Recovering Light Directions and Camera Poses from a Single Sphere

Kwan-Yee K. Wong, Dirk Schnieders, and Shuda Li

Department of Computer Science,
The University of Hong Kong,
Pokfulam Road, Hong Kong.
[kykwong,sdirk,sdli]@cs.hku.hk

Abstract. This paper introduces a novel method for recovering both the light directions and camera poses from a single sphere. Traditional methods for estimating light directions using spheres either assume both the radius and center of the sphere being known precisely, or they depend on multiple calibrated views to recover these parameters. It will be shown in this paper that the light directions can be uniquely determined from the specular highlights observed in a single view of a sphere without knowing or recovering the exact radius and center of the sphere. Besides, if the sphere is being observed by multiple cameras, its images will uniquely define the translation vector of each camera from a common world origin centered at the sphere center. It will be shown that the relative rotations between the cameras can be recovered using two or more light directions estimated from each view. Closed form solutions for recovering the light directions and camera poses are presented, and experimental results on both synthetic and real data show the practicality of the proposed method.

1 Introduction

The recovery of light directions from images is an important problem in both computer vision and computer graphics. For instance, light directions can be used to infer shape from images using shape-from-shading algorithms. In image-based rendering and augmented reality, light information of a scene is exploited to render virtual/real objects into the scene seamlessly.

In the literature, there exists a relative large number of work dealing with light estimation. Early works [1–4] focused on estimating a single distant point light source. However, multiple light sources are often present in a natural environment, and the problem of estimating multiple illuminants is even more challenging. In [5], Yang and Yuille showed that shading information along the occluding boundaries imposes strong constraints on the light directions. They also noted that without extra information, a unique solution for more than four light sources cannot be computed from a single image under the Lambertian model. In [6], Zhang and Yang estimated multiple illuminants from a sphere of known physical size by identifying the critical points, which are often difficult

to detect due to their sensitivity to noise. Takai et al. [7] proposed an approach based on two spheres to estimate near light sources. Wang and Samaras [8] combined shading and shadow cues to improve light estimation. The aforementioned methods are mostly based on the Lambertian model, and they all require the prior knowledge of the projections of some reference objects with known geometry to give the relationship between surface orientations and image intensities.

The specular reflection component of light is known to work in a very predictable manner, and it can be exploited for light estimation. A mirror ball was utilized in [9] to estimate the global illumination in a real world scene. Using such a mirror ball might, however, change the scene illumination due to its strong reflection property. Instead of using a purely specular sphere, spherical objects which exhibit both specular and diffuse reflections have been utilized in [10] and [11]. Powell et al. [10] used three spheres with known relative positions as a calibration object to triangulate the positions of light sources from specular highlights. Zhou and Kambhamettu [11] proposed an iterative method to recover the location and radius of a sphere from a pair of calibrated images, and used the recovered sphere to estimate the light directions from the specular highlights on the sphere. In [12], Li et al. combined multiple cues like shading, shadows and specular reflections to estimate illumination in a textured scene.

Similar to [11], this paper considers the problem of recovering multiple distance point light sources from a single sphere with unknown radius and location. Unlike [11] which requires multiple fully-calibrated views for recovering the radius and location of the sphere via an iterative method, it will be shown in this paper that light directions can be recovered directly from a scaled sphere estimated from a single view. Given multiple views of the sphere, a simple method is introduced to estimate the relative positions and orientations of the cameras using the recovered light directions. Hence, both the light directions and camera poses can be recovered using a single sphere. The proposed method will work under the assumption of a perspective camera with known intrinsics observing a sphere that reflects the specular component of multiple distant point light sources. As will be shown later, at least two distant point light sources are needed to uniquely determine the extrinsic parameters of the cameras.

The rest of this paper is organized as follows. Section 2 addresses the problem of sphere reconstruction from a single image. It is shown that with unknown radius, a one-parameter family of solutions will be obtained with all the sphere centers lying on the line joining the camera center and the true sphere center. Section 3 briefly reviews the standard technique for recovering light directions from the observed highlights of a sphere with known radius and location. It then proves that any sphere from the family of solutions recovered from a single image can be used to estimate the light directions. Section 4 presents a simple method for recovering the camera poses using the recovered sphere and light directions. Finally, experimental results on both synthetic and real data are given in Sect. 5, followed by conclusions in Sect. 6.

2 Scaled Reconstruction of Sphere

Consider a pinhole camera P viewing a sphere S. Without loss of generality, let the radius and center of the sphere be R and (X_c, Y_c, Z_c) respectively, and the camera coordinate system be coincide with the world coordinate system. The sphere S can be represented by a 4×4 symmetric matrix $\mathbf{Q}_s$ given by

$$\mathbf{Q}_s = \begin{bmatrix} \mathbf{I}_3 & -\mathbf{S}_c \\ -\mathbf{S}_c^T & (\mathbf{S}_c^T\mathbf{S}_c - R^2) \end{bmatrix}, \tag{1}$$

where $\mathbf{S}_c = [\, X_c \; Y_c \; Z_c \,]^T$ is the sphere center. Any 3D point X lying on S will satisfy the equation $\tilde{\mathbf{X}}^T\mathbf{Q}_s\tilde{\mathbf{X}} = 0$ where $\tilde{\mathbf{X}}$ represents its homogeneous coordinates. Suppose the 3×3 calibration matrix $\mathbf{K}$ of P is known, the projection matrix for P can be written as $\mathbf{P} = \mathbf{K}[\mathbf{I}\,\mathbf{0}]$. The image of S under P will be a conic C. This conic C can be represented by a 3×3 symmetric matrix $\mathbf{C}$, given by [13]

$$\begin{aligned} \mathbf{C} &= (\mathbf{P}\mathbf{Q}_s^*\mathbf{P}^T)^* \\ &= (\mathbf{K}\mathbf{K}^T - (\mathbf{K}\mathbf{S}_c/R)(\mathbf{K}\mathbf{S}_c/R)^T)^*, \end{aligned} \tag{2}$$

where $\mathbf{Q}^*$ denotes the dual to the quadric $\mathbf{Q}$, and is equal to $\mathbf{Q}^{-1}$ if $\mathbf{Q}$ is invertible. Any 2D point x lying on C will satisfy the equation $\tilde{\mathbf{x}}^T\mathbf{C}\tilde{\mathbf{x}} = 0$ where $\tilde{\mathbf{x}}$ represents its homogeneous coordinates.

The conic image C and the camera center will define a right circular cone which will be tangent to S, and the axis of this cone will pass through the sphere center $\mathbf{S}_c$ (see Fig. 1). If the radius R of the sphere S is known, $\mathbf{S}_c$ can be uniquely determined along this axis. In the next paragraph, a closed form solution for $\mathbf{S}_c$ will first be derived under a special case. The method for estimating $\mathbf{S}_c$ under the general case will then be discussed.

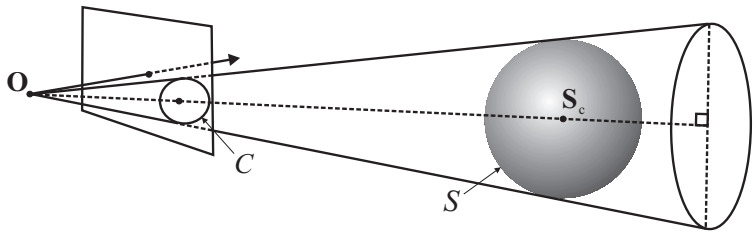

Fig. 1. The conic image C of the sphere S and the camera center will define a right circular cone. This cone is tangent to S and its axis passes through the sphere center $\mathbf{S}_c$.

Consider the special case where the sphere center lies along the positive Z-axis, and the camera calibration matrix is given by the identity matrix $\mathbf{I}_3$. Under this configuration, the sphere center will have coordinates $\mathbf{S}_c' = [\, 0 \; 0 \; d \,]$.

Note that d is also the distance between the camera center and the sphere center. The image of the sphere can be obtained using (2), and is given by

$$\mathbf{C}' = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{R^2}{R^2-d^2} \end{bmatrix}. \tag{3}$$

Note that $\mathbf{C}'$ represents a circle with radius $r = \sqrt{-\frac{R^2}{R^2-d^2}}$. The center of $\mathbf{C}'$ is at the origin (i.e., $(0,0)$), which is also the image of the sphere center. Given the radius r of $\mathbf{C}'$, the distance d between the camera center and the sphere center can be recovered as

$$d = R\frac{\sqrt{1+r^2}}{r}, \tag{4}$$

and the location of the sphere center follows.

Consider now the general case where the sphere center and the camera calibration matrix are given by $\mathbf{S}_\mathrm{c}$ and $\mathbf{K}$ respectively. Generally, the image of the sphere will no longer be a circle centered at the origin, but a conic $\mathbf{C}$ centered at an arbitrary point $\mathbf{x}_\mathrm{a}$. Note that $\mathbf{x}_\mathrm{a}$ is in general *not* the image of $\mathbf{S}_\mathrm{c}$. In order to recover $\mathbf{S}_\mathrm{c}$ from $\mathbf{C}$, the effect of $\mathbf{K}$ is first removed by normalizing the image using $\mathbf{K}^{-1}$. The conic $\mathbf{C}$ will be transformed to a conic $\hat{\mathbf{C}} = \mathbf{K}^\mathrm{T}\mathbf{C}\mathbf{K}$ in the normalized image. This conic $\hat{\mathbf{C}}$ can be diagonalized into

$$\hat{\mathbf{C}} = \mathbf{M}\mathbf{D}\mathbf{M}^\mathrm{T} = \mathbf{M}\begin{bmatrix} a & 0 & 0 \\ 0 & a & 0 \\ 0 & 0 & b \end{bmatrix}\mathbf{M}^\mathrm{T}, \tag{5}$$

where $\mathbf{M}$ is an orthogonal matrix whose columns are the eigenvectors of $\hat{\mathbf{C}}$, and $\mathbf{D}$ is a diagonal matrix consisting of the corresponding eigenvalues. The matrix $\mathbf{M}^\mathrm{T}$ defines a rotation that will transform $\hat{\mathbf{C}}$ to the circle $\mathbf{D}$ with radius $r = \sqrt{-\frac{b}{a}}$ centered at the origin. This transformation corresponds to rotating the camera about its center until its principle axis passes through the sphere center. This reduces the general case to the previously described special case, and the distance d between the camera center and the sphere center can be recovered in terms of r and R. Finally, the sphere center can be recovered as

$$\begin{aligned} \mathbf{S}_\mathrm{c} &= \mathbf{M}[0\ 0\ d]^\mathrm{T} \\ &= d\mathbf{m}_3, \end{aligned} \tag{6}$$

where $\mathbf{m}_3$ is the third column of $\mathbf{M}$.

3 Estimation of Light Directions

Suppose the center $\mathbf{S}_\mathrm{c}$ of a sphere with known radius R have been estimated using the method described in the previous section, it is then straightforward to recover the light directions from the observed highlights on the sphere. Standard

techniques begin by first constructing a ray from the camera center through a pixel corresponding to a highlight. The intersections of this ray with the sphere are then located to determine the point on the sphere giving rise to the highlight. By using the property that the angle of the incoming light must equal the angle of the outgoing light to the camera at a surface point with highlight, the light direction $\mathbf{L}$ can be recovered as

$$\mathbf{L} = (2\mathbf{N} \cdot \mathbf{V})\mathbf{N} - \mathbf{V}, \tag{7}$$

where $\mathbf{V} = \frac{\mathbf{X}-\mathbf{O}}{|\mathbf{X}-\mathbf{O}|}$ is the unit viewing vector, $\mathbf{N} = \frac{\mathbf{X}-\mathbf{S}_\mathrm{c}}{|\mathbf{X}-\mathbf{S}_\mathrm{c}|}$ is the unit surface normal vector at $\mathbf{X}$, $\mathbf{X}$ is a point with specular highlight on the sphere and $\mathbf{O}$ is the camera center.

Now suppose the radius R of the sphere is unknown, it has been shown in Sect. 2 that there exists a one-parameter family of solutions for the sphere center $\mathbf{S}_\mathrm{c}$ which all lie on the straight line joining the camera center and the true sphere center. It will now be shown that the light direction recovered from an observed highlight using any of these scaled spheres will be identical. In other words, light directions can be recovered from the highlights observed in the image of a sphere without knowing its size and location.

Proposition 1. *Consider a ray casted from the camera center and the family of spheres with varying radius recovered from the conic image of a sphere. If this ray intersects any of these spheres, it will intersect all the spheres and the first point of intersection with each sphere will all have the same unit surface normal.*

Proof. Since the cone constructed from the camera center and the conic image of the sphere will be tangent to all the recovered spheres. Any ray lying within this cone will intersect all these spheres, whereas any ray lying outside this cone will intersect none of them.

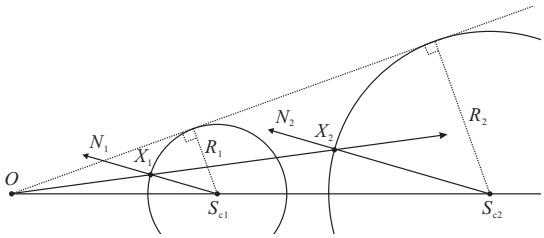

Fig. 2. The first intersection point of the ray with each sphere from the family of solutions will have the same unit surface normal.

To prove that the intersection points have the same unit surface normal, it is sufficient to consider only the cross-section containing both the ray and the line defined by the sphere centers (see Fig. 2). Without loss of generality, consider a sphere S_1 from the family, and let its radius and center be r_1 and

S_{c1} respectively. Suppose the ray intersect S_1 at X_1. The surface normal N_1 at X_1 is given by the vector $S_{c1}X_1$. Now consider a second sphere S_2 from the family, and let its radius and center be r_2 and S_{c2} respectively. A line being parallel to N_1 can be constructed from S_{c2}, and let the intersection point between this line and S_2 be X_2. By construction, the surface normal N_2 at X_2 will be parallel to N_1. Consider the two triangles $\triangle OS_{c1}X_1$ and $\triangle OS_{c2}X_2$. Obviously, $|X_1S_{c1}| : |X_2S_{c2}| = r_1 : r_2$. It follows from (4) that $|OS_{c1}| : |OS_{c2}| = r_1 : r_2$. Finally by construction, $\angle OS_{c1}X_1 = \angle OS_{c2}X_2$. Hence $\triangle OS_{c1}X_1$ and $\triangle OS_{c2}X_2$ are similar and $\angle S_{c1}OX_1 = \angle S_{c2}OX_2$. It follows that the ray will intersect S_2 at X_2 at which the surface normal N_2 is parallel to the surface normal N_1 at X_1. Since the two spheres being considered are chosen arbitrarily, the same argument applies to all spheres in the family, and the proof is completed. □

From (7), the light direction $\mathbf{L}$ only depends on the unit viewing vector $\mathbf{V}$ and the unit surface normal $\mathbf{N}$. The following corollary therefore follows immediately from Proposition 1:

Corollary 1. *The light direction estimated from an observed specular highlight in an image of a sphere will be independent of the radius used in recovering the location of the sphere center.*

4 Estimation of Camera Poses

Suppose two images of a sphere are captured at two distinct viewpoints. By applying the method described in Sect. 2 to each image independently, the sphere center can be recovered in each of the two camera-centered coordinate systems respectively. By assuming an arbitrary but fixed radius for the sphere in both views, it is possible to relate the two cameras in a common coordinate system. Without loss of generality, let the sphere center in the camera-centered coordinate system of the first view be $\mathbf{S}_{c1}$ and that of the second view be $\mathbf{S}_{c2}$ respectively. By considering a common world coordinate system centered at the sphere center, the projection matrices for the two views can be written as

$$\begin{aligned} \mathbf{P}_1 &= \mathbf{K}_1[\,\mathbf{I}\ \mathbf{S}_{c1}\,] \\ \mathbf{P}_2 &= \mathbf{K}_2[\,\mathbf{I}\ \mathbf{S}_{c2}\,]. \end{aligned} \tag{8}$$

Note that the above projection matrices are not unique. Due to the symmetry exhibited in the geometry of the sphere, an arbitrary rotation about the sphere center (i.e., the world origin) can be applied to the camera without changing the image of the sphere. This corresponds to rotating the camera around the sphere while keeping the cone constructed from the image tangent to the sphere. Hence, by choosing the first camera as a reference view, a more general form of the projection matrices for the two views is given by

$$\begin{aligned} \mathbf{P}_1 &= \mathbf{K}_1[\,\mathbf{I}\ \mathbf{S}_{c1}\,] \\ \mathbf{P}_2 &= \mathbf{K}_2[\,\mathbf{R}\ \mathbf{S}_{c2}\,], \end{aligned} \tag{9}$$

where $\mathbf{R}$ is a 3×3 rotation matrix with three degrees of freedom.

By assuming the light directions being fixed (globally) in both views, the highlights observed in the two images can be exploited to uniquely determine the relative rotation between the two cameras. Note that the location of the highlight on the sphere surface will depend on both the light direction and the viewpoint. Hence the locations of the highlights due to the same light direction will be different under two distinct viewpoints, and their projections on the two images do not provide a pair of point correspondences. Nonetheless, using the method described in Sect. 3, the light direction can be recovered in each of the two camera-centered coordinate systems. Without loss of generality, let the (unit) light direction in the camera-centered coordinate system of the first view be $\mathbf{L}_1$ and that of the second view be $\mathbf{L}_2$ respectively. Since these two directions are parallel in the common world coordinate system, the rotation matrix $\mathbf{R}$ relating the two cameras should bring $\mathbf{L}_1$ to $\mathbf{L}_2$, i.e.,

$$\mathbf{R}\mathbf{L}_1 = \mathbf{L}_2. \tag{10}$$

The above equation places two independent constraints on $\mathbf{R}$. Hence observing two highlights produced by two distinct light directions in two images will provide four constraints which are enough to determine $\mathbf{R}$ uniquely. Reader may refer to [14] for a robust method of determining a rotation from two of more pairs of directions using quaternion representation.

5 Experimental Results

The closed form solutions described in the previous sections for recovering light directions and camera poses have been implemented. Experiments on both synthetic and real data were carried out and the results are presented in the following sections.

5.1 Synthetic Data

The experimental setup consisted of a synthetic sphere being viewed by two synthetic cameras under two distinct directional lights. The conic images of the sphere were obtained analytically using (2). To locate the specular highlights in the images, the sphere was rendered by OpenGL with a Phong Shader using the viewing parameters of the two synthetic cameras (see Fig. 3). The specular highlights were then picked and matched manually, and a region growing technique was applied to extract the highlight regions from the images. The centroid of each region was then used as the highlight location for recovering the light directions.

In order to evaluate the robustness of the proposed method, uniformly distributed random noise was added to the conic images as well as to the locations of the specular highlights. To add noise to a conic, points were first sampled from the conic and perturbed in a radial direction from the conic center. A noisy conic

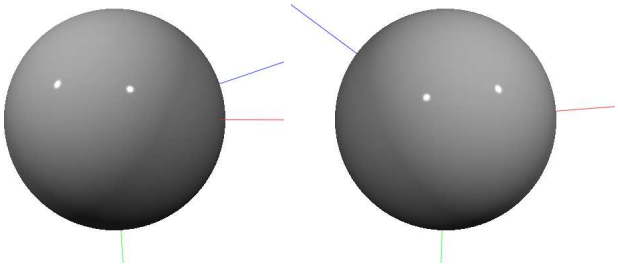

Fig. 3. Synthetic sphere rendered by OpenGL with a Phong Shader using the viewing parameters of the two synthetic cameras.

was then obtained as a conic robustly fitted to these noisy points using a direct least squares method [15]. For the location of a specular highlight, noise was added directly to its pixel coordinates.

Experiments on synthetic data with noise levels ranging from 0.1 to 3.0 pixels were carried out. For each noise level, 200 independent trials were conducted to estimate both the light directions and the camera poses from the noisy conics and highlights. Figure 4 shows a plot of the average angular error (in degrees) in the estimated light directions against the noise level (in pixels). It can be seen that the average error increases linearly with the noise level. For a noise level of 1.0 pixel, the average angular error in the estimated light directions is only about 0.5°. Figure 5 shows a plot of the average angular errors (in degrees) in the estimated rotations of the cameras against the noise level (in pixels). The rotation is represented here using a rotation axis, parameterized by the two angles azimuth and elevation, and a rotation angle theta around the axis. It can be seen again that the average angular errors increase linearly with the noise level. This is expected as the computation of the rotation depends directly on the estimated light directions. For a noise level of 1.0 pixels, the average angular errors are less than 0.75° for the angles defining the rotation axis and less than 0.5° for the rotation angle around the axis.

5.2 Real Data

In the real data experiment, five light sources were placed approximately 3 meters away from a red snooker ball which has a radius of around 54mm. Seven images of the snooker ball were then taken at distinct viewpoints (see Fig. 6). The ground truth projection matrices for these seven views were obtained using a planar calibration pattern [16]. The intrinsic parameters of the camera were obtained by decomposing the ground truth projection matrices. Alternatively, the intrinsic parameters of the camera may also be recovered using existing techniques based on the conic images of the sphere [17, 18]. Cubic B-spline snake was applied to extract the contours of the sphere in the images, and conics were then fitted to these contours using a direct least squares method [15]. Like in the synthetic experiment, the specular highlights were picked and matched manually

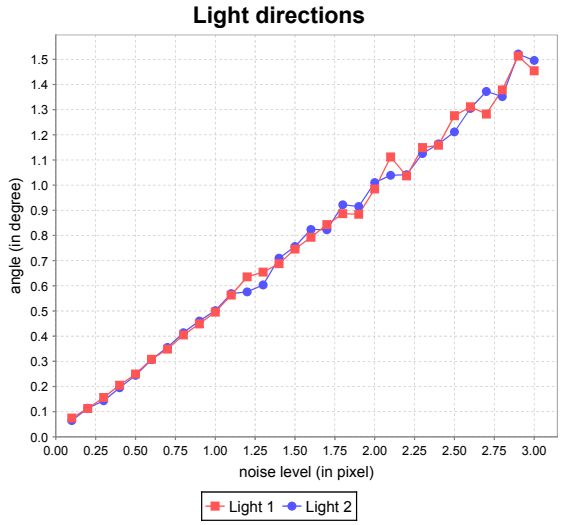

Fig. 4. A plot of the average angular error (in degrees) in the estimated light directions against the noise level (in pixels).

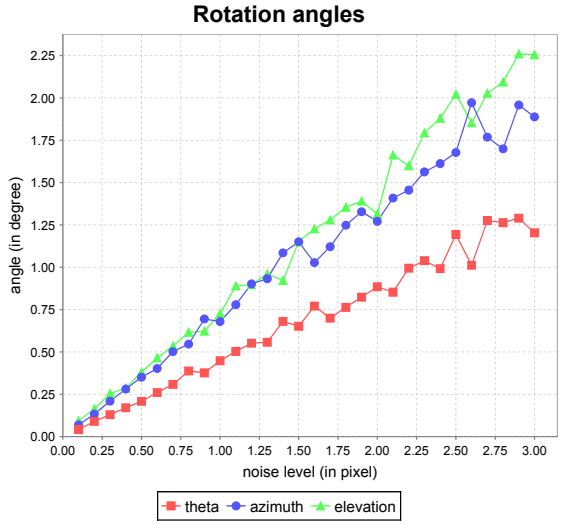

Fig. 5. A plot of the average angular errors (in degrees) in the estimated rotations of the cameras against the noise level (in pixels). The rotation is represented here using a rotation axis, parameterized by the two angles azimuth and elevation, and a rotation angle theta around the axis.

from the images and a region growing technique was applied to extract the highlight regions from the images. The centroid of each region was then used as the highlight location for recovering the light directions. Since the camera poses were estimated up to an unknown scale, it is not very meaningful to directly compare the relative translations between the estimated cameras with those of the ground truth. Instead, the relative rotations between the estimated cameras were compared against those of the ground truth, and the angular errors in the estimated rotations are listed in Table. 1. Similar to the synthetic experiment, the rotation is represented here using a rotation axis, parameterized by the two angles azimuth and elevation, and a rotation angle theta around the axis. It can be seen from the table that the maximum angular error is about 0.33°, while most errors lie between 0.1° and 0.2°.

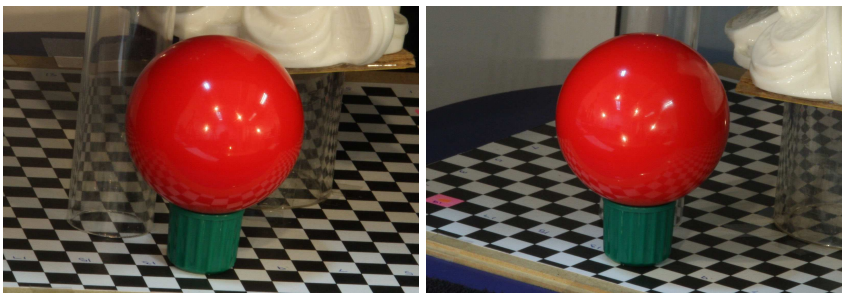

Fig. 6. Two images of a specular sphere taken under five distant point light sources.

Table 1. Angular errors (in degrees) in the estimated rotations.

view	theta	azimuth	elevation
1-2	0.2228	0.0178	0.1170
1-3	0.3264	0.0464	0.0089
1-4	0.1930	0.2250	0.1556
1-5	0.1587	0.1389	0.0197
1-6	0.0900	0.1079	0.0572
1-7	0.1020	0.1178	0.0110

6 Conclusions

This paper addresses the problem of recovering both the light directions and camera poses from a single sphere. The two main contributions of this paper are

1. a closed form solution for recovering light directions from the specular highlights observed in a single image of a sphere with unknown size and location; and
2. a closed form solution for recovering the relative camera poses using the estimated sphere and light directions.

It is shown that given the intrinsic parameters of a camera, a scaled sphere can be reconstructed from its image. The translation direction of the sphere center from the camera center can be determined uniquely, but the distance between them will be scaled by the unknown radius of the sphere. It is then proved that the light directions can be recovered independent of the radius chosen in locating the sphere. If the sphere is observed by multiple views, the sphere center recovered using a common fixed radius will fix the translations of the cameras from the sphere center. The relative rotations between the cameras can then be determined by aligning the relative light directions recovered in each view. As there exists closed form solutions for all the computation steps involved, the proposed method is extremely fast and efficient. Experiments on both synthetic and real images show promising results. With the proposed method, both the light directions and camera poses can be estimated simultaneously. This greatly eases the work of multiple views light estimation.

Acknowledgement

This project is supported by a grant from the Research Grants Council of the Hong Kong Special Administration Region, China, under Project HKU 7180/06E.

References

1. Horn, B.K.P., Brooks, M.J.: Shape and source from shading. In: International Joint Conference on Artificial Intelligence. 932–936
2. Pentland, A.P.: Finding the illuminant direction. Journal of Optical Soc. of America **72**(4) (April 1982) 448–455
3. Ikeuchi, K., Sato, K.: Determining reflectance properties of an object using range and brightness images. IEEE Trans. on Pattern Analysis and Machine Intelligence **13**(11) (November 1991) 1139–1153
4. Zheng, Q., Chellappa, R.: Estimation of illuminant direction, albedo, and shape from shading. IEEE Trans. on Pattern Analysis and Machine Intelligence **13**(7) (July 1991) 680–702
5. Yang, Y., Yuille, A.L.: Sources from shading. In: Proc. Conf. Computer Vision and Pattern Recognition. (1991) 534–539
6. Zhang, Y.F., Yang, Y.H.: Multiple illuminant direction detection with application to image synthesis. IEEE Trans. on Pattern Analysis and Machine Intelligence **23**(8) (August 2001) 915–920
7. Takai, T., Niinuma, K., Maki, A., Matsuyama, T.: Difference sphere: An approach to near light source estimation. In: Proc. Conf. Computer Vision and Pattern Recognition. Volume I. (2004) 98–105

8. Wang, Y., Samaras, D.: Estimation of multiple directional light sources for synthesis of augmented reality images. Graphical Models **65**(4) (2003) 185–205
9. Debevec, P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with gobal illumination and high dynamic range photography. In: Proc. ACM SIGGRAPH. (1998) 189–198
10. Powell, M.W., Sarkar, S., Goldgof, D.: A simple strategy for calibrating the geometry of light sources. **23**(9) (September 2001) 1022–1027
11. Zhou, W., Kambhamettu, C.: Estimation of illuminant direction and intensity of multiple light sources. In: Proc. 7th European Conf. on Computer Vision. Volume IV. 206–220
12. Li, Y.Z., Lin, S., Lu, H.Q., Shum, H.Y.: Multiple-cue illumination estimation in textured scenes. In: Proc. 9th Int. Conf. on Computer Vision. (2003) 1366–1373
13. Hartley, R.I., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, Cambridge, UK (2000)
14. Horn, B.: Closed-form solution of absolute orientation using unit quaternions. Journal of Optical Soc. of America A **4**(4) (1987) 629–642
15. Fitzgibbon, A.W., Pilu, M., Fisher, R.B.: Direct least square fitting of ellipses. IEEE Trans. on Pattern Analysis and Machine Intelligence **21**(5) (May 1999) 476–480
16. Zhang, Z.: A flexible new technique for camera calibration. IEEE Trans. on Pattern Analysis and Machine Intelligence **22**(11) (November 2000) 1330–1334
17. Agrawal, M., Davis, L.S.: Camera calibration using spheres: a semi-definite programming approach. In: Proc. 9th Int. Conf. on Computer Vision. (2003) 782–789
18. Zhang, H., Wong, K.Y.K., Zhang, G.: Camera calibration from images of spheres. IEEE Trans. on Pattern Analysis and Machine Intelligence **29**(3) (March 2007) 499–503