Camera calibration & radiometry

- Reading:
  - Chapter 2, and section 5.4, Forsyth & Ponce
  - Chapter 10, Horn
- Optional reading:
  - Chapter 4, Forsyth & Ponce

Sept. 17, 2002
MIT 6.801/6.866
Profs. Freeman and Darrell

6.801/6.866 Machine Vision

Syllabus

#	Date	Description	Readings	Assignments	Materials
1	9/5	Course Introduction		Pset #0 (not collected)	Freeman Slides Darrell Slides Matlab Tutorial Diary
2	9/10	Cameras, Lenses, and Sensors	Req: FP 1 Opt: H 2.1, 2.3		Freeman Slides
3	9/12	Radiometry and Shading Models	Req: FP 2, 5.4; H 10 Opt: FP 4	Pset #1 Assigned	Freeman Slides
4	9/17	Color	Req: FP 6.1-6.4		
5	9/19	Multiview Geometry	Req: FP 10		
6	9/24	Stereo	Req: FP 11; H 13	Pset #1 Due	
7	9/26	Shape from Shading		Pset #2 Assigned	

6.801/6.866 Machine Vision

Syllabus

#	Date	Description	Readings	Assignments	Materials
1	9/5	Course Introduction		Pset #0 (not collected)	Freeman Slides Darrell Slides Matlab Tutorial Diary
2	9/10	Cameras, Lenses, and Sensors	Req: FP 1 Opt: H 2.1, 2.3		Freeman Slides
3	9/12	Radiometry and Shading Models	Req: FP 2, 5.4; H 10 Opt: FP 4	Pset #1 Assigned	Freeman Slides
4	9/17	Radiometry and Shading Models			
5	9/19	Multiview Geometry	Req: FP 10		
6	9/24	Stereo	Req: FP 11; H 13	Pset #1 Due	
7	9/26	Color	Req: FP 6.1-6.4	Pset #2 Assigned	

Camera calibration

- Geometric: how *positions* in the image relate to 3-d positions in the world.
- Photometric: how the *intensities* in the image relate surface and lighting properties in the world.

Calibration target

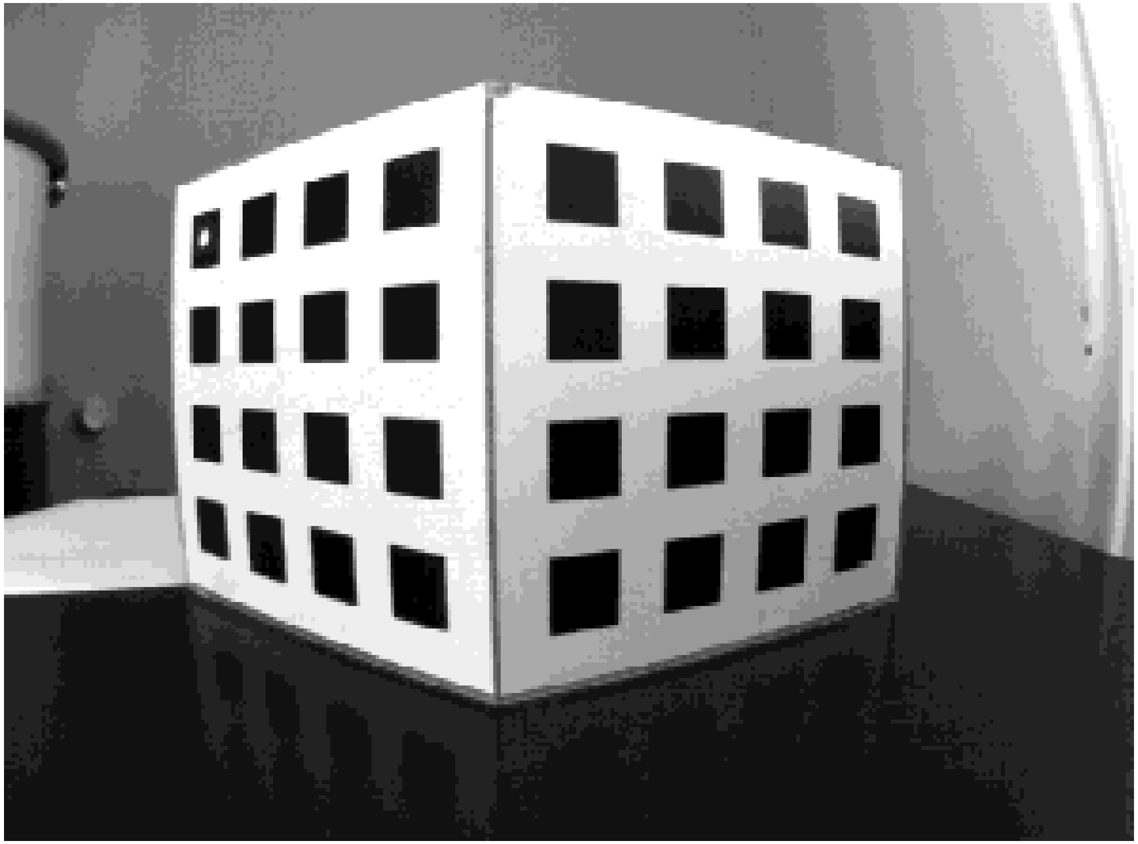

The Opti-CAL Calibration Target Image

http://www.kinetic.bc.ca/CompVision/opti-CAL.html

From last lecture: camera calibration

pixel coordinates

world coordinates

$$\vec{p} = \frac{1}{z} M \vec{P}$$

$$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \frac{1}{z} \begin{pmatrix} . & m_1^T & . & . \\ . & m_2^T & . & . \\ . & m_3^T & . & . \end{pmatrix} \begin{pmatrix} W_x \\ W_y \\ W_z \\ 1 \end{pmatrix}$$

$$\begin{cases} u = \dfrac{m_1 \cdot \vec{P}}{m_3 \cdot \vec{P}} \\[2ex] v = \dfrac{m_2 \cdot \vec{P}}{m_3 \cdot \vec{P}} \end{cases}$$

z is in the *camera* coordinate system, but we can solve for that, since $1 = \dfrac{m_3 \cdot \vec{P}}{z}$, leading to:

Camera calibration

Because of these relations,

$$u = \frac{m_1 \cdot \vec{P}}{m_3 \cdot \vec{P}}$$

$$v = \frac{m_2 \cdot \vec{P}}{m_3 \cdot \vec{P}}$$

For each feature point, i, we have:

$$(m_1 - u_i m_3) \cdot \vec{P}_i = 0$$
$$(m_2 - v_i m_3) \cdot \vec{P}_i = 0$$

Camera calibration

$$\begin{pmatrix} P_{1x} & P_{1y} & P_{1z} & 1 & 0 & 0 & 0 & 0 & -u_1 P_{1x} & -u_1 P_{1y} & -u_1 P_{1z} & -u_1 \\ 0 & 0 & 0 & 0 & P_{1x} & P_{1y} & P_{1z} & 1 & -v_1 P_{1x} & -v_1 P_{1y} & -v_1 P_{1z} & -v_1 \\ & & & & & & \cdots & \cdots & \cdots & & & \\ P_{nx} & P_{ny} & P_{nz} & 1 & 0 & 0 & 0 & 0 & -u_n P_{nx} & -u_n P_{ny} & -u_n P_{nz} & -u_n \\ 0 & 0 & 0 & 0 & P_{nx} & P_{ny} & P_{nz} & 1 & -v_n P_{nx} & -v_n P_{ny} & -v_n P_{nz} & -v_n \end{pmatrix} \begin{pmatrix} m_{11} \\ m_{12} \\ m_{13} \\ m_{14} \\ m_{21} \\ m_{22} \\ m_{23} \\ m_{24} \\ m_{31} \\ m_{32} \\ m_{33} \\ m_{34} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}$$

P m = 0

We want to solve for the unit vector m (the stacked one) that minimizes $|Pm|^2$

The eigenvector corresponding to the minimum eigenvalue of the matrix $P^T P$ gives us that (see Forsyth&Ponce, 3.1).

What makes a valid M matrix?

A projection matrix can be written explicitly as a function of its five intrinsic parameters (α, β, u_0, v_0, and θ) and its six extrinsic ones (the three angles defining $\mathcal{R}$ and the three coordinates of $\boldsymbol{t}$), namely,

$$\mathcal{M} = \begin{pmatrix} \alpha \boldsymbol{r}_1^T - \alpha \cot\theta \boldsymbol{r}_2^T + u_0 \boldsymbol{r}_3^T & \alpha t_x - \alpha \cot\theta t_y + u_0 t_z \\ \dfrac{\beta}{\sin\theta} \boldsymbol{r}_2^T + v_0 \boldsymbol{r}_3^T & \dfrac{\beta}{\sin\theta} t_y + v_0 t_z \\ \boldsymbol{r}_3^T & t_z \end{pmatrix}, \tag{2.17}$$

where $\boldsymbol{r}_1^T$, $\boldsymbol{r}_2^T$, and $\boldsymbol{r}_3^T$ denote the three rows of the matrix $\mathcal{R}$ and t_x, t_y, and t_z are the coordinates of the vector $\boldsymbol{t}$.

- $M = \begin{pmatrix} \mathrm{A} & \vec{\mathrm{b}} \end{pmatrix} = \begin{pmatrix} \vec{a}_1^T & \\ \vec{a}_2^T & \vec{b} \\ \vec{a}_3^T & \end{pmatrix}$ defined only up to a scale;

  normalize M so that $\left|\vec{a}_3^T\right| = \left|\vec{r}_3^T\right| = 1.$

- M is a perspective projection matrix iff $Det(A) \neq 0$

Camera calibration

- Geometric: how *positions* in the image relate to 3-d positions in the world.
- Photometric: how the *intensities* in the image relate surface and lighting properties in the world.

Irradiance, E

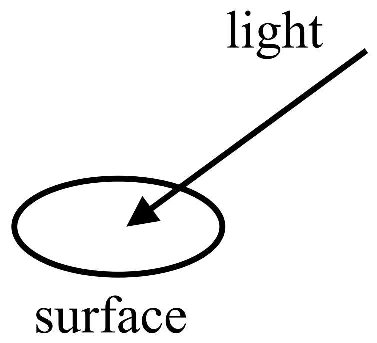

- Light power per unit area (watts per square meter) incident on a surface.

Radiance, L

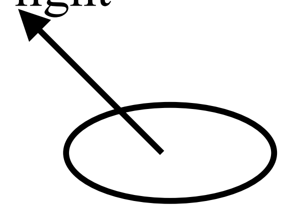

- Amount of light radiated from a surface into a given solid angle per unit area (watts per square meter per steradian).
- Note: the area is the foreshortened area, as seen from the direction that the light is being emitted.

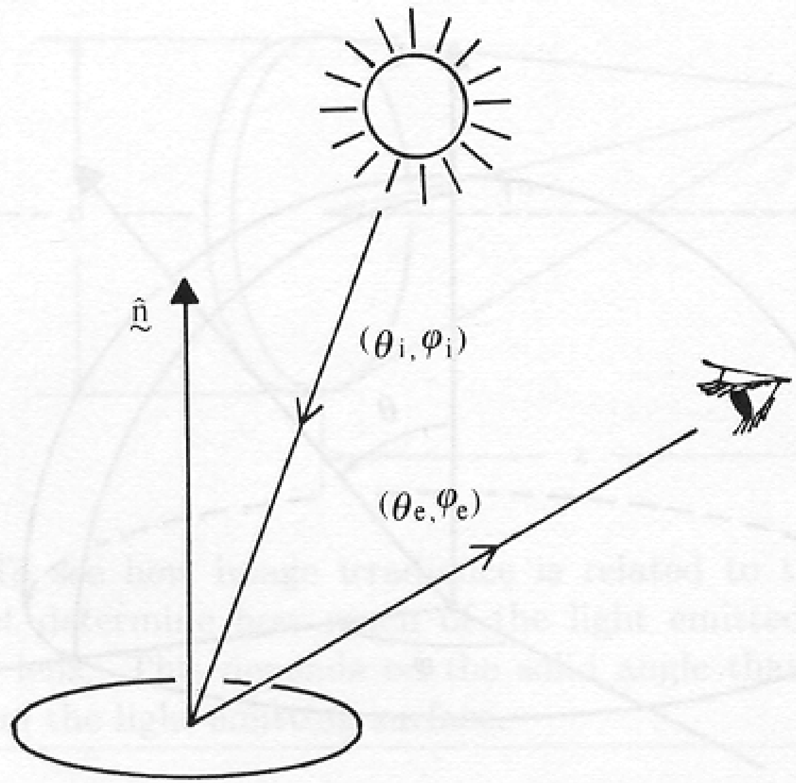

Horn, 1986

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{L(\theta_e, \phi_e)}{E(\theta_i, \phi_i)}$$

How does the world give us the brightness we observe at a point?

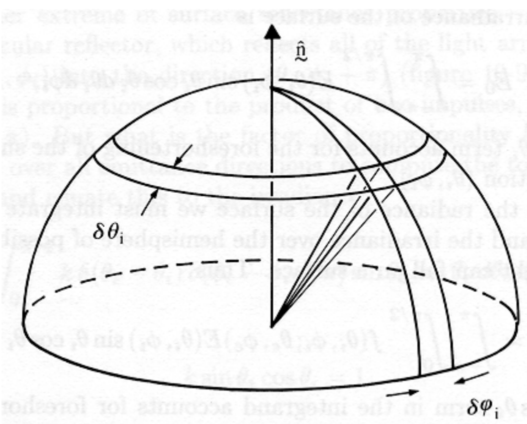

The total irradiance of the surface is:

radiance per solid angle

$$E_0 = \int_{-\pi}^{\pi} \int_{0}^{\pi/2} E(\theta_i, \phi_i) \sin(\theta_i) \cos(\theta_{\mathrm{i}}) \, d\theta_i \, d\phi_i$$

Accounting for the foreshortened area of center patch relative to illuminant.

The total radiance reflected from the surface patch is:

$$L(\theta_e, \phi_e) = \int_{-\pi}^{\pi} \int_{0}^{\pi/2} f(\theta_i, \phi_i, \theta_e, \phi_e) \, E(\theta_i, \phi_i) \sin(\theta_i) \cos(\theta_{\mathrm{i}}) \, d\theta_i \, d\phi_i$$

What you'd like to pull out from L

Pixel intensities may be proportional to radiance reflected from the surface patch:

$$L(\theta_e, \phi_e) = \int_{-\pi}^{\pi} \int_{0}^{\pi/2} f(\theta_i, \phi_i, \theta_e, \phi_e)\, E(\theta_i, \phi_i) \sin(\theta_i) \cos(\theta_\mathrm{i})\, d\theta_i\, d\phi_i$$

θ_e, ϕ_e **surface orientation relative to camera**

$f(\theta_i, \phi_i, \theta_e, \phi_e)$ **surface BRDF**

$E(\theta_i, \phi_i)$ **illumination conditions**

θ_i, ϕ_i **surface orientation relative to illumination**

That's hard, so let's focus on special cases for the rest of this lecture.

Special case BRDF: Lambertian reflectance

BRDF is a constant. These surfaces look equally bright from all viewing directions.

$$f(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{1}{\pi}$$

Radiance reflected from Lambertian surface illuminated by point source:

$$L(\theta_e, \phi_e) = \int_{-\pi}^{\pi} \int_{0}^{\pi/2} \frac{1}{\pi} \delta(\theta_i - \theta_0) \delta(\phi_i - \phi_0) \sin(\theta_i) \cos(\theta_i)\, d\theta_i\, d\phi_i$$

$$\propto \cos(\theta_0)$$

Reflectance map

- For orthographic projection, and light sources at infinity, the reflectance map is a useful tool for describing the relationship of surface orientation to image intensity.
- Describes the image intensity for a given surface orientation.
- Parameterize surface orientation by the partial derivatives p and q of surface height z.

Relate surface normal to p & q

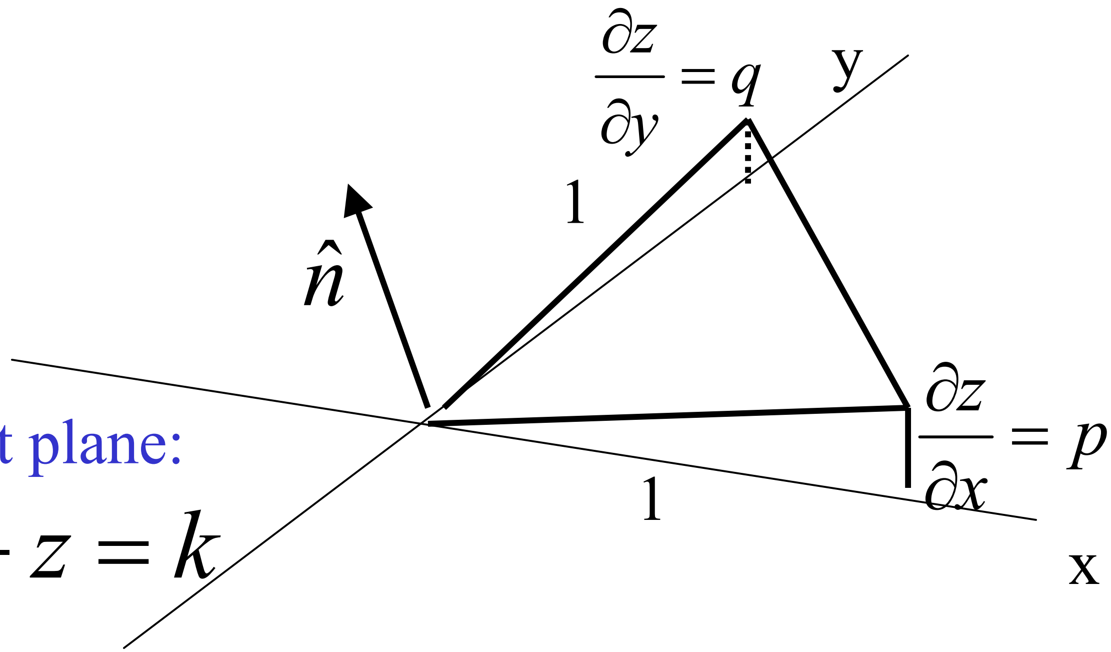

Local tangent plane:

$$px + qy - z = k$$

Unit normal to surface:

$$\hat{n} = \frac{\vec{n}}{|\vec{n}|} = \frac{\begin{pmatrix} -p & -q & 1 \end{pmatrix}^T}{\sqrt{1 + p^2 + q^2}}$$

Refl map for point source (in direction $\hat{s}$) Lambertian surface

For a Lambertian surface,

$$R(p,q) \propto \hat{n} \cdot \hat{s} = \cos(\theta_i)$$

$$= k \frac{1 + p_s p + q_s q}{\sqrt{1 + p^2 + q^2}\sqrt{1 + p_s^2 + q_s^2}}$$

Unit vector to source:

$$\hat{s} = \frac{\begin{pmatrix} -p_s & -q_s & 1 \end{pmatrix}^T}{\sqrt{1 + p_s{}^2 + q_s{}^2}}$$

Picture of Lambertian refl map

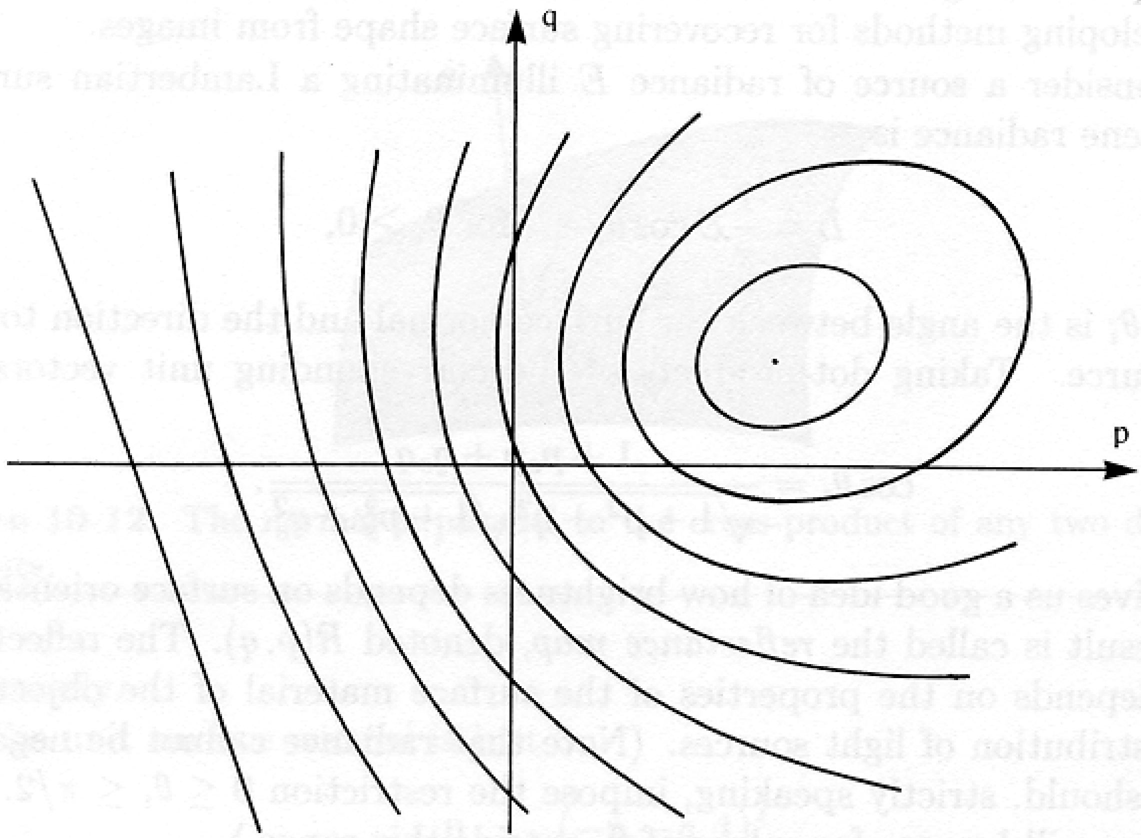

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p,q)$ occurs at 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.

Horn, 1986

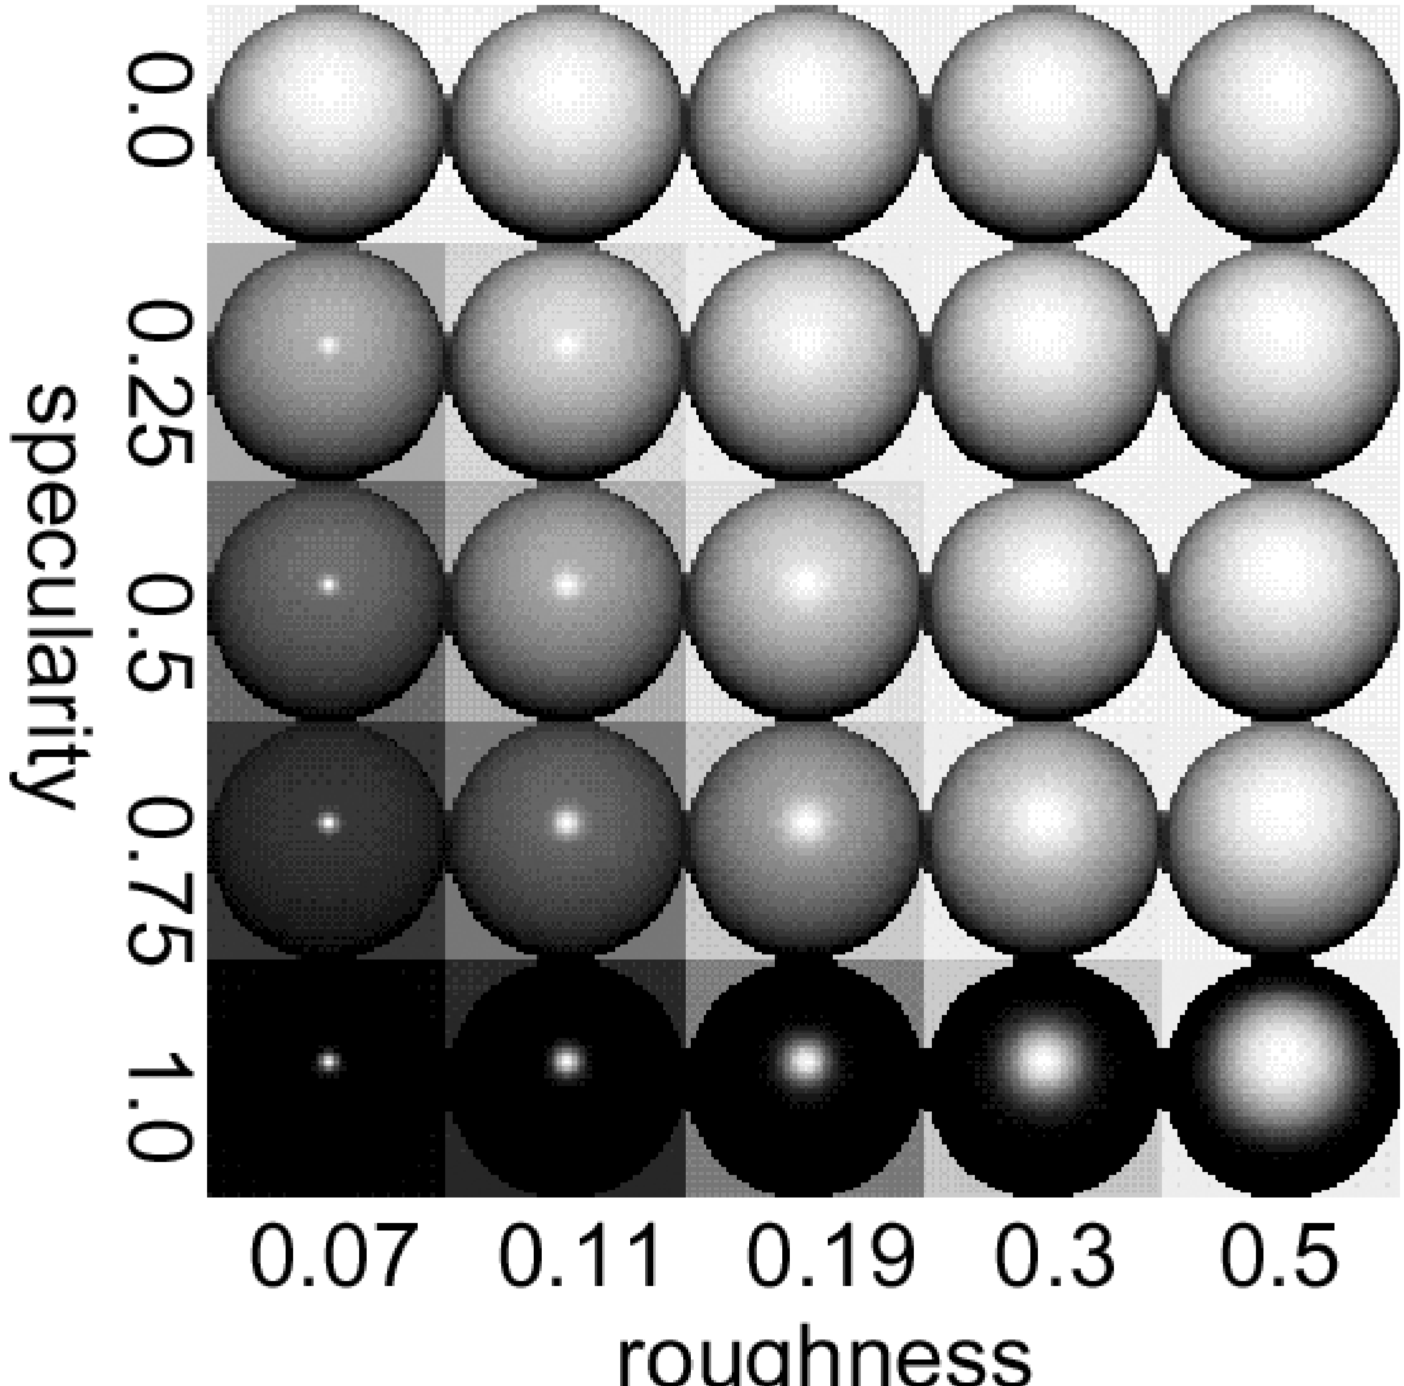

W. T. Freeman, *Exploiting the generic viewpoint assumption*, International Journal Computer Vision, 20 (3), 243-261, 1996

What constraints are there for form of reflectance map?

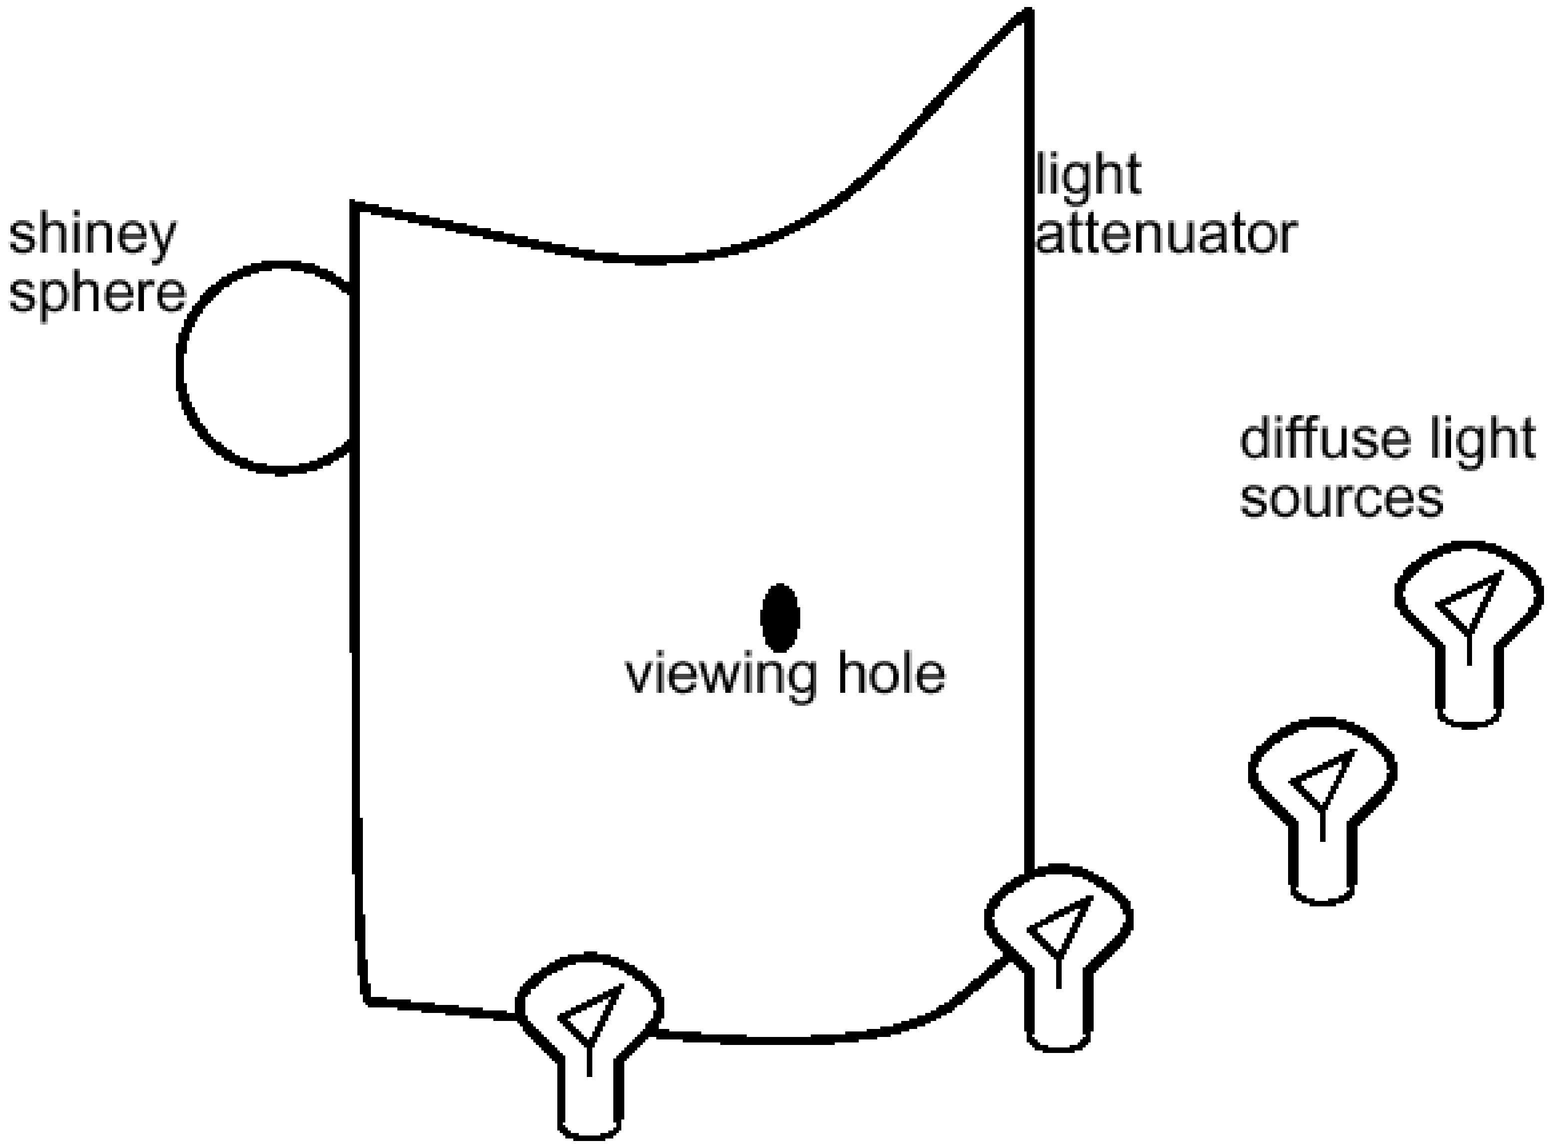

Freeman, 1994

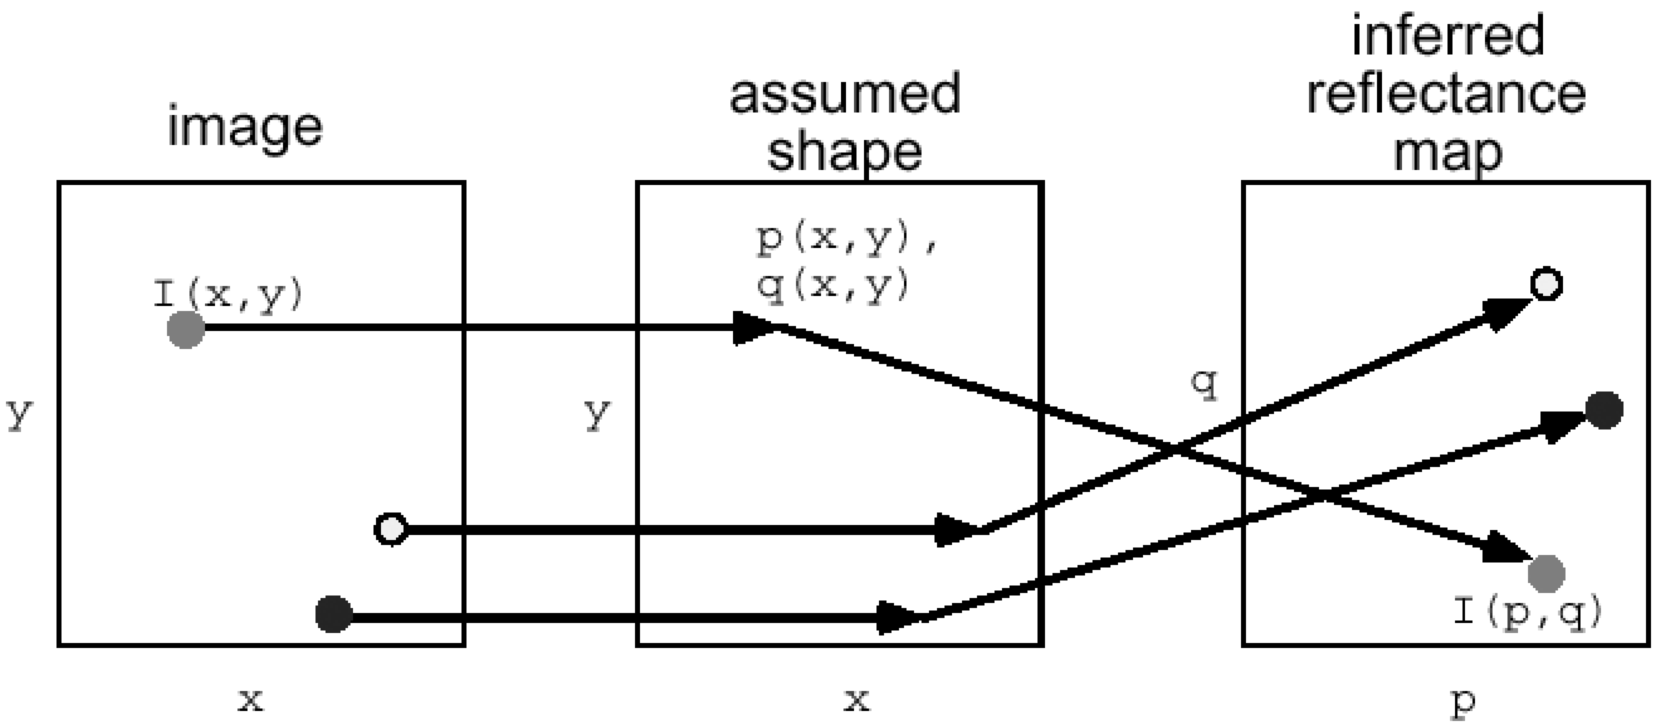

How to construct a feasible solution to the uncalibrated shape from shading problem.

Freeman, 1994

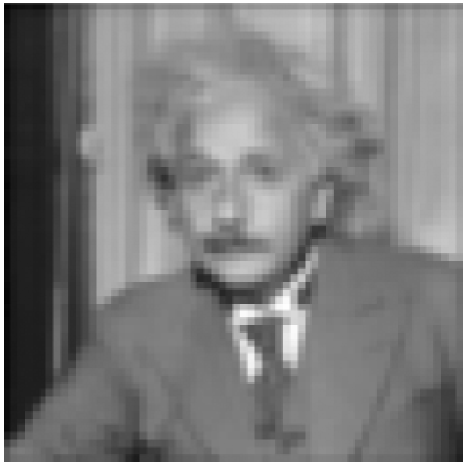

(a) image

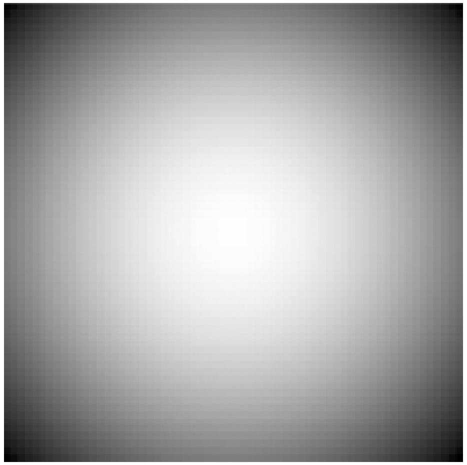

(b) assumed shape

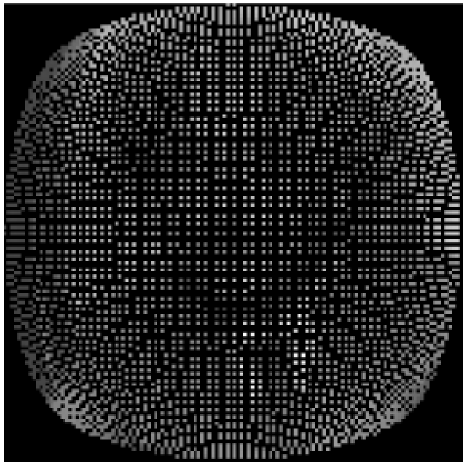

(c) inferred reflectance map

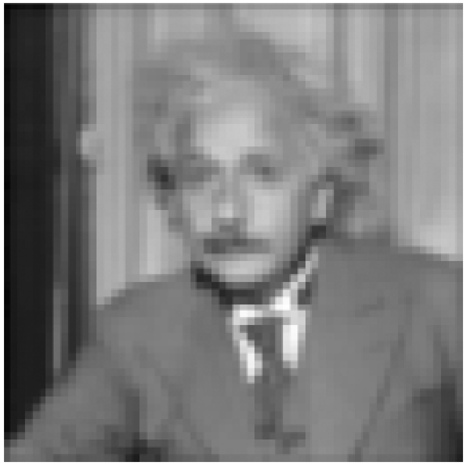

(d) re-rendered image

Figure 3: (a) Image. (b) Assumed shape which created the image, (a). (c) Reflectance map which, when applied to the shape (b), yields the image (a). Note that because the assumed shape has a nearly sphereical shape, the inferred reflectance map is a distorted replia of the original image. (d) Interpolated reflectance map. (e) Numerical rendering of shape (b) with reflectance map (d).

Freeman, 1994

Shape 1

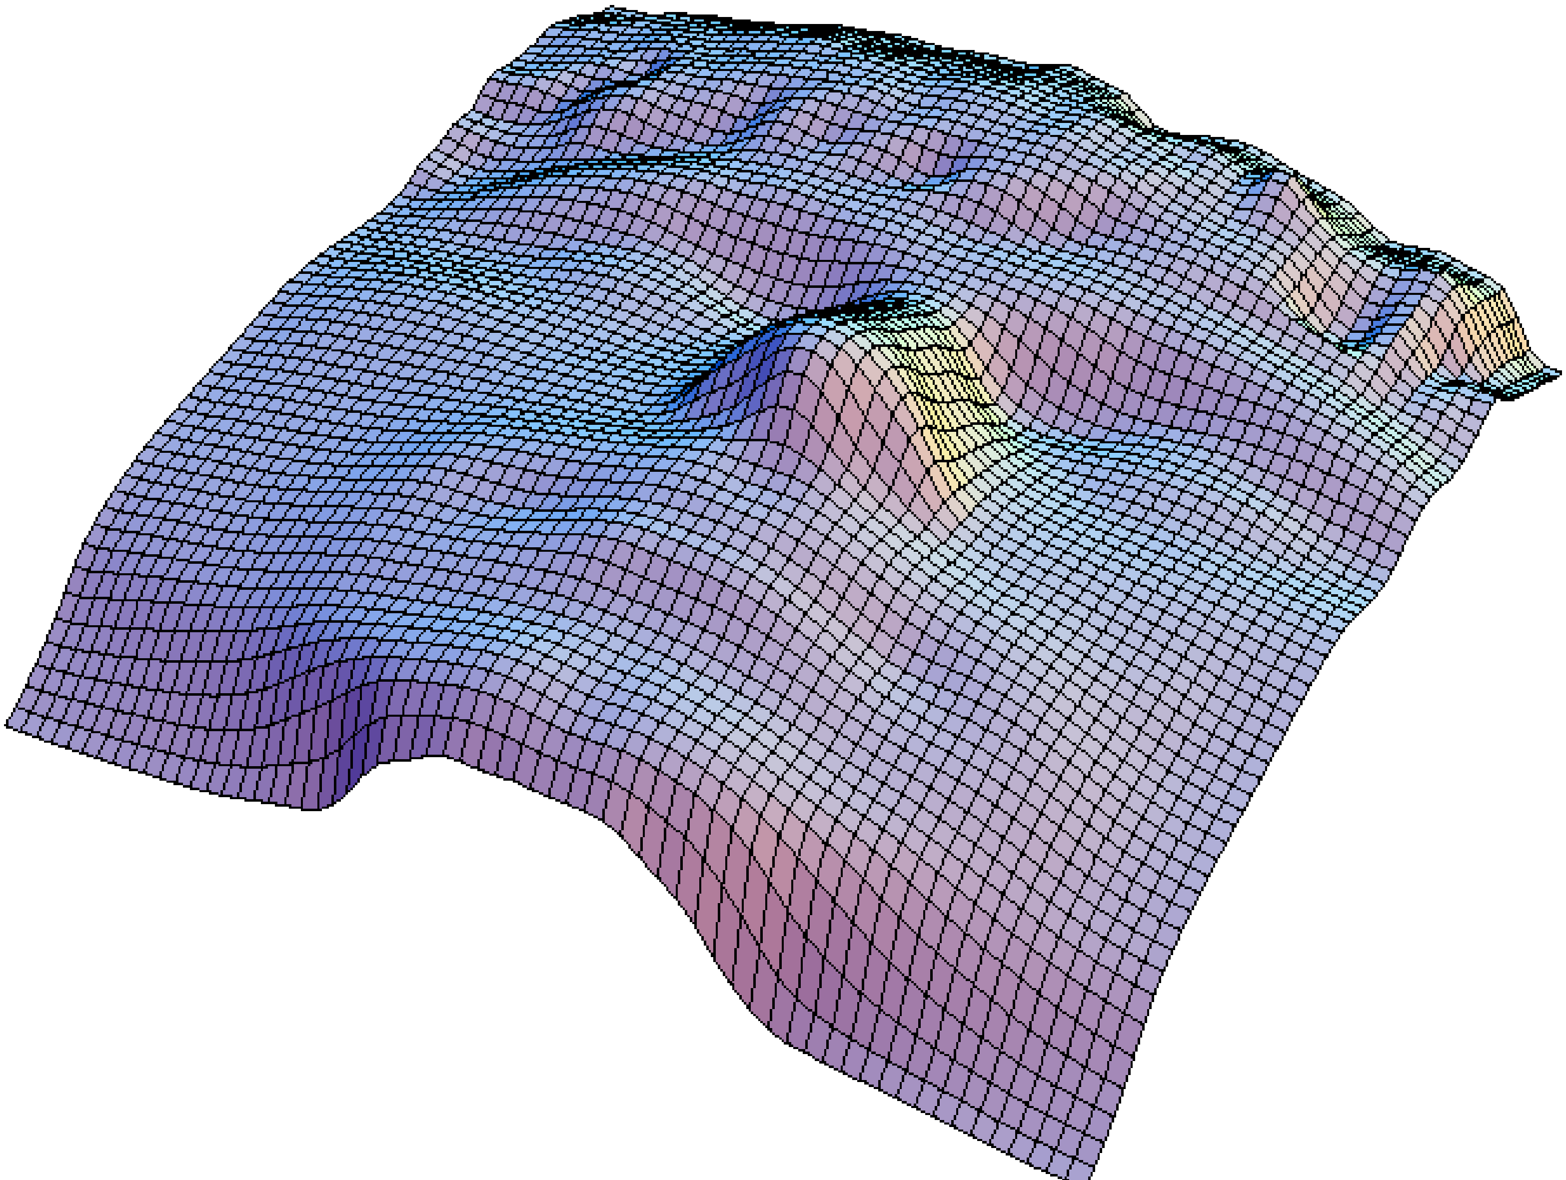

Freeman, 1994

Shape 2

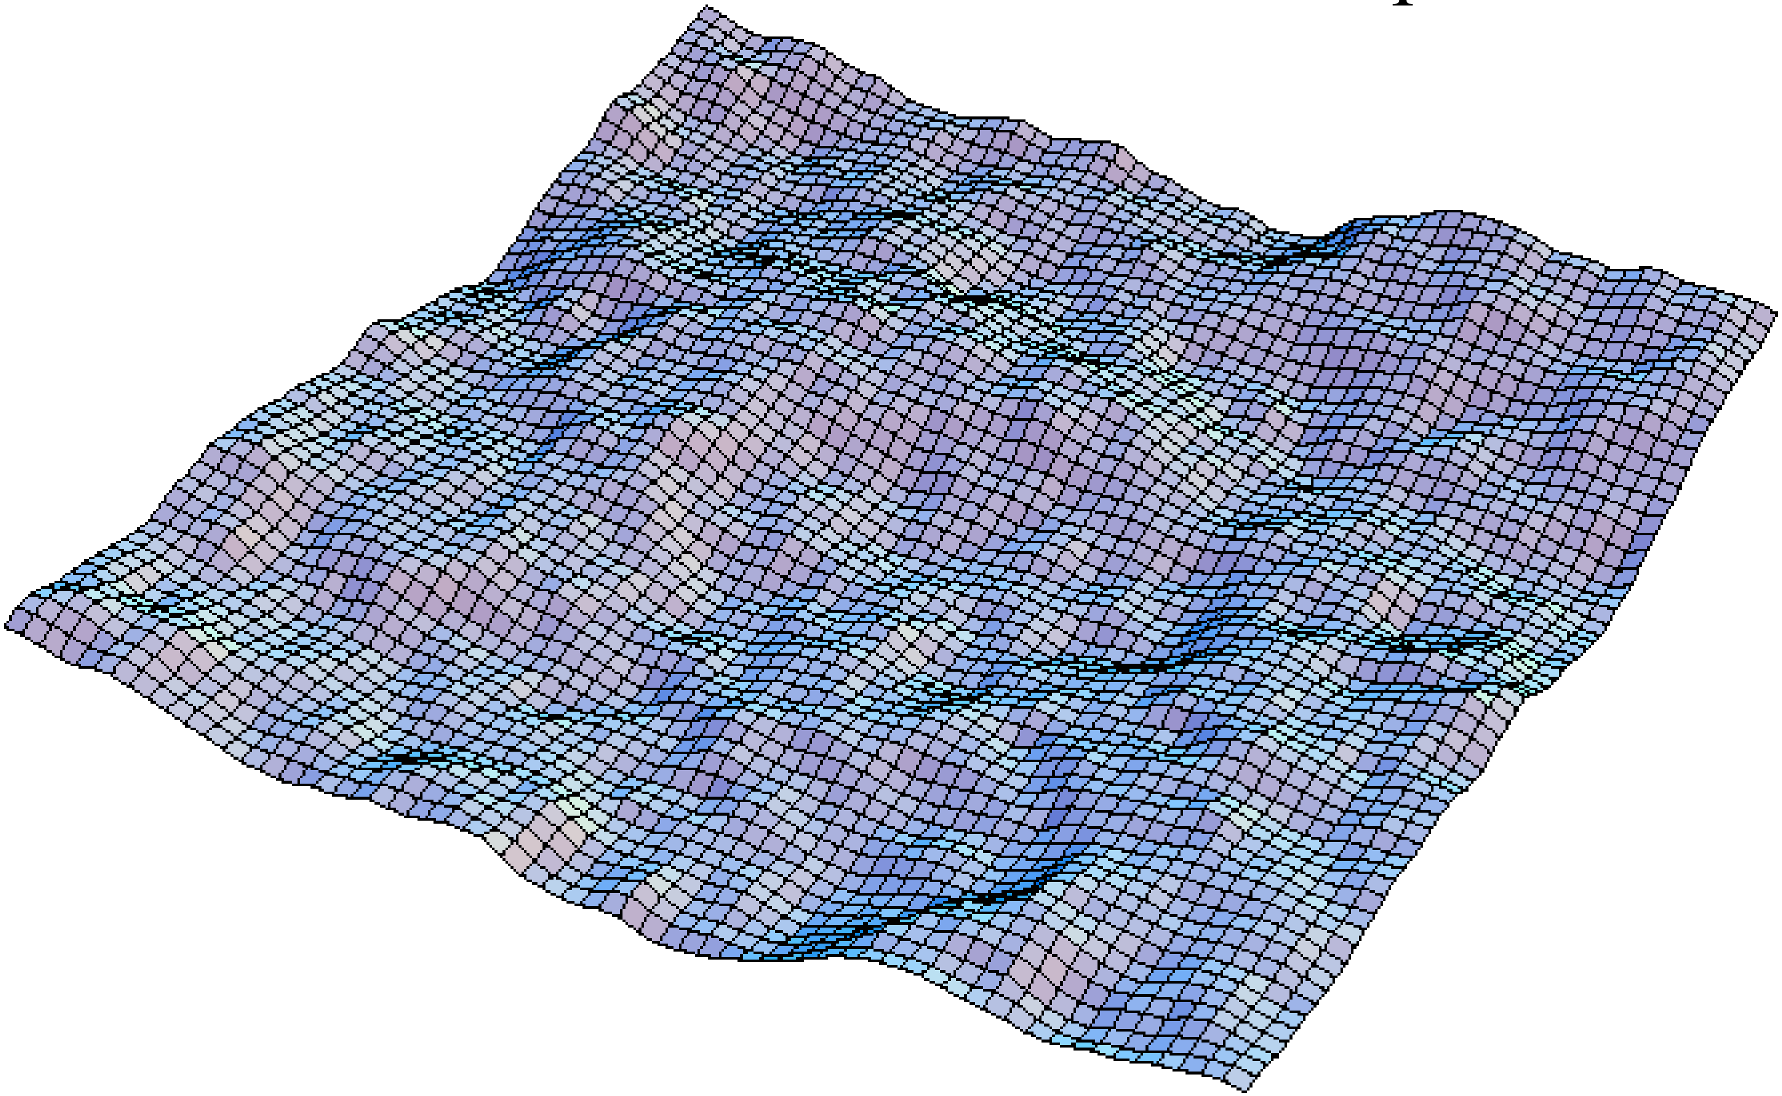

Freeman, 1994

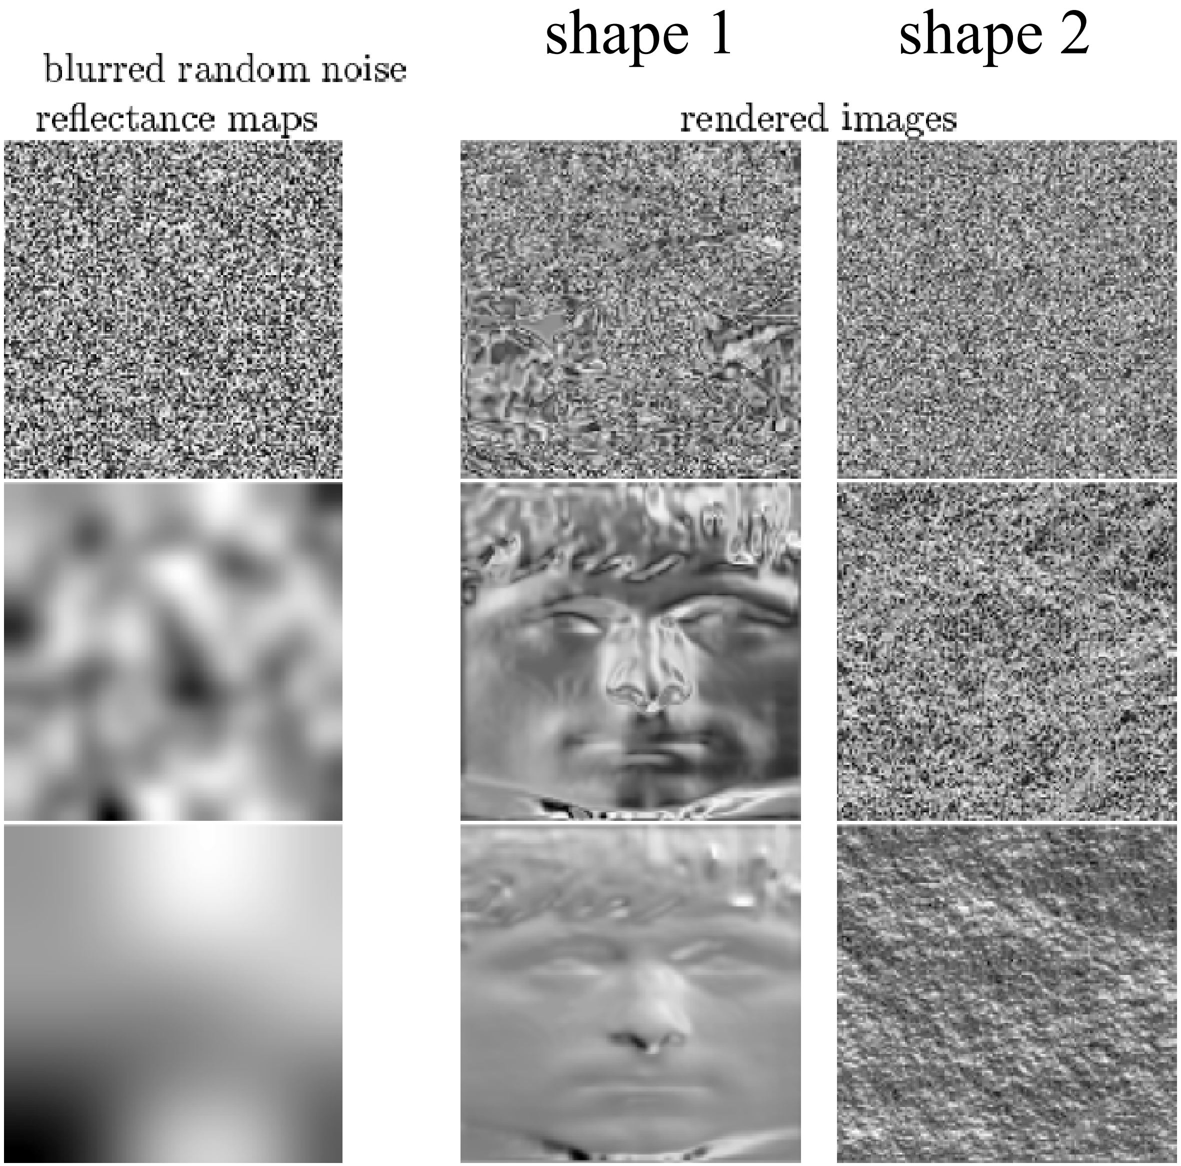

Freeman, 1994

Figure 6: Showing that blurred reflectance maps lead to shapes which are easier to interpret.

Freeman, 1994

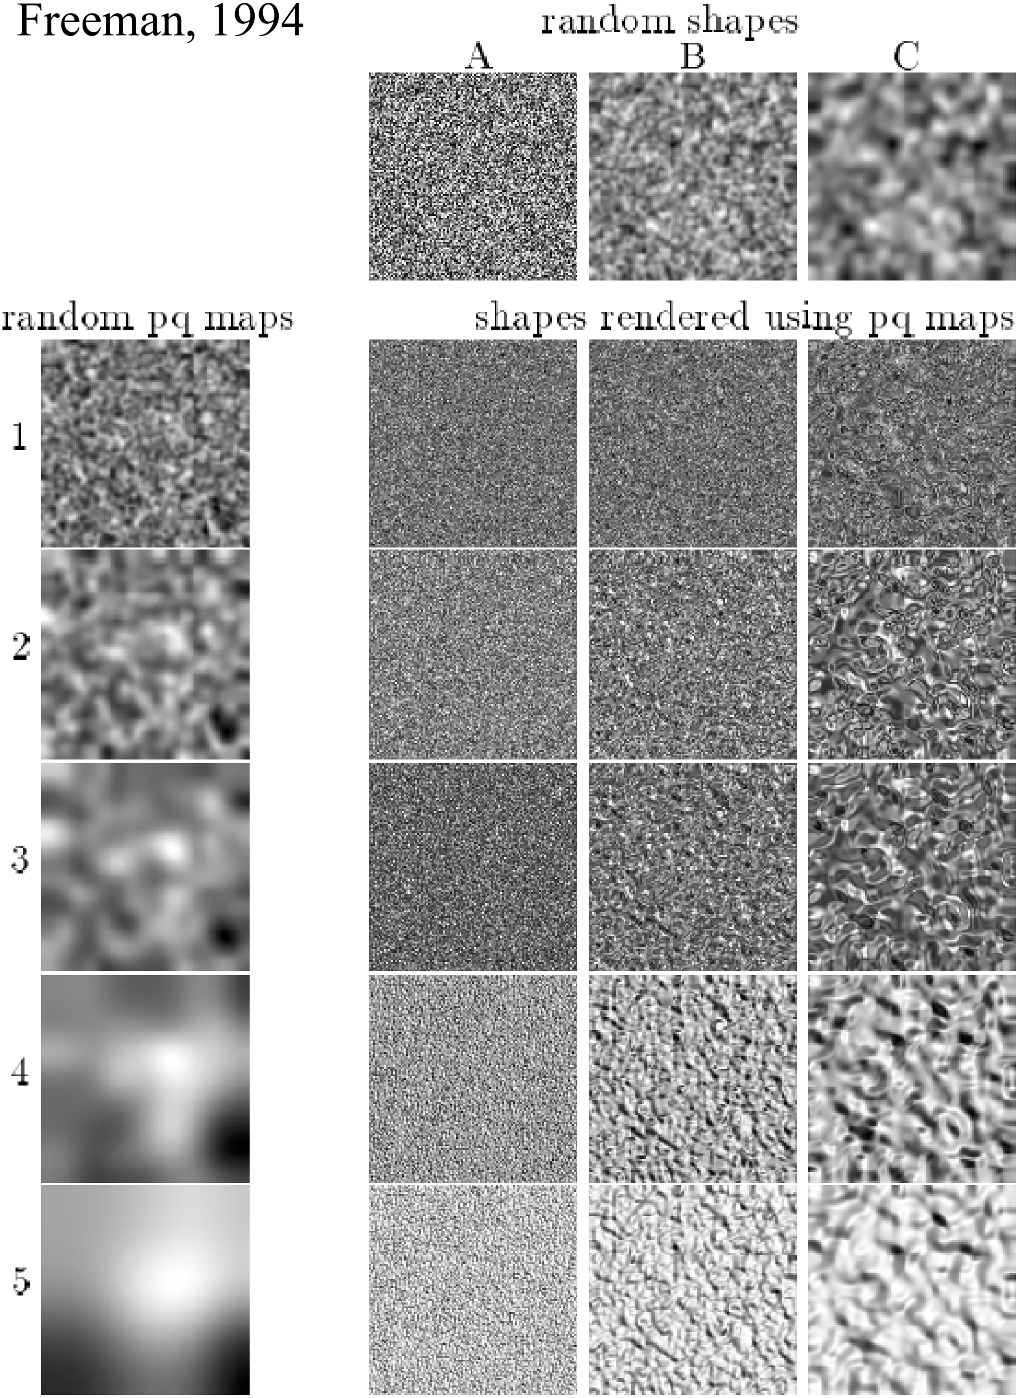

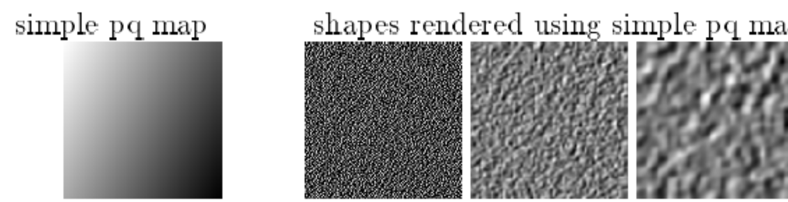

Figure 1: 3 range images (successively blurred versions of the same range image) are rendered with different pq reflectance maps. 1–5 are successively blurred versions of the same original random noise pq map. A Lambertian (linear-shading approximation valid) reflectance function and its rendered images are shown for comparison. Shapes are scaled up so that the range of slopes fits the pq reflectance map.

2

Freeman, 1994

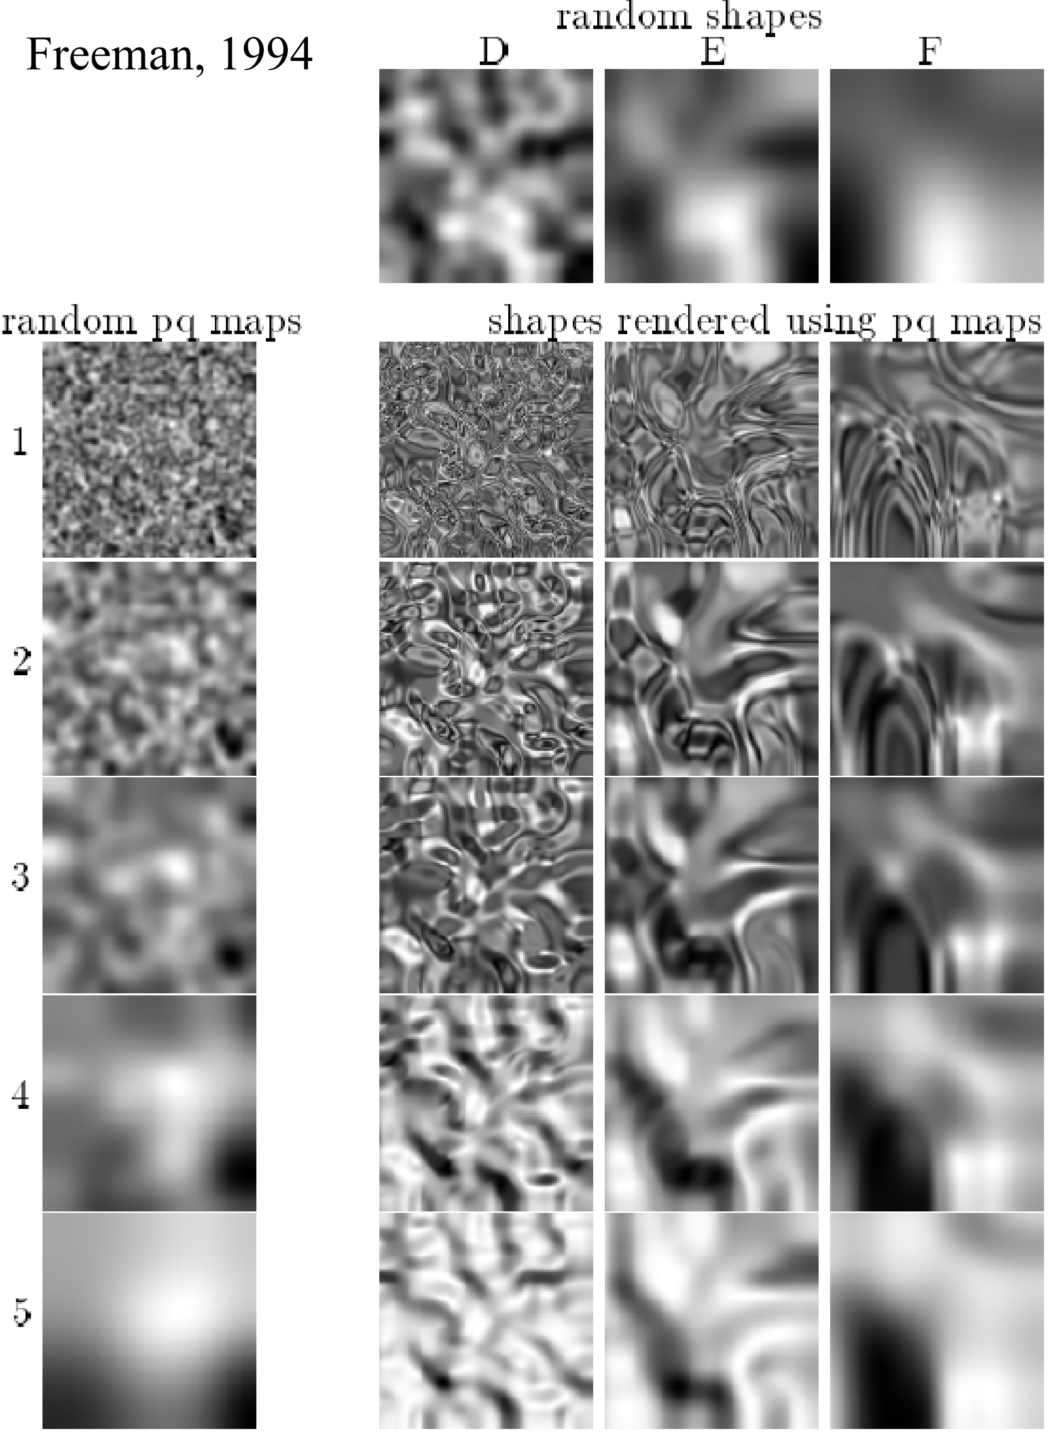

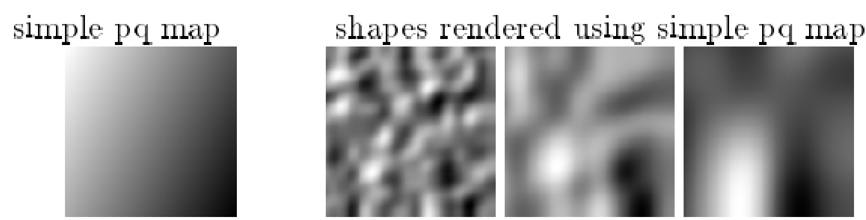

Figure 2: Same as Fig. 1, with still more blurring of the original random noise shape.

Freeman, 1994

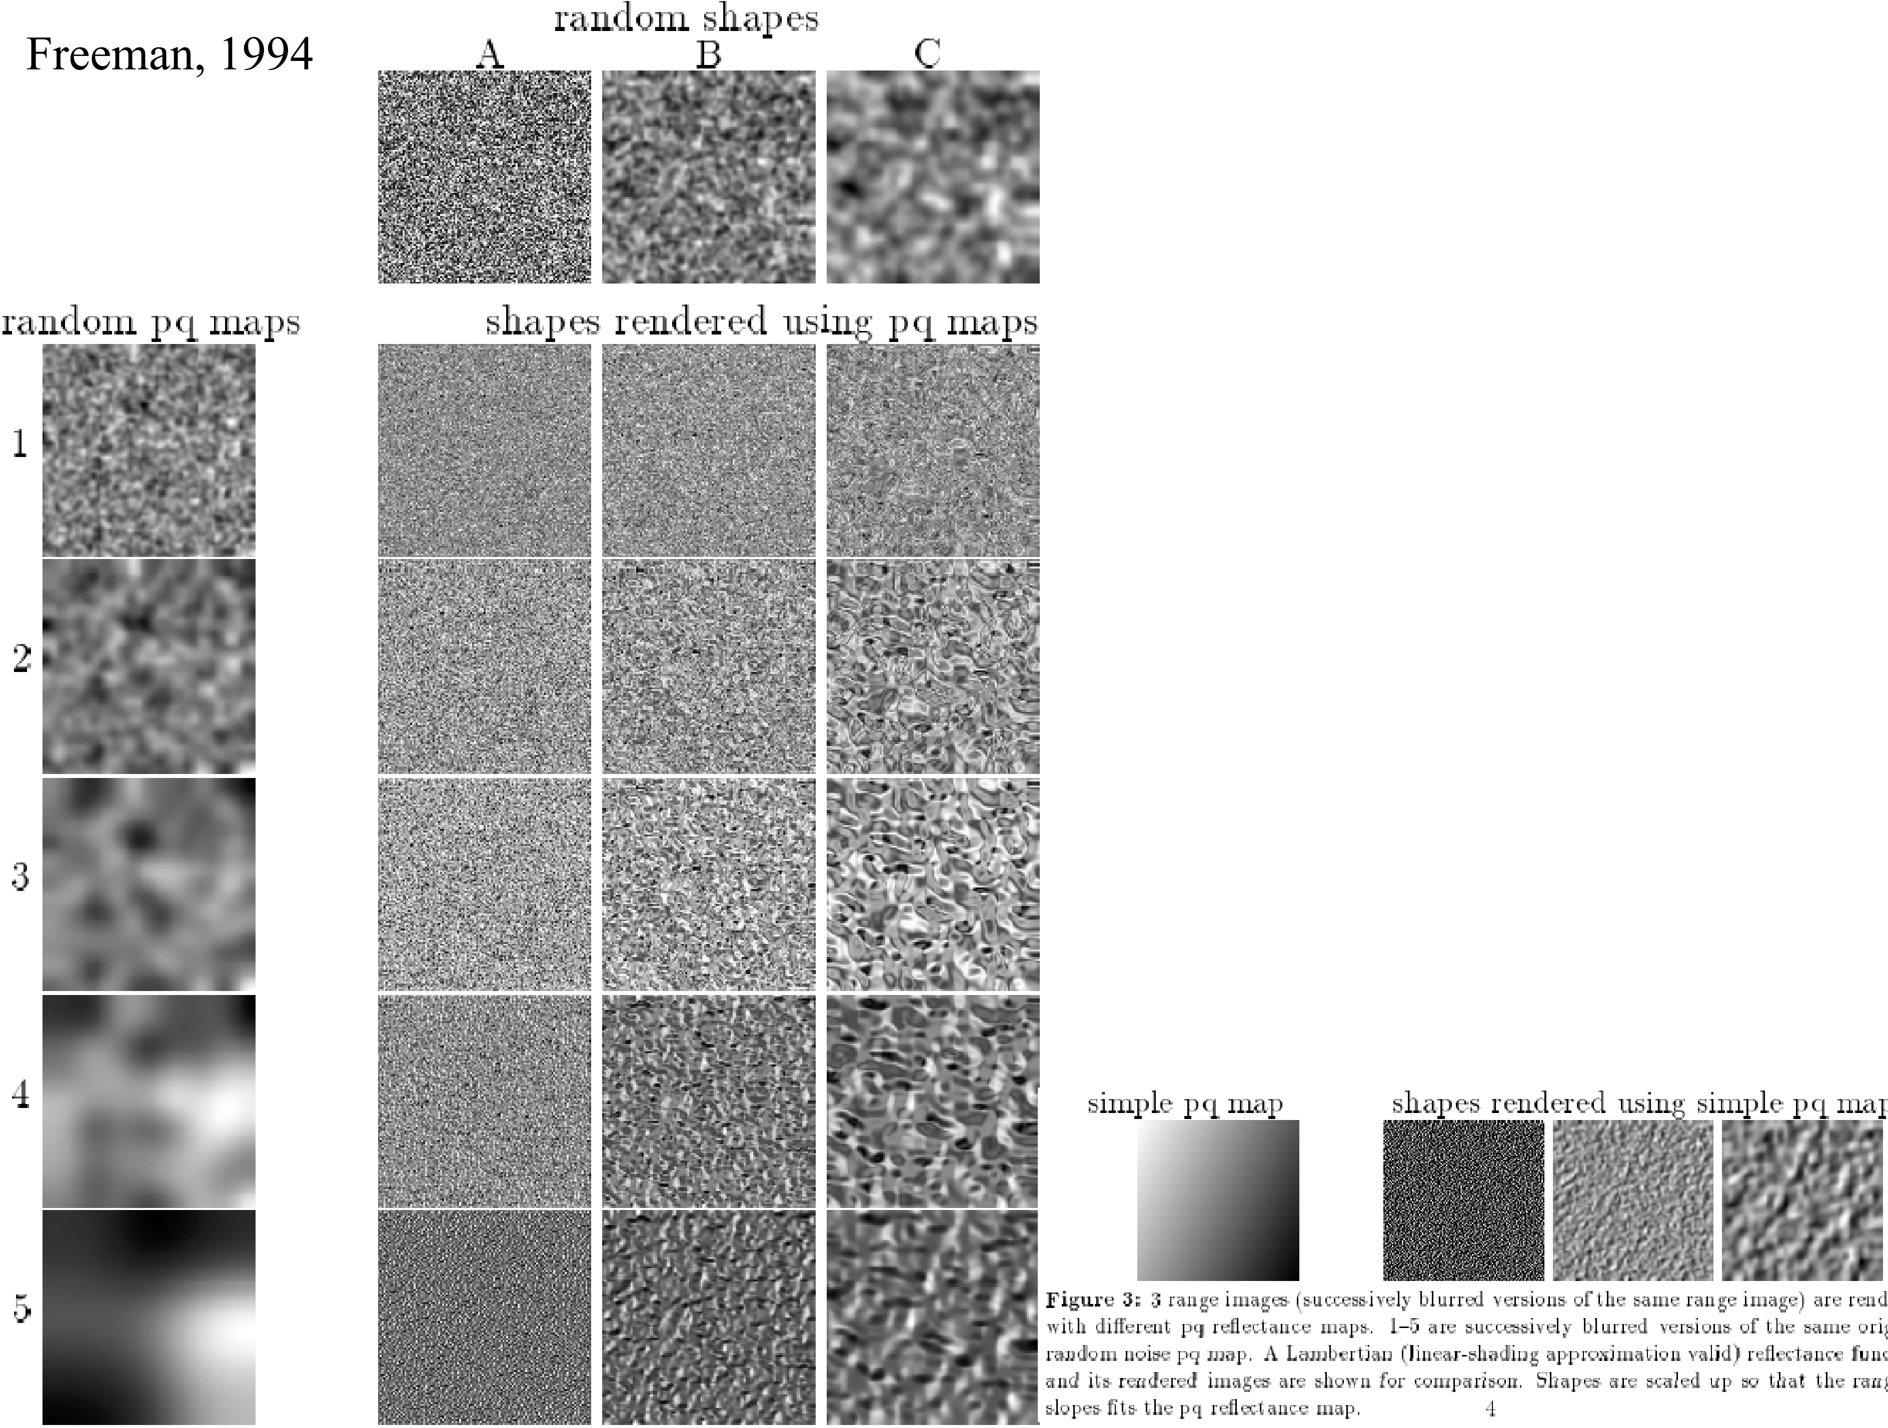

Figure 3: 3 range images (successively blurred versions of the same range image) are rendere with different pq reflectance maps. 1–5 are successively blurred versions of the same origin random noise pq map. A Lambertian (linear-shading approximation valid) reflectance functio and its rendered images are shown for comparison. Shapes are scaled up so that the range o slopes fits the pq reflectance map. 4

Freeman, 1994

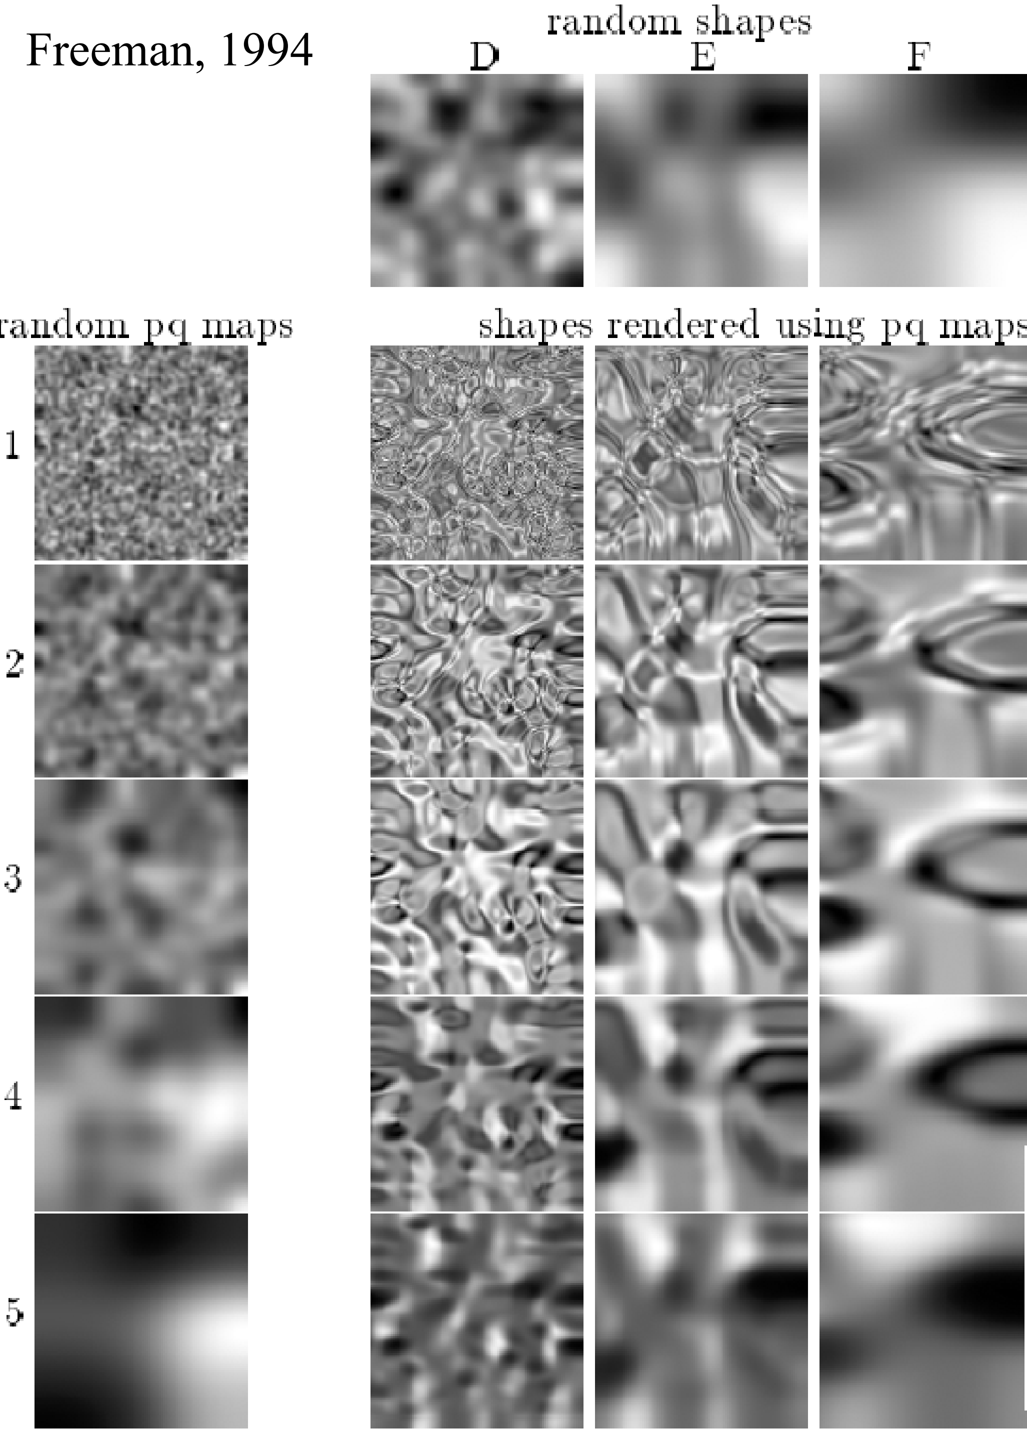

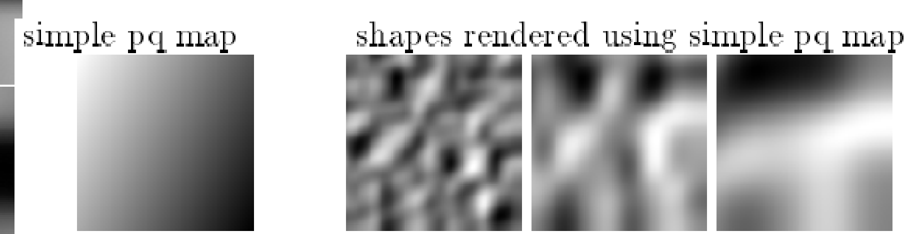

Figure 4: Same as Fig. 3, with still more blurring of the original random noise shape.

Let’s list the things this model doesn’t handle properly

- Occluding edges
- Albedo changes
- Perspective effects (small)
- Interreflections
- Material changes across surfaces in the image

Linear shading map

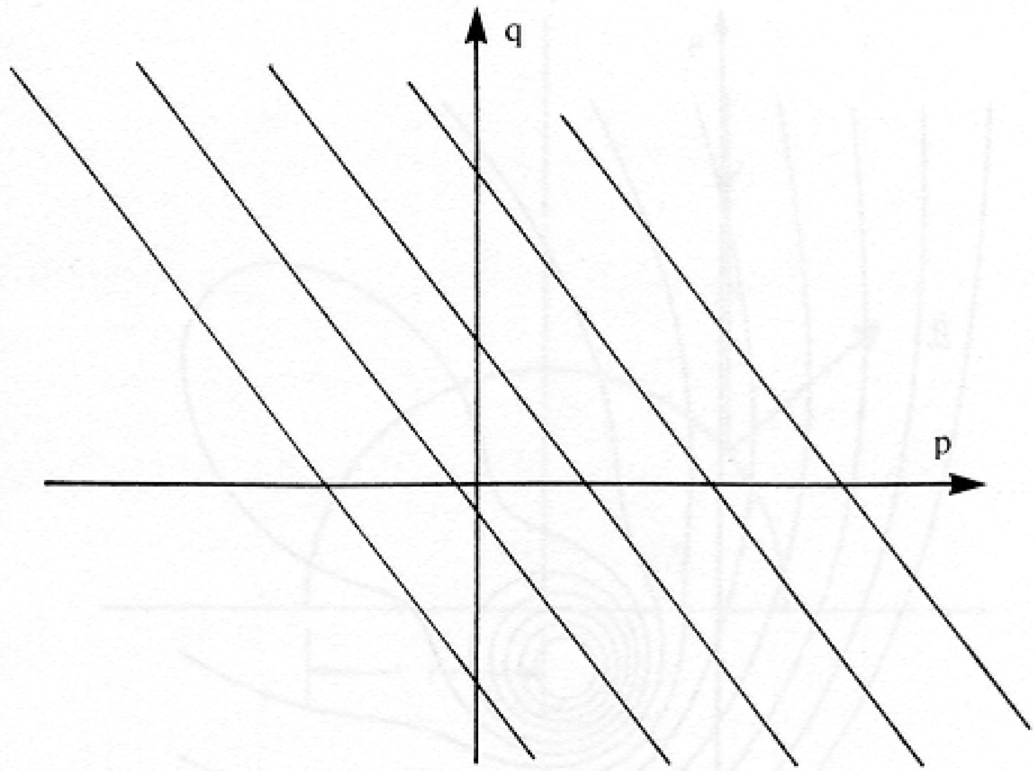

Figure 10-14. In the case of the material in the maria of the moon, the reflectance map can be closely approximated by a function of a linear combination of the components of the gradient. The contours of constant brightness are parallel straight lines in gradient space.

Horn, 1986

Linear shading: 1st order terms of Lambertian shading

Lambertian point source

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

1st order Taylor series about p=q=0

$$\approx k_2 + \left.\frac{\partial R(p,q)}{\partial p}\right|_{p=0,q=0} p + \left.\frac{\partial R(p,q)}{\partial q}\right|_{p=0,q=0} q$$

$$= k_2(1 + p_s p + q_s q)$$

See Pentland, IJCV vol. 1 no. 4, 1990.

Linear shading

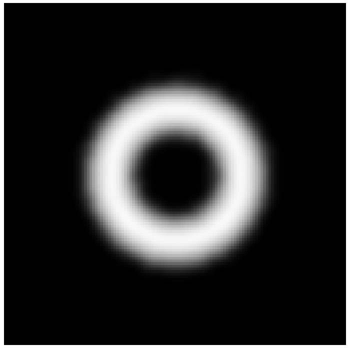

range image

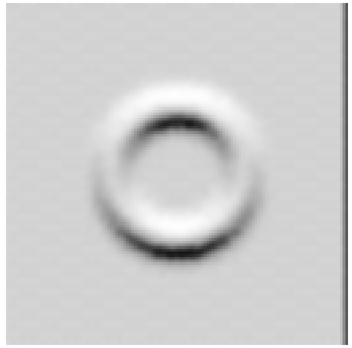

Lambertian shading

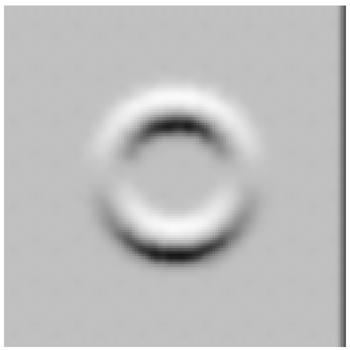

linear shading

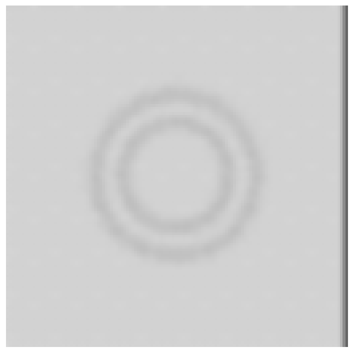

quadratic terms

higher-order terms

Pentland 1990, Adelson&Freeman, 1991

Advantages of linear shading

- Linear relationship between surface range map and rendered image.
- Rendering is easy: differentiate with respect to azimuthal light source direction.
- Applies: linear sources, or shallow illumination angles and Lambertian surface.
- Allows for very simple inverse transformation from rendered image to surface range map, which we'll discuss later with shape-from-shading material.

Knowing the reflectance map, can we infer the gradient at any point?

Generic reflectance map

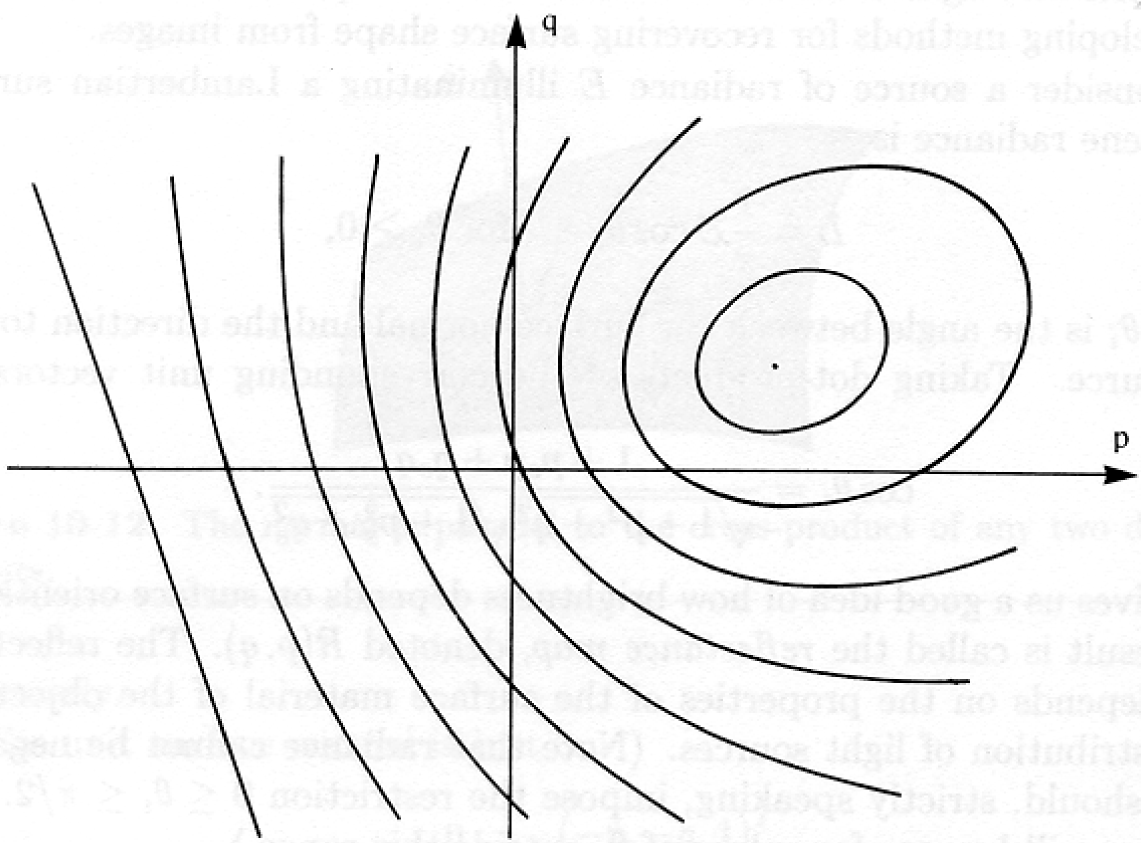

Figure 10-13. The reflectance map is a plot of brightness as a function of surface orientation. Here it is shown as a contour map in gradient space. In the case of a Lambertian surface under point-source illumination, the contours turn out to be nested conic sections. The maximum of $R(p, q)$ occurs 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.

Horn, 1986

Photometric stereo

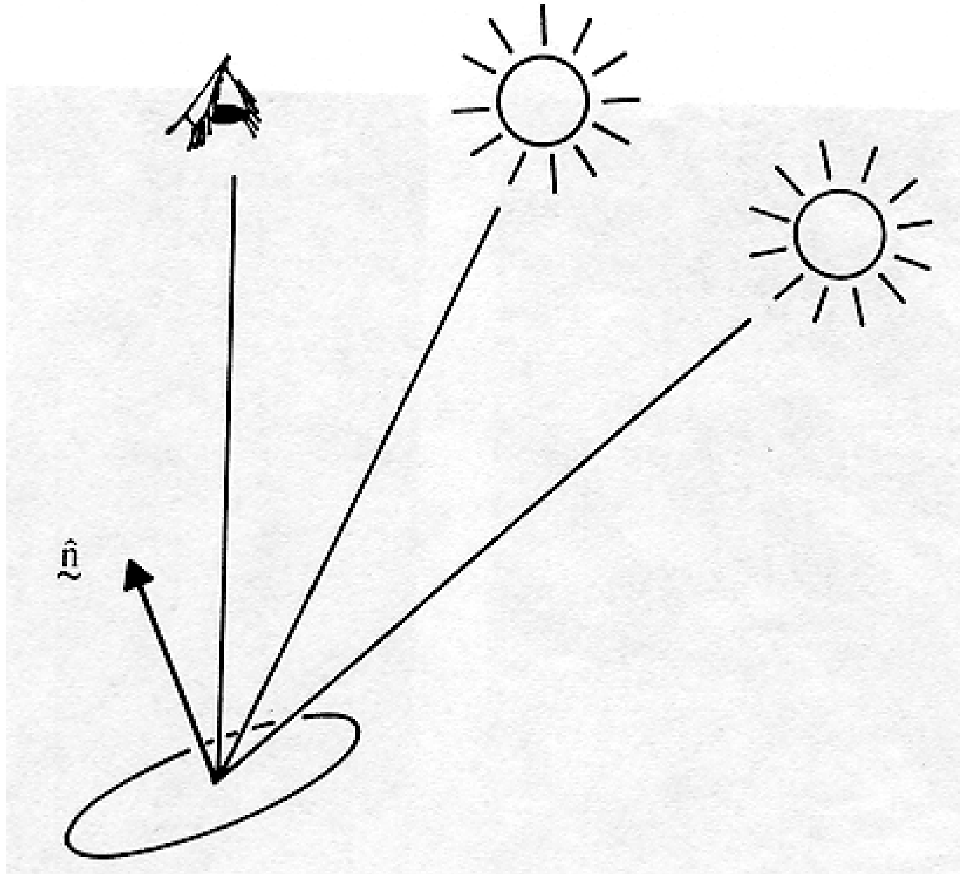

Horn, 1986

Fixed camera and object positions.
Take two or more images under different lighting conditions.

frame 2

frame 11

. . .

. . .

Approach 1

- Photograph the object with a calibration object in the picture, or available in another photograph.
- Use the multiple responses of the calibration object to the different light sources to form a look-up table.
- Index into that table using the multiple responses of the unknown object.
- Handles arbitrary BRDF.

Approach 2

- Assume a particular functional form for the BRDF (Lambertian). Assume known light source positions (point sources at infinity as specified locations).
- Analytically determine the surface slope for each location's collection of image intensities.

Photometric stereo

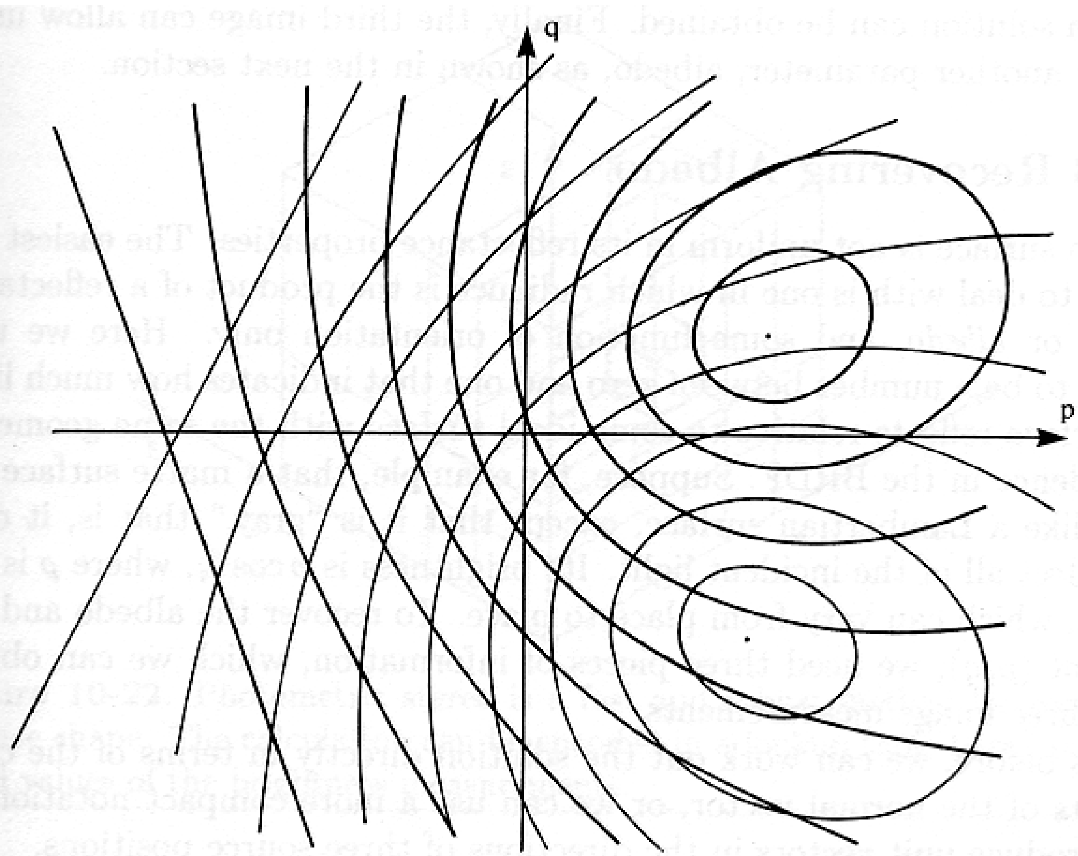

Figure 10-21. In the case of a Lambertian surface illuminated successively by two different point sources, there are at most two surface orientations that produce a particular pair of brightness values. These are found at the intersection of the corresponding contours in two superimposed reflectance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

See Forsyth&Ponce sect. 5.4 for procedure. In HW: don't need to integrate the surface normals to get the shape.

From the image under the ith lighting condition (Lambertian)

Pixel intensity at position x,y in ith image.

surface radiance

$$I_i(x,y) = kL_i(x,y)$$

$$= k\rho(x,y)\hat{N}(x,y)\cdot\hat{S}_i$$

surface albedo

surface normal

ith light source direction

Combining all the measurements

$$\begin{pmatrix} I_1(x,y) \\ I_2(x,y) \\ \vdots \\ I_n(x,y) \end{pmatrix} = \begin{pmatrix} \hat{S}_1^T \\ \hat{S}_2^T \\ \vdots \\ \hat{S}_n^T \end{pmatrix} \rho(x,y)\hat{N}(x,y)$$

Solve for g(x,y). May be ill-conditioned

$$\begin{pmatrix} I_1(x,y) \\ I_2(x,y) \\ \vdots \\ I_n(x,y) \end{pmatrix} = \begin{pmatrix} \hat{S}_1^T \\ \hat{S}_2^T \\ \vdots \\ \hat{S}_n^T \end{pmatrix} \vec{g}(x,y)$$

A fix to avoid problems in dark areas: pre-multiply both sides by the image intensities

$$\begin{pmatrix} I_1(x,y) & 0 & \cdots & 0 \\ 0 & I_2(x,y) & \cdots & 0 \\ \vdots & \vdots & \vdots & 0 \\ 0 & 0 & 0 & I_n(x,y) \end{pmatrix} \begin{pmatrix} I_1(x,y) \\ I_2(x,y) \\ \vdots \\ I_n(x,y) \end{pmatrix}$$

$$= \begin{pmatrix} I_1(x,y) & 0 & \cdots & 0 \\ 0 & I_2(x,y) & \cdots & 0 \\ \vdots & \vdots & \vdots & 0 \\ 0 & 0 & 0 & I_n(x,y) \end{pmatrix} \begin{pmatrix} \hat{S}_1^T \\ \hat{S}_2^T \\ \vdots \\ \hat{S}_n^T \end{pmatrix} \vec{g}(x,y)$$

Recovering albedo and surface normal

$$\rho(x, y) = \left|\vec{g}(x, y)\right|$$

$$\hat{N}(x, y) = \frac{\vec{g}(x, y)}{\left|\vec{g}(x, y)\right|}$$

Surface shape from surface gradients

- Can you do it?
- What are the ambiguities?
- What are the constraints?
- Method of Weiss, to be discussed in linear filtering section.
- So for your homework, we’ll leave the computation at the gradients.