

Diffusion of Geometric Affinity for Surface Integration

Roberto Fraile and Edwin Hancock

Department of Computer Science
University of York, YO10 5DD UK

Abstract. A combinatorial method is used to reconstruct a surface by integrating a field of surface normals. An affinity function is defined over pairs of adjacent locations. This function is based on the surface's principal curvature directions, which are intrinsic and can be estimated from the surface normals. The values of this locally supported function are propagated over the field of surface normals using a diffusion process. The surface normals are then regularised, by computing the weighted sum of the affinity evolved over time. Finally, the surface is reconstructed by integrating along integration paths that maximise the total affinity. Preliminary experimental results are shown for different degrees of evolution under the presence of noise.

1 Introduction

Directional information about surfaces, in the form of surface normals or gradients, is involved in several computer vision problems such as Shape-from-Shading and Photogrammetric Stereo, or, more recently, diffusion tensor magnetic resonance (DT-MRI). Integration of a field of surface normals can be exact if the vector field is integrable, that is, if the measured curl is zero. Since this is not the case for most applications, due to measurement errors, it is necessary to develop methods to estimate the most likely surface from which the surface normals have been obtained. Figure 1 illustrates surface integration over facial data.

Most surface integration methods use a variational approach [1,2,3,4], which consist of defining a suitable functional

$$J(S) = \int \int E(S, \nabla S, \mathbf{n})$$

where S is the surface to be estimated, and \mathbf{n} is the surface normal information provided. Frankot-Chellappa [1] project the gradient field on integrable Fourier basis functions, and variations of this method use other families of integrable basis functions [4].

Alternatively, considering a discrete field of surface normals as a labelled grid graph, surface height estimates can be obtained by integrating along paths of the graph. Local information can be used to construct space-filling integration paths, for example, an affinity function can be defined for graph edges corresponding to affinity between the vertexes' surface normals.

The path integration approach is optimal under the assumption that at least there is one path joining every pair of surface locations, over which the height increments can be estimated. Graph-spectral methods are used to find those paths within the set of all possible paths.

Robles-Kelly [5] and Klette [6] describe path integration methods. In [5] a path is constructed using a graph-spectral analysis over the affinity matrix, defined using estimates of the sectional curvature. Several non-overlapping paths are used to cover the entire graph.




Fig. 1. An example of surface integration from directional data

The algorithm presented in this paper uses an intrinsic geometric property as affinity. This affinity function is used for regularisation of the surface normals, and for making a choice of integration paths. The affinity function is based on the surface’s principal curvature directions, which can be estimated from the surface normals and are representation invariant. First, the affinity is propagated according to a diffusion process; second, the surface normals are modified according to those local reliability estimates; and third, the surface is integrated, from the modified surface normals. The paths overlap, thus avoiding the need for segmentation.

The values of this locally supported function are propagated over the field of surface normals following a diffusion procedure described in [7,8]. A discrete Laplacian operator is used to construct the heat kernel, which can then be evaluated at different times.

The surface normals are then regularised, by computing the weighted sum of the affinity evolved over time. Integration paths that span the entire surface and minimise the total affinity are computed using a Minimum Spanning Tree algorithm. The surface is reconstructed applying the trapezium rule piecewise along each integration path. In this approach the surface S is modelled by a spanning tree T , and the following cost function is minimised:

$$J(T) = \sum_{(i,j) \in T} \frac{1}{A_{ij}}$$

where A_{ij} is the affinity of edge (i, j) .

We now proceed by describing a combinatorial method for surface integration given a local affinity function. In the following sections, the geometric affinity

function is motivated, and the affinity is propagated to non-neighbouring surface locations. The field of surface normals is updated according to this extended affinity, and the combinatorial surface integration method is applied to the modified surface normals.

2 Surface Integration

In this section a combinatorial method to integrate a surface from a field of surface normals is described. This method only requires a definition of affinity between adjacent locations, and assumes that, for every two locations, the path that minimises the affinity function also minimises the integration error.

Given two adjacent locations i and j , and their surface normals \mathbf{n}_i and \mathbf{n}_j , it is possible to estimate the height difference by applying the trapezium rule. An affinity function assigns a number between 0 and 1 to each pair of adjacent surface normals i and j , estimating the reliability of integration along the edge joining i and j . Therefore an affinity matrix is defined $A = \{a_{ij}\}_{i,j \in V}$.

Consider the graph $G = (E, V)$ whose vertexes V are the surface normal locations, and whose edges E are all the pairs of adjacent locations (typically the 4-neighbours). Let W be the array whose entries are inverse of the entries in A . In this way we have assigned weights to the edges of the graph.

A first requirement for a path integration method is that every two locations are connected by a path, so that the height can be estimated for all locations. A second requirement is that there is only one path between every two points, so that surface height estimates are unique. Therefore the paths will form a spanning tree T over G . A spanning tree that minimises the total weights given by W , or, equivalently, maximises the total affinity, can be obtained by using a Minimum Spanning Tree algorithm [9].

The surface can then be reconstructed by integration along the edges of T , applying the trapezium rule to the surface normals. This optimisation procedure is independent of the affinity function, and the problem of obtaining adequate height estimates is reduced to defining an appropriate affinity function. Figure 2 illustrates the notion of integration tree.

For this path-based integration method, it requires that minimising the affinity function is equivalent to finding a path over which integration is valid. Note




Fig. 2. Paths defined over the set of locations. Each location is labeled with a 3D surface normal vector. The surface is integrated over a path using the trapezium rule.

that this method does not require the field of surface normals to have zero curl, which would imply that *every* path joining every two locations provides correct height estimates.

The practical use of such a path-based integration method is limited by the extent to which the affinity function can capture global structure of the graph, and also by the quality of the surface normals. The field of surface normals contains redundant information because the trapezium rule only uses one of the two components of the gradient. It seems therefore reasonable to integrate over a modified field of surface normals in which the redundant information has been propagated along the graph G . The proposed way of doing this involves a local geometric affinity function, and a diffusion process to take into account the global structure of the graph G .

3 Geometric Affinity: Principal Curvature Directions

The affinity function needs to be defined for each pair of adjacent locations. In order to assign greater affinity to lower risk of integration errors, it should be monotonic with any distance defined over the surface normals as Euclidean vectors. It is also desirable that the affinity function does not depend on the way the surface normals have been sampled from the surface.

The field of surface normals is sampled from a surface, it is therefore possible to estimate intrinsic properties of the surface from the surface normals. One such intrinsic property is the principal curvature and its directions. A measure of geometric affinity between two surface locations is given by how close the direction linking them is to the minor principal curvature direction. (In regions where the principal curvature directions are not defined, the affinity can be considered the same in all directions). Figure 3 illustrates the principal curvature directions for a simple surface.

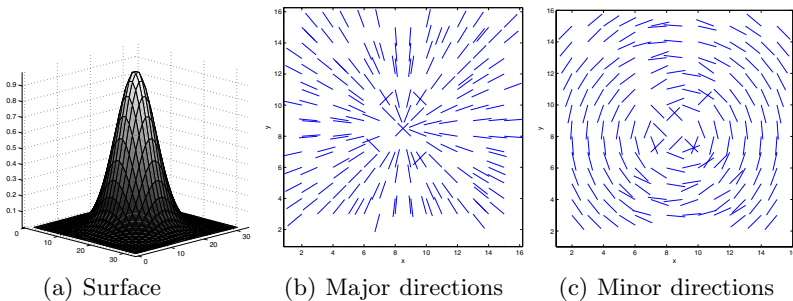


Fig. 3. Principal curvature directions estimated from a field of surface normals

To estimate the principal curvature directions from the surface normals, consider S , the surface function. The Hessian matrix, which can be used to calculate the principal curvature directions, is constructed using the second derivatives at each point.

The Hessian matrix at each location,

$$H = \begin{pmatrix} \partial_{xx} & \partial_{xy} \\ \partial_{yx} & \partial_{yy} \end{pmatrix} S$$

can be obtained by estimating the second partial derivatives of S at each point using finite differences over the first partial derivatives.

The first derivatives $\partial_x S$ and $\partial_y S$ are obtained from the normal vector (n_x, n_y, n_z)

$$\begin{pmatrix} n_x \\ n_y \\ n_z \end{pmatrix} = -n_z \begin{pmatrix} \partial_x S \\ \partial_y S \\ -1 \end{pmatrix}$$

The partial derivatives $\partial_x S$ and $\partial_y S$ are defined whenever n_z is nonzero. This is the case, for example, when data is available in shape-from-shading.

The second derivatives can be approximated using a finite difference operator. Let \mathbf{v}_i be the coordinates of location i , and \mathbf{e} a unit 2D vector.

$$\partial_{\mathbf{e}} f(\mathbf{v}_i) \simeq \Delta_{\mathbf{e}} f(\mathbf{v}_i) = \frac{f(\mathbf{v}_i + \mathbf{e}) - f(\mathbf{v}_i - \mathbf{e})}{2}$$

The eigensystem of H consists of the principal curvatures and their directions. For most surfaces the eigenvectors \mathbf{h}_{\min} and \mathbf{h}_{\max} of H are not linearly dependent, and form a basis of the 2D space.

The edge joining two adjacent locations l_1 and l_2 , corresponds to a direction vector \mathbf{e} , which can be represented in the eigensystem of H :

$$\mathbf{e} = \alpha \mathbf{h}_{\min} + \beta \mathbf{h}_{\max}$$

Therefore an geometric measure of affinity α to each pair of adjacent locations l_1 and l_2 . Greater geometric affinity corresponds to directions closer to the minor principal curvature direction.

We have obtained an affinity function that can be used to calculate an affinity matrix A , which is symmetric with size $|V| \times |V|$. Without loss of generality, we can assume that A is normalised so that each row adds up to 1. This measure of affinity is only defined locally, we now proceed to propagate affinity across the field of surface normals.

4 Propagating Affinity Using the Heat Kernel

The affinity matrix A embodies local information, in the form of geometric affinity between adjacent locations. The matrix A can be considered as a transition probability between locations in the graph of locations. A path joining two locations i and j is a subset of E and therefore has an probability induced by A .

We would like to calculate the affinity between every two locations, not only those which are neighbouring in the graph. To do so, let us consider a random

walk over G whose transition probability is given by A , and calculate the probability of the random walk joining locations i and j in t steps. This corresponds to a diffusion process in which the conductivity is given by A [7,8].

Let \hat{L} be the normalised Laplacian associated to A , $\hat{L} = D^{-\frac{1}{2}}LD^{\frac{1}{2}}$, where $L = D - A$ and D is the diagonal matrix such that $D_{ii} = \sum_j A_{ij}$. The matrix \hat{L} can be seen as a discrete approximation of the continuous Laplacian operator in the following diffusion equation [7]

$$\partial_t H = -\hat{L}H \quad (1)$$

The solution $H(t)$ can be calculated by matrix exponentiation

$$H(t) = e^{-t\hat{L}}$$

Each entry (i, j) of the matrix $H(t)$ can be interpreted as the probability of a random walk joining locations i and j , after t steps, given the transition probability A . Note that the size of both A and $H(t)$ is $|V| \times |V|$.

The presence of a path in G of high affinity between two nodes i and j increases the probability $H(t)_{ij}$. The matrix $H(t)$ can now be used to modify the surface normals.

5 Updating the Surface Normals

While the field of surface normals contains redundant information, integration along paths, which effectively are subsets of E , discards all surface normal information not in the direction of the edges that form the path. In order to make use of redundant information before the integration step, let us modify the surface normals using the transition probability $H(t)$ obtained in the previous section. This is a generalisation of a simple average of neighbouring surface normals, and is similar to subjecting the surface normals to a process of anisotropic diffusion [10].

The updated normals corresponding to a random walk of length t , $N(t) = \{\mathbf{n}(t)_i\}_{i \in V}$, are defined as a weighted sum of the surface normals:

$$\mathbf{n}(t)_i = \sum_j H(t)_{ij} \mathbf{n}_j \quad (2)$$

where the indexes i and j visit all locations.

As a result, the field of surface normals $N(t)$ is the weighted sum over all locations, with weights given by the probability of a random walk of length t joining those two locations. The transition probability for the random walks was given by the geometric affinity matrix A . We will use the modified surface normals $N(t)$ to perform path integration.

6 Experiments

Experiments have been performed with surface normal data corresponding to a human face. A field of surface normals was the input, and the output was the

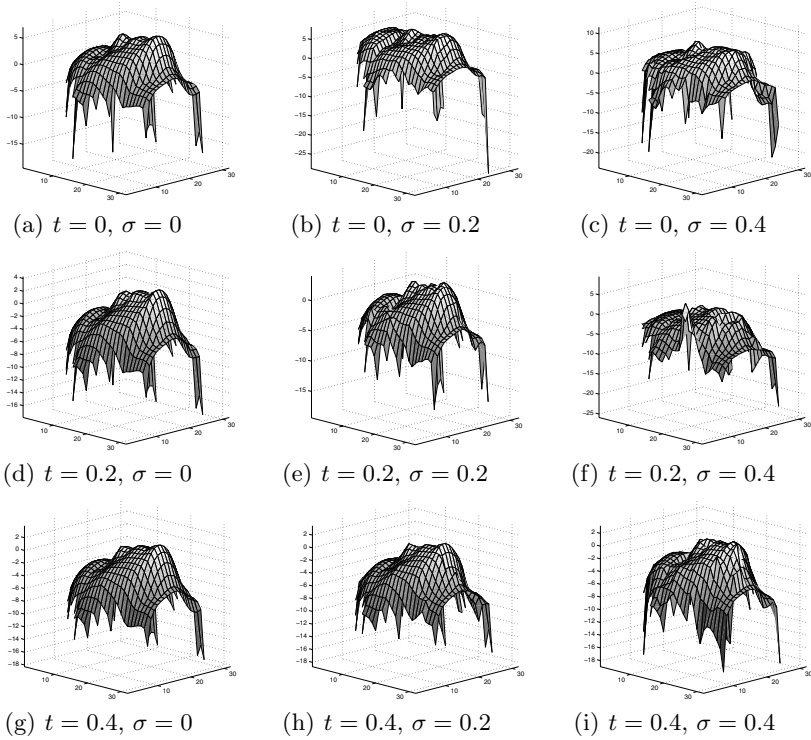


Fig. 4. Surface reconstruction for varying levels of Gaussian noise σ (columns) and time t (rows)

reconstructed surfaces. The performance was assessed with Gaussian noise added to the field of surface normals, in order to simulate a source of measurement noise. The experiment parameters are therefore the standard deviation of the Gaussian noise, and the time parameter t used to evaluate the heat kernel.




Fig. 5. Reconstruction error for $t = 0.4$, measured in RMS as a function of σ . The minimum corresponds to the bottom-right reconstruction in Figure 4.

The diffusion parameter t produced a reasonable smoothing of the field of surface normals for values below 1. For example, using diffusion parameter $t = 0.4$, the reconstruction is not affected by noise of parameter $\sigma < 0.5$. Figure 4 illustrates the reconstruction from a field of size 32×32 , for varying noise and time. (A raised chin usually corresponds to a lower reconstruction error, even when detail is not recovered). Figure 5 shows the RMS error between the ground truth and the reconstructed surface, for this value of t .

7 Conclusion

We have presented a method for path-based surface integration whose only parameters are a local affinity function and a time parameter t . We have also presented an affinity function based on an intrinsic geometric property of the surface being reconstructed, namely the principal curvature directions. The affinity is propagated over the graph of surface locations using a diffusion process. The result is used to re-weight the field of surface normals in order to make use of its spatial redundancy.

A future direction for this work is to interpret the diffusion process in the anisotropic diffusion framework presented by [4], and to state explicitly the relationship between the affinity function and the error model for the field of surface normals in a probabilistic setting.

References

1. Frankot, R.T., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence* **10** (1988) 439–451
2. Zhang, R., Tsai, P.S., Cryer, J., Shah, M.: Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21** (1999) 690–706
3. Agrawal, A., Chellappa, R.: An algebraic approach to surface reconstruction from gradient fields. In: *Proceedings of ICCV 2005*. (2005) 23–114
4. Agrawal, A., Chellappa, R.: What is the range of surface reconstructions from a gradient field. To appear in the proceedings of *ECCV 2006* (2006)
5. Robles-Kelly, A., Hancock, E.: Steady state random walks for path estimation. In: *Structural, Syntactic, and Statistical Pattern Recognition*. Volume 3138 of LNCS., Springer (2004) 143–152
6. Klette, R., Schläins, K.: Height data from gradient fields. In: *Proceedings of SPIE*. Number 2908 (1996) 204–215
7. Kondor, R.I., Lafferty, J.: Diffusion kernels on graphs and other discrete structures. *ICML 2002* (2002)
8. Zhang, F., Qiu, H., Hancock, E.: Evolving spanning trees using the heat equation. In: *Int. Conf. on Computer Analysis of Images and Patterns*. (2005)
9. Gibbons, A.: *Algorithmic Graph Theory*. Cambridge University Press (1985)
10. Perona, P., Malik, J.: Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12** (1990) 629–639