Photometric Stereo, Shape from Shading SfS Chapter 12.1.1. Szelisky

Guido Gerig

CS 6320, Spring 2012

Credits: M. Pollefey UNC CS256, Ohad Ben-Shahar CS BGU, Wolff JUN (http://www.cs.jhu.edu/~wolff/course600.461/week9.3/index.htm)

Photometric Stereo

Depth from Shading?

First step: Surface Normals from Shading

Second step: Re-integration of surface from Normals

Examples

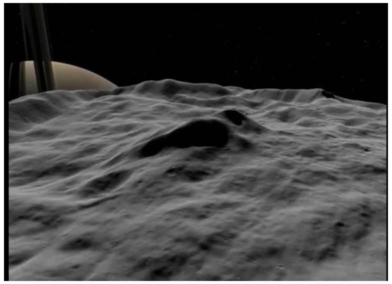

http://www.youtube.com/watch?v=sfCQ7f7PMbc&feature=related

Simulated voyage over the surface of Neptune's large moon Triton

http://www.youtube.com/watch?v=nwzVrC2GQXE

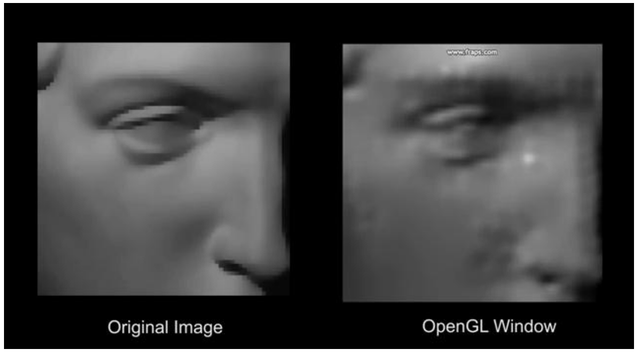

http://www.youtube.com/watch?v=KiTA6ftyQuY

Shape from Shading

Inverting the image formation process

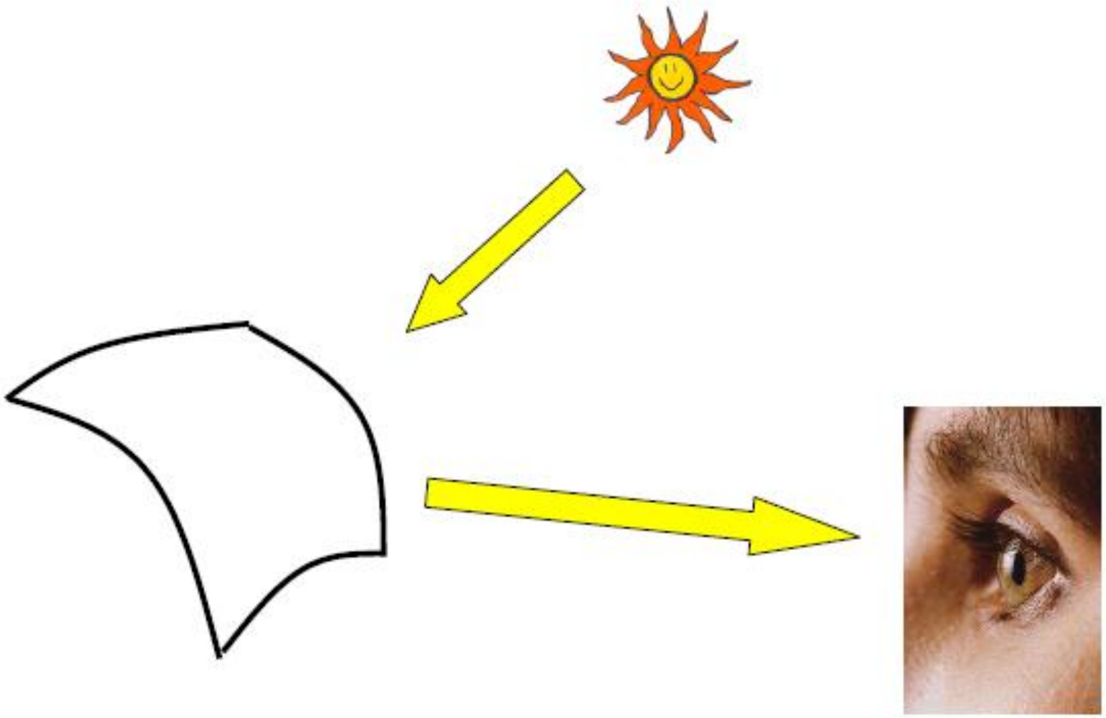

Image formation = “Shading from shape” (and light sources)

Credit: Ohad Ben-Shahar CS BGU

Shape from Shading

Authors: Emmanuel Prados and Olivier Faugeras

CVPR'2005, International Conference on Computer Vision and Pattern Recognition, San Diego, CA, USA, June 2005.

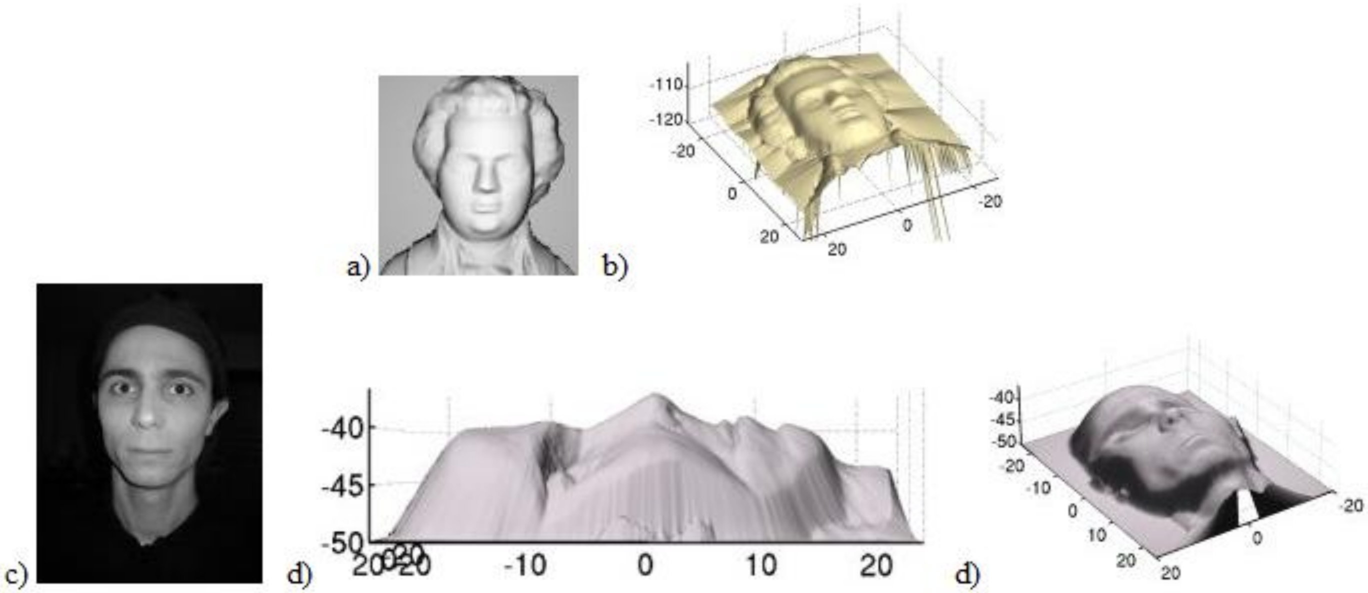

a) Synthetic image generated from the classical Mozart's face [Zhang-Tsai-etal:99]; b) reconstructed surface from a) by new algorithm; c) real image of a face; d)-e) reconstructed surface from c) by new algorithm.

Photometric Stereo

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

Setting for Photometric Stereo

Multiple images with different lighting (vs binocular/geometric stereo)

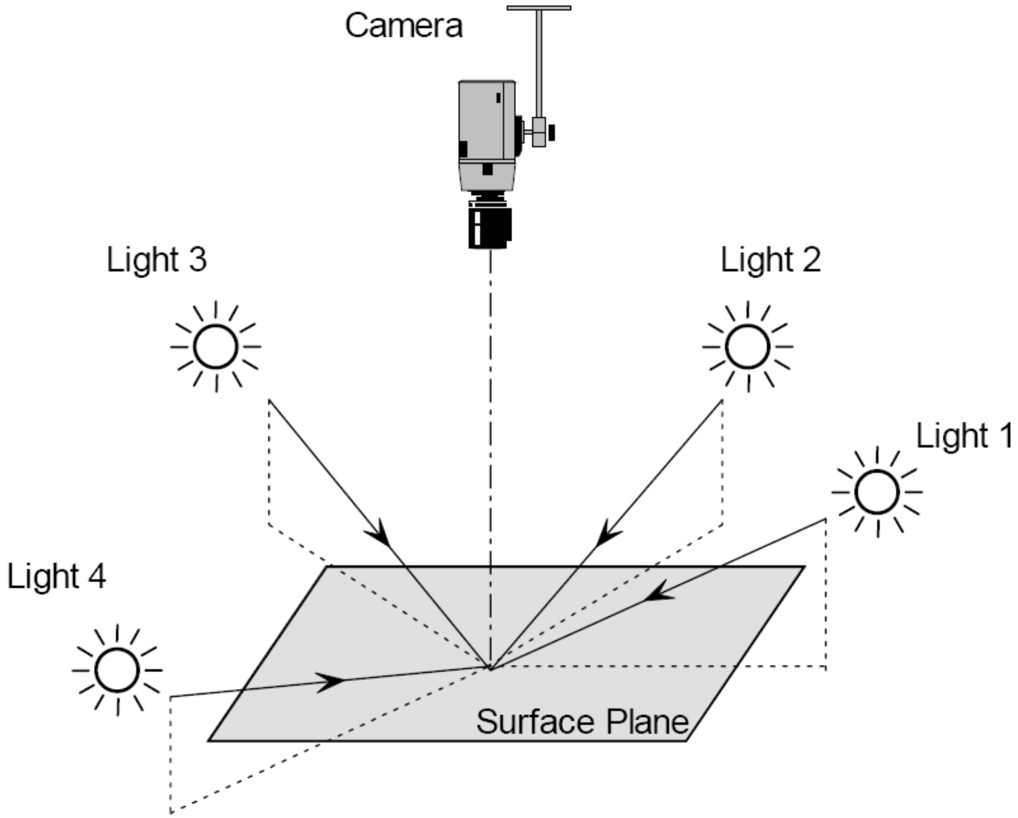

Goal: 3D from One View and multiple Source positions

Input images

Usable Data Mask

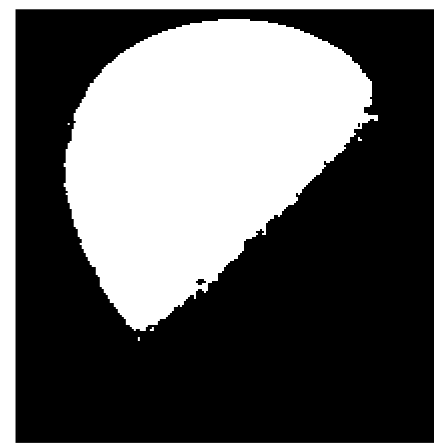

Scene Results

Needle Diagram:
Surface Normals

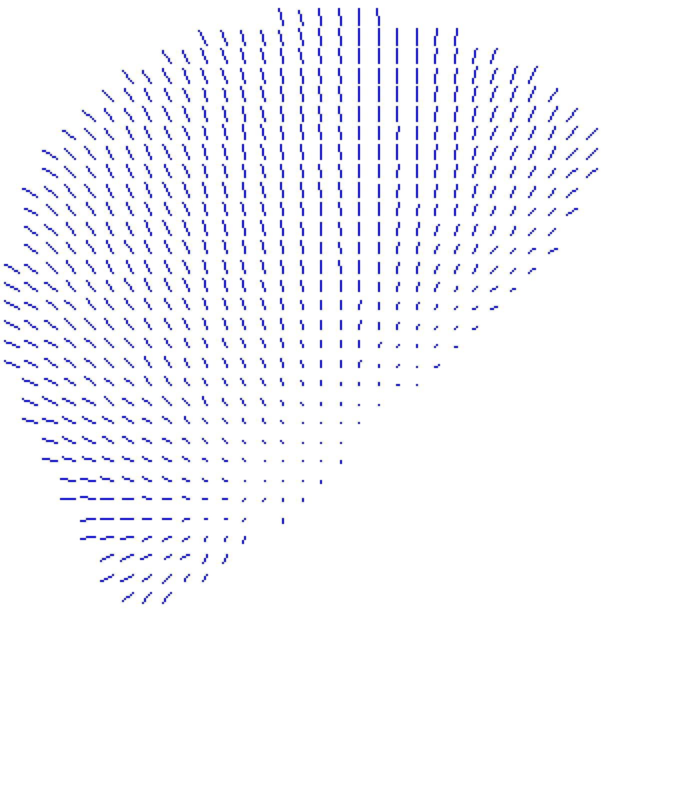

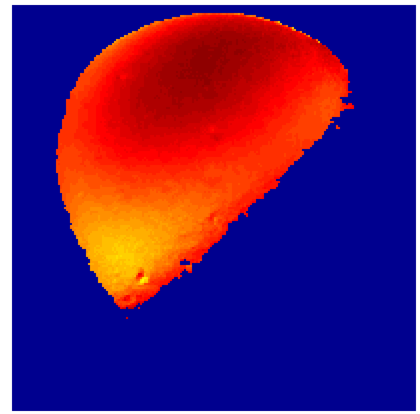

Albedo

Re-lit:

Projection model for surface recovery - usually called a Monge patch

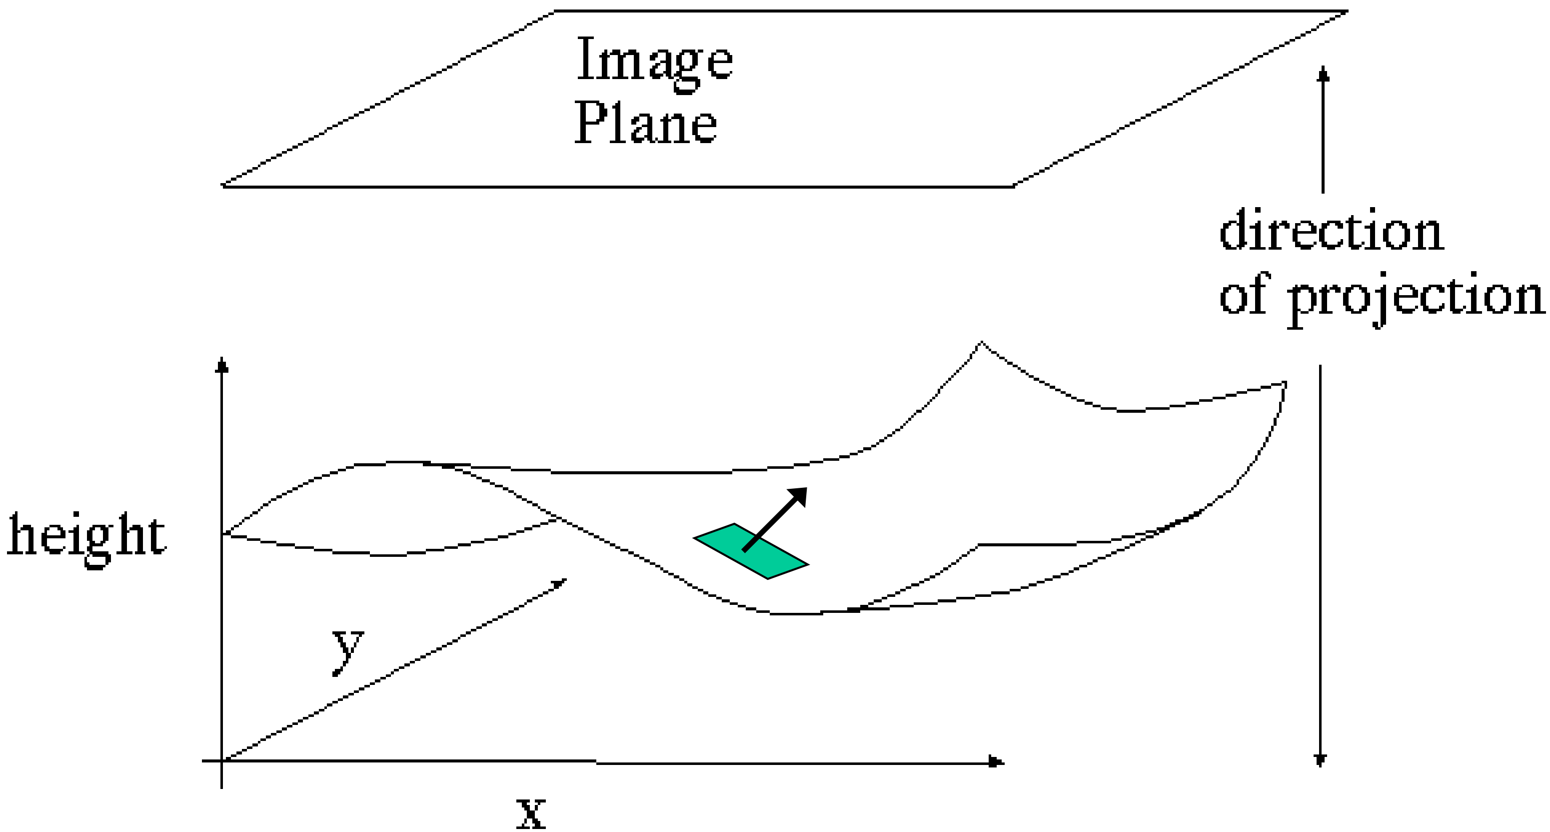

Lambertian Reflectance Map

LAMBERTIAN MODEL

$$E = \rho \langle n, n_s \rangle = \rho \, COS \, \theta$$

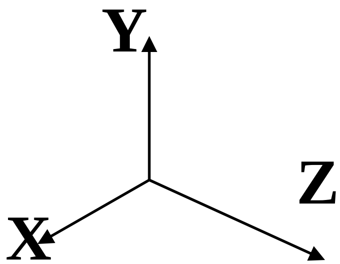

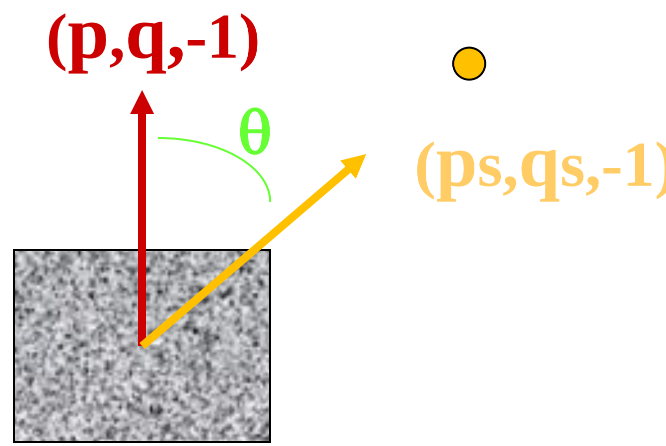

$$COS\theta = \frac{1 + pp_L + qq_L}{\sqrt{1+p^2+q^2}\sqrt{1+p_L{}^2+q_L{}^2}}$$

REFLECTANCE MAP IS A VIEWER-CENTERED REPRESENTATION OF REFLECTANCE

$(f_x, f_y, -1) = (0,1,f_x) \times (1,0,f_y)$

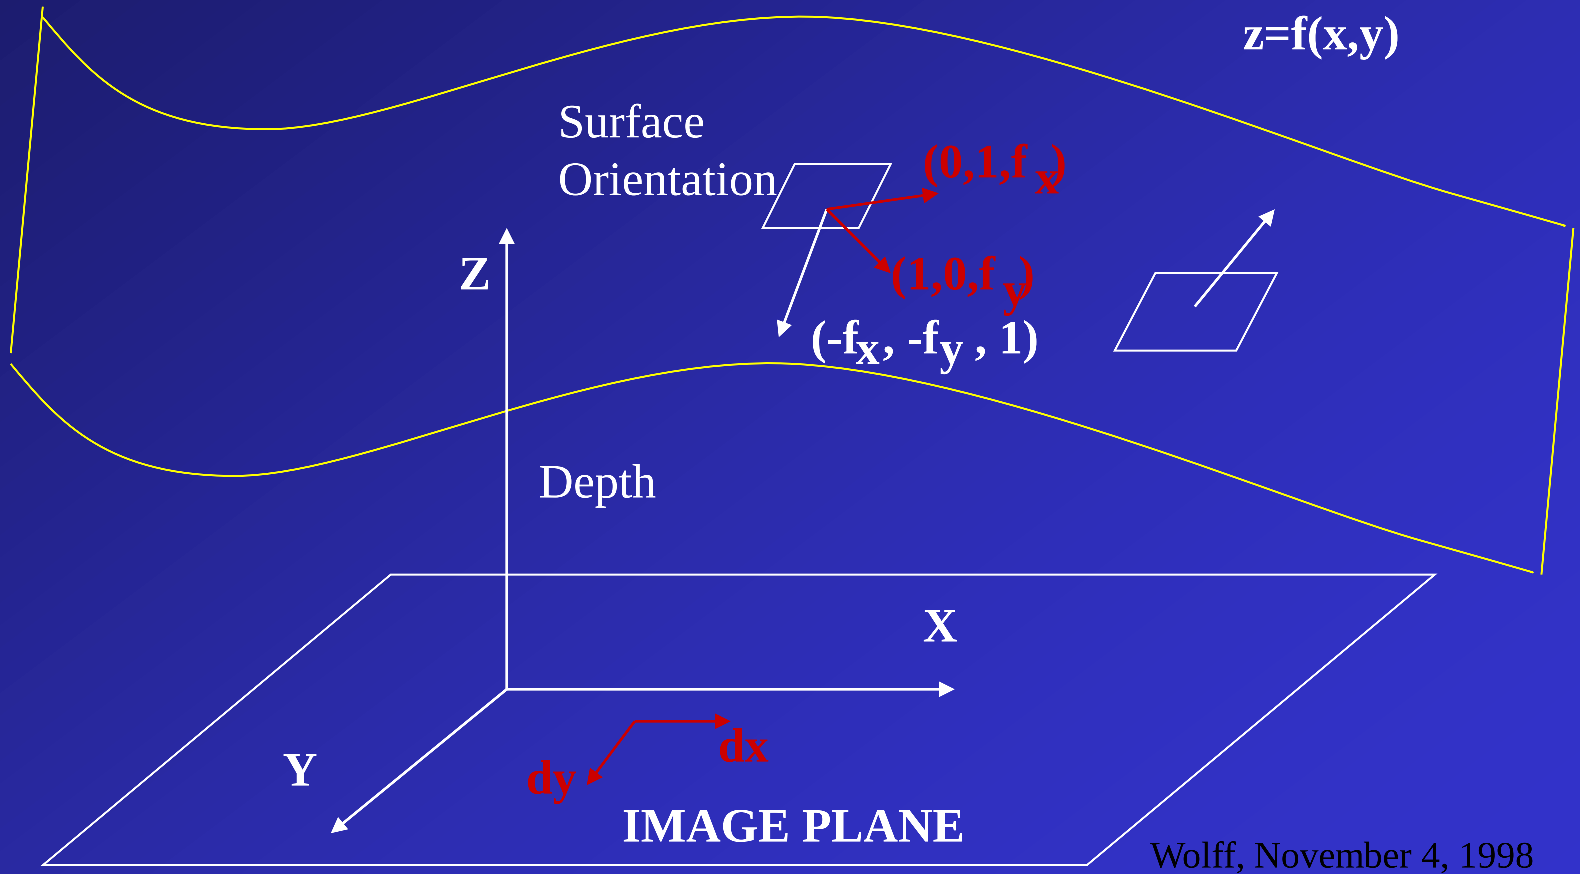

REFLECTANCE MAP IS A VIEWER-CENTERED REPRESENTATION OF REFLECTANCE

$$(-f_x, -f_y, 1) = (-p, -q, 1)$$

p, q comprise a gradient or gradient space representation for local surface orientation.

Reflectance map expresses the reflectance of a material directly in terms of viewer-centered representation of local surface orientation.

Reflectance Map (ps=0, qs=0)

The Reflectance Map – Lambertian surface from overhead source position

$$R(p,q) = \frac{1}{\sqrt{p^2 + q^2 + 1}}$$

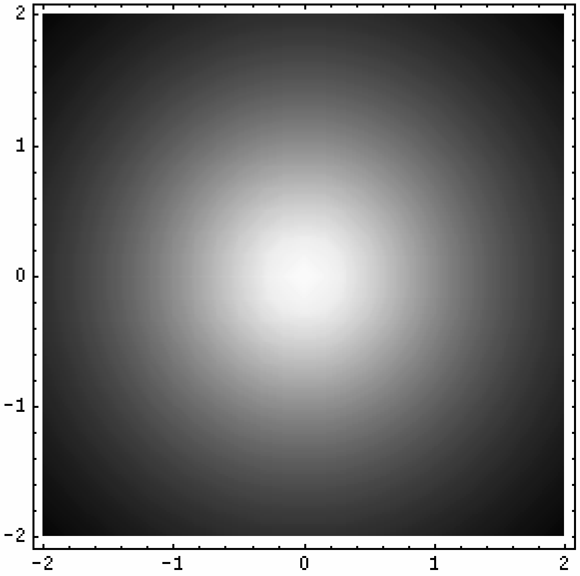

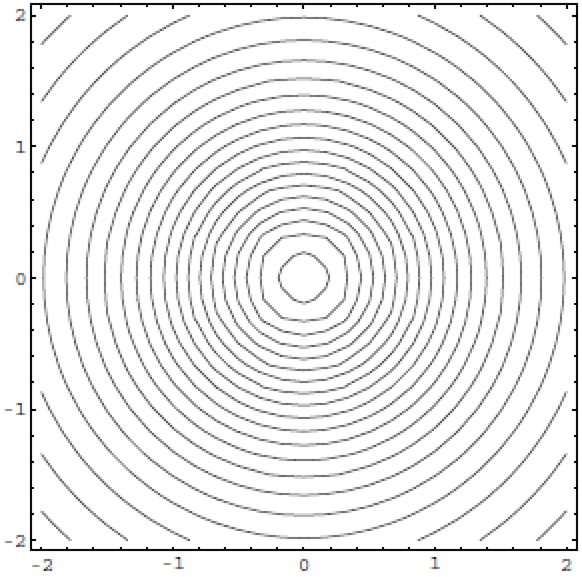

Reflectance Map

Shading on Lambertian surface – Overhead point source

$$I(x,y) = \rho(\hat{N}\cdot[0,0,1]) = \rho\frac{1}{\sqrt{p^2+q^2+1}} = R(p,q)$$

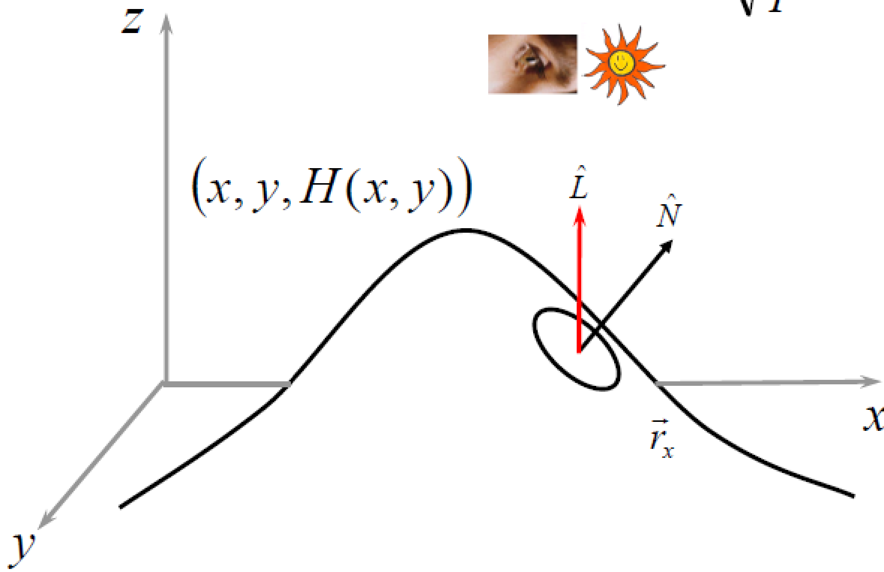

Credit: Ohad Ben-Shahar CS BGU

Reflectance Map

Shape from Shading

Shading on Lambertian surface – General point source

$$I=\rho(\hat{N}\cdot\hat{L})=\rho\frac{-p\cdot L_x-q\cdot L_y+L_z}{\sqrt{p^2+q^2+1}\sqrt{L_x^{\,2}+L_y^{\,2}+L_z^{\,2}}}=\rho\frac{p\cdot p_L+q\cdot q_L+1}{\sqrt{p^2+q^2+1}\sqrt{p_L^{\,2}+q_L^{\,2}+1}}$$

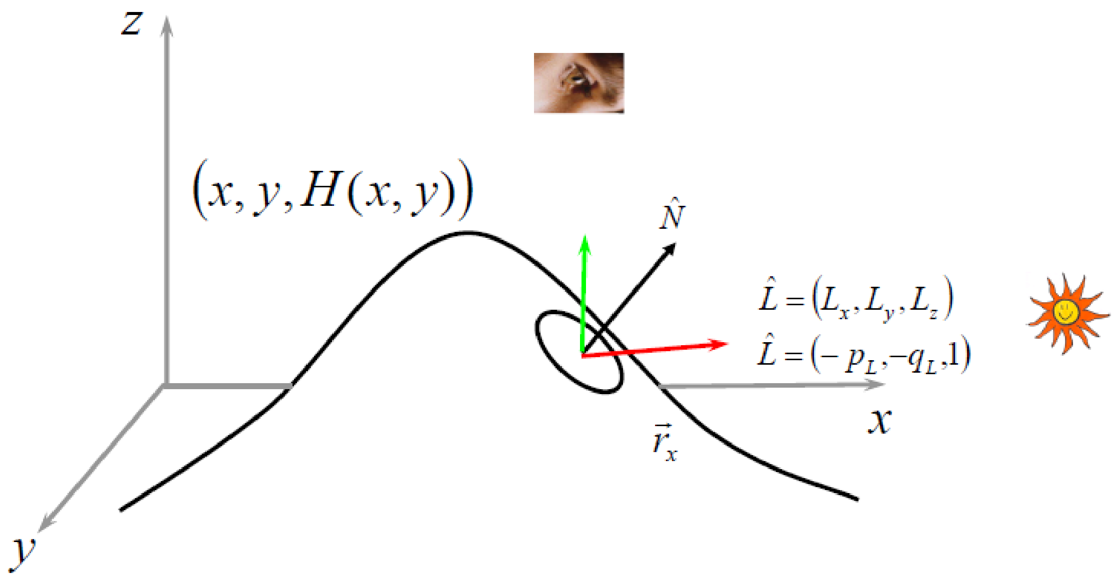

Credit: Ohad Ben-Shahar CS BGU

Reflectance Map

The Reflectance Map – Lambertian surface from general source position

$$R(p,q) = \frac{p \cdot p_L + q \cdot q_L + 1}{\sqrt{p^2 + q^2 + 1}\sqrt{p_L^{\,2} + q_L^{\,2} + 1}}$$

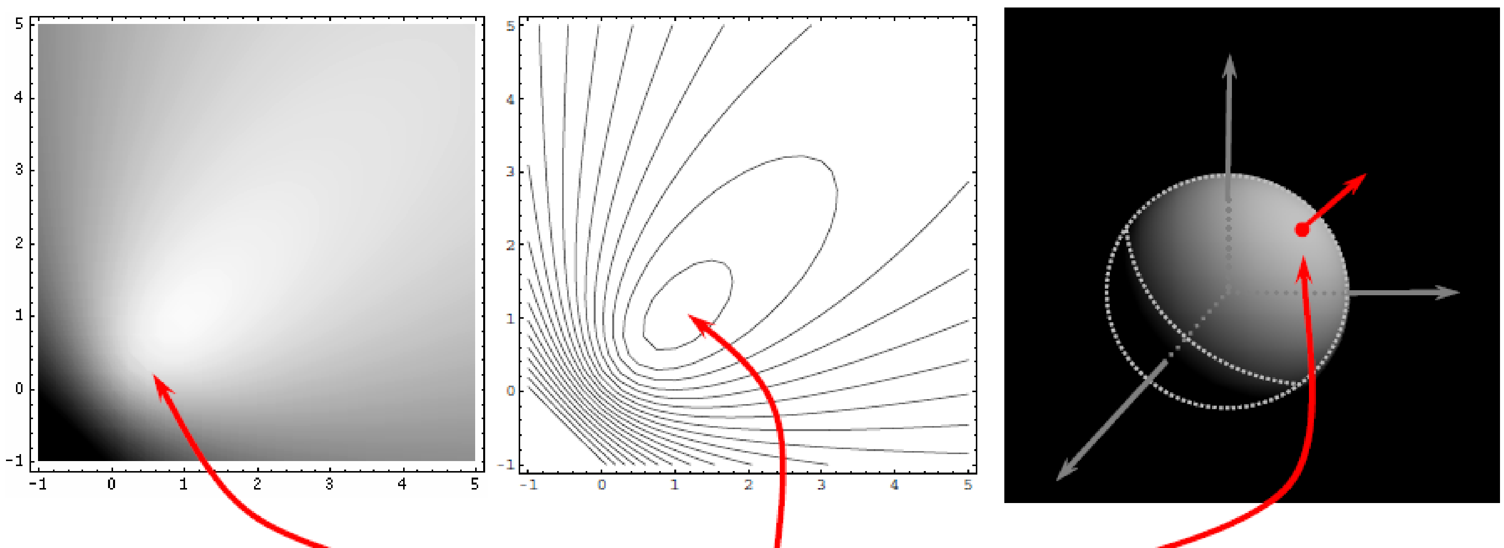

Gradient point of maximum brightness

Reflectance Map (General)

$$p = \frac{\partial z}{\partial x}$$

$$q = \frac{\partial z}{\partial y}$$

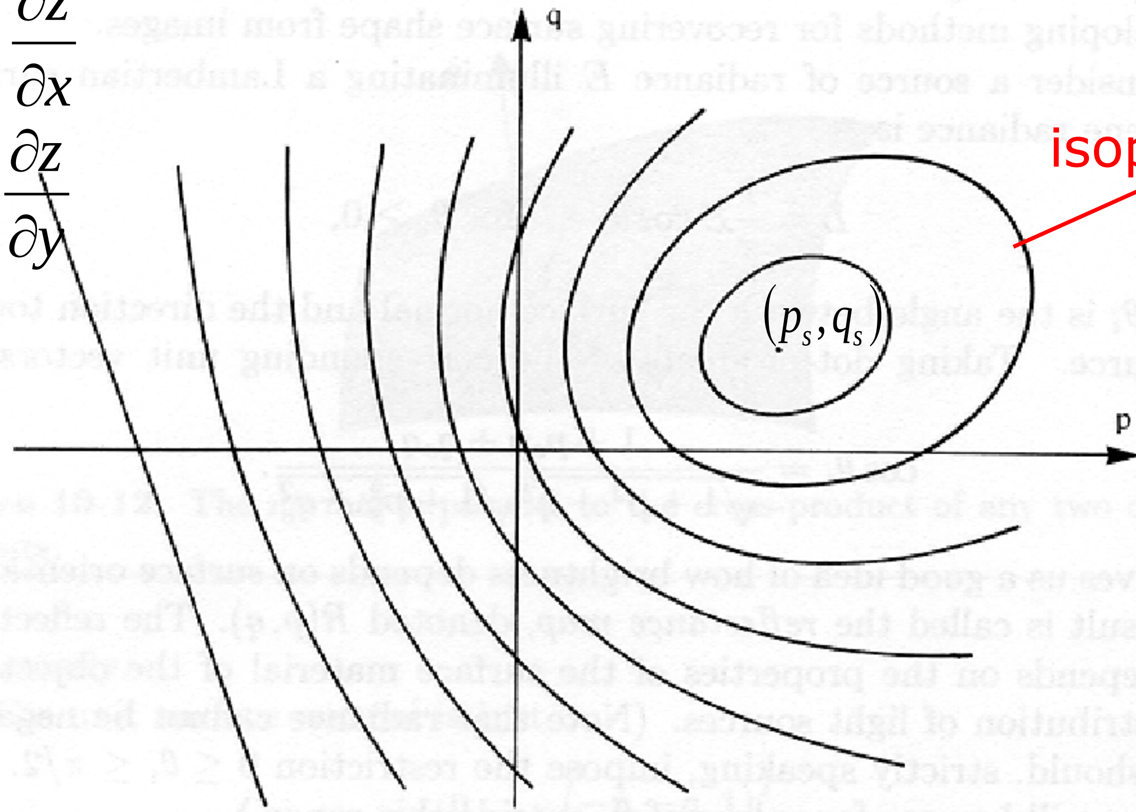

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p, q)$ occurs at the point $(p, q) = (p_s, q_s)$, found inside the nested conic sections, while $R(p, q) = 0$ all along the line on the left side of the contour map.

Reflectance Map

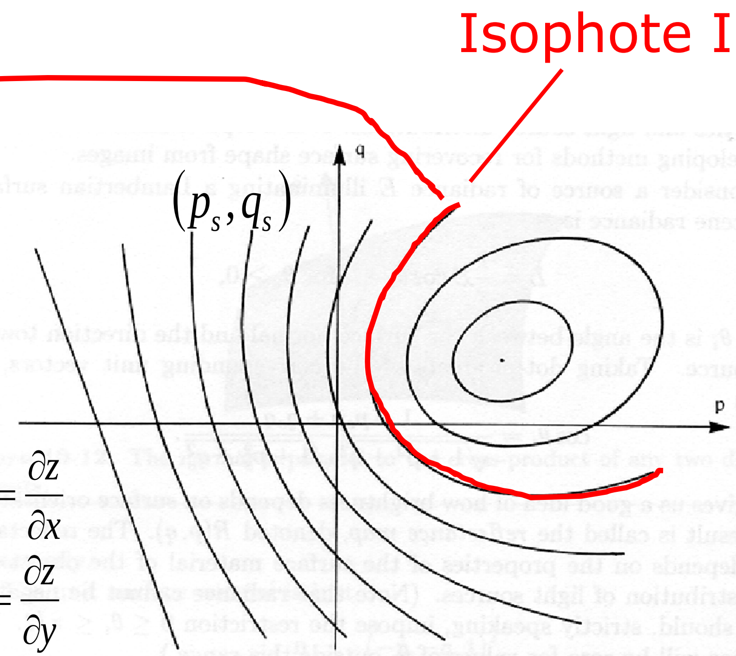

Given Intensity I in image, there are multiple (p,q) combinations (= surface orientations).

$$p = \frac{\partial z}{\partial x}$$

$$q = \frac{\partial z}{\partial y}$$

⇨ Use multiple images with different light source directions.

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p,q)$ occurs at the point $(p,q) = (p_s, q_s)$, found inside the nested conic sections, while $R(p,q) = 0$ all along the line on the left side of the contour map.

Multiple Images = Multiple Maps

Can isolate p, q as contour intersection

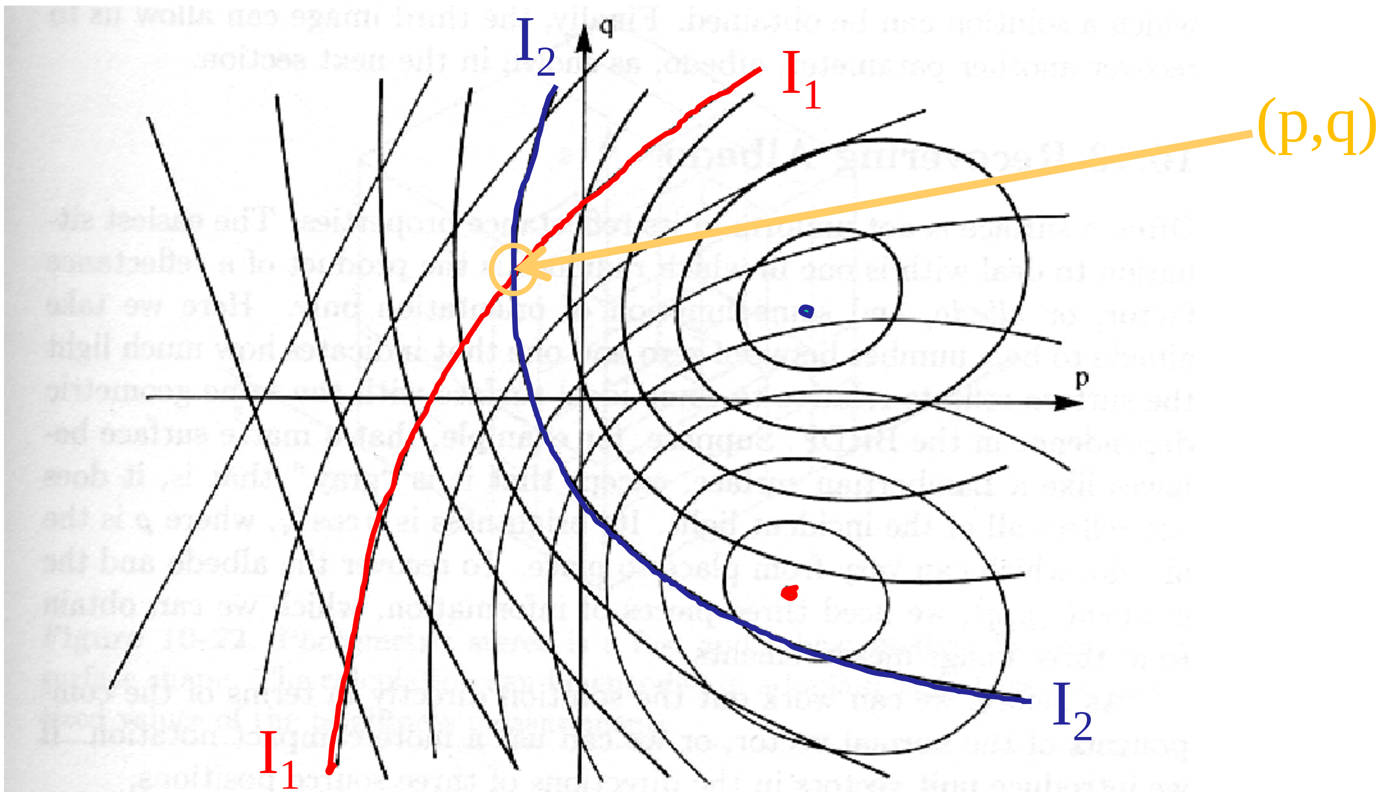

Figure 10-21. In the case of a Lambertian surface illuminated successively by two different point sources, there are at most two surface orientations that produce a particular pair of brightness values. These are found at the intersection of the corresponding contours in two superimposed reflectance maps.

Example: Two Views

$$I_1(x, y) = R_1(p, q)$$
$$I_2(x, y) = R_2(p, q)$$

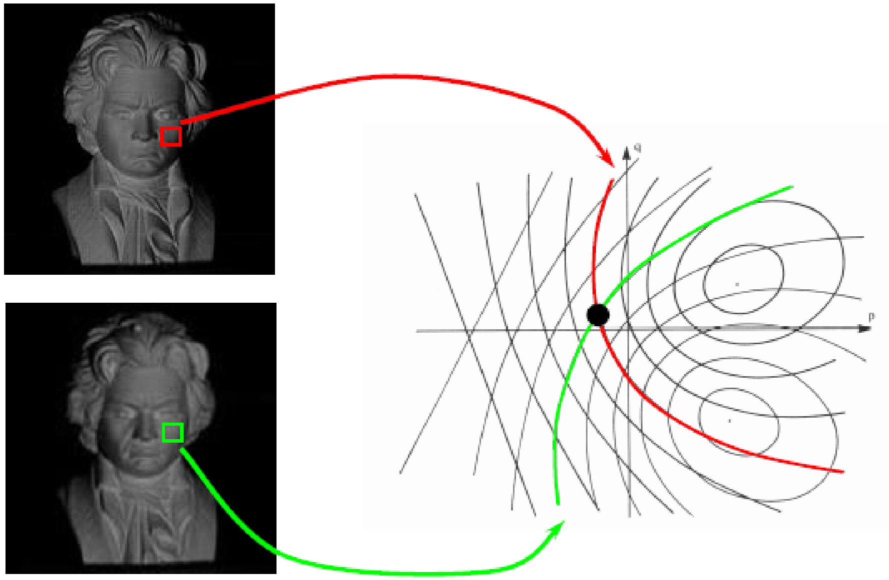

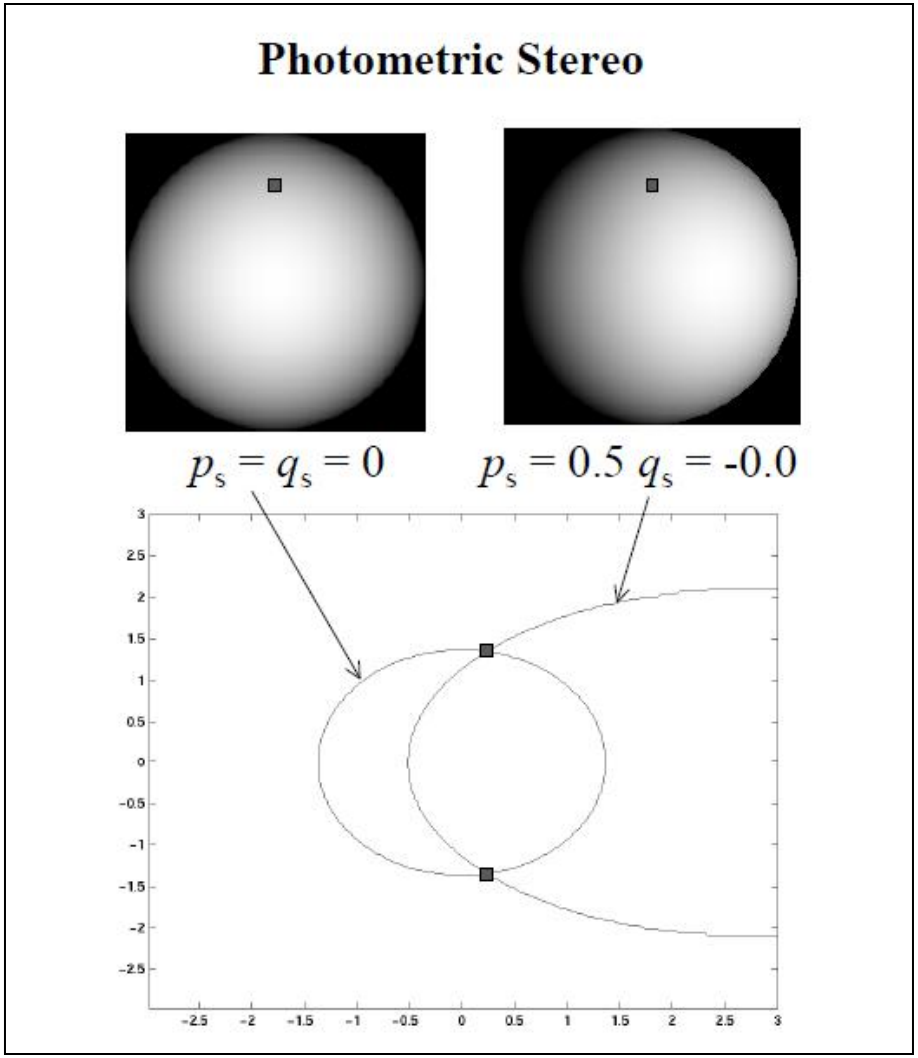

Still not unique for certain intensity pairs.

Constant Albedo

Photometric Stereo

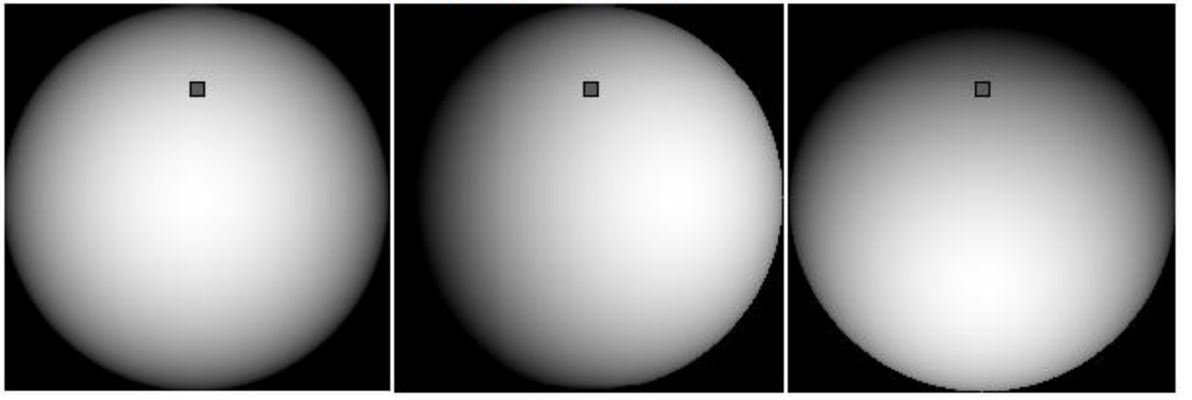

$I_1 = \rho \mathbf{S}_1 . \mathbf{N}$ $\quad$ $I_2 = \rho \mathbf{S}_2 . \mathbf{N}$ $\quad$ $I_3 = \rho \mathbf{S}_3 . \mathbf{N}$

$$\downarrow$$

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = \begin{bmatrix} \mathbf{S}_1^T \\ \mathbf{S}_2^T \\ \mathbf{S}_3^T \end{bmatrix} \rho \mathbf{N}$$

$$\downarrow$$

$$\rho \mathbf{N} = \mathbf{S}^{-1} \mathbf{I}$$

Solve linear equation system to calculate $\bar{N}$.

Varying Albedo

Solution Forsyth & Ponce:

For each point source, we know the source vector (by assumption). We assume we know the scaling constant of the linear camera (k). Fold the normal (**N**) and the reflectance ($\boldsymbol{\rho}$(x,y)) into one vector **g**, and the scaling constant and source vector into another $\mathbf{V_j}$.

- Out of shadow:

$$
\begin{aligned}
I(x,y) &= kB(\boldsymbol{x}) \\
&= kB(x,y) \\
&= k\rho(x,y)\boldsymbol{N}(x,y)\cdot \boldsymbol{S}_1 \\
&= \boldsymbol{g}(x,y)\cdot \boldsymbol{V}_1
\end{aligned}
$$

- In shadow:

$$I(x,y) = 0$$

where $\boldsymbol{g}(x,y) = \rho(x,y)\boldsymbol{N}(x,y)$ and $\boldsymbol{V}_1 = k\boldsymbol{S}_1$, where k is the constant connecting the camera response to the input radiance.

Multiple Images: Linear Least Squares Approach

- Combine albedo and normal
- Separate lighting parameters
- More than 3 images => overdetermined system

$$\mathcal{V} = \begin{pmatrix} \boldsymbol{V}_1^T \\ \boldsymbol{V}_2^T \\ \cdots \\ \boldsymbol{V}_n^T \end{pmatrix} \qquad \boldsymbol{i}(x,y) = \{I_1(x,y), I_2(x,y), \ldots, I_n(x,y)\}^T$$

$$\boldsymbol{i}(x,y) = \mathcal{V}\boldsymbol{g}(x,y)$$

$\boldsymbol{g}$ is obtained by solving this linear system: $\bar{g}(\mathrm{x,y}) = V^{-1}\mathrm{i}(\mathrm{x,y})$

- How to calculate albedo ρ and $\bar{N}$?

$$\bar{g}(x,y) = \rho(x,y)\bar{N}(\mathrm{x,y})$$

$$\rightarrow \bar{N} = \frac{\bar{g}}{|\bar{g}|}, \quad \rho(x,y) = |\bar{g}|$$

Example LLS Input

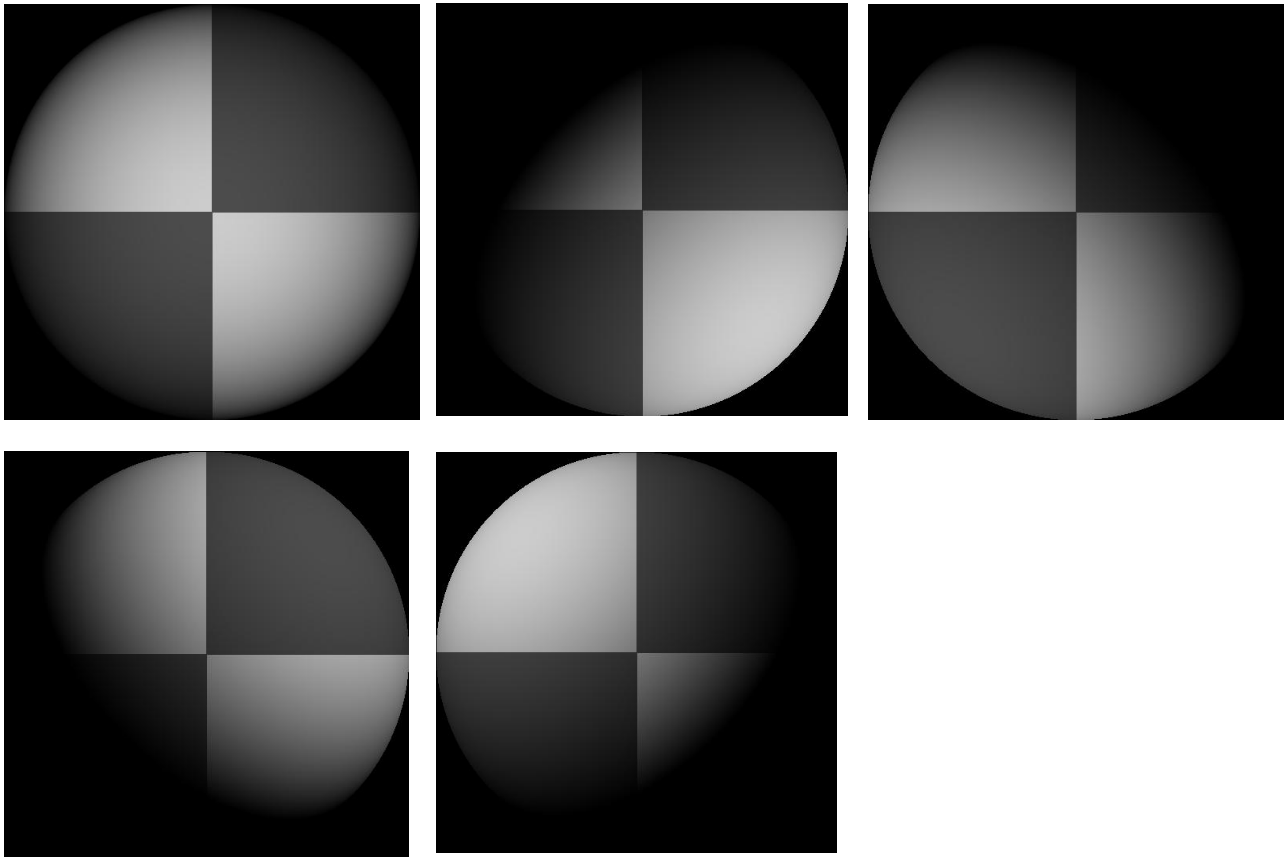

Problem: Some regions in some images are in the shadow (no image intensity).

Dealing with Shadows (Missing Info)

- For each point source, we know the source vector (by assumption). We assume we know the scaling constant of the linear camera. Fold the normal and the reflectance into one vector g, and the scaling constant and source vector into another Vj

- Out of shadow:

$$\begin{aligned} I_j(x,y) &= kB(x,y) \\ &= k\rho(x,y)\left(\mathbf{N}(x,y)\bullet\mathbf{S}_j\right) \\ &= \mathbf{g}(x,y)\bullet\mathbf{V}_j \end{aligned}$$

- In shadow:

$$\mathrm{I_j(x,y)} = 0$$

No partial shadow

Matrix Trick for Complete Shadows

- Matrix from Image Vector:

$$\mathcal{I}(x,y) = \begin{pmatrix} I_1(x,y) & \dots & 0 & 0 \\ 0 & I_2(x,y) & \dots & 0 \\ \dots & & & \\ 0 & 0 & \dots & I_n(x,y) \end{pmatrix}$$

- Multiply LHS and RHS with diag matrix

$$\mathcal{I}\boldsymbol{i} = \mathcal{I}\mathcal{V}\boldsymbol{g}(x,y)$$

$$\underbrace{\begin{pmatrix} I_1^2(x,y) \\ I_2^2(x,y) \\ .. \\ I_n^2(x,y) \end{pmatrix}}_{\text{Known}} = \underbrace{\begin{pmatrix} I_1(x,y) & 0 & .. & 0 \\ 0 & I_2(x,y) & .. & .. \\ .. & .. & .. & 0 \\ 0 & .. & 0 & I_n(x,y) \end{pmatrix}}_{\text{Known}} \underbrace{\begin{pmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \\ .. \\ \mathbf{V}_n^T \end{pmatrix}}_{\text{Known}} \underbrace{\mathbf{g}(x,y)}_{\text{Unknown}}$$

⇨ Relevant elements of the left vector and the matrix are zero at points that are in shadow.

Obtaining Normal and Albedo

- Given sufficient sources, we can solve the previous equation (most likely need a least squares solution) for **g**(x, y).
- Recall that **N**(x, y) is the unit normal.
- This means that $\rho(x,y)$ is the magnitude of **g**(x, y).
- This yields a check
  - If the magnitude of **g**(x, y) is greater than 1, there's a problem.
- And **N**(x, y) = **g**(x, y) / $\rho(x,y)$.

Example LLS Input

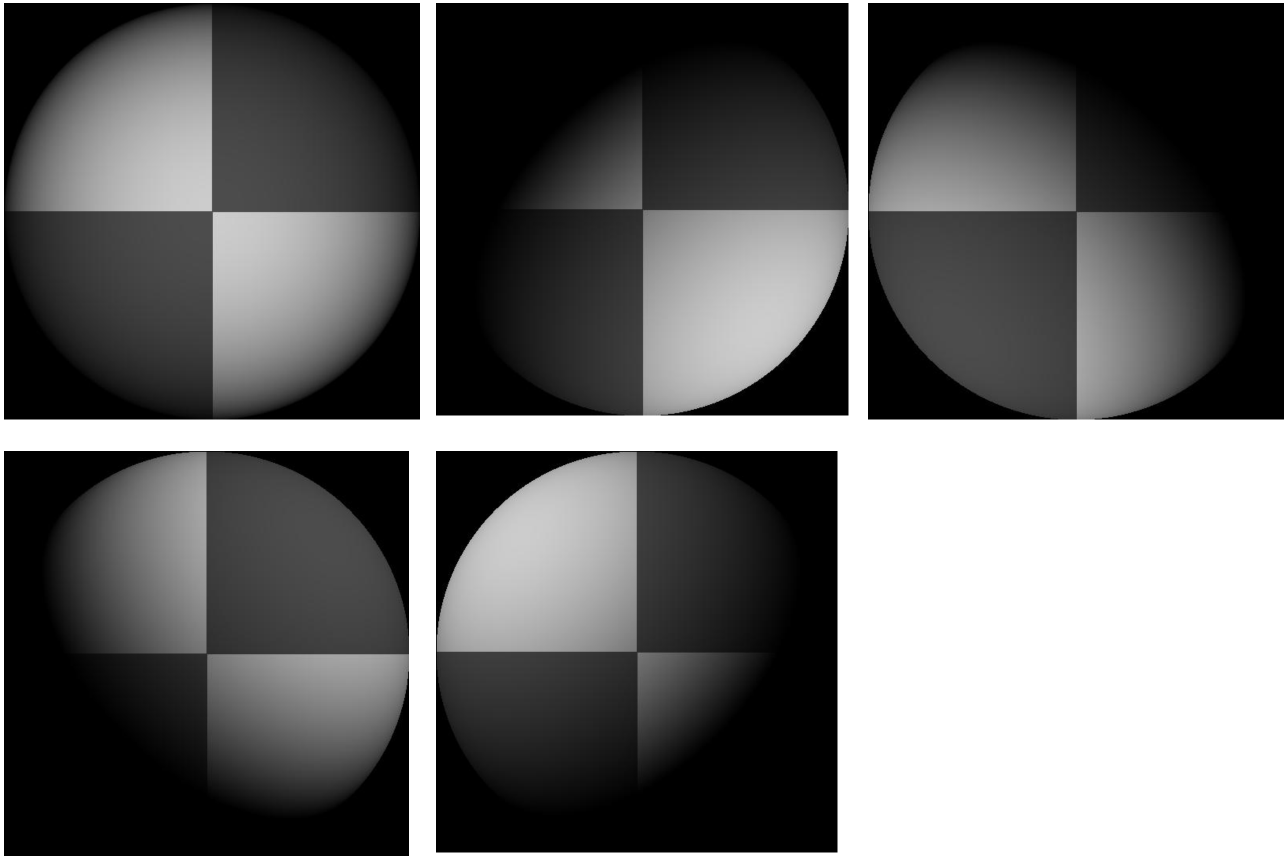

Example LLS Result

- Reflectance / albedo:

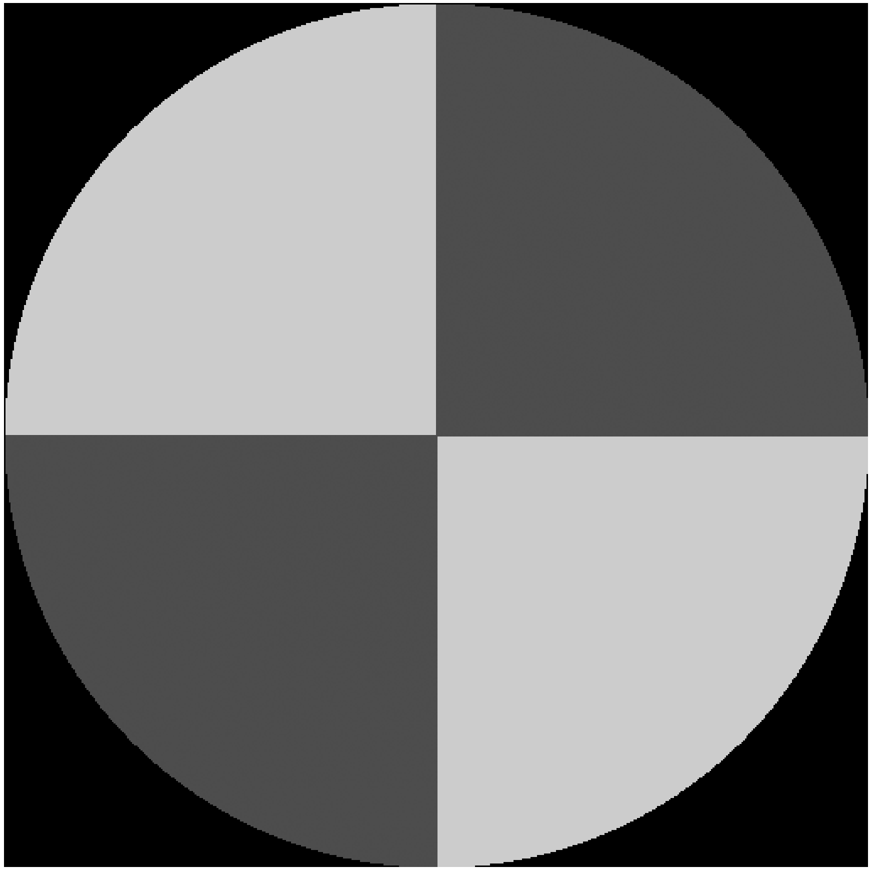

Recap

- Obtain normal / orientation, no depth

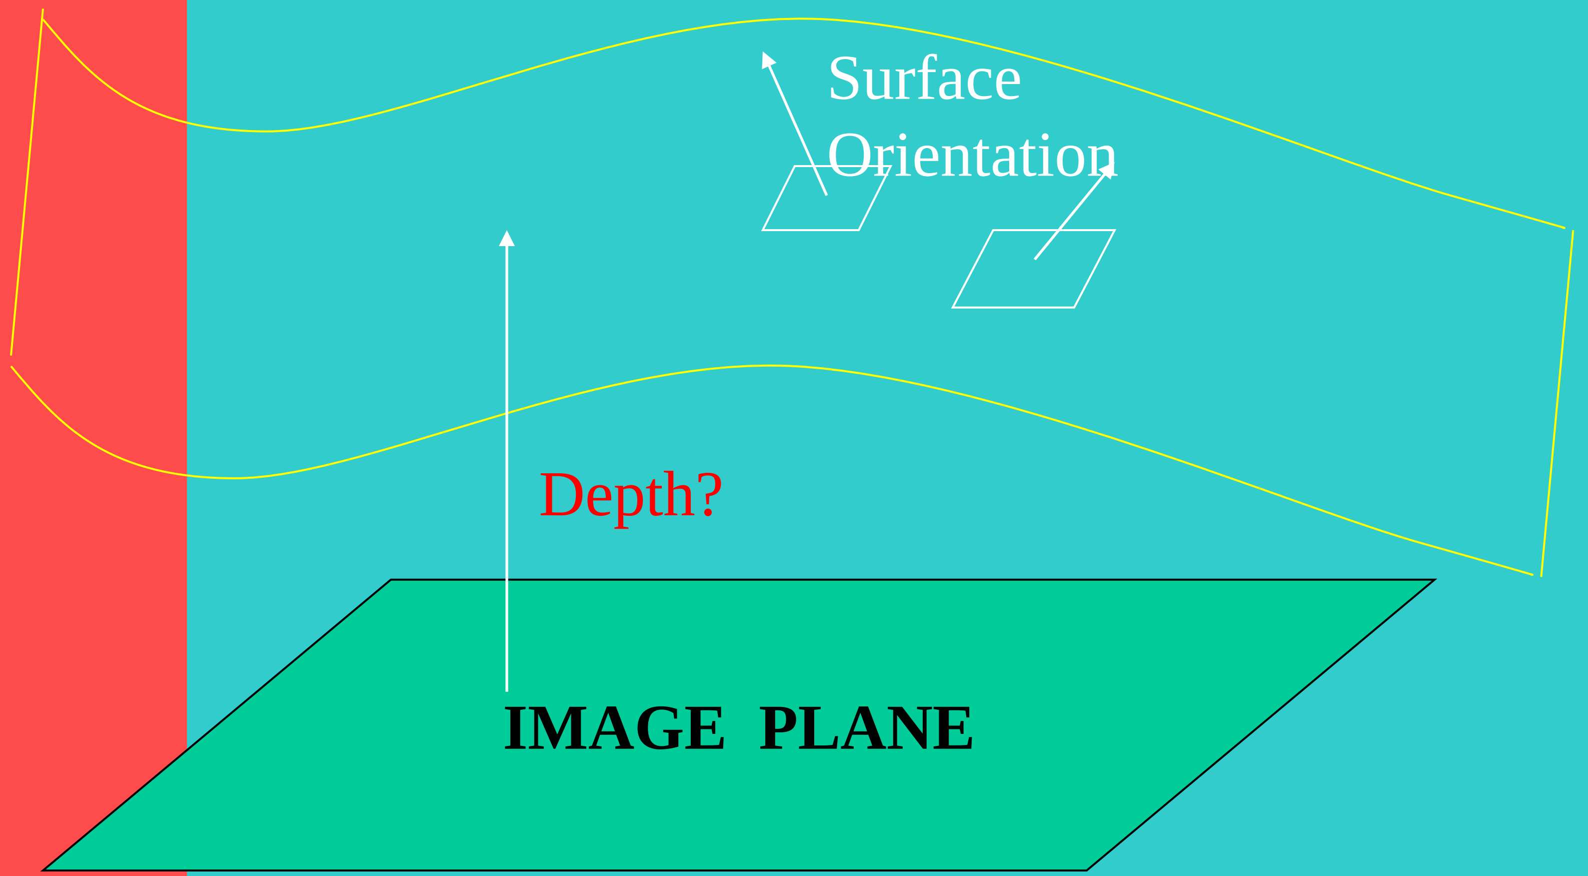

Goal

Shape as surface with depth and normal

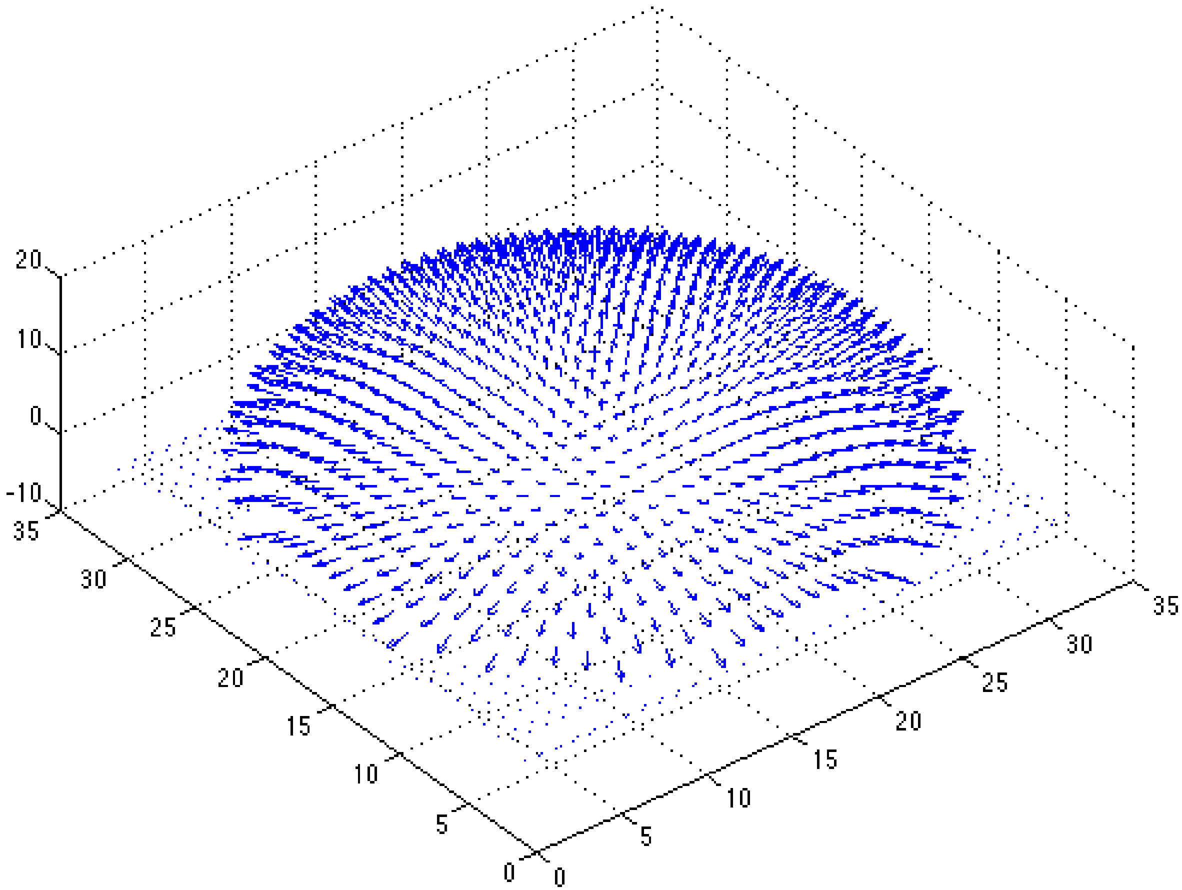

Recovering a surface from normals - 1

- Recall the surface is written as

$$(x, y, f(x, y))$$

- This means the normal has the form:

$$N(x,y) = \left(\frac{1}{\sqrt{f_x^2 + f_y^2 + 1}} \right) \begin{pmatrix} -f_x \\ -f_y \\ 1 \end{pmatrix}$$

- If we write the known vector **g** as

$$\mathbf{g}(x,y) = \begin{pmatrix} g_1(x,y) \\ g_2(x,y) \\ g_3(x,y) \end{pmatrix}$$

- Then we obtain values for the partial derivatives of the surface:

$$f_x(x,y) = \left(g_1(x,y) / g_3(x,y) \right)$$

$$f_y(x,y) = \left(g_2(x,y) / g_3(x,y) \right)$$

Recovering a surface from normals - 2

- Recall that mixed second partials are equal --- this gives us an integrability **check.** We must have:

$$\frac{\partial\left(g_1(x,y)/g_3(x,y)\right)}{\partial y} = \frac{\partial\left(g_2(x,y)/g_3(x,y)\right)}{\partial x}$$

- We can now recover the surface height at any point by integration along some path, e.g.

$$f(x,y) = \int_0^x f_x(s,y)ds + \int_0^y f_y(x,t)dt + c$$

Height Map from Integration

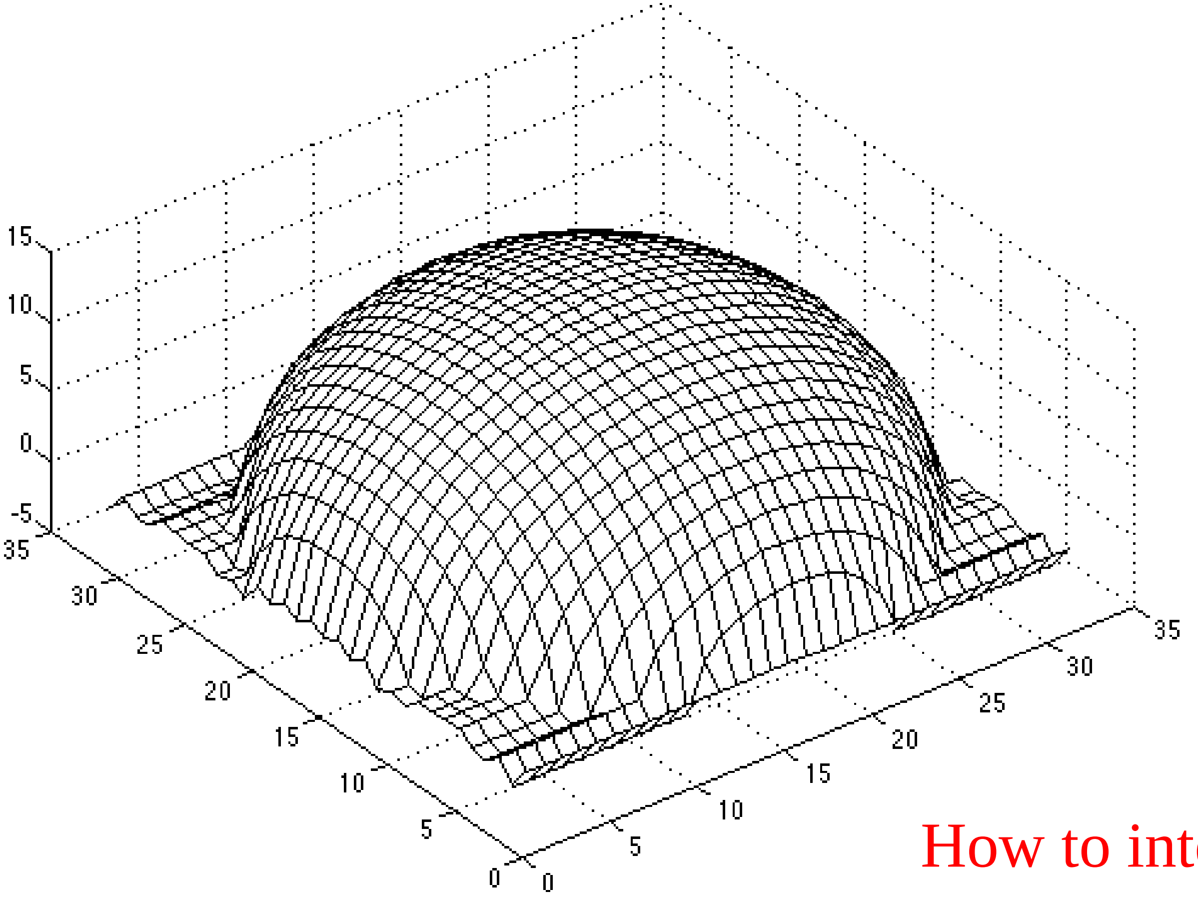

How to integrate?

Possible Solutions

- Engineering approach: Path integration (Forsyth & Ponce)
- In general: Calculus of Variation Approaches
- Horn: Characteristic Strip Method
- Kimmel, Siddiqi, Kimia, Bruckstein: Level set method
- Many others ….

Shape by Integation (Forsyth&Ponce)

- The partial derivative gives the change in surface height with a small step in either the x or the y direction
- We can get the surface by summing these changes in height along some path.

$$f(x,y) = \oint_C \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right) \cdot d\boldsymbol{l} + c$$

For example, we can reconstruct the surface at (u, v) by starting at $(0, 0)$, summing the y-derivative along the line $x = 0$ to the point $(0, v)$, and then summing the x-derivative along the line $y = v$ to the point (u, v)

$$f(u,v) = \int_0^v \frac{\partial f}{\partial y}(0,y)dy + \int_0^u \frac{\partial f}{\partial x}(x,v)dx + c$$

Simple Algorithm
Forsyth & Ponce

Obtain many images in a fixed view under different illuminants

Determine the matrix $\mathcal{V}$ from source and camera information

Create arrays for albedo, normal (3 components),
  p (measured value of $\frac{\partial f}{\partial x}$) and
  q (measured value of $\frac{\partial f}{\partial y}$)

For each point in the image array
  Stack image values into a vector $\boldsymbol{i}$
  Construct the diagonal matrix $\mathcal{I}$
  Solve $\mathcal{I}\mathcal{V}\boldsymbol{g} = \mathcal{I}\boldsymbol{i}$
    to obtain $\boldsymbol{g}$ for this point

  albedo at this point is $|\boldsymbol{g}|$
  normal at this point is $\frac{\boldsymbol{g}}{|\boldsymbol{g}|}$
  p at this point is $\frac{N_1}{N_3}$
  q at this point is $\frac{N_2}{N_3}$
end

Check: is $(\frac{\partial p}{\partial y} - \frac{\partial q}{\partial x})^2$ small everywhere?

top left corner of height map is zero

for each pixel in the left column of height map
  height value=previous height value + corresponding q value
end

for each row
  for each element of the row except for leftmost
    height value = previous height value + corresponding p value
  end
end

Problem: Noise and numerical (in)accuracy are added up and result in distorted surface.

Solution: Choose several different integration paths, and build average height map.

Mathematical Property: Integrability

- Smooth, C2 continuous surface:

$$Z(x,y)_{xy} = Z(x,y)_{yx}$$

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

$\Rightarrow$ check if $\left(\frac{\partial p}{\partial y} - \frac{\partial q}{\partial x}\right)^2$ is small

SHAPE FROM SHADING (Calculus of Variations Approach)

- First Attempt: Minimize error in agreement with Image Irradiance Equation over the region of interest:

$$\iint_{object} (I(x,y) - R(p,q))^2 \, dxdy$$

SHAPE FROM SHADING (Calculus of Variations Approach)

- Better Attempt: Regularize the Minimization of error in agreement with Image Irradiance Equation over the region of interest:

$$\iint_{object} {p_x}^2 + {p_y}^2 + {q^2}_x + {q^2}_y + \lambda(I(x,y) - R(p,q))^2 \, dxdy$$

Horn: Characteristic Strip Method

Horn,
Chapter11,
pp. 250-255

Small step in x,y → change in depth:

$$\delta z = p\,\delta x + q\,\delta y$$

New values of p,q at this new point (x,y):

$$\delta p = r\,\delta x + s\,\delta y \quad \text{and} \quad \delta q = s\,\delta x + t\,\delta y$$

(r, s, t: second partial derivatives of z(x,y) w.r.t. x and y)

$$\begin{pmatrix} \delta p \\ \delta q \end{pmatrix} = \mathbf{H} \begin{pmatrix} \delta x \\ \delta y \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} r & s \\ s & t \end{pmatrix}$$

Hessian: curv. of surface

$$r = \frac{\partial p}{\partial x} \qquad s = \frac{\partial p}{\partial y} = \frac{\partial q}{\partial x} \qquad t = \frac{\partial q}{\partial y}$$

Horn: Characteristic Strip Method

Horn,
Chapter11,
pp. 250-255

Irradiance Equation, Reflectance Map:

$$E(x,y) = R(p,q)$$

Derivatives (chain rule):

$$E_x = r\,R_p + s\,R_q \qquad \text{and} \qquad E_y = s\,R_p + t\,R_q,$$

$$\begin{pmatrix} E_x \\ E_y \end{pmatrix} = \mathbf{H} \begin{pmatrix} R_p \\ R_q \end{pmatrix},$$

Relationship between gradient in the image and gradient in the reflectance map

Horn: Characteristic Strip Method

Horn,
Chapter11,
pp. 250-255

2 Equations for 3 unknowns (r,s,t): We can't continue in artibrary direction.

→ Trick: Specially chosen direction

$$\begin{pmatrix} \delta x \\ \delta y \end{pmatrix} = \begin{pmatrix} R_p \\ R_q \end{pmatrix} \delta \xi$$

Step in image E(x,y) parallel to gradient in R

Horn: Characteristic Strip Method

Horn,
Chapter11,
pp. 250-255

2 Equations for 3 unknowns (r,s,t): We can't continue in artibrary direction.

→ Trick: Specially chosen direction

$$\begin{pmatrix} \delta x \\ \delta y \end{pmatrix} = \begin{pmatrix} R_p \\ R_q \end{pmatrix} \delta\xi$$

Step in image E(x,y) parallel to gradient in R

Solving for new values for p,q:

$$\begin{pmatrix} \delta p \\ \delta q \end{pmatrix} = \mathbf{H} \begin{pmatrix} \delta x \\ \delta y \end{pmatrix} = \mathbf{H} \begin{pmatrix} R_p \\ R_q \end{pmatrix} \delta\xi$$

$$\longrightarrow \begin{pmatrix} \delta p \\ \delta q \end{pmatrix} = \begin{pmatrix} E_x \\ E_y \end{pmatrix} \delta\xi.$$

Change in (p,q) can be computed via gradient of image

Horn: Characteristic Strip Method

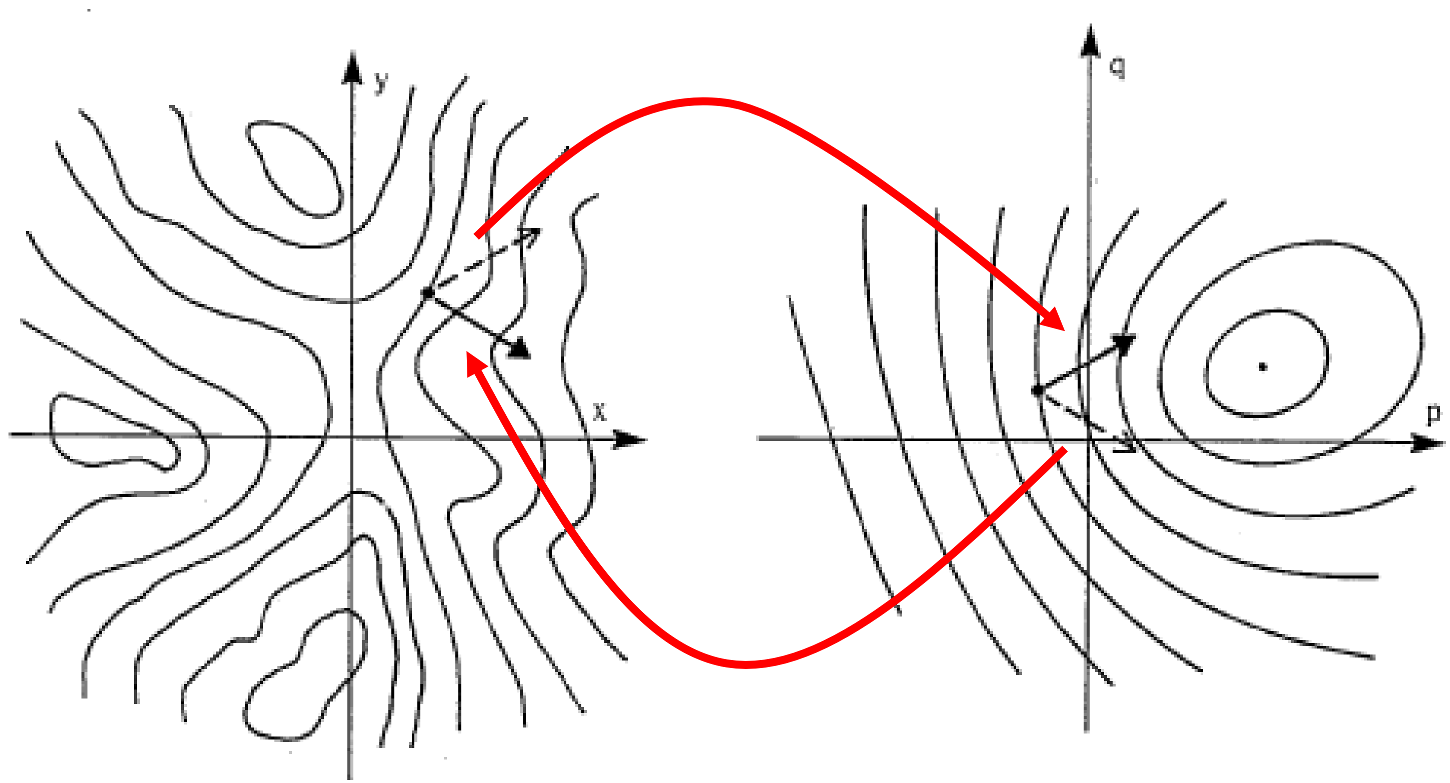

Figure 11-6. Curiously, the step taken in pq-space is parallel to the gradient of $E(x, y)$, while the step taken in xy-space is parallel to the gradient of $R(p, q)$.

Horn: Characteristic Strip Method

$$\dot{x} = R_p, \qquad \dot{y} = R_q, \qquad \dot{z} = p\,R_p + q\,R_q,$$

$$\dot{p} = E_x, \qquad \dot{q} = E_y,$$

dots denote differentiation with respect to ξ

Solution of differential equations: Curve on surface

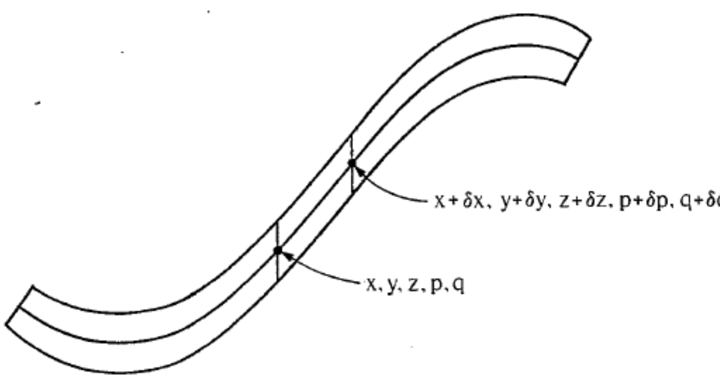

Figure 11-5. The solution of the shape-from-shading problem is determined by solving five differential equations for x, y, z, p, and q. The result is a characteristic strip, a curve in space along which surface orientation is known.

Horn: Characteristic Strip Method

Shape recovery via characteristic strips

Shape from Shading via Characteristic Curves

Given

- $I(x,y)$ of an (orthographic) projection of unknown $H(x,y)$
- The reflectance map $R(p,q)$
- Initial data $x_0, y_0, H(x_0,y_0), p(x_0,y_0), q(x_0,y_0)$

Develop a curve solution on $H(x,y)$ by taking small steps of size δs via the system

$$\delta x = R_p \delta s$$

$$\delta y = R_q \delta s$$

$$\delta H = (pR_p + qR_q)\delta s$$

$$\delta p = I_x \delta s$$

$$\delta q = I_y \delta s$$

Horn: Characteristic Strip Method

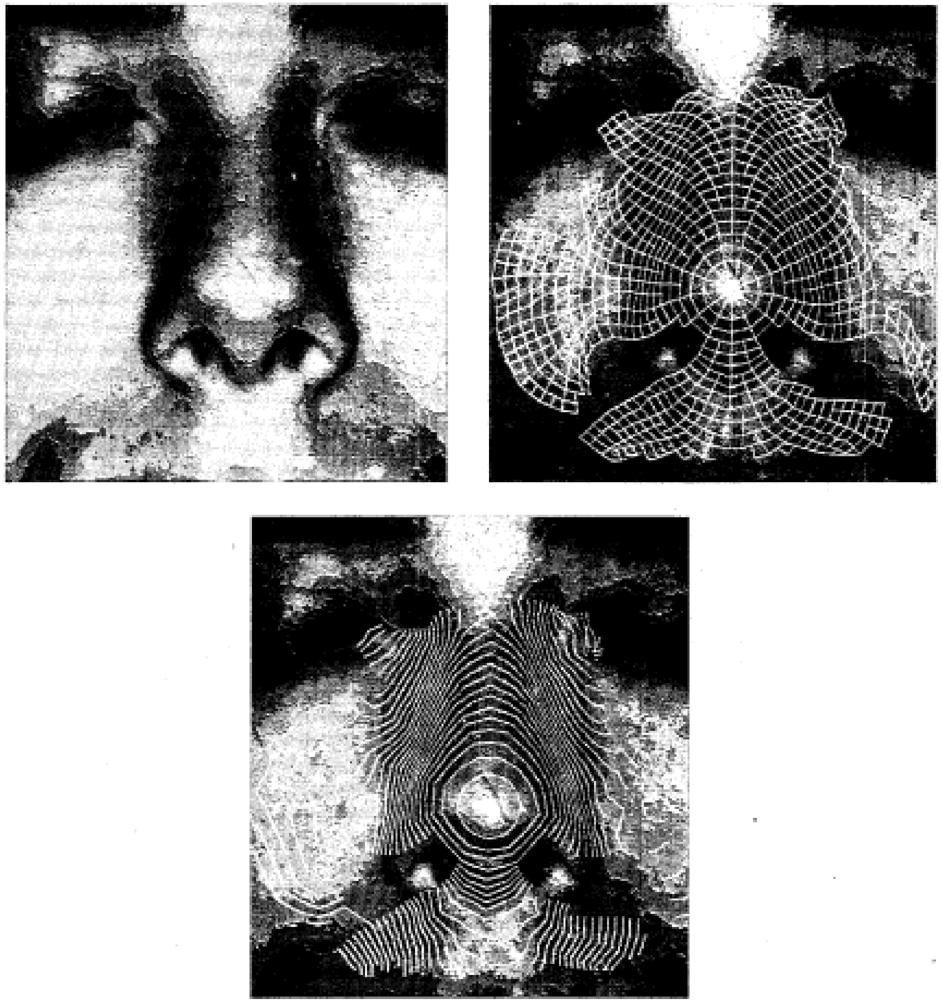

Figure 11-7. The shape-from-shading method is applied here to the recovery of the shape of a nose. The first picture shows the (crudely quantized) gray-level image available to the program. The second picture shows the base characteristics superimposed, while the third shows a contour map computed from the elevations found along the characteristic curves.

Horn, Chapter11, pp. 250-255

Another Solution to SFS: Kimmel, Siddiqi, Kimia, Bruckstein

Proposed Solution: Evolve a curve such that it tracks the height contours of $z(x, y)$.
[Kimmel *et al.*, IJCV95]

Height climbed while progressing a distance $|\Delta C|$ in the direction $\hat{n}$ <u>in the (x, y) plane</u> is given by $|\Delta C| = |\Delta z| \cot(\alpha)$.

Let z denote time in the course of evolution, *i.e.*, $z = t$. Since $E = \rho\lambda \cos(\alpha)$, we have

$$\left|\frac{\Delta C}{\Delta t}\right| = cot(\alpha) = \frac{E/\rho\lambda}{\sqrt{1 - (E/\rho\lambda)^2}}. \qquad (11)$$

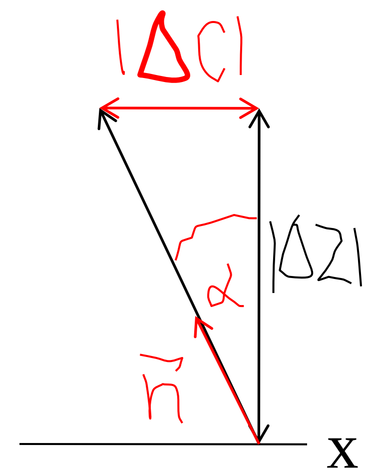

pdf document

Kimmel, Siddiqi, Kimia, Bruckstein

Proposed Solution: Evolve a curve such that it tracks the height contours of $z(x, y)$.
[Kimmel *et al.*, IJCV95]

The curve evolution equation is:

$$\begin{cases} \frac{\partial \mathcal{C}}{\partial t} & = \frac{E/\rho\lambda}{\sqrt{1-E^2/(\rho\lambda)^2}}.\hat{n}, \\ \mathcal{C}(s,0) & = \mathcal{C}_0(s). \end{cases}$$

Kimmel, Siddiqi, Kimia, Bruckstein

Examples - Pyramids

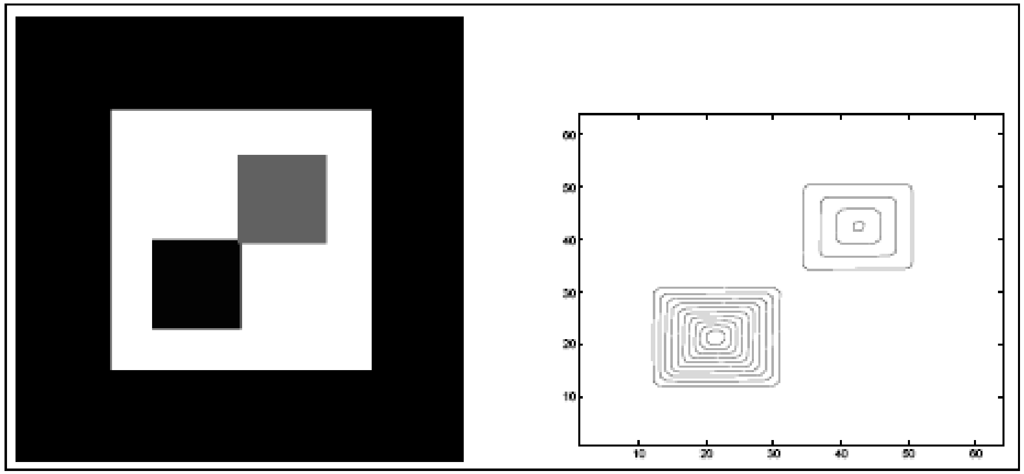

shaded image — equal height contours

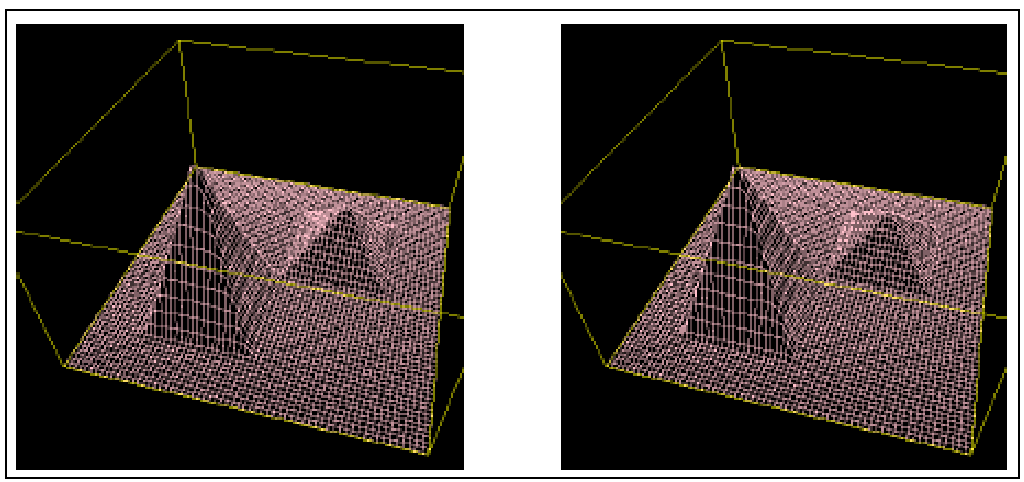

numerical solution — true surface

Kimmel, Siddiqi, Kimia, Bruckstein

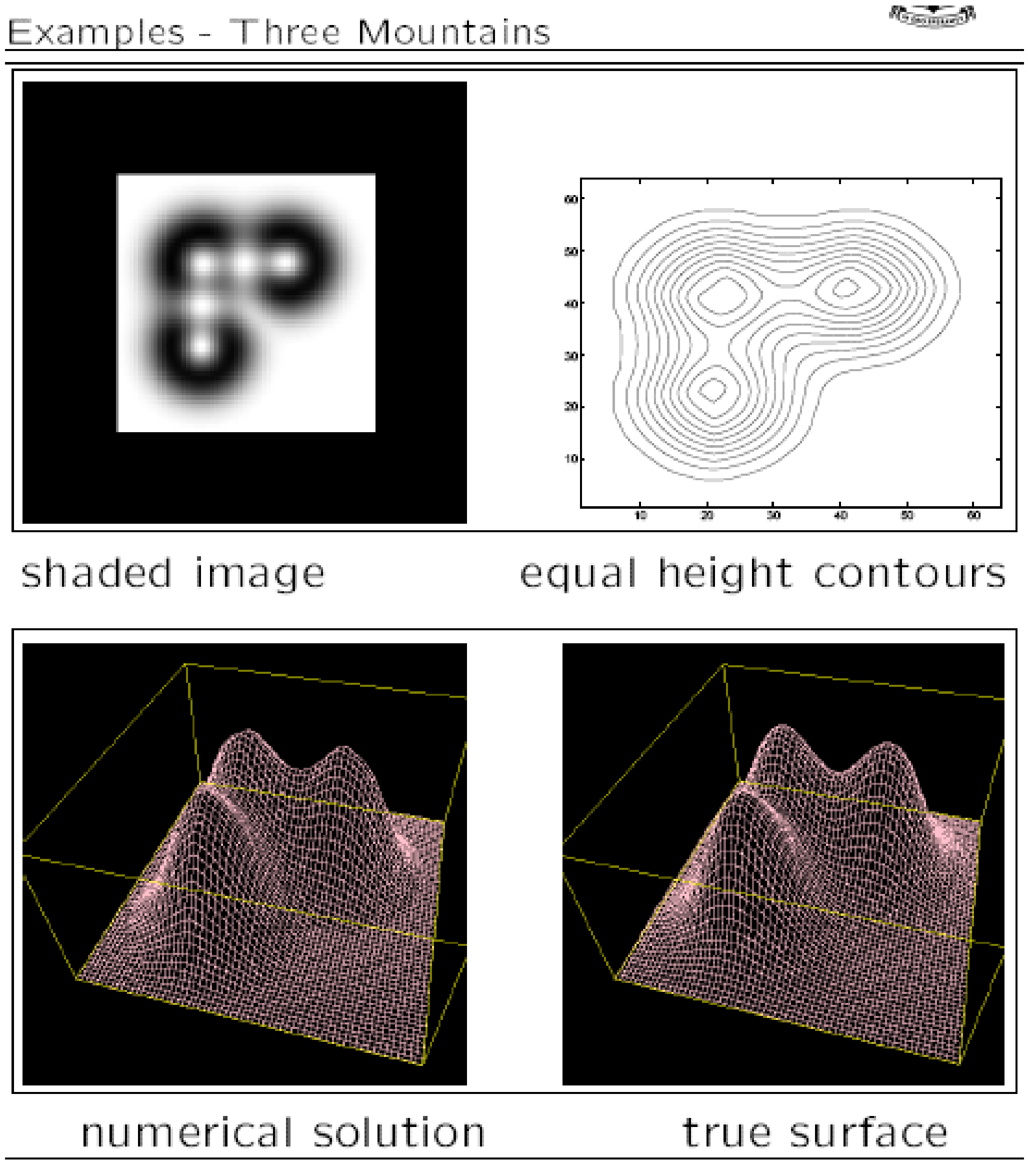

Application Area: Geography

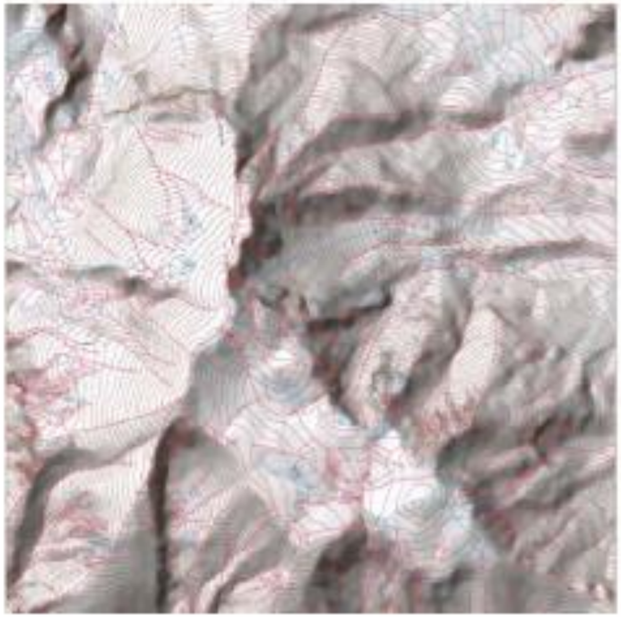

Abb. 13: Schräglicht-Schräglichtschattierung des Untersuchungsgebiets auf Basis des verbesserten DGM (Auflösung = 10m), überlagert mit den harten (=rot) und weichen (=blau) Strukturen und den Höhenlinien (=braun). Abbildungsmaßstab 1:35.000.

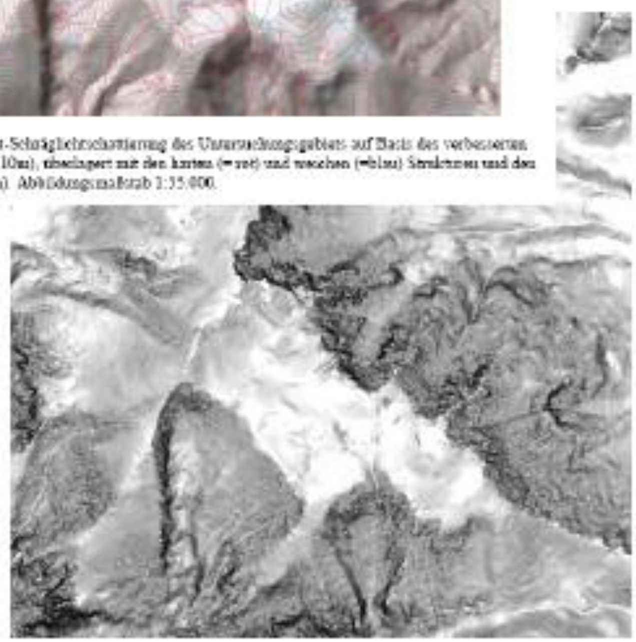

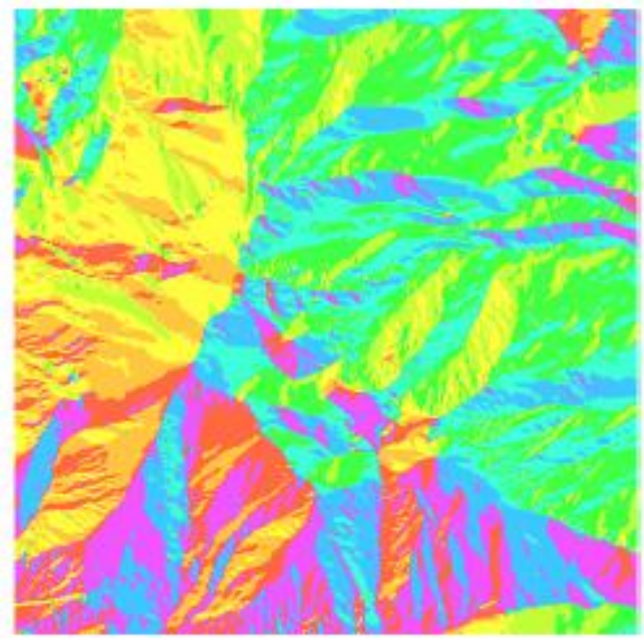

Application: Braille Code

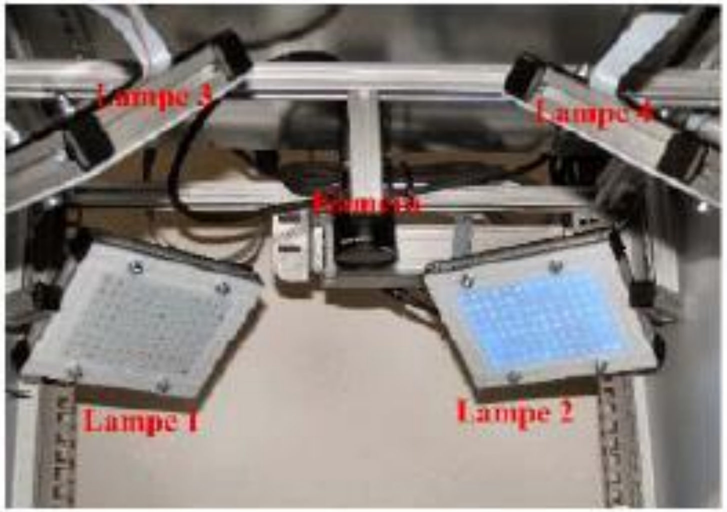

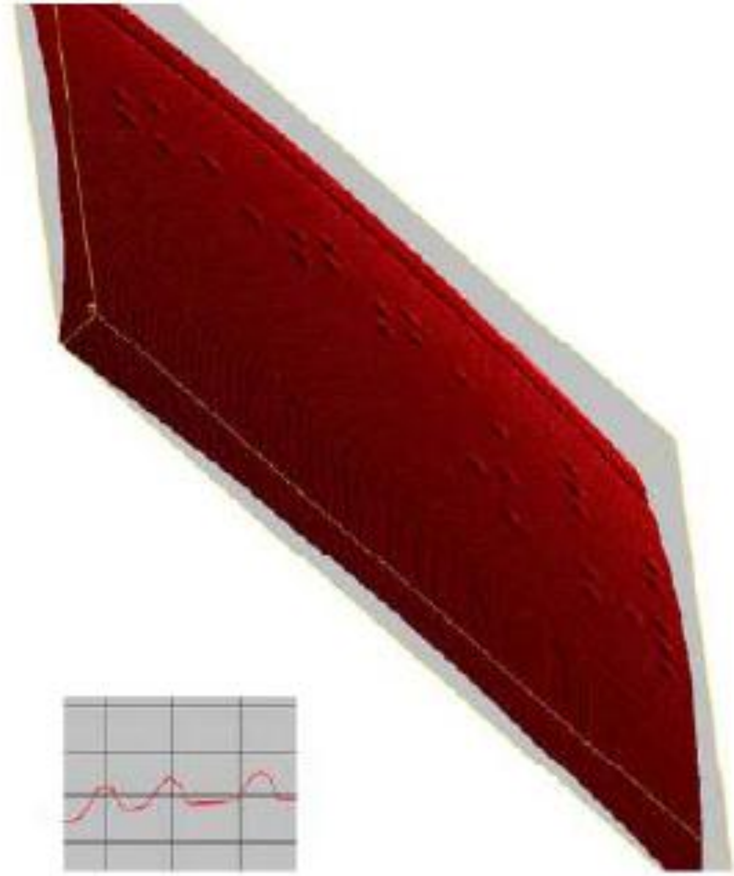

Abbildung 3:

Oben links: Messanordnung mit einer Kamera und vier blauen LED-Leuchtfeldern.
Unten links: Ausschnitt einer Faltschachtel mit Blindenschrift-Prägung.
Rechts: 3D-Bild nach SFS-Analyse. Darunter ist ein Höhenprofil durch drei Braille-Punkte dargestellt.

[pdf document]

Mars Rover Heads to a New Crater NYT Sept 22, 2008

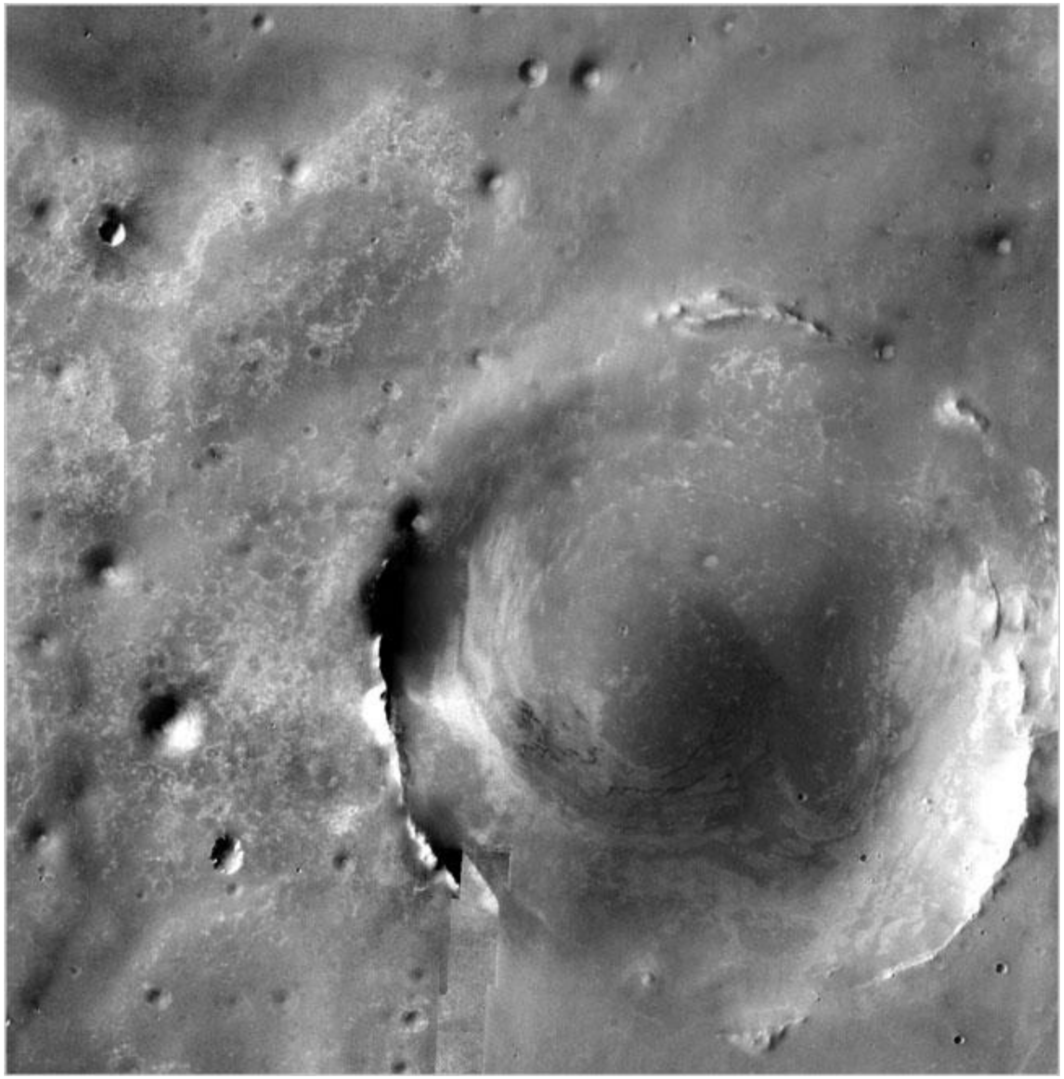

Limitations

- Controlled lighting environment
  - Specular highlights?
  - Partial shadows?
  - Complex interrreflections?

- Fixed camera
  - Moving camera?
  - Multiple cameras?

=> Another approach: binocular / geometric stereo