

Depth from Gradient Fields and Control Points: Bias Correction in Photometric Stereo

Itsik Horovitz and Nahum Kiryati*

Department of Electrical Engineering–Systems
Faculty of Engineering
Tel Aviv University
Tel Aviv 69978, Israel
nk@eng.tau.ac.il

Abstract

Photometric stereo is capable of high quality reconstruction of fine shape details, but is prone to systematic errors due to nonideal illumination or imperfect calibration. We present methods for correcting the bias, using sparse control points of known 3D location. An easy way to obtain control points is via the projection of a dot-matrix pattern using a laser pointer with a suitable adapter, and triangulation. Straightforward incorporation of control points as constraints in the computation of depth from the gradient field leads to singularities. We propose two well-behaved methods for bias correction using control points. One is based on constrained weighted least squares extension of depth from gradient-field computation. The other adds an interpolation surface to the reconstructed shape. Practical computation of depth from a gradient field requires an efficient numerical scheme. We employ full-multigrid computation with successive over-relaxation and show how to propagate the gradient field and the control points through the pyramid. Experimental results demonstrate significant bias reduction in photometric stereo, allowing high reconstruction quality even in the presence of severe setup errors.

Keywords: photometric stereo, control points, depth from gradient field, weighted least squares (WLS), thin plate spline interpolation, full multigrid (FMG).

*Corresponding author. Phone: +972 3 640 7767. Fax: +972 3 640 7095.

1 Introduction

Photometric Stereo [13] is a monocular 3D shape recovery method, based on a few images of the same scene taken under different illumination conditions. Standard photometric stereo is a two-stage process: First, the normal to the surface (or gradient) is estimated at each visible point [35]. The set of normals, or the gradient field, is then used to determine the 3D surface itself, using either the calculus of variations [13, 14], direct integration [36] or frequency-domain methods [11, 18].

For a Lambertian surface, the gradient and the albedo of the surface at each point can be determined locally using three images, obtained with different illumination directions. For neutral-colored surfaces, the use of color allows photometric stereo recovery from a single image; for extensions see [7, 8]. Without using color, reducing the number of images is possible only by sacrificing the locality of the gradient computation, taking spatial considerations into account, see [21]. Shape from shading (SFS) [13, 15] requires a single image, but then solving for the gradient or relative depth at any point requires global integration of data over the image. Interreflection effects are considered in [4].

The local estimation of the gradients in photometric stereo is the key to its excellent numerical stability. In determining the 3D shape of the surface, each visible point contributes one unknown (depth) but two constraints, the components of the normal vector. The problem is thus overdetermined and the solution is very robust with respect to noise.

To demonstrate the quality of 3D recovery possible with three-image photometric stereo, compare the two reconstructions shown in Fig. 1. Shown at the left is the visualization of a mannequin head, acquired with a high precision laser-based Cyberware range scanner, regarded here as ground truth. Fig. 1 (right) is a 3D reconstruction of the same object obtained using photometric stereo, with careful adjustment of the angles and intensities of the light sources. It is very similar to the ground truth.

The quality of 3D recovery possible with photometric stereo, its fundamental noise immunity, the absence of the correspondence problem, the need for only three input frames and the low cost of the setup are all important advantages. They are contrasted by the phenomenon of bias (distortion) that often appears in the reconstruction. Photometric stereo can be used to accurately recover fine shape details, but cannot always be trusted for global shape measurements.



Figure 1: *Left*: A mannequin head shape, acquired using a Cyberware laser-based range scanner, used in this study as “ground truth”. *Right*: Reconstruction by (carefully tuned) photometric stereo.




Figure 2: *Left*: Photometric stereo reconstruction, with bias due to an error of 5 degrees in the assumed direction of one of the three light sources. *Right*: Biased reconstruction due to a 5% error in the intensity of one of the sources.

3D shape recovery using photometric stereo is based on knowing the directions of the light sources¹ and also requires some assumptions regarding their relative intensity, beam cross-sectional uniformity and presence of ambient light. The bias in photometric stereo is mostly due to systematic error build-up, caused by light source calibration errors or by other deviations from the assumed illumination. Fig. 2 (left) demonstrates the distortion caused by an error of 5 degrees in the direction of one of the light sources. The bias shown in Fig. 2 (right) is due to a 5% error in the intensity of one of the sources. The systematic errors in Fig. 3 (left) are due to poor cross-sectional uniformity of one of the light sources. Fig. 3 (right) presents the effect of uncompensated ambient light in one of the images. Note, however, that since photometric stereo is overconstrained, moderate deviations of the surface from the Lambertian reflectance model [22] lead to fairly small errors in the reconstruction,

¹Lambertian photometric stereo with uncalibrated light directions and intensities determines the surface normals only up to an invertible linear transformation [6]. Additional aspects of uncalibrated photometric stereo are discussed in [5, 9].




Figure 3: *Left*: Biased reconstruction due to poor cross-sectional uniformity of one of the light sources. *Right*: Bias due to the presence of uncompensated ambient light in one of the three images.




Figure 4: Photometric stereo reconstruction of the mannequin head, covered with a thin layer of hair shampoo to increase its specularity. The arrows point at errors. The good overall quality confirms that photometric stereo is not very sensitive to deviations from the Lambertian surface model.

see Fig. 4.

One way to eliminate bias in photometric stereo is to carefully assemble a high quality, well calibrated image acquisition setup. This is neither easy nor cheap and is in many cases impractical. Our aim is to relax the requirements for laboratory-grade illumination setup and calibration and help turn photometric stereo into a useful, low cost technique.

The approach taken in this paper is to correct the reconstruction bias in photometric stereo after the data had already been acquired, using sparse control points of known 3D location on the reconstructed surface. It is motivated by the fact that the bias is spatially smooth: it is a global, rather than local phenomenon. The following questions arise:

- How to obtain control points?
- How to incorporate control points in photometric stereo reconstruction?

- How to accommodate control points in an efficient numerical scheme for photometric stereo reconstruction?

Fortunately, sparse control points are easy to obtain. A simple and cheap way is by illuminating the scene using a laser pointer with a pattern adapter tip. We used a Clearline Concepts CL2042 keychain laser pointer, that costed less than US\$ 20 and included 42 different interchangeable pattern tips. The projected pattern, say a dot matrix, creates sparse, brightly illuminated points on the surface that are usually easy to identify and register. The 3D position of each of these points is simple to determine, since it is at the intersection of the projection ray and the line of sight to the camera. Calibrating this minimal structured-light system is a straightforward process.

As a direct outcome of the robust and overdetermined nature of photometric stereo, the incorporation of control points is nontrivial². The excessive number of constraints presented by the gradient field leads to an extremely rigid reconstruction. Imposing the control points as hard constraints creates singularities in the surface. Giving them weights similar to those assigned to gradient field data minimizes their corrective effect on the reconstruction. Various interesting approaches to combining geometric and photometric data have been suggested (in different contexts) in [2, 3, 12, 16, 17, 19, 20, 24, 26, 33, 34], but these papers do not address the basic problem of surface recovery from gradient field data and sparse control points. The combination of orientation data with sparse depth constraints has been considered within the comprehensive computational framework of Terzopoulos [29, 30, 31] and in [25], but, as will be seen in the sequel, the fundamental dilemma between the rigid photometric stereo reconstruction and the hard control points has not been resolved.

Two well-behaved methods for bias correction in photometric stereo using control points are presented in this paper. One is based on constrained weighted least squares extension of depth from gradient field computation. The other adds an interpolation surface to the reconstructed shape. Note that straightforward computation of depth from the gradient field is very expensive computationally. An efficient numerical scheme is therefore required for practical photometric stereo reconstruction. We employ full-multigrid computation with successive over-relaxation and show how to propagate control points through the pyramid.

²Note that isolated control points cannot be regarded as boundary conditions since they do not form a closed contour.

Experimental results demonstrate significant bias reduction in photometric stereo, allowing high reconstruction quality even in the presence of severe setup errors.

2 Depth from Gradient Field and Control Points

2.1 Depth from gradient field

Given a gradient field $\{(p,q)\}$ within some image domain \mathcal{D} , Horn and Brooks [14] propose to reconstruct the surface³ $z(x, y)$ using the calculus of variations, minimizing the functional

$$\iint_{\mathcal{D}} [(z_x - p)^2 + (z_y - q)^2] dx dy . \quad (1)$$

The corresponding Euler equation is

$$\nabla^2 z = p_x + q_y . \quad (2)$$

If $z(x, y)$ is known along the perimeter $\partial\mathcal{D}$ of \mathcal{D} , it can be used as a boundary condition. Otherwise, one should use the natural boundary condition

$$(z_x, z_y) \cdot \vec{n} = (p, q) \cdot \vec{n} ,$$

where

$$\vec{n} = \left(-\frac{dy}{ds}, \frac{dx}{ds} \right)$$

is the normal vector to $\partial\mathcal{D}$ in the (x, y) plane and s denotes arc length along $\partial\mathcal{D}$. For the discrete case, Horn and Brooks [14] suggest an iterative numerical scheme, based on discrete approximations of partial derivative operators.

A numerical scheme similar to that of Horn and Brooks can be obtained using purely discrete reasoning. Denote the gradient field data $\{(\tilde{p}_{i,j}, \tilde{q}_{i,j})\}$ and define

$$p_{i,j} = \frac{\tilde{p}_{i,j} + \tilde{p}_{i+1,j}}{2} \quad q_{i,j} = \frac{\tilde{q}_{i,j} + \tilde{q}_{i+1,j}}{2} .$$

³The surface is usually assumed to be C^2 smooth (or piecewise smooth), see [14].

Within the image domain, the surface $\{z_{i,j}\}$ should satisfy

$$z_{i+1,j} - z_{i,j} = p_{i,j} \quad (3)$$

$$z_{i,j+1} - z_{i,j} = q_{i,j} . \quad (4)$$

Optimal surface recovery in the least squares sense requires the minimization of

$$\sum_{i,j} (z_{i+1,j} - z_{i,j} - p_{i,j})^2 + (z_{i,j+1} - z_{i,j} - q_{i,j})^2 . \quad (5)$$

Differentiating with respect to z_{ij} and equating to zero leads to the iterative scheme

$$z_{i,j}^{k+1} = \frac{1}{4} [z_{i+1,j}^k + z_{i-1,j}^k + z_{i,j+1}^k + z_{i,j-1}^k - p_{i,j} + p_{i-1,j} - q_{i,j} + q_{i,j-1}] . \quad (6)$$

The convergence of this algorithm for computing depth from gradient field data is very slow; fast numerical schemes are considered in section 3.

2.2 Control points as constraints

Suppose that, in addition to the gradient field, we are given several control points, i.e., sparse depth constraints of the form

$$z_{i_k,j_k} = z_k . \quad (7)$$

How can constraints of this type be imposed in computing depth from the gradient field?

Two approaches come to mind, but both unfortunately fail. One way is to add the control point equation (7) to the set of equations of types (3) and (4), i.e., to include the control point equation as another term within the least squares minimization objective (5). This procedure is ineffective, because the solution is not forced to satisfy the constraint and the corrective effect of the control point is small. An alternative approach is to enforce the constraint by *replacing* Eq. 6 with Eq. 7 at $i = i_k$, $j = j_k$ in the iterative scheme. The inadequacy of this method is best understood by an example in the continuous domain.

Let the image domain \mathcal{D} be a circle of radius R and the gradient field be identically zero, i.e.,

$$p(x, y) = q(x, y) = 0 \quad \forall (x, y) \in \mathcal{D} .$$

Assume the boundary condition $z(x, y) = 0$ at $(x, y) \in \partial\mathcal{D}$, i.e., at $x^2 + y^2 = R^2$. In this trivial case, the solution to the depth from gradient field problem is obviously $z(x, y) = 0 \quad \forall (x, y) \in \mathcal{D}$.

Let us now add a control point, $z(0, 0) = 1$, and solve the problem formally starting with Eq. 1. Since p and q are identically zero within \mathcal{D} , the functional to minimize takes the form

$$\iint_{\mathcal{D}} (z_x^2 + z_y^2) dx dy \quad (8)$$

and the Euler equation is

$$\nabla^2 z = 0 \quad .$$

Assuming a circularly symmetric solution and converting to polar coordinates, the functional becomes

$$\int \left(\frac{\partial z}{\partial r} \right)^2 r dr \quad (9)$$

and the Euler equation leads to

$$\nabla^2 z = \frac{\partial^2 z}{\partial r^2} + \frac{1}{r} \cdot \frac{\partial z}{\partial r} = 0 \quad . \quad (10)$$

Generally, the solution of Eq. 10 is of the form

$$z = c_1 \ln(r) + c_2 \quad (11)$$

where c_1 and c_2 are constants. Here, the boundary condition dictates that $z = 0$ at $r = R$. The control point requires $z = 1$ at $r = 0$. Thus, the solution is

$$z(r) = \begin{cases} 0 & 0 < r \leq R \\ 1 & r = 0 \end{cases}$$

and in Cartesian coordinates

$$z(x, y) = \begin{cases} 0 & (x, y) \in \mathcal{D}, (x, y) \neq (0, 0) \\ 1 & (x, y) = (0, 0) \end{cases} \quad (12)$$

Eq. 12 demonstrates that imposing a control point as a constraint in the classic method



Figure 5: *Left*: In the discrete case, straightforward use of a control point as a constraint in computing depth from normal directions creates a tent-pole. *Right*: Reconstruction using the Terzopoulos formulation. In both cases, as dictated by the control point, the height of the peak above the base plane is 1.

for computing depth from gradient fields (Eq. 2) is useless, since the control point has no effect on the recovered surface beyond its own location! The discrete version of this example, shown in Fig. 5 (left), is slightly better behaved. However, it is clear that in the discrete case too, straightforward use of control points as constraints does not lead to the intended goal: correcting the spatially global bias while maintaining the local shape.

The combination of orientation data with isolated depth constraints has been considered in [25]. It is also a special case within the computational framework of Terzopoulos [29, 30, 31]. In both studies, the basic approach is to minimize a functional that, in addition to measuring deviations from the orientation data (Eq. 1), severely penalizes depth errors at the control points. For example, the functional suggested in [25] is

$$\iint_{\mathcal{D}} [(z_x - p)^2 + (z_y - q)^2 + \lambda(x, y)(z - z^c)^2] dx dy \quad , \quad (13)$$

where z^c represents the known depth at the control points, and

$$\lambda(x, y) = \lambda_0 \sum_k \delta(x - x_k, y - y_k) \quad (14)$$

represents the positions of the control points. The functional suggested in [29, 30, 31] contains an additional smoothness term. Discrete solution of the above example using the Terzopoulos formulation is shown in Fig. 5 (right). The result still resembles a tent-pole, which is clearly not the desired influence of a control point in photometric stereo.

2.3 Reconstruction via weighted least squares (WLS)

The phenomenon shown in Eq. 12 for the continuous case and in Fig. 5 for the discrete case is related to the robustness and rigidity of the algorithm for computing depth from gradient field: the effect of a disturbance (in this case, in the form of a control point) is tightly localized, leading to singularity in the continuous case and to singular-like behavior in the discrete setup. Based on this observation, when imposing control points as hard depth constraints, we suggest to also enforce compliance of the reconstructed surface with the gradient field near the control points. This can be accomplished by moving to a weighted least squares formulation, generalizing from Eq. 5 to

$$\sum_{i,j} w_{i,j} [(z_{i+1,j} - z_{i,j} - p_{i,j})^2 + (z_{i,j+1} - z_{i,j} - q_{i,j})^2], \quad (15)$$

where $w_{i,j}$ is a decreasing function of the distance r to the control point nearest to (i, j) . More accurately, since the discrete orientation components $p_{i,j}$ and $q_{i,j}$ refer to slightly different locations, $c_{i,j}^x = (i + 1/2, j)$ and $c_{i,j}^y = (i, j + 1/2)$ respectively, we use

$$\sum_{i,j} w_{i,j}^x (z_{i+1,j} - z_{i,j} - p_{i,j})^2 + w_{i,j}^y (z_{i,j+1} - z_{i,j} - q_{i,j})^2, \quad (16)$$

where $w_{i,j}^x$ is a decreasing function of the distance between the nearest control point and $c_{i,j}^x$ and $w_{i,j}^y$ is a decreasing function of the distance from the control point to $c_{i,j}^y$. Here, distance between points is taken to be the length of the shortest path between them in the (i, j) domain. For a convex region, we can simply calculate the Euclidean distance; for non-convex domains we compute distance using the constrained distance transform of Verwer [32].

Consider again the continuous example presented in subsection 2.2. Suppose that compliance of the reconstructed surface with the gradient field is enforced near the control point at $(x, y) = (0, 0)$ ($r = 0$), by incorporating a weight function $w(r)$ within the functional (9). This leads to the minimization problem

$$\arg \min_z \int w(r) \left(\frac{\partial z}{\partial r} \right)^2 r dr \quad (17)$$



Figure 6: *Left*: Shape recovery with $w(r) = 1/r$ for the example discussed in subsection 2.2. *Right*: $w(r) = 1/r^2$. The difference between the two cases is in the distribution of the discrepancy between the gradient field data and the reconstructed gradients within the domain \mathcal{D} , necessary in order to comply with the hard constraint $z(0,0) = 1$. In both cases, as dictated by the control point, the height of the peak above the base plane is 1.

and its solution takes the form

$$z = c_1 \int \frac{1}{w(r)r} dr + c_2 . \quad (18)$$

The choice of $w(r)$ dictates the behavior of the reconstructed surface near the control point. For example, with $w(r) = 1/r$, the discrepancy between the gradient field data and the reconstructed gradients, caused by the control point, is distributed in \mathcal{D} and its maximal value is minimized. However, as seen in Fig. 6 (left), the reconstructed surface is still not smooth at the control point. With $w(r) = 1/r^2$ we obtain a paraboloid, shown in Fig. 6 (right). Generally, compliance of the reconstructed surface with the gradient field near *all* the control points is enforced by making the weight function depend on the distance \tilde{r} to the *nearest* control point. In this example, compliance with the gradients near the perimeter of \mathcal{D} is achieved with $\tilde{r} = \min\{r, R - r\}$. See Fig. 7, and compare with Fig. 5.

Consider again Fig. 2 (right), that shows the effect of an error of 5% in the estimated intensity of one of the three light sources on the photometric stereo reconstruction. The difference between the distorted reconstruction and the “ground truth” acquired with the Cyberware laser range scanner is shown in Fig. 8 as a set of equal-height contours, superimposed on the mannequin image⁴.

Suppose now that we obtain ten control points in which the true coordinates of the surface are known. Integrating these constraints with the gradient field data, using the weighted least squares approach, reduces the bias and leads to an improved reconstruction,

⁴Without control points, the recovery of depth from a gradient field is up to an additive constant. This means that the reference (0) level in Fig. 8 is arbitrary.



Figure 7: *Left:* Reconstruction with $w(\tilde{r}) = 1/\tilde{r}$. *Middle:* $w(\tilde{r}) = 1/\tilde{r}^{1.5}$. *Right:* $w(\tilde{r}) = 1/\tilde{r}^2$. Here $\tilde{r} = \min\{r, R - r\}$; the reconstructed surface complies with the gradient field data near the perimeter of \mathcal{D} . In all three cases, as dictated by the control point, the height of the peak above the base plane is 1.

as shown in Fig. 9. The reconstruction depends, of course, on the weight functions $w_{i,j}^x, w_{i,j}^y$ used. Quantitative analysis, using several reconstruction error measures and several weight functions, has shown that $w(r) = 1/r^{1.5}$ is a good choice. It has been used in all the WLS experiments shown.

2.4 Reconstruction using an interpolation surface

Photometric stereo captures fine shape details, but bias leads to global distortion. Motivated by this observation, we devised an alternative method for augmenting gradient field data with control points, in which the correction of the smooth bias is decoupled from the photometric stereo recovery of the fine details.

The surface is reconstructed as the sum of two surfaces independently computed. One is obtained by computing depth from the gradient field without using the control points. The other is a smooth interpolation surface that compensates for the bias. The interpolation surface uses relative control point depth values $\{v_i\}$ that are the difference between the true control point values $\{t_i\}$ and the corresponding depth values $\{s_i\}$ of the surface computed from the gradient field at the control point locations. This ensures that the sum of the two surfaces passes through all the control points. The relative control point depth values $v_i = t_i - s_i$ are packed as a column vector \mathbf{v} .

Finding a smooth surface that passes through a given set of points is a classical interpolation problem [10]. Solutions to the problem differ in their smoothness criteria and in the permissible distance between the surface and the points, see e.g. [27]. We use thin plate spline interpolation, a specific global basis function type method [10]. It yields a surface that

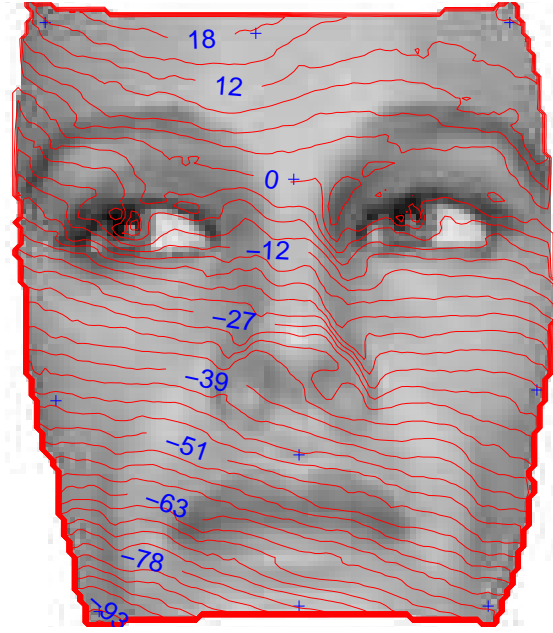


Figure 8: The difference between the distorted reconstruction caused by an error of 5% in the estimated intensity of one light source (Fig. 2-right), and the “ground truth” obtained using a Cyberware laser range finder (Fig. 1-left).

passes through the relative control points while minimizing the bending energy functional

$$\int_{R^2} [(\frac{\partial^2 f}{\partial x^2})^2 + 2(\frac{\partial^2 f}{\partial x \partial y})^2 + (\frac{\partial^2 f}{\partial y^2})^2] dx dy .$$

The interpolation surface is of the form

$$z(x, y) = \sum_i w_i G(d_i(x, y)) + a_0 + a_x x + a_y y , \quad (19)$$

where $d_i(x, y)$ is the distance⁵ between the i -th control point and (x, y) , $G(x) \equiv x^2 \log x$ and $\{w_i, a_0, a_x, a_y\}$ are coefficients that have to be determined.

The coefficients are computed as described in [1]. Let r_{ij} denote the distance between the i and j control points. Assuming N control points, create an $N \times N$ matrix $\mathbf{K} = [G(r_{ij})]$. The coordinates of the control points are placed in an $N \times 3$ matrix \mathbf{P} , with the i 'th row

⁵When dealing with non-convex domains \mathcal{D} , the straight line between a point (x, y) and some control point may cross the boundary $\partial\mathcal{D}$ of \mathcal{D} . In such cases we measure the distance using Verwer's 2D constrained distance transform [32].




Figure 9: By augmenting the gradient field data with ten control points (marked +) via the weighted least squares approach, the bias is reduced.

being $[1, x_i, y_i]$. These matrices are blocks within the $(N + 3) \times (N + 3)$ matrix \mathbf{L} ,

$$\mathbf{L} = \begin{pmatrix} \mathbf{K} & \mathbf{P} \\ \mathbf{P}^T & \mathbf{0} \end{pmatrix}$$

where $\mathbf{0}$ is a 3×3 matrix of zeros. An $N + 3$ column vector \mathbf{y} is obtained by extending \mathbf{v} with three zeroes: $\mathbf{y} = [\mathbf{v}^T | 000]^T$. Regarding the coefficients as the components of a $N \times 3$ vector, it can be shown that

$$[w_1, w_2, \dots, w_N, a_0, a_x, a_y]^T = \mathbf{L}^{-1} \mathbf{y} .$$

For an $n \times n$ image, computing the interpolation surface requires the inversion of the matrix \mathbf{L} and the evaluation of the interpolation surface at each of the $n \times n$ points. Since the number of control points N is usually small, sophisticated matrix inversion methods need not be used and the complexity of matrix inversion can be taken as $O(N^3)$. The evaluation of the interpolation surface is an $O(n^2 \cdot N)$ operation and is thus the computationally dominant stage in the process.

Results of bias correction using control points via the interpolation surface approach are




Figure 10: The difference between the bias-corrected reconstruction using ten control points (marked +) via the interpolation surface approach, and the “ground truth” obtained with a Cyberware laser range finder.

presented in Fig. 10. For the case of photometric stereo with an error of 5% in the estimated intensity of one light source, it shows the difference between the bias-corrected reconstruction and the “ground truth” obtained using the Cyberware laser range finder. Comparison with the uncompensated reconstruction (Fig. 8) reveals the remarkable improvement obtained by adding information from only ten control points.

2.5 Quantitative evaluation

As seen in Fig. 2 (right) and in Fig. 8, an error in the intensity of one of the three light sources leads to large bias in the photometric stereo reconstruction. Fig. 11 (top) is a histogram of the depth error, i.e., the difference between the ground-truth and the photometric-stereo reconstruction, without using control points. For bias correction with 10 control points, using the WLS and interpolation-surface approaches, the corresponding histograms are respectively shown in Fig. 11 (left) and Fig. 11 (right). The mean and standard deviation of the depth errors are compared in Table 1.

Additional insight can be gained by comparing the difference between the *gradients* of the surfaces reconstructed using photometric stereo, with or without control points, and




Figure 11: Histogram of the depth difference between a distorted reconstruction caused by an error of 5% in the intensity of one light source, and the “ground truth” obtained using a Cyberware laser range finder. *Top*: Without bias correction. *Bottom Left*: With bias correction using the WLS approach. *Bottom Right*: With bias correction using the interpolation-surface technique.

Depth errors	Mean	STD
Without bias correction	-	28.53
With bias correction (WLS)	1.18	3.73
With bias correction (interp. surf.)	0.57	2.34

Table 1: The mean and standard deviation of the depth errors in photometric stereo reconstruction. *Top row*: without bias correction (the mean is arbitrary, because depth in this case is recovered only up to an additive constant). *Middle row*: with bias correction using the WLS approach. *Bottom row*: with bias correction using the interpolation surface approach.

those of the ground truth. Histograms of the errors in the y components of the gradients are presented in Fig. 12 (the errors in the x component are much smaller in this particular example). Fig. 12 (top left) shows the histogram of the difference between the raw gradient data and that of the ground truth - before surface reconstruction (integration). It is biased (the mean is not zero) and noisy (large standard deviation). Fig. 12 (top right) refers to the surface reconstructed without bias correction. The over-determined surface integration process clearly reduces the standard deviation of the error (noise), but maintains the bias. Fig. 12 (bottom-left) and Fig. 12 (bottom-right) correspond to reconstruction with bias correction, using the WLS and interpolation-surface approaches respectively. Using both methods, the bias is eliminated. The mean and standard deviation of the errors in the y component of the gradient are compared in Table 2.

2.6 Benchmark: facial feature editing

Consider the 3D digitized face model visualized in Fig. 13 (courtesy of the Institute for Information Technology, National Research Council of Canada). Starting from a depth image (150×256 16-bit depth values), we simulated three images of the face as it would appear if illuminated from three different directions, see Fig. 14. Without adding errors and noise, photometric stereo reconstruction of the face based on the three images is virtually identical to the original model (Fig. 13).

We then provided a few artificial control points. All but one of them coincided with the original surface. The depth of remaining control point, located at the tip of the nose, was modified to extend the nose. The surface was reconstructed using a standard method (subsection 2.2, Eq. 13), the weighted least squares approach (subsection 2.3) and the interpolation surface technique (subsection 2.4). The result obtained via Eq. 13 is shown in Fig. 15. The spike at the tip of the nose demonstrates the limitation of the previous approaches, and is generally similar to the synthetic example presented in Fig. 5. Excellent results were obtained both with the weighted least squares technique (Fig. 16 - left) and the interpolation surface approach (Fig. 16 - right).




Figure 12: Histograms of the difference between the y components of the gradient fields obtained using photometric stereo, and the y component of the gradient field derived from the ground-truth data. *Top Left*: Prior to surface reconstruction. *Top Right*: Following surface reconstruction, without bias correction. *Bottom Left*: With bias correction, using the WLS method. *Bottom Right*: With bias correction, using the interpolation surface approach.

Gradient errors (y component)	Mean	STD
Prior to surface reconstruction	0.859	0.839
Without bias correction	0.871	0.568
With bias correction (WLS)	0.074	0.628
With bias correction (interp. surf.)	-0.013	0.526

Table 2: The mean and standard deviation of the errors in the y component of the gradient.




Figure 13: The original face model.




Figure 14: Three images created from the face model.




Figure 15: Incorporating a (synthetic) control point in photometric stereo reconstruction using standard methods creates a large spike, which is not the desired effect.




Figure 16: Spike-free nose extension using the methods presented in this paper. *Left:* Reconstruction with depth constraints and weighted orientation compliance via WLS. *Right:* Reconstruction using the interpolation surface technique.

3 Numerical Scheme

3.1 Fundamentals

The convergence of the generic algorithm for computing depth from a gradient field, as presented in [14] or as defined by Eq. 6, is slow. Fast numerical schemes are considered in this section, with emphasis on compatibility with the specific requirements arising in photometric stereo reconstruction with control points. These include

- Operation on non-square image domains (as in the examples shown in this paper).
- Accommodation of weights and pointwise constraints.

The numerical scheme presented in this section is applicable to both the WLS method and the interpolation surface method for bias removal in photometric stereo. In the latter, it is used for computing depth from the gradient field without control points, which is one component of the reconstructed surface.

At the basis of the methods used in this paper is the need to solve a large overdetermined system of linear equations, of the forms (3) and (4). This set of equations can be written in matrix form as

$$\mathbf{H} \cdot \mathbf{z} = \mathbf{g} ,$$

where the vector \mathbf{z} holds the (unknown) depth values $z_{i,j}$, the vector \mathbf{g} holds the gradient components $p_{i,j}$ and $q_{i,j}$ and \mathbf{H} is a huge sparse matrix. The optimal solution in the least

squares sense

$$\arg \min_{\mathbf{z}} (\mathbf{g} - \mathbf{H} \cdot \mathbf{z})^T (\mathbf{g} - \mathbf{H} \cdot \mathbf{z})$$

can be written explicitly as

$$\hat{\mathbf{z}} = (\mathbf{H}^T \mathbf{H})^{-1} \cdot (\mathbf{H}^T \mathbf{g}) .$$

However, direct solution by inverting $\mathbf{H}^T \mathbf{H}$ is very expensive.

Iterative methods for solving large sparse linear systems are reviewed in [23]. In fact, the iterative algorithm defined by Eq. 6 is the Jacobi scheme for our problem. The Gauss-Seidel scheme, defined by

$$z_{i,j}^{k+1} = \frac{1}{4} [z_{i+1,j}^k + z_{i-1,j}^{k+1} + z_{i,j+1}^k + z_{i,j-1}^{k+1} - p_{i,j} + p_{i-1,j} - q_{i,j} + q_{i,j-1}] \quad (20)$$

is faster. In successive over-relaxation (SOR), a relaxation parameter $1 < \omega < 2$ is included, and the algorithm takes the form

$$z_{i,j}^{k+1} = (1 - \omega) z_{i,j}^k + \omega \frac{1}{4} [z_{i+1,j}^k + z_{i-1,j}^{k+1} + z_{i,j+1}^k + z_{i,j-1}^{k+1} - p_{i,j} + p_{i-1,j} - q_{i,j} + q_{i,j-1}]. \quad (21)$$

Degenerate forms of Eq. 21 can be used at the boundaries⁶. For example, at the left boundary $z_{i-1,j}$ and $p_{i-1,j}$ are excluded and the equation becomes

$$z_{i,j}^{k+1} = (1 - \omega) z_{i,j}^k + \omega \frac{1}{3} [z_{i+1,j}^k + z_{i,j+1}^k + z_{i,j-1}^{k+1} - p_{i,j} - q_{i,j} + q_{i,j-1}]. \quad (22)$$

The convergence of SOR is faster than that of the Gauss-Seidel scheme, especially if ω is close to its optimal value.

In surface reconstruction via the weighted least-squares (WLS) approach, the objective function is generalized from Eq. 5 to Eq. 16, leading to changes in the numerical schemes. It is easy to show that in the WLS case the Gauss-Seidel iterative equation becomes

$$z_{i,j}^{k+1} = \frac{1}{w_{i,j}^x + w_{i-1,j}^x + w_{i,j}^y + w_{i,j-1}^y} [w_{i,j}^x (z_{i+1,j}^k - p_{i,j}) + w_{i-1,j}^x (z_{i-1,j}^{k+1} + p_{i-1,j}) + w_{i,j}^y (z_{i,j+1}^k - q_{i,j}) + w_{i,j-1}^y (z_{i,j-1}^{k+1} + q_{i,j-1})] \quad (23)$$

⁶They can be described as degenerate computational molecules [29].

and the SOR scheme is

$$z_{i,j}^{k+1} = (1 - \omega)z_{i,j}^k + \omega \frac{1}{w_{i,j}^x + w_{i-1,j}^x + w_{i,j}^y + w_{i,j-1}^y} [w_{i,j}^x(z_{i+1,j}^k - p_{i,j}) + w_{i-1,j}^x(z_{i-1,j}^{k+1} + p_{i-1,j}) + w_{i,j}^y(z_{i,j+1}^k - q_{i,j}) + w_{i,j-1}^y(z_{i,j-1}^{k+1} + q_{i,j-1})] . \quad (24)$$

Significant acceleration can be obtained by combining iterative schemes with full multi-grid (FMG) computation [23]. Terzopoulos [30, 31] used FMG with Gauss-Seidel iterations. We used SOR at the coarsest level and Gauss-Seidel at all other levels (ω set to 1). The 100×120 photometric stereo reconstructions presented in Figs. 1-4 were obtained using this numerical scheme. Computing time was about 2-3 seconds on a Pentium 3 500MHz personal computer with 128MB RAM, compared to about 20 *minutes* with the Jacobi scheme (Eq. 6).

Unlike the FMG method described in [23] (section 19) or in [28], in our implementation of FMG the grid points at a coarse level do not coincide with grid points in the fine level. This is necessary for the propagation of control points up the pyramid, since we did not find a way to do that using the restriction operator suggested in [28]. Another difference, discussed below, is that the *fine to coarse* process is performed separately on p and q .

3.2 Propagating the gradient field up the pyramid

The rule for propagating depth up the pyramid, that defines the restriction operator used, is

$$z_{i,j}^{\text{coarse}} = \frac{1}{8} [z_{2i,2j} + z_{2i+1,2j} + z_{2i,2j+1} + z_{2i+1,2j+1}] \quad (25)$$

Based on this restriction operator, we proceed to calculate the corresponding restriction operator for the gradient fields. Applying Eq. 3 at the coarse level, and substituting Eq. 25, we obtain

$$\begin{aligned} p_{i,j}^{\text{coarse}} &= z_{i+1,j}^{\text{coarse}} - z_{i,j}^{\text{coarse}} \\ &= \frac{1}{8} [p_{2i,2j} + 2p_{2i+1,2j} + p_{2i+2,2j} + p_{2i,2j+1} + 2p_{2i+1,2j+1} + p_{2i+2,2j+1}] \end{aligned} \quad (26)$$

Similarly, for q we get

$$q_{i,j}^{\text{coarse}} = \frac{1}{8} [q_{2i,2j} + 2q_{2i,2j+1} + q_{2i,2j+2} + q_{2i+1,2j} + 2q_{2i+1,2j+1} + q_{2i+1,2j+2}] . \quad (27)$$

The coarse to fine propagation of depth is defined by

$$z_{2i+m,2j+n} = 2z_{i,j}^{\text{coarse}} \quad \forall m, n \in \{0, 1\}$$

At this point one might wonder why we bother to create coarse versions of p, q . After all, if the problem to solve is the discrete case of the Poisson equation (Eq. 2), all we need to pass is $p_x + q_y$. However, had we done that we would have lost the ability to define the iterative equations at the boundary. For example, without having the coarse value of $p_{i-1,j}$ by itself, Eq. 22 cannot be calculated at the coarse level.

3.3 Propagating the control points up the pyramid

The control point is reduced using the point depth data and the gradient values between the four points that produce the reduced point. This is consistent with the restriction operator described in [23, 28]. Assume for example that the control point is located at $(2i, 2j)$. In this case we calculate $z_{2i+1,2j}$, $z_{2i,2j+1}$, $z_{2i+1,2j+1}$ and use them to derive $z_{i,j}^{\text{coarse}}$.

$$\begin{aligned} z_{2i+1,2j} &= z_{2i,2j} + p_{2i,2j} \\ z_{2i,2j+1} &= z_{2i,2j} + q_{2i,2j} \\ z_{2i+1,2j+1} &= z_{2i,2j} + p_{2i,2j} + q_{2i+1,2j} \end{aligned}$$

To the estimated depth we can add $k(p_{2i,2j} + q_{2i+1,2j} - p_{2i,2j+1} - q_{2i,2j})$ (a closed loop), where k is a constant. We select k that minimizes the variance of $z_{i,j}^{\text{coarse}}$ for the case of $p_{i,j}, q_{i,j}$ that are independent and identically distributed normal random variables, with zero mean and variance σ^2 .

$$\begin{aligned} z_{i,j}^{\text{coarse}} &= \frac{1}{8}(z_{2i,2j} + z_{2i+1,2j} + z_{2i,2j+1} + z_{2i+1,2j+1}) \\ &= \frac{1}{8}(4z_{2i,2j} + 2p_{2i,2j} + q_{2i,2j} + q_{2i+1,2j} + k(p_{2i,2j} + q_{2i+1,2j} - p_{2i,2j+1} - q_{2i,2j})) \\ \text{var}(z_{i,j}^{\text{coarse}}) &= \frac{1}{64}\sigma^2((2+k)^2 + (1-k)^2 + (1+k)^2 + k^2) = \frac{1}{64}\sigma^2(4k^2 + 4k + 6) \end{aligned}$$

The variance is minimized by $k = 1/2$, for which

$$z_{i,j}^{\text{coarse}} = \frac{1}{8} \left(4z_{2i,2j} + \frac{3}{2}p_{2i,2j} + \frac{3}{2}q_{2i,2j} + \frac{1}{2}q_{2i+1,2j} + \frac{1}{2}p_{2i,2j+1} \right). \quad (28)$$

3.4 Convergence speed

We experimentally compared the convergence speed of the numerical schemes considered. The test was performed on a 128×128 region, with gradient data

$$\begin{aligned} p_{i,j} &= 0 & \forall i, j \\ q_{i,j} &= 0 & \forall i, j \end{aligned}$$

and initial values

$$z_{i,j}^0 = \begin{cases} 0 & (i, j) \neq (i_0, j_0) \\ 1 & (i, j) = (i_0, j_0) \end{cases}$$

The results are shown in Table 3. Note that SOR is based on Gauss-Seidel. FMG refers to a three-level pyramid, with SOR at the coarsest level and Gauss-Seidel at the other levels. Time measurements refer to a 500MHz Pentium 3 computer, with 128MB RAM and double precision calculations. Note that time (rather than the number of iterations) is used for comparison, since FMG operates on grids of varying sizes.

Method	Time (sec)
Jacobi	374
Gauss-Seidel	189
SOR	5.4
FMG	0.8

Table 3: Convergence time comparison

4 Experiment

Since the popularization of photometric stereo is an important motivation for this research, we demonstrate the results using a very inexpensive experimental setup, assembled at home. The camera used was a Sony Mavica MVC-FD71 that acquires 640×480 still images in (lossy)



Figure 17: The imaged object and the calibration accessories: the illumination direction is derived from the position of the shadow of the vertical rod on the lego wall.

JPEG format on a floppy disk. The (lossless) BMP storage option was not selected, since it would have required to replace the floppy disk between images, thus moving the camera. The illumination source was a clear pear shaped E27 40W light bulb. It is obviously neither a point source nor a parallel source. The illumination direction was measured via the sundial principle, using the shadow of a small plastic rod on a lego wall, see Fig. 17. Control points were obtained using a Clearline Concepts CL2042 keychain laser-pointer, with a tip that created a complex dot pattern. Their depth was derived by comparing their position on the scene to their position on the wall behind, with the imaged object removed. The object used was a plastic mannequin head; its true 3D shape was acquired in advance using a high precision Cyberware laser-based range scanner⁷

Many error sources come to mind, most of them directly related to the intentionally primitive experimental setup. In particular, since the illumination source was weak, it had to be positioned close to the object. Thus, the illumination angle varied by more than five degrees over the object, in contrast to the parallel illumination assumed. Obviously, the illumination intensity was not uniform either. Errors were also due to JPEG loss, self shadows and deviations from the Lambertian surface model, especially in the eyes.

⁷Registration between the range data and the photometric stereo reconstructions was obtained by least-squares optimization of the affine transformation between corresponding landmarks.

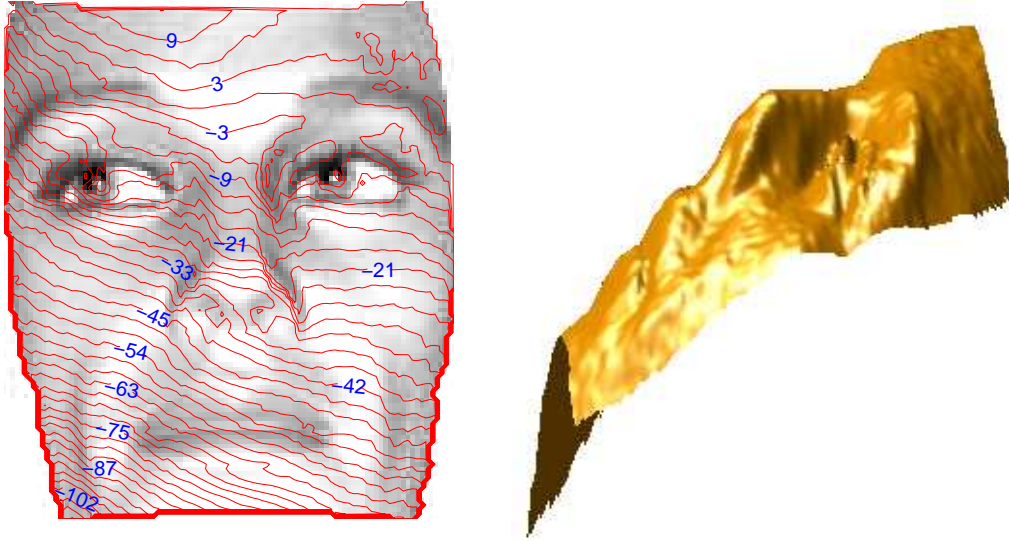


Figure 18: *Left:* The difference between straightforward, uncompensated photometric stereo surface reconstruction, and ground truth data for the same surface obtained with the Cyberware range scanner. *Right:* Visualization of the uncompensated photometric stereo reconstruction. The bias is evident.

These errors lead to bias in 3D reconstruction of the scene using photometric stereo. Fig. 18 (left) shows the difference between the photometric stereo reconstruction and the ground-truth obtained using the Cyberware range scanner (Fig. 1-left). Fig. 18 (right) is a visualization of the photometric stereo reconstruction itself. It can be seen that the bias is significant and quite smooth. Using the laser pointer, we obtained 20 control points and reconstructed the scene using both the *WLS* method and the interpolation surface technique. Fig. 19 (left and right) show the difference between the reconstruction and the ground truth, with the *WLS* and interpolation surface techniques respectively. The corresponding control-point corrected photometric stereo reconstructions are visualized in Fig 20. With both methods the bias is dramatically reduced. In fact it is nearly eliminated, with residual bias remaining mostly near the boundary of the domain when using the *WLS* technique.

5 Conclusions

The incorporation of sparse control points as constraints in recovering depth from the gradient field is not as simple as it may initially seem. The over-determined characteristic of computing depth from a gradient field makes the reconstructed surface extremely rigid. En-




Figure 19: The difference between control-point corrected photometric stereo reconstructions and the ground truth. *Left*: Bias correction with depth constraints and weighted orientation compliance via WLS. *Right*: Bias correction using an interpolation surface. In both cases the bias is greatly reduced. Control point locations are indicated by the + signs.




Figure 20: Control-point corrected photometric stereo reconstructions. *Left*: Bias correction with depth constraints and weighted orientation compliance via WLS. Residual bias remains only near the boundary of the domain. *Right*: Bias correction using an interpolation surface. Compare with Fig. 18 (right).

forcing pointwise depth constraints leads in the continuous case to pointwise singularities and in the discrete case to tent-pole effects, rather than to global bias correction.

We propose two methods for bias correction using control points. One is based on constrained weighted least squares extension to depth from gradient field computation. It forces the reconstruction to comply with the gradient field in the vicinity of control points. This has the desired effect of maintaining the smoothness of the surface, and distributing the information provided by a control point away from its own location.

An alternative approach decouples the recovery of the fine surface details from global bias correction. A smooth interpolation surface is added to the usual depth from gradient field reconstruction, thus ensuring that the control point constraints are satisfied.

Experimental results are provided for both methods. To allow comparisons with “ground truth” data, the surface under test was also acquired using a high accuracy Cyberware laser scanner. It is seen that both approaches significantly reduce bias and, in contrast to previous methods, reconstruct surfaces without spikes at the control point locations (see subsection 2.6).

There is a significant difference between the two methods presented. With the *WLS* technique, the shape tends to return to its original form (without control points) away from the control points. This is not the case in the interpolation surface approach. For best results with both methods, it is recommended to acquire some of the control points near the boundary of the domain.

An efficient numerical scheme is the key to practical computation of depth from the gradient field. Standard solutions to the Poisson equation are not directly applicable due to the presence of control points and non-rectangular domains. We have adapted an efficient solver to satisfy these requirements and obtained a fast algorithm. In typical examples, with 100×120 data, computing time went down from about 20 minutes (with the standard iterative scheme [14]) to 2-3 seconds, on a Pentium 3 500Mhz computer with 128MB RAM.

A problem left for future research is the determination of the optimal over-relaxation parameter ω in our numerical scenario. This is important, since successive over-relaxation (SOR) is most efficient in a fairly narrow region around the optimal ω value.

The methods developed in this research have been motivated and presented in connection with photometric stereo, but are not limited to that context. They can be applied wherever sparse depth constraints can be available to augment and correct dense gradient field data,

as in certain shape from texture problems.

References

- [1] F.L. Bookstein, “Principal warps: thin plate splines and the decomposition of deformations”, *IEEE Trans. Pattern Analysis Machine Intelligence*, Vol. 11, pp. 567-585, 1989.
- [2] Y. Chen and R. Sara, “Integration of photometric stereo and shape from occluding contours by fusing orientation and depth data”, Technical report CITR-TR-68, Computer Science Department, University of Auckland, New Zealand, 2000.
- [3] J.E. Cryer, P.S. Tsai and M. Shah, “Integration of shape from shading and stereo”, *Pattern Recognition*, Vol. 28, pp. 1033-1043, 1995.
- [4] M.S. Drew and B.V. Funt, “Variational approach to interreflection in color images”, *J. Opt. Soc. Am. A*, Vol. 9, pp. 1255-1265, 1992.
- [5] O. Drbohlav and R. Šára, “Unambiguous determination of shape from photometric stereo with unknown light sources”, *Proc. IEEE Int. Conf. Computer Vision*, Vol. 1, pp. 581-586, Vancouver, Canada, 2001.
- [6] O. Drbohlav and R. Šára, “Specularities reduce ambiguity of uncalibrated photometric stereo”, *Proc. European Conf. Computer Vision*, Vol. 2, pp. 46-60, Copenhagen, Denmark, 2002.
- [7] M.S. Drew, “Photometric stereo without multiple images”, *SPIE* Vol. 3016, pp. 369-380, 1997.
- [8] M.S. Drew and M.H. Brill, “Color from shape from color: a simple formalism with known light sources”, *Journal of the Optical Society of America-A*, Vol. 17, pp. 1371-1381, 2000.
- [9] J. Fan and L.B. Wolff, “Surface curvature and shape reconstruction from unknown multiple illumination and integrability”, *Computer Vision and Image Understanding*, Vol. 65, pp. 347-359, 1997.
- [10] R. Franke, “Scattered data interpolation: tests of some methods”, *Mathematics of Computation*, Vol. 38, pp. 181-200, 1982.

- [11] R.T. Frankot and R. Chellappa, “A Method for enforcing integrability in shape from shading algorithms”, *IEEE Trans. Pattern Analysis Machine intelligence*, Vol. 10, pp. 439-451, 1988.
- [12] P. Fua and Y.G. Leclerc, “Object-centered surface reconstruction: combining multi-image stereo and shading”, *International Journal of Computer Vision*, Vol. 16, pp. 35-56, 1995.
- [13] B.K.P. Horn, *Robot Vision*, MIT Press, 1986.
- [14] B.K.P. Horn and M.J. Brooks, “The variational approach to shape from shading”, *Computer Vision, Graphics and Image Processing*, Vol. 33, pp. 174-208, 1986.
- [15] B.K.P. Horn, “Height and gradient from shading”, *Int. J. Computer Vision*, Vol. 5, pp. 37-75, 1990.
- [16] K. Ikeuchi, “Reconstructing a depth map from intensity maps”, *Proc. ICPR*, Montreal, Canada, 1984, pp. 736-738.
- [17] K. Ikeuchi, H.K. Nishihara, B.K.P. Horn, P. Sobalvarro and S. Nagata, “Determining grasp configurations using photometric stereo and the PRISM binocular stereo system”, *International Journal of Robotics Research*, Vol. 5, pp. 46-65, 1986.
- [18] R. Klette, K. Schluns and A. Koschan, *Computer Vision: Three-Dimensional Data from Images*, Springer, 1998.
- [19] Y.G. Leclerc and A.F. Bobick, “The direct computation of height from shading”, *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 552-558, Maui, Hawaii, 1991.
- [20] K.M. Lee and C.-C. J. Kuo, “Shape from photometric ratio and stereo”, *Journal of Visual Communication and Image Representation*, Vol. 7, pp. 155-162, 1996.
- [21] R. Onn and A.M. Bruckstein, “Integrability disambiguates surface recovery in two-image photometric stereo”, *International Journal of Computer Vision*, Vol. 5, pp. 105-113, 1990.
- [22] S.K. Nayar and M. Oren, “Generalization of the Lambertian model and implications for machine vision”, *International Journal of Computer Vision*, Vol. 14, pp. 227-251, 1995.

- [23] W.H. Press, S.A. Teukolsky, W.T. Vetterling and B.P. Flannery *Numerical Recipes in C*, 2nd edition, Cambridge University Press, 1992.
- [24] J. Shah, H.H. Pien and J.M. Gauch, “Recovery of surfaces with discontinuities by fusing shading and range data within a variational framework”, *IEEE Trans. Image Processing*, Vol. 5, pp. 1243-1251, 1996.
- [25] M. Shao, R. Chellappa and T. Simchony, “Reconstructing a 3-D depth map from one or more images”, *Computer Vision, Graphics and Image Processing*, Vol. 53, pp. 219-226, 1991.
- [26] I. Shimshoni, Y. Moses and M. Lindenbaum, “Shape reconstruction of 3D bilaterally symmetric surfaces”, *International Journal of Computer Vision*, Vol. 39, pp. 97-110, 2000.
- [27] R. Szeliski, “Fast surface interpolation using hierarchical basis functions”, *IEEE Trans. Pattern Analysis Machine Intelligence*, Vol. 12, pp. 513-528, 1990.
- [28] M. Spiegelman, Myths and Methods in Modeling: section 9, Lamont-Doherty Earth Observatory, Columbia University. <http://www.ldeo.columbia.edu/~mspieg/mmm> .
- [29] D. Terzopoulos, “The role of constraints and discontinuities In visible-surface reconstruction”, *Proc. IJCAI83*, pp. 1019-1022, 1983.
- [30] D. Terzopoulos, *Multiresolution computation of visible-surface representations*, MIT Ph.D. Thesis, 1984.
- [31] D. Terzopoulos, “Multilevel Computational processes for visual surface reconstruction”, *CVGIP*, Vol. 24, pp. 52-96, 1983.
- [32] B.J.H. Verwer, “Local distances for distance transformations in two and three dimensions”, *Pattern Recognition Letters*, Vol. 12, pp. 671-682, 1991.
- [33] G.Q. Wei and G. Hirzinger, “Parametric shape from shading by radial basis functions”, *IEEE Trans. Pattern Analysis Machine Intelligence*, Vol. 19, pp. 353-365, 1997.
- [34] L.B. Wolff and E. Angelopoulou, “Three dimensional stereo by photometric ratios”, *Journal of the Optical Society of America A*, Vol. 11, pp. 3069-3078, 1994.

- [35] R.J. Woodham, “Photometric method for determining surface orientations from multiple images”, *Optical Engineering*, Vol. 19, pp. 139-144, 1980.
- [36] Z. Wu and L. Li, “A line integration based method for depth recovery from surface normals”, *Computer Vision, Graphics and Image Processing*, Vol. 43, pp. 53-66, 1988.