Short Papers

A Simple Strategy for Calibrating the Geometry of Light Sources

Mark W. Powell, *Student Member, IEEE*,
Sudeep Sarkar, *Member, IEEE*, and
Dmitry Goldgof, *Senior Member, IEEE*

Abstract—We present a methodology for calibrating multiple light source locations in 3D from images. The procedure involves the use of a novel calibration object that consists of three spheres at known relative positions. The process uses intensity images to find the positions of the light sources. We conducted experiments to locate light sources in 51 different positions in a laboratory setting. Our data shows that the vector from a point in the scene to a light source can be measured to within 2.7 .4 at α .05 (6 percent relative) of its true direction and within 0.13 .02 m at α .05 (9 percent relative) of its true magnitude compared to empirically measured ground truth. Finally, we demonstrate how light source information is used for color correction.

Index Terms—Light source calibration, constancy, color correction.

1 Introduction

Our motivation in this research is to use physics to improve our understanding of the surfaces and objects we would like to reason about. Traditionally, most work in computer vision is done without correcting for illumination artifacts like shading, highlights and shadows, which limit effectiveness. For example, intensity changes that result from surface shading are a major problem for a segmentation algorithm. In one particular application in which color segmentation is used to recognize psoriasis discolorations on skin, the authors note: "...the cylindrical shape of the limbs resulted in some difficulties in assessing involved skin area [1]." These kinds of problems are due in part to the lack of knowledge and control over the illumination source. It is intuitive that correcting for variation of illumination would improve segmentation, but the extent of improvement remains to be demonstrated.

One class of work seeks to bypass the illumination calibration problem by formulating illumination-invariant measures, e.g., [2], [3]. In this paper, we adopt a more direct approach and estimate the light source parameters. To enable us to correct for illumination we present a novel strategy for *calibrating light source* direction and location using specular spheres. To our knowledge, this is the first work that presents a technique for calibrating the light source geometry for a source that is near and uses real images. Possible applications of these techniques include controlled scenarios such as those observed in human-computer interaction using face and gesture, teleconferencing, and medical applications such as quantification of burn scars and skin cancer. The knowledge of light source position can also be used in vision techniques such as shape-from-shading, photometric stereo, shape-from-shadows, and shape-from-specularities.

• *The authors are with the Department of Computer Science and Engineering, University of South Florida, Tampa, FL 33620. E-mail: {mpowell, sarkar, goldgof}@csee.usf.edu.*

Manuscript received 22 Mar. 2000; revised 25 Aug. 2000; accepted 8 Jan. 2001.
Recommended for acceptance by E. Hancock.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 111753.

2 Related Work

Color constancy is a difficult problem, which is made apparent by the number of theories and algorithms that have been conjectured to solve it [4], [5], [6], [7], [8], [9], [10], [12]. Because the problem is underconstrained, quite often researchers make assumptions about the scene that are highly impractical. One common assumption is that the scene is two-dimensional and evenly illuminated [5], [7], [8], [9]. The Mondrian, for example, is a set of homogeneously-colored surface patches arranged on a flat surface that is sometimes used to validate color constancy methods. Even if such large assumptions are made, the problem of color constancy is still underconstrained since for n homogeneously-colored regions, there are $3n$ knowns from a color image and $3\ n\quad 1$ unknowns, where the n regions are each described by three surface reflectance coefficients and the light source color is described by three parameters. This formulation is still a simplification of the ideal, which is to obtain a finer sampling of the surface reflectance function of each region and the spectral power distribution of the illuminant, which would require the estimation of 30 or more parameters per region and illuminant. The complexity of this problem is exacerbated by the added complexity of variations in intensity due to shading in real images. Through explicit calibration of 3D light source position and estimation of surface orientation such as surface normal estimation from a range image, we can transform an image so that it appears uniformly illuminated, factoring out intensity variation from shading. By accounting for the added complexity introduced in real images, the assumptions made to simplify color constancy that are criticized as impractical can be made more reasonable.

Among previous works in light source calibration, Mancini and Wolff triangulate the position of the light source given intensity and range information for any illuminant position near the scene, initially using surface normal estimation from a range image. The position estimate is then refined using a shape-from-shading technique [13]. The performance of this algorithm was shown on synthetic images but not real ones. Other works in estimation of light source direction has been in the context of shape from shading, where a point source that is far away is assumed [14], [15], [16], [17], [18], [19]. The assumption of an infinitely distant light source is not practical for an indoor scene such as a laboratory or production facility, where the direction from each point in the image to the light source varies over the entire image.

3 Methodology

We can triangulate the positions of any number of light sources that give rise to an image if we can place a pair of calibration spheres in the scene which produce visible highlights for each light source. Our strategy is to estimate the positions and surface normals at the highlights in order to triangulate the illuminants.

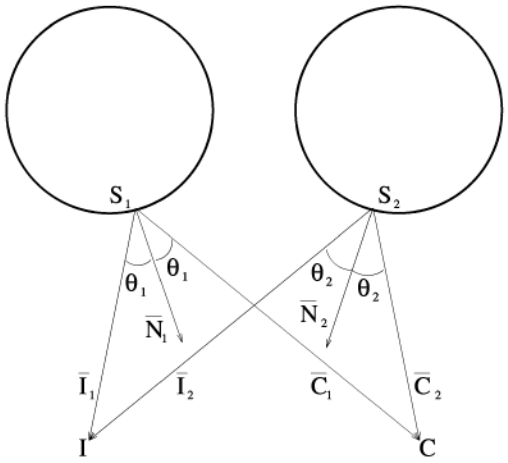

Fig. 1. Setup for light source position estimation. (a) Two spheres are placed in the scene which exhibit specular reflections at points S_1 and S_2 from the illuminant at unknown point I captured in an image by a camera of known position at point C. The surface normals to the sphere at the specularity points are denoted $\vec{N}_1$ and $\vec{N}_2$. The incoming and outgoing angles (which are equal by the law of reflection) are labeled θ_1 and θ_2. Note: for multiple light sources $I_1, I_2, \ldots, I_n$, this diagram can be replicated n times.

3.1 Light Source Calibration

First, we would like to estimate the center positions of the calibration spheres. Consider two spheres with centers located at x_c, y_c, z_c and x_c, y_c, z_c, separated by distance D and located relatively near to and far from the camera, respectively. Let the points u, v $f\frac{x_c}{z_c}, f\frac{y_c}{z_c}$ and u, v $f\frac{x_c}{z_c}, f\frac{y_c}{z_c v}$ in the image plane be the perspective-projected locations of the two sphere center points, where f is the focal length. We know the sphere radii R and R a priori. Let us assume that the spheres project as circles onto the image plane. This approximation is valid if the spheres are far away from the camera relative to the deviation from the optic axis. We will handle the most general case in the next section. We can relate the sphere radii to the projected circle radii r and r by r $f\frac{R}{z_c}$ and r $f\frac{R}{z_c}$. By substitution, the 3D coordinates of the sphere centers are given by x_c $\frac{uR}{r}$, y_c $\frac{vR}{r}$, z_c $\frac{fR}{r}$, x_c $\frac{uR}{r}$, x_c $\frac{uR}{r}$, y_c $\frac{vR}{r}$, and z_c $\frac{fR}{r}$.

To obtain the z-coordinates, it is necessary to estimate the focal length f. We can estimate this by using the known distance D between the center points of two spheres:

$$x_c \quad x_c{}^2 \quad y_c \quad y_c{}^2 \quad z_c \quad z_c{}^2 \quad D.$$

By isolating the z_c terms and substituting the z terms, we have

$$f^2\left(\frac{R^2}{r^2} \quad \frac{R^2}{r^2}\right)^2 \quad D^2 \quad x_c \quad x_c{}^2 \quad y_c \quad y_c{}^2$$

and solving for f, we get

$$f \quad \sqrt{\frac{D^2 \quad x_c \quad x_c{}^2 \quad y_c \quad y_c{}^2}{\left(\frac{R^2}{r^2} \quad \frac{R^2}{r^2}\right)^2}}.$$

Given the estimate of the focal length, we can estimate the center of the spheres. Each point light source that lies in front of a specular sphere will produce a visible highlight on its surface. Next, we wish to estimate the positions of the specular highlights that appear on each calibration sphere. The solution, given in [11], is as follows: Let the known sphere center position be a, b, c and the unknown highlight point be a, b, c $r\, i, j, k$, where r is the known radius of the sphere. The vector of unknown direction cosines for the line segment between the center of the sphere and the highlight point is i, j, k. Let u, v be the perspective projection of the highlight on the image plane. Therefore, u $f\frac{a}{c}\frac{ri}{rk}$ and v $f\frac{a}{c}\frac{rj}{rk}$. So, i $\frac{urk}{fr}\frac{uc}{fr}\frac{fa}{fr}$ and j $\frac{vrk}{fr}\frac{vc}{fr}\frac{fb}{fr}$. We also know that i^2 j^2 k^2 1, so substituting for i and j gives us:

$$r^2\, u^2 \quad v^2 \quad f^2\, k^2 \quad 2r\, u\, uc \quad fa \quad v\, vc \quad fb$$
$$k \quad uc \quad fa^2 \quad vc \quad fb^2 \quad rf^2 \quad 0.$$

We can solve this quadratic for two possible values of k. Given k, the values of i and j are also determined. We choose the highlight position from the two solutions by taking the root where the k component is negative, or pointing toward the camera.

Once we have computed the locations of the highlights that a light source creates on two spheres of known radius, we can triangulate its position. The locations of the light source I, camera C, and spheres are shown in Fig. 1. Points S_1 and S_2 are the specularities on each sphere from the light source. Point C is also the origin. We have the surface normals $\vec{N}_1$ and $\vec{N}_2$ at the points S_1 and S_2 because we have previously estimated the sphere's center and highlight positions. The normalized vector from the center to the highlight is the normal. The angles θ_1 and θ_2 are the angles between the vectors $\vec{N}_1$ and $\vec{C}_1$ and $\vec{N}_2$ and $\vec{C}_2$, respectively. We first wish to find the direction of vector $\vec{I}_1$ (and $\vec{I}_2$). We exploit the fact that the angle of incidence is equal to the angle of reflection, which is denoted by θ_1 (and θ_2). We need to scale the vectors $\vec{N}_1$ and $\vec{N}_2$ such that $\vec{N}_1$ $\vec{C}_1$ and $\vec{N}_2$ $\vec{C}_2$ are parallel to $\vec{I}_1$ and $\vec{I}_2$, respectively, or $a_1\vec{N}_1$ $\vec{C}_1$ $\vec{I}_1$ 0 and $a_2\vec{N}_2$ $\vec{C}_2$ $\vec{I}_2$ 0. Expanding the first equation, we have $a_1\vec{N}_1$ $\vec{I}_1$ $\vec{C}$ $\vec{I}_1$ 0. We evaluate the vector product to get

$$a_1\, \vec{N}_1 \quad \vec{I}_1 \sin\theta_1 \quad \vec{C}_1 \quad \vec{I}_1 \sin 2\theta_1 \quad 0.$$

Therefore, a_1 $2\,\vec{C}_1 \cos\theta_1$, and, similarly, a_2 $2\,\vec{C}_2 \cos\theta_2$.

Finally, we triangulate the light source position I. Consider the triangle in Fig. 2b, where the vertices are I and the two highlight positions S_1 and S_2. We know a priori the length of b, the distance between points S_1 and S_2. We can compute ι_1 opposite I_1 and ι_2 opposite I_2 by finding the angles between $\vec{I}_1$ and $\overline{S_1S_2}$ and between $\vec{I}_2$ and $\overline{S_2S_1}$. Finally, the law of sines gives us I_1 $\frac{b\sin\iota_1}{\sin\beta}$, and, similarly, I_2 $\frac{b\sin\iota_2}{\sin\beta}$. When we rescale the vectors $\vec{I}_1$ and $\vec{I}_2$ to have magnitudes I_1 and I_2, respectively, we have vectors from each highlight point to the light source. The point I, where $\vec{I}_1$ and $\vec{I}_2$ come nearest to intersecting, may either be found by least-squares estimation, or by computing the midpoint of the line segment of minimum distance that connects the two rays, which is equivalent.

3.2 Implementation

We use calibration spheres that are specularly reflective on one hemisphere and produce highlights for every light source in front of them (from the camera's perspective). The other hemisphere is matte, allowing us to rotate the sphere in place with no translation. The matte surface is painted very brightly relative to the image background, simplifying sphere boundary segmentation. The calibration object consists of three spheres (see Fig. 3). The two front-most spheres are used to produce the highlights needed to triangulate the illuminant positions. The third sphere is farther away and used to estimate the focal length.

The positions of the specular regions on the spheres are automatically determined from the image. Since the regions are specular, they are by definition among the brightest in the image because they reflect back most of the illumination they receive. It is

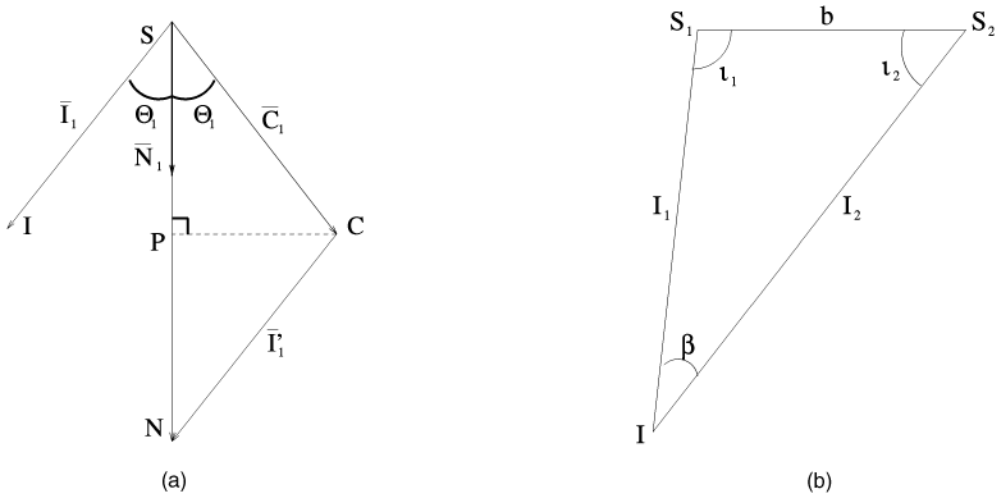

Fig. 2. (a) Geometry used to find the direction of the light source. $\vec{N}_1$ is a unit vector parallel to $\overline{SN}$. I_1 is estimated by $\vec{I}_1 \quad a_1\vec{N}_1 \quad \vec{C}_1$, where $a_1 \quad 2\,\vec{C}_1 \cos\theta_1$. (b) Triangulating I, the light source position.

therefore possible to threshold out all the specular regions from the intensity image. We find the maximum intensity value of the image, set a threshold value to 90 percent of this value and use it to threshold the image, producing a binary image where the highlights are foreground regions. We use empirically set maximum (300) and minimum (25) size thresholds to choose only the highlight regions from the set of foreground regions.

The next stage of the algorithm automatically determines the correspondence between the specular regions. For each source of illumination, one specular region on each of the two spheres is created. We can determine which pairs of regions result from the same illuminant by grouping the highlight regions into two sets by their relative locations. For our implementation, it is sufficient to determine whether they lie on the left or right half of the image, since each sphere lies completely on one side of the image or the other. We then sort the two sets of regions first by their horizontal coordinates and then by their vertical coordinates, which effectively pairs the highlight regions together according to light source across the two sets. More sophisticated correspondence algorithms might be used to match up the specular points in the general case when the calibration spheres are not horizontally aligned.

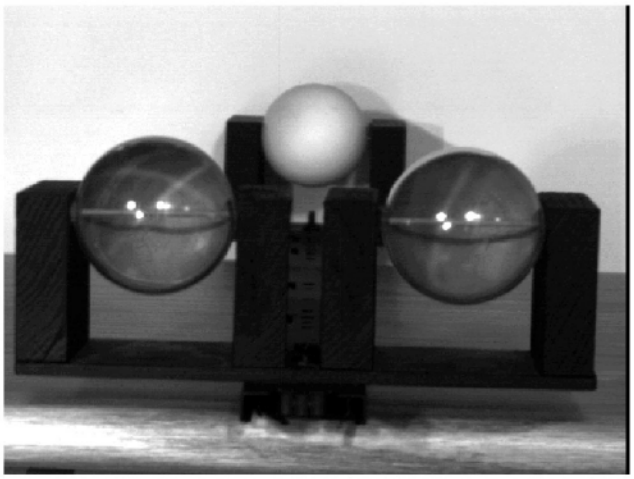

Fig. 3. Light source calibration setup illuminated by three different light sources simultaneously. The three spheres are specular on one side (shown on the front-most two spheres) and Lambertian on the other (as shown on the rear-most sphere). The 3D positions of the specularities visible on the surface are used to triangulate for the light source positions.

We have thus far assumed that the radius of the circle in the image is known. In order to know this, we must be able to find the boundary of the circle in the image and accurately estimate the ellipse that best fits it. The matte halves of the spheres used in our experiments are painted bright gray for the two near to the camera and near white for the third that is farther away. The farthest sphere is painted brighter so that it will not appear darker than the other two spheres, since the brightness of illumination diminishes over distance. The mounting device for the three spheres was painted matte black and placed on a table top covered in dull black fabric in front of a black felt backdrop. This was an ideal setup to threshold out the spheres from the background. To automatically estimate the parameters of the ellipse in the image, the raw greyscale image was binary thresholded, then edges were computed with a Sobel operator, which in turn were input to Maclaughlin's Hough transform based ellipse finder [21].

4 Experimental Results

4.1 Accuracy

A set of 17 images each with three illuminants in different positions lighting the calibration object were captured. The 51 3D locations of the light sources in these images were empirically measured and recorded as ground truth for comparison. For the 51 light source locations we used, the distance from the sphere varied from 0.5 to 2.0 m, and the angle between the vector from the sphere to the light source and from the sphere to the camera varied from 10 to 55 , using the front, center point of the leftmost sphere as a reference point. We express the calibration error as the difference in the magnitude and direction of the actual vector from the point in question to the light source, and the estimated light source position vector. The ground truth light source locations were empirically measured in 3D with respect to the walls and floor of our lab with a measuring tape and a plumb bob (lead weight on string) suspended from the object to be measured and lowered until it touched the floor. Table 1 shows the results for the light source calibration procedure.

Fig. 4 plots the variation of the errors with respect to the distance of the light source from the scene. As expected, the error

TABLE 1
Error in Calibrating Light Source Direction and Distance from a Point in the Scene Measured over 51 Trials (α .05)

	Position estimation error
Mean magnitude error	.13±.02 m
Mean magnitude error (relative)	9%
Median magnitude error	.11 m
Median magnitude error (relative)	7%
Mean direction error	2.7±.4°
Mean direction error (relative)	6%

Light sources were placed over a range of 0.5 to 2.0 m distance and 10 *to* 55 *direction with respect to the optical axis.*

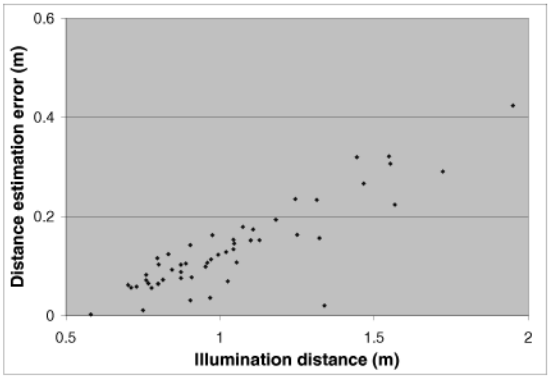

Fig. 4. Light source distance estimation error versus light source distance.

in magnitude of the light source vector with respect to the scene increases with distance from the scene. The reason for this is intuitive: the position of the illuminant must be triangulated from the two highlight positions on the sphere, which were 0.11 m apart. Like stereo imaging, triangulation is more accurate for positions closer to the optic axis and nearer to the scene rather than farther, or if the baseline distance between the calibration spheres is increased. We supposed that the error in estimated distance would also increase with respect to the angle between the light source vector and the optical axis for the same reason, but there was no apparent correlation. This may be because we were varying light source distance and direction at the same time in this experiment and the error with increasing distance outweighed the error with increasing angle from the optical axis. The error in estimated direction was not correlated with either distance of the light source or the angle between the light source vector and the optical axis.

4.2 Color Correction

As stated in Section 1, our motivation for calibrating light sources is to extract objective color from images. We can use known light source positions and surface geometry to remove the shading on surfaces which get darker in intensity as the angle between the surface normal and the lighting direction increases. To compute intrinsic surface color, we apply a correction factor to intensity values based on the light source and surface normal direction. We assume a Lambertian surface, where the reflected light varies directly with the cosine of the angle θ between the direction of incoming light and the normal of the surface at a given point. We estimate the surface normal at each pixel using a range image and the local least-squares planar fit. For every point in the image, the intensity value is divided by $\cos\theta$. Dividing the intensity values by this directionality factor results in an image that is very close to how the surface would appear if the entire surface were flat and oriented perpendicularly with respect to the direction of illumination. Although there are other methods of color correction that do not require knowledge of light source locations such as photometric normalization, explicitly modeling light sources and surface geometry allows us to correct color in a more physically objective and comprehensive manner. For instance, in addition to shading removal, we can use this information to correct for illumination variation over distance from the source, correct for intensity variation from perspective foreshortening, or identify the location of possible specular highlights by testing for equal incoming and outgoing angles of surface irradiance and radiance with respect to the surface normal.

Fig. 5 shows a person's face before and after color correction. The face is illuminated by two tungsten lamps. One light was positioned slightly to the left of the face (from the camera's point of view) about one meter away (measured from the center of the rightmost calibration sphere), and the other was positioned far to the right and about 1.5 meters away. In order to correct the color of this image, we first identify the regions where each light will contribute to the color image. In an area where both lights contribute to color, the correction is computed for both lights. However, in areas where just one light illuminates the surface, only the appropriate correction is applied there. Overall, the color correction procedure removes much of the shading and corrects for the varying distance from the surface to the light. There is a perceivable brightness difference between the regions of the face that are singly illuminated and those that are doubly illuminated. When corrected, the doubly illuminated regions are slightly

(a) (b)

Fig. 5. Image of a face before and after color correction.

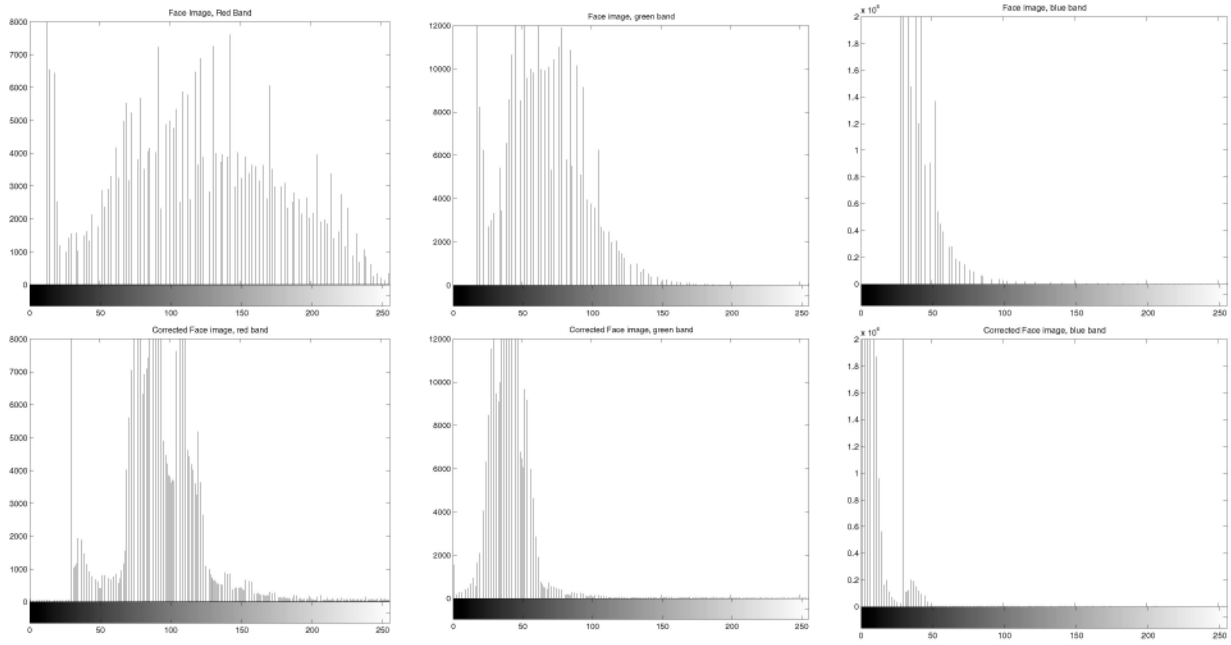

Fig. 6. From left to right, histograms of the red, green, and blue bands of the original (above) and corrected (below) face images (see Fig. 5), illustrating the reduction in intensity variance from color correction.

darker. Since the correction procedure is, in effect, applied twice at these areas, the error in the surface normal estimation from the range data is also doubled, and this may account for the slight difference in appearance. There are also several shadows in this image, the most noticeable of which are on either side of the nose. Fig. 6 shows the histograms of the original and corrected face image, illustrating the reduction in the intensity variance.

One practical application of color correction is in the objective measurement of burn scar color to assist in medical diagnosis. Fig. 7 shows a person's right cheek and neck with burn scars. After color correction, the contrast of the scar area is enhanced by the removal of shading and we arrive at a more objective measurement of scar color. Fig. 8 shows the histograms of the original and corrected burn scar image. For this image, we could have achieved a similar contrast enhancement by using histogram equalization. However, color correction from light calibration has the advantage that it can be performed across images taken at different times. Besides, color correction gives us a more objective, absolute measure of scar color. Histogram enhancement would just enhance the subjective quality of the image.

5 Conclusions

We demonstrated a methodology for locating light source positions using both range and intensity images and using only intensity images. We introduced a novel calibration device consisting of three specular spheres at known relative locations. We conducted experiments using both variations of the technique, to locate light sources in 51 different positions in a laboratory setting. Our data confirms that the vector from a point in the scene to a light source can be estimated to within 3 (6 percent) of its true direction and 0.13 m (9 percent) of its true magnitude compared to empirical ground truth. Furthermore, we demonstrated how light source information can be used to extract intrinsic surface color. Our future work will involve looking at alternative methods of surface normal estimation and color correction that also accounts for non-Lambertian reflection from surfaces.

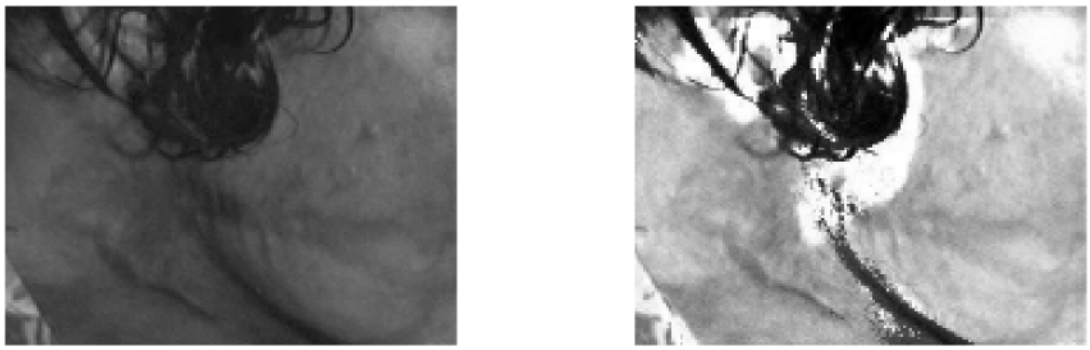

(a) (b)

Fig. 7. Image of right side of neck and cheek with burn scars before and after color correction.

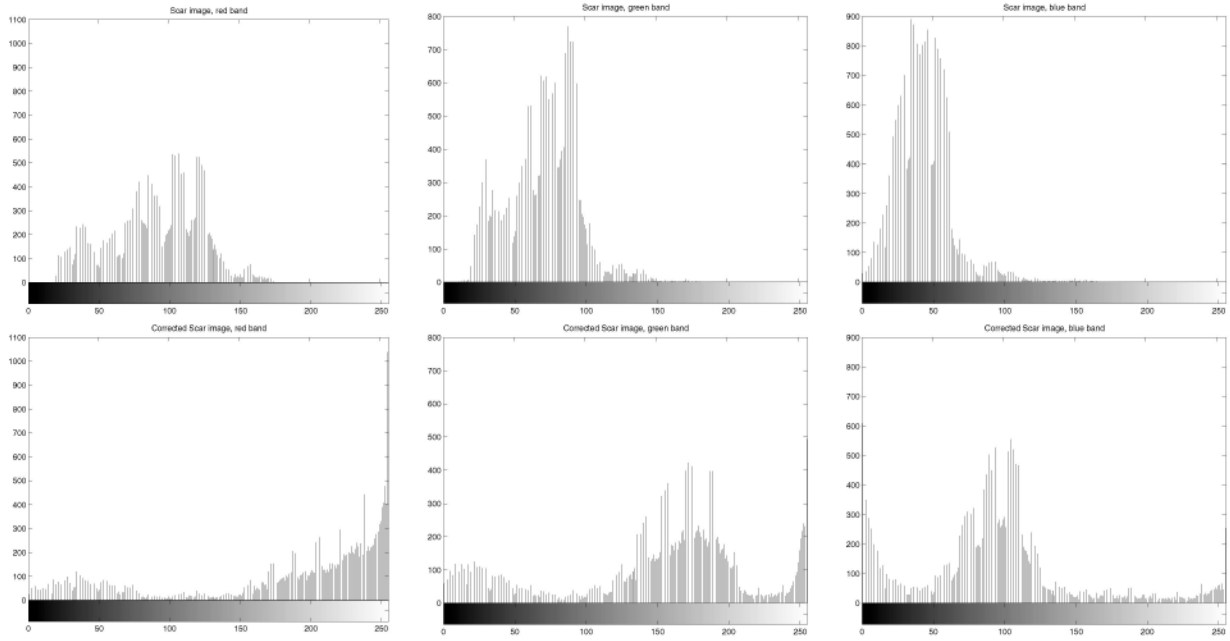

Fig. 8. Histograms of the red, green, and blue bands of the original (above) and corrected (below) burn scar images, illustrating the increased contrast after color correction.

ACKNOWLEDGMENTS

This work was supported by the Whitaker Foundation Biomedical Research Grant and the US National Science Foundation Grant EIA 9724422. A preliminary version of this work was presented at the IEEE Conference on Computer Vision and Pattern Recognition, 2000.

REFERENCES

[1] L. Savolainen, J. Kontinen, J. Roning, and A. Oikarinen, "Application of Machine Vision to Assess Involved Surface in Patients with Psoriasis," *British J. Dermatology*, vol. 137, no. 3, pp. 395-400, Sept. 1997.

[2] G. Healey and D Slater, "Computing Illumination-Invariant Descriptors of Spatially Filtered Color Image Region," *IEEE Trans. Image Processing*, vol. 6, no. 7, pp. 1002-1013, 1997.

[3] D.K. Panjwani and G Healey, "Markov Random Field Models for Unsupervised Segmentation of Textured Color Images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 10, pp. 939-954, Oct. 1995.

[4] D.H. Brainard and W.T. Freeman, "Bayesian Color Constancy," *J. Optical Soc. Am.*, vol. 14, no. 7, pp.1393-1411, 1997.

[5] M.M. D'Zmura and G. Iverson, "Probabilistic Color Constancy," *Geometric Representations of Perceptual Phenomena: Papers in Honor of Tarow Indow's 70th Birthday*, M.M. D'Zmura, D. Hoffman, G. Iverson, and K. Romeny, eds. Laurence Relbaum Assoc., 1994.

[6] G.D. Finlayson, "Color in Perspective," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 10, pp. 1034-1038, Oct. 1996.

[7] D.A .Forsyth, "A Novel Algorithm for Colour Constancy," *Int'l J. Computer Vision*, vol. 5, no. 1, pp. 5-36, 1990.

[8] E.H. Land, "Recent Advances in Retinex Theory and Some Implications for Cortical Computations: Color Vision and the Natural Image," *Proc. Nat'l Academy Science USA*, vol. 80, pp. 5163-5169, 1983.

[9] L.T. Maloney and B.A. Wandell, "Color Constancy: A Method for Recovering Surface Spectral Reflectance," *J. Optical Soc. Am.* vol. 3, no. 1, pp. 29-33, 1986.

[10] S.A. Shafer, "Using Color to Separate Reflection Components," *Color Research and Application*, vol. 10, pp. 210-218, 1985.

[11] S. Tominaga, "Surface Reflectance Estimation by the Dichromatic Model," *Color Research and Application*, vol. 21, no. 2, pp. 104-114, 1996.

[12] S. Tominaga and B.A. Wandell, "Standard Surface-Reflective Model and Illuminant Estimation," *J. Optical Soc. Am.*, vol. 6, no. 4, pp. 576-584, 1996.

[13] T.A. Mancini and L.B. Wolff, "3D Shape and Light Source Location from Depth and Reflectance," *IEEE Computer Vision Pattern Recognition*, pp. 707-709, 1992.

[14] D.R. Hougen and N. Ahuja, "Estimation of the Light Source Distribution and Its Use in Integrated Shape Recovery from Stereo and Shading," *Proc. Fourth IEEE Int'l Conf. Computer Vision*, pp. 148-155, 1993.

[15] M.J. Brooks and B.K. P Horn, "Shape and Source from Shading," *Proc Sixth Int'l Joint Conf. Artificial Intelligence*, pp. 932-936, 1985.

[16] K. Ikeuchi and K. Sato, "Determining Reflectance Parameters Using Range and Brightness Images," *Proc. Third IEEE Int'l Conf. Computer Vision*, pp. 12-20, 1990.

[17] C.H. Lee and A. Rosenfeld, "Improved Methods of Estimating Shape from Shading Using the Light Source Coordinate System," *Artifical Intelligence*, vol. 26, no. 2, pp. 125-143, 1985.

[18] A. Pentland, "Linear Shape from Shading," *Int'l J. Computer Vision*, vol. 4, pp. 153-162, 1990.

[19] Q. Zheng, R. Chellappa, "Estimation of Illuminant Direction, Albedo and Shape from Shading," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 680-702, July 1991.

[20] R.M. Haralick and L.G. Shapiro, *Computer and Robot Vision*, vol. 2, Reading, Mass.: Addison-Wesley, 1993.

[21] R.A. McLaughlin, "Randomized Hough Transform: Improved Ellipse Detection with Comparison," *PR Letters*, vol. 19, pp. 299-305, 1998.