Radiometry and Photometric Cues

radiometric units, reflectance, photometric stereo

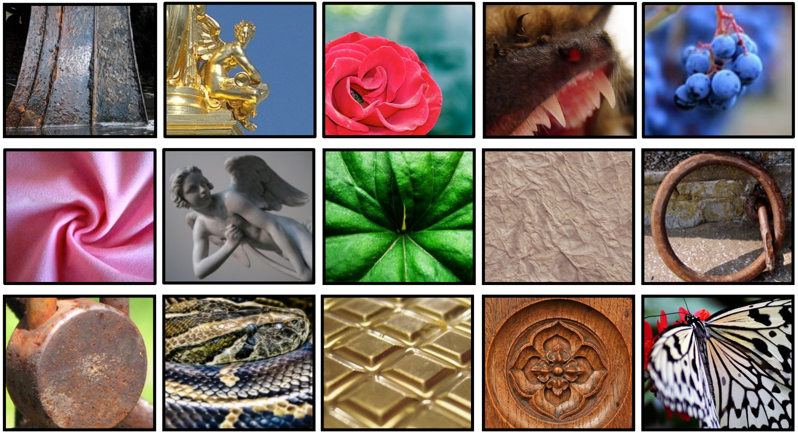

CSC420
David Lindell
University of Toronto
cs.toronto.edu/~lindell/teaching/420

Slide credit: Srinivasa Narasimhan, Todd Zickler, Steven Gortler, Yannis Gkioulekas (CMU 16-385)

Announcements

- HW 2 is out (due Oct 16)

Course Overview

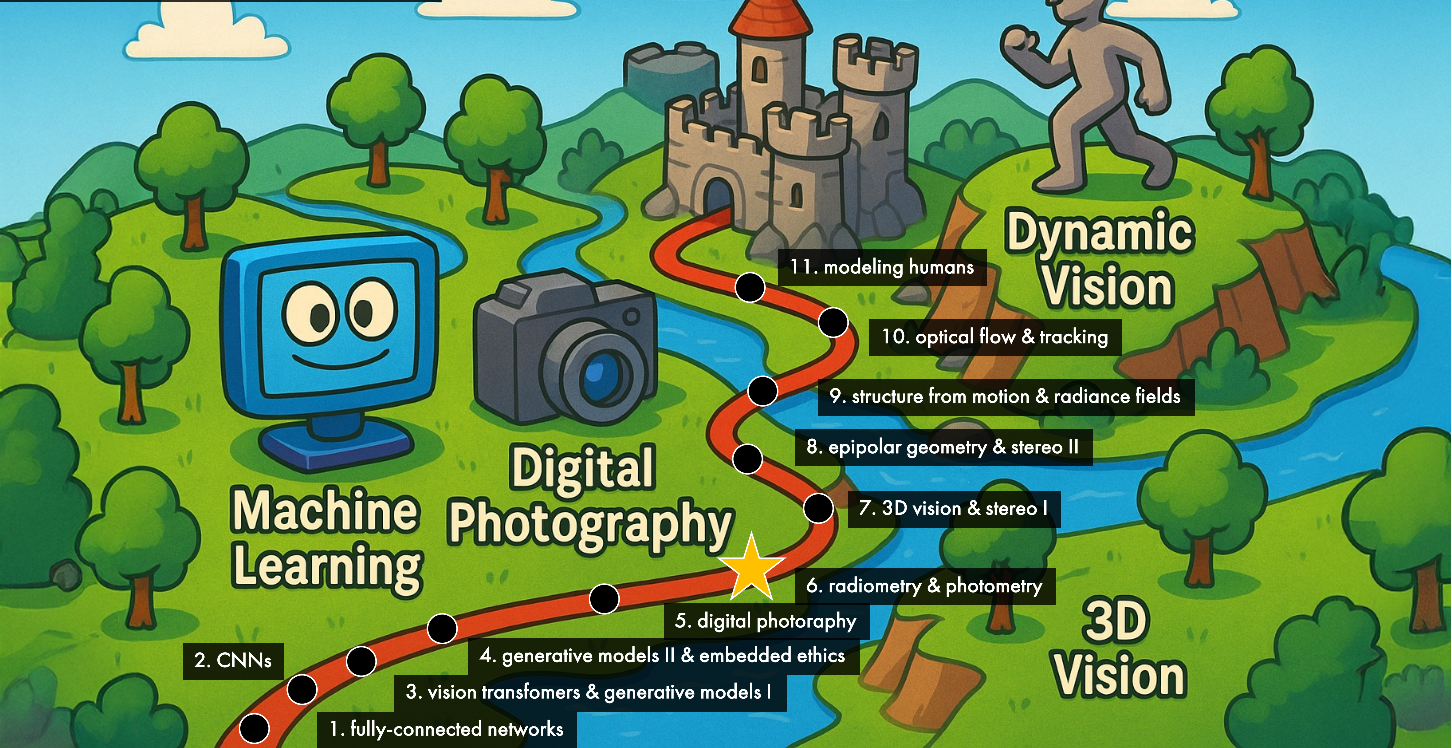

Overview of today’s lecture

- Appearance phenomena
- Measuring light and radiometry
- Reflectance and BRDF
- Light sources
- Photometric stereo

Appearance

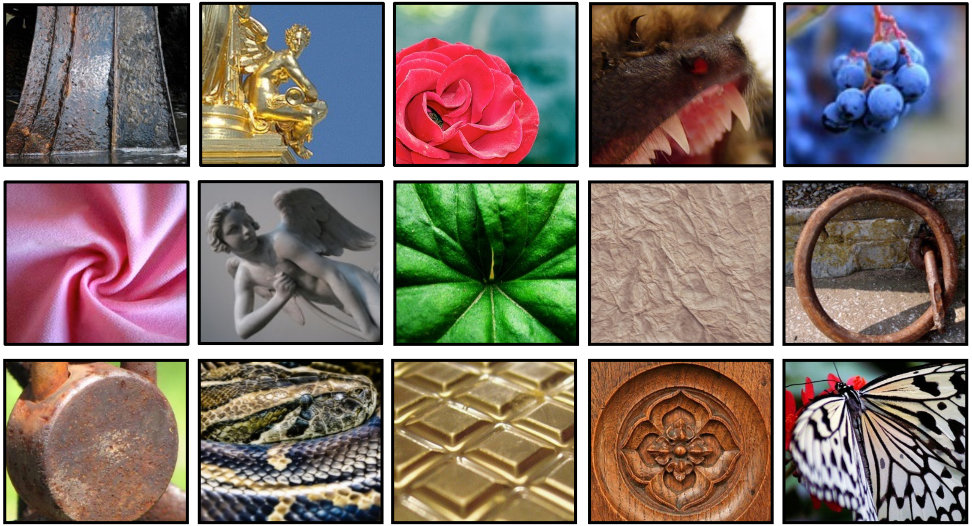

“Physics-based” computer vision
(a.k.a “inverse rendering”)

“Physics-based” computer vision
(a.k.a “inverse rendering”)

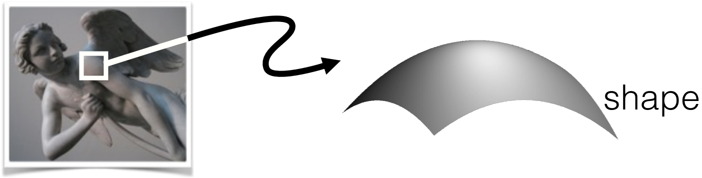

“Physics-based” computer vision
(a.k.a “inverse rendering”)

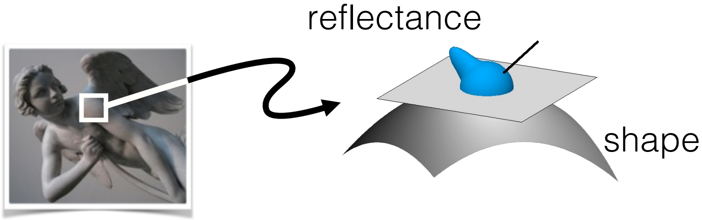

“Physics-based” computer vision (a.k.a “inverse rendering”)

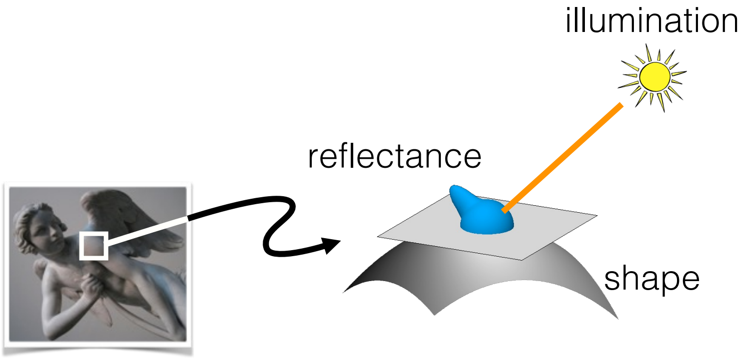

“Physics-based” computer vision
(a.k.a “inverse rendering”)

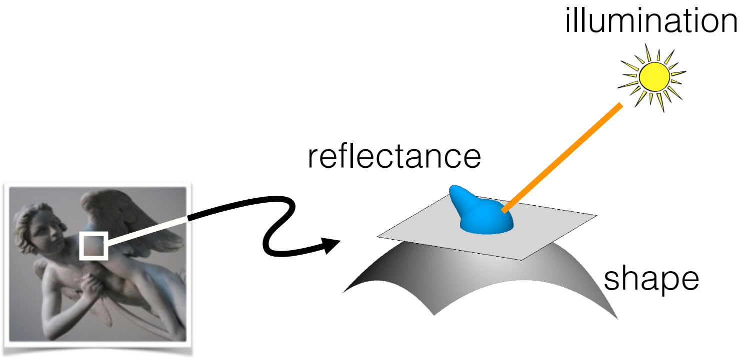

$\mathbf{I} \Longrightarrow$ shape, illumination, reflectance

“Physics-based” computer vision
(a.k.a “inverse rendering”)

Our challenge: Invent representations of shape, lighting, and reflectance that are efficient, yet general enough to capture the world’s most important phenomena

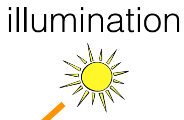

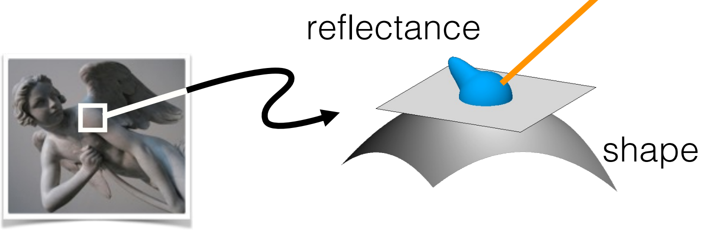

$\mathbf{I} \Longrightarrow$ shape, illumination, reflectance

Example application: Photometric Stereo

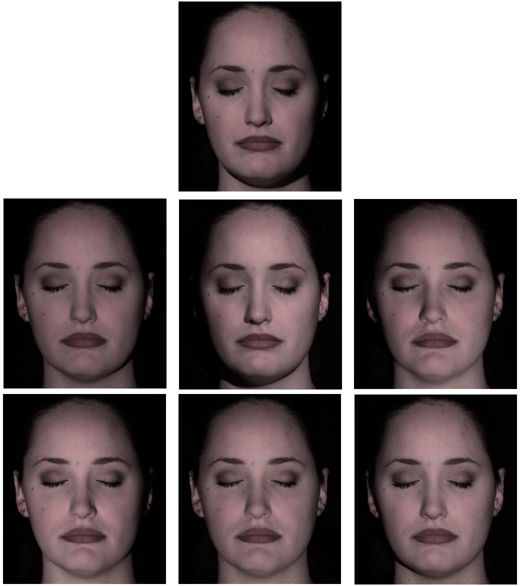

Example application: Photometric Stereo

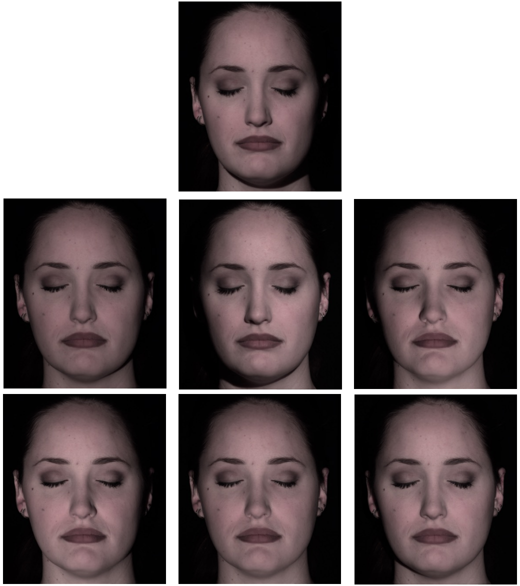

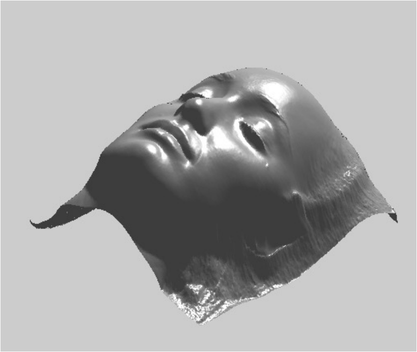

Why study the physics (optics) of the world?

Lets see some pictures!

Light and Shadows


Reflections


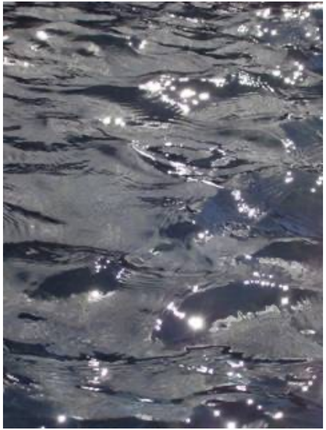

Refractions


HENRIK WANN JENSEN 1995
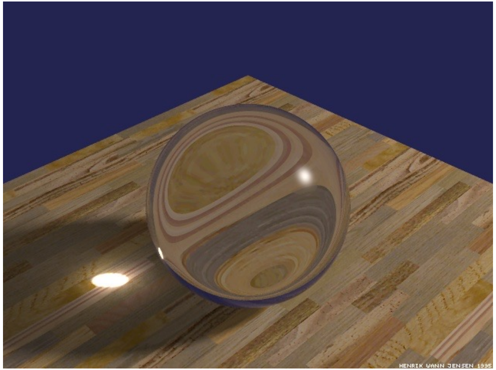
HENRIK WANN JENSEN 1995

Interreflections


Mies Courtyard House with Curved Elements

Scattering

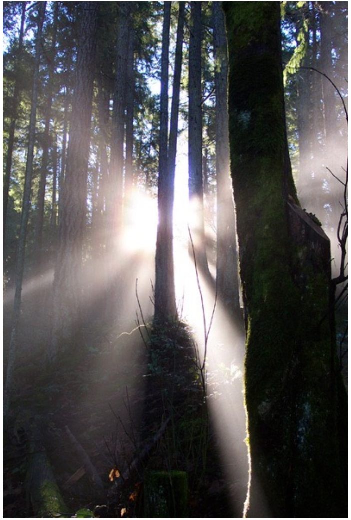

More Complex Appearances


RENDERED BY HENRIK WANN JENSEN - 2001

RENDERED BY HENRIK WANN JENSEN - 2001











Measuring light and radiometry

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

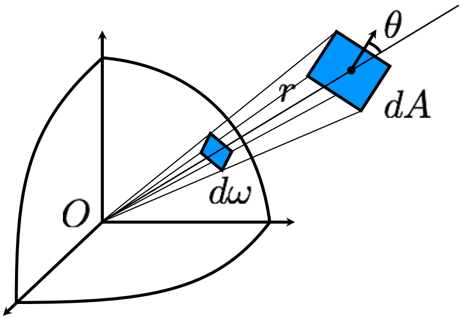

Depends on:

- orientation of patch
- distance of patch

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

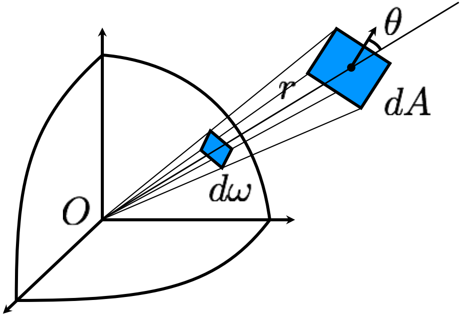

Depends on:

- orientation of patch
- distance of patch

$$d\omega = \; ? \quad dA$$

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

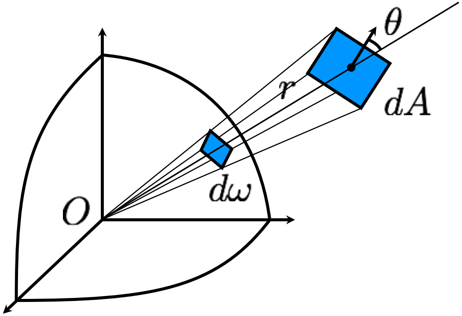

Depends on:

- orientation of patch
- distance of patch

$$d\omega = \ ? \quad dA$$

how does area change with r?

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

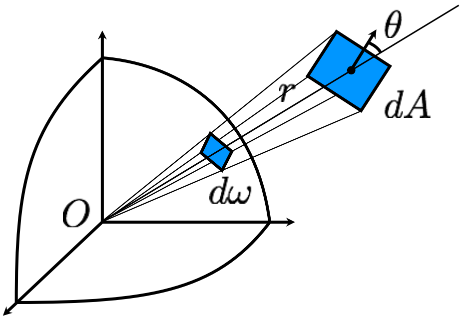

Depends on:

- orientation of patch
- distance of patch

$$d\omega = \;?\; dA$$

how does solid angle change with r?

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

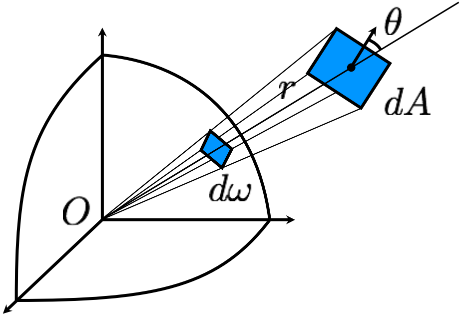

Depends on:

- orientation of patch
- distance of patch

$$d\omega = \ ? \quad dA$$

how does solid angle change with theta?

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

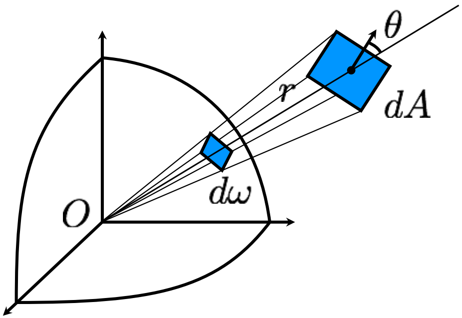

Depends on:

- orientation of patch
- distance of patch

One can show:

$$\boxed{d\omega = \frac{dA \cos\theta}{r^2}}$$

Units: steradians [sr]

Solid angle

- The solid angle subtended by a small surface patch with respect to point O is the area of its central projection onto the unit sphere about O

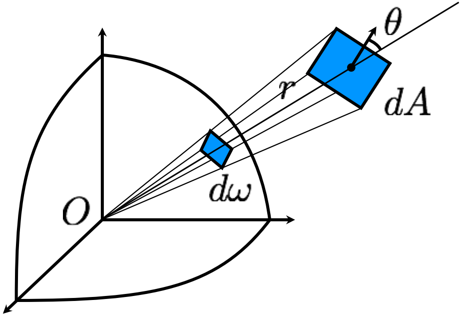

Depends on:

- orientation of patch
- distance of patch

One can show:

"surface foreshortening"

$$\boxed{d\omega = \frac{dA \cos\theta}{r^2}}$$

Units: steradians [sr]

Solid angle

- To calculate solid angle subtended by a surface S relative to O you must add up (integrate) contributions from all tiny patches (nasty integral)

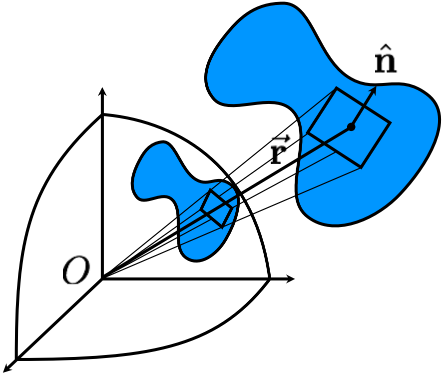

$$\Omega = \iint_S \frac{\vec{\mathbf{r}} \cdot \hat{\mathbf{n}}\ dS}{|\vec{\mathbf{r}}|^3}$$

One can show:

“surface foreshortening”

$$\boxed{d\omega = \frac{dA \cos\theta}{r^2}}$$

Units: steradians [sr]

Solid angle

- To calculate solid angle subtended by a surface S relative to O you must add up (integrate) contributions from all tiny patches (nasty integral)

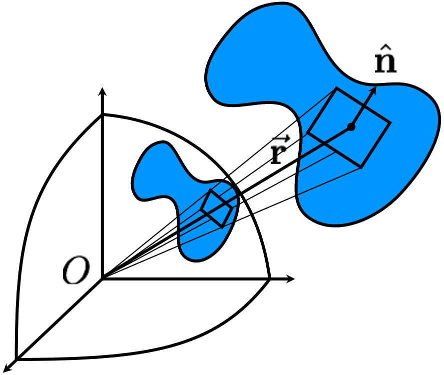

$$\Omega = \iint_S \frac{\vec{\mathbf{r}} \cdot \hat{\mathbf{n}} \; dS}{|\vec{\mathbf{r}}|^3}$$

combination of r^2 falloff and normalization

One can show:

"surface foreshortening"

$$\boxed{d\omega = \frac{dA \cos\theta}{r^2}}$$

Units: steradians [sr]

Question

- Suppose surface S is a hemisphere centered at O. What is the solid angle it subtends?

Question

- Suppose surface S is a hemisphere centered at O. What is the solid angle it subtends?

  Hint: what is the area of a sphere?

Question

- Suppose surface S is a hemisphere centered at O. What is the solid angle it subtends?

  Hint: what is the area of a sphere?

  Hint: what is theta?

Question

- Suppose surface S is a hemisphere centered at O. What is the solid angle it subtends?

  Hint: what is the area of a sphere?

  Hint: what is theta?

- Answer: 2 pi
  - area of sphere is 4 pi*r^2;
  - area of unit sphere is 4 pi
  - half of that is 2 pi

Quantifying light: flux, irradiance, and radiance

- Imagine a sensor that counts photons passing through planar patch X in directions within angular wedge W

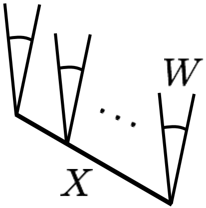

* shown in 2D for clarity; imagine three dimensions

radiant flux $\Phi(W, X)$

Quantifying light: flux, irradiance, and radiance

- Imagine a sensor that counts photons passing through planar patch X in directions within angular wedge W
- **It measures radiant flux [watts = joules/sec]**: rate of photons hitting sensor area

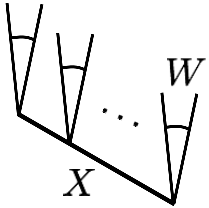

* shown in 2D for clarity; imagine three dimensions

radiant flux $\Phi(W, X)$

Quantifying light: flux, irradiance, and radiance

- Imagine a sensor that counts photons passing through planar patch X in directions within angular wedge W
- It measures radiant flux [watts = joules/sec]: rate of photons hitting sensor area
- Measurement depends on sensor area |X| and angular wedge W

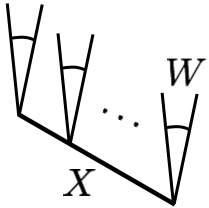

* shown in 2D for clarity; imagine three dimensions

radiant flux $\Phi(W, X)$

Quantifying light: flux, irradiance, and radiance

- Irradiance:
  A measure of incoming light that is independent of sensor area |X|

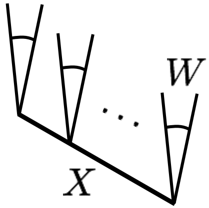

$$\frac{\Phi(W, X)}{|X|}$$

Quantifying light: flux, irradiance, and radiance

- **Irradiance:**
  A measure of incoming light that is independent of sensor area |X|

- Units: ?

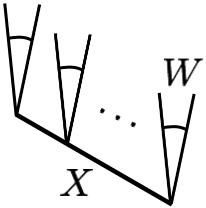

$$\frac{\Phi(W, X)}{|X|}$$

Quantifying light: flux, irradiance, and radiance

- Irradiance:
  A measure of incoming light that is independent of sensor area |X|
- Units: watts per square meter [W/m²]

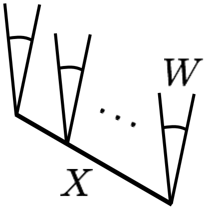

$$\frac{\Phi(W, X)}{|X|}$$

Quantifying light: flux, irradiance, and radiance

- **Irradiance:**
  A measure of incoming light that is independent of sensor area |X|
- Units: watts per square meter [W/m^2]

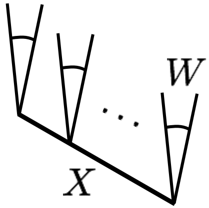

$$\frac{\Phi(W, X)}{|X|}$$

how do we define irradiance at a point?

Quantifying light: flux, irradiance, and radiance

- Irradiance:
  A measure of incoming light that is independent of sensor area |X|
- Units: watts per square meter [W/m^2]

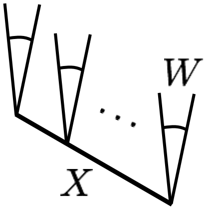

$$\frac{\Phi(W, X)}{|X|}$$

$$\lim_{X \longrightarrow x}$$

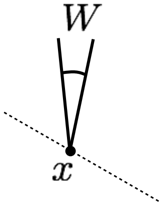

$$E_{\hat{\mathbf{n}}}(W, x)$$

Quantifying light: flux, irradiance, and radiance

- **Irradiance:**
  A measure of incoming light that is independent of sensor area |X|
- **Units: watts per square meter [W/m^2]**
- Depends on sensor direction normal.

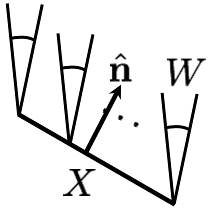

$$\frac{\Phi(W, X)}{|X|}$$

$$\lim_{X \longrightarrow x}$$

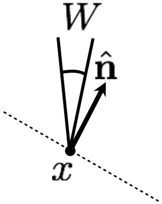

$$E_{\hat{\mathbf{n}}}(W, x)$$

- We keep track of the normal because a planar sensor with distinct orientation would converge to a different limit
- In the literature, notations n and W are often omitted, and values are implied by context

Quantifying light: flux, irradiance, and radiance

- Radiance:
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|

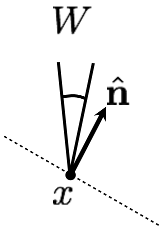

$$\frac{E_{\hat{\mathbf{n}}}(W, x)}{|W|}$$

Quantifying light: flux, irradiance, and radiance

- **Radiance:**
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|

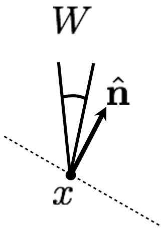

$$\frac{E_{\hat{\mathbf{n}}}(W, x)}{|W|}$$

how do we define
radiance at a single direction
(not wedge)?

Quantifying light: flux, irradiance, and radiance

- Radiance:
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|

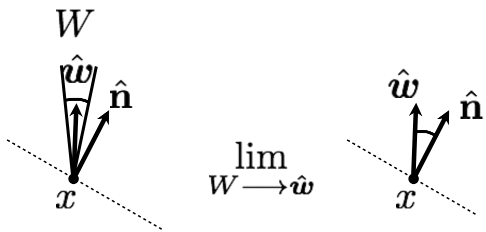

$$\frac{E_{\hat{\mathbf{n}}}(W, x)}{|W|} \qquad L_{\hat{\mathbf{n}}}(\hat{\boldsymbol{\omega}}, x)$$

Quantifying light: flux, irradiance, and radiance

- **Radiance:**
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|
- **Units: watts per steradian per square meter [W/(m^2·sr)]**

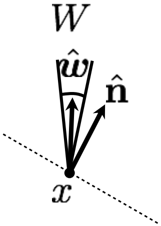

$$\lim_{W \longrightarrow \hat{\boldsymbol{w}}}$$

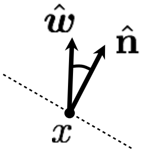

$$\frac{E_{\hat{\mathbf{n}}}(W, x)}{|W|} \qquad L_{\hat{\mathbf{n}}}(\hat{\boldsymbol{\omega}}, x)$$

- Has correct units, but still depends on sensor orientation

Quantifying light: flux, irradiance, and radiance

- **Radiance:**
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|
- **Units: watts per steradian per square meter [W/(m²·sr)]**

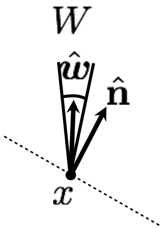

$$\lim_{W \longrightarrow \hat{w}}$$

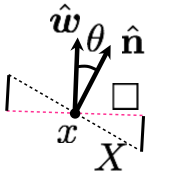

$$\cos\theta = \frac{\square/2}{|X|/2}$$

$$\longrightarrow \square = |X|\cos\theta$$

"foreshortened area"

$$\frac{E_{\hat{\mathbf{n}}}(W, x)}{|W|}$$

$$L_{\hat{\mathbf{n}}}(\hat{\omega}, x)$$

- Has correct units, but still depends on sensor orientation
- To correct this, convert to measurement that would have been made if sensor was perpendicular to direction ω

Quantifying light: flux, irradiance, and radiance

- **Radiance:**
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|

- **Units: watts per steradian per square meter [W/(m²·sr)]**

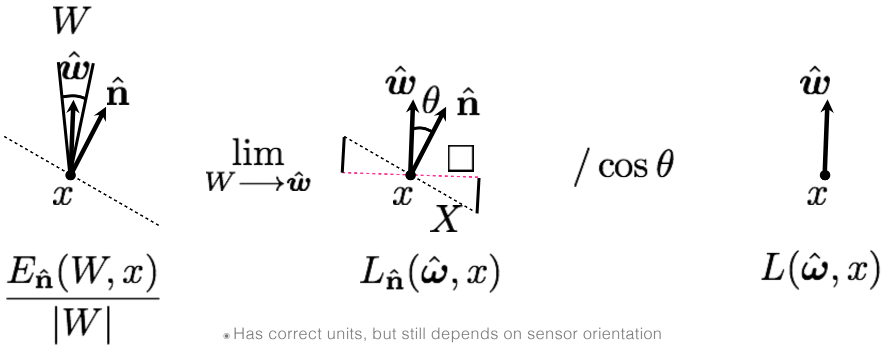

- Has correct units, but still depends on sensor orientation
- To correct this, convert to measurement that would have been made if sensor was perpendicular to direction ω

Quantifying light: flux, irradiance, and radiance

- **Radiance:**
  A measure of incoming light that is independent of sensor area |X|, orientation n, and wedge size (solid angle) |W|

- **Units: watts per steradian per square meter [W/(m²·sr)]**

"foreshortened in the direction of travel"

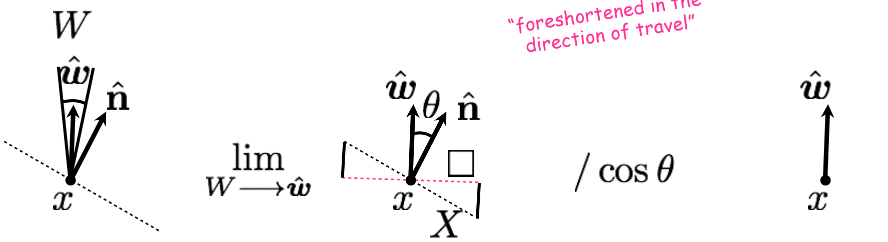

$$\frac{E_{\hat{\mathbf{n}}}(W, x)}{|W|} \qquad L_{\hat{\mathbf{n}}}(\hat{\omega}, x) \qquad L(\hat{\omega}, x)$$

- Has correct units, but still depends on sensor orientation
- To correct this, convert to measurement that would have been made if sensor was perpendicular to direction ω

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

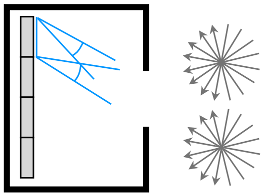

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

$$\Phi(W, X) = \int_X \int_W L(\hat{\omega}, x) \cos\theta \, d\omega \, dA$$

?

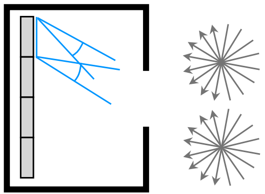

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

$$\Phi(W, X) = \int_X \int_W L(\hat{\omega}, x) \cos\theta d\omega dA$$

?

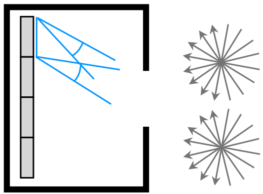

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

$$\Phi(W, X) = \int_X \int_W L(\hat{\omega}, x) \cos\theta \, d\omega \, dA$$

integral over which dimensions?

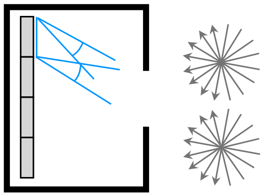

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

$$\Phi(W, X) = \int_X \int_W L(\hat{\omega}, x) \cos\theta d\omega dA$$

units?

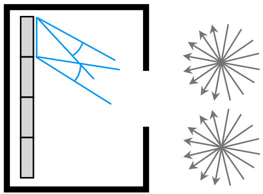

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

$$\Phi(W, X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta d\boldsymbol{\omega} dA$$

  - Constant along a ray in free space

$$L(\hat{\boldsymbol{\omega}}, x) = L(\hat{\boldsymbol{\omega}}, x + \hat{\boldsymbol{\omega}})$$

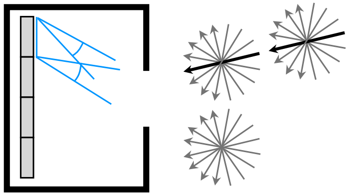

Quantifying light: flux, irradiance, and radiance

- Attractive properties of radiance:
  - Allows computing the radiant flux measured by any finite sensor

$$\Phi(W, X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta \, d\boldsymbol{\omega} \, dA$$

  - Constant along a ray in free space

$$L(\hat{\boldsymbol{\omega}}, x) = L(\hat{\boldsymbol{\omega}}, x + \hat{\boldsymbol{\omega}})$$

  - A camera measures radiance (after a one-time radiometric calibration). So RAW pixel values are proportional to radiance.

    - "Processed" images (like PNG and JPEG) are not linear radiance measurements!!

Question

- Most light sources, like a heated metal sheet, follow Lambert's Law

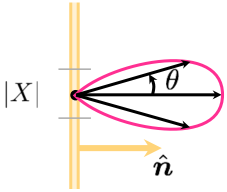

"Lambertian area source"

$$J(\hat{\boldsymbol{\omega}}) = J_o \langle \hat{\boldsymbol{\omega}}, \hat{\boldsymbol{n}} \rangle = J_o \cos\theta$$

radiant intensity [W/sr]

- What is the radiance $L(\hat{\boldsymbol{\omega}}, \boldsymbol{x})$ of an infinitesimal patch [W/sr・m2]?

Question

- Most light sources, like a heated metal sheet, follow Lambert's Law

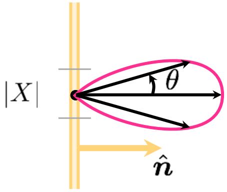

"Lambertian area source"

$$J(\hat{\boldsymbol{\omega}}) = J_o \langle \hat{\boldsymbol{\omega}}, \hat{\boldsymbol{n}} \rangle = J_o \cos\theta$$

radiant intensity [W/sr]

- What is the radiance $L(\hat{\boldsymbol{\omega}}, \boldsymbol{x})$ of an infinitesimal patch [W/sr・m2]?

hint: how do units of radiant intensity differ from radiance

Question

- Most light sources, like a heated metal sheet, follow Lambert's Law

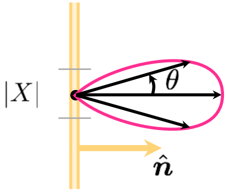

"Lambertian area source"

$$J(\hat{\boldsymbol{\omega}}) = J_o \langle \hat{\boldsymbol{\omega}}, \hat{\boldsymbol{n}} \rangle = J_o \cos\theta$$

radiant intensity [W/sr]

- What is the radiance $L(\hat{\boldsymbol{\omega}}, \boldsymbol{x})$ of an infinitesimal patch [W/sr・m2]?

hint: how do units of radiant intensity differ from radiance
hint: what is the projected area?

Question

- Most light sources, like a heated metal sheet, follow Lambert's Law

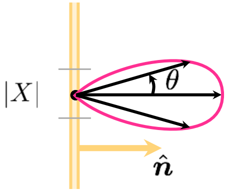

"Lambertian area source"

$$J(\hat{\boldsymbol{\omega}}) = J_o \langle \hat{\boldsymbol{\omega}}, \hat{\boldsymbol{n}} \rangle = J_o \cos\theta$$

radiant intensity [W/sr]

- What is the radiance $L(\hat{\boldsymbol{\omega}}, \boldsymbol{x})$ of an infinitesimal patch [W/sr・m2]?

Answer: $L(\hat{\boldsymbol{\omega}}, \boldsymbol{x}) = \dfrac{J_0 \cos\theta}{|X| \cos\theta} = J_0 / |X|$ (independent of direction)

Question

- Most light sources, like a heated metal sheet, follow Lambert's Law

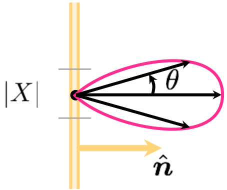

"Lambertian area source"

$$J(\hat{\boldsymbol{\omega}}) = J_o \langle \hat{\boldsymbol{\omega}}, \hat{\boldsymbol{n}} \rangle = J_o \cos\theta$$

radiant intensity [W/sr]

- What is the radiance $L(\hat{\boldsymbol{\omega}}, \boldsymbol{x})$ of an infinitesimal patch [W/sr・m2]?

Answer: $$L(\hat{\boldsymbol{\omega}}, \boldsymbol{x}) = \frac{J_0 \cos\theta}{|X| \cos\theta} = J_0 / |X|$$ (independent of direction)

"Looks equally bright when viewed from any direction"

Radiometric concepts – boring…but, important!

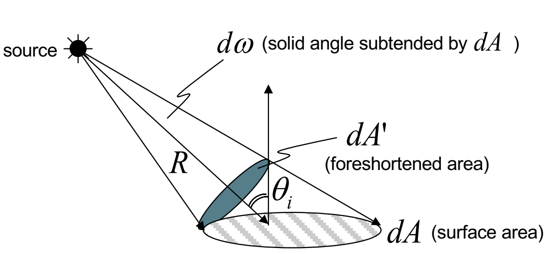

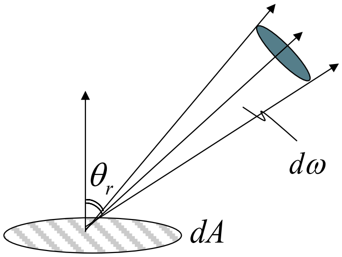

(1) Solid Angle : $d\omega = \frac{dA'}{R^2} = \frac{dA \cos\theta_i}{R^2}$ (steradian)

What is the solid angle subtended by a hemisphere?

(2) Radiant Intensity of Source : $J = \frac{d\Phi}{d\omega}$ (watts / steradian)

Light Flux (power) emitted per unit solid angle

(3) Surface Irradiance : $E = \frac{d\Phi}{dA}$ (watts / m)

Light Flux (power) incident per unit surface area.

Does not depend on where the light is coming from!

(4) Surface Radiance (tricky) :

$$L = \frac{d^2\Phi}{(dA \cos\theta_r)\ d\omega}$$ (watts / m^2 steradian)

- Flux emitted per unit foreshortened area per unit solid angle.
- L depends on direction θ_r
- Surface can radiate into whole hemisphere.
- L depends on reflectance properties of surface.

Appearance

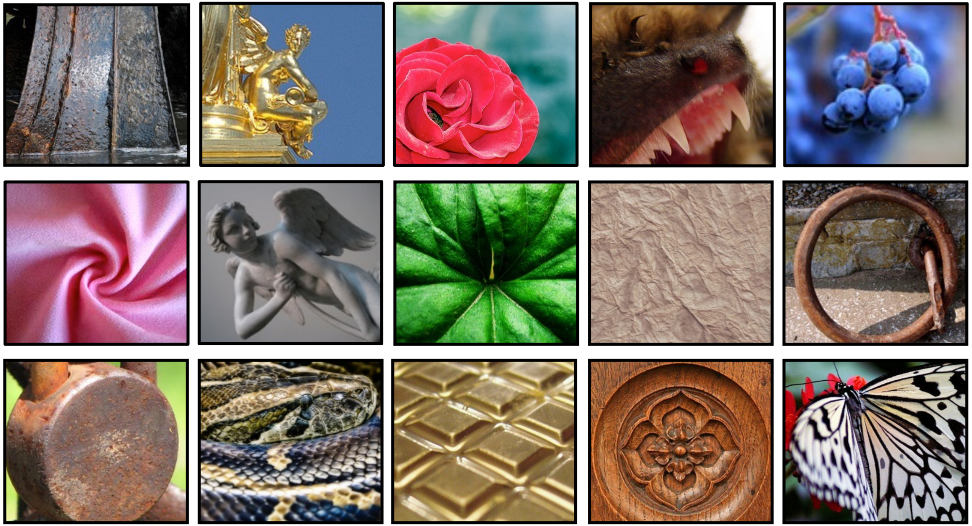

“Physics-based” computer vision
(a.k.a “inverse optics”)

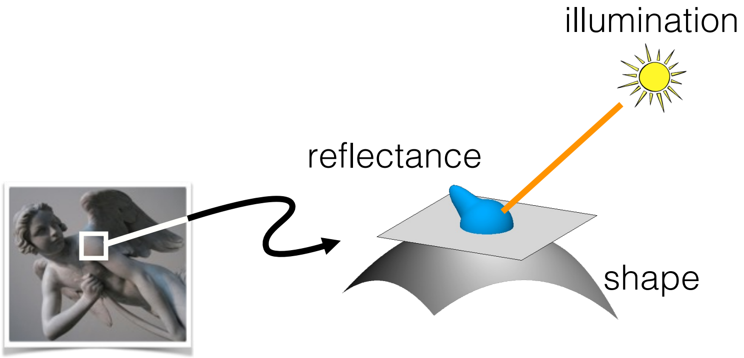

$\mathbf{I} \Longrightarrow$ shape, illumination, reflectance

Reflectance and BRDF

Reflectance

- Ratio of outgoing energy to incoming energy at a single point

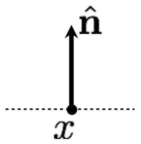

Reflectance

- Ratio of outgoing energy to incoming energy at a single point

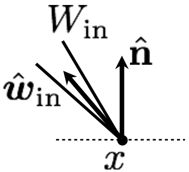

Reflectance

- Ratio of outgoing energy to incoming energy at a single point

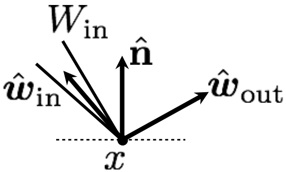

Reflectance

- Ratio of outgoing energy to incoming energy at a single point

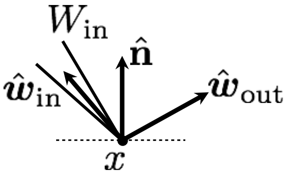

$$\frac{L^{\text{out}}(x, \hat{\boldsymbol{\omega}}_{\text{out}})}{E^{\text{in}}_{\hat{\mathbf{n}}}(W_{\text{in}}, x)}$$

Reflectance

- Ratio of outgoing energy to incoming energy at a single point
- Want to define a ratio such that it:
  - converges as we use smaller and smaller incoming and outgoing wedges

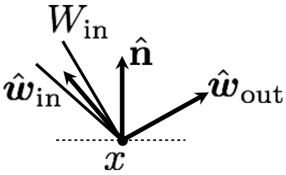

$$\frac{L^{\text{out}}(x, \hat{\boldsymbol{\omega}}_{\text{out}})}{E^{\text{in}}_{\hat{\mathbf{n}}}(W_{\text{in}}, x)}$$

Reflectance

- Ratio of outgoing energy to incoming energy at a single point
- Want to define a ratio such that it:
  - converges as we use smaller and smaller incoming and outgoing wedges
  - does not depend on the size of the wedges (i.e. is intrinsic to the material)

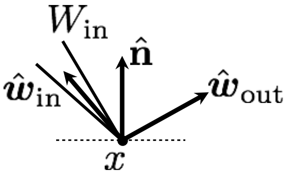

$$\frac{L^{\text{out}}(x, \hat{\boldsymbol{\omega}}_{\text{out}})}{E^{\text{in}}_{\hat{\mathbf{n}}}(W_{\text{in}}, x)}$$

Reflectance

- Ratio of outgoing energy to incoming energy at a single point
- Want to define a ratio such that it:
    - converges as we use smaller and smaller incoming and outgoing wedges
    - does not depend on the size of the wedges (i.e. is intrinsic to the material)

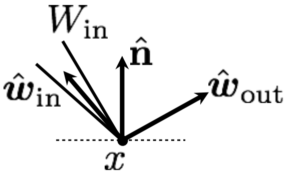

$$\lim_{W_{\text{in}} \longrightarrow \hat{\boldsymbol{\omega}}_{\text{in}}}$$

$$f_{x,\hat{\mathbf{n}}}(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}})$$

$$\frac{L^{\text{out}}(x, \hat{\boldsymbol{\omega}}_{\text{out}})}{E^{\text{in}}_{\hat{\mathbf{n}}}(W_{\text{in}}, x)}$$

- Notations x and n often implied by context and omitted; directions \omega are expressed in local coordinate system defined by normal n (and some chosen tangent vector)
- Units: sr^{-1}
- Called Bidirectional Reflectance Distribution Function (BRDF)

BRDF: Bidirectional Reflectance Distribution Function

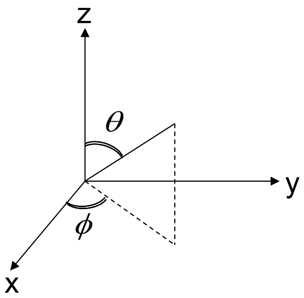

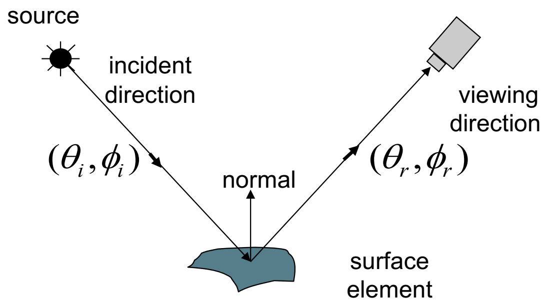

$E^{surface}(\theta_i, \phi_i)$ Irradiance at Surface in direction (θ_i, ϕ_i)

$L^{surface}(\theta_r, \phi_r)$ Radiance of Surface in direction (θ_r, ϕ_r)

$$\text{BRDF}: f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{L^{surface}(\theta_r, \phi_r)}{E^{surface}(\theta_i, \phi_i)}$$

Reflectance: BRDF

- Units: sr^{-1}
- Real-valued function defined on the double-hemisphere
- Has many useful properties

Important Properties of BRDFs

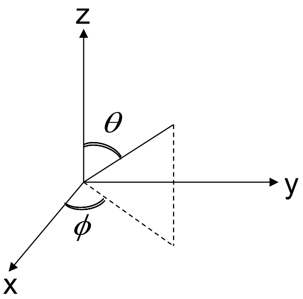

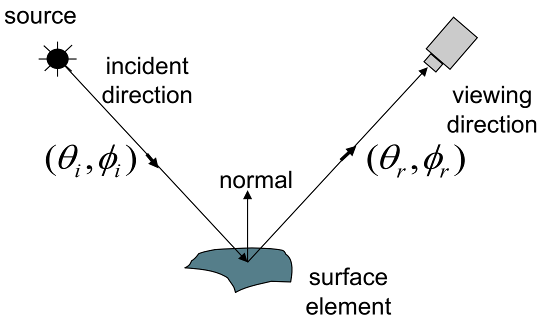

- Conservation of Energy:

$$\forall \hat{\boldsymbol{\omega}}_{\text{in}}, \quad \int_{\Omega_{\text{out}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) \cos\theta_{\text{out}} d\hat{\boldsymbol{\omega}}_{\text{out}} \leq 1$$

Important Properties of BRDFs

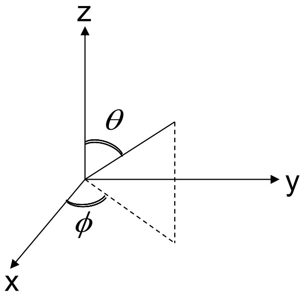

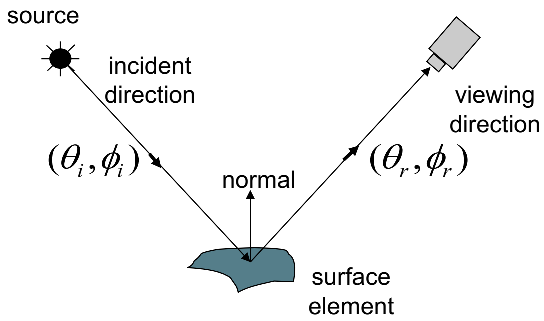

- Conservation of Energy:

$$\forall \hat{\boldsymbol{\omega}}_{\text{in}}, \quad \int_{\Omega_{\text{out}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) \cos\theta_{\text{out}} d\hat{\boldsymbol{\omega}}_{\text{out}} \leq 1$$

Why?

Property: “Helmholtz reciprocity”

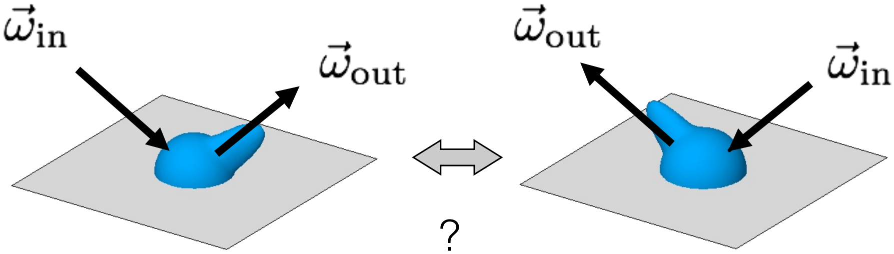

Property: "Helmholtz reciprocity"

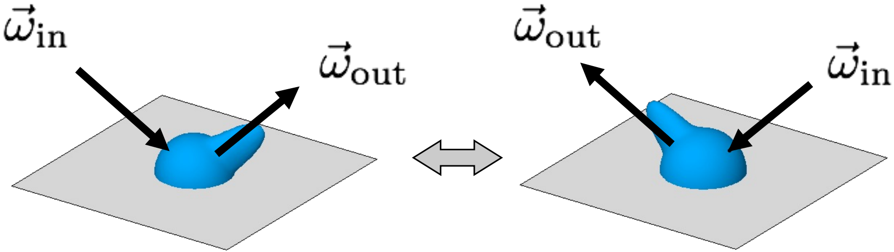

- Helmholtz Reciprocity: (follows from 2nd Law of Thermodynamics)

BRDF does not change when source and viewing directions are swapped.

$$f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}}) = f_r(\vec{\omega}_{\text{out}}, \vec{\omega}_{\text{in}})$$

Common assumption: Isotropy

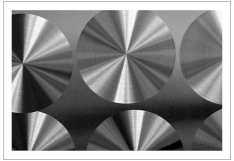

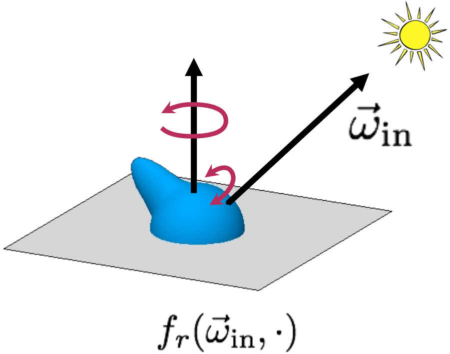

Common assumption: Isotropy

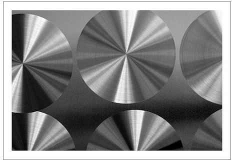

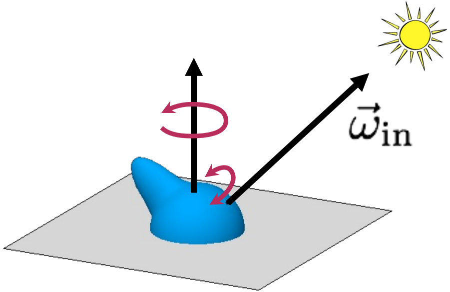

$$f_r(\vec{\omega}_{\text{in}}, \cdot)$$

BRDF does not change when surface is rotated about the normal.

4D → 3D

$$\boxed{f_r(\cancel{\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

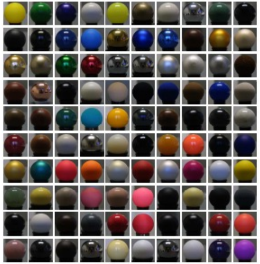

[Matusik et al., 2003]

Can be written as a function of 3 variables : $f(\theta_i, \theta_r, \phi_i - \phi_r)$

Reflectance: BRDF

- Units: sr^{-1}
- Real-valued function defined on the double-hemisphere
- Has many useful properties
- Allows computing output radiance (and thus pixel value) for any configuration of lights and viewpoint

Reflectance: BRDF

- Units: sr^{-1}
- Real-valued function defined on the double-hemisphere
- Has many useful properties
- Allows computing output radiance (and thus pixel value) for any configuration of lights and viewpoint

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

reflectance equation

Reflectance: BRDF

- Units: sr^{-1}
- Real-valued function defined on the double-hemisphere
- Has many useful properties
- Allows computing output radiance (and thus pixel value) for any configuration of lights and viewpoint

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

reflectance equation

- Why is there a cosine in the reflectance equation?

Derivation of the Reflectance Equation

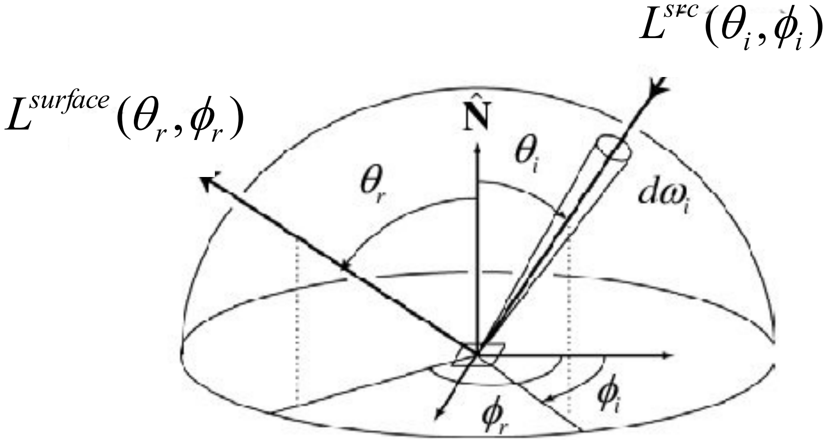

Derivation of the Reflectance Equation

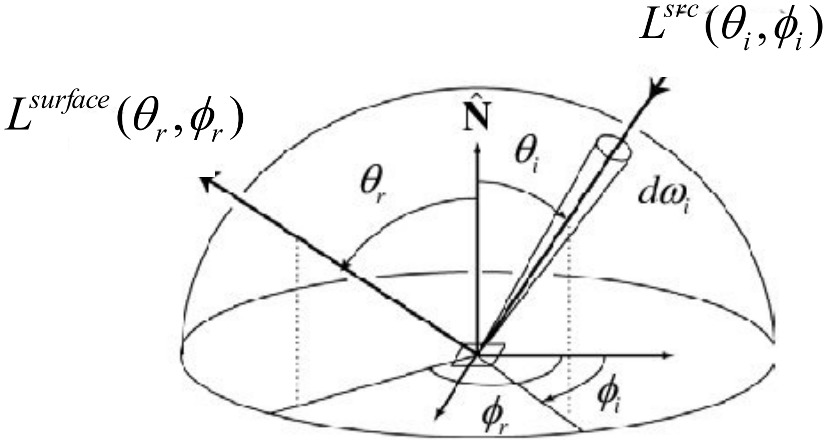

From the definition of BRDF:

$$L^{surface}(\theta_r, \phi_r) \quad = \quad E^{surface}(\theta_i, \phi_i) f(\theta_i, \phi_i; \theta_r, \phi_r)$$

Derivation of the Scene Radiance Equation

From the definition of BRDF:

$$L^{surface}(\theta_r,\phi_r) = \underline{E^{surface}(\theta_i,\phi_i)}\, f(\theta_i,\phi_i;\theta_r,\phi_r)$$

Derivation of the Scene Radiance Equation

From the definition of BRDF:

$$L^{surface}(\theta_r, \phi_r) \quad = \quad \underline{E^{surface}(\theta_i, \phi_i)} f(\theta_i, \phi_i; \theta_r, \phi_r)$$

Write Surface Irradiance in terms of Source Radiance:

$$L^{surface}(\theta_r, \phi_r) \quad = \quad ?$$

Derivation of the Scene Radiance Equation

From the definition of BRDF:

$$L^{surface}(\theta_r,\phi_r) \quad = \quad \underline{E^{surface}(\theta_i,\phi_i)}\,f(\theta_i,\phi_i;\theta_r,\phi_r)$$

Write Surface Irradiance in terms of Source Radiance:

$$L^{surface}(\theta_r,\phi_r) \quad = \quad \underline{L^{src}(\theta_i,\phi_i)}\,f(\theta_i,\phi_i;\theta_r,\phi_r)\underbrace{\cos\theta_i d\omega_i}_{\text{foreshortening factor}}$$

Derivation of the Scene Radiance Equation

From the definition of BRDF:

$$L^{surface}(\theta_r, \phi_r) \quad = \quad \underline{E^{surface}(\theta_i, \phi_i)}\, f(\theta_i, \phi_i; \theta_r, \phi_r)$$

Write Surface Irradiance in terms of Source Radiance:

$$L^{surface}(\theta_r, \phi_r) \quad = \quad \underline{L^{src}(\theta_i, \phi_i)}\, f(\theta_i, \phi_i; \theta_r, \phi_r)\, \underline{\cos\theta_i d\omega}_i$$

Integrate over entire hemisphere of possible source directions:

$$L^{surface}(\theta_r, \phi_r) \quad = \quad \int_{2\pi} L^{src}(\theta_i, \phi_i)\, f(\theta_i, \phi_i; \theta_r, \phi_r) \cos\theta_i \, \underline{d\omega_i}$$

Derivation of the Scene Radiance Equation

From the definition of BRDF:

$$L^{surface}(\theta_r,\phi_r) = \underline{E^{surface}(\theta_i,\phi_i)}\,f(\theta_i,\phi_i;\theta_r,\phi_r)$$

Write Surface Irradiance in terms of Source Radiance:

$$L^{surface}(\theta_r,\phi_r) = \underline{L^{src}(\theta_i,\phi_i)}\,f(\theta_i,\phi_i;\theta_r,\phi_r)\,\underline{\cos\theta_i d\omega_i}$$

Integrate over entire hemisphere of possible source directions:

$$L^{surface}(\theta_r,\phi_r) = \int_{2\pi} L^{src}(\theta_i,\phi_i)\,f(\theta_i,\phi_i;\theta_r,\phi_r)\cos\theta_i\,\underline{d\omega_i}$$

Convert from solid angle to theta-phi representation:

$$L^{surface}(\theta_r,\phi_r) = \int_{-\pi}^{\pi}\int_{0}^{\pi/2} L^{src}(\theta_i,\phi_i)\,f(\theta_i,\phi_i;\theta_r,\phi_r)\cos\theta_i\,\underline{\sin\theta_i d\theta_i d\phi_i}$$

Differential Solid Angles

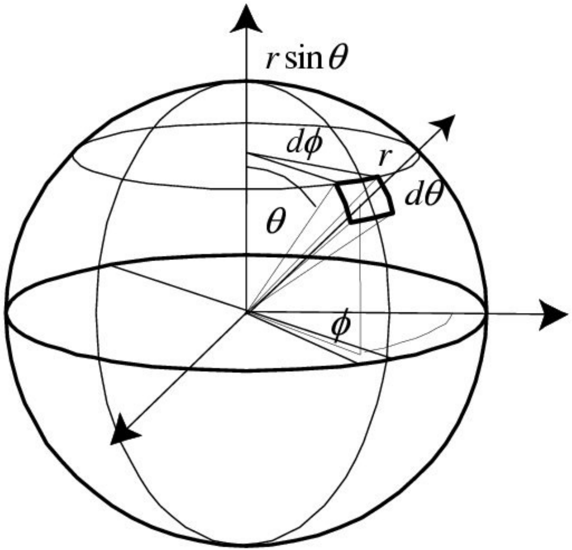

$$dA = (r\,d\theta)(r\sin\theta\,d\phi)$$
$$= r^2\sin\theta\,d\theta\,d\phi$$

$$d\omega = \frac{dA}{r^2} = \sin\theta\,d\theta\,d\phi$$

$$S = \int_0^{\pi}\int_0^{2\pi} \sin\theta\,d\theta\,d\phi = 4\pi$$

BRDF

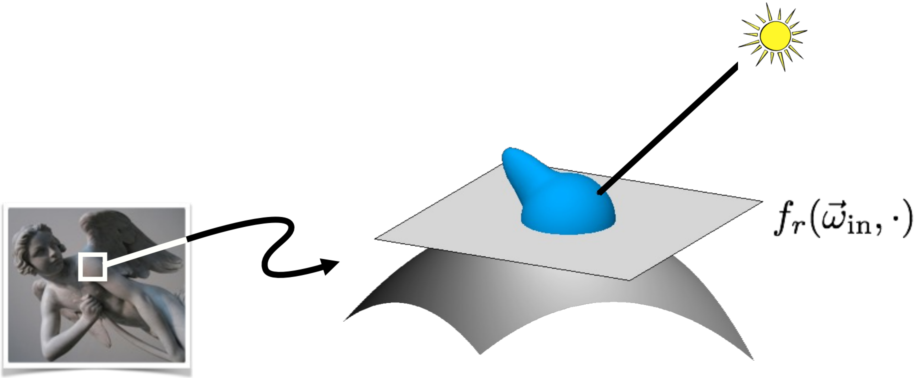

$$\boxed{f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

BRDF

Lambertian (diffuse) BRDF: energy equally distributed in all directions

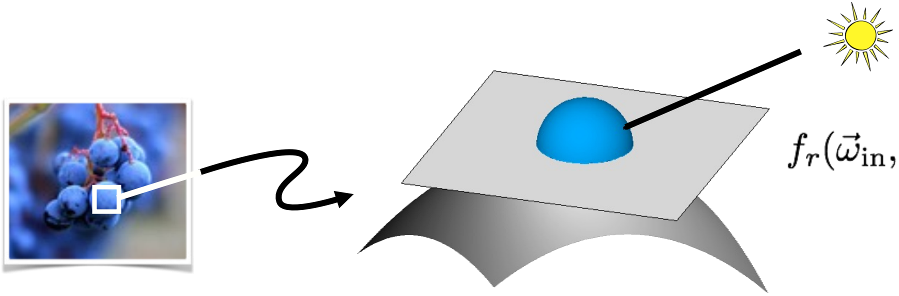

$$f_r(\vec{\omega}_{\text{in}}, \cdot)$$

$$\boxed{f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

BRDF

Lambertian (diffuse) BRDF: energy equally distributed in all directions

What does the BRDF equal in this case?

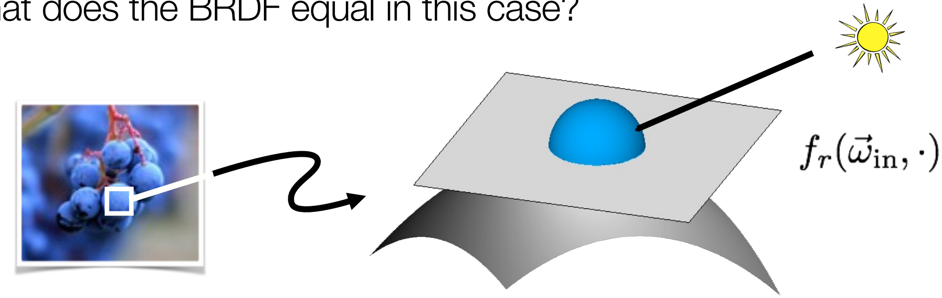

$$\boxed{f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

Diffuse Reflection and Lambertian BRDF

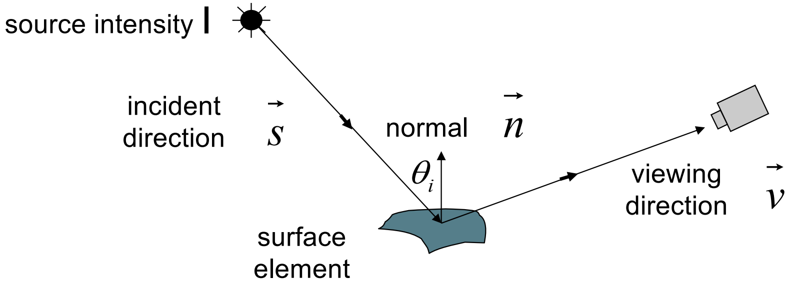

- Surface appears equally bright from ALL directions! (independent of $\vec{v}$)

Diffuse Reflection and Lambertian BRDF

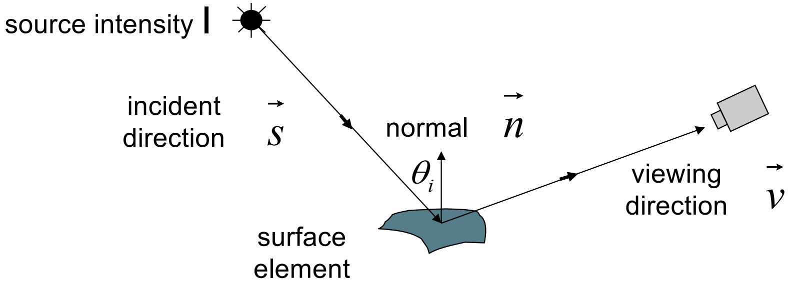

- Surface appears equally bright from ALL directions! (independent of $\vec{v}$)

- Lambertian BRDF is simply a constant : $f(\theta_i, \phi_i; \theta_r, \phi_r) = \dfrac{\rho_d}{\pi}$ albedo

Diffuse Reflection and Lambertian BRDF

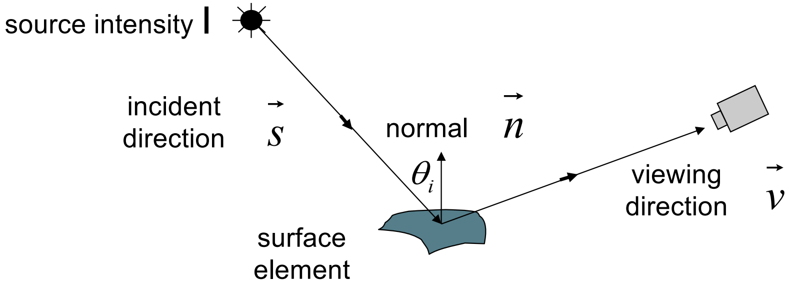

- Surface appears equally bright from ALL directions! (independent of $\vec{v}$)

- Lambertian BRDF is simply a constant : $f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\rho_d}{\pi}$ albedo

- Most commonly used BRDF in Vision and Graphics!

BRDF

Specular BRDF: all energy concentrated in mirror direction

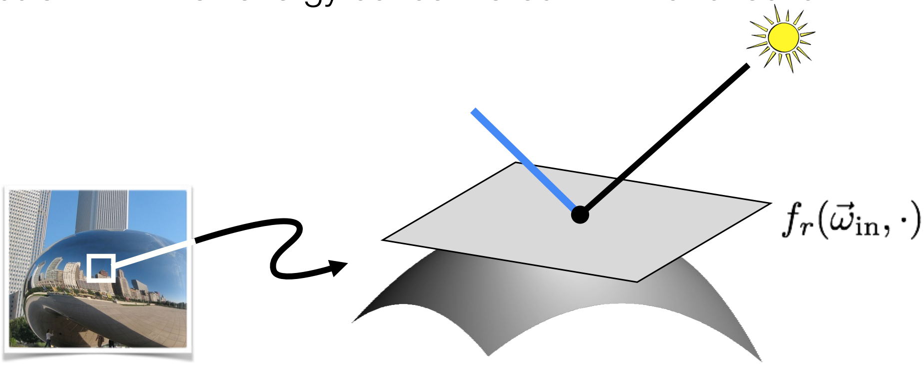

$$\boxed{f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

BRDF

Specular BRDF: all energy concentrated in mirror direction

What does the BRDF equal in this case?

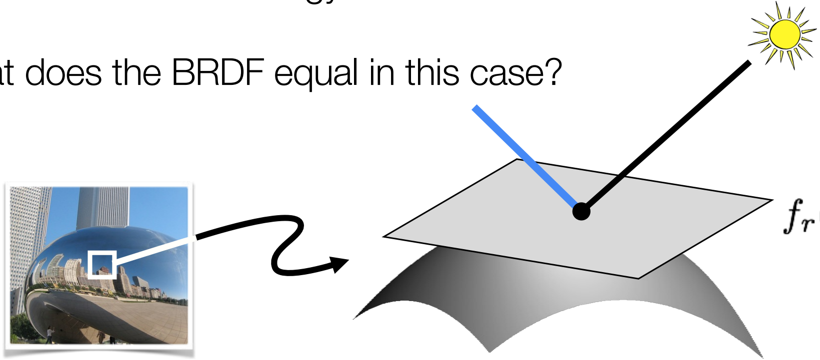

$$\boxed{f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

Specular Reflection and Mirror BRDF

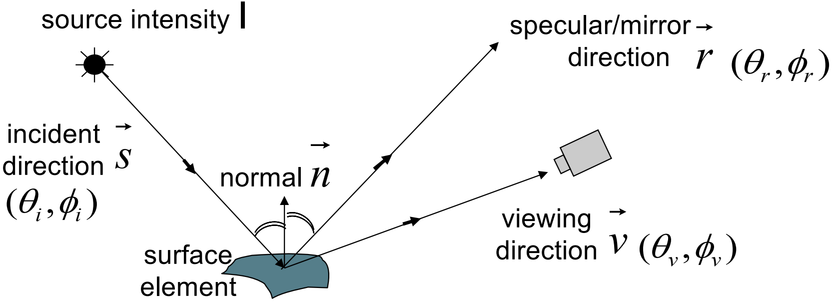

- Valid for very smooth surfaces.
- All incident light energy reflected in a SINGLE direction (only when $\vec{v} = \vec{r}$).

Specular Reflection and Mirror BRDF

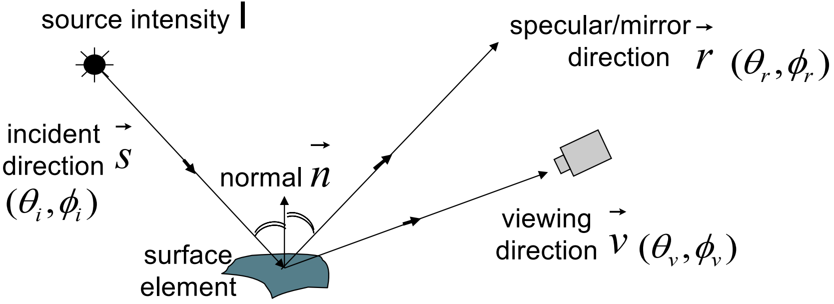

- Valid for very smooth surfaces.

- All incident light energy reflected in a SINGLE direction (only when $\vec{v} = \vec{r}$).

- Mirror BRDF is simply a double-delta function :

specular albedo

$$f(\theta_i, \phi_i; \theta_v, \phi_v) = \rho_s \; \delta(\theta_i - \theta_v)\, \delta(\phi_i + \pi - \phi_v)$$

Specular Reflection and Mirror BRDF

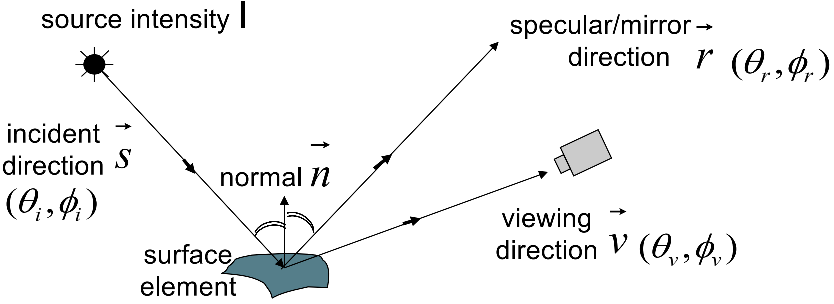

- Valid for very smooth surfaces.

- All incident light energy reflected in a SINGLE direction (only when $\vec{v} = \vec{r}$).

- Mirror BRDF is simply a double-delta function :

specular albedo

$$f(\theta_i, \phi_i; \theta_v, \phi_v) = \rho_s \; \delta(\theta_i - \theta_v) \; \delta(\phi_i + \pi - \phi_v)$$

Example Surfaces

Body Reflection:

Diffuse Reflection
Matte Appearance
Non-Homogeneous Medium
Clay, paper, etc

Example Surfaces

Body Reflection:

Diffuse Reflection
Matte Appearance
Non-Homogeneous Medium
Clay, paper, etc

Surface Reflection:

Specular Reflection
Glossy Appearance
Highlights
Dominant for Metals

Example Surfaces

Body Reflection:

Diffuse Reflection
Matte Appearance
Non-Homogeneous Medium
Clay, paper, etc

Surface Reflection:

Specular Reflection
Glossy Appearance
Highlights
Dominant for Metals

Many materials exhibit both Reflections:

BRDF

Glossy BRDF: more energy concentrated in mirror direction than elsewhere

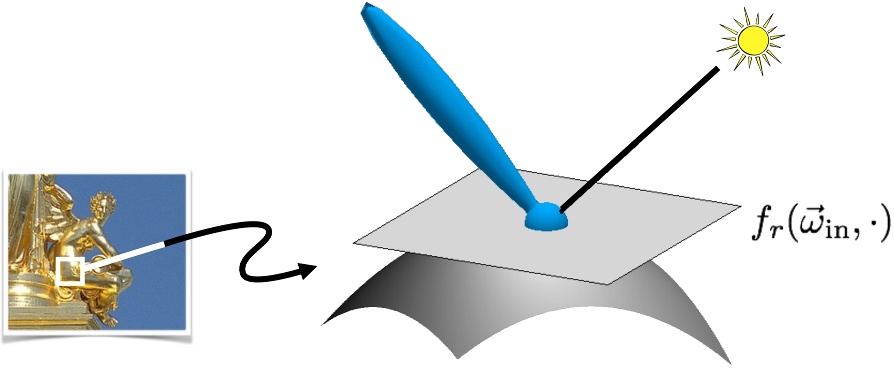

$$\boxed{f_r(\vec{\omega}_{\text{in}}, \vec{\omega}_{\text{out}})}$$

Bi-directional Reflectance Distribution Function (BRDF)

Trick for dielectrics (non-metals)

BRDF is a sum of a Lambertian diffuse component and non-Lambertian specular components

The two components differ in terms of color and polarization, and under certain conditions, this can be exploited to separate them.

$$f(\vec{\omega}_i, \vec{\omega}_o) = f_d + f_s(\vec{\omega}_i, \vec{\omega}_o)$$

Trick for dielectrics (non-metals)

BRDF is a sum of a Lambertian diffuse component and non-Lambertian specular components

The two components differ in terms of color and polarization, and under certain conditions, this can be exploited to separate them.

$$f(\vec{\omega}_i, \vec{\omega}_o) = f_d + f_s(\vec{\omega}_i, \vec{\omega}_o)$$

Often called the dichromatic BRDF:

- Diffuse term varies with wavelength, constant with polarization
- Specular term constant with wavelength, varies with polarization

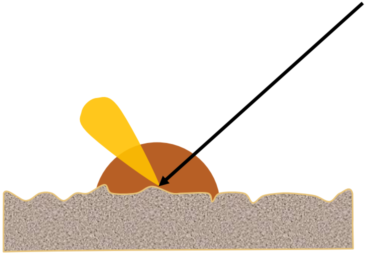

Trick for dielectrics (non-metals)

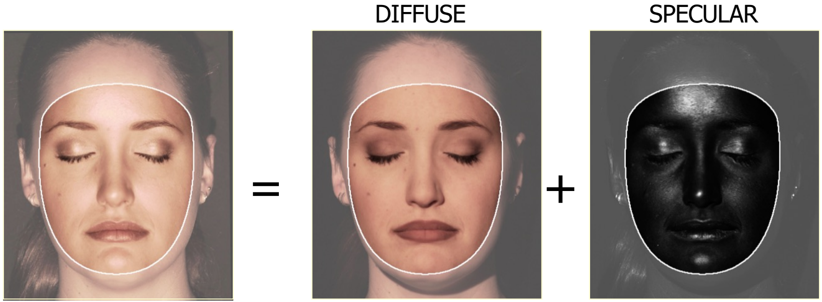

In this example, the two components were separated using linear polarizing filters on the camera and light source.

Trick for dielectrics (non-metals)

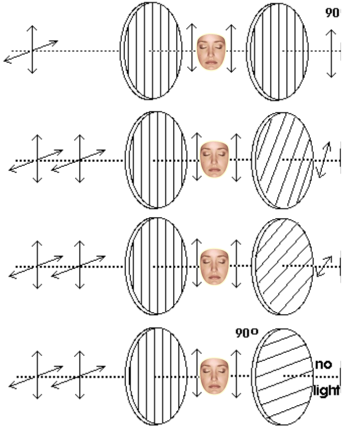

$$I_{\parallel} = \frac{1}{2} I_{\text{diffuse}} + I_{\text{specular}}$$

$$I_{\perp} = \frac{1}{2} I_{\text{diffuse}}$$

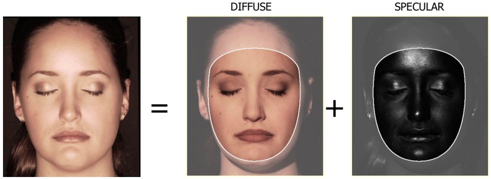

Tabulated 4D BRDFs (hard to measure)

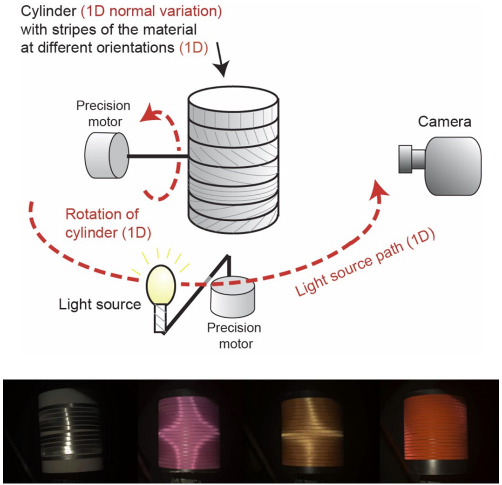

[Ngan et al., 2005]

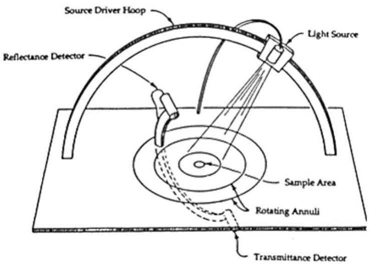

Gonioreflectometer

Low-parameter (non-linear) BRDF models

A small number of parameters define the (2D,3D, or 4D) function

Except for Lambertian, the BRDF is non-linear in these parameters

Examples:

Lambertian: $f(\omega_i, \omega_o) = \frac{a}{\pi}$

Phong: $f(\omega_i, \omega_o) = \frac{a}{\pi} + b\cos^c\left(2\langle\omega_i, n\rangle\langle\omega_o, n\rangle - \langle\omega_i, \omega_o\rangle\right)$

Blinn: $f(\omega_i, \omega_o) = \frac{a}{\pi} + b\cos^c b(\omega_i, \omega_o)$

Lafortune: $f(\omega_i, \omega_o) = \frac{a}{\pi} + b(-\omega_i^\top A\omega_o)^k$

Ward: $f(\omega_i, \omega_o) = \frac{a}{\pi} + \frac{b}{4\pi c^2\sqrt{\langle n, \omega_i\rangle\langle n, \omega_o\rangle}}\exp\left(\frac{-\tan^2 b(\omega_i, \omega_o)}{c^2}\right)$

α is called the albedo

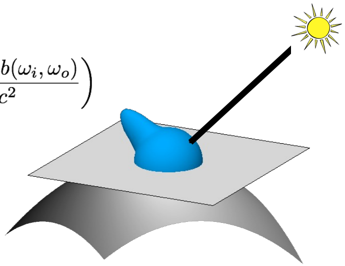

Low-parameter (non-linear) BRDF models

A small number of parameters define the (2D,3D, or 4D) function

Except for Lambertian, the BRDF is non-linear in these parameters

Examples:

Where do these constants come from?

Lambertian: $f(\omega_i, \omega_o) = \frac{a}{\pi}$

Phong: $f(\omega_i, \omega_o) = \frac{a}{\pi} + b \cos^c \left(2\langle \omega_i, n \rangle \langle \omega_o, n \rangle - \langle \omega_i, \omega_o \rangle\right)$

Blinn: $f(\omega_i, \omega_o) = \frac{a}{\pi} + b \cos^c b(\omega_i, \omega_o)$

Lafortune: $f(\omega_i, \omega_o) = \frac{a}{\pi} + b(-\omega_i^\top A \omega_o)^k$

Ward: $f(\omega_i, \omega_o) = \frac{a}{\pi} + \frac{b}{4\pi c^2 \sqrt{\langle n, \omega_i \rangle \langle n, \omega_o \rangle}} \exp\left(\frac{-\tan^2 b(\omega_i, \omega_o)}{c^2}\right)$

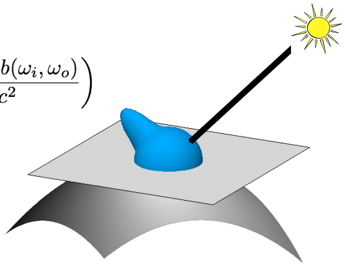

ɑ is called the albedo

Reflectance Models

Reflection: An Electromagnetic Phenomenon

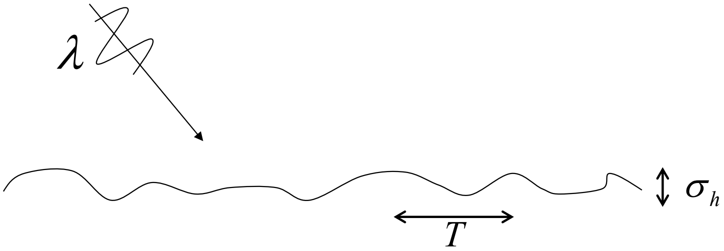

Two approaches to derive Reflectance Models:

– Physical Optics (Wave Optics)

– Geometrical Optics (Ray Optics)

Geometrical models are approximations to physical models

But they are easier to use!

Reflectance that Require Wave Optics

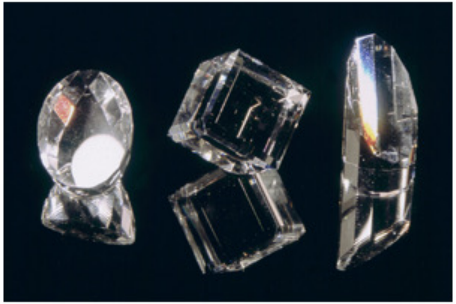




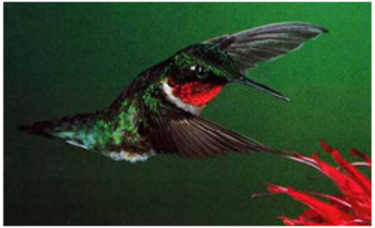

Light sources

Lighting models: Plenoptic function

- Radiance as a function of position and direction
- Radiance as a function of position, direction, and time
- Spectral radiance as a function of position, direction, time and wavelength

$$L(x, \omega, t, \lambda)$$

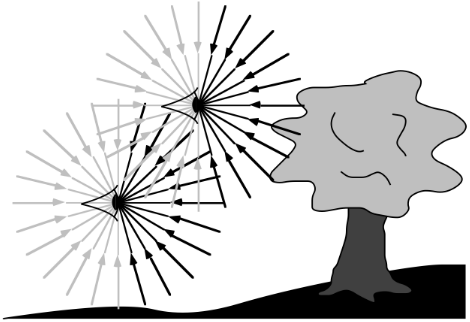

Fig.1.3
The plenoptic function describes the information available to an observer at any point in space and time. Shown here are two schematic eyes-which one should consider to have punctate pupils-gathering pencils of light rays. A real observer cannot see the light rays coming from behind, but the plenoptic function does include these rays.

[Adelson and Bergen, 1991]

Lighting models: far-field (or directional) approximation

- Assume that, over the observed region of interest, all source of incoming flux are relatively far away

$$L(x, \omega, t, \lambda) \longrightarrow L(\omega, t, \lambda)$$

$$L(x, \omega) \longrightarrow L(\omega)$$

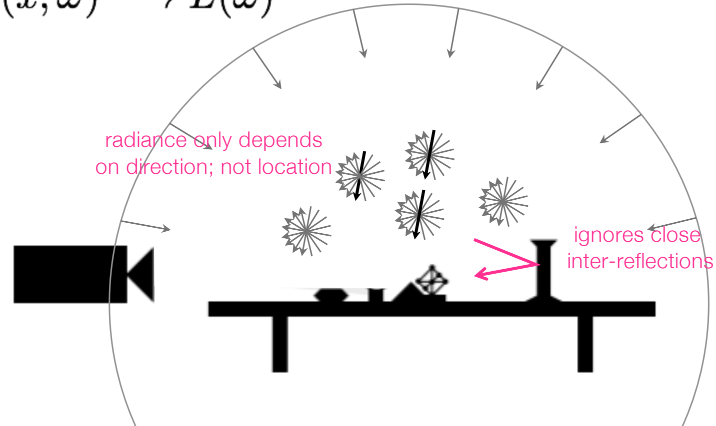

[Debevec, 1998]

Application: augmented reality

$L(\omega)$

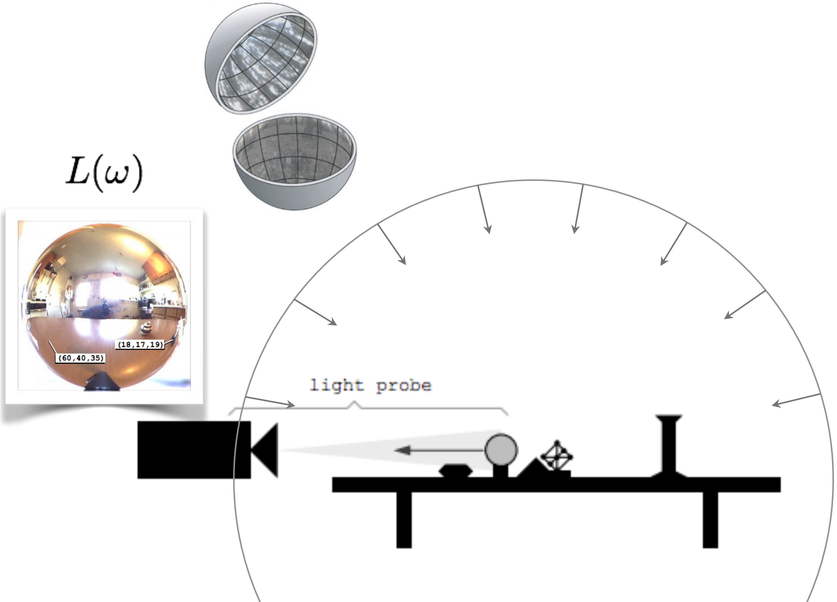

[Debevec, 1998]

Application: augmented reality

(a) Background photograph

(b) Camera calibration grid and light probe

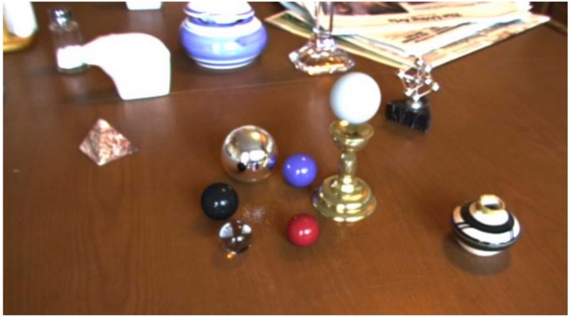

(g) Final result with differential rendering

[Debevec, 1998]

Application: augmented reality

http://gl.ict.usc.edu/LightStages/

Lighting models: far-field approximation

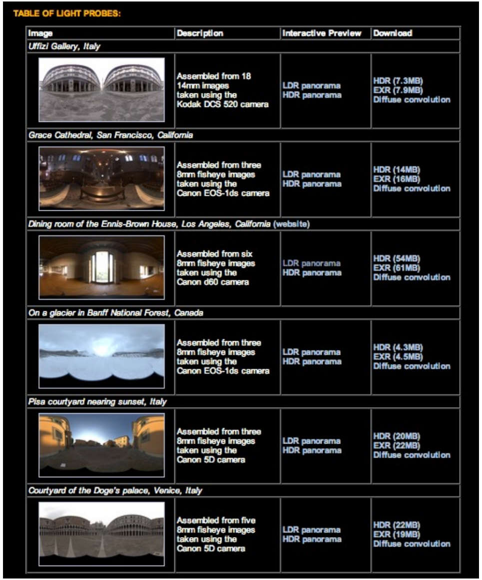

TABLE OF LIGHT PROBES:

Image	Description	Interactive Preview	Download
Uffizi Gallery, Italy			
	Assembled from 18 14mm images taken using the Kodak DCS 520 camera	LDR panorama HDR panorama	HDR (7.3MB) EXR (7.9MB) Diffuse convolution
Grace Cathedral, San Francisco, California			
	Assembled from three 8mm fisheye images taken using the Canon EOS-1ds camera	LDR panorama HDR panorama	HDR (14MB) EXR (16MB) Diffuse convolution
Dining room of the Ennis-Brown House, Los Angeles, California (website)			
	Assembled from six 8mm fisheye images taken using the Canon d60 camera	LDR panorama HDR panorama	HDR (54MB) EXR (61MB) Diffuse convolution
On a glacier in Banff National Forest, Canada			
	Assembled from three 8mm fisheye images taken using the Canon EOS-1ds camera	LDR panorama HDR panorama	HDR (4.3MB) EXR (4.5MB) Diffuse convolution
Pisa courtyard nearing sunset, Italy			
	Assembled from three 8mm fisheye images taken using the Canon 5D camera	LDR panorama HDR panorama	HDR (20MB) EXR (22MB) Diffuse convolution
Courtyard of the Doge's palace, Venice, Italy			
	Assembled from five 8mm fisheye images taken using the Canon 5D camera	LDR panorama HDR panorama	HDR (22MB) EXR (19MB) Diffuse convolution

- One can download far-field lighting environments that have been captured by others

  [http://gl.ict.usc.edu/Data/HighResProbes/]

- A number of apps and software exist to help you capture capture your own environments using a light probe

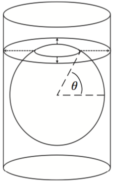

Figure 6. To produce the equal-area cylindrical projection of a spherical map, one projects each point on the surface of the sphere horizontally outward onto the cylinder, and then unwraps the cylinder to obtain a rectangular "panoramic" map.

Application: inferring outdoor illumination

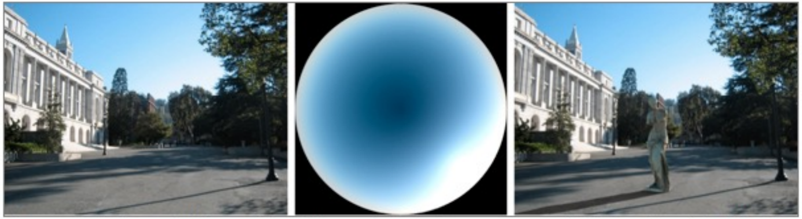

From a single image (left), we estimate the most likely sky appearance (middle) and insert a 3-D object (right). Illumination estimation was done entirely automatically.

[Lalonde et al., 2009]

A further simplification: Low-frequency illumination

$$L(\omega) = \sum_i a_i Y_i(\omega)$$

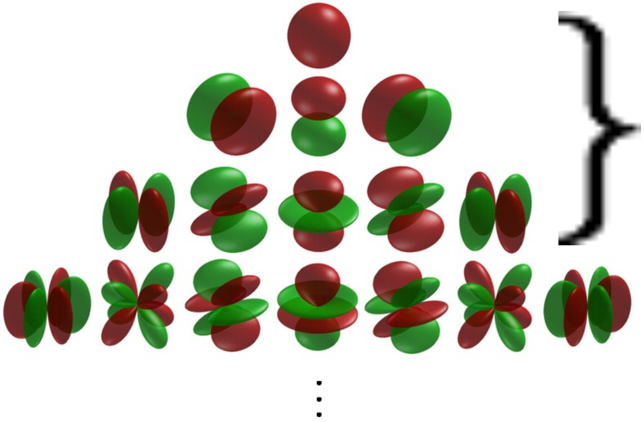

First nine basis functions are sufficient for re-creating Lambertian appearance

[Ramamoorthi and Hanrahan, 2001; Basri and Jacobs, 2003]

Low-frequency illumination

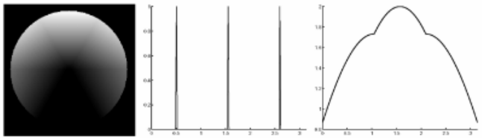

Fig. 2. On the left, a white sphere illuminated by three directional (distant point) sources of light. All the lights are parallel to the image plane, one source illuminates the sphere from above and the two others illuminate the sphere from diagonal directions. In the middle, a cross-section of the lighting function with three peaks corresponding to the three light sources. On the right, a cross-section indicating how the sphere reflects light. We will make precise the intuition that the material acts as a low-pass filtering, smoothing the light as it reflects it.

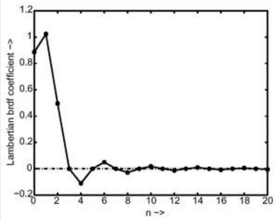

Figure 3. Plot of spherical harmonic terms in Lambertian BRDF filter.

[Ramamoorthi and Hanrahan, 2001; Basri and Jacobs, 2003]

Low-frequency illumination

$$L(\omega) = \sum_i a_i Y_i(\omega)$$

$$\Big\Downarrow \text{ Truncate to first 9 terms}$$

$$\vec{\ell} = (\ell_1, \ldots, \ell_9)$$

[Ramamoorthi and Hanrahan, 2001; Basri and Jacobs, 2003]

Application: Trivial rendering

Capture light probe

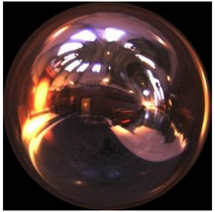

Low-pass filter (truncate to first nine SHs)

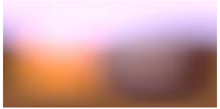

Rendering a (convex) diffuse object in this environment simply requires a lookup based on the surface normal at each pixel

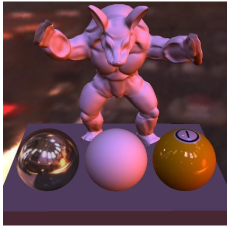

Even simpler:
Directional lighting

- Assume that, over the observed region of interest, all source of incoming flux is from one direction

$$L(x, \omega, t, \lambda) \longrightarrow L(x, t, \lambda) \longrightarrow s(t, \lambda)\delta(\omega = \omega_o(t))$$

$$L(x, \omega) \longrightarrow L(\omega) \longrightarrow s\delta(\omega = \omega_o)$$

- Convenient representation

$$\vec{\boldsymbol{\ell}} = (\ell_x, \ell_y, \ell_z)$$

"light direction" $\hat{\boldsymbol{\ell}} = \dfrac{\vec{\boldsymbol{\ell}}}{||\vec{\boldsymbol{\ell}}||}$

"light strength" $||\vec{\boldsymbol{\ell}}||$

Simple shading

ASSUMPTION 1:
LAMBERTIAN

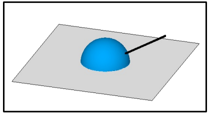

ASSUMPTION 2:
DIRECTIONAL LIGHTING

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

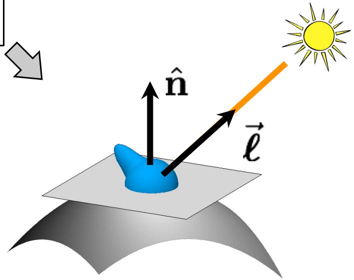

Simple shading

ASSUMPTION 1:
LAMBERTIAN

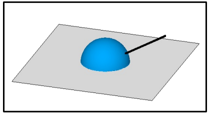

ASSUMPTION 2:
DIRECTIONAL LIGHTING

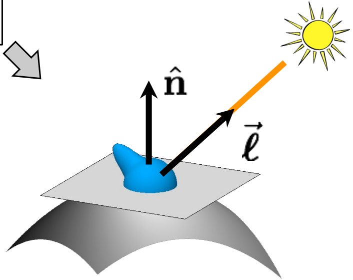

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

$$\boxed{I = a\hat{\mathbf{n}}^\top \vec{\boldsymbol{\ell}}}$$

An ideal point light source

$$L(\boldsymbol{x}, \boldsymbol{\omega}) = \frac{s}{||\boldsymbol{x} - \boldsymbol{x}_o||^2} \; \delta\left(\boldsymbol{\omega} = \frac{\boldsymbol{x} - \boldsymbol{x}_o}{||\boldsymbol{x} - \boldsymbol{x}_o||}\right)$$

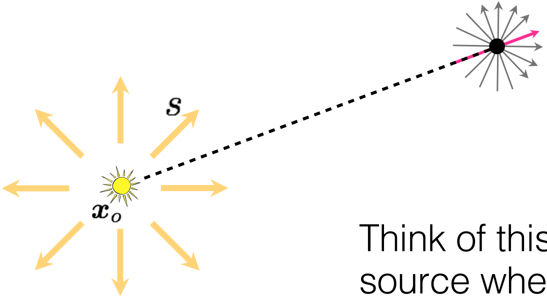

Think of this as a spatially-varying directional source where

1. the direction is away from x_o
2. the strength is proportional to 1/(distance)^2

Summary of some useful lighting models

- plenoptic function (function on 5D domain)
- far-field illumination (function on 2D domain)
- low-frequency far-field illumination (nine numbers)
- directional lighting (three numbers = direction and strength)
- point source (four numbers = location and strength)

Some notes about radiometry

What about color?

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

retinal color

$$\mathbf{c}(\ell(\lambda)) = (c_s, c_m, c_l)$$

$$c_s = \int k_s(\lambda)\ell(\lambda)d\lambda$$

$k_s(\lambda)$ $k_m(\lambda)$ $k_l(\lambda)$

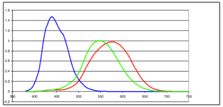

LMS senstivity functions

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

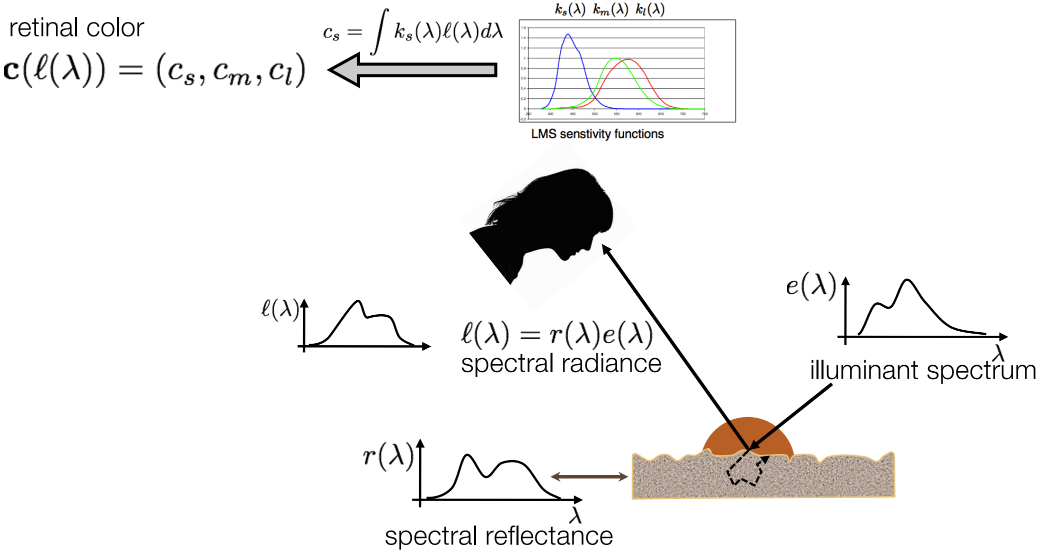

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

Does this view of color ignore any important phenomena?

Spectral radiance

- Distribution of radiance as a function of wavelength.
- All of the phenomena we described in this lecture can be extended to take into account color, by considering separate radiance and BRDF functions independently for each wavelength.

Does this view of color ignore any important phenomena?

- Things like fluorescence and any other phenomena where light changes color.

Spectral Sensitivity Function (SSF)

- *Any* light sensor (digital or not) has different sensitivity to different wavelengths.
- This is described by the sensor's *spectral sensitivity function*

Spectral Sensitivity Function (SSF)

- *Any* light sensor (digital or not) has different sensitivity to different wavelengths.
- This is described by the sensor's *spectral sensitivity function*
- When measuring light of a some SPD $\Phi(\lambda)$, the sensor produces a *scalar* $f(\lambda)$ response:

sensor response $\longrightarrow$

$$R = \int_\lambda \underbrace{\Phi(\lambda)}_{\text{light SPD}} \underbrace{f(\lambda)}_{\text{sensor SSF}} d\lambda$$

Weighted combination of light's SPD: light contributes more at wavelengths where the sensor has higher sensitivity.

Spectral Sensitivity Function (SSF)

- *Any* light sensor (digital or not) has different sensitivity to different wavelengths.
- This is described by the sensor's *spectral sensitivity function*
- When measuring light of a some SPD $\Phi(\lambda)$, the sensor produces a *scalar* $f(\lambda)$ response:

sensor response $\longrightarrow$ $$R = \int_\lambda \underbrace{\Phi(\lambda)}_{\text{light SPD}} \underbrace{f(\lambda)}_{\text{sensor SSF}} d\lambda$$

Weighted combination of light's SPD: light contributes more at wavelengths where the sensor has higher sensitivity.

The spectral sensitivity function converts radiometric units (radiance, irradiance) defined per wavelength, to radiometric quantities where wavelength has been averaged out.

Radiometry versus photometry

- **All radiometric quantities have equivalents in photometry**
- **Photometry: accounts for response of human visual system to electromagnetic radiation**
- **Luminance (Y) is photometric quantity that corresponds to radiance: integrate radiance over all wavelengths, weight by eye's luminous efficacy curve, e.g.:**

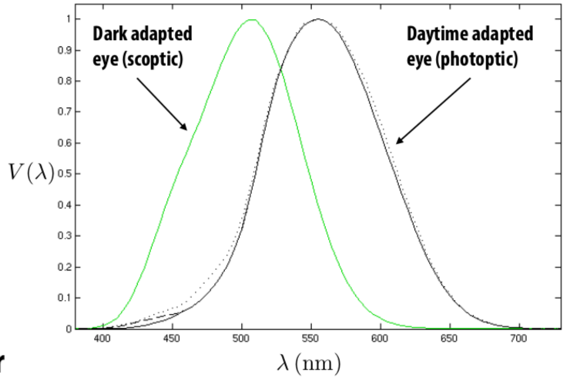

$$Y(\mathrm{p}, \omega) = \int_0^\infty L(\mathrm{p}, \omega, \lambda)\, V(\lambda)\, \mathrm{d}\lambda$$

Radiometric vs. photometric units

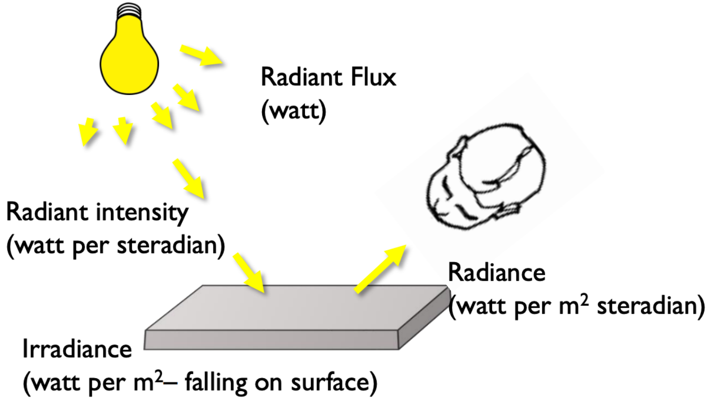

Radiometric values

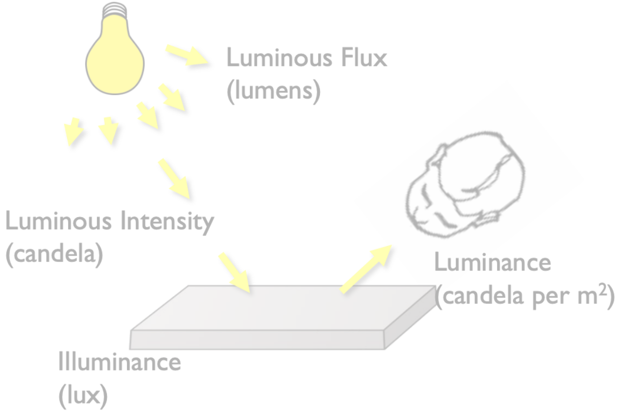

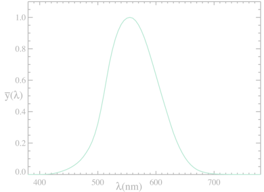

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

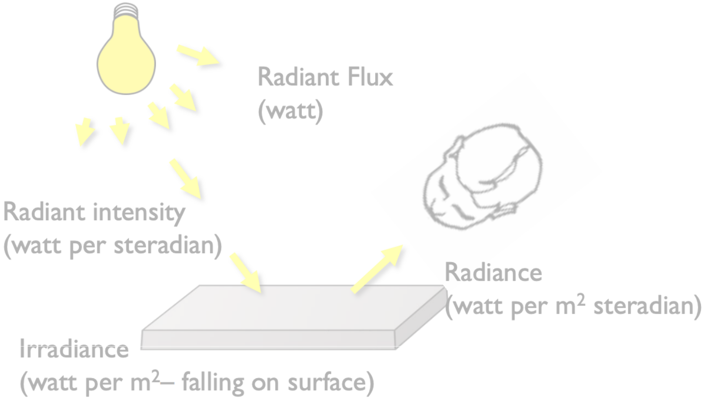

Radiometric values

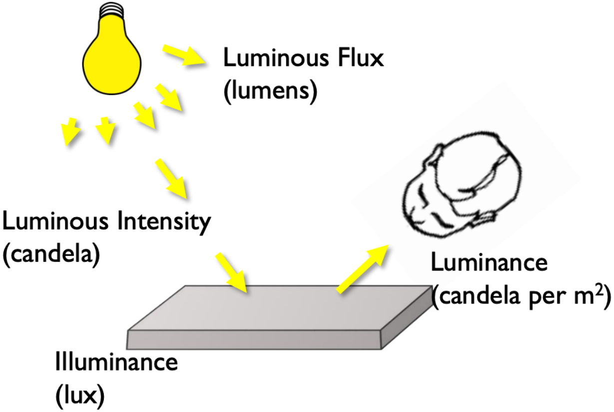

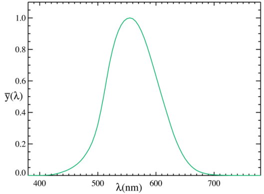

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

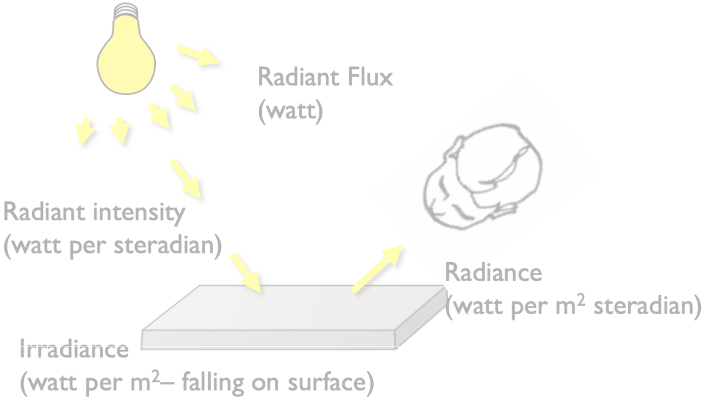

Radiometric values

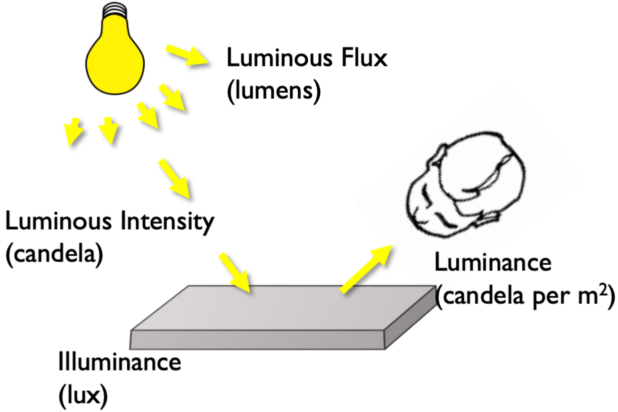

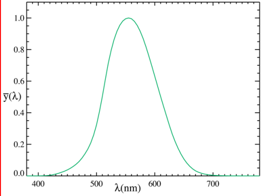

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

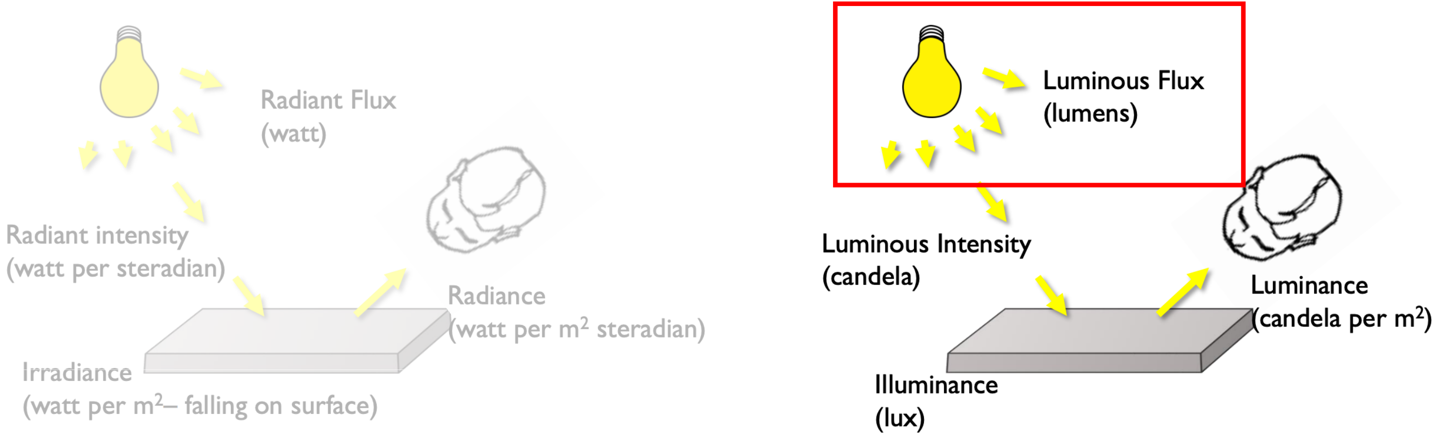

Radiometric values

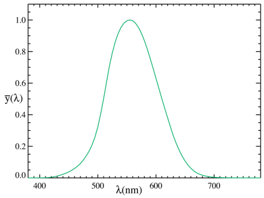

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

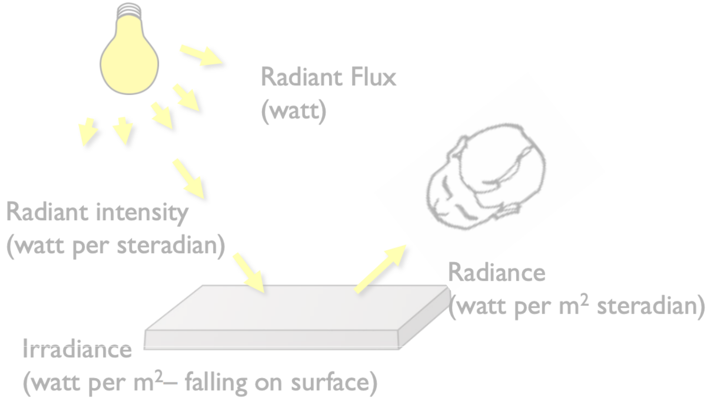

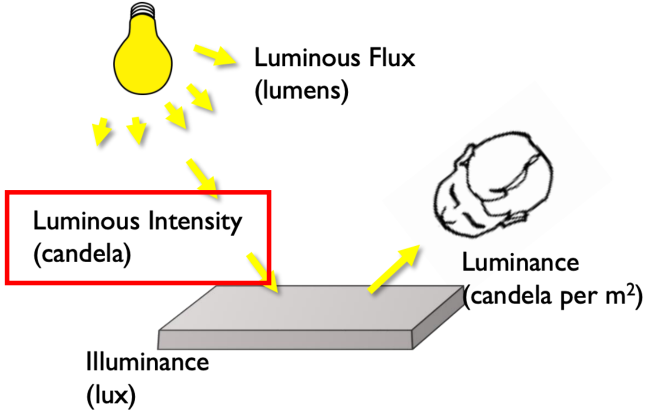

Radiometric values

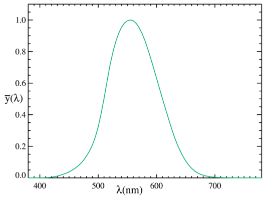

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

- “candlepower” used to be given by a “standard candle” made from spermaceti (whale oil)
- now defined as a source emitting monochromatic radiation at 555 nm and 1/683 W/sr
- also 1 lumen/sr

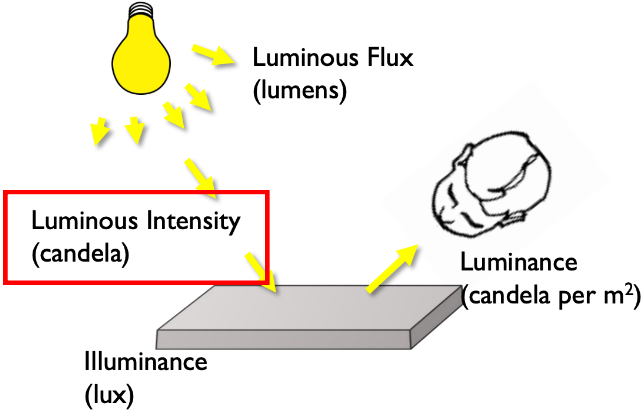

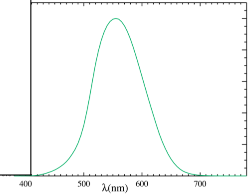

Photometric values
(Radiometric values weighted by the Luminosity Function)

Radiometric vs. photometric units

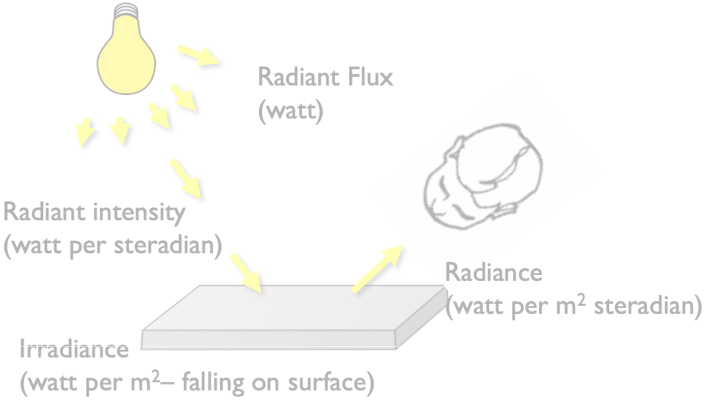

Radiometric values

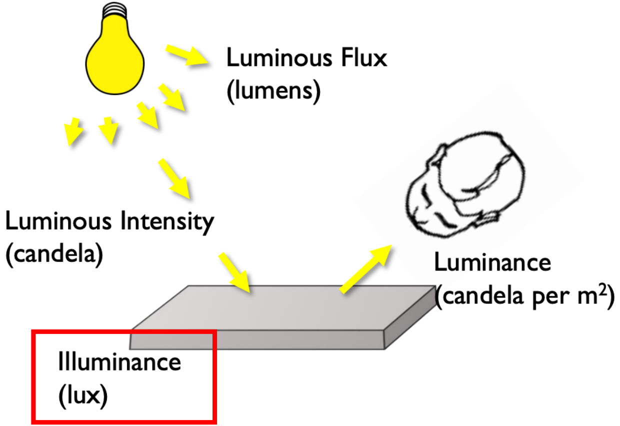

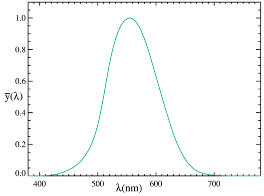

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

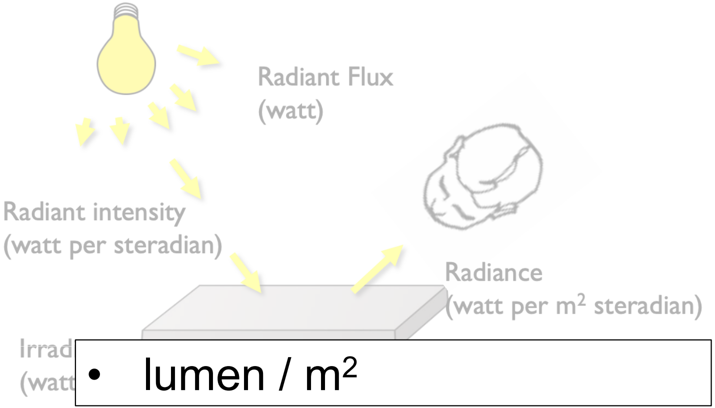

- lumen / m²

Radiometric values

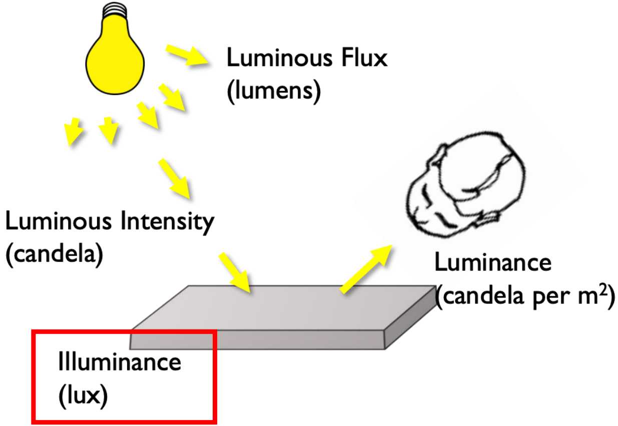

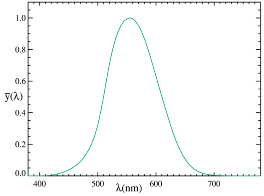

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Radiometric vs. photometric units

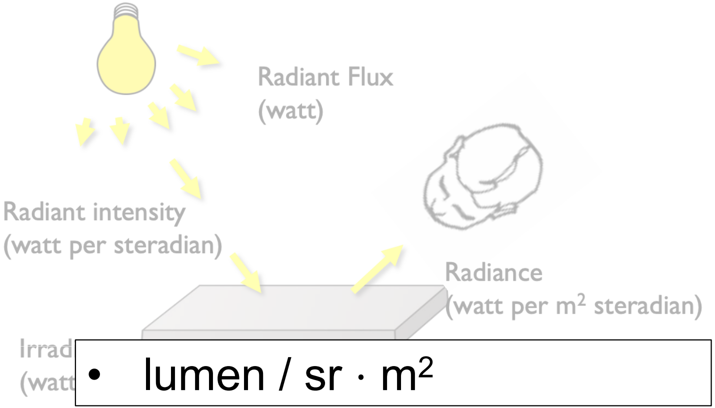

- lumen / sr · m²

Radiometric values

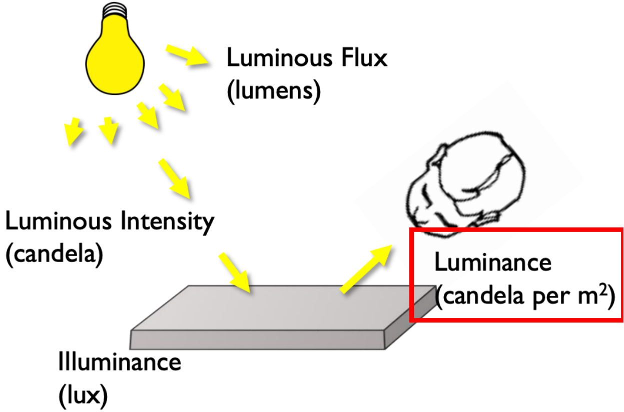

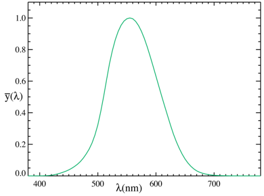

Photometric values
(Radiometric values weighted by the Luminosity Function)

slide: Michael Brown

Modern LED light

Input power: 11 W

Output: 815 lumens

(~ 80 lumens / Watt)

Incandescent bulbs:
~15 lumens / Watt)

Quiz 1: Measurement of a sensor using a thin lens

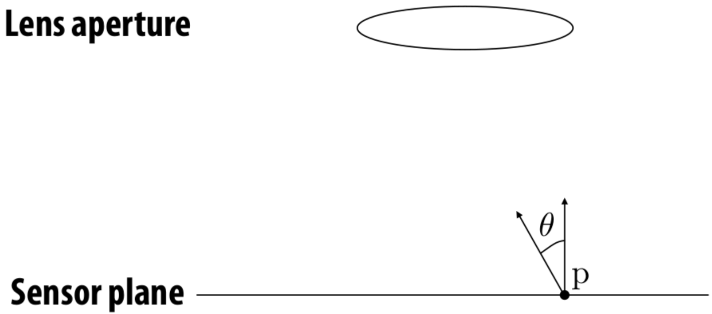

What integral should we write for the power measured by infinitesimal pixel p?

Quiz 1: Measurement of a sensor using a thin lens

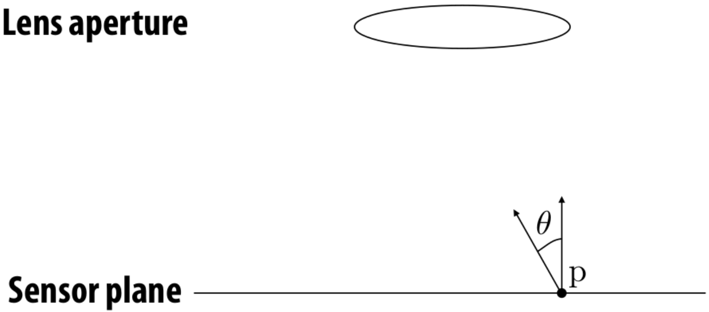

What integral should we write for the power measured by infinitesimal pixel p?

$$E(\mathrm{p}, t) = \int_{H^2} L_i(\mathrm{p}, \omega', t)\ \cos\theta \, \mathrm{d}\omega'$$

Can I transform this integral over the hemisphere to an integral over the aperture area?

Quiz 1: Measurement of a sensor using a thin lens

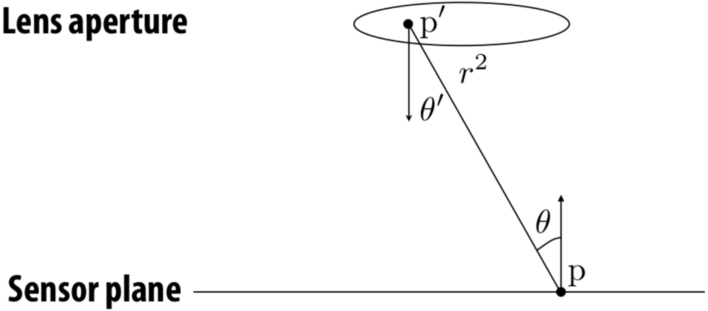

What integral should we write for the power measured by infinitesimal pixel p?

$$E(\mathrm{p}, t) = \int_{H^2} L_i(\mathrm{p}, \omega', t)\ \cos\theta \,\mathrm{d}\omega'$$

Can I transform this integral over the hemisphere to an integral over the aperture area?

$$E(\mathrm{p}, t) = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos\theta \cos\theta'}{||\mathrm{p}' - \mathrm{p}||^2} \,\mathrm{d}A'$$

Transform integral over solid angle to integral over lens aperture

Quiz 1: Measurement of a sensor using a thin lens

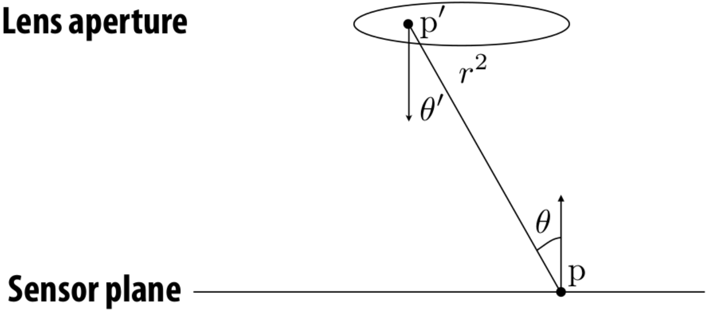

$$E(\mathrm{p}, t) = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos\theta \cos\theta'}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Transform integral over solid angle to integral over lens aperture

$$= \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos^2\theta}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Assume aperture and film plane are parallel: $\theta = \theta'$

Can I write the denominator in a more convenient form?

Quiz 1: Measurement of a sensor using a thin lens

Lens aperture

$$||\mathrm{p}' - \mathrm{p}|| = \frac{d}{\cos\theta}$$

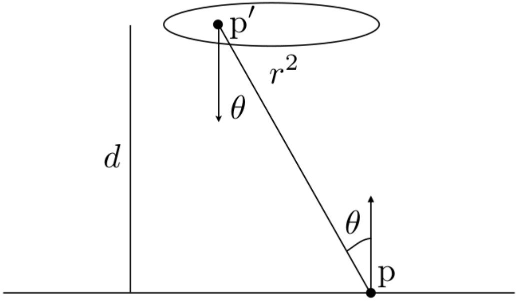

Sensor plane

$$E(\mathrm{p}, t) = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos^2\theta}{||\mathrm{p}' - \mathrm{p}||^2}\, \mathrm{d}A'$$

$$= \frac{1}{d^2} \int_A L(\mathrm{p}' \to \mathrm{p}, t)\ \cos^4\theta \, \mathrm{d}A'$$

What does this say about the image I am capturing?

Vignetting

Fancy word for: pixels far off the center receive less light

white wall under uniform light

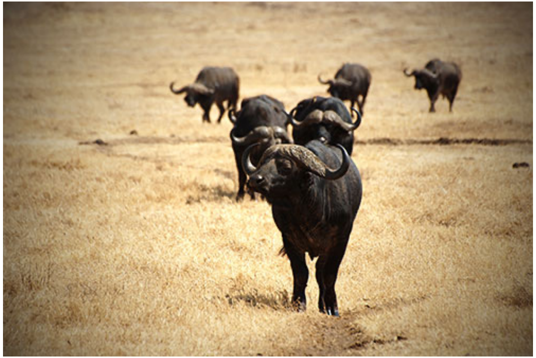

more interesting example of vignetting

Vignetting

Fancy word for: pixels far off the center receive less light

white wall under uniform light

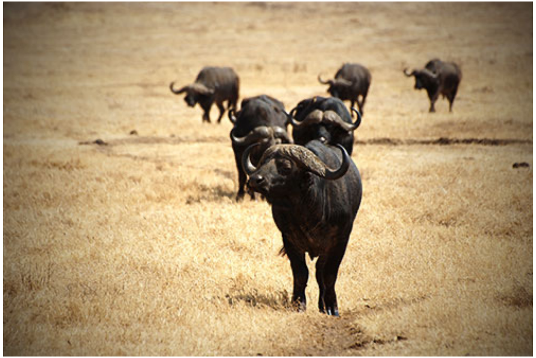

more interesting example of vignetting

Four types of vignetting:

- Mechanical: light rays blocked by hoods, filters, and other objects.
- Lens: similar, but light rays blocked by lens elements.
- <u>Natural: due to radiometric laws ("cosine fourth falloff").</u>
- Pixel: angle-dependent sensitivity of photodiodes.

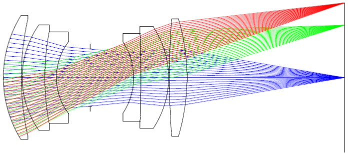

Quiz 2: BRDF of the moon

What BRDF does the moon have?

Quiz 2: BRDF of the moon

What BRDF does the moon have?

- Can it be diffuse?

Quiz 2: BRDF of the moon

What BRDF does the moon have?

- Can it be diffuse?

Even though the moon appears matte, its edges remain bright.

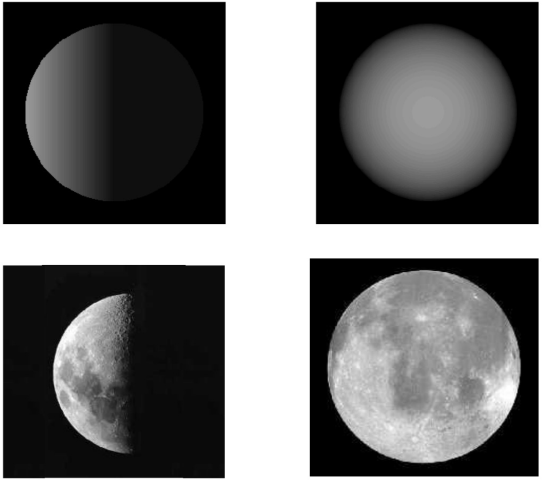

Rough diffuse appearance

Surface Roughness Causes Flat Appearance

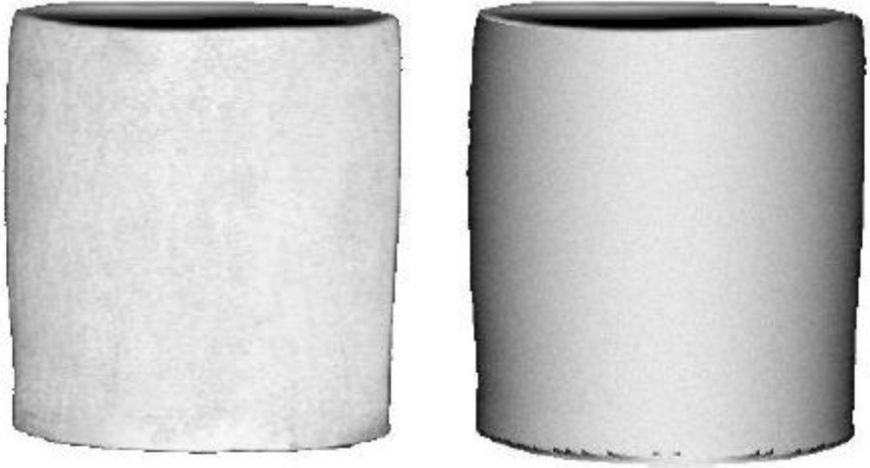

Actual Vase Lambertian Vase

Five important equations/integrals to remember

Flux measured by a sensor of area X and directional receptivity W:

$$\Phi(W, X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta d\boldsymbol{\omega} dA$$

Five important equations/integrals to remember

Flux measured by a sensor of area X and directional receptivity W:

$$\Phi(W, X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta d\boldsymbol{\omega} dA$$

Reflectance equation:

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

Five important equations/integrals to remember

Flux measured by a sensor of area X and directional receptivity W:

$$\Phi(W, X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta d\boldsymbol{\omega} dA$$

Reflectance equation:

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

Radiance under directional lighting and Lambertian BRDF (“n-dot-l shading”):

$$L^{\text{out}} = a \hat{\mathbf{n}}^\top \vec{\boldsymbol{\ell}}$$

Five important equations/integrals to remember

Flux measured by a sensor of area X and directional receptivity W:

$$\Phi(W, X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta d\boldsymbol{\omega} dA$$

Reflectance equation:

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

Radiance under directional lighting and Lambertian BRDF ("n-dot-l shading"):

$$L^{\text{out}} = a\hat{\mathbf{n}}^\top \vec{\boldsymbol{\ell}}$$

Conversion of a (hemi)-spherical integral to a surface integral:

$$\int_{H^2} L_i(\mathrm{p}, \omega', t) \ \cos\theta \, \mathrm{d}\omega' \ = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos\theta \cos\theta'}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Five important equations/integrals to remember

Flux measured by a sensor of area X and directional receptivity W:

$$\Phi(W,X) = \int_X \int_W L(\hat{\boldsymbol{\omega}}, x) \cos\theta \, d\boldsymbol{\omega} \, dA$$

Reflectance equation:

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} \, d\hat{\boldsymbol{\omega}}_{\text{in}}$$

Radiance under directional lighting and Lambertian BRDF ("n-dot-l shading"):

$$L^{\text{out}} = a \hat{\mathbf{n}}^\top \vec{\boldsymbol{\ell}}$$

Conversion of a (hemi)-spherical integral to a surface integral:

$$\int_{H^2} L_i(\mathrm{p}, \omega', t) \, \cos\theta \, \mathrm{d}\omega' = \int_A L(\mathrm{p}' \to \mathrm{p}, t) \frac{\cos\theta \cos\theta'}{||\mathrm{p}' - \mathrm{p}||^2} \, \mathrm{d}A'$$

Computing (hemi)-spherical integrals:

$$d\omega = \frac{dA}{r^2} = \sin\theta \, d\theta \, d\phi$$

and

$$\int d\omega = \int_0^{\pi} \int_0^{2\pi} \sin\theta \, d\theta \, d\phi$$

Photometric stereo

Image Intensity and 3D Geometry

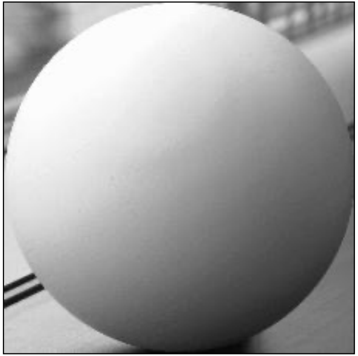

- *Shading* as a cue for shape reconstruction
- What is the relation between intensity and shape?

"N-dot-l" shading

ASSUMPTION 1:
LAMBERTIAN

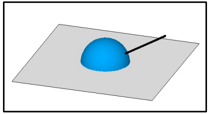

ASSUMPTION 2:
DIRECTIONAL LIGHTING

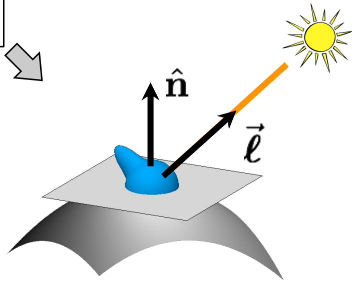

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

$$\boxed{I = a\hat{\mathbf{n}}^\top \vec{\boldsymbol{\ell}}}$$

“N-dot-l” shading

ASSUMPTION 1:
LAMBERTIAN

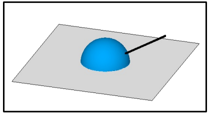

ASSUMPTION 2:
DIRECTIONAL LIGHTING

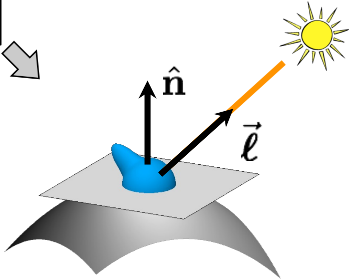

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

$$\boxed{I = a\hat{\mathbf{n}}^{\top}\vec{\boldsymbol{\ell}}}$$

Why do we call these normal “shape”?

Surfaces and normals

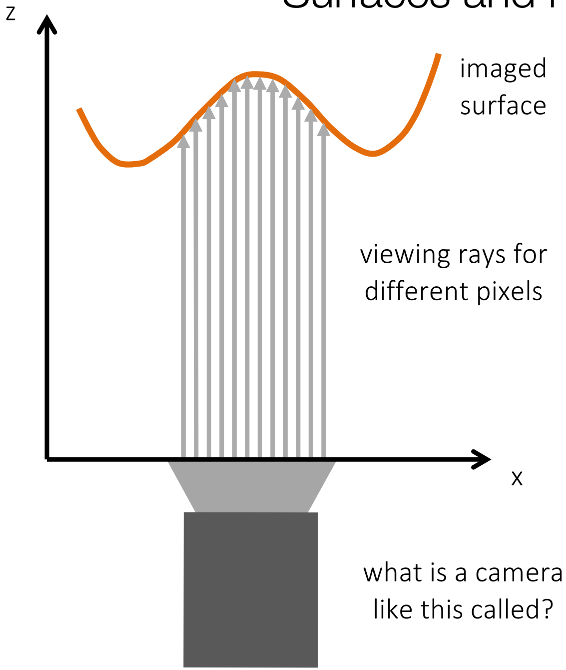

Surfaces and normals

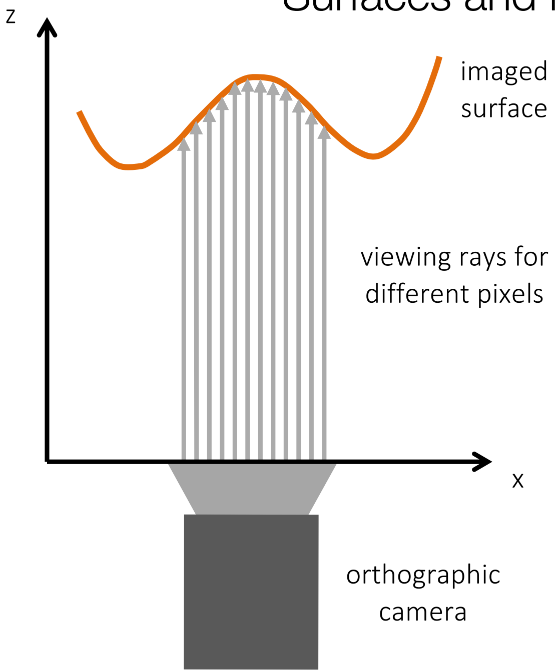

Surfaces and normals

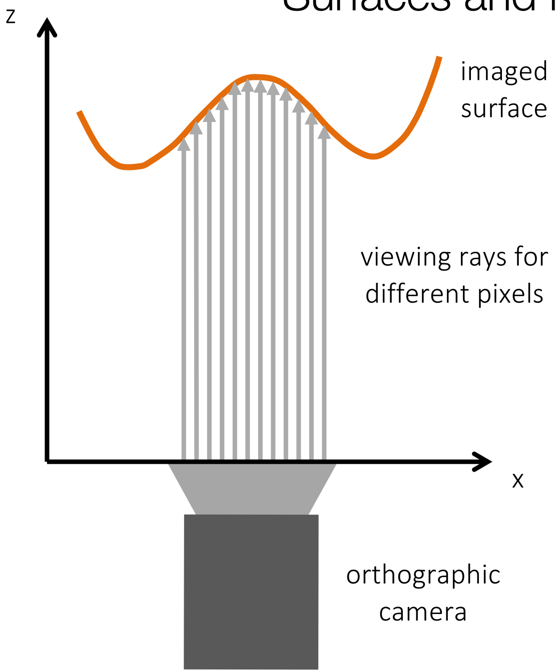

Surface representation as a depth image (also known as Monge surface):

$$z = f(\underbrace{x, y}_{\text{pixel coordinates on image place}})$$

depth at each pixel

Surfaces and normals

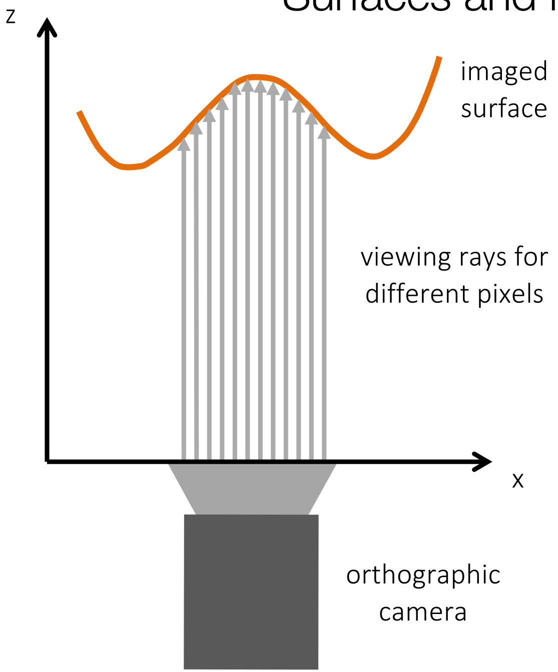

Surface representation as a depth image (also known as Monge surface):

$$z = f\underbrace{(x, y)}_{\text{pixel coordinates on image place}}$$

depth at each pixel

How does surface normal relate to this representation?

Surfaces and normals

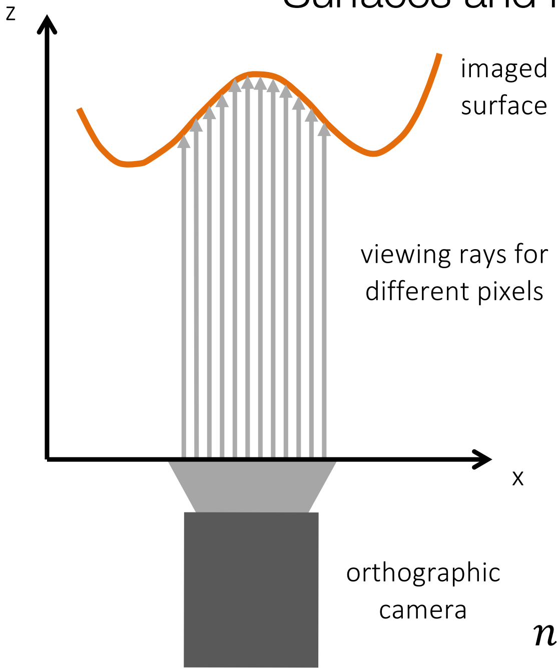

Surface representation as a depth image (also known as Monge surface):

$$\underset{\substack{\uparrow \\ \text{depth at each pixel}}}{z} = f\underbrace{(x, y)}_{\text{pixel coordinates on image place}}$$

Unnormalized normal:

$$\tilde{n}(x, y) = \left(\frac{df}{dx}, \frac{df}{dy}, -1\right)$$

Actual normal:

$$n(x, y) = \tilde{n}(x, y) / \|\tilde{n}(x, y)\|$$

Normals are scaled spatial derivatives of depth image!

Shape from a Single Image?

- Given a single image of an object with known surface reflectance taken under a <u>known light source</u>, can we recover the shape of the object?

Human Perception

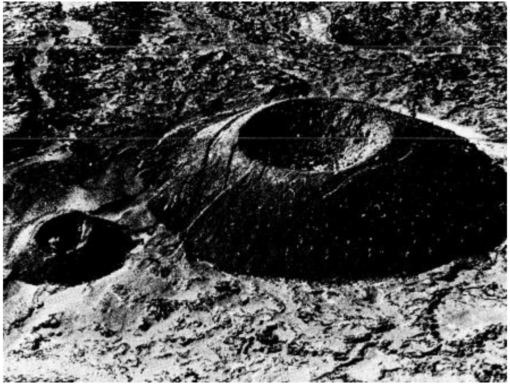

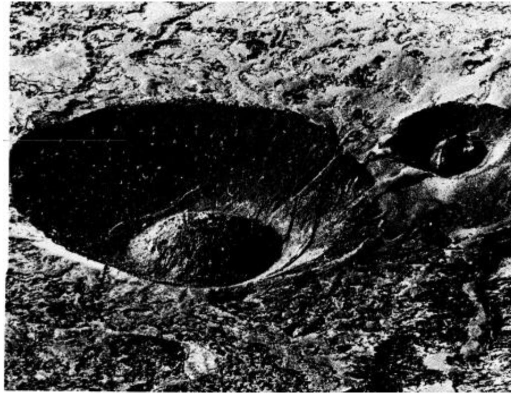

Examples of the classic bump/dent stimuli used to test lighting assumptions when judging shape from shading, with shading orientations (a) 0° and (b) 180° from the vertical.

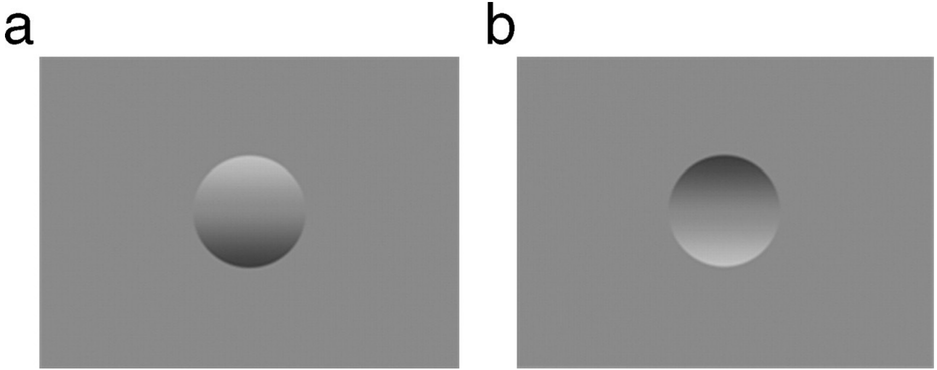

Thomas R et al. J Vis 2010;10:6

journal of VISION

Human Perception

- Our brain often perceives shape from shading.
- Mostly, it makes many assumptions to do so.
- For example:

    Light is coming from above (sun).

    Biased by occluding contours.

by V. Ramachandran

Single-lighting is ambiguous

ASSUMPTION 1:
LAMBERTIAN

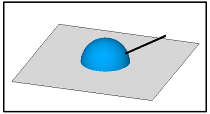

ASSUMPTION 2:
DIRECTIONAL LIGHTING

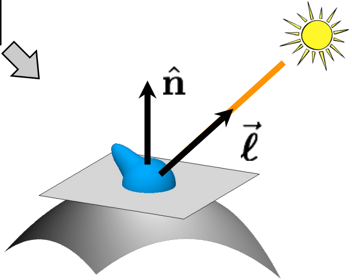

$$L^{\text{out}}(\hat{\boldsymbol{\omega}}) = \int_{\Omega_{\text{in}}} f(\hat{\boldsymbol{\omega}}_{\text{in}}, \hat{\boldsymbol{\omega}}_{\text{out}}) L^{\text{in}}(\hat{\boldsymbol{\omega}}_{\text{in}}) \cos\theta_{\text{in}} d\hat{\boldsymbol{\omega}}_{\text{in}}$$

$$\boxed{I = a\hat{\mathbf{n}}^{\top}\vec{\boldsymbol{\ell}}}$$

Lambertian photometric stereo

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

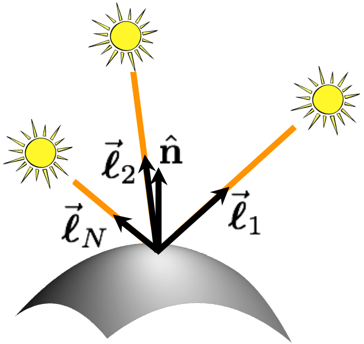

Assumption: We know the lighting directions.

Lambertian photometric stereo

$$\begin{aligned} I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\ I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\ &\vdots \\ I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N \end{aligned}$$

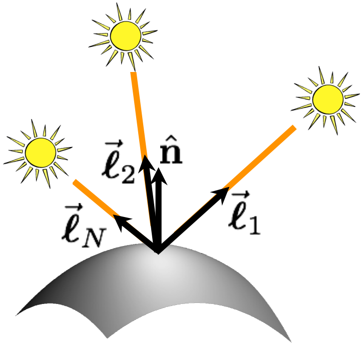

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}$$

What are the dimensions of these matrices?

Lambertian photometric stereo

$$\begin{aligned} I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\ I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\ &\vdots \\ I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N \end{aligned}$$

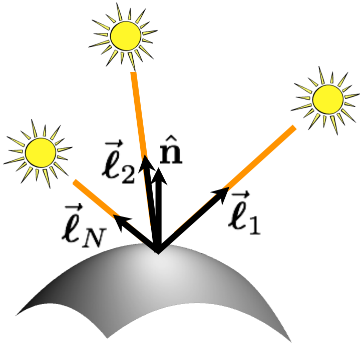

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}$$

What are the knowns and unknowns?

Lambertian photometric stereo

$$\begin{aligned} I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\ I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\ &\vdots \\ I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N \end{aligned}$$

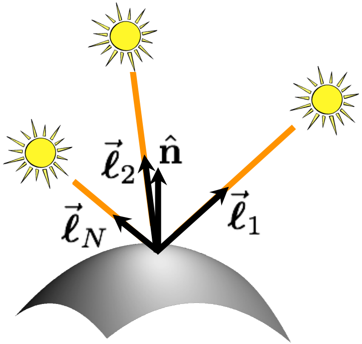

define "pseudo-normal" $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

How many lights do I need for unique solution?

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}$$

Lambertian photometric stereo

$$
\begin{aligned}
I_1 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_1 \\
I_2 &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_2 \\
&\vdots \\
I_N &= a\hat{\boldsymbol{n}}^\top \vec{\boldsymbol{\ell}}_N
\end{aligned}
$$

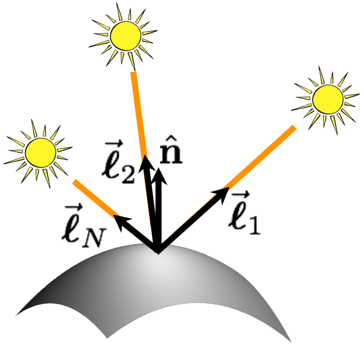

define “pseudo-normal” $\vec{\boldsymbol{b}} \triangleq a\hat{\boldsymbol{n}}$

solve linear system for pseudo-normal

$$
\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{bmatrix}_{N\times 1} = \begin{bmatrix} \vec{\boldsymbol{\ell}}_1^\top \\ \vec{\boldsymbol{\ell}}_2^\top \\ \vdots \\ \vec{\boldsymbol{\ell}}_N^\top \end{bmatrix}_{N\times 3} \begin{bmatrix} \vec{\boldsymbol{b}} \end{bmatrix}_{3\times 1}
$$

once system is solved, b gives normal direction and albedo

How do we solve this system?

Solving the Equation with three lights

$$\underbrace{\begin{bmatrix} I_1 \\ I_2 \\ I_2 \end{bmatrix}}_{\underset{3\times 1}{\mathbf{I}}} = \underbrace{\begin{bmatrix} \mathbf{s}_1^T \\ \mathbf{s}_2^T \\ \mathbf{s}_3^T \end{bmatrix}}_{\underset{3\times 3}{\mathbf{S}}} \underbrace{\rho \mathbf{n}}_{\underset{3\times 1}{\tilde{\mathbf{n}}}}$$

$$\tilde{\mathbf{n}} = \mathbf{S}^{-1}\mathbf{I}$$

inverse

$$\rho = |\tilde{\mathbf{n}}|$$

$$\mathbf{n} = \frac{\tilde{\mathbf{n}}}{|\tilde{\mathbf{n}}|} = \frac{\tilde{\mathbf{n}}}{\rho}$$

Is there any reason to use more than three lights?

More than Three Light Sources

- Get better SNR by using more lights

$$\begin{bmatrix} I_1 \\ \vdots \\ I_N \end{bmatrix} = \begin{bmatrix} \mathbf{s}_1^T \\ \vdots \\ \mathbf{s}_N^T \end{bmatrix} \rho \mathbf{n}$$

- Least squares solution:

$$\mathbf{I} = \mathbf{S}\tilde{\mathbf{n}} \quad \longleftarrow \quad N \times 1 = \underline{(N \times 3)}(3 \times 1)$$

$$\mathbf{S}^T \mathbf{I} = \mathbf{S}^T \mathbf{S} \tilde{\mathbf{n}}$$

$$\tilde{\mathbf{n}} = \boxed{\left(\mathbf{S}^T \mathbf{S}\right)^{-1} \mathbf{S}^T} \mathbf{I}$$

Moore-Penrose pseudo inverse

- Solve for $\rho, \mathbf{n}$ as before

Computing light source directions

- Trick: place a chrome sphere in the scene

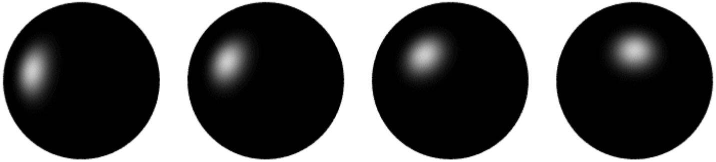

  – the location of the highlight tells you the source direction

Limitations

- Big problems
  - Doesn’t work for shiny things, semi-translucent things
  - Shadows, inter-reflections
- Smaller problems
  - Camera and lights have to be distant
  - Calibration requirements
    - measure light source directions, intensities
    - camera response function

Depth from normals

- Solving the linear system per-pixel gives us an estimated surface normal for each pixel

Input photo

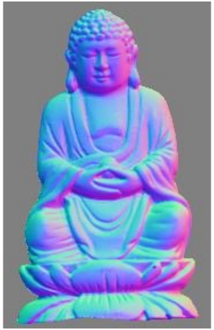

Estimated normals

Estimated normals
(needle diagram)

- How can we compute depth from normals?
  - Normals are like the “derivative” of the true depth

Surfaces and normals

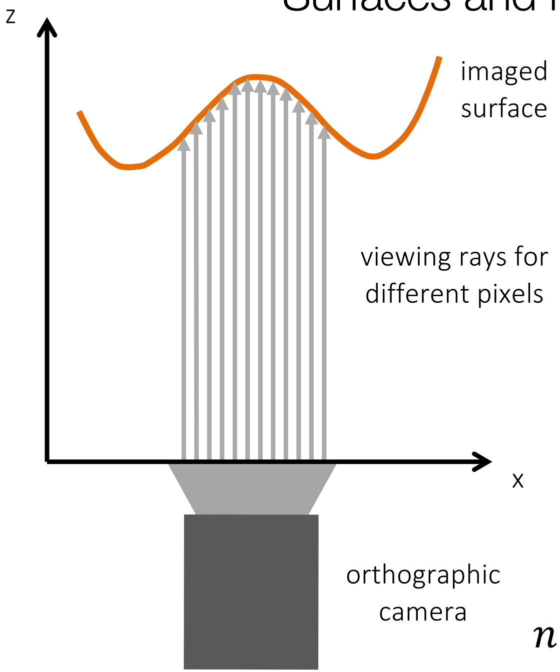

Surface representation as a depth image (also known as Monge surface):

$$z = f\underbrace{(x, y)}_{\text{pixel coordinates in image space}}$$

depth at each pixel

Unnormalized normal:

$$\tilde{n}(x, y) = \left(\frac{df}{dx}, \frac{df}{dy}, -1\right)$$

Actual normal:

$$n(x, y) = \tilde{n}(x, y) / \|\tilde{n}(x, y)\|$$

Normals are scaled spatial derivatives of depth image!

Normal Integration

- Integrating a set of derivatives is easy in 1D
  - (similar to Euler's method from diff. eq. class)

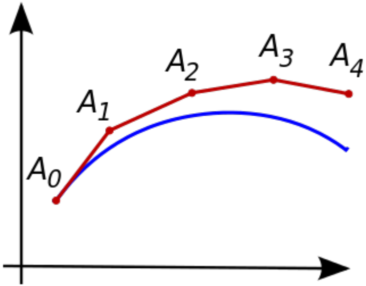

- Could just integrate normals in each column / row separately
- Instead, we formulate as a linear system and solve for depths that *best agree with the surface normals*

Depth from normals

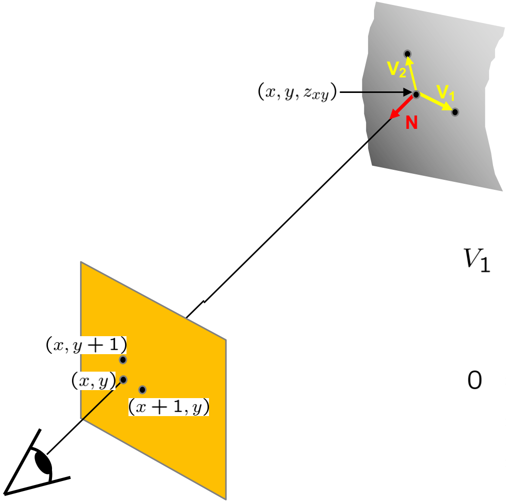

$$\begin{aligned} V_1 &= (x+1, y, z_{x+1,y}) - (x, y, z_{xy}) \\ &= (1, 0, z_{x+1,y} - z_{xy}) \end{aligned}$$

$$\begin{aligned} 0 &= N \cdot V_1 \\ &= (n_x, n_y, n_z) \cdot (1, 0, z_{x+1,y} - z_{xy}) \\ &= n_x + n_z(z_{x+1,y} - z_{xy}) \end{aligned}$$

Get a similar equation for $\mathbf{V_2}$

- Each normal gives us two linear constraints on z
- compute z values by solving a matrix equation

Results

1. Estimate light source directions
2. Compute surface normals
3. Compute albedo values
4. Estimate depth from surface normals
5. Relight the object (with original texture and uniform albedo)

Results: Lambertian Sphere

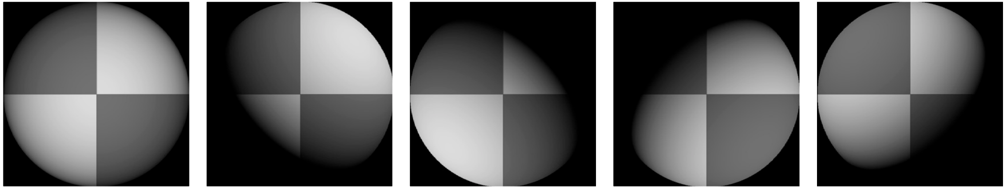

Input Images

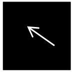

Needles are projections of surface normals on image plane

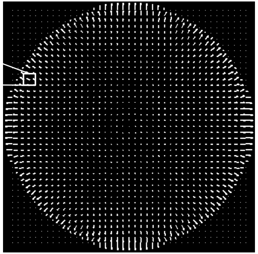

Estimated Surface Normals

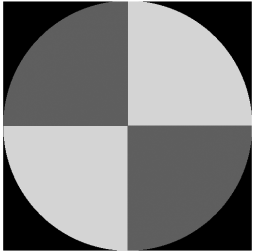

Estimated Albedo

Lambertian Mask

Results – Albedo and Surface Normal

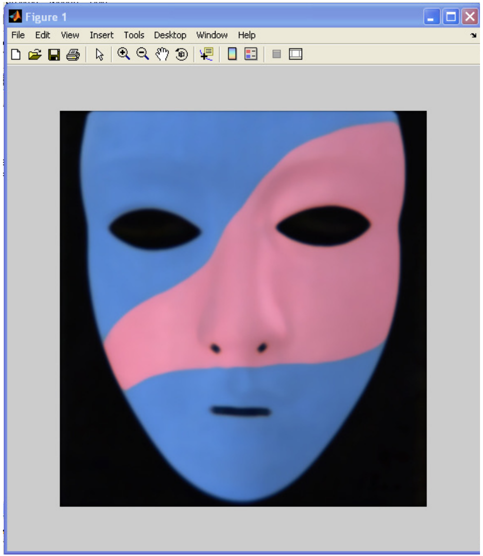

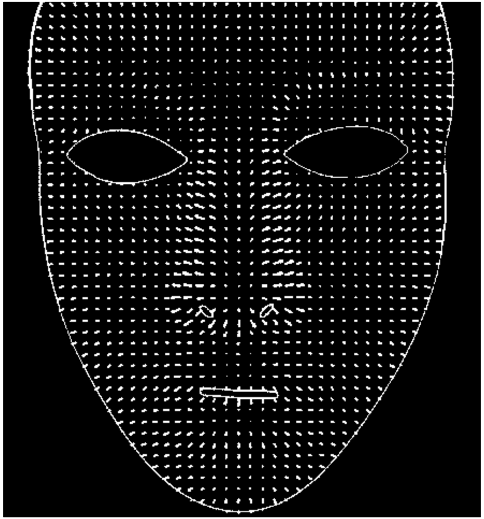

Results – Shape of Mask

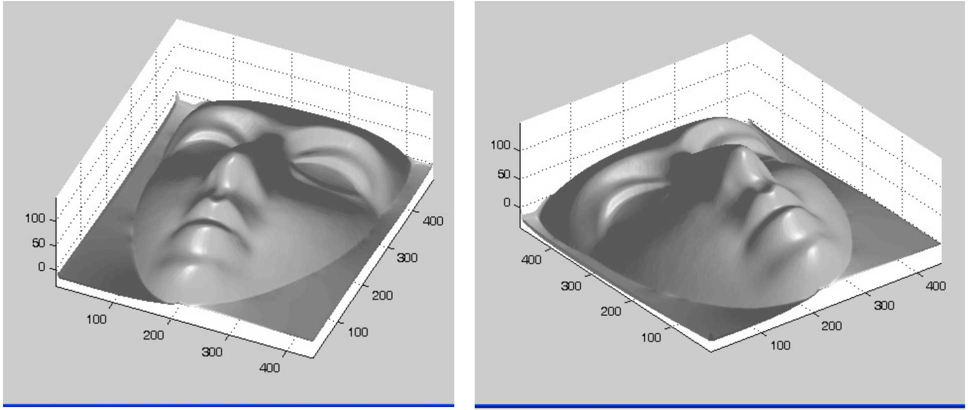

Results: Lambertian Toy





Non-idealities: interreflections

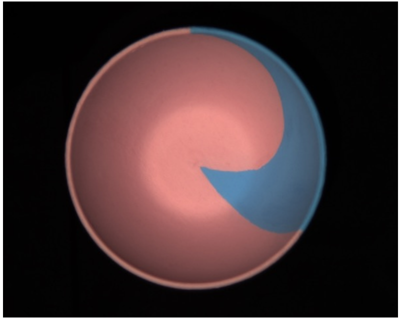
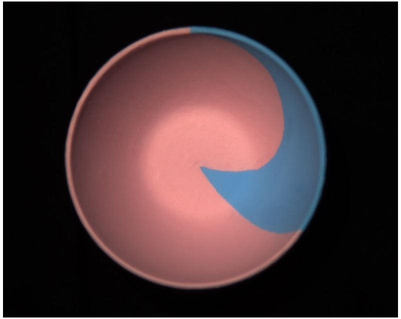
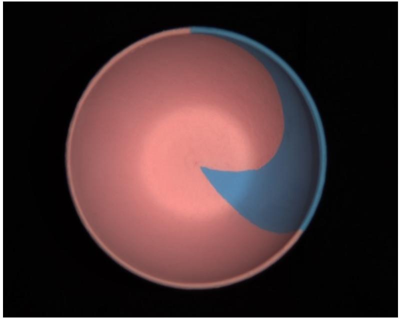

Non-idealities: interreflections

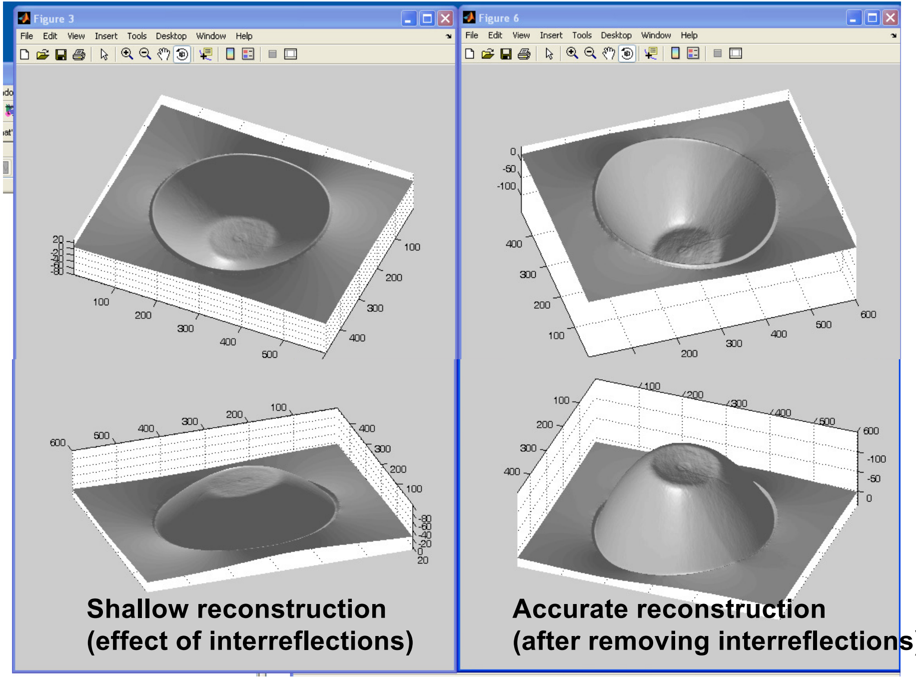

Shallow reconstruction (effect of interreflections)

Accurate reconstruction (after removing interreflections)

Next: 3D Vision & Stereo

References

Basic reading:

- Szeliski, Section 2.2.
- Gortler, Chapter 21.

This book by Steven Gortler has a great introduction to radiometry, reflectance, and their use for image formation.

Additional reading:

- Arvo, “Analytic Methods for Simulated Light Transport,” Yale 1995.
- Veach, “Robust Monte Carlo Methods for Light Transport Simulation,” Stanford 1997.

These two thesis are foundational for modern computer graphics. Among other things, they include a thorough derivation (starting from wave optics and measure theory) of all radiometric quantities and associated integro-differential equations. You can also look at them if you are interested in physics-based rendering.

- Dutre et al., “Advanced Global Illumination,” 2006.

A book discussing modeling and simulation of other appearance effects beyond single-bounce reflectance.

- Weyrich et al., “Principles of Appearance Acquisition and Representation,” FTCGV 2009.

A very thorough review of everything that has to do with modeling and measuring BRDFs.

- Walter et al., “Microfacet models for refraction through rough surfaces,” EGSR 2007.

This paper has a great review of physics-based models for reflectance and refraction.

- Matusik, “A data-driven reflectance model,” MIT 2003.

This thesis introduced the largest measured dataset of 4D reflectances. It also provides detailed discussion of many topics relating to modelling reflectance.

- Rusinkiewicz, “A New Change of Variables for Efficient BRDF Representation,” 1998.
- Romeiro and Zickler, “Inferring reflectance under real-world illumination,” Harvard TR 2010.

These two papers discuss the isotropy and other properties of common BRDFs, and how one can take advantage of them using alternative parameterizations.

- Shafer, “Using color to separate reflection components,” 1984.

The paper introducing the dichromatic reflectance model.

- Stam, “Diffraction Shaders,” SIGGRAPH 1999.
- Levin et al., “Fabricating BRDFs at high spatial resolution using wave optics,” SIGGRAPH 2013.
- Cuypers et al., “Reflectance model for diffraction,” TOG 2013.

These three papers describe reflectance effects that can only be modeled using wave optics (and in particular diffraction).

References (cont.)

Additional reading:

- Oren and Nayar, "Generalization of the Lambertian model and implications for machine vision," IJCV 1995.

The paper introducing the most common model for rough diffuse reflectance.

- Debevec, "Rendering Synthetic Objects into Real Scenes," SIGGRAPH 1998.

The paper that introduced the notion of the environment map, the use of chrome spheres for measuring such maps, and the idea that they can be used for easy rendering.

- Lalonde et al., "Estimating the Natural Illumination Conditions from a Single Outdoor Image," IJCV 2012.

A paper on estimating outdoors environment maps from just one image.

- Basri and Jacobs, "Lambertian reflectance and linear subspaces," ICCV 2001.
- Ramamoorthi and Hanrahan, "A signal-processing framework for inverse rendering," SIGGRAPH 2001.
- Sloan et al., "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments," SIGGRAPH 2002.

Three papers describing the use of spherical harmonics to model low-frequency illumination, as well as the low-pass filtering effect of Lambertian reflectance on illumination.

- Zhang et al., "Shape-from-shading: a survey," PAMI 1999.

A review of perceptual and computational aspects of shape from shading.

- Woodham, "Photometric method for determining surface orientation from multiple images," Optical Engineering 1980.

The paper that introduced photometric stereo.

- Yuille and Snow, "Shape and albedo from multiple images using integrability," CVPR 1997.
- Belhumeur et al., "The bas-relief ambiguity," IJCV 1999.
- Papadhimitri and Favaro, "A new perspective on uncalibrated photometric stereo," CVPR 2013.

Three papers discussing uncalibrated photometric stereo. The first paper shows that, when the lighting directions are not known, by assuming integrability, one can reduce unknowns to the bas-relief ambiguity. The second paper discusses the bas-relief ambiguity in a more general context. The third paper shows that, if instead of an orthographic camera one uses a perspective camera, this is further reduced to just a scale ambiguity.

- Alldrin et al., "Resolving the generalized bas-relief ambiguity by entropy minimization," CVPR 2007.

A popular technique for resolving the bas-relief ambiguity in uncalibrated photometric stereo.

- Zickler et al., "Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction," IJCV 2002.

A method for photometric stereo reconstruction under arbitrary BRDF.

- Nayar et al., "Shape from interreflections," IJCV 1991.
- Chandraker et al., "Reflections on the generalized bas-relief ambiguity," CVPR 2005.

Two papers discussing how one can perform photometric stereo (calibrated or otherwise) in the presence of strong interreflections.

- Frankot and Chellappa, "A method for enforcing integrability in shape from shading algorithms," PAMI 1988.
- Agrawal et al., "What is the range of surface reconstructions from a gradient field?," ECCV 2006.

Two papers discussing how one can integrate a normal field to reconstruct a surface.