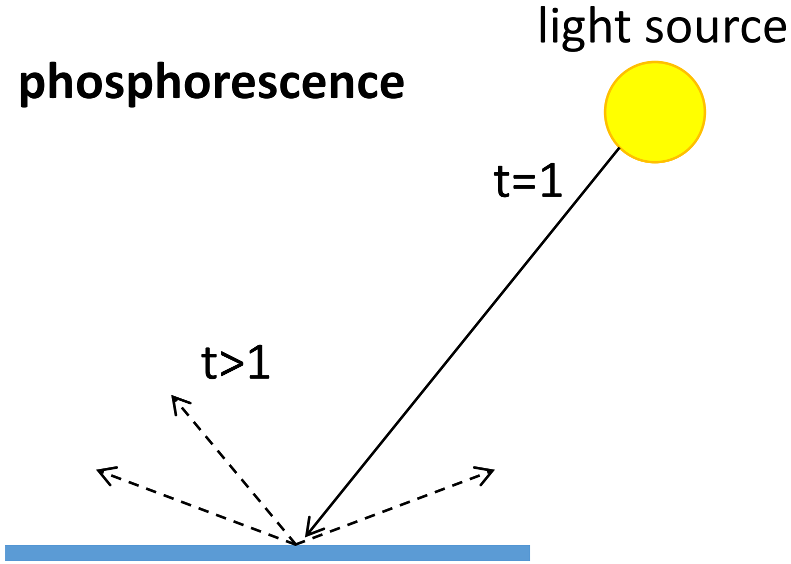

Slide credit: Derek Hoiem

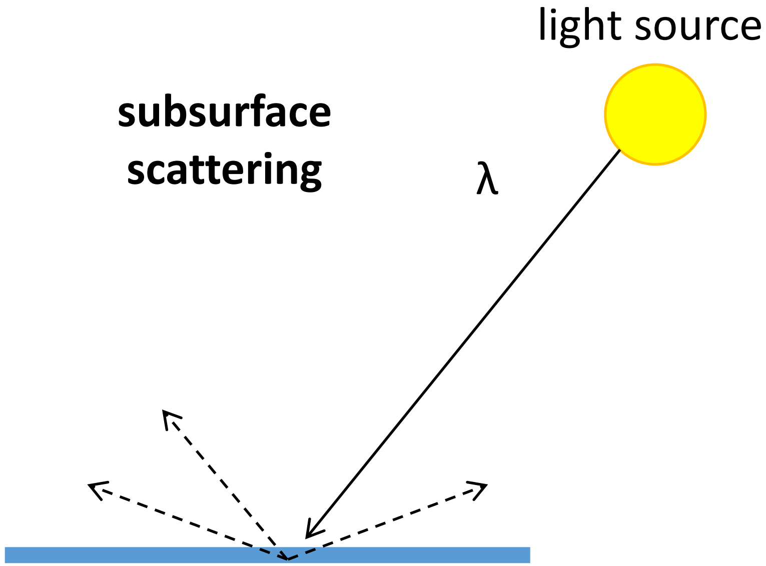

Slide credit: Derek Hoiem

BRDF: Bidirectional Reflectance Distribution Function

- Model of local reflection that tells how bright a surface appears when viewed from one direction when light falls on it from another

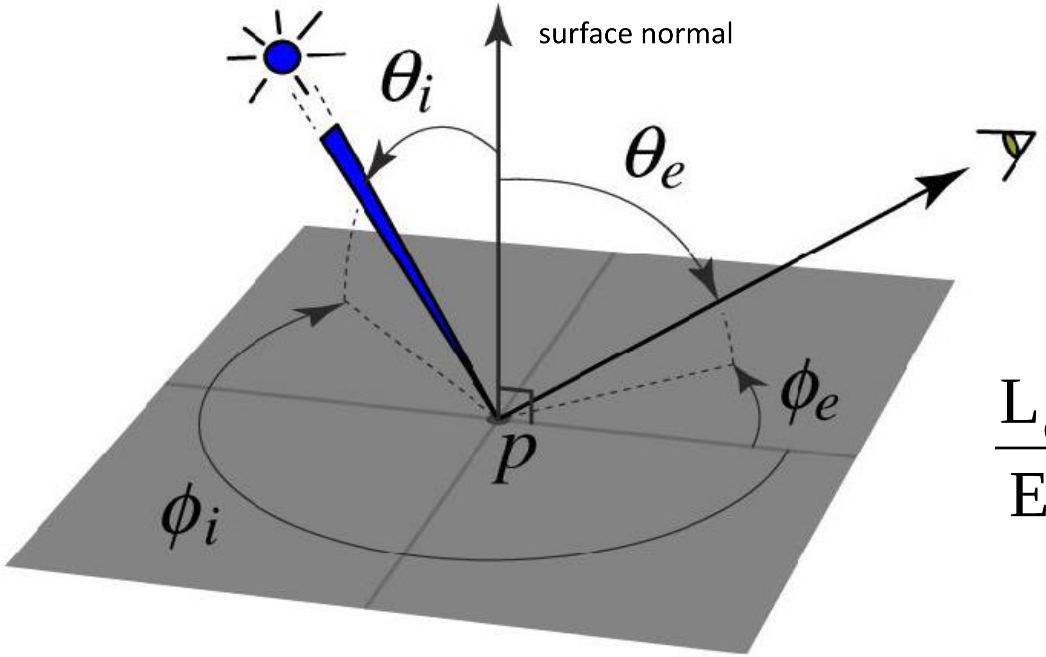

$$\rho(\theta_i, \phi_i, \theta_e, \phi_e; \lambda) =$$

$$\frac{L_e(\boldsymbol{\theta}_e, \boldsymbol{\phi}_e)}{E_i(\boldsymbol{\theta}_i, \boldsymbol{\phi}_i)} = \frac{L_e(\boldsymbol{\theta}_e, \boldsymbol{\phi}_e)}{L_i(\boldsymbol{\theta}_i, \boldsymbol{\phi}_i) \cos \boldsymbol{\theta}_i d\boldsymbol{\omega}}$$

Reflection models

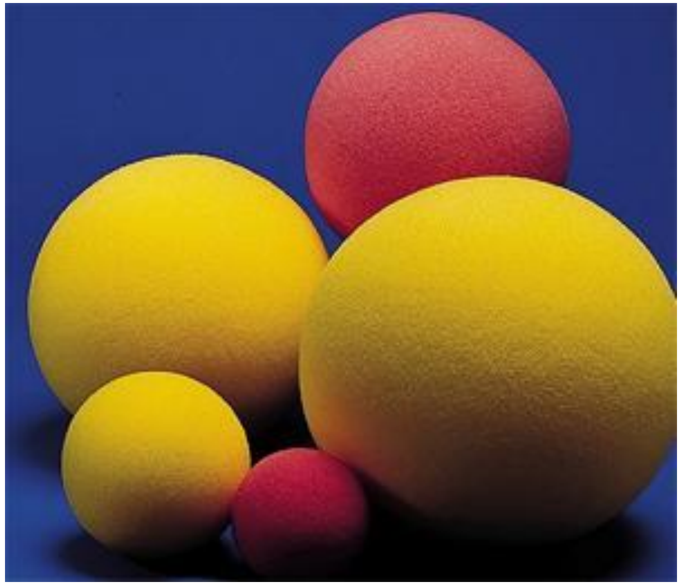

Lambertian:
reflection all diffuse

Mirrored: reflection
all specular

Glossy: reflection mostly
diffuse, some specular

Dynamic range and camera response

- Typical scenes have a huge dynamic range
- Camera response is roughly linear in the mid range (15 to 240) but non-linear at the extremes
  - called saturation or undersaturation

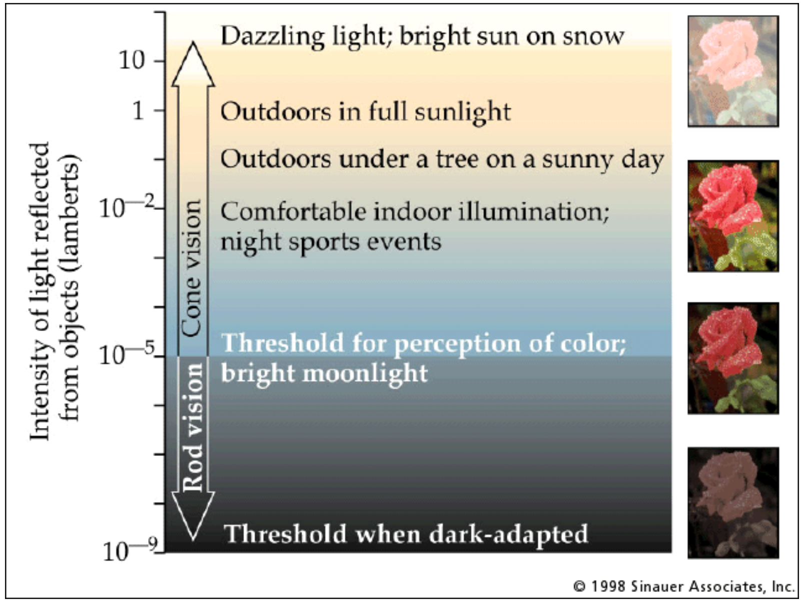

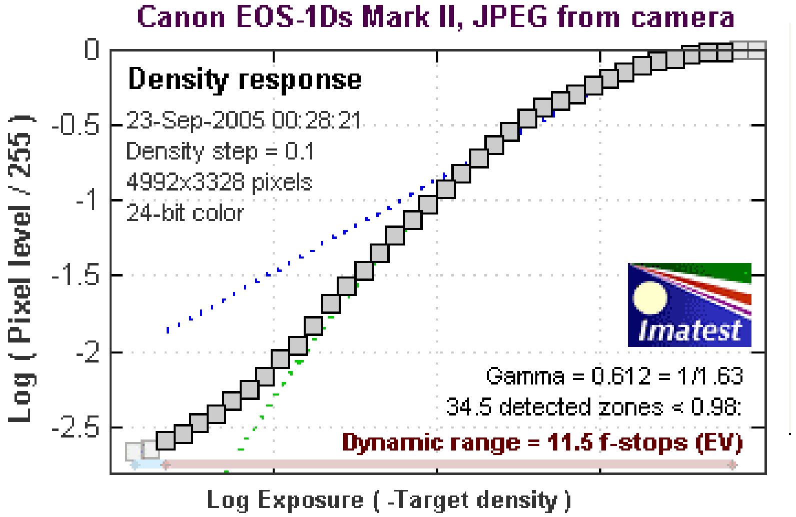

What determines pixels' color?

https://upload.wikimedia.org/wikipedia/commons/b/b1/Colouring_pencils.jpg

The Eye

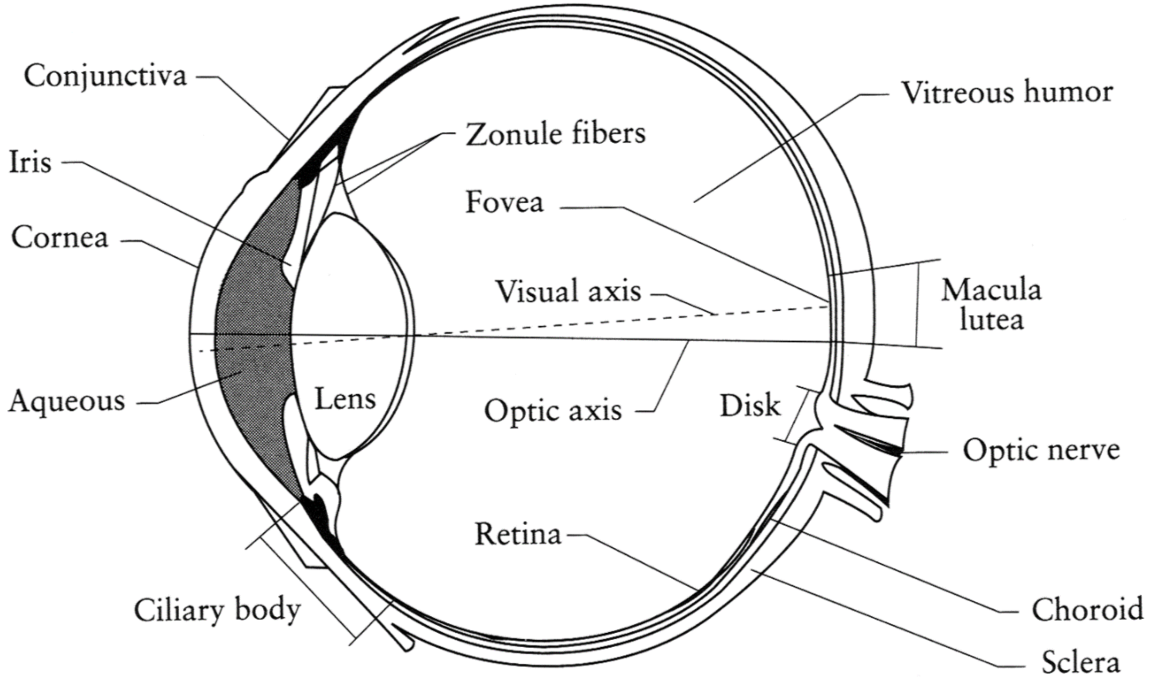

- The human eye is a camera!
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the "film"?
    - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

Retina up-close

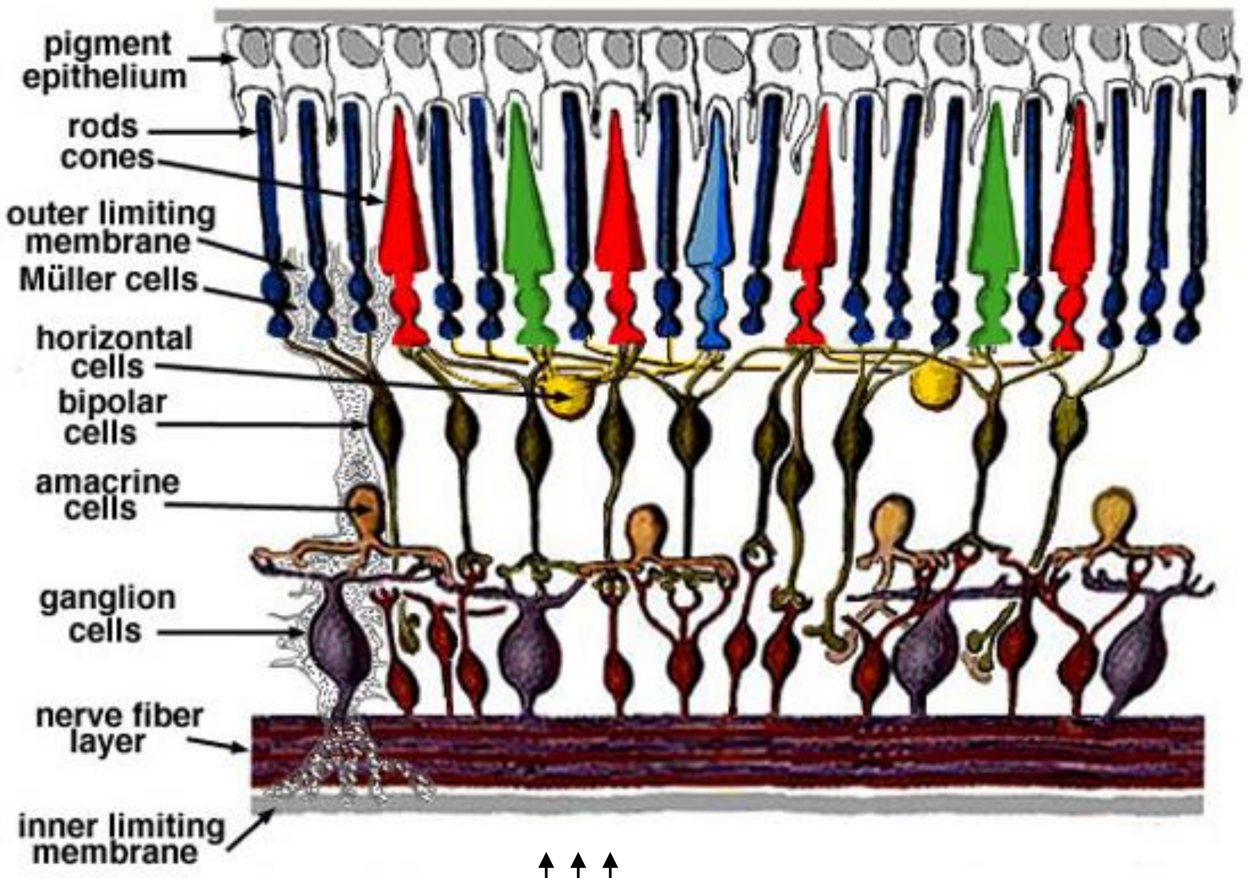

↑↑↑

Light

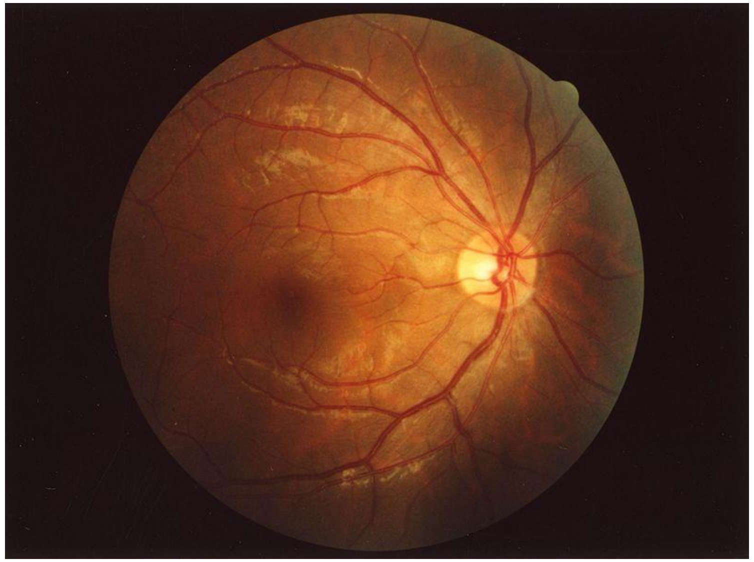

Two types of light-sensitive receptors

Cones

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

Rods

- rod-shaped
- highly sensitive
- operate at night
- gray-scale vision
- slower to respond

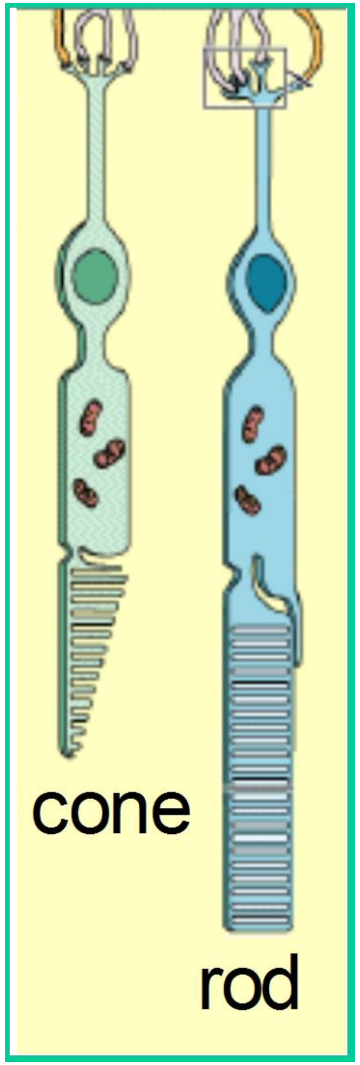

Slide Credit: Efros

Distribution of Rods and Cones

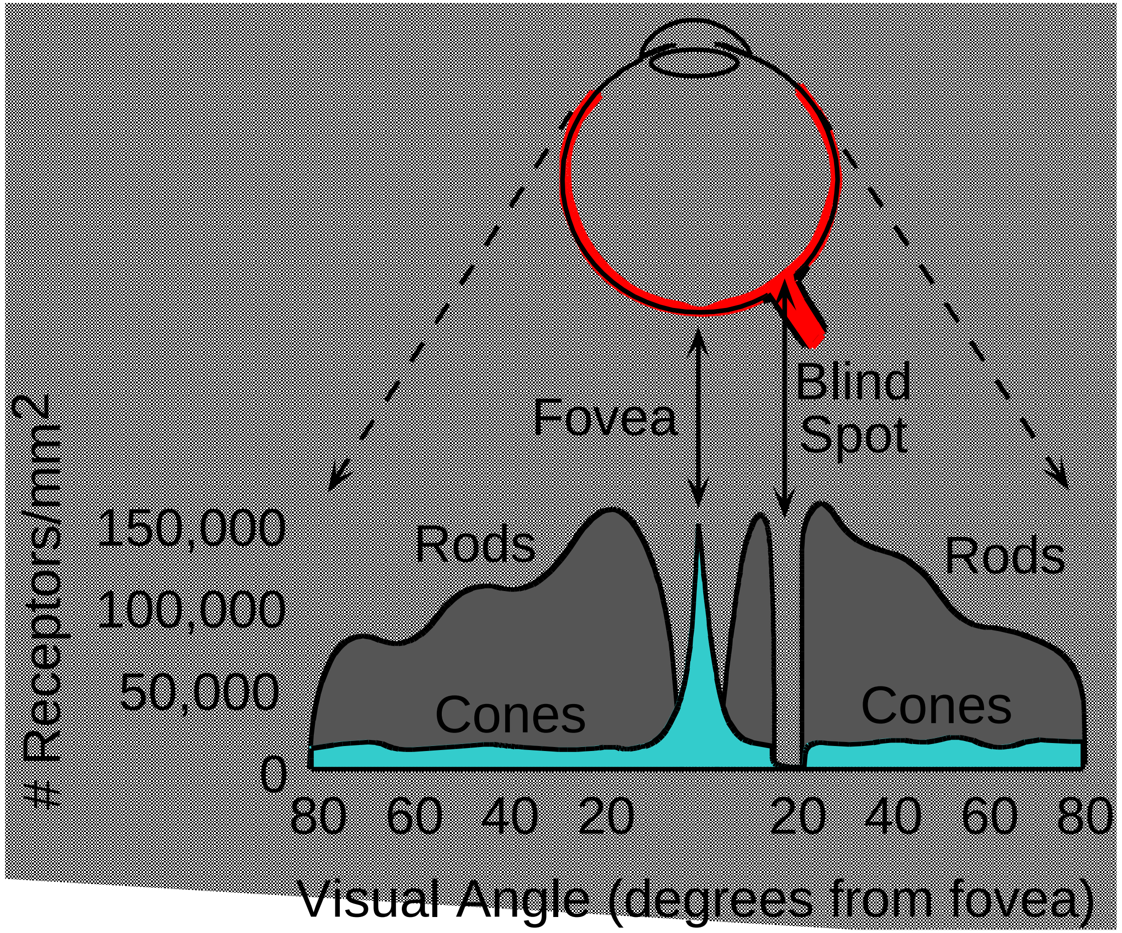

Night Sky: why are there more stars off-center?

Slide credit: Efros

Find your blind spot

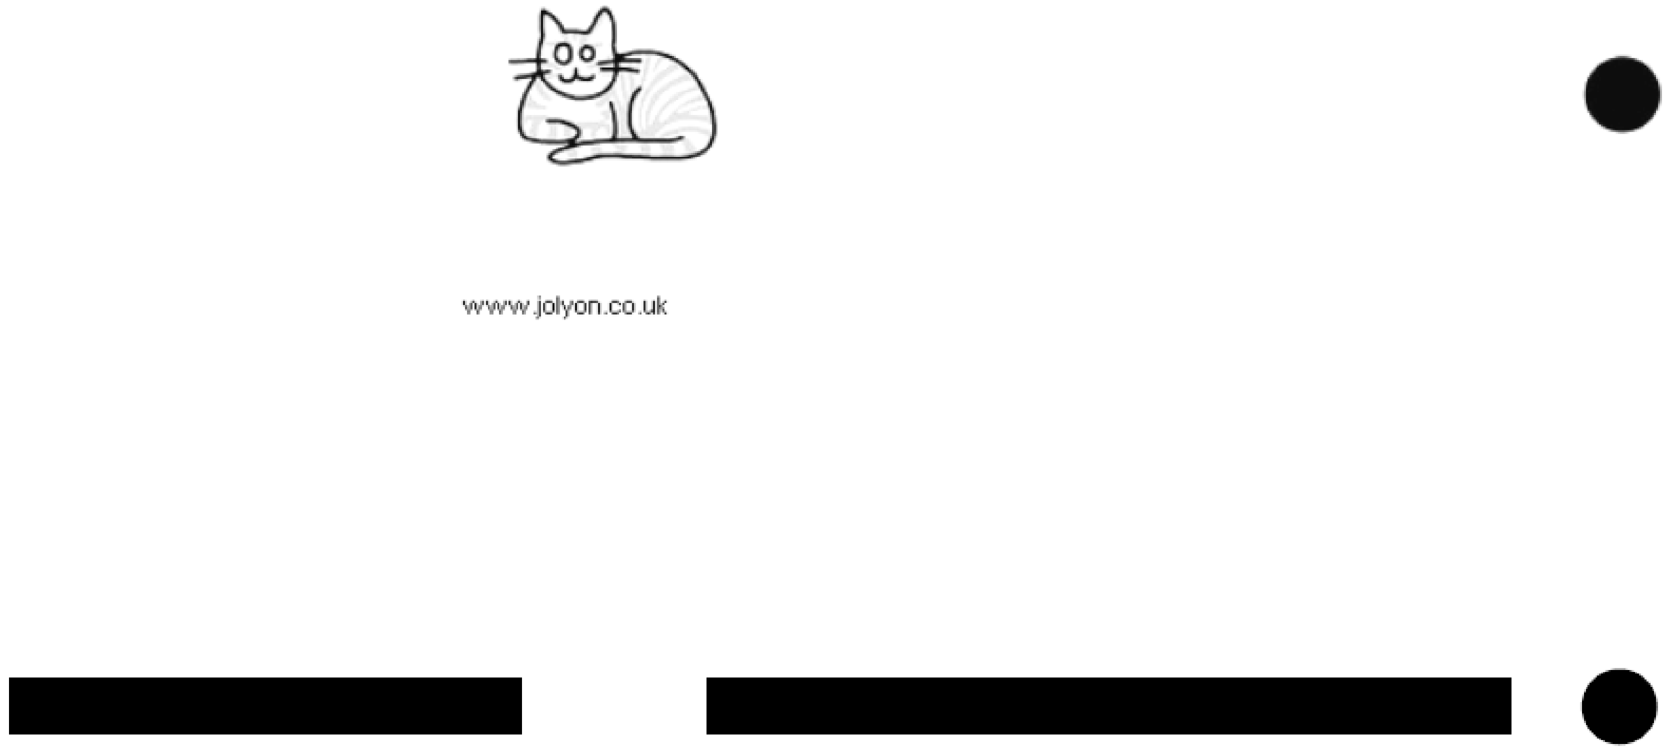

The Physics of Light

Light: Electromagnetic energy whose wavelength is between 400 nm and 700 nm. (1 nm = 10 -9 meter)

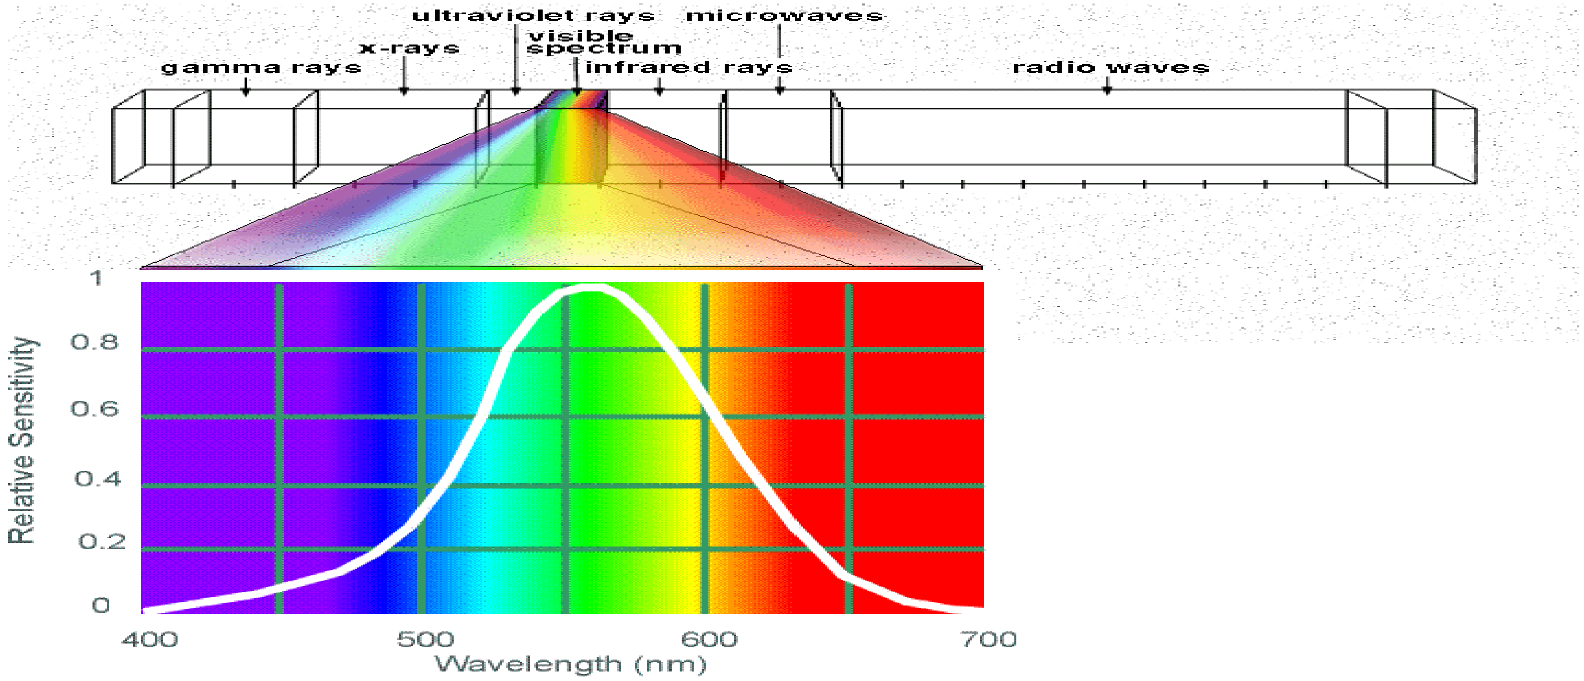

Human Luminance Sensitivity Function

Slide Credit: Efros

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

Visible Light

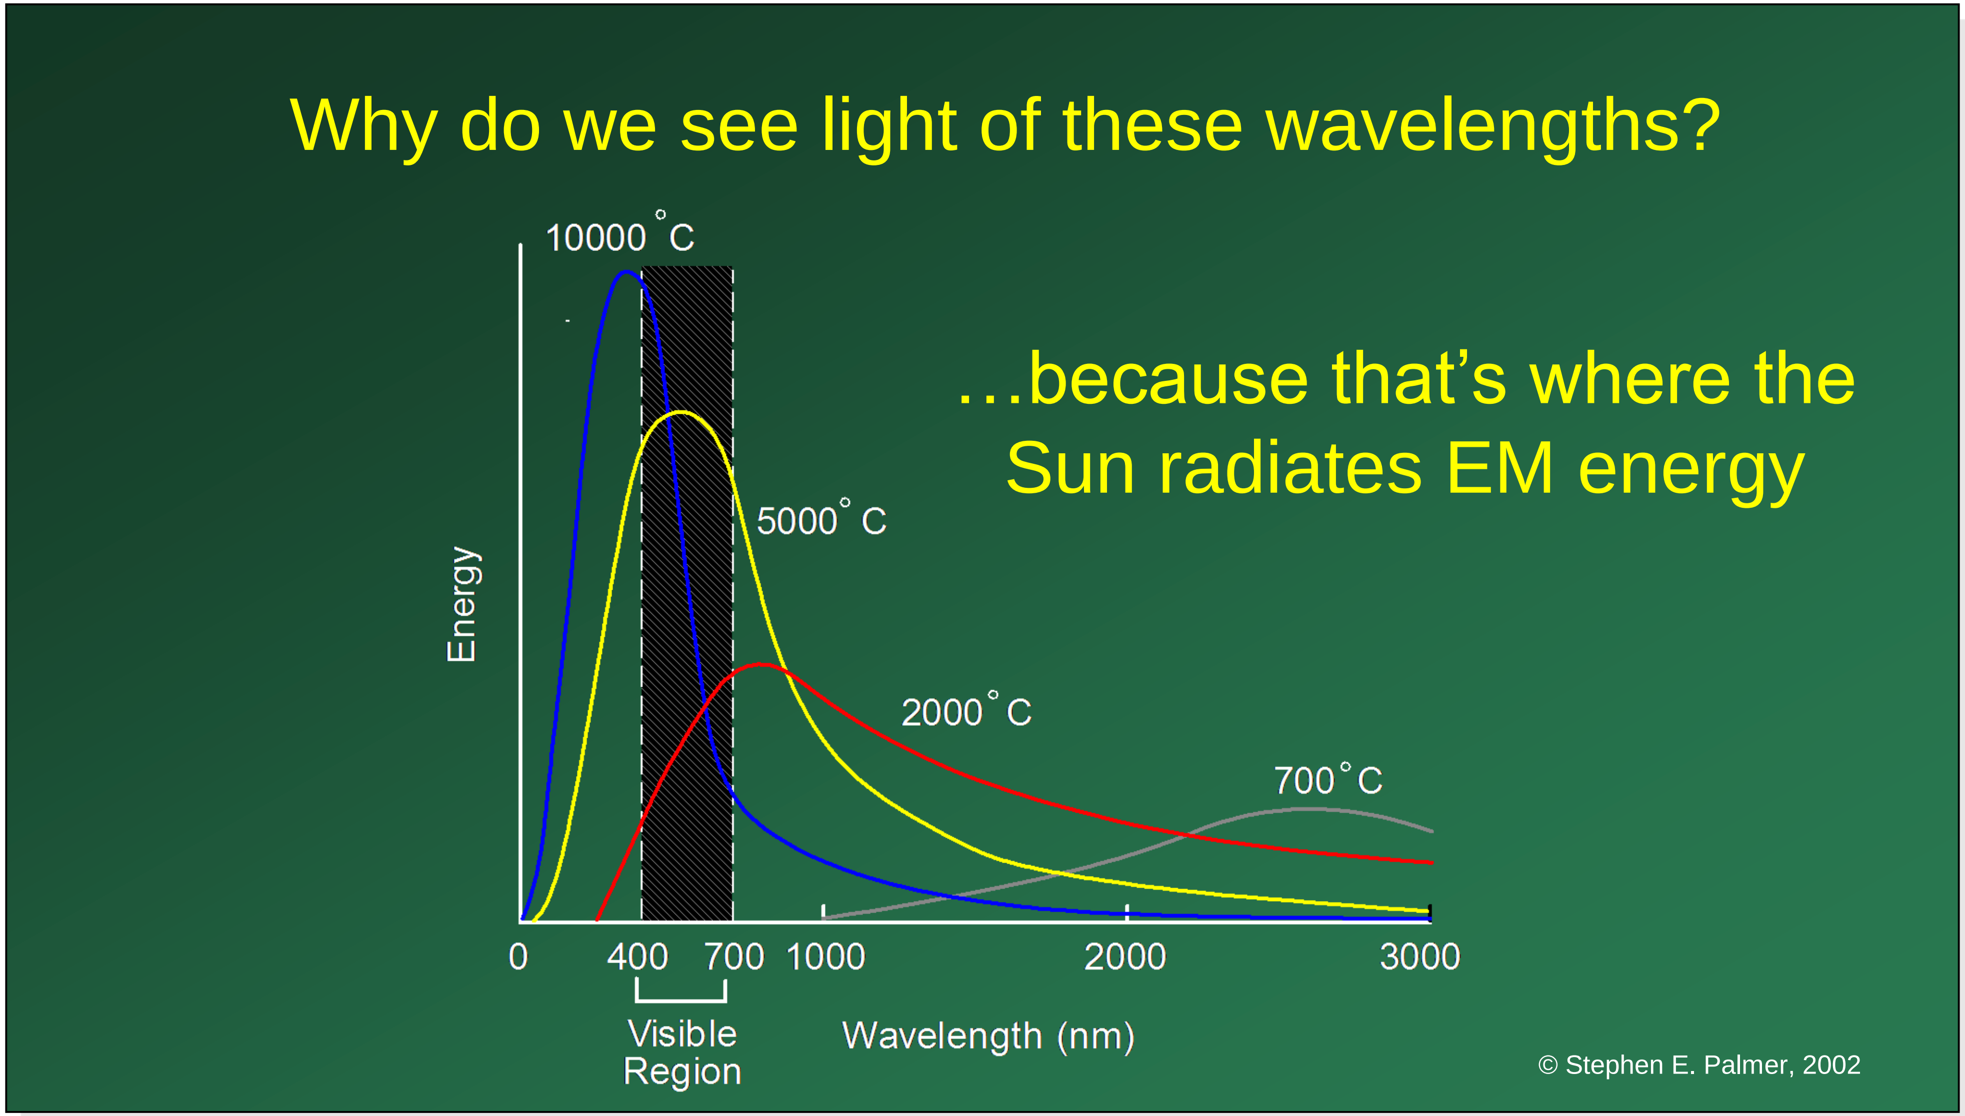

The Physics of Light

Any patch of light can be completely described physically by its spectrum: the number of photons (per time unit) at each wavelength 400 - 700 nm.

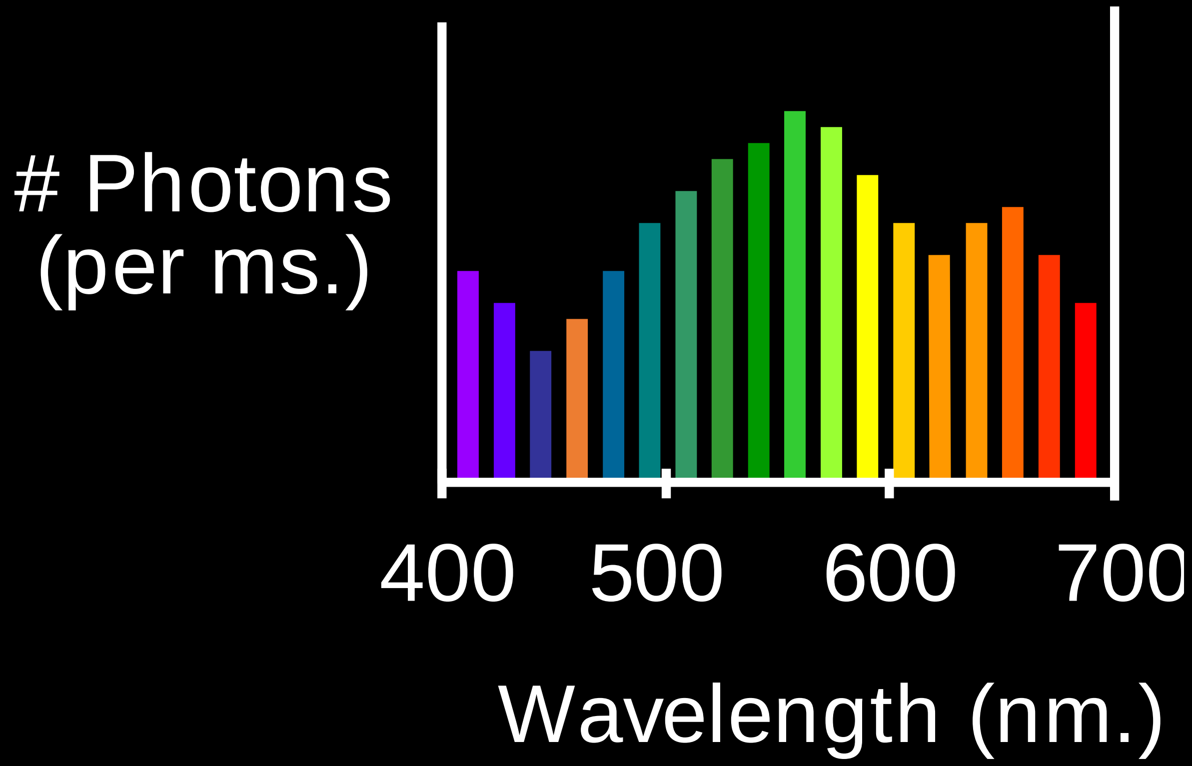

The Physics of Light

Some examples of the spectra of light sources

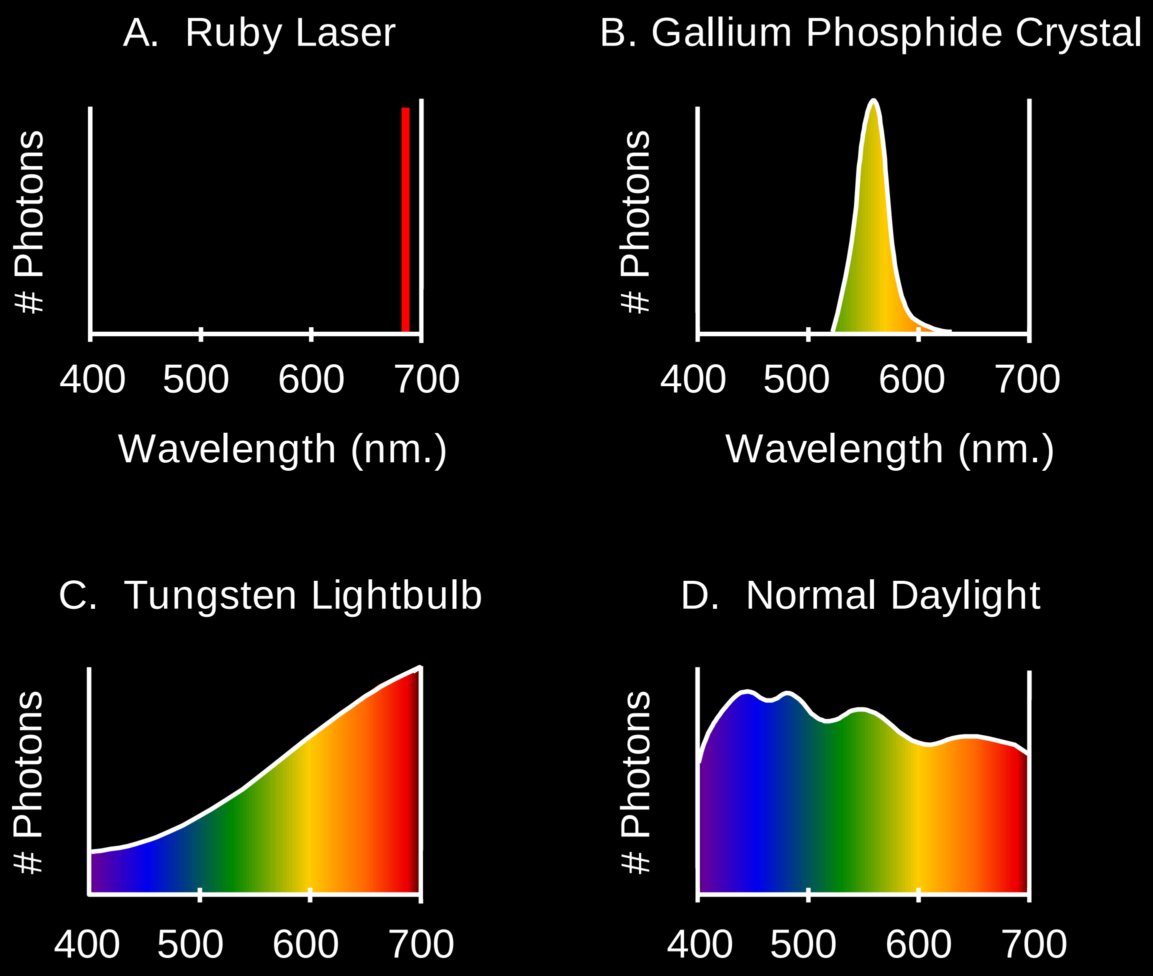

The Physics of Light

Some examples of the reflectance spectra of surfaces

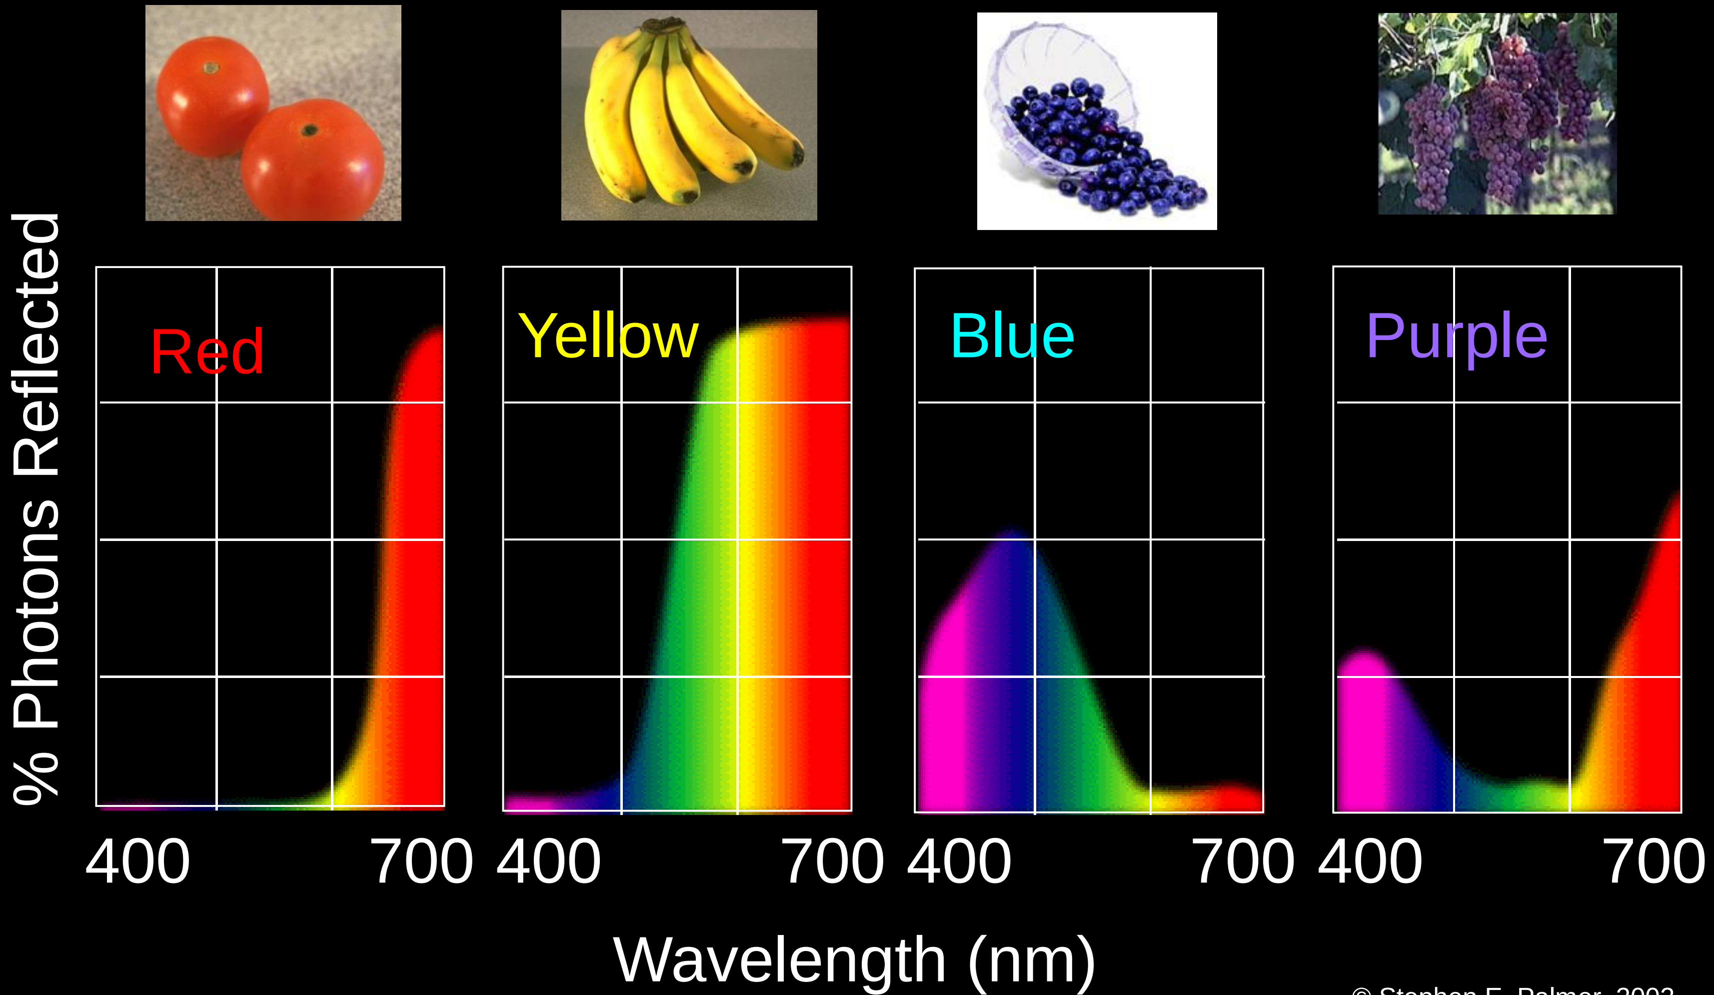

The Psychophysical Correspondence

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

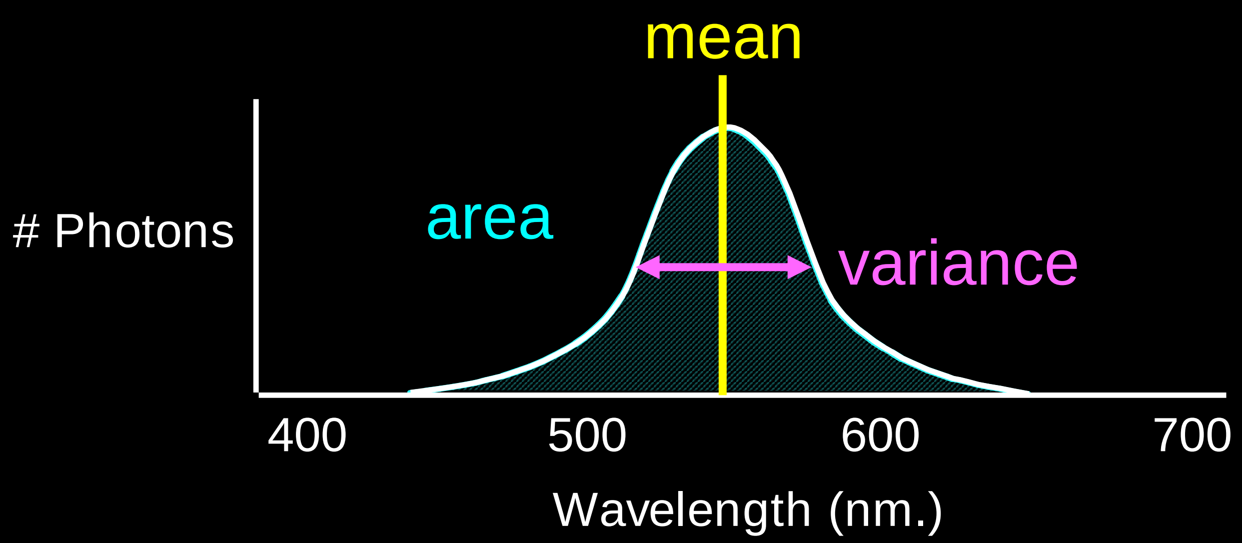

The Psychophysical Correspondence

Mean ⟷ Hue

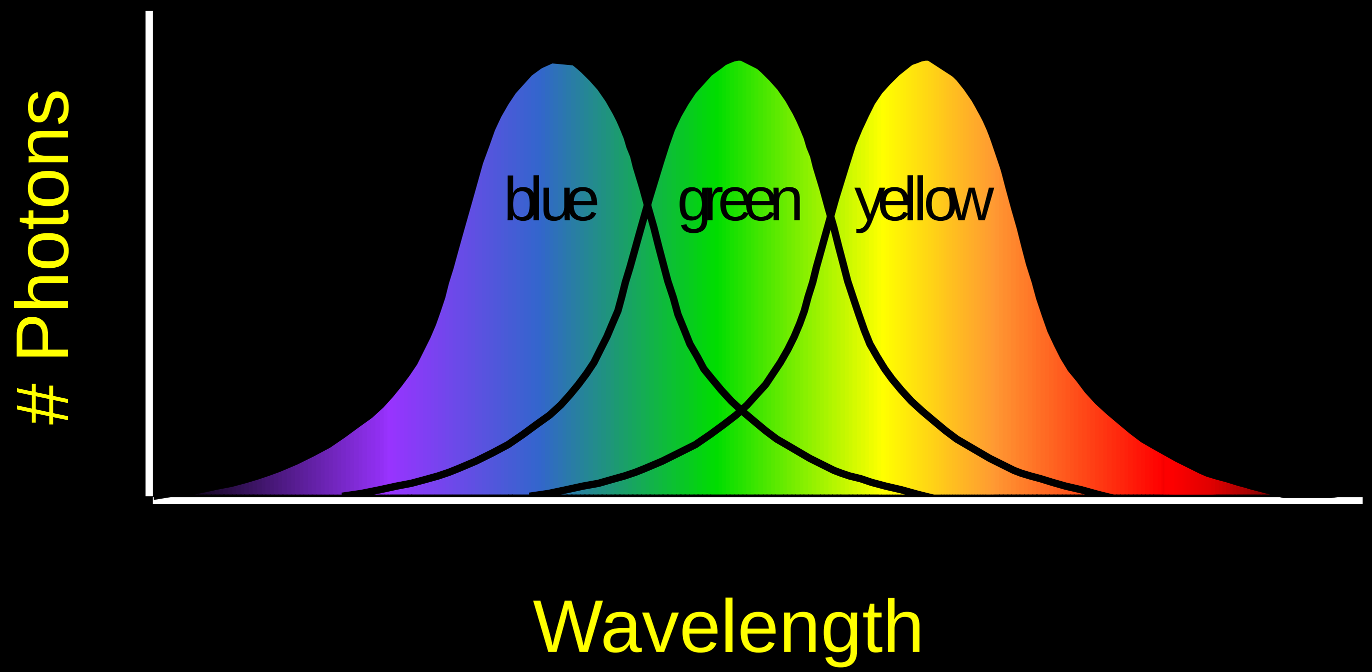

The Psychophysical Correspondence

Variance ⟷ Saturation

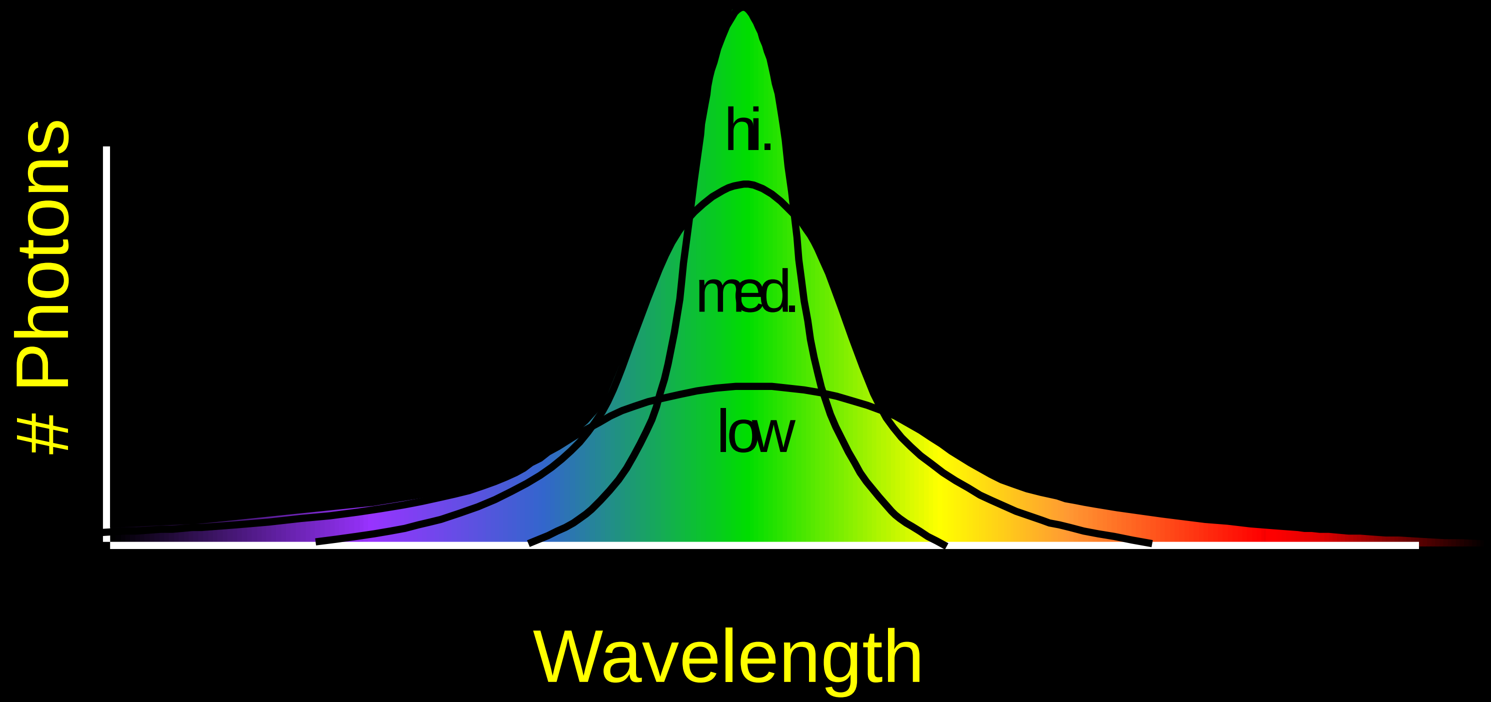

The Psychophysical Correspondence

Area ⟷ Brightness

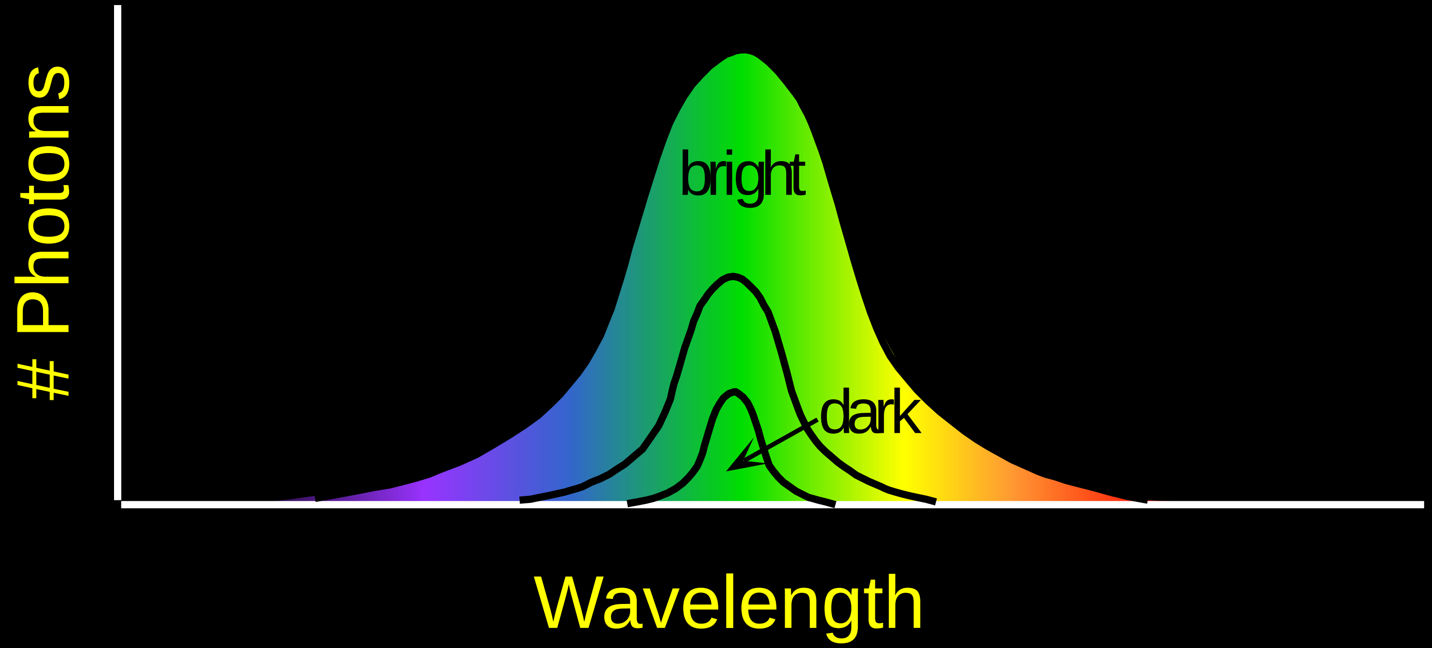

Question: draw a “pink” light

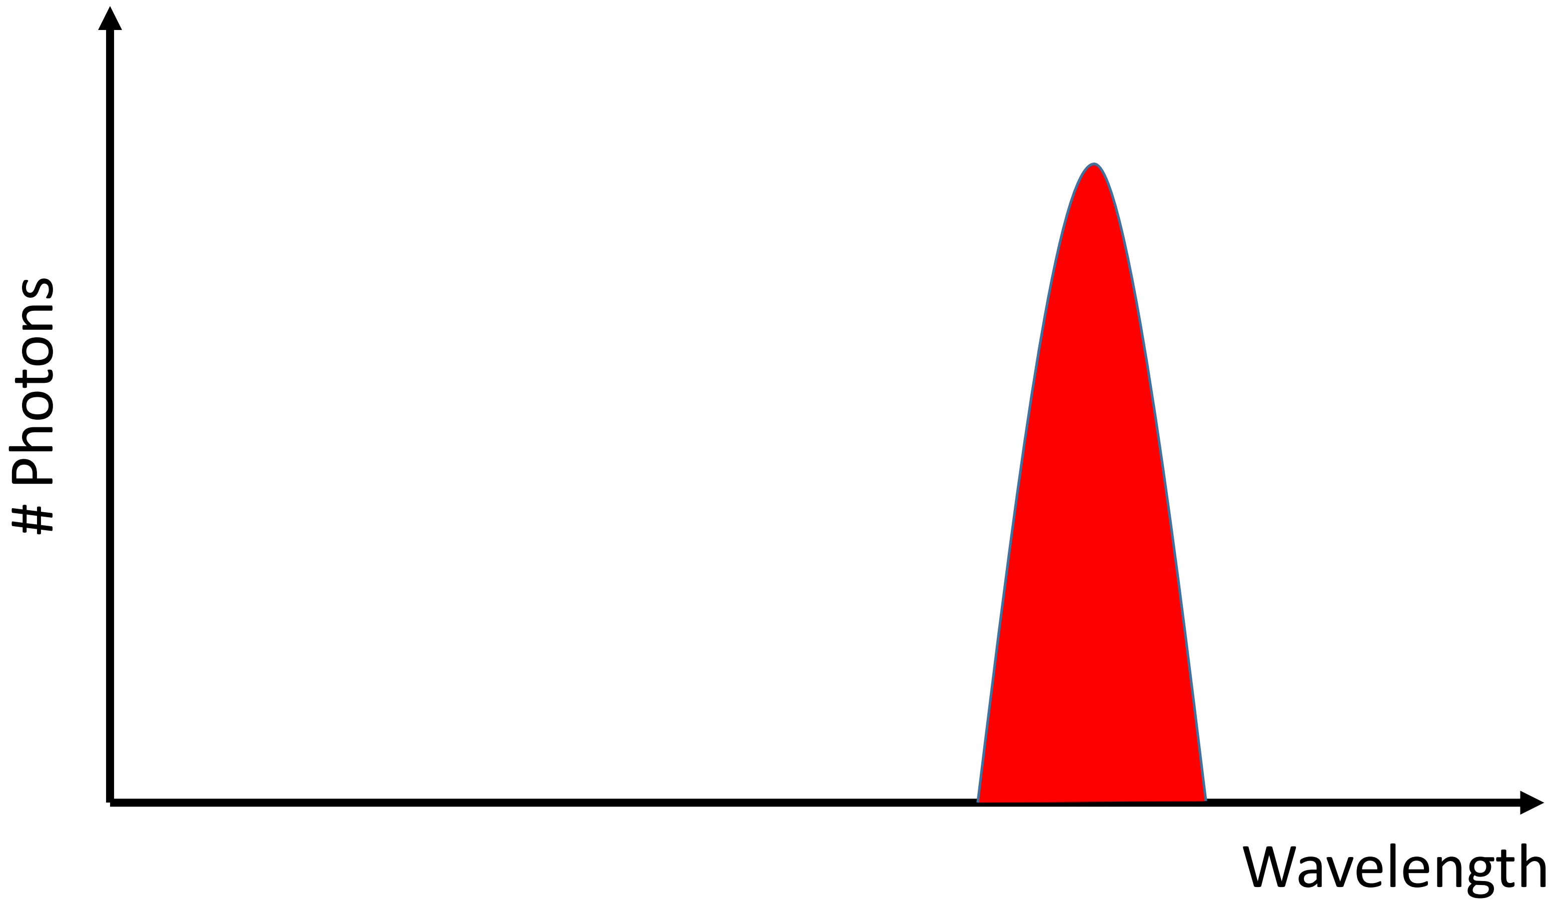

Physiology of Color Vision

Three kinds of cones:

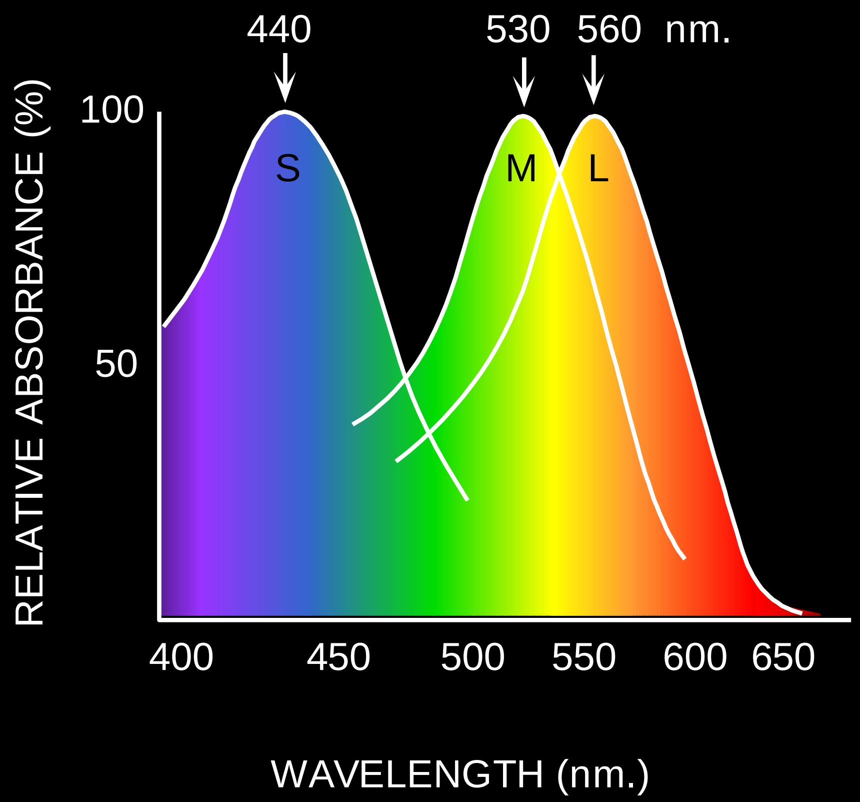

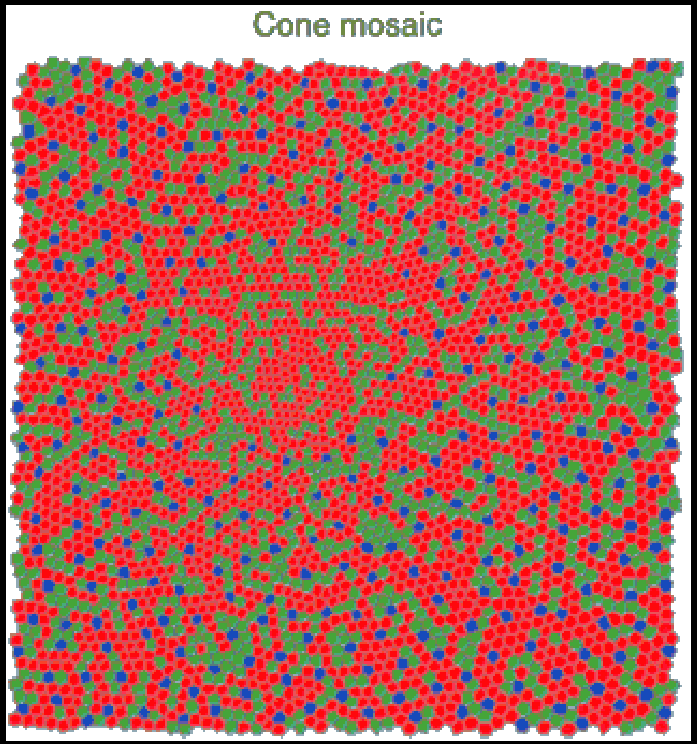

- Why are M and L cones so close?
- Why are there 3?

Trichromacy

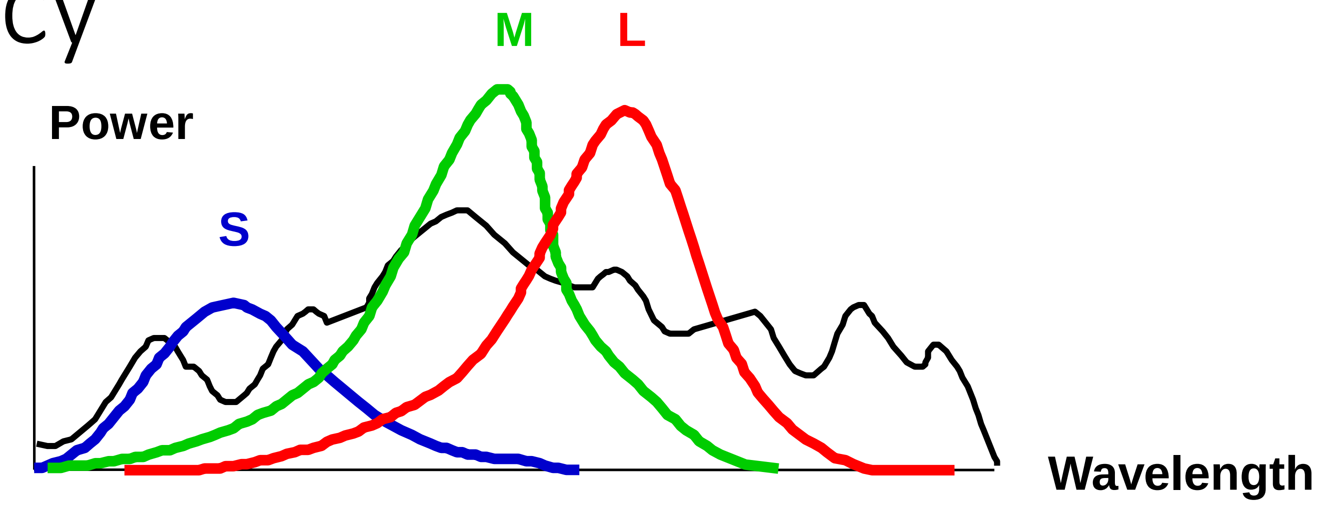

Rods and cones act as *filters* on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number

- How can we represent an entire spectrum with 3 numbers?
- We can't! Most of the information is lost
  - As a result, two different spectra may appear indistinguishable
    - » such spectra are known as **metamers**

Slide by Steve Seitz

Physiology of Color Blindness

Protanopia: Lack of L-cones

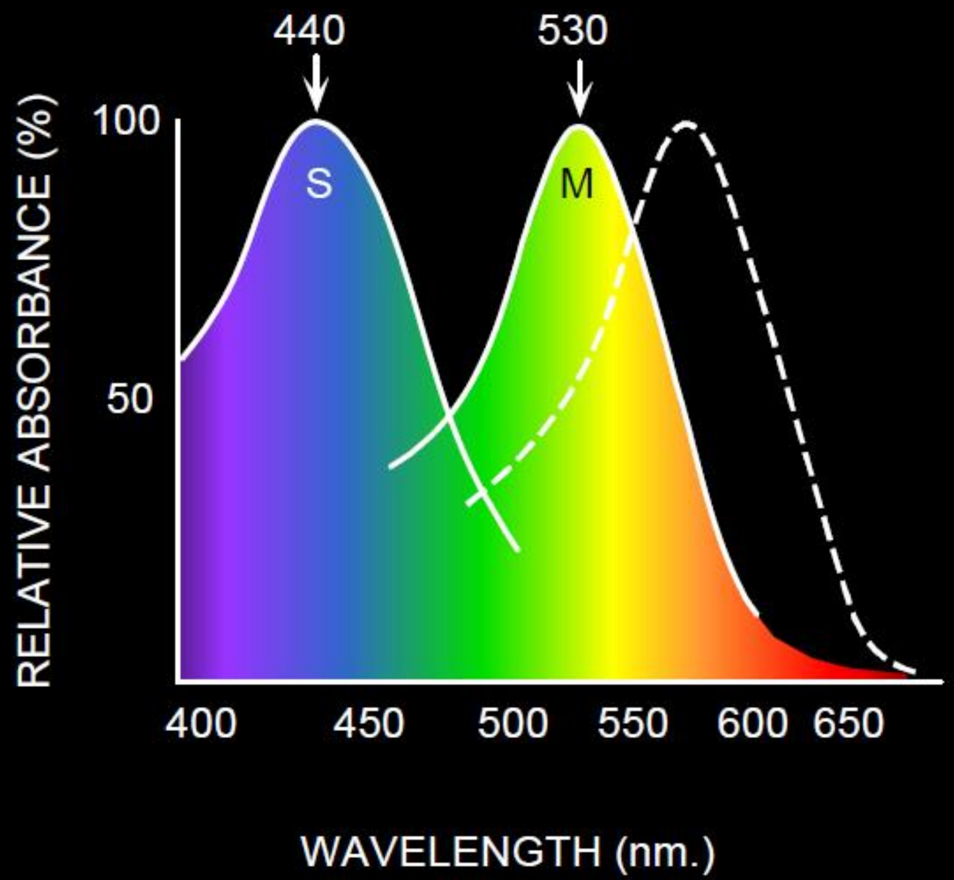

Normal Trichromat

Physiology of Color Blindness

Deuteranopia: Lack of M-cones

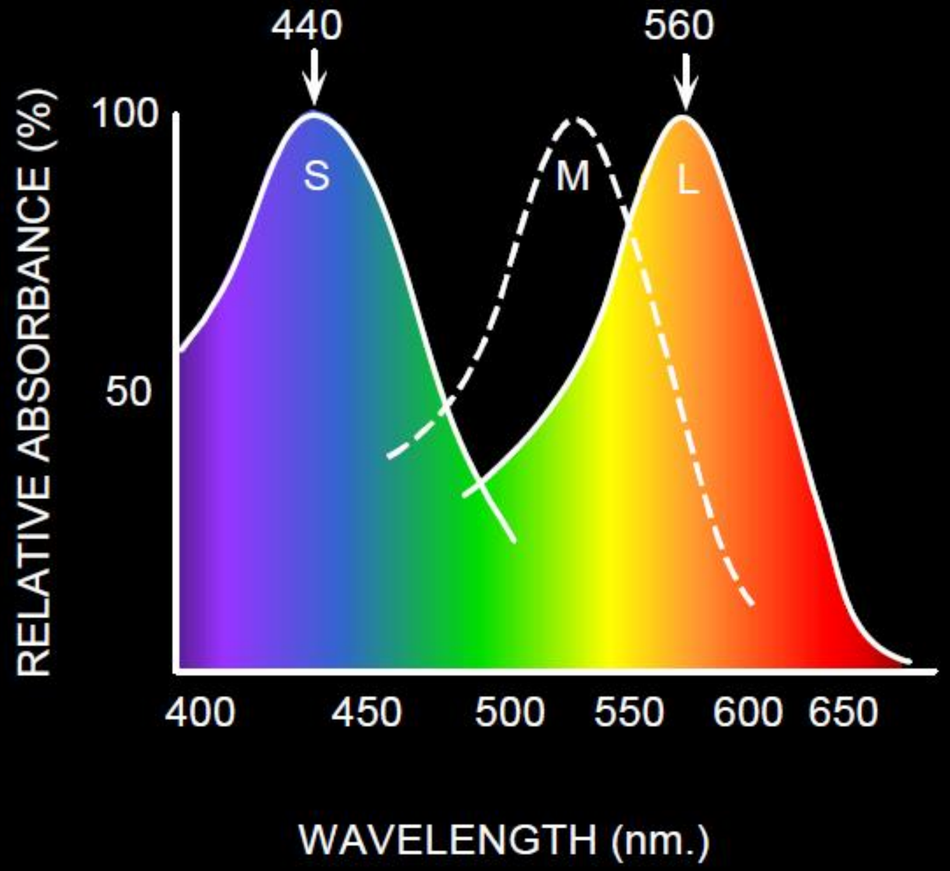

Normal Trichromat

Physiology of Color Blindness

Tritanopia: Lack of S-cones

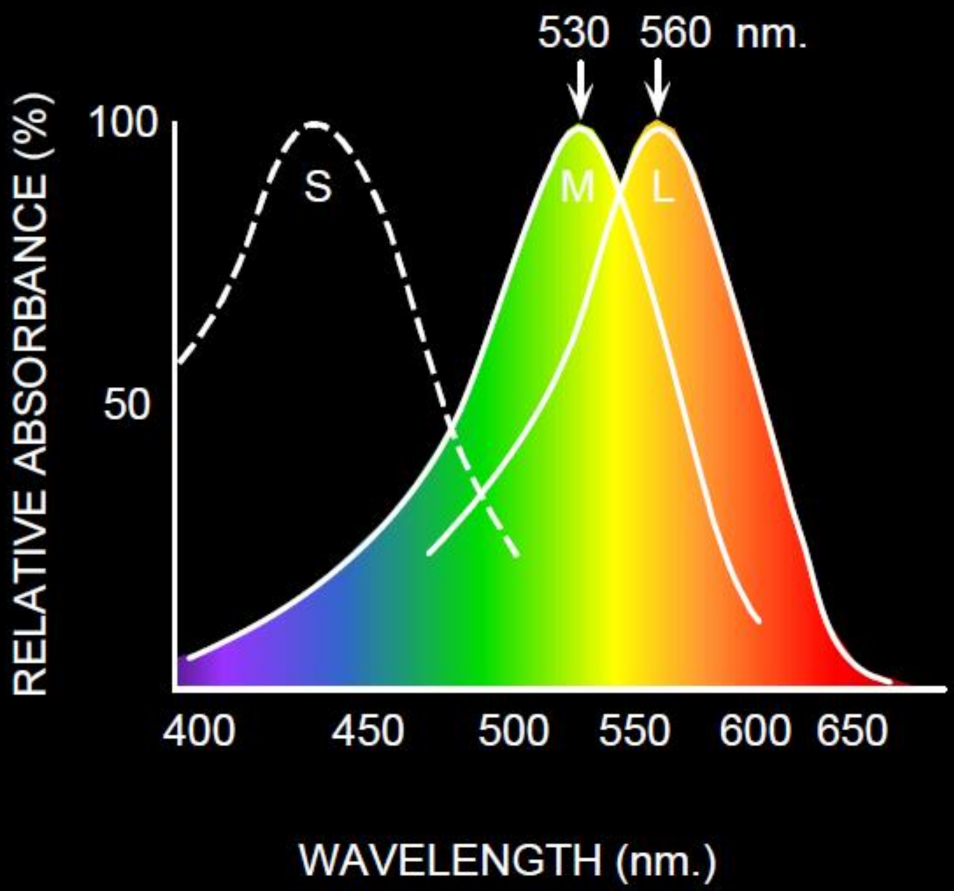

Normal
Trichromat

Correcting Colorblind?

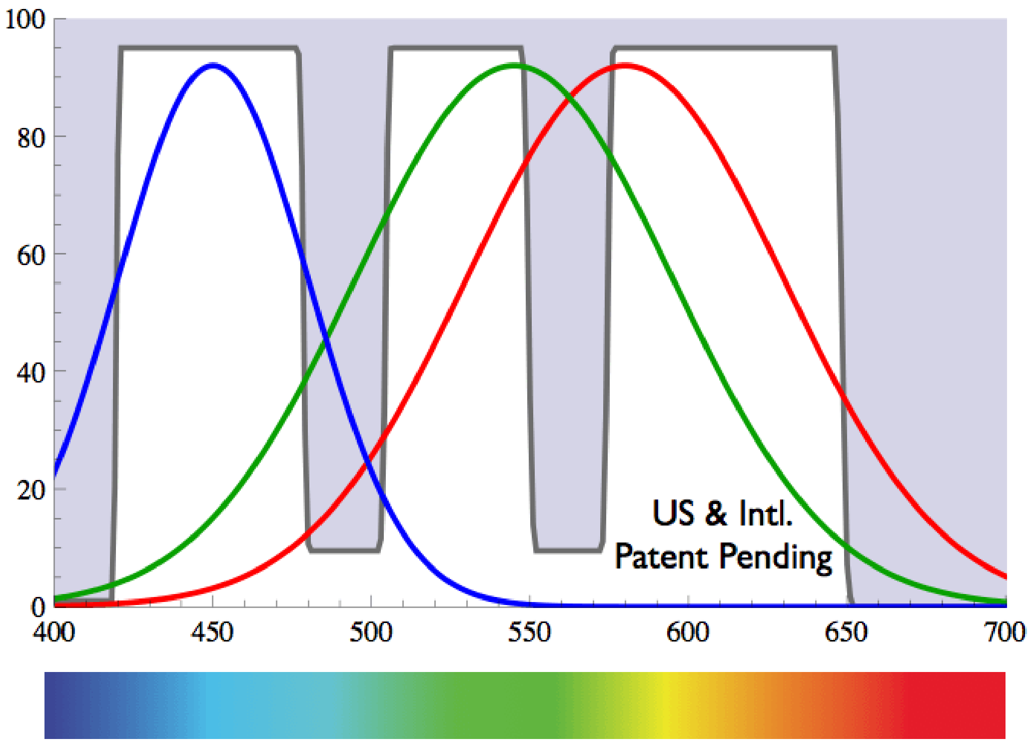

EnChroma Cx

http://enchroma.com/

Color Constancy

<u>The "photometer metaphor" of color perception</u>:
Color perception is determined by the spectrum of light on each retinal receptor (as measured by a photometer).

Color Constancy

The "photometer metaphor" of color perception:
Color perception is determined by the spectrum of light on each retinal receptor (as measured by a photometer).

Color Constancy

~~The "photometer metaphor" of color perception:~~
~~Color perception is determined by the spectrum of light on each retinal receptor (as measured by a photometer).~~

Color Constancy

~~Do we have constancy over~~
~~all global color transformations?~~

60% blue filter

Complete inversion

Color Constancy

Color Constancy: the ability to perceive the invariant color of a surface despite ecological Variations in the conditions of observation.

Another of these hard inverse problems:

Physics of light emission and surface reflection underdetermine perception of surface color

Practical Color Sensing: Bayer Grid

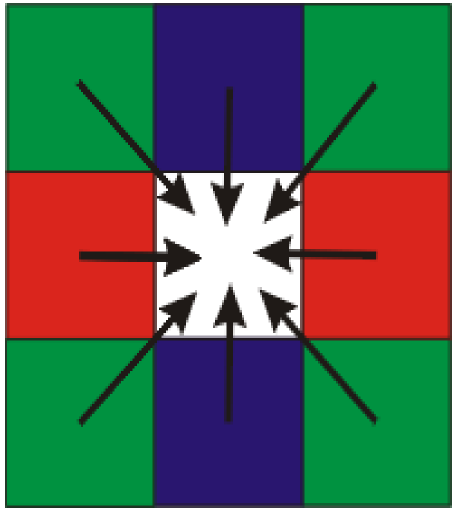

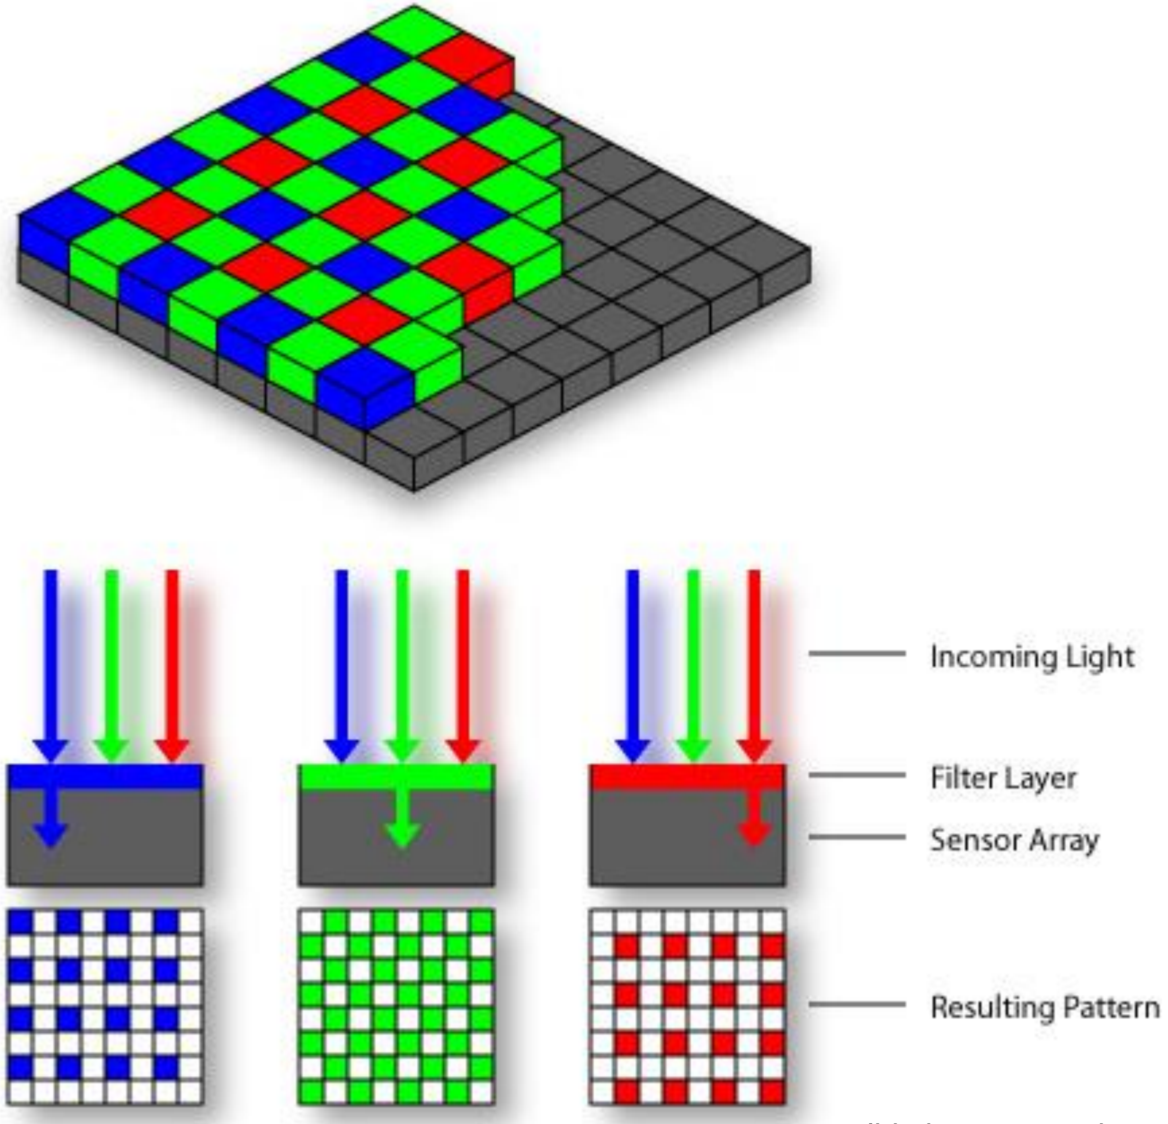

- Estimate RGB at 'G' cels from neighboring values

http://www.cooldictionary.com/words/Bayer-filter.wikipedia

Slide by Steve Seitz

Color Image

Images in Matlab

- Images represented as a matrix
- Suppose we have a NxM RGB image called “im”
  - im(1,1,1) = top-left pixel value in R-channel
  - im(y, x, b) = y pixels down, x pixels to right in the bth channel
  - im(N, M, 3) = bottom-right pixel in B-channel
- imread(filename) returns a uint8 image (values 0 to 255)
  - Convert to double format (values 0 to 1) with im2double

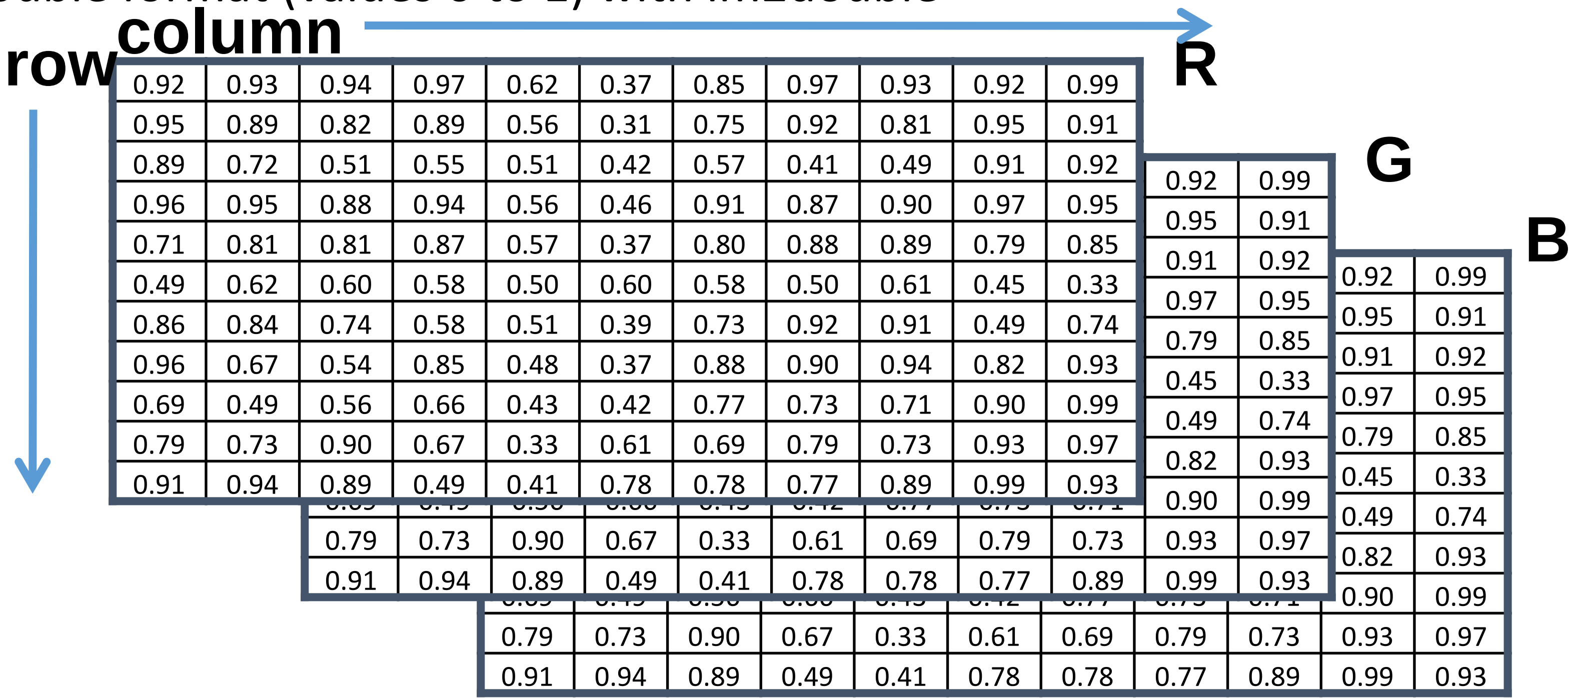

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Color spaces

- How can we represent color?

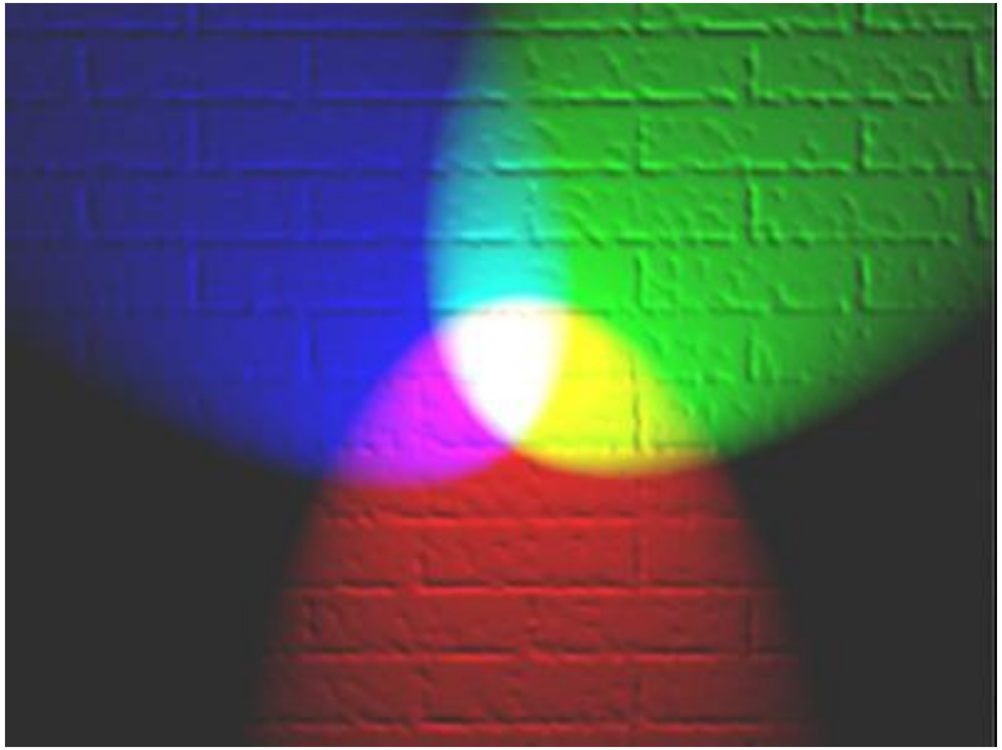

http://en.wikipedia.org/wiki/File:RGB_illumination.jpg

Color spaces: RGB

Default color space

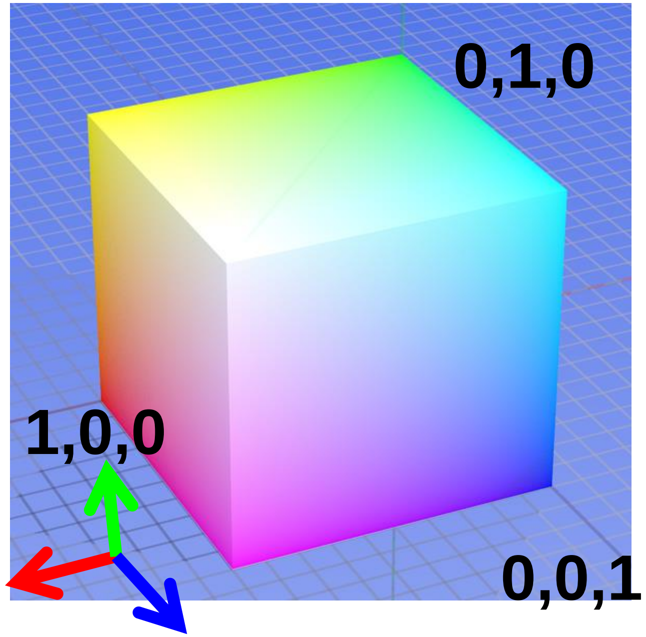

RGB cube

- Easy for devices
- But not perceptual
- Where do the grays live?
- Where is hue and saturation?

R
(G=0,B=0)

G
(R=0,B=0)

B
(R=0,G=0)

Image from: http://en.wikipedia.org/wiki/File:RGB_color_solid_cube.png

HSV

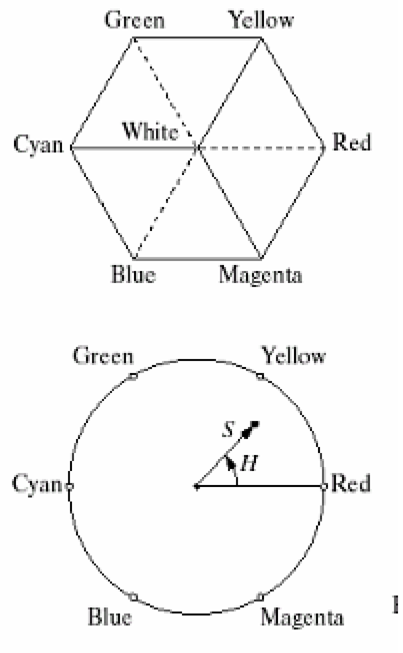

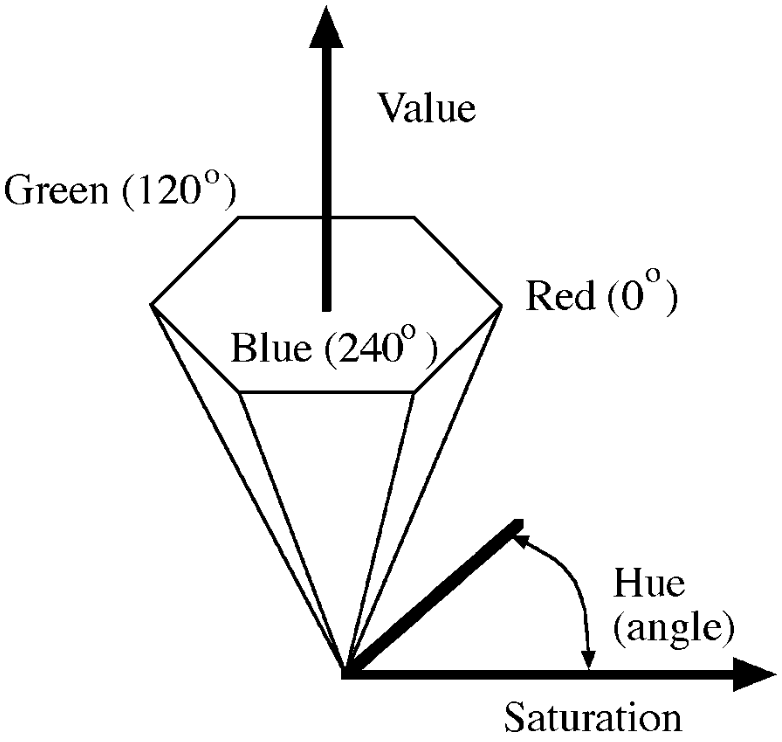

- Hue, Saturation, Value (Intensity)
  - RGB cube on its vertex
- Decouples the three components (a bit)
- Use rgb2hsv() and hsv2rgb() in Matlab

Slide by Steve Seitz

Color spaces: HSV

Intuitive color space

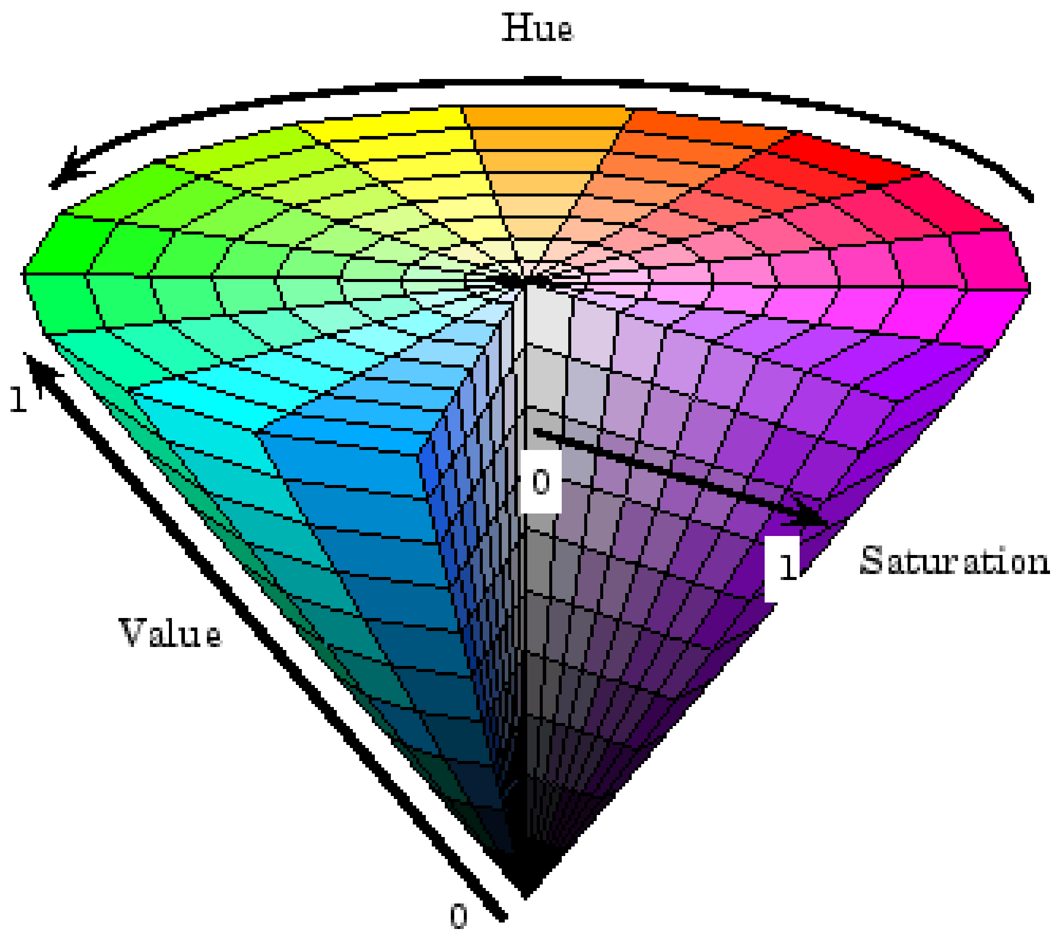

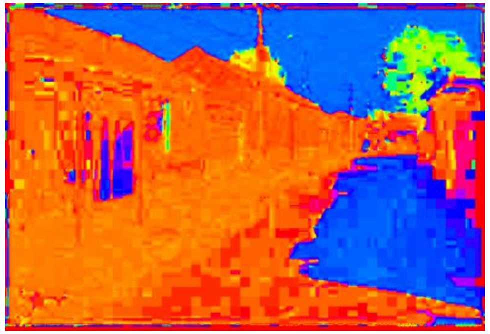

H
(S=1,V=1)

S
(H=1,V=1)

V
(H=1,S=0)

Color spaces: L*a*b*

“Perceptually uniform” color space

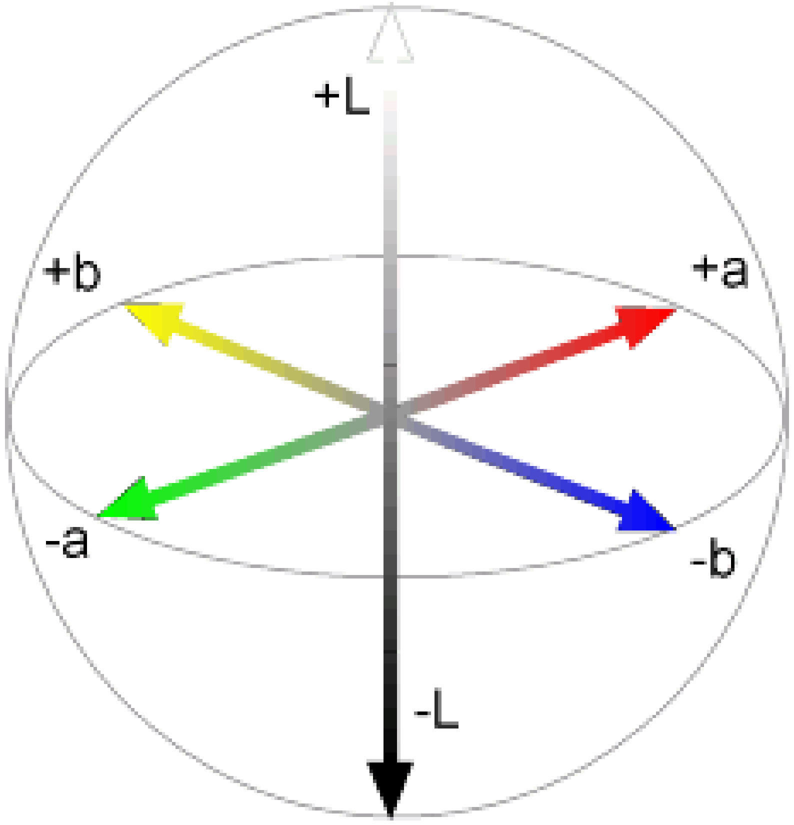

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

So far: light→surface→camera

- Called a local illumination model
- But much light comes from surrounding surfaces

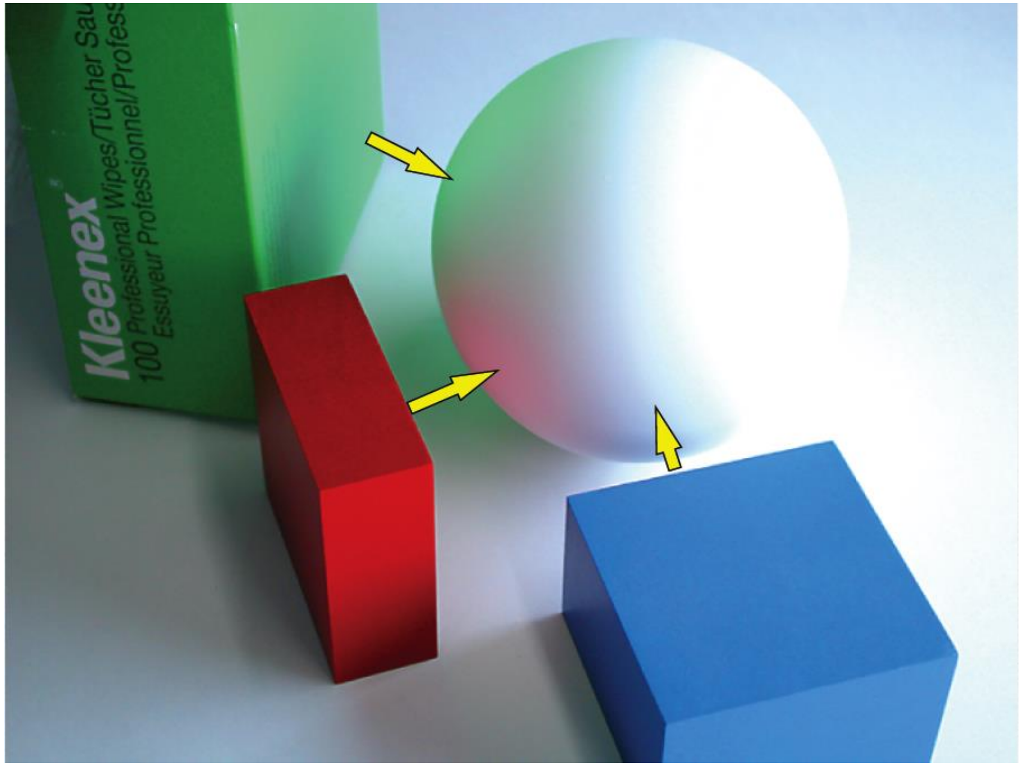

From Koenderink slides on image texture and the flow of light

Inter-reflection is a major source of light

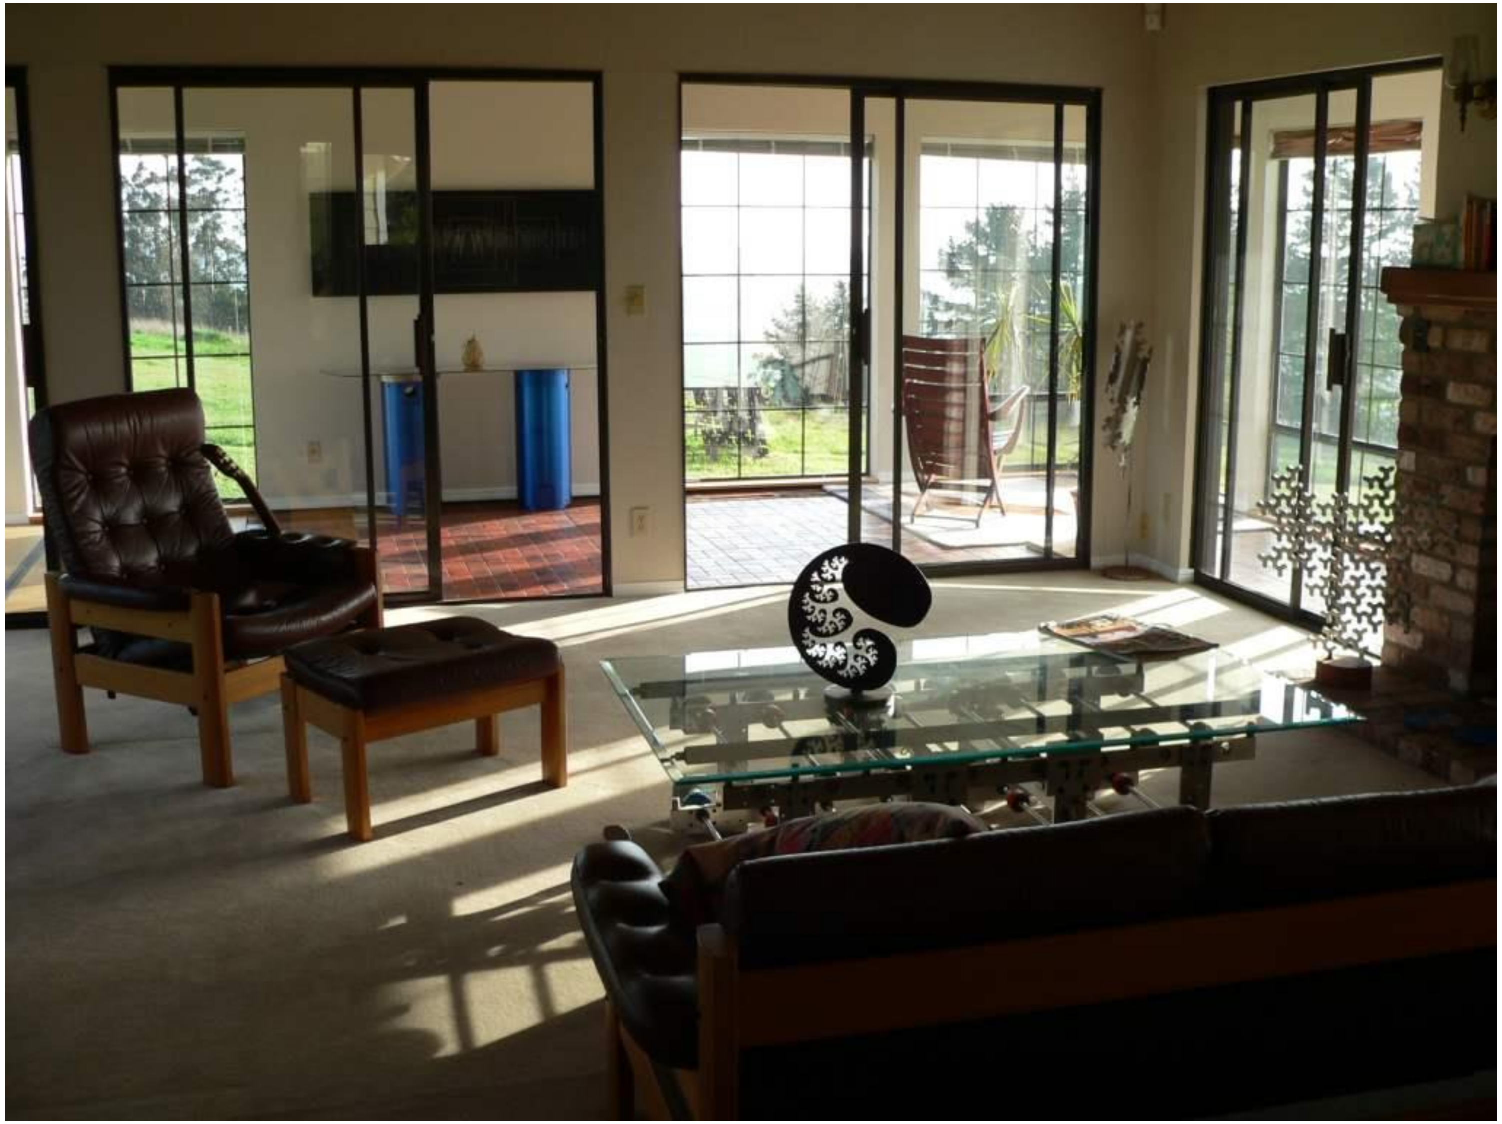

Inter-reflection affects the apparent color of objects

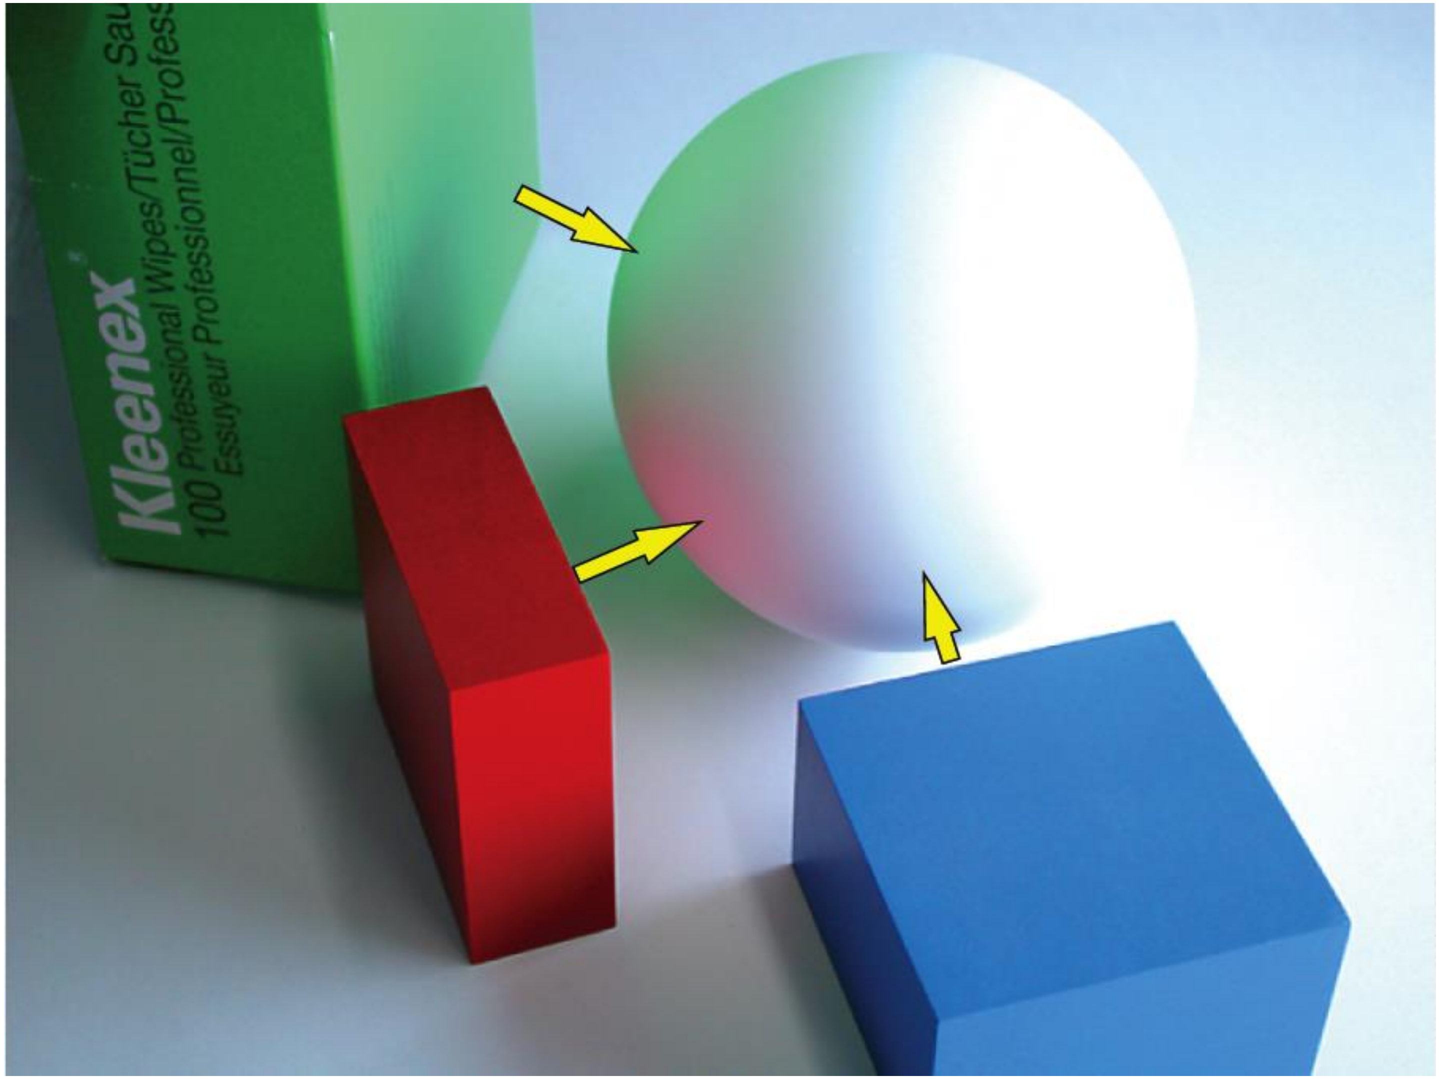

From Koenderink slides on image texture and the flow of light

Scene surfaces also cause shadows

- Shadow: reduction in intensity due to a blocked source

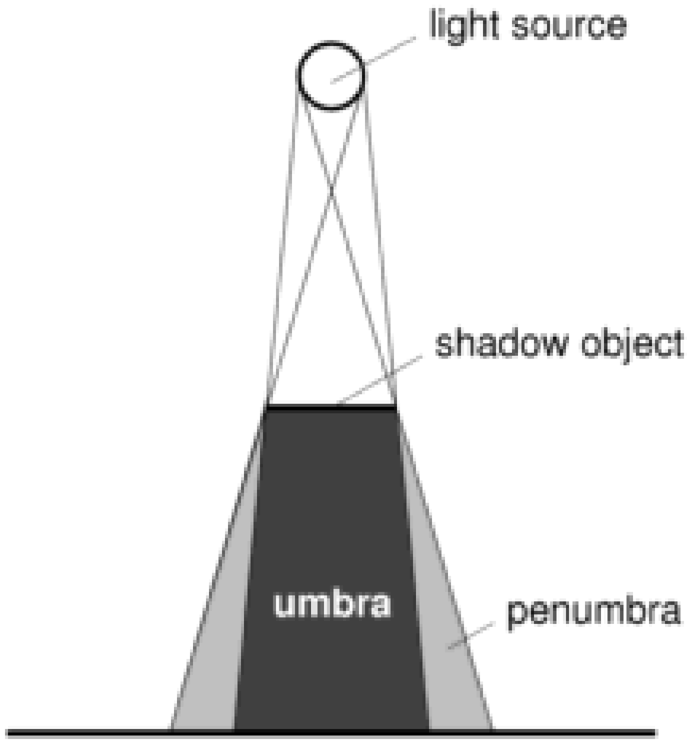

Shadows

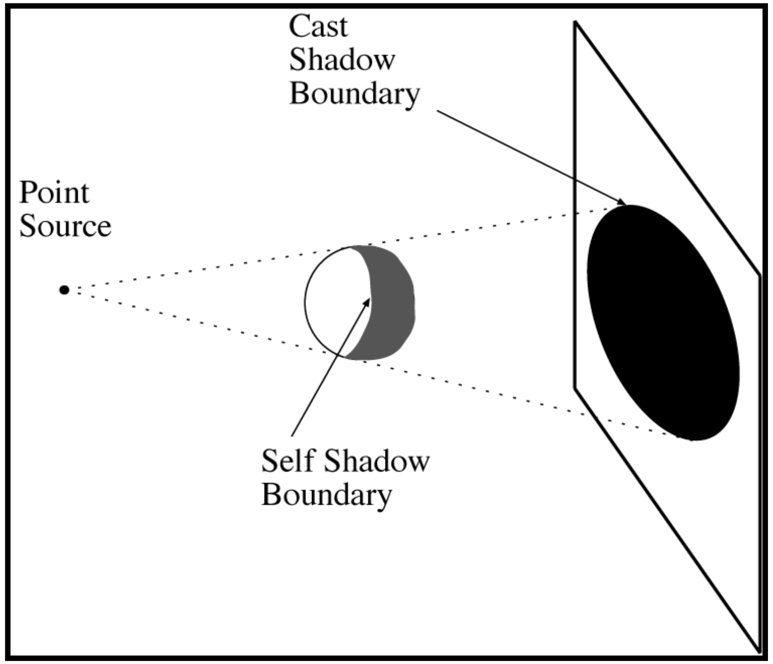
Cast
Shadow
Boundary
Point
Source
Self Shadow
Boundary

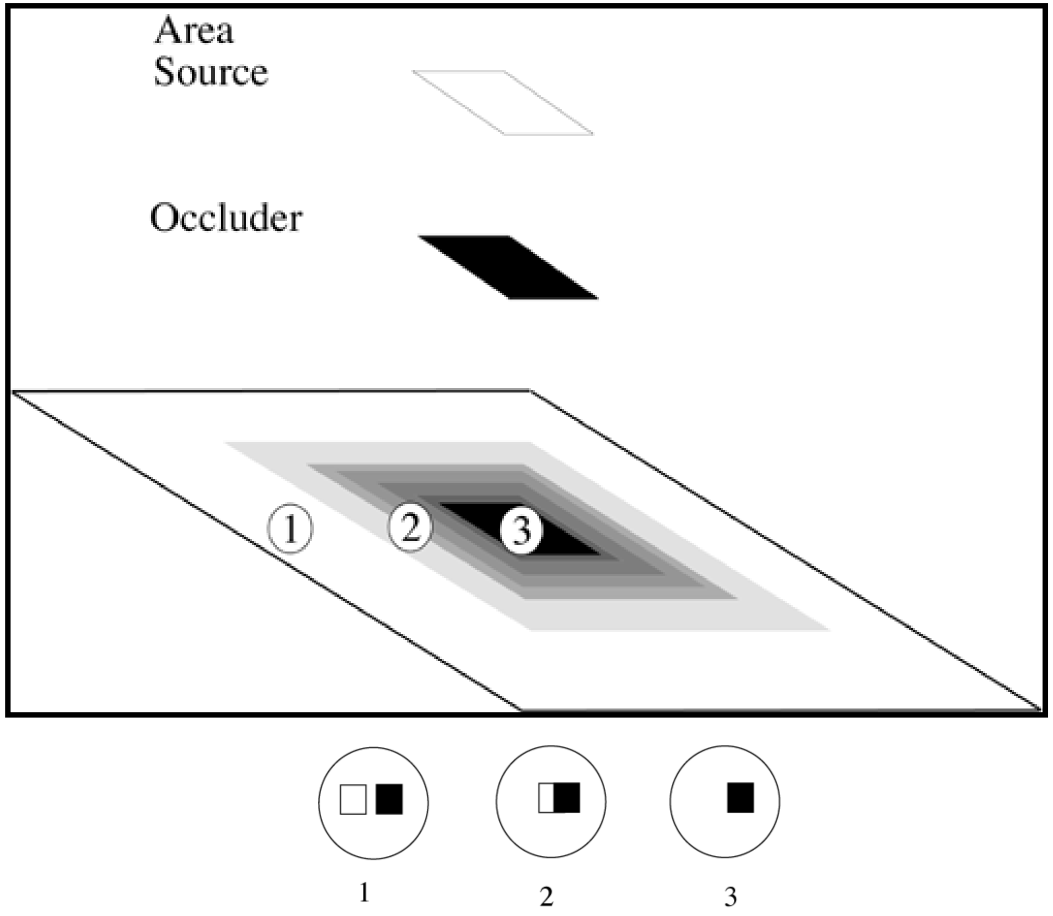
Area
Source
Occluder
1
2
3
1
2
3

Models of light sources

- Distant point source
  - One illumination direction
  - E.g., sun
- Area source
  - E.g., white walls, diffuser lamps, sky
- Ambient light
  - Substitute for dealing with interreflections
- Global illumination model
  - Account for interreflections in modeled scene

Questions

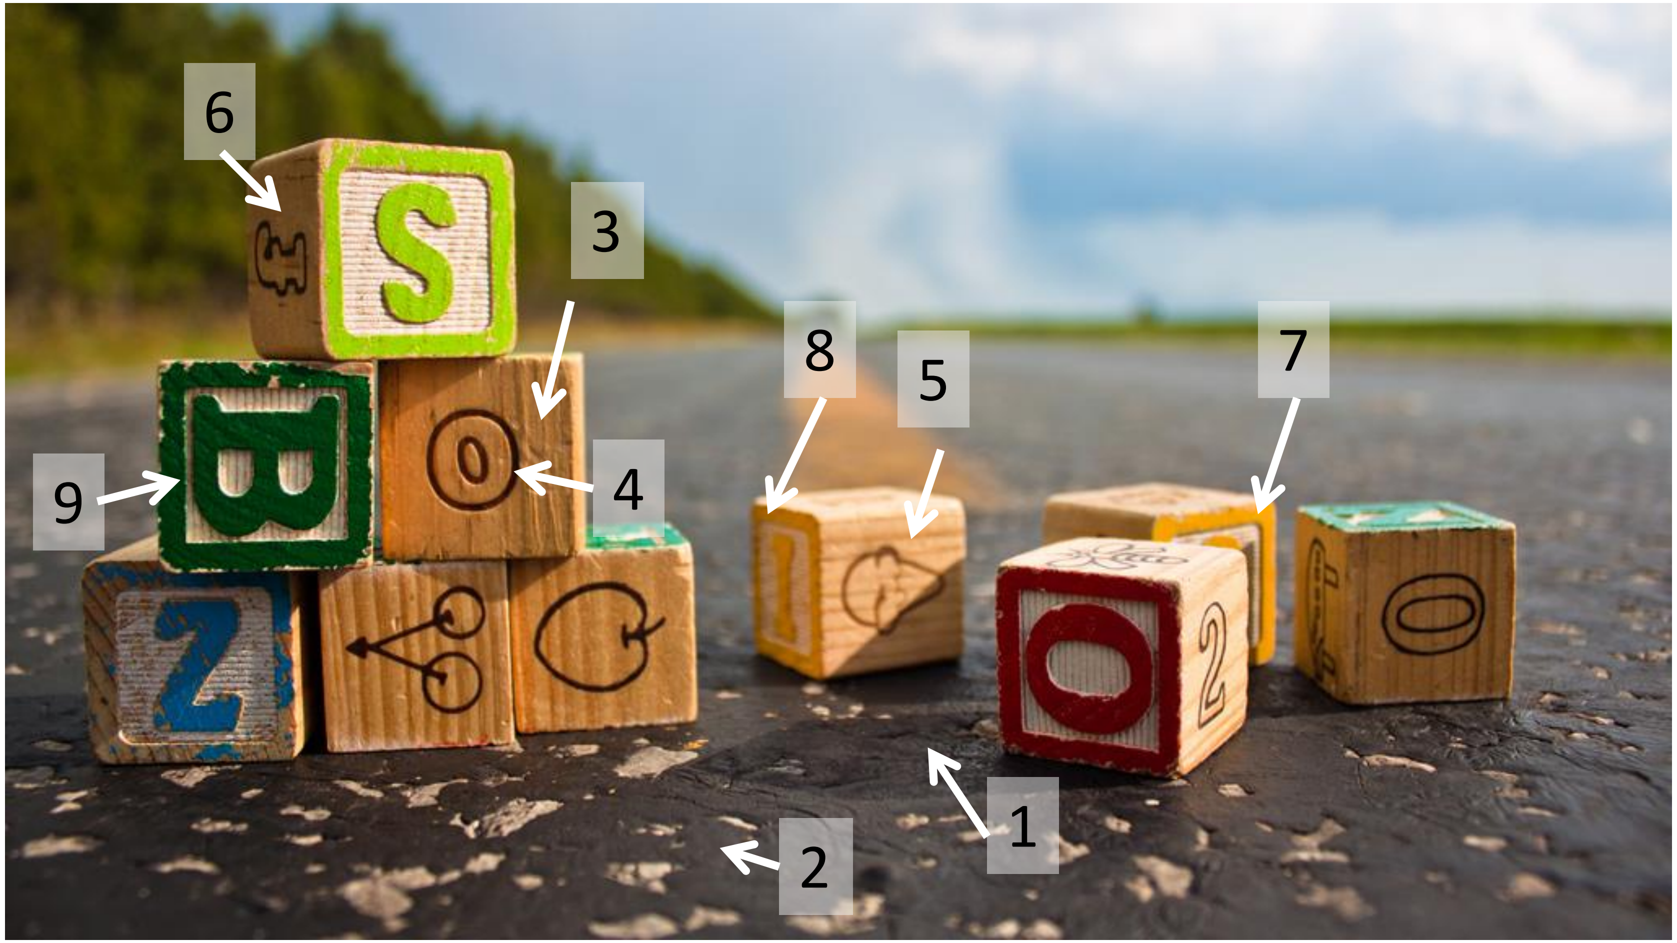

A. Why is (2) brighter than (1)? Each points to the asphalt.
B. Why is (4) darker than (3)? (4) points to the marking.
C. Why is (5) brighter than (3)? Each points to the side of the wooden block.
D. Why isn't (6) black, given that there is no direct path from it to the sun?
E. Why (7) brighter than (8)? Both point to the yellow paints.
F. Why is (9) green, given that the sun light contains all visible wavelengths?

What does the intensity of a pixel tell us?

im(234, 452) = 0.58

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

The plight of the poor pixel

- A pixel’s brightness is determined by
  - Light source (strength, direction, color)
  - Surface orientation
  - Surface material and albedo
  - Reflected light and shadows from surrounding surfaces
  - Gain on the sensor

- A pixel’s brightness tells us nothing by itself

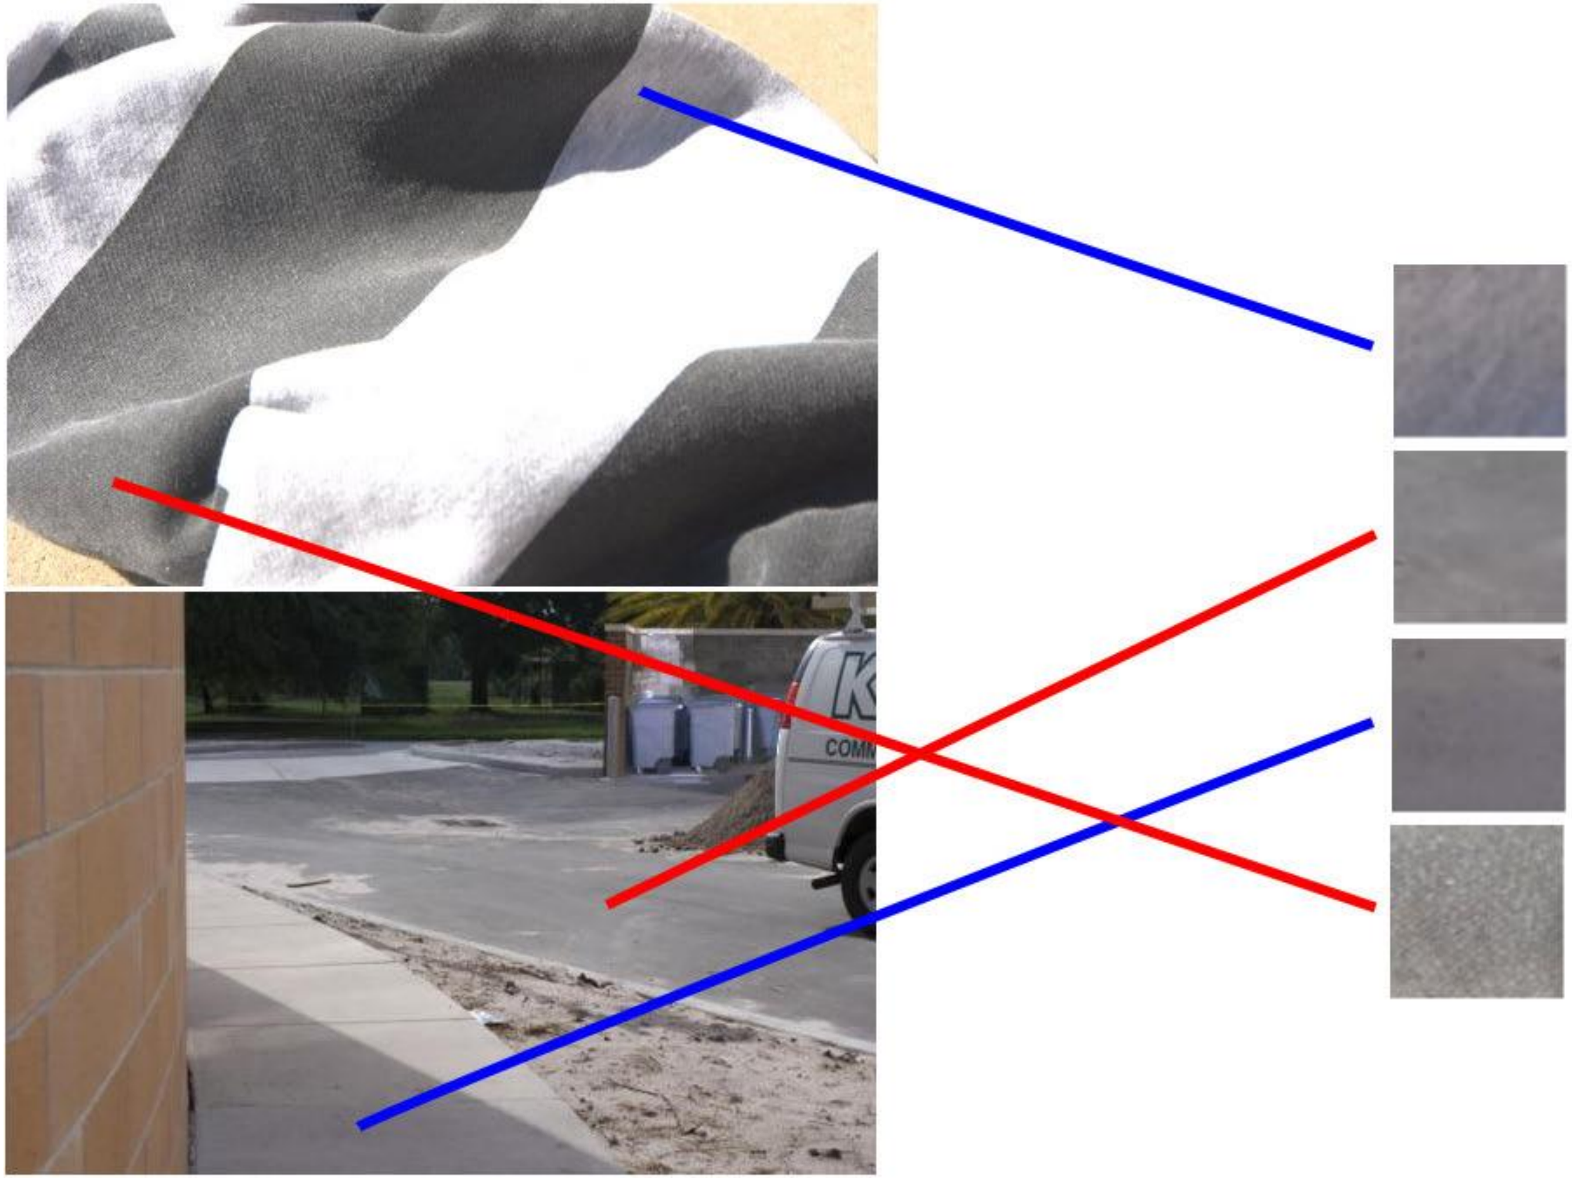

And yet we can interpret images…

- Key idea: for nearby scene points, most factors do not change much
- The information is mainly contained in *local differences* of brightness

Darkness = Large Difference in Neighboring Pixels

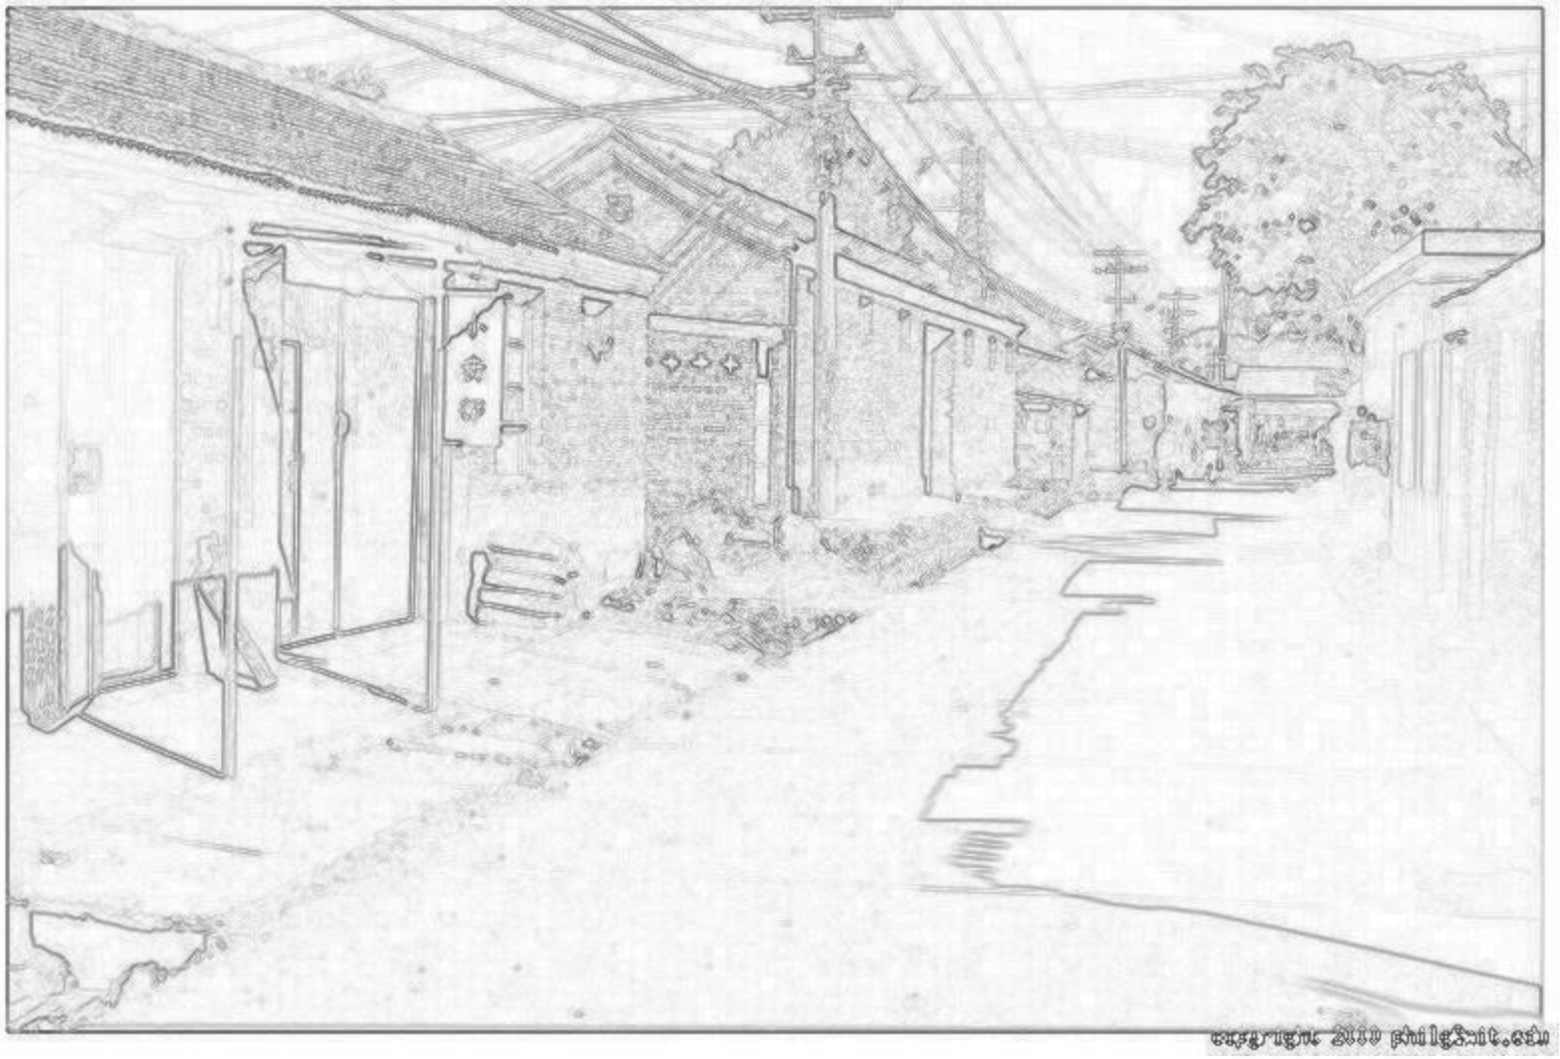

What is this?


HONDA

What differences in intensity tell us about shape?

- Changes in surface normal
- Texture
- Proximity
- Indents and bumps
- Grooves and creases

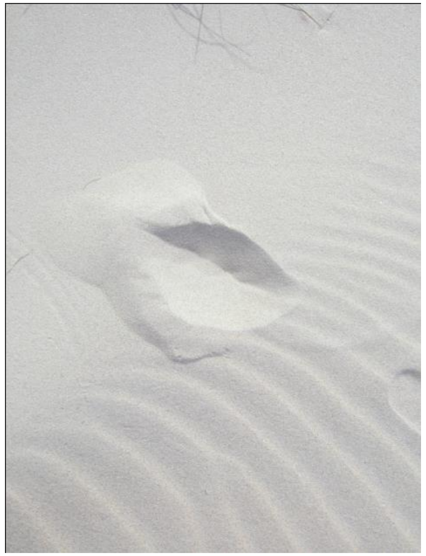

Photos Koenderink slides on image texture and the flow of light

Shadows as cues

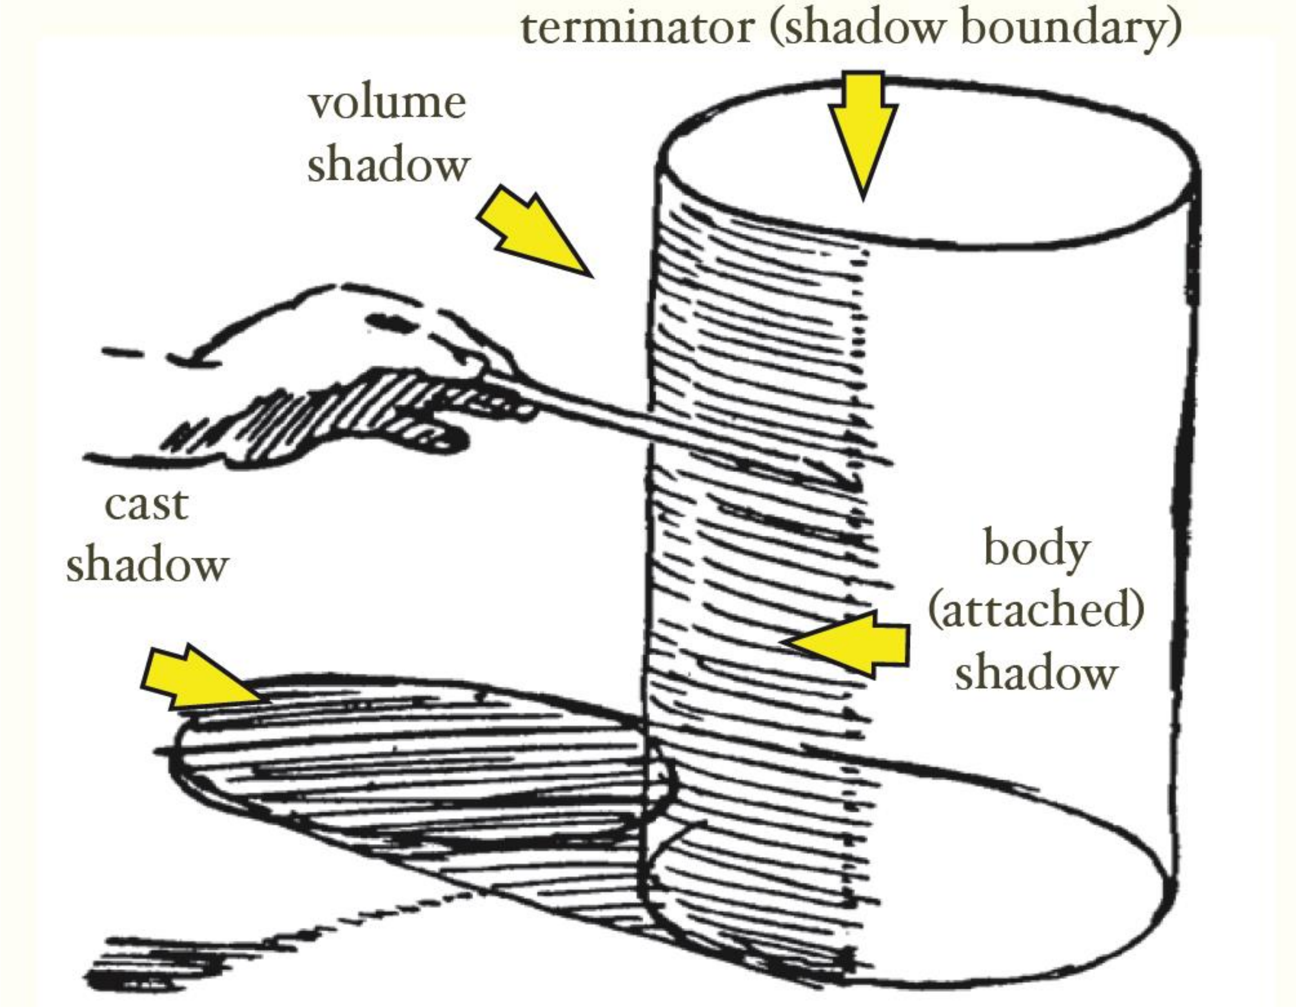

From Koenderink slides on image texture and the flow of light

Slide: Forsyth

Color constancy

- Interpret surface in terms of albedo or “true color”, rather than observed intensity
  - Humans are good at it
  - Computers are not nearly as good

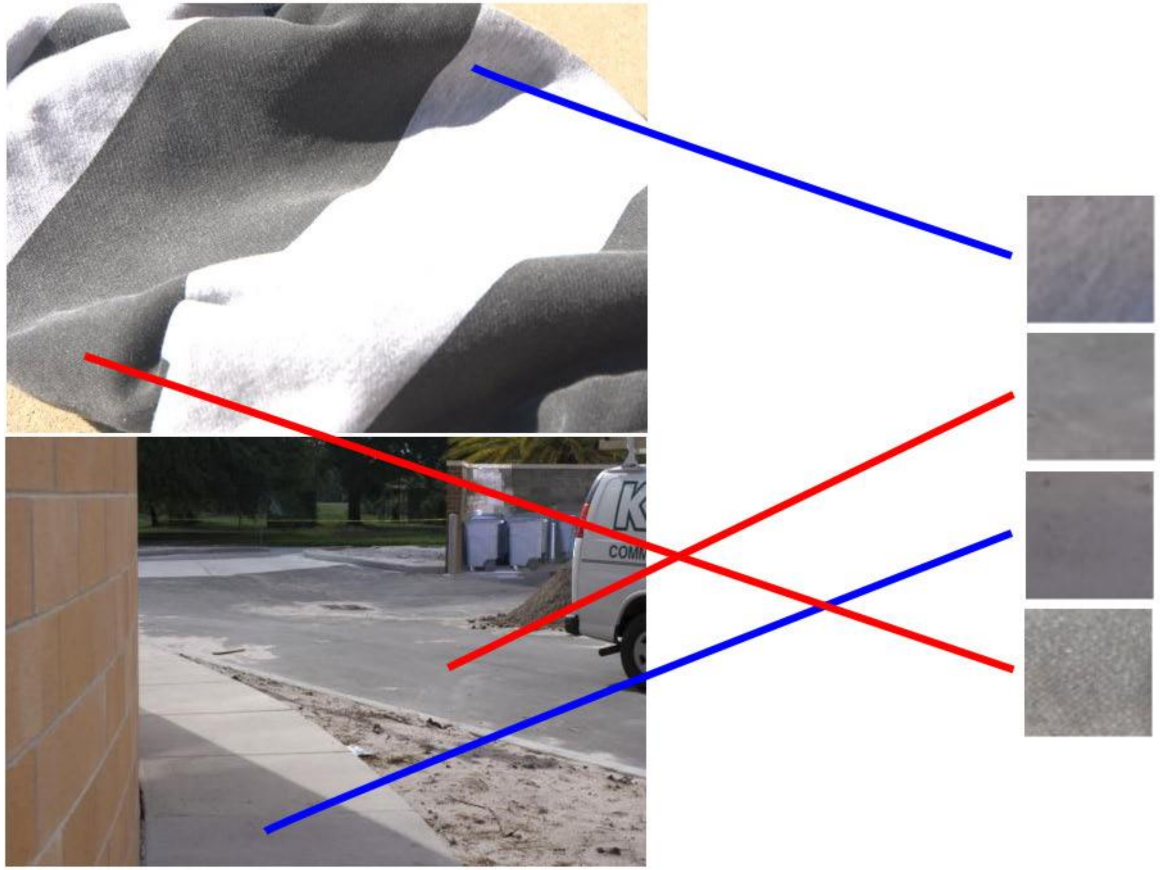

One source of constancy: local comparisons

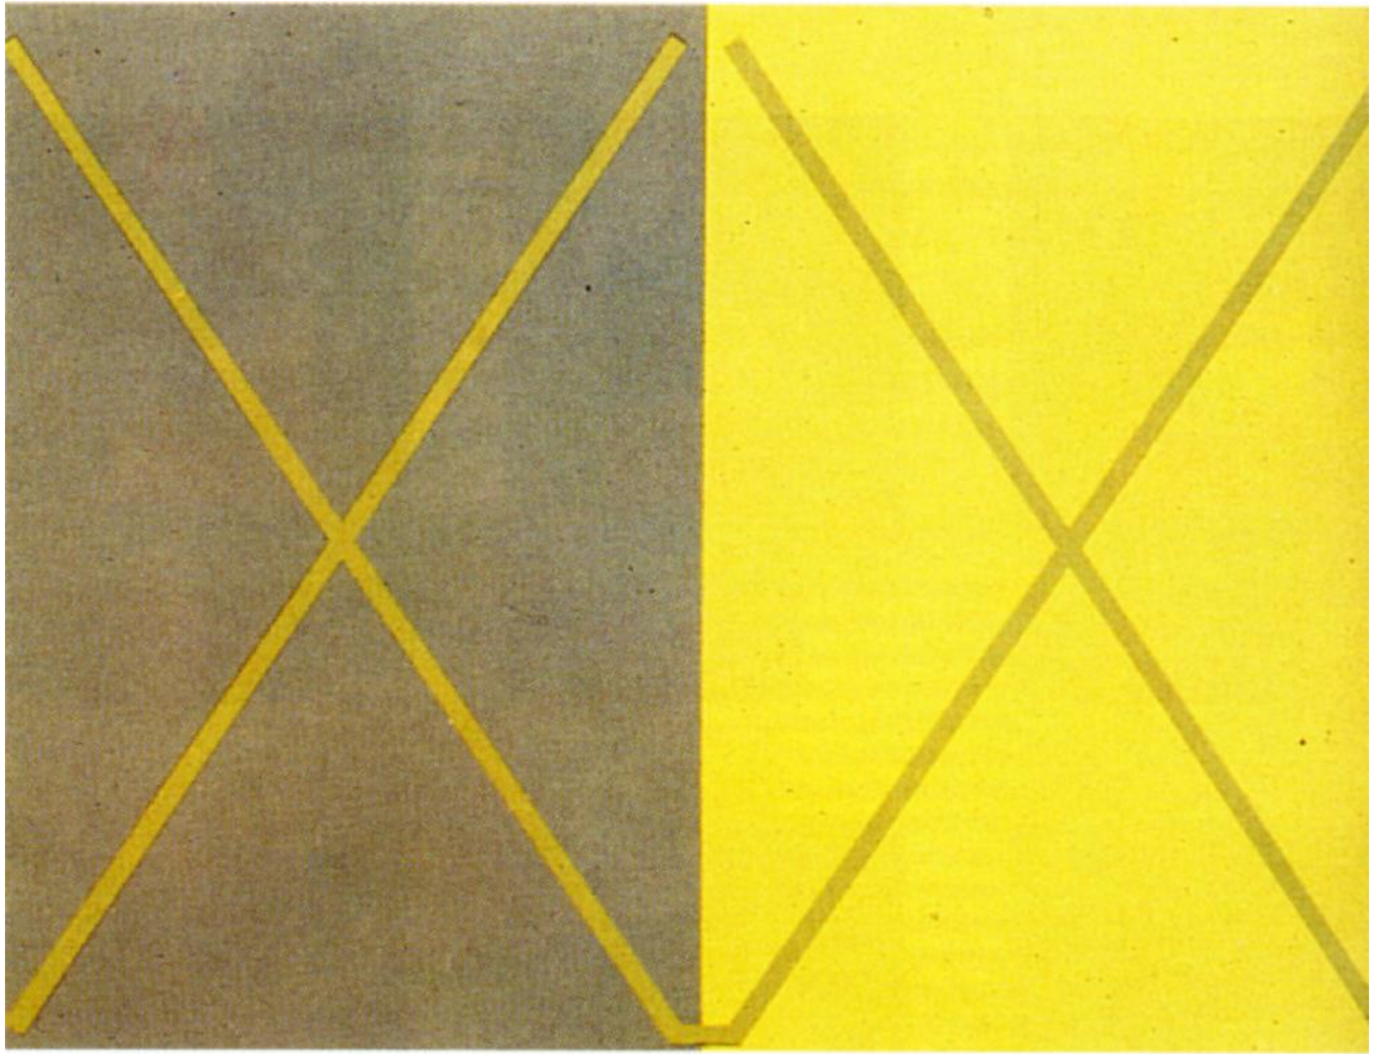

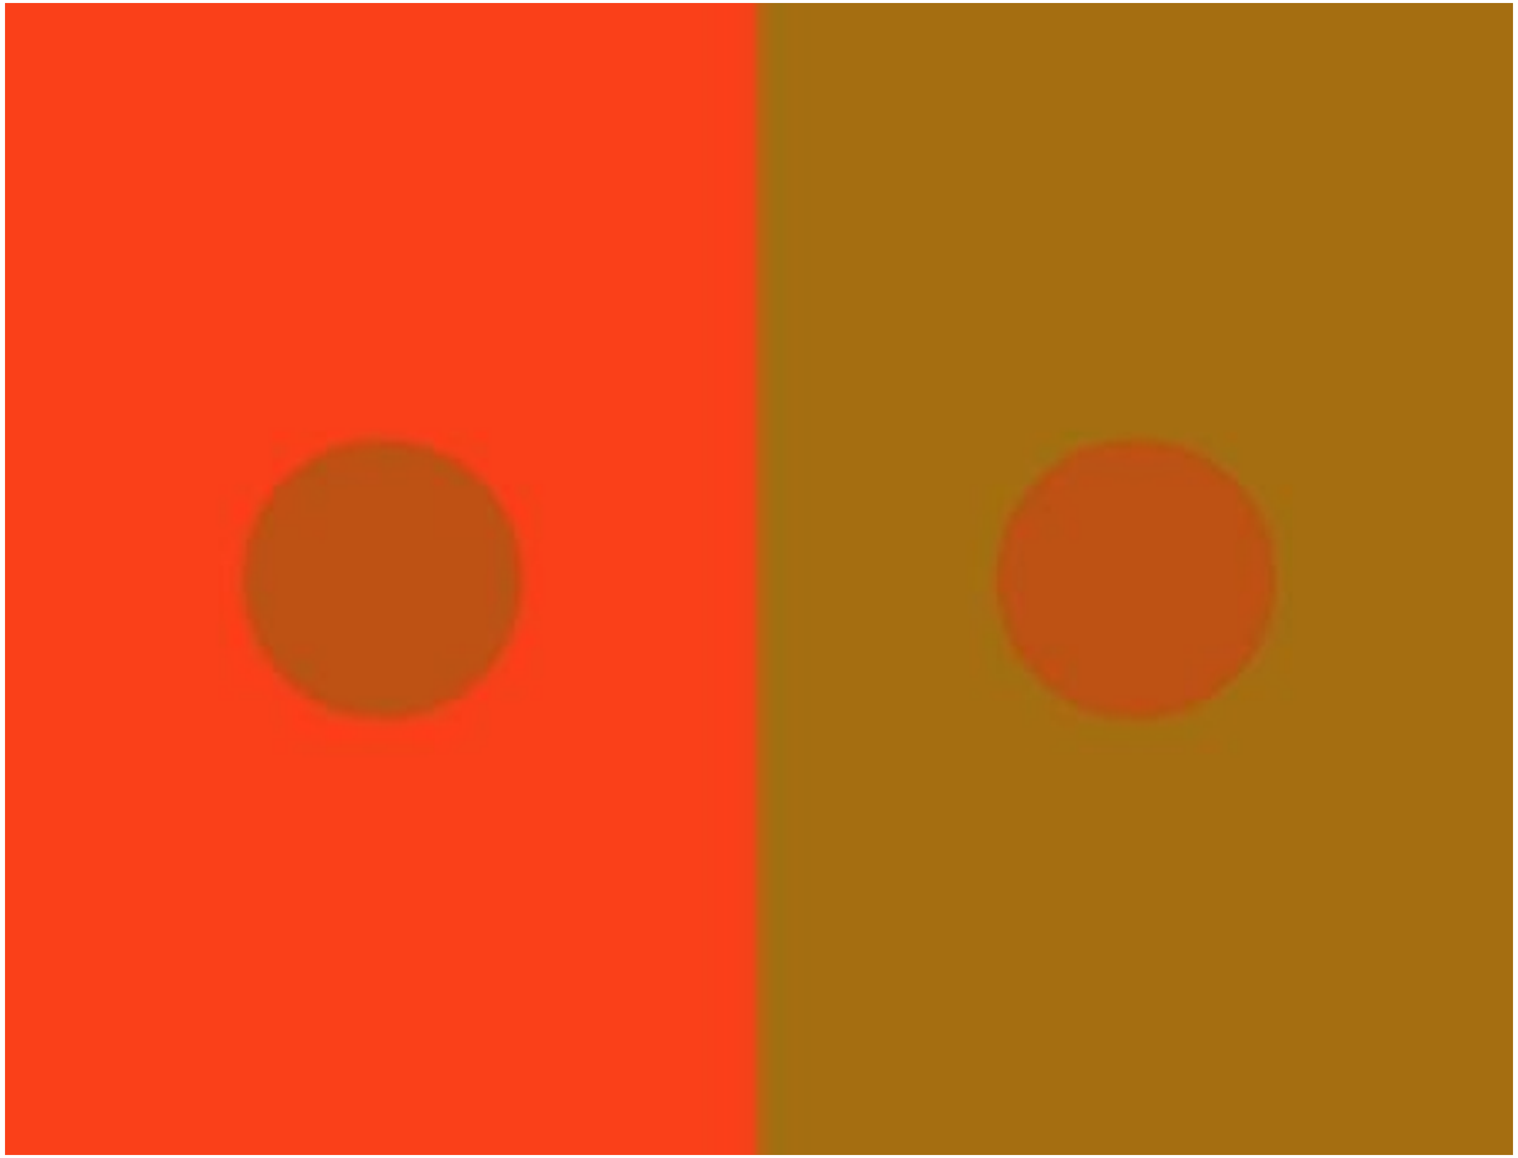

http://www.echalk.co.uk/amusements/OpticalIllusions/colourPerception/colourPerception.html

Perception of Intensity

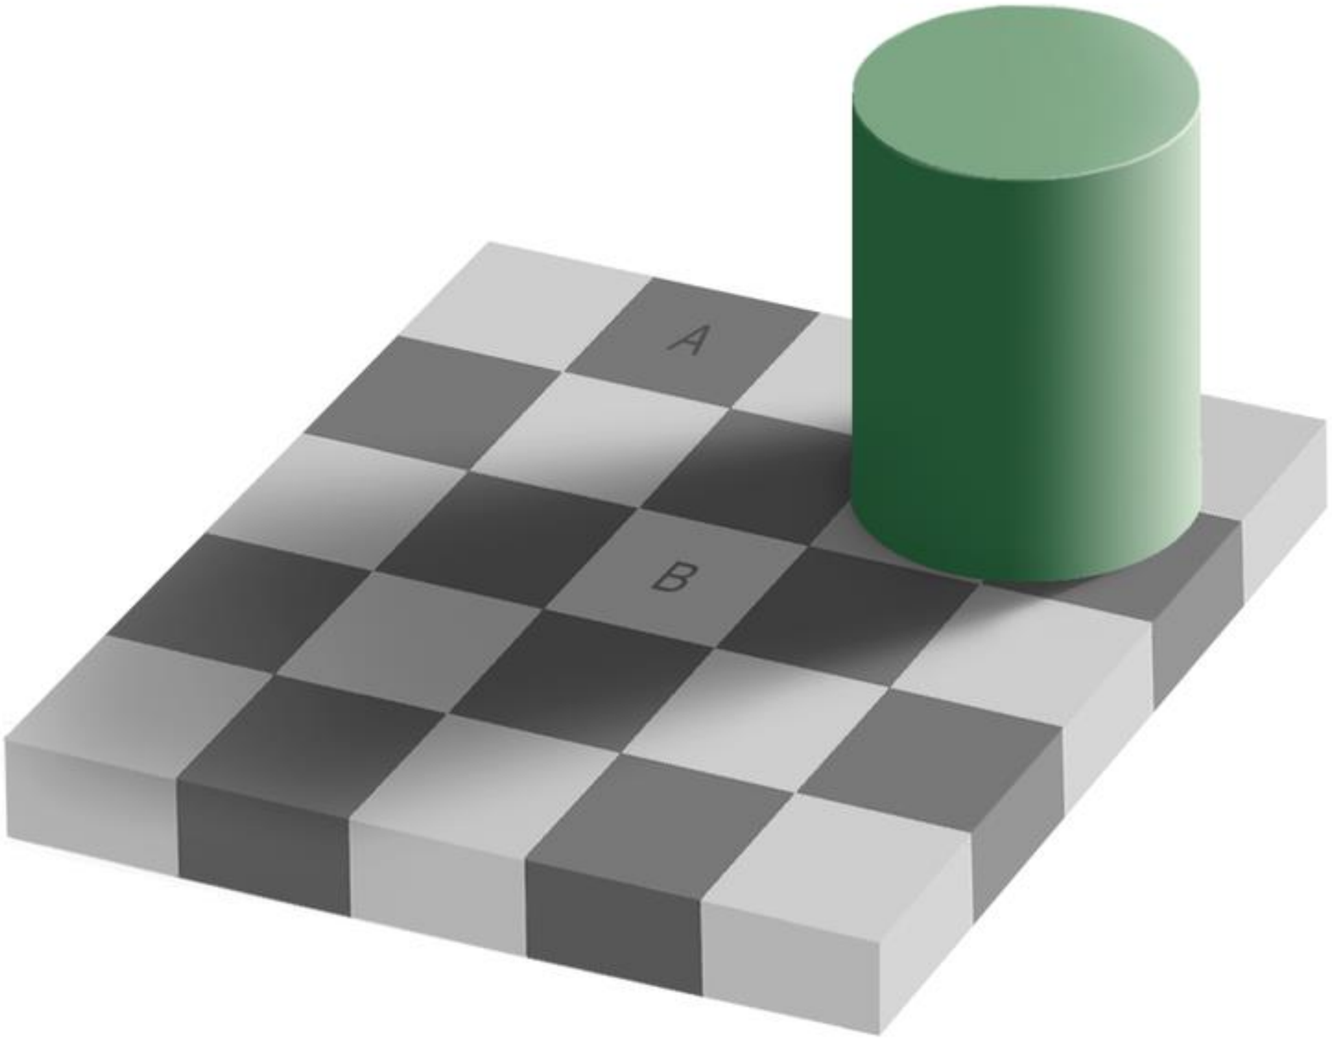

from Ted Adelson

Perception of Intensity

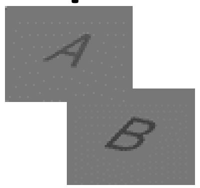

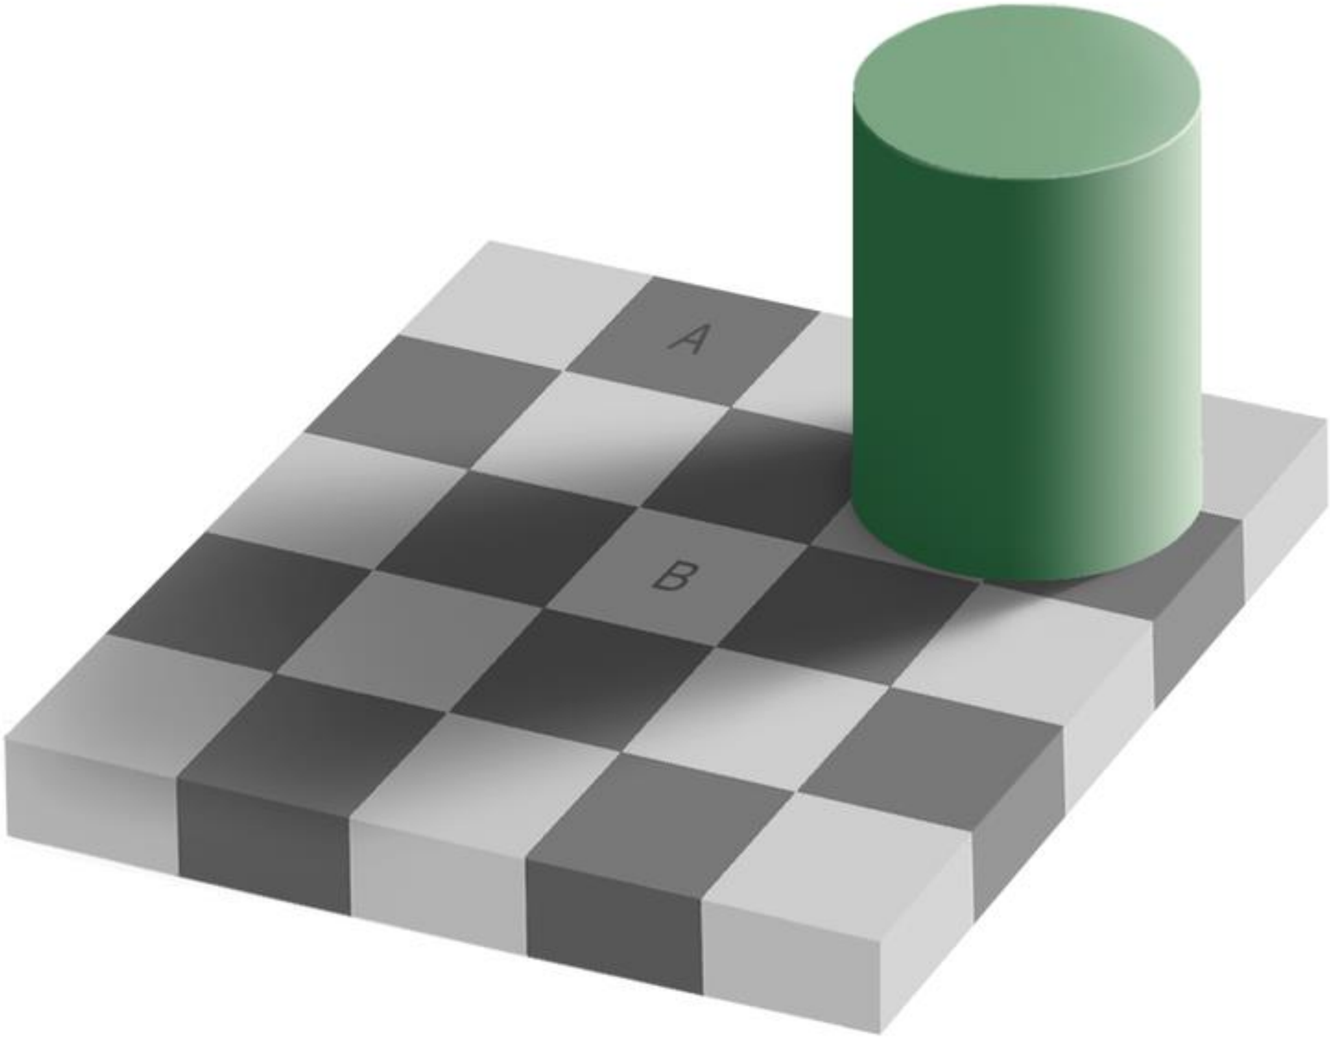

from Ted Adelson

Color Correction

- Simple idea: multiply R, G, and B values by separate constants

$$\begin{bmatrix} \tilde{r} \\ \tilde{g} \\ \tilde{b} \end{bmatrix} = \begin{bmatrix} \alpha_r & 0 & 0 \\ 0 & \alpha_g & 0 \\ 0 & 0 & \alpha_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

- How to choose the constants?
  - "White world" assumption: brightest pixel is white
    - Divide by largest value
  - "Gray world" assumption: average value should be gray
    - E.g., multiply r channel by avg(r) /avg((r+g+b)/3)
  - White balancing: choose a reference as the white or gray color

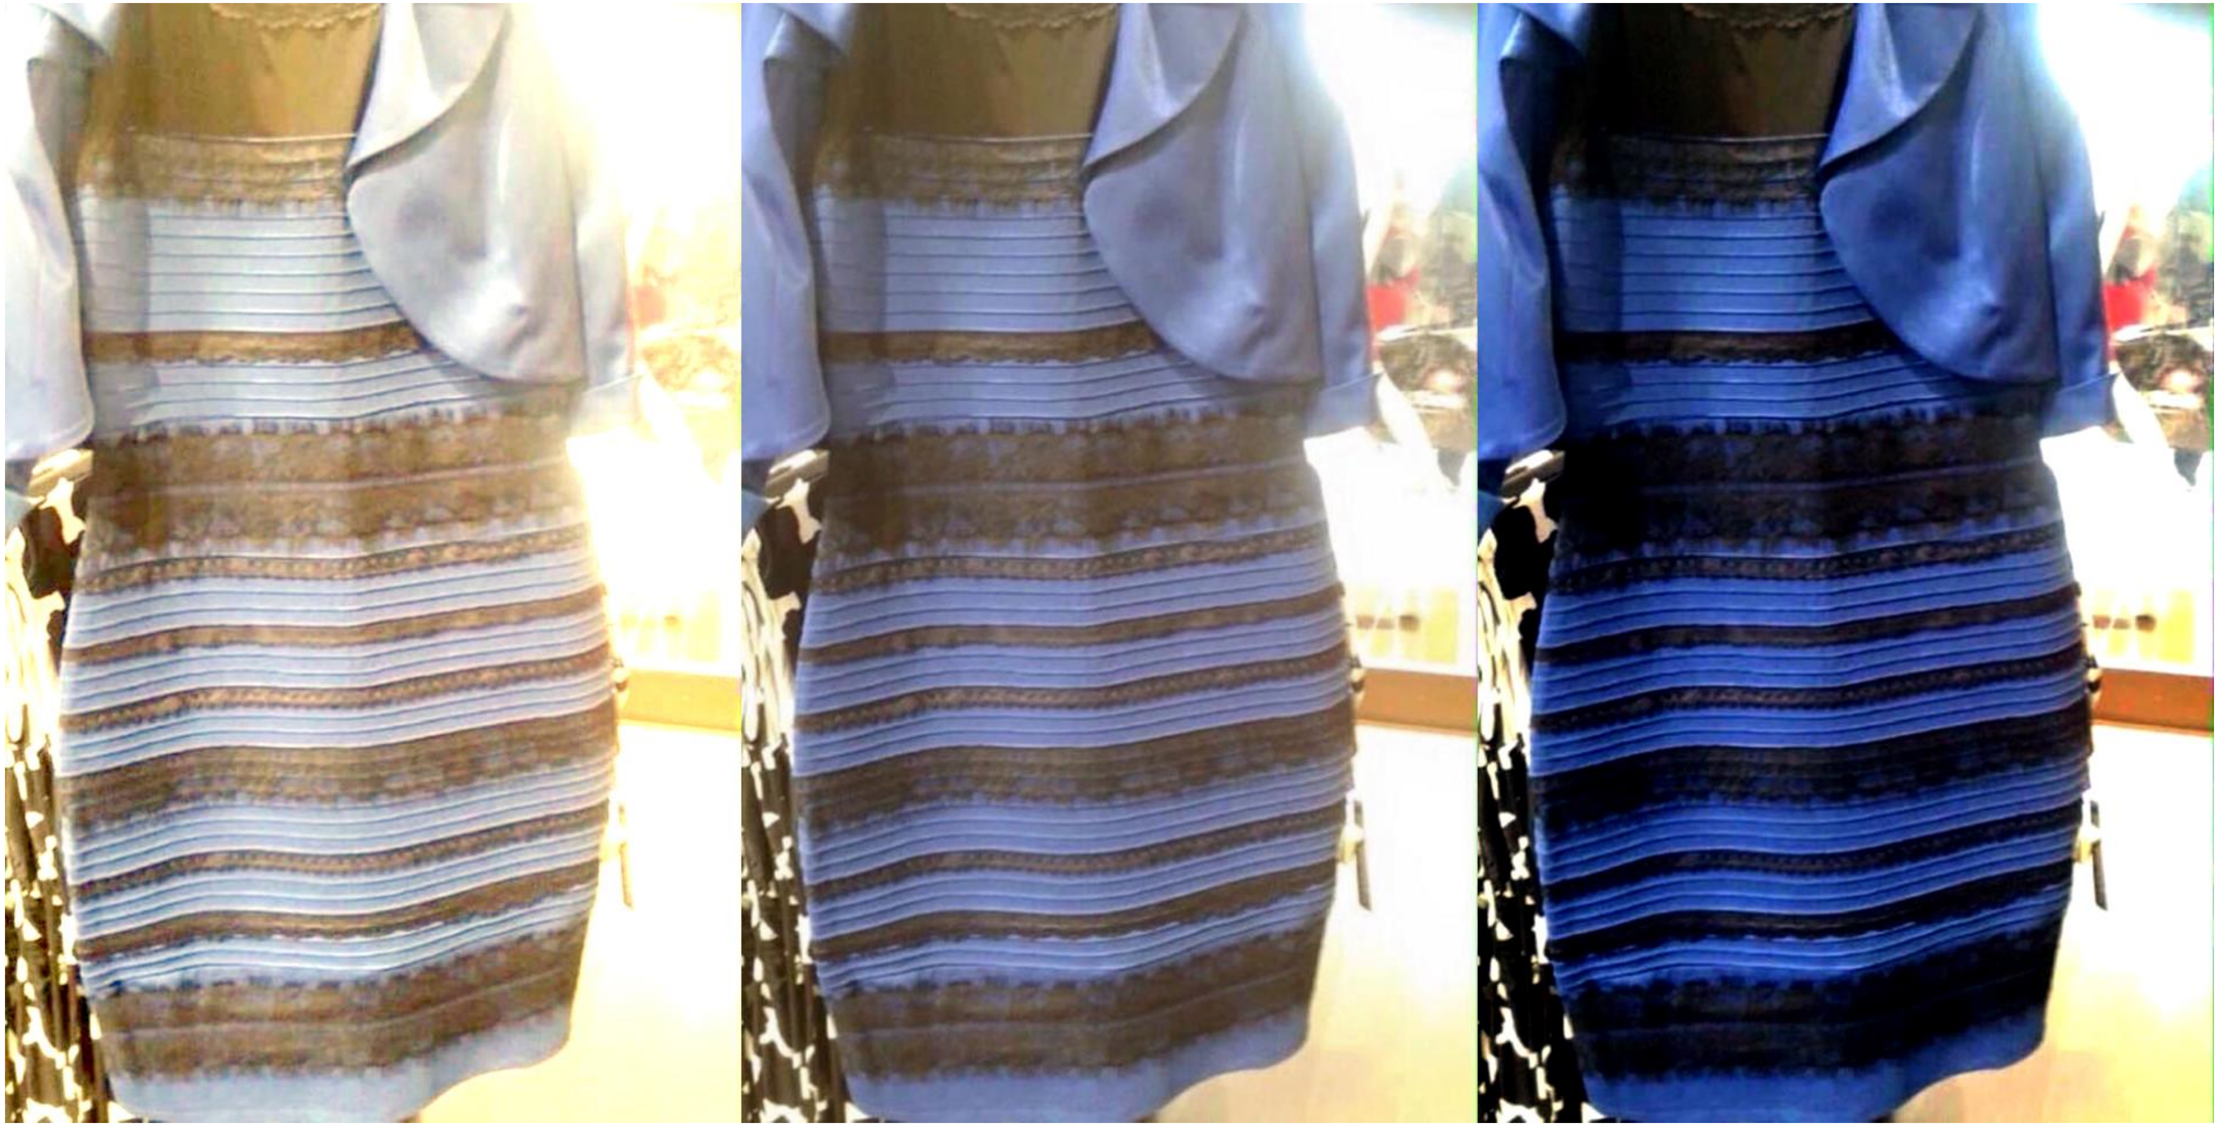

Discount the blue side

Discount the gold side

Things to remember

- Important terms: diffuse/specular reflectance, albedo
- Color vision: physics of light, trichromacy, color consistency, color spaces (RGB, HSV, Lab)
- Observed intensity depends on
  - light sources,
  - geometry/material of reflecting surface,
  - surrounding objects,
  - camera settings
- Objects cast light and shadows on each other
- Differences in intensity are primary cues for shape

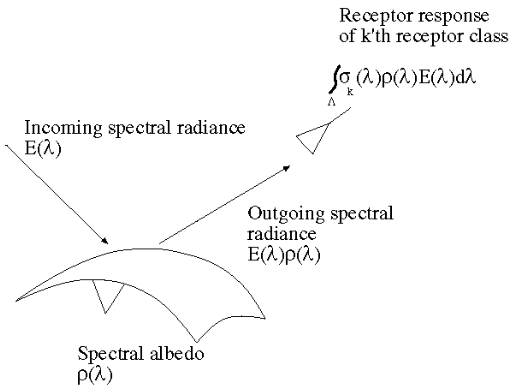

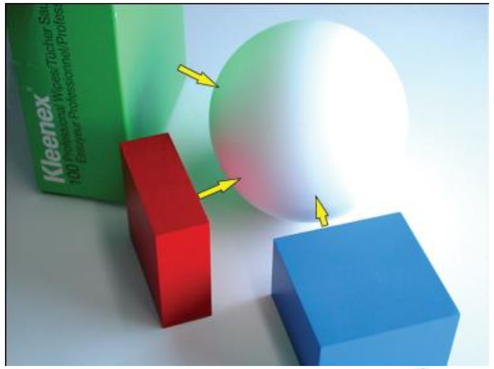

Thank you

- Next class: Image Filters