

Least Squares Surface Reconstruction from Measured Gradient Fields

Matthew Harker and Paul O'Leary
University of Leoben, Institute for Automation
A-8700 Leoben, Austria

matthew.harker@stud.unileoben.ac.at

Abstract

This paper presents a new method for the reconstruction of a surface from its x and y gradient field, measured, for example, via Photometric Stereo. The new algorithm produces the unique discrete surface whose gradients are equal to the measured gradients in the global vertical-distance least-squares sense. We show that it has been erroneously believed that this problem has been solved before via the solution of a Poisson equation. The numerical behaviour of the algorithm allows for reliable surface reconstruction on exceedingly large scales, e.g., full digital images; moreover, the algorithm is direct, i.e., non-iterative. We demonstrate the algorithm with synthetic data as well as real data obtained via photometric stereo. The algorithm does not exhibit a low-frequency bias and is not unrealistically constrained to arbitrary boundary conditions as in previous solutions. In fact, it is the first algorithm which can reconstruct a surface of polynomial degree two or higher exactly. It is hence the first viable algorithm for online industrial inspection where real defects (as opposed to phantom defects) must be identified in a robust manner.

1. Introduction

Photometric stereo provides a means of measuring the x and y gradient field of a surface at discrete points. Inasmuch, it holds great promise for industrial applications such as geometric surface inspection. What is rarely discussed in photometric stereo literature, is the reconstruction of the surface from said gradient measurements. Often, analysis is attempted based on the gradients alone. Existing algorithms provide reconstructions with significant systematic (as opposed to stochastic) errors, which make them unusable for industrial applications where real (and not phantom) flaws must be detected in a robust manner. It is for this reason that photometric stereo has remained a rather theoretical nicety, and not yet found its way prolifically into industrial environments. In this paper, we present the unique solution to least squares surface reconstruction from its measured gradient.

Previous methods in the literature work with either local or global cost functions, and can be summarized as follows. Line integral methods [12], entail choosing an arbitrary height at an arbitrary point, then evaluating and averaging (arbitrarily chosen) line integrals to compute the height at a nearby point. The underlying principle is that all possible line integrals between two points should be the same. This can be considered to be a local method, and hence the error distribution is highly irregular, starting from zero at the initial point and propagating irregularly outward to the most distant points.

Grid-based basis function methods such as [9, 7]. As demonstrated in the following, introducing basis functions to represent the surface is redundant, and, in fact, only introduces computational error into the reconstruction process. Reconstruction results reflect this numerical ill-conditioning.

The variational approach [6, 4, 10, 2], based on the calculus of variations is the root of the most common algorithms for surface reconstruction from gradient fields. Formulating the cost function as a continuous integral, a stationary point of the integral should satisfy the corresponding Euler-Lagrange differential equation. However, since there is a family of non-unique solutions that satisfy said equation, none can be the least-squares minimizer proper; that is, a true least-squares solution of a linear system is unique for fully-determined problems. This approach is discussed further in comparison to the method presented here in Section 5.

The contributions of the paper are as follows:

- A straightforward derivation of a numerically sound solution to the least-squares surface reconstruction from gradient fields. It has been falsely believed that this has been solved before; however, we show that the new method is the first strictly correct and unique solution. This is first proven, then verified by numerical testing; a simple evaluation of the least squares cost function shows that the previous methods are not least squares minimizers.

- Simple parameterization of the discrete surface as a matrix leads to a computationally efficient solution by avoiding Fourier and Cosine transformations, whilst showing that introducing basis functions to represent the surface is redundant. An appropriate circumstance for the introduction of basis functions is briefly discussed in Section 8.
- Since the surface is discrete, we can do no better than numerical derivatives (i.e. there is no analytical form of the derivatives of an unknown surface). The algorithm presented here allows for numerical derivatives of arbitrary polynomial accuracy (in practice up to 16th order accurate). Previous methods are limited to forward/backward differences which are only first order accurate; hence, the new method is the first that can reconstruct surfaces of degree two and higher.
- We show that most existing methods are based on a subtly false premise, whereby tacit assumptions pertaining to the Euler-Lagrange equation have gone unnoticed by previous authors.

In the following, we formulate the reconstruction problem as that of reconstruction of a discrete surface Z , whose derivatives are equal to the measured derivatives in the least-squares sense. With a matrix definition of the cost function, we make use of common linear algebra to derive the solution of the unique least-squares minimum; this solution is shown to satisfy a Lyapunov equation. It is shown that a previous approach based on solving a Poisson equation has been mistakenly believed to be the least-squares solution to the problem at hand. The above notions are confirmed via numerical testing with synthetic data, as well as real data obtained via photometric stereo.

2. Gradient Measurement via Photometric Stereo

Photometric stereo [11] is a method for measuring the x and y gradient field (equivalently the orientation) of a surface at discrete points. The principal assumption is that light reflects off the surface according to a mathematical reflectance model. Assuming that the image is obtained via an orthographic projection (camera centrepoint at infinity) and the surface has the explicit form $z = z(x, y)$, we define the parameters

$$p \triangleq \frac{\partial z(x, y)}{\partial x} \quad \text{and} \quad q \triangleq \frac{\partial z(x, y)}{\partial y}, \quad (1)$$

such that we may work in a gradient space $(p, q) \in \mathbb{R}^2$. A typical reflectance model in said space is,

$$R(p, q) = \frac{\rho(1 + pp_s + qq_s)}{\sqrt{1 + p^2 + q^2} \sqrt{1 + p_s^2 + q_s^2}}, \quad (2)$$

where (p_s, q_s) is the direction of the light source. This particular model assumes that the surface is Lambertian, although, many other possible reflectance maps exist [11]. Given three images with three different light sources we may solve uniquely for p and q ; however, for improved measurement results this may be generalized to n sources as follows. A set of n images give the following set of equations for each point (x_i, y_j) ,

$$\begin{bmatrix} I_1(j, i) \\ \vdots \\ I_n(j, i) \end{bmatrix} = \begin{bmatrix} s_{1x} & s_{1y} & s_{1z} \\ \vdots & \vdots & \vdots \\ s_{nx} & s_{ny} & s_{nz} \end{bmatrix} \begin{bmatrix} u_x \\ u_y \\ u_z \end{bmatrix} \quad (3)$$

where, (s_{kx}, s_{ky}, s_{kz}) is the direction of the k^{th} light source, $I_k(j, i)$ is the intensity at pixel (j, i) in the k^{th} image, and $[u_x, u_y, u_z]^T$ is the surface normal vector at the point (x_i, y_j) . This equation, written as,

$$\mathbf{I} = \mathbf{S}\mathbf{u}, \quad (4)$$

may be solved in the least squares sense via the Moore-Penrose pseudoinverse of \mathbf{S} ,

$$\mathbf{u} = \mathbf{S}^+ \mathbf{I}, \quad (5)$$

The measured surface gradients are consequently given as

$$p = \frac{u_x}{u_z} \quad \text{and} \quad q = \frac{u_y}{u_z}. \quad (6)$$

Taken over the whole image, we have the following measured quantities in the form of $m \times n$ matrices,

$$\hat{Z}_x(j, i) \triangleq \left. \frac{\partial z}{\partial x} \right|_{(x, y) = (x_i, y_j)} \quad (7)$$

$$\hat{Z}_y(j, i) \triangleq \left. \frac{\partial z}{\partial y} \right|_{(x, y) = (x_i, y_j)} \quad (8)$$

for $i = 1, \dots, n, j = 1, \dots, m$, obtained from p and q respectively. Note that for the sake of consistency the indices i and j appear juxtaposed so as to define a right-handed coordinate system. The advantage of using more than three light sources is the suppression of noise in the gradient measurements; furthermore, gross outliers such as saturated pixels may be omitted from the measurements.

3. Least-Squares Surface Reconstruction

The premise of the solution is that we have the measured gradient of a surface, that is, we have a gradient field that is corrupted by noise. Assuming the noise is Gaussian, the appropriate cost function is the least-squares cost function, which in its continuous form is,

$$\begin{aligned} \epsilon &= \int_c^d \int_a^b (\hat{z}_x(x, y) - z_x(x, y))^2 \\ &+ (\hat{z}_y(x, y) - z_y(x, y))^2 dx dy, \end{aligned} \quad (9)$$

where \hat{z}_x and \hat{z}_y represent the measured gradient. This cost function is the volume of the squared differences of the functions; hence, the equivalent discrete form of the cost function over a rectangular grid is,

$$\epsilon = \left\| \hat{Z}_x - Z_x \right\|_F^2 + \left\| \hat{Z}_y - Z_y \right\|_F^2, \quad (10)$$

where the subscript ‘‘F’’ denotes the Frobenius norm [5]. Since differentiation is a linear operator, we may write the derivatives as a simple matrix multiplication (i.e. a linear transformation),

$$Z_x = ZL_x^T \quad \text{and} \quad Z_y = L_y Z, \quad (11)$$

where L_x and L_y are respectively $n \times n$ and $m \times m$ matrices such that Z is generally $m \times n$. With this substitution, we now have the cost function parameterized in terms of the unknown surface, Z , i.e.,

$$\epsilon = \left\| \hat{Z}_x - ZL_x^T \right\|_F^2 + \left\| \hat{Z}_y - L_y Z \right\|_F^2. \quad (12)$$

By definition of the Frobenius norm, the cost function reads,

$$\begin{aligned} \epsilon = & \text{trace} \left\{ \left(\hat{Z}_x - ZL_x^T \right) \left(\hat{Z}_x - ZL_x^T \right)^T \right\} \\ & + \text{trace} \left\{ \left(\hat{Z}_y - L_y Z \right) \left(\hat{Z}_y - L_y Z \right)^T \right\}, \quad (13) \end{aligned}$$

which upon expanding yields,

$$\begin{aligned} \epsilon = & \text{trace} \left\{ \hat{Z}_x \hat{Z}_x^T - \hat{Z}_x L_x Z^T - ZL_x^T \hat{Z}_x^T + ZL_x^T L_x Z^T \right\} \\ & + \text{trace} \left\{ \hat{Z}_y \hat{Z}_y^T - \hat{Z}_y Z^T L_y^T - L_y Z \hat{Z}_y^T + L_y Z Z^T L_y^T \right\} \quad (14) \end{aligned}$$

To find the minimum of the cost function, we differentiate with respect to the matrix Z and equate to zero, yielding the following $m \times n$ matrix equation,

$$L_y^T L_y Z + ZL_x^T L_x - L_y^T \hat{Z}_y - \hat{Z}_x L_x = 0. \quad (15)$$

This equation, typically known as a Lyapunov Equation, is clearly linear in the elements of Z . There exist stable solutions for Z based on orthogonal matrices [3]. The solution is unique, provided that the eigenvalues, α_i of $L_y^T L_y$ and β_j of $L_x^T L_x$, satisfy $\alpha_i + \beta_j \neq 0$ for all pairs (i, j) . For Equation (15) as it stands, this condition does not hold; however, we do have a-priori knowledge of the system which allows us to find the unique solution to the equation. Firstly, we cannot solve for the constant of integration, which eliminates a degree of freedom from Z , meaning the actual problem has fewer unknowns. Secondly, the derivative matrices L_x and L_y are rank-1 deficient, but the span of their null-spaces is known analytically. Specifically, a constant function has a zero derivative, which means that,

$$L_x \mathbf{1} = \mathbf{0}, \quad (16)$$

and hence the vector $\mathbf{1}$ spans the null-space of L_x . The Lyapunov equation has a more convenient form if the first row and first column of both $L_y^T L_y$ and $L_x^T L_x$ are all zeros. This can be accomplished by transforming the null-vector $\mathbf{1}$ to the vector $[1 \ 0 \ \dots \ 0]^T$. Specifically, the Householder reflection [5],

$$P_x = I_{n \times n} - \frac{2\mathbf{v}\mathbf{v}^T}{\mathbf{v}^T \mathbf{v}}, \quad (17)$$

where,

$$\mathbf{v} \triangleq [1 + \sqrt{n} \quad 1 \quad \dots \quad 1]^T \quad (18)$$

will transform the matrix L_x to a matrix of the form,

$$\hat{L}_x \triangleq L_x P_x = [\mathbf{0} \quad M], \quad (19)$$

and similarly a matrix P_y yields \hat{L}_y . Since the Householder matrices P_x and P_y are orthogonal, they are perfectly conditioned and do not affect the numerical accuracy of the solution. Hence, we premultiply Equation (15) by P_y^T , postmultiply by P_x , and make the substitutions $P_x P_x^T = I$ and $P_y P_y^T = I$, yielding,

$$\begin{aligned} P_y^T L_y^T L_y P_y P_y^T Z P_x + P_y^T Z P_x P_x^T L_x^T L_x P_x \\ - P_y^T L_y^T \hat{Z}_y P_x - P_y^T \hat{Z}_x L_x P_x = 0. \quad (20) \end{aligned}$$

Making the substitution $W = P_y^T Z P_x$, yields a Lyapunov Equation in W ,

$$\hat{L}_y^T \hat{L}_y W + W \hat{L}_x^T \hat{L}_x - \hat{L}_y^T \hat{Z}_y P_x - P_y^T \hat{Z}_x \hat{L}_x = 0. \quad (21)$$

This equation can be shown to have the form,

$$\begin{aligned} \begin{bmatrix} 0 & \mathbf{0}^T \\ \mathbf{0} & A \end{bmatrix} \begin{bmatrix} w_{00} & \mathbf{w}_{01}^T \\ \mathbf{w}_{10} & W_{11} \end{bmatrix} + \begin{bmatrix} w_{00} & \mathbf{w}_{01}^T \\ \mathbf{w}_{10} & W_{11} \end{bmatrix} \begin{bmatrix} 0 & \mathbf{0}^T \\ \mathbf{0} & B \end{bmatrix} \\ + \begin{bmatrix} 0 & \mathbf{c}_{01}^T \\ \mathbf{c}_{10} & C \end{bmatrix} = \begin{bmatrix} 0 & \mathbf{0}^T \\ \mathbf{0} & 0 \end{bmatrix}, \quad (22) \end{aligned}$$

which is equivalently the following system of equations,

$$\mathbf{w}_{01}^T B + \mathbf{c}_{01}^T = \mathbf{0}^T \quad (23)$$

$$A \mathbf{w}_{10} + \mathbf{c}_{10} = \mathbf{0} \quad (24)$$

$$A W_{11} + W_{11} B + C = 0. \quad (25)$$

Equation (25) is again a Lyapunov equation, but now of dimension $(m-1) \times (n-1)$; the two remaining equations are simple linear equations which can be trivially solved since A and B are invertible¹. Note that, quite correctly, we cannot solve for one parameter in the set of equations, i.e. the

¹Invertibility can be proven as follows: for an appropriate derivative matrix, only the derivative of a constant function should vanish, i.e. $L_x \mathbf{1} = \mathbf{0}$ and $\text{rank} \{L_x\} = n-1$. Since P_x is orthogonal, $\text{rank} \{L_x P_x\} = n-1$ and so $\text{rank} \{[\mathbf{0} \ M]\} = n-1$. Finally, this implies $\text{rank} \{M\} = n-1$, and since M is $(n-1) \times (n-1)$, $A = M^T M$ is invertible.

constant of integration (w_{00} can be set to zero). Solving Equations (23) – (25) for the elements of W and recomposing the matrix, we obtain the reconstructed surface as,

$$Z = P_y W P_x^T. \quad (26)$$

Since the matrices A and B are by necessity positive definite, the conditions for the uniqueness of the Lyapunov equation in (25) hold, and hence Z is the unique minimizing surface.

It can be shown that the missing degree of freedom is the constant of integration. Recall that $L_x \mathbf{1} = \mathbf{0}$ and $L_y \mathbf{1} = \mathbf{0}$. Consequently, if Z is a solution to Equation (15), then so is $Z(\alpha) = Z + \alpha \mathbf{1} \mathbf{1}^T$. Substituting yields

$$L_y^T L_y (Z + \alpha \mathbf{1} \mathbf{1}^T) + (Z + \alpha \mathbf{1} \mathbf{1}^T) L_x^T L_x - L_y^T \hat{Z}_y - \hat{Z}_x L_x = 0, \quad (27)$$

and clearly all terms pertaining to the constant of integration, α , vanish.

4. Matrix Based Numerical Differentiation

Note that in the preceding derivation, there was no assumption as to the particular form of the derivatives beyond the fact that they are linear transformations. This generalization enables the use of numerical derivative formulas that are accurate to any particular polynomial degree. All previous methods [4, 10, 2] have used simple forward/backward differences, which are only first-order accurate (i.e. exact only for linear functions). Here we briefly describe the derivation of N -point derivative formulas, which yield exact derivatives for functions up to degree $d = N - 1$.

For equally spaced nodes, we expand the function in a Taylor series of degree $N - 1$ about the point x_k in the N -point sequence. This yields $N - 1$ equations of the form,

$$f(x_j) - f(x_k) = f'(x_k)(k - j)h + \dots + \frac{f^{(N-1)}(x_k)(k - j)^{(N-1)}h^{(N-1)}}{(N - 1)!} \quad (28)$$

for $j = 1, \dots, N$, $j \neq k$, which are linear in the unknown derivatives $f'(x_k)$ through $f^{(N-1)}(x_k)$. Using Cramer's rule, we may solve for $f'(x_k)$ in this system of equations. Applying this approach for $k = 1, \dots, N$ yields an N point derivative formula for each point in the N -point sequence of function evaluations, $f(x_k)$, $k = 1, \dots, N$. Thus, numerical differentiation takes the form of a matrix multiplication, e.g. $Z_y = L_y Z$. If the N^{th} derivative of the function vanishes, then the Taylor expansion is exact, and hence so are the computed derivatives. For example, for a five point sequence, the corresponding three-point derivatives

with $h = 1$ read,

$$\begin{bmatrix} f'(x_1) \\ f'(x_2) \\ f'(x_3) \\ f'(x_4) \\ f'(x_5) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} -3 & 4 & -1 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 \\ 0 & -1 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 & 1 \\ 0 & 0 & 1 & -4 & 3 \end{bmatrix} \begin{bmatrix} f(x_1) \\ f(x_2) \\ f(x_3) \\ f(x_4) \\ f(x_5) \end{bmatrix}. \quad (29)$$

Note the end-point formulas; the derivatives are degree-two accurate at all five points. Note that if N is chosen to be odd, we yield a symmetric formula for the internal points in the sequence.

5. On the Variational Approach to Surface Reconstruction

An approach to surface reconstruction based on the calculus of variations was first addressed by Horn and Brooks [6]. The premise is that if the continuous form of the cost function in Equation (9) has a stationary point, then it satisfies its corresponding Euler-Lagrange equation,

$$\nabla^2 z(x, y) = \frac{\partial}{\partial x} \hat{z}_x(x, y) + \frac{\partial}{\partial y} \hat{z}_y(x, y), \quad (30)$$

which in this case is a Poisson equation. The trouble with formulating the problem as a differential equation of this form, is that it does not uniquely specify the solution. This can be seen by discretizing Equation (30) while, as before, taking differentiation to be a linear operator, i.e.

$$L_y L_y Z + Z L_x^T L_x^T - L_y \hat{Z}_y - \hat{Z}_x L_x^T = 0. \quad (31)$$

Although this equation is strikingly similar to Equation (15), it is the subtleties which are its undoing. It should be noted that discretizing the Euler-Lagrange equation in this form is essentially new. Any alternative approach leads to a problem of integrability [4], whereby a constraint upon the solution is required. Whereas this new form solves the integrability problem, it unfortunately demonstrates that any approach based on the Euler-Lagrange equation is only valid for a few special cases. Clearly the matrix $L_y L_y$ is a second derivative operator and hence must be at least rank-two deficient (i.e. second derivatives of both linear and constant terms must vanish, hence the null-space is two-dimensional). This implies that in Equation (31) there are at least four unknowns which we cannot solve for, and hence there is a parametric family of non-unique solutions. A specific solution in this case is usually found by imposing boundary conditions. For example, the Frankot-Chellappa method [4], and similarly Kovessi's method [8] both assume the surface is periodic, which is a largely unrealistic assumption for a real surface.

There is, however, a case in which Equations (31) and (15) are identical. If both L_x and L_y are skew-symmetric, then

we may substitute $L_x = -L_x^T$, $L_y = -L_y^T$ into Equation (31) and multiply through by -1 to yield Equation (15). Skew-symmetric derivative matrices are indeed valid, but only for the special case of periodic functions. The underlying assumption of the variational approach is therefore that the surface is periodic. That is, methods relying on the Euler-Lagrange equation and forward difference formulas can only yield exact reconstruction for surfaces that are of degree-one or less, and periodic; the only such surface is $Z = \alpha \mathbf{1}\mathbf{1}^T$. Such methods include the Poisson solver [10], and the “continuum of solutions” proposed by Agrawal [2]. The Poisson method actually overcomes the lack of skew-symmetry near the edges by padding the measured derivatives with zeros and simply ignoring the accuracy of the derivatives near the boundary. It can therefore reconstruct surfaces up to degree-one, i.e. planes. When working with real data, however, the fact the cost function is global means this error at the boundary propagates throughout the solution. It should be noted that all solutions based on this “variational approach” tend to exhibit systematic, as opposed to stochastic, error in their reconstructions. This can be attributed to the fact that they are non-unique solutions to an equation which is only valid in a few special cases.

6. Conditions Yielding Exact Reconstruction

In stark contrast to previous solutions, there is a range of conditions under which the algorithm proposed here yields exact reconstruction up to a constant of integration. Specifically, these are: Polynomial surfaces of degree $d \leq N - 1$, where N is the number of points used in the derivative formulas; and transcendental functions with rapidly decaying derivatives, in which case the reconstruction is numerically exact (i.e. the absolute errors are close to machine accuracy).

7. Numerical Testing

To demonstrate the new approach in comparison to previous solutions, we propose the following tests:

1. Reconstruction of a surface with piecewise continuous derivatives.
2. Reconstruction of an analytic, but non-polynomial, surface in the presence of synthetic noise. We present a Monte-Carlo simulation for noise standard deviations ranging from 0 – 5% of the data amplitude.
3. Reconstruction from real photometric stereo data. We compare the results to “ground truth” data obtained with a laser scanner.

For the previous solutions, we have made use of the code available from [1]. Note that all previous methods only reconstruct up to an offset as well as overall scale; therefore,

for the purpose of comparison, we have computed the least-squares scale and offset to evaluate cost functions, etc. We have done the same with the new results in spite of the fact the overall scale is one to one up to a factor of noise. For the first test, we have computed the analytic derivatives of the function,

$$z(x, y) = \max \left[\frac{1}{20} (x + y), \sqrt{x^2 + y^2 - 1} \right], \quad (32)$$

which is the maximum height between a plane and a sphere. The derivatives are therefore only piecewise continuous (i.e. discontinuous). The closest comparable method to the Frankot-Chellappa method is to use three-point derivatives with the proposed algorithm; results are shown in Figure 1. The new method is reasonably close to a perfect reconstruc-




Figure 1. (TOP-LEFT) the proposed solution with three-point derivatives (TOP-RIGHT) the Frankot-Chellappa method (which has been post-scaled). (BOTTOM) The residual surface of the new method. Note the Gibbs Phenomenon near the discontinuities. This is due the fact that polynomial approximation is implicit in the use of numerical derivatives, whereby the phenomenon arises when making continuous approximations of discontinuities.

tion, with the exception of the Gibbs phenomena that occur near the discontinuities, which can be seen in the residual surface. This can be attributed to the fact that polynomial approximation is implicit in the use of numerical differentiation; the Gibbs phenomenon arises when approximations are made to functions with discontinuities via basis functions. The Frankot-Chellappa method clearly has difficulty

to account for the simple, albeit small, gradient. The fact that the solution is constrained to be periodic leads to a systematic bias in the solution. Such a small deviation might represent a misaligned workpiece in an industrial environment; such a solution is therefore unsuitable for real data. The second experiment is to reconstruct a transcendental surface of the form,

$$z(x, y) = \sum_{k=1}^n A_k \exp\left(-\frac{(x - x_k)^2 + (y - y_k)^2}{a_k}\right). \quad (33)$$

Since the function is infinitely differentiable, its numerical derivatives will only tend to its analytic derivatives for large N . Figure 2 shows the reconstructions using the new algorithm with three-point derivatives and the Poisson solver; the gradients have added synthetic Gaussian noise with a standard deviation of 2% of the amplitude. The residual surfaces of these particular reconstructions are shown in Figure 3. Note that the residual surface for the new method is largely stochastic in nature; this is to be expected of a proper least squares solution of data subjected to Gaussian noise. In contrast, the residual surface of the Poisson solution is largely systematic in nature, which is typical of when a geometric model (in this case the Poisson equation) is inappropriate for modeling the data. A Monte-Carlo simulation of the reconstruction is shown for noise ranging from $\sigma = 0 - 5\%$, comparing the Poisson solver, the new method with three-point derivatives, as well as the new method with eleven-point derivatives. The relative reconstruction error and cost function evaluations for the simulation are shown in Figure 4. This shows that even for exact data, the Poisson solver has a relative reconstruction error upwards of 8%. For the same exact data, the three-point solution has a relative error less than 1%, whereas the eleven point solution is (numerically) exact. Note that the error for the new solution has a linear upward trend, which is expected from a proper least-squares solution. The cost function evaluations confirm that the Poisson solution is not the least-squares minimizer. Note that the cost function evaluations for the new method are smaller than the actual noise added to the exact gradients. This can be attributed to the fact that the noise is, to some degree, integrable.

The final experiment is to reconstruct a surface from the gradient field measured via photometric stereo, shown in Figure 5. Figure 6 shows the result of the reconstruction using the new method, along with a laser scanned depth map of the actual surface. The results show that the least squares solution from the photometric data produces a credible reconstruction of a surface with significant depth variation.

8. Outlook

Future work is to extend the method to incorporate linear filtering simultaneously in the reconstruction algorithm.

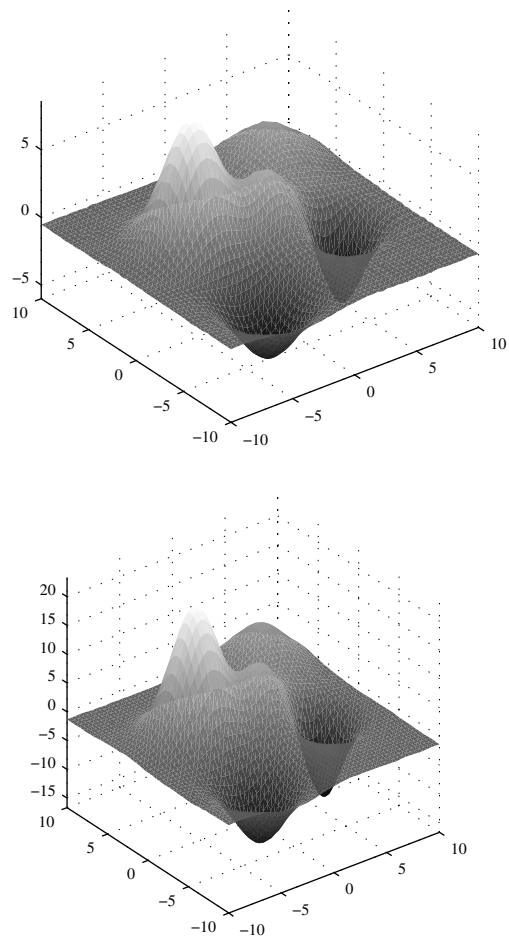


Figure 2. (TOP) Surface reconstruction with the proposed solution (BOTTOM) Reconstruction via the Poisson solver. The new solution is correctly scaled whereas in the Poisson the overall scaling is arbitrary. Indeed the two surfaces appear similar, but are systematically different.

This would, in fact, reduce the computational load, while suppressing noise at the same time. This entails modeling the surface as, $Z = YMX^T$, where the matrices X and Y are sampled basis functions, and M is the matrix of associated moments (or the spectrum of Z). For example, if the basis functions were chosen to be complex exponentials (a Fourier basis), then we could incorporate Fourier transform (i.e., FFT) based filtering into the reconstruction. Although it would be equivalent to post-processing, combining the two steps would be more efficient since it leads to a Lyapunov equation of smaller dimension.

9. Conclusions

In this paper, we presented the true global least-squares solution to the reconstruction of a surface from its gradient




Figure 3. Residual surfaces with respect to the analytic surface which have been reconstructed from gradients with 2% synthetic noise for (LEFT) the proposed solution with three-point derivatives and (RIGHT) the Poisson solver. Note that the residuals of the new solution are largely stochastic in nature, as to be expected of an appropriate least squares solution, whereas the residuals of the Poisson solution are systematic in nature.




Figure 4. (LEFT) Relative error of the reconstructed surface with respect to the noise level for Poisson (—), the new method with 3-point derivatives (— —), and the new method with 11-point derivatives (—). (RIGHT) The cost function evaluations with respect to the noise level. The line (· · ·) represents the actual error added to the error free gradients.




Figure 5. Gradient of the end of a steel slab with a crack measured via photometric stereo.

field. It had been erroneously believed that this problem has been solved before. This fact was proven, then confirmed via numerical testing. The new method is the first viable algorithm for photometric stereo in industrial environments.

References

[1] A. Agrawal. Robust surface reconstruction from 2D gradient fields. <http://www.cfar.umd.edu/~aagrwal>,




Figure 6. (LEFT) The depth map of a surface reconstructed from real photometric stereo data using the new method. (RIGHT) A laser range scan of the same surface. Note in the range scan the depth data of the crack is missing; the reconstruction indicates the crack depth is small with respect to the overall height.

2006. 5

- [2] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In *European Conference on Computer Vision*, volume 3951/2006, pages 578–591, Graz, Austria, 2006. LNCS. 1, 4, 5
- [3] R. Bartels and G. Stewart. Solution of the matrix equation $AX+XB=C$. *Communications of the ACM*, 15(9):820–826, 1972. 3
- [4] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10(4):439–451, 1988. 1, 4
- [5] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, third edition, 1996. 3
- [6] B. Horn and M. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33:174–208, 1986. 1, 4
- [7] B.-H. Kim and R.-H. Park. Shape from shading and photometric stereo using surface approximation by Legendre polynomials. *Computer Vision and Image Understanding*, 66(3):255–270, 1997. 1
- [8] P. Kovessi. Shapelets correlated with surface normals produce surfaces. In *10th IEEE International Conference on Computer Vision*, pages 994–1001, Beijing, 2005. 4
- [9] K. Lee and C.-C. Kuo. Surface reconstruction from photometric stereo images. *J. Opt. Soc. Am. A*, 10(5):855–868, 1993. 1
- [10] T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving Poisson equations in computer vision problems. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(5):453–446, 1990. 1, 4, 5
- [11] R. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980. 2
- [12] Z. Wu and L. Li. A line-integration method for depth recovery from surface normals. In *9th International Conference on Pattern Recognition*, volume 1, pages 591–595, Rome, Italy, 1988. 1