Photometric Stereo

Introduction to Computer Vision I

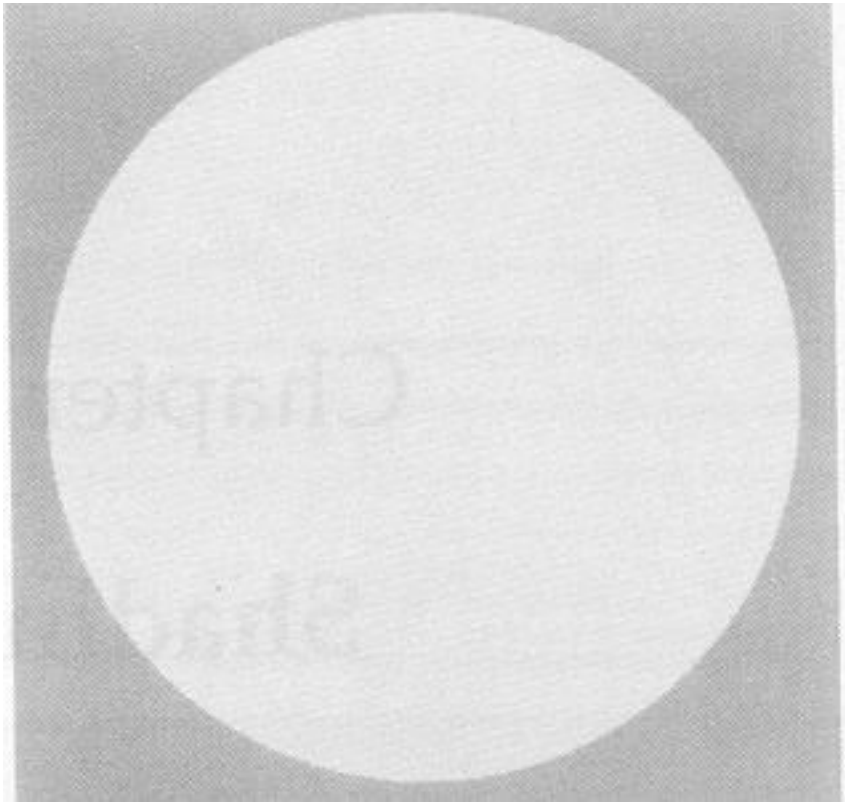
CSE 152A

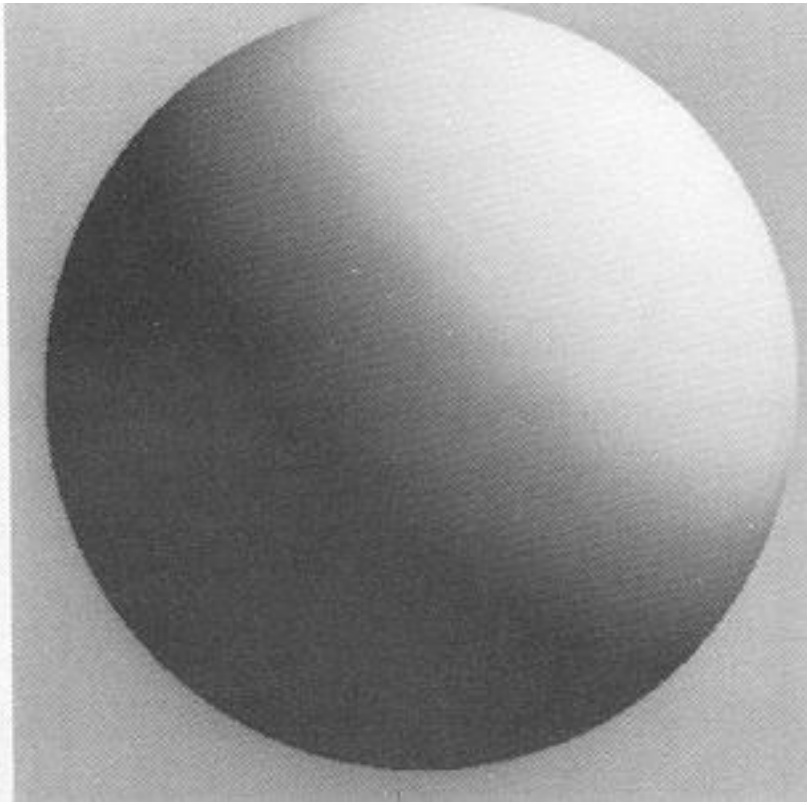
Lecture 4

Announcements

- Assignment 0 is due today, 11:59 PM
- Assignment 1 will be released today
  - Due Oct 21, 11:59 PM

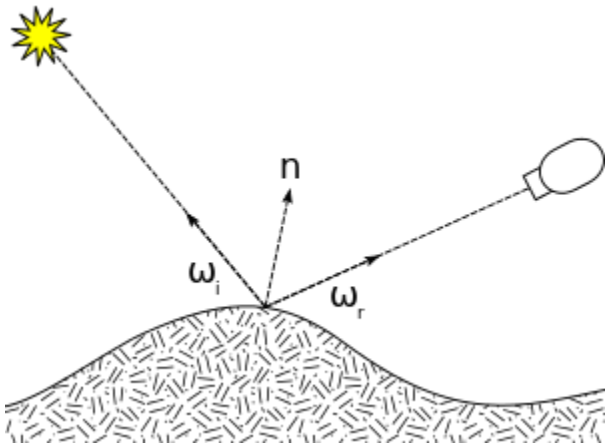
Shading reveals 3D surface geometry





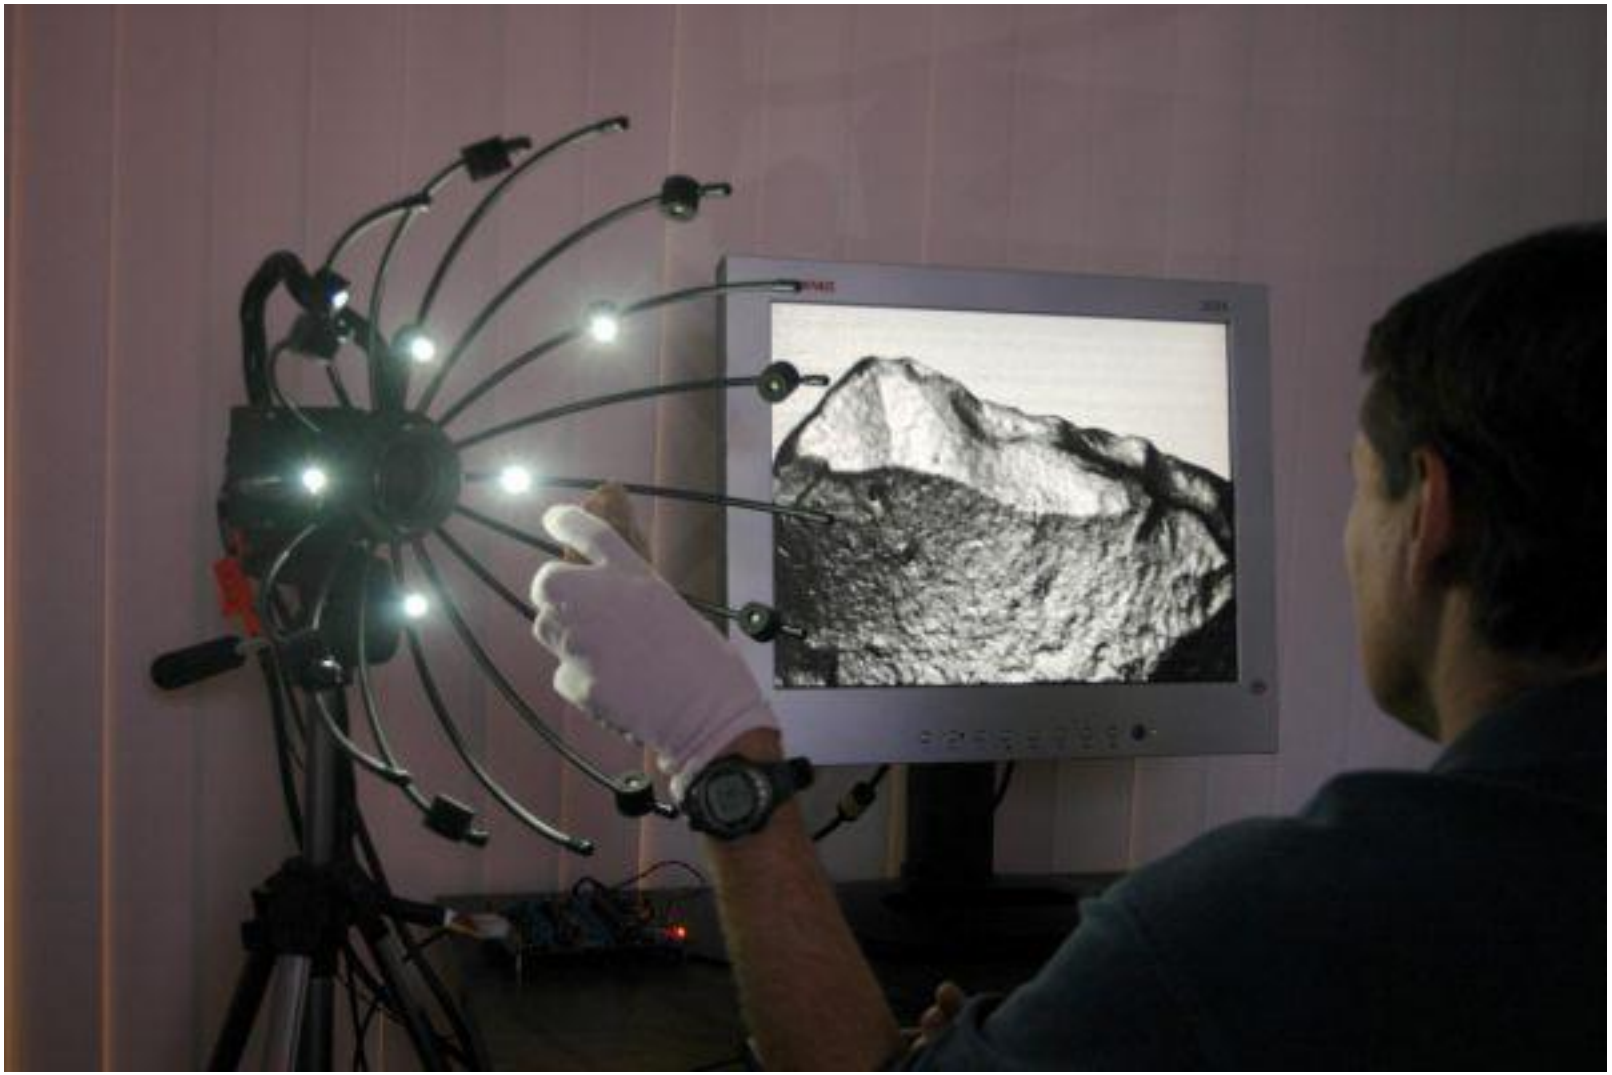
Two shape-from-X methods that use shading

- Shape-from-shading: Use just one image to recover shape. Requires knowledge of light source direction and BRDF everywhere. Too restrictive (good candidate for deep learning)
- Photometric stereo: Single viewpoint, multiple images under different lighting.

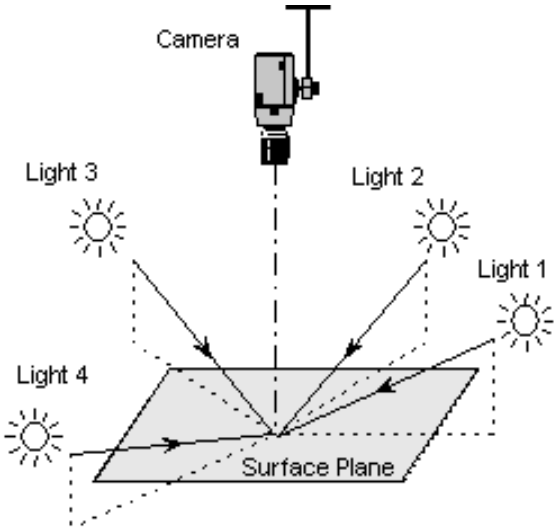


BRDF
(four dimensional function)

Photometric Stereo Rigs: One viewpoint, changing lighting



An example of photometric stereo





surface
(albedo texture map)

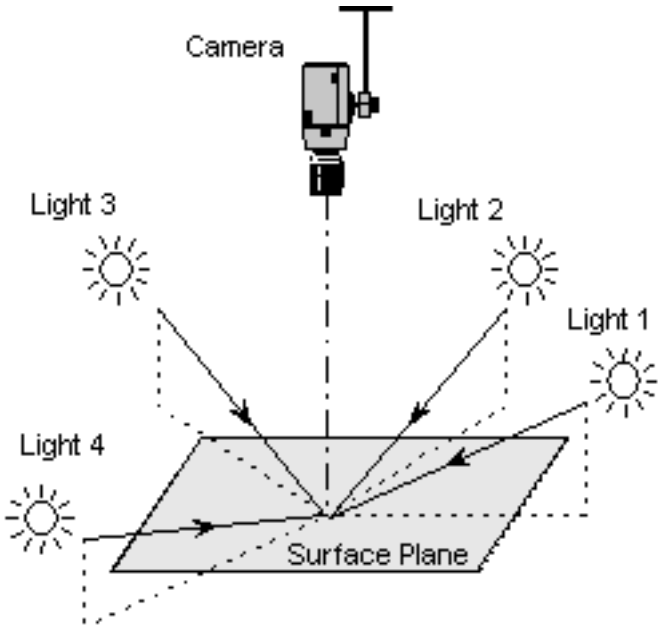




albedo

+ surface normals

Photometric stereo



- Single viewpoint, multiple images under different lighting
  1. General BRDF, known lighting
  2. Lambertian BRDF, known lighting
  3. Lambertian BRDF, unknown lighting

1. Photometric Stereo: General BRDF and Reflectance Map

BRDF

- Bi-directional Reflectance Distribution Function

$$\rho(\theta_{in}, \phi_{in}; \theta_{out}, \phi_{out})$$

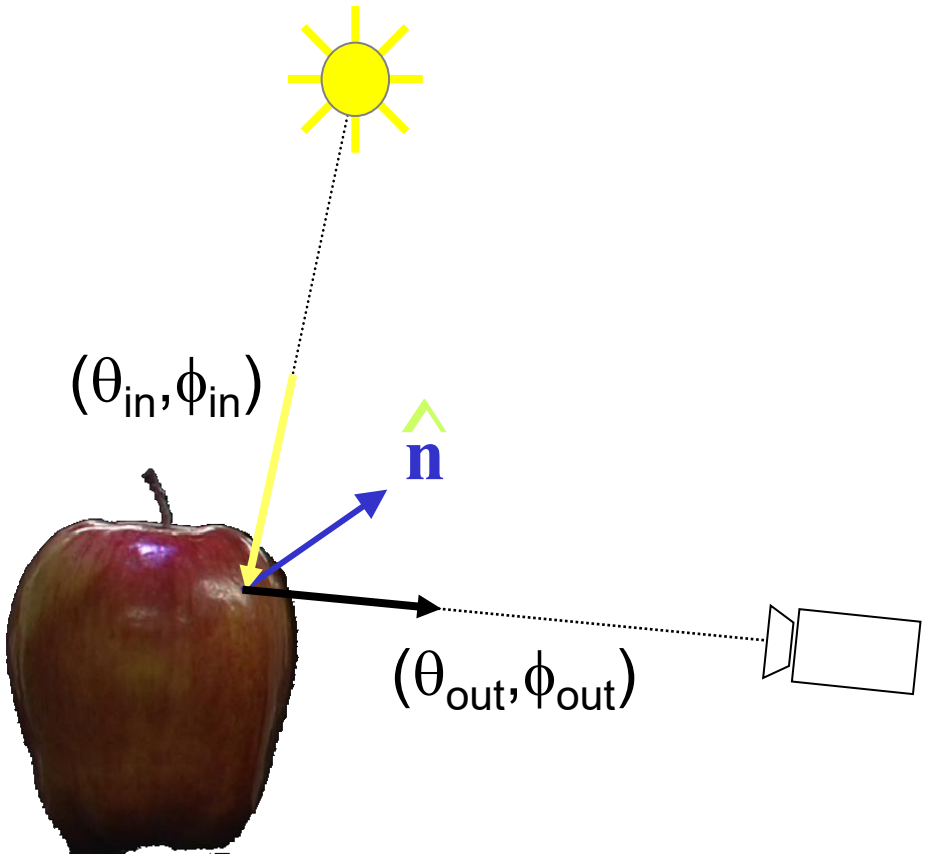
- Function of
  - Incoming light direction:

$$\theta_{in}, \phi_{in}$$

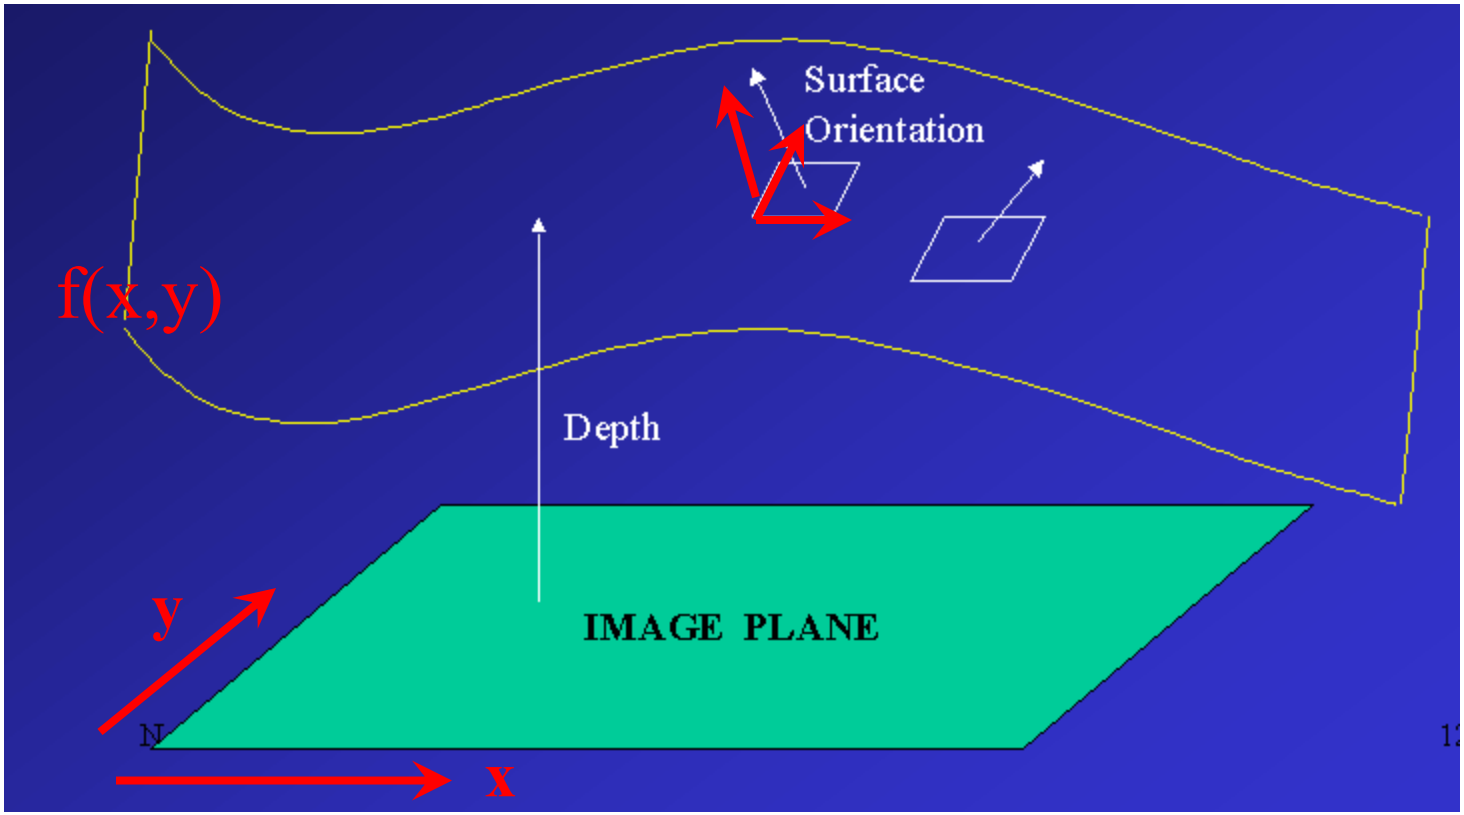
  - Outgoing light direction:

$$\theta_{out}, \phi_{out}$$

- Ratio of emitted radiance to incident irradiance



Coordinate system



Surface: **s**(x,y) =(x,y, f(x,y))

Tangent vectors: $\frac{\partial s(x,y)}{\partial x} = \left(1, 0, \frac{\partial f}{\partial x}\right)^\top$

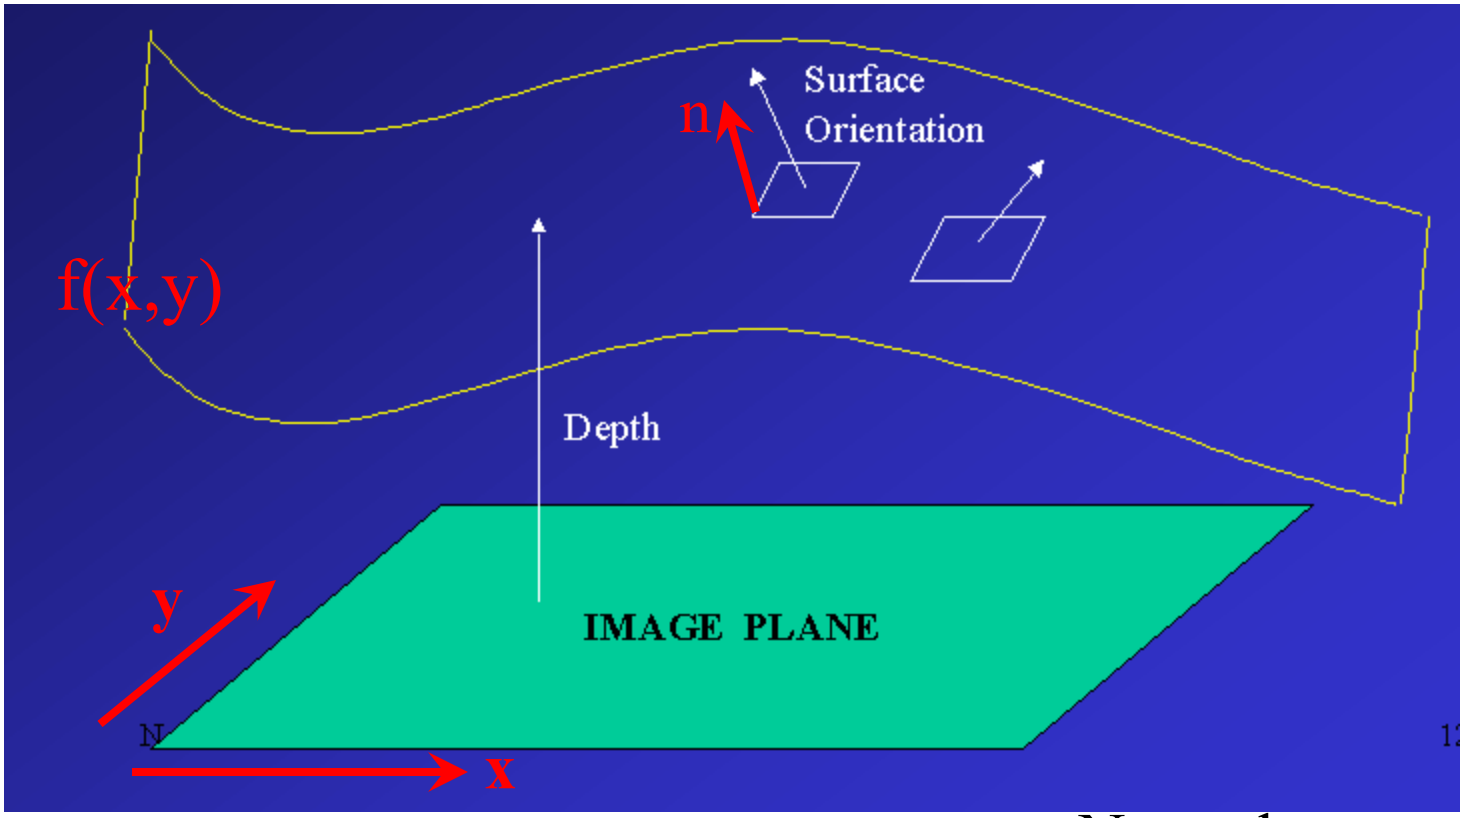
$\frac{\partial s(x,y)}{\partial y} = \left(0, 1, \frac{\partial f}{\partial y}\right)^\top$

Normal vector

$$\mathbf{n} = \frac{\partial s}{\partial x} \times \frac{\partial s}{\partial y}$$

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

Gradient Space (p,q)



Gradient Space : (p,q)

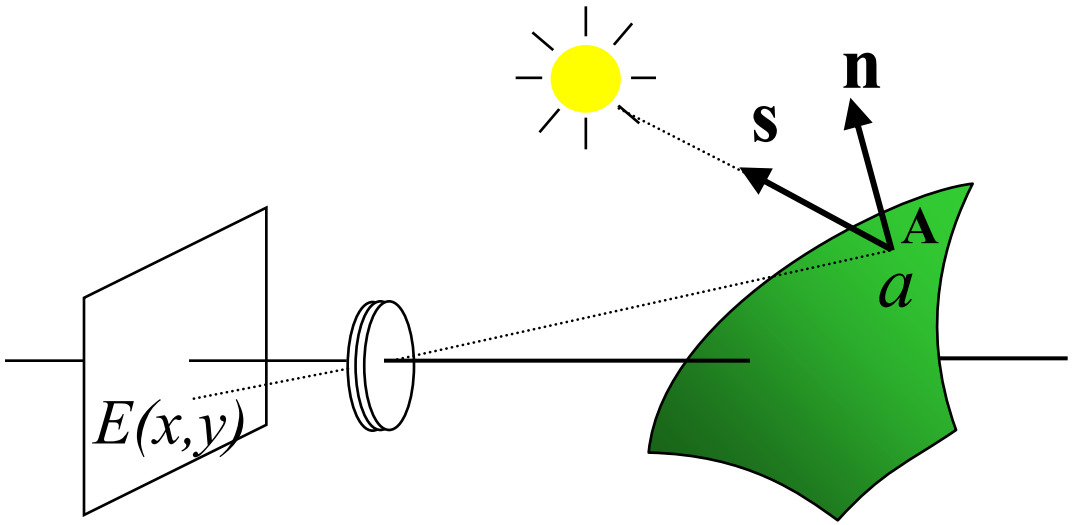
$$p = \frac{\partial f}{\partial x}, \quad q = \frac{\partial f}{\partial y}$$

Normal vector

$$\mathbf{n} = \frac{\partial s}{\partial x} \times \frac{\partial s}{\partial y} = \left(-\frac{\partial f}{\partial x}, -\frac{\partial f}{\partial y}, 1 \right)^T$$

$$\hat{\mathbf{n}} = \frac{1}{\sqrt{p^2 + q^2 + 1}} (-p, -q, 1)^T$$

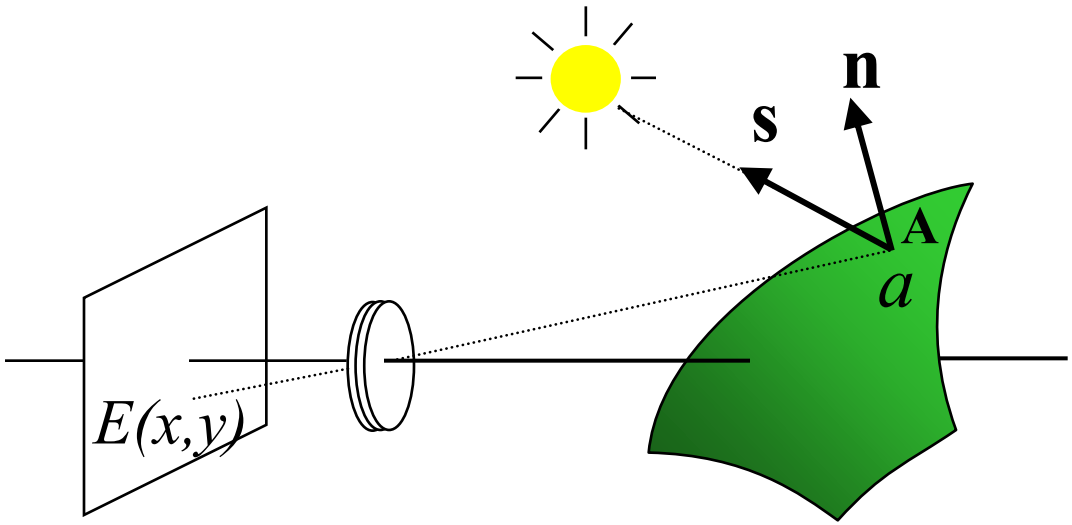
Image Formation



For a given point A on the surface *a*, the image irradiance E(x,y) is a function of

1. The BRDF at A
2. The surface normal at A
3. The direction of the light source

Reflectance Map

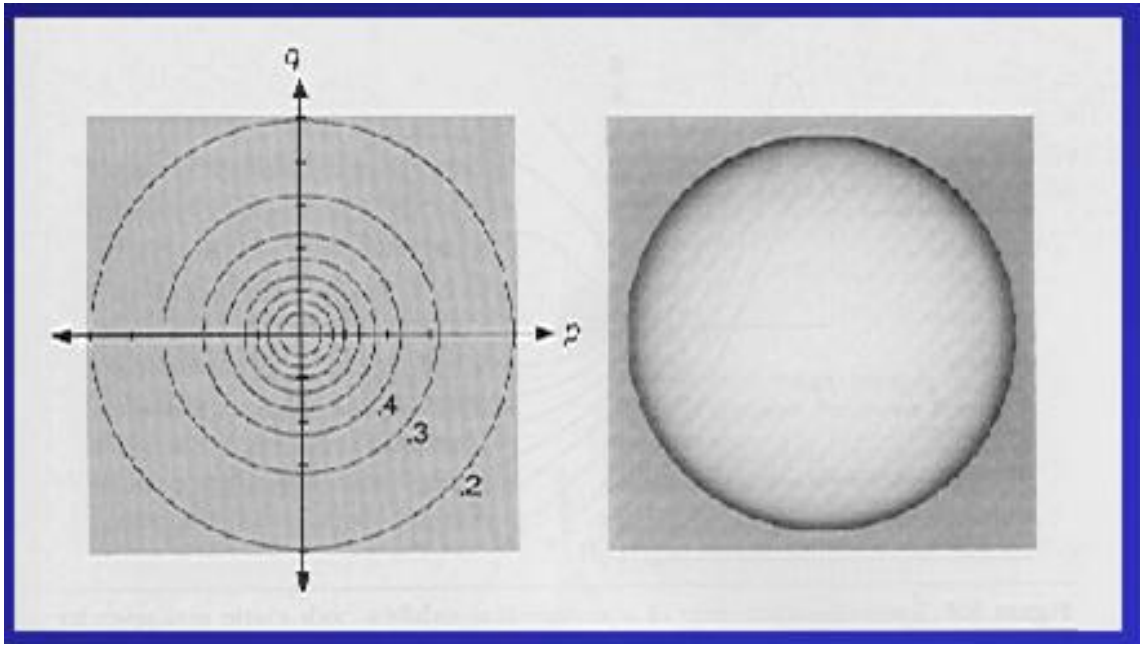


Let the BRDF be the same at all points on the surface, and let the light direction **s** be a constant

1. Then image irradiance is a function of only the direction of the surface normal
2. In gradient space, we have E(p,q)

Example Reflectance Map: Lambertian surface

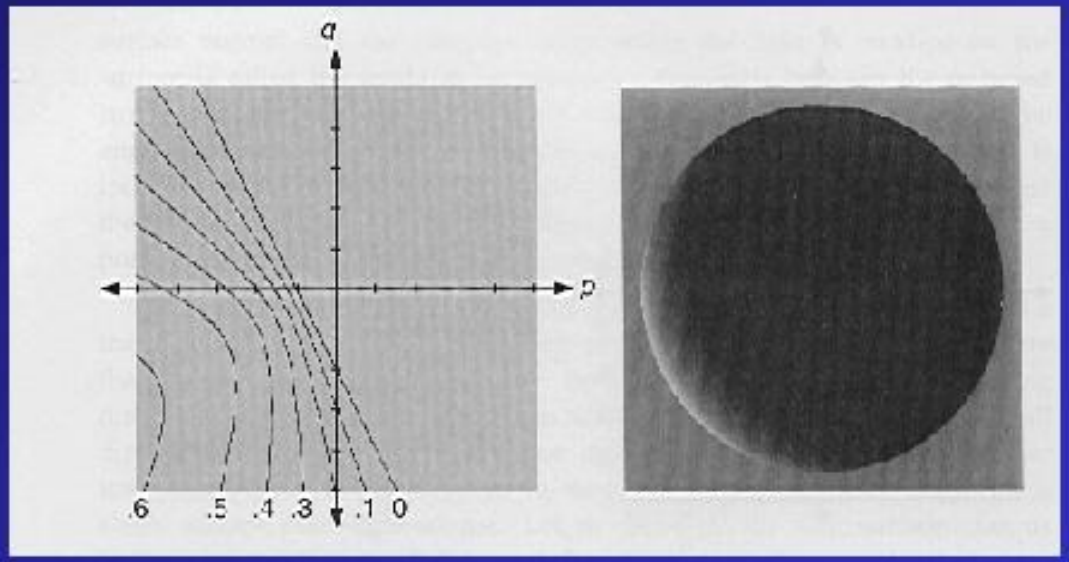
E(p,q)



For lighting from front

LAMBERTIAN REFLECTANCE MAP

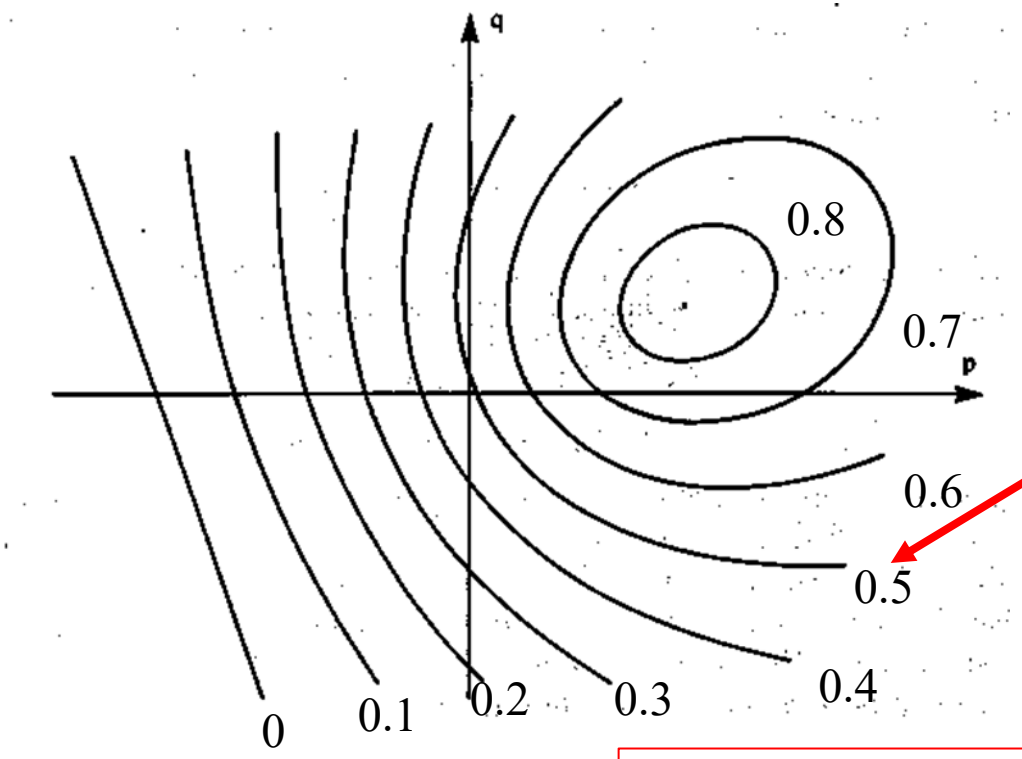
$$E = L\rho \frac{1 + pp_s + qq_s}{\sqrt{1 + p^2 + q^2}\sqrt{1 + p_s^2 + q_s^2}}$$

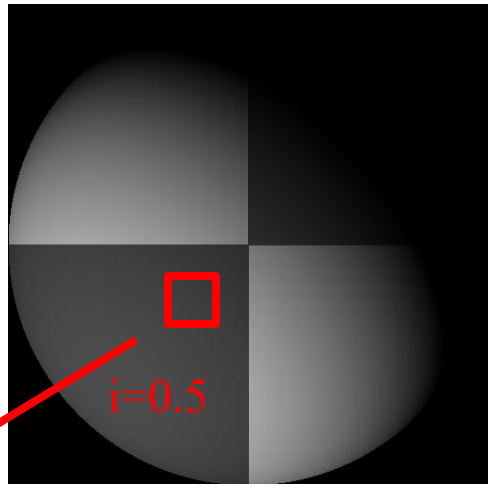


$p_s = -2 \quad q_s = -1$

Light Source Direction, expressed in gradient space.

Reflectance Map of Lambertian Surface





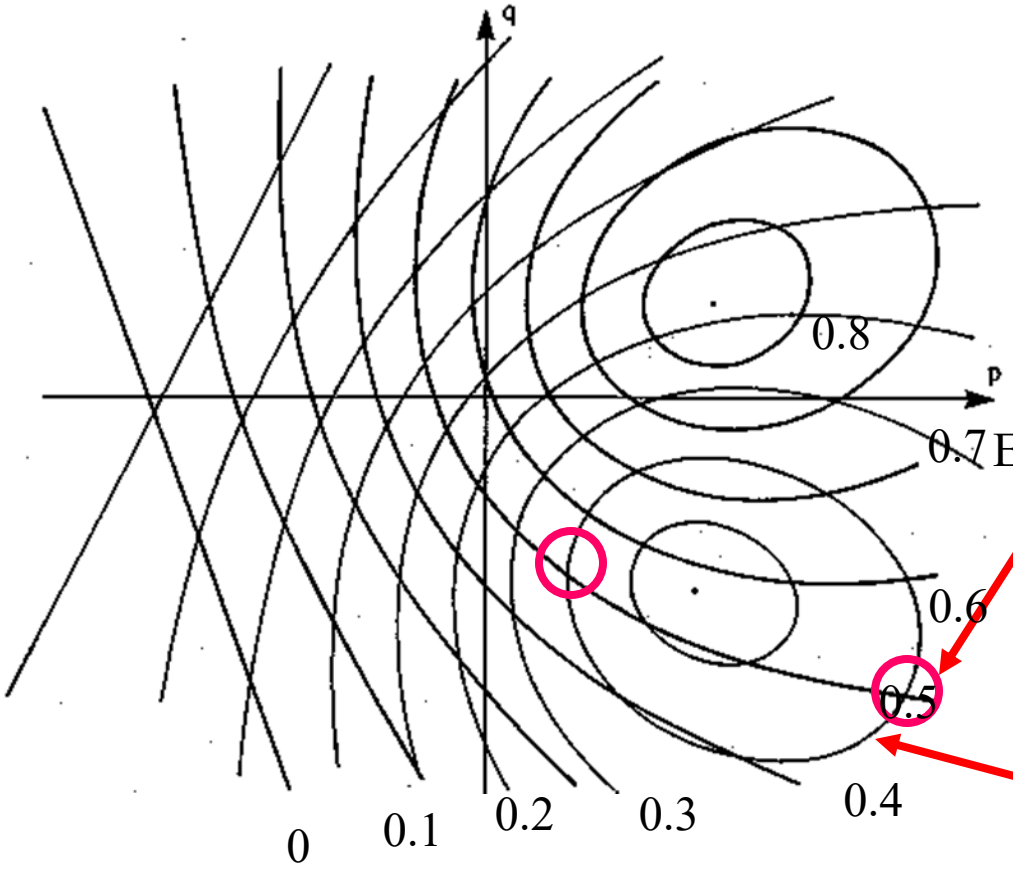
E.g., Normal lies on this curve

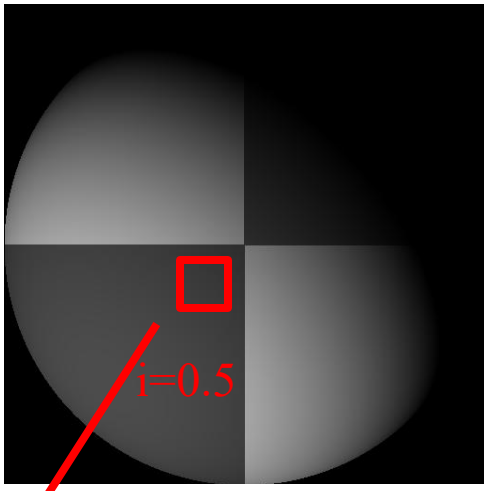
Curves are not really for this image

What does the intensity (irradiance) of one pixel in one image tell us?

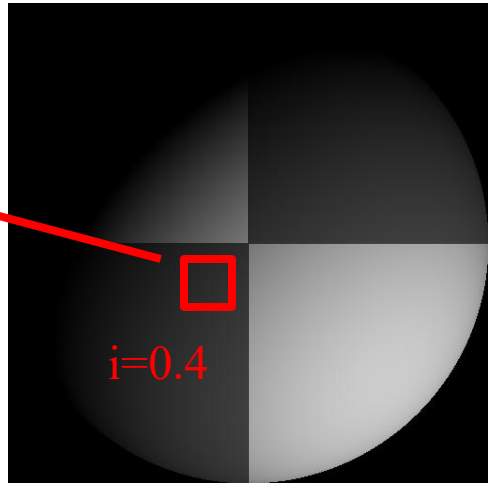
It constrains the surface normal projecting to that point to a curve

Two Light Sources, Two reflectance maps





E.g., Normal lies on this curve



Curves are not really for these images

A third image would disambiguate match

Three Source Photometric stereo: Step1

Offline:

Using source directions & BRDF, construct reflectance map for each light source direction. $R_1(p,q)$, $R_2(p,q)$, $R_3(p,q)$

Online:

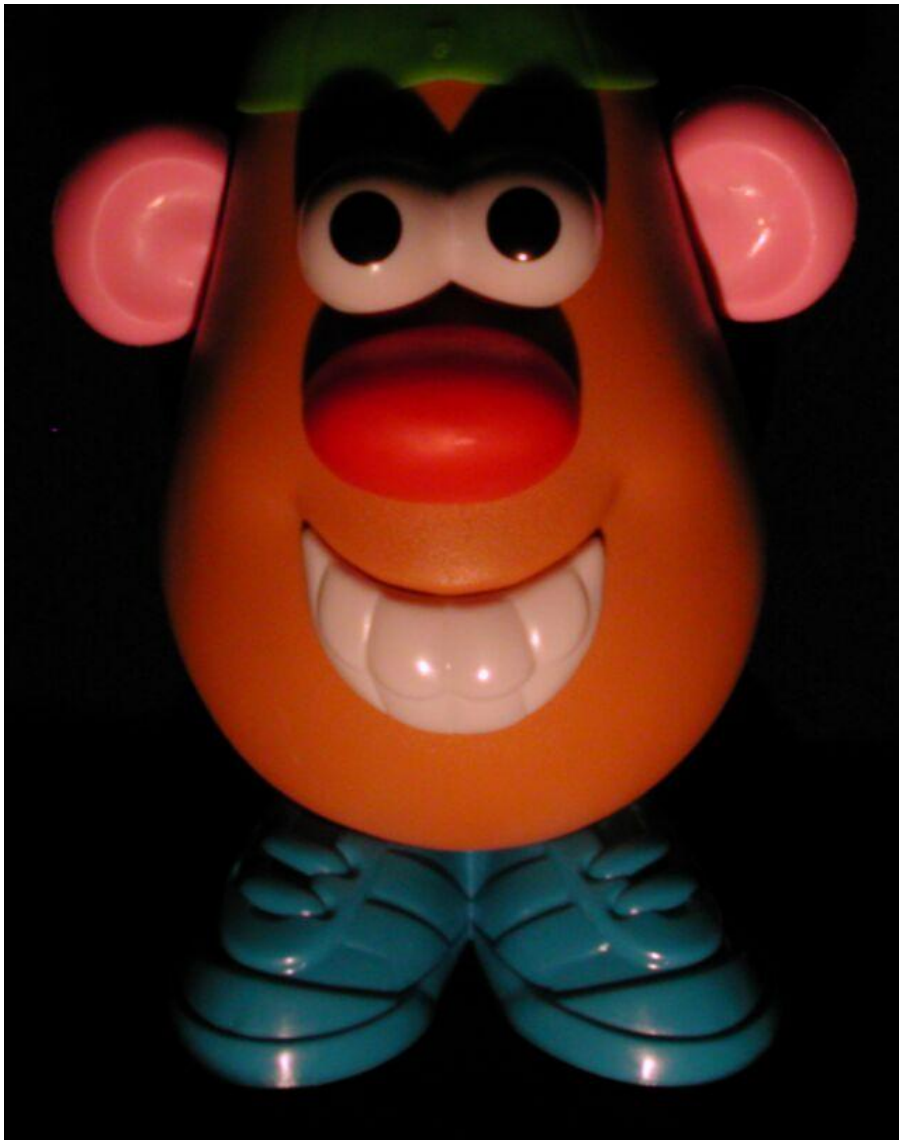
1. Acquire three images with known light source directions. $E_1(x,y)$, $E_2(x,y)$, $E_3(x,y)$
2. For each pixel location (x,y), find (p,q) as the intersection of the three curves

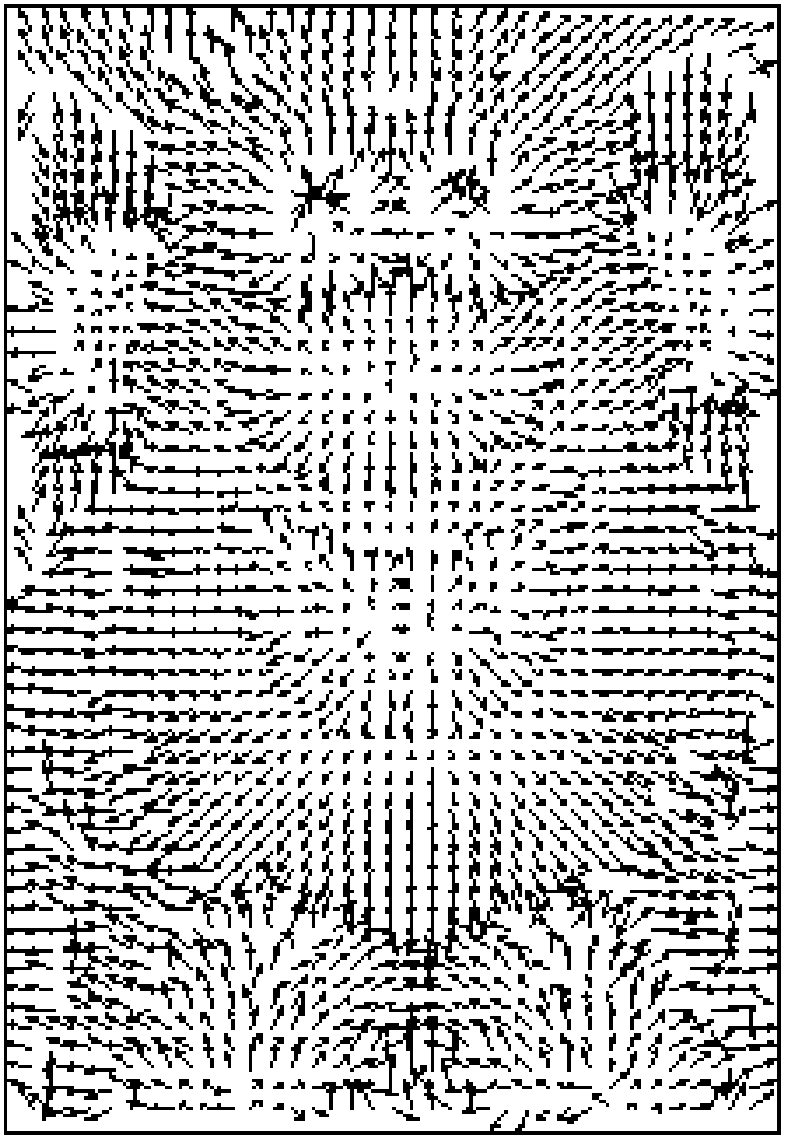
   $R_1(p,q)=E_1(x,y)$

   $R_2(p,q)=E_2(x,y)$

   $R_3(p,q)=E_3(x,y)$
3. This is the surface normal at pixel (x,y). Over image, the normal field is estimated

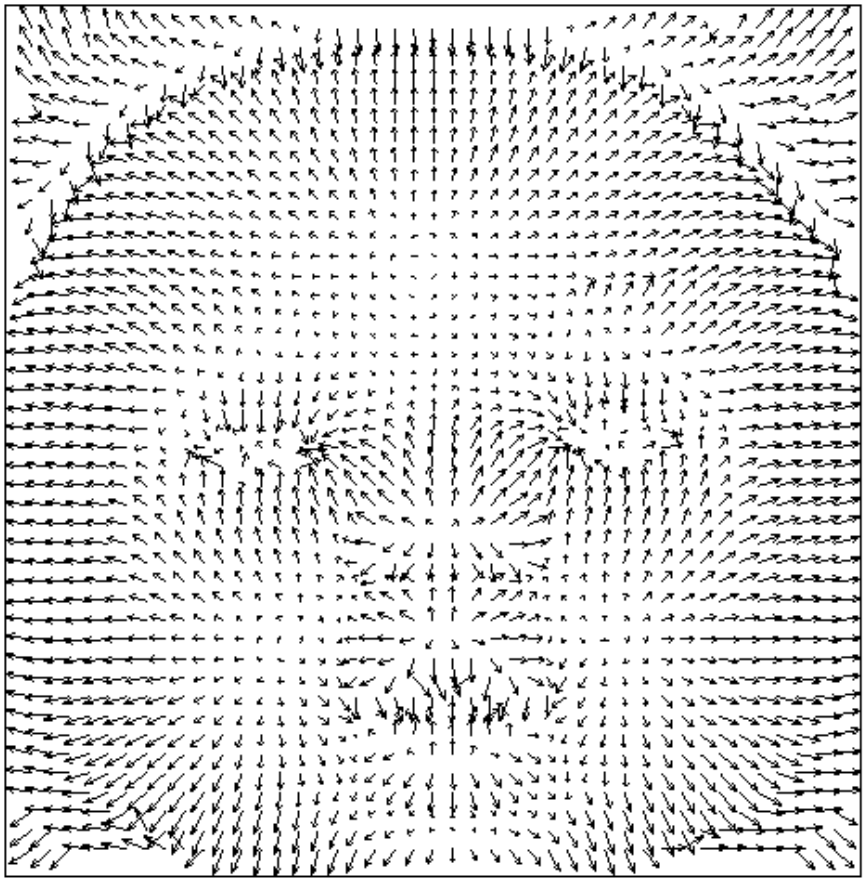
Normal Field





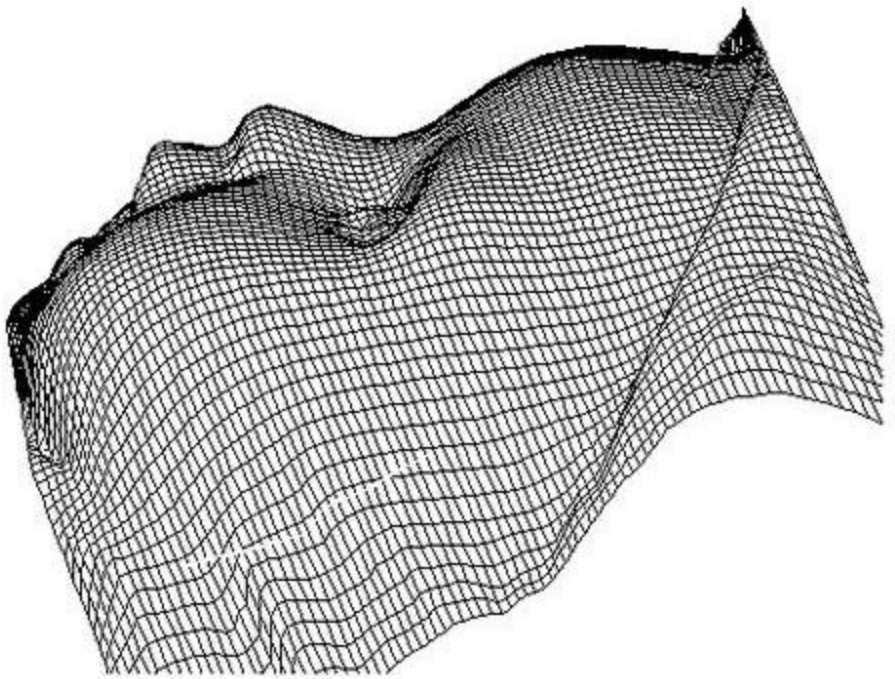
Plastic Baby Doll: Normal Field





Next step:
Go from normal field to surface



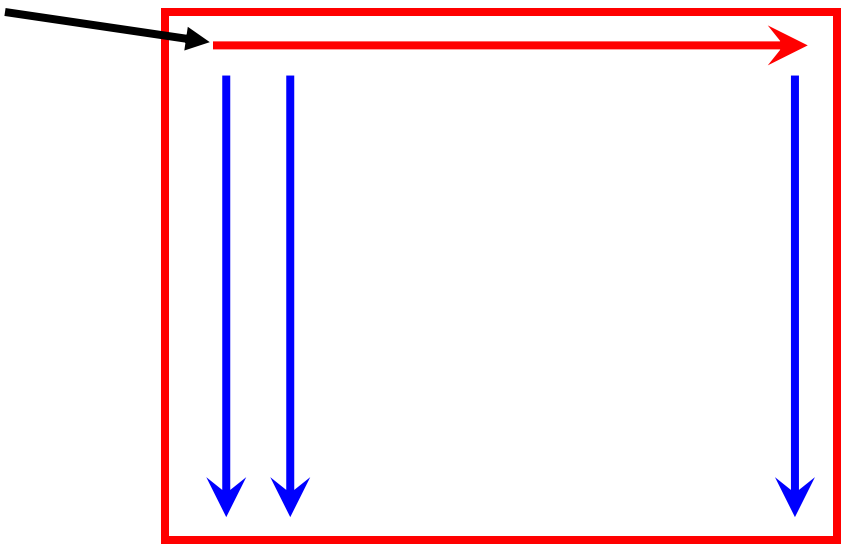


Recovering the surface f(x,y)

Many methods: Simplest approach

1. From estimate $\mathbf{n} = (n_x, n_y, n_z)$, $p = -n_x/n_z$, $q = -n_y/n_z$
2. Integrate $p = df/dx$ along row $(x,0)$ to get $f(x,0)$
3. Then integrate $q = df/dy$ along each column starting with value of the first row

Start here, setting f(0,0) = 0

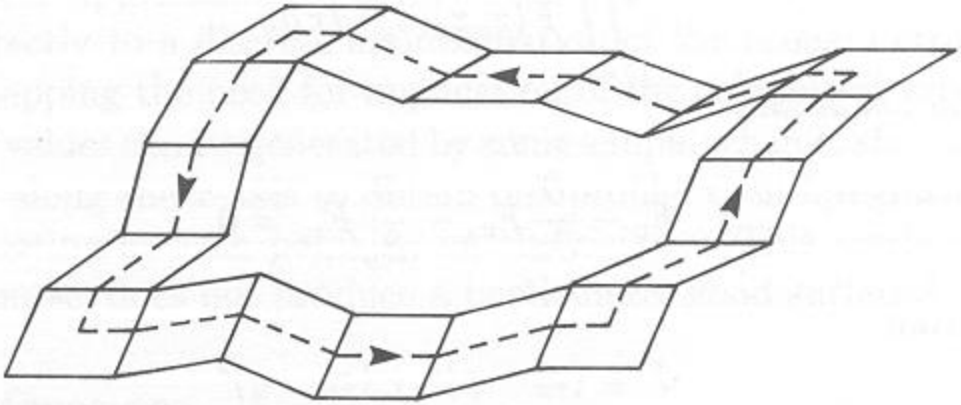


f(x,0)

Use a mask to only include object

When done, some values may be negative. Subtract all values from the minimum value.

What might go wrong?



• Height z(x,y) is obtained by integration along a curve from (x_0, y_0).

$$z(x,y) = z(x_0,y_0) + \int_{(x_0,y_0)}^{(x,y)} (pdx + qdy)$$

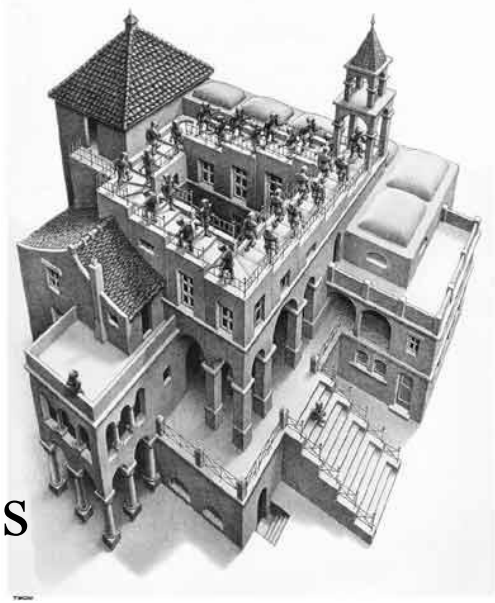
• If one integrates the derivative field along any closed curve, one expects to get back to the starting value.

• Might not happen because of noisy estimates of (p,q)

What might go wrong?

Integrability. If f(x,y) is the height function, we expect that

$$\frac{\partial}{\partial y}\frac{\partial f}{\partial x} = \frac{\partial}{\partial x}\frac{\partial f}{\partial y}$$



In terms of estimated gradient space (p,q), this means:

$$\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x}$$

But since p and q were estimated indpendently at each point as intersection of curves on three reflectance maps, equality is not going to exactly hold

Horn's Method

["Robot Vision", B.K.P. Horn, 1986]

- Formulate estimation of surface height z(x,y) from gradient field by minimizing cost functional:

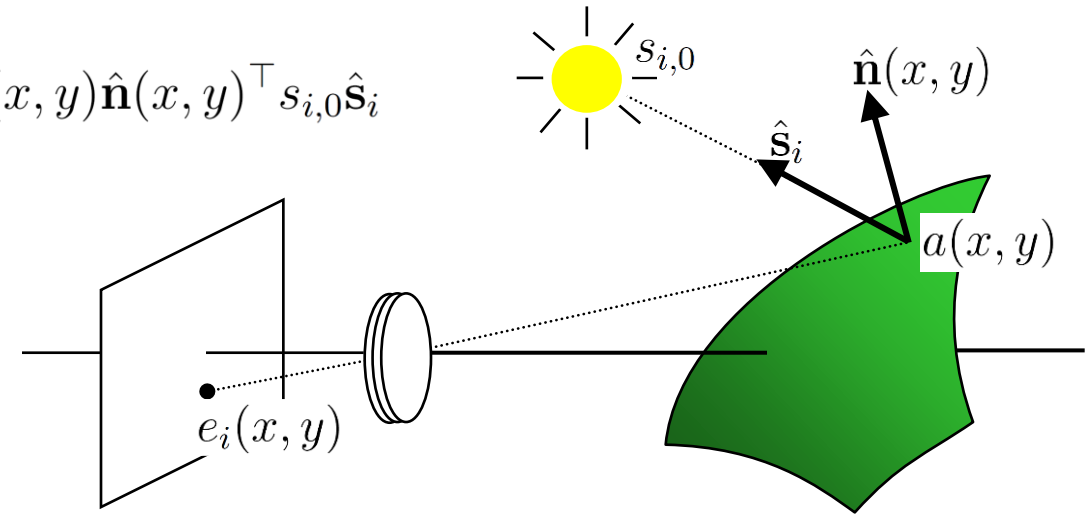
$$\iint_{\text{Image}} (z_x - p)^2 + (z_y - q)^2 \, dxdy$$

  where (p,q) are estimated components of the gradient while z_x and z_y are partial derivatives of best fit surface
- Solved using calculus of variations – iterative updating
- z(x,y) can be discrete or represented in terms of basis functions
- Integrability is naturally satisfied

2. Photometeric Stereo: Lambertian Surface, Known Lighting

Photometric stereo, Lambertian surface

$$e_i(x,y) = a(x,y)\hat{\mathbf{n}}(x,y)^\top s_{i,0}\hat{\mathbf{s}}_i$$



$s_{i,0}$ is the intensity of the i-th light source

$\hat{\mathbf{s}}_i$ is the i-th light source direction at the surface point

$\hat{\mathbf{n}}(x,y)$ is the normal at the surface point projected to the image coordinates (x,y)

$a(x,y)$ is the albedo of the surface point projected to the image coordinates (x,y)

$e_i(x,y)$ is the intensity of the light reflected from the surface point projected to the i-th image coordinates (x,y)

Photometric stereo, Lambertian surface

For each pixel

$$e_i(x,y) = a(x,y)\hat{\mathbf{n}}(x,y)^\top s_{i,0}\hat{\mathbf{s}}_i, \text{solve for } a(x,y) \text{ and } \hat{\mathbf{n}}(x,y)$$

$$e_i(x,y) = \mathbf{b}(x,y)^\top \mathbf{s}_i, \text{where } \mathbf{b}(x,y) = a(x,y)\hat{\mathbf{n}}(x,y) \text{ and } \mathbf{s}_i = s_{i,0}\hat{\mathbf{s}}_i$$

$$e_i(x,y) = \mathbf{s}_i^\top \mathbf{b}(x,y)$$

$$\begin{bmatrix} e_1(x,y) \\ e_2(x,y) \\ \vdots \\ e_n(x,y) \end{bmatrix} = \begin{bmatrix} \mathbf{s}_1^\top \\ \mathbf{s}_2^\top \\ \vdots \\ \mathbf{s}_n^\top \end{bmatrix} \mathbf{b}(x,y)$$

$n \geq 3$ light sources and associated images

$$\mathbf{e}(x,y) = \mathtt{S}\mathbf{b}(x,y), \text{solve for } \mathbf{b}(x,y)$$

$$\mathtt{S}^\top \mathbf{e}(x,y) = \mathtt{S}^\top \mathtt{S}\mathbf{b}(x,y)$$

$$(\mathtt{S}^\top \mathtt{S})^{-1}\mathtt{S}^\top \mathbf{e}(x,y) = \mathbf{b}(x,y)$$

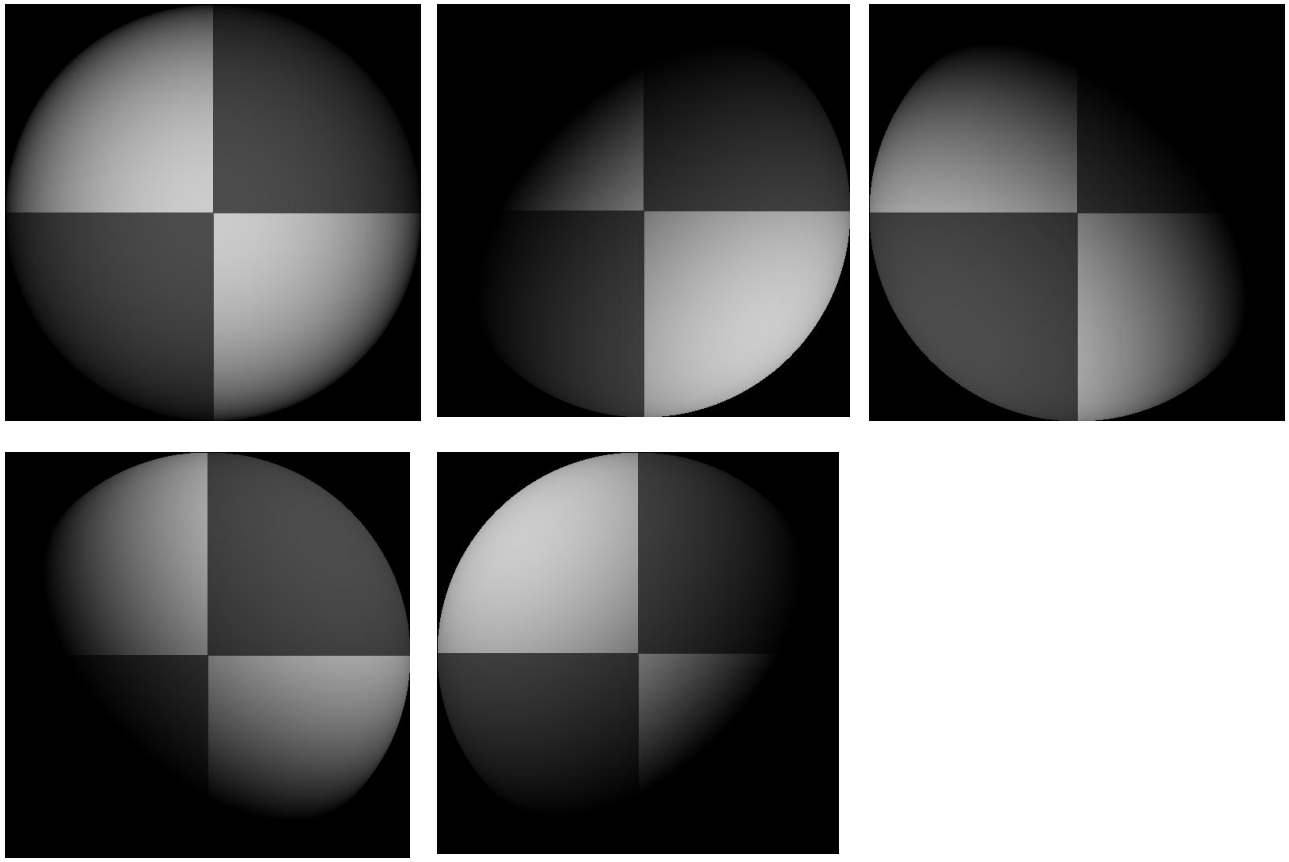
$$\mathtt{S}^+ \mathbf{e}(x,y) = \mathbf{b}(x,y), \text{where } \mathtt{S}^+ = (\mathtt{S}^\top \mathtt{S})^{-1}\mathtt{S}^\top$$

Same for all pixel coordinates

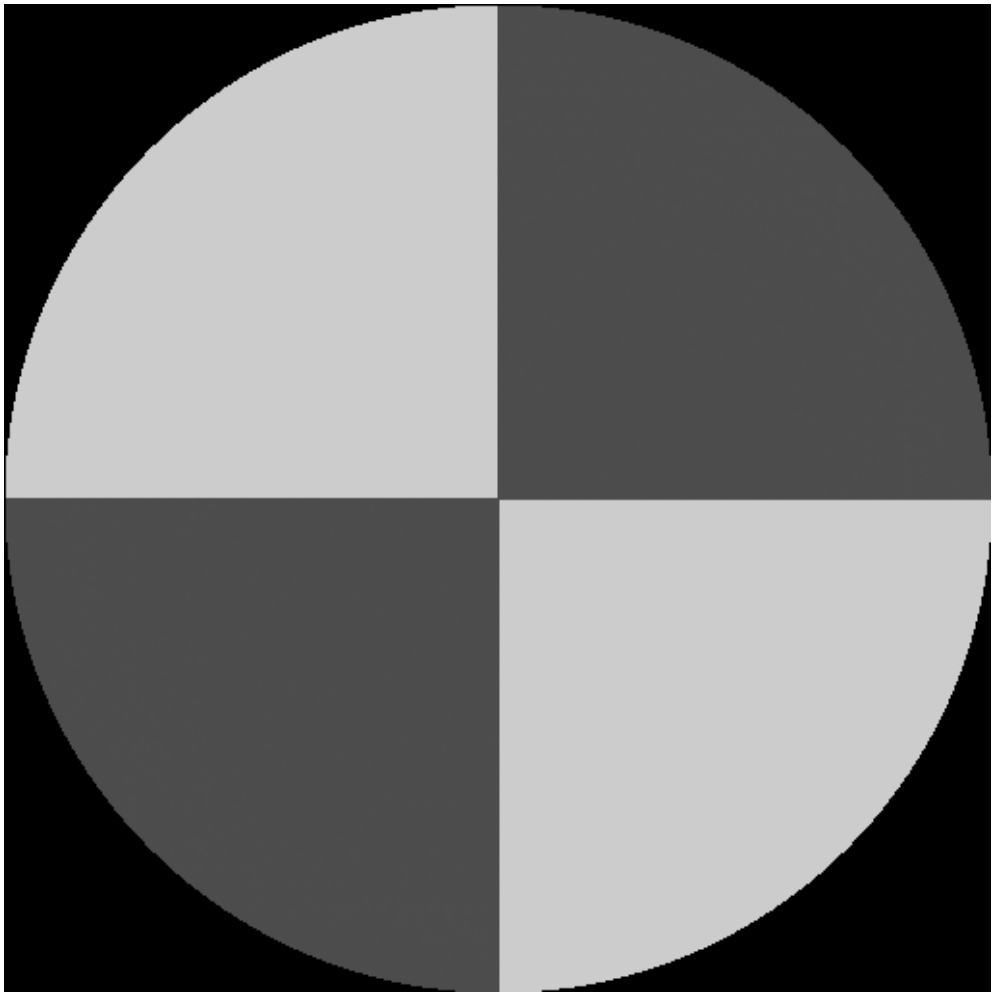
$$\mathbf{b}(x,y) = a(x,y)\hat{\mathbf{n}}(x,y), \text{where } a(x,y) = \|\mathbf{b}(x,y)\| \text{ and } \hat{\mathbf{n}}(x,y) = \frac{\mathbf{b}(x,y)}{\|\mathbf{b}(x,y)\|}$$

albedo  + surface normals

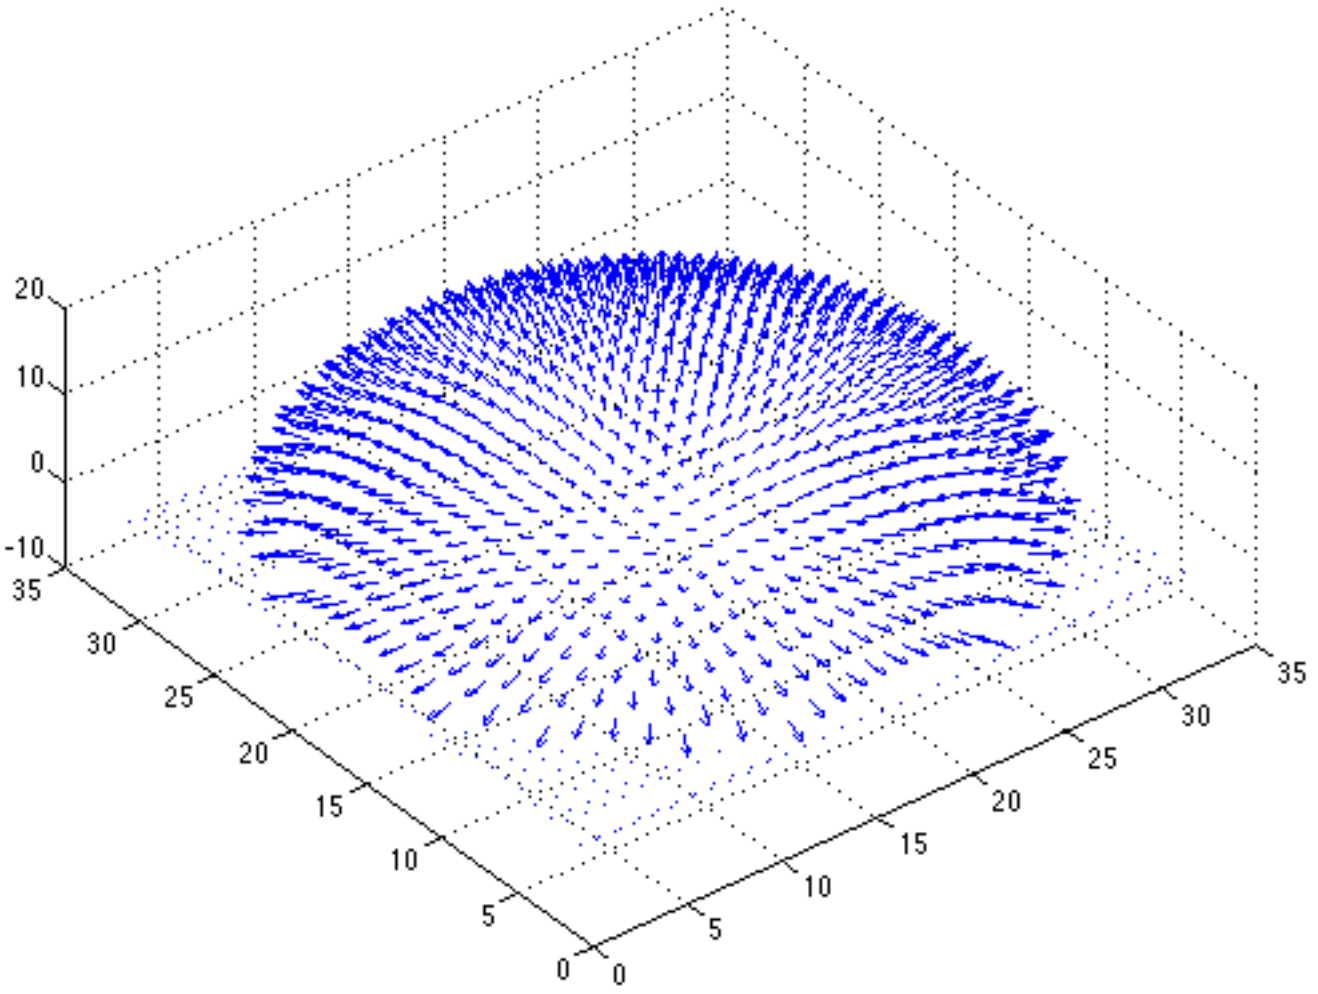
Input Images



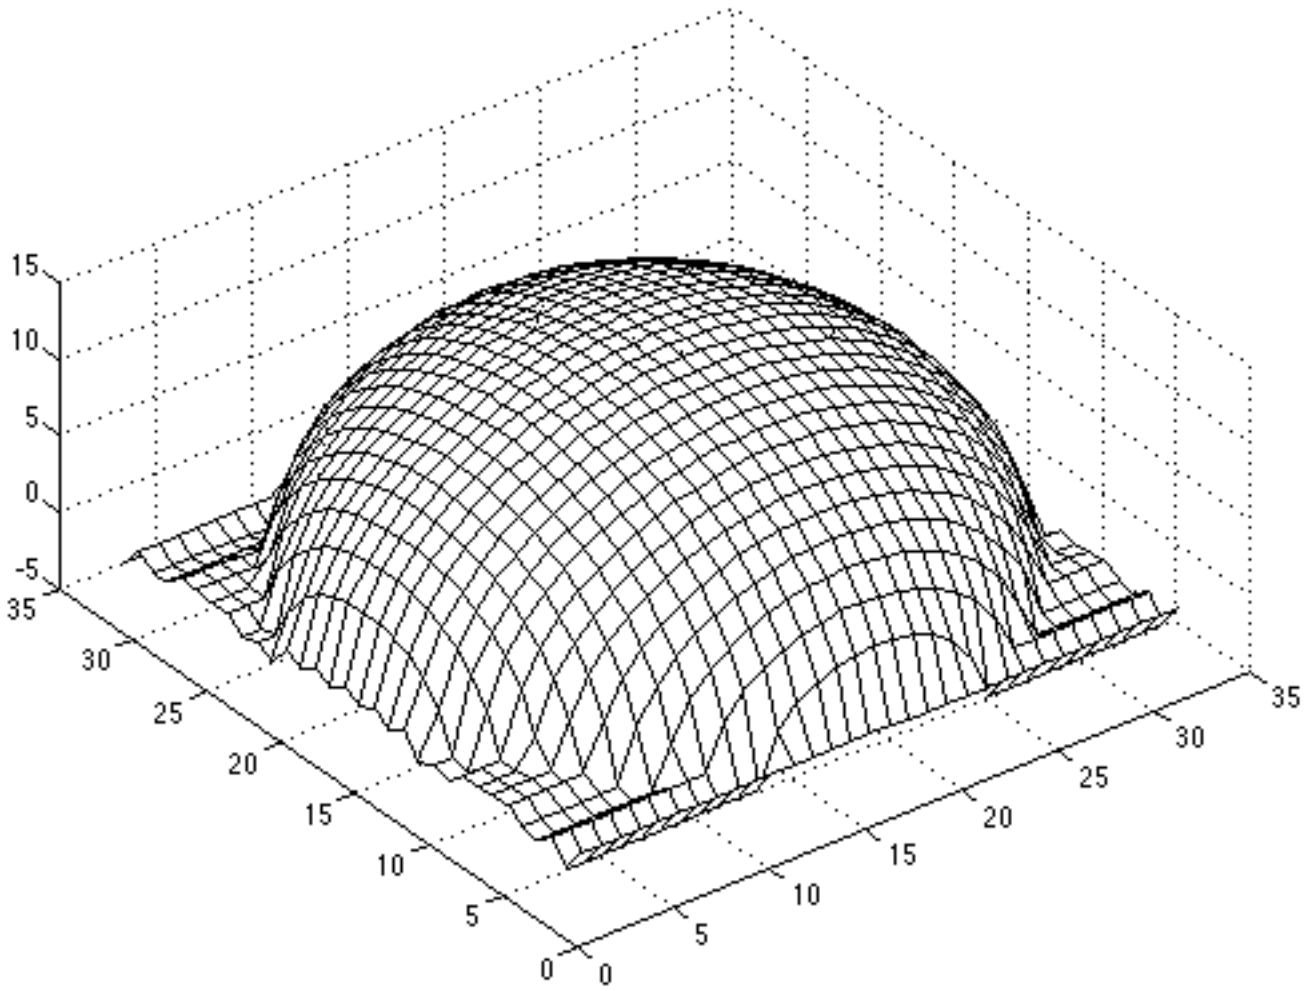
Recovered albedo



Recovered normal field

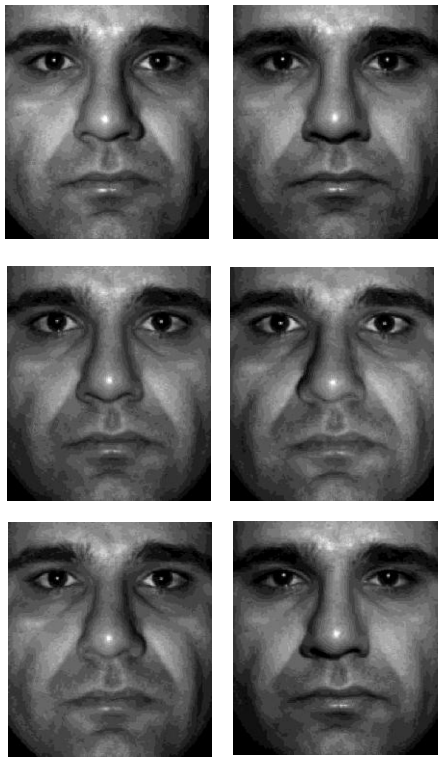


Surface recovered by integration

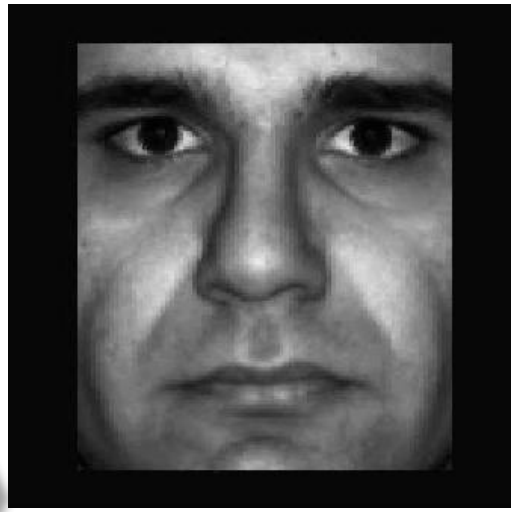


Use a mask to only include object

An example of photometric stereo



Images with known associated light sources



Albedo



Surface (from normals)



Surface (albedo texture map)

Use a mask to only include object

Next Lecture

- Image filtering