

Enforcing Integrability for Surface Reconstruction Algorithms Using Belief Propagation in Graphical Models

Nemanja Petrovic

Ira Cohen

Brendan J. Frey
Beckman Institute
University of Illinois
Urbana, IL 61801

Ralf Koetter

Thomas S. Huang

Abstract

Accurate calculation of the three dimensional shape of an object is one of the classic research areas of computer vision. Many of the existing methods are based on surface normal estimation, and subsequent integration of surface gradients. In general, these methods do not produce valid surface due to violation of surface integrability. We introduce a new method for shape reconstruction by integration of valid surface gradient maps. The essence of the new approach is in the strict enforcement of the surface integrability via belief propagation across graphical model. The graphical model is selected in such a way to extract information from underlying, possibly noisy, surface gradient estimators, utilize the surface integrability constraint, and produce the maximum a-posteriori estimate of a valid surface. We demonstrate the algorithm for two classic shape reconstruction techniques; shape-from-shading and photometric stereo. On a set of real and synthetic examples the new approach is shown to be fast and accurate, in the sense that shape can be rendered even in the presence of high levels of noise and sharp occlusion boundaries.

1 Introduction

Recovering 3D shape of objects, classified as shape-from-X techniques, is a classic and fundamental computer vision research area. Shape-from-X refers to the recovery of shape from stereo, motion, texture, shading, etc. In this paper we focus our attention on the shape recovery from images of a static object made under different lighting conditions, also known as photometric stereo (PMS) and shape-from-shading (SFS).

Shape-from-shading was first studied by Horn in the early 70's and since then there has been a substantial literature dealing with the problem. The principal idea in SFS and PMS algorithms is to invert the mapping from surface gradients to image intensity. The mapping is referred to as the reflectance map. The reflectance map combines information about the light source, surface material and viewing

geometry to form the generally non-linear relationship between the image intensity and surface gradients. Because the mapping is not invertible locally using a single measurement of image intensity, SFS algorithms solve for the surface gradients using several different methods. A class of algorithms use additional constraints on the surface and minimize an energy function [6, 9]. Other algorithms use propagation of shape information from reference points to iteratively solve for the surface gradients, local surface assumptions or linear approximations to the reflectance map. A good reference comparing and describing different methods can be found in [19].

Photometric stereo is a shape-from-shading algorithm using several images to invert the reflectance map. It has been first introduced by Woodham in 1980 [17], and has been an ongoing research problem in the computer vision community. The basic algorithm for PMS estimates the surface gradients locally for each pixel without using global constraints. Among the more advanced methods are those that use local confidence measures to account for surface inter-reflections and shadowing [18, 10].

Not all SFS methods are based on surface normal reconstruction followed by gradient integration. There are methods that directly estimate the absolute height [4]. For these methods there is no need to estimate the normal, and no need to enforce integrability. However, normal based methods are still being used, and we think that problem of accurate reconstruction still remained unsolved at large.

In this paper we describe the use of belief propagation in factor graphs to enforce integrability on valid surfaces. Belief propagation in factor graphs has been shown to be successful in dealing with complicated problems involving functions of many variables that could be factored into local functions. The message passing in factor graphs, referred to as the sum-product algorithm, is an efficient algorithm for computing marginals or posterior of variables in the graph. Various algorithms such as the forward-backward algorithm, the Viterbi algorithm, Kalman filtering, and Pearl's belief propagation in Bayesian networks can be viewed as instances of the sum-product algorithm [12]. In cases of loopy factor

graph, the marginalization is not guaranteed to converge to the correct marginal or maximum *a-posteriori* (MAP) of the variables of interest. However, belief propagation in loopy factor graphs showed promising results in several problems such as turbo-codes implementation [3, 2], 2D phase unwrapping [11], super-resolution [7] and others [13]. Although the successful experimental results have not been rigorously explained yet, recent studies have begun to theoretically analyze and prove convergence properties of belief propagation in loopy graphs [15, 16, 1].

The integrability constraint in many SFS and most PMS algorithms is not strictly enforced. One exception is the work of Frankot and Chellappa [6], in which the integrability is enforced at each iteration of the Horn and Brooks SFS algorithm [9] by orthogonal projection to the subspace of valid integrable surfaces. This method is not easily applicable to a large class of SFS problems and is highly dependent on the choice of basis functions. Enforcing the constraint using factor graphs propagates probabilities across the loopy graph to find a valid integrable surface, which is the MAP solution given the estimated surface gradients and constraints. Enforcing integrability constraints across graphs was first introduced by Koetter et al. [11] for solving the 2D phase unwrapping problem. For phase unwrapping, the most likely discrete jumps in the phase are estimated in such manner to satisfy integrability constraint. For SFS and PMS, we are enforcing integrability on continuous surface gradients.

The rest of this paper is organized as follows: in Section 2, we give a brief description and notations of shape-from-shading and photometric stereo problems. In Section 3, we describe and analyze the factor graphs and messages passing schedule used to enforce the integrability constraint on surface gradients. In Section 4, we show experiments for both PMS and SFS algorithms on real and synthetic images. Finally we summarize our contributions and discuss directions for future work.

2 Photometric stereo and Shape-From-Shading - Description and Notation

Shape reconstruction in PMS and SFS algorithms is achieved by estimating the surface gradients in the x and y direction, noted as $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$ respectively, from the image(s) pixel intensities. To recover the depth map $z(x, y)$, estimated gradients are integrated for each (x, y) along some path starting from (x_0, y_0) . Figure 1 shows the basic configuration of the surface, light and image plane used in the SFS and PMS algorithms. Normally, the image plane is assumed to be parallel to the $x - y$ plane and the image projection is orthographic, although this assumption can be relaxed. The image irradiance is given as a function of p




Figure 1: Basic configuration of the image acquisition in SFS and PMS algorithms

and q ,

$$I(x, y) = R(p, q), \quad (1)$$

where $R(p, q)$ is the reflectance map, combining all information about the surface material, illumination and viewing geometry into a mapping from the surface orientation to image intensity. In many cases, the surface is assumed to be Lambertian, the lighting to be diffused and there are no inter-reflections between surface elements. Using the reflectance map based on these assumptions leads to the following equation for image intensity:

$$I(x, y) = kg(x, y) \cdot \mathbf{S}, \quad (2)$$

where k is a known constant, $\mathbf{g}(x, y) = \rho(x, y)\mathbf{N}(x, y)$, $\rho(x, y)$ is the albedo of the surface, $\mathbf{N}(x, y)$ is the surface unit length normal and \mathbf{S} is the direction of the light source. Note that $\mathbf{N}(x, y) = \frac{\mathbf{g}}{\|\mathbf{g}\|}$ is the surface normal at point (x, y) . The surface normal is related to the partial derivatives $\frac{\partial z}{\partial x}$ and $\frac{\partial z}{\partial y}$ through the formula:

$$\mathbf{N}(x, y) = \frac{1}{\sqrt{1 + (\frac{\partial z}{\partial x})^2 + (\frac{\partial z}{\partial y})^2}} \left\{ \frac{\partial z}{\partial x}, \frac{\partial z}{\partial y}, 1 \right\}^T. \quad (3)$$

The problem in PMS is to estimate \mathbf{g} given multiple images $I(x, y)$. Because we are able to control the direction of the light sources \mathbf{S} , with sufficient number of light sources (two or three depending on the prior knowledge of albedo), we would be able to estimate \mathbf{g} using a least squares (LS) solution for over-constrained system [5]. This is done by stacking Eq. 2 for all light sources in matrix form:

$$\mathbf{I}(x, y) = k\mathbf{V}\mathbf{g}(x, y), \quad (4)$$

where $\mathbf{I}(x, y) = [I_1(x, y), \dots, I_n(x, y)]^T$, and \mathbf{V} is a $n \times 3$ matrix whose i^{th} row is \mathbf{S}_i , and n is the number of light sources. Given noiseless data, this solution is expected to give a good approximation of the surface. Forsythe and Ponce [5] describe a slightly modified representation of Eq. 4

that takes care of singular points (points in which the intensity is zero for all images). In our experiments we use the modified equations.

For the SFS problem, Eq. 2, or its equivalent under different reflectance map assignments, places a single constraint on the two degrees of freedom of the problem. Typically, this problem is overcome by imposing some form of smoothness constraint on the field of surface normals.

Let us denote the partial derivative with respect to a variable with the variable name in the subscript. If at some point $p_y \neq q_x$ the line integral along any contour around that point will give nonzero value, which is, in terms of vector fields, a violation of 'zero curl' of a gradient field. A direct consequence is that the reconstructed surface, $z(x, y)$, will depend on the integration path. We will refer to the equality of the second order partial derivatives as the integrability constraint, i.e.,

$$\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x} \text{ or,} \\ p(x, y) - p(x, y + 1) = q(x, y) - q(x + 1, y). \quad (5)$$

For most of the SFS and PMS algorithms the estimated surface gradients, p and q , need not satisfy the integrability constraint. However, for most algorithms, the true surface gradient can not be too far from these estimated values. Belief propagation in the factor graph would try to correct all zero curl violations, while taking into account the original estimates. We will demonstrate that imposing the integrability over elementary loops in the graphical model will correct the irregularities in the data caused by either shadowing, occlusion boundaries or noise.

Although we describe the algorithm for the PMS and SFS problems only, there are no formal reasons why our approach would not be used in any shape reconstruction algorithm that estimates surface gradients and integrates over the gradients to recover the surface.

3 Using Factor Graphs to Enforce Integrability for Surface Reconstruction

Factor graph is a graphical representation of a function that could be factored into product of simpler terms. Factor graph consists of two types of nodes: function nodes (black circles in Figure 2), and variable nodes. Variable nodes, in turn, could be hidden (white circles) or observable (gray circles). Observable nodes are surface gradient estimates obtained from a shape-from-X algorithm ($\hat{\Delta}_x, \hat{\Delta}_y$), which might not satisfy the integrability constraint. The hidden nodes represent the unknown true values of partial derivatives (gradients) of the surface height. We will refer to

them as the gradient nodes. Each hidden node is connected to his observable node via a function node, which is the conditional probability density function (pdf) of the hidden node given the observation. The conditional pdf is estimated from training data, knowledge of the problem, or previous experience. The function nodes connecting four hidden nodes enforce the zero curl constraint of each elementary loop. They are indicators whether Eq. 5 is satisfied. Hence, we name them check nodes, the term borrowed from error correcting codes literature. The task is to infer the values of gradient nodes, given the observable nodes, and known constraints. Inference is carried out by iterative propagation of messages between check and gradient nodes.

This graphical model represents the joint probability density function of all unknown surface gradients given their initial noisy estimates and indicators whether the integrability constraint is satisfied for the whole surface. The sum-product algorithm provides an efficient way to estimate the marginals over the gradient nodes.

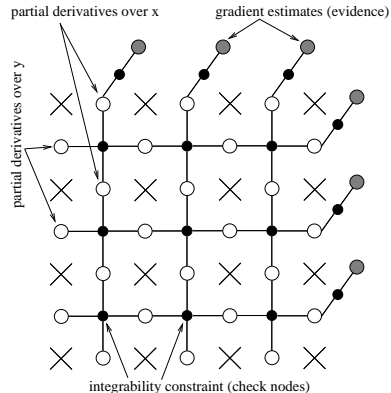


Figure 2: Factor graph used for surface reconstruction. White nodes are the hidden variables related to partial derivatives (gradients) to be estimated. Gray nodes are the observed nodes - inputs from PMS or SFS algorithm. Black nodes in the interior of the grid are the check nodes.

3.1 Message Passing in the Factor Graphs

All messages passed in the factor graph are in the form of conditional probability density functions. The sum-product algorithm provides two simple rules how the messages are handled in the variable and check nodes. For the check node with m neighboring variable nodes (Figure 2(b)) messages from $m - 1$ variable nodes are multiplied mutually and multiplied with the value of function node for the given configuration of variables. This product is then summed over all configurations of $m - 1$ variables, and the result is sent to remaining m^{th} variable node. For the variable node with n neighboring check nodes (Figure 2(c)) messages from $n-1$

check nodes are multiplied and sent to remaining n^{th} check node. In the last iteration all messages arriving to variable node are combined (Figure 2(d)).

Initially, all gradient nodes receive a single message, ν_0 from their observable nodes. Observable nodes do not receive any messages and their values do not alter. The message passed from an observable node to its gradient node is the conditional probability density function of the hidden variable given its local observation. For ease of notation, we refer to either p or q as h , and to local observation as lo . We assume Gaussian density, i.e.,

$$\nu_0 = f(h|lo) \sim \mathcal{N}(h; \mu_{lo} = \hat{\Delta}_x, \sigma_{lo}^2), \text{ or} \quad (6)$$

$$\mathcal{N}(h; \mu_{lo} = \hat{\Delta}_y, \sigma_{lo}^2)$$

where σ_{lo} is level of uncertainty *equal* for all observations.

At each iteration of the sum-product algorithm, the gradient nodes send their messages to all the check nodes they are connected to. As shown in Figure 2(b), the check node receives messages ν_1, ν_2 and ν_3 from its three neighbors and, after processing, sends the message ν_4 to its fourth neighbor. This process is repeated for each of the gradient nodes connected to a check node. The message being passed to the fourth gradient node can be thought of as an evidence of the value of the fourth variable given the values of other three variables when the integrability constraint is satisfied in the elementary loop, i.e., $p(x, y + 1) = -q(x, y) + p(x, y) + q(x + 1, y)$. Therefore ν_4 is the convolution of the three Gaussian pdfs, in effect it is a Gaussian with,

$$\mu_4 = -\mu_1 + \mu_2 + \mu_3$$

$$\sigma_4^{-2} = \sigma_1^{-2} + \sigma_2^{-2} + \sigma_3^{-2}, \quad (7)$$

The variable nodes receive two types of messages: (1) message from their local observation, and (2) messages from neighboring check nodes. In each iteration of sum-product algorithm, apart from the last, the gradient nodes multiply the message from one check node with the local observation and send the product to the other check node (Figure 2(c)). This message passing schedule forces messages to traverse a long path before eventually returning to originating node. Because the product of two Gaussians is again a Gaussian, in each iteration of the sum-product algorithm all messages are Gaussians. In the final iteration of the sum-product algorithm all messages sent to a gradient node are multiplied (Figure 2(d)). The posterior of the gradient node in terms of three final messages is

$$f(h|all\ observations) = \nu_0 \cdot \nu_1 \cdot \nu_2 \quad \Rightarrow$$

$$f(h|all\ observations) \sim \mathcal{N}(\mu_h, \sigma_h^2) \quad \text{where,}$$

$$\mu_h = \left(\frac{\mu_{lo}}{\sigma_{lo}^2} + \frac{\mu_{ch1}}{\sigma_{ch1}^2} + \frac{\mu_{ch2}}{\sigma_{ch2}^2} \right) \cdot \frac{1}{\sigma_{lo}^{-2} + \sigma_{ch1}^{-2} + \sigma_{ch2}^{-2}} \quad \text{and,}$$

$$\sigma_h^2 = \frac{1}{\sigma_{lo}^{-2} + \sigma_{ch1}^{-2} + \sigma_{ch2}^{-2}}, \quad (8)$$

and where lo relates to the local observation ν_0 , and ch_1, ch_2 relate to messages from the check nodes, ν_1, ν_2 . The mean of a resulting Gaussian, μ_h , is the MAP estimate of the unknown surface gradient.




Figure 3: Message passing for an elementary loop. (a) shows the messages sent from three gradient nodes to the check node, and the message passed to the fourth node. (b) shows the messages sent from local observation and one check node, and the message passed to second check node. (c) shows combination of all messages in the final iteration.

4 Experiments

To illustrate the advantages of our method we show results for shape-from-shading and photometric stereo. Because there is a large variety of algorithms for both PMS and SFS, we decided to choose just one algorithm for each method. The main purpose of the experiments is to illustrate how post-processing the output of shape-from-X algorithms using belief propagation successfully enforces the integrability. The experiments also demonstrate that enforcing the constraint yields a better surface reconstruction, both visually and analytically.

4.1 Photometric Stereo Experiment

We use depth map images of three objects available from [8] to generate three images of each object using different illuminants, and add significant amount of white Gaussian noise to the original noiseless depth map (approximately 10% noise variance). The 3D reconstruction of the objects from the noiseless depth maps can be seen in Figure 4(a).

The basic algorithm for estimating the surface normals and albedo is described in [5] and in Eq. 4. Because the least-square estimation does not necessarily produce a valid smooth surface, the estimated first order derivatives serve as the local observations to the gradient nodes, as shown in Figure 2. Integrating the finite differences from the pixel (1,1) yields the reconstructed surfaces as seen in Figure 4(b). Because the direction of integration is fixed, the errors propagate and accumulate, and the reconstructed surfaces are noisy. The first column of Table 1 shows the percentage of elementary loops violating the integrability. An elementary loop violates the integrability constraint if the absolute




Figure 4: Surface reconstruction. Objects from top to bottom: Mozart, Sombrero, Penny, Lenna.

value of the integral of the four neighboring gradients is greater than ϵ .

The number of iterations needed to achieve zero number of curl violations is shown in the left column of Table 1. For the Sombrero and Mozart surfaces, it takes few iterations, due to the fact that the surface is very smooth. The Mozart surface is more complex, as it includes a sharp occlusion boundary and fine surface details. At close inspection, it is apparent that discontinuities are maintained, and there is a minimal loss of the details in the surface. This is a compelling feature of the method - at each iteration the value of the original measurement is always used as the local observation, and its effect does not fade away. For the Penny example more iterations are needed to correct the curl violations. The Penny example has sharp boundaries, and the bowl shape of the depth map causes occlusions. Nevertheless, most errors are corrected in the first few iterations, and a small percentage of errors at the outer region of the surface require further iterating. Using the real depth maps of the surfaces we compute the mean squared error (MSE) for each example. The MSE and the variance of the MSE for the reconstructed surface before and after the belief propagation are shown in Table 1.

4.2 Shape-from-shading Experiment

We use an algorithm proposed and implemented by Tsai and Shah [14]. For the tests, we use a real photo of Lenna. The initial reconstruction has a high percentage of integrability violations on the surface (Table 1), though not large in magnitude, so that the initial results are visibly good (Figure 4(b)). Only a few iterations were needed to achieve a visibly better reconstruction, shown in Figure 4(c).

5 Summary and Conclusions

In this paper we show how enforcement of the integrability constraint on estimates of surface gradients is efficiently applied using belief propagation across factor graphs. This is an important step in surface reconstruction, both for visualization and recognition applications using 3D object information.

We demonstrate the algorithm for two related shape-from-X methods, namely, shape-from-shading and photometric stereo. Good results are shown for real and synthetic images. Enforcing the integrability handles the noise without smoothing out fine details. The result is also reflected in the significant improvement of the MSE between the original surface and the reconstructed one. Apart from the visual

Table 1: Initial percentage of elementary loops violating the curl constraint, number of iterations needed to achieve zero violations and Mean Square error of surface reconstruction for the PMS synthetic examples. σ_{MSE} is the standard deviation of the MSE

Object	Initial % violations	# of iterations	Initial MSE + σ_{MSE}	Final MSE + σ_{MSE}
Sombrero	95%	27	$3.4 \cdot 10^{-3}$, $\sigma = 6.8 \cdot 10^{-3}$	$0.48 \cdot 10^{-3}$, $\sigma = 1.7 \cdot 10^{-3}$
Mozart	94%	28	$1.7 \cdot 10^{-3}$, $\sigma = 5 \cdot 10^{-3}$	$0.7 \cdot 10^{-3}$, $\sigma = 5 \cdot 10^{-3}$
Penny	81%	124	$5 \cdot 10^{-3}$, $\sigma = 18.1 \cdot 10^{-3}$	$0.95 \cdot 10^{-3}$, $\sigma = 0.78 \cdot 10^{-3}$
Lenna	96%	59	N/A	N/A

improvement, it is fair to assume that pattern recognition algorithms would benefit from a more accurate surface estimation. For future work it remains to be seen if similar results can be achieved by applying the method to other shape-from-X algorithms. Also, learning the variance of the local observation pdf from the data would assign higher variance to areas that have larger violations. Another improvement would be to enforce integrability constraint on larger loops than the elementary quads, possibly improving the robustness of the algorithm.

Acknowledgments

We would like to thank Prof. Yi Ma for valuable discussions, encouragements and Prof. Robert Fossum for his editorial help. This work has been supported by NSF Grants IIS-00-85980, EIA-99-75019, CDA-96-24396 and Hewlett-Packard laboratories fellowship.

References

- [1] Frey B.J. and R. Koetter. Exact inference using the attenuated max-product algorithm. in *Advanced Mean Field Methods: Theory and Practice*, 2000.
- [2] Berrou C. and Glavieux A. Near optimum error correcting coding and decoding: Turbo-codes. *IEEE Transactions on Communications*, 44:1261–1271, 1999.
- [3] MacKay D.J.C. Good error-correcting codes based on very sparse matrices. *IEEE Transactions on Information Theory*, 45(2):399–431, 1999.
- [4] P. Dupuis and J. Oliensis. Direct method for reconstructing shape from shading. In *CVPR92*, pages 453–458, 1992.
- [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. *PAMI*, 10(4):439–451, July 1988.
- [7] W.T. Freeman and E.C. Pasztor. Markov networks for super-resolution. In *Proceedings of CISS*, 2000.
- [8] Vision Group. A fast linear shape from shading and sample images. In <ftp://eustis.cs.ucf.edu/pub>.
- [9] B.K.P. Horn and M.J. Brooks. The variational approach to shape from shading. *CVGIP*, 33(2):174–208, February 1986.
- [10] Y. Iwahori, R.J. Woodham, M. Shoaib Bhuiyan, and N. Ishii. Neural network based photometric stereo for object with non-uniform reflectance factor. In *Proceedings of ICONP*, pages 1213–1218, 1999.
- [11] R. Koetter, B. J. Frey, N. Petrovic, and D. C. Munson Jr. Unwrapping phase images by propagating probabilities across graphs. In *ICASSP*, 2001.
- [12] F. R. Kschischang and B. J. Frey. Factor graphs and the sum-product algorithm. *IEEE Journal on Information Theory*, 47:498–519, 2001.
- [13] K.P. Murphy, Y. Weiss, and M.I. Jordan. Loopy belief propagation for approximate inference: and empirical study. In *Proceedings of UAI*, 1999.
- [14] P.S. Tsai and M. Shah. Shape from shading using linear-approximation. *IVC*, 12(8):487–498, 1994.
- [15] Y. Weiss. Correctness of local probability propagation in graphical models with loops. *Neural Comp*, 12:1–41, 2000.
- [16] Y. Weiss and W.T. Freeman. On the optimality of solutions of the max-product belief propagation algorithm in arbitrary graphs. In *MERL TR99-39*, 2000.
- [17] R.J. Woodham. Photometric method for determining surface orientation from multiple images. *OptEng*, 19(1):139–144, January 1980.
- [18] R.J. Woodham. Gradient and curvature from the photometric-stereo method, including local confidence estimation. *JOSA-A*, 11(11):3050–3068, November 1994.
- [19] R. Zhang, P.S. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Transactions on PAMI*, 21(8):690–706, August 1999.