

Light and Shading




“Empire of Light”, Magritte

Computer Vision

Derek Hoiem, University of Illinois

How light is recorded




Digital camera



A digital camera replaces film with a sensor array

- Each cell in the array is light-sensitive diode that converts photons to electrons
- Two common types: Charge Coupled Device (CCD) and CMOS
- <http://electronics.howstuffworks.com/digital-camera.htm>

Sensor Array



a b


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




CMOS sensor

Each sensor cell records amount of light coming in at a small range of orientations

The raster image (pixel matrix)



The raster image (pixel matrix)




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

Today's class: Light and Shading




- What determines a pixel's intensity?
- What can we infer about the scene from pixel intensities?

How does a pixel get its value?




How does a pixel get its value?

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

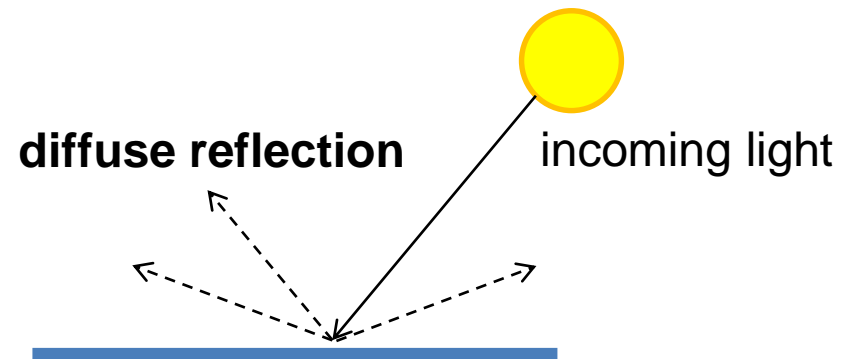


Basic models of reflection

- 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



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 $\cos(\theta)$



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

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

Intensity and Surface Orientation

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


ρ = albedo

\mathbf{S} = directional source

\mathbf{N} = surface normal

I = reflected intensity

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








1

2

Recap

- When light hits a typical surface
 - Some light is absorbed ($1-\rho$)
 - 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




transparency

light source




refraction

light source







fluorescence




phosphorescence





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





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

$$\frac{L_e(\theta_e, \phi_e)}{E_i(\theta_i, \phi_i)} = \frac{L_e(\theta_e, \phi_e)}{L_i(\theta_i, \phi_i) \cos \theta_i d\omega}$$

Application: photometric stereo

- Assume:
 - a set of point sources that are infinitely distant
 - a set of pictures of an object, obtained in exactly the same camera/object configuration but using different sources
 - A Lambertian object (or the specular component has been identified and removed)





Intensity for pixel x


Source i direction and strength

Each image is: $I_i(x) = \mathbf{S}_i \cdot (p(x)\mathbf{N}(x))$

So if we have enough images with known

sources, we can solve for $\mathbf{B}(x) = p(x)\mathbf{N}(x)$

albedo times 3D
normal vector





$$\mathbf{B}(x) = p(x)\mathbf{N}(x)$$

And the albedo (shown here) is given by:

$$p(x) = \sqrt{\mathbf{B}(x) \cdot \mathbf{B}(x)}$$



(the normal is a unit vector)






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



Color


Light is composed of a spectrum of wavelengths




Human Luminance Sensitivity Function

Some examples of the spectra of light sources


A. Ruby Laser




B. Gallium Phosphide Crystal



C. Tungsten Lightbulb







D. Normal Daylight



Some examples of the reflectance spectra of surfaces



% Photons Reflected



400

700

400

700

400

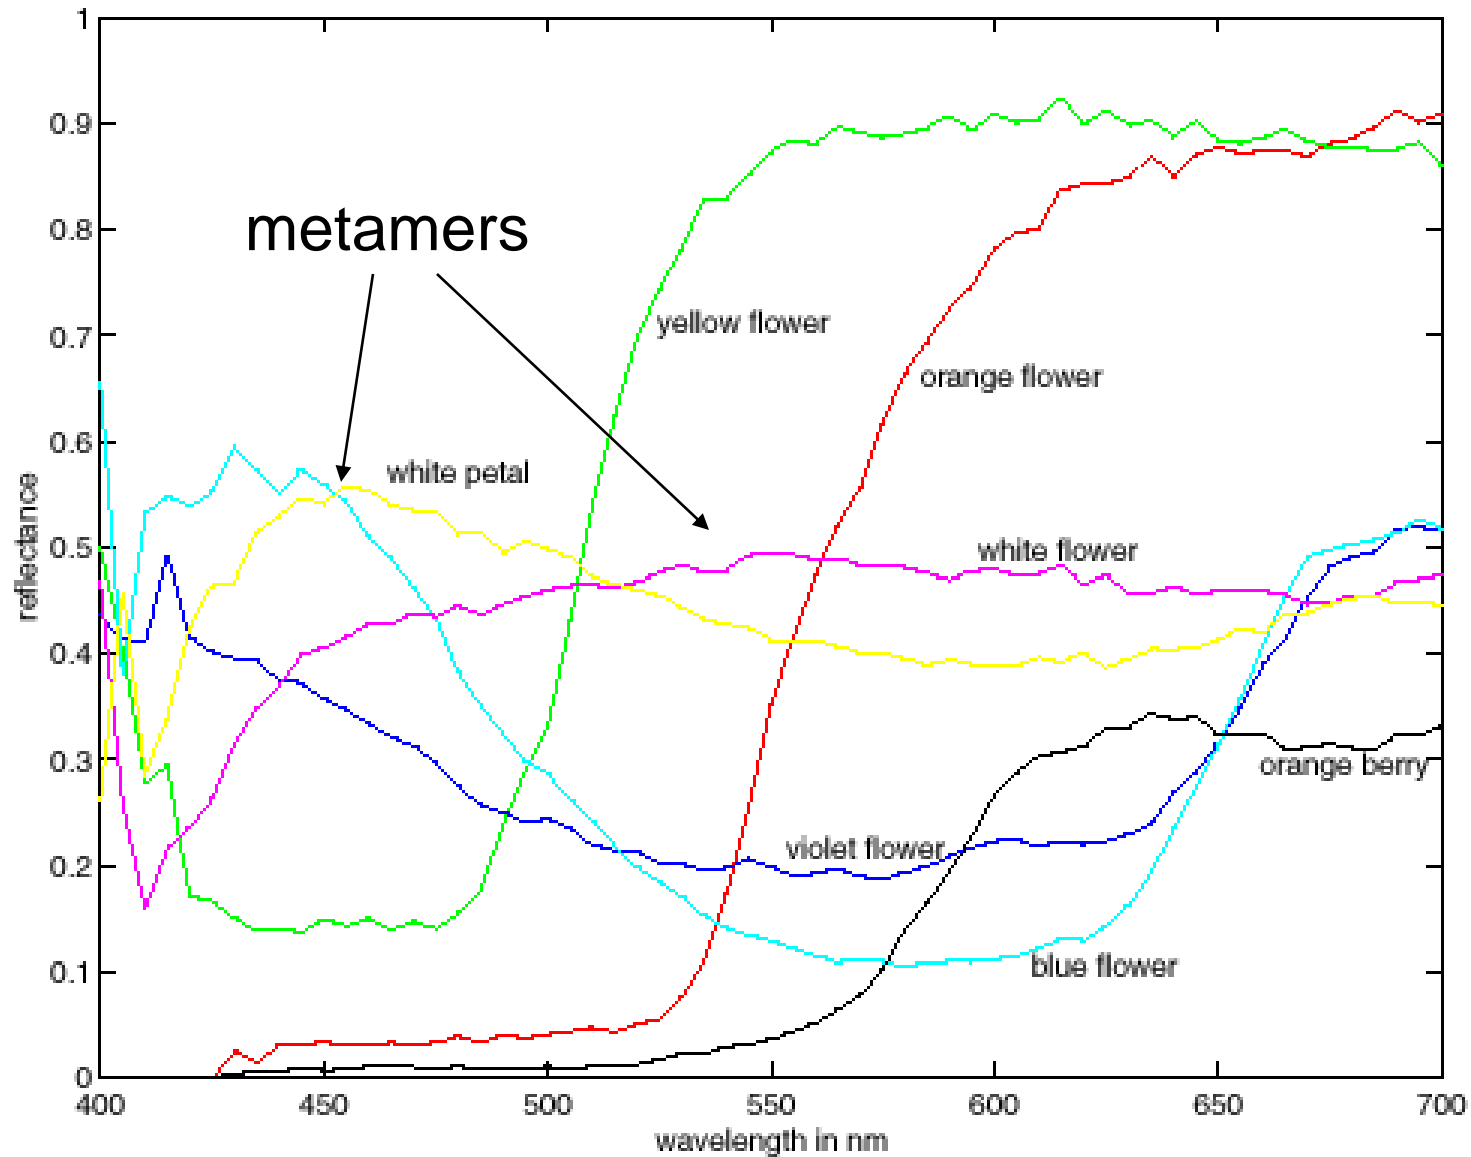
700

400

700


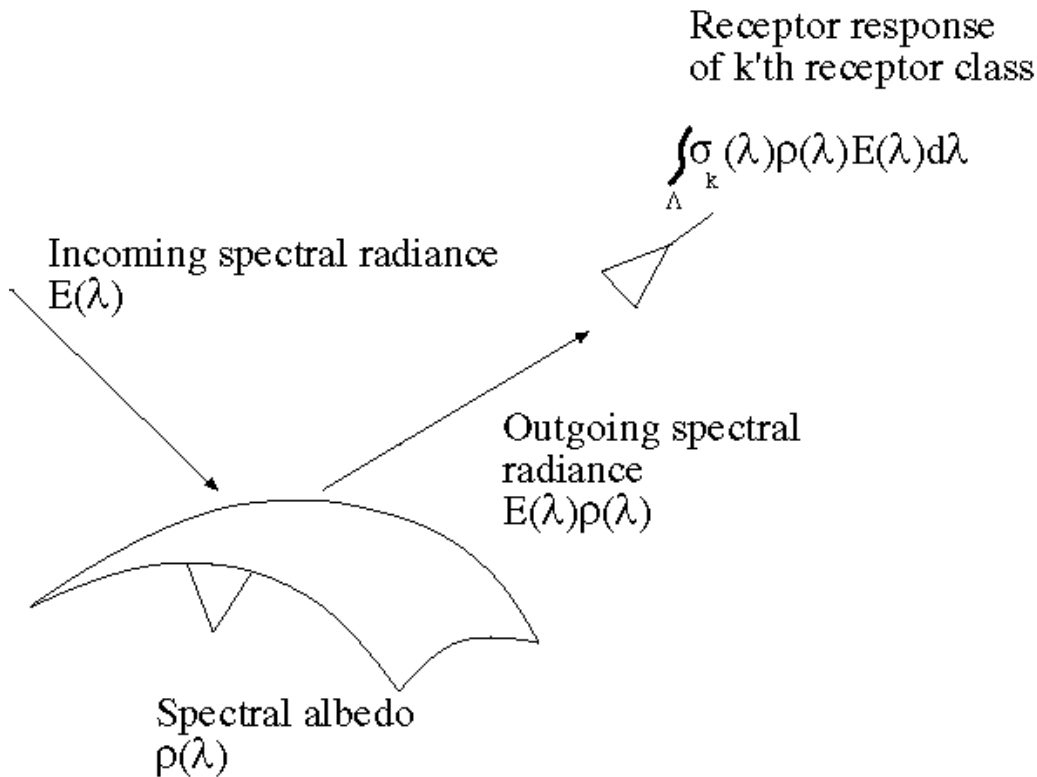
Wavelength (nm)

More spectra




The color of objects


- Colored light arriving at the camera involves two effects
 - The color of the light source (illumination + inter-reflections)
 - The color of the surface



Color Sensing: Bayer Grid



Estimate RGB at each cell from neighboring values



Color Image



R



G




B



Why RGB?

If light is a spectrum, why are images RGB?


Human color receptors



- Long (red), Medium (green), and Short (blue) cones, plus intensity rods
- Fun facts
 - “M” and “L” on the X-chromosome
 - That’s why men are more likely to be color blind (see what it’s like: <http://www.vischeck.com/vischeck/vischeckImage.php>)
 - “L” has high variation, so some women are tetrachromatic
 - Some animals have 1 (night animals), 2 (e.g., dogs), 4 (fish, birds), 5 (pigeons, some reptiles/amphibians), or even 12 (mantis shrimp) types of cones


So far: light \rightarrow surface \rightarrow 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



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

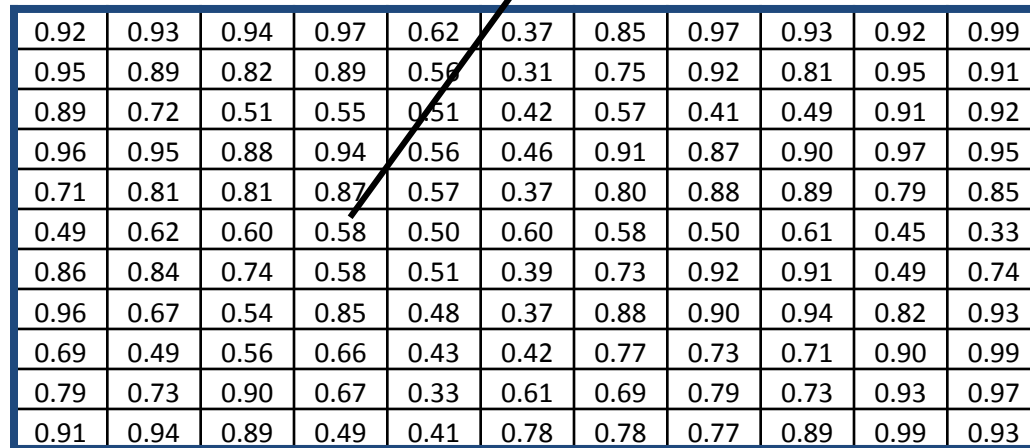
Recap



Possible factors: albedo, shadows, texture, specularities, curvature, lighting direction

What does the intensity of a pixel tell us?


$$\text{im}(234, 452) = 0.58$$



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






Photo by nickwheeleroz, Flickr

Slide: Forsyth

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?






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

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






Perception of Intensity



Perception of Intensity



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 $\text{avg}(r) / \text{avg}((r+g+b)/3)$
 - White balancing: choose a reference as the white or gray color

Things to remember

- Important terms: diffuse/specular reflectance, albedo, umbra/penumbra
- 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