Photometric Stereo with Nearby Planar Distributed Illuminants

James J. Clark
Centre for Intelligent Machines
McGill University

Abstract

This paper considers the problem of shape-from-shading using nearby planar distributed illuminants. It is shown that a rectangular planar nearby distributed uniform isotropic illuminant shining on a small Lambertian surface patch is equivalent to a single isotropic point light source at infinity. A closed-form solution is given for the equivalent point light source direction in terms of the illuminant corner locations. Equivalent point light sources can be obtained for multiple rectangular illuminants allowing standard photometric stereo algorithms to be used. An extension is given to the case of a rectangular planar illuminant with arbitrary radiance distribution. It is shown that a Walsh function approximation to the arbitrary illuminant distribution leads to an efficient computation of the equivalent point light source directions. A search technique employing a solution consistency measure is presented to handle the case of unknown depths. Applications of the theory presented in this paper include visual user-interfaces using shape-from-shading algorithms making use of the illumination from computer monitors, or movie screens.

1. Introduction

Methods for the extraction of surface shape information from images of the surface have been the focus of much research since the foundational work of Horn and coworkers [4]. Much of this research has concentrated on situations involving point light sources at large distances from the surfaces in question (i.e. the point-light-source-at-infinity model). There have been a number of studies, however, on the shape-from-shading problem with nearby distributed light sources, beginning with the work of Ikeuchi [5]. Ikeuchi used a nearby planar light source, and used it to illuminate specular surfaces. The use of distributed light sources has a number of advantages over point light sources, the main one being that the radiance of the illuminant need not be as high. As noted by Schechner *et al* [10], a problem with practical shape-from-shading systems using point light sources is that the illuminant might not be bright enough to provide adequate signal-to-noise ratios in the cameras, except perhaps at specularities. Schechner *et al* suggest using controllable distributed light sources, and integrate multiple images acquired under differing patterns of illumination. Distributed light sources are also easier to construct than point light sources, as they do not have as serious problems with heat dissipation. Distributed light sources do have issues with uniformity, but recent advances in lcd projection technology have alleviated this concern somewhat.

The use of nearby light sources also carries advantages over distant illuminants. First, the use of nearby light sources permits the imaging setups to be relatively compact, an important issue in many applications. In addition, as pointed out in work by Clark [1], Iwahori *et al* [6, 7, 8], and Kim and Burger, the shading induced by nearby illuminants is dependent on the distance between the illuminant and the surface. While this dependence is usually nonlinear and complicates the shading equations, it also provides the possibility of extracting absolute depth information from the shading. In [2] Clark and Pekau showed that using a distributed nearby illuminant can improve the robustness of a differential photometric stereo approach to obtaining depth from shading.

In this paper we consider further the problem of determining the shape of a surface from multiple images obtained under illumination from nearby planar distributed light sources. Specifically, we show that, for a small Lambertian patch of known position relative to a rectangular illuminant, the illuminant can be replaced by an equivalent point light source at infinity, and closed form equations are given for the equivalent light source direction vector and radiance. Given such equivalent point light sources at infinity standard photometric techniques such as 3-image photometric stereo can be employed to compute the surface patch's surface normal.

2. Shape-from-shading with Rectangular Planar Illuminants

We will consider the special case of planar rectangular distributed illuminants. Such illuminants are readily available, in the form of computer video monitors, TV displays,

Proceedings of the 3rd Canadian Conference on Computer and Robot Vision (CRV'06)
0-7695-2542-3/06 $20.00 © 2006 IEEE

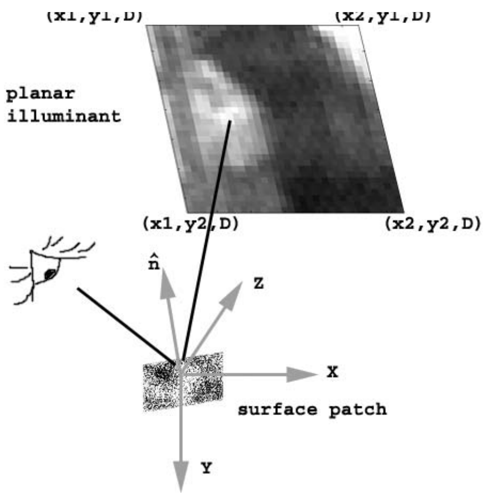

Figure 1. Geometry of the planar distributed illumination scheme.

and LCD projectors (when projected onto a planar surface, such as a wall). Such illuminants were used by Schechner *et al*. Using a computer monitor as an illumination source may be useful in visual user interface applications, where the operating system of the computer obtains information about the user of the computer via visual means.

As depicted in figure 1, let us take as our illuminant a rectangular planar segment, oriented perpendicular to the z-axis of the world coordinate system, and with corners located at $(x_1,y_1,D),(x_2,y_1,D),(x_1,y_2,D),(x_2,y_2,D)$. If we ignore shadowing effects and assume that the position of the surface patch is known (set to the origin of the world coordinate system, $X=0,Y=0,Z=0$), then the reflected light from a Lambertian surface patch with albedo ρ and unit surface normal $\hat{n}=\frac{(p,q,-1)}{\sqrt{1+p^2+q^2}}$ illuminated by a uniform isotropic rectangular illuminant with unit radiance is:

$$I(p,q,\rho;x_1,x_2,y_1,y_2,D) = \int_{y_1}^{y_2}\int_{x_1}^{x_2} \frac{\rho(px+qy-D)}{\sqrt{1+p^2+q^2}\sqrt{(x^2+y^2+D^2)^3}}dxdy \qquad (1)$$

Evaluating the inner integral gives us:

$$I(p,q,\rho;x_1,x2,y_1,y_2) = \frac{\rho}{\sqrt{1+p^2+q^2}}\int_{y_1}^{y_2}(qy-D) \left(\frac{x_2}{(D^2+y^2)\sqrt{D^2+y^2+x_2^2}} - \frac{x_1}{(D^2+y^2)\sqrt{D^2+y^2+x_1^2}}\right)dy + \frac{\rho}{\sqrt{1+p^2+q^2}}\int_{y_1}^{y_2} p\left(\frac{1}{\sqrt{D^2+y^2+x_1^2}} - \frac{1}{\sqrt{D^2+y^2+x_2^2}}\right)dy \qquad (2)$$

These integrals can be evaluated in closed form, but are quite complicated. In particular, it can be seen that the resulting equations are nonlinear with respect to the variables x_1,x_2,y_1,y_2,D. If we assume that these variables are known, the equation in terms of the remaining unknowns is quite simple.

$$I(p,q,\rho;D,x_1,x_2,y_1,y_2) = \frac{\rho[qF_1+F_2+pF_3]}{\sqrt{1+p^2+q^2}} \qquad (3)$$

where

$$F_1 = \log\left(\frac{(x_1+\sqrt{D^2+y_2^2+x_1^2})(x_2+\sqrt{D^2+y_1^2+x_2^2})}{(x_1+\sqrt{D^2+y_1^2+x_1^2})(x_2+\sqrt{D^2+y_2^2+x_2^2})}\right) \qquad (4)$$

$$F_2 = \tan^{-1}\left(\frac{x_1y_2}{D\sqrt{D^2+y_2^2+x_1^2}} - \frac{x_1y_1}{D\sqrt{D^2+y_1^2+x_1^2}}\right) - \tan^{-1}\left(\frac{x_2y_2}{D\sqrt{D^2+y_2^2+x_2^2}} - \frac{x_2y_1}{D\sqrt{D^2+y_1^2+x_2^2}}\right) \qquad (5)$$

$$F_3 = \log\left(\frac{(y_1+\sqrt{D^2+y_1^2+x_2^2})(y_2+\sqrt{D^2+y_2^2+x_1^2})}{(y_1+\sqrt{D^2+y_1^2+x_1^2})(y_2+\sqrt{D^2+y_2^2+x_2^2})}\right) \qquad (6)$$

Equation (3) can be interpreted as expressing the intensity of reflected light obtained from a single isotropic point light source at infinity, with light source direction vector:

$$\hat{s} = \frac{(F_3,F_1,-F_2)}{\sqrt{F_1^2+F_2^2+F_3^2}} \qquad (7)$$

and illuminant radiance:

$$R = \sqrt{F_1^2+F_2^2+F_3^2} \qquad (8)$$

We can use this equivalent point-light-source at infinity in standard shape-from-shading techniques. In particular, if we have three different illuminant patterns, we obtain three different equivalent point light sources, and if their direction vectors are not co-planar, we can use the images generated by them in a 3-image photometric stereo algorithm. This algorithm solves a system of 2 linear equations in p and q [12]. Woodham points out (first derived in [11]) that any spatial distribution of distant illuminants can be replaced by a single distant point illuminant and still produce the same reflectance map (assuming no part of the surface is shadowed for any portion of the illuminants). We see that this extends to the case of nearby illuminant distributions as well, although the particular equivalent point light source direction vector depends on the location of the surface patch (in the case of distant illuminants the equivalent point light source direction vector is independent of the surface patch location).

For completeness, we will present here the solution process for the 3-image photometric stereo problem, which is quite straightforward. We begin by eliminating ρ:

$$\frac{\rho}{\sqrt{1+p^2+q^2}} = \frac{I_0}{(qF_1^0 + F_2^0 + pF_3^0)} \tag{9}$$

$$\rho = \frac{I_0\sqrt{1+p^2+q^2}}{(qF_1^0 + F_2^0 + pF_3^0)} \tag{10}$$

where the superscript 0 indicates a specific configuration of the illuminant. With this expression for ρ we can express the image irradiance for a general configuration as:

$$I_i = I_0 \frac{(qF_1^i + F_2^i + pF_3^i)}{(qF_1^0 + F_2^0 + pF_3^0)} \tag{11}$$

which can be rewritten as:

$$I_i(qF_1^0 + F_2^0 + pF_3^0) = I_0(qF_1^i + F_2^i + pF_3^i) \tag{12}$$

or

$$p(I_0F_3^i - I_iF_3^0) + q(I_0F_1^i - I_iF_1^0) = I_iF_2^0 - I_0F_2^i \tag{13}$$

If we take measurements using two other illuminant configurations, l_1, l_2 we get a system of two linear equations for p, q:

$$\begin{pmatrix} (I_0F_3^1 - I_1F_3^0) & (I_0F_1^1 - I_1F_1^0) \\ (I_0F_3^2 - I_2F_3^0) & (I_0F_1^2 - I_2F_1^0) \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix} = \begin{pmatrix} I_1F_2^0 - I_0F_2^1 \\ I_2F_2^0 - I_0F_2^2 \end{pmatrix} \tag{14}$$

So, given three measurements with different illumination patterns, we can compute the slant, tilt, and albedo of a surface patch with known 3-D position by simply solving a system of two linear equations. If the illumination patterns are known beforehand, the F_j^i's can be precomputed. Thus the solution process can be very fast. If we have more than 3 illuminant configurations, and their associated images, we can use a Least-Squares approach to solve for p and q.

3. Arbitrary Illuminants

The technique that we have proposed involves illuminating a scene with three different constant radiance planar rectangular illuminant patterns. To obtain robustness to noise we can extend the algorithm to use more than 3 images, via a least-squares solution process. This process will take some time, as the illumination patterns are presented sequentially.

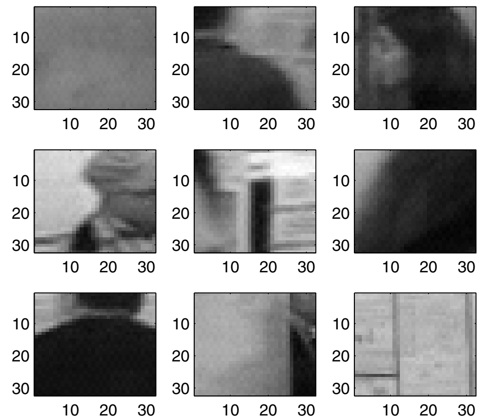

Figure 2. A sampling of the frames in the example video sequence.

It could be useful if the presentation of the illumination patterns also served other purposes at the same time. For example, the planar illuminant could be showing a movie that people are watching. The light cast on the audience could then be used by our algorithm to produce depth maps or surface normal maps of the audience, perhaps for use in a visual user interface.

The question is then raised as to whether it is possible to adapt our algorithm for arbitrary illuminants. In general the integrals involved in computing the reflected light from the arbitrary distributed illuminants are not expressible in closed form, thereby complicating the solution process. If the illuminant, however, consists of an $N \times M$ element array of rectangular planar segments, as is the case for a computer monitor or a lcd video projection, we can represent the resulting light field reflecting from the surface patch as

$$I_k = \sum_{i=1}^{N}\sum_{j=1}^{M} L_{ij}^k \frac{\rho[qF_1^{ij} + F_2^{ij} + pF_3^{ij}]}{\sqrt{1+p^2+q^2}} \tag{15}$$

where L_{ij}^k represents the radiance of the i, jth illuminant patch for illuminant configuration k.

This representation of an arbitrary illuminant is somewhat unappealing, since to compute its equivalent point-light-source-at-infinity direction requires computing the F_is for every illuminant patch. It would be nice if there was a way to obtain these directions without this high computational load. Figure 3 shows the the distribution of equivalent

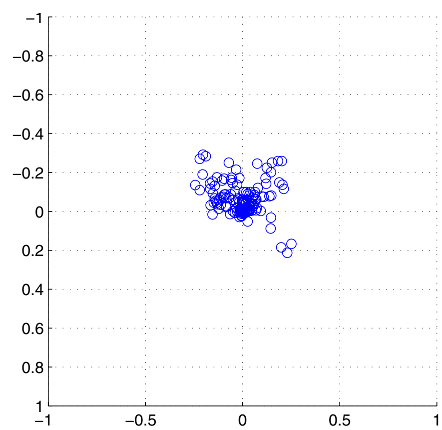

Figure 3. The distribution of equivalent point light source directions (projected onto the illuminant plane) for the projected video sequence, assuming $D = 1, x_1 = -1, x_2 = +1, y_1 = -1, y_2 = +1$.

light source directions (projected onto the illuminant plane) for a test video sequence of 152 different 32x32 pixel images used as illuminant patterns projected onto a rectangular screen with $D = 1, x_1 = -1, x_2 = +1, y_1 = -1, y_2 = +1$ (a sampling of the images in the video sequence is shown in figure 2). Inspection of this distribution shows that most of the directions are clustered about the Z-axis (i.e. perpendicular to the illuminant plane). This is due to the global nature of the computation of the equivalent light source directions, which tends to average out contributions from different parts of the illuminant array. This suggests that we could get a good approximation to the equivalent light source directions just by looking at low-spatial-frequency aspects of the illuminant field. Now, there are many spatial-frequency domain representations of the illuminant field that could be used for this purpose, but the most useful in the present context is the so-called *Walsh* function representation [3]. The Walsh functions form an orthogonal basis for square-integrable functions and can therefore represent arbitrary illumination patterns:

$$L_k(x,y) = \sum_{i=1}^{N}\sum_{j=1}^{M} c_{i,j}^{k} W^{(i,j)}(x,y) \tag{16}$$

The use of Walsh functions to describe extended light source distributions was proposed by Schechner *et al* [10]. Most importantly for the purposes of this paper, the Walsh functions are all piece-wise constant, with values of either -1 or +1. As such, a single Walsh function as an illumination pattern will result in a reflected light intensity that is a sum of terms having the form given in equation (3). Thus, the reflected light from an arbitrary illuminant can likewise be expressed as a weighted sum of terms of the form (3), where the weights depend on the coefficients of the Walsh transform of the arbitrary illuminant pattern.

$$I_k = \sum_{i=1}^{N}\sum_{j=1}^{M} c_{ij}^{k} \frac{\rho[qF_1^{ij} + F_2^{ij} + pF_3^{ij}]}{\sqrt{1+p^2+q^2}} \tag{17}$$

where the c_{ij}^k are the coefficients of the Walsh basis function representation of the arbitrary illuminant distribution $l_k(x,y)$ and the F_k^{ij} are the F_k for the illuminant corresponding to the ij^{th} Walsh function.

As we have seen earlier each pattern of distributed illumination can be replaced by a single distant illuminant. This means that for every Walsh function there is a corresponding distant illuminant. The arbitrary planar illuminant can therefore be replaced with an equivalent single distant illuminant that is the vector sum of the individual Walsh equivalent point light source directions, scaled by the equivalent radiance, weighted by the coefficient of the corresponding Walsh function in the Walsh transform decomposition of the

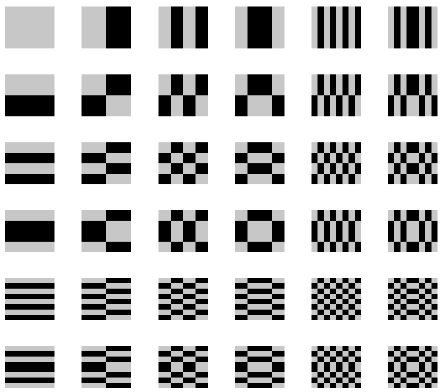

Figure 4. The Walsh functions of order 6 and less.

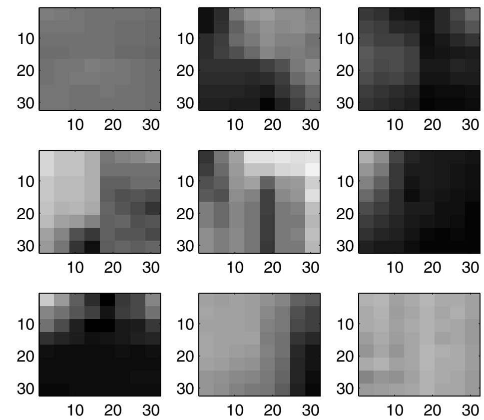

Figure 5. The approximation of the video frames shown in figure 2 using Walsh functions of order 6 and lower.

arbitrary illuminant:

$$\vec{s}_k = \sum_{i=1}^{N} \sum_{j=1}^{M} c_{ij}^k \vec{s}^{ij} \tag{18}$$

where

$$\vec{s}^{ij} = (F_3^{ij}, F_1^{ij}, -F_2^{ij}) \tag{19}$$

Figure 4 shows the Walsh functions of order 6 and less. In figure 5 we show the approximation of the video frames shown in figure 2, using only the Walsh functions of order 6 and less. It is expected that a good approximation to the equivalent light source direction for an arbitrary illuminant field can be obtained by a superposition of the Walsh functions of order 6 or lower, multiplied by the corresponding coefficient of the Walsh transform of the illuminant field. The equivalent light source directions for each of the included Walsh functions can be pre-computed, so the only significant computational expenses is incurred by the computation of the Walsh transform of the arbitrary illuminant field. But this cost is not very high since we only need to compute the first few coefficients, and not the complete Walsh transform, and the computation of the transform only involves addition operations.

In figure 6 we show the equivalent light source directions for the same video sequence as used for figure 3, save that we have used the truncated Walsh function approximation. It can be seen that they are quite similar to those obtained by

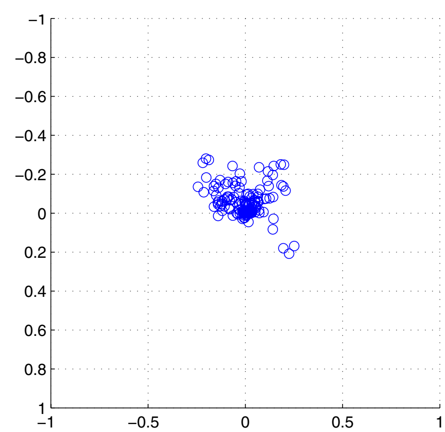

Figure 6. The distribution of equivalent point light source directions for the projected video sequence, as computed using the truncated Walsh function approximation of the video images.

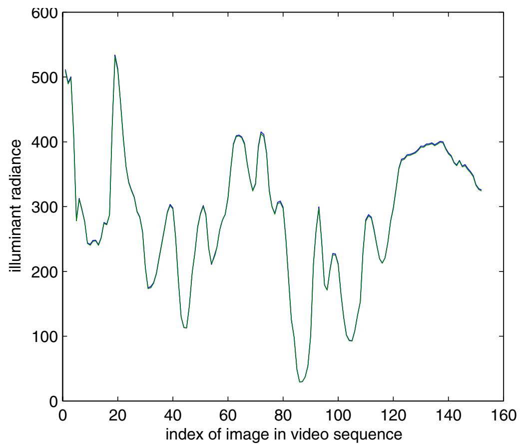

Figure 7. The estimated radiance for the equivalent point-light-source illuminants for both the original and Walsh function approximations of the video sequence images.

looking at each pixel of the video illuminant. The standard deviation of the difference in the actual versus approximate equivalent light source unit direction vectors over the 152 video frames is $(37.2, 34.0, 8.93) \times 10^{-4}$.

In figure 7, we plot the estimated image radiances for the equivalent point light sources corresponding to both the original video sequence images, and to the Walsh function approximations. The two plots almost completely overlap each other, indicating that the coarse Walsh function approximation of the video images does not significantly affect the global illumination characteristics.

4. Handling Unknown Depths

The technique described above requires that the location of the surface patch be known. If the camera viewing the patch is far from the patch, with its image plane aligned with the X-Y world coordinate plane, then the X and Y coordinates of the patch can usually be determined from the image, assuming proper calibration. The Z coordinate of the patch is more difficult to compute from the image, however. Treating D as an unknown and trying to solve for it from the equations is out of the question due to the nonlinearity. But we can try to do a search for the proper value of D using an optimization criteria based on consistency between solutions.

One approach to consistency testing is to take a number of illuminant sets consisting of three illuminant patterns each, and for each of these sets measure the reflected light from the surface patch. Then, using the equations developed earlier compute the equivalent point light source directions and radiances for each of the illuminants, under the assumption of a given value for D. With this information we can use the 3-image photometric stereo technique to obtain a solution for p and q. For each value of D in our search space, we can compute the variance of the derived values of p and q taken over the different illuminant sets. We then ascribe the optimum D value to that which produces the smallest p and q variance.

For the consistency testing approach to be practical, we must be able to limit the search space for D. In many applications, such as in the case of a computer monitor illuminating a person sitting in front of it, the depth of the surface patch may be known to lie within a certain range. We can then constrain our search for the optimum depth value to this limited range. To test this approach we carried out a simulation of the situation where a small surface patch with $p = 0, q = 0$ was illuminated by a square planar illuminant with $D = 1, x_1 = y_1 = -1, x_2 = y_2 = +1$. The illuminant patterns were taken from the 152-image video sequence described earlier. The measurements were simulated using equations (3) (and using equations (4,5,6) to compute the F's), with added zero mean Gaussian noise. Three different measurement noise variances were tested: 0.1, 0.25, and 0.5. The average noise free measurement value over all of the images was 284, so the signal to noise ratios for the simulations were quite high.

The resulting variances of the derived quantities p and q are shown in figure 8 as a function of the assumed D values, where the search interval for the D values was from 0.9 to 1.1.

For a measurement noise variance of 0.1 the minimum variance in the computed p and q values occurs at $D = 1$, for which the mean values of p and q are $p = -0.035$ and $q = 0.0023$. At higher measurement noise levels the minimum computed variance point can shift away from $D = 1$. For example, at a measurement noise variance of 1, the optimal D is 1.1 (at one extreme of the search range). So it would appear that this method for finding D does not work well for noisy measurements. It must be kept in mind, however, that only 3 images were used in the estimation of each p and q value. If a least squares solution was used involving a larger number of images, it is expected that the noise sensitivity would decrease.

5. Discussion

In this paper we have looked at the possibility of estimating surface shape from images of the surface illuminated by nearby rectangular planar illuminants. We showed that for a

Proceedings of the 3rd Canadian Conference on Computer and Robot Vision (CRV'06)

surface patch at a known location, the illuminant can be replaced by an equivalent point light source at infinity, and provided close form expressions for the equivalent light source direction vectors and radiances. This equivalence permits the application of standard shape-from-shading techniques that assume point light sources at infinity. In particular, if we have at least three illuminant patterns, with non-coplanar equivalent point light source directions we can perform photometric stereo.

We also considered the case in which the illumination field on the planar rectangular illuminant was not uniform. This case can be handled by simply decomposing the illuminant field into an array of small rectangular sub-regions (or pixels). This approach has the drawback of requiring the computation of the equivalent point light source for each of the illuminant pixels, which can be quite numerous. To address this problem we used the intuition that the equivalent point light source direction is dominated by the low spatial frequency characteristics of the illuminant field. Thus, a low-dimensional approximation in terms of spatial basis functions may suffice. For this purpose the Walsh function representation fits very well. The Walsh functions are piecewise constant, so that we can directly apply our closed form equations for the equivalent point light source directions for each of them, and a good approximation of the equivalent point light source for an arbitrary illuminant field can be had just by summing the point light source vectors for each Walsh function, weighted by the respective coefficients of the Walsh transform of the arbitrary illuminant.

There are a number of issues that remain to be properly addressed before this technique will be practically applicable. In the above analysis we have assumed that the surface patch is centered on the origin of the illuminant pattern (e.g. $X = Y = 0$). We can extend the analysis to non-zero values of X, Y by shifting the origin of the illuminant pattern. This will necessitate a re-calculation of the Walsh transform of the illuminant patterns, thereby increasing the amount of computation required. It may be possible to pre-compute much of this.

Perhaps the most significant problem with the proposed approach, as in all methods that employ distributed illuminants is that of shadowing. In the algorithm presented above, we assumed that the extent of the distributed illuminant is such that no part of the illuminant is self-shadowed by the planar patch. Clearly, this restricts both the size of the illuminant that we can use, as well as the range of patch tilts that we can measure. The problem becomes worse once we allow our surface patch to become part of an extended surface, as remote parts of the surface can possible cause shadowing of portions of the illuminant.

We could handle the shadowing problem with an active approach, whereby a limited extent illuminant is *scanned* across the illuminant plane. This approach has two main drawbacks - increased measurement time due to the scan-

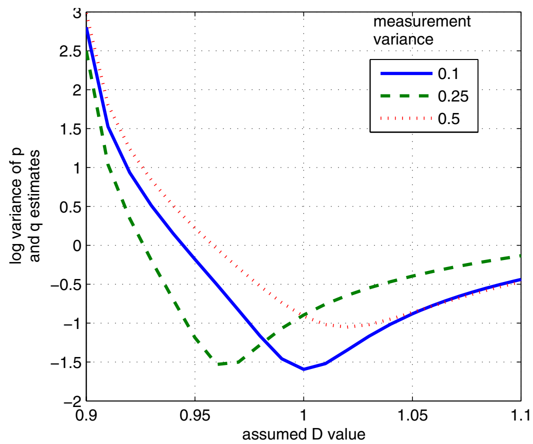

Figure 8. The log of the variance of the computed p, q values for a range of hypothesized D values (true value of D is 1). The different curves correspond to different noise levels in the simulated intensity measurements.

ning, and reduced signal to noise ratio due to the reduction in the illuminant's spatial extent.

Another possible remedy is to use the discontinuous nature of the shadowing process. We could make small changes to the illuminant distribution and look for discontinuous jumps in the light reflected from the surface as indications of the presence of shadowing. The magnitude of the jumps could perhaps be used to compensate for the effect of the shadowing.

Another issue is the nature of the surface's reflectance model. In this paper we assume a Lambertian model. Much of the theory applies to non-Lambertian models, however. The details will vary, and in most cases the equations will be nonlinear in p and q, making the solution difficult. For mirror-like specular surfaces the equations may be quite simple, as shown by Ikeuchi [5] in his work employing a planar light source to illuminate a specular surface.

Perhaps the most intriguing potential application of this theory lies in obtaining shape information of people's faces as they view a computer monitor, a television show, or movie screen. This shape information could be used, for example, to recognize a viewer's identity, for purposes of authentication, or providing viewing material tailored to the actual audience.

References

[1] Clark, J.J., "Active photometric stereo", Proceedings of the Computer Vision and Pattern Recognition conference, Champaign, Illinois, pp 29-34, 1992.

[2] Clark, J.J. and Pekau, H., "An integral formulation of differential photometric stereo", Proceedings of the 1999 IEEE Computer Vision and Pattern Recognition Conference, Fort Collins, Colorado, June 1999, pp 119-124.

[3] Harmuth, H.F. "Applications of Walsh Functions in Communications." IEEE Spectrum, Vol. 6, pp 82-91, 1969.

[4] Horn, B.K.P., "Obtaining shape from shading information", Chapter 4 in *The Psychology of Computer Vision,* P.H. Winston, ed., McGraw-Hill: New York, 1975.

[5] Ikeuchi, K., "Determining surface orientations of specular surfaces by using the photometric stereo method", IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. 3(6), pp 661–669, 1981.

[6] Iwahori, Y., Sugie, H., and Ishii, N., "Reconstructing shape from shading images under point light source illumination", Proceedings of the 10th International Conference on Pattern Recognition, Atlantic City, pp 83-87, 1990.

[7] Iwahori, Y., Woodham, R.J., and Ishii, N., "Shape from shading with a nearby moving point light source", Proceedings of the 2nd International Conference on Automation, Robotics and Computer Vision, pp CV.5.5.1-5.5.5, Singapore, 1992.

[8] Iwahori, Y., Woodham, R.J., Tanaka, H., and Ishii, N., "Moving point light source photometric stereo", IEICE Transactions on Information and Systems, Vol. E77-D, No. 8, pp 925-929, 1994.

[9] Kim, B. and Burger, P., "Depth and shape from shading using the photometric stereo method," Comput. Vision Graph. Image Proc. 54, 416-427 (1991).

[10] Schechner, Y.Y., Nayar, S.K, and Belhumeur, P.N., "A Theory of Multiplexed Illumination", Proc. 2003 ICCV, Nice, France, pp 808-813, 2003.

[11] Silver, W.M., "Determining shape and reflectance using multiple images", SM thesis, MIT Dept. Electrical Engineering and Computer Science, Cambridge, MA, 1980.

[12] Woodham, R.J., "Photometric method for determining surface orientation from multiple images", Optical Engineering, Vol. 19, No. 1, pp 139-144, 1980.

[13] Woodham, R.J., "Analysing images of curved surfaces", Artificial Intelligence, vol. 17, pp. 117-140, 1981.

[14] Woodham, R.J., "Gradient and curvature from photometric stereo including local confidence estimation", Journal of the Optical Society of America A, Vol. 11, pp 3050-3068, 1994.

Proceedings of the 3rd Canadian Conference on Computer and Robot Vision (CRV'06)
0-7695-2542-3/06 $20.00 © 2006 IEEE