Vision and color

Reading

Good resources:

Glassner, *Principles of Digital Image Synthesis*, pp. 5-32.

Palmer, *Vision Science: Photons to Phenomenology*.

Wandell. *Foundations of Vision.*

Lenses

The human eye employs a lens to focus light.

To quantify lens properties, we'll need some terms from *optics* (the study of sight and the behavior of light):

- **Focal point** - the point where parallel rays converge when passing through a lens.
- **Focal length** - the distance from the lens to the focal point.

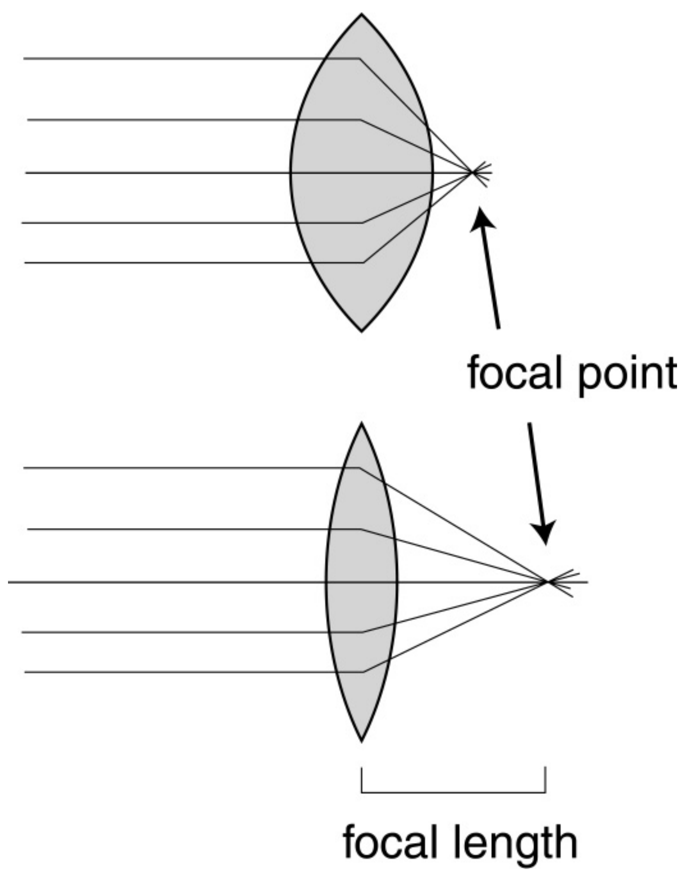

Lenses, cont'd

By tracing rays through a lens, we can generally tell where an object point will be focused to an image point:

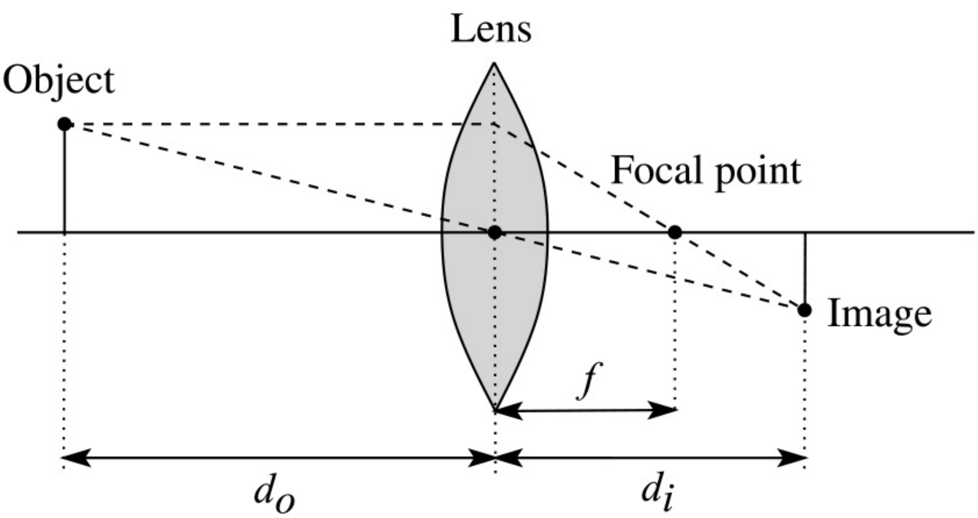

This construction leads to the Gaussian lens formula for simple, thin lenses:

$$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$$

Q: Given these three parameters, how does the human eye keep the world in focus?

Structure of the eye

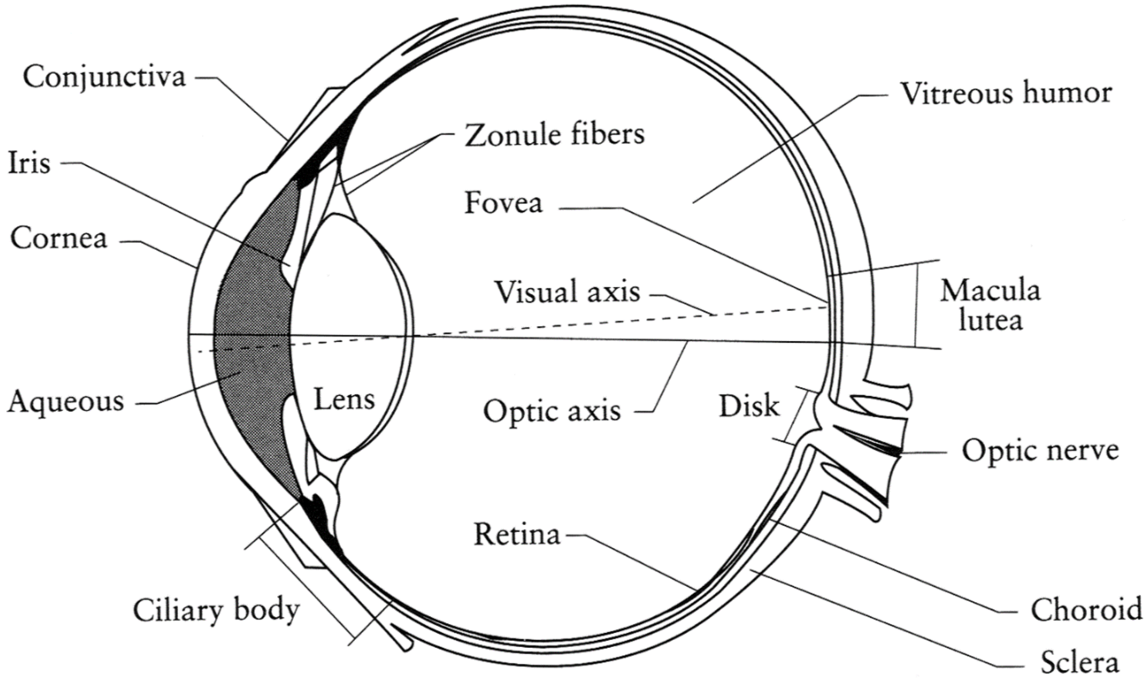

Physiology of the human eye (Glassner, 1.1)

The most important structural elements of the eye are:

- **Cornea** - a clear coating over the front of the eye:
  - Protects eye against physical damage.
  - Provides initial focusing (40D).
- **Crystalline lens** – provides additional focusing
- **Retina** – layer of photosensitive cells lining the back of the eye.

Structure of the eye

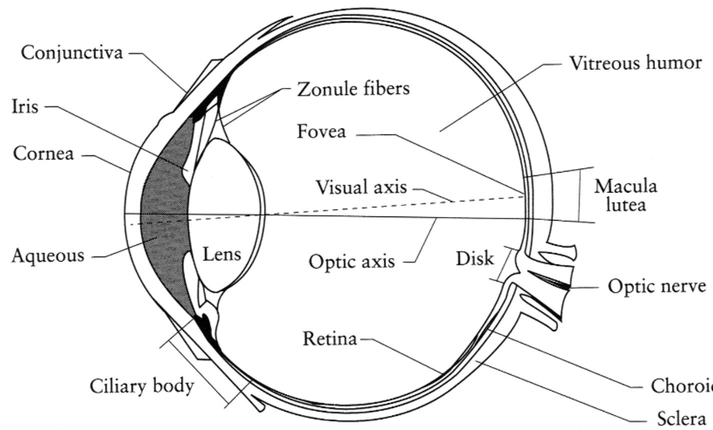

We can treat the cornea + crystalline lens as a compound lens, which roughly follows the Gaussian lens formula:

$$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$$

Q: Given these three parameters, how does the human eye keep the world in focus?

Structure of the eye, cont.

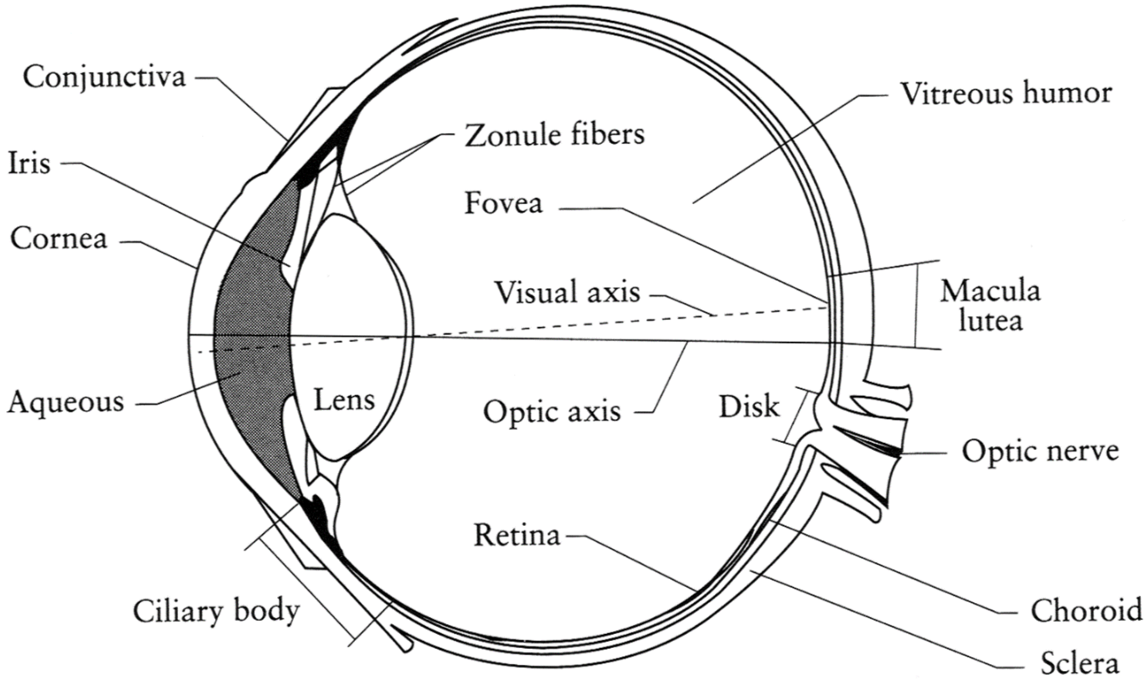

Physiology of the human eye (Glassner, 1.1)

- **Crystalline lens** - controls the focal distance:
  - Power ranges from 10 to 30D in a child.
  - Power and range reduces with age.
- **Ciliary body** - The muscles that compress the sides of the lens, controlling its power.

Q: As an object moves closer, do the ciliary muscles contract or relax to keep the object in focus?

Structure of the eye

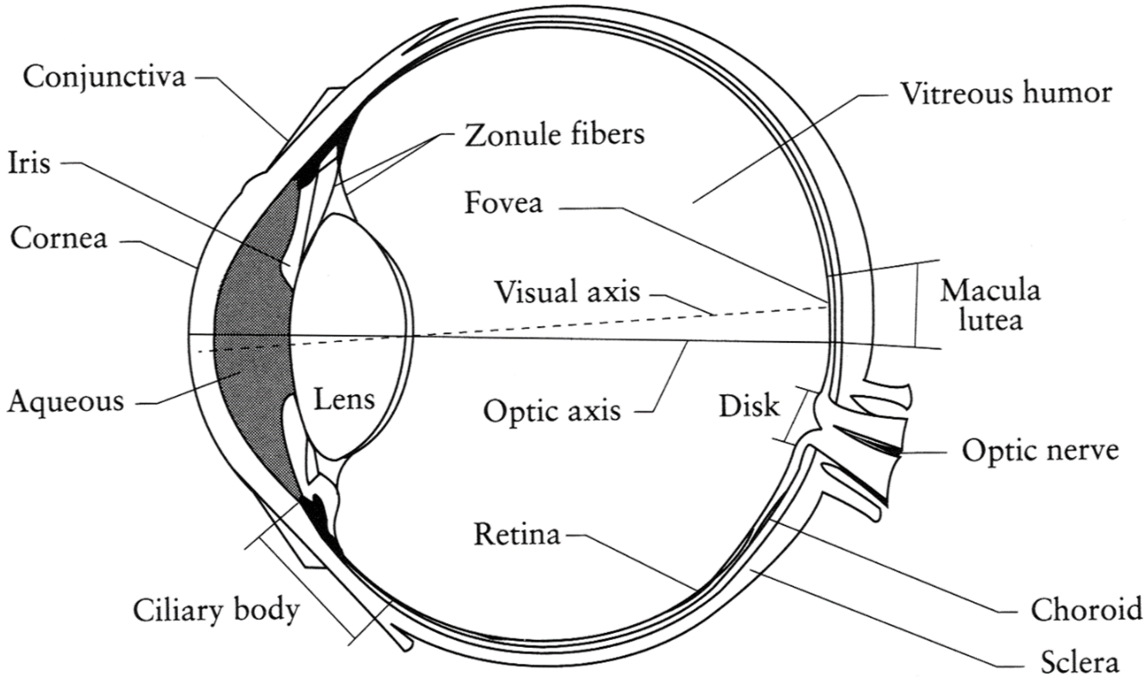

Physiology of the human eye (Glassner, 1.1)

The remaining important elements are:

- **Iris** - Colored annulus with radial muscles.
- **Pupil** - The hole whose size is controlled by the iris.

The iris adjusts the size of the pupil according to the light levels in front of the subject.

Eye geometry

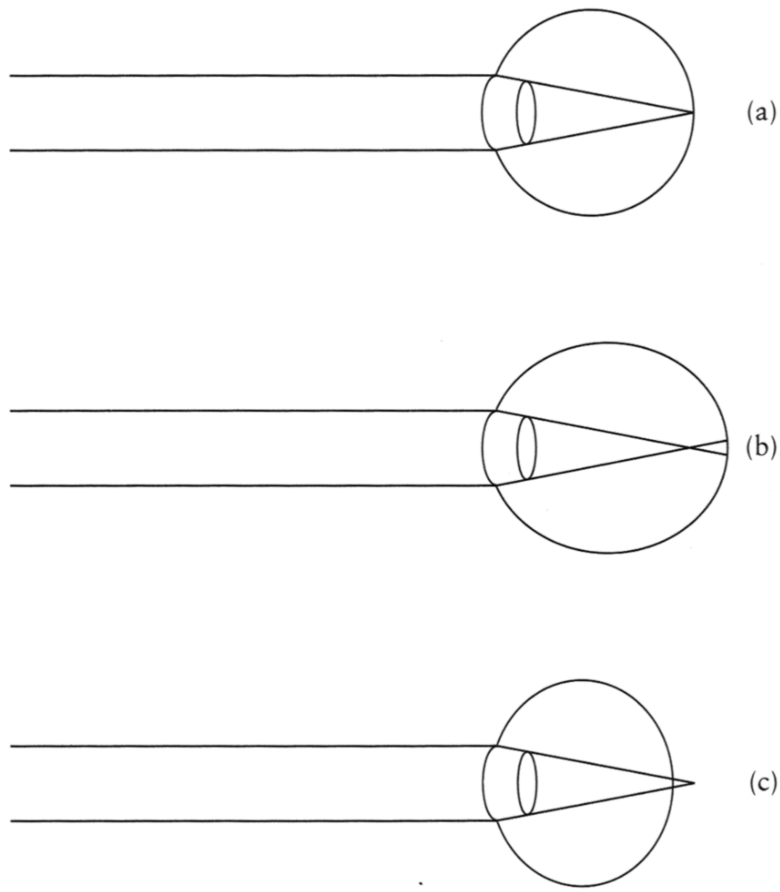

Eye geometry can account for near- and far-sightedness.

- **Emmetropic eye** - resting eye has focal point on retina.
- **Myopic eye** - eye too long.
- **Hyperopic eye** - eye too short.

Near- and far-sightedness can also result from deficiencies in focusing at the cornea or through the lens.

Retina

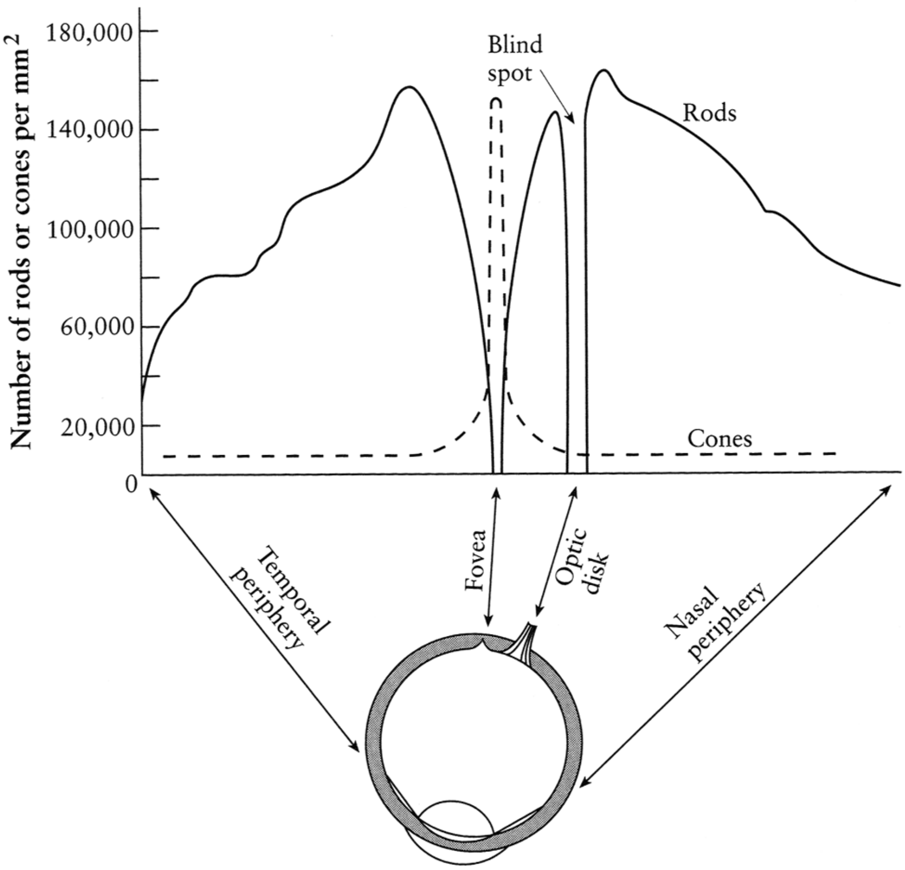

Density of photoreceptors on the retina (Glassner, 1.4)

- **Retina** - a layer of photosensitive cells covering 200° on the back of the eye.
  - **Cones** - responsible for color perception.
  - **Rods** - Limited to intensity (but 10x more sensitive).
- **Fovea** - Small region (1 or 2°) at the center of the visual axis containing the highest density of cones (and no rods).

The human retina

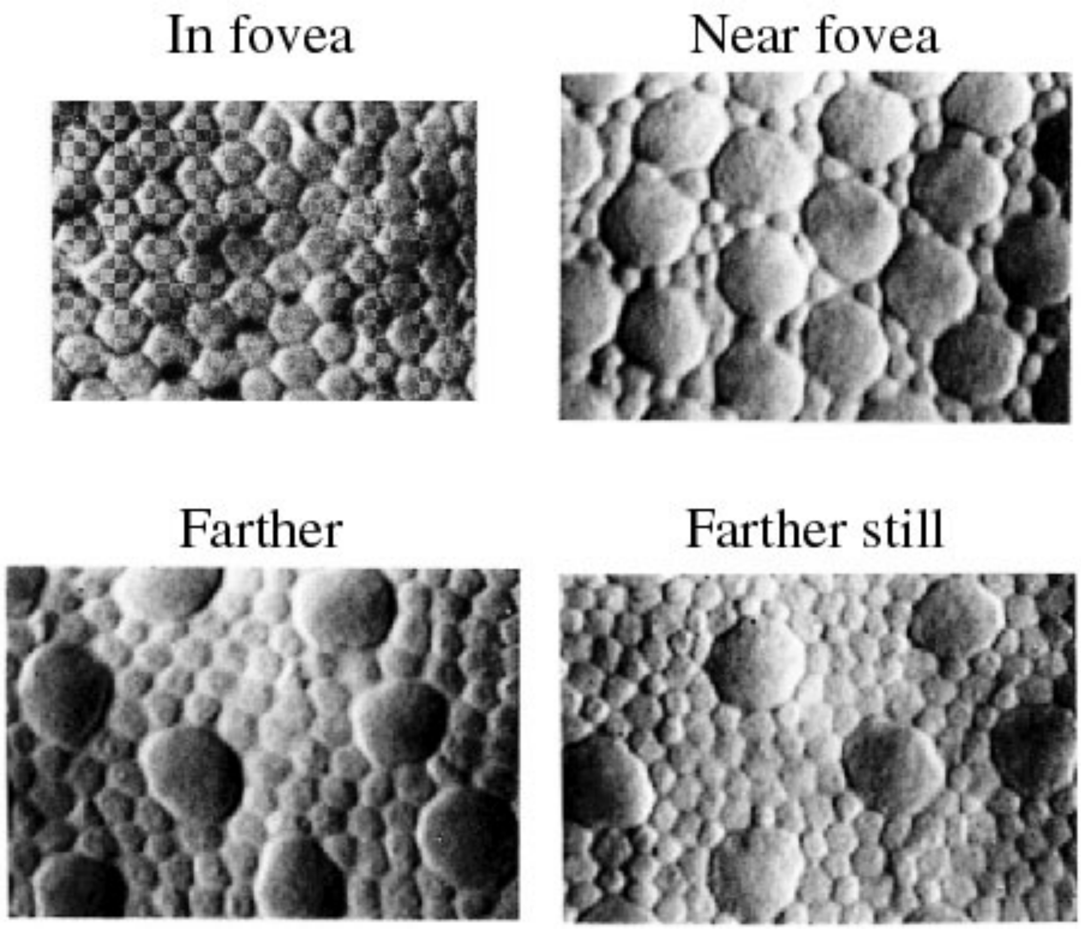

Photomicrographs at incresasing distances from the fovea. The large cells are cones; the small ones are rods. (Glassner , 1.5 and Wandell, 3.4).

Photomicrographs at increasing distances from the fovea. The large cells are cones; the small ones are rods.

The human retina, cont'd

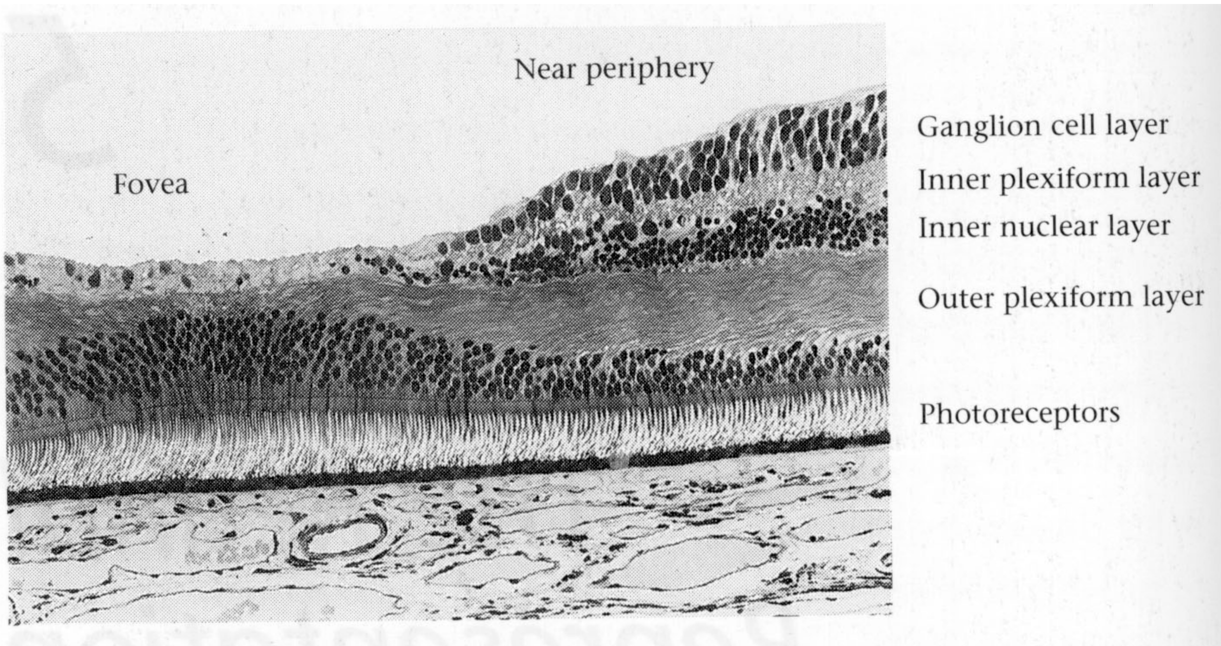

Photomicrograph of a cross-section of the retina near the fovea (Wandell, 5.1).

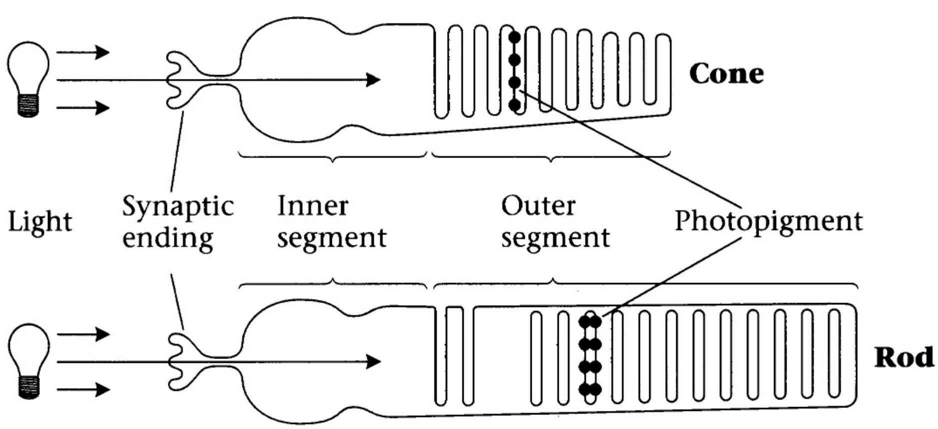

Light gathering by rods and cones (Wandell, 3.2)

Neuronal connections

Even though the retina is very densely covered with photoreceptors, we have much more acuity in the fovea than in the periphery.

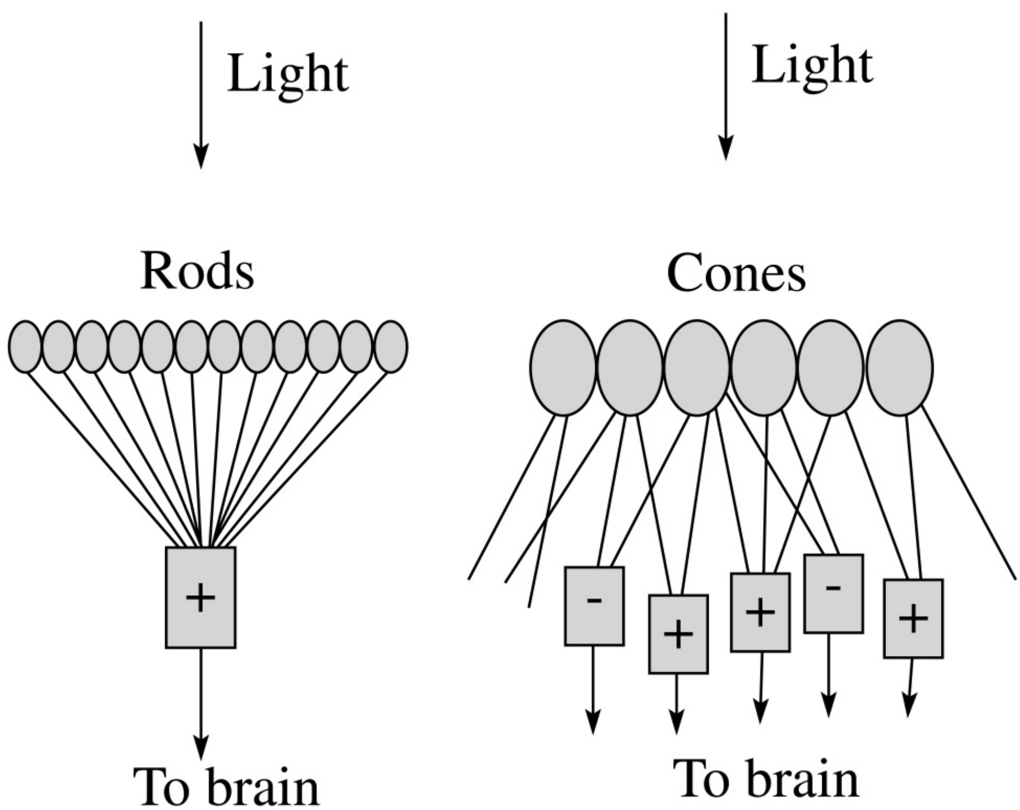

In the periphery, the outputs of the photoreceptors are averaged together before being sent to the brain, decreasing the spatial resolution. As many as 1000 rods may converge to a single neuron.

Demonstrations of visual acuity

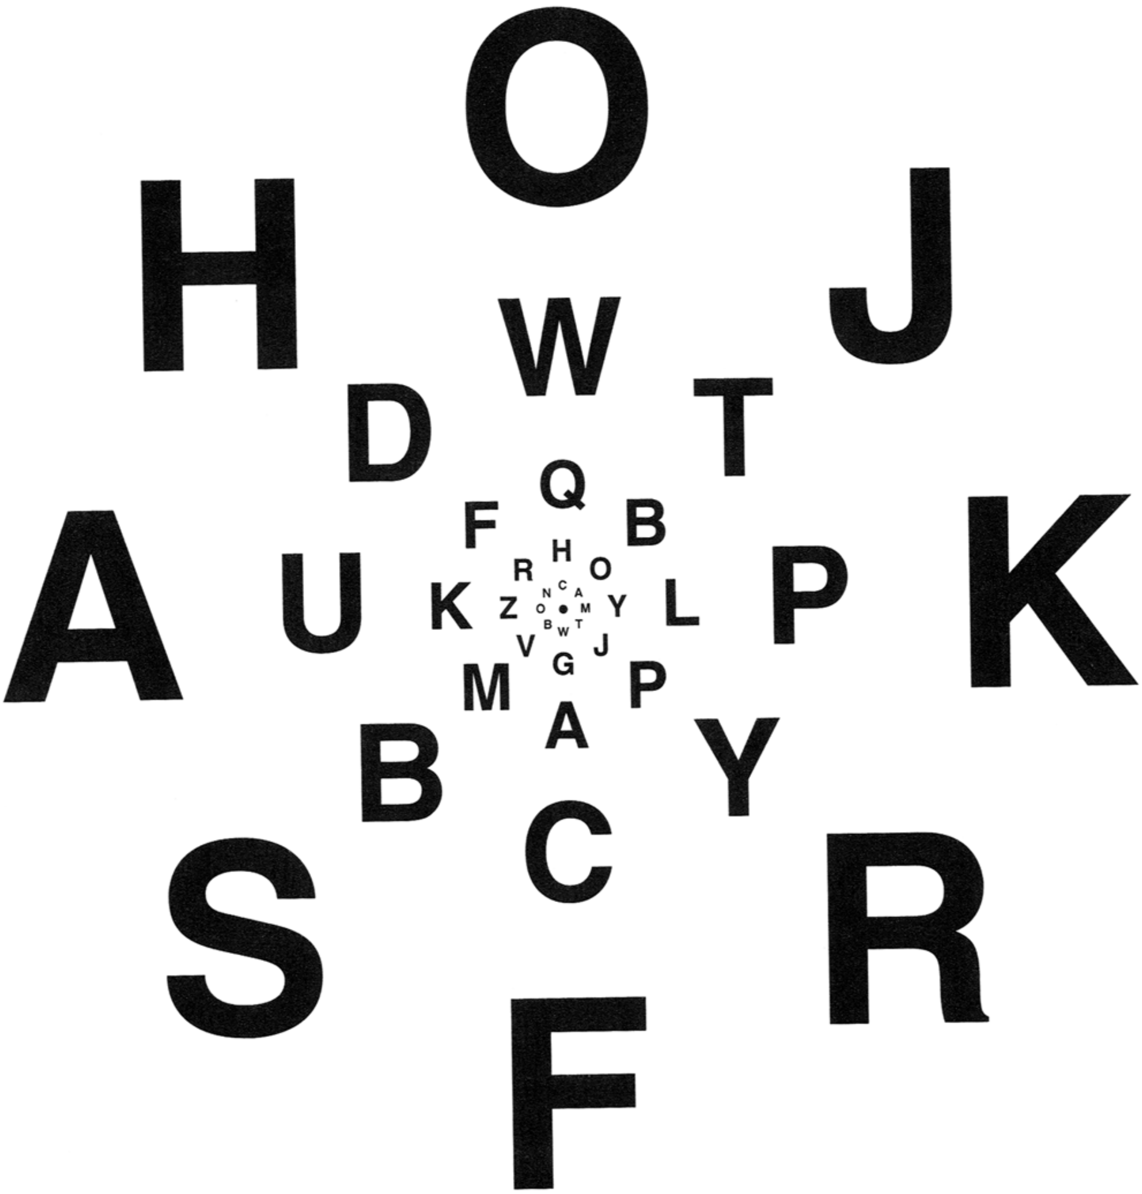

With one eye shut, at the right distance, all of these letters should appear equally legible (Glassner, 1.7).

Blind spot demonstration (Glassner, 1.8)

High resolution imaging?

Given that our vision is only high resolution over a very small range of our visual field…

…how do we manage to see “everything” at high resolution?

Saccades

By scanning our eyes over a scene, we build a composite high resolution image in our brain.

The scanning behavior, known as **saccades**, is very jittery.

Through gaze tracking, scientists can study how we look at the world.

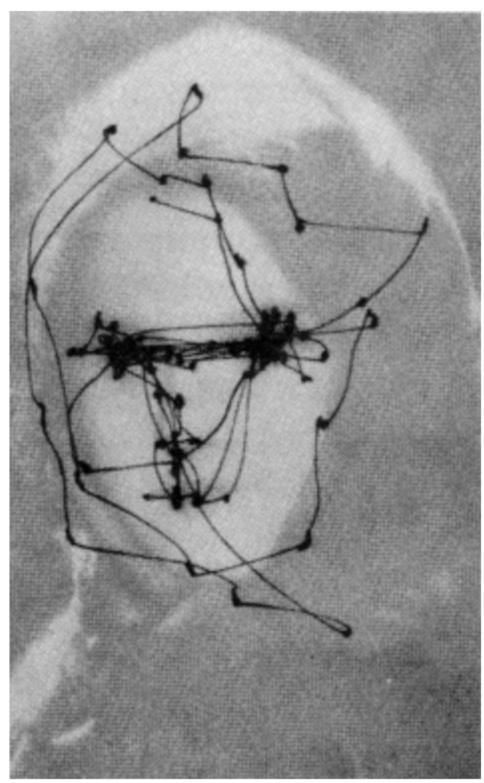

Yarbus, 1965

Saccades, cont’d

The saccidic behavior is task-specific:

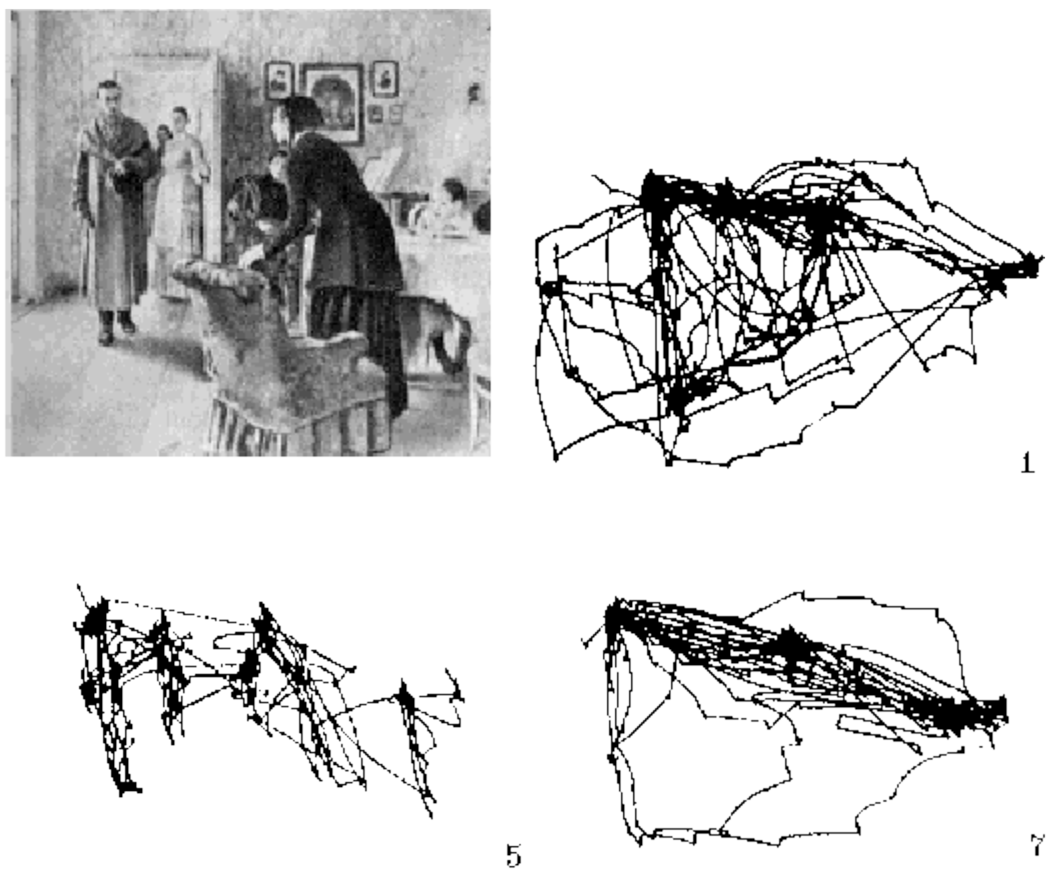

1. Free examination.

5. Remember the clothes worn by the people

7. Estimate how long the "unexpected visitor" had been away from the family

Yarbus, 1965

Perceptual light intensity

The human eye is highly adaptive to allow us a wide range of flexibility.

One consequence is that we perceive light intensity as we do sound, I.e., on a *relative* or *logarithmic* scale.

Example: The perceived difference between 0.20 and 0.22 is the same as between 0.80 and ______.

Lightness contrast and constancy

The apparent brightness of a region depends largely on the surrounding region.

The **lightness contrast** phenomenon makes a constant colored region seem lighter or darker depending on the surround:

The **lightness constancy** phenomenon makes a surface look the same under widely varying lighting conditions.

Adaptation

Adaptive processes can adjust the base activity (“bias”) and scale the response (“gain”).

Through **adaptation**, the eye can handle a large range of illumination:

Background	Luminance (cd/m^2)
Moonless overcast night	0.00003
Moonlit covercast night	0.003
Twilight	3
Overcast day	300
Day with sunlit clouds	30,000

Some of our ability to handle this range comes from our ability to control the iris (aperture) of our eyes, and the fact that we have different types of photoreceptors.

However, much of the range comes from the adaptability of the photoreceptors themselves.

Noise

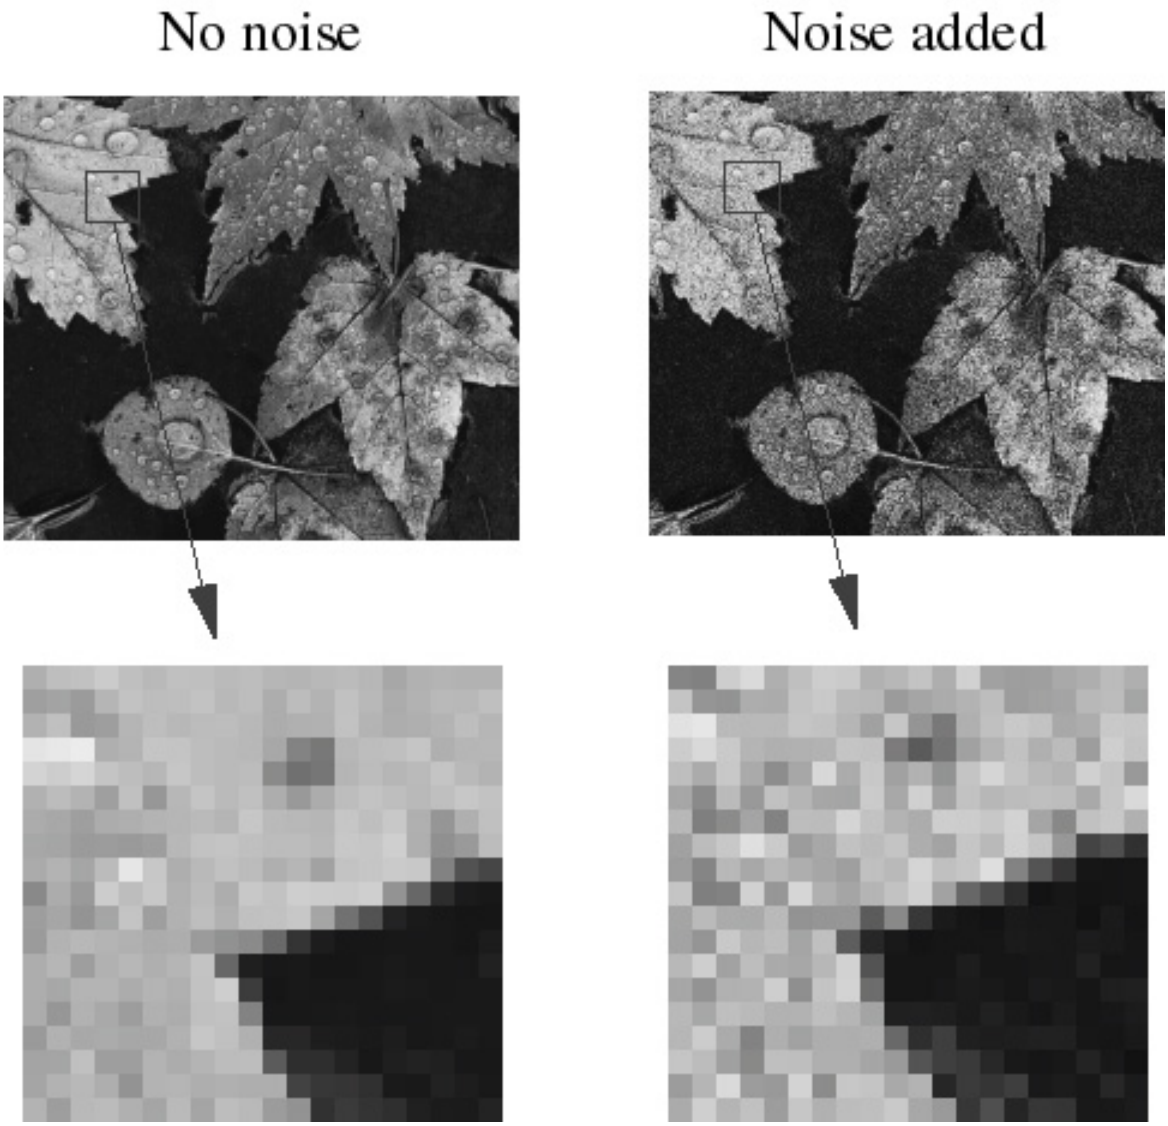

Noise can be thought of as randomness added to the signal.

The eye is relatively insensitive to noise.

Mach bands

Mach bands were first dicussed by Ernst Mach, an Austrian physicist.

Appear when there are rapid variations in intensity, especially at C^0 intensity discontinuities:

And at C^1 intensity discontinuities:

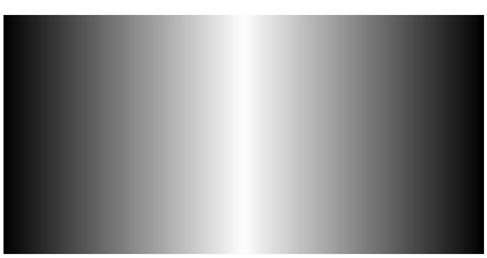

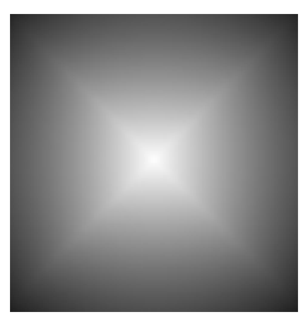

Mach bands, cont.

Possible cause: lateral inhibition of nearby cells.

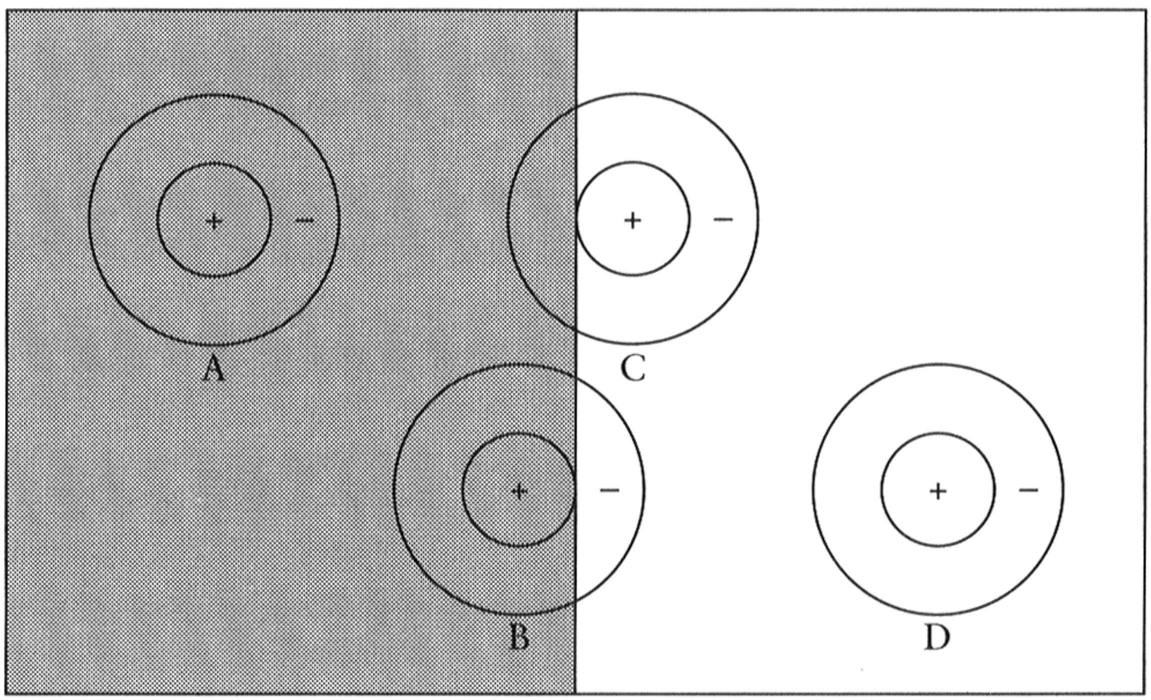

Lateral inhibition effect (Glassner, 1.25)

Q: What image processing filter does this remind you of?

The radiant energy spectrum

We can think of light as waves, instead of rays.

Wave theory allows a nice arrangement of electromagnetic radiation (EMR) according to wavelength:

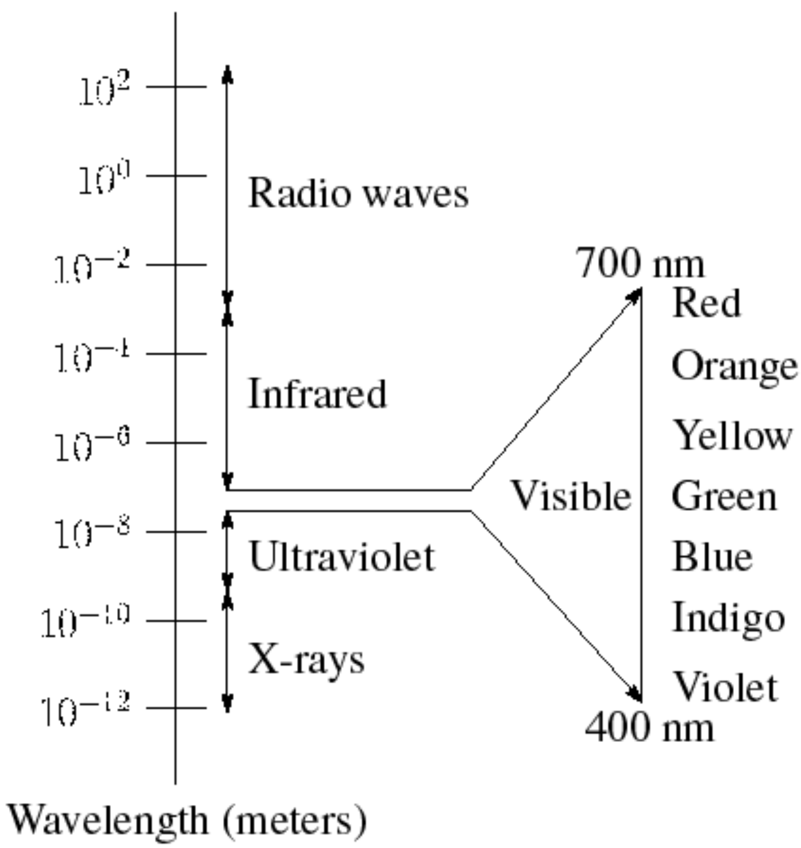

Emission spectra

A light source can be characterized by an emission spectrum:

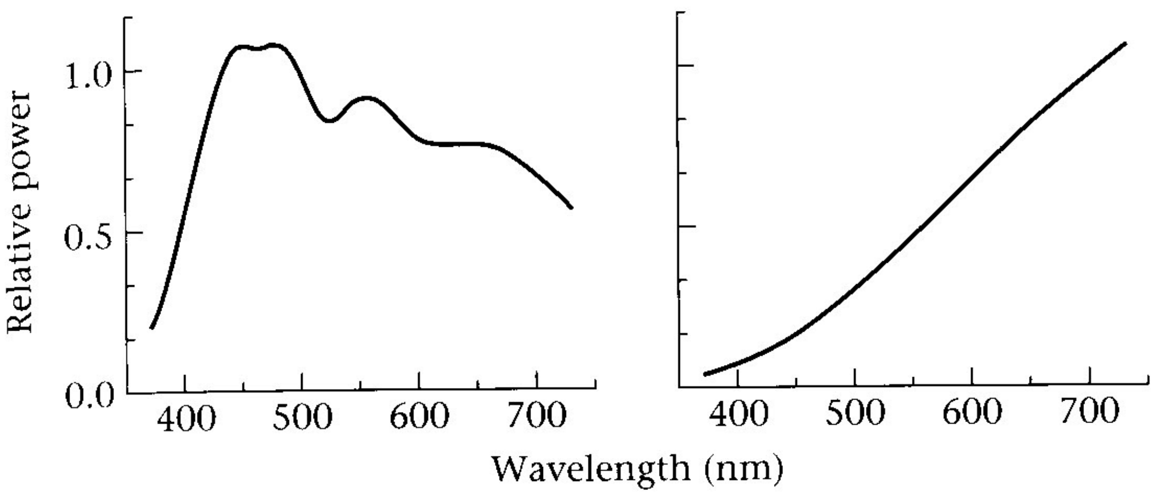

Emission spectra for daylight and a tungsten lightbulb (Wandell, 4.4)

The spectrum describes the energy at each wavelength.

What is color?

The eyes and brain turn an incoming emission spectrum into a discrete set of values.

The signal sent to our brain is somehow interpreted as *color*.

Color science asks some basic questions:

- When are two colors alike?
- How many pigments or primaries does it take
  to match another color?

One more question: why should we care?

Photopigments

Photopigments are the chemicals in the rods and cones that react to light. Can respond to a single photon!

Rods contain **rhodopsin**, which has peak sensitivity at about 500nm.

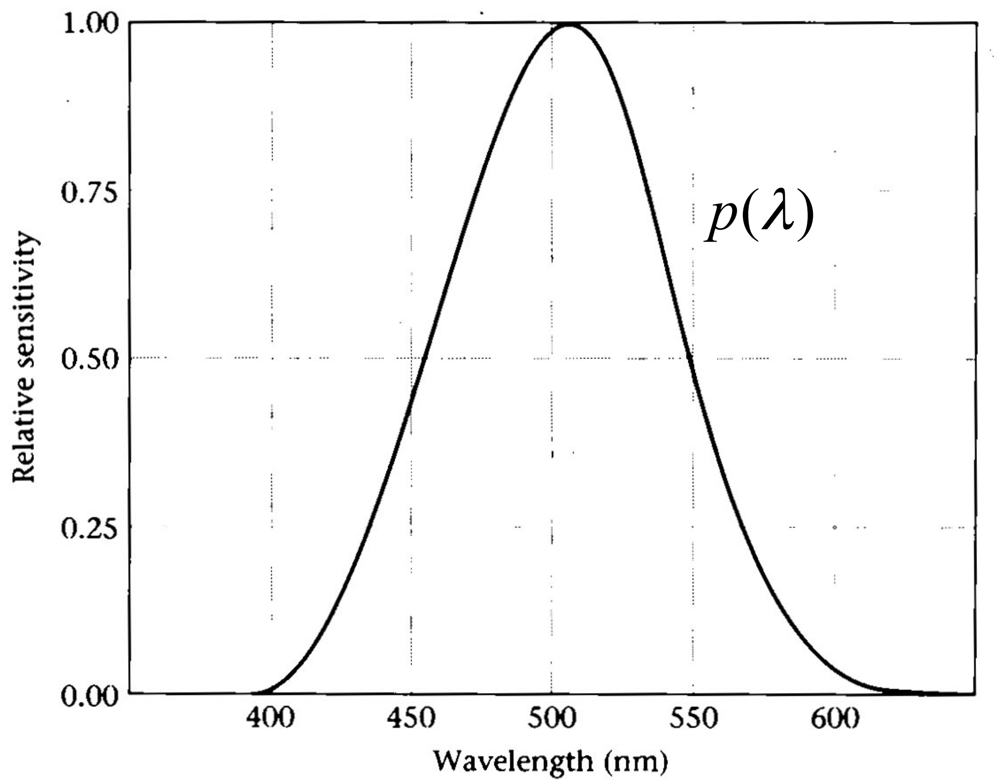

Rod sensitivity (Wandell ,4.6)

Rods are active under low light levels, i.e., they are responsible for **scotopic** vision.

Univariance

Principle of univariance: For any single photoreceptor, no information is transmitted describing the wavelength of the photon.

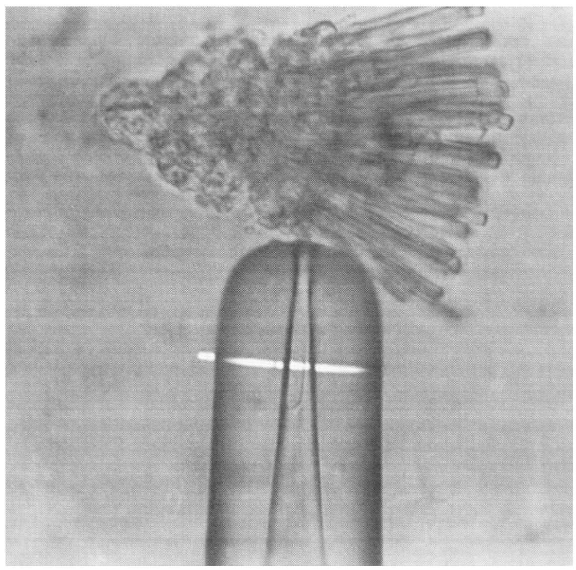

Measuring photoreceptor photocurrent (Wandell, 4.15)

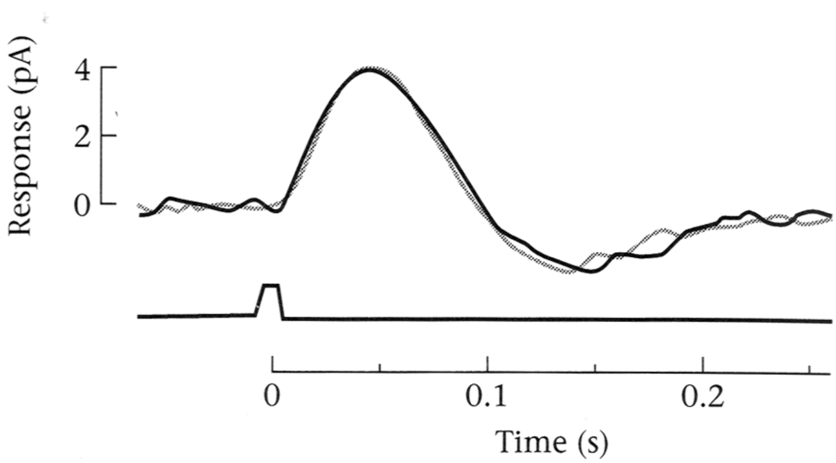

Photocurrents measured for two light stimuli: 550nm (solid) and 659 nm (gray). The brightnesses of the stimuli are different, but the shape of the response is the same. (Wandell 4.17)

What rods measure

A rod responds to a spectrum through its spectral sensitivity function, *p(λ)*.

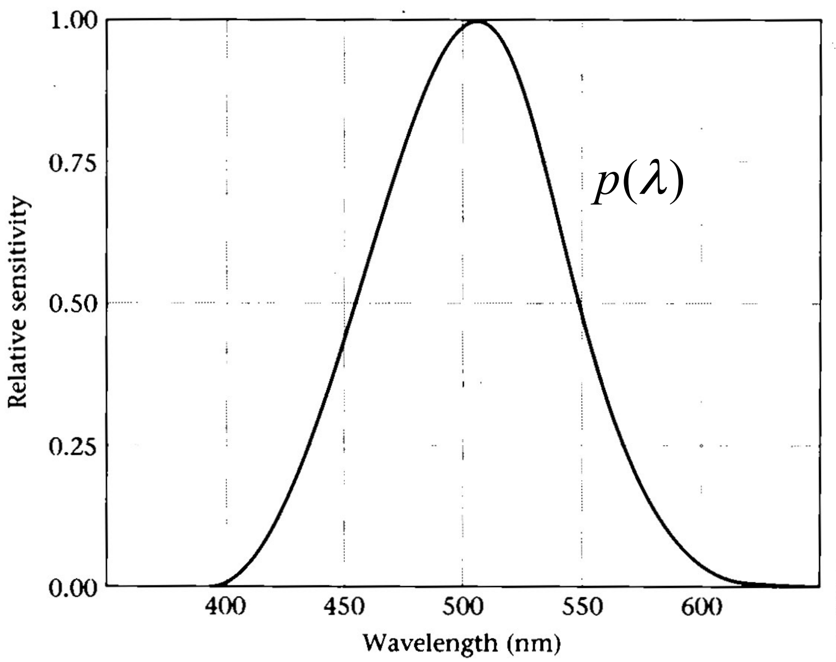

The response to a test light, *t*(λ), is simply:

$$P_t = \int t(\lambda)p(\lambda)d\lambda$$

Suppose a rod sees three light spots:

455nm blue laser of amplitude 1.0

505nm green laser of amplitude 0.5

550nm yellow laser of amplitude 1.0

Will these spots look different?

Cone photopigments

Cones come in three varieties: L, M, and S.

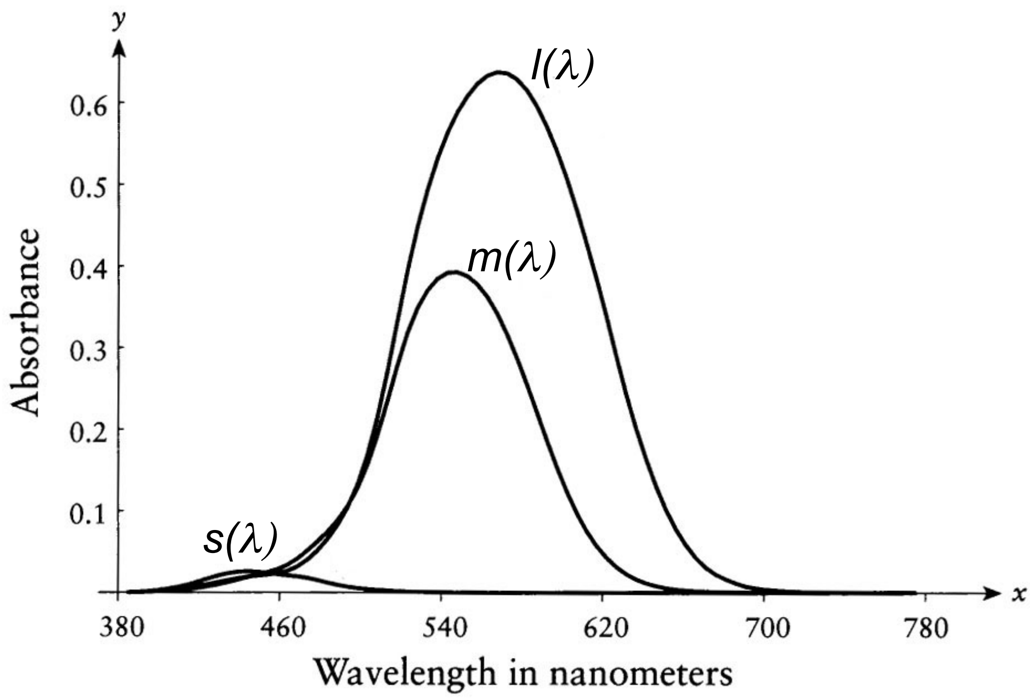

Cone photopigment absorption (Glassner, 1.1)

Cones are active under high light levels, i.e., they are responsible for **photopic** vision.

What cones measure

Color is perceived through the responses of the cones to light.

The response of each cone can be written simply as:

$$L_t = \int t(\lambda) l(\lambda) d\lambda$$

$$M_t = \int t(\lambda) m(\lambda) d\lambda$$

$$S_t = \int t(\lambda) s(\lambda) d\lambda$$

These are the only three numbers used to determine color.

What cones measure

Consider the sensitivity spectra again:

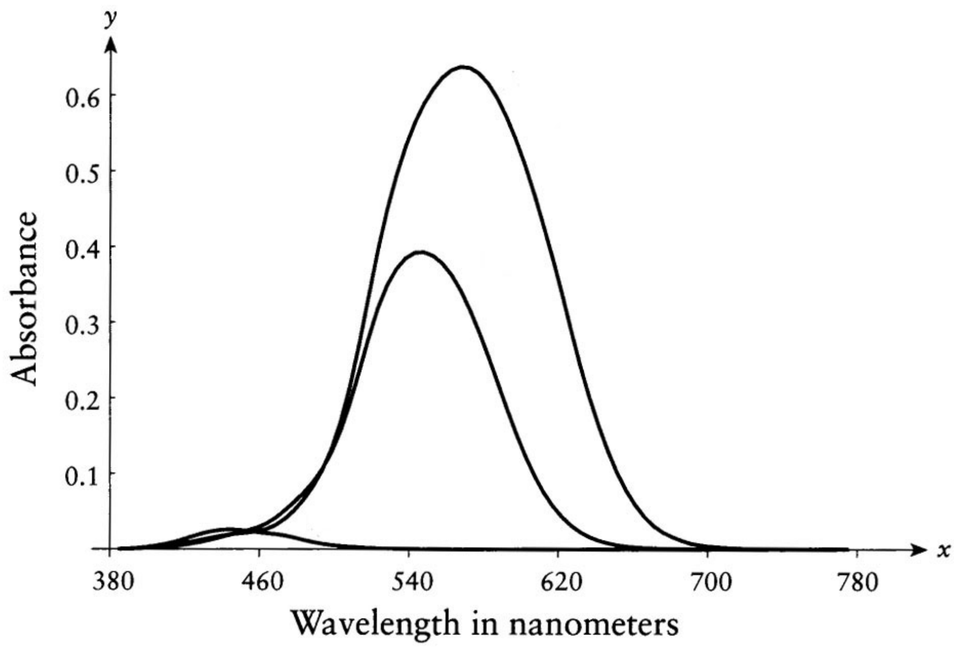

Suppose we show three light spots with unit intensity lasers at 460nm, 540nm, and 620nm. What will the cones measure?

What if I put all of these together in one spot?

The color matching experiment

We can actually distinguish all of the individual wavelengths as different colors. Does this mean our eyes are full spectral sensors?

Unfortunately, no. To show this, we can perform a color matching experiment.

The idea is to see if we can match a given test light using a finite number of lights called **primaries** with power control knobs.

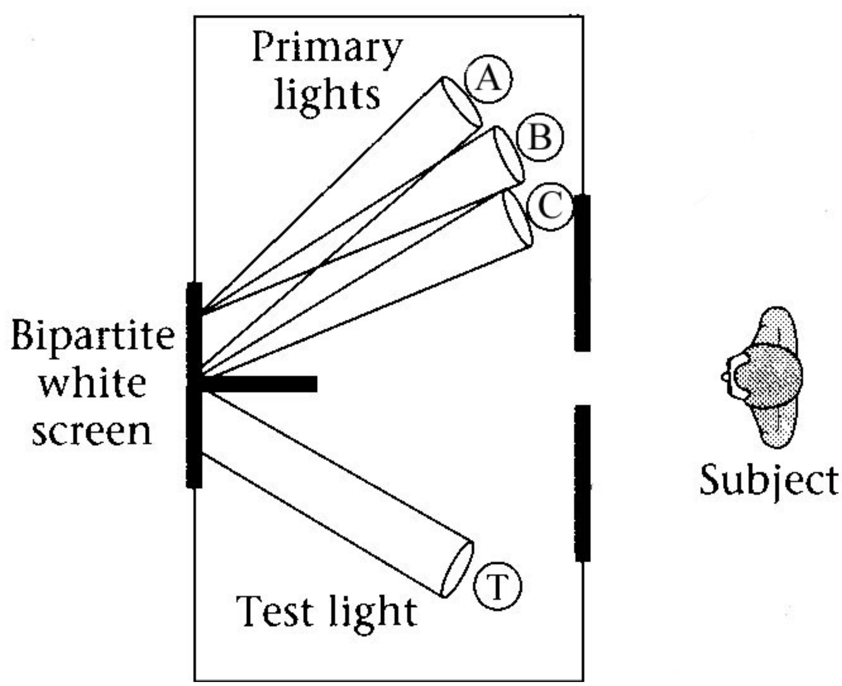

The color matching experiment (Wandell, 4.10)

The primary spectra are $a(\lambda)$, $b(\lambda)$, $c(\lambda)$, …

The power knob settings are A, B, C, …

What cones measure

Again, for our test light, the cone responses are:

$$L_t = \int t(\lambda) l(\lambda) d\lambda$$

$$M_t = \int t(\lambda) m(\lambda) d\lambda$$

$$S_t = \int t(\lambda) s(\lambda) d\lambda$$

Suppose we use the laser lights noted earlier. Can we match the test light?

Choosing Primaries

The primaries could be three color (monochromatic) lasers.

But, they can also be non-monochromatic, e.g., monitor phosphors:

$$e(\lambda) = Rr(\lambda) + Gg(\lambda) + Bb(\lambda)$$

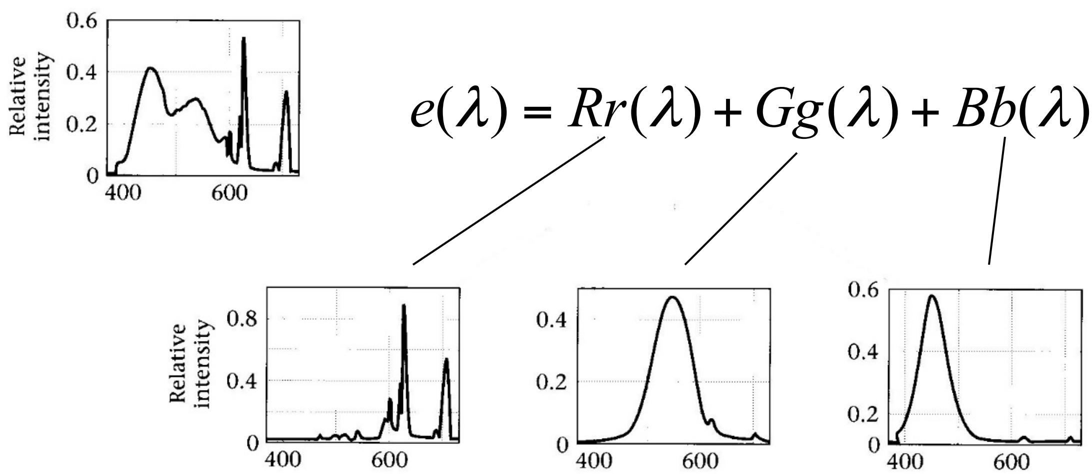

Emission spectra for RGB monitor phosphors (Wandell B.3)

Emission Spectrum is not Color

Clearly, spectral information is lost when measured by cones…

Different light sources can evoke exactly the same colors. Such lights are called **metamers**.

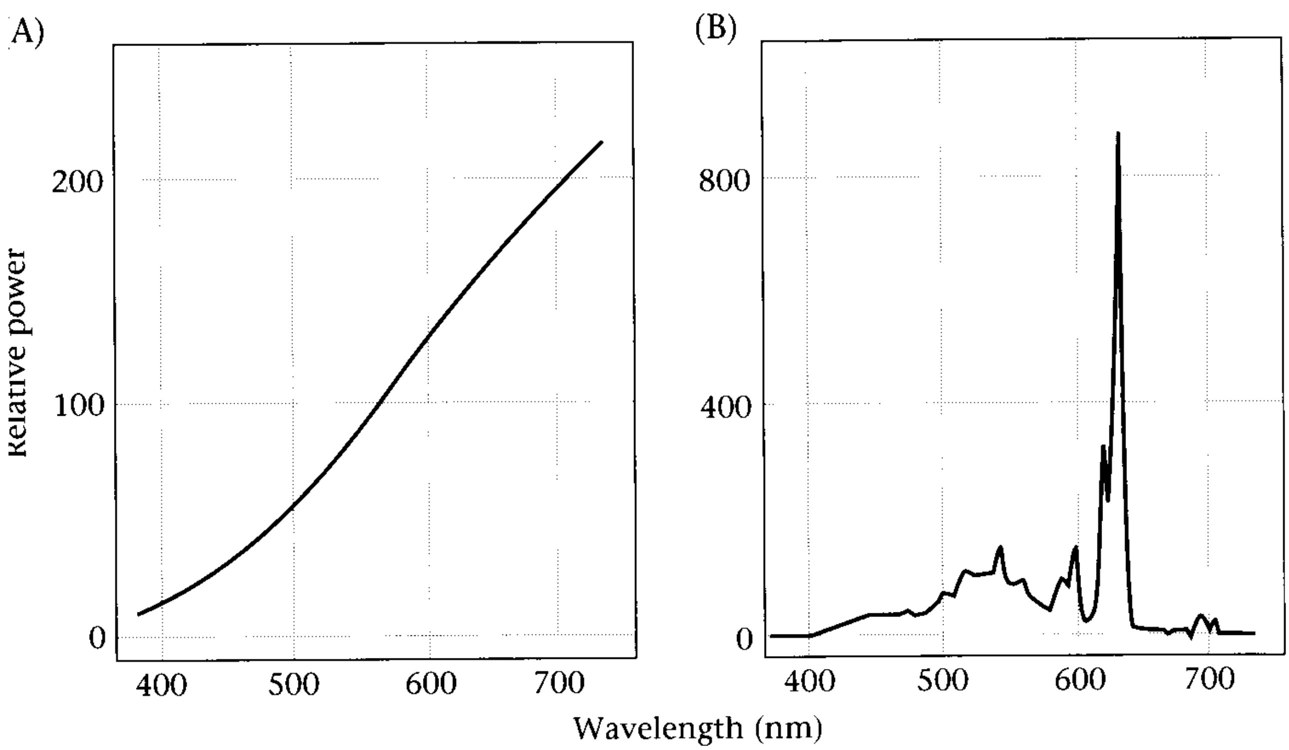

A dim tungsten bulb and an RGB monitor set up to emit a metameric spectrum (Wandell 4.11)

Colored Surfaces

So far, we've discussed the colors of lights. How do *surfaces* acquire color?

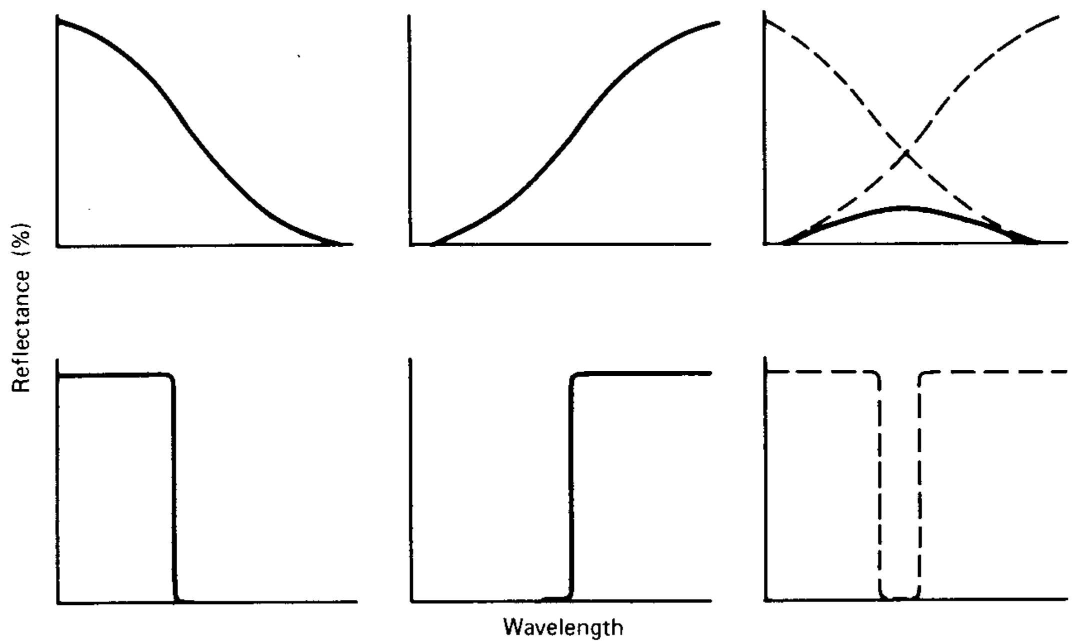

Subtractive colour mixing (Wasserman 2.2)

A surface's **reflectance**, $\rho(\lambda)$, is its tendency to reflect incoming light across the spectrum.

Reflectance is combined "**subtractively"** with incoming light. Actually, the process is *multiplicative*:

$$I(\lambda) = \rho(\lambda)t(\lambda)$$

Subtractive Metamers

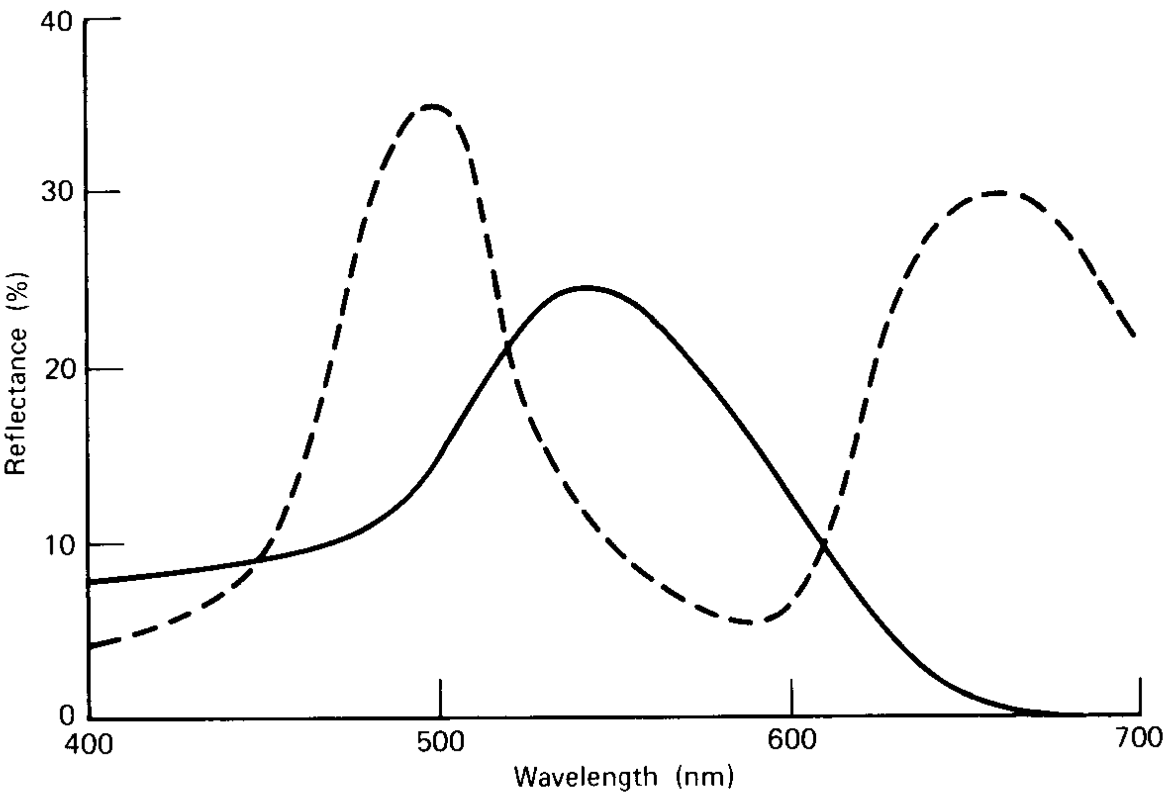

Surfaces that are metamers under only some lighting conditions (Wasserman 3.9)

Reflectance adds a whole new dimension of complexity to color perception.

The solid curve appears green indoors and out. The dashed curve looks green outdoors, but brown under incandescent light.

Illustration of Color Appearance

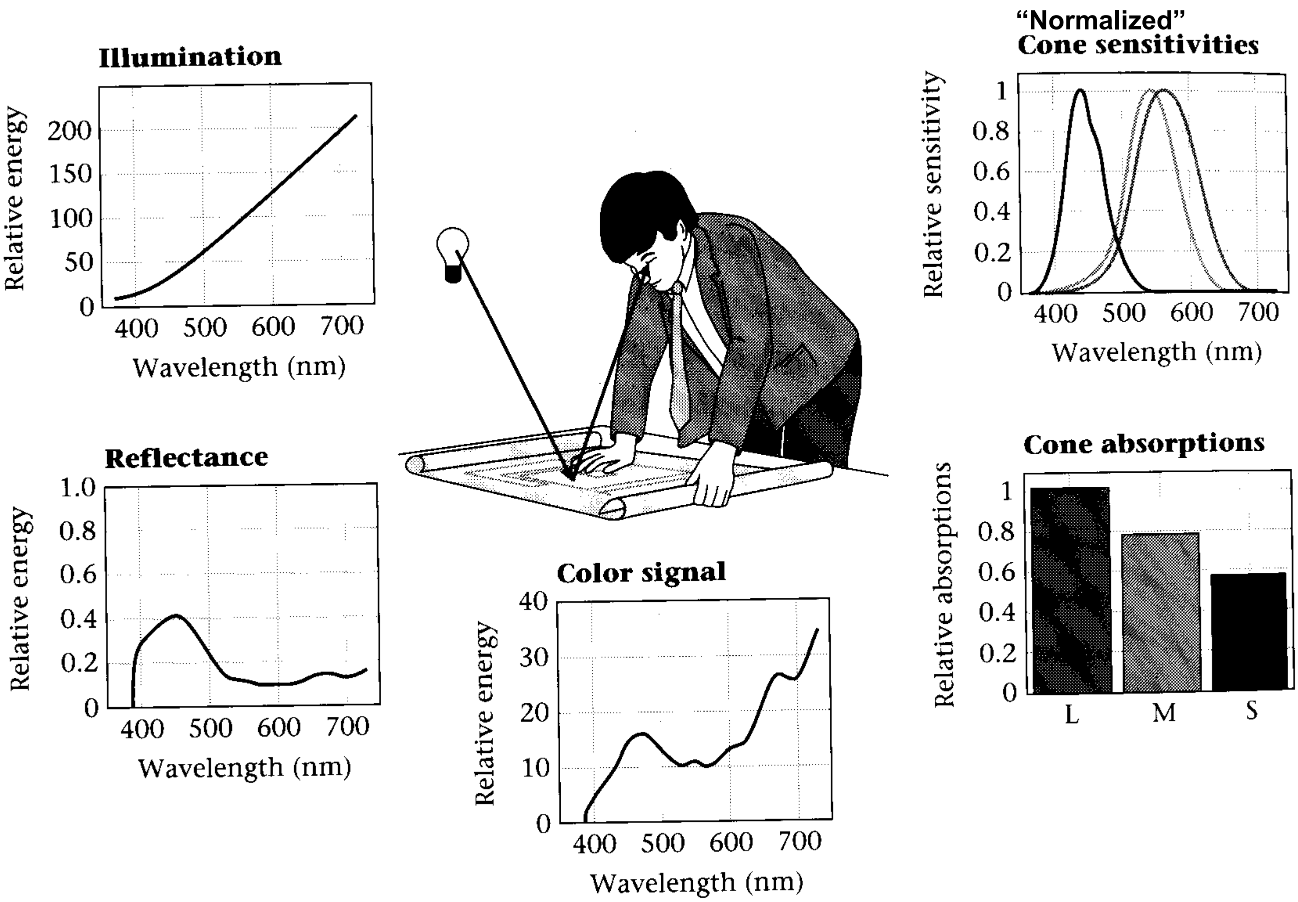

How light and reflectance become cone responses (Wandell, 9.2)