SURFACE HEIGHT RECOVERY FROM SURFACE NORMALS USING MANIFOLD EMBEDDING

Antonio Robles-Kelly and Edwin R. Hancock
Department of Computer Science,
University of York,York YO1 5DD, UK.
{arobkell,erh}@cs.york.ac.uk

Abstract

In this paper we show how surface height recovery from the field of surface normals can be posed as one of low dimensional embedding. To do this, we make use of the surface normals to compute the surface height increments corresponding to each location on the pixel lattice. With the height increments to hand, we can estimate the surface height difference between each pair of pixel-locations. We pose the problem of surface height recovery as that of embedding the surface normals on a manifold in a 3D space that preserves both the pattern of surface height differences and the lattice footprint of the field of surface normals. We present a sensitivity study on synthetic imagery and perform experiments on a variety of real world image data, where it produces qualitatively good reconstructed surfaces.

1. INTRODUCTION

The analysis of the literature on the topic of surface height recovery is not a straightforward task. The reason for this is that surface recovery is frequently viewed as an integral part of the shape-from shading or shape-from-texture process. Horn and Brooks [5] realise surface height recovery as a post-processing step. The process proceeds from the occluding boundary and involves incrementing the surface height by an amount determined by the distance traversed and the slope angle of the local tangent plane. In some of the earliest work, Wu and Li [11] average the surface normal directions to obtain a height estimate. A more elegant solution is proposed by Frankot and Chellappa [4] who project the surface normals into the Fourier domain to impose integrability constraints and surface height is recovered using an inverse Fourier transform. Klette and his co-workers have enhanced this approach by showing how more complex regularisation constraints can be formulated in the Fourier domain [9]. Dupuis and Oliensis [3] have developed a method which draws on differential geometry and involves propagation in the direction of the steepest gradient from singular points. A fast variant of this algorithm is described by Bichsel and Pentland [1] who compute the relative height of the surface with respect to the highest intensity point.

In this paper, our contribution is to pose the problem of recovering the surface height function as that of embedding the field of surface normals into a manifold. There is of course a considerable body of tangentially related work on this topic. Since manifold learning attempts to embed pattern vectors in a low dimensional space, it is potentially useful for a number of tasks. Proof of this are the results obtained making use of isometric embedding methods [8] and unidimensional scaling [2].

2. SURFACE HEIGHT RECOVERY

In this section, we pose the problem of recovering the surface height function from the field of surface normals as that of embedding the field of normals into a manifold that resides in a Euclidean space. To do this, we use the field of surface normals to estimate the height difference between each pair of locations in the field of normals. Once the height estimates have been computed, the surface may be then recovered by embedding them into the single-dimensional Euclidean space perpendicular to the lattice footprint.

2.1. Height Difference Approximation

Let d_{s_k,s_l} be the estimate of the surface height difference between the pair of points s_k and s_l on the surface under study S whose x and y coordinates correspond to the row and column indexes of the pixels p_k and p_l on the image plane. An estimate of d_{s_k,s_l} may be recovered by making use of the increments along the pixel-sites falling on the path Γ_{p_k,p_l} that best describes the projection of the geodesic connecting the points onto the image or pixel lattice. To compute the quantity d_{s_k,s_l}, we traverse the path Γ_{p_k,p_l} and compute the height increments associated with the pixel-site transitions along the path. The quantity d_{s_k,s_l} is then given by the sum of these height increments. The approximation to the height increment associated with the transition from the pixel indexed $p_i \in \Gamma_{p_k,p_l}$ to the pixel indexed $p_j \in \Gamma_{p_k,p_l}$ can be computed by assuming that the two pixel sites are connected by a plane whose slope is determined by the normal vectors $\vec{N}(s_i)$ and $\vec{N}(s_j)$. The height increment

0-7803-8554-3/04/$20.00 ©2004 IEEE.

is given by

$$h_{s_i,s_j} = \frac{r_{p_i,p_j}}{2}\left\{\frac{\vec{N}_x(s_i)}{\vec{N}_y(s_i)} + \frac{\vec{N}_x(s_j)}{\vec{N}_y(s_j)}\right\} \quad (1)$$

where $r_{p_i,p_j} = || p_i - p_j ||$ is the Euclidean distance between the site-centres on the lattice of surface normals.

2.2. Surface Integration Via Unidimensional Embedding

The problem of recovering the surface from a set of pairwise height difference estimates may be viewed as one of a unidimensional embedding which preserves the footprint of the surface-normal lattice. We commence by rewriting the height difference estimate d_{s_k,s_l} between a pair of points on the surface S as

$$d_{s_k,s_l} = z_{s_k} - z_{s_l} + e_{s_k,s_l} \quad (2)$$

where z_{s_k} is the surface height at a point s_k corresponding to the pixel-site indexed p_k and e_{s_k,s_l} is the error of representation in the height difference.

Here, we consider the errors of representation to be a set of independent random variables of size $N = | N_{rows} \times N_{cols} |$, where N_{rows} and N_{cols} are the lattice dimensions. Furthermore, the central limit theorem [6] states that, if $N \rightarrow \infty$, the distribution of normalised errors $\delta = e_{s_k,s_l}$ will be Gaussian with mean μ and variance σ, i.e.

$$P(\delta) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left\{-\frac{(\delta-\mu)^2}{2\sigma^2}\right\} \quad (3)$$

For N samples, the error in the mean μ is

$$\epsilon = \frac{\sigma}{\sqrt{N}} \quad (4)$$

As a result, as $N \rightarrow \infty$, the uncertainty of the estimator tends to zero.

With these ingredients, we define the column vector $\mathbf{B}$ whose k^{th} coefficient is given by

$$\begin{aligned}\mathbf{B}_k &= \frac{1}{N}\sum_{l=1}^{l=N} d_{s_k,s_l} \\ &= \frac{1}{N}\sum_{l=1}^{l=N}(z_{s_k} - z_{s_l} + e_{s_k,s_l}) \\ &= z_{s_k} - \frac{1}{N}\sum_{l=1}^{l=N} z_{s_l} + \frac{1}{N}\sum_{l=1}^{l=N} e_{s_k,s_l} \end{aligned} \quad (5)$$

To take our analysis further, we turn our attention to the last term of the equation above. The quantity

$$\bar{\mu}_k = \frac{1}{N}\sum_{l=1}^{l=N} e_{s_k,s_l} \quad (6)$$

is the mean height error at the location l on the surface. If we assume that the distribution of errors is stationary, then we can write

$$\bar{\mu}_i = \frac{1}{N}\sum_{l=1}^{l=N} e_{s_i,s_l} \approx \bar{\mu}_j = \frac{1}{N}\sum_{l=1}^{l=N} e_{s_j,s_l} \quad (7)$$

As a result, noting that the quantity $\frac{1}{N}\sum_{l=1}^{l=N} z_{s_l}$ in Equation 5 does not depend on the index k, it follows that

$$\begin{aligned}\mathbf{B}_i - \mathbf{B}_1 &= z_{s_i} + \frac{1}{N}\sum_{l=1}^{l=N} e_{s_i,s_l} - z_{s_1} - \frac{1}{N}\sum_{l=1}^{l=N} e_{s_1,s_l} \\ &= z_{s_i} - z_{s_1} + \bar{\mu}_i - \bar{\mu}_1 \approx z_{s_i} - z_{s_1} \end{aligned} \quad (8)$$

Finally, if we set $z_{s_1} = 0$ and solve for z_{s_i}. The result is

$$z_{s_i} \approx \mathbf{B}_i - \mathbf{B}_1 = \frac{1}{N}\left(\sum_{l=1}^{l=N} d_{s_i,s_l} - \sum_{l=1}^{l=N} d_{s_1,s_l}\right) \quad (9)$$

It is important to stress that, whereas setting $z_{s_1} = 0$ may alter the relative position of the embedding with respect to the coordinate system, the configuration described by the relative height between each pair of points remains unchanged.

3. EXPERIMENTS

We have experimented with our new method on a variety of real world and synthetic data. To this end, we have used gray-scale imagery and recovered the field of surface normals making use of the method of Worthington and Hancock [10]. We have also compared our results with those delivered by the algorithm of Bichsel and Pentland [1].

3.1. Synthetic Data

Here, we investigate the effect of varying the light source direction. To this end, we have generated a set of 6 images of a teapot illuminated with a single light source positioned in the direction $\vec{L}_D = [\sin(\theta_{L_D}), 0, \cos(\theta_{L_D})]^T$. For our experiments, we have set $\vec{V} = [0, 0, 1]^T$ and varied the angle θ_{L_D} in increments of 6^o between 0^o and 30^o.

In the top row of Figure 1, we show three of the images in our dataset. Here, we have ordered the example images, from left-to-right, in increasing θ_{L_D}, i.e. $\theta_{L_D} = 0^o, 18^o, 30^o$. The second and third rows in the figure show, again, the recovered height and, the difference between the recovered and the ground-truth heights. The results for the Bichsel and Pentland's method are shown in the two bottom rows. The errors are largest at highly inclined locations on the surface. Whereas the effect of varying the light source direction is to magnify the errors for both methods, our algorithm is much more robust and delivers much better results. Figure 2 provides a quantitative investigation of this

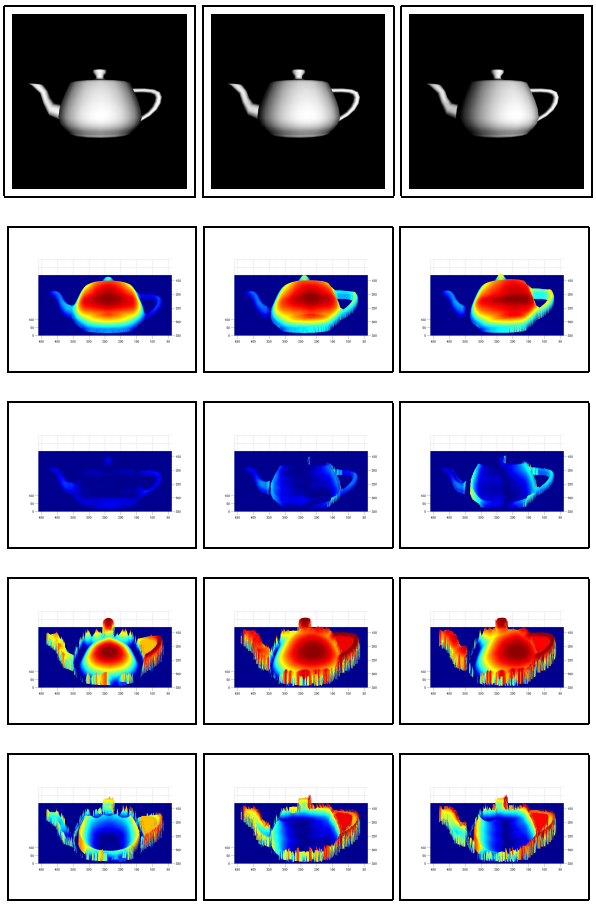

Fig. 1. Surface height recovery results. Top row: input images; Second and third rows: recovered height and mean-squared error for our method; Fourth and fifth rows: recovered height and mean-squared error for the method of Bichsel and Pentland

effect. Here we plot the mean squared height error for both methods as a function of the angle between viewer and light

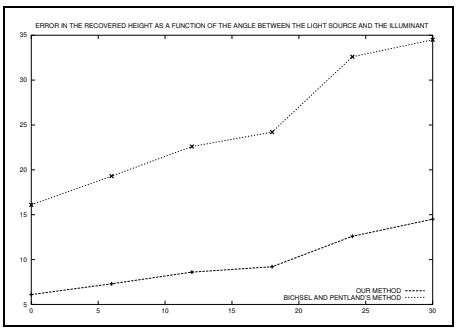

Fig. 2. Mean-squared error as a function of the angle difference between the illuminant and the viewer direction

source directions. The error for the Bichsel and Pentland's algorithm is larger.

3.2. Real World Data

We have experimented with the height recovery method on a variety of real world imagery of statuary. These are a detail of the Michelangelo's "Moses", a fragment of the "Three Graces" relief, an image of Barthel Melchiorre's "Malinconia", a marble table-centre, a porcelain vase and an image of a toy duck from the COIL database. In Figure 3, the images in the top row show the input images. In the second and third rows we show the result of re-illuminating the objects from different directions. For visualisation purposes, we have rendered the recovered surfaces making use of a Phong BRDF [7] with a shininess factor of 10 and illuminated them from the directions $[\sin(20^o), 0, \cos(20^o)]^T$ and $[\sin(-20^o), 0, \cos(20^o)]^T$. The two views show the surface rotated 20^o degrees about the y-axis. The two bottom rows show the views for the surfaces recovered by the method of Bichsel and Pentland. In contrast with the surfaces yielded by Bichsel and Pentland's algorithm, the objects recovered making use of our method all have quite complex surface detail, with a mix of concave and convex structure. Furthermore, the re-illuminations capture the main surface features.

4. CONCLUSIONS

In this paper, we have presented a novel approach for recovering the 3D height data from the field of surface normals. We posed this process as that of a lower dimensional embedding. To do this, we computed height increments making use of the field of surface normals and used these to estimate the surface height difference between each pair of locations on the lattice footprint. With the increments at hand, we recovered the depth information by embedding the surface normals on a manifold in a 3D space. Our experiments on synthetic and real-world image data produced good reconstructed surfaces.

5. REFERENCES

[1] M. Bichsel and A. P. Pentland. A simple algorithm for shape from shading. In *Proc. of the IEEE Conf. on Comp. Vision and Pattern Recognition*, pages 459–465, 1992.

[2] I. Borg and P. Groenen. *Modern Multidimensional Scaling, Theory and Applications*. Springer Series in Statistics. Springer, 1997.

[3] P. Dupuis and J. Oliensis. Direct method for reconstructing shape from shading. In *Proc. of the IEEE Conf. on Comp. Vision and Pattern Recognition*, pages 453–458, 1992.

[4] R. T. Frankot and R. Chellappa. A method of enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4(10):439–451, 1988.

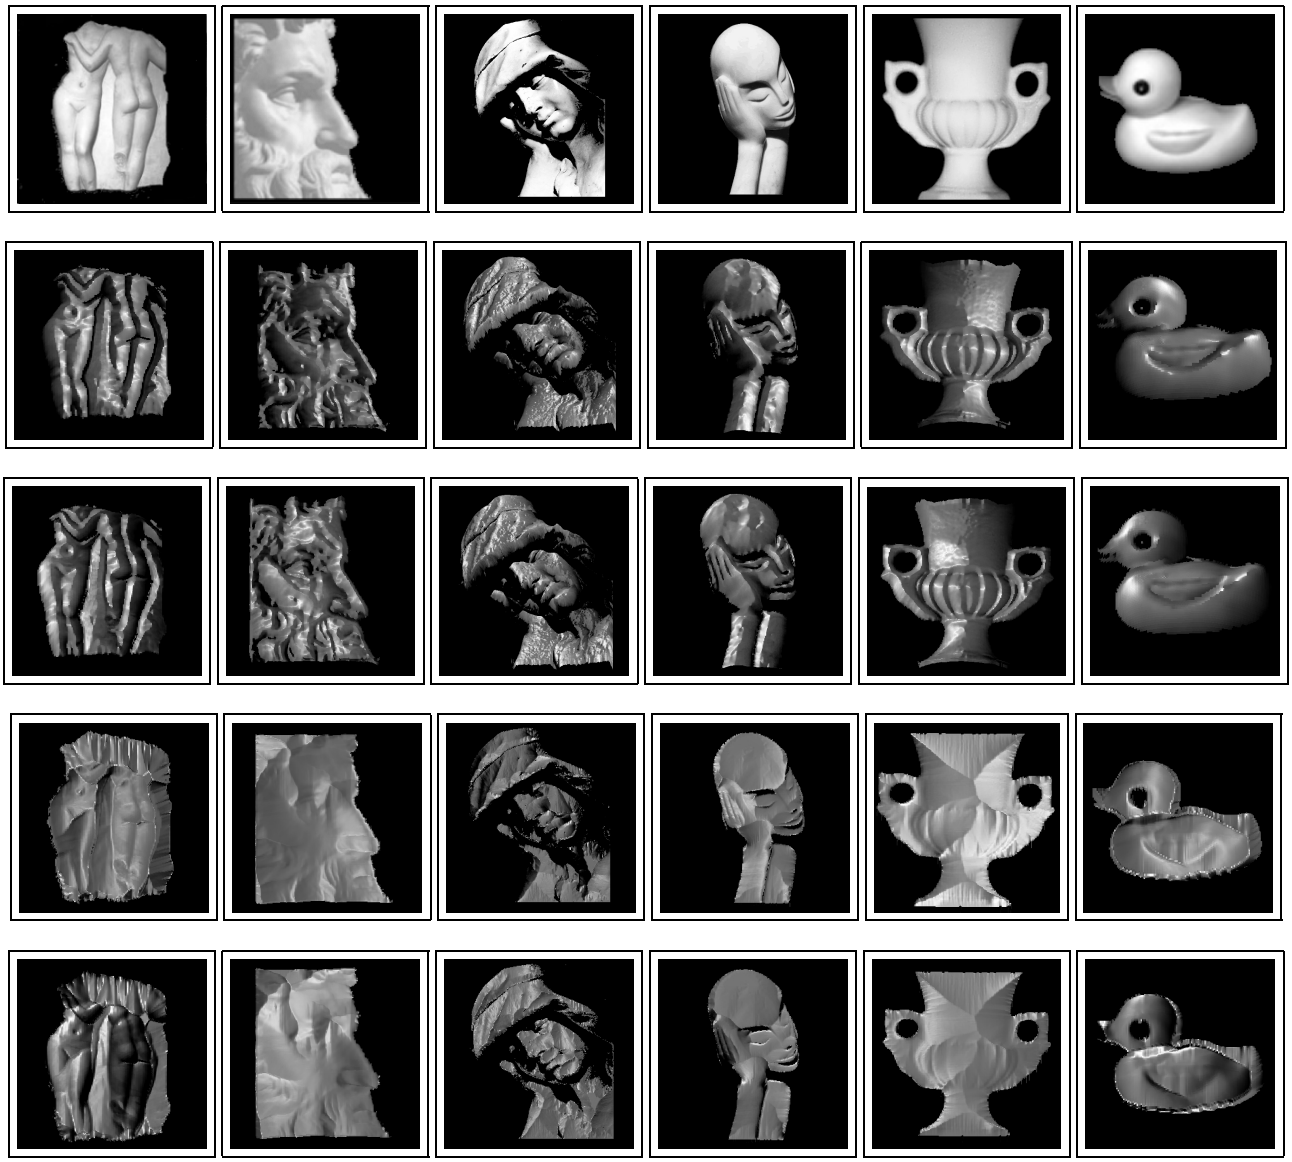

Fig. 3. Surface height recovery results on real-world data. Top row: input images; Second and third rows: two re-illuminated views of the surface recovered making use of our method; Fourth and fifth rows: two re-illuminated views of the surface recovered making use of Bichsel and Pentland's method

[5] B. K. P. Horn. Height and gradient from shading. *International Journal of Computer Vision*, 5(1):37–75, 1990.

[6] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. Mc Graw Hill, 1984.

[7] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.

[8] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[9] T. Wei and R. Klette. Theoretical analysis of finite difference algorithms for linear shape from shading. In *CAIP*, pages 638–645, 2001.

[10] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(12):1250–1267, 1999.

[11] Z. Wu and L. Li. A line-integration based method for depth recovery from surface normals. *CVGIP*, 43(1):53–66, July 1988.