Tackles questions like what is appearance? How do pixels get their value?

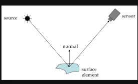

Image intensities = f (normal, surface reflectance, illumination)
Surface Reflection depends on both the viewing and illumination direction.

BRDF: Bidirectional Reflectance Distribution Function

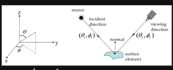

ϕ θ
azimuth and polar angles

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

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

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

Properties: (applicable to most)

- Rotational symmetry (Isotropy):
  - BRDF does not change when surface is rotated about the normal
  - BRDF is only a function of 3 variables:
    $$f(\theta_i, \theta_r, \phi_i - \phi_r)$$
- Helmholtz Reciprocity: follows from 2nd law of thermodynamics
  BRDF does not change when source and viewing directions are swapped.
  $$f(\theta_i, \phi_i ; \theta_r, \phi_r) = f(\theta_r, \phi_r ; \theta_i, \phi_i)$$

Mechanisms of Reflection

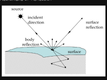

Body Reflection:
Diffuse Reflection
Matte Appearance
Non-Homogeneous Medium
Clay, paper, skin, etc

Surface Reflection:
Specular Reflection
Glossy Appearance
Highlights
Dominant for Metals

Many materials exhibit both Reflections

Color is because of body reflection,
highlights/gloss is because of surface reflection.

Image Intensity = Body Reflection + Surface Reflection

Lambert's Cosine Law:

The radiant intensity observed from a diffusely reflecting or emitting surface is directly proportional to the cosine of the angle between the surface normal and the observation direction.

Explains how bright a surface looks depending on the angle you're viewing it.
- If you're looking straight at the surface (at a 90° angle, directly above it), it looks the brightest. $\theta = 0^\circ$
- As you move to the side and view it at an angle, it becomes dimmer.
- The dimming happens because the light spreads out more as the angle increases.

It's like shining a flashlight on a wall. When the flashlight is straight on, the spot is bright and focused. When you tilt the flashlight, the light spreads out, and the spot becomes dimmer.

This rule is often used to describe how dull, matte surfaces reflect light evenly in all directions, making them appear less shiny. It's also used in computer graphics to make objects look more realistic.

$I = I_0 \cos\theta$

I = observed intensity
I_0 = intensity when the observer is directly perpendicular to the surface
θ = The angle between the surface normal (a perpendicular line to the surface) and the direction of observation.

If $\theta >= 90^\circ$, we get a shadow.

Diffuse Reflection and Lambertian BRDF

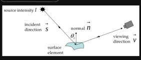

- Surface appears equally bright from all directions. This means that the viewing direction $\bar{v}$ does not matter.
- Incident light is reflected equally in all directions.
- Lambertian BRDF is simply a constant:
  $$f(\theta_i, \phi_i ; \theta_r, \phi_r) = \frac{\rho_d}{\pi}$$
  albedo
- Surface Radiance:
  $$L = \frac{\rho_d}{\pi} I \cos\theta_i = \frac{\rho_d}{\pi} I \, \bar{n} \cdot \bar{s}$$
  source intensity

Photometric Stereo

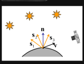

Lambertian case:

$$I = \frac{\rho}{\pi} k c \cos\theta_i = \rho \bar{n} \cdot \bar{s} \qquad \left[\frac{kc}{\pi} = 1\right]$$

$$I = \max\left(\frac{\rho}{\pi} kc \cos\theta, 0\right)$$ for cases when $\theta > 90^\circ$

Image irradiance:
$I_1 = \rho \bar{n} \cdot \bar{s}_1$
$I_2 = \rho \bar{n} \cdot \bar{s}_2$
$I_3 = \rho \bar{n} \cdot \bar{s}_3$

$\bar{n}$ has 2 DOF
ρ is 1 unknown
I is 3 unknowns

$$\underset{3\times1}{\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix}} = \rho \underset{3\times3}{\begin{bmatrix} s_1^T \\ s_2^T \\ s_3^T \end{bmatrix}} \bar{n}$$

$$I = S \underbrace{\rho n}_{\tilde{n}} \Rightarrow \tilde{n} = S^{-1} I$$

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

surface normal $n = \frac{\tilde{n}}{\rho}$

To get better results, we use more light sources

$$\underset{k\times1}{\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_k \end{bmatrix}} = \underset{k\times3}{\begin{bmatrix} s_1^T \\ s_2^T \\ \vdots \\ s_k^T \end{bmatrix}} \rho n$$

There have to be at least 3 linearly independent light sources out of the k

$$I = S \tilde{n}$$
$$S^T I = S^T S \tilde{n}$$
$$\tilde{n} = \underbrace{(S^T S)^{-1} S^T}_{\text{Moore-Penrose pseudo inverse}} I$$

S is non-invertible if it has less rank. This means that 2 or more light sources lie on the same plane, or one light source is a linear combination of the other two.

So for S to be invertible, all rows have to be linearly independent. This means that all light sources have to lie on different planes.

Depth from Normals

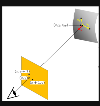

$$V_1 = (x+1, y, z_{x+1,y}) - (x, y, z_{xy})$$
$$= (1, 0, z_{x+1,y} - z_{xy})$$

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

Get a similar equation for V2
– Each normal gives us two linear constraints on z
– compute z values by solving a matrix equation

Computing Source Direction

Calibration problem

Use an object of know shape/BRDF. This means that we know the surface normal.
Find the region on the object with highest intensity.
At such a spot, $\cos\theta = 1$, which means that the vectors $\bar{n}$ and $\bar{s}$ are the same.
This way we compute the source direction.

If there are 2 (or more) light sources,

$$I = n s_1 + n s_2 = n \cdot (s_1 + s_2)$$

The two light sources can be replaced by a single effective light source.

Uncalibrated Photometric Stereo

Bas Relief Ambiguity

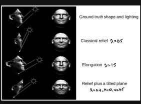

What is Bas-Relief Ambiguity?
Bas-relief ambiguity refers to the challenge of distinguishing between different shapes that can produce the same shading patterns under the same lighting conditions. Essentially, slightly distorted or stretched versions of an object's surface can look identical in terms of how light interacts with them.

Breakdown of the Image:

Ground Truth (Top Row):
The actual shape of the face and the lighting setup are shown here. This is the correct 3D model we ideally want to reconstruct.

Classical Relief (λ=0.5):
The surface is "flattened" (scaled down along the depth axis). Even though the depth is halved, the lighting and shading appear the same, making it hard to distinguish from the ground truth.

Elongation (λ=1.5):
The surface is "stretched" (scaled up along the depth axis). Again, the shading looks identical to the ground truth, even though the depth is exaggerated.

Relief with a Tilted Plane (λ=0.7,μ=0,ν=0.5):
Here, the surface is scaled and also tilted, adding even more ambiguity. Despite the distortions, the shading pattern remains consistent with the original lighting.

Why Does This Happen?
The problem arises because shading depends on the angles between the light, the surface normal, and the viewer. Certain transformations of the object's shape (like scaling or tilting) don't change these angles, leading to similar shading.

For one pixel,

$$\underset{3\times1}{\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix}} = \rho \underset{3\times3}{\begin{bmatrix} s_1^T \\ s_2^T \\ s_3^T \end{bmatrix}} \bar{n}$$

For p pixels,

$$\underset{M}{\underset{p\times3}{\begin{bmatrix} I_1^1 & I_2^1 & I_3^1 \\ I_1^2 & I_2^2 & I_3^2 \\ \vdots & \vdots & \vdots \\ I_1^p & I_2^p & I_3^p \end{bmatrix}}} = \underset{N^T}{\underset{p\times3}{\begin{bmatrix} n_1^{1T} & n_2^{1T} & n_3^{1T} \\ n_1^{2T} & n_2^{2T} & n_3^{2T} \\ \vdots & \vdots & \vdots \\ n_1^{pT} & n_2^{pT} & n_3^{pT} \end{bmatrix}}} \; A^{-1}A \; \underset{S}{\underset{3\times3}{\begin{bmatrix} s_1^T \\ s_2^T \\ s_3^T \end{bmatrix}}}$$

$$M = N^T (A^{-1} A) S$$
$$M = (A^{-T} N)^T (AS)$$

A is the ambiguity matrix

This means that if the surface S is transformed by A, the normal must be transformed by A^{-T}

We want to estimate both the lighting directions $\hat{L}$ and the pseudonormals $\hat{B}$ from the intensity matrix I without prior knowledge of either L or B. To solve this, we first perform matrix factorization on I via SVD. For a matrix I of size $m \times n$, SVD is given as

$$I = U \Sigma V^T$$

For the best rank 3 approximation of I, we truncate its singular values Σ to keep only the top 3 singular values and set all others to 0. Reconstructing this, we get

$$\hat{I} = U_3 \Sigma_3 V_3^T$$

where

- $\hat{I}$: $m \times n$ is the estimated intensity matrix
- Σ_3 : 3×3 is a diagonal matrix of top 3 singular values
- U_3 : $m \times 3$ is the first 3 columns of U
- V_3^T : $3 \times n$ is the first 3 rows of V^T

We truncate them to 3 because rank of $\hat{I}$ must be 3 to reflect the physical constraint that the surface normals are 3D vectors.

$$\Sigma_3 = \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & \sigma_3 \end{bmatrix}$$
$$U_3 = U[:, :3]$$
$$V_3^T = V^T[:3, :]$$

We compute the square root of Σ_3 to split its scaling factors.

$$\Sigma_3^{1/2} = \begin{bmatrix} \sqrt{\sigma_1} & 0 & 0 \\ 0 & \sqrt{\sigma_2} & 0 \\ 0 & 0 & \sqrt{\sigma_3} \end{bmatrix}$$

We then factorize $\hat{I}$ as

$$\hat{L}^T = U_3 \Sigma_3^{1/2}$$
$$\hat{B} = \Sigma_3^{1/2} V_3^T$$

These estimated $\hat{L}^T$ and $\hat{B}$ satisfy the dimensions and rank constraints:

- $\hat{L}^T$: $m \times 3$ represents the lighting directions
- $\hat{B}$: $3 \times n$ represents the pseudonormals

```python
def estimatePseudonormalsUncalibrated(I):
    '''
    Question 2 (b)

    Estimate pseudonormals without the help of light source directions.

    Parameters
    ----------
    I : numpy.ndarray
        The 7 x P matrix of loaded images

    Returns
    -------
    B : numpy.ndarray
        The 3 x P matrix of pesudonormals

     L : numpy.ndarray
         The 3 x 7 array of lighting directions

    '''

    B = None
    L = None
    ### YOUR CODE HERE
    U, Sigma, Vt = np.linalg.svd(I, full_matrices=False)

    U_3 = U[:, :3]
    Sigma_3 = np.diag(Sigma[:3])
    Vt_3 = Vt[:3, :]

    L = U_3 @ np.sqrt(Sigma_3)
    B = np.sqrt(Sigma_3) @ Vt_3

    ### END YOUR CODE
    return B, L

# Part 2 (b)
I, L, s = loadData(data_dir)
B, LEst = estimatePseudonormalsUncalibrated(I)
albedos, normals = estimateAlbedosNormals(B)
albedoIm, normalIm = displayAlbedosNormals(albedos, normals, s)
plt.imsave('2b-a.png', albedoIm, cmap = 'gray')
plt.imsave('2b-b.png', normalIm, cmap = 'rainbow')
```

Normal, Depth and Integrability

Explicit form of surface:

$$z = f(x,y)$$

Implicit form of surface:

$$F(x,y,z) = z - f(x,y) = 0$$

Surface normal

$$\bar{n} = \nabla F(x,y,z) = \left(-\frac{\partial f}{\partial x}, -\frac{\partial f}{\partial y}, 1\right)$$

If $\bar{n} = (n_1, n_2, n_3)$

$$\frac{\partial f}{\partial x} = \frac{-n_1}{n_3} \qquad \frac{\partial f}{\partial y} = \frac{-n_2}{n_3}$$

Integrability

$$\frac{\partial^2 f}{\partial x \partial y} = \frac{\partial^2 f}{\partial y \partial x} \qquad \frac{\partial}{\partial y}\left(\frac{n_1}{n_3}\right) = \frac{\partial}{\partial x}\left(\frac{n_2}{n_3}\right)$$

Taken from HW5 F'24

The shape of the face as a 3D depth map is given by the function

$$z = f(x, y)$$

This is the explicit form of the surface.

The implicit form of the surface is written as

$$F(x, y, z) = z - f(x, y) = 0$$

The surface normal vector $\vec{n}$ is equal to the gradient of $F(x, y, z)$.

$$\vec{n} = \nabla F(x, y, z)$$

This gradient is equal to

$$\nabla F(x, y, z) = \left(\frac{\partial f(x,y)}{\partial x}, \frac{\partial f(x,y)}{\partial y}, -1\right)$$

Hence

$$\vec{n} = \left(\frac{\partial f(x,y)}{\partial x}, \frac{\partial f(x,y)}{\partial y}, -1\right)$$

If we represent the surface normal $\vec{n}$ as

$$\vec{n} = (n_1, n_2, n_3)$$

then

$$n_1 \propto \frac{\partial f(x,y)}{\partial x}$$
$$n_2 \propto \frac{\partial f(x,y)}{\partial y}$$
$$n_3 \propto -1$$

This implies

$$\frac{n_1}{n_3} = -\frac{\partial f(x,y)}{\partial x}$$
$$\frac{n_2}{n_3} = -\frac{\partial f(x,y)}{\partial y}$$

Hence

$$\frac{\partial f(x,y)}{\partial x} = f_x = -\frac{n_1}{n_3}$$
$$\frac{\partial f(x,y)}{\partial y} = f_y = -\frac{n_2}{n_3}$$

The partial derivatives f_x and f_y describe how the surface height $z = f(x, y)$ changes in the x and y directions respectively.
The ratio $-\frac{n_1}{n_3}$ gives the slope of the surface in the x direction. The ratio $-\frac{n_2}{n_3}$ gives the slope of the surface in the y direction.

If we assume the surface is integrable, then the ambiguity A reduces to

$$A = \lambda f(x,y) + \mu x + \nu y \qquad \text{3 DoF} \qquad \lambda > 0$$

$$A = G = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{bmatrix} \qquad G^{-1} = \frac{1}{\lambda}\begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ -\mu & -\nu & 1 \end{bmatrix}$$

Bas relief ambiguity

λ This scales the surface normals along the x- and y-directions.
- $\lambda > 1$ Stretches the depth (elongates the surface).
- $\lambda < 1$ Compresses the depth (flattens the surface).

μ This introduces a tilt in the xz-plane, shifting the surface normals to "lean" in the x-direction.

ν This introduces a tilt in the yz-plane, shifting the surface normals to "lean" in the y-direction.

Surface roughness causes flat appearance

The Lambertian model is valid only for smooth matte surfaces, and bad for rough matte surfaces.

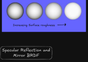

Specular Reflection and Mirror BRDF

- Very smooth and shiny objects
- All incident light energy is reflected in a single direction (only when $\bar{v} = \bar{r}$)

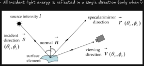

- Mirror BRDF is a double-delta function

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

Ensures the incident angle equals the viewing angle

Ensures the azimuthal angle of reflection is the mirror of the incident azimuthal angle (shifted by π).
This means the reflection "flips" the azimuthal direction of the incident light around the surface normal.

Without this, the reflected light could "veer off" horizontally in any direction, breaking the symmetry of mirror-like reflection. Reflections in computer-generated scenes could appear in random directions, breaking realism.

- Surface radiance

$$L = I \rho_s \, \delta(\theta_i - \theta_v) \, \delta(\phi_i + \pi - \phi_v)$$

You only see reflected light if you're looking directly along the reflected direction, making specular reflections very directional and sharp.

Glossy Surfaces

- Delta Function too harsh a BRDF model - valid only for highly polished mirrors and metals
- Many glossy surfaces show broader highlights in addition to mirror reflection.
- Surfaces are not perfectly smooth – they show micro-surface geometry (roughness).
- Example Models : Phong model, Torrance Sparrow model

Blurred Highlights and Surface Roughness