

A Method for Re-illuminating Faces from a Single Image

Mario Castelán* and Edwin R. Hancock

Dept. of Computer Science, University of York, York YO10 5DD, UK
{mario, erh}@cs.york.ac.uk

Abstract. In this paper we propose a local-shape based method for correcting shape-from-shading information for raw surface height recovery of faces aiming to generate differently illuminated versions of a face. The underpinning idea comes from the observation that subtle changes in the elements of a gradient field can cause notable changes in its integrated surface. A new gradient field is calculated by modifying its orientations in accordance with critical points on the surface and local shape indicators. A raw height map is then calculated by integrating such field. Experiments show that altering the directions of a surface normal field of a face can be enough to generate differently illuminated fixed-pose versions of a face from a single image.

1 Introduction

The problem of acquiring surface models of faces is an important one with potentially important applications in biometrics, computer games and production graphics. There are many ways in which surface models can be acquired, and these include the use of range-scanners, stereoscopic cameras and structured light sensors. However, one of the most appealing methods is to use shape-from-shading (SFS), since this is a non-invasive process which mimics the capabilities of the human vision system. Unfortunately, the use of SFS for face analysis has proved to be an elusive task, since the concave-convex ambiguity can result in the inversion of prominent features such as the nose. To overcome this problem, domain specific constraints must be used. For instance Zhao and Chellappa [9] have exploited facial symmetry to overcome the problem.

In general, though, SFS is an under-constrained problem since the two degrees of freedom for surface orientation (slant and tilt), must be recovered from a single measured intensity value. Hence, it is frequently posed as that of minimizing cost functionals that capture constraints on the gradient field. This is usually carried out through iterative schemes in the discrete domain. Despite sustained research activity in the field for some three decades[7, 10], no SFS scheme has been demonstrated to work as accurately as the *specialty* constrained SFS

* Supported by National Council of Science and Technology (CONACYT), Mexico, under grant No. 141485.

cases of photometric stereo[3, 5] and statistical SFS[1]. The first of these methods requires at least three images of the same object illuminated from different view-points, while the second uses a database of accurate surface information (height or gradient) belonging to objects of the same class, i.e. faces. When it comes to the original single-image SFS problem, the resulting data is seriously affected by the noise caused by factors such as inaccurately calculated illumination direction and unknown reflectance properties, and both of these are difficult to obtain from a single image. The restrictions imposed by most SFS schemes on the gradient field (smoothness, irradiance, integrability, unit length) are insufficient to overcome with these problems. Moreover, if misused, they may result in either oversmoothing of the gradient field, loss of fine surface detail, or over-dependence on the intensity information. As a result, the recovered height surface may represent an oversmoothed reconstruction of the object under study.

However, although a precise height map is difficult to obtain by integration of the field of surface normals delivered by SFS, due to local errors in direction, this does not imply that the entire field of surface normals is in error. In fact, some regions on the image do provide surface normal information that is sufficiently faithful for good surface reconstruction. The problem originates from those locations where the combined effects of image noise, and overreliance on the consistency constraints, results in effects such as inversion of the surface convexity or concavity. When applied to the problem of face reconstruction, for instance, the effect can be to cause high curvature features such as the nose to become imploded with respect to the remainder of the surface. If such regions can be identified and the surface normal directions corrected, then the result is improved overall surface topography.

The outline of the this paper is as follows. We commence with a brief overview of the Frankot and Chellappa's global surface integration method and local shape indicators, then we develop a method that can be used to reassign the surface gradient orientations. We provide experiments to evaluate the method on a human face. Finally, we present some conclusions and identify directions for future work.

2 Integrating a Gradient Field

The integrability condition in SFS ensures that the recovered surface satisfies the following condition on the partial derivatives of the height function: $Z_{xy} = Z_{yx}$. In [4] Frankot and Chellappa proposed a method to project a gradient field to the nearest integrable solution. They suggested to use a set of integrable basis functions to represent the surface slopes so as to minimize the distance between an ideally integrable gradient field and a non integrable one.

Following [4], if the surface Z is given by

$$\tilde{Z}(x, y) = \sum_{\omega \in \Omega} \tilde{C}(\omega) \phi(x, y, \omega)$$

where ω is a two dimensional index belonging to a domain Ω , and $\phi(x, y, \omega)$ is a set of basis functions which are not necessarily mutually orthogonal, the partial

derivatives of \tilde{Z} can also be expressed in terms of this set of basis functions using the formulae

$$\tilde{Z}_x(x, y) = \sum_{\omega \in \Omega} \tilde{C}(\omega) \phi_x(x, y, \omega) \quad \text{and} \quad \tilde{Z}_y(x, y) = \sum_{\omega \in \Omega} \tilde{C}(\omega) \phi_y(x, y, \omega)$$

Given that $\phi_x(x, y, \omega)$ and $\phi_y(x, y, \omega)$ are integrable, then so are the mixed partial derivatives of $\tilde{Z}(x, y)$.

In the same way, the possibly non integrable gradient field (which, indeed, is the only information we have) can be represented as

$$\hat{Z}_x(x, y) = \sum_{\omega \in \Omega} \hat{C}_1(\omega) \phi_x(x, y, \omega) \quad \text{and} \quad \hat{Z}_y(x, y) = \sum_{\omega \in \Omega} \hat{C}_2(\omega) \phi_y(x, y, \omega)$$

Note that, as $\hat{C}_1 \neq \hat{C}_2$, then $\hat{Z}_{xy} \neq \hat{Z}_{yx}$.

The goal then is to find the set of coefficients that minimize the quantity

$$d\{(\hat{Z}_x, \hat{Z}_y), (\tilde{Z}_x, \tilde{Z}_y)\} = \iint \left\| \tilde{Z}_x - \hat{Z}_x \right\|^2 + \left\| \tilde{Z}_y - \hat{Z}_y \right\|^2 dx dy$$

In [4], details are given about how to solve this equation globally, in the Fourier domain.

3 Local Shape Indicators

A local shape indicator is a scalar that gives information about the local topography of a surface using its principal curvatures. The principal curvatures may be estimated using the surface normal directions to compute the Hessian matrix. They may also be estimated by fitting a local quadric surface patch $P(u, v) = \frac{1}{2}(\kappa_1 u^2, \kappa_2 v^2)$ having κ_1 and κ_2 as the principal curvatures with directions u and v respectively. κ_1 and κ_2 are the calculated eigenvalues of the Hessian matrix.

Local shape indicators are usually coupled. For instance, the *HK* classification[2] uses the Gaussian and mean curvatures $H = \frac{(\kappa_1 + \kappa_2)}{2}$ and $K = \kappa_1 * \kappa_2$ respectively. By distinguishing between the cases in which *H* and *K* are individually negative, zero or positive, it is possible on the basis of the joint behavior to assign topographic labels to points on a surface.

A different and slightly more convenient set of attributes is the shape-index and curvedness representation developed by Koenderink and Von Doorn *SC*[8]. Here the principal curvatures are used to compute the shape index $S = -\frac{2}{\pi} * \arctan(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2})$ for $\kappa_1 \geq \kappa_2$; and the curvedness, $C = (\frac{\kappa_1^2 + \kappa_2^2}{2})^{1/2}$. The shape index is an angular variable that relates to the local surface topography. It varies continuously from -1 to $+1$ as the surface changes through cup, rut, saddle-rut, saddle, saddle-ridge, ridge and dome, and cup again. The curvedness relates to the degree of curvature of the surface.

4 Using the Curvedness Indicator to Redirect SFS Gradient Fields

Inevitably, any surface gradient field delivered by SFS will be inaccurate due to noise or albedo changes which cause variations in the intensities of the input image. SFS works well for objects that are uniformly concave or convex. However, if the object under study is more complex, with both concave and convex regions, then SFS can fail. In these situations although the recovered surface normal direction is consistent with the measured image brightness, the recovered surface does not reflect the structure of the object under study. In particular, there may be inversions of the sign of the surface curvature with convex regions appearing concave and vice-versa. However, in the case of faces (and many other objects) the surface under study is largely convex.

Based on this above observation, in this paper we present a method for enforcing the convexity of the integrated surface while ensuring a global maximum on a particular position of it.

Formally stated, suppose S is a smooth surface immersed in \mathbb{R}^3 . Let p and U_p be a critical point of S and a neighborhood of p respectively. Suppose that S is locally concave over U_p . Then, the new surface \tilde{S} constructed from S by reversing the sign of all its partial derivatives, S_x and S_y , is locally convex in U_p . Besides, a local maximum on \tilde{S} will be located on that point where the function stops increasing and starts decreasing¹; if all the partial derivatives of \tilde{S} with respect to x , \tilde{S}_x , have negative sign before reaching the position of the critical point p in the x axis and have positive sign after reaching it, and the same occurs for \tilde{S}_y , then the critical point p on U_p will be the position of the global maximum² of \tilde{S} .

The main idea underlying this paper is to enforce the condition that the integrated surface has a global height maximum where it is suitable for our problem, i.e. the tip of the nose. Such a point will serve as a division for the positive and negatively signed areas of the needle map. To enforce this condition we follow the simple rule:

$$\check{Z}_x(x, y) = \begin{cases} abs(\hat{Z}_x(x, y)) & \text{if } x \leq a \text{ and } C(x, y) \geq \tau_x \\ -abs(\hat{Z}_x(x, y)) & \text{if } x > a \text{ and } C(x, y) \geq \tau_x \\ \hat{Z}_x(x, y) & \text{otherwise} \end{cases}$$

$$\check{Z}_y(x, y) = \begin{cases} abs(\hat{Z}_y(x, y)) & \text{if } y \leq b \text{ and } C(x, y) \geq \tau_y \\ abs(\hat{Z}_y(x, y)) & \text{if } y > b \text{ and } C(x, y) \geq \tau_y \\ \hat{Z}_y(x, y) & \text{otherwise} \end{cases}$$

¹ Of course, \tilde{S} will present many local maxima for a face-like surface.

² It might be a maximum or a minimum depending on the integration method.

where \check{Z}_x and \check{Z}_y are the updated gradients, \hat{Z}_x and \hat{Z}_y are the original gradients, a and b are the coordinates for the position of the desired global height maximum, on the x and y axis respectively. $C(x, y)$ is the curvedness indicator which is compared to the thresholds τ_x and τ_y for deciding whether the element of the gradient field at the location (x, y) will be altered or not.

To illustrate the global height maximum enforcement procedure, Figure 1 (left) shows the results of applying the method to the derivatives of a sphere with radius 75 units. From left to right, transversal plots of the recovered surface on the x axis are shown. The peak coordinates (a, b) are both set to 75, 50 and 20 respectively, and τ_x and τ_y are set to zero: all the derivatives are taken into account to expose the extreme case. The convexity strengthening is clearer in Figure 1 (right), where the method is applied to the Mexican hat function. Cross sections of the recovered surface are shown, and from left to right they show the original surface and recovered surface after applying the method taking as a peak point the center of the surface with $\tau_x = \tau_y = 0$. Note how the concave parts of the hat become convex.




Fig. 1. Applying the method to a sphere and mexican hat.

It is evident that the peak-enforcement procedure will segment the recovered surface into four quadrants. As a result the curvedness of the recovered surface will be reduced. This is not desirable from the point of view of surface recovery from a face. The net effect will be to make the surface structure pyramidal. This problem is overcome by using the thresholding procedure to either force the normals to change direction or to allow them to remain unchanged.

In the following section some experiments will be presented in order to illustrate these points on an application involving raw face reconstruction using SFS and its use for calculating differently illuminated images of a face.

5 Experiments

The tests were carried out on the image shown in Figure 3(top row - left)³. To compute the surface gradients from the raw image brightness we followed a geometric method. The method involves first computing the image gradient. The surface normals are positioned on a cone whose axis is in the light source direction and whose opening angle is given by $\theta = \arccos E$. The position of the

³ The face database was provided by the Max-Planck Institute for Biological Cybernetics in Tuebingen, Germany.

surface normal on the cone is such at its projection onto the image plane is the same as the direction of the measured image gradient. This construction ensures the image irradiance equation to be satisfied. For the surface integration step we used the global method proposed in [4] which recovers surface height using the inverse Fourier transform of the field of surface normal directions.

The first analysis is shown in figure 2. From left to right: curvedness, first partial derivative with respect to x , first partial derivative with respect to y , $[0, 0, 1]$ re-illumination of the gradient field and recovered height map. The first row corresponds to the originally obtained needle map whilst the second row shows the case for the modified gradient field with $\tau_x = 0.4$ and $\tau_y = 0.3$. It is important to mention that a post-processing step of smoothing was realized in order to soften the recovered surface for better illumination effects. Note how the curvedness (first column) is diminished in the second row, as a consequence of the convexity enforcement, which becomes evident when analyzing the partial derivatives of each case (second and third columns), when the imploded regions of the face seem to emerge for the modified case. The perpendicular re-illumination of the needle maps (fourth column) also shows the convexity enforcement effect, since the second row case seems to naturally fit the illumination of the *surface* of a face rather than the re-illumination of the *image* of a face. The fifth column makes evident the improvement in the recovered height map, showing a profile plot of the surface. Note how the formerly imploded features of the face (first row) become salient (second row). Special comment deserves the choice of an appropriate threshold, since we are able to enhance the salient features of a particular face while maintaining the overall face composition.

It is important to remark that, if considerable, the improvement of the recovered height map depends on the accuracy of the input data, therefore the obtained raw height information still lacks the preciseness of that delivered by invasive methods. However, our method could be used for tasks such as generation of unseen illuminated views with fixed pose, as suggested by the following set of experiments.




Fig. 2. Curvedness, gradient and surface recovery analysis (see text).

Figure 3 presents re-illumination experiments for a couple of images of faces. The first column corresponds to the input image. The next columns show the generated re-illuminations after applying the method. The light source vector is nearly parallel to the x axis in both negative in positive directions, for the second and third columns, while nearly parallel to the y axis, for the fourth and fifth columns. The input images were assumed to be the albedo maps, and negative values for the re-illuminations (zero values under the lambertian model) were forced to be positive for illustrative clarity.



Fig. 3. Re-illumination tests for two different images (see text).

A further analysis is shown in figure 4, where an unprocessed image of a face (first column, middle of the rows) was used for the experiments. The first row presents the results for the modified gradient field with both $\tau_x = \tau_y = 0.3$. The second row exhibits the case for the original gradient field. Note how in the top row the quality of the re-illuminations is improved, contrasting the evidently imploded features presented in the second row case. It is necessary mentioning, however, that some incorrectness appears in some areas of the face, i.e. those



Fig. 4. Comparison of re-illuminations (see text).

surrounding the mouth. This can be explained as the consequence of the change in reflectance properties on the lips. Of course, the eyes and the mustache area also present different reflectance properties. This suggests that more attention should be paid on the reflectance model for the re-illumination, and also for the recovery of the gradient field.

6 Conclusions

We have proposed an algorithm for correcting a gradient field of a face aiming to generate a height map with a global maximum on a critical point and whose concave regions are constrained to become convex in accordance with curvedness indicators. The recovered height maps from the new gradient fields can be used to generate re-illuminated versions of the face in a straightforward way. Future lines of work are proposed towards the use of a special reflectance model for the skin as well as taking into account face-symmetry constraints for the recovery of a better height map.

References

1. Atick, J., Griffin, P. and Redlich, N. (1996), Statistical Approach to Shape from Shading: Reconstruction of Three-Dimensional Face Surfaces from Single Two-Dimensional Images, *Neural Computation*, Vol. 8, pp. 1321-1340.
2. Besl, P.J. and Jain, R.C. (1986) Invariant surface characteristics for 3-d object recognition in range images, *Comput. Vision Graphics Image Proc.*, Vol. 33, pp. 33-80.
3. Forsythe, D. and Ponce, J. (2001), *Computer Vision: a Modern Approach*, Prentice-Hall.
4. Frankot, R.T. and Chellapa, R. (1988), A Method for Enforcing Integrability in Shape from Shading Algorithms, *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 10, No. 4, pp. 438-451.
5. Georghiadis, A., Belhumeur, D. and Kriegman, D. (2001), From Few to Many: Illumination Cone Models for Face Recognition under Variable Lighting and Pose, *IEEE Trans. Pattern Analysis and Machine Intelligence*, pp. 643-660.
6. Horn, B.K.P. (1977), Understanding Image Intensities, *Artificial Intelligence*, Vol. 8, pp. 201-231.
7. B.K.P. Horn and M.J. Brooks (1989), *Shape from Shading*. MIT Press, Cambridge, MA.
8. Koenderink, J.J., and Van Doorn, A.J. (1992), Surface Shape and Curvature Scales, *Image en Vision Computing*, Vol. 10, pp. 557-565.
9. Zhao, W. and Chellapa, R. (2000), Illumination-insensitive Face Recognition Using Symmetric Shape-from-Shading. *Conference on Computer Vision and Pattern Recognition*, pp. 286-293.
10. R. Zhang, P.S. Tsai, J.E. Cryer and M. Shah (1999), Shape from Shading: A Survey, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(8):690-706.