Light & Perception

CS5670

March 16, 2020

Remote Lectures

- Today we are trying out Zoom lecture
  - If all goes well, this will likely be the main method of holding class moving forward
  - Some lectures may end up being pre-recorded, but we will have office hours where you can ask questions

- You will be able to complete the course remotely (including exams)
  - Specifics to come

Announcements

- Assignment 3 deadline extended to Friday March 27, 11:59pm EDT
- Artifact due at the same time

Road map

- Today: Light & Perception
  - Readings: Szeliski, 2.2, 2.3.2
- Finishing up geometry
  - ~3 more lectures
- Next:
  - Recognition
  - Classification
  - Deep learning,
  - A sample of hot topics in computer vision

CS5670: Computer Vision

Abe Davis (most slides from Noah Snavely)

Light & Perception

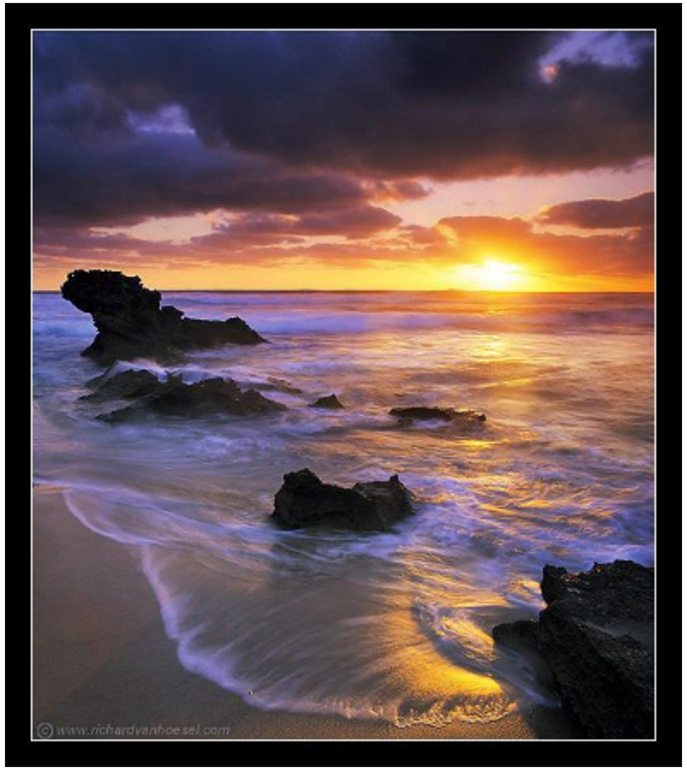

Properties of light

- Today
  - What is light?
  - How do we measure it?
  - How does light propagate?
  - How does light interact with matter?

Reflectance, Illumination, and Perception

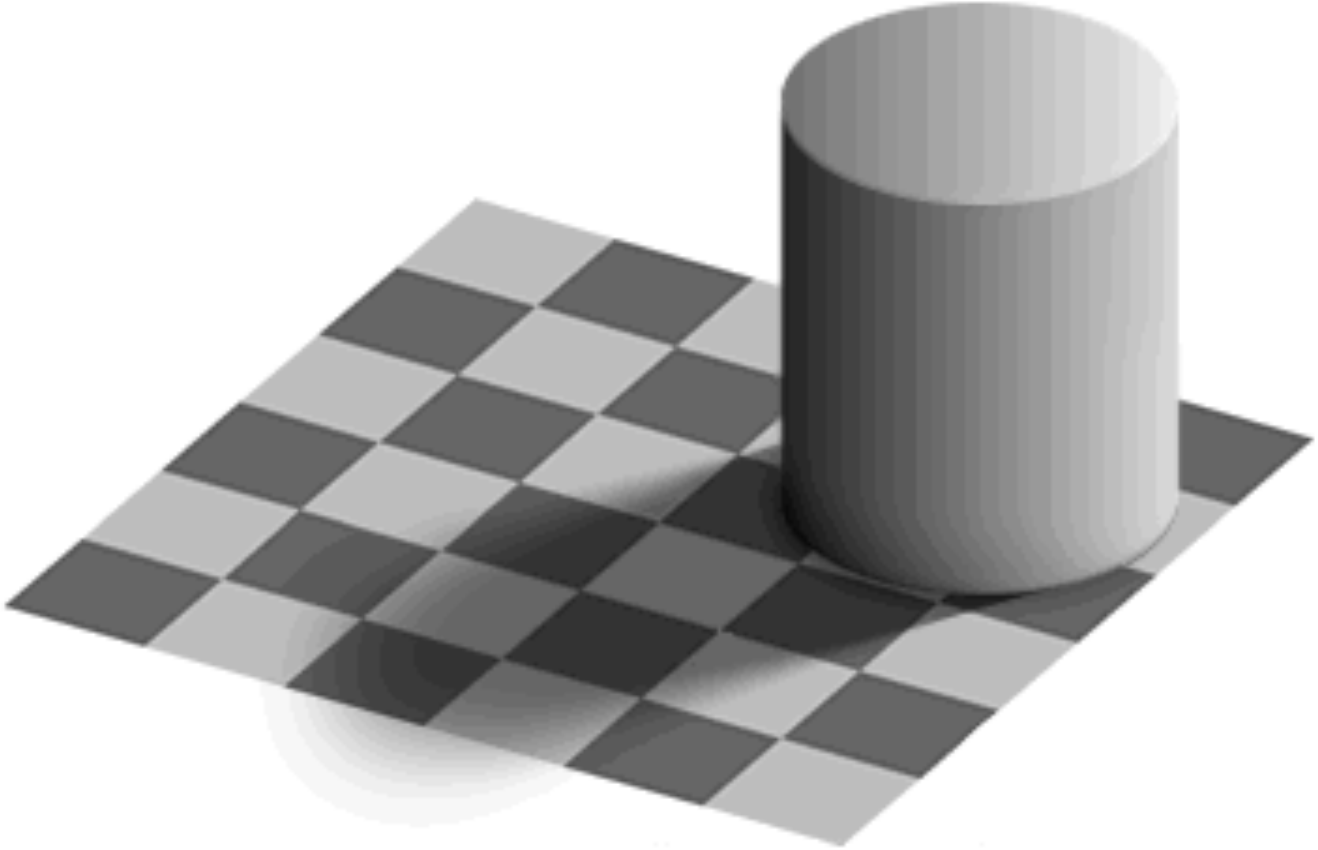

by Ted Adelson

Reflectance, Illumination, and Perception

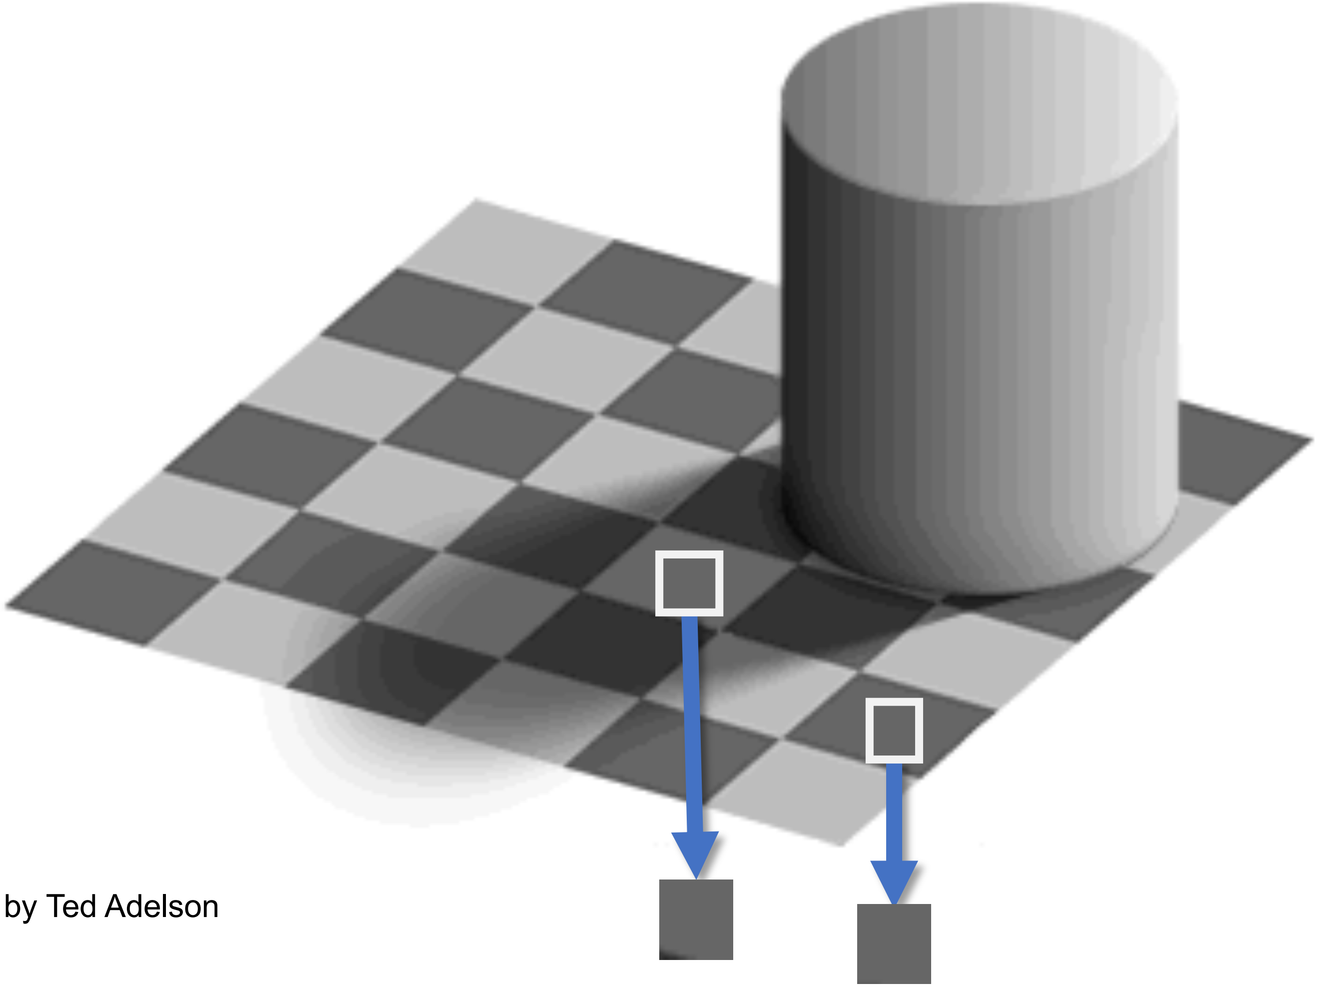

by Ted Adelson

Reflectance, Illumination, and Perception

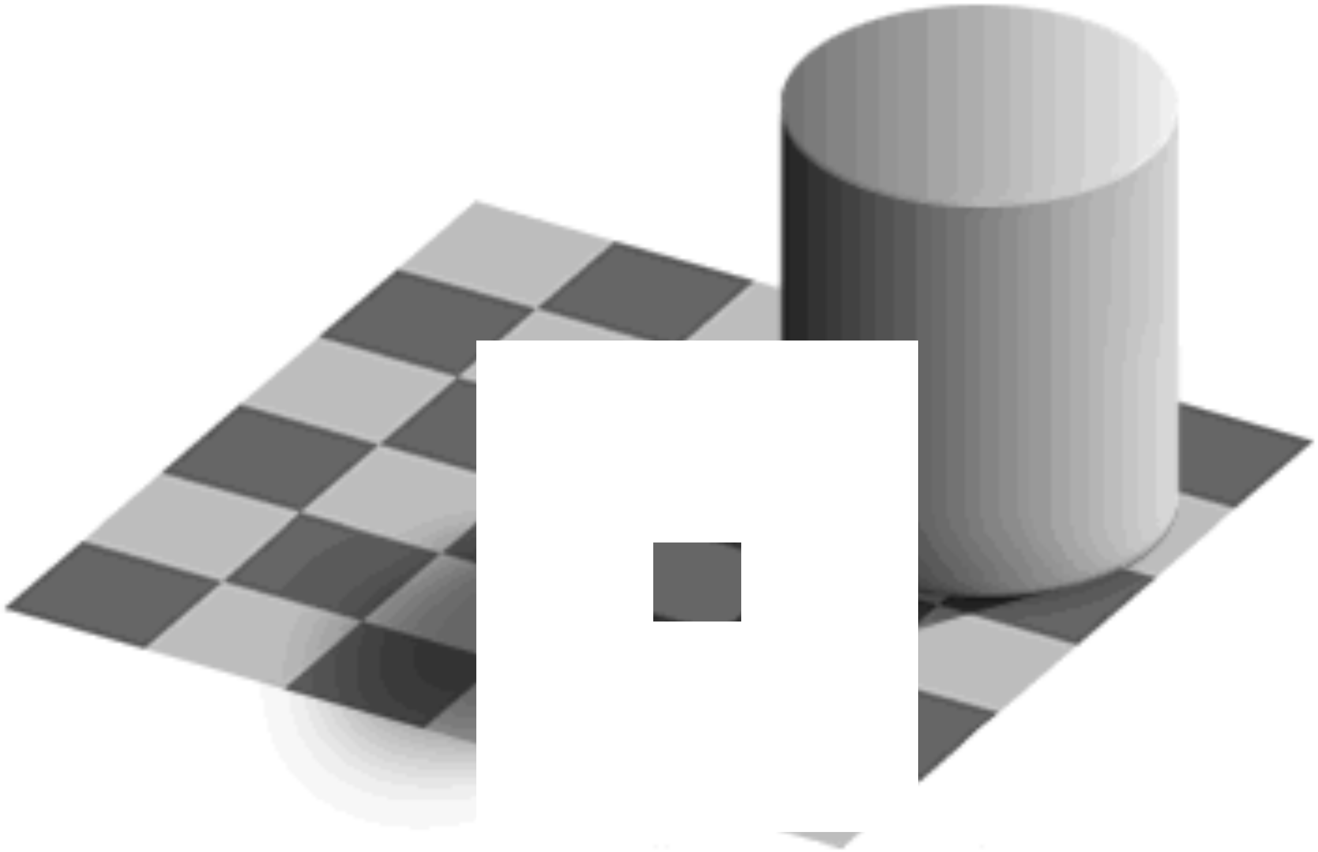

by Ted Adelson

From Lecture 1:

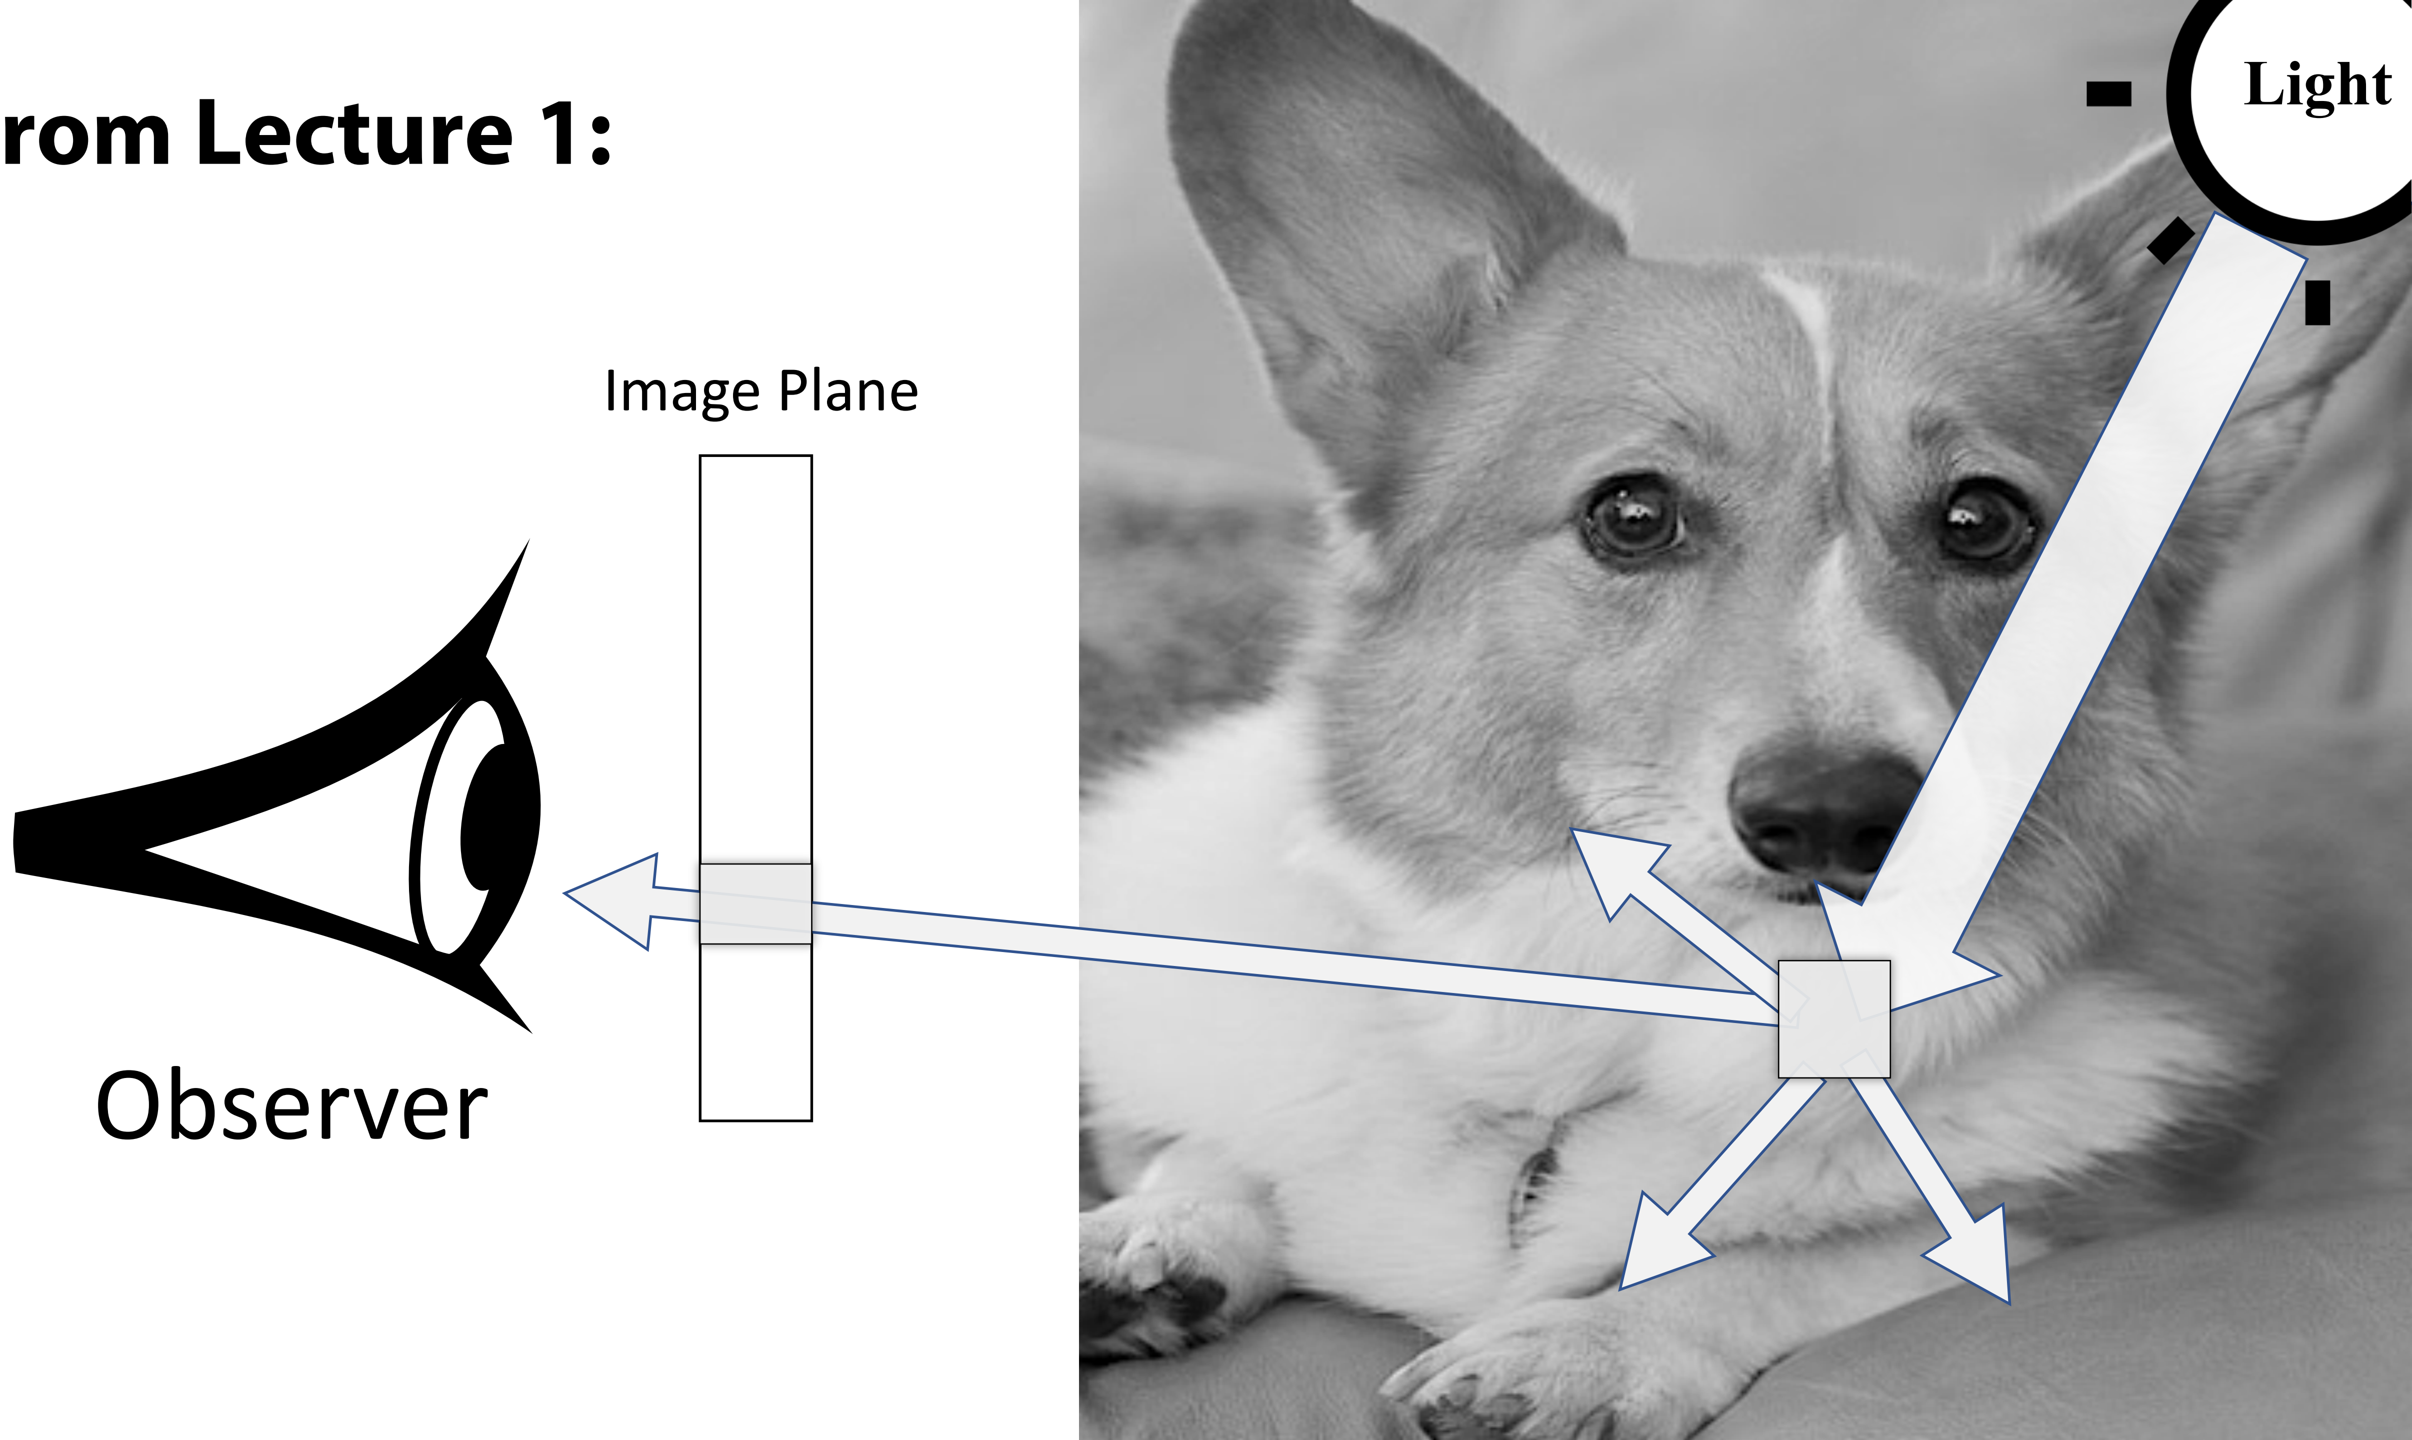

Reflectance

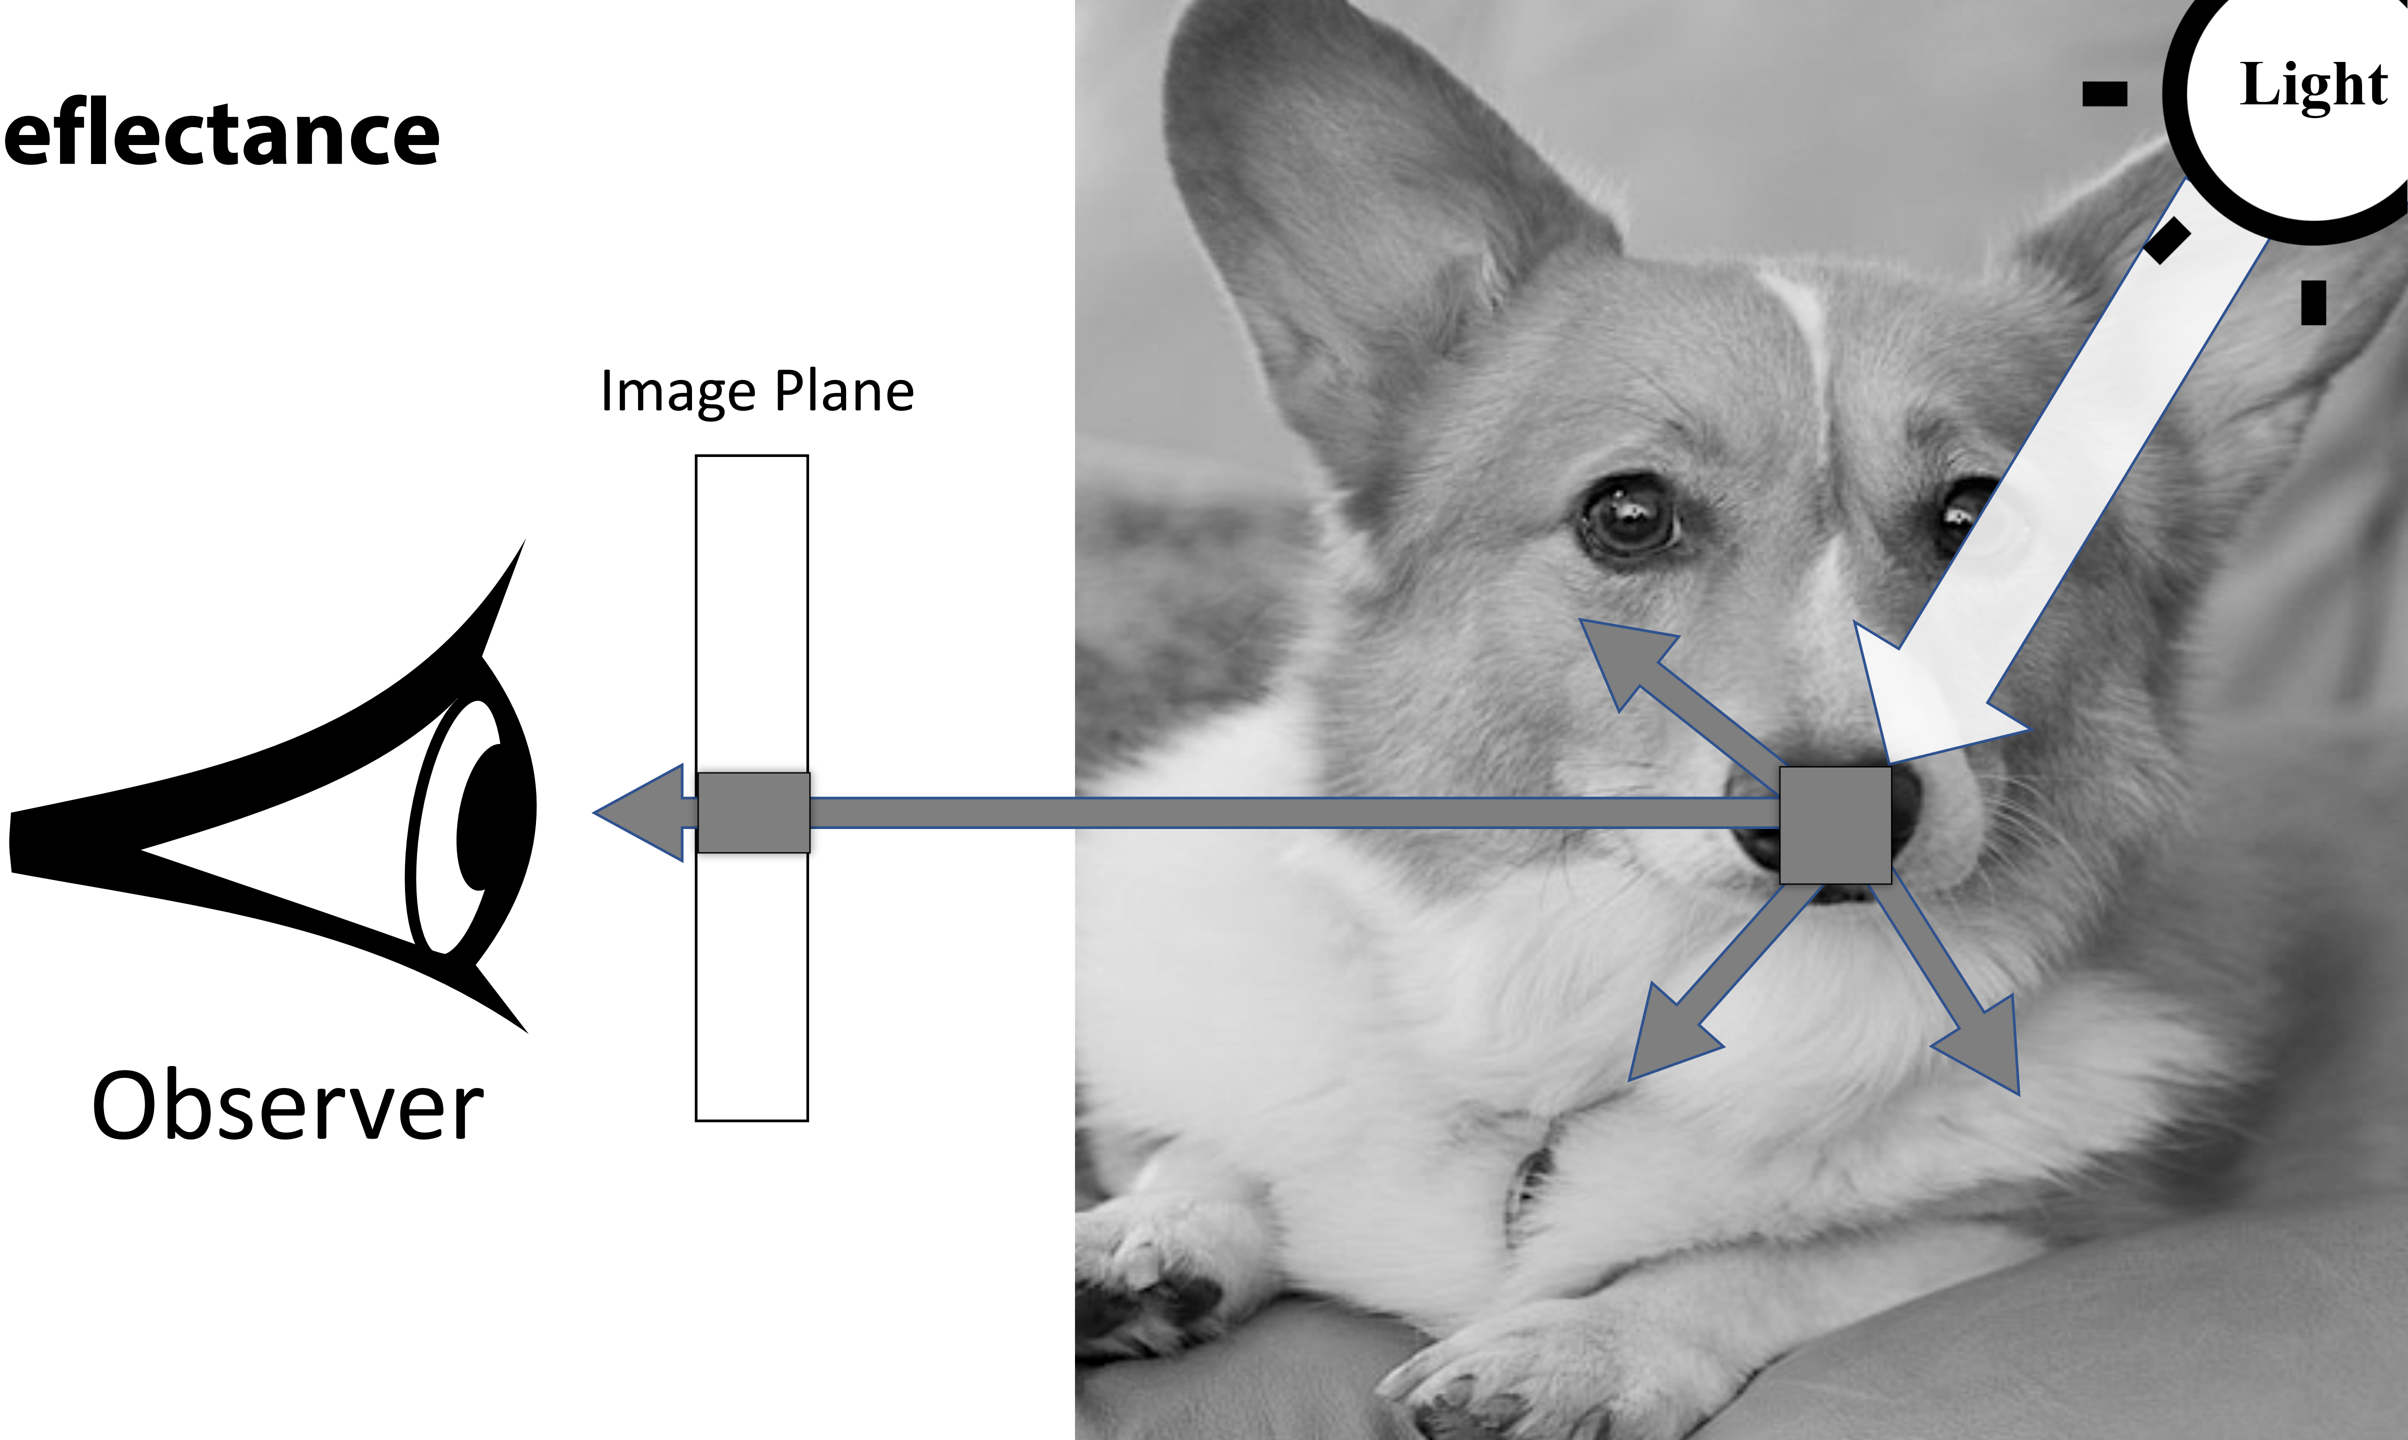

Illumination

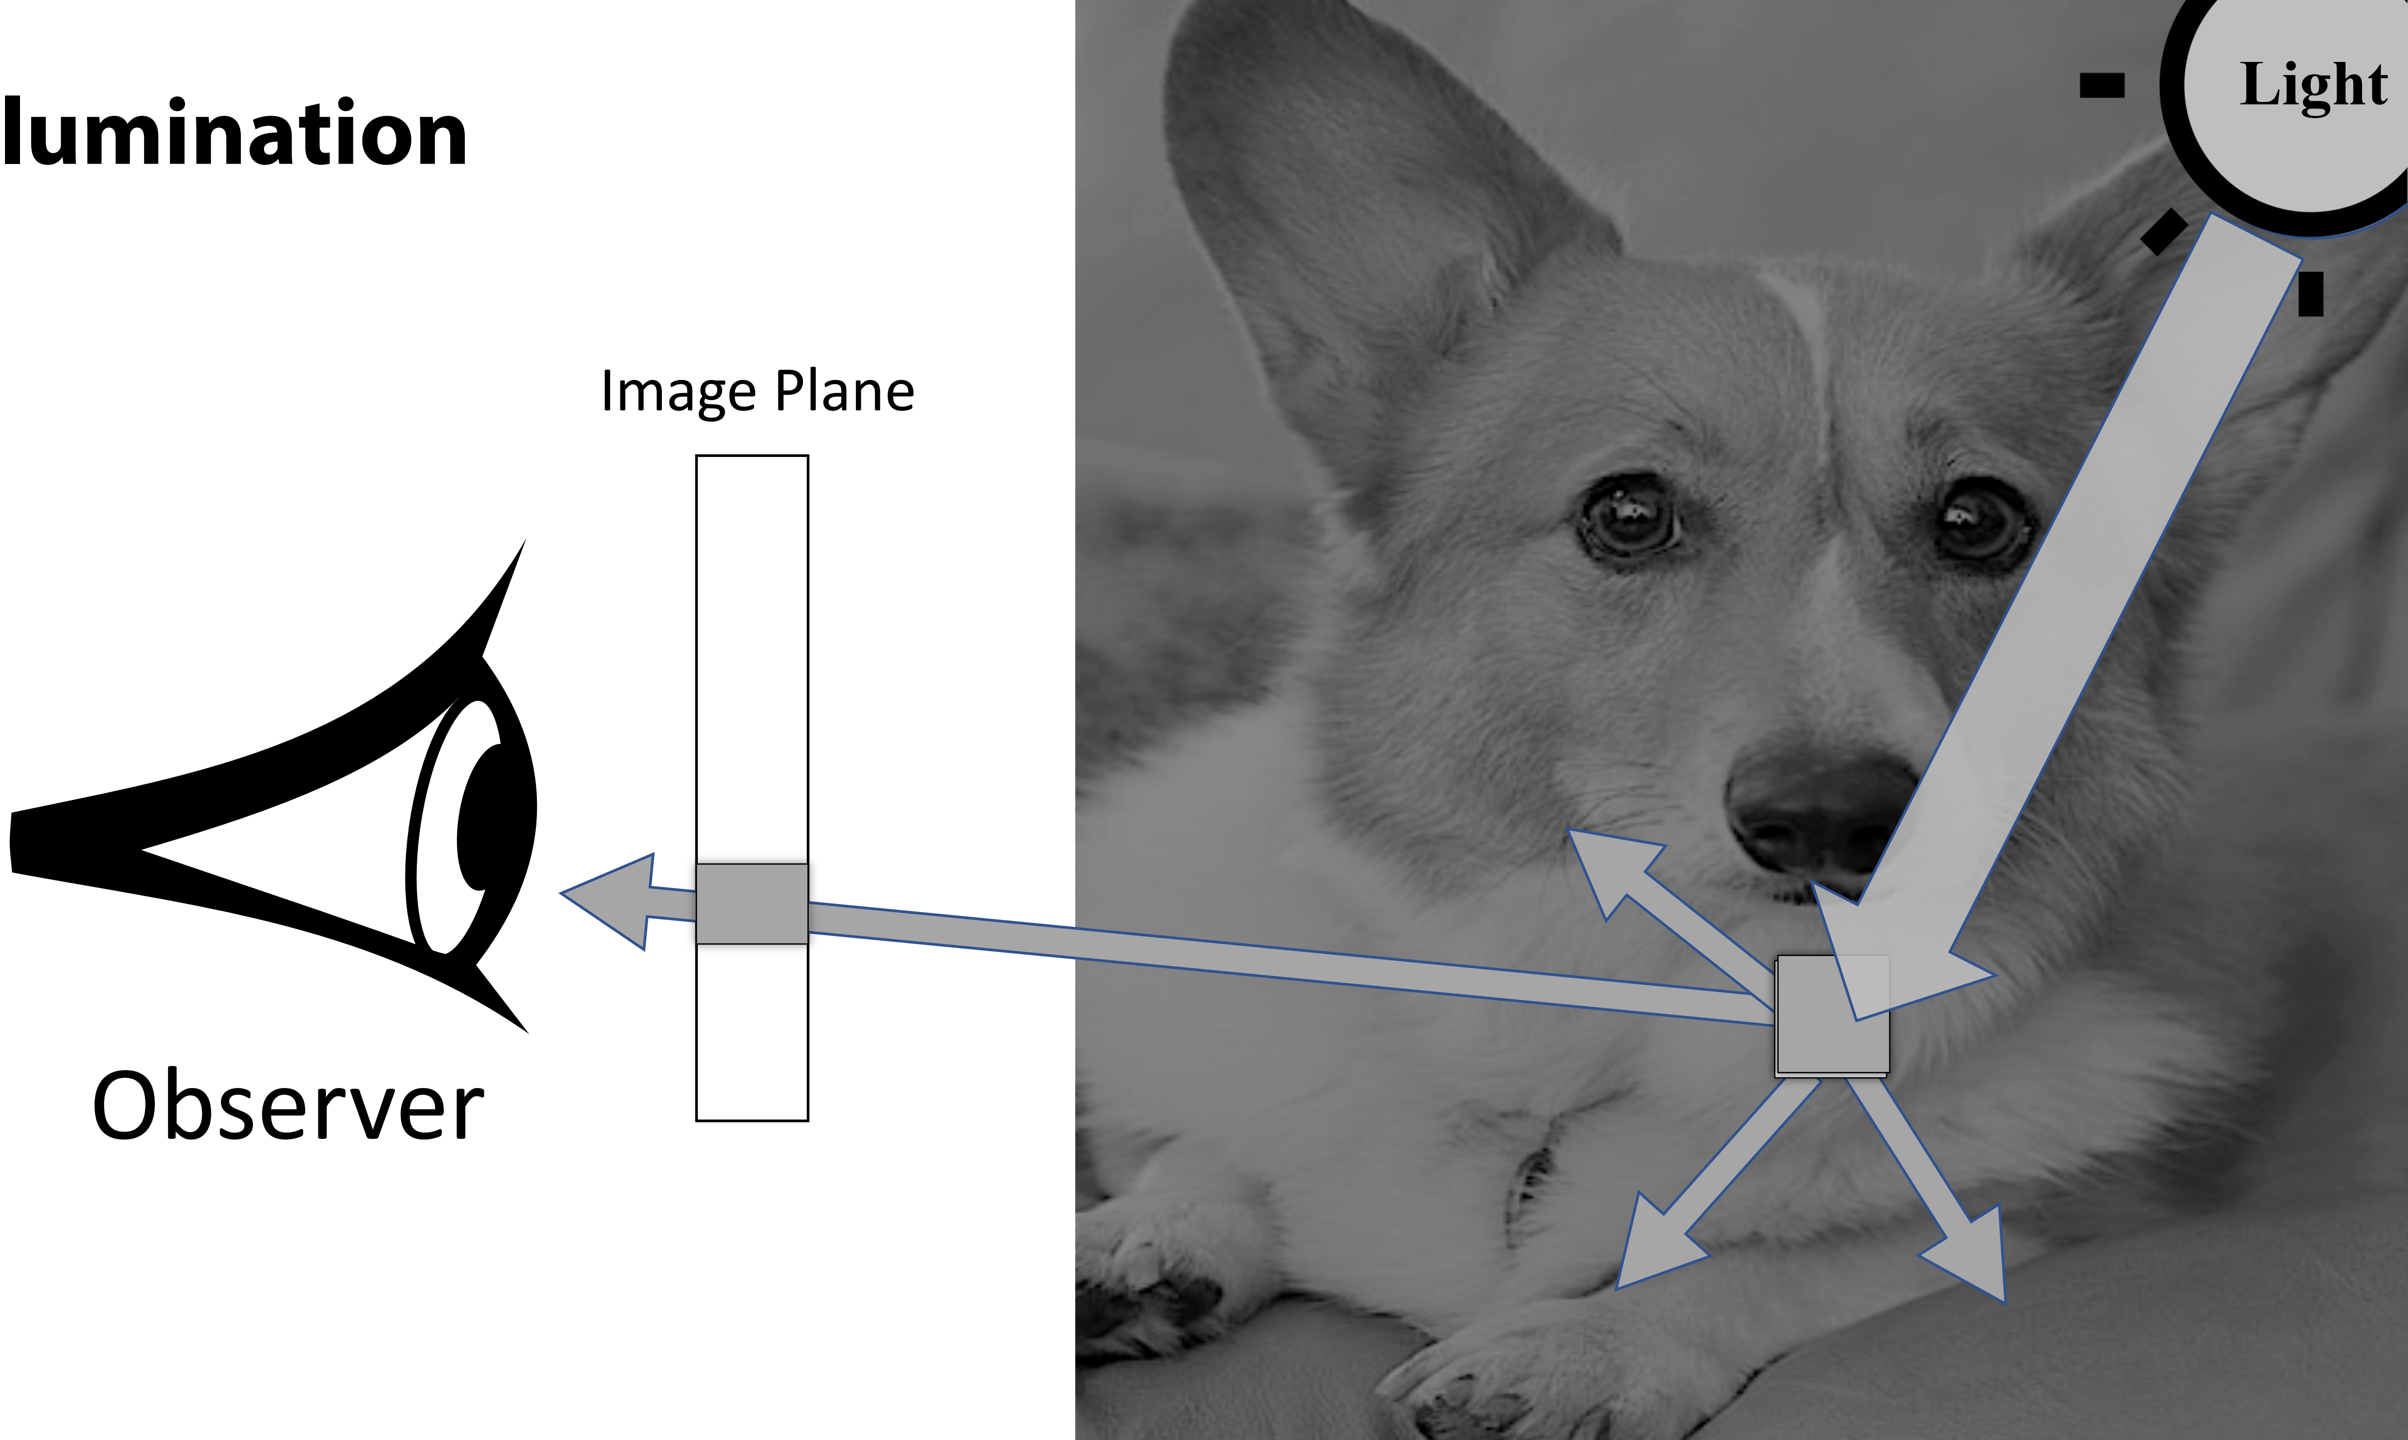

Illumination

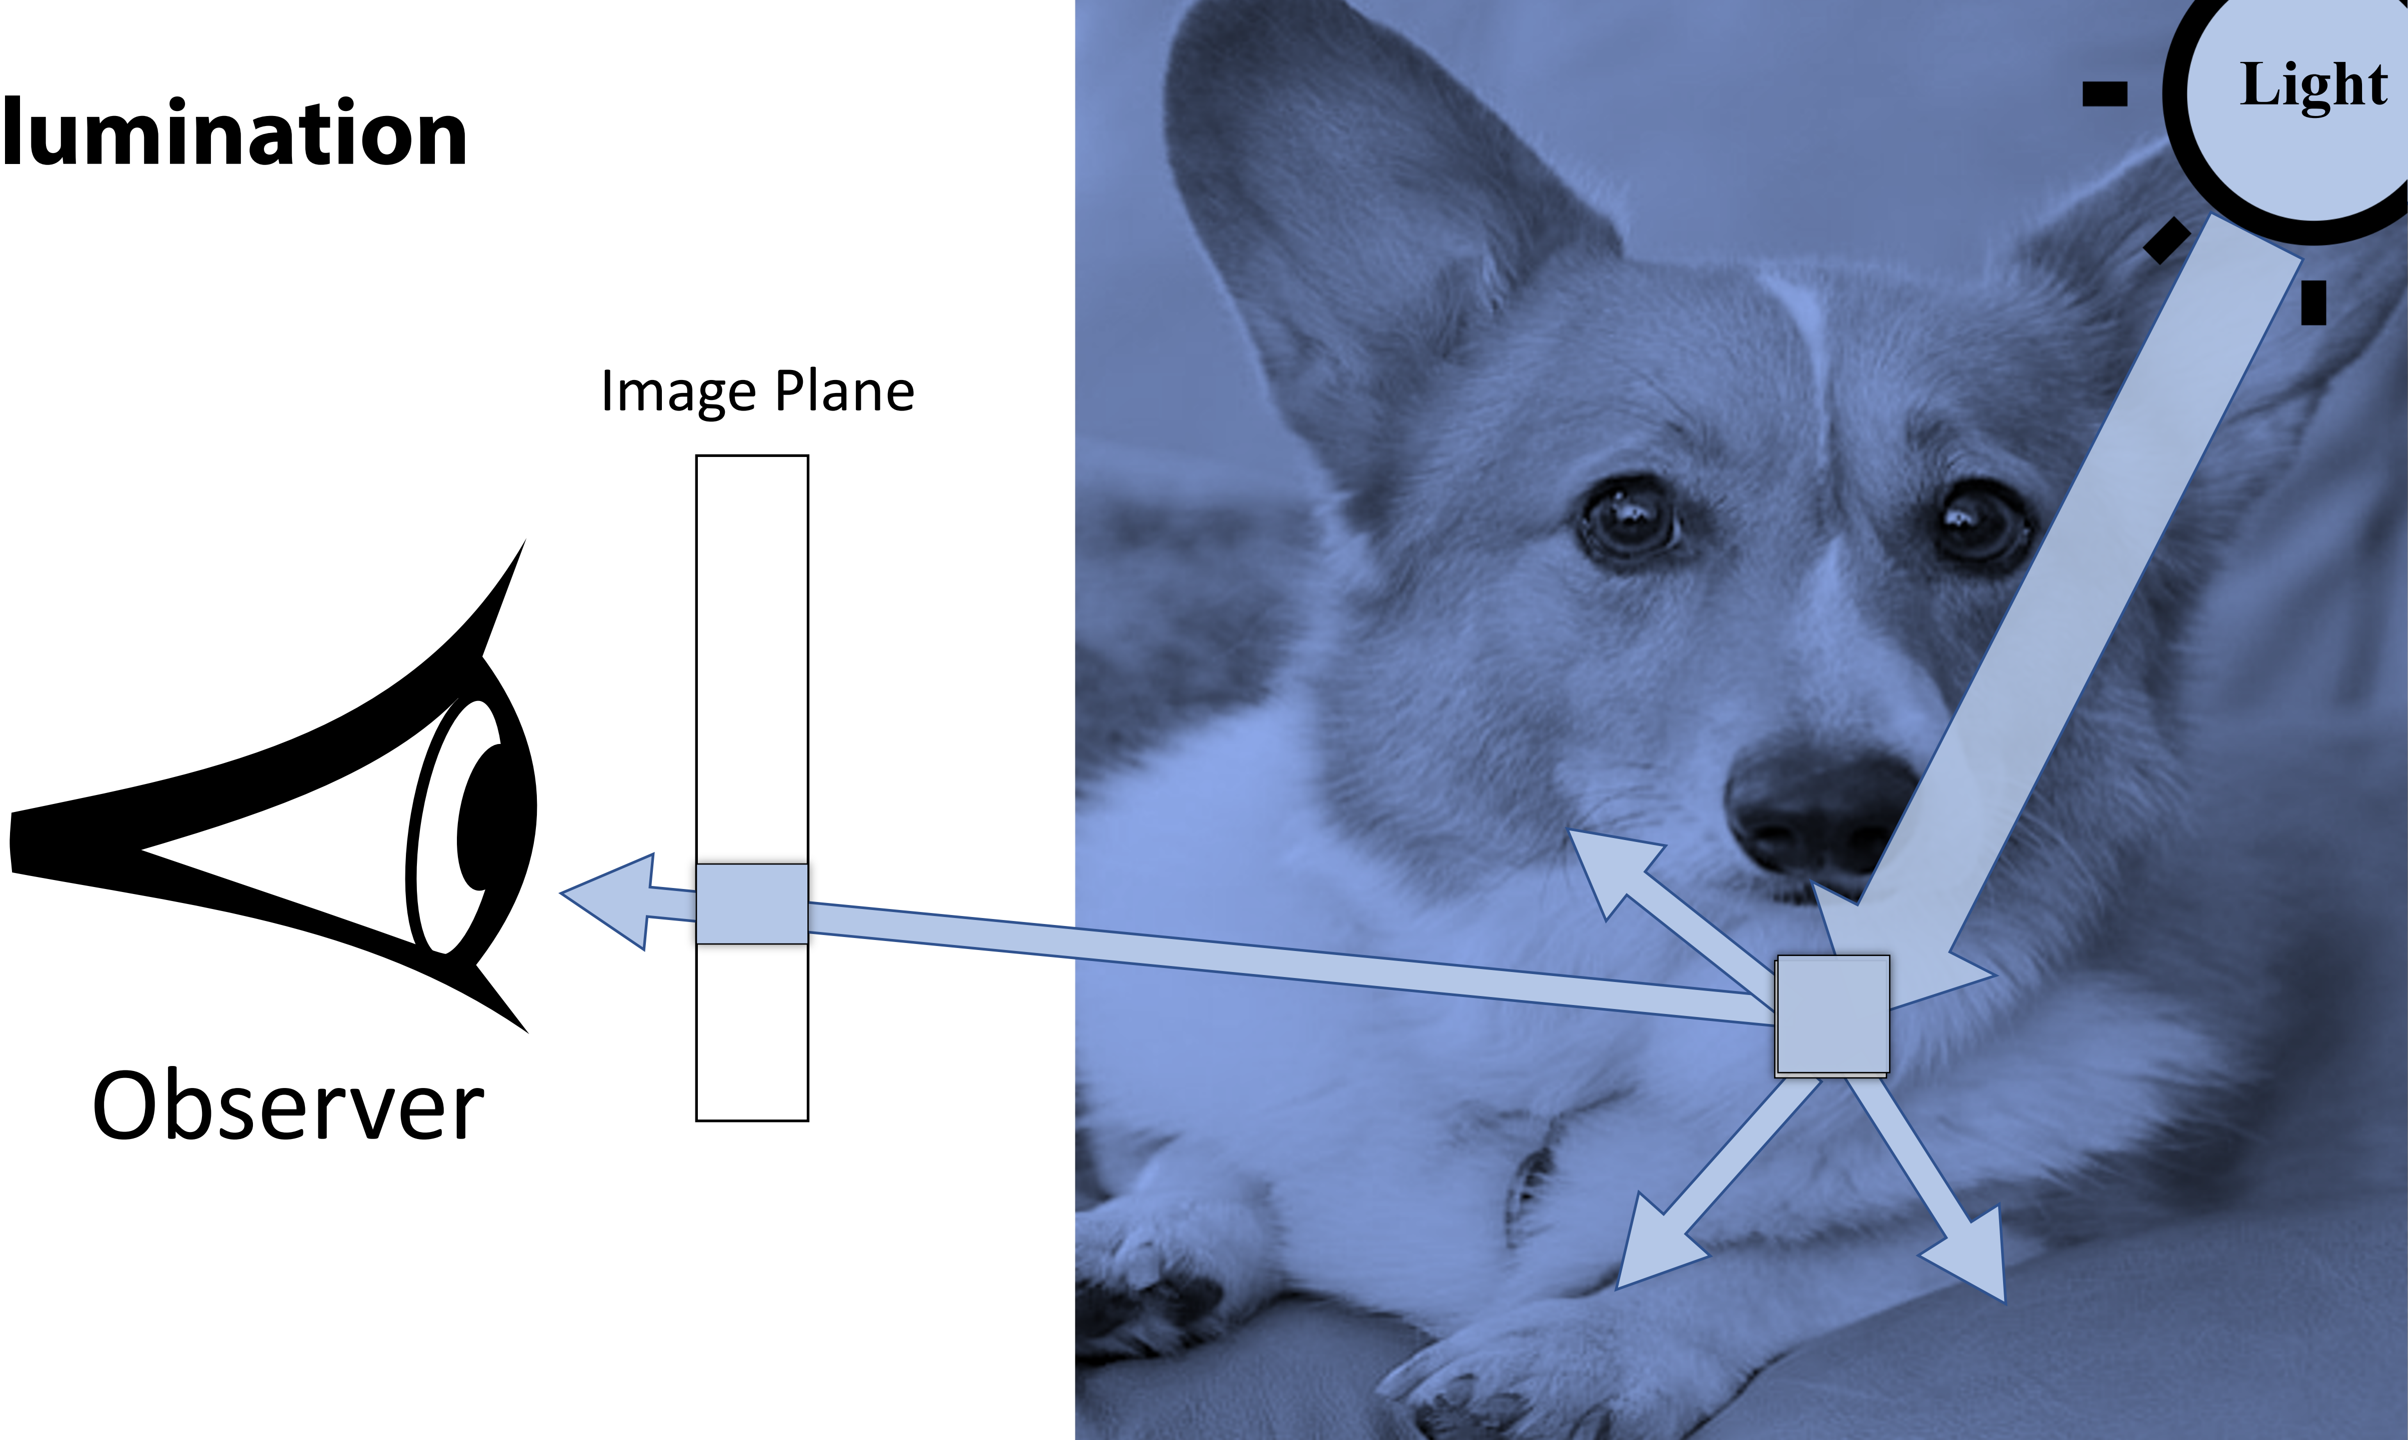

What Determines the Brightness of a Surface?

- Illumination
  - What light hits a surface
- Reflectance
  - How the surface reflects light

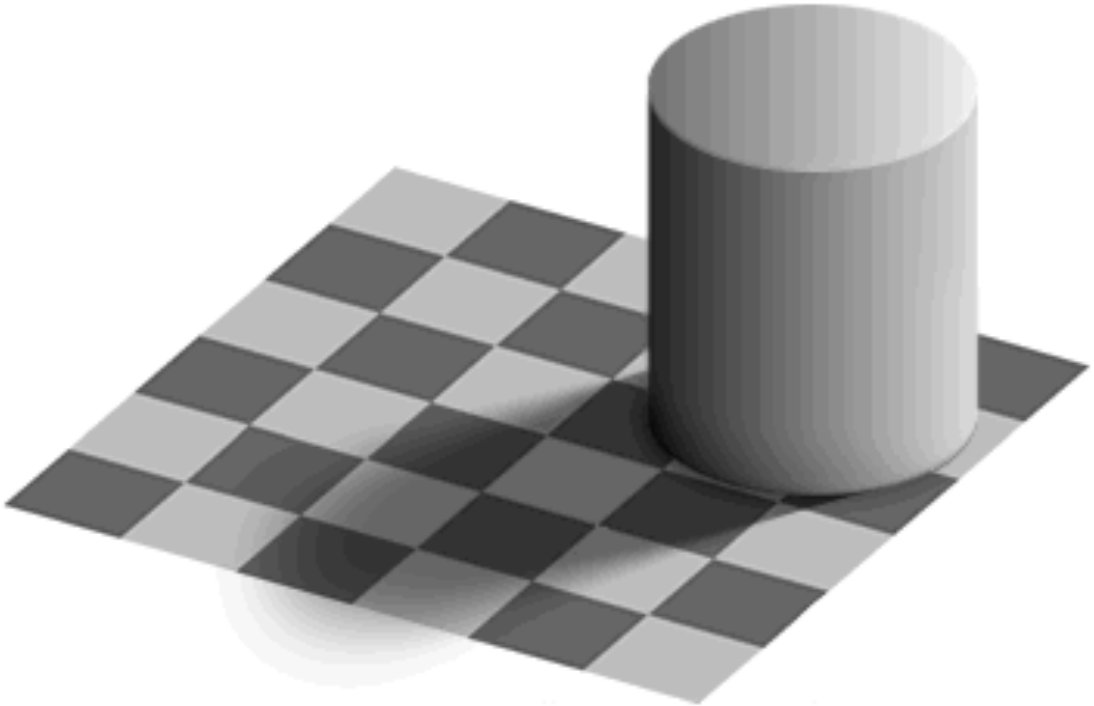

The brightness of a pixel is an ambiguous combination of illumination and reflectance…

$$A * B = C$$

A → Illumination; B → Reflectance; C → Reflected Light

Reflectance and Illumination In Popular Culture…

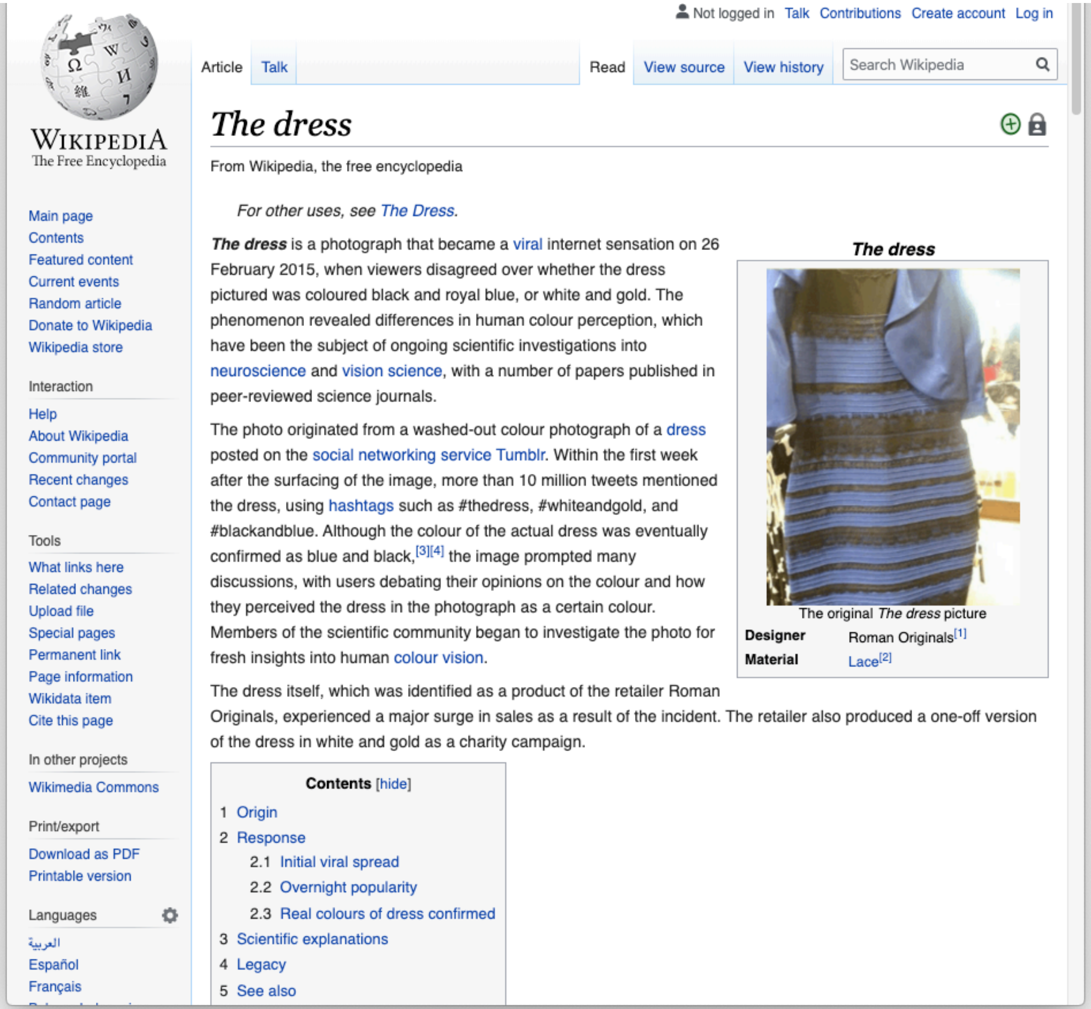
Not logged in Talk Contributions Create account Log in

Article Talk Read View source View history Search Wikipedia

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia
Wikipedia store

Interaction
Help
About Wikipedia
Community portal
Recent changes
Contact page

Tools
What links here
Related changes
Upload file
Special pages
Permanent link
Page information
Wikidata item
Cite this page

In other projects
Wikimedia Commons

Print/export
Download as PDF
Printable version

Languages
العربية
Español
Français

The dress

From Wikipedia, the free encyclopedia

For other uses, see The Dress.

The dress is a photograph that became a viral internet sensation on 26 February 2015, when viewers disagreed over whether the dress pictured was coloured black and royal blue, or white and gold. The phenomenon revealed differences in human colour perception, which have been the subject of ongoing scientific investigations into neuroscience and vision science, with a number of papers published in peer-reviewed science journals.

The photo originated from a washed-out colour photograph of a dress posted on the social networking service Tumblr. Within the first week after the surfacing of the image, more than 10 million tweets mentioned the dress, using hashtags such as #thedress, #whiteandgold, and #blackandblue. Although the colour of the actual dress was eventually confirmed as blue and black,[3][4] the image prompted many discussions, with users debating their opinions on the colour and how they perceived the dress in the photograph as a certain colour. Members of the scientific community began to investigate the photo for fresh insights into human colour vision.

The dress

The original *The dress* picture

Designer	Roman Originals[1]
Material	Lace[2]

The dress itself, which was identified as a product of the retailer Roman Originals, experienced a major surge in sales as a result of the incident. The retailer also produced a one-off version of the dress in white and gold as a charity campaign.

Contents [hide]

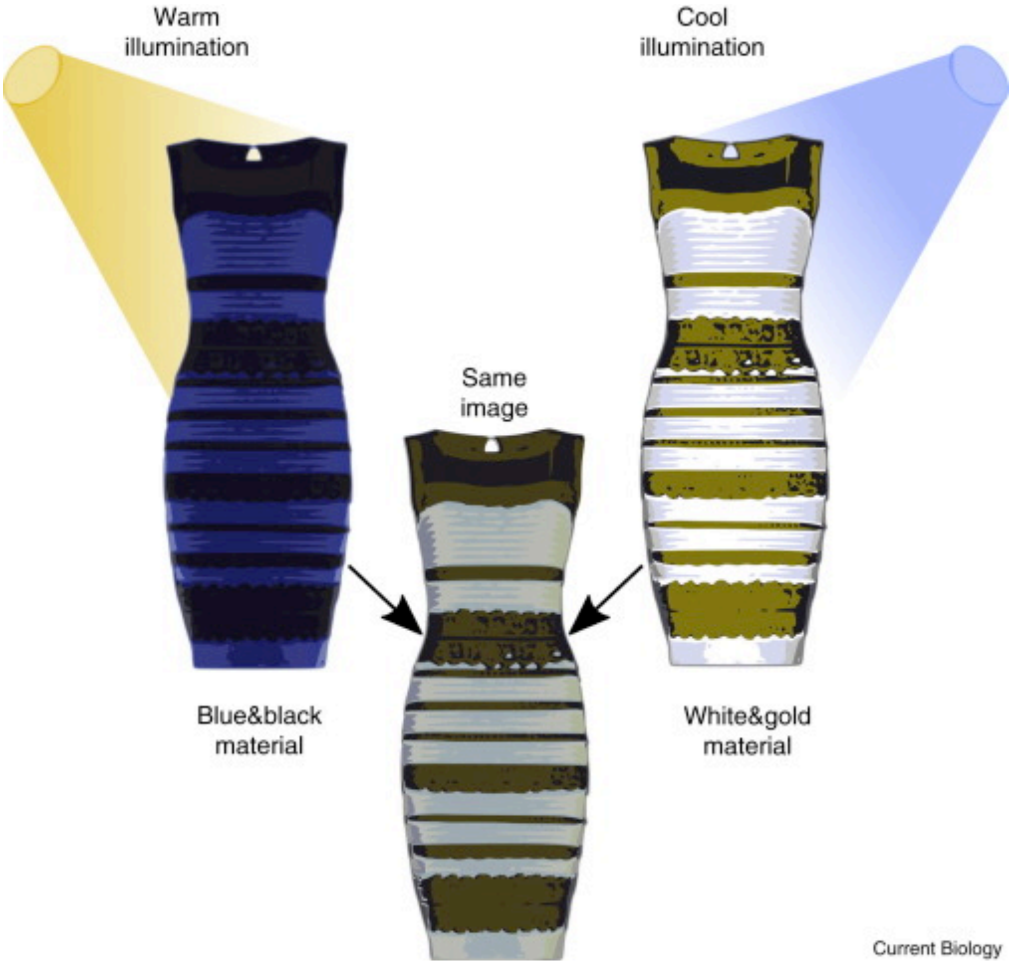

Perception is Solving an Under-Constrained Problem

- Are these spheres, or just flat discs painted with varying albedo?
  - Probably spheres…
  - The flat disk trick would only work from one perspective, so it seems pretty unlikely…

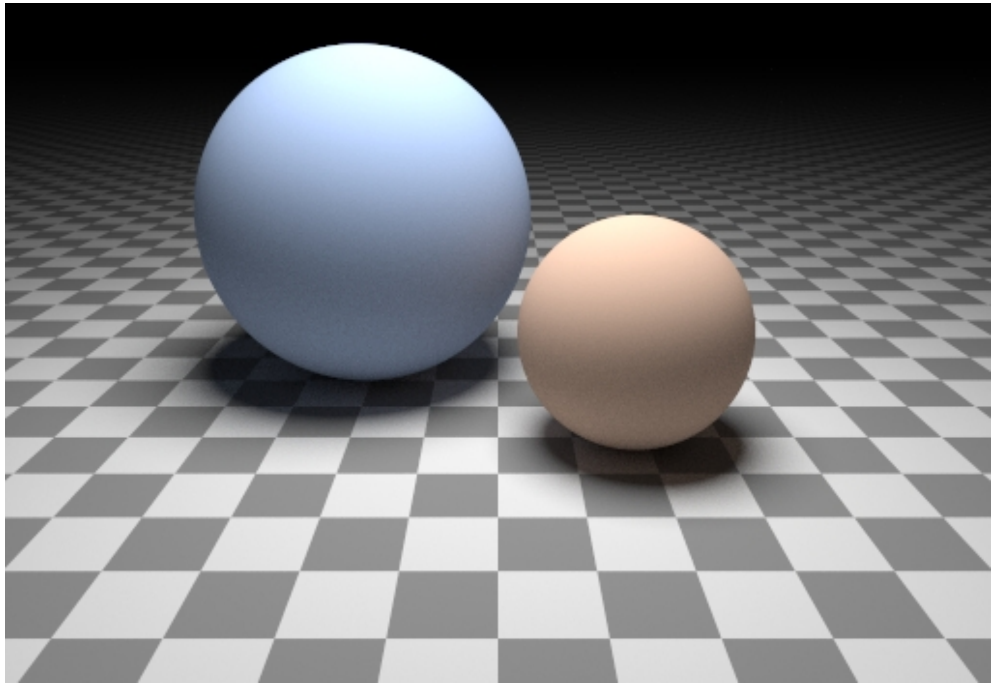

What we know: Stereo

Key Idea: use feature motion to understand shape

Next: Photometric Stereo

Key Idea: use pixel brightness to understand shape

Next: Photometric Stereo

Key Idea: use pixel brightness to understand shape

Photometric Stereo

What results can you get?

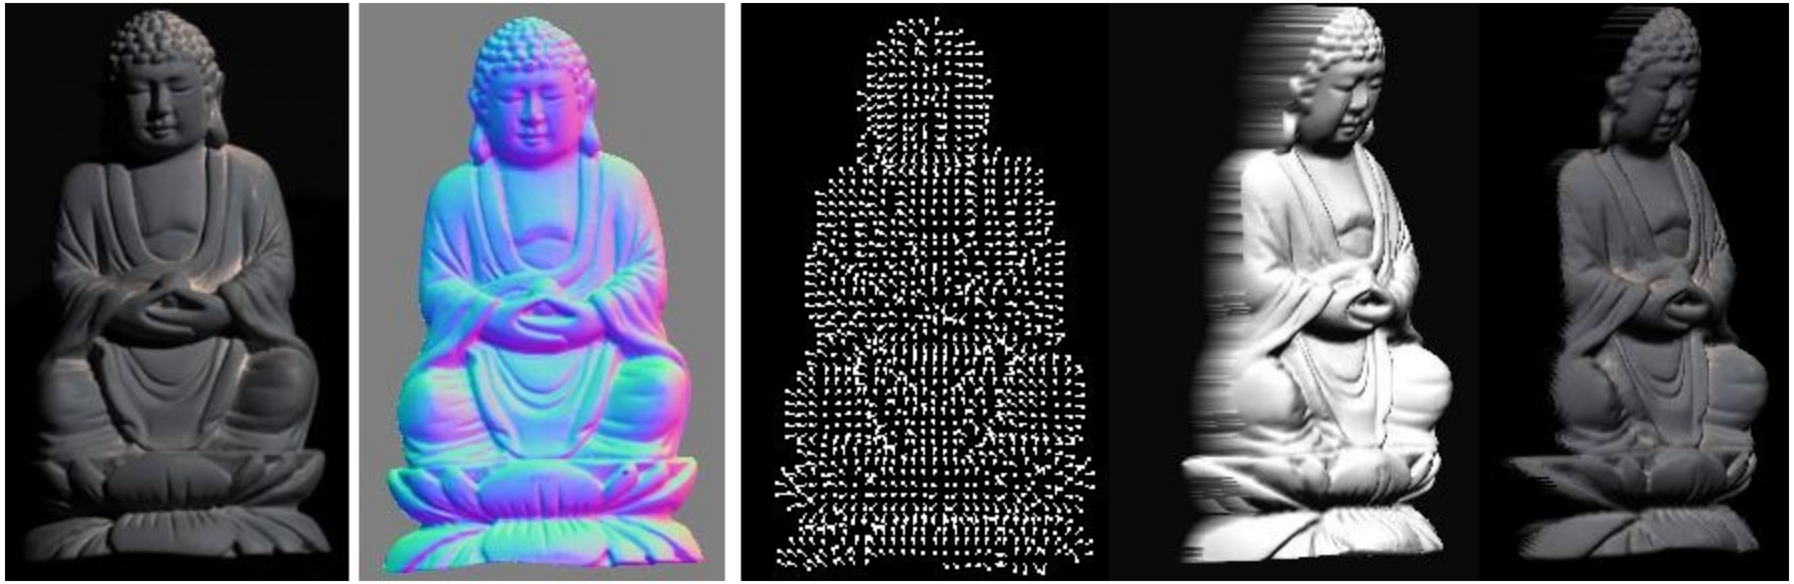

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Radiometry

What determines the brightness of an *image pixel*?

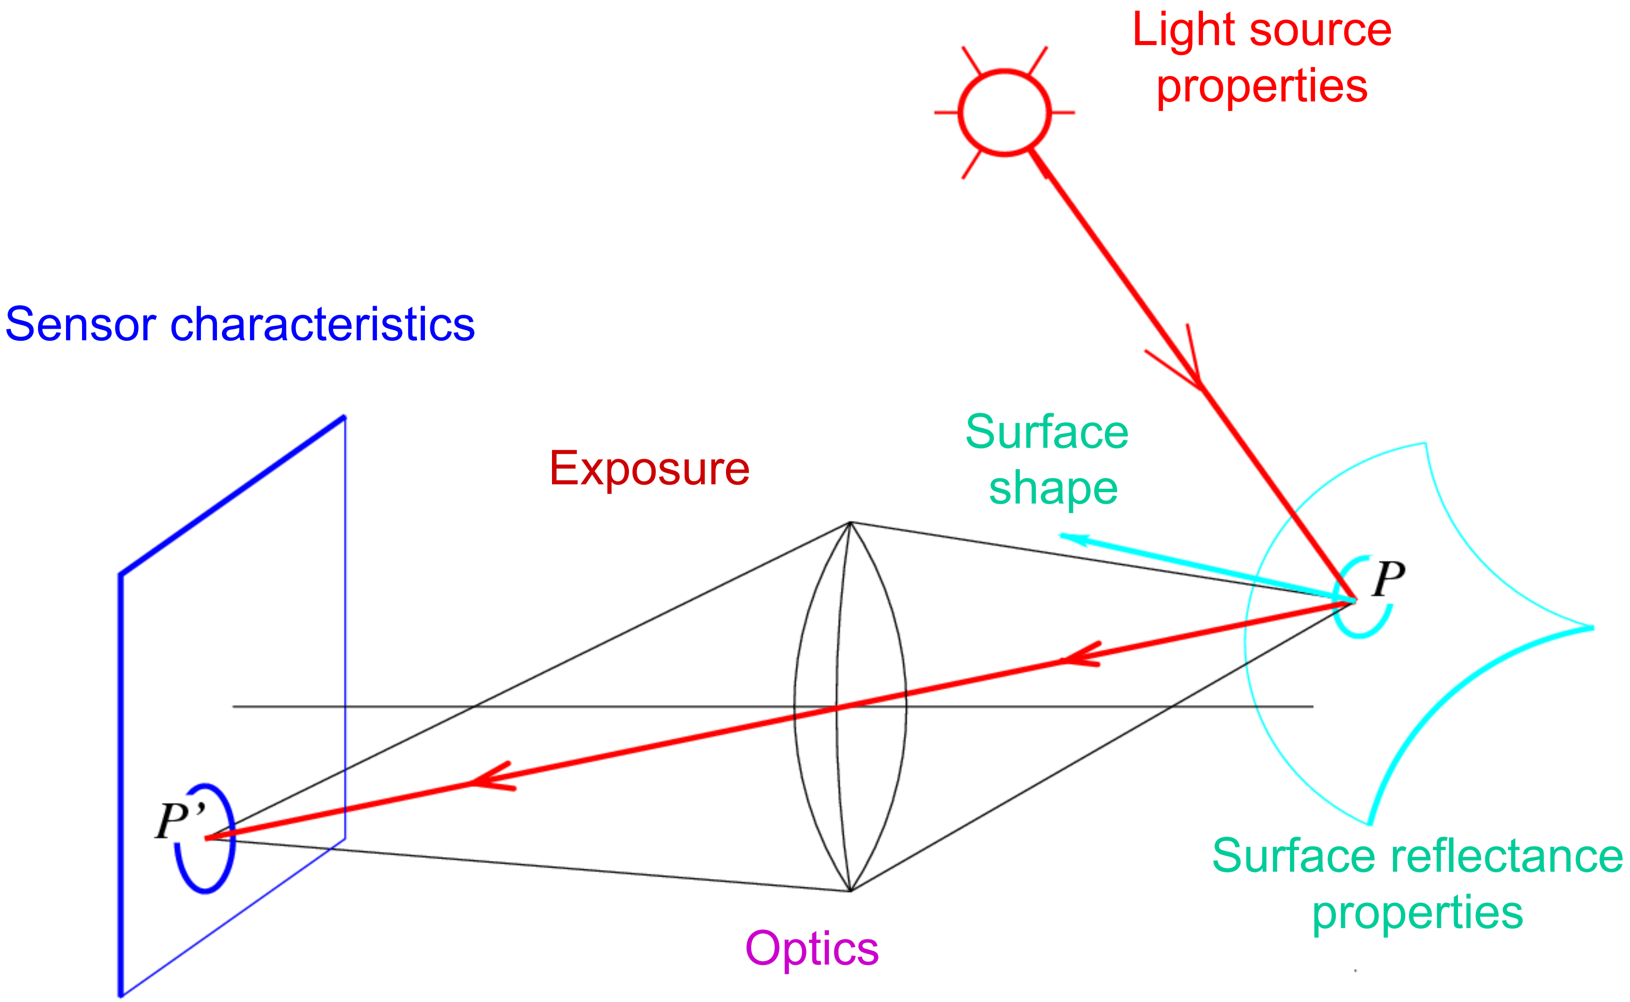

Slide by L. Fei-Fei

What is light?

Electromagnetic radiation (EMR) moving along rays in space

- R(λ) is EMR, measured in units of power (watts)
  - λ is wavelength

Light field

- We can describe all of the light in the scene by specifying the radiation (or **“radiance”** along all light rays) arriving at every point in space and from every direction

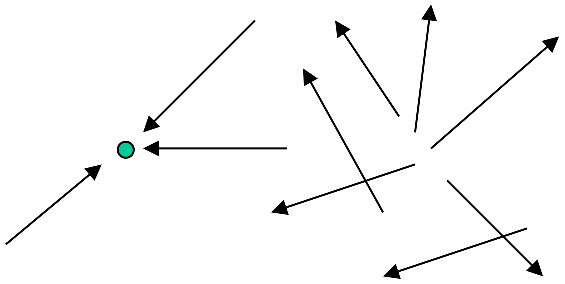

$$R(X, Y, Z, \theta, \phi, \lambda, t)$$

Color perception

Electromagnetic radiation (EMR) moving along rays in space

- R(λ) is EMR, measured in units of power (watts)
  - λ is wavelength

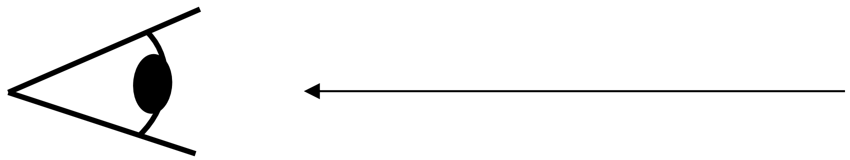

Perceiving light

- How do we convert radiation into “color”?
- What part of the spectrum do we see?

Visible light

We "see" electromagnetic radiation in a range of wavelengths

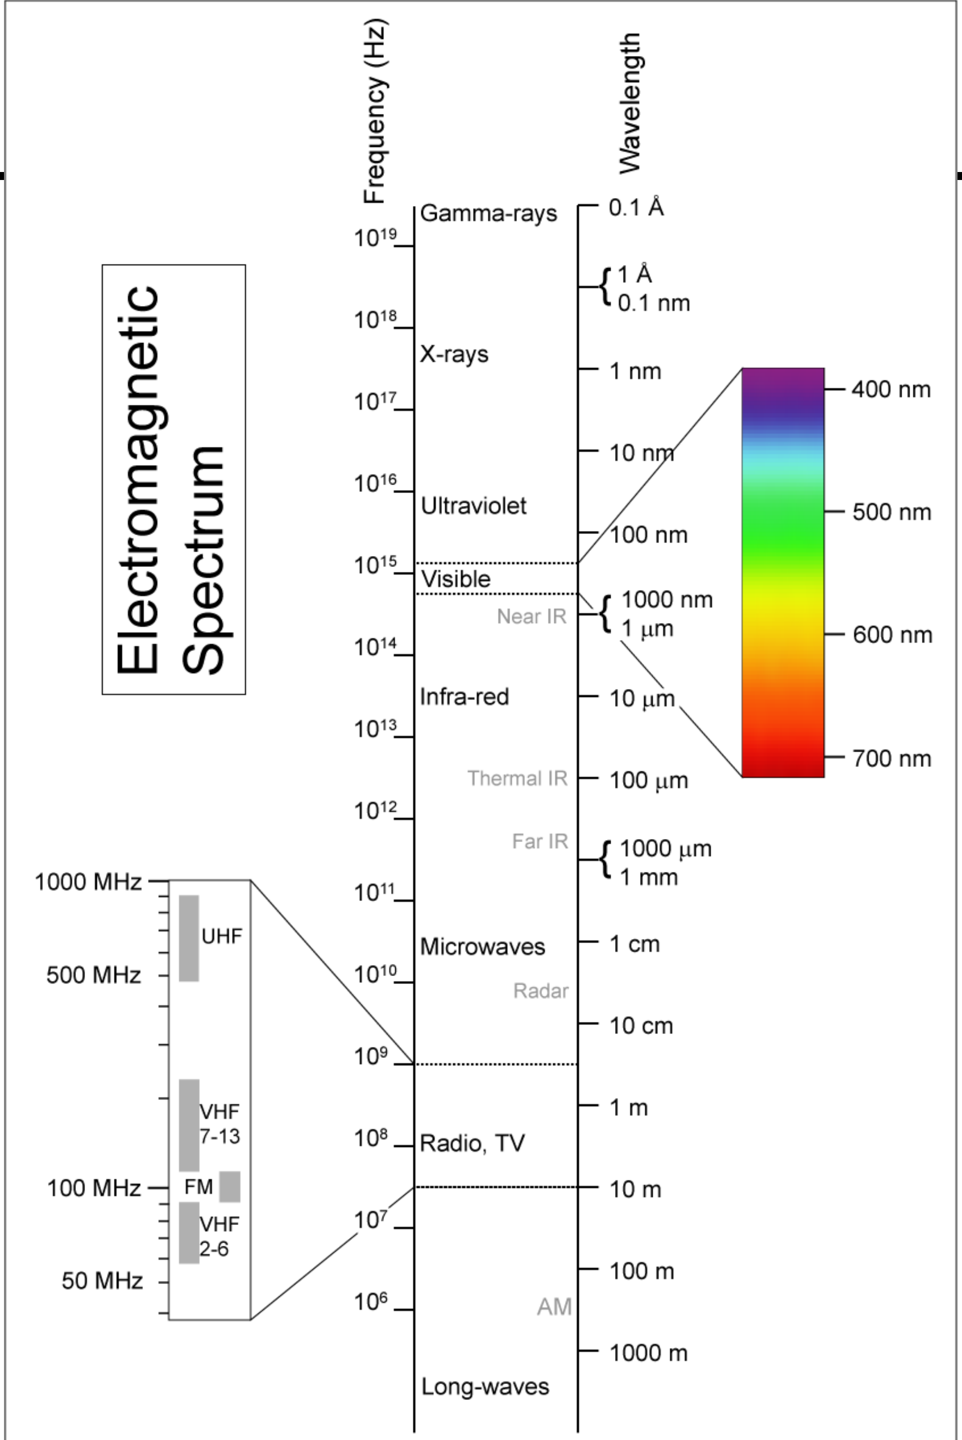

Light spectrum

The appearance of light depends on its power **spectrum**

- How much power (or energy) at each wavelength

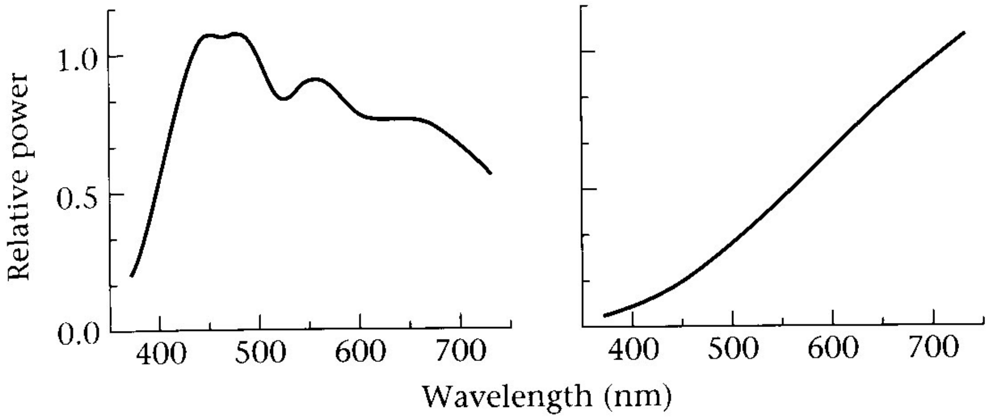

daylight

tungsten bulb

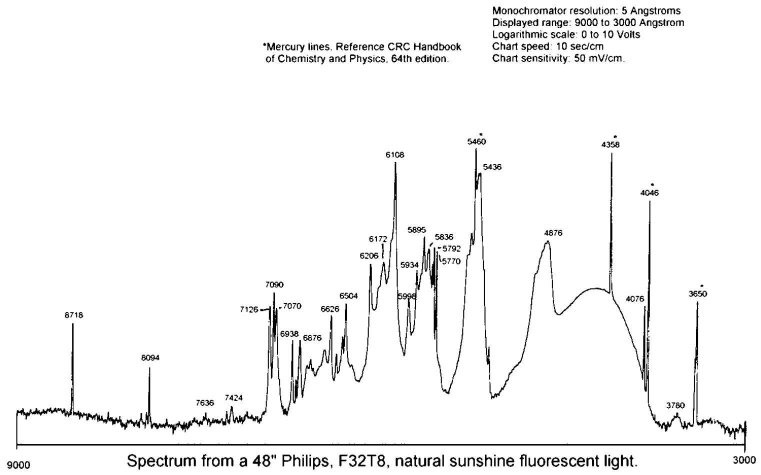

fluorescent bulb

Our visual system converts a light spectrum into “color”

- This is a rather complex transformation

The human visual system

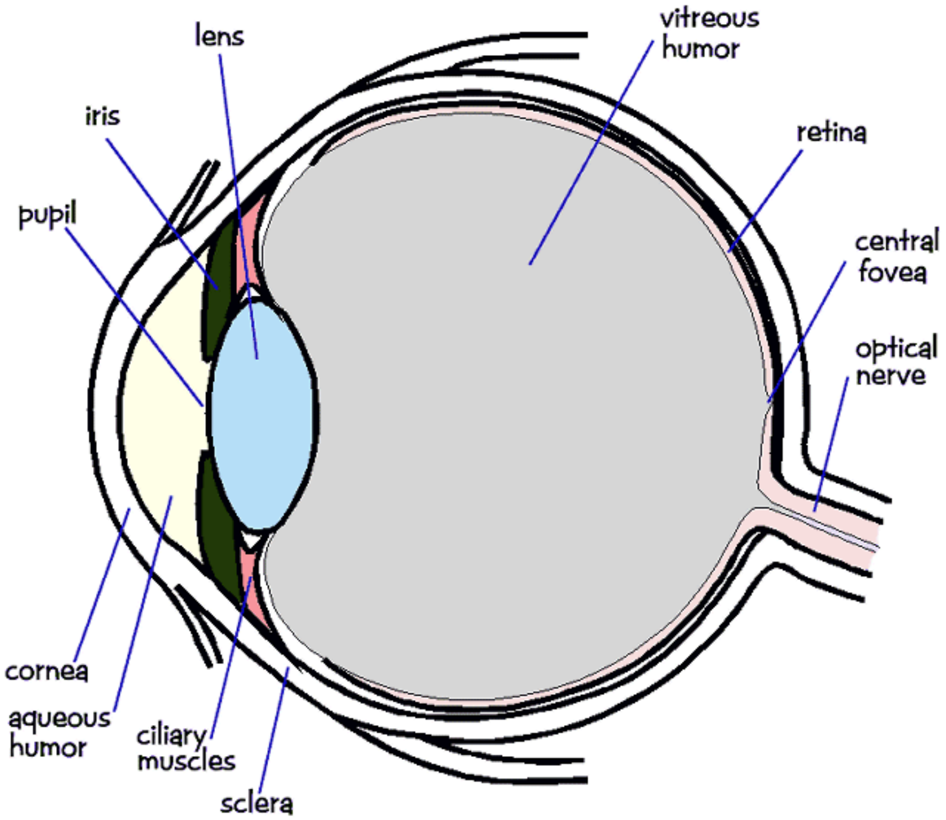

Color perception

- Light hits the retina, which contains photosensitive cells
  - rods and cones
- These cells convert the spectrum into a few discrete values

Density of rods and cones

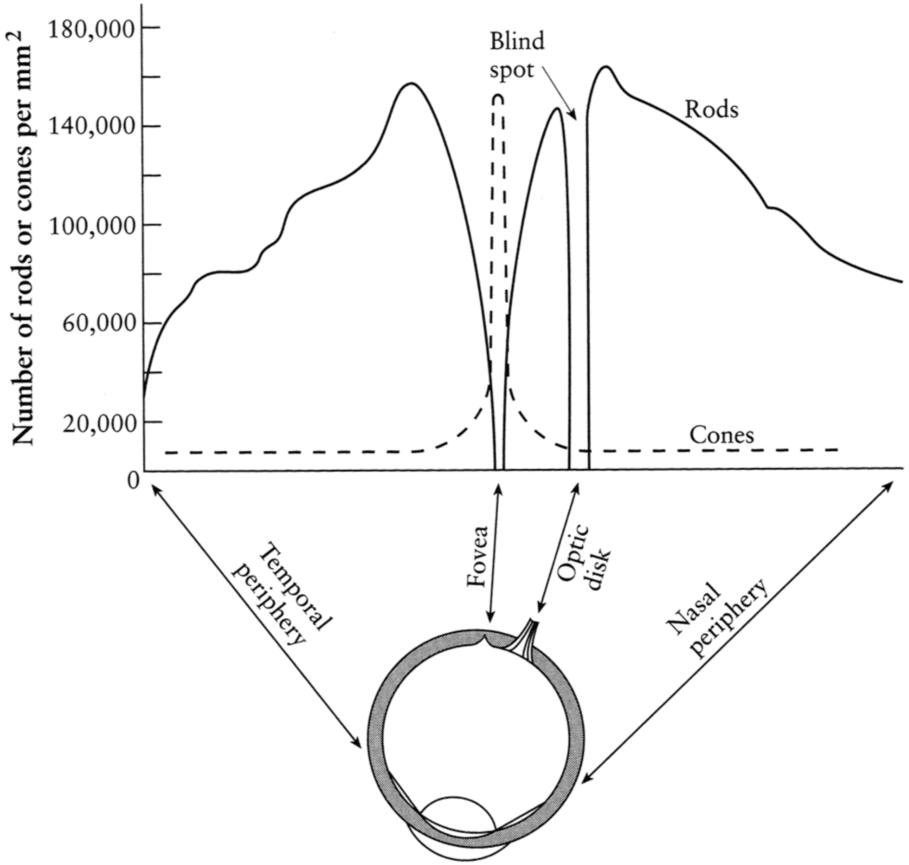

Rods and cones are *non-uniformly* distributed on the retina

- Rods responsible for intensity, cones responsible for color
- **Fovea** - Small region (1 or 2°) at the center of the visual field containing the highest density of cones (and no rods).
- Less visual acuity in the periphery—many rods wired to the same neuron

Demonstrations of visual acuity

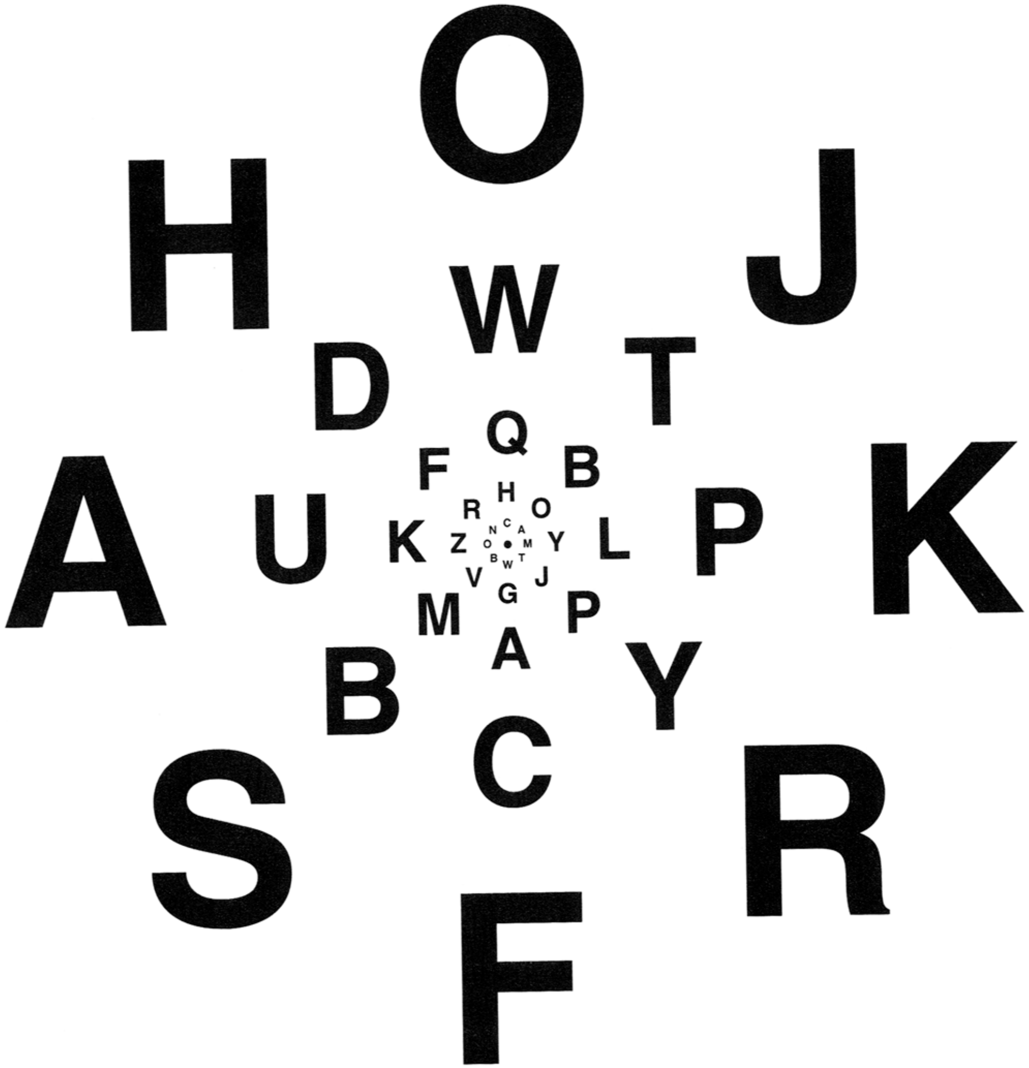

With one eye shut, at the right distance, all of these letters should appear equally legible (Glassner, 1.7).

Demonstrations of visual acuity

With left eye shut, look at the cross on the left. At the right distance, the circle on the right should disappear (Glassner, 1.8).

Brightness contrast and constancy

The apparent brightness depends on the surrounding region

- **brightness contrast**: a constant colored region seems lighter or darker depending on the surrounding intensity:

  – http://www.sandlotscience.com/Contrast/Checker_Board_2.htm

- **brightness constancy**: a surface looks the same under widely varying lighting conditions.

Light response is nonlinear

Our visual system has a large *dynamic range*

- We can resolve both light and dark things at the same time
- One mechanism for achieving this is that we sense light intensity on a *logarithmic scale*
  - an exponential intensity ramp will be seen as a linear ramp
- Another mechanism is *adaptation*
  - rods and cones adapt to be more sensitive in low light, less sensitive in bright light.

Visual dynamic range

Background	Luminance (candelas per square meter)
Horizon sky	
Moonless overcast night	0.00003
Moonless clear night	0.0003
Moonlit overcast night	0.003
Moonlit clear night	0.03
Deep twilight	0.3
Twilight	3
Very dark day	30
Overcast day	300
Clear day	3,000
Day with sunlit clouds	30,000
Daylight fog	
Dull	300–1,000
Typical	1,000–3,000
Bright	3,000–16,000
Ground	
Overcast day	30–100
Sunny day	300
Snow in full sunlight	16,000

FIGURE 1.13
Luminance of everyday backgrounds. *Source:* Data from Rea, ed., *Lighting Handbook 1984 Reference and Application*, fig. 3-44, p. 3-24.

A piece of white paper can be 1,000,000,000 times brighter in outdoor sunlight than in a moonless night.

BUT in a given lighting condition, light ranges over only about two orders of magnitude.

Learning to See in the Dark

Chen Chen, Qifeng Chen, Jia Xu and Vladlen Koltun

IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018

(a) Camera output with ISO 8,000 (b) Camera output with ISO 409,600 (c) Our result from the raw data of (a)

Figure. Extreme low-light imaging by a Sony a7S II camera using ISO 8000, f/5.6, 1/30 second. Dark indoor environment. The illuminance at the camera is <0.1 lux.

http://cchen156.web.engr.illinois.edu/SID.html

Visual dynamic range

If we were sensitive to this whole range all the time, we wouldn't be able to discriminate lightness levels in a typical scene.

The visual system solves this problem by restricting the 'dynamic range' of its response to match the current overall or 'ambient' light level.

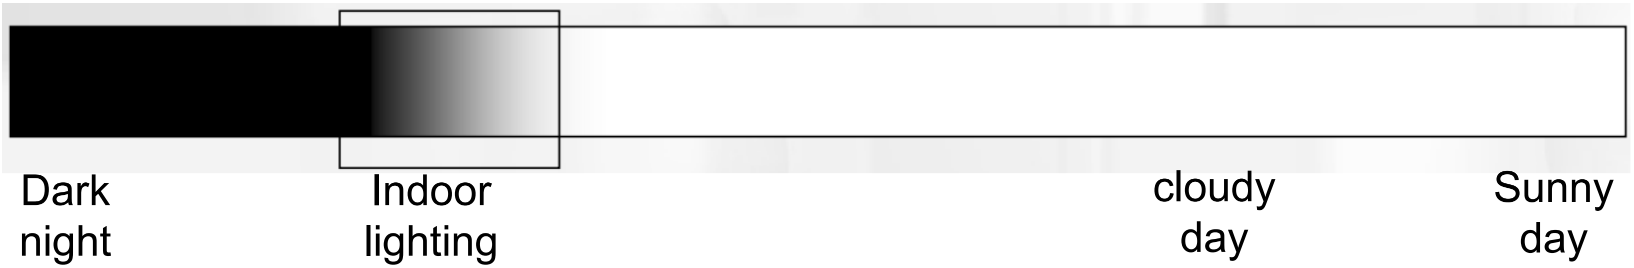

Color perception

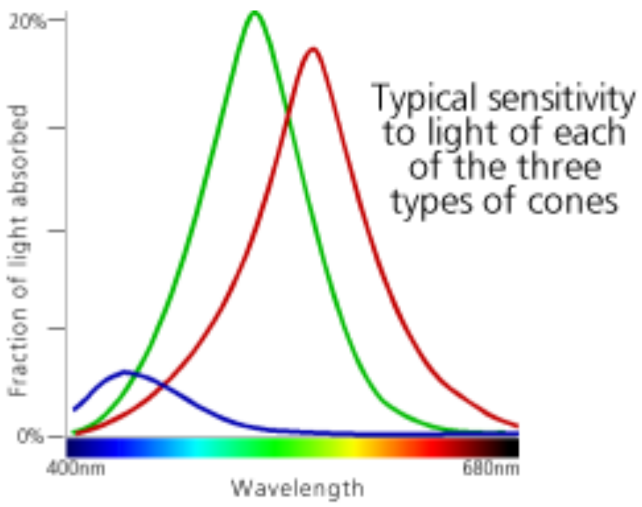

L response curve

Three types of cones

- Each is sensitive in a different region of the spectrum
  - but regions overlap
  - Short (S) corresponds to blue
  - Medium (M) corresponds to green
  - Long (L) corresponds to red
- Different sensitivities: we are more sensitive to green than red
  - varies from person to person (and with age)
- Colorblindness—deficiency in at least one type of cone

Color perception

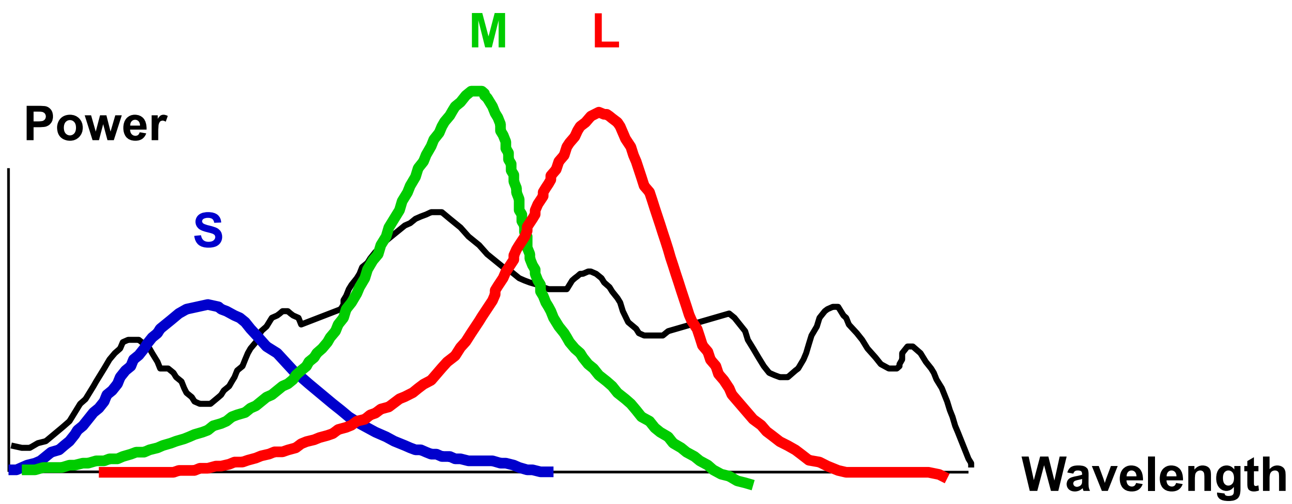

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number
- Q: How can we represent an entire spectrum with 3 numbers?
- A: We can't! Most of the information is lost.
  - As a result, two different spectra may appear indistinguishable
    - » such spectra are known as **metamers**
    - » http://www.cs.brown.edu/exploratories/freeSoftware/repository/edu/brown/cs/exploratories/applets/spectrum/metamers_guide.html

What kind of bulb is it?

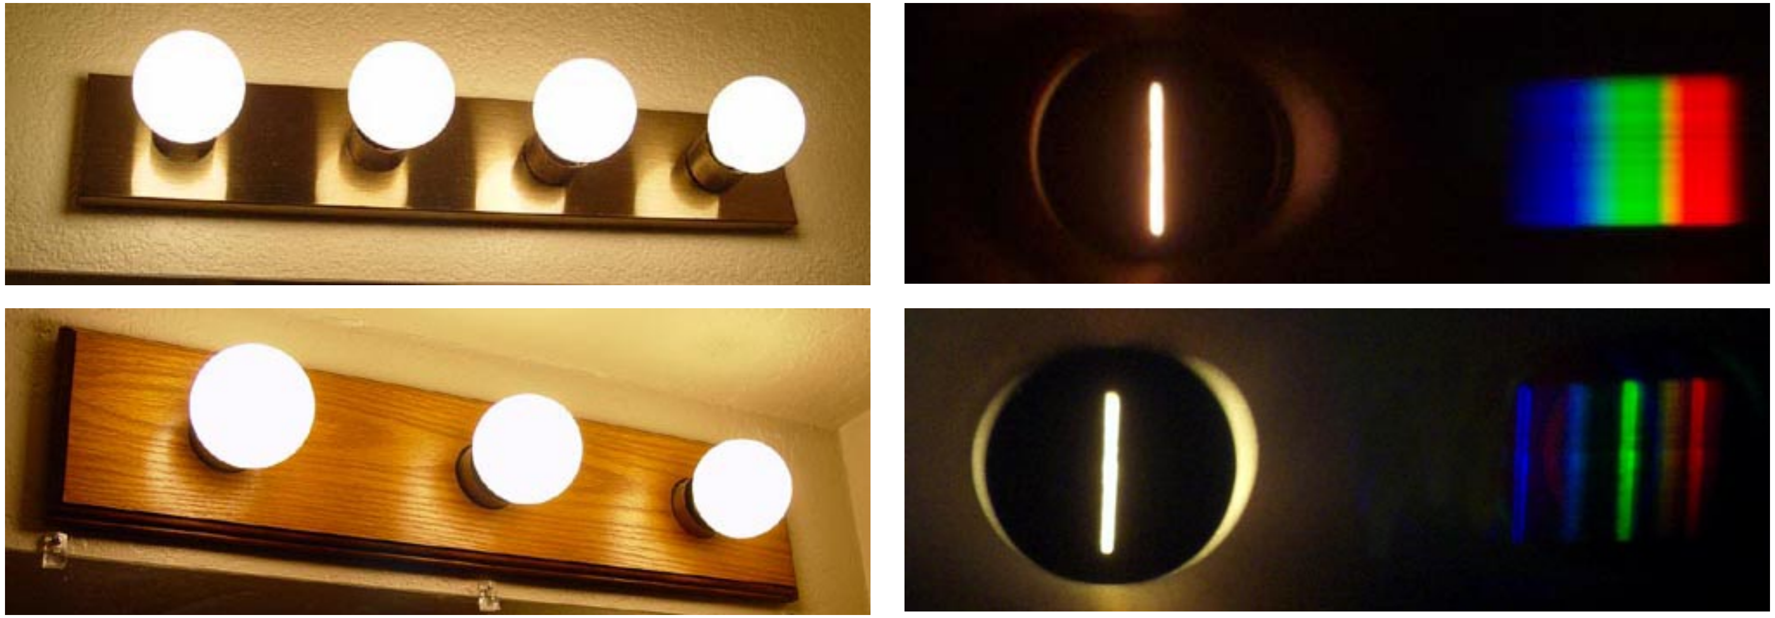

http://www.chemistryland.com/CHM107Lab/Exp7/Spectroscope/Spectroscope.html

Perception summary

The mapping from radiance to perceived color is quite complex!

- We throw away most of the data
- We apply a logarithm
- Brightness affected by pupil size
- Brightness contrast and constancy effects

The same is true for cameras

- But we have tools to correct for these effects
  - (Computational Photography)

Light transport

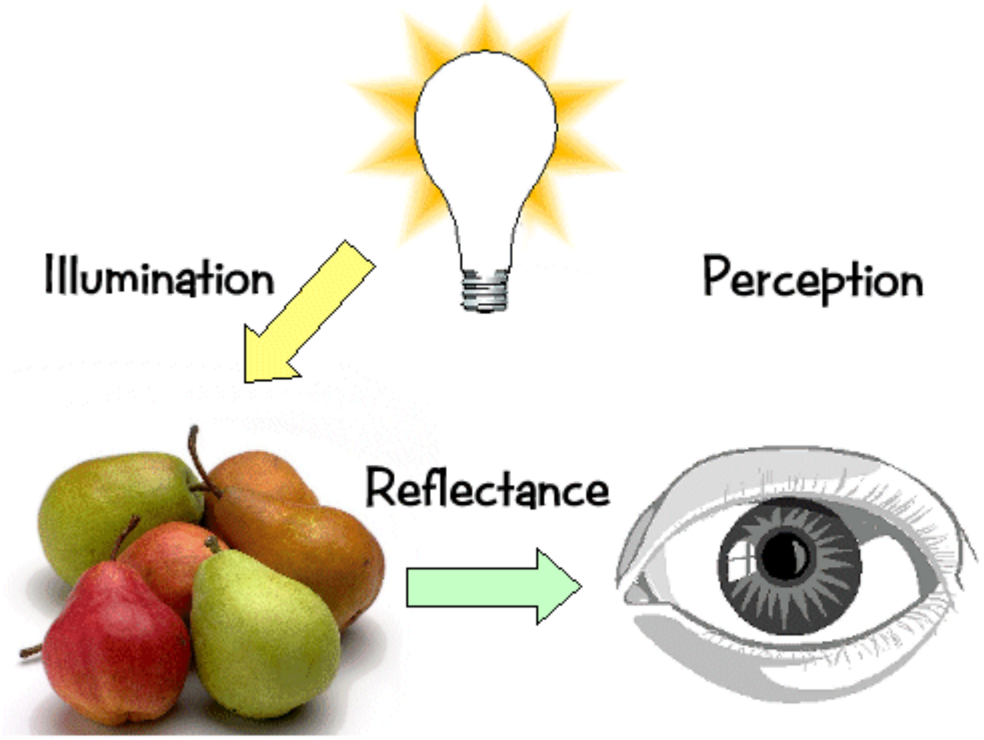

Light sources

Basic types

- point source
- directional source
  - a point source that is infinitely far away
- area source
  - a union of point sources

More generally

- a light field can describe *any* distribution of light sources

What happens when light hits an object?

Modeling Image Formation

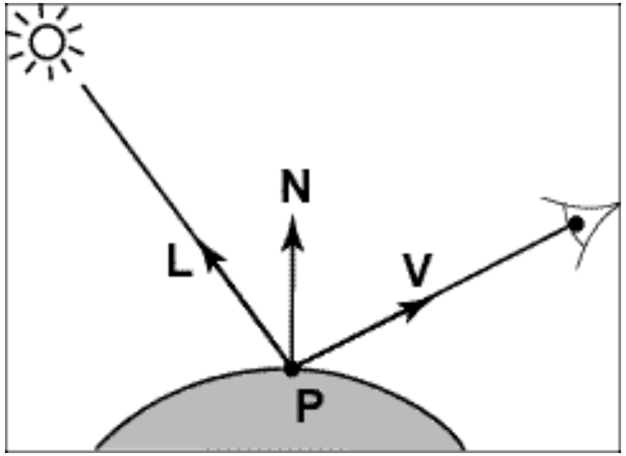

We need to reason about:

- How light interacts with the scene
- How a pixel value is related to light energy in the world

Track a “ray” of light all the way from light source to the sensor

Directional Lighting

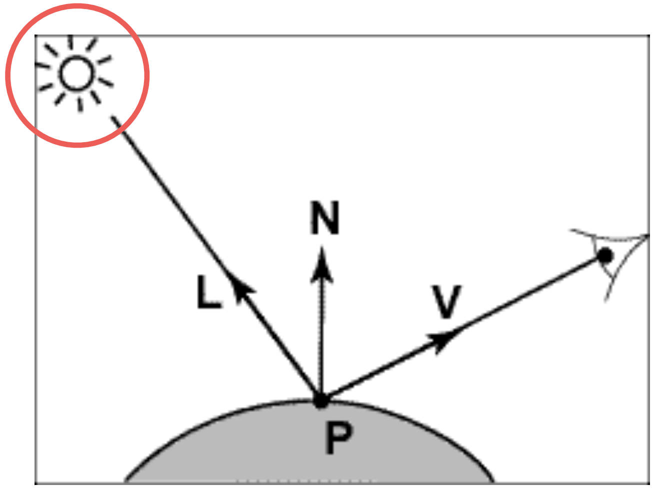

- Key property: all rays are parallel
- Equivalent to an infinitely distant point source

Lambertian Reflectance

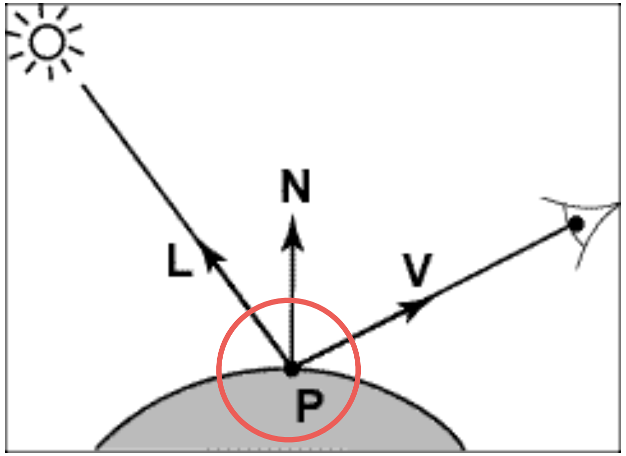

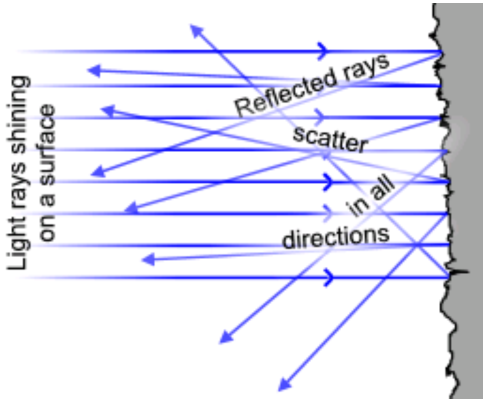

$$I = N \cdot L$$

Image intensity = Surface normal • Light direction

Image intensity $\propto$ cos(angle between N and L)

Materials - Three Forms

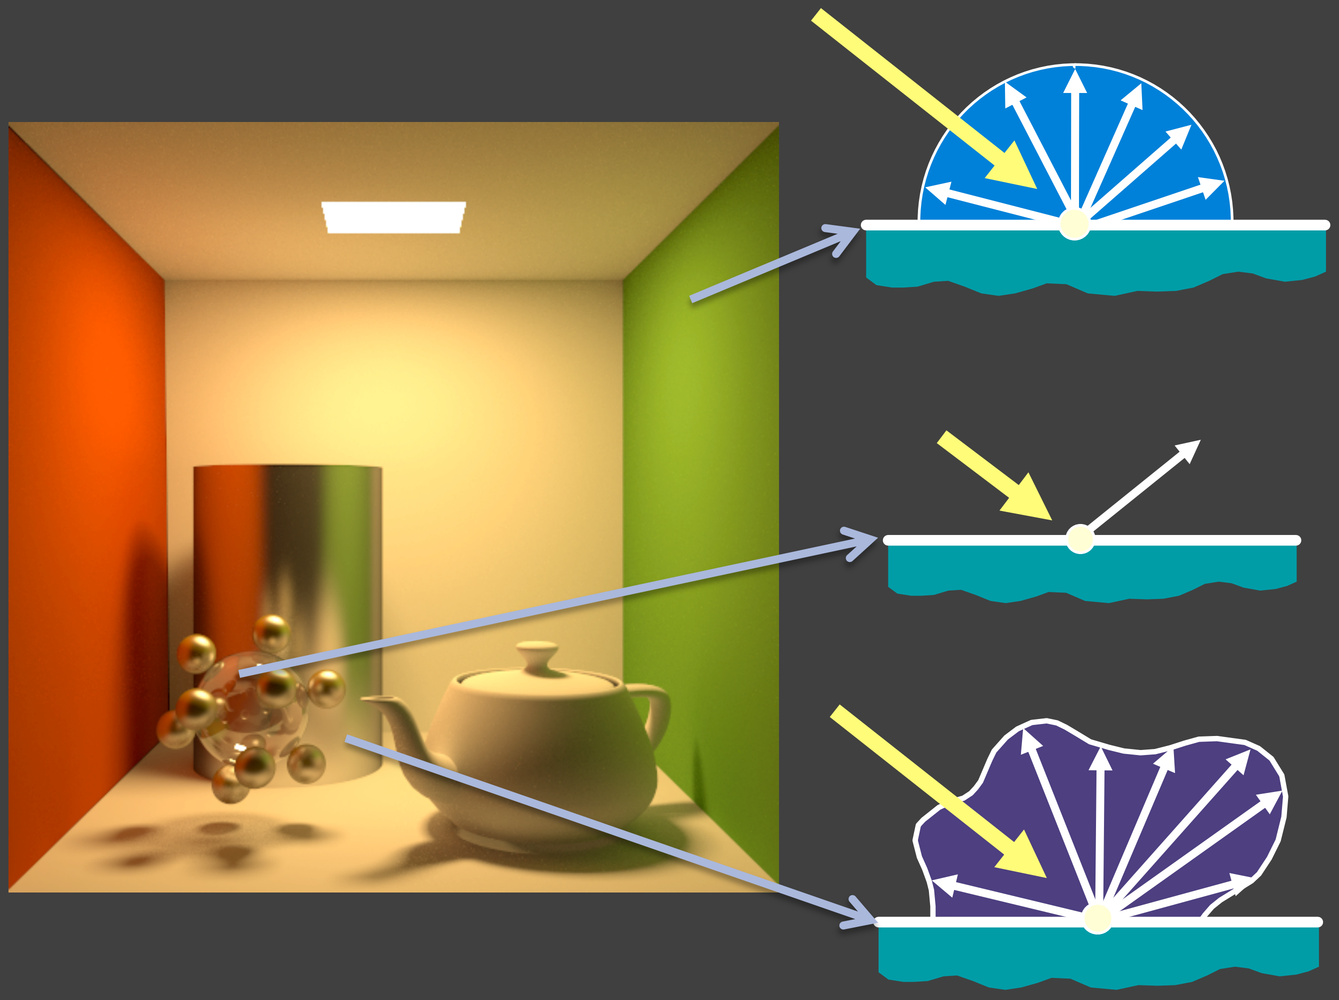

Ideal diffuse (Lambertian)

Ideal specular

Directional diffuse

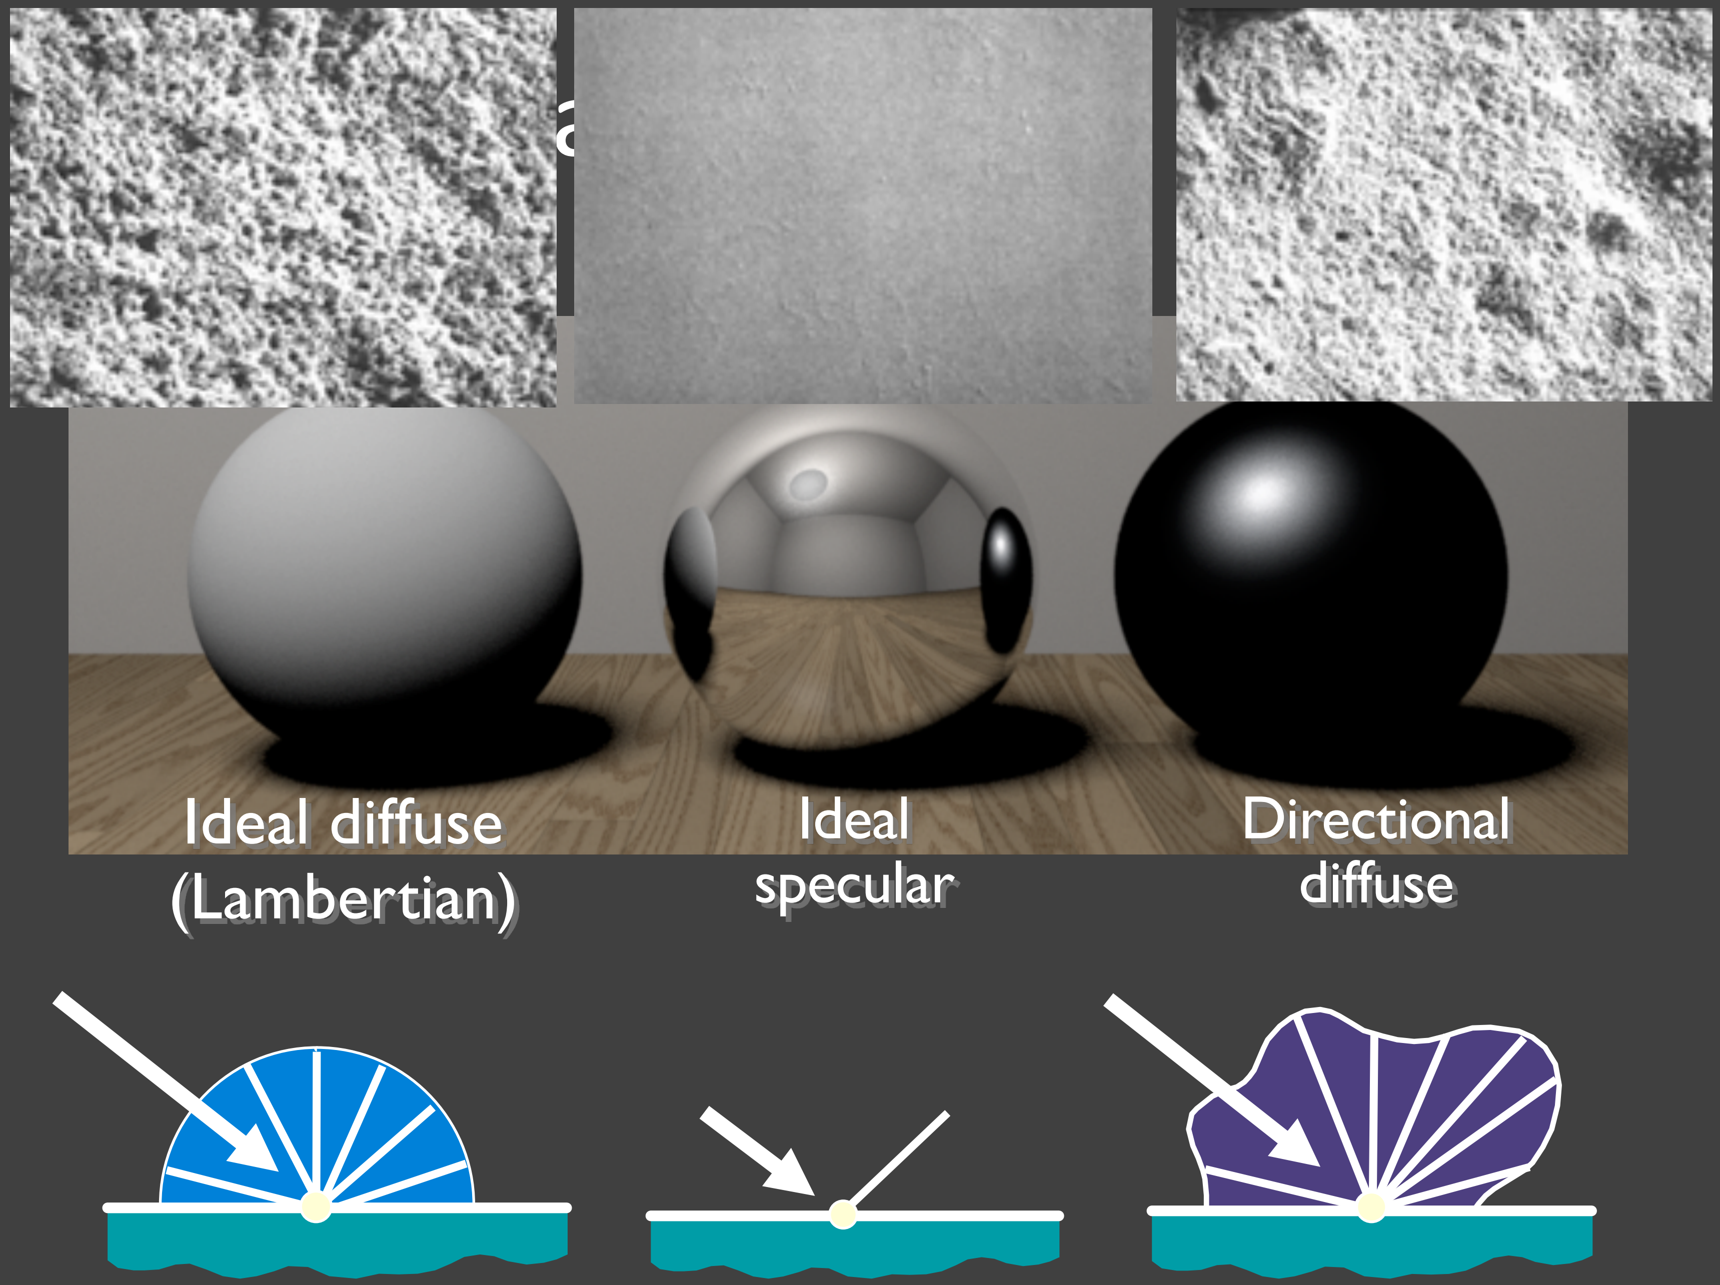
Ideal diffuse (Lambertian)
Ideal specular
Directional diffuse

Ideal Diffuse Reflection

- Characteristic of multiple scattering materials
- An idealization but reasonable for matte surfaces

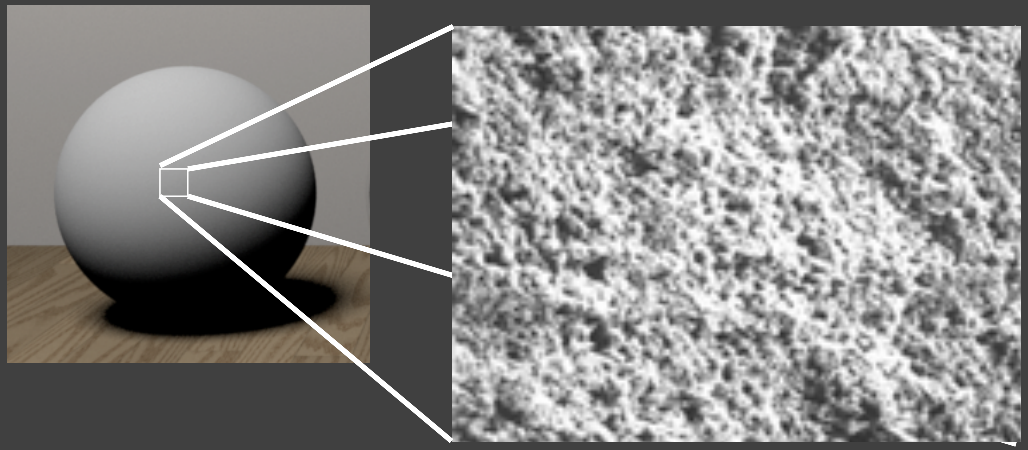

Lambertian Reflectance

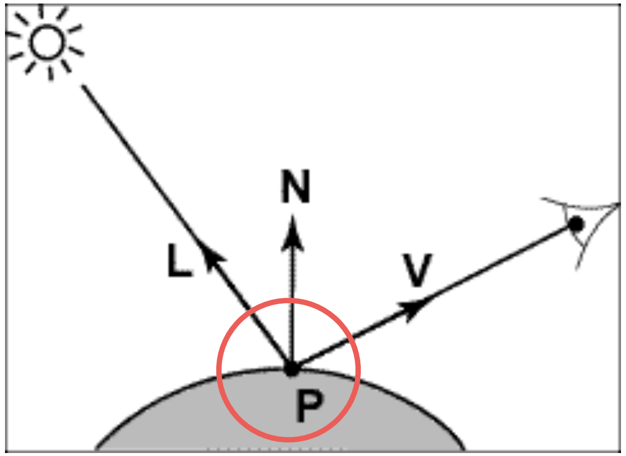

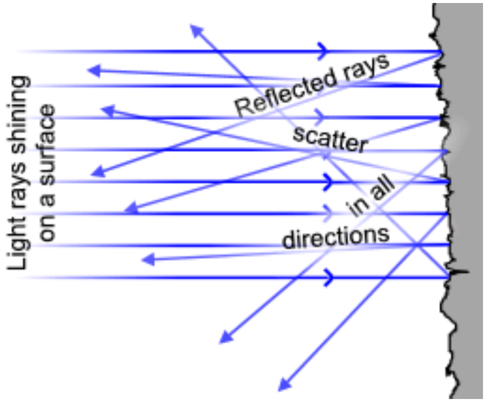

1. Reflected energy is proportional to cosine of angle between L and N **(incoming)**

2. Measured intensity is viewpoint-independent **(outgoing)**

Lambertian Reflectance: Incoming

1. Reflected energy is proportional to cosine of angle between L and N

Lambertian Reflectance: Incoming

1. Reflected energy is proportional to cosine of angle between L and N

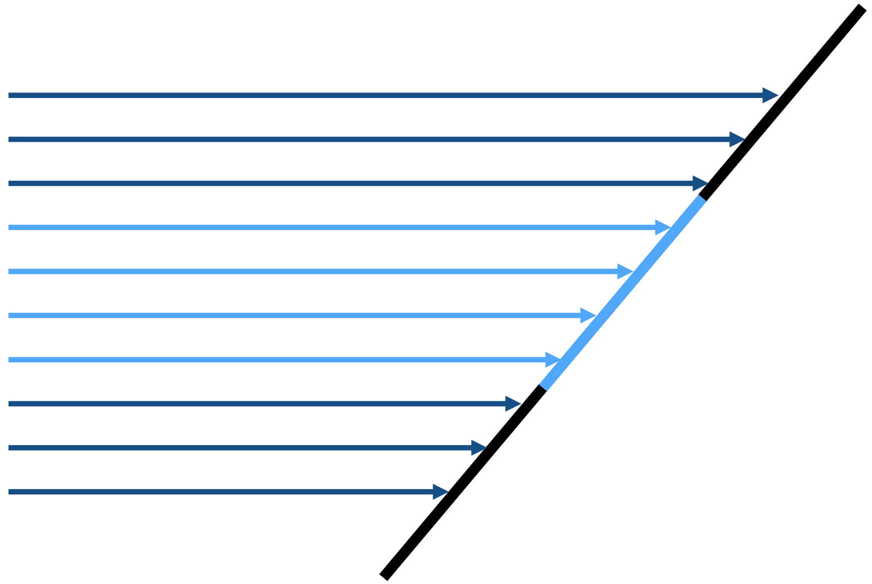

Lambertian Reflectance: Incoming

1. Reflected energy is proportional to cosine of angle between L and N

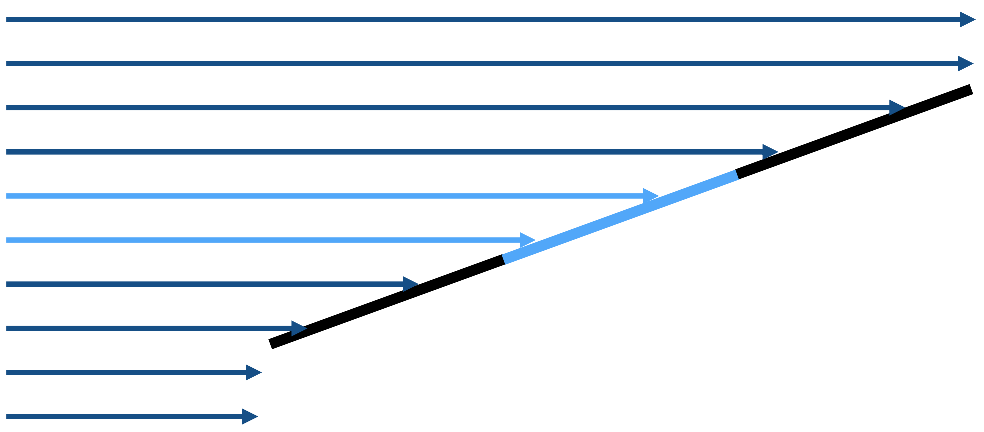

Light hitting surface is proportional to the **cosine**

Lambertian Surfaces: Appearance vs Reflected Photons

- Appearance is the same from every angle

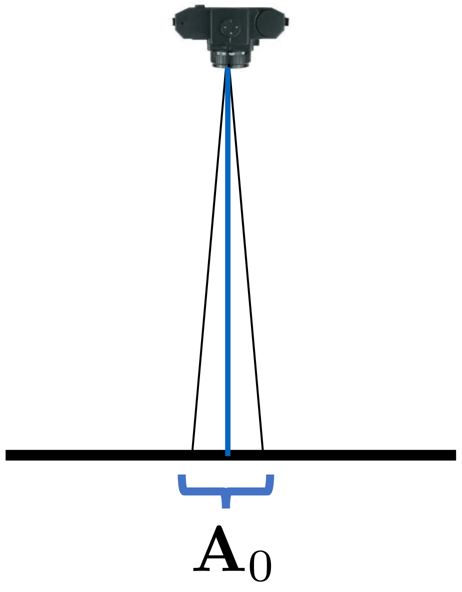

$$Pixel = \mathbf{B}_0\mathbf{A}_0 = \mathbf{B}_\theta\mathbf{A}_\theta$$

- Radiant Intensity? (how many photons reflected per angle)

Note for those following the 2020 videos:

This and the following two slides were changed after the initial recording of the lecture. These slides were then presented in the following lecture on photometric stereo

Lambertian Surfaces: Appearance vs Reflected Photons

- Appearance is the same from every angle
- Radiant Intensity? (how many photons reflected per angle)

$$\mathbf{A}_\theta \propto \mathbf{A}_0 \frac{1}{\cos\theta}$$

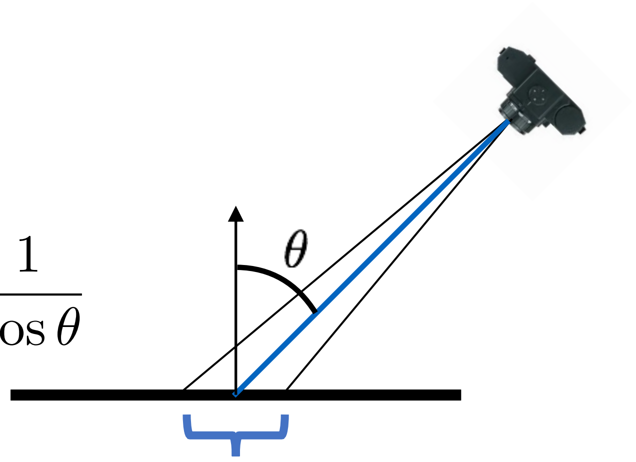

$$Pixel = \mathbf{B}_0\mathbf{A}_0 = \mathbf{B}_\theta\mathbf{A}_\theta$$

Lambertian Surfaces: Appearance vs Reflected Photons

- Appearance is the same from every angle
- Radiant Intensity? (how many photons reflected per angle)

$$\mathbf{A}_\theta \propto \mathbf{A}_0 \frac{1}{\cos\theta}$$

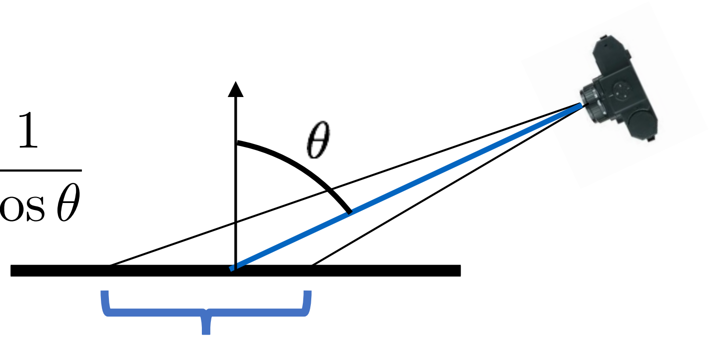

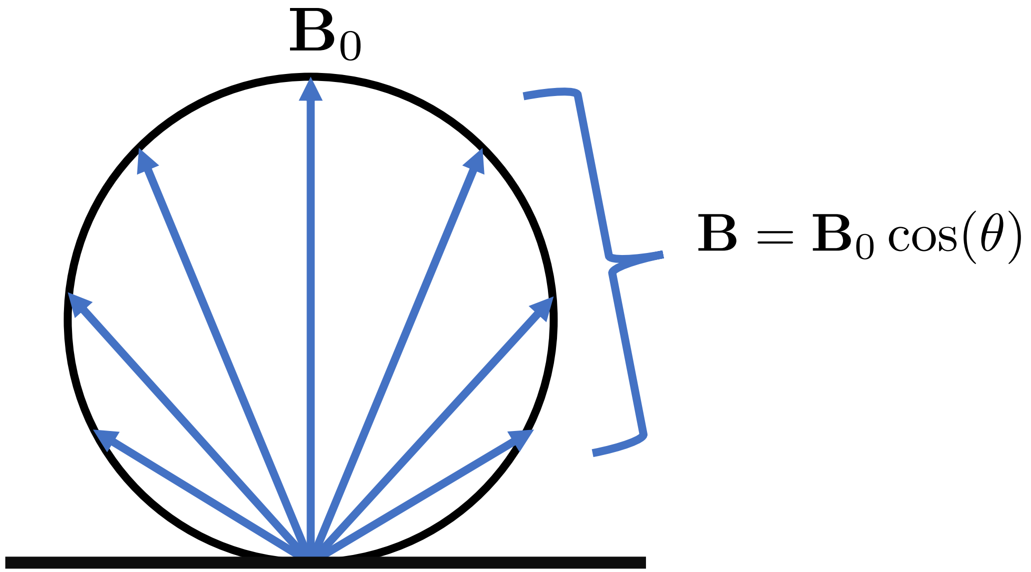

$$Pixel = \mathbf{B}_0\mathbf{A}_0 = \mathbf{B}_\theta\mathbf{A}_\theta$$

Lambert's cosine law: $\mathbf{B} = \mathbf{B}_0 \cos(\theta)$

$$Pixel = \mathbf{B}_0\mathbf{A}_0 \frac{\cos\theta}{\cos\theta}$$

Image Formation Model: Final

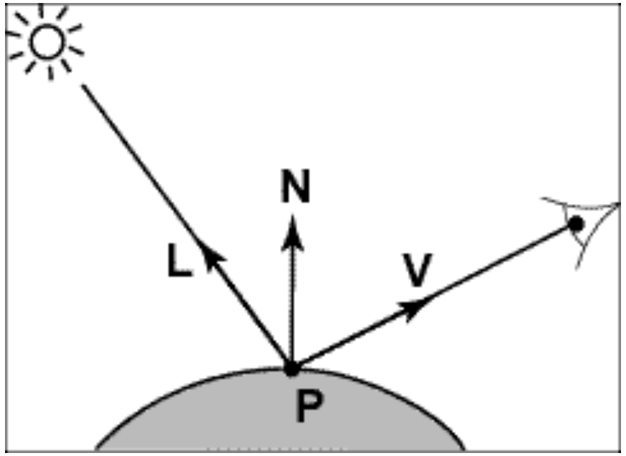

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

1. Diffuse albedo: what fraction of incoming light is reflected?
   - Introduce scale factor k_d
2. Light intensity: how much light is arriving?
   - Compensate with camera exposure (global scale factor)
3. Camera response function
   - Assume pixel value is linearly proportional to incoming energy (perform radiometric calibration if not)

A Single Image: Shape from Shading

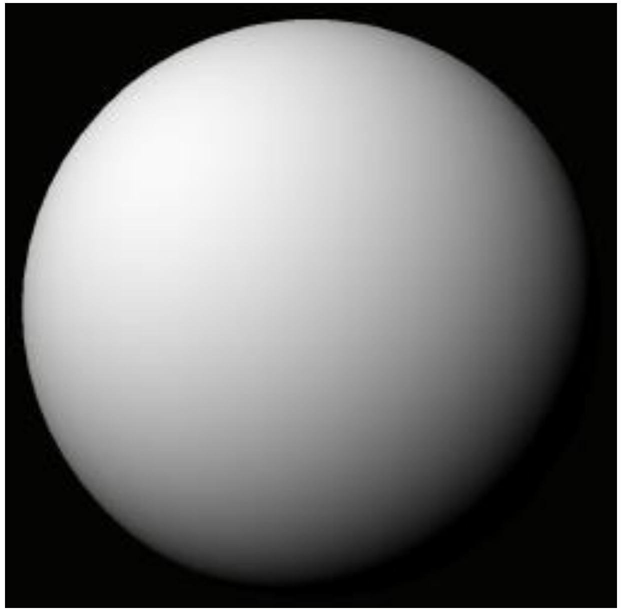

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

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, SFS doesn't work very well:
assumptions are too restrictive,
too much ambiguity in nontrivial scenes.

Application: Detecting composite photos

Fake photo

Real photo

Questions?