

Vector Field Smoothing Via Heat Flow

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 develop a new method for recovering and smoothing fields of surface normals in shape-from-shading. We show how transform the problem of recovering surface normals that satisfy Lambert's law into that of solving the steady state heat equation for a scalar potential. According to this picture, the smoothed field of surface normals is found by taking the gradient of the scalar field. The heat equation for the scalar field can be solved using simple finite difference methods, and leads to an iterative procedure for vector estimation. We illustrate the utility of the new method on real world imagery and compare our results with those delivered by an alternative.

1. Introduction

The recovery of a consistent field of surface normals is one that arises in a number of areas in computer vision. Examples include shape-from-shading, shape-from-texture, optical flow and stereopsis. One of the problems that is of pivotal importance in the processing of vector fields recovered from noisy image data is that of smoothing, and practical smoothing or regularisation methods have been reported in the contexts of optical flow computation [4], curvature-based surface shape representation [6] and the smoothing of stereo disparity fields [3].

An important issue underpinning the smoothing process is that of preserving the differential structure of the underlying vector-field, while satisfying constraints on data-closeness. In fact, the data-closeness constraints can sometimes be re-cast as constraints on the individual components of the vectors. For instance, in shape-from-shading, when Lambertian reflectance is assumed, the component of the surface normal in the light source direction is constrained to be equal to the inverse cosine of the normalised image brightness [7]. For the purposes of optical flow computation, brightness constraints are frequently used [1].

The analysis of the literature on the topic of vector field smoothing is not a straightforward task. The reason for this is that smoothing is frequently viewed as an integral part of the relevant vision module, rather than one of generic utility. However, in shape-from-shading simple regularisation methods have been shown to oversmooth the field of surface normals, and this problem can be overcome using either statistically motivated robust regularisation methods or gradient consistency constraints [8]. In optical flow estimation, existing algorithms have been greatly influenced by the anisotropic diffusion approach to image regularisation developed by Perona and Malik [5]. There have also been a number of attempts to extend the variational optical flow estimation algorithm of Horn and Schunk [2].

In this paper, we focus in detail on the problem of shape-from-shading. We pose the problem of smoothing a vector field as that of solving the heat equation subject to the constraints imposed by Lambert's law. We recast the problem of estimating the field of surface normals as that of recovering a scalar potential field. The z-component of the field of surface normals is constrained by Lambert's law to be equal to the inverse cosine of the normalised image brightness. The remaining components of the vector are given by the gradient of the scalar potential. A simple variational analysis using the divergence theorem shows that the time derivative of the scalar potential is given by the average of the image intensity. As a result, a simple finite element method can be used to compute the components of the field of vectors.

2. Heat Flow

The problem studied in shape-from-shading is that of estimating the direction of the surface normal vector $\vec{N}(s_k) = [N_x(s_k), N_y(s_k), N_z(s_k)]^T$ at a point s_k on a surface $S \in \mathbb{R}^3$ from a measurement of the brightness $I(p_k)$ at the corresponding image point p_k . The surface normal satisfies Lambert's law, i.e. $I(p_k) = \vec{L} \cdot N_z(s_k)$, where \vec{L} is light source direction. Unfortunately, Lambert's law allows only the recovery of the zenith angle between the light source




Figure 1. Ground truth for our synthetic data

vector and the surface normal. The recovery of the azimuth angle remains undetermined, and requires the use of smoothness constraints. Here we use a heat-flow analogy to perform the smoothing process.

The analogy is as follows. We assume that there is a time-dependant planar scalar field $Q(\vec{x}, t)$, where \vec{x} are the x, y and z-coordinates of the point \hat{p}_k and t is the time epoch. As a result of the constraint imposed by the scalar field, the z-component of the vector $\vec{N}(s_k)$ is equal to the measured intensity $I(p_k)$ at the point p_k . The component on the vector $\vec{N}(s_k)$ on the plane is found by taking the gradient of the scalar field $Q(\vec{x}, t)$. As a result, we can write $\vec{N}(s_k) = (\frac{\partial Q(\vec{x}, t)}{\partial x}, \frac{\partial Q(\vec{x}, t)}{\partial y}, I(p_k))^T$. Our aim is to find the scalar field $Q(\hat{p}_k, t)$ that is the steady state solution of the heat equation

$$\nabla^2 Q(\vec{x}, t) = \frac{\partial^2 Q(\vec{x}, t)}{\partial x^2} + \frac{\partial^2 Q(\vec{x}, t)}{\partial y^2} = \frac{\partial Q(\vec{x}, t)}{\partial t} \quad (1)$$

subject to the constraint $N_z(s_k) = I(p_k)$. This involves the analysis of the heat flow on the plane $\hat{\Pi}$. Once the steady state solution to the heat equation is found, then the vector may be computed by taking its gradient. However, here we are only interested in the azimuth angle of the vector $\vec{N}(s_k)$, and this can be computed by taking the arc-tangent of the ratio of the first-differences of $Q(\vec{x}, t)$ in the x and y directions.

To analyse the behaviour of this system, we turn to the divergence theorem. Consider an area A on the plane $\hat{\Pi}$ which is enclosed by a volumetric region G . The flux associated with the field of vectors is given by

$$\int \int_A \vec{N}(s_k) \cdot \vec{L} dA = \int \int \int_G \nabla \cdot \vec{N}(s_k) dG \quad (2)$$

Since, $\vec{N}(s_k)$ is the divergence of $Q(\vec{x}, t)$, we have that

$$\int \int_A \vec{N}(s_k) \cdot \vec{L} dA = \int \int \int_G \left\{ \frac{\partial^2 Q(\vec{x}, t)}{\partial x^2} + \frac{\partial^2 Q(\vec{x}, t)}{\partial y^2} + \frac{\partial I(p_k)}{\partial z} \right\} dG \quad (3)$$

Noting that $\frac{\partial I(p_k)}{\partial z} = 0$, i.e. the z-component of the vector $\vec{N}(s_k)$ is constrained to be $I(p_k)$, and substituting for the

heat equation, we find that

$$\int \int_A \vec{N}(s_k) \cdot \vec{L} dA = \int \int \int_G \frac{\partial Q(\vec{x}, t)}{\partial t} dG \quad (4)$$

Finally, we use the dot product and note that $I(p_k) = \vec{N}(s_k) \cdot \vec{L} = N_z(s_k)$, and hence the divergence equation

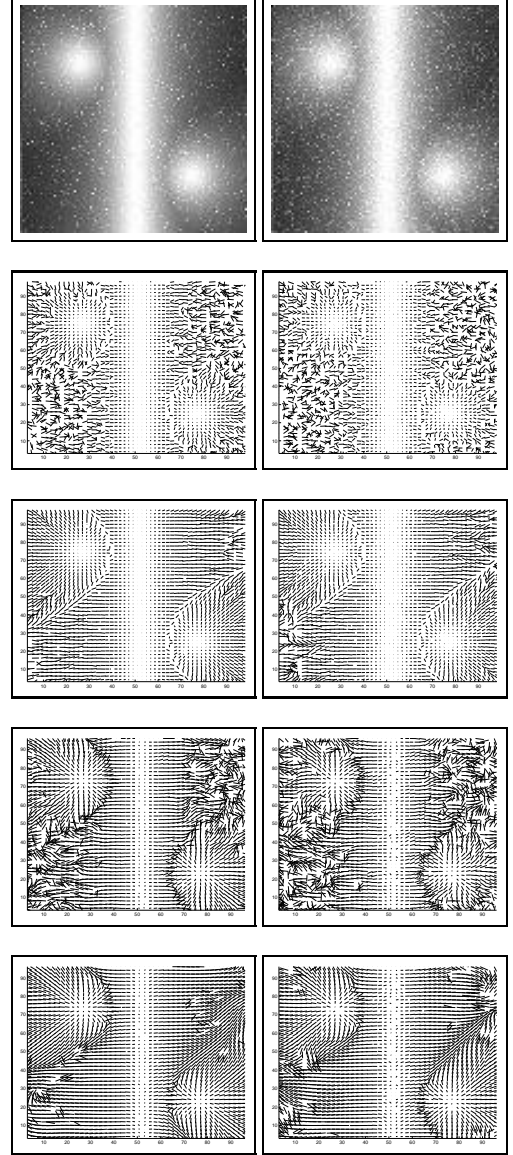


Figure 2. Smoothing results. Top row: input images; Second and third rows: smoothed field of surface normals after one and eight iterations of our algorithm; Bottom rows: field of surface normals delivered by the algorithm of Worthington and Hancock after 10 and 100 iterations.

becomes

$$\int \int_A I(p_k) dA = \int \int \int_G \frac{\partial Q(\vec{x}, t)}{\partial t} dG \quad (5)$$

In other words, the average rate of change of $Q(\vec{x}, t)$ with time is equal to the total potential enclosed within the surface patch A .

We can solve this equation using a simple finite element approach. We commence by setting $Q(\hat{p}_k, t = 0) = 1$. At time step n , the update equation is

$$Q(\hat{p}_k, n) \approx \frac{1}{|\Omega_{\hat{p}_k}|} \sum_{p_l \in \Omega_{\hat{p}_k}} (1 - I(p_k)) \hat{Q}(p_l, n - 1) \quad (6)$$

where $\Omega_{\hat{p}_k}$ is the set of points adjacent to \hat{p}_k . The quantity \hat{Q} is the averaged and normalised scalar potential, computed using the formula

$$\hat{Q}(\hat{p}_k, n) = \frac{\frac{1}{|\Omega_{\hat{p}_k}|} \sum_{p_l \in \Omega_{\hat{p}_k}} Q(p_l, n)}{\sum_{k \in \hat{\Pi}} \frac{1}{|\Omega_{\hat{p}_k}|} \sum_{p_l \in \Omega_{\hat{p}_k}} Q(p_l, n)} \quad (7)$$

This process is repeated until the scalar field $Q(\hat{p}_k, n)$ stabilises. On convergence, the azimuth angle of the vector $\vec{N}(s_k)$ at the location indexed k is given by

$$|\alpha_{s_k}| = \left| \arctan \left(\frac{\Delta Q(\hat{p}_k, n)_x}{\Delta Q(\hat{p}_k, n)_y} \right) \right| \quad (8)$$

In the above equation $\Delta Q(\hat{p}_k, n)_x$ and $\Delta Q(\hat{p}_k, n)_y$ are the first difference approximations to the x and y components of the gradient of Q , and are given by

$$\begin{aligned} \Delta Q(\hat{p}_k, n)_x &= Q(u - 1, v, w, n) - Q(u + 1, v, w, n) \\ \Delta Q(\hat{p}_k, n)_y &= Q(u, v - 1, w, n) - Q(u, v + 1, w, n) \end{aligned} \quad (9)$$

Since the computation of $|\alpha_{s_k}|$ constrains the azimuth of the vector to the first quadrant, we make use of the sign of the flow on the x and y -directions to determine the sign of the vector components on the plane $\hat{\Pi}$. Thus, the rotated vector $\vec{N}(s_k)$ is computed making use of the following rule

$$\begin{aligned} \vec{N}(s_k) &= \begin{bmatrix} \text{sign}(\Delta Q(\hat{p}_k, n)_x) \cos(|\alpha_{s_k}|) \sin(\theta_{s_k}) \\ \text{sign}(\Delta Q(\hat{p}_k, n)_y) \sin(|\alpha_{s_k}|) \sin(\theta_{s_k}) \\ \cos(\theta_{s_k}) \end{bmatrix} \\ &= \begin{bmatrix} N_x(s_k) \\ N_y(s_k) \\ N_z(s_k) \end{bmatrix} \end{aligned} \quad (10)$$

where $\theta_{s_k} = \arccos(I(p_k))$.

3. Experiments

In this section, we illustrate the utility of our new method for purposes of smoothing fields of surface normals estimated from shading images which exhibit large amounts of

noise and corruption. To this end, we have added controlled levels of Gaussian noise to synthetic and real-world imagery of objects that exhibit Lambertian reflectance. In all our experiments, we follow Worthington and Hancock [7] and use the gray-scale gradient as an initial estimate of the surface normal field directions

In Figure 1b, we show a parabolic dome with two elliptic domes superimposed. We have used this object to investigate the effect of adding noise to the raw image data. In the top row of Figure 2, we show two images obtained by rendering the surface in Figure 1b with a Lambertian reflectance model and adding controlled levels of Gaussian noise with $\sigma = 0.3, 0.5$ to the resultant image. The second and third rows show the field of surface normals at the first and last iterations of our smoothing algorithm. For both images, the algorithm converged in eight iterations. In the fourth and fifth rows, we show the fields of normals delivered by the algorithm of Worthington and Hancock [7] after 10 and 100 iterations. In contrast with the algorithm of Worthington and Hancock, our method preserves the shape-edges while being able to cope with large amounts of added noise.

Next, we present results on real-world imagery of statuary. These are a detail of the Michelangelo's "Moses" and a fragment of the "Three Graces" relief. In the left-hand column of Figure 3, we show the noise-free images for both objects. In the second column, we show the resultant images after having added Gaussian noise whose variance is of 0.2. The third and fourth columns show the field of surface normals at the first and last iterations of the algorithm. In the case of the Michelangelo's "Moses", the field of surface normals stabilised at the sixth iteration. For the "Three Graces", the algorithm took five iterations to converge. In Figure 4, we show the field of surface normals smoothed making use of the Worthington and Hancock's algorithm for the input images in Figure 3. The top row shows the fields of surface normals after 10 iterations. The fields of normals after 100 iterations are shown in the bottom row. From Figure 3, we can note that the fields of surface normals smoothed making use of our method all have quite complex detail, with a mix of concave and convex structure. Furthermore, the topology of the surface has been preserved, specially near the edges. Hence, our results lack the oversmoothing exhibited by the fields of normals delivered by the algorithm of Worthington and Hancock [7].

4. Conclusions

In this paper, we have shown how a vectorial field can be smoothed by making use of a heat flow process. This approach leads to a parameter-free model which constraints the field of vectors to satisfy a scalar potential field as a hard constraint. Furthermore, the differential equations that char-

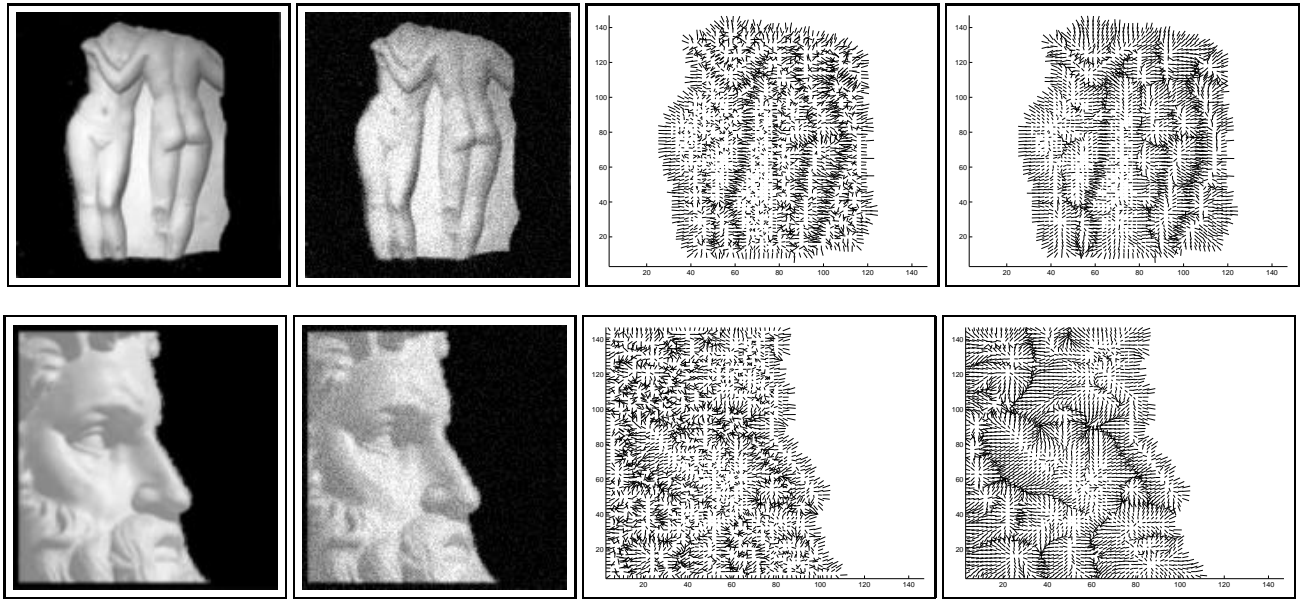


Figure 3. Smoothing results for real-world imagery. Left-hand column: noise-free images; Second column: images with added Gaussian noise; Third and fourth columns: smoothed field of surface normals at the first and last iterations of our algorithm.

acterise the heat flow process may be solved making use of a computationally efficient finite element method. We have performed experiments on synthetic and real-world imagery which suggest the algorithm can remove large amounts of added noise.

References

- [1] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *Int. Journal of Computer Vision*, 12(1):43–77, 1994.
- [2] B. K. P. Horn and B. Schunk. Determining optical flow. *Artificial Intelligence*, 17:185–204, 1981.
- [3] D. Marr and T. Poggio. A computational theory of human stereo vision. In *Proceedings of the Royal Society of London*, volume 204 of *Series B, Biological Sciences*, pages 301–328, 1979.
- [4] H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
- [5] P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [6] D. Terzopoulos. Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics and Image Understanding*, 24:52–96, 1983.
- [7] 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.

- [8] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):680–702, 1991.

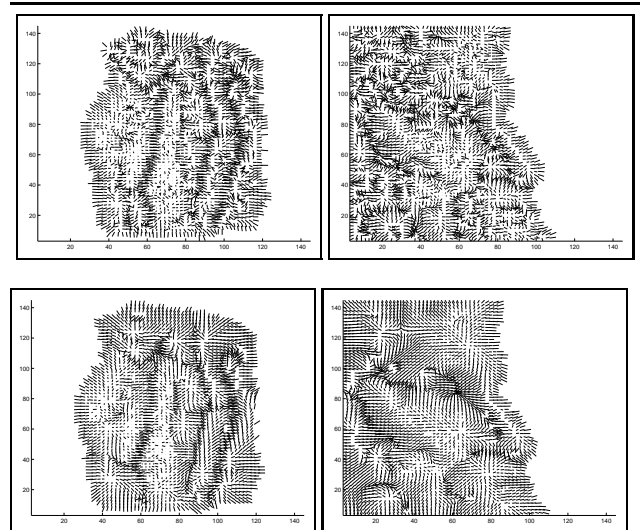


Figure 4. Field of surface normals delivered by the algorithm of Worthington and Hancock after 10 and 100 iterations for the input images in Fig. 3.