Image and Vision Computing xxx (2008) xxx–xxx

Contents lists available at ScienceDirect

Image and Vision Computing

journal homepage: www.elsevier.com/locate/imavis

Photometric stereo using LCD displays

James J. Clark *

Centre for Intelligent Machines, McGill University, 3480 University Street, Montreal, Que., Canada H3A2A7

ARTICLE INFO

Article history:
Received 4 January 2007
Received in revised form 17 October 2008
Accepted 27 October 2008
Available online xxxx

Keywords:
Shape-from-shading
Photometric stereo
Distributed illuminant

ABSTRACT

This paper considers the problem of shape-from-shading using nearby extended light sources. The paper reviews a number of methods that employ nearby illuminants, and describes a new technique that assumes a rectangular planar nearby distributed uniform isotropic illuminant. It is shown that such a light source illuminating a small Lambertian surface patch is equivalent to a single isotropic point light source at infinity, in the absence of shadowing. 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 distinct illuminant patterns allowing standard photometric stereo algorithms to be used. An extension is given to the case of a rectangular planar illuminant with arbitrary radiance distribution. Experimental results are shown demonstrating the application of the theory to photometric stereo using illumination from a LCD computer monitor. Details on the photometric calibration of the illumination source and image acquisition device are provided.

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 [7]. 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 are two significant paths by which the classical point-source-at-infinity methods can be extended. The first is to bring the light source close to the object being illuminated, while the other is to distribute the light source across an extended area of space.

Using nearby light sources provides a number of advantages, but gives rise to serious difficulties as well. The primary advantages of using nearby light sources are the possibility of obtaining absolute depth information and the increased flexibility of illuminant positioning. The use of nearby light sources permits the imaging setups to be relatively compact, an important issue in many applications. Clark [2], Iwahori et al. [10–12] and Kim and Burger [13] point out that the shading induced by nearby illuminants is dependent on the distance between the illuminant and the surface. While this dependence is usually non-linear and complicates the shading equations, it also provides the possibility of extracting absolute depth information from the shading. The inherent difficulties with nearby illuminants are many, however. They include increased shadowing effects (since the range of incident angles is typically larger than for sources at infinity), increased dynamic range (due to the intensity variation induced by depth variations and the $1/r^2$ falloff with distance), and the non-linearity of the resulting shading equations.

The use of distributed light sources has a number of inherent advantages, and attendent difficulties, over point illuminants. Primary among the positive aspects of distributed illuminants is that they require less intense light sources with reduced power density. This lowers illuminant temperature and simplifies construction. As noted by Schechner et al. [15], 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 integrating multiple images acquired under differing patterns of illumination. As shown by Clark and Pekau [3], distributed light sources can be used to alleviate the noise-sensitivity of some shape-from-shading algorithms by transforming gradient-based algorithms into integral-based algorithms. On the down side, distributed light sources have issues with lack of uniformity and isotropy. Recent advances in LCD projection technology have alleviated this concern somewhat.

2. Extending photometric stereo

Following in the footsteps of Horn's pioneering work on extracting surface gradient and height information from image irradiance [7], Woodham [18] developed a method, which he termed *photometric stereo*, for obtaining surface gradient information from a series of images acquired with the light source in different positions. Woodham showed that if images of a planar Lambertian surface

* Tel.: +1 514 398 2654; fax: +1 514 398 7348.
E-mail address: clark@cim.mcgill.ca

0262-8856/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.imavis.2008.10.011

patch are acquired under three different illumination conditions consisting of point source illuminants at infinity, the slant and tilt can be obtained by solving a simple linear system of equations. Woodham points out (first derived by Silver [16]) 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). This suggests an extension to the classical photometric stereo which employs distributed illuminants. In the remainder of this section of the paper we will review a number of techniques that have been proposed which extend or generalize photometric stereo, leading step-by-step towards our approach involving nearby distributed light sources.

One of the first modifications of photometric stereo involved a differential form of it. Wolff [17] examined the optical flow induced by the motion of a point light source at infinity, and provided an algorithm for extracting the slant and tilt of surfaces from the flow images. Wolff's algorithm can be considered as a differential version of photometric stereo, wherein the different light source positions are very close together. For Lambertian surfaces, Wolff's algorithm requires solving a pair of quadratic equations. Neither Woodham's photometric stereo or Wolff's differential method provided absolute shape information, and both algorithms required solution of systems of highly non-linear equations when faced with non-Lambertian surfaces.

The use of a nearby illuminant yields the possibility to compute absolute depth, since the scene radiance is then a function of the distance of the illuminant to the scene point. Perhaps the earliest shape-from-shading method to make use of nearby illuminants was one proposed by Ikeuchi [8]. He considered the illumination of specular surfaces with a nearby planar light source. Ikeuchi's method was not a photometric stereo method, however, since only one light source position was used. In [10], Iwahori et al. presented a photometric stereo technique that did provide absolute depth information. This technique used a series of images of an object illuminated by a moving point source *near* to the object. The attainment of absolute depth measurement came at a cost, however. Even in the case of Lambertian reflectance their algorithm required the solution a system of non-linear equations in three unknowns – depth, surface slant and tilt, while iteratively updating the albedo value.

The nearby light source photometric stereo method of Iwahori [10] involved solutions of highly non-linear equations, which caused problems in practice. Following on from this work, Iwahori [11,12] and Clark [2] independently developed a *differential* photometric stereo method for obtaining absolute depth values from a series of images of a surface illuminated by a point source near the object. In this approach, the light source is moved in small increments along three orthogonal axes, allowing the computation of an approximate 3D gradient of the image with respect to light source position. The Iwahori–Clark method effectively combines the differential photometric stereo approach of Wolff with the nearby light source approach of Iwahori et al. The advantages of this approach over the earlier photometric shape recovery techniques is threefold: it provides absolute depth, it works for a very general class of reflectance models, and it requires solution of only a linear equation for the depth. Their theory also provides a means for computing the surface normal vector and for extracting the nominal position of the point illumination source. These methods assume a constant albedo. The albedo value, however, does not need to be known.

More recently, in a rather different approach, Magda et al. [14] presented a photometric stereo method based on Helmholtz reciprocity that also used nearby light sources. In this technique, a set of cameras (or a single moving one) is co-located with a set of nearby point light sources. Images are acquired one camera at a time with all of the illuminants associated with the other cameras turned on. Because of the simplifications permitted by the Helmholtz reciprocity principle the surface normals can be computed using a straightforward singular value decomposition. This method works for arbitrary (and unknown) forms of the surface's bidirectional reflectance distribution function (BRDF). The same paper also presents a light-field based approach which uses a set of images acquired from a fixed camera with a single point illuminant placed at various points on two concentric (hemi)spheres about the surface. This is effectively a differential photometric stereo approach, but unlike the Iwahori–Clark technique, it works for arbitrary (and unknown) forms of the surface's bidirectional reflectance distribution function. While this approach can provide a depth estimate with only two light sources, the implementation given in the paper used 143 different light source positions, presumably to deal with the sensitivity to noise of differential approaches. These different illuminant positions, however, do provide a sampling of a 2D slice of the 4D BDRF as a bonus.

The differential photometric stereo method of Iwahori and Clark requires measurement of the gradient of the image with respect to the light source position coordinates. As such, it can be very sensitive to noise. In [3], Clark and Pekau showed that the particular form of the equation for absolute depth in the Iwahori–Clark approach allows an integral formulation of the solution process to be constructed. This integral formulation results in an algorithm which does not require differentiation. The particular approach described by Clark and Pekau involved using point light sources that can be moved about the faces of, or throughout the interior of, a cubic volume, and required the computation of volume and surface integrals over the light source position space. The required integrals can be computed simply by using planar extended light sources corresponding to the faces of the cube and a volume light source corresponding to the interior of the cube.

The reformulation in terms of distributed illuminants makes the technique at least possible in principle, if not completely practical. Clark and Pekau [3] implemented an experimental system using a planar square illuminant under the control of a robotic manipulator. The robot could be commanded to move the planar illuminant to trace out a cubical light source position volume. In this case, only six images are needed to obtain the required surface integrals. Creation of volume illuminants that do not self-occlude is difficult, if not impossible. But the volume illuminant can be approximated by scanning a planar light source through the cube volume. The number of images required for computation of the volume integral depends on the smoothness of the image intensity as a function of light source position and on the particular integration method being used.

3. Shape-from-shading with rectangular planar illuminants

The Clark–Pekau approach provides a shape-from-shading technique that makes use of an extended nearby illuminant. It is, however, rather cumbersome to implement in practice due to the need to move a rectangular planar illuminant through a cubic volume in space in small increments, as well as to move and orient this planar illuminant sequentially to the six faces of a cube. The process requires a robotic manipulator under tight position control, and can take a significant amount of time due to the large number of images that must be acquired.

In [4], Clark proposed a new approach to shape-from-shading that makes use of a single planar rectangular distributed illuminant, fixed in space but with controllable illuminant distribution. Such illuminants are readily available, in the form of computer video monitors, TV displays, and LCD projectors (when projected onto a planar surface, such as a wall). Such planar controlled extended illuminants have been used to good effect previously in

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

computer vision systems, such as the high quality image acquisition system developed by Schechner et al. [15]. 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. Funk [5,6] used a LCD monitor as a controllable light source, and employed it to provide multiple point illuminants for a photometric stereo implementation. His approach used the LCD screen to display small rectangular areas of light, to approximate point light sources. The fact that the illuminant was not at infinity implies that the illuminant direction vectors will vary, and that there is a $1/r^2$ variation in the irradiance. Funk's approach handled the computation of the illuminant direction with an iterative estimation process. This consists of assuming an initial depth and illuminant direction followed by a standard photometric stereo algorithm to compute surface normals, followed by integration of the normals to provide an updated depth estimate.

The work by Funk [5,6] assumes point illuminants, but clearly a computer controlled LCD monitor or other similar planar display permits the implementation of more general extended illuminants. This paper describes in detail how to use such a display to perform photometric stereo. First, we will outline the theory underlying the use of planar illuminants in photometric stereo, as described in [4].

As depicted in Fig. 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 (and set for convenience to the origin of the world coordinate system, $X = 0, Y = 0, Z = 0$), then the radiance of a Lambertian surface patch with albedo ρ and unit surface normal $\hat{n} = (p, q, -1)/\sqrt{1+p^2+q^2}$ illuminated by a uniform isotropic rectangular illuminant with radiance R_I is

$$R_S(p,q,\rho;x_1,x_2,y_1,y_2,D) = \rho R_I \int_{y_1}^{y_2}\int_{x_1}^{x_2} \times \frac{(px+qy-D)}{\sqrt{(1+p^2+q^2)(x^2+y^2+D^2)^3}}\,dx\,dy. \quad (1)$$

Evaluating the inner integral gives us

$$R_S(p,q,\rho;x_1,x_2,y_1,y_2,D) = R_I\frac{\rho}{\sqrt{1+p^2+q^2}}\int_{y_1}^{y_2}(qy-D) \times\left(\frac{x_2}{(D^2+y^2)\sqrt{D^2+y^2+x_2^2}} - \frac{x_1}{(D^2+y^2)\sqrt{D^2+y^2+x_1^2}}\right)dy + R_I\frac{\rho}{\sqrt{1+p^2+q^2}}\int_{y_1}^{y_2} p\left(\frac{1}{\sqrt{D^2+y^2+x_1^2}} - \frac{1}{\sqrt{D^2+y^2+x_2^2}}\right)dy. \quad (2)$$

These integrals can be evaluated in closed form, but are quite complicated. In particular, it can be seen that the resulting equations are non-linear 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.

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

where

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

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

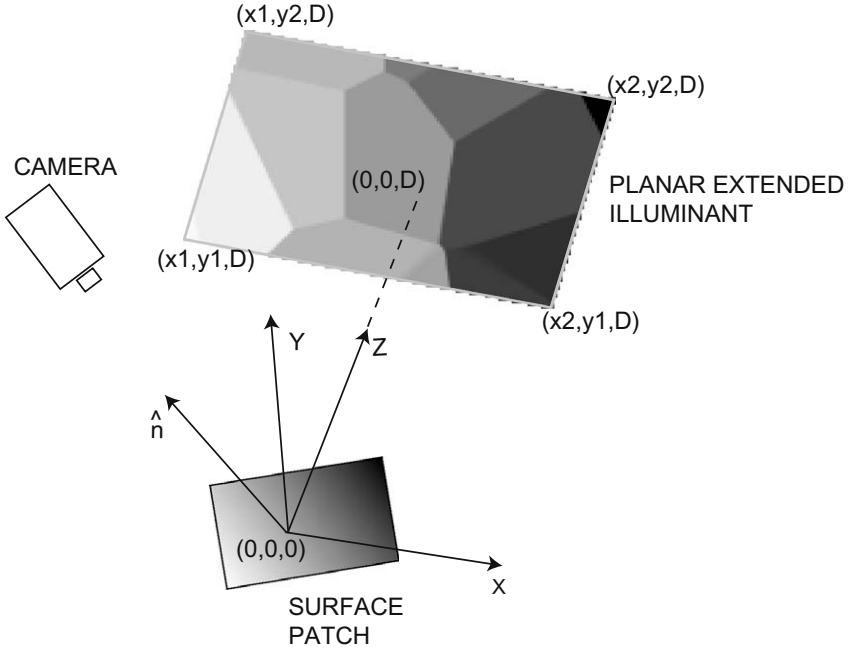

Fig. 1. Geometry of the planar distributed illumination scheme.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

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

Eq. (6) can be interpreted as expressing the intensity of reflected light obtained from a single isotropic point light source at infinity,

$$R_s = R_I^E \rho \hat{s}^E \cdot \hat{n}, \quad (7)$$

where $\hat{s}^E$ is the equivalent light source direction vector

$$\hat{s}^E = \frac{(F_3, F_1, -F_2)}{\sqrt{F_1^2+F_2^2+F_3^2}} \quad (8)$$

and R_I^E is the equivalent light source radiance

$$R_I^E = R_I\sqrt{F_1^2+F_2^2+F_3^2}. \quad (9)$$

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 [18]. Woodham points out (first derived in [16]) 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. Begin by rewriting Eq. (7), combining the albedo and the unit surface normal to give

$$R_s = R_I^E \hat{s}^E \cdot \vec{n}, \quad (10)$$

where

$$\vec{n} = (n_1, n_2, n_3) = \rho\hat{n} = \rho\frac{(p, q, -1)}{\sqrt{1+p^2+q^2}}. \quad (11)$$

If we take three measurements with different light sources we can form a matrix equation

$$\begin{pmatrix} R_I^{E1}\hat{s_1}^E \\ R_I^{E2}\hat{s_2}^E \\ R_I^{E3}\hat{s_3}^E \end{pmatrix}\begin{pmatrix} n_1 \\ n_2 \\ n_3 \end{pmatrix} = \begin{pmatrix} R_I^1(F_3^1, F_1^1, -F_2^1) \\ R_I^2(F_3^2, F_1^2, -F_2^2) \\ R_I^3(F_3^3, F_1^3, -F_2^3) \end{pmatrix}\begin{pmatrix} n_1 \\ n_2 \\ n_3 \end{pmatrix} = \begin{pmatrix} R_s^1 \\ R_s^2 \\ R_s^3 \end{pmatrix}. \quad (12)$$

Given the measurements we can solve this linear system of equations to find (n_1, n_2, n_3). We can then find ρ, p and q:

$$\rho = \sqrt{n_1^2+n_2^2+n_3^2}, \quad (13)$$

$$p = -\frac{n1}{n3}, \quad (14)$$

$$q = -\frac{n2}{n3}. \quad (15)$$

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 three illuminant configurations, and their associated images, we can use a least-squares approach to solve for n_1, n_2, n_3.

Note that this solution assumes that the depth D of the surface patch is known. If it is not known then a more complicated set of non-linear equations result, which will be challenging to solve. Alternatively, a separate process could be used for estimating D for use in the above approach. For example, [4] showed that if a lot of images are available, a solution consistency technique can be used to estimate D.

3.1. Arbitrary illuminants

The technique that we have proposed involves illuminating a scene with three different isotropic planar rectangular illuminants with uniform radiance. To obtain robustness to noise we can extend the algorithm to use more than three images, via a least-squares solution process. This process will take some time, as the illumination patterns are presented sequentially. 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 (non-uniform radiance) 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, as proposed by Schechner et al. [15], 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

$$R_S = \sum_{i=1}^{N}\sum_{j=1}^{M} R_I^{ij}\frac{\rho\left(qF_1^{ij}+F_2^{ij}+pF_3^{ij}\right)}{\sqrt{1+p^2+q^2}} \quad (16)$$

or

$$R_s = \frac{\rho\left(qS([R_I^{ij}]*[F_1^{ij}])+S([R_I^{ij}]*[F_2^{ij}])+pS([R_I^{ij}]*[F_3^{ij}])\right)}{\sqrt{1+p^2+q^2}}, \quad (17)$$

where $[R_I^{ij}]$ represents the matrix of radiances of the i,jth illuminant patch (or pixel), and where the notation $S([\,]*[\,])$ indicates the sum over the element-by-element products of two matrices. We will refer to the matrices $[F_1^{ij}]$, $[F_2^{ij}]$ and $[F_3^{ij}]$ as the *F-images* for the extended illuminant. Comparing this equation to Eq. (3), we see that we can get an equivalent F_1, F_2, F_3 for a given illuminant radiance pattern matrix, $[R^{ij}]$:

$$\widetilde{F}_1 = S([R_I^{ij}]*[F_1^{ij}]), \quad (18)$$

$$\widetilde{F}_2 = S([R_I^{ij}]*[F_2^{ij}]), \quad (19)$$

$$\widetilde{F}_3 = S([R_I^{ij}]*[F_3^{ij}]). \quad (20)$$

Thus, the extended illuminant's equivalent point light source direction vector is given by

$$\tilde{s}^E = \frac{(\widetilde{F}_3, \widetilde{F}_1, -\widetilde{F}_2)}{\sqrt{\widetilde{F}_1^2+\widetilde{F}_2^2+\widetilde{F}_3^2}} \quad (21)$$

with equivalent radiance

$$\widetilde{R}_I^E = \sqrt{\widetilde{F}_1^2+\widetilde{F}_2^2+\widetilde{F}_3^2}. \quad (22)$$

For a given planar distributed illuminant, the *F*-images can be precomputed. Then, given an illuminant radiance matrix $[R^{ij}]$ to be displayed, the equivalent point light source direction can be straightforwardly computed by summing over the element by element products of the *F*-images with the display image and substituting these into Eq. (21). An example of the *F*-images for an actual display device is shown in Fig. 2.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

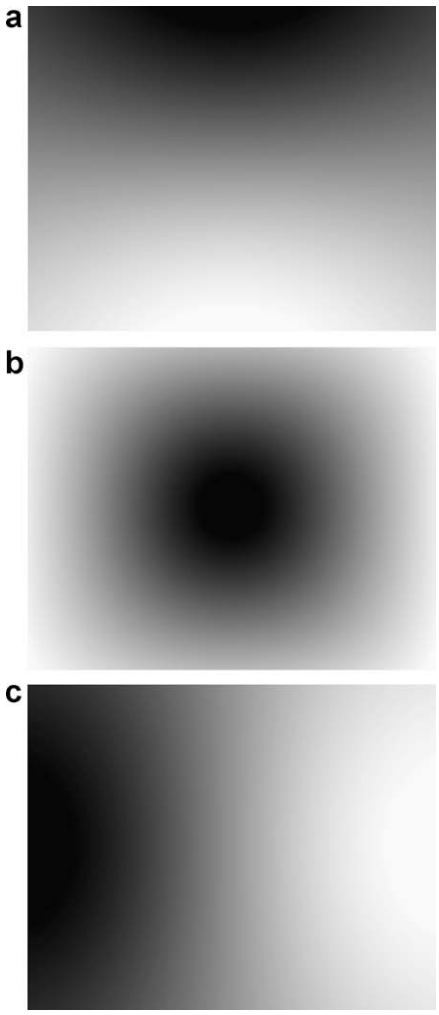

Fig. 2. The *F*-images for the Samsung 931BF monitor, for an object distance of 291 mm. (a) F1, (b) F2 and (c) F3.

3.2. Radiance, spectral radiance, and luminance

The equations given earlier involve the *radiance* of the illuminant, and of the illuminated surface. These equations, however, are strictly correct only if the illuminant spectrum and the surface albedo are independent of wavelength. To be more general, we must use the *spectral radiance* in place of the radiance in these equations. The spectral radiance expresses the radiance at different wavelengths. In particular, Eq. (16) above should be expressed in a more general form as

$$R_S(\lambda) = \sum_{i=1}^{N}\sum_{j=1}^{M} R_I^{ij}(\lambda)\frac{\rho(\lambda)\left(qF_1^{ij} + F_2^{ij} + pF_3^{ij}\right)}{\sqrt{1+p^2+q^2}}, \tag{23}$$

where the illuminant's and the surface's spectral radiance $R_k^{ij}(\lambda)$ and $R_S^{ij}(\lambda)$ are now used, and where the surface albedo is now taken to be a function of wavelength. If we take the radiance to be the integral of the spectral radiance with respect to wavelength, it can be seen that if we take the surface albedo to be independent of wavelength, and integrate both sides of the above equation with respect to λ, we recover Eq. (16). If the surface albedo is not independent of wavelength we cannot write a closed-form expression that involves radiances – we must use spectral radiances instead.

In general, the spectral radiance formulation results in a large number (potentially infinite) of unknowns as the albedo must be determined for every wavelength which affects the camera. In addition, the illuminant's spectral distribution must be known, resulting in a difficult modeling and calibration problem. To simplify matters, we can aggregate the spectral variation of the surface and illuminant radiances into a single measure, called *luminance*. Luminance is a perceptually weighted version of spectral radiance, where the weighting depends on wavelength, as specified by the *CIE photopic luminous efficiency function*, $V(\lambda)$ [21]. This function models the apparent brightness of a surface patch with unit radiance at a given wavelength, as perceived by a *standard* observer. The luminance of an illuminant or surface patch with spectral radiance $R(\lambda)$ is then given by

$$L = K\int V(\lambda)R(\lambda)d\lambda, \tag{24}$$

where K is a scale factor dependent on the units chosen for radiance and luminance. Thus, the observed luminance, L_S, of a surface patch being illuminated by an extended planar illuminant, with a spectral radiance pattern $R_I^{ij}(\lambda)$ is given by

$$L_S = \sum_{i=1}^{N}\sum_{j=1}^{M} K\int V(\lambda)R_I^{ij}(\lambda)\frac{\rho(\lambda)\left(qF_1^{ij} + F_2^{ij} + pF_3^{ij}\right)}{\sqrt{1+p^2+q^2}}d\lambda. \tag{25}$$

This formulation is convenient because the camera need only measure the surface patch luminance rather than its spectral radiance, but it still requires knowledge of the illuminant's spectral radiance. In this most general form it is not possible to express the right-hand side in terms of the illuminant's luminance. Because of this difficulty, we will not attempt to solve the most general problem, and instead just look at two special cases.

The first special case is that of *white* surfaces, in which the surface's albedo is independent of wavelength, $\rho(\lambda) = \rho$. In this case, the observed surface luminance is given by

$$L_S = \rho\sum_{i=1}^{N}\sum_{j=1}^{M} L_I^{ij}\frac{\left(qF_1^{ij} + F_2^{ij} + pF_3^{ij}\right)}{\sqrt{1+p^2+q^2}}. \tag{26}$$

Thus, we need only know the luminance of the illuminant, and need not model its spectral radiance.

The second special case arises when the illuminant is white, that is, the illuminant radiance is independent of wavelength, $R_I^{ij}(\lambda) = R_I^{ij}$. In this case, the observed illuminant luminance is given by

$$L_I^{ij} = KR_I^{ij}\int V(\lambda)d\lambda. \tag{27}$$

Thus, it can be seen that the observed surface luminance is

$$L_S = \left(\frac{\int\rho(\lambda)d\lambda}{\int V(\lambda)d\lambda}\right)\sum_{i=1}^{N}\sum_{j=1}^{M} L_I^{ij}\frac{\left(qF_1^{ij} + F_2^{ij} + pF_3^{ij}\right)}{\sqrt{1+p^2+q^2}}d\lambda. \tag{28}$$

The quantity $(\int\rho(\lambda)d\lambda/\int V(\lambda)d\lambda)$ can be thought of as an effective albedo, or as a perceptually weighted albedo.

4. Experiments

One of the goals of the experimental demonstration is to show that the approach will work with commonly available illuminant sources and cameras. To this end, the experiments used a Samsung 931BF LCD monitor was used as the programmable planar extended illuminant, and the luminance was measured by a Canon A95 digital camera.

The LCD monitor has a pixel pitch of 0.294 mm, and has a resolution of 1280×1024 pixels. The *F*-images for this monitor, assuming an object distance of 291 mm, are shown in Fig. 2.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

4.1. Monitor calibration

The monitor has a non-linear relationship between the commanded R, G, B values supplied by the program displaying the image and the actual luminance of the pixels. To determine this relationship, a monitor calibration was performed using the *Datacolor Spyder3 Elite* colorimeter. This provides the monitor luminance as a function of the commanded R, G, B values. The result of the monitor calibration is shown in Fig. 3. This shows the relationship between the commanded digital value (assuming $R = G = B =$ digital value) and the measured monitor luminance. The measurements are shown as small circles in the graph. The solid line in the graph represents a curve fit to the data

$$L = 0.045 + 21.36\left(\frac{V}{255}\right)^{\gamma}, \tag{29}$$

where V is the digital value and $\gamma = 2.27$.

We assume that the LCD pixels are isotropic illuminants, which is not the case. There is a significant off-axis falloff in emitted light, which is wavelength dependent. The assumption of isotropy is made more palatable by the limited range of angles in our experimental setup, where the object is small and centered with respect to the monitor screen.

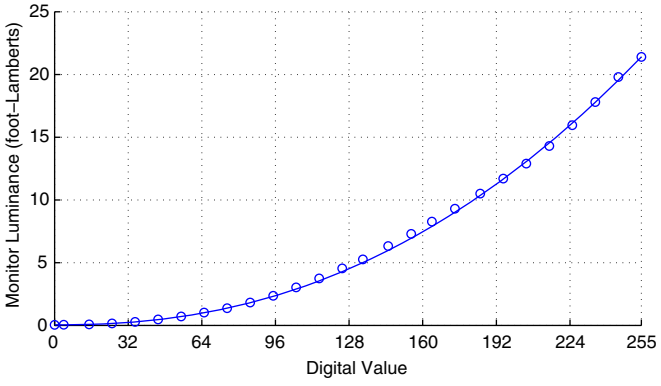

Fig. 3. The relationship between the commanded intensity value and the measured luminance for the Samsung 931BF monitor, measured using a DataColor Spyder3 Elite colorimeter.

4.2. Camera photometric calibration

There is a also a non-linear relationship between the imaged luminance and the measurement, normally referred to as the opto-electronic conversion function (OECF). This relationship is determined by imaging of a standard test chart (as defined by the ISO standard document ISO14524 [9]) illuminated by an isotropic diffuse illuminant [20]. The test pattern chart is shown in Fig. 4. It consists of 12 small square patches, arranged in a circle, each with a different density, superimposed on a neutral background.

In the experiments described in this paper the camera settings were as follows:

- Manual mode.
- Shutter speed: 0.1 s (0.4 s for the experiment involving the flashlight object).
- Aperture: f2.2.
- Equivalent ISO sensor speed: 400.
- White balance mode: daylight.

These settings provided a good dynamic range when viewing the ISO14524 test sheet as illuminated by a full-screen mid-grey from the monitor. The "daylight" setting for the color balance provided digital output levels closest to neutral for the test chart background.

Following the procedure defined in the ISO14524 standard, ten images were captured of the test pattern. The test pattern was centered on (or near) the optical axis of the camera before acquiring the images. As prescribed in the standard, the test pattern was oriented at 60° with respect to the illumination source to avoid specular reflections, and was located at a distance of 291 mm from the illuminant. In all cases the illumination was provided by the monitor displaying a full-screen full-intensity ($R = G = B = 255$). The camera was aligned with its optical axis perpendicular to the test chart, at a distance of roughtly 250 mm.

As a luminance meter was not available, the test chart luminances were estimated, as suggested in the ISO14524 standard, using the formula $L = 10^{-D}E/\pi$ where D is the test patch density and E is the illuminance, in lux, seen by the test patch. The illuminance at each test patch location was measured with a cosine-corrected photometer, model TES-1339. The test patch densities are as given in the ISO14524 standard document [9].

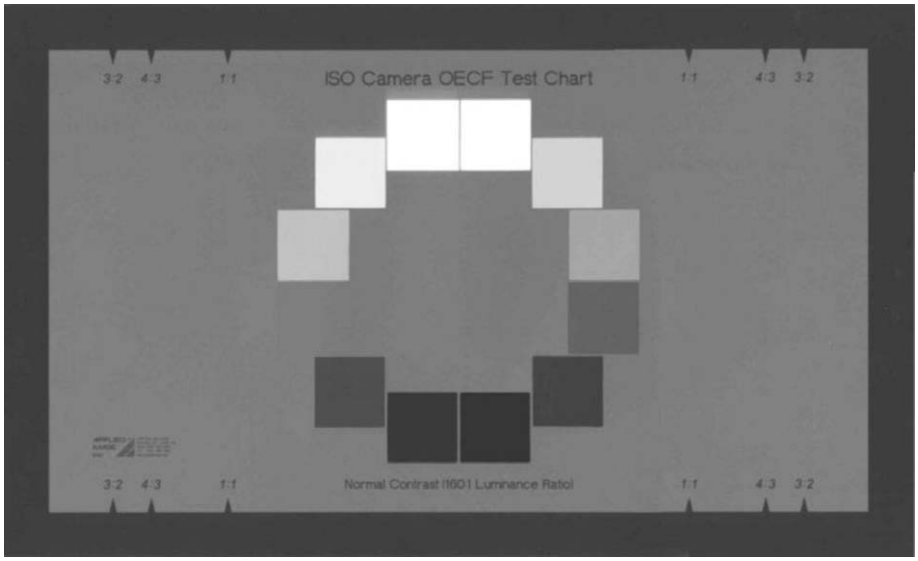

Fig. 4. The ISO 14524 standard contrast test chart used to determine the opto-electronic conversion function of the Canon A95 digital camera.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

For each image, the mean values of the $(R+G+B/3)$ values in 64×64 pixel windows centered on each of the twelve test pattern patches were computed. The result of the calibration process is summarized in Fig. 5. In this figure is shown the logarithm of the mean measured camera $(R+G+B)$ values as a function of the logarithm of the estimated log scene luminance, averaged over all 10 acquired images. The circles show the measured values and the solid line is a least-squares fit to the data. It can be seen that the data are well modeled by a power-law relation

$$L = 0.004V^{1.32}. \tag{30}$$

4.3. Experimental results – planar white surface, white illuminant

In the first set of experiments, the ISO14524 testchart was used as a test surface. The ISO14524 standard requires that the test charts be spectrally neutral, with a variation of not more than 0.1 density units over the wavelength range of 420–680 nm. Thus, we assumed the surface albedo of the test chart to be wavelength independent and applied Eq. (26). The center of the testchart was aligned with a laser level to fall on the perpendicular axis through the center of the monitor, and at a distance of 291 mm. The testchart plane was oriented with its normal vector in a horizontal plane perpendicular to the monitor plane, and rotated to various angles about the vertical axis. The rotation about the vertical axis was measured to an accuracy of about 2° with a laser level protractor. The rotation about the horizontal axis, which was nominally set to zero, was difficult to measure accurately in our setup and may have deviated by as much as 20° from zero. Images were acquired of the testchart oriented at three different angles about the vertical axis: −15°, 0° and 60°. Angles between 20° and 50° resulted in significant specular reflections being observed by the camera and so were not considered.

The algorithm was run using two different sets of illumination patterns, each produced by 1280×1024 pixel grayscale (8 bits/pixel) images displayed on the monitor. The first set of illumination patterns is shown in Fig. 6. This "blocks" set consists of nine images containing simple rectangular regions of constant luminance. These patterns were chosen to provide high equivalent radiances with a relatively wide spread of equivalent light source directions. Fig. 7 shows the second set, the "Grand Place" set, which consists of nine images taken from a video sequence. The bottom three images

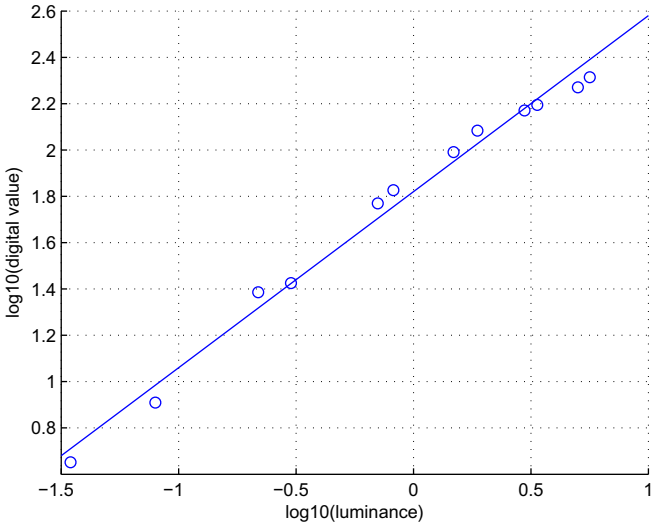

Fig. 5. The relationship between the log of the measured intensity value $(R+G+B/3)$ and the log of the estimated luminance of a test pattern patch, for a Canon A95 digital camera.

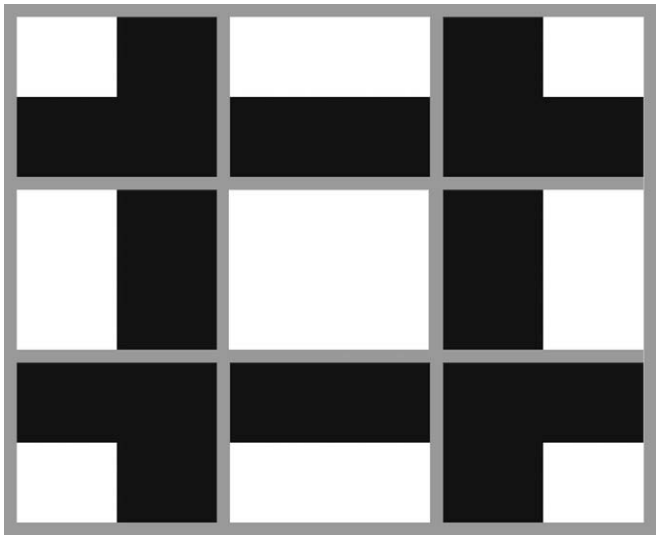

Fig. 6. The first illumination pattern set, "blocks", consisting of nine different images with constant rectangular luminance regions.

in this set were actually mirrored versions of the top three images. This was done to provide a wider spread of equivalent light source directions. The distribution of equivalent light source directions (projected onto the illuminant plane) for the two illuminant pattern sets are shown in Fig. 8. The locations on the illuminant plane are given in units of monitor pixels (0.294 mm).

The results of running the least squares photometric stereo algorithm on the images acquired from the two sets of illuminant patterns are summarized in Table 1. The table shows the nominal surface rotation angles compared with the ones measured by the algorithm (X–Z angle = $\arctan(p)$, Y–Z angle = $\arctan(q)$) for the two sets of illuminant patterns. The nominal X–Z angles were measured to an accuracy of about 2°, while the nominal Y–Z angles varied around zero between experiments (as the testpattern was rotated) by about ±20°. The Y–Z angle was relatively uncontrolled and could not be measured accurately with our setup. The X–Z angles were estimated by the algorithms to within 4° in the case of the "blocks" illuminant pattern, while the estimates in the case of the "Grand Place" illuminant set were off by as much as 10°. Presumably, this is in part due to the lower effective radiance of the "Grand Place" illuminant patterns as compared to the "blocks" illuminants, resulting in a lower signal-to-noise ratio in the former case. The spread in equivalent light source directions was similar in the two cases.

4.4. Experimental results – curved colored surface, white illuminant

The second experiment involved running the algorithm on a curved plastic object, in this case a flashlight. The object was a bright yellow color, and white illuminant patterns were used, so Eq. (28) was employed. The surface of the object was slightly specular. The part of the object that was considered by the algorithm can be seen in Fig. 9. This part is cylindrical with two indented regions. The radius of the cylinder is 25 mm. The algorithm is run on non-overlapping 10×10 pixel regions of the acquired images, under the assumption of a fixed distance (assumed to be 291 mm, which was the distance of the closest point on the object to the monitor plane). The axis perpendicular to the monitor through the center of the display screen intersected the object at the closest point. The same set of F-images was used for each 10×10 image block. In theory, a new set of F-images should be computed for each 10×10 block, since the illuminant patterns are effectively shifted by a displacement equal to the centroid of the 10×10 regions as projected onto the object surfaces. In practice, this dis-

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

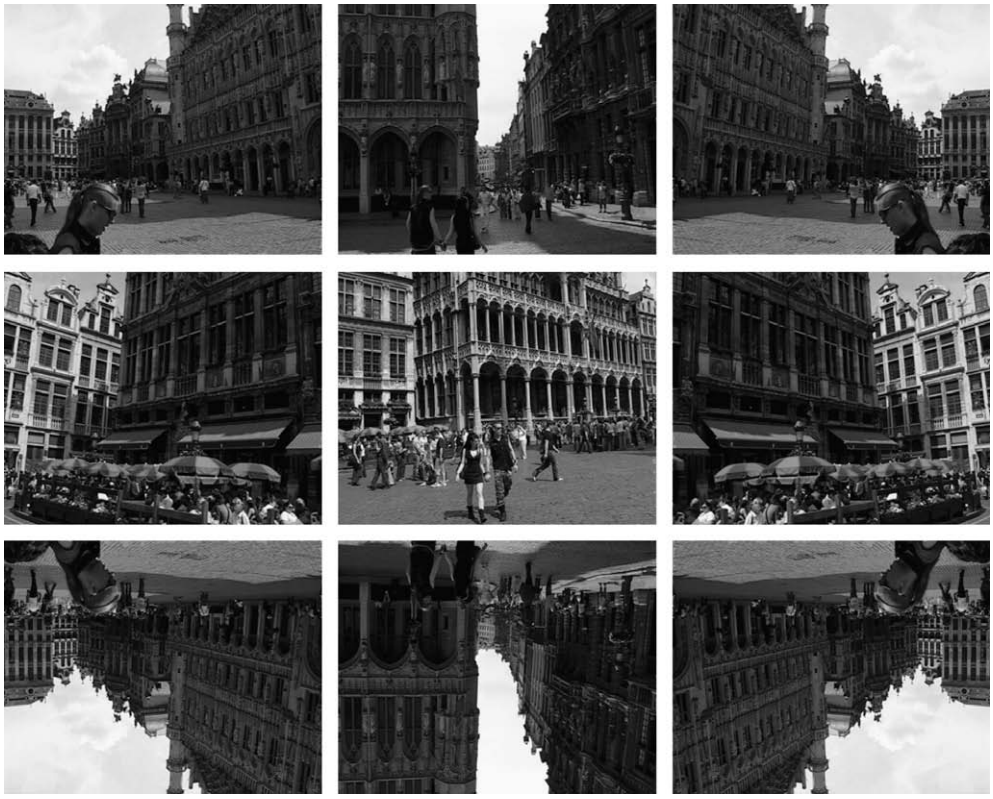

Fig. 7. The second illumination pattern set, "Grand Place", consisting of nine frames taken from a video sequence. Some of the frames are repeated with mirroring about the *x*- and *y*-axes to improve the distribution of equivalent illuminant directions.

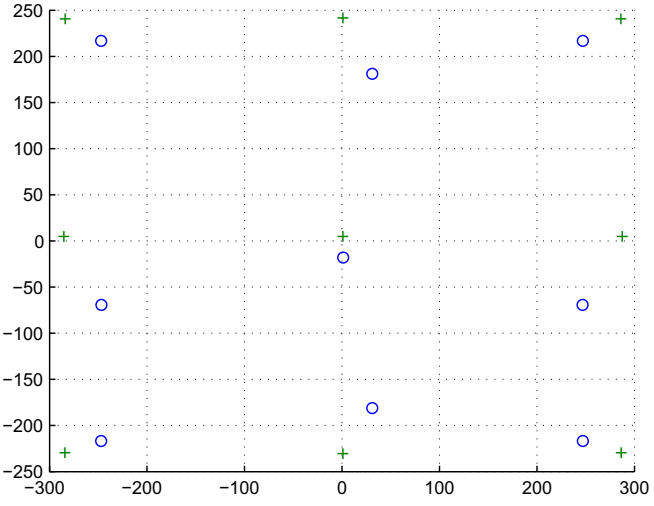

Fig. 8. The distribution of equivalent point light source directions (projected onto the illuminant plane, with units of pixels). The circles indicate the equivalent illuminant directions for the "Grand Place" image sequence, and the crosses indicate those for the "blocks" image sequence.

placement is quite small (about 10 mm) and, because of the smoothness of the *F*-images, made little difference to the computed result, while greatly reducing computation. For larger objects these displacements would be big enough to require computation of individual *F*-images for each block.

The images acquired of the object as illuminated by the nine different illuminant patterns in the "blocks" set are shown in Fig. 9.

The resulting surface normal map for the part of the image corresponding to the dashed box in Fig. 9 is shown in Fig. 10.

Table 1
Experimental results.

	Nominal *X*–*Z* angle	Computed *X*–*Z* angle	Computed albedo
"blocks"	−15	−16.4	0.30
	0	3.93	0.42
	60	62.5	0.51
"grand place"	−15	−5.1	0.34
	0	7.76	0.41
	60	58.8	0.41

The depthmap obtained by integrating the surface normal field is shown in Figs. 11 and 12. Fig. 11 shows the surface obtained using the surface normals from the entire image, while Fig. 12 shows the surface obtained using the surface normals only from the region outlined by the dashed box in Fig. 9. The "M-estimator" method described by Agrawal et al. in [1] was used to do the integration (making use of the Matlab code kindly supplied by Agrawal on his web site). The surface shown in Fig. 11 demonstrates the effect of shadowing – the left hand part of the surface flattens out where it should continue curving. Note that the algorithm picks out the two slight indentations in the surface. There is also a slight tilt computed by the algorithm. This represents the slight tilt of the monitor plane relative to the vertical axis of the flashlight object. Fig. 11 also shows that the algorithm produces a curving along the vertical axis of the object near the top of the image. The object surface should be flat in this direction. This effect cannot be explained by shadowing, nor to specular reflection, it could be a result of vignetting, or due to the invalidity of the assumed *F*-images away from the perpendicular axis through the center the monitor. Fig. 12, which shows the depth map in a small region about the monitor axis, does not exhibit this vertical curvature.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

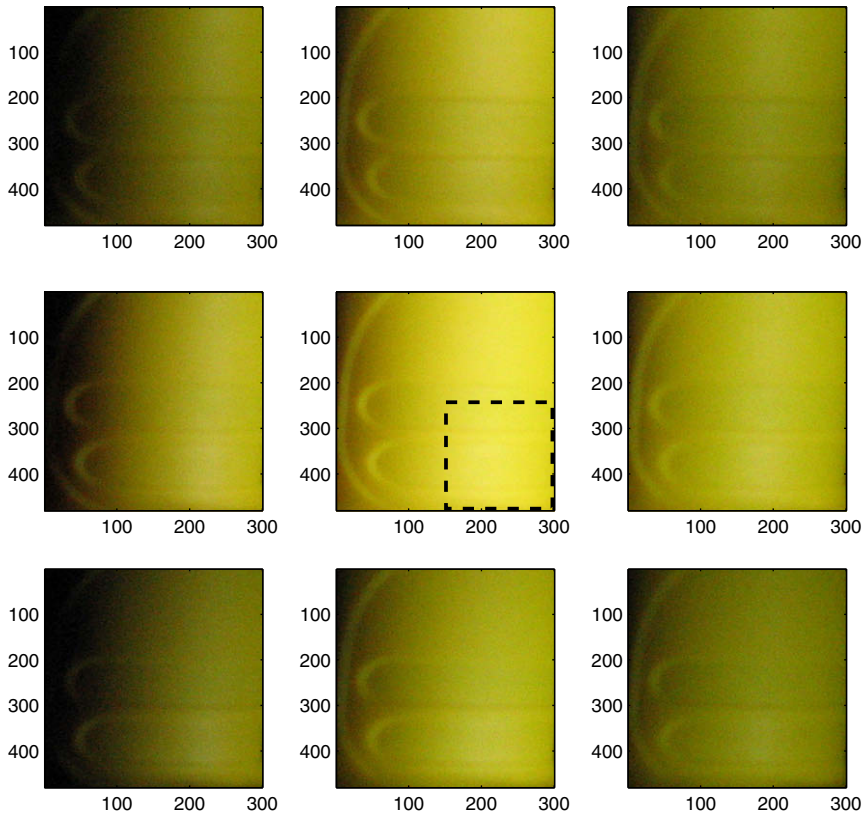

Fig. 9. The images of the flashlight object taken with the "blocks" illumination pattern set.

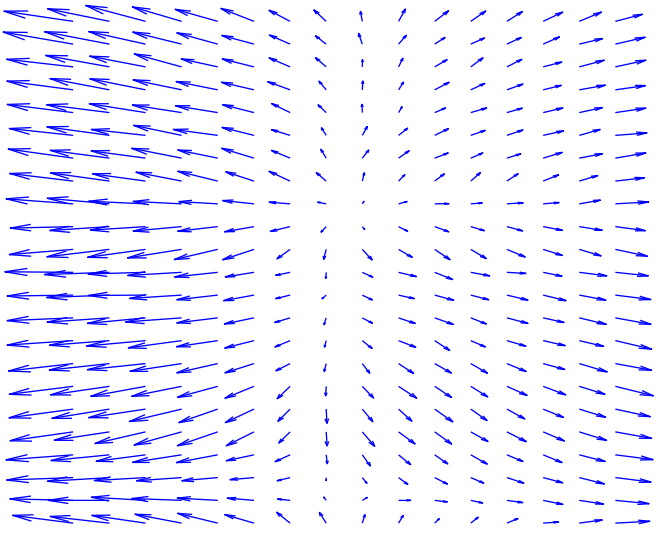

Fig. 10. The surface normal map obtained for a small region of the flashlight object, centered on the perpendicular axis through the monitor center.

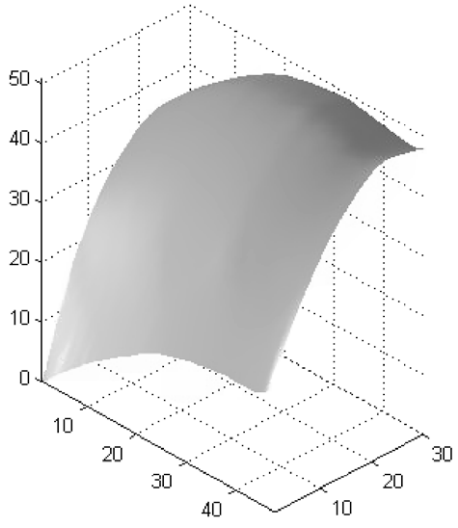

Fig. 11. The depth map for images in Fig. 9, obtained by integrating the surface normal field.

The variation of the computed effective albedo over the image is shown in Fig. 13. The average of the effective albedo over the image was 0.17, with a maximum of 0.23 and a minimum of 0.04. The higher albedo values are most likely due to the slight specular component and the lower albedo values are most likely due to shadowing.

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 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 illumi-

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

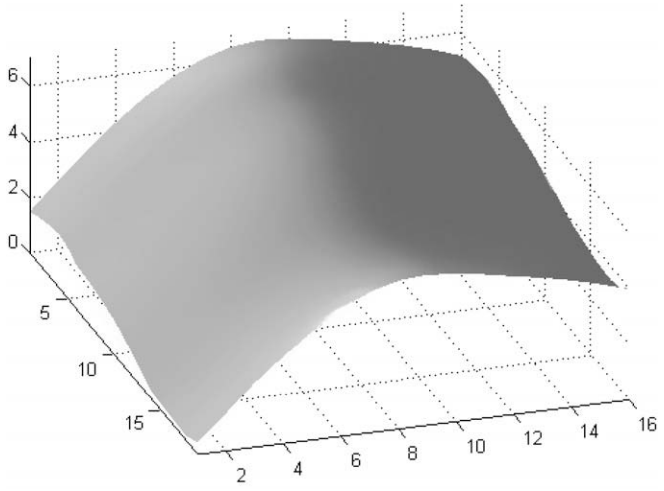

Fig. 12. The depth map for a small region of the image indicated by the dashed box in Fig. 9.

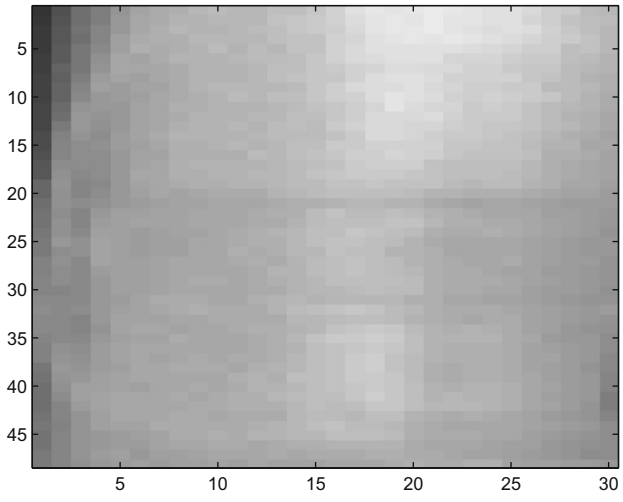

Fig. 13. The effective albedo map for the flashlight object.

nant was not uniform. This case can be handled by simply decomposing the illuminant field into an array of small rectangular subregions (or pixels).

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 and recomputing the F-images.

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. Another difficulty arises in practical situations in which objects are non-convex. In such cases, there can be significant mutual illumination, or interreflection, between different surface patches of the object(s). Woodham [19] proposes using the distance between triples of intensity measurements, resulting from different illuminants, and a 2D calibration surface dependent on the surface albedo. If the surface is Lambertian and there is no shadowing or interreflection, all such triples of intensity measurements will fall on the 2D calibration surface. Shadowing or interreflection will provide either a decrease or increase in the intensity measurement, causing the measurements to move off of the calibration surface. This provides a way of detecting shadowing or interreflection. One can use this detection technique to discard measurements which do not satisfy the assumptions of the approach. Woodham's implementation projects the measurements on the calibration surface, and uses the projection to estimate the surface gradient (p, q). Woodham notes that this projection tends to reduce errors in his algorithm due to interreflections, but this decrease is not guaranteed. He points out that the distance between the measurement point (in 3D measurement space) and the calibration surface can, however, be usefully employed as a confidence measure.

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 non-linear in p and q, making the solution difficult. For mirror-like specular surfaces the equations may be quite simple, as shown by Ikeuchi [8] in his work employing a planar light source to illuminate a specular surface.

LCD display technology is ever marching forward and recent advances can enhance the approach described in this paper. For example, autostereoscopic displays have the ability to control the amount of light emitted in narrow angular ranges. This has the potential of relieving problems with shadowing. The use of high-power LED backlights are increasing the brightness of displays. Combined with high-dynamic range digital cameras (to avoid saturation), this will increase the signal-to-noise ratio of the measurements used in our shape-from-shading algorithm.

References

[1] A. Agrawal, R. Raskar, R. Chellappa, What is the range of surface reconstructions from a gradient field? in: European Conference on Computer Vision (ECCV), 2006.
[2] J.J. Clark, Active photometric stereo, in: Proceedings of the Computer Vision and Pattern Recognition Conference, Champaign, IL, 1992.
[3] J.J. Clark, H. Pekau, An integral formulation of differential photometric stereo, in: Proceedings of the 1999 IEEE Computer Vision and Pattern Recognition Conference, Fort Collins, CO, 1999, pp. 119–124.
[4] J.J. Clark, Photometric stereo with nearby planar distributed illuminants, in: Proceedings of the Third Canadian Conference on Computer and Robot Vision, Quebec City, June 2006, p. 16.
[5] N. Funk, Using a raster display for controlled illumination, M.S. Thesis, University of Alberta Computer Science Department, 2005.
[6] N. Funk, Y.H. Yang, Using a raster display for photometric stereo, Technical Report, University of Alberta, Computer Science Department, TR06-13, Jun 2006.
[7] B.K.P. Horn, Obtaining shape from shading information, in: P.H. Winston (Ed.), The Psychology of Computer Vision, McGraw-Hill, New York, 1975 (Chapter 4).
[8] K. Ikeuchi, Determining surface orientations of specular surfaces by using the photometric stereo method, IEEE Transactions on Pattern Analysis and Machine Intelligence 3 (6) (1981) 661–669.
[9] ISO 14524:1999(E), Photography – electronic still-picture cameras method for measuring opto-electronic conversion functions (OECFs), International Standards Organization, 1999.
[10] Y. Iwahori, H. Sugie, N. Ishii, Reconstructing shape from shading images under point light source illumination, in: Proceedings of the 10th International Conference on Pattern Recognition, Atlantic City, 1990, pp. 83–87.
[11] Y. Iwahori, R.J. Woodham, N. Ishii, Shape from shading with a nearby moving point light source, in: Proceedings of the 2nd International Conference on Automation, Robotics and Computer Vision, Singapore, 1992, pp. CV.5.5.1–CV.5.5.5.
[12] Y. Iwahori, R.J. Woodham, H. Tanaka, N. Ishii, Moving point light source photometric stereo, IEICE Transactions on Information and Systems E77-D (No. 8) (1994) 925–929.
[13] B. Kim, P. Burger, Depth and shape from shading using the photometric stereo method, Computer Vision Graph Image Proceedings 54 (1991) 416–427.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011

[14] S. Magda, D.J. Kriegman, T. Zickler, P.N. Belhumeur, Beyond Lambert: reconstructing surfaces with arbitrary BRDFs, in: Proceedings of Eighth IEEE International Conference on Computer Vision, vol. 2, 2001, pp. 391–398.
[15] Y.Y. Schechner, S.K. Nayar, P.N. Belhumeur, A Theory of multiplexed illumination, in: Proceedings of the 2003 ICCV, Nice, France, 2003, pp. 808–813.
[16] W.M. Silver, Determining shape and reflectance using multiple images, SM Thesis, MIT Dept. Electrical Engineering and Computer Science, Cambridge, MA, 1980.
[17] L.B. Wolff, Shape understanding from Lambertian photometric flow fields, in: Proceedings of the 1989 Computer Vision and Pattern Recognition Conference, 1989, pp. 46–53.
[18] R.J. Woodham, Photometric method for determining surface orientation from multiple images, Optical Engineering 19 (No 1) (1980) 139–144.
[19] R.J. Woodham, Gradient and curvature from photometric stereo including local confidence estimation, Journal of the Optical Society of America A 11 (1994) 3050–3068.
[20] D. Wüller, H. Gabele, The usage of digital cameras as luminance meters, in: R.A. Martin, J.M. DiCarlo, N. Sampat, Digital Photography III, Proceedings of SPIE-IS&T Electronic Imaging, SPIE vol. 6502.
[21] G. Wyszecki, W.S. Stiles, Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd ed., Wiley, New York, 1982.

Please cite this article in press as: J.J. Clark, Photometric stereo using LCD displays, Image Vis. Comput. (2008), doi:10.1016/j.imavis.2008.10.011