

# Visible Surface Reconstruction from Normals with Discontinuity Consideration\*

Tai-Pang Wu and Chi-Keung Tang  
Vision and Graphics Group  
The Hong Kong University of Science and Technology  
Clear Water Bay, Hong Kong  
{pang|cktang}@cs.ust.hk

## Abstract

*Given a dense set of imperfect normals obtained by photometric stereo or shape from shading, this paper presents an optimization algorithm which alternately optimizes until convergence the surface integrabilities and discontinuities inherent in the normal field, in order to derive a segmented surface description of the visible scene without noticeable distortion. In our Expectation-Maximization (EM) framework, we enforce discontinuity-preserving integrability so that fine details are preserved within each output segment while the occlusion boundaries are localized as sharp surface discontinuities. Using the resulting weighted discontinuity map, the estimation of a discontinuity-preserving height field can be formulated into a convex optimization problem. We compare our method and present convincing results on synthetic and real data.*

## 1. Introduction

Over the past decades, many methods have been proposed to recover surface normals of a scene from images. The class of methods commonly known as shape-from-X provides different alternatives to achieve the goal. Given a static scene captured by a camera with fixed viewpoint, photometric stereo [16] and shape-from-shading [7] are two methods widely used in computer vision to recover surface gradients of the visible scene. To reconstruct a surface description from the estimated normals, the integrability constraint or the zero-curl constraint is often applied. However, the underlying surface of the captured scene may not be uniformly integrable, for example, see Figure 1(a).

Recently, many methods [6, 17, 14, 18] have been proposed in which a very good normal field can be produced. Surface discontinuities or abrupt changes in normal orientations are well-preserved in the output field, where the sur-

face integrability constraint may *not* be uniformly applicable. Given such a discontinuity-preserving normal map in which noise may still be present, the problem we want to address in this paper is: *How to reconstruct a discontinuity-preserving surface where fine geometric details are preserved and occlusion boundaries are respected?* This is not a new problem: the classical method [15] provides a compact framework for visible surface reconstruction, which considers simultaneously surface normal, height field, orientation discontinuity and depth discontinuity. The resulting error surface is very complex and so incomplete height field is needed as input. One may apply recent methods such as [9, 10, 11, 13] to perform surface reconstruction. However, almost all the available methods produce somewhat distorted results. Some apply the integrability constraint uniformly over the whole normal map. Although partial integrability was used in [9, 10], the binary decision made by manual labeling or thresholding by edge detection may not discern occlusion boundaries from local fine details resulted by sharp changes in surface orientation.

In this paper, we study and model the properties of the class of visible surfaces where smoothness and discontinuities may co-exist in the same scene. Based on the observation and cues inherent in the normal data, we propose an Expectation-Maximization (EM) scheme to simultaneously consider smoothness and discontinuities in an alternating optimization framework which is guaranteed to converge [1] to an optimal solution. Using the resulting weighted discontinuity map of the scene, a segmented description of the visible surface is obtained via convex optimization, where fine details within each segment and surface discontinuities across segments are preserved.

## 2. Related Work

This paper concerns with the derivation of a segmented visible surface description from a set of normals which may be imperfect. The traditional approach enforces the surface integrability or the zero-curl constraint across the whole

---

\*This research was supported by the Research Grant Council, Hong Kong under grant nos 620005 and AoE/E-01/99.

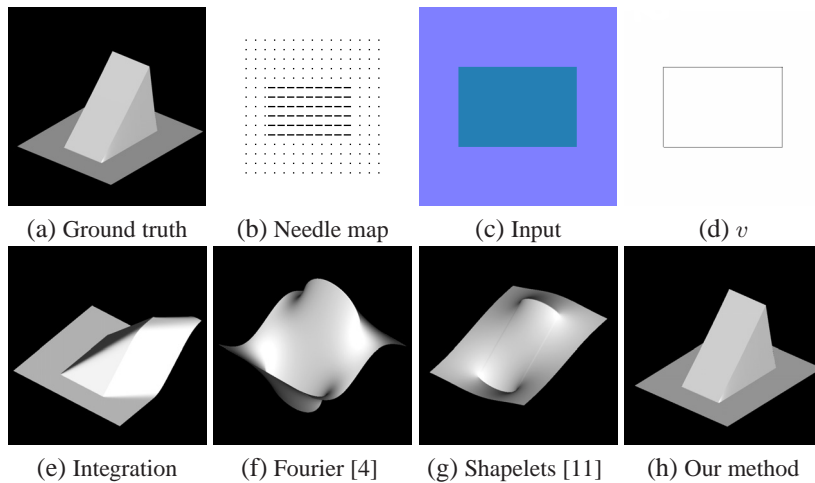


Figure 1. Slope: (a) A partially integrable surface (ground truth). (b) the corresponding needle map. (c) Color coded input normal map. (d) The visualized weighted discontinuity map  $v$  found by our EM algorithm. (e) Surface reconstructed by using integration. (f) Surface reconstructed by using [4]. (g) Surface reconstructed by using [11]. (h) Surface reconstructed by using our method.

normal field [3], so that the height at any given point can be computed by summing up the changes in height along some integration path.

More sophisticated techniques can be categorized into regularization-based or basis-function-based. *Regularization* methods attempt to smooth the input normal gradients under certain criteria (e.g., a data fidelity term, a smoothness term and possibly a penalty term for orientation discontinuities) so that integrability can be enforced or the height field can be estimated directly. The classical approach proposed by Terzopoulos [15] uses a variational formulation that was translated into an energy minimization problem, which is one of the regularization-based methods consider discontinuity information. In [2], other regularization-based methods are described. Different surface models such as weak membrane are discussed. Line-processes [12, 5] can be used to deal with discontinuities. However, these methods estimate discontinuity and height at the same time, which may complicate the underlying error surface and so rough input height is sometimes needed. Later, Petrovic et al [13] formulated the problem using a Markov Random Field network, and proposed to use belief propagation to enforce the integrability constraint using factor graphs.

One classical approach that makes use of *basis functions* was proposed by Frankot and Chellppa [4]. They projected the surface gradients onto a set of integrable basis functions for enforcing surface integrability. Fourier basis functions were used. Variants such as wavelet basis functions have been proposed [8]. Karacali and Snyder [9] made use of an orthonormal set of gradient fields. They projected the measured gradient field onto a feasible subspace to generate an integrable surface with surface gradients resembling the measured ones. Recently, shape from shapelets was pro-

posed in [11]. This method correlates the measured surface gradients with the gradients of a bank of shapelet basis functions which satisfy a set of constraints such as admissibility, minimal ambiguity, uniform coverage and preservation of phase information. One common drawback of all these basis-function-based methods is that all occlusion boundaries are smoothed out in their generated visible surface description. This is due to the global enforcement of the surface integrability constraint. Recently, a kernel-based approach was used in [9] which enforces partial integrability to preserve discontinuities which are manually labeled. Later, they used edge detection to label discontinuity via multi-scale analysis [10]. The examples presented in [10] are however simple.

The rest of the paper is organized as follow. We discuss the issues pertinent to *discontinuity-preserving* integrability in section 3. Section 4 details the EM approach for computing the weighted discontinuity map. Section 5 describes our height field estimation algorithm subject to the discontinuity-preserving integrability constraint. Experimental results and comparison are presented in section 6.

### 3. Issues

In this section, we study the issues in visible surface reconstruction from surface normals in the presence of sharp surface discontinuities, by first providing an illustrative discussion on a simple case. Then, we compare the surface outputs generated by the state-of-the-art basis-function techniques in order to motivate our study (Regularization-based method sometimes require initial heights which are absent here). Based on our observation, we identify the necessary considerations and translate them into a pertinent

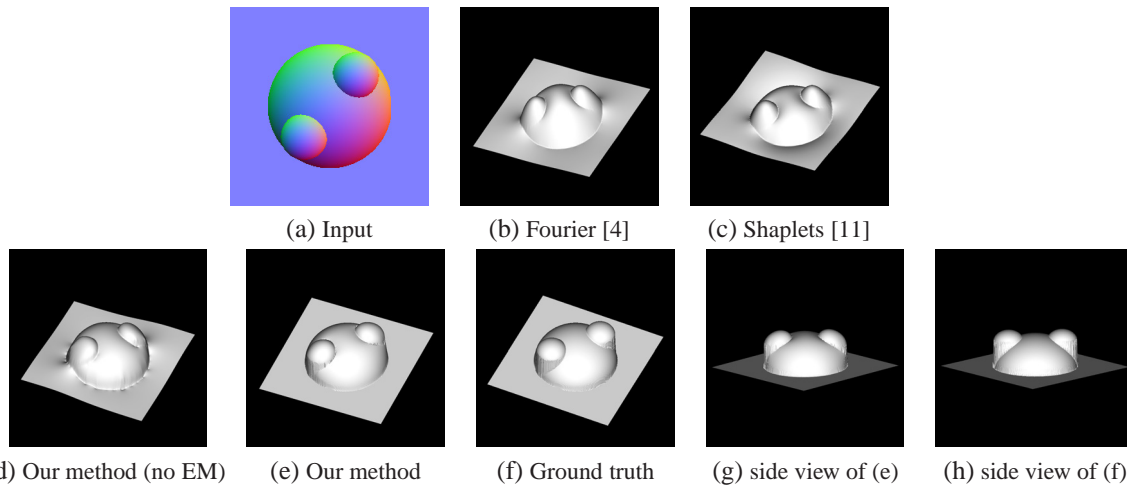


Figure 2. Three Spheres: (a) Color coded input normal map. (b) Surface reconstructed by using [4]. (c) Surface reconstructed by using [11]. (d) Surface reconstructed by our height field generation without weighted discontinuity map. (e) Surface reconstructed by using our method. (f) The ground truth surface. (g) Surface reconstructed by using our method (side view). (h) The ground truth surface (side view).

objective function so that they can be incorporated into our EM framework for localizing and weighting discontinuities (Section 4) and convex optimization for solving for the visible surface or height field (Section 5).

Given a 2D surface normal map, our goal is to recover a segmented surface description where fine details on smooth surfaces as well as surface discontinuities are preserved. Such surface is represented by a height field, represented as a set of triples  $(x, y, f(x, y))$ , where  $x$ ,  $y$  and  $f(x, y)$  are the  $x$ -,  $y$ -coordinate and the height at  $(x, y)$  respectively. So, a normal at  $(x, y)$  is given by:  $\mathbf{N}(x, y) = \frac{1}{\sqrt{1 + \frac{\partial f^2}{\partial x^2} + \frac{\partial f^2}{\partial y^2}}}(-\frac{\partial f}{\partial x}, -\frac{\partial f}{\partial y}, 1)^T$ . Very often, the measured normal at  $(x, y)$  is given in the form of  $(e_1(x, y), e_2(x, y), e_3(x, y))^T$ , then we have

$$\frac{\partial f}{\partial x}(x, y) = -\frac{e_1(x, y)}{e_3(x, y)} \quad \frac{\partial f}{\partial y}(x, y) = -\frac{e_2(x, y)}{e_3(x, y)} \quad (1)$$

The uniform enforcement of integrability (or zero-curl) constraint globally in the whole normal map allows the surface (height field) be estimated by integration, so that the recovered surface does not depend on the choice of the integration path [3]. The constraint can be implemented by minimizing the following

$$\left\| \frac{\partial f}{\partial x \partial y} - \frac{\partial f}{\partial y \partial x} \right\|. \quad (2)$$

### 3.1. Integrable smooth surfaces with discontinuities

Many approaches uniformly enforce the surface integrability in a global way. This is fundamentally problematic because “partially integrable surface” such as one shown in Figure 1 is ubiquitous. Figure 1(a) and Figure 1(b) respectively show a partially integrable surface and the cor-

responding normal map. If we integrate the surface along different paths, it will give different results. Moreover, the reconstructed surface is not acceptable (Figure 1(e)).

If global surface integrability constraint is enforced, which may severely alter the surface normal orientations to satisfy the zero-curl constraint, a distorted surface is obtained by [4] as shown in Figure 1(f). Another result produced by applying the method of [11] is shown in Figure 1(g). Both surface results suffer severe distortion, indicating that adequate consideration of discontinuity is essential.

To understand how such consideration contributes to significant improvement in surface reconstruction, we study a case shown in Figure 2. First, by enforcing the global integrability constraint, a surface shown in Figure 2(b) and (c) is produced. Note the distortion of the two small spheres and the ground. This is due to the fact that the smaller spheres are partially (or tangentially) occluding the ground without contact, while the uniform integrability constraint makes the ground contribute to the surface of the two small spheres, which is incorrect (The same also happens in 1(f) and (g)). One better way to reconstruct the two small spheres is to perform surface integration *on* the small spheres only.

Naturally, the above discussion indicates that segmentation should be performed. One *plausible* method to localize the segmentation curves is to identify where the integrability constraint is violated, i.e., identify locations where Eqn. 2 gives a large value. Not surprisingly, in real case, applying thresholding to Eqn. 2 may produce very poor discontinuity maps rendering them unusable at all. Instead of producing a binary discontinuity map, or blocking the contribution of neighbors completely or using a thresh-

old [9, 10], we turn to estimate a *weighted discontinuity map* to weigh the relative contribution of the pixel neighbors by optimizing surface normals and surface discontinuities alternately.

Let  $p(x, y) = \frac{\partial f(x, y)}{\partial x}$  and  $q(x, y) = \frac{\partial f(x, y)}{\partial y}$  be the partial derivatives in  $x$  and  $y$  directions respectively. The discrete version of Eqn. 2 is:

$$||p(x, y) - p(x, y + 1) + (q(x + 1, y) - q(x, y))|| \quad (3)$$

We intend to find a configuration for the weighted discontinuity map that minimizes Eqn. 3.

In addition, to refine the input normal map, we introduce the neighborhood inference by minimizing the following equation:

$$||p(x, y) - p(x + 1, y)|| + ||q(x, y) - q(x + 1, y)|| \quad (4)$$

In this paper, the estimation of discontinuity weights and refinement of normals are performed alternately by EM under the Markov Random Field (MRF) condition, because we consider a pixel’s neighborhood information when optimizing for the solution at the pixel. Our EM framework will be described in detail in Section 4.

### 3.2. Height field in the presence of discontinuities

Given a weighted discontinuity map, we need a method for computing a height field (surface) that utilizes the discontinuity weights so that fine details and surface discontinuities are preserved.

Since the loci of discontinuities can exist in many forms such as closed and opened loops, the ideal method should consider all possible integration paths. Instead of doing this explicitly, we perform the following integration recursively: Consider the first-order neighbors (up, bottom, left and right) of  $(x, y)$ , we minimize the integration error induced by these four neighbors where the error is weighted by the discontinuity map. Since there are only at most 4 neighbors for each pixel location, at most 4 possible integration paths exist within the scope. If the 4 neighbors have the correct height values,  $(x, y)$  will also obtain the correct height. Recursively, all possible paths can be considered.

In implementation, we can formulate the above by a convex function so that standard convex optimization tools are applicable. The details will be described in Section 5.

## 4. An EM Approach for Discontinuity Map Estimation

Surface discontinuities do not exist without smoothness in its vicinity: For example, surface orientation discontinuity is resulted when two smooth surfaces intersect, and discontinuous surfaces are observed when one surface occludes the other. We propose to cooperate integrability

and discontinuity in an alternating optimization framework: Expectation-Maximization (EM). In this section, we define our objective function which forms the basis of our EM algorithm for estimating the discontinuity map.

### 4.1. Objective function

Let  $i$  be the location index of the normal map, where  $i \in \{1, \dots, N\}$  and  $N$  is the size of the normal map. We use  $x_i$  to denote the attribute  $x$  at pixel location  $i$ . For example,  $p_i$  and  $q_i$  are the partial derivatives at location  $i$  in  $x$  and  $y$  directions respectively. Without loss of generality, we let  $a, b, c$  and  $d$  be the respective pixel location indices of the left, right, up and bottom neighbors of pixel  $i$ .

We define  $\mathbf{O} = \{o_i = (\tilde{p}_i, \tilde{q}_i)\}$  to be the set of observations, where  $\tilde{p}_i$  and  $\tilde{q}_i$  are the partial derivatives at location  $i$  in  $x$  and  $y$  direction which are kept constant as observations. They are given from the input normal map.

Given  $\mathbf{O}$  and under the MRF assumptions, our goal is to estimate the optimal partial derivatives  $p_i$  and  $q_i$ . Mathematically, we estimate the following:

$$\lambda^* = \arg \max_{\lambda} P(\mathbf{O}, \mathbf{R}|\lambda) \quad (5)$$

where  $P(\mathbf{O}, \mathbf{R}|\lambda)$  is the complete-data likelihood we want to maximize,  $\lambda = \{\{p_i, q_i\}, \beta, \sigma_1, \sigma_2\}$  is the set of parameters to be estimated,  $\beta, \sigma_1$  and  $\sigma_2$  are the parameters of certain distributions that will be introduced later and  $\mathbf{R} = \{r_i\}$  is a set of hidden variables on the connectivity labels. Suppose that  $i$  and  $j$  are first-order neighbors, they are connected if  $r_i = r_j$ , otherwise,  $r_i \neq r_j$ .

To estimate Eqn. 5, the EM algorithm computes the expected value of the complete-data log-likelihood  $\log P(\mathbf{O}, \mathbf{R}|\lambda)$  w.r.t.  $\mathbf{R}$  given the observation  $\mathbf{O}$  and the current estimated parameters  $\lambda'$ :

$$Q(\lambda, \lambda') = \sum_{\mathbf{R} \in \varphi} \log P(\mathbf{O}, \mathbf{R}|\lambda) P(\mathbf{R}|\mathbf{O}, \lambda') \quad (6)$$

where  $\varphi$  is a space containing all  $\mathbf{R}$  of size  $N$ .

Because we are dealing with MRFs, we need to understand the dependency of the neighboring pixels. Here we review the two standard MRF assumptions which make the associated EM algorithm tractable.

- 1 Given the MRF network, the hidden variable  $r_i$  depends only on the hidden variables of its first-order neighbors.
- 2 The observation at  $i$  depends only on the hidden variable at  $i$ .

### 4.2. Expectation Estimation

In this section, we describe how to estimate  $P(\mathbf{R}|\mathbf{O}, \lambda')$  so that we can maximize the expectation  $Q$  defined by

Eqn. 6 by proceeding to the next iteration given the current estimated parameters.

Without loss of generality, our marginal probability  $p(r_i|o_i, \lambda')$  always equals to 1 since there is no connectivity consideration. Since we want to locate discontinuities, we need to define  $p(r_i, r_j|o_i, o_j, \lambda')$  (or  $p(r_i, r_j|\lambda')$  because of assumption (1)), where  $i$  and  $j$  are first-order neighbors. To know whether discontinuity exists between  $i$  and  $j$ , we use the measurement of integrability (Eqn. 3) and similarity of neighbors (Eqn. 4). However, Eqn. 3 is difficult to be embedded into a MRF network together with the neighborhood inference given by Eqn. 4. So, instead of using Eqn. 3 directly, by triangle inequality, we consider its upper bound as the approximation:

$$\|p_i - p_d\| + \|q_b - q_i\| \quad (7)$$

Therefore, if the surface connecting locations  $i$  and  $j$  is smooth (i.e.  $r_i = r_j$ ), Eqn. 4 and Eqn. 7 should be minimized, which can now be combined nicely into a single equation. In contrast, if there exists a discontinuity between them (i.e.  $r_i \neq r_j$ ), the orientation of the two neighboring normals is arbitrary. So, we have:

$$p(r_i, r_j, \lambda') \propto \begin{cases} \exp(-\frac{\|p_i - p_j\|^2 + \|q_i - q_j\|^2}{2\sigma_1^2}), & \text{if } r_i = r_j; \\ \frac{1}{C}, & \text{if } r_i \neq r_j. \end{cases} \quad (8)$$

where  $C$  is a constant and  $\sigma_1$  is the standard deviation of the difference in partial derivative. In this paper,  $C$  is the only parameter input by user, which controls the overall amount of smoothness in the resulting normal map. Large  $C$  increases the influence of smoothness and preserves less discontinuity, and vice versa. For a fully automatic method, one may make  $C$  equal to the mean of  $\|p_i - p_j\|^2 + \|q_i - q_j\|^2$ , which gives reasonable result in all our experiments.

Let  $\beta$  be the proportion of the neighborhood pairs that are connected. Then, the mixture probability is defined as:

$$p(r_i = r_j) = \beta \quad (9)$$

So, given  $\lambda'$  and by considering  $i$  and  $j$  only, we have:

$$p(\lambda') \propto \beta \exp(-\frac{\|p_i - p_j\|^2 + \|q_i - q_j\|^2}{2\sigma_1^2}) + \frac{1-\beta}{C} \quad (10)$$

Let  $v_{ij} = p(r_i = r_j|\lambda')$ , The marginal probability of the connectivity between  $i$  and  $j$  is:

$$\begin{aligned} v_{ij} &= p(r_i = r_j|\lambda') = \frac{p(r_i = r_j, \lambda')}{p(\lambda')} \\ &= \frac{\beta \exp(-\frac{\|p_i - p_j\|^2 + \|q_i - q_j\|^2}{2\sigma_1^2})}{\beta \exp(-\frac{\|p_i - p_j\|^2 + \|q_i - q_j\|^2}{2\sigma_1^2}) + \frac{1-\beta}{C}} \quad (11) \end{aligned}$$

Hence, in the E-Step, we solve  $v_{ij}$  for every first-order pair of neighboring pixels  $\{i, j\}$ . Note that  $v$  encodes the discontinuity probability for every pair of pixels, thus giving the weighted discontinuity map we desire.

### 4.3. Maximization

Recall in Eqn. 6 that the likelihood  $P(\mathbf{O}, \mathbf{R}|\lambda)$  is very complicated due to the use of MRF. Therefore, we break down  $P(\mathbf{O}, \mathbf{R}|\lambda)$  by rewriting it into a combination of simple elements first. Based on the two MRF assumptions, the likelihood can be broken down in this way:

$$P(\mathbf{O}, \mathbf{R}|\lambda) = \prod_i \prod_{j \in \mathcal{N}_i} p(r_i|r_j, \lambda)p(o_i|r_i, \lambda) \quad (12)$$

where  $\mathcal{N}_i$  is a set of right and bottom neighbors of  $i$ .  $p(r_i|r_j, \lambda) = 1$  if  $j$  does not exist.

Now, the  $Q$  function in Eqn. 6 can be rewritten as:

$$\begin{aligned} Q(\lambda, \lambda') &= \sum_{\mathbf{R} \in \varphi} \log \left( \prod_i \prod_{j \in \mathcal{N}_i} p(r_i|r_j, \lambda)p(o_i|r_i, \lambda) \right) P(\mathbf{R}|\mathbf{O}, \lambda') \\ &= \sum_{\mathbf{R} \in \varphi} \sum_i \sum_{j \in \mathcal{N}_i} \log(p(r_i|r_j, \lambda)) P(\mathbf{R}|\mathbf{O}, \lambda') \\ &+ \sum_{\mathbf{R} \in \varphi} \sum_i \log(p(o_i|r_i, \lambda)) P(\mathbf{R}|\mathbf{O}, \lambda') \quad (13) \end{aligned}$$

Now, we model the likelihood  $p(o_i|r_i, \lambda)$  by minimizing the difference between  $o_i$  and the estimated output. So we have

$$p(o_i|r_i, \lambda) \propto \exp(-\frac{\|p_i - \tilde{p}_i\|^2 + \|q_i - \tilde{q}_i\|^2}{2\sigma_2^2}) \quad (14)$$

where  $\sigma_2$  is the standard deviation of the difference between the observation and the estimated output.

Since we only have 2 states  $\mathcal{G}_{ij} = \{r_i = r_j, r_i \neq r_j\}$ , Eqn. 13 can be further simplified:

$$\begin{aligned} Q(\lambda, \lambda') &= \sum_{g \in \mathcal{G}_{ij}} \sum_i \sum_{j \in \mathcal{N}_i} \log(p(g|r_j, \lambda))p(g|\lambda') \\ &+ \sum_i \log(p(o_i|r_i, \lambda))p(r_i|o_i, \lambda') \\ &= \sum_i \sum_{j \in \mathcal{N}_i} \log \left( \frac{\beta}{\sigma_1 \sqrt{2\pi}} \exp(-\frac{k(i, j)}{2\sigma_1^2}) \right) v_{ij} \\ &+ \sum_i \sum_{j \in \mathcal{N}_i} \log \left( \frac{1-\beta}{C} \right) (1 - v_{ij}) \\ &+ \sum_i \log \left( \frac{1}{\sigma_2 \sqrt{2\pi}} \exp(-\frac{m(i)}{2\sigma_2^2}) \right) \quad (15) \end{aligned}$$

where  $k(i, j) = \|p_i - p_j\|^2 + \|q_i - q_j\|^2$  and  $m(i) = \|p_i - \tilde{p}_i\|^2 + \|q_i - \tilde{q}_i\|^2$ .

To maximize Eqn. 15, we set the first derivative of  $Q$  w.r.t.  $\beta$ ,  $\sigma_1$ ,  $\sigma_2$ ,  $p_i$  and  $q_i$  respectively equal to zero and finally we obtain the parameter updating rule for  $\lambda$  as:

$$\begin{aligned}\beta &= \frac{1}{|\mathcal{H}|} \sum_{\{i,j\} \in \mathcal{H}} v_{ij} \\ \sigma_1 &= \frac{\sum_{\{i,j\} \in \mathcal{H}} k(i,j) v_{ij}}{\sum_{\{i,j\} \in \mathcal{H}} v_{ij}} \\ \sigma_2 &= \frac{1}{N} \sum_i m(i) \\ p_i &= \frac{1}{\sum_{j \in \mathcal{K}_i} v_{ij}} \sum_{j \in \mathcal{K}_i} p_j v_{ij} \\ q_i &= \frac{1}{\sum_{j \in \mathcal{K}_i} v_{ij}} \sum_{j \in \mathcal{K}_i} q_j v_{ij}\end{aligned}\quad (16)$$

where  $\mathcal{H}$  is a set of neighboring pixel pairs and  $\mathcal{K}_i$  is the set of first-order neighbors (up, bottom, left and right) of  $i$ .

Hence, in the M-Step, the updating rule (Eqn. 16) is applied to estimate a new set of parameter values. The E-Step and M-Step are iterated alternately until convergence. The marginal probability  $v_{ij}$  computed using the converged parameters (Eqn. 11) gives the weighted discontinuity map.

## 5. Convex Optimization for Height Field Estimation

Given the discontinuity map  $v$ , we now give a method to compute the height field. We apply the idea of recursive integration outlined in Section 3.2.

Let  $h_i$  be the height at location  $i$ . Given a first-order neighbor pair  $i$  and  $j$ , the residual of the height  $h_i$  w.r.t.  $h_j$  is defined by the difference between the estimated  $h_i$  and the height integrated starting from  $j$ . With the discontinuity map, the residual is weighted by  $v_{ij}$ . Mathematically, the residual is defined as follows:

$$\begin{cases} v_{ij}(h_i - h_j + p_i)^2, & \text{if } j = a; \\ v_{ij}(h_i - h_j - p_j)^2, & \text{if } j = b; \\ v_{ij}(h_i - h_j + q_i)^2, & \text{if } j = c; \\ v_{ij}(h_i - h_j - q_j)^2, & \text{if } j = d. \end{cases}\quad (17)$$

where the  $p$ 's and  $q$ 's are the output of the EM algorithm. Thus, the overall residual  $E$  of the reconstructed surface is defined by the summation of all the residuals in Eqn. 17 for each pair of  $i$  and  $j$ :

$$E(h) = \sum_i (v_{ib}(h_i - h_b - p_b)^2 + v_{id}(h_i - h_d - q_d)^2) \quad (18)$$

Since each residual is a convex function, the summation of the residuals, i.e.  $E(h)$ , is also a convex function making the optimization guaranteed to converge. One can use a

standard optimization package to minimize Eqn. 18 to extract the height field. In our implementation, we solve it by setting the first derivative of  $E$  w.r.t.  $h_i$  equals to zero and solve  $h$  by applying an iterative approach. By setting  $\partial E(h)/\partial h_i = 0$ , we have:

$$\begin{aligned}h_i &= \frac{1}{\sum_{j \in \mathcal{K}(i)} v_{ij}} (v_{ai}(h_a + p_i) + v_{ib}(h_b - p_b) \\ &\quad + v_{ci}(h_c + q_i) + v_{id}(h_d - q_d))\end{aligned}\quad (19)$$

In each iteration, for each  $i$ , we estimate  $h_i$  by solving Eqn. 19 until the algorithm converges or the maximum number of iterations is reached.

## 6. Experimental Results

We demonstrate and analyze the results produced by our method. For a  $256 \times 256$  normal map, the EM algorithm converges in 30–40 iterations to estimate the weighted discontinuity map, taking 4–5 secs only. The height field generation takes longer, requiring about 5 minutes. The experiments were performed on a Pentium D 2.8Hz machine with 1GB memory.

A synthetic case will be shown first, which is considered a very difficult case in some previous approaches. Next, a more complicated synthetic case is used to test the effects of the two main estimation components of our method – the weighted discontinuity map  $v$  and the height field  $h$ . The robustness against noisy data will then be demonstrated. Finally, we present our result on real data.

For visualization, we color code the input normal map. Given a normal  $(e_1(x, y), e_2(x, y), e_3(x, y))^T$ , the color of the normal is defined as  $(R = \frac{e_1(x, y) + 1}{2}, G = \frac{e_2(x, y) + 1}{2}, B = e_3(x, y))$ , where  $R$ ,  $G$  and  $B$  are the values of red, green and blue channels respectively.

**Slope.** Figure 1 shows a synthetic example *Slope*. Note that this example is different from and more difficult than the one shown in [11] where the surface transition is smooth along the slab boundary whereas ours consists of  $C^0$  and/or  $C^1$  discontinuities (Figure 1(b) and (c)). Figure 1(e)–(g) are the surfaces produced by integration, [4] and [11] respectively. The Matlab implementation of [4] and [11] were obtained from the author of [11] and is gratefully acknowledged. All these methods completely fail to produce a surface faithful to the ground truth shown in Figure 1(a), because it is incorrect to consider the surface as a single connected one, or use global integrability constraint in typical situation like this. Specifically, the floor should not contribute to the inclined plane during surface reconstruction. By using our method, a weighted discontinuity map  $v$  is estimated (Figure 1(d)) which is used to weight the contribution of neighbors. Our result shown in Figure 1(h)) is faithful to the ground truth.

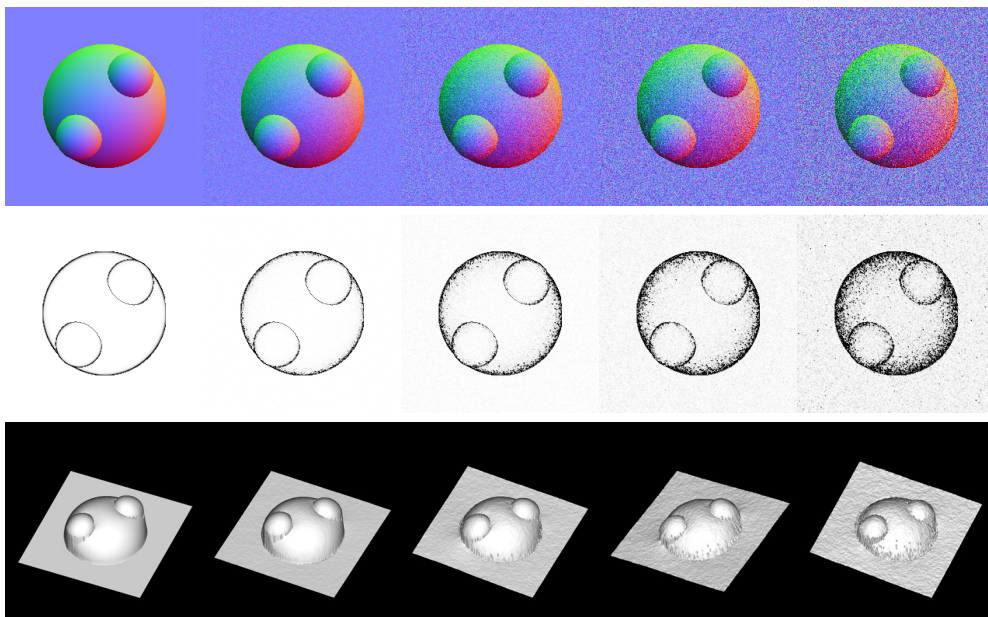


Figure 3. Noise robustness. First row: Color coded input normal map with different noise levels - (from left to right) 0%, 5%, 10%, 15% and 20% Gaussian noise. Second row: the weighted discontinuity maps estimated with the inputs shown in first row. Third row: surfaces reconstructed by using our method with the inputs shown in first row. The experiments show the graceful degradation of our method.

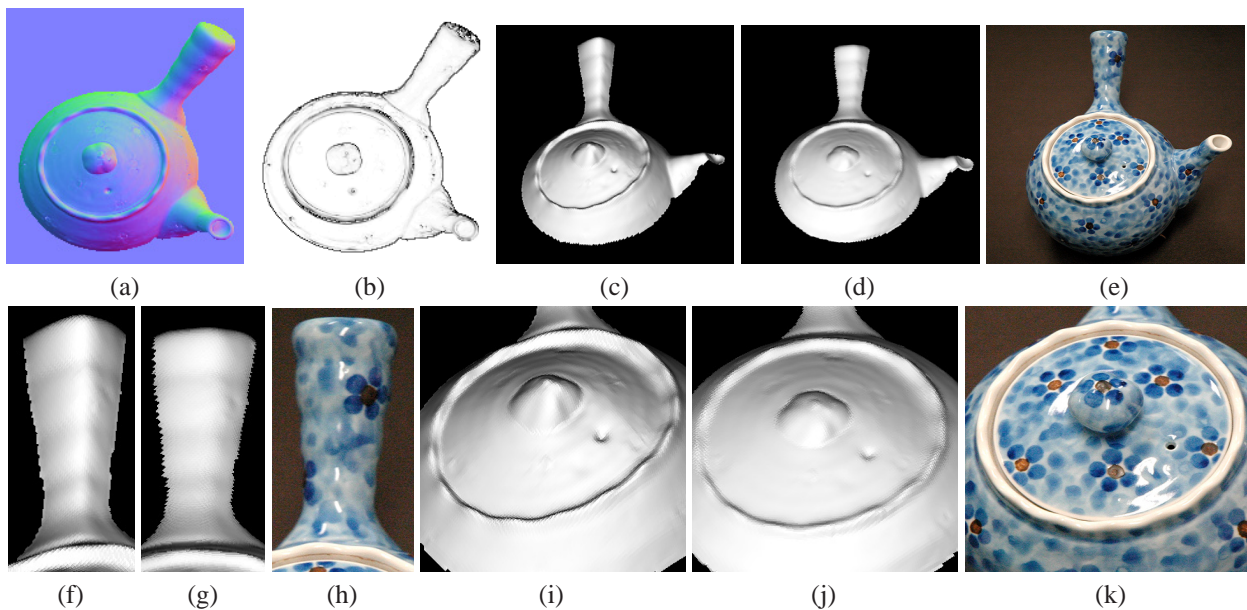


Figure 4. Teapot: (a) Color coded input normal map. (b) The visualized weighted discontinuity map found by our EM algorithm. (c) The surface reconstructed by using [11]. (d) The surface reconstructed by using our method. (e) The real teapot captured with a similar view point as (c) and (d). (f) The zoom-in view of the handle of (c). (g) The zoom-in view of the handle of (d). (h) The zoom-in view of the real teapot handle. (i) The zoom-in view of the lid of (c). (j) The zoom-in view of the lid of (d). The zoom-in view of the real teapot.

**Three Spheres.** Figure 2 shows another example. Again, the methods in [4] and [11] produce distorted results. In our case, we first test the usefulness of the weighted discontinuity map  $v$ . By setting all  $v_{ij} = 1$  and applying only

our height field estimation, which is an analogy of [6], we still produce a decent surface shown in Figure 2(d). Although distortion can be observed around the small spheres and the contact between the plane and the large sphere, the

salient features are still retained. Figure 2(e) shows the result produced by our complete method where less distortion is observed and our result is comparable to the ground truth (Figure 2(f)), showing the necessity of our weighted discontinuity map. Note that a binary discontinuity map is inadequate: the small spheres observed have  $C^0$  continuities and  $C^1$  discontinuities (sharp changes in surface gradients) along the loci of contact with the large sphere, except along the  $C^0$  discontinuities (occlusion boundary) where the small spheres tangentially occlude the large one. Our *weighted* discontinuity map is capable of dealing with this complex situation, restricting the influence of the largest sphere’s contribution to the small spheres along the loci of contact to recover the faithful depths for the small spheres. Figure 2(g) and (h) respectively show the side views of our result and the ground truth for comparison.

**Noise robustness.** From left to right, the first row in Figure 3 shows the input respectively corrupted with 0%, 5%, 10%, 15% and 20% Gaussian noise. The second row shows the corresponding  $v$  estimated. The third row shows the corresponding reconstructed height fields. Note the distortion of the small spheres becomes apparent starting from the 10% case. However, even in the presence of 20% Gaussian noise, salient features are still maintained despite the observable distortion. Observe that we do not smooth out but preserve the “noise” on the reconstructed surfaces in the 5% case. This implies that our approach can indeed preserve fine details or subtle geometry found in real-life objects.

**Teapot.** Finally, we demonstrate our method using a real case. The input normal map is obtained by [18]. We compare our result side by side with the one produced by [11] (Figure 4(c)–(d)). Our result exhibits significantly less distortion. Figure 4(f)–(k) show the zoom-in views of selected regions. Compared with the real object, our result is more reasonable. Note the lid in Figure 4(i) is clearly distorted, where the lower-left and upper-right portion of the ripple-shaped teapot opening are stretched producing an asymmetric appearance. Overall, our method faithfully maintains the shape and preserves the subtle geometry of the teapot without noticeable distortion.

## 7. Conclusion

In this paper, we propose to enforce discontinuity-preserving integrability in the given normal field. This is achieved by our new Expectation-Maximization (EM) algorithm operating on a Markov Random Field (MRF) network, which alternately optimizes for surface smoothness and discontinuities until convergence. The refined normal map and the weighted discontinuity maps are used in our height field generation, which can be formulated into a convex optimization problem. No initial height is required. Our method is shown to be very robust to measurement noise in the input. Our output preserves fine details within each seg-

ment and maintains surface discontinuities between surface segments without noticeable distortion. We believe that our new methods are useful to computer vision algorithms that output an imperfect normal map from which a segmented surface description should be inferred.

## References

- [1] J. Bilmes. A gentle tutorial on the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models. Technical Report ICSI-TR-97-021, ICSI, 1997.
- [2] A. Blake and A. Zisserman. *Visual reconstruction*. MIT Press, Cambridge, MA, USA, 1987.
- [3] D. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice-Hall, 2003.
- [4] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. In *ICCV87*, pages 118–128, 1987.
- [5] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *PAMI*, 6(6):721–741, 1984.
- [6] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *ICCV05*, October 2005.
- [7] B. Horn. Height and gradient from shading. *IJCV*, 5(1):37–76, August 1990.
- [8] J. Hsieh, H. Liao, M. Ko, and K. Fan. Wavelet-based shape from shading. *GMIP*, 57(4):343–362, July 1995.
- [9] B. Karacali and W. Snyder. Partial integrability in surface reconstruction from a given gradient field. In *ICIP02*, pages II: 525–528, 2002.
- [10] B. Karacali and W. Snyder. Reconstructing discontinuous surfaces from a given gradient field using partial integrability. *CVIU*, 92(1):78–111, October 2003.
- [11] P. Kovessi. Shapelets correlated with surface normals produce surfaces. In *ICCV05*, pages 994–1001, 2005.
- [12] D. Mumford and J. Shah. Boundary detection by minimizing functionals, I. In *CVPR1985*, pages 22–26, 1985.
- [13] N. Petrovic, I. Cohen, B. Frey, R. Koetter, and T. Huang. Enforcing integrability for surface reconstruction algorithms using belief propagation in graphical models. In *CVPR01*, pages I:743–748, 2001.
- [14] K. Tang, C. Tang, and T. Wong. Dense photometric stereo using tensorial belief propagation. In *CVPR2005*, volume 1, pages 132–139, June 2005.
- [15] D. Terzopoulos. The computation of visible-surface representations. *PAMI*, 10(4):417–438, July 1988.
- [16] R. Woodham. Photometric method for determining surface orientation from multiple images. *OptEng*, 19(1):139–144, January 1980.
- [17] T. Wu and C. Tang. Dense photometric stereo using a mirror sphere and graph cut. In *CVPR2005*, volume 1, pages 140–147, June 2005.
- [18] T. Wu and C. Tang. Dense photometric stereo by expectation maximization. In *ECCV2006*, volume IV, pages 159–172, May 2006.