ISTA 352

Lecture 5

Images from Light

The big picture

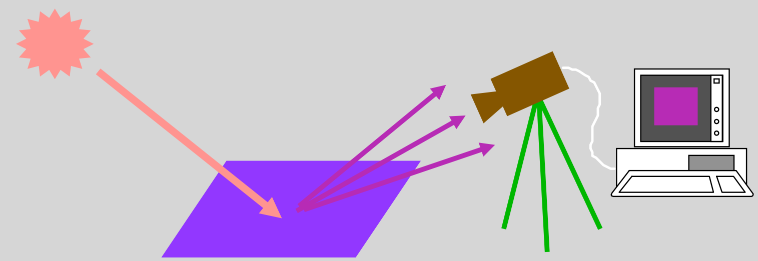

Light from a single point

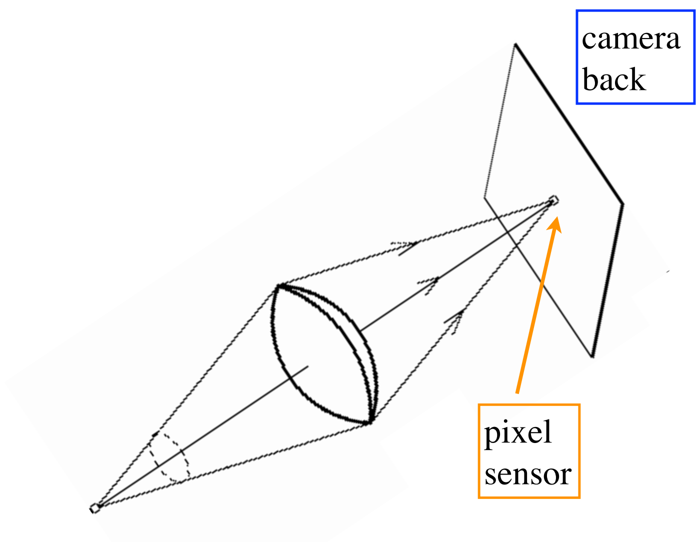

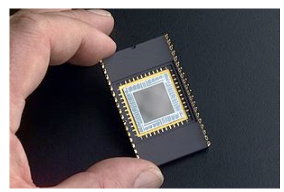

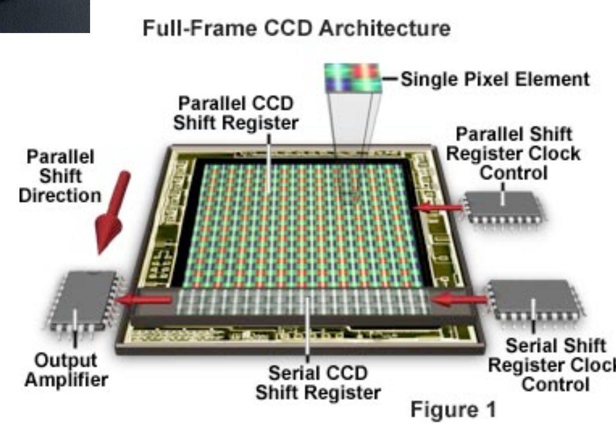

A few facts about light

- Bits of light are called photons
- Photons have different wavelengths
  - Energy per photon is inversely related to wavelength
- Visible light is consists of light with wavelengths of 400 to 700 nanometers

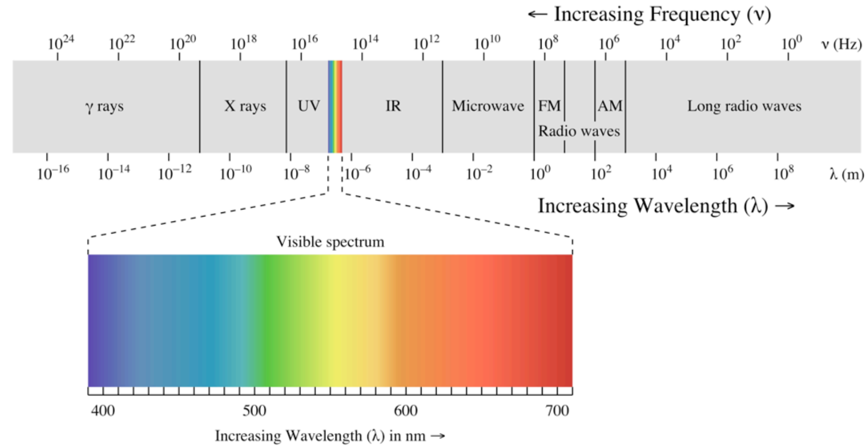

(from WikiPedia commons)

A few facts about light (2)

- The light around us is a mix of photons of different wavelengths, directions, and polarization.
- Light coming towards the camera in a specific direction has a distribution over wavelengths (light energy spectrum).

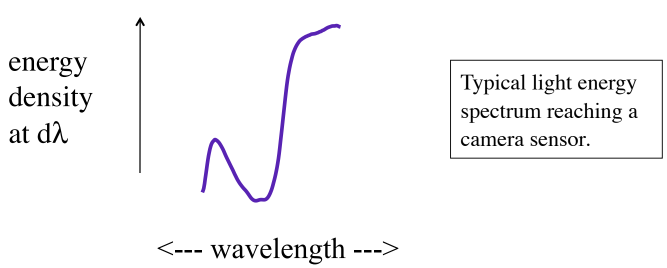

Two disparate source spectra

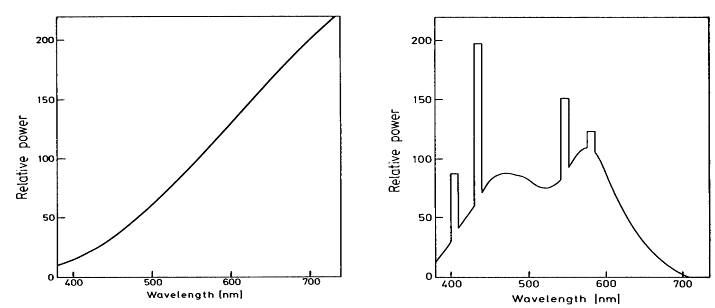

Fig. 4.1. Wavelength composition of light from a tungsten-filament lamp [typified by CIE ILL A (Sect. 4.6)]. Relative spectral power distribution curve. Color temperature: 2856 K

Fig. 4.2. Wavelength composition of light from a daylight fluorescent lamp. Typical relative spectral power distribution curve. Correlated color temperature: 6000 K. (Based on data of Jerome reported in [Ref. 3.14, p. 37])

Spectra of many sources

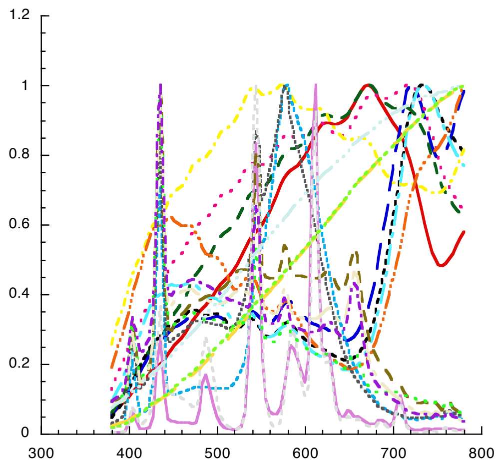

Camera sensor cartoon

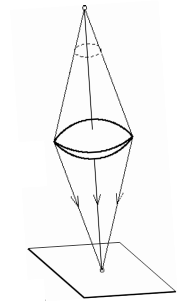

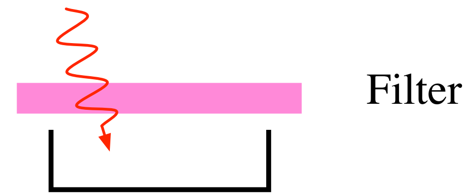

Camera sensor cartoon

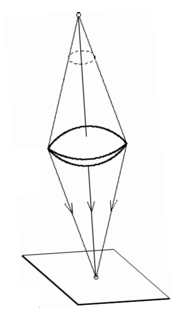

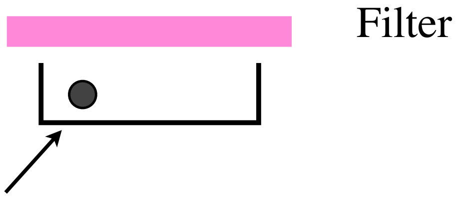

(electron dislodged and saved in well)

Camera sensor cartoon

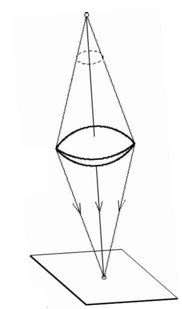

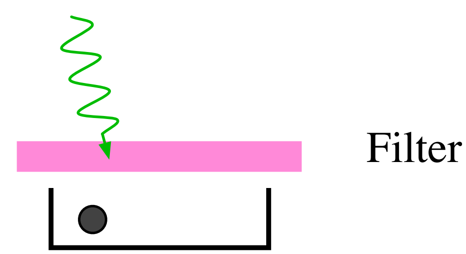

(absorbed by filter, no electrons freed)

Camera sensor cartoon

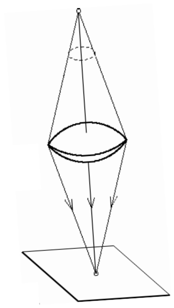

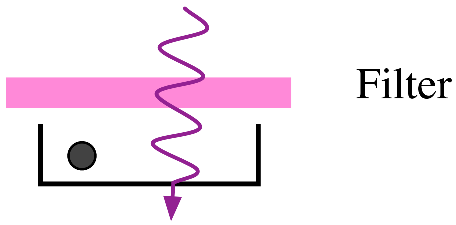

(no electron, nothing is perfect)

Camera sensor cartoon

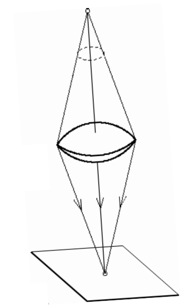

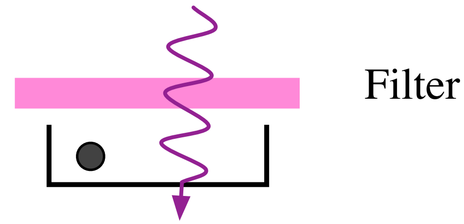

(no electron, nothing is perfect)

Camera sensor cartoon

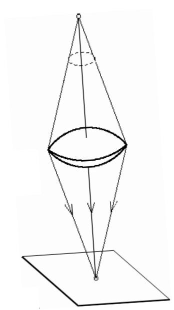

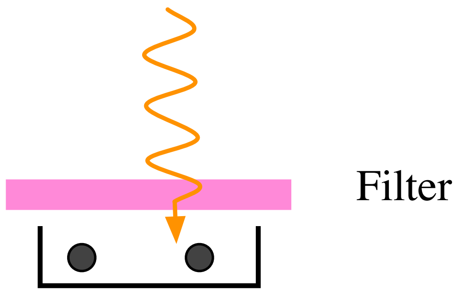

(yeah, another electron)

Camera sensor cartoon

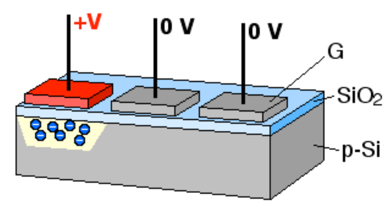

(time’s up! time to measure the electrons)

Discussion on light capture

Summary on light capture

- Light capture is **linear**
  - scale the light by a factor ==> scale the reading by that factor,
  - have two lights together ==> sum the readings

- Mathematically (**definition of linear function**)

$$f(a*x+b*y)=a*f(x)+b*f(y)$$

(a and b are constants)

Summary on light capture (2)

- Different photon wavelengths have different capture efficiency
  - (Partly due to filters needed to make a color camera)
  - We express this relative sensitivity by sensitivity curves
    - The absolute scale of the curve is set by the units used in calibration and is not very interesting

Sensitivity at a wavelength is indicated by the height of the curve

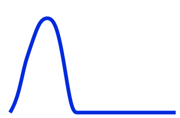

<--- wavelength --->

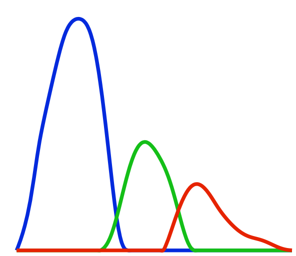

<--- wavelength --->

The sensor sensitivity spectra for the three kinds of pixels in a particular color camera (they vary).

Sensor response mechanism (linear)

For **each** wavelength, the sensor responds in proportion to the signal, AND the sensor sensitivity

Thus the response for a **wavelength** is the product of L(λ)*R(λ)).

The response to the **entire** signal is the sum of the above for all wavelengths, λ.

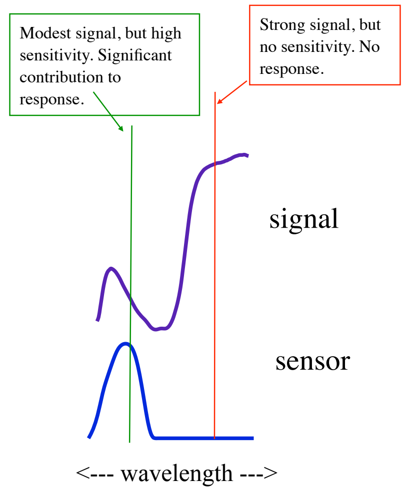

Representing spectra

- Two representations of spectra
  - A continuous function from 400nm to 700nm
  - A vector of discrete samples
    - e.g., 400nm, 404nm, 408nm, … . 696nm, 700nm.
    - (101 evenly spaced samples)
- We will use the vector representation
  - Let **L** be the light energy spectrum
  - Let $\mathbf{R}^{k}$ be the sensitivity of the k'th channel (R,G, or B)
- Then the camera response (k'th channel) is given by

$$\rho^{(k)} = \mathbf{L} \bullet \mathbf{R}^{(\mathbf{k})}$$

Important

Sensor/light interaction example

R=(0,0,1,3,7,4,2,0,0,0)

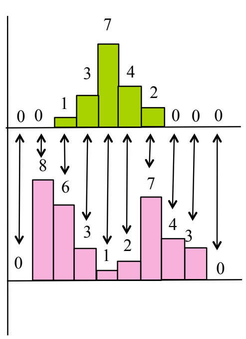

L=(0,8,6,3,1,2,7,4,3,0)

Multiply lined up pairs of numbers and then sum up

Important

Sensor/light interaction example

R=(0,0,1,3,7,4,2,0,0,0)

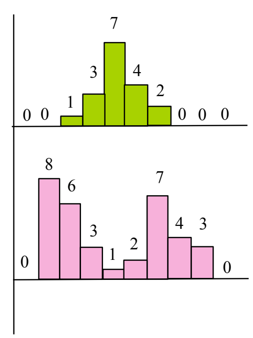

L=(0,8,6,3,1,2,7,4,3,0)

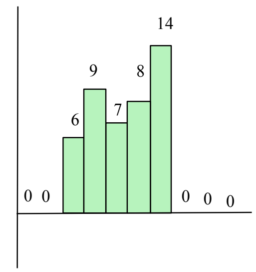

L .* R=
(0*0, 0*8, 1*6, 3*3, 7*1, 4*2, 2*7, 0*4, 0*3, 0*0)
=(0,0,6,9,7,8,14,0,0,0)

L • R = 0 + 0 + 6 + 9 + 7 + 8 + 14 + 0 + 0 + 0
= 44

- Our dot product gives the response of **one** channel to **one** light spectra

- To compute R, G, B for many light spectra (in pictures)

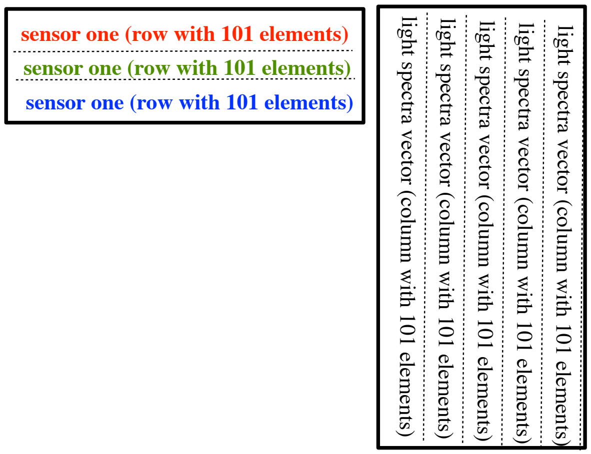

Gamma correction

- For complicated reasons, the final output of a camera is a **non-linear transformation** of the RGB described so far.

- Usually the same transformation is used for R, G, and B

- A typical "gamma correction" transformation is

$$F(x) = 255 * \left(\frac{x}{255}\right)^{1/2.2} \qquad \text{(roughly square root)}$$

- For more information
  - Some supplementary slides follow
  - Perhaps we can cover it together with the demo on Sep 7?
  - Perhaps we will say more when we do color.

Supplemental material

More formally,

The response of an image capture system to a light signal $L(\lambda)$ associated with a given pixels is modeled by

$$\upsilon^{(k)} = F^{(k)}(\rho^{(k)}) = F^{(k)}\left(\underbrace{\int L(\lambda) \;\; R^{(k)}(\lambda) \;\; d\lambda}_{\text{estimated by the dot-product}} \right)$$

where $R^{(k)}(\lambda)$ is the sensor response function for the k^{th} channel and $\upsilon^{(k)}$ is the k^{th} channel result.

$R^{(k)}(\lambda)$ includes the contributions due to the aperture, focal length, sensor position in the focal plane.

$F^{(k)}$ absorbs typical non-linearities such as gamma. Typically F is the same for all values of *k*.

Supplemental material

Image Formation (non-linear transform)

$F^{(k)}$ is often ignored (assumed to be the identity), but this is not a safe assumption, especially when color or radiometric measurements matter.

Commonly images are "gamma" corrected by raising the RGB values (normalized to [0,1]) to the power 1/(2.2).

Note that in such an image, a number twice as large does not mean that the light had twice the power!

To linearize RGB's from such a signal we compute:

$p=F^{-1}(v)=255*(v/255)$^2.2

Supplemental material

Image Formation (non-linear transform)

The non-linear transformation is added by captured devices **after** the raw capture (which is typically linear).

Because it is a single function applied to responses, it is easy to measure and compensate for.

Supplemental material

Image Formation (non-linear transform)

Why are images typically encoded in this way?

Historically, images have been gamma corrected on the assumption that their values drive a CRT (cathode ray tube) monitor which are non-linear devices. Their theoretical response to a voltage is energy output proportional to that voltage raised to the (5/2) power. Appropriately gamma corrected images display as linear on such devices.

Supplemental material

Image Formation (non-linear transform)

Coincidentally, this typically gamma correction is a sensible way to encode image data into a limited number of values (e.g. 256) due to the noise sensitivity of the human vision system.

Hence, while CRT displays are now obsolete, images are still typically non-linear, and the signal to modern displays (which are linear) are typically adjusted assuming typical incoming non-linear in images.

Supplemental material

Image Formation (non-linear transform)

If you have access to a Mac, then you can play with this under System Preferences --> Displays --> Color --> Calibrate (may need to select “expert”)

Demo!