

Surface-from-Gradients with Incomplete Data for Single View Modeling*

Heung-Sun Ng Tai-Pang Wu Chi-Keung Tang
Vision and Graphics Group
The Hong Kong University of Science and Technology
Clear Water Bay, Hong Kong

Abstract

Surface gradients are useful to surface reconstruction in single view modeling, shape-from-shading, and photometric stereo. Previous algorithms minimize a complex, non-linear energy functional, or require dense surface gradients to perform integration to generate 3D locations, or require user-input heights to constrain the solution space, or produce severe distortion and smooth out surface details. Most single-view algorithms output a Monge patch (height-field), which may introduce further surface distortion along object silhouettes and surface orientation discontinuities. Our proposed algorithm operates on a single view of complete or incomplete data. The data can be gradients without 3D locations, or 3D locations without gradients. The output surface, which is not necessarily a height-field, preserves salient depth and orientation discontinuities. Experimental comparisons on both simple and complex data show that our method produces better surfaces with significantly less distortion and more details preserved. The implementation of our closed-form solution is very straightforward.

1. Introduction

Single view modeling [25, 15], photometric stereo [7, 6, 10, 4, 22] and geometric stereo [18, 17] are important topics in computer vision. These algorithms typically produce a sparse or dense normal/disparity map, from which a surface is reconstructed. While existing surface estimation algorithms already produce good surfaces, they have one or more of the following limitations:

- A complex, non-linear energy functional is minimized, so initialization and convergence rate are issues.
- Initial user-input heights are necessary to constrain the solution space.
- When no heights are available, dense surface gradients are required to integrate the underlying surface.

*This research was supported by the Research Grant Council of Hong Kong Special Administrative Region, China, under grant no. 620006.

- Surface details are either smoothed out in the output, or detected as crease curves and surface-curve integration is performed as a separate process.
- Many surface-from-gradients methods output a Monge patch ($h = f(x, y)$), that is, a height-field), which cannot handle normals lying on the image plane. To handle these normals, their orientations must be altered to satisfy the integrability constraint, which introduces distortion to the computed surface.

Given a single view of dense or sparse data, which can be surface gradients without 3D locations, or 3D locations without surface gradients, or a combination of both, we propose a uniform algorithm to reconstruct the underlying surface from the single view data. Our closed-form solution offers good numerical stability, where initialization and convergence rate are not issues. The output surface preserves important surface details, and does not suffer surface distortion common to height-field output. Such output is highly desirable for tasks in computer vision, image-based modeling, and computer graphics.

2. Related work

Surface reconstruction for single view modeling has specific challenges (e.g., heights are usually unavailable, or a few user-input heights are required), properties (e.g., the output is a visible surface) and requirements (e.g., the reconstructed surface should be highly-detailed for graphics-related tasks); all of which can be actively exploited to produce results better than those produced by general 3D surface fitting algorithms, which are targeted at generating generic surface descriptions for vision-related tasks.

Terzopoulos et al. [21] inflated a surface using thin plate energy given sparse geometric cues. Later, the method presented in [20] provides a compact framework for visible surface reconstruction, which considers simultaneously surface normal, height field, orientation discontinuity and depth discontinuity. The energy functional is non-linear and complex, so incomplete height field is needed for initialization. Other regularization approaches also minimize a




Figure 1. Normals lie on the image plane. For the methods proposed in [5, 9], these normals have to be altered before running the algorithms. Our method produces a faithful ‘staircase’ surface without manually changing input normals.

functional to derive an optimal surface while preserving discontinuities. For instance, the Mumford-Shah model represents discontinuity using a discrete set. Blake and Zisserman [2] developed a weak membrane model that explicitly deals with discontinuities, by iteratively solving for the labeling of data and the model in a functional optimization framework. [1] focused on discontinuity and introduced anisotropic weights for better surface estimation. The complexity and/or the heuristic nature inherent in these approaches have inspired non-iterative approaches to address the surface inference problem.

A non-iterative method, tensor voting [11], was developed by Medioni et al. to infer a layered surface description, from a combination of points, points with tangents, and points with normals, which can be sparse or dense. Surfaces, curves and junctions are inferred. Given sparse input, since the detected curves are often not localized well on the inferred surface, an integration process should be performed [19].

When dense surface gradients are available, the integrability constraint can be applied to obtain the height-field (e.g., notably [6] where discontinuity and height constraints can be imposed readily). One classical approach that makes use of basis functions was proposed by Frankot and Chellappa [5], where Fourier basis functions were used. More recently, shape from shapelets was proposed [9]. This method accepts dense normals as input, and correlates the measured surface gradients with the gradients of a bank of shapelet basis functions which satisfy a set of constraints. Discontinuity and height constraints, however, cannot be imposed readily in these methods based on bases. When both dense surface gradients and dense position constraints are present, one can obtain decent 3D surfaces by using [13].

In [25], an interactive approach on single-view modeling was presented. The method takes as input a combination of user-specified constraints, such as surface positions, normals, silhouettes, and creases to generate a height-field that satisfies the supplied constraints. The associated energy functional is based on thin-plate-spline. Typical of

height-fields, surface distortion is particularly severe along object silhouettes. In [15], a closed-form solution was proposed for single-view reconstruction of curved 3D surfaces. The formulation is discrete in nature. In particular, the user needs to supply the inflation constraints, by guessing surface heights at sample image locations. The input normals are specified in the discrete parameter space, which is unintuitive compared with directly specifying normals in the spatial domain.

3. Motivation

In this section, we present cases to illustrate that our closed-form solution has advanced previous results, using a classical approach [5] and a new approach [9] for comparison on surface-from-gradients, where no heights are available.

Height-field and integrability. Figure 1 shows that we can handle case where normals are nearly lying on the image plane, i.e. $\frac{\partial f(x,y)}{\partial x}$ and/or $\frac{\partial f(x,y)}{\partial y} \rightarrow \infty$. For [5, 9], which were designed to handle height-fields, fail to do so because these normals introduce severe numerical instability. In order to generate the suboptimal output for [5, 9] shown in Figure 1, the normals in question must be perturbed. Albeit such manual perturbation, the output surface is still unsatisfactory. The observed distortion is in fact due to enforcement of the integrability constraint, which requires changing the orientations of the normals. This is, however, problematic because this ‘staircase’ surface is characterized by these normals, which define the surface orientation discontinuity. Unfortunately, this discontinuity exactly violates the integrability constraint.

Dense normals with no heights. Figure 2 compares the surfaces generated using an input dense ‘needle map’ of a sphere with ‘ICCV’ engraved on the spherical surface. Our method produces a hemispherical shape, where all observable surface details are faithfully preserved. For [5, 9] to work properly, we perturbed the normals along the sphere’s outline in order to avoid division-by-zero. Even this, needle-like spikes are produced as shown, when the z

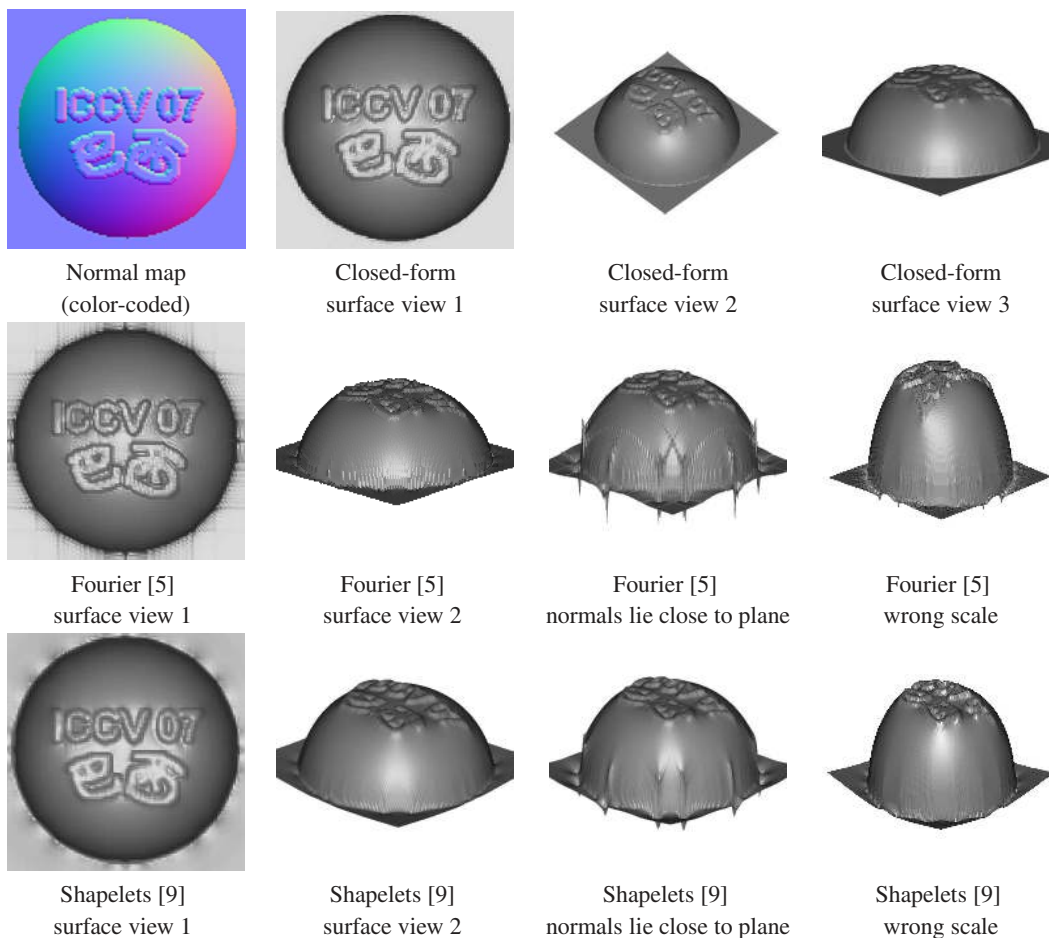


Figure 2. Sphere. The x, y, z components of the input dense normals are color-coded in RGB. Top: Closed-form solution. Middle: [5]. Bottom: [9]. For our closed-form solution, normals lying on image plane need not be altered, and no user-input parameter is required.




Figure 3. Button. From left to right: the sparse needle map, the discontinuity map, results generated by [5], [9], our closed-form solution. For [5, 9], the missing normals are assumed to be $[0\ 0\ 1]^T$.

components of normals are close to zero. Also, these methods rely on a good input scale, which is guessed by trial-and-error. Our method, on the other hand, does not require such user-input parameter.

Sparse normals, with discontinuity but no heights. The third example shown in Figure 3 demonstrates the ability of different approaches in incorporating discontinuities into consideration. Such sparse normals and discontinuity

map on a single view are typically available in single-view modeling [25, 15, 24]. Different from manually guessing heights, we can specify normals and discontinuities easily via a 2D sketching interface [24]. Because the input is quite sparse, integration approaches cannot be applied to this case. Our closed-form solution, on the other hand, produces a surface faithful to the input gradients where the specified discontinuity is preserved. For [5, 9], there is

no easy way to impose the given discontinuity constraint, where the enforcement of integrability constraint further aggravates the distortion problem.

4. A Closed-Form Solution to Surface-from-Gradients with Incomplete Data

In this section, we describe how we can generate a 3D surface given the complete or incomplete 3D and first-order gradient information, within a unified framework implemented as a closed-form solution. The following lists the information that can be obtained from a single view.

Normals/Gradients They can be either sparse or dense, residing on a 2D regular or irregular grid space.

3D points They can be either sparse or dense, scattering in a finite space.

Discontinuities They describe the occlusion boundaries or orientation discontinuities between surface regions.

Note that, as mentioned, previous methods can either handle a subset of the situations listed above, or produce results that exhibit severe distortion, or use a complex problem formulation.

Now, we show that, by formulating the above into pertinent equations, we can derive a closed-form solution to solve for the corresponding 3D point observed at an image location. The collection of the (sparse) 3D points solved will be employed to compute the underlying surface.

Given the information listed above, to avoid producing a Monge patch and to generate a true 3D surface, we want to find a mapping function $\mathbf{f} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ that maps an image point to the corresponding 3D location. The form of the function can be defined as:

$$\mathbf{f}(\mathbf{t}) = (f_x(\mathbf{t}), f_y(\mathbf{t}), f_z(\mathbf{t}))^T \quad (1)$$

where $\mathbf{t} = (u, v)^T \in \mathbb{R}^2$ in image coordinate, $f_a : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $a \in \{x, y, z\}$. In terms of parametric surface representation, one may consider that the underlying 3D surface is parameterized by \mathbf{t} .

4.1. Representation

By Eqn. 1, \mathbf{f} is characterized by f_a . So, the question we want to ask is *what is the form of f_a ?* Since f_a can be any function, we need a representation that provides as much flexibility as possible. One feasible and convenient way to model f_a is by using *kernel functions* [16] or *radial basis functions* [3]. Mathematically,

$$f_a(\mathbf{t}) = \sum_i^N \alpha_{ai} k(\mathbf{t}_i, \mathbf{t}) \quad (2)$$

where $k : \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ is a chosen kernel, N is the total number of kernel sites (i.e., input and non-input sites

depending on image resolution), and $\{\alpha_{ai}\}$ is a set of unknowns to be estimated.

Theoretically, any differentiable function can be modeled by Eqn. 2, given a sufficient number and a correct distribution of kernel sites, and a chosen kernel function. While the discussion is out of the scope of this paper, we will show that in our experiments that a simple arrangement is sufficient to produce very decent results. We can set the kernel sites to be resided either on the image grid, or in a best resolution that contains the sparse input data. Although the kernel sites are discrete, the resultant function f_a is continuous. We will discuss the choice of kernel later.

4.2. Data Constraints

Here we translate the available information (complete or incomplete) provided by a single view into constraints for the estimation of $\{\alpha_{ai}\}$.

Let us first assume the input is formed by orthographic projection, which is usually assumed in single-view modeling. The u and v components of the image points must coincide with the x and y components of the corresponding 3D points, which still allows for estimating a 3D surface but not a Monge patch. Suppose we have some incomplete heights (they can be obtained from object silhouettes where $z = 0$, disparity, user-supplied cues, depending on applications), the z components of the corresponding image points can be set to the known values. To translate this paragraph into mathematical description, we let $\mathcal{H} = \{\mathbf{s}_j \in \mathbb{R}^2 | j = 1, 2, \dots, m\}$ be the set of image locations where the heights are known. The estimated surface must satisfy:

$$\mathbf{f}(\mathbf{s}_j) = (\mathbf{s}_j^T, h_j)^T \quad (3)$$

where h_j is the known height value at \mathbf{s}_j . In general, $\mathbf{f}(\mathbf{s}_j) = \mathbf{v}_j$, where \mathbf{v}_j is the 3D location that corresponds to \mathbf{s}_j , can be used to encode known 3D locations, without the assumption of orthographic projection.

Besides, we may have some known normals at image locations $\mathcal{V} = \{\mathbf{c}_j \in \mathbb{R}^2 | j = 1, 2, \dots, n\}$ (note that $\mathcal{H} \cap \mathcal{V}$ can be non-empty). To force the estimated surface to align with the known normals corresponding to \mathcal{V} , the tangents along u and v directions in the set \mathcal{V} in the parameter space must be perpendicular to the given normals. Let $\mathbf{n}(\mathbf{t}) \in \mathbb{R}^3$ be the surface normal at \mathbf{t} and $\mathbf{q}_p(\mathbf{t}) = (\frac{\partial f_x(\mathbf{t})}{\partial p}, \frac{\partial f_y(\mathbf{t})}{\partial p}, \frac{\partial f_z(\mathbf{t})}{\partial p})^T$ be the tangent along direction $p \in \{u, v\}$ at \mathbf{t} , the estimated surface must satisfy:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \mathbf{f}(\mathbf{c}_j) = \mathbf{c}_j \quad (4)$$

$$\mathbf{q}_p(\mathbf{c}_j)^T \mathbf{n}(\mathbf{c}_j) = 0 \quad (5)$$

When \mathbf{c}_j is in \mathcal{H} , Eqn. 4 is duplicated and can be removed.




Figure 4. Noise robustness. Top: color-coded normal map. Bottom: the reconstructed surface. From left to right: The normals are perturbed with additive Gaussian noise with standard deviations of 0, 0.3, 0.6, 0.9 radians respectively.

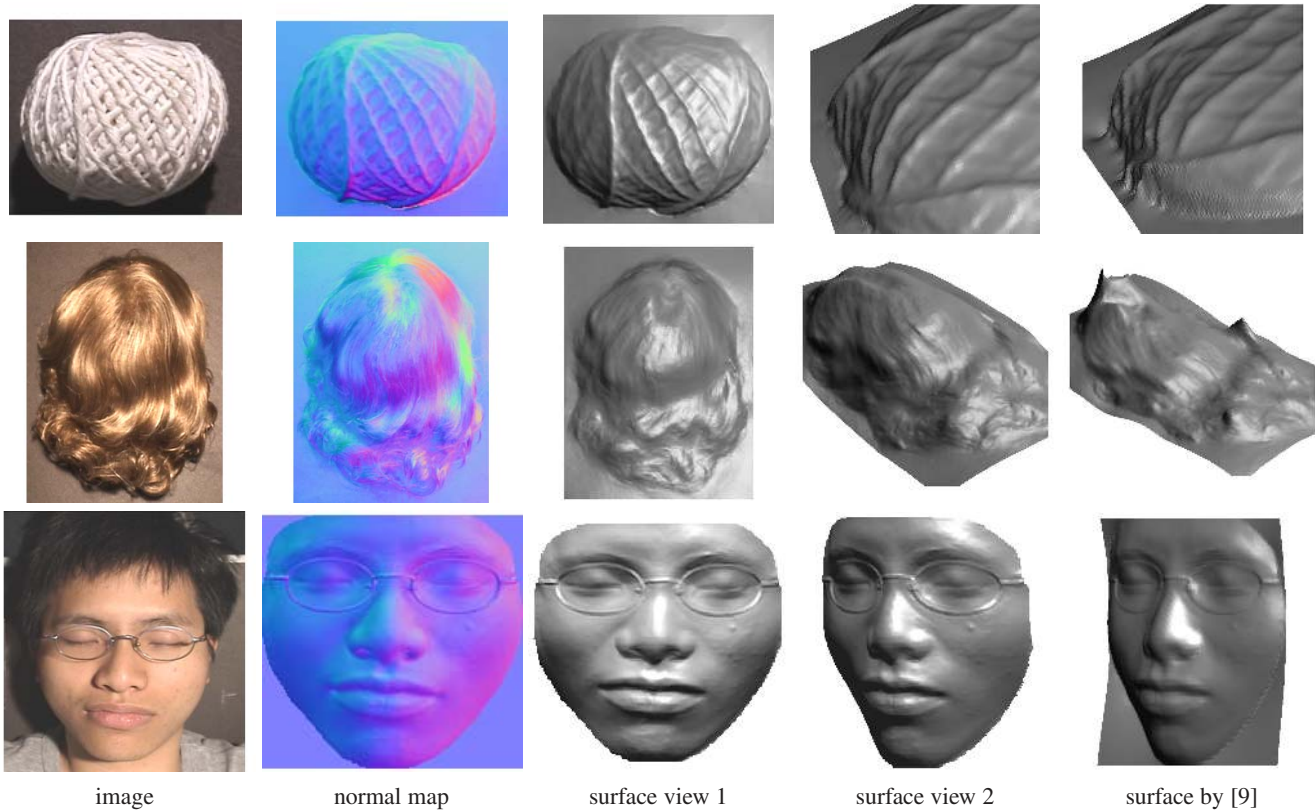


Figure 5. Results on dense gradients. The three examples are *Rope*, *Hair* and *Face*. From left to right: image, color-coded normal map obtained by photometric stereo, two views of the reconstructed surfaces. The corresponding zoom-in views of the surface generated by [9] are also shown.

Let

$$\Lambda = (\alpha_{x1}, \dots, \alpha_{xN}, \alpha_{y1}, \dots, \alpha_{yN}, \alpha_{z1}, \dots, \alpha_{zN})^T \quad (6)$$

we can rewrite the above three sets of constraints (Eqn. 3, 4 and 5) into the following form:

$$\mathbf{A}\Lambda = \mathbf{b} \quad (7)$$

Now, all the known normals and 3D information are encapsulated in Eqn. 7, which is a set of linear equations. Except for normals and incomplete 3D, one may incorporate any constraint into Eqn. 7, as long as the constraints can be modeled by linear equations. This provides an interface for future extension.

4.3. Inference Capability

Since the available information can be incomplete, we need to define a term to incorporate inference or interpolation capability to the algorithm. The idea is based on minimizing the smoothness over the estimated surface.

To propagate the shape information from one image location to another, we can consider first-order neighborhood and minimize their difference. In order to produce smooth surfaces for non-input sites with no information, we minimize the surface gradients or tangents between kernel sites while preserving discontinuities. Let $(s, t) \in \mathcal{N}$ be the set of first order neighbor locations of site p , we minimize the following:

$$w(\mathbf{t}_s, \mathbf{t}_t) \|\mathbf{q}_p(\mathbf{t}_s) - \mathbf{q}_p(\mathbf{t}_t)\|^2 \quad (8)$$

where $w(\mathbf{t}_s, \mathbf{t}_t) \in \{0, 1\}$. If normal or depth discontinuity exists between \mathbf{t}_s and \mathbf{t}_t , $w(\mathbf{t}_s, \mathbf{t}_t) = 0$. Otherwise, $w(\mathbf{t}_s, \mathbf{t}_t) = 1$. In general, $w(\mathbf{t}_s, \mathbf{t}_t)$ can be a soft label in the range of $[0, 1]$ (user-supplied or computed by, for instance, [23]). By considering all $(s, t) \in \mathcal{N}$, we can rewrite the set of Eqn. 8 into matrix form:

$$\|\mathbf{S}\mathbf{A}\|^2 \quad (9)$$

and this term will be incorporated during estimation and thus 3D location can be inferred at each site.

4.4. Solution

With Eqn. 7 and Eqn 8, we can obtain our energy function as the following:

$$E(\mathbf{f}) = \|\mathbf{A}\mathbf{f} - \mathbf{b}\|^2 + \lambda \|\mathbf{S}\mathbf{A}\|^2 \quad (10)$$

where λ is the weight indicating the importance of the smoothness term (The typical value of λ is 0.1). We can see that the surface estimation problem is now translated into a regression problem. To minimize the energy function, we set the first derivative of the energy function w.r.t. \mathbf{f} to zero, and obtain a set of linear equations:

$$(\mathbf{A}^T \mathbf{A} + \lambda \mathbf{S}^T \mathbf{S}) \mathbf{f} = \mathbf{A}^T \mathbf{b} \quad (11)$$

To solve Eqn. (11), we can apply Gauss-Seidel method with successive over-relaxation to obtain \mathbf{f} , and so we can find the 3D location for each image point.

At this point, one may notice that the size of matrix \mathbf{A} and \mathbf{S} can be very large for a dense input, for example, a normal map produced by photometric stereo. To make the estimation practical, the choice of the kernel function k is very important. We choose $\exp(-\frac{\|\mathbf{t}-\mathbf{t}'\|^2}{\sigma})$ as our kernel because Gaussian is simple and it leads to a sparse system so that both the computational and storage requirements can be reduced in implementation.

To take the depth discontinuity into account, we can multiply one weighting function to the Gaussian kernel and obtain our final kernel function as the following:

$$k(\mathbf{t}, \mathbf{t}') = d(\mathbf{t}, \mathbf{t}') \exp(-\frac{\|\mathbf{t}-\mathbf{t}'\|^2}{\sigma}) \quad (12)$$

where $d(\mathbf{t}, \mathbf{t}') \in \{0, 1\}$. If depth discontinuity exists between the straight line joining \mathbf{t} and \mathbf{t}' , $d(\mathbf{t}, \mathbf{t}') = 0$. Otherwise, $d(\mathbf{t}, \mathbf{t}') = 1$. So, d , which can be obtained by image snapping tools or from disparity map, indicates and controls the influence between kernel sites. Finally, to generate a 3D surface mesh for visualization, we use the solved \mathbf{f} (Eqn. 1) to find the 3D point for every image location, and then generate the surface by applying [14], for instance.

5. Results

Surface from dense surface gradients. We first demonstrate the robustness of our system by adding different amount of Gaussian noise to the input. Figure 4 shows a *Torus*, a genus-one object, where the orientations of the normals are perturbed. The results show that our closed-form solution is quite robust and degrades gracefully with increasing amount of noise. Figure 5 shows several input dense normal maps obtained by photometric stereo (courtesy of [22]). *Rope* is a complex object with a lot of structures. For *Hair*, the normal map is very noisy because some hairs of the wig were always under shadows regardless the lighting direction, when the object was captured under variable illumination. *Face* is complex with typical facial features. Interestingly, the subject has a few pimples on the face, and wears eye glasses which introduce surface orientation discontinuities to the resulting dense normal map. Note the faithfulness of the surface results generated by our closed-form solution.

Surface from very sparse surface gradients. Figure 6 shows several views of the sparse normals of *Torsos*. For each view, the normals along each pair of closely-spaced curves are pointing away from each other. They are used to model crease curves, which incorporate surface details on the torsos. Note that even without any user-input height, our surface results are very good considering only sparse normals are available. We tested the *Male Torso* using [5, 9] and show the results in Figure 7. The resultant surfaces exhibit terrace-like structures. For [9], we tested the input with different input parameters. Better results with less distortion are produced, but the visual quality of the surfaces is still far from ours.

Surface from incomplete surface gradients and incomplete heights. Finally, we propose an application on disparity editing, by assigning sparse normals to the disparity map based on a monocular view. State-of-the-art stereo algorithms already produce very good disparity map for modeling the global structure of a real scene. However, local

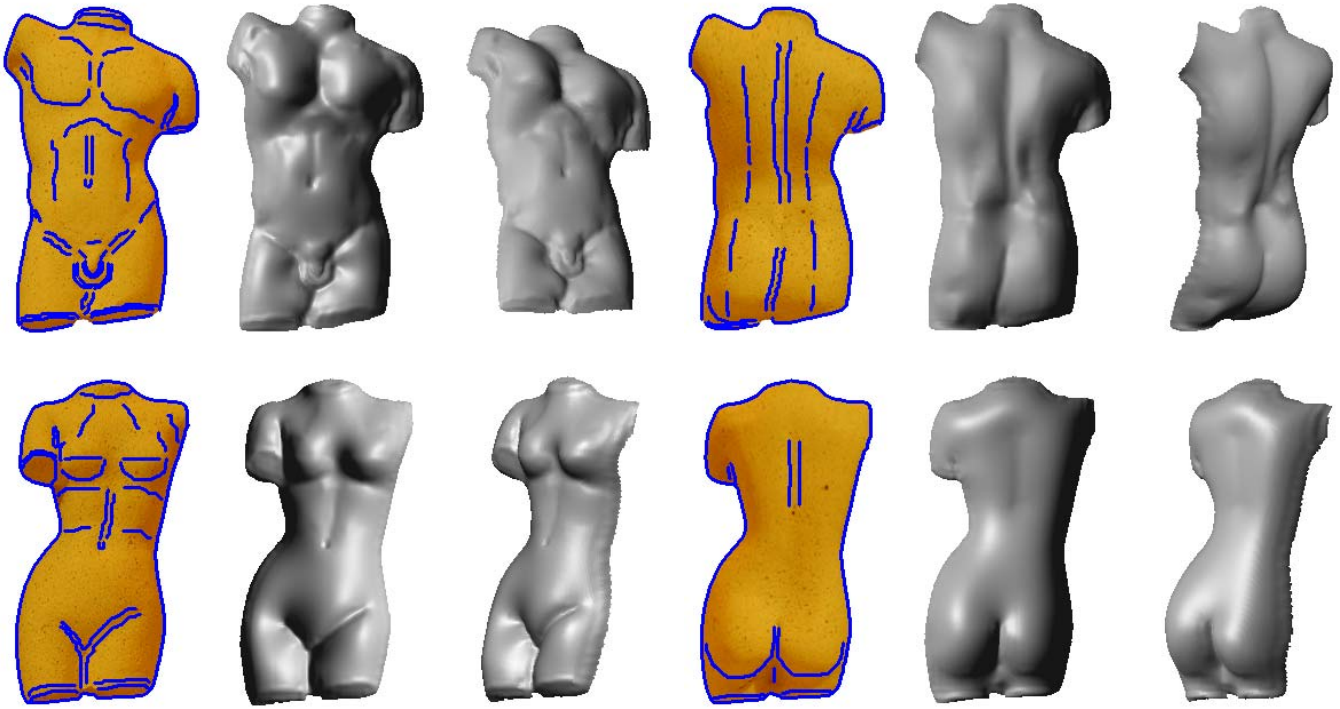


Figure 6. Results on very sparse gradients. From left to right: Input sparse normals, obtained from the sketched-based single-view modeling system [24], are available along the blue curves, which are overlaid onto the images of the statuettes. Two views of the reconstructed surface are shown. Top: *Male Torso* – front and back. Bottom: *Female Torso* – front and back.




Figure 7. Surfaces from sparse gradients generated by [5] and [9]. The missing normals are filled with $[0 \ 0 \ 1]^T$ as applied in [9]. In particular [9] reconstructs the surface by correlating multiple shapelet scales, while an over-smoothed surface is reconstructed using a large minimum σ . Increasing the number of scales used will produce an over-flattened surface.


structures are usually lost during stereo reconstruction, due to limited image resolution. For the *Tsukuba* dataset (Figure 8), for instance, while the global structure of the scene is excellent, the details on the bust and the curvature of the lamp-shade are both lost.

We propose to improve the disparity map by taking only the disparity along the object silhouettes (incomplete heights without normals), and associate in the interior of each chosen disparity region sparse normals (incomplete normals without heights), based on a given stereo image.


The principle behind the specification of local structures using normals can be traced back to *Pictorial Relief* by Koenderick et al. [8], which concludes that humans are very good at assigning local surface normals for specifying local shape. To specify the normals, we make use of intelligent scissor [12] to follow the salient curves on the image, followed by interactive normal transfer using a sketching interface [24]. Given the incomplete heights and incomplete normals, our closed-form solution generates a highly-detailed disparity surface as shown in Figure 8, where the bust suf-



input image



incomplete heights



incomplete normals



frontal view of disparity map, without texture



zoom-in oblique view of disparity map, with texture



Figure 8. Surface from incomplete heights and incomplete normals. The blue curves indicate the loci where normals are available. The left shows the original disparity map. The right shows the disparity surface generated using our closed-form solution.

fers significantly less distortion after it is rotated.

6. Conclusion

We have described a closed-form solution to surface-from-gradients from incomplete data on a single view. The proposed algorithm can operate on dense or sparse input. The input can consist of normals without corresponding 3D locations, or 3D locations without corresponding normals. We tested our implementation in a combination of different input scenarios. The good results we obtained demonstrate that our method is ready to be “plugged” into

single-view modeling, photometric and geometric stereo systems, where numerically-stable surface generation, capable of producing surfaces faithful to the input while being robustness to noise, is essential.

References

- [1] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In *ECCV (1)*, pages 578–591, 2006.
- [2] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987.
- [3] M. D. Buhmann and M. J. Ablowitz. *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003.
- [4] O. Drbohlav and M. Chantler. Can two specular pixels calibrate photometric stereo? In *ICCV05*, pages II: 1850–1857, 2005.
- [5] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. In *ICCV87*, pages 118–128, 1987.
- [6] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *ICCV05*, pages I: 341–348, 2005.
- [7] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *PAMI*, 27(8):1254–1264, August 2005.
- [8] J. Koenderink. Pictorial relief. *Royal*, 356(1740):1071–1086, 1998.
- [9] P. Kovsi. Shapelets correlated with surface normals produce surfaces. In *ICCV05*, pages 994–1001, 2005.
- [10] D. Kriegman. Passive photometric stereo from motion. In *ICCV05*, pages II: 1635–1642, 2005.
- [11] G. Medioni, M. S. Lee, and C. K. Tang. *A Computational Framework for Segmentation and Grouping*. Elsevier, 2000.
- [12] E. Mortensen and W. Barrett. Intelligent scissors for image composition. In *ACM SIGGRAPH 1995*, pages 191–198, 1995.
- [13] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH 2005)*, 24(3), Aug. 2005.
- [14] Y. Ohtake, A. Belyaev, and H.-P. Seidel. Ridge-valley lines on meshes via implicit surface fitting. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH 2004)*, 23(3):609–612, 2004.
- [15] M. Prasad, A. Zisserman, and A. Fitzgibbon. Single view reconstruction of curved surfaces. In *CVPR06*, pages II: 1345–1354, 2006.
- [16] B. Schoelkopf and A. Smola. *Learning with kernels*. MIT Press, 2002.
- [17] J. Sun, Y. Li, S. B. Kang, and H. Y. Shum. Symmetric stereo matching for occlusion handling. In *CVPR05*, pages II: 399–406, 2005.
- [18] J. Sun, N. N. Zheng, and H. Y. Shum. Stereo matching using belief propagation. *PAMI*, 25(7):787–800, July 2003.
- [19] C. K. Tang and G. Medioni. Inference of integrated surface, curve, and junction descriptions from sparse 3d data. *PAMI*, 20(11):1206–1223, November 1998.
- [20] D. Terzopoulos. The computation of visible-surface representations. *PAMI*, 10(4):417–438, July 1988.
- [21] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models for 3D object recognition. *IJCV*, 1(3):211–221, Oct 1987.
- [22] T. P. Wu and C. K. Tang. Dense photometric stereo by expectation maximization. In *ECCV06*, pages IV: 159–172, 2006.
- [23] T. P. Wu and C. K. Tang. Visible surface reconstruction from normals with discontinuity consideration. In *CVPR06*, pages II: 1793–1800, 2006.
- [24] T. P. Wu, C. K. Tang, M. Brown, and H. Y. Shum. Shapepalettes: Interactive normal transfer via sketching. In *ACM SIGGRAPH 2007*, August 2007.
- [25] L. Zhang, G. Dugas-Phocion, J. Samson, and S. Seitz. Single view modeling of free-form scenes. In *CVPR01*, pages I:990–997, 2001.