Minimal Image Sets for Robust Spherical Gradient Photometric Stereo

Abhishek Dutta and William A.P. Smith
Department of Computer Science, University of York *

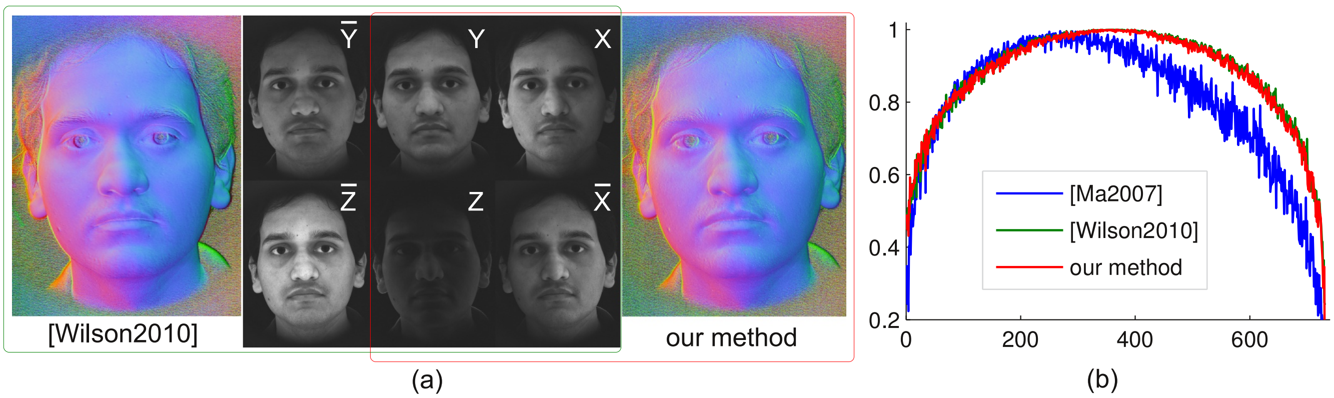

Figure 1: *(a) Facial normal map capture using 6 images [Wilson et al. 2010] (left) and using the proposed 4 image approach (right). (b) z-component of estimated surface normals of a white cylinder (1 pixel wide horizontal stripe).*

Abstract

State-of-the-art photometric shape and reflectance analysis uses polarized spherical gradient illumination images captured in a light stage. The quality of the estimated surface geometry degrades with interreflection, rotationally asymmetric reflectance lobes and light discretisation. We show that introducing additional parameters to model distortions in the diffuse reflectance lobe results in an underdetermined linear system. Results from existing approaches can be refined to satisfy this system using quadratic programming. We also describe how robust shape recovery can be achieved using a minimal four image set.

CR Categories: I.4.1 [IMAGE PROCESSING AND COMPUTER VISION]: Digitization and Image Capture—Reflectance, Radiometry;

Keywords: light stage, ambient occlusion, interreflection

1 Introduction

Ma et al. [2007] have shown how high resolution shape and reflectance information can be measured using an extended version of photometric stereo. An object is placed at the centre of a 'light stage' which uses polarized spherical gradient illumination arranged such that the plane of polarisation of all illuminants is equal. This allows specular and diffuse reflections to be separated by acquiring parallel and cross polarised images. The key observation underpinning the approach is that the centroid of the diffuse or specular reflectance lobe coincides with the surface normal or reflection vector respectively. The insight of Ma et al. [2007] was to show how to estimate the reflectance centroids using gradient illumination conditions. When integrated with an illumination gradient in the X, Y or Z direction, the corresponding component of the reflectance centroid, and hence surface normal, can be recovered.

The quality of the recovered surface geometry is affected by the extent to which the following assumptions are satisfied: (a) no shadowing of light sources, i.e. object is convex; (b) no interreflections; and (c) light sources closely approximate a continuous illumination environment. The last assumption can be addressed by maximising the number of light sources in the light stage (Ma et al. [2007] used 156 LEDs). The method also ignores light source attenuation effects, which is equivalent to assuming all points on the object lie exactly at the centre of the light stage.

More recently, Wilson et al. [2010] showed how use of complementary gradient images (i.e. reversed illumination, giving 6 gradient images plus one under constant illumination) can be used for photometric alignment. This exploits the constraint that the sum of a gradient image and its complement should equal the constant image. In this sketch, we show that by relaxing the assumptions described above, we can post-process a shape estimate by computing distortion parameters using quadratic programming. As an alternative, we show how a minimal image set can be used to estimate shape whilst still preserving the distortion invariance made possible by our relaxed assumptions.

2 Background

For any Lambertian surface patch, the value of radiance is given by

$$r = \int_\Omega P(\omega)R(\omega, n)d\omega = \int_\Omega P(\omega')R(\omega', [0,0,1])d\omega', \quad (1)$$

where, $P(\omega)$ and $P(\omega')$ represent the intensity of light incoming from direction ω (global coordinates) and ω' (local coordinates) respectively and R is the Lambertian reflectance function. In the case of X-gradient illumination, the intensity of light incident from direction $\omega' \in \Omega$ is proportional to the X-component of $\omega \in \Omega$ (the corresponding incident direction represented in global coordinates). As it is not possible to emit light with negative intensity, we cannot realise an X-gradient illumination with $P(\omega') \in [-1, 1]$. Hence, we rescale as follows:

$$P(\omega') = \frac{P_x(\omega) + 1}{2} = \frac{(\omega'_u u_x + \omega'_v v_x + \omega'_n n_x) + 1}{2} \in [0, 1]. \quad (2)$$

We can write the radiance equation for X-gradient illumination as:

$$r_x = \int_\Omega V_{p,\omega'} \left(\frac{\omega'_u u_x + \omega'_v v_x + \omega'_n n_x + 1}{2} \right) R(\omega', [0,0,1])d\omega', \quad (3)$$

*{abhishek, wsmith}@cs.york.ac.uk

where $V_{p,\omega'}$ is the binary occlusion term. Both [Ma et al. 2007] and [Wilson et al. 2010] assume that the surface is convex (i.e. $V_{p,\omega'} = 1, \forall \omega' \in \Omega$) and that the diffuse reflectance is symmetric about the surface normal. Hence, the integral over the hemisphere along u_x and v_x axes becomes 0 and the X-gradient radiance simplifies to

$$r_x = \frac{\pi \rho_D}{2} \left\{ \frac{1}{3} n_x + \frac{1}{2} \right\}, \tag{4}$$

where ρ_D is the diffuse albedo. Similar expressions can be derived for Y and Z gradients. Under the same assumptions, constant illumination (i.e. $P(\omega') = 1$) can be used to recover the albedo: $r_c = \pi \rho_D / 2$, and hence the surface normal components can be obtained by taking ratios between gradient and constant images:

$$n_{\{x,y,z\}} = \frac{1}{N} \left(\frac{r_{\{x,y,z\}}}{r_c} - \frac{1}{2} \right). \tag{5}$$

where, N is a normalizing constant. A similar analysis can be applied to the case where complementary gradient conditions [Wilson et al. 2010] are also used (i.e. where the coordinate system is reversed). In this case, the X component of the centroid is simply given by: $r_x - r_{\bar{x}}$, where $r_{\bar{x}}$ is the radiance under X complement illumination, and the surface normal components by:

$$n_{\{x,y,z\}} = \frac{1}{N} \left(r_{\{x,y,z\}} - r_{\{\bar{x},\bar{y},\bar{z}\}} \right). \tag{6}$$

3 Relaxing Assumptions

Interreflections and discretisation of the illumination environment deforms the diffuse reflectance function along the u_x and v_x axes. Hence, we cannot ignore the contribution of integrals along u_x and v_x. However, it is not possible to evaluate this integral and hence we define quantities $\delta_{\{x,y,z\}}$ which scale the diffuse albedo $\pi \rho_D$ to represent the overall deformation in the diffuse reflectance function in the x, y and z directions (similarly for the complement conditions). We also consider shadowing by introducing the ambient occlusion term $V_p \in [0, 1]$. Our generalised term for radiance under gradient illumination is therefore:

$$r_{\{x,y,z,\bar{x},\bar{y},\bar{z}\}} = \frac{\pi \rho_D V_p}{2} \left\{ \delta_{\{x,y,z,\bar{x},\bar{y},\bar{z}\}} + \frac{1}{3} n_{\{x,y,z,\bar{x},\bar{y},\bar{z}\}} + \frac{1}{2} \right\}, \tag{7}$$

where $n_{\{\bar{x},\bar{y},\bar{z}\}} = -n_{\{x,y,z\}}$ follows from the complementary image constraint $r_c = r_{\{x,y,z\}} + r_{\{\bar{x},\bar{y},\bar{z}\}}$. A similar relaxation for the constant illumination case yields $r_c = \pi \rho_D c_p V_p / 2$, where c_p is the intensity of light incident for $\omega' \in \Omega$, i.e. $P(\omega') = c_p$ (assumed equal to unity by [Ma et al. 2007] and [Wilson et al. 2010]). Substituting the modified radiance terms into (5), we observe that the ambient occlusion terms cancel. This gives the surprising result that the original method of Ma et al. is invariant to ambient occlusion. However, the correction terms (δ_x etc) do not cancel and the method is therefore still affected by deformed diffuse lobes.

The 6 linear equations of (7) form an underdetermined system in the 10 unknowns $\mathbf{x} = (\delta_x, \delta_y, \delta_z, \delta_{\bar{x}}, \delta_{\bar{y}}, \delta_{\bar{z}}, c_p, n_x, n_y, n_z)$, which we write as $\mathbf{Ax} = \mathbf{b}$ where $\mathbf{A} \in \mathbb{R}^{6 \times 10}$. However, we can regularise the problem by finding parameters which satisfy our linear system which are closest to an initial solution. For example, if we take the diffuse centroid $(n_x^{\text{Wil}}, n_y^{\text{Wil}}, n_z^{\text{Wil}})$ estimated by the method of Wilson et al. and define $\mathbf{x}_0 = (0, 0, 0, 0, 0, 0, 1, n_x^{\text{Wil}}, n_y^{\text{Wil}}, n_z^{\text{Wil}})$, then we can correct for deformed diffuse lobes by solving the quadratic programming problem: minimise $\|\mathbf{x} - \mathbf{x}_0\|^2$ subject to $\mathbf{Ax} = \mathbf{b}$.

The visual improvement in the normal maps obtained by applying this optimisation to the estimates of the method of Wilson et al. are negligible. The reason for this can be seen by substituting our radiance equation with relaxed assumptions (7) into (6), giving:

$$r_{\{x,y,z\}} - r_{\{\bar{x},\bar{y},\bar{z}\}} = \frac{\pi \rho_D V_p}{2} \left\{ \delta_{\{x,y,z\}} - \delta_{\{\bar{x},\bar{y},\bar{z}\}} + \frac{2}{3} n_{\{x,y,z\}} \right\}. \tag{8}$$

The interesting observation here is that when the deformation to the reflectance lobe is approximately symmetric, i.e. $\delta_x \approx \delta_{\bar{x}}$, then the deformation parameters cancel. In other words, symmetric deformations in the reflection lobe are averaged out and $\mathbf{x}_0$ is therefore already close to satisfying the linear system of equations. With this observation to hand, it is possible to derive a minimal four image solution in which symmetric deformations to the reflectance lobes still cancel. To do so, we exploit the X complementary image constraint and rewrite (6) as:

$$\mathbf{n}_{(\mathbf{x},\mathbf{y},\mathbf{z},\bar{\mathbf{x}})} = \frac{[r_x - r_{\bar{x}}, 2r_y - (r_x + r_{\bar{x}}), 2r_z - (r_x + r_{\bar{x}})]^T}{\|[r_x - r_{\bar{x}}, 2r_y - (r_x + r_{\bar{x}}), 2r_z - (r_x + r_{\bar{x}})]^T\|} \tag{9}$$

Similarly, Y and Z base complement pairs can also be used to obtain $\mathbf{n}_{(\mathbf{x},\mathbf{y},\mathbf{z},\bar{\mathbf{y}})}$ and $\mathbf{n}_{(\mathbf{x},\mathbf{y},\mathbf{z},\bar{\mathbf{z}})}$. This combines the advantage of the original method of Ma et al. (reduced data capture requirement) with that of Wilson et al. (improved robustness). The intuition behind why this approach is an improvement over that of Ma et al. is that information obtained in the X and X complement gradient conditions is exploited in the estimate of the Y and Z components whereas the terms in the method of Ma et al. are independent.

4 Results and Discussion

Our results use a light stage with only 41 LEDs (a 74% reduction over [Ma et al. 2007] and [Wilson et al. 2010]). In this case the lighting environment is highly discretised and the performance of the method of Ma et al. degrades dramatically. In Figure 1(b) we plot the z-component of the surface normals estimated from a horizontal slice through a white cylinder. The result of Ma et al. is significantly distorted and noisy, whilst the results of Wilson et al. and our 4 image solution are both very close to ground truth. In Figure 1(a), we show results for a face. The mean angular deviation of $\mathbf{n}_{(\mathbf{x},\mathbf{y},\mathbf{z},\{\bar{\mathbf{x}},\bar{\mathbf{y}},\bar{\mathbf{z}}\})}$ with the normal map estimates from Wilson et al. and Ma et al. was $(5.2°, 4.6°, 7.3°)$ and $(42.9°, 42.4°, 46.6°)$ respectively. Further results are shown in our accompanying video.

We have presented a quadratic programming approach which estimates parameters describing the distortion in the diffuse reflectance lobes measured by spherical gradient illumination photometric stereo. However, the improvement in estimated normal maps over the method of Wilson et al. is small. We used our analysis to propose a method in which we use a minimal set of four images with very little degradation in the quality of the recovered normals. Reducing image capture requirements is extremely important if such approaches are to be used for realtime performance capture.

Acknowledgments: We acknowledge EURECA Project for funding A. Dutta. EURECA is funded by the Erasmus Mundus External Cooperation Window (EMECW) of the European Commission.

References

MA, W. C., HAWKINS, T., PEERS, P., CHABERT, C. F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *Eurographics Symposium on Rendering*, 183–194.

WILSON, C. A., GHOSH, A., PEERS, P., CHIANG, J.-Y., BUSCH, J., AND DEBEVEC, P. 2010. Temporal upsampling of performance geometry using photometric alignment. *ACM Trans. Graph. 29*, 2, 1–11.