Ring-Light Photometric Stereo

Zhenglong Zhou and Ping Tan

Department of Electrical & Computer Engineering, National University of Singapore

Abstract. We propose a novel algorithm for uncalibrated photometric stereo. While most of previous methods rely on various assumptions on scene properties, we exploit constraints in lighting configurations. We first derive an ambiguous reconstruction by requiring lights to lie on a view centered cone. This reconstruction is upgraded to Euclidean by constraints derived from lights of equal intensity and multiple view geometry. Compared to previous methods, our algorithm deals with more general data and achieves high accuracy. Another advantage of our method is that we can model weak perspective effects of lighting, while previous methods often assume orthographical illumination. We use both synthetic and real data to evaluate our algorithm. We further build a hardware prototype to demonstrate our approach.

1 Introduction

Photometric stereo algorithms [1] reconstruct local surface orientations (i.e. normal directions) from multiple images captured at a fixed viewpoint and variant illumination conditions. Most of these algorithms assume the illumination conditions are recorded during data capturing so that the normal directions are uniquely determined. However, capturing illumination conditions is often tedious, requiring the insertion of additional calibration objects such as mirrored spheres into the scene. These calibration objects can further cause inter-reflections that are often not modeled in photometric stereo algorithms and therefore increase reconstruction error.

There is a series of works [2,3,4,5,6,7,8,9,10,11] studying this problem without recording illumination conditions, known as uncalibrated photometric stereo. Almost all these methods rely on various assumptions about scene properties such as integrable surface, non-Lambertian materials, inter-reflections, six normals of equal albedo or small albedo entropy.[1] Hence, these methods can work for certain types of scenes that meet their assumptions, but cannot handle other types. For example, the gift box shown in Figure 1 (a) contains a few discrete planes with only three different normal directions and no significant non-Lambertian reflection. Notice that a plane does not provide integrable constraint as a linearly transformed plane is also integrable. Hence, all these previous methods will fail on this simple example. Figure 1 (b) is another challenging data which contains many depth discontinuities. Methods based on integrability must first identify

[1] An exception is Hayakawa's work [2] that used six lights with equal intensity to partially solve the problem.

K. Daniilidis, P. Maragos, N. Paragios (Eds.): ECCV 2010, Part II, LNCS 6312, pp. 265–279, 2010.

these discontinuities which is a non-trivial task. Indeed previous uncalibrated photometric stereo algorithms mainly focused on a single segmented, smoothly curved object. Little work has been proposed to handle challenging data like those shown in Figure 1.

We propose to study uncalibrated photometric stereo by exploiting constrains in lighting configurations such that our method can be applied to more general data. We consider the case where a scene is illuminated by directional lights located on a view centered cone as illustrated in Figure 2 (a). We show that with at least five lights on such a cone, surface normal directions of a Lambertian scene can be recovered up to two kinds of rotations, and a scaling compounded with a mirror ambiguity. These ambiguities can be resolved if additional constraints are available, such as three lights of equal interval, five lights of equal intensity, surface integrability, non-Lambertian reflectance or corresponding normals from multiple viewpoints. To handle more general data, we choose to combine constraints derived from lighting configurations to achieve an Euclidean reconstruction. All we require about the scene is that two corresponding normals can be identified from two views, a constraint which can be easily satisfied for most inputs. We use synthetic and real data to evaluate our algorithm and build a prototype device to demonstrate potential applications.

2 Related Work

We first briefly review uncalibrated photometric stereo methods. Hayakawa [2] showed that surface normals can be recovered up to a general linear transformation if lighting directions are unknown. If one can identify six lights with equal intensity, or six normals with equal albedo, this general linear ambiguity can be reduced to a 3D rotation ambiguity. This approach can hardly handle surfaces with smooth varying texture or scenes with only a few different normals like the gift box example in Figure 1.

Most of the works in uncalibrated photometric stereo follow the seminal work by Belhumeur et al.[3] that proved the linear ambiguity can be reduced to a generalized bas-relief (GBR) ambiguity by surface integrability. Since then, many

(a) (b)

Fig. 1. Challenging data for uncalibrated photometric stereo. (a) is too simple and (b) is too complicate for most of existing methods.

works have been proposed to study and resolve this ambiguity. Drbohlav and Chantler [4,5] showed spike-specular reflectance can resolve the GBR ambiguity. Tan et al. [9,10] further proved any homogenous isotropic reflectance can resolve it. The GBR ambiguity can also be resolved by inter-reflections [6] and minimizing the entropy of surface albedos [8]. All these methods share a common limitation that depth discontinuities must be identified before integrability can be applied to obtain a reconstruction up to the GBR ambiguity. However, this identification of depth discontinuities is nontrivial in practice. Typically, a mask image is provided to separate the object from its background and the whole object surface is assumed to be integrable. This approach cannot handle complicated scenes like the one in Figure 1 (b). Furthermore, a piecewise planar scene, like Figure 1 (a), does not provide integrability constraints, because a plane is always integrable after any linear transformation. Hence, these algorithms often require a pre-segmented, smoothly curved surface.

Different from these previous works, we exploit partial information in the lighting conditions to resolve the shape ambiguity. Our method makes little assumption about the scene property. Hence, our method can be applied to more general data which cannot be handled by previous methods. Similar illumination configuration has been used [12] to minimize the reconstruction error due to camera sensor noise when lighting directions are known. Alldrin and Kriegman [13,14] also used the same configuration with known lighting. However, [13] recovers only partial surface geometry and [14] requires much more (about 100) input images. In comparison, our method requires only five images and our lighting directions are unknown. Our method is also related to those works that combine photometric stereo and structure-from-motion[15,16,17]. These methods assume the surface is differentiable and are difficult to be applied to complicated shapes like Figure 1 (b).

3 Ring-Light Photometric Stereo

Uncalibrated photometric stereo algorithms typically do not assume any prior knowledge about lighting conditions. In this section, we show that if the illumination is partially known, i.e. directional lights lying on a view centered cone, the problem can be significantly simplified. We first briefly review the shape ambiguity in uncalibrated photometric stereo. Then we show that lights lying on a view centered cone significantly reduce the ambiguity. At last, we describe several ways to resolve the remaining ambiguities.

3.1 Uncalibrated Photometric Stereo

We first briefly review the factorization based formulation of uncalibrated photometric stereo. Suppose F images are captured for a Lambertian surface under a variant directional lighting and each image contains P pixels. Ignoring shadows, inter-reflections and non-Lambertian effects, we can formulate the image intensity matrix I as $I = NL$. Here, I is a $P \times F$ matrix formed by pixel intensities. N and L are $P \times 3$ and $3 \times F$ matrices respectively. Each row of N indicates the

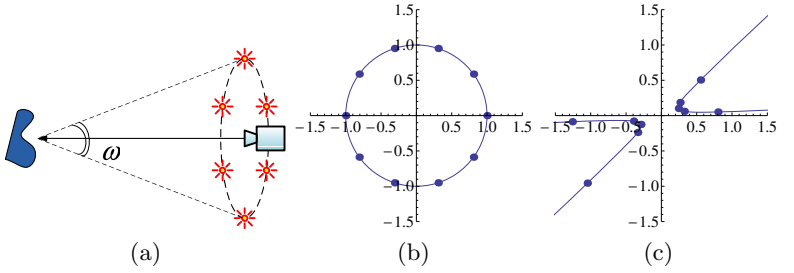

Fig. 2. Ring-light photometric stereo. (a) Lighting directions lie on a view centered cone. The term ω denotes the cone opening angle. (b) In the projective plane, these lights lie on a ring centered at origin (i.e. viewing direction). (c) When there is a linear ambiguity, these lights lie on a general planar conic. Our algorithm resolves this linear ambiguity by mapping lights back to their canonic positions.

scaled surface normal (unit surface normal multiplied with albedo), and each column of L is the scaled lighting direction (unit lighting direction multiplied with its intensity). In uncalibrated photometric stereo, only I is known and both N and L are unknown. Applying singular value decomposition (SVD), the matrix I can be decomposed as:

$$I = UDV^\top = (UD^{1/2})(D^{1/2}V^\top) = \hat{N}\hat{L}. \tag{1}$$

$\hat{N}, \hat{L}$ could differ from their true values by an arbitrary 3×3 invertible matrix A since $\hat{N}\hat{L} = \hat{N}A^{-1}A\hat{L}$. The autocalibration of photometric stereo amounts to recover A. Once A is estimated, the true surface normals and lighting directions can be computed as $N = \hat{N}A^{-1}$, $L = A\hat{L}$.

3.2 Constraints from a Ring-Light

Suppose the lights are distributed on a cone centered at the viewing direction as shown in Figure 2 (a). We follow the work [10] to analyze the problem in the projective plane where a lighting direction (l_x, l_y, l_z) is considered as a point $(l_x/l_z, l_y/l_z)$. We choose a world coordinate system such that the viewing direction is $(0, 0, 1)$ and corresponds to the origin in the projective plane. In the projective plane, the true lighting directions should lie on a circle centered at origin as shown in Figure 2 (b). This circle can be denoted by a diagonal matrix $C = diag(s^2, s^2, -1)$ and $C = S^\top C_u S$. Here, $C_u = diag(1, 1, -1)$ is the unit circle and $S = diag(s, s, 1)$ is a uniform scaling matrix. The SVD based reconstruction Equation (1) recovers lighting and normal directions up to an arbitrary invertible linear transformation A. The estimated lights form a general conic $\hat{C} = A^\top C A$ in the projective plane as shown in Figure 2 (c). Hence, we can resolve the ambiguity A by mapping $\hat{C}$ back to C. In this subsection, we first reduce the ambiguity by mapping $\hat{C}$ to the unit circle C_u. The remaining ambiguities are resolved in Section 3.3.

It is well known [18] that a conic can be computed from five points on it. Hence, we first use five estimated lighting directions to fit the conic $\hat{C}$ which

is a 3×3 symmetric matrix. We can apply SVD again to compute a linear transformation B that maps $\hat{C}$ to C_u, i.e.

$$\hat{C} = UDU^\top = (UD_1^{1/2})C_u(D_1^{1/2}U^\top) = B^\top C_u B.$$

Here, $D_1^{1/2} C_u D_1^{1/2} = D$. Then the lighting and surface normal directions can be updated accordingly by $\tilde{L} = B\hat{L}, \tilde{N} = \hat{N}B^{-1}$. Now, the general linear ambiguity is reduced and the estimated lights $\tilde{L}$ are on a view centered ring in the projective plane. But two kinds of ambiguities remain. First, the scaling matrix S between C and C_u is still unknown. Second, B can only be estimated up to a circle invariant transformation P that maps C_u to C_u. In other words, there could be an ambiguity matrix P such that $B^\top C_u B = B^\top P^\top C_u P B$. The following proposition specifies the structure of P.

Proposition 1: If a 3×3 linear transformation P maps the unit circle C_u to itself, i.e. $P^\top C_u P = C_u$, then P can be decomposed as $P = M^n R_\phi H_t R_\theta$, $n = 1 \text{ or } 2$. Here, M is a mirror transformation about the y axis, R_ϕ, R_θ are rotations in the plane (centered at origin), and H_t is a hyperbolic rotation, i.e.

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, R_\theta = \begin{pmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, H_t = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cosh t & \sinh t \\ 0 & \sinh t & \cosh t \end{pmatrix}. \quad (2)$$

R_ϕ has the same form as R_θ. Please refer to the appendix for a proof of this proposition. By this proposition, P is a compounded ambiguity that includes ordinary and hyperbolic rotations and a mirror transformation.

In the next section, we will discuss these ambiguities in more detail and propose methods to resolve them. Here we summarize these ambiguities by the following equation. The general conic $\hat{C}$ can be decomposed as:

$$\begin{aligned} \hat{C} &= B^\top C_u B = B^\top P^\top C_u P B \\ &= B^\top P^\top S^{-\top} C S^{-1} P B = A^\top C A \end{aligned} \quad (3)$$

Here, B is known, P and S are unknown transformations. Once P, S are determined, we can resolve the general linear ambiguity A. In the following, we refer to the compounded ambiguity $S^{-1}P$ as the *ring-light ambiguity*. It is also called the *ring-light transformation* depending on the context. The auto-calibration of ring-light photometric stereo amounts to estimate this compound transformation to upgrade the reconstruction $\tilde{L}, \tilde{N}$ to Euclidean as: $L = S^{-1}P\tilde{L}, N = \tilde{N}P^{-1}S$.

3.3 Ring-Light Ambiguities

We first briefly study each component of the *ring-light ambiguity* and later propose methods to solve it. Figure 3 summarizes these components and their geometric implications. The ambiguity S is a scaling in the projective plane which corresponds to the classic bas-relief ambiguity. M flips the normal and lighting directions vertically. It corresponds to the convex vs. concave ambiguity along the vertical direction. R_θ rotates the lighting and normal directions around the

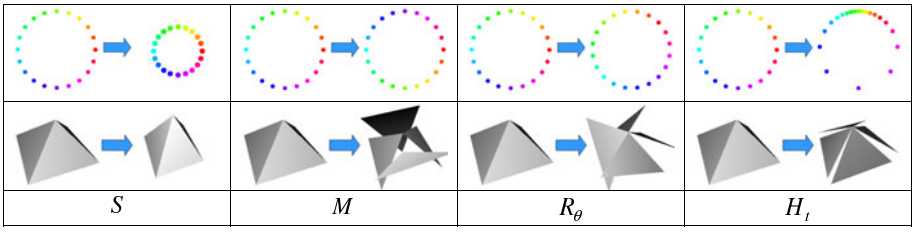

Fig. 3. Geometric explanations of the components of the 'ring-light ambiguity'. The first row shows the transformations induced to lighting directions in the projective plane. The second row illustrates the corresponding transformations to a 3D shape.

origin. It preserves all origin centered circles and could map a continuous shape to a discontinuous one. H_t is a hyperbolic rotation that preserves the unit circle. The relative positions of points on the unit circle are changed after a hyperbolic rotation as shown in Figure 3. It could also map continuous shapes to discontinuous ones.

In the following, we show various priors that resolve these ambiguities. We first discuss some widely used priors and later introduce three novel priors.

Integrability: Surface integrability is a widely used scene prior to resolve the ambiguity in uncalibrated photometric stereo. If the scene is known to be integrable, the linear ambiguity A can be reduced to a GBR ambiguity [3]. The intersection of the GBR transformation group with the ring-light transformation contains only the classic bas-relief transformation. Hence, if applicable, integrability resolves all the other components except the scaling S.

Points with Equal Albedo: Hayakawa [2] showed that six general normals with the same albedo can reduce the linear ambiguity to a 3D rotation compounded with a mirror reflection. The intersection of this ambiguity with the ring-light transformation contains only the planar rotation R_ϕ compounded with M. Hence, this prior reduces the ring-light ambiguity to a planar rotation with a mirror reflection.

Lights with Equal Intensity: Hayakawa's method [2] can also be applied to six general lights with equal intensity. However, since our lights lie on a view centered cone, constraints derived this way are degenerated. Both S, M and a 3D rotation cannot be resolved (explained in the next section). Hence, it reduces the ring-light ambiguity to a planar rotation R_ϕ compounded with a scaling S and a mirror M.

Lights with Equal Interval: If lights are uniformly distributed over the view centered cone, all lighting directions are determined up to a planar rotation (about the cone axis) and a scaling (corresponding to the unknown cone opening angle). Hence this constraint can reduce the linear ambiguity to a planar rotation R_ϕ compounded with a scaling S.

Multiple Viewpoint: Suppose a surface is observed from two different viewpoints with known relative motion and some corresponding points can be identified among these views. If the surface normals of both views are reconstructed up to some ambiguity, these corresponding points give constraints to resolve these ambiguities. In next section, we show that two corresponding normals from two views can resolve a planar rotation R_ϕ and a scaling S in both views.

Clockwise/Counter-Clockwise Lighting: M causes a vertical flipping of the estimated lighting and normal directions. If the lights on the ring are turned on one by one in clockwise or counter-clockwise, M reverses this order. Hence, M can be resolved if the order of lighting is known beforehand.

4 A Complete Stratified Reconstruction

We combine some of the discussed priors to achieve a Euclidean reconstruction. Those priors derived from lighting configurations are favored to handle more general scenes. We propose two methods to reduce the linear ambiguity to a planar rotation compounded with a scaling. In the next, we employ constraints derived from corresponding normals in different views to resolve the remaining ambiguities. For this stratified reconstruction, all we need are images from two viewpoints and five lights of equal interval/intensity distributed clockwise (or counterclockwise) on a view centered cone for each viewpoint.

4.1 Lights with Equal Interval

Suppose we know the order of lights (clockwise or counterclockwise). All lighting directions are determined up to the unknown cone opening angle and a planar rotation. We can assume arbitrary values of these two parameters to get pseudo lighting directions $\check{L}$ up to a scaling S (corresponding to the cone opening angle) and a planar rotation R_ϕ (corresponding to the rotation about the cone axis). We can recover normal directions up to the same ambiguity according to $\check{N} = I\check{L}^{-1}$. However, as we will see in experiments, this approach generates larger errors. Hence, we derive a more sophisticated approach in the following.

4.2 Lights with Equal Intensity

We first apply the ring-light constraint described in Section 3.2 to reconstruct normal directions up to a *ring-light ambiguity*. Then we apply the equal lighting intensity constraint to reduce the remaining ambiguities to a mirror transformation M, a planar rotation R_θ compounded with a scaling S. Afterwards, we use the known lighting order (clockwise in our experiments) to resolve M.

After applying the ring-light constraint, the estimated lighting direction $\tilde{\mathbf{l}}$ lies on the unit circle in the projective plane and is related to the true lighting direction $\mathbf{l}$ by $\mathbf{l} = S^{-1}P\tilde{\mathbf{l}}$. Suppose 5 lights are known to have equal intensity, we obtain

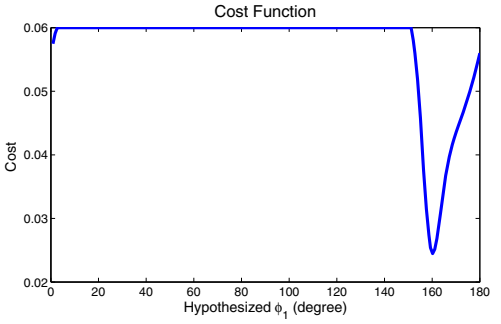

Fig. 4. The cost as a function of the hypothesized ϕ_1. This functions has a clear global minimum because ϕ_1, ϕ_2, s_1, s_2 are uniquely determined in principle.

$$\begin{aligned} k_1 &= \mathbf{l}_i{}^\top \mathbf{l}_i = \tilde{\mathbf{l}}_i{}^\top P^\top S^{-\top} S^{-1} P \tilde{\mathbf{l}}_i \\ &= \tilde{\mathbf{l}}_i{}^\top R_\theta{}^\top H_t{}^\top S^{-2} H_t R_\theta \tilde{\mathbf{l}}_i \qquad i = 1, 2, \cdots 5. \end{aligned} \tag{4}$$

It is easy to verify that M and R_ϕ are both eliminated from the equation. Here, k_1 is an unknown constant indicating the lighting intensity and i is an index of the lights. Let $F = R_\theta{}^\top H_t{}^\top S^{-2} H_t R_\theta$. Then Equation (4) is a linear equations about F, i.e. $\tilde{l}_i{}^\top F \tilde{l}_i = k_1$.

Hayakawa [2] used six such equations from different lighting directions to solve F. However, in our problem there are at most five independent linear equations because of the special configuration of lights. More specifically, $\tilde{\mathbf{l}}' \doteq H_t R_\theta \tilde{\mathbf{l}}$ must lie on the unit circle on the projective plane, because $\tilde{\mathbf{l}}$ lie on the unit circle which is invariant under H_t and R_θ. Hence, no matter what $S = diag(s, s, 1)$ is the expression, $\tilde{\mathbf{l}}_i{}^\top F \tilde{\mathbf{l}}_i = \tilde{\mathbf{l}}_i'{}^\top S^{-2} \tilde{\mathbf{l}}_i'$ is always a constant. In other words, S cannot be recovered from Equation (4) if these lights all lie on a view centered cone. To provide an experimental validation, we uniformly sample 360 lights on the unit circle. The six singular values of all these 360 equations are 17.72, 6.70, 6.70, 0.89, 0.63, 0.00. This suggests one degree of freedom of F cannot be determined.

Hence, we can only solve the 1D null space of F as $k_1 F_1 + k_2 F_2$. Here, F_1, F_2 satisfy $\tilde{\mathbf{l}}_i{}^\top F_1 \tilde{\mathbf{l}}_i = 1$ and $\tilde{\mathbf{l}}_i{}^\top F_2 \tilde{\mathbf{l}}_i = 0$ respectively, k_1 is the unknown but fixed constant and k_2 can vary to generate the whole 1D null space. We substitute $F = k_1 F_1 + k_2 F_2$ into $F = R_\theta{}^\top H_t{}^\top S^2 H_t R_\theta$. We solve s, t, θ, k_1 for any given k_2 according to the formulas provided in Appendix B. It can be verified that the solutions of t and θ are independent of k_2, while k_1 and s vary according to k_2. Hence, we obtain a unique solution of H_t and R_θ but cannot determine S, and the original *ring-light ambiguity* is reduced to M, S and R_ϕ. The result of this subsection is summarized into the following proposition.

Proposition 2: If five lights with equal intensity can be identified, the *ring-light ambiguity* can be reduced to a mirror transformation, a planar rotation compounded with a scaling.

4.3 Two Corresponding Normals in Two Views

We further exploit the constraints from multiple views. Suppose $\mathbf{n}_1$ and $\mathbf{n}_2$ are two corresponding normals in different views. They are defined in their local camera coordinate system and are related by the relative rotation between the two cameras, i.e. $\mathbf{n}_1 = T\mathbf{n}_2$. The relative rotation T can be computed separately, for example, by structure-from-motion. Suppose $\tilde{\mathbf{n}}_1, \tilde{\mathbf{n}}_2$ are the estimated normals which are subject to a planar rotation R_ϕ and scaling S. We have the following equations:

$$\mathbf{n}_1 \simeq S_1 R_{-\phi 1} \tilde{\mathbf{n}}_1 \qquad \mathbf{n}_2 \simeq S_2 R_{-\phi 2} \tilde{\mathbf{n}}_2 \qquad \mathbf{n}_1 = T\mathbf{n}_2. \tag{5}$$

Here, $\simeq$ means equal up to a scale. Hence,

$$\tilde{\mathbf{n}}_1 \simeq R_{\phi 1} S_1^{-1} T S_2 R_{-\phi 2} \tilde{\mathbf{n}}_2. \tag{6}$$

Let $E = R_{\phi 1} S_1^{-1} T S_2 R_{-\phi 2}$. We get $\tilde{\mathbf{n}}_1 \simeq E\tilde{\mathbf{n}}_2$. This equation provides two independent constraints. Hence, the four ambiguities $S_1, S_2, R_{\phi 1}, R_{\phi 2}$ can be resolved from two corresponding normals in two views.

Equation (6) can be written as $\tilde{\mathbf{n}}_1 \times E\tilde{\mathbf{n}}_2 = 0$, where $\times$ is the vector cross product. This vector equation expands to the following three equations:

$$s_2 \mathcal{A}^{(1)}(\phi_1, \phi_2) + \mathcal{B}^{(1)}(\phi_1) + s_1 s_2 \mathcal{C}^{(1)}(\phi_2) + s_1 \mathcal{D}^{(1)} = 0 \tag{7}$$

$$s_2 \mathcal{A}^{(2)}(\phi_1, \phi_2) + \mathcal{B}^{(2)}(\phi_1) + s_1 s_2 \mathcal{C}^{(2)}(\phi_2) + s_1 \mathcal{D}^{(2)} = 0 \tag{8}$$

$$s_2 \mathcal{A}^{(3)}(\phi_1, \phi_2) + \mathcal{B}^{(3)}(\phi_1) = 0. \tag{9}$$

Here, $\mathcal{D}^{(i)}$ are constants and $\mathcal{A}^{(i)}, \mathcal{B}^{(i)}$ and $\mathcal{C}^{(i)}$ are polynomials of trigonometrical functions of ϕ_1, ϕ_2.

$$\mathcal{A}^{(i)}(\phi_1, \phi_2) = a_1^{(i)} cos\phi_1 cos\phi_2 + a_2^{(i)} sin\phi_1 cos\phi_2 + a_3^{(i)} cos\phi_1 sin\phi_2 + a_4^{(i)} sin\phi_1 sin\phi_2$$

$$\mathcal{B}^{(i)}(\phi_1) = b_1^{(i)} cos\phi_1 + b_2^{(i)} sin\phi_1 \qquad \mathcal{C}^{(i)}(\phi_2) = c_1^{(i)} cos\phi_2 + c_2^{(i)} sin\phi_2$$

Here, $a_j^{(i)}, b_j^{(i)}$ and $c_j^{(i)}$ are all constants. These constants are provided in Appendix C.

Given two pairs of corresponding normals, it is nontrivial to derive an analytic solution for s_1, s_2, ϕ_1 and ϕ_2. We apply a 1D search for ϕ_1. For each hypothesized value of ϕ_1, ϕ_2 and s_2 can be easily solved from Equation (9) of both pairs. Then Equation (7) and Equation (8) from both pairs give totally 4 results for s_1. We use the consistency of these four values to choose the optimal ϕ_1 and its associated ϕ_2, s_2, s_1. In principle these four parameters are uniquely determined, so this 1D search has a global minimum and is robust as indicated in Figure 4. The result of this subsection is summarized in the following proposition.

Proposition 3: Given partial reconstructions of surface normals up to a planar rotation and a scaling from two views, if two pairs of corresponding normals can be identified, the reconstructions in both views can be upgraded to Euclidean.

5 Experiments

We apply our method to the challenging data shown in Figure 1. As explained earlier, these two examples cannot be handled by previous methods because they are either too simple (too few normals and planar surfaces) or too complicated (too many depth discontinuities). Our method first recovers a normal map up to the *ring-light ambiguity* as shown in the left column of Figure 5. Here, the x,y,z components of a normal direction are linearly encoded into the R,G,B color channels. This result is then upgraded to Euclidean by constraints derived from equal lighting intensity as shown in the middle. The right is a validation computed by calibrated photometric stereo where a metal sphere is used to record lighting directions. The difference between our results and the calibrated method is small. Some artifacts of the recovered normals on the box surface are due to the inaccuracy in radiometric calibration and inter-reflections.

Some additional results are shown in Figure 6. From left to right, we show one of the input images, our reconstructed surface normals and ground truth (obtained by calibrated photometric stereo). Some of the artifacts are due to non-Lambertian effects like shadow and highlight which are not modeled in our method. To handle shadows and highlights, we use simple intensity thresholding to exclude points with non-Lambertian effects. Our method is applied to Lambertian pixels to calibrate lighting directions. Then non-Lambertian pixels are processed with recovered lighting directions.

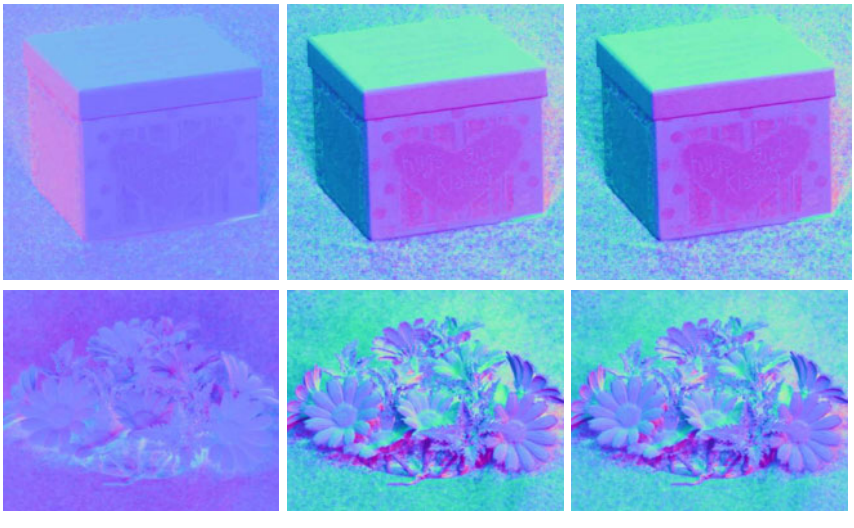

Fig. 5. Results for the challenging data in Figure 1. On the left are results up to the ring-light ambiguity. In the middle is our reconstructed surface normals. For a validation, we calibrate all incident lighting directions with a metal sphere and use calibrated photometric stereo to compute a ground truth in the right. Our result is very consistent to the ground truth.

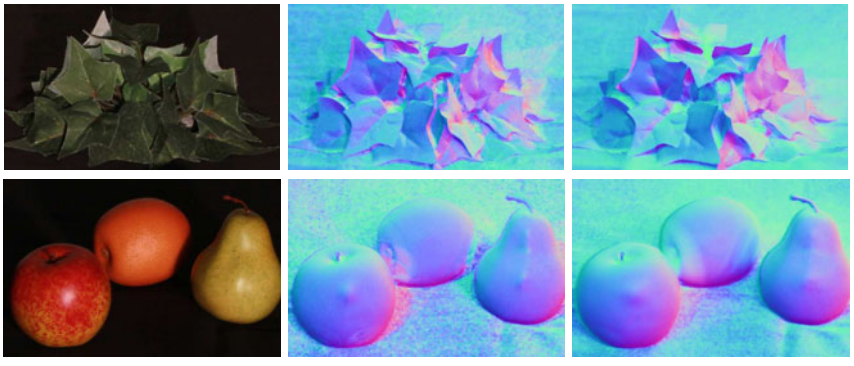

Fig. 6. Additional results. From left to right, they are one of the input images, our reconstructed surface normals, ground truth (by calibrated photometric stereo). Some of the artifacts are due to non-Lambertian effects like shadow and highlight.

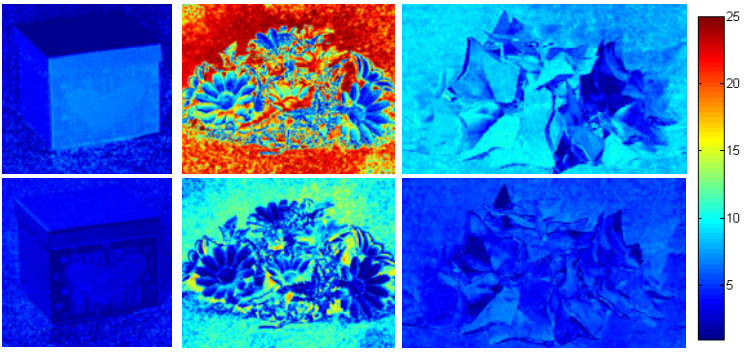

Fig. 7. The first and the second row are the angular errors of reconstructed normal directions by using equal lighting interval and equal lighting intensity constraint respectively. Typically, equal lighting intensity constraint generates more accurate results.

We also compare the two approaches to reduce the general linear ambiguity to a planar rotation and a scaling. The first and the second row of Figure 7 show the angular errors of reconstructed normal directions by equal lighting interval and equal lighting intensity constraints respectively. In the first row, the average angular errors are 5.8, 16.4 and 7.5 degrees from left to right. In the second row, these errors are 3.0, 6.0 and 4.4 degrees respectively. Please notice that normals in the background (a black cloth on table to reduce inter-reflection) are very noisy which increase the average angular error by 0.5-1 degrees in general. In our experiments, we find the constraints derived from lights with equal intensity are often more reliable. The flower example has larger error in both methods due to its strong shadowing and inter-reflection.

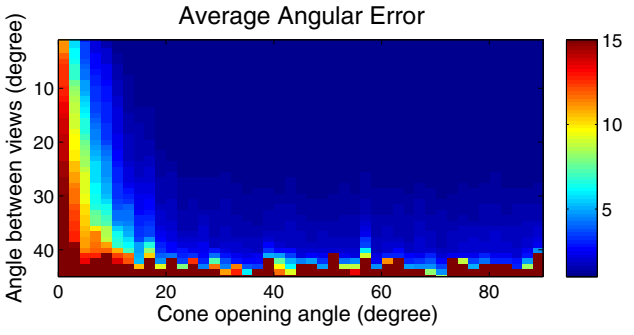

Fig. 8. Averaged angular error in the recovered normal directions as a function of the cone opening angle and the angle between two viewpoints. In most of time, the reconstruction error is smaller than 5 degrees.

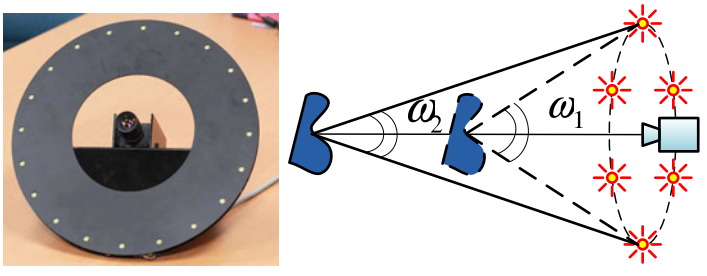

(a) (b)

Fig. 9. Shown in (a) is a prototype device. 20 LED bulbs lie on a circle with radius of 150 millimeters centered at the viewing direction. Our method allows us to consider the weak perspective effects of the lighting which is critical for a handheld photometric stereo setup operating at relatively small distance. This weak perspective effects is illustrated in (b). To ensure the opening angle of the cone is larger than 10 degrees, the distance between the camera and captured objects should be within 1.7 meters.

Next, we use a synthetic scene containing two spheres to evaluate our system under various conditions. Images are synthesized at 840×560 resolution. The images are contaminated by Gaussian noise with zero mean and standard deviation 0.01 (pixel values are within [0,1]). We synthesize the scene from two viewpoints and at each viewpoint 10 lighting directions are generated on a view centered cone. Zero mean Gaussian noise with standard deviation of 0.5 pixel is added to the true corresponding pixel positions. We evaluate the reconstruction accuracy with respect to different values of the cone opening angle ω and the angle between the two viewing directions. The average angular error in reconstructed normal directions is shown in Figure 8. In most of the cases, the reconstruction is quite good with average error smaller than 5 degrees.

5.1 A Prototype Device

We manufacture a prototype device for ring-light photometric stereo according to our evaluation on synthetic data. The device is shown in Figure 9 (a) and consists 20 LED bulbs that are synchronized with a video camera to capture photometric stereo image sequences. The radius of the plate is 150 millimeters. Hence, according to Figure 8, the operation distance of the device should be less than 1.7 meters (cone opening angle larger than 10 degrees) to ensure the reconstruction accuracy. This device is similar to those handheld photometric stereo setups proposed in [15,16,17]. The advantage of our method is that our algorithm handles more general data and allows us to consider the weak perspective effects of lighting as illustrated in Figure 9 (b). We consider the lighting directions depend on the operation distance, e.g. $\omega_1 \neq \omega_2$. This effect is important when the operation distance is relatively small. A consequence is that this device cannot be pre-calibrated, because the incident lighting directions changes when the operation distance changes. For example, we pre-calibrate lighting directions for an operation distance of about 0.6 meters and apply it to the operation distance of about 0.8 meters. This incorrect pre-calibration causes average angular error on the box scene as large as 8.5 degrees. Almost three times larger than the 3.0 degrees error when our method is applied.

6 Conclusion

We have presented an stratified method for ring-light photometric stereo. We have shown that five lights on a view centered cone reduce the general linear ambiguity to two rotations, one mirror reflection compounded with a scaling. If these lights have equal intensity or equal interval, this compound *ring-light ambiguity* can be reduced to a planar rotation plus a scaling. If two corresponding normals from two viewpoints can be identified, Euclidean reconstruction can be obtained. Different from previous works on uncalibrated photometric stereo, we minimize the restriction on scene properties. Hence, our method can be applied to more general scenes. We also built a prototype device to demonstrate our method.

References

1. Woodham, R.: Photometric stereo: A reflectance map technique for determining surface orientation from image intensities. In: Proc. SPIE 22nd Annual Technical Symposium, pp. 136–143 (1978)
2. Hayakawa, H.: Photometric stereo under a light source with arbitrary motion. J. Optical Society of America A 11, 3079–3089 (1994)
3. Belhumeur, P.N., Kriegman, D.J., Yuille, A.L.: The bas-relief ambiguity. Int. Journal of Computer Vision 35, 33–44 (1999)
4. Drbohlav, O., Sara, R.: Specularities reduce ambiguity of uncalibrated photometric stereo. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2351, pp. 46–60. Springer, Heidelberg (2002)

5. Drbohlav, O., Chantler, M.: Can two specular pixels calibrate photometric stereo? In: Proc. ICCV, vol. 2, pp. 1850–1857 (2005)
6. Chandraker, M.K., Kahl, F., Kriegman, D.: Reflections on the generalized bas-relief ambiguity. In: Proc. of CVPR (2005)
7. Basri, R., Jacobs, D., Kemelmacher, I.: Photometric stereo with general, unknown lighting. Int. J. Comput. Vision 72, 239–257 (2007)
8. Alldrin, N., Mallick, S.P., Kriegman, D.J.: Resolving the generalized bas-relief ambiguity by entropy minimization. In: Proc. of CVPR (2007)
9. Tan, P., Mallick, S.P., Quan, L., Kriegman, D.J., Zickler, T.: Isotropy, reciprocity and the generalized bas-relief ambiguity. In: Proc. of CVPR (2007)
10. Tan, P., Zickler, T.: A projective framework for radiometric image analysis. In: Proc. of CVPR (2009)
11. Shi, B., Matsushita, Y., Wei, Y., Tan, P.: Self-calibrating photometric stereo. In: Proc. of CVPR (2010)
12. Drbohlav, O., Chantler, M.: On optimal light configurations in photometric stereo. In: Proc. ICCV (2005)
13. Alldrin, N., Kriegman, D.: Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In: Proc. ICCV (2007)
14. Alldrin, N., Zickler, T., Kriegman, D.: Photometric stereo with non-parametric and spatially-varying reflectance. In: Proc. of CVPR (2008)
15. Lim, J., Ho, J., Yang, M.H., Kriegman, D.: Passive photometric stereo from motion. In: Proc. ICCV (2005)
16. Joshi, N., Kriegman, D.: Shape from varying illumination and viewpoint. In: Proc. ICCV (2007)
17. Higo, T., Matsushita, Y., Joshi, N., Ikeuchi, K.: A hand-held photometric stereo camera for 3-d modeling. In: Proc. ICCV (2009)
18. Coxeter, H.S.M.: Introduction to Geometry, 2nd edn. Wiley, Chichester (1989)

A Appendix A: Proof of Proposition 1

Proposition 1: If a 3×3 linear transformation P maps the unit circle C_u to itself, i.e. $P^\top C_u P = C_u$, then P can be decomposed as $P = M^n R_\phi H_t R_\theta$,$n = 1$ or 2.

Proof: Our proof is based on the following two lemmas:

Lemma 1: If a conic C is mapped to another conic C' by a projective transformation P, then P maps the interior/exterior of C to the interior/exterior of C'.

Lemma 2: Suppose A and A' are two points on two different conics C and C'. B, B' lies inside of C, C' respectively. Then there are precisely two projective transformations which map C to C', A to A', and B to B'.

These lemmas can be found in [18]. In the following, for a general linear transformation P that maps C_u to C_u, we assume the pre-images of $(1, 0, 1)$ and $(0, 0, 1)$ are A and B respectively. We explicitly derive two transformations P_1, P_2, $P_1 \neq P_2$, with the form $M^n R_\phi H_t R_\theta$ that maps A, B to $(1, 0, 1)$ and $(0, 0, 1)$ respectively. Then according to Lemma 2, we know Proposition 1 is true.

According to the Lemma 1, B is a point within C_u. So we can denote B as $(rcos\theta, rsin\theta, 1)$, where $0 < r < 1$. It is easy to verify that $H_t R_{\pi/2-\theta}$ maps the

point B to the origin. Here, t is uniquely decided by $r = -sinh(t)/cosh(t)$. It is also easy to verify that $H_t R_{\pi/2-\theta}$ maps A to another point A' on the circle. We can denote A' as $(cos\phi, sin\phi, 1)$. Then a rotation $R_{-\phi}$ will maps A' to the point $(1, 0, 1)$ and keep the origin invariant. As a result, we get the following transformation $P_1 = R_{-\phi} H_t R_{\pi/2-\theta}$, that maps B to $(0, 0, 1)$ and A to $(1, 0, 1)$. Note that, we can define$P_2 = M R_{-\phi} H_t R_{\pi/2-\theta}$. P_2 should also maps B to origin and A to $(1, 0, 1)$. Further, $P_1 \neq P_2$. Hence, according to Lemma 2, they are the only two transformations that map A, B to $(1, 0, 1)$ and $(0, 0, 1)$ respectively.

B Appendix B: Determine t, s from F

θ can be directly computed from F, $\quad \theta = \arctan\left(-F_{13}/F_{23}\right)$
$k_1$can be solved from equation $(a^2 - b^2 - c^2)k_1^2 - (a + 3c)k_1 - 2 = 0$
where $a = \frac{1}{2}\left(F_{11} + F_{22}\right) + \frac{3}{2}F_{33} \quad b = \frac{1}{2}\left(F_{11} + F_{22} - F_{33}\right) \quad c = \frac{2F_{23}}{\cos\theta} = -\frac{2F_{13}}{\sin\theta}$
$s^{-2} = \frac{1}{2}\left(k_1(F_{11} + F_{22} - F_{33}) + 1\right)$
$t = \frac{1}{2} arcsinh\left(\frac{2k_1 F_{23}}{\cos\theta(s^{-2}+1)}\right) = \frac{1}{2} arccosh\left(\frac{k_1(F_{11}+F_{22}+F_{33})-s^{-2}}{s^{-2}+1}\right)$

C Appendix C: Constants in Equation 7–9

$$
\begin{aligned}
&T = \{t_{ij}\}_{3\times 3} \\
&a_1^{(1)} = -t_{21}n_{21}n_{13} - t_{22}n_{22}n_{13} \quad a_1^{(2)} = +t_{11}n_{21}n_{13} + t_{12}n_{22}n_{13} \\
&a_2^{(1)} = +t_{11}n_{21}n_{13} + t_{12}n_{22}n_{13} \quad a_2^{(2)} = +t_{21}n_{21}n_{13} + t_{22}n_{22}n_{13} \\
&a_3^{(1)} = -t_{22}n_{21}n_{13} + t_{21}n_{22}n_{13} \quad a_3^{(2)} = +t_{12}n_{21}n_{13} - t_{11}n_{22}n_{13} \\
&a_4^{(1)} = +t_{12}n_{21}n_{13} - t_{11}n_{22}n_{13} \quad a_4^{(2)} = +t_{22}n_{21}n_{13} - t_{21}n_{22}n_{13} \\
&a_1^{(3)} = +t_{21}n_{21}n_{11} + t_{22}n_{22}n_{11} - t_{11}n_{21}n_{12} - t_{12}n_{22}n_{12} \\
&a_2^{(3)} = -t_{11}n_{21}n_{11} - t_{12}n_{22}n_{11} - t_{21}n_{21}n_{12} - t_{22}n_{22}n_{12} \\
&a_3^{(3)} = +t_{22}n_{21}n_{11} - t_{21}n_{22}n_{11} - t_{12}n_{21}n_{12} + t_{11}n_{22}n_{12} \\
&a_4^{(3)} = -t_{12}n_{21}n_{11} + t_{11}n_{22}n_{11} - t_{22}n_{21}n_{12} + t_{21}n_{22}n_{12} \\
&b_1^{(1)} = -t_{23}n_{23}n_{13} \quad b_2^{(1)} = +t_{13}n_{23}n_{13} \quad b_1^{(2)} = +t_{13}n_{23}n_{13} \quad b_2^{(2)} = +t_{23}n_{23}n_{13} \\
&b_1^{(3)} = +t_{23}n_{23}n_{11} - t_{13}n_{23}n_{12} \quad b_2^{(3)} = -t_{13}n_{23}n_{11} - t_{23}n_{23}n_{12} \\
&c_1^{(1)} = +t_{31}n_{21}n_{12} + t_{32}n_{22}n_{12} \quad c_1^{(2)} = -t_{31}n_{21}n_{11} - t_{32}n_{22}n_{11} \\
&c_2^{(1)} = +t_{32}n_{21}n_{12} - t_{31}n_{22}n_{12} \quad c_2^{(2)} = -t_{32}n_{21}n_{11} + t_{31}n_{22}n_{11} \\
&c_1^{(3)} = c_2^{(3)} = \mathcal{D}^{(3)} = 0 \quad \mathcal{D}^{(1)} = +t_{33}n_{23}n_{12} \quad \mathcal{D}^{(2)} = -t_{33}n_{23}n_{11}
\end{aligned}
\tag{10}
$$