Lecture 28:
3-D Reconstruction III:
Shape-from-Shading

Contents

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Introduction

The Shape-from-Shading Problem

- A 3-D surface is illuminated and photographed.
  Is is possible to recover the 3-D scene, if one knows the position of the light source and the reflectance properties of the surface?
- important e.g. in astronomy:
  A planet is too far away in order to use stereo reconstruction techniques.
- also a classical computer vision problem:
  Already in 1970, Berthold Horn wrote his Ph.D. thesis on this topic.
- Related problem:
  The same surface is observed under different illumination conditions.
  Is is possible to exploit this additional information in order to obtain a better reconstruction?

Image of the surface of the planet Mars. Does it really depict a face? Shape-from-shading has the goal to reconstruct the 3-D profile from such a single image. Source: NASA.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Illumination and Reflectance

Two factors determine the radiation properties of a surface:

- illumination direction:
  position of the surface relative to the light source
- reflectance properties of the surface:
  depend on the optical properties of the material

Illumination and Reflectance (2)

MIA
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Bidirectional Reflectance Distribution Function

- Consider a surface element that is illuminated by a single light source.
  Its direction is given by the polar coordinates (θ_i, ϕ_i); see next slide.
  The corresponding incoming energy is denoted by $E(\theta_i, \phi_i)$.
- The brightness that is emitted by the surface under the polar angles (θ_e, ϕ_e) is denoted by $L(\theta_e, \phi_e)$.
- The *bidirectional reflectance distribution function (BDRF, bidirektionale Reflexionsverteilung)* $f_r(\theta_i, \phi_i, \theta_e, \phi_e)$ decribes the ratio between the emitted brightness $L(\theta_e, \phi_e)$ and the incoming energy $E(\theta_i, \phi_i)$:

$$L(\theta_e, \phi_e) \;=\; f_r(\theta_i, \phi_i, \theta_e, \phi_e)\, E(\theta_i, \phi_i).$$

- Usually it only depends on the difference angles:

$$f_r(\theta_i, \phi_i, \theta_e, \phi_e) \;=\; f_r(\theta_i - \theta_e,\ \phi_i - \phi_e)$$

Illumination and Reflectance (3)

MIA
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

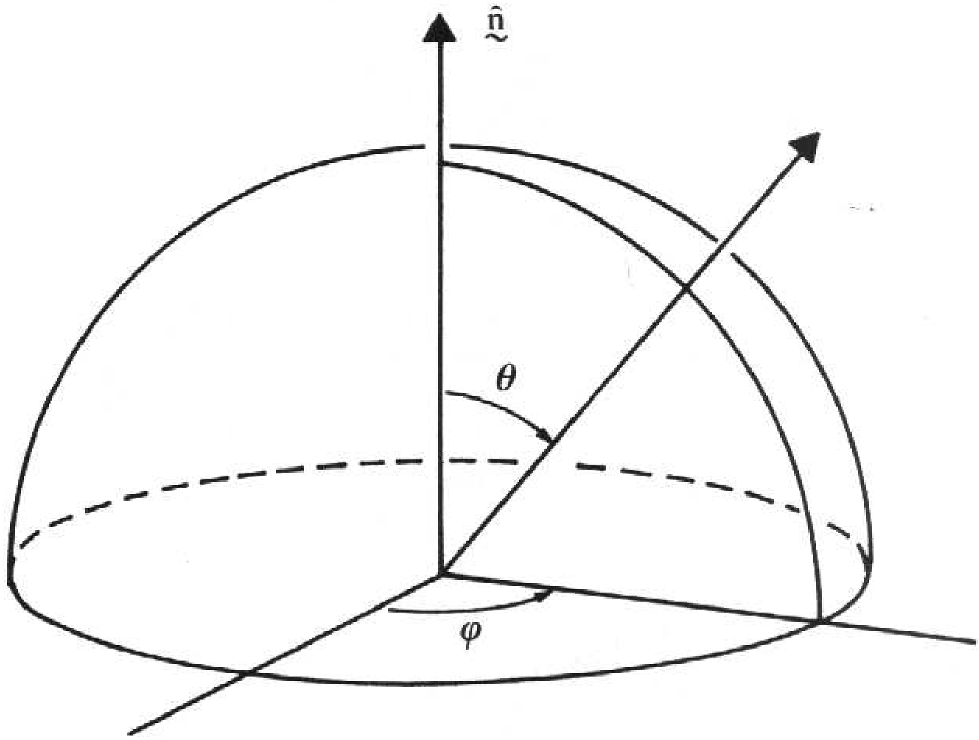

3-D polar coordinates. ϕ is called azimuth angle, and θ is the zenith angle. Author: B. K. P. Horn (1986).

Illumination and Reflectance (4)

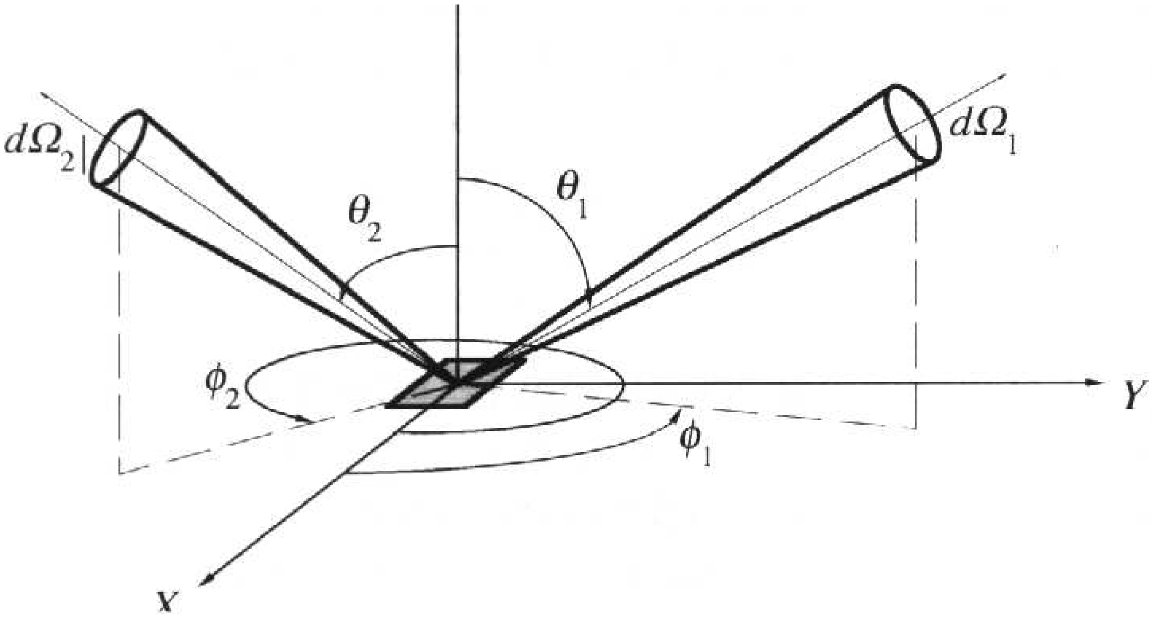

Angles of the BDRF. Authors: R. Klette, K, Schlüns, A. Koschan (1998).

Illumination and Reflectance (5)

Extremal Cases of Reflectance

- *Ideally Reflecting Surface:*
  - reflects the incoming ray in exactly one direction:
    An incoming ray with polar angles (θ_i, ϕ_i) is reflected in direction $(\theta_e, \phi_e) = (\theta_i, \phi_i + \pi)$.
  - requires a perfectly smooth surface
- *Lambertian Surface (Lambert'sche Oberfläche):*
  - reflects light energy in all directions *(diffuse reflection)*.
  - frequently used as a model for rough, unsmooth surfaces
  - Often one also requires that the Lambertian surface does not absorb energy.
  - However, we do permit absorption. The nonabsorbed (i.e. reflected) fraction is determined by the *albedo (reflectance coefficient, Albedofaktor)* $\rho \in [0, 1]$.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Basic Relations for Shape-from-Shading

- The reflected radiance R (abgestrahlter Anteil über alle Richtungen) of a Lambertian surface in a point (X_w, Y_w, Z_w) is proportional to the cosine of the angle between the surface normal vector $\mathbf{n}$ and the light source direction $\mathbf{s}$:

$$R_{\rho,\mathbf{s}}(X_w, Y_w, Z_w) \;=\; \rho\, \mathbf{s}^\top \mathbf{n}$$

  where the albedo ρ is the corresponding constant of this proportionality.

- Thus, R is largest, if the light source direction $\mathbf{s}$ is in the normal direction $\mathbf{n}$.
- There exist methods for estimating ρ and $\mathbf{s}$, but we do not consider them here.

M I A

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Assumptions

- The reflected radiance $R_{\rho,\mathbf{s}}$ in a scene point (X_w, Y_w, Z_w) can be identified with the grey value f in the corresponding image point (x, y):

$$f(x, y) \;=\; R_{\rho,\mathbf{s}}(X_w, Y_w, Z_w).$$

  (image irradiance equation)

- The depth Z_w can be described as a function $Z_w = Z_w(X_w, Y_w)$.
- The image plane is coplanar to the (X_w, Y_w) plane,
  and the observed surface is sufficiently far away.

Consequences

- We may approximate the projective pinhole camera model by a specific affine model, where the image coordinates (x, y) and the world coordinates (X_w, Y_w) are proportional *(orthographic projection)*.
  Thus we identify (x, y) and (X_w, Y_w).
- In this way the shape-from-shading problem comes down to the computation of the depth map $Z_w = Z_w(x, y)$.

Basic Relations for Shape-from-Shading (3)

Reconstruction of the Depth Map

- For recovering the depth map $Z_w(x, y)$, we first investigate the connection between the surface normal vector $\mathbf{n}$ and the partial derivatives $p := \partial_x Z_w$ and $q := \partial_y Z_w$.
- Differentiating the surface $(x, y, Z_w(x, y))^\top$ with respect to x and y creates two vectors in the tangential plane of $Z_w(x, y)$:

$$\begin{pmatrix} 1 \\ 0 \\ p \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 1 \\ q \end{pmatrix}$$

- The outer product (Vektorprodukt, Kreuzprodukt)

$$\begin{pmatrix} 1 \\ 0 \\ p \end{pmatrix} \times \begin{pmatrix} 0 \\ 1 \\ q \end{pmatrix} = \begin{pmatrix} -p \\ -q \\ 1 \end{pmatrix}$$

creates a vector that is orthogonal to both tangential vectors.

Basic Relations for Shape-from-Shading (4)

- Thus the unit normal vector is given by

$$\mathbf{n} = \frac{1}{\sqrt{1 + p^2 + q^2}} \begin{pmatrix} -p \\ -q \\ 1 \end{pmatrix},$$

- By means of $f = R_{\rho,\mathbf{s}} = \rho\,\mathbf{s}^\top\mathbf{n}$ we obtain

$$f(x, y) = \frac{\rho}{\sqrt{1 + p^2 + q^2}}\, \mathbf{s}^\top \begin{pmatrix} -p \\ -q \\ 1 \end{pmatrix}.$$

- similar situation as in the optic flow problem:
  - one equation with two unknowns (p and q)
  - For obtaining a unique solution, one needs additional constraints.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

A Variational Method by Ikeuchi and Horn (1981)

Assumption:

- p and q vary smoothly in space.

Energy Functional:

$$E(p,q) \;=\; \int_{\Omega} \Big(\big(f(x,y) - R(p,q)\big)^2 + \alpha \,\big(|\nabla p|^2 + |\nabla q|^2\big) \Big)\, dx\, dy.$$

We saw that in the case of a Lambertian surface, one has

$$R(p,q) \;=\; \frac{\rho}{\sqrt{1+p^2+q^2}}\, \mathbf{s}^\top \begin{pmatrix} -p \\ -q \\ 1 \end{pmatrix}.$$

Problem

- The data term does not guarantee convexity.
  The energy functional may have multiple local minima.
- For encouraging convergence to a global minimum, one can use similar methods as for stereo reconstruction (Lecture 27):
  - Smooth the data f on a coarse scale in order to reduce the number of local minima.
  - Reduce the scale gradually, while using the solution of the previous scale as intialisation for the iterations at the present scale.

Euler–Lagrange Equations

$$\Delta p - \tfrac{1}{\alpha}\,(f - R(p,q))\,\partial_p R(p,q) \;=\; 0,$$
$$\Delta q - \tfrac{1}{\alpha}\,(f - R(p,q))\,\partial_q R(p,q) \;=\; 0.$$

M I A 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

A Variational Method (3)

Discretisation of the Euler–Lagrange Equations

- usual approximation of the Laplacian (cf. Lecture 25, grid size 1):

$$\Delta p|_i \;\approx\; \sum_{j \in \mathcal{N}(i)} (p_j - p_i)$$

$\mathcal{N}(i)$: 4 neighbours of pixel i (boundary pixels: 3 neighbours; corner pixels: 2)

- creates the difference equations

$$\begin{aligned} 0 &= \sum_{j \in \mathcal{N}(i)} (p_j - p_i) - \tfrac{1}{\alpha} \left(f_i - R(p_i, q_i)\right) \partial_p R(p_i, q_i), \\ 0 &= \sum_{j \in \mathcal{N}(i)} (q_j - q_i) - \tfrac{1}{\alpha} \left(f_i - R(p_i, q_i)\right) \partial_q R(p_i, q_i) \end{aligned}$$

for all pixels $(i = 1, ..., N)$

A Variational Method (4)

- simple iterative scheme (Jacobi-like):

$$\begin{aligned} p_i^{(k+1)} &= \frac{\sum\limits_{j \in \mathcal{N}(i)} p_j^{(k)} - \frac{1}{\alpha} \left(f_i - R(p_i^{(k)}, q_i^{(k)})\right) \partial_p R(p_i^{(k)}, q_i^{(k)})}{|\mathcal{N}(i)|}, \\ q_i^{(k+1)} &= \frac{\sum\limits_{j \in \mathcal{N}(i)} q_j^{(k)} - \frac{1}{\alpha} \left(f_i - R(p_i^{(k)}, q_i^{(k)})\right) \partial_q R(p_i^{(k)}, q_i^{(k)})}{|\mathcal{N}(i)|}. \end{aligned}$$

A Variational Method (5)

Removal of Inconsistencies

- *Problem:*
  $p = \partial_x Z_w$ and $q = \partial_y Z_w$ are not independent. Since this dependency has not been taken into account in the numerical scheme, one cannot expect to find a function $Z(x, y)$ with $\partial_x Z_w = p$ and $\partial_y Z_w = q$.

- *Remedy 1:*
  Supplement the energy functional with an additional consistency term $(p_y - q_x)^2$.
  This complicates the situation.

- *Remedy 2:*
  After each iteration, project the results $p^{(k)}$ and $q^{(k)}$ onto the closest integrable pair of functions.
  There is no guarantee that this is stable.

- *Remedy 3:*
  Use a variational model for Z_w instead of its derivatives p and q.
  This creates a more complicated data term.
  It may also be advisable to have higher order smoothness constraints for Z_w then, e.g. by penalising the Frobenius norm of the Hessian of Z_w.

A Variational Method (6)

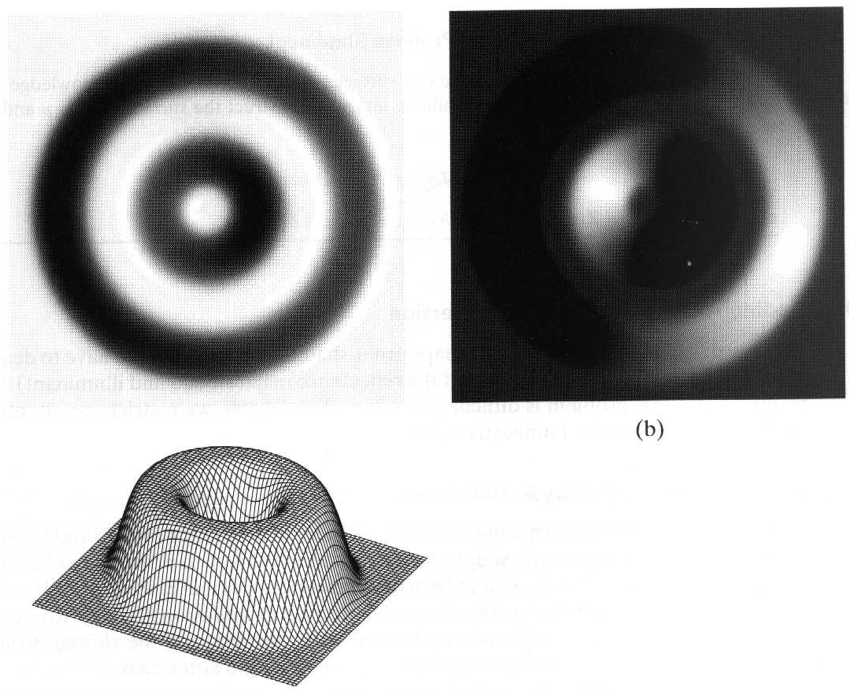

(b)

(a)

Top: Two images of the same Lambertian surface, but illuminated from different directions. **Bottom:** 3-D model. Authors: E. Trucco, A. Verri (1998).

MIA 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

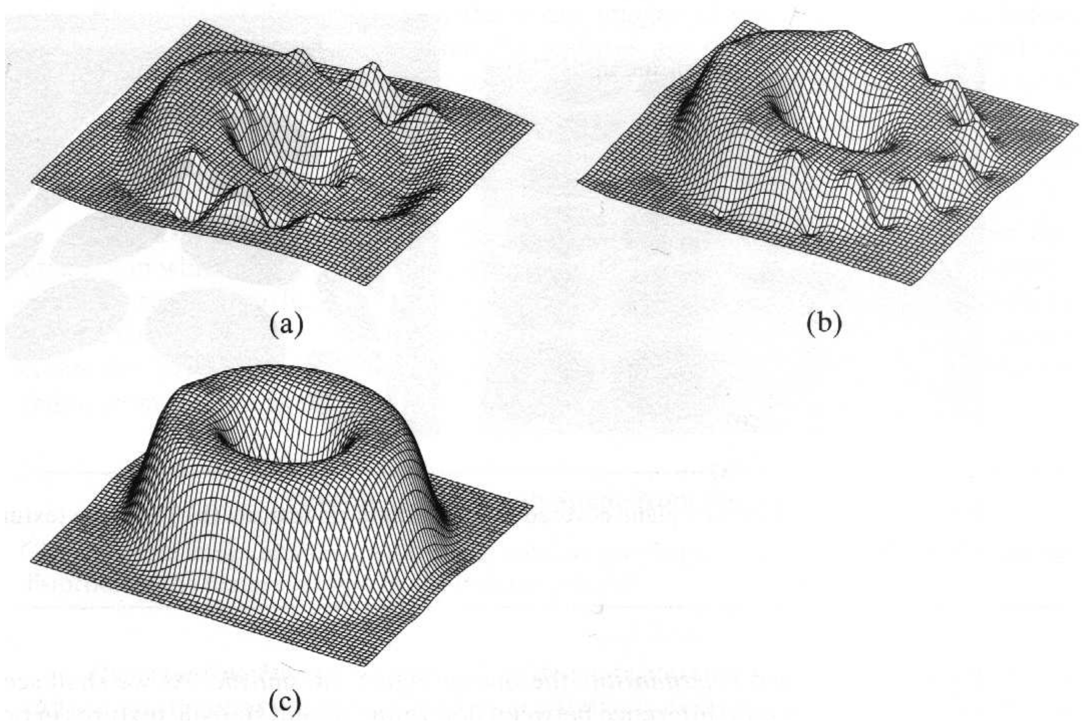

Reconstructed surface after (a) 100 iterations, (b) 1000 iterations, (c) 2000 iterationen. The asymmetry in (a) und (b) has been caused by the illumination direction. Authors: E. Trucco, A. Verri (1998).

Photometric Stereo Reconstruction

- The same object is photographed several times under different illumination conditions.
- Advantage compared to usual stereo reconstruction:
  There is no correspondence problem between the images.
- Advantage compared to usual shape-from-shading:
  The method is more robust and precise, since errors are averaged.

Energy Functional for m Images $f_1,\dots,f_m$:

$$E(p,q) \;=\; \int_{\Omega} \left(\sum_{l=1}^{m} \Big(f_l(x,y) - R_l(p,q) \Big)^2 + \alpha \,\left(|\nabla p|^2 + |\nabla q|^2 \right) \right) \, dx\, dy.$$

If the accuracy of the measurements differs, one can assign different weights to the individual data terms.

Photometric Stereo Reconstruction (2)

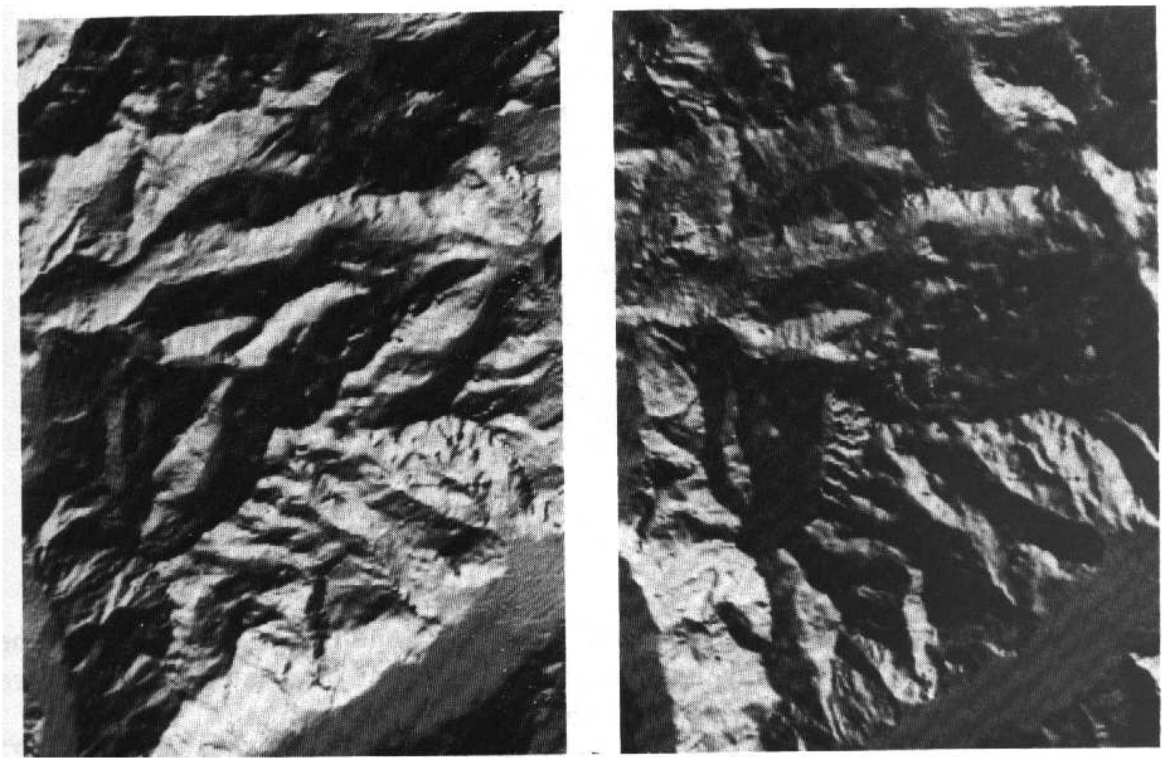

Two photos of the same mountains, but with different illumination. Author: B. K. P. Horn (1986).

Photometric Stereo Reconstruction (3)

MIA 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

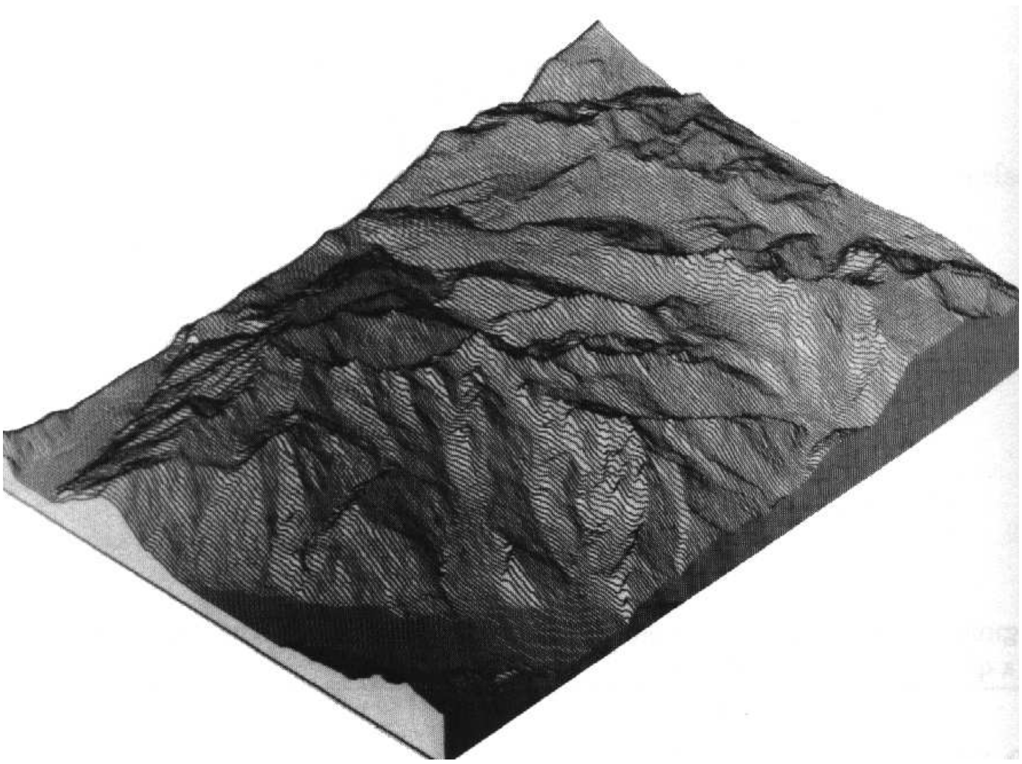

Digital terrain model reconstructed from the preceding two images. Author: B. K. P. Horn (1986).

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Other Shape-from-X Methods

- *Shape-from-Motion, Structure-from-Motion*
  Disparities between subsequent frames are used in a similar way as for stereo reconstruction in order to estimate the depth map. Important for robotics.

- *Shape-from-Focus*
  Depth information is recovered from varying the focus setting of the lens system.

- *Shape-from-Texture*
  Density, size and orientation of the texture are used for 3-D reconstruction.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

A nice example for shape-from-texture: The logo of the theorem prover SPASS. Source: `http://spass.mpi-sb.mpg.de/graphics/bigspasslogo.jpg`.

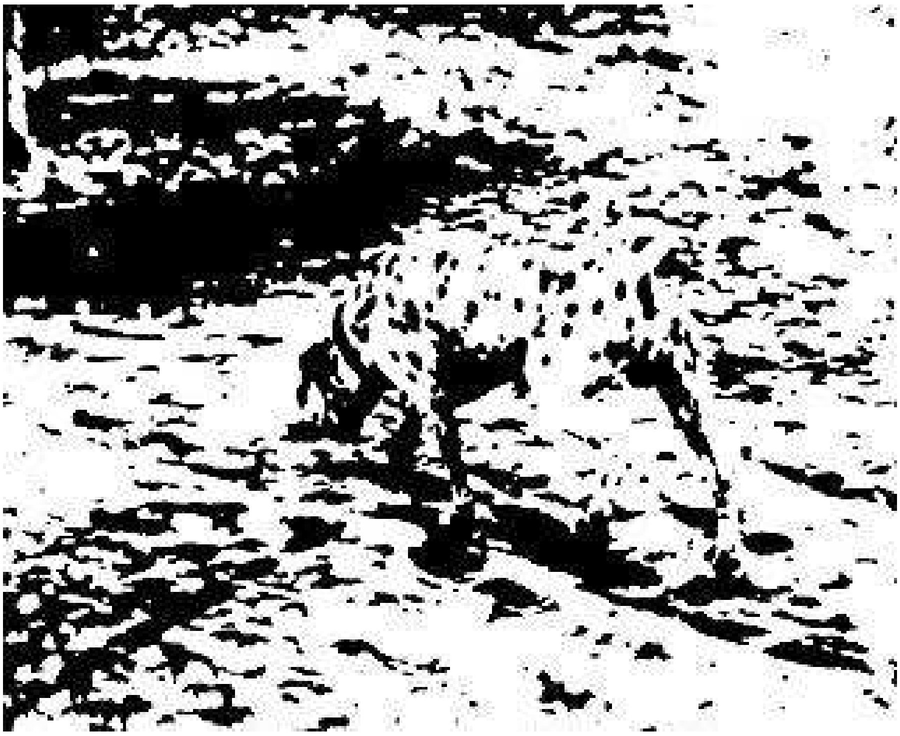

The depicted Dalmatian dog shows the difficulty of the shape-from-texture problem. Source: www.unesco.org/courier/1999_06/photoshr/41_2.htm. Based on a photo by Ronald C. James (1966).

Summary

- A Lambertian surface emits radiation in all directions.
- It leads to an analytical description between the observed grey value und the surface orientation.
  However, it is not sufficient for recovering the full depth information.
- To this end, additional smoothness assumptions are introduced within a variational framework.
- The energy functional may have local minima, and there may be inconsistencies between the different partial derivatives of the depth map.
- For photometric stereo reconstructions one photographs the same object under different illuminations.

Literature

- E. Trucco, A. Verri: *Introductory Techniques for 3-D Computer Vision*. Prentice–Hall, Upper Saddle River, 1998.
  (Chapter 9)
- K. Ikeuchi, B. K. P. Horn: Numerical shape from shading and occluding boundaries. *Artifical Intelligence*, Vol. 17, 141–184, 1981.
  (introduced the variational approach from this lecture)
- B. K. P. Horn, M. J. Brooks: The variational approach to shape from shading. *Computer Vision, Graphics and Image Processing*, Vol. 33, 174–208, 1986.
  (nice survey article)
- B. K. P. Horn: *Robot Vision*. MIT Press, Cambridge, MA, 1986.
  (Chapters 10 and 11)
- R. Klette, K. Schlüns, A. Koschan: *Computer Vision: Three-Dimensional Data from Images*. Springer, Singapore, 1998.
  (Chapters 6–8)
- R. Jain, R. Kasturi, B. G. Schunck: *Machine Vision*. McGraw-Hill, New York, 1995.
  (Chapter 9)
- B. K. P. Horn, M. J. Brooks (Eds.): *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
  (entire book on shape-from-shading)

Assignment C7 – Classroom Work

Problem 1 (Projection Matrices)

Compute a transformation matrix in homogeneous coordinates which describes a rotation around the x-axis through an angle of 30 degrees followed by a rotation around the z-axis through 60 degrees and a translation by the vector $(1, 2, -3)^\top$.

(*This problem shows you how to set up transformation matrices for rigid body motion.*)

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29

Assignment STP7 – Self Test Problem

From today on you can download an IPCV self test problem sheet from our webpage. It

- contains 6 problems covering different topics
- is similar in style and difficulty to a 180-minute written exam
- is neither submitted to nor corrected by the tutors

Use these problems to test yourself. Opportunity to ask questions related to the problems of this test is given in the final theoretical classroom exercise on

Tuesday, February 19 / Wednesday, February 20, 2008.