

Single image facial view synthesis using SFS

W.A.P. Smith and A. Robles-Kelly and E.R. Hancock
Department of Computer Science, The University of York, UK
{wsmith, arobkell, erh}@cs.york.ac.uk

Abstract

We present a facial view synthesis technique based on explicit shape and reflectance information extracted from a single image. The technique combines an image based reflectance estimation process with a novel method of interpolating between needle-maps recovered using shape from shading. This allows images of a face to be synthesised under novel lighting, pose and skin reflectance given only one example image.

1 Introduction

Synthesising images of a face under novel lighting, viewpoint and expression from one or more example images is an area that has attracted considerable attention in recent years [3, 15]. Such a process has a number of useful applications, most notably pose and lighting correction for use in face recognition systems and reanimation of faces from images. However, existing methods have tended to be entirely model or parameter driven and few have addressed the issues of lighting and reflectance [3, 1, 2].

In face recognition the problem of variation in pose has proven to be the greatest obstacle to robust performance. The FRVT2000 evaluation found that the best recognition rate fell to 68% for rotations in depth of 40° in a database of 200 subjects. There have been numerous attempts to overcome the problem of variation in pose without explicitly synthesising new views. Pentland et al. [8] extended the Eigenfaces technique of Turk and Pentland [10] in order to achieve pose invariant recognition, although their method required many example images of each subject to cover all possible views. A feature based approach was proposed by Wiskott et al. [11] which used elastic bunch graph matching. For rotations of 22° they achieved a recognition rate of 88%, though varying illumination would make feature location more difficult. Cootes et al. [5] extended the concept of the “Active Appearance Model” to account for variation in pose by building an AAM for a number of different poses, for example 0° , 45° and 90° . They then proposed learning the relationship between models at different views allowing them to represent a face in a particular view, given an image of the face in another view. This essentially 2D model-based approach was shown to give a reasonable prediction of novel views.

An alternative to developing a pose invariant recognition system is to introduce a preprocessing stage in which a pose corrected image of the input face is generated. This is known as recognition by synthesis. Previous work on facial view synthesis can be divided into those techniques which require more than one example image and those that require only one.

Given more than one image, a generic stereo technique could be applied to face images, for example the trilinear tensor method of Avidan and Shashua [1]. Georgiades et

al. [6] developed a variant of the photometric stereo technique specifically for faces which was able to synthesise views of face under novel lighting and pose given as few as three images of the face taken under variable lighting. However, for many applications it is desirable to be able to synthesise novel views of a face given only a single example view. This requires assumptions to be made about occluded areas of the face and therefore the majority of single view techniques use a model which encapsulates prior knowledge of face structure and appearance. Beymer and Poggio [2] proposed such a technique based on image warping. However, any method based on image warping alone takes no account of reflectance and lighting, and moreover this technique cannot synthesise occluded areas of the face. More recently, an interesting model-based approach was proposed by Blanz et al. [3]. They construct a morphable model of appearance based on principal components analysis of 3D head shape and texture collected using a laser range finder. Given a single input image they simultaneously estimate pose, lighting, camera properties, face shape and face texture. Once the model has been recovered any novel pose can be rendered under any arbitrary lighting conditions. However, the technique requires manual initialisation, uses the Phong model to capture all skin reflectance and is reliant on the model successfully capturing all possible face shapes.

All the techniques mentioned can be criticised on account of their lack of an adequate model of reflectance and a tendency to be purely model driven. In contrast, shape from shading holds out the possibility of recovering 3D shape from a single image. This potentially allows the synthesis of new poses given a single example view without requiring a prior model of head shape. Previous attempts to use shape from shading for novel facial view synthesis have been limited. The most well known work is that of Zhao and Chellappa [15]. They exploited the symmetry in a frontal view of a face to constrain the shape from shading process and cancel the effects of variation in albedo, although they did assume Lambertian reflectance. They used this to synthesise novel lighting for frontal images. However, images of rotated faces are not symmetric and hence the constraint cannot be applied. They therefore used image warping based on a single 3D model to generate a frontal view before applying symmetric shape from shading. The only results presented were based on pure image warping of very low resolution images, though they did report improved recognition performance.

In the remainder of this paper we use shape from shading to develop a non-model-based, non-parametric facial view synthesis technique which allows us to synthesise images of a face under novel pose, lighting and skin reflectance given only one example image. The technique consists of two processes: image based reflectance estimation followed by view synthesis using needle-map correspondence. We describe these two processes in the next two sections after which we present the results of applying the technique to real world images.

2 Image based reflectance estimation

The physics of light reflection from a surface can be captured using the bi-directional reflectance distribution function (BRDF), which is the ratio of the incident radiance to the outgoing radiance per solid angle:

$$\rho(\theta_i, \phi_i, \theta_r, \phi_r) = \frac{f_o(\theta_r, \phi_r)}{f_i(\theta_i, \phi_i) \cos \theta_i d\omega}$$

where θ and ϕ are the zenith and azimuth angles respectively of the incident (i) and outgoing (r) radiance. Our aim is to estimate the BRDF of a subject's skin from a single image of their face. This will prove useful to the process of novel view synthesis in two ways. Firstly, by using the technique for photometric correction prior to applying shape from shading, it allows more accurate shape information to be recovered and hence improves the synthesised needle-map. Secondly, the synthesised needle-map can be rendered with the BRDF estimated from the input image resulting in improved realism in the final image or with a novel BRDF, increasing the variety of novel views available.

The problem is simplified considerably if the viewer and light source directions are identical, i.e. we work with the retroreflection condition $\theta_i = \theta_r$ and $\phi_i = \phi_r$. In addition, we assume constant albedo. Perhaps a more realistic approach would be to segment the face into areas of approximately constant albedo, for example the lips, the forehead etc, and estimate the BRDF separately for each area.

2.1 Representing the BRDF on the Gauss sphere

For an orientable surface $S \in \mathfrak{R}^3$, the Gauss map $G : S \mapsto \hat{S}$ maps points on the surface S onto locations on the unit Gauss sphere \hat{S} which have identical surface normal directions. Our aim is to use correspondences between surface normal directions to map brightness values from the image onto the unit sphere. The polar distribution of brightness on the unit sphere \hat{S} is the radiance function for the surface.

Assuming that the reflectance function is isotropic, circles of latitude on the image of the unit sphere will have the same brightness. In other words, given that the viewer and light source direction are equal, the brightness at a particular point is dependent only on the zenith angle θ of the surface normal at that point. Brightness does not vary with the azimuth angle ϕ . As a result, the observed brightness values mapped onto the unit sphere by the Gauss map G can be generated by revolving the function $g(\theta) = f_o(\theta, 0)$ in azimuth angle ϕ about the axis defined by the viewer and light source directions. The problem of describing the observed brightness distribution over the Gauss sphere hence reduces to that of approximating the function $g(\theta)$ and computing its trace of revolution.

Unfortunately, since the surface normals on the surface S are not known, the correspondences between locations on the surface and the Gauss sphere are unknown. In the next section we show how the relationship between the differential structure of the observed brightness in the image can be used to estimate $g(\theta)$.

2.2 Using the intensity gradients

We commence by rewriting $g(\theta)$ as the integral of the partial derivative of the observed brightness with respect to the angular variable θ . As a result, $g(\theta)$ is given by

$$g(\theta) = \frac{1}{2\pi} \int_0^{2\pi} \int_0^\theta \frac{\partial f_o(\theta, \phi)}{\partial \theta} d\theta d\phi \quad (1)$$

In other words, the generating function $g(\theta)$ on the unit sphere can be expressed in term of the cumulative distribution of the derivatives of the radiance function or alternatively the derivatives of the image brightness.

The relationship between the image gradient and the radiance is given by

$$\nabla I = \frac{1}{\det(J)} \begin{bmatrix} \frac{\partial f_{\phi}(\theta, \phi)}{\partial \theta} \\ \frac{\partial f_{\phi}(\theta, \phi)}{\partial \phi} \end{bmatrix} \quad (2)$$

where the Jacobian matrix is

$$J = \begin{bmatrix} \frac{\partial x}{\partial \theta} & \frac{\partial x}{\partial \phi} \\ \frac{\partial y}{\partial \theta} & \frac{\partial y}{\partial \phi} \end{bmatrix} \quad (3)$$

We can therefore relate the image gradient to the derivative of the function $g(\theta)$ with respect to the zenith angle θ

$$|\nabla I| = \frac{1}{\cos \theta} \frac{\partial g(\theta)}{\partial \theta} = \frac{\partial g(\theta)}{\partial \sin \theta} \quad (4)$$

2.3 Approximating the BRDF

On the unit sphere \hat{S} , it is always possible to choose points to be sampled so that the difference in brightness is a constant τ . We can estimate $\sin \theta$ using the cumulative gradient distribution as follows

$$\sin \theta = \int \frac{\tau}{|\nabla I|} dI = \sum \frac{\tau}{|\nabla I|} + \kappa \quad (5)$$

where κ is the integration constant. Hence, we can use the cumulative distribution of inverse gradients to index the zenith angle on the unit sphere. This indexation property means that we can approximate the function $g(\theta)$ by tabulation.

To pursue this idea, in principle, we only require a single image gradient corresponding to each of the distinct brightness levels in the image. In practice, we make use of the cumulative distribution of image gradients in order to minimise the approximation error by averaging. Let $H(l) = \{s \mid I_s = l\}$ be the set of pixels with brightness value l . For the brightness value $l = g(\theta)$, the average gradient is given by $h(l) = \frac{\sum_{s \in H(l)} |\nabla I_s|}{|H(l)|}$. We can therefore derive a tabular representation of $g(\theta)$ as a set of Cartesian pairs

$$g = \left\{ \left(\tau \sum_{s=1}^l \frac{1}{h(s)} + \kappa, l \right), l = 1, \dots, n_{max} \right\} \quad (6)$$

where n_{max} is the maximum brightness value in the image.

2.4 Lambertian Re-mapping

The technique described above allows us to estimate the reflectance function of a subject's skin from a single image, assuming the light source and viewer direction are co-incident. With the reflectance function to hand, it is straightforward to re-render the image using Lambertian reflectance. This is achieved by noting the observed brightness l at a pixel and identifying the associated value of $\sin \theta$ according to g . The corresponding Lambertian intensity is simply $\cos \theta$. This process removes specularities in images of faces due to perspiration and improves image contrast across the face. Figure 1 shows the effect of this process on a real world image. The raw image is shown on the left and the Lambertian corrected image is shown on the right.



Figure 1: Lambertian re-mapping




Figure 2: Two views from one

3 View interpolation using shape from shading

If a 3D representation of an object is available, view synthesis from any viewpoint is purely a rendering problem. Whilst it appears clear that one of the motivations of early shape from shading research was to enable such a 3D representation to be derived from a single image, the difficulties encountered in achieving accurate and robust needle-map recovery have proved a serious obstacle to progress in this direction, in particular given the variation in albedo present in real world images.

These difficulties have led to the development of direct view synthesis techniques which do not involve explicitly constructing a 3D model. More recently, Worthington and Hancock [14] have proposed an interpolation technique based on needle-maps. This has the advantage over traditional stereo techniques that any synthesised needle map can be rendered with novel lighting and reflectance properties, increasing the variety of novel views possible. Worthington [12] has extended this work by developing a true shape from shading based view synthesis approach by making use of the dense correspondence between needle-maps.

In the remainder of this section we extend this technique to account for the rotation between views in order to allow new face images to be synthesised by interpolating between the needle-map of a rotated face and its vertical reflection. We use the geometric shape from shading technique of Worthington and Hancock [13] to recover our input needle map.

3.1 Two views from one

The observation underpinning this paper is that, under the assumption of facial symmetry, the vertical reflection of an image of a face rotated an angle θ about the Y-axis, yields an image of the face rotated by an angle $-\theta$ about the Y-axis. In Figure 2 the image on the right is real, whereas the image on the left was generated synthetically by reflection. Interpolating towards the mid-view between the two images will produce a frontal view of the face. This idea also extends to rotations about the Z-axis and combined rotations about the Y and Z axes. This observation was noted by Gordon [7] but has not been exploited in the context of faces.

However, in typical stereo pairs the rotation is considered small enough to be ignored and correspondences can be calculated on the basis of image intensity or intensity gradients. This is not the case with rotated faces, where pose correction from three quarter view (45 degree rotation) or even greater is desirable. Therefore, an approach which accounts for reflectance and lighting is needed.

3.2 Needle-map correspondence

In [12] Worthington introduced an extension to the view interpolation technique of Scharstein [9]. Rather than using image gradients to find dense correspondences between two im-

ages, Worthington proposed using the surface normals generated by shape from shading. The mid view is generated by transporting normals to an interpolated position. The resulting needle map can then be rendered with frontal or novel lighting.

The dense correspondence between two needle maps, \mathbf{N}_L and \mathbf{N}_R , is established by accumulating evidence for each surface normal in \mathbf{N}_L with each surface normal in \mathbf{N}_R within the allowable disparity range. This process yields a disparity map \mathbf{d}_{LR} and is repeated to find \mathbf{d}_{RL} . In order to improve local consistency, the evidence values are smoothed for a given disparity. This encourages neighboring surface normals to have similar disparities.

Rather than performing shape from shading on both face images, in our case the needle map \mathbf{N}_R is simply the reflection of \mathbf{N}_L about the Y-axis. This means shape from shading need only be executed once.

3.3 Evidence measure

The evidence measure, e , proposed by Worthington [12] utilised the inner, or dot, product between two surface normals to measure similarity:

$$e_{LR}(i, j) = \max_{\Delta x, \Delta y} \frac{e^{\mathbf{N}_L(i, j) \cdot \mathbf{N}_R(i + \Delta x, j + \Delta y)} - e^{0.5}}{e^1 - e^{0.5}} \quad (7)$$

where \mathbf{N}_L and \mathbf{N}_R are a pair of needle maps. The exponential aimed to encourage good matches while penalising poor matches. However, this formulation still assumes that the rotation is small enough to be discounted and directly compares surface normals in the two needle maps. We therefore propose a reformulation of Worthington's evidence measure which accounts for the rotation between views.

We assume that the angle of rotation from a frontal pose is known. Robust techniques exist for estimating subject pose from an image of a face, for example, the shape from shading based facial pose estimation technique of Choi et al [4] which would be particularly suitable, since the needle map of the face will already be to hand.

In order to account for the rotation between views, we explicitly rotate surface normals from the left hand needle map by twice the estimated rotation angle before matching with surface normals in the right hand needle map. To do this, we propose the following revised evidence measure:

$$e_{LR}(i, j) = \max_{\Delta x, \Delta y} \frac{e^{(\mathbf{R}_y(-2\theta)\mathbf{N}_L(i, j)) \cdot \mathbf{N}_R(i + \Delta x, j + \Delta y)} - e^{0.5}}{e^1 - e^{0.5}} \quad (8)$$

where θ is the estimated rotation about the Y-axis in a clockwise direction and $\mathbf{R}_y(\alpha)$ is the matrix which rotates a normal an angle α about the Y-axis. Likewise, e_{RL} is given by:

$$e_{RL}(i, j) = \max_{\Delta x, \Delta y} \frac{e^{(\mathbf{R}_y(2\theta)\mathbf{N}_R(i, j)) \cdot \mathbf{N}_L(i + \Delta x, j + \Delta y)} - e^{0.5}}{e^1 - e^{0.5}} \quad (9)$$

Clearly, these evidence measures could be extended to also include rotations about the Z-axis.

3.4 View Synthesis

The above correspondence process will recover a pair of disparity maps \mathbf{d}_{LR} and \mathbf{d}_{RL} , which describes the displacement in pixels from the normal in \mathbf{N}_L to the corresponding




Figure 3: The forward mapping process

normal in \mathbf{N}_R and vice-versa. From these disparity maps we can interpolate any view between the original views or in fact even outside the two original views. As well as transporting each surface normal to its intermediate position, we must again account for the rotation between views. We do this by rotating the normals through a fraction of the total rotation defined by the intermediate viewpoint required. Given that the intermediate needle-map is denoted \mathbf{N}_C , we therefore reformulate Worthington’s forward mapping as follows:

$$\mathbf{R}_y(-2X_s\theta)\mathbf{N}_L(i, j) \rightarrow \mathbf{N}_{C,L}(i + X_s\mathbf{d}_{LR_x}(i, j), j + Y_s\mathbf{d}_{LR_y}(i, j)) \quad (10)$$

$$\mathbf{R}_y(2(1 - X_s)\theta)\mathbf{N}_R(i, j) \rightarrow \mathbf{N}_{C,R}(i + (1 - X_s)\mathbf{d}_{RL_x}(i, j), j + (1 - Y_s)\mathbf{d}_{RL_y}(i, j)) \quad (11)$$

where X_s and Y_s define an intermediate viewpoint. The geometry of this process is illustrated in Figure 3. For the frontal view $X_s = Y_s = 0.5$. Views outside the original views are generated when X_s or Y_s do not lie in the range $(0, 1)$.

The two intermediate needle maps, $\mathbf{N}_{C,L}$ and $\mathbf{N}_{C,R}$ need to be combined to produce \mathbf{N}_C . We chose to select the surface normal for which the evidence value is greater, though averaging the two surface normals produces similar results.

Holes are present in the synthesised view due to errors in the disparity maps and scene points which are invisible in the two input views. We choose to fill these holes in a similar way to Scharstein, using simple scan-line interpolation. However, rather than interpolating intensity values across the holes, and therefore making assumptions about reflectance, we interpolate surface normals, yielding a complete synthesised needle-map which can be rendered. However, there is clearly scope for the use of more accurate techniques such as surface fitting.

3.5 Summary of the Process

- Estimate BRDF from input image and re-render using Lambert’s law
- Apply shape from shading to recover \mathbf{N}_L and reflect this to recover \mathbf{N}_R
- Find dense correspondence between needle-maps to determine \mathbf{d}_{LR} and \mathbf{d}_{RL}
- Transport and rotate surface normals to intermediate positions using the disparity maps to produce $\mathbf{N}_{C,L}$ and $\mathbf{N}_{C,R}$
- Combine the pair of needle maps by selecting the normals with highest evidence at each point
- Render the new needle-map with original BRDF or any other and desired lighting

It is worth noting that since this technique uses no prior class knowledge, it could be applied to any symmetric object, not only faces.




Figure 4: Synthesised novel poses

4 Experiments

In this section we report the results of applying our facial view synthesis procedure to real world images of faces. In all the face images, areas in which shape from shading cannot be meaningfully applied, such as the eyes and eyebrows, have been manually removed. In these areas traditional stereo matching techniques could be used to find correspondences, resulting in complete synthesised images without the gaps present in the images presented in this paper. Some of the images used are from the Yale Face Database B [6] and all are illuminated with a single point light source situated close to the camera.

4.1 Novel Pose Synthesis

Figure 4 shows the results of applying our needle-map interpolation technique to images of 3 subjects. Columns 2 and 6 show the input and reflected images. The synthesised needle maps were rendered with Lambertian reflectance and a single light source situated at the viewpoint. Columns 3, 4 and 5 show synthesised views within the viewing cone of the input images, with $X_s = 0.25, 0.5$ and 0.75 respectively, while columns 1 and 7 show views outside the viewing cone, where $X_s = -0.25$ and 1.25 respectively.

Given the simplicity of the hole filling scheme employed and the lack of any smoothing or noise removal, the results of the view synthesis are very encouraging. Certainly, the quality of the synthesis is comparable with that of Zhao and Chellappa [15], without having resorted to using a 3D model. The area in which the greatest difficulties have been encountered are the edges of the nose. The occlusion of these areas in the input images means the correspondences are likely to be incorrect and therefore the surface normals have been transported to the wrong location. The poses outside the original views appear as accurate as those within, perhaps because the fact that no occluded areas are being brought into view compensates for the extrapolation of the surface normals. To provide a more quantitative assessment of the pose synthesis process, we calculated the percentage relative error of the synthesised images by aligning a range scanned 3D model of the subject in the third row with the synthesised image and finding the error. For $X_s = \pm 0.25$, i.e. $\pm 7^\circ$, the error was less than 2%, though this increased to 18% when $X_s = 0.5$.

In order to assess the quality of the combined pose and reflectance synthesis, Figure 5(a) compares a real frontal view of each of the subjects with the synthesised frontal view. For this comparison, the synthesised needle map has been rendered with the BRDF




Figure 5: (a) comparison of real with synthetic view (b) re-illumination of synthetic views estimated from the input image. There is a good agreement in the placement of specularities, the shape of the synthesised and original faces appears similar and, noise aside, the process as a whole appears reasonably accurate.

4.2 Re-illumination

One of the motivations behind interpolating surface normals as opposed to intensity is that the synthesised needle map can be rendered with arbitrary lighting. Figure 5(b) shows the synthesised frontal views of the three subjects under varying lighting rendered with Lambertian reflectance. From left to right the light source is at 0° , 10° , 20° and 30° to the original light source along the x -axis. The needle maps are reasonably stable under re-illumination, in fact more so than the original needle maps extracted from the input images. This is perhaps because the synthesised needle maps are interpolated from two other needle maps, reducing the errors introduced by the shape from shading process.

5 Conclusions and Future Work

In this paper we have presented a novel approach to the problem of facial view synthesis. By combining a reflectance estimation process with a technique for interpolating between needle maps we are able to synthesise novel poses given only one image without recovering the 3D shape of the face. Since we interpolate between needle maps, the synthesised view can be rendered with novel lighting or reflectance.

This has a number of potential applications. In face recognition the technique can be used to perform recognition by synthesis. Alternatively, the technique could be used to reduce the data storage and collection requirements of a view-based pose invariant recognition system, by generating additional views of a subject in novel pose and lighting. The technique could be used in graphics for reanimation of faces from a single image.

There are a number of ways in which the work described in this paper can be further developed. The process of finding correspondences could be improved by using rotation invariant evidence measures such as local curvature or relative normal direction. This would remove the dependence on accurate pose estimation. In addition, where there is insufficient evidence to match a surface normal with any other, the disparity could be

estimated from neighboring normals. Finally, the smoothing step applied to the evidence values should also be improved so as to avoid boundary blurring.

References

- [1] S. Avidan and A. Shashua. Novel view synthesis by cascading trilinear tensors. *IEEE Transactions on Visualization and Computer Graphics*, 4(4), 1998.
- [2] D. Beymer and T. Poggio. Face recognition from one example view. In *Proc. 5th Intl. Conf. on Computer Vision*, pages 500–507, 1995.
- [3] V. Blanz, C. Basso, T. Poggio, and T. Vetter. Reanimating faces in images and video. In *Proceedings of EUROGRAPHICS*, 2003.
- [4] K. N. Choi, P. L. Worthington, and E. R. Hancock. Estimating facial pose using shape-from-shading. *Pattern Recognition Letters*, 23(5):533–548, 2002.
- [5] T. F. Cootes, K. N. Walker, and C. J. Taylor. View-based active appearance models. In *Proc. Int. Conf. on Face and Gesture Recognition*, pages 227–232, 2000.
- [6] A. S. Georgiades, P.N. Belhumeur, and D.J. Kriegman. Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. In *IEEE Workshop MMAVS 99*, pages 47–54, 1999.
- [7] G. G. Gordon. Shape from symmetry. In *Proc. of SPIE, Intelligent Robots and Computer Vision VIII: Algorithms and Techniques*, 1989.
- [8] A. Pentland, B. Moghaddam, and T. Starner. View-based and modular eigenspaces for face recognition. In *Proc. of IEEE Conf. CVPR*, 1994.
- [9] D. Scharstein. *View Synthesis Using Stereo Vision. Lecture Notes in Computer Science (LNCS), volume 1583*. Springer Verlag, 1999.
- [10] M. Turk and A. Pentland. Face recognition using eigenfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–591, 1991.
- [11] Laurenz Wiskott, Jean-Marc Fellous, Norbert Krüger, and Christoph von der Malsburg. Face recognition by elastic bunch graph matching. In *Intelligent Biometric Techniques in Fingerprint and Face Recognition*, pages 355–396, 1999.
- [12] P. L. Worthington. Novel view synthesis using needle-map correspondence. In *13th British Machine Vision Conference*, 2002.
- [13] 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.
- [14] P. L. Worthington and E. R. Hancock. Coarse view synthesis using shape-from-shading. *Pattern Recognition*, 36(2):439–449, 2003.
- [15] W. Y. Zhao and R. Chellappa. SFS based view synthesis for robust face recognition. In *4th IEEE Conf. AFGR*, pages 285–292, 2000.