

An Algebraic Approach to Surface Reconstruction from Gradient Fields

Amit Agrawal, Rama Chellappa, Ramesh Raskar

TR2005-145 October 2005

Abstract

Several important problems in computer vision such as Shape from Shading (SFS) and Photometric Stereo (PS) require reconstructing a surface from an estimated gradient field, which is usually non-integrable, i.e. have non-zero curl. We propose a purely algebraic approach to enforce integrability in discrete domain. We first show that enforcing integrability can be formulated as solving a single linear system $Ax=b$ over the image. In general, this system is under-determined. We show conditions under which the system can be solved and a method to get to those conditions based on graph theory. The proposed approach is non-iterative, has the important property of local error confinement and can be applied to several problems. Results on SFS and PS demonstrate the applicability of our method.

IEEE International Conference on Computer Vision (ICCV)

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

An Algebraic Approach to Surface Reconstruction from Gradient Fields

Amit Agrawal, Rama Chellappa
Center for Automation Research
University of Maryland
College Park, MD 20742 USA
(aagraval,rama)@cfar.umd.edu

Ramesh Raskar
Mitsubishi Electric Research Labs
201 Broadway
Cambridge, MA 02139 USA
raskar@merl.com

Abstract

Several important problems in computer vision such as Shape from Shading (SFS) and Photometric Stereo (PS) require reconstructing a surface from an estimated gradient field, which is usually non-integrable, i.e. have non-zero curl. We propose a purely algebraic approach to enforce integrability in discrete domain. We first show that enforcing integrability can be formulated as solving a single linear system $Ax = b$ over the image. In general, this system is under-determined. We show conditions under which the system can be solved and a method to get to those conditions based on graph theory. The proposed approach is non-iterative, has the important property of local error confinement and can be applied to several problems. Results on SFS and PS demonstrate the applicability of our method.

1. Introduction

The notion of integrability arises whenever a surface has to be reconstructed from a gradient field. In several core computer vision problems such as Shape from Shading and Photometric Stereo, an estimate of the gradient field is available. The gradient field is then integrated to obtain the desired 2D surface (shape). However, the estimated gradient field often has non-zero curl making it non-integrable. In this paper, we address the problem due to curl and present a method to enforce integrability. The approach is non-iterative and has the important property that the errors due to non-zero curl do not propagate across the image.

We present an algebraic approach where integrability is enforced by finding a residual gradient field. Adding the computed residual gradients to the specified gradients produces an integrable gradient field. Our main contributions are as follows:

- We present a method to exploit the information contained in the curl of the given non-integrable field. In discrete domain, the residual gradient field and the curl

form a linear system which we use to achieve a better reconstruction.

- When curl is non-zero at several pixels, the linear system is potentially under-determined with more unknowns (residual gradient values) than the number of equations (corresponding to known curl values). Using a graph analogy, we derive conditions under which the residual field can be recovered and show a method to achieve those conditions.
- Unlike least square approaches, our approach has the property of local error confinement and the error in reconstructed surface does not spread throughout the gradient field. Due to its global nature, our method is non-iterative compared to iterative techniques [9] which may have convergence issues.

1.1. Related work

Researchers have addressed the issue of enforcing integrability typically specific to the problem at hand. In **Shape from Shading** algorithms such as [17][8], integrability was enforced as a constraint in the minimization routine. Frankot & Chellappa [6] enforce integrability by orthogonally projecting the non-integrable field on to a vector subspace spanning the set of integrable slopes. However, their method is dependent on the choice of basis functions. Simchony et. al. [11] find the integrable gradient field closest to the given gradient field in a least squares sense by solving the Poisson equation. One can show that (see section 2) their method ignores the information in the curl and finds a zero-curl field which has the same divergence as the given non-integrable field. The method also lacks the property of error confinement. For a survey of SFS algorithms see [16].

Photometric stereo [13][5] uses multiple images obtained under different illumination directions to recover the surface gradients. In [9], belief propagation in graphical networks was used to enforce integrability for SFS and PS problems. In [14], the integrability constraint was used to




Figure 1. Photometric Stereo on Yale database (a) 4 out of 64 input images (b) Reconstructed depth map using SCS method [11]. Global distortions are present (c) Curl from estimated gradient field (d) Reconstructed depth map using our method. Our method exploits information in curl often ignored in gradient reconstruction. This brings in high gradients that are smoothed out in SCS method. In addition, the reconstruction has a local error confinement property so that errors due to curl do not create global distortions.

remove the ambiguity in the estimation of shape and albedo from multiple images. Calibration information about the illumination geometry is obtained by imposing integrability constraint in [3]. In [18], the integrability constraint was used along with rank constraints in a minimization routine to estimate surface albedo and normals. Other examples include [4][15]. The idea of enforcing integrability has also been used in flow field visualization [12][10] by decomposing the given field into curl-free and divergence-free parts.

We show the effectiveness of our technique in the context of the above two problems. The rest of the paper is organized as follows. In Section 2, we formally state the integrability problem. Section 3 describes the discrete domain formulation for enforcing integrability in terms of solving a linear system. Section 4 provides a graph analogy, and derive conditions under which the system can be solved. We then show a method to reach these conditions. Results on Shape from Shading and Photometric Stereo are presented in Section 5.

2. Problem definition

Let $S(x, y)$ be a 2D real valued scalar function defined on a rectangular grid $\{x = 0, \dots, W - 1; y = 0, \dots, H - 1\}$. Let $(p^0 = \frac{\partial S}{\partial x}, q^0 = \frac{\partial S}{\partial y})$ denote the integrable gradient field of S . Given proper boundary conditions, the surface S can be exactly recovered by integrating the gradient field (p^0, q^0) without any error by solving a Poisson equation. Let div denotes the divergence operator and curl denotes the curl operator on gradient field (p, q) , i.e.

$$\begin{aligned} L &= \text{div}(p, q) = \frac{\partial p}{\partial x} + \frac{\partial q}{\partial y} = p_x + q_y \\ C &= \text{curl}(p, q) = \frac{\partial p}{\partial y} - \frac{\partial q}{\partial x} = p_y - q_x \end{aligned} \quad (1)$$

Let L^0 and C^0 denotes the divergence and curl of (p^0, q^0) . Thus $C^0 = 0$. Often, an estimate of the gradient field (p, q) is known, which is non-integrable. The goal is to integrate the gradient field (p, q) to obtain the surface (shape).

A quick review of enforcing integrability using Simchony, Chellappa and Shao's (referred to as the SCS) method described in [11] is as follows. Given a non-integrable gradient field (p, q) , the SCS method finds the surface \hat{S} which minimizes the following least square cost function

$$J(\hat{S}) = (\hat{S}_x - p)^2 + (\hat{S}_y - q)^2 \quad (2)$$

The Euler-Lagrange equation gives the Poisson equation: $\nabla^2 \hat{S} = \text{div}(p, q)$. The integrable field is found by differentiating the estimated surface \hat{S} . Thus (\hat{S}_x, \hat{S}_y) is the corresponding integrable field. One should notice that $\text{div}(\hat{S}_x, \hat{S}_y) = \frac{\partial \hat{S}_x}{\partial x} + \frac{\partial \hat{S}_y}{\partial y} = \nabla^2 \hat{S} = \text{div}(p, q)$. Thus the SCS method enforce integrability by finding a zero curl gradient field which has the same divergence as the given non-integrable gradient field.

2.1. Curl-divergence space

One can visualize a two dimensional curl-divergence space as shown in Figure 2. A vector in this space represents a gradient field. All integrable (zero-curl, irrotational) gradient fields lie along the real (X) axis. All divergence free (solenoidal) gradient fields lie along the imaginary (Y) axis.

Let \vec{OA}_1 denotes an integrable gradient field with divergence L . Suppose the estimated gradient field is given by \vec{OA}_2 with divergence L_1 and curl C_1 . The residual (error) gradient field is then given by $\vec{A_1A_2}$. As discussed earlier, reconstruction using Poisson solver will give a solution which has the same divergence as the given non-integrable field \vec{OA}_2 . Thus the integrable field given by the SCS method will be \vec{OA}_3 . It is important to note that the divergence free part (curl) of \vec{OA}_2 is completely ignored during the reconstruction. The method that we propose tries to




Figure 2. The curl divergence space: A vector in this space represents a gradient field. Vectors along X axis (OA_1) represents zero curl (integrable) gradient fields. Usually a non-integrable field OA_2 is given as an estimate of OA_1 . The residual gradient field is given by A_1A_2 . Enforcing integrability using Poisson reconstruction gives OA_3 as the solution. The proposed algorithm tries to estimate the residual gradient field A_1A_2 using the information in the curl to move from OA_2 to OA_1 .

estimate the residual field $\overrightarrow{A_1A_2}$ using the curl to move from $\overrightarrow{OA_2}$ to $\overrightarrow{OA_1}$, thus giving a better estimate of the underlying field.

2.2. Information in curl

When does the curl of a non-integrable field have information about the divergence of the underlying integrable field? Consider a very simple example: suppose due to some physical process, the estimated gradient field (p, q) is equivalent to the true gradient field rotated by an angle θ , i.e.

$$\begin{bmatrix} p \\ q \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} p^0 \\ q^0 \end{bmatrix} \quad (3)$$

From (1), the curl and divergence of the estimated gradient field is

$$\begin{aligned} L &= \frac{\partial}{\partial x}(p^0 \cos(\theta) - q^0 \sin(\theta)) \\ &\quad + \frac{\partial}{\partial y}(q^0 \cos(\theta) + p^0 \sin(\theta)) \\ &= L^0 \cos(\theta) + C^0 \sin(\theta) = L^0 \cos(\theta) \\ C &= \text{curl}(p, q) = -L^0 \sin(\theta) \end{aligned} \quad (4)$$

Thus, we see that the curl of the estimated gradient field is directly related to the divergence of the true gradient field. In fact at $\theta = \pi/2$, the divergence will be zero all over the image and all information about L^0 is in the curl. Thus, the Poisson reconstruction in that case will just depend on the boundary conditions and will fail to reconstruct the underlying surface. However, if we define the magnitude of new divergence field \hat{L} as

$$\hat{L} = (L^2 + C^2)^{1/2}$$

then $\hat{L} = L^0$ and the original surface can be recovered.

The above example may seem a bit far stretched but it shows that there can be cases where the curl contains important information about the underlying surface. In practical scenarios however, one can't expect to find analytic expressions for the curl. Even if such equations may be known, they might be extremely complicated to solve. Thus we want to have an approach that does not depend on the analytical form of curl.

It is important to note that there can be cases where curl does not provide any useful information. For example, if noise is added to an integrable gradient field, the curl of the resultant field is just the curl of the noise and is independent of the underlying gradient field. However, in many problems such as SFS, PS etc., noise does not get added to the gradients directly. Usually, the errors in estimating the gradient field depends on the underlying surface. Thus the curl can be expected to have information about the surface.

3. Enforcing integrability

We enforce integrability in discrete domain by formulating the problem as solving a linear system. We use the information in the curl of the given non-integrable field to achieve better reconstruction. In discrete domain, the gradients of a surface S along x and y directions at any pixel (y, x) are defined as simple forward differences

$$\begin{aligned} p(y, x) &= S(y, x+1) - S(y, x) \\ q(y, x) &= S(y+1, x) - S(y, x) \end{aligned} \quad (5)$$

The equation for curl can be written in discrete domain as (using (1))

$$C(y + \frac{1}{2}, x + \frac{1}{2}) = p(y+1, x) - p(y, x) + q(y, x) - q(y, x+1) \quad (6)$$

This correspond to the loop integral around a box of four pixels. Note that the notation $C(y + \frac{1}{2}, x + \frac{1}{2})$ is used to emphasize that the curl value is for that particular loop. In addition, the loop integral around any other bigger loop can be written as linear combinations of these elementary loop integrals and thus will not provide additional information.

3.1. Residual field and curl

Let (p, q) be the given non-integrable gradient field. One can always write

$$(p, q) = (p^0, q^0) + (p_\epsilon, q_\epsilon) \quad (7)$$

Thus

$$\text{curl}(p, q) = \text{curl}(p^0, q^0) + \text{curl}(p_\epsilon, q_\epsilon) = \text{curl}(p_\epsilon, q_\epsilon) \quad (8)$$

One can compute the curl given (p, q) . Thus we wish to estimate the *residual gradient field* (p_ϵ, q_ϵ) given $\text{curl}(p, q)$.

Once the residual gradient field is obtained, it can be subtracted from the given gradient field to get an estimate of (p^0, q^0) . The following observations can be inferred from (6)

§ 1. If curl is non-zero for a loop at $(y + \frac{1}{2}, x + \frac{1}{2})$, one or more of the following gradients $\{p_\epsilon(y + 1, x), p_\epsilon(y, x), q_\epsilon(y, x), q_\epsilon(y, x + 1)\}$ is non-zero.

§ 2. If the x gradient, $p_\epsilon(y, x)$ is non-zero, it affects the curl at loops $(y + \frac{1}{2}, x + \frac{1}{2})$ and $(y - \frac{1}{2}, x + \frac{1}{2})$ only.

§ 3. Similarly, if the y gradient, $q_\epsilon(y, x)$ is non-zero, it affects the curl at loops $(y + \frac{1}{2}, x + \frac{1}{2})$ and $(y + \frac{1}{2}, x - \frac{1}{2})$ only.

Suppose we know that the estimated x gradients are in error in K_p locations and estimated y gradients are in error in K_q locations. If we do not have this information, for example in SFS, we first compute the curl given the estimated gradient field using (6). Then using §1, we find all such gradients that can potentially be in error. Let $\mathbf{x} = [p_\epsilon^1 \dots p_\epsilon^{K_p}, q_\epsilon^1 \dots q_\epsilon^{K_q}]^T$ denote the vector containing all the erroneous gradients.

Now using §2 and §3, we find the loops for which (6) can have a term corresponding to any of the erroneous gradients in \mathbf{x} . Let K_b be the number of all such loops and $[(y_i^1 + \frac{1}{2}, x_i^1 + \frac{1}{2}) \dots (y_i^{K_b} + \frac{1}{2}, x_i^{K_b} + \frac{1}{2})]$ be those loop locations.

For each such loop, we can write (6). In fact, for any loop at $(y + \frac{1}{2}, x + \frac{1}{2})$, (6) can be written as

$$C(y + \frac{1}{2}, x + \frac{1}{2}) = [1 \dots -1 \dots 1 \dots -1] \mathbf{x} \quad (9)$$

where the entries 1 and -1 occur according to position of $p_\epsilon(y + 1, x), p_\epsilon(y, x), q_\epsilon(y, x), q_\epsilon(y, x + 1)$ in vector \mathbf{x} and rest of the entries are zero. If any of $p_\epsilon(y + 1, x), p_\epsilon(y, x), q_\epsilon(y, x), q_\epsilon(y, x + 1)$ is not present in \mathbf{x} , its corresponding row entry will be zero.

Stacking all such equations on top of each other one obtains

$$b = A\mathbf{x} \quad (10)$$

where $b = [C(y_i^1, x_i^1) \dots C(y_i^{K_b}, x_i^{K_b})]^T$ denotes the curl values. Thus the matrix A is a sparse matrix of size $K_b \times (K_p + K_q)$. Each row of A is according to (6) and has at most 4 non-zero values, two $+1$'s corresponding to $p_\epsilon(y + 1, x), q_\epsilon(y, x)$ and two -1 's corresponding to $p_\epsilon(y, x), q_\epsilon(y, x + 1)$.

3.2. Rank considerations

In SFS, the curl will be non-zero almost everywhere. This means that almost all of the estimated p and q values are erroneous. However, the number of curl equations can

be at most the number of basic 4 pixel loops in the image. For an image of size $H \times W$, there will be $W - 1$ loops along each column and $H - 1$ loops along each row. Thus the total number of equations in $Ax = b$ can be $(H - 1) \times (W - 1)$. However, the sum of all the 4 pixel loops will be equal to the loop integral around the boundary (which we assume is zero due to Dirichlet boundary conditions). Thus one equation will be dependent. Hence, the maximum number of independent equations in $Ax = b$ can be $(H - 1) \times (W - 1) - 1$.

On the other hand, the maximum number of unknowns (gradient values) will be approximately twice the number of pixels in the image. Thus, in general, we expect the number of unknowns (gradients) to be much larger than the number of known curl equations and hence $Ax = b$ will be under-determined. To solve the system, the rank of A must be equal to the number of unknowns $K_p + K_q$. Instead of solving for a minimum norm least square solution (as in Poisson solver), we propose to reduce the number of unknowns. Intuitively, the gradient values corresponding to the low curl loops will not be in much error and can be considered error free. But one would like to make that assumption at minimum number of places to recover the maximum number of erroneous gradients. In the next section, we give a graph analogy to derive a formal way for reducing the number of unknowns. Notice that K_b depends on the which p 's and q 's are in error and thus K_b can change when reducing the number of unknowns.

4. Graph analogy

We now show when the system of equation $Ax = b$ is solvable using graph theory. The image is represented as a graph where nodes corresponds to pixels and edges corresponds to gradients. If the edges corresponding to non-zero residual gradients are considered to be broken, the graph will not remain as a single connected piece. Intuitively, if we assume some edges to be known such that the graph becomes connected, $Ax = b$ can be solved. The idea is to find a minimal set of such edges. Graph analogy has also been used for residue-cut algorithm [7] for phase unwrapping.

We assume Dirichlet *boundary conditions* for all the problems. In SFS, often the use of gradient space is critiqued as it is not well-defined at the boundary but as noted in [6], for discrete data p and q are bounded for all practical purposes.

4.1. Defining a graph

We define an undirected graph $G = (V, E)$ on the image plane. Each node in the graph correspond to a pixel in the image (including boundary nodes). Each node in the interior (not on boundary) has 4 edges, connecting it to nodes (pixel) in north, south, east, west direction. For nodes on boundary, those at the corners of the image have 2 edges and rest have 3 edges each. Each edge in the graph between two nodes represent the gradient p or q between the nodes.

With this configuration, for a $H \times W$ image, one have $W - 1$ edges (or x gradients) for each row and $H - 1$ edges (or y gradient) for each column. Figure 3(a) shows a sample graph.

Thus with the above terminology, given an estimated non-integrable gradient field, one knows the value of the graph edges and the value of graph nodes at the boundary points (from Dirichlet conditions). The goal is to integrate the gradient field or to find the value of nodes (pixels) in the interior of the image.

If the estimated gradient p or q is erroneous at a node, we **break** the corresponding edge in the graph. Thus, when curl is non-zero everywhere, all the edges in the graph (except those between boundary points) will be broken. One should note that we assume Dirichlet boundary conditions or that the pixel values are know at the boundary. This implies that the gradients along the boundary (top/bottom row, first/last column) are always known without any error. Thus a configuration such as Figure 3(b) cannot happen.

If the broken edges are such that graph remains connected, the corresponding system of equation $Ax = b$ can be solved. If however, the graph breaks into n pieces, rank deficit of A is equal to $n - 1$, i.e. $rank(A) = (K_p + K_q) - (n - 1)$

Here is an intuitive argument. If the graph remains connected, one can reach any node of the graph starting from some boundary node. Thus for each node (pixel), one can always find at least one integration path from some boundary pixel. Thus, the node value can be obtained for all the pixels. This in turn implies that the value of broken edges can be obtained. Thus the system $Ax = b$ has to be full rank.

If however the graph breaks into n pieces, the minimum number of edges that are needed to make it connected is $n - 1$. Thus if we have known the edge value of $n - 1$ more edges, we could have solved the system. Thus the rank of A will be $n - 1$ less than required $(K_p + K_q)$. Figures 3(c) and 3(d) show an example.

Which edges to join How can we find which edges to join to make the system solvable ?. Two criteria need to be satisfied

1. A minimum number of edges should be joined so that the maximum number of erroneous edges could be corrected.
2. Edges joined first should be along the loops with low curl values.

To meet the above criterions, we define an edge weight for each edge. For edges corresponding to the x gradient between the nodes $V_{y,x+1}$ and $V_{y,x}$ and corresponding to the y gradient between the nodes $V_{y,x}$ and $V_{y+1,x}$, the edge weight is defined as the curl value of the loop. We then find that *minimal* set of edges which connects the graph and have the minimum total weight.

4.2. Algorithm Outline

The complete algorithm can be specified as follows. Given a non-integrable gradient field (p, q) ,

1. Find curl all over the image. Form the image graph. For each edge in graph, assign a weight as described in the previous paragraph. Put all the boundary nodes in set B_2 .
2. Identify nodes corresponding to loops where curl is greater than some threshold ($\tau = 10^{-2}$). Put all such nodes in set B_1 if it is not a boundary node. Put rest of the nodes in set B_2 . Note that when curl is non-zero all over, B_2 will contain just the boundary nodes and B_1 will contain the rest of the nodes.
3. Break all the edges connecting any node in B_1 to any node in B_2 . Also break all the edges between the nodes in B_1 .
4. **Finding minimal set of edges to join**
While B_1 is not empty
 - (a) Find the shortest path from any node in B_2 to a node in B_1 . Let V_{i_1,j_1} be the node in B_1 and V_{i_2,j_2} be the node in B_2 corresponding to this path.
 - (b) Remove V_{i_1,j_1} from B_1 and put in B_2 .
 - (c) Join the edge connecting V_{i_1,j_1} and V_{i_2,j_2} .
5. Now the graph will be connected. For all the edges still broken, form the equation $Ax = b$. Solve for x and subtract x from the corresponding locations in the given gradient field.
6. Now we have obtained an integrable gradient field. The underlying 2D surface can be obtained by integrating using multi-grid or Poisson solvers.

5. Applications

We present results on Shape from Shading and Photometric Stereo and show that our algorithm gives significantly better solutions. The approach is however not restricted to applications shown and can be used for other imaging applications which requires gradient field reconstruction such as [1]. To evaluate the quality of the reconstruction, we define the percentage depth error between the true depth map Z_t and the estimated depth map Z_{est} as $100 * \sum ((Z_t - Z_{est})/Z_t)^2$, where the summation is over all the pixels. In all experiments, the curl of the estimated gradient field is shown as gray-color coded image (white=positive, black=negative, gray=0).




Figure 3. Graph analogy and rank considerations (a) A sample graph on a 8×8 grid. Lines denotes edges and dots denotes nodes. (b) A configuration such as this cannot happen as boundary edges are known (c) A particular configuration of broken edges where the graph is broken into $n = 4$ parts. $K_p = 18$, $K_p = 16$, total number of broken links = 34, $\text{rank}(A) = 31$. rank deficit = $34 - 31 = 3 = n - 1$ (d) If 3 more links are known to connect the graph (dashed edges), number of unknowns reduces to $31 = \text{rank}(A)$

5.1. Photometric Stereo

For photometric stereo, we present synthetic results using the Mozart depth map and on real data using Yale face database B (*YaleB01_{p00}*). For Mozart, five images were generated assuming Lambertian reflectance model using the depth map. For Yale database, all 64 images corresponding to the frontal pose for the first subject were used. We assume that the light source directions are known. A simple least squares approach was used to estimate the surface normal (n_x, n_y, n_z) at each pixel. The gradient field was obtained as $p = -n_x/n_z$, $q = -n_y/n_z$.

Figure 4(a) shows 4 of the 5 input images generated using the Lambertian model for Mozart whose true depth map is shown in Figure 4(e). Figure 4(b) shows the curl of the estimated gradient field. Figure 4(e)(center) shows the reconstruction from the initial estimated gradient field using the *SCS* method. As described in section 4.2, we find those edges which can potentially give rise to non-zero curl values. Nodes connecting these edges (set B_1) are shown in Figure 4(c) in white. It is clear that the resultant graph will not be connected. We then find the minimal set of edges to join so as to connect the graph. Figure 4(d) shows the final configuration of the connected edges. Note that in this graph, all nodes can be reached using the edges present in the graph. For the edges not present between the nodes, the system of equation $Ax = b$ was solved. The gradient field was then updated using the x values. Figure 4(e) shows the reconstruction using curl corrected gradient field which is much better than using the *SCS* method. The percentage depth error between the true depth map and the reconstructed depth map was 4.26 using the *SCS* method and 2.7 using our method.

Figure 1(a) shows 4 (out of 64) input images for one of the subjects in frontal pose from Yale database. Figure 1(b) shows the curl of the estimated gradient field. Fig-

ures 1(c) and 1(d) shows the reconstruction using the *SCS* method and using our method. Note that the reconstruction using our method is much better especially along the face boundaries. *SCS* method can have global distortions since the gradients at any location affect the reconstruction all over the image. In contrast, our method can locally confine errors, does not have any global distortions and all features (sharp gradients) are preserved. Also notice the transitions between right cheek and back plane which is smoothed out in the reconstruction using the *SCS* method.

5.2. Shape from Shading

For shape from shading, we implemented the algorithm by Brooks and Horn [2] and extended it to incorporate the proposed integrability method. This algorithm assumes a Lambertian reflectance model for the surfaces. In this method, at each iteration, new estimates of the surface gradients $\begin{bmatrix} \hat{p} \\ \hat{q} \end{bmatrix}_{k+1}$ are obtained from the previous estimates

$\begin{bmatrix} \hat{p} \\ \hat{q} \end{bmatrix}_k$ as

$$\begin{bmatrix} \hat{p} \\ \hat{q} \end{bmatrix}_{k+1} = \begin{bmatrix} \hat{\hat{p}} \\ \hat{\hat{q}} \end{bmatrix}_k + \lambda(I - R) \begin{bmatrix} R_x \\ R_y \end{bmatrix} \quad (11)$$

where $\hat{\hat{p}}_k$ and $\hat{\hat{q}}_k$ denotes the smoothed values of \hat{p}_k and \hat{q}_k respectively, I is the input image, R is the reflectance map and R_x , R_y denotes its derivatives. We impose integrability as follows. At each iteration, we first find the new update using the above equation. Then we find the integrable field using our method and use the integrable field in the next iteration.

Figure 5 shows synthetic example for Vase and real example for the Lena image. For vase, we use the illumination direction $(0.05, 0.05, 1)$ to generate the image using a




Figure 4. Photometric Stereo for Mozart (a) Sample images (b) Curl values for initial estimated gradient field (c) Initial nodes in set B_1 (section 4.2 step 3) (d) Final image graph which is connected. Nodes are represented by dots and edges by lines. Edges not present can be solved for by forming $Ax = b$ (e) Original depth map, Reconstructed surface using the *SCS* method and using proposed algorithm from updated gradient field.

Lambertian reflectance model. For Lena, the illumination direction $(1.5, 0.866, 1)$ was used as suggested in [16]. For both examples, 50 iterations of Brooks & Horn algorithm were performed. Figures 5(b) and 5(d) compares the reconstructed surfaces using the *SCS* method and the proposed method. Note that in vase image, enforcing integrability using the *SCS* method does not properly maintains the boundaries of the object, whereas in the reconstruction using our method, the boundaries are well maintained. The percentage depth errors using the *SCS* method and our method were 183.0 and 87.6 respectively. For the Lena image, a better visual reconstruction is obtained as shown in Figure 5(d).

6. Conclusions

Discussions: The only tunable parameter in our algorithm is the threshold τ for deciding whether a loop is curl free or not. With increasing τ , more gradients will be used in the estimation and the solution will move towards the solution given by the *SCS* method. As τ is decreased, we selectively start choosing gradients for integration. Thus one can have a tradeoff between smoothness and local error confinement by adjusting this parameter.

We have proposed a method for surface reconstruction from gradients for problems commonly encountered in Shape from Shading and Photometric Stereo. We presented

an algebraic method for enforcing integrability. Our approach corrects for the curl of the given non-integrable gradient field by solving a linear system. Using a graph analogy, we derived conditions under which this system can be solved. Our method is non-iterative and has the important property that the errors due to non-zero curl do not propagate across the surface. In comparison with conventional methods such as the *SCS* method, the extra computation cost for our approach is in determining the minimal set of edges to construct a connected graph. However, the worst case computational expense corresponds to finding the minimum spanning tree of the image graph for which standard algorithms are available.

References

- [1] A. Agrawal, R. Raskar, S. Nayar, and Y. Li. Removing photography artifacts using gradient projection and flash-exposure sampling. *to appear in ACM Trans. Graph.*, 2005.
- [2] M. Brooks and B. Horn. Shape and source from shading. In *Proc. Intl. Joint Conf. Artif. Intelli.*, pages 932–936, 1985.
- [3] J. Fan and L. Wolff. Surface curvature from integrability. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 520–526, 1994.
- [4] D. Forsyth. Shape from texture and integrability. In *Proc. Int'l Conf. Computer Vision*, pages 447–452, July 2001.




Figure 5. Shape from Shading (a) Synthetic vase image generated using $(0, 0, 1)$ illumination direction (b) Original depth map, Reconstructed surface using the *SCS* method and using the proposed algorithm. Note that in the *SCS* method, the surface is globally distorted and boundaries are not preserved. (c) Lena Image (f) Reconstructed surface using the *SCS* method and using the proposed algorithm. There are distortions along the lips and top left of the reconstructed surface using the *SCS* method.

- [5] D. Forsythe and J. Ponce. *Computer Vision: A modern approach*. Prentice Hall, 2001.
- [6] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. Pattern Anal. Machine Intell.*, 10(4):439–451, July 1988.
- [7] R. M. Goldstein, H. A. Zebker, and C. L. Werner. Satellite radar interferometry: Two-dimensional phase unwrapping. *Radio Science*, 23:713–720, 1988.
- [8] B. Horn. Height and gradient from shading. *Int'l J. Computer Vision*, 5(1):37–75, 1990.
- [9] N. Petrovic, I. Cohen, B. Frey, R. Koetter, and T. Huang. Enforcing integrability for surface reconstruction algorithms using belief propagation in graphical models. In *Proc. Conf. Computer Vision and Pattern Recognition*, 2001.
- [10] K. Polthier and E. Preuss. Variational approach to vector field decomposition. In *Proc. Eurographics Workshop on Scientific Visualization*, 2000.
- [11] T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving poisson equations in computer vision problems. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(5):435–446, May 1990.
- [12] Y. Tong, S. Lombeyda, A. Hirani, and M. Desbrun. Discrete multiscale vector field decomposition. *ACM Trans. Graph.*, 22(3):445–452, 2003.
- [13] R. Woodham. Photometric method for determining surface orientation from multiple images. *OptEng*, 19(1):139–144, 1980.
- [14] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 158–164, 1997.
- [15] A. L. Yuille, D. Snow, R. Epstein, and P. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability. *Int'l J. Computer Vision*, 35(3):203–222, 1999.
- [16] R. Zhang, P. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Trans. Pattern Anal. Machine Intell.*, pages 690–706, 1991.
- [17] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(7):680–702, 1991.
- [18] S. Zhou, R. Chellappa, and D. Jacobs. Characterization of human faces under illumination variations using rank, integrability, and symmetry constraints. In *Proc. European Conf. Computer Vision*, volume 1, pages 588–601, 2004.