Height data from gradient fields

Reinhard Klette[+] **and Karsten Schlüns**[++]

[+] Auckland University, Computer Science Department, Tamaki Campus
Private Bag 92019, NZ-1005 Auckland, New Zealand
[++] Berlin Technical University, Computer Science Department, FR 3-11
D-10587 Berlin, Germany

ABSTRACT

The paper starts with a review of integration techniques for calculating height maps from dense gradient fields. There exist a few proposals of local integration methods (Coleman/Jain 1982, Healey/Jain 1984, Wu/Li 1988, Rodehorst 1993), and two proposals for global optimization (Horn/Brooks 1986 and Frankot/ Chellappa 1988). Several experimental evaluations of such integration techniques are discussed in this paper. The examined algorithms received high markings on curved objects but low markings on polyhedral shapes. Locally adaptive approaches are suggested to cope with complex shapes in general.

Keywords: surface reconstruction, shape from shading, photometric stereo, evaluation of algorithms, computer vision, gradient fields, potentials

1 INTRODUCTION

Several methods in the field of shape reconstruction [5,7,9] (most shading based methods) lead to gradient data that still have to be transformed into (scaled) height or depth maps, or into surface data for many applications. Thus the reconstruction accuracy also depends upon the performance of such a transformation module. Surprisingly, not much work has been done so far in this area. This paper starts with a review of the state of the art and discusses two approaches in detail. Several experimental evaluations of both methods for transforming gradient data into height data are reported. The studied (synthetic and real) object classes are curved and polyhedral objects. General qualitative evaluations of the compared transformation procedures are possible in relation to these object classes and in relation to different types of noise simulated for synthetic objects.

The outline of the paper is as follows: Fundamental definitions are given in this introductory Section. Section 2 contains a description of different methods for obtaining height from gradients. Methodologies for comparing selected algorithms are discussed in Section 3. A concluding discussion of evaluation results is given in Section 4.

1.1 Vector fields and potentials

We consider the reconstruction of a surface function $Z(x,y)$, defined in Cartesian xyz-space $\Re^3$, where only gradient field information $\mathbf{K}(\mathbf{p}): \Re^2 \to \Re^2$,

$$\mathbf{K}(\mathbf{p}) = grad\,(Z(x,y)) = \Big(p(x,y),q(x,y)\Big) = \left(\frac{\partial Z(x,y)}{\partial x}, \frac{\partial Z(x,y)}{\partial y}\right) \tag{1}$$

about this surface function is available as input data. The function Z is the *potential* of the vector field $\mathbf{K}$ in terms of vector analysis [1]. The gradient field information is given at discrete points $\mathbf{p} = (x,y)$ in com-

puter vision. The task consists in inverting differentiation, i.e., a certain anti-derivative has to be calculated by (discrete) integration up to a certain additive constant.

A vector field $\mathbf{K}$ that has a potential, i.e., that is integrable, is called *conservative*. For conservative vector fields any curve integral on a closed curve is equal to zero. This property means that we are not restricted to specific integration paths if conservative vector fields may be assumed.

Furthermore, if both functions K_x, K_y of a vector field $\mathbf{K}(\mathbf{p}) = (K_x, K_y)$ have continuous partial derivatives, then $\mathbf{K}$ is conservative if and only if the equation

$$\frac{\partial K_x}{\partial y} = \frac{\partial K_y}{\partial x} \tag{2}$$

holds, i.e., Z satisfies the integrability condition

$$\frac{\partial^2 Z}{\partial x \partial y} = \frac{\partial^2 Z}{\partial y \partial x}, \text{ or } p_y = q_x \text{ using the } pq\text{-notation.} \tag{3}$$

This property can be used for constraining a solution process if conservative vector fields may be assumed.

1.2 Classification of techniques

Essentially there are two types of known approaches in literature, local integration along paths [2,4,8,10] or global techniques [3,5,6].

Path integration techniques are directed on local calculations of height increments by curve integrals. Specific curves, i.e., scan lines are suggested by such techniques as well as special procedures for local approximations of height increments. These techniques are easy to implement and very efficient in computing speed. However, the locality of calculations causes a high dependency on data accuracy, and the propagation of height increments along paths means also propagation of errors. The known local integration techniques are not based on any assumption about the integrability condition.

Surface integration is treated as an optimization problem in global techniques. Surface integration can be considered to be a variational problem where a certain functional has to be minimized [5,6]. In general this does not lead to a global analytical solution because local iterative schemes are used for calculating discrete height data. Also a certain representation of the unknown surface, e.g. the Fourier base functions, and the integrability condition (3) can be used to constrain the global optimization process [3]. This method defines global analytical functions as final solution. It can be expected that both global techniques should be more robust against noise in comparison to local path integration because the surface gradient data have global impact on the solution process.

1.3 Evaluation methodology

A local technique [8] and a global technique [3] were selected for this paper as representatives for both classes of approaches. These two techniques were chosen based on theoretical comparisons and on the availability of complete documentations for implementation.

Discrete 256×256 gradient fields resulting from shading based surface analysis of real objects as well as synthetic 256×256 gradient fields are used as input data for surface integration. If a gradient value is not defined then value zero is assumed. Height maps are initialized with a constant at all 256×256 positions.

Curved and polyhedral object surfaces are studied. A synthetic sphere, a synthetic Mozart statue, and a plaster statue ("man in chair") were selected as curved objects. The sphere is very useful because visible orientations cover a Gaussian hemisphere, and because it allows evaluations of reconstructions based on

noisy gradient data. The Mozart statue is of "medium shape complexity", and the plaster statue is of "high shape complexity".

A synthetic K-shaped polyhedron and a milk box (TetraPak) were selected as polyhedral objects. Ideal polyhedrons do not satisfy surface conditions for integration at certain surface locations due to existing surface edges or occluding edges. However, these objects are especially of interest for industrial shape reconstruction.

Gradient fields of synthetic objects are generated either by analytic differentiation (sphere) or by approximate numerical differentiation if the surface data are given as height values on a discrete array (Mozart, K-polyhedron). In the latter case three different methodologies were applied: backward difference quotient, central difference quotient, and a regularized numerical differentiation by convolution with the derivatives of a Gaussian kernel with standard deviation $\sigma = \sqrt{2}$. The gradient fields of real objects (plaster statue, TetraPak) were obtained using photometric stereo analysis [7,9].

Different types of noise can be simulated for synthetic objects for studying the robustness of the implemented techniques. Gaussian noise with different standard deviations can be used for modeling image acquisition errors. In this short note, only outliers in gradient fields are modeled by spike noise.

Quantitative and qualitative evaluations are used for comparing the two selected integration techniques. The absolute differences in height between synthetic model (ground truth) and reconstructed surface define quantitative errors. These height differences are statistically characterized by histograms, maximum differences, averaged differences and standard deviations. Synthetic objects (i.e., given ground truth) allow the measurement of height differences. For obtaining discrete histogram data a certain interval width ΔZ is assumed for counting real height data Z. Percentages are calculated with respect to the total height of the object (i.e., maximum height minus minimum height).

Integration techniques lead to relative height values. The differences between maximum and minimum height could be used for adjusting model and reconstructed data into the same range of height values. However, this approach would be very sensitive to noise. Therefore least-square optimization was used for shifting (not scaling!) of height values of calculated height maps into the proper range of height values.

2 INTEGRATION TECHNIQUES

Local techniques are based on curve integrals and differ in specifying an integration path and a local neighborhood. Coleman and Jain [2] start in the middle of the gradient field and their initial path forms a cross in the array. Then the integration is performed in all four quadrants in column direction. For two points in sequence the averaged surface normal is calculated defining a surface tangent from the previous point to the next location. Their technique is known as *two-point method* in the computer vision literature. Healey and Jain [4] have extended this to an *eight-point method*. Wu and Li [10] also suggested paths parallel to the *x*-axis or the *y*-axis, but averaging gradient values for obtaining increments in height. Rodehorst [8] used four different scans through the gradient field starting at the four corners. Averaging the results of all four scans for obtaining the final height data is based on the assumption that the same initial height value is adequate in all four corners of the gradient data array. In each scan the local calculation of increments is based on a 2×2 neighborhood. This multiple-scan method was selected to represent local techniques in the comparison with global techniques.

Fig. 1 shows the integration paths (four scans) as used by Rodehorst [8]. During each scan a height map is reconstructed. The same initial height value is assumed in each start point **s**. The final result is defined by the arithmetic average of the four intermediate height maps.

The local approximation scheme is explained for a scan starting in the uppermost row from left to right, row by row to the bottom. The corresponding neighborhood scheme is also shown in Fig. 1. The

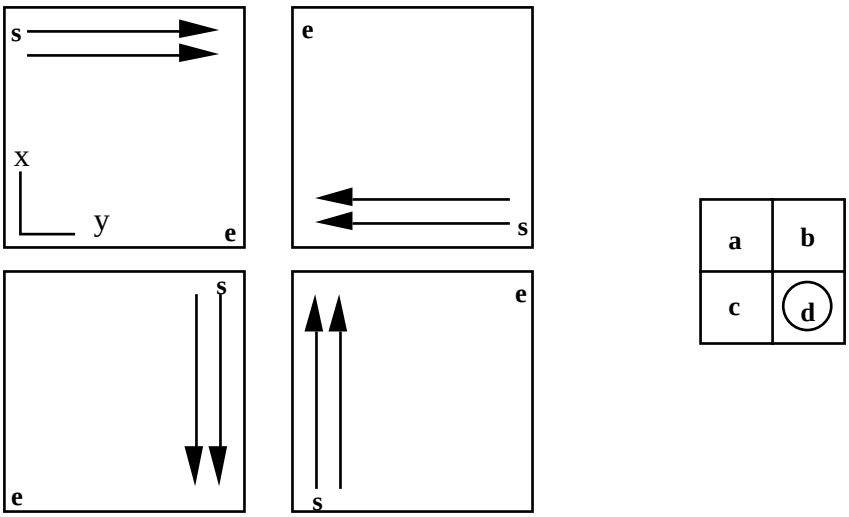

Figure 1: Four scans are performed, each one from start point **s** to end point **e**. The neighborhood scheme corresponds to the scan above on the left.

new height value Z has to be calculated at position **d**. The height values at positions **b** and **c** are available. The formula

$$Z(\mathbf{d}) = \frac{Z(\mathbf{b}) + Z(\mathbf{c})}{2} + \frac{\Delta x + \Delta y}{2} \tag{4}$$

is applied where the height increments Δx and Δy in x- and in y-direction are calculated based on the gradient data at positions **a**, **b**, **c**, and **d** as follows: Let $\mathbf{n}_i^\circ = (n_{i,x}, n_{i,y}, n_{i,z})$ be the unit normal at position $i = \mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$, where $n_{i,z} < 0$ is assumed to specify one of the two possible normal vectors, i.e., the following holds

$$\mathbf{n}^\circ = \frac{(p, q, -1)}{\sqrt{p^2 + q^2 + 1}} \tag{5}$$

for the different locations. Let

$$\overline{\mathbf{n}}^\circ = (\overline{n}_x, \overline{n}_y, -\sqrt{1 - \overline{n}_x^2 - \overline{n}_y^2}) \tag{6}$$

with

$$\overline{n}_x = \frac{n_{\mathbf{a},x} + n_{\mathbf{b},x} + n_{\mathbf{c},x} + n_{\mathbf{d},x}}{4}, \quad \overline{n}_y = \frac{n_{\mathbf{a},y} + n_{\mathbf{b},y} + n_{\mathbf{c},y} + n_{\mathbf{d},y}}{4} \quad . \tag{7}$$

Then the values

$$\Delta x = \frac{-\overline{n}_x}{\overline{n}_z} \quad \text{and} \quad \Delta y = \frac{\overline{n}_y}{\overline{n}_z} \tag{8}$$

are used in formula (4). Four simple scans through the gradient field are sufficient to calculate the resulting height map.

Two different global techniques are known in the computer vision literature. The method suggested by Horn and Brooks [6] is directed on minimizing the error functional

for $0 \le x,y \le N-1$ **do**
  if ($|p(x,y)| < max\_pq$ and $|q(x,y)| < max\_pq$) **then begin**
    $P1(x,y) := p(x,y)$; $P2(x,y) := 0$;
    $Q1(x,y) := q(x,y)$; $Q2(x,y) := 0$;
  end {then}
  else begin
    $P1(x,y) := 0$; $P2(x,y) := 0$;
    $Q1(x,y) := 0$; $Q2(x,y) := 0$;
  end {else};
$FFT(P1, P2, N \times N, \text{forward})$; $FFT(Q1, Q2, N \times N, \text{forward})$;
for ($0 \le u,v \le N-1$ and $(u,v) \ne (0,0)$) **do begin**
  $H1(u,v) := \left(u \cdot P2(u,v) + v \cdot Q2(u,v)\right) \Big/ \left(u^2 + v^2\right)$;
  $H2(u,v) := \left(-u \cdot P1(u,v) - v \cdot Q1(u,v)\right) \Big/ \left(u^2 + v^2\right)$
end {for};
$H1(0,0) :=$ something; $H2(0,0) := 0$;
$FFT(H1, H2, N \times N, \text{backward})$;
for $0 \le x,y \le N-1$ **do**
  $Z(x,y) := H1(x,y) \cdot ScalingFactor$

Figure 2: Height map generation based on Fourier expansion of the gradient field. The value $\max\_pq = 4$ was used in the experiments that are described in Section 3.

$$\mathbf{F}(\tilde{p},\tilde{q}) = \int\int \left| p(x,y) - \tilde{p}(x,y) \right|^2 + \left| q(x,y) - \tilde{q}(x,y) \right|^2 \, dx\, dy \tag{9}$$

where p, q denote the given gradient field components and

$$\tilde{p}(x,y) = \tilde{Z}_x(x,y) = \frac{\partial \tilde{Z}(x,y)}{\partial x}, \quad \tilde{q}(x,y) = \tilde{Z}_y(x,y) = \frac{\partial \tilde{Z}(x,y)}{\partial y} \tag{10}$$

denote the unknown (ideal) gradient field components that have to be reconstructed. A surface is calculated by minimization of **F** which ensures a maximum consistency of the reconstructed surface with the given data array. The difficulty with this method consists in selecting proper initial values at the boundary of the integration process. Horn [5] suggested some boundary conditions.

Frankot and Chellappa [3] assume that the unknown surface function Z satisfies the integrability condition (3). Furthermore they assumed a Fourier coefficient representation

$$Z(x,y) = \tfrac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} Z^{(F)}(u,v) \cdot e^{-j \cdot (u \cdot x + v \cdot y)} du\, dv \tag{11}$$

of this function, where

$$Z^{(F)}(u,v) = \tfrac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} Z(x,y) \cdot e^{j \cdot (u \cdot x + v \cdot y)} dx dy \tag{12}$$

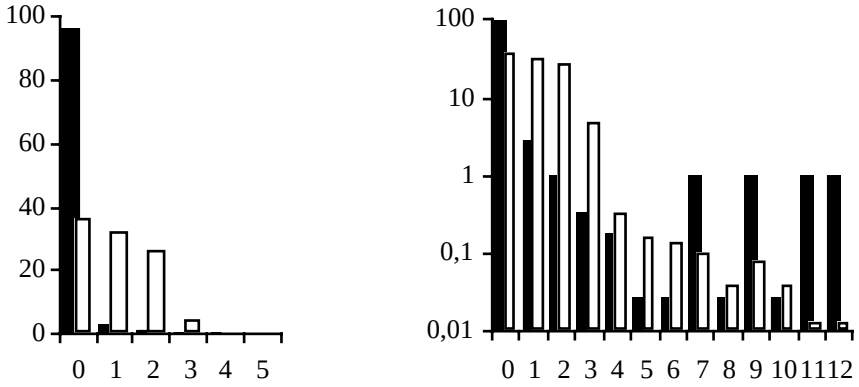

Figure 3: A linear scale of percentage of array positions is used in the histogram on the left and a logarithmic scale in the histogram on the right (white columns: **fpi**).

denote the Fourier coefficients of Z. Based on these assumptions they could prove a theorem allowing the reconstruction of function Z in the Fourier space. Then an inverse Fourier transform leads to the desired surface data, see Fig. 2. This method was selected to represent global techniques in the comparison with local techniques.

3 EVALUATION

Several experiments were performed to evaluate the quality of surface integration applying the local four-path integration algorithm (**fpi**) and the global Fourier expansion algorithm (**FE**).

3.1 Experiment SPHERE

The continuous gradient field of a synthetic sphere is digitized as a 256×256 gradient field. The radius of the sphere is equal to 100 grid units. The reconstructed height values are shifted in the range of the original surface using LSE optimization. Then the height differences between the ideal surface and the reconstructed surface were calculated.

Fig. 3 shows the histograms of these height differences. The height interval used is $\Delta Z = 1$. The white columns characterize **fpi** and the black columns **FE**. Missing columns relate to values close to zero. The left histogram suggests that nearly all **FE** height values seem to be correct. However, the logarithmic scale allows a more detailed insight, and in fact there are more errors in the range between 7 and 12 than in the **fpi** case.

The **fpi** allows reconstructions with errors less than 3% (in relation to the maximum height of the visible part of the sphere, i.e., 100 grid units) at 94.4% of array positions. The averaged error is 1.5%, and the maximum error is 12.8%. In case of the **FE** there were 95.6% positions below 1% of error. The averaged error is 0.21% and the maximum error is 10.4%. Maximum errors appear for both methods at the border of the sphere, i.e., for long gradient vectors.

The **fpi** produces a certain distortion of the sphere (see Fig. 4) due to the propagation of errors effected by long gradients, and due to insufficient local approximations within the discrete integration process.

Figure 4: Some isoheight-lines in the calculated depth maps (**fpi** on the left).

3.2 Experiment MOZART

A synthetic Mozart statue was given as a discrete array of height data, see. Fig. 5. Three different methods (see Section 1.3) were used to calculate the input data arrays for surface integration. Then both algorithms used these generated gradient fields as input. The reconstructed height values are shifted in the range of the original surface using LSE optimization. The resulting errors, i.e., the absolute differences in height, are shown in Tab. 1 (maximum error, averaged error, and standard deviation). Positions with long gradients (length greater than 4) were excluded from this statistical analysis of resulting height values.

The regularization through Gaussian kernels leads to the worst results for **fpi** (see averaged error and standard deviation). The simple differentiation schemes correspond to the simple local integration techniques. The reconstructed **FE** height data are more robust, see maximum errors and standard deviations in Tab. 1. The **FE** values were essentially improving against the **fpi** values if the regularized differential scheme was used. However, an explicit numerical differentiation is not contained in the **FE** algorithm. Therefore the regularized differential scheme was used for further comparisons.

The quality of the results of the simple **fpi** algorithm as shown in Tab. 1 are quite surprising in comparison with the **FE** results. The competitive behavior of the **fpi** algorithm is also illustrated by the error histogram in Fig. 5. The percentage is calculated with respect to the maximum height of the statue. The width of the 15 classes in this histogram is about 1.54%. About 80% of the **fpi** results have errors $\leq$ 7.7%, and about 80% of the **FE** results have errors $\leq$ 4.6%. If all gradient data, i.e., gradient vectors of

gradient method	maximum error	averaged error	σ_z
backward difference quotient	**18.9** 14.3	**2.6** 2.9	**2.3** 2.0
central difference quotient	**33.0** 13.2	**2.9** 3.0	**2.3** 2.0
Gaussian kernels $\sigma = \sqrt{2}$	**12.4** 14.7	**4.0** 2.7	**2.5** 1.9

Table 1: Different discrete differentiation schemes were used to generate the input gradient fields. The upper rows (bold numbers) denote the resulting errors using **fpi**, and the lower rows correspond to **FE**.

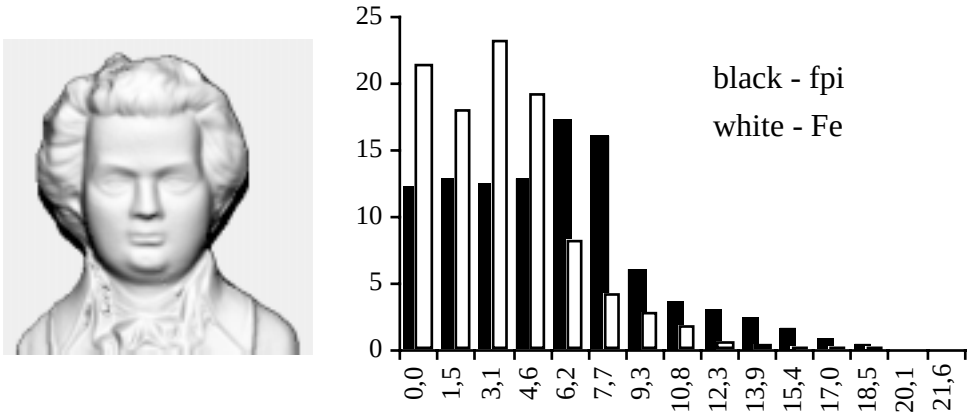

Figure 5: Mozart statue (shaded representation of input data) and histograms of height differences (white - **FE**).

any length are used as input to the **FE** algorithm then the reconstructed surface is heavily distorted. This problem does not appear with the **fpi** algorithm.

Error distributions over the surface range are also very descriptive besides global error statistics. Fig. 6 illustrates the distribution of absolute height errors using a floating horizon representation. The **FE** values are nearly without error within the face region. Some waves of errors were generated in this region using the **fpi** approach. As expected there are large errors at the object boundary (that were excluded from the statistical analysis above), i.e., both methods are very sensitive to such "abrupt changes in orientation".

3.3 Experiment K

A synthetic polyhedral K-shaped object was given as a discrete 256×256 array of height data, see. Fig. 7. The height extent is 162.43 grid units. This object turned out to be a very complicated one for height data integration based on gradient information.

There are locations for this surface function where integration is not possible. At surface edges there is no C^1-continuity, and at occluding edges there is even no C^0-continuity. Some faces of the polyhedron are not visible that would be needed to reconstruct the diagonally into the view space placed polyhedron. A further difficulty arises from the fact that the object border has very different height values. The concavity of the object also renders more difficulties.

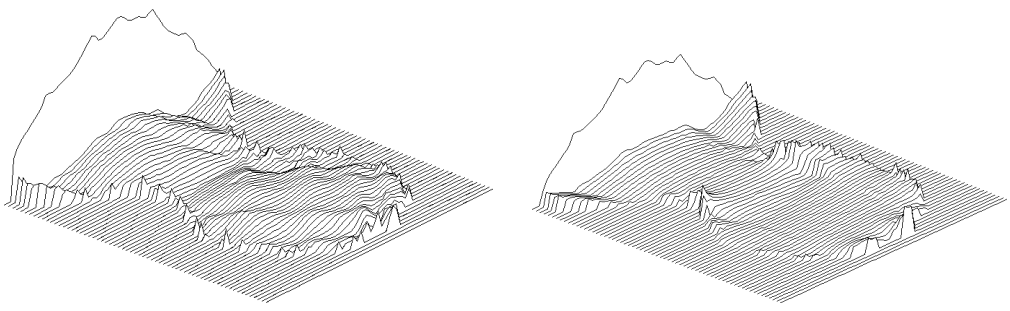

Figure 6: Floating horizon representation of absolute height differences (**fpi** on the left).

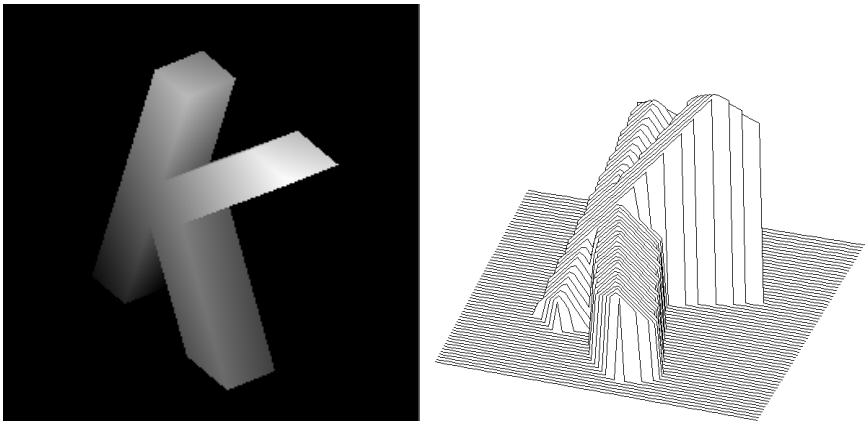

Figure 7: A K-shaped synthetic polyhedron with a maximum height of 162.43 grid units: height map - on the left, floating horizon representation - on the right.

The gradient field was generated using the backward difference quotient method. This method possesses the advantage of less smoothing at object edges in comparison to the numerical approximation (Gaussian kernel). The property of edge preserving should be used to evaluate integration techniques.

Fig. 8 (on the left) illustrates the results using local path integration. The visual comparison with Fig. 7 shows roughness for originally planar faces. The four-path method assumes the same height values at the four array corners. This was about adequate for the sphere and the Mozart statue but not for this K-shaped polyhedron. These objects had about the same height value at the object border. The K-shaped polyhedron takes its minimum and its maximum height value at certain border points! The initialization of the height map with a uniform constant is extremely disadvantageous for this polyhedron as for polyhedrons in general.

Fig. 8 (on the right) illustrates the results using the global **FE** integration technique. This global technique was not able to cope with the large height differences. However, the general surface structure is much closer to the original planar surfaces than in case of the **fpi** results.

The statistical analysis shows similar results for both techniques **fpi** / **FE**: maximum errors 90.94 / 77.15, averaged error 23.03 / 21.71, and standard deviation 18.25 / 18.82. From the error histograms (not shown here) it follows that both techniques reconstruct 73% of height values with errors ≤15%.

In case of undefined or long gradients (length at least equal to 4) the value 0 (i.e., no impact on height changes) was used instead. This was reasonable for the sphere and the Mozart statue but not for the

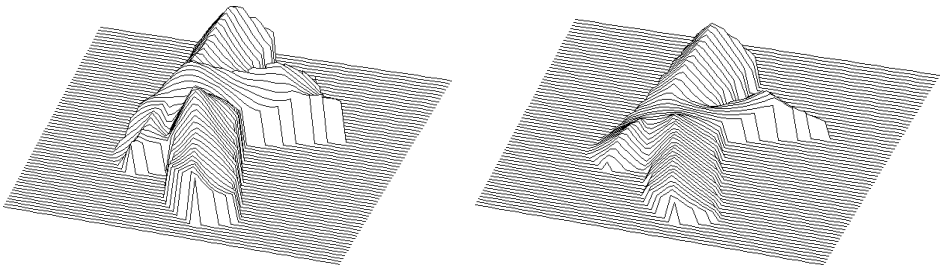

Figure 8: The maximum error of the reconstructed **fpi** surface (on the left) is equal to 90.94 grid units, and 77.15 grid units for the reconstructed **FE** surface.

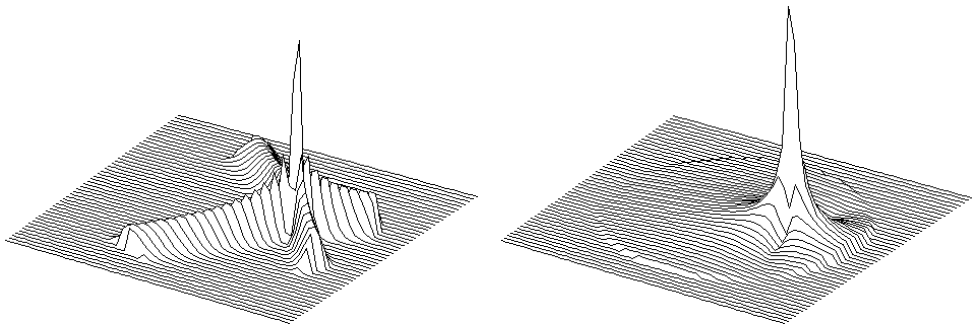

Figure 9: Absolute height differences after spike noise for the reconstructed sphere (**fpi** on the left: maximum error 3.58, **FE** on the right: maximum error 4.43).

polyhedron. Altogether, both techniques are not adequate to reconstruct this polyhedron as well as polyhedrons in general within reasonable quality limits. However, it can be suggested that a locally adaptive path integration technique should be sensitive to local patterns in the gradient map to obtain adequate reconstruction results. The use of the edge map of the given irradiance image could also support the proper choice of the integration path, and of the locally adaptive integration procedure.

3.4 Experiment SPIKE_NOISE

For the synthetic sphere in Section 3.1, with the projected sphere centroid at point $\mathbf{p} = (0,0)$, a single point noise is simulated by modifying slant $30°$ at point $\mathbf{p} = (50,0)$ into $arctan(4) \approx 76°$. The same slant value is assigned to three neighbors of point $\mathbf{p} = (50,0)$ - on the left, below, below on the left. All the remaining gradient values remain unchanged including all the tilt values. The resulting absolute height differences between the ground truth (ideal sphere) and the reconstructed sphere are shown in Fig. 9. The four-path integration generates an adequate error pattern (propagation into the direction of the four array corners). The **FE** maximum error is higher than the **fpi** maximum error, but the **FE** error function is faster decreasing than the **fpi** error function. In both cases the spike noise did result into a spike distortion of the reconstructed shape. If smoothness constraints could be assumed for reconstructed surfaces then such distortions could be suppressed.

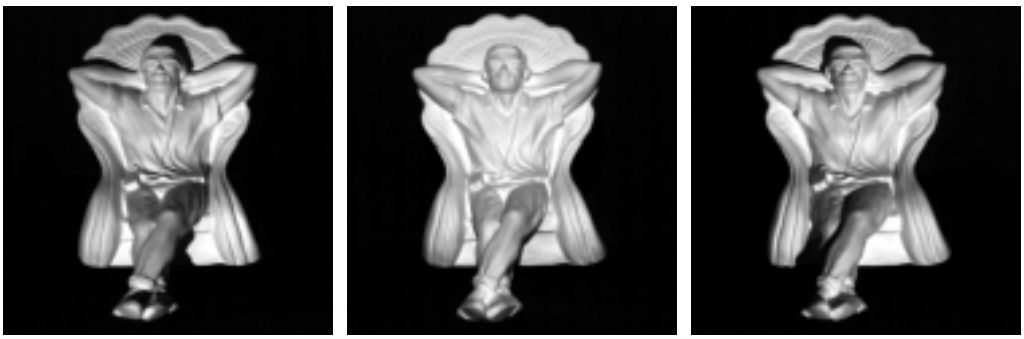

Figure 10: Input images for a three-source photometric stereo analysis (plaster statue "man in chair").

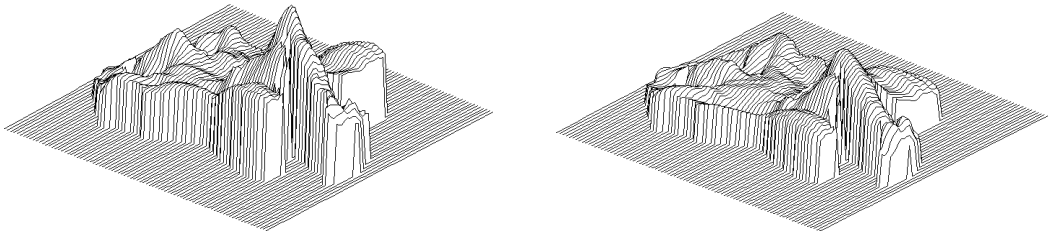

Figure 11: Reconstructed "man in a chair" (**fpi** is on the left).

3.5 Experiment QUALITATIVE_EVALUATION

A large number of real objects were used to compare the behavior of both integration techniques. The gradient maps were generated using photometric stereo analysis [7,9]. Here we illustrate two typical examples. Fig. 10 shows three input images for photometric stereo analysis of the "man in chair". This plaster statue is of "very complex shape". A serious difficulty is that the object border has different height values, similar to the synthetic K-shaped polyhedron. The reconstruction results are illustrated in Fig. 11.

Both methods produce reasonable results for such curved objects in general. The **fpi** algorithm generated erroneous height values in the head region and in the knee region for this specific object.

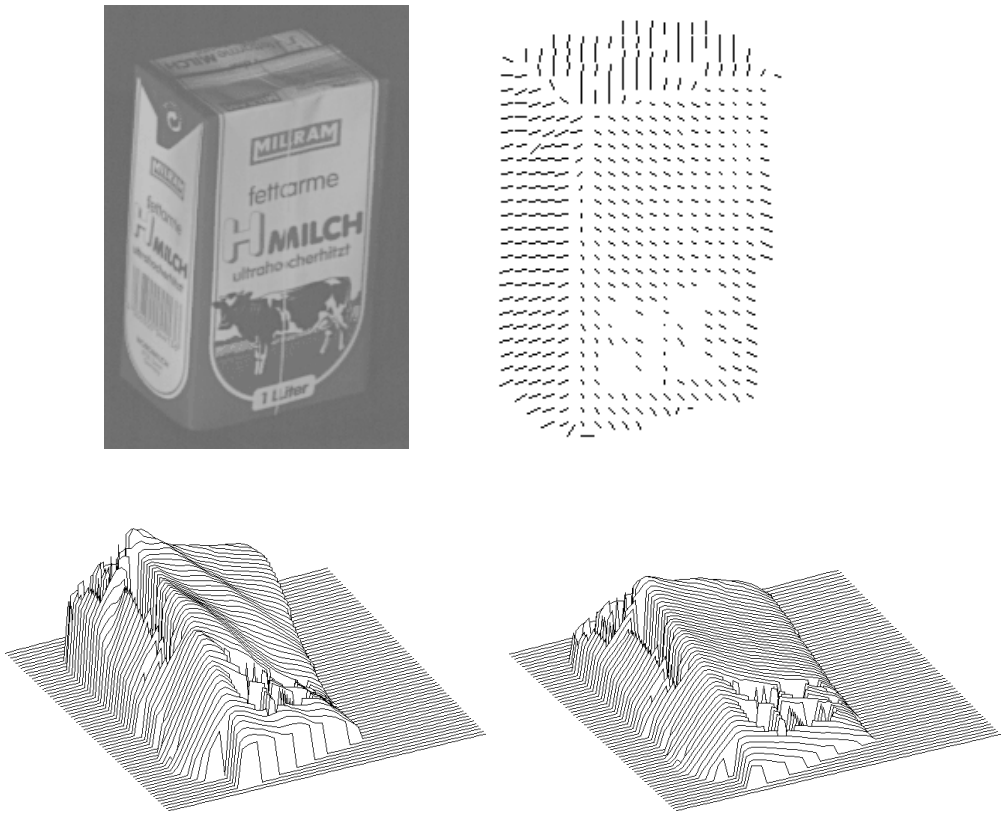

Figure 12: TetraPak, calculated gradient field in needle map representation, and reconstructed surfaces (**fpi** on the left).

The reconstruction results for polyhedron-like real world objects had the same essential distortions as demonstrated in Section 3.3 for a synthetic object. Fig. 12 shows a TetraPak, its gradient field as reconstructed from photometric stereo analysis, and both surface reconstruction results using the floating horizon representation. Especially remarkable are the strong **fpi** distortions of planar faces (wave-like distortion pattern - as for the Mozart statue). Errors in the generated gradient map (especially at locations of very low albedo) result in reconstruction errors for both integration techniques. This also points out that further improvements of shading based gradient calculations are required to ensure applications in 3-D object modeling.

4 CONCLUSIONS

The advantages of the Fourier expansion in comparison to **fpi** are a better robustness with respect to noise (besides spike noise this was also studied for additive Gaussian noise), and better qualitative evaluations for real scenes, especially for curved objects. The advantage of the four-path integration in comparison to **FE** is robustness with respect to high gradient values (however, this seems to be of low importance in real scenes). The most difficult reconstruction situation is defined by polyhedral objects as the synthetic K-shaped polyhedron, or the milk box. The computing time of the **FE** algorithm was about 1.6 times the computing time of the **fpi** algorithm for 256×256 maps.

More generally speaking, the reconstruction results depend upon the object shape. Polyhedral objects (e.g. especially relevant to industrial applications) are the most critical shapes for the studied local and global methods. The proposed error measures should be discussed in relation to applications of shape reconstruction techniques. In the ideal case, certain "data sheets" about transformation modules should support the decision that one can be chosen for a specific application context.

There is a remarkable deficiency of literature about integration techniques, at least in computer vision. Further research could be directed on refinement of methods (esp. for serving polyhedral objects). The global technique seems to be slightly in favor in relation to the local technique so far. However, it can be expected that the local techniques will have restrictions to deal with complex shapes. The development of locally adaptive methods (e.g. relaxation) could be appropriate. There is also a need in developing quantitative evaluation measures characterizing the qualitative appearance.

5 REFERENCES

1. D. E. Bourne and P. C. Kendall, *Vektoranalysis*, Teubner Verlagsgesellschaft, Stuttgart, 1988.
2. N. E. Coleman, Jr. and R. Jain, "Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry", *CGIP*, 18 (1982), pp. 309-328.
3. R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms", *IEEE Trans. on PAMI*, 10 (1988), pp. 439-451.
4. G. Healey and R. Jain, "Depth recovery from surface normals", *ICPR´84*, Montreal, Canada, Jul. 30-Aug. 2, 1984, pp. 894-896.
5. B. K. P. Horn, "Height and gradient from shading", *Int. J. Computer Vision*, 5 (1990), pp. 37-75.
6. B. K. P. Horn and M. J. Brooks, "The variational approach to shape from shading", *CVGIP*, 33 (1986), pp. 174-208.
7. R. Klette, A. Koschan, and K. Schlüns, *Computer Vision - Räumliche Information aus digitalen Bildern*. Vieweg, Braunschweig 1996.
8. V. Rodehorst, "Vertiefende Analyse eines Gestalts-Constraints von Aloimonos und Shulman", Technischer Bericht, CV-Bericht 8, Institut für Technische Informatik, TU Berlin, 1993.
9. R. J. Woodham, "Photometric method for determining surface orientations from multiple images", *Optical Engineering*, 19 (1980), pp. 139-144.
10. Z. Wu and L. Li, "A line-integration based method for depth recovery from surface normals", *CVGIP*, 43 (1988), pp. 53-66.