

Calibrating a Non-isotropic Near Point Light Source using a Plane

Jaesik Park^{1*} Sudipta N. Sinha² Yasuyuki Matsushita³ Yu-Wing Tai¹ In So Kweon¹

¹Korea Advanced Institute of Science and Technology, Republic of Korea

²Microsoft Research Redmond, USA ³Microsoft Research Asia, China

Abstract

We show that a non-isotropic near point light source rigidly attached to a camera can be calibrated using multiple images of a weakly textured planar scene. We prove that if the radiant intensity distribution (RID) of a light source is radially symmetric with respect to its dominant direction, then the shading observed on a Lambertian scene plane is bilaterally symmetric with respect to a 2D line on the plane. The symmetry axis detected in an image provides a linear constraint for estimating the dominant light axis. The light position and RID parameters can then be estimated using a linear method. Specular highlights if available can also be used for light position estimation. We also extend our method to handle non-Lambertian reflectances which we model using a biquadratic BRDF. We have evaluated our method on synthetic data quantitatively. Our experiments on real scenes show that our method works well in practice and enables light calibration without the need of a specialized hardware.

1. Introduction

Calibration of light sources is often an important first step in various photometric methods in computer vision. For example, shape-from-shading [11], photometric stereo [24] and reflectometry [16] methods often require accurate knowledge of light sources for recovering surface normal maps and reflectances. While estimating an incident light direction from a distant source [27, 23, 1] is well studied and the problem of *near* light position estimation has also been analyzed [2, 12, 28], the topic of modeling a near point light source with a *non-isotropic* radiance distribution has received less attention in our field. Accurate calibration of a non-isotropic near light is important for disambiguating causes behind light attenuation in a scene and can reduce

unmodeled errors that may be present when simpler light models are used.

Near point light source calibration is usually performed in controlled scenes using specialized equipments. In particular, accurate recovery of luminous or radiant intensity distribution of non-isotropic lights requires special-purpose equipments such as an Imaging Sphere [17] or a goniophotometer [20]. However, such hardware based approaches are generally expensive in both cost and time.

In this paper, we explore an image-based approach for calibrating a near point light source that is rigidly attached to a conventional camera, that does not require special calibration objects or a special acquisition phase. We address both geometric and photometric calibration issues. The geometric aspects include estimating the position and orientation of the light whereas photometric calibration involves estimating the radiant intensity distribution (RID) of the light, which we assume is radially symmetric with respect to its dominant direction, *i.e.*, axis of the light source.

We show that under this assumption, the shading produced by a non-isotropic point light on a Lambertian plane is bilaterally symmetric whereas the shading produced by an isotropic point light is radially symmetric. Given a calibrated camera and a known plane, detecting the axis of bilateral symmetry or the center of radial symmetry in an image provides linear constraints for estimating the light axis and position, respectively. We model the RID using a low-order polynomial and show that if the light position and axis are known, the unknown RID coefficients can be efficiently estimated. Also, non-Lambertian surfaces of a single material can be handled by our method using a more a flexible reflectance model, and we show such an example using a biquadratic BRDF model [19].

Based on these observations, we propose a practical technique to calibrate a light-camera rig using a plane that recovers all unknown parameters of the light model by analyzing the shading and specular reflections observed in images captured from multiple viewpoints. When combined with structure-from-motion [9] for calibrating the camera and estimating the plane, our method can perform a full self-calibration of the light-camera rig in an unknown scene.

*Part of this work was done while the first author was visiting Microsoft Research Redmond as a research intern. This work was also supported in part by the National Research Foundation grant funded by the Korea government (MSIP) (No. 2010-0028680). Yu-Wing Tai was supported by the MSIP, Korea, under the IT/SW Creative research program supervised by the NIPA (NIPA-2013-H0503-13-1011)

2. Related Work

Various aspects of light calibration have been studied in the computer vision literature. Debevec [7] used a mirror sphere to capture an environment map for photo-realistic rendering and relighting where all incoming distant illumination was modeled. For a single distant light source as in photometric stereo [24], the problem reduces to that of estimating the light direction. This is possible in a lab setup [5], using reflective objects [25, 8] or by self-calibration [4]. Recently, Abrams *et al.* [1] treat the sun as a distant light source and extend photometric stereo to outdoor cameras.

Unlike distant lights, near light sources cause the incident light directions to vary within a scene. Although some methods model area light sources [27], it is common to assume that the light source is a point thereby reducing the problem to that of estimating the 3D point position. Objects of known geometry e.g. spheres are used to further simplify the problem [15, 21, 26]. Spherical probes have also been used for multiple, mixed (directional and near) light sources [28]. Ackermann *et al.* [2] proposes a new technique for accurately modeling reflections on a sphere. Other methods also use either reflective spheres [12, 23, 18] or specially designed planar light probes [3] or analyze shading of known geometric objects [22]. The case of a near point light rigidly attached to a camera has been analyzed for near light photometric stereo [10] and the light fall-off effect due to distance from the light is modeled. Light fall-off can provide cues for depth estimation [13]. However, such methods will suffer in accuracy with non-isotropic lights unless their intensity distribution is also modeled.

Certain aspects of non-isotropic near lights have been studied for light source design and optimization [20, 14]. The geometric light source model now requires an additional parameter - the dominant axis of the light; a direction that provides a reference point for storing the non-uniform radiant intensity distribution on a sphere centered at the light position. However, most existing geometric methods that use reflective calibration objects cannot be used for this model. In our method, we propose a general approach for recovering this light axis from shading cues. Our approach does not use a specific calibration object but instead requires multiple images of a plane. We propose a simple method based on detecting symmetries in shading as well as specular highlights if available. Once the light source geometry is known, we show that it is possible to estimate the radiant intensity distribution (RID) from the same input images.

To the best of our knowledge, our work is the first to address the estimation of geometric (position and orientation of light source w.r.t. camera) and photometric (RID curve) properties of a non-isotropic near point light source using a pure image-based approach. We show that our method holds promise for calibrating regular cameras that contain a built-in flash.




Figure 1. Polar graphs of $E(\phi)$ radiant intensity distribution curves. Some common shapes – (a) circular, (b) petal, (c) bell, (d) heart, that can be modeled using cubic polynomials.

3. Shading Model and Properties

Consider the case of a homogeneous textureless plane illuminated by a light source observed by a calibrated camera as shown in Fig. 2a. Let $\Pi = (\mathbf{n}, b)$ denote the plane, where $\mathbf{n} \in \mathbb{R}^3$ is its surface normal and $b \in \mathbb{R}$ is an offset, such that any 3D point $\mathbf{P} \in \mathbb{R}^3$ on the plane satisfies the equation $\mathbf{n}^\top \mathbf{P} + b = 0$. Let \mathbf{C} denote the camera projection center and I denote the observed image. We denote the light position as $\mathbf{L} \in \mathbb{R}^3$, its dominant direction as \mathbf{v} , and the angle between the light vector \mathbf{l} and \mathbf{v} as ϕ . Vectors \mathbf{n} , \mathbf{l} and \mathbf{v} are unit vectors. This is true for other vectors used elsewhere in the paper unless specified otherwise.

Assuming Π to be a Lambertian surface that is illuminated by a nearby, isotropic point light source, the intensity at the pixel p in the image I is given by

$$I_p = \rho E \frac{\mathbf{n}^\top \mathbf{l}}{d^2}, \quad (1)$$

where ρ is the diffuse albedo, E is the light radiance emitted towards the point \mathbf{P} that is observed at pixel p and d is the distance from the light source to \mathbf{P} . The image I is related to I' , the image projected on the plane Π via the homography induced by Π [9]. We assume that vignetting, gamma correction and radial distortion correction have been performed, and image intensities are in a linear space.

3.1. Isotropic and non-isotropic point lights

An *isotropic* point light source emits light equally in all direction. Hence, the term E in Eq. (1) is a constant. The shading at a scene point depends only on the incident light direction and the squared distance to the light.

In general, point lights are often non-isotropic, *i.e.*, they have a dominant light direction or axis and emit different amounts of light in different directions. *Radiant intensity* is defined as the radiant power of the light source emitted in a particular direction and it is measured in Watts per steradians. Varying radiant intensities of a light source are represented by its *radiant intensity distribution* (RID) and this is often visualized as a polar graph to display the angular light




Figure 2. (a) Geometry of a camera observing a plane illuminated by a near, directional point light source. (b,c) Supporting figures for explaining the symmetry properties (in Sec. 3.2).

fall-off from the light's axis. Figure 1 shows common RID curves as described in detail in [20].

For point lights, it is common to assume that the RID curve is radially symmetric (see [14]). As a result the term E in Eq. (1) becomes a function of the angle ϕ between the light vector l and the light's dominant axis v , and will hence be denoted as $E(\phi)$. The observed image intensities are now a function of the angular variation of radiance, varying incident light directions and the distance to the light.

3.2. Symmetry Properties

We now prove a key result concerning symmetry in the shading observed on a plane illuminated by a point light.

Lemma 1 (Bilateral Symmetry) *Let X be the point on the plane Π closest to the light L and Y be the point where the light axis meets Π . If the radiant intensity distribution is radially symmetric about the light axis, then the shading image generated on the plane will be bilaterally symmetric with respect to the line through X and Y .*

Proof of Lemma 1. Let Π_{LXY} denote the plane determined by the three points L , X and Y , and let l_{XY} denote the line where Π_{LXY} intersects Π . Since X is the closest point on Π to L , the line l_{LX} must be perpendicular to Π . Therefore, Π_{LXY} must be orthogonal to Π (Fig. 2b).

Consider any two points P_1 and P_2 on Π which are equidistant from l_{XY} and the line segment $\overline{P_1P_2}$ is perpendicular to l_{XY} . Let, M be the midpoint on $\overline{P_1P_2}$ (Fig. 2c). Since Π_{LXY} and Π are orthogonal to each other, we have $\overline{P_1P_2} \perp \overline{LM}$. Then, triangles $\triangle LMP_1$ and $\triangle LMP_2$ must be congruent (using the SAS rule for congruent triangles). This gives two equality constraints.

$$\|\overline{P_1L}\| = \|\overline{P_2L}\|, \quad (2)$$

$$\angle LP_1M = \angle LP_2M. \quad (3)$$

Similarly, triangles $\triangle YMP_1$ and $\triangle YMP_2$ are also congruent (result of the SAS rule). Therefore, we have $\|\overline{P_1Y}\| = \|\overline{P_2Y}\|$. This leads to the fact that triangles $\triangle LYP_1$ and $\triangle LYP_2$ are also congruent (follows from the SSS rule for congruent triangles). Therefore we have,

$$\angle P_1LY = \angle P_2LY. \quad (4)$$

Now, consider the intensities at P_1 and P_2 according to the shading model (Eq. (1)). Condition (2) implies that P_1 and P_2 are equidistant from L , so the denominators $d_{P_1}^2$ and $d_{P_2}^2$ are equal. Condition (3) implies that the light vectors l_{P_1} and l_{P_2} at P_1 and P_2 form the same angle with the plane normal, so the dot products must be equal, *i.e.*, $n^\top l_{P_1} = n^\top l_{P_2}$. Condition (4) implies that the light vectors at P_1 and P_2 form the same angle with the light axis, and as the RID is radially symmetric, we have $E_{P_1} = E_{P_2}$.

From this, we conclude that $I'_{P_1} = I'_{P_2}$. As this holds true for any pair of points located equidistant from l_{XY} and its line segment is perpendicular to l_{XY} , the image I' is bilaterally symmetric with respect to that line. \square

Lemma 2 (Radial Symmetry) *The shading from the light L on the plane Π is radially symmetric with respect to X , the closest point on Π from L when either the light is isotropic or the light is non-isotropic but is oriented such that its dominant axis is perpendicular to the plane.*

Proof of Lemma 2. For an isotropic light, consider points on Π at a constant distance from L . Any point P in this set lies on a circle C centered at X . The light vector at P *i.e.* the line l_{LP} traces out a right circular cone with axis l_{LX} and C as its base. This implies that the terms $n^\top l_P$ and d_P are constant for all points P on C . As E is a global constant, it implies image intensities at all points on C must be equal making the image radially symmetric w.r.t. X .

For a non-isotropic light, when the light axis v is perpendicular to the plane, the point Y , where the axis meets the plane coincides with X . Points on a circle centered at X are again equidistant from L and their light vectors form the same angle with the plane normal. Hence, $E(\phi)$ is constant for all points on the circle which implies that the observed intensities at these points are equal and the image is radially symmetric with respect to the center X . \square

3.3. Geometric Constraints

The symmetry properties imply that for a calibrated camera and a known plane illuminated by an unknown light




Figure 3. (a) Half-vector \mathbf{h} and angles θ_h and θ_d at the point \mathbf{X} on the plane Π . (b) Light camera configuration for the case when \mathbf{C} is not coplanar with \mathbf{L} , \mathbf{X} and \mathbf{Y} . The lines l_{LC} and l_{XY} meet on Π at an angle ψ (see Sec. 3.4).

source, detecting the axis of bilateral symmetry in the image provides a *plane constraint* that is equivalent to one linear constraint on the coordinates of \mathbf{L} . Alternatively, the center of radial symmetry if detected, provides a *line constraint* that provides two linear constraints on \mathbf{L} . In Sec. 4.2, we describe our algorithm that uses these constraints to estimate the light position and orientation.

3.4. Non-Lambertian surfaces

We now address the non-Lambertian case and analyze how the symmetry deviates from the perfect case depending on the camera-light geometry or the type of BRDF.

We replace $\rho(\mathbf{n}^T \mathbf{I})$ in Eq. (1) with $\rho(\theta_h, \theta_d)$, an isotropic BRDF model which is a function of θ_d , the angle between \mathbf{h} and l_{XL} and θ_h , the angle between \mathbf{n} and \mathbf{h} , where \mathbf{h} is the half-way vector [16, 19] (Fig. 3a). This biquadratic BRDF model is known to be appropriate for modeling low-frequency components of isotropic BRDFs [19].

Consider the plane Π_{LC} passing through \mathbf{L} and \mathbf{C} orthogonal to Π . Let $l_{XC'}$ denote the line where Π_{LC} intersects Π (Fig. 3b). For any two points on a line perpendicular to $l_{XC'}$ on Π , equidistant from $l_{XC'}$ and on opposite sides of it, their half-vectors form the same angle with the plane normal and their difference angles are also equal. So the BRDF map will be bilaterally symmetric with respect to $l_{XC'}$. Therefore, if lines $l_{XC'}$ and l_{XY} coincide, symmetry is preserved about this line for any isotropic BRDF. This occurs when the four points \mathbf{L} , \mathbf{X} , \mathbf{Y} and \mathbf{C} are coplanar. If there are not coplanar, the lines $l_{XC'}$ and l_{XY} do not coincide anymore but intersect at some arbitrary angle ψ at \mathbf{X} (see Fig. 3b). The degree of asymmetry due to the BRDF with respect to l_{XY} will be maximum when $\psi = \frac{\pi}{2}$.

The degree of asymmetry also depends on the ratio $\omega = \frac{\|\mathbf{LC}\|}{\|\mathbf{CC}'\|}$ where $\|\mathbf{LC}\|$ and $\|\mathbf{CC}'\|$ are distances between \mathbf{L} and \mathbf{C} and between \mathbf{C} and its closest point on Π . Consider the case where the distance between the camera and light position is similar to the distance of the camera from the plane. In this case, θ_d will vary more and this may cause a large variance in the observed BRDF.

To quantify the deviation in symmetry for a reasonable

Algorithm: Light Calibration

Input: Images $\{I_k\}$, Cameras $\{P_k\}$, plane Π

Output: Light position \mathbf{L} , direction \mathbf{v} in camera coordinates, parameters of $E(\phi)$.

1. Rectify $\{I_k\}$ using plane-induced homographies H .
 2. Detect specular highlights.
 3. Recover shading images.
 4. Detect axes of symmetry in shading images.
 5. Estimate \mathbf{L} and \mathbf{v} using output of steps 2 and 4.
 6. Estimate coefficients of $E(\phi)$ using estimated \mathbf{L} , \mathbf{v} .
-

range of configurations values for ψ and ω , we analyzed pose parameters for the camera-light rig *i.e.* \mathbf{C} and \mathbf{L} for values of $\psi \in (0, \frac{\pi}{2})$ and $\omega \in (\frac{1}{50}, 1)$. We considered realistic BRDF values from the database used in [19] in this analysis and empirically selected different BRDF functions to measure the deviation from symmetry in the average case. In our analysis, the deviation in symmetry for many BRDFs is small enough when ω is less than 0.1 *i.e.* the distance between \mathbf{L} and \mathbf{C} is less than one-tenth the approximate distance to the plane. The details of this analysis can be found in the supplementary material. Thus, we conclude that except for BRDFs with extreme values, we expect the deviation in symmetry to be small and that our method will work accurately in most configurations.

4. Proposed method

In this section, we describe our method for calibrating a point light rigidly attached to a camera using multiple images of a plane. See Algorithm 1 for an overview. We refer to the input images as $\{I_k\}$, the camera projection matrices as $\{P_k\}$ and camera centers as $\{\mathbf{C}_k\}$ respectively. For each camera position \mathbf{C}_k , the corresponding light position in the world coordinate system is denoted \mathbf{L}_k .

Camera Calibration. We use structure from motion [9] to calibrate the cameras and then estimate the plane by robustly fitting a plane to the reconstructed 3D points. We assume that the scene is planar, has a single material and has enough texture for automatic feature matching to work.

Specular Highlight Detection. We detect highlights in images by analyzing pixel intensity distributions. Images containing specular highlights tend to have long-tail distributions as pixels at and around the highlight are much more brighter than pixels in the rest of the image. On the other hand, image without highlights tend to have short-tail distributions. Hence, we construct a histogram of image intensities using all image pixels and fit a 1D Gaussian $N(\mu, \sigma)$ to the histogram values. Bright pixels, whose intensities exceed the value 6σ are classified as specular highlights and we fit a 2D ellipse to those pixels and use the ellipse center for light position estimation.

Gradient Filtering. Since our method requires a shading image, we first need to remove high frequency gradients in the input image that correspond to albedo edges. We recover the shading image \hat{I} by minimizing the following objective function

$$\hat{I} = \operatorname{argmin}_I \sum_{\mathbf{p} \in \mathcal{P}} \left\{ \left(\frac{\partial I_{\mathbf{p}}}{\partial \mathbf{x}} - f\left(\frac{\partial O_{\mathbf{p}}}{\partial \mathbf{x}}\right) \right)^2 + \lambda w_{\mathbf{p}} (I_{\mathbf{p}} - O_{\mathbf{p}})^2 \right\}, \quad (5)$$

where $f(\cdot)$ is gradient clipping function defined as

$$f(\mathbf{x}) = \begin{cases} \mathbf{x} & \text{if } \|\mathbf{x}\|_2 < \tau \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

and O is the input image. This problem can be formulated as a discrete Poisson problem which can be solved efficiently using a sparse linear solver. In Eq. (5), the first term encourages the gradients in I to match the clipped gradient of O whereas the second term encourages the intensities to be similar. The second term is also modulated by a per-pixel weight $w_{\mathbf{p}} \in [0, 1]$ that accounts for outliers arising from strong gradients in the albedo map or specular highlights. We define $w_{\mathbf{p}}$ as follows.

$$w_{\mathbf{p}} = s_{\mathbf{p}} (1 - |O_{\mathbf{p}} - G * O_{\mathbf{p}}|), \quad (7)$$

where G is a $\rho \times \rho$ gaussian kernel ($\rho = 100$), $*$ is the convolution operator and $s_{\mathbf{p}}$ is a binary variable (zero when the pixel is classified as a specular highlight). $w_{\mathbf{p}}$ is inversely proportional to per-pixel values in $|O_{\mathbf{p}} - G * O_{\mathbf{p}}|$, the high-pass filter response on O . Parameters λ and τ are set to 0.1 and 0.5 respectively for $I \in [0, 255]$. An example of a recovered shading image is shown in Fig. 4.

Ortho-Rectification: Using the known plane and camera parameters, we compute the plane-induced homographies for each image and rectify the shading image \hat{I}_k computed in the previous step to remove all perspective effects and create rectified images required by our method.

4.1. Bilateral Symmetry Detection

We now describe our method for detecting the symmetry axis in the rectified shading image using RANSAC in a typical hypothesize-and-verify framework. We use the fact, that for a least squares quadratic fit $g(z) = az^2 + bz + c$ to samples of a 1D function $f(z)$ symmetric about z^* , evaluated at points sampled at uniform steps in $z^* \pm D$, the function $g(z)$ will attain an extremum at z^* . However, this is not always true for an asymmetric 1D function.

To evaluate any line l for symmetry, we select a point x_j on l and sample points along l_j^\perp , the line perpendicular to l through x_j at uniform steps along l_j^\perp . We consider the intensity profile I_j^\perp at these sampled points between points z_j^+ and z_j^- that are equidistant from x_j but in opposite directions, such that either z_j^+ or z_j^- is on the image boundary.




Figure 4. (a) Input image. (b) Corresponding shading image computed by gradient filtering (Image contrast stretched for visualization). (c) Intensity plots for the red and blue scanlines. (d) Per-pixel weights $w_{\mathbf{p}}$ used in Eq. (7).

Using least squares, we then fit a quadratic function $Q_j(z)$ to I_j^\perp and compute the distance

$$d_j = \|z_j^* - x_j\|, \quad \text{where } z^* = \operatorname{argmax}_z Q_j(z). \quad (8)$$

To measure the overall degree of asymmetry, we perform quadratic fits at n positions along l and calculate the average distance $\bar{d}(l) = \frac{\sum_j d_j}{n}$ for all the positions. After generating random line hypotheses¹, we select the hypothesis \hat{l} with the smallest value of $\bar{d}(l)$. We further refine the estimate \hat{l} using local optimization. In near degenerate cases, bilateral symmetry degrades to radial symmetry and then, multiple line candidates may have small values of $\bar{d}(l)$. The center of symmetry can then be found by robustly intersecting multiple such symmetric lines found using RANSAC.

4.2. Light Position and Axis Estimation

Each detected symmetry axis provides a *plane* constraint on the light position in world coordinates (C_W). We transform these planes to camera coordinates C_C and derive a linear constraint on the coordinates of \mathbf{L} in C_C . These linear constraints $\pi_k \mathbf{L} = 0$ can be stacked into a matrix of the form $\mathbf{A}\mathbf{L} = \mathbf{b}$. It may appear that a solution of this linear system will provide an estimate of \mathbf{L} . Unfortunately, these linear equations are not independent even when the cameras are in general positions. This is because the planes Π_{LXY} for different light positions always form a bundle or pencil of planes and pass through a 3D line. In fact, that 3D line is the axis of the light, denote by \mathbf{R} in C_C . This implies that we cannot directly compute \mathbf{L} but it is possible to estimate \mathbf{R} by finding the null space of the linear system described

¹In practice, a 1-point RANSAC is used when a prior guess for the light position L is available, since the projection of L on Π lies on line l and only one other point is needed.

above. Let \mathbf{L}_1 and \mathbf{L}_2 denote 3D points in C_C whose linear combination defines points on the line \mathbf{R} .

In practice, we need at least 2 images, each of which provides one plane constraint. As long as these planes do not all coincide, there is no degeneracy. Multiple input images with random 3d camera orientations makes the degenerate case very unlikely. The presence of camera translation avoids the degeneracy in the initial SfM stage and the 3D plane fitting step.

If specular highlights are present, we can use them to estimate \mathbf{L} . Using the ellipse center as the image of the reflection of \mathbf{L} , we compute the corresponding incident light ray in C_W using the law of reflection. We transform this ray into C_C and obtain a *line* constraint on \mathbf{L} , which gives two linear constraints on its coordinates. We estimate \mathbf{L} using least squares assuming cameras in general position.

Alternate method. When specular highlights are absent (as in Lambertian scenes), it is possible to estimate \mathbf{L} indirectly. If the 3D line \mathbf{R} is known, we can parameterize the unknown position $\mathbf{L} = \mathbf{L}_0 + \lambda \mathbf{u}$, where \mathbf{L}_0 is a point on \mathbf{R} closest to the camera center \mathbf{C} , \mathbf{u} is a unit vector parallel to \mathbf{R} and λ is an unknown scalar. For each k -th plane, we can find the point \mathbf{Y}_k where \mathbf{R} intersects it. Then, the distance from \mathbf{L} to \mathbf{Y}_k denoted by d_k , can be written in terms of λ .

To compute λ , we can substitute various terms for points \mathbf{Y}_k into Eq. (1), namely the intensities $I_{\mathbf{Y}_k}$, known values of $\mathbf{n}^\top \mathbf{l}$ and expressions for d_k in terms of λ . As these points lie on the light axis, *i.e.*, $\phi = 0$, the term $E(\phi = 0)$ is unknown but constant. Taking ratios between pairs of equations eliminates ρ and $E(\phi = 0)$. Multiple equations provide a set of sample estimates of λ after which the median value λ^* is selected and used to compute $\mathbf{L} = \mathbf{L}_0 + \lambda^* \mathbf{u}$.

4.3. Radiant Intensity Distribution Estimation

Once the light position and axis is known, for each image pixel, we can substitute I and d^2 as well as values for ϕ , θ_h , θ_d and d into Eq. (1) where the product of the polynomials $E(\phi)$ and $\rho(\theta_h, \theta_d)$ is expanded in terms of the unknown coefficients q_1, \dots, q_6 and p_1, \dots, p_4 .

$$\rho(\theta_h, \theta_d)E(\phi) = (q_1 c_h^2 + q_2 c_h + q_3)(q_4 c_d^2 + q_5 c_d + q_6) / (p_1 \phi^3 + p_2 \phi^2 + p_3 \phi + p_4). \quad (9)$$

Here, $c_h = \cos(\theta_h)$ and $c_d = \cos(\theta_d)$ as proposed in [19]. On expanding $\rho(\theta_h, \theta_d)E(\phi)$, we get a higher-order polynomial $f(\phi, \theta_h, \theta_d)$ with 36 unknown coefficients. These coefficients of f , denoted by vector \mathbf{f} can be estimated by solving a least squares problem of the following form.

$$A\mathbf{f} = I d^2, \quad (10)$$

where A has 36 columns and a row corresponding to a pixel. To recover p_1, \dots, p_4 from \mathbf{f} , we create a 9×4 matrix F ,

$$F = [\mathbf{f}_1 p_1 \quad \mathbf{f}_2 p_2 \quad \mathbf{f}_3 p_3 \quad \mathbf{f}_4 p_4], \quad (11)$$




Figure 5. (a) Synthetic dataset example. (b) RID curves with different angular fall-offs used for evaluation (best seen in color).

where $\mathbf{f}_i \in \mathbb{R}^9$ are coefficients for p_i . On noise-free data, F would be a rank-1 matrix. Hence, we factorize $F = UDV^\top$ using singular value decomposition (SVD) and the values of p_1, \dots, p_4 are obtained from the first column of V .

5. Experiments

We have tested our method with both synthetic and real data. These results are now described.

5.1. Synthetic evaluation

We generated 50 synthetic datasets with 20 images each. Each set has randomly generated cameras approximately on a hemisphere oriented roughly towards the origin. The scene plane is fixed and passes through the origin. Suitable intrinsic parameters are chosen randomly as well. The light position and orientation was generated randomly in the camera coordinates. Each dataset was assigned one of the RID curves shown in Fig. 5b. Images were rendered at 320×240 resolution and a Blinn-Phong model [6] was used to generate specular highlights². Image noise was added to the synthetic images. The level of noise added was varied from 1 to 8 on a scale of $[0, 255]$ on all 50 datasets to generate results for 400 different runs; these are now discussed.

We first present the symmetry axis detection results in Fig. 6a where the detected line is compared with the ground truth line using the Hausdorff distance. In general the estimation is stable but the median errors across different runs increases slightly with noise. As expected the detection is consistently more accurate for lights that have a stronger angular fall-off in their RID (blue curves).

The light axis estimation which depends on the accuracy of the previous step also appears to be quite stable in general (Fig. 6b). The estimated direction is within 1° of ground truth for the blue curves and is at most 2.3 degrees for the light which has the most uniform distribution. The position estimation using specular highlights clearly outperforms the alternative method that seems to be affected by

²We use a very high value of the shininess parameter (= 50) and whenever the specular highlight is observed within the image, the view vectors are not at grazing angles to the plane. We confirmed that the error introduced by the Blinn-Phong model which is approximate, is quite small.




Figure 6. Median errors across multiple runs for (a) symmetry axis detection and for estimating (b) light direction, (c) light position using specular highlights and (d) light position using the alternative method. Results are color-coded by the RID curves shown in Fig. 5b.




Figure 7. Ground truth and estimated RID for each of the five light models (Fig. 5b). The medians of the estimated coefficients from 80 runs are used to plot the curve. The error bars show the deviation in the curve due to variance of each estimated coefficient.

Dataset	I_0	I_1	Est.	X	Y	Z	MAE
CAMERA-LED	39	3	Pos.	5.57	9.02	0.94	8.97
			Dir.	0.09	-0.05	0.99	
SLR-FLASH	17	3	Pos.	0.29	26.35	17.38	15.35
			Dir.	0.02	-0.25	0.97	

Table 1. Statistics for real datasets: #images used for calibration (I_0) and held-out for evaluation (I_1). Estimated light position and orientations are listed (position reported in cm. from camera center and direction reported as a unit vector).

noise levels (Fig. 6c and d). The results of the first method also does not depend on the RID curve, which is another advantage.

Figure 7 shows the ground truth RID curves and the estimated curves. In each case, the coefficients of the RID curve were estimated from multiple runs and the curve corresponding to the median of these estimates is shown. The deviation due to the 25-th and 75-th percentiles is indicated using error bars. In general, the synthetic evaluation shows that our method successfully estimates the RID curves with high accuracy. Difficult random camera configurations was a source of error in less than 5% of the cases – either the optical axis was at a grazing angle to the plane or the angle between the light vector and the optical axis was quite large. In the second case, our method suffers in accuracy due to the limited field of view (FoV).

5.2. Evaluation on Real datasets

To test on real data, we captured two datasets with 42 and 20 images of whiteboards (some of which contained specular highlights) in two different dark rooms. The first dataset CAMERA-LED was captured using a bright Cree XM-L LED, rigidly attached to the top of a point-and-shoot camera (Panasonic DMC-LX5). The second dataset SLR-FLASH was captured with a Canon Mark III-1D camera

with an auxiliary flash. Even though the flash violates the assumption of a point light source, we evaluate the performance of our method on this dataset to assess to what extent it can handle traditional camera flashes. After obtaining the SfM reconstruction, we resolved the scale ambiguity using prior knowledge of the scale. For these datasets, we randomly held out three images for evaluating estimated parameters. We generated images using estimated parameters and compared them to the held-out images using the mean absolute error (MAE)³ as the accuracy metric.

Figure 8 shows results on the CAMERA-LED and SLR-FLASH datasets. For the Cree XM-L LED, our RID estimate appears to be in reasonable agreement with the manufacturer’s profile within a cone of 25° from the light axis. For angles greater than 25°, the insufficient horizontal field-of-view (70° approx.) of the camera causes fewer pixels with larger angular values to be observed. Having the dominant axis of the light parallel to the camera’s optical axis also adds to the error for the same reason. Therefore, the full range of light emission may not be calibrated. However, calibrating the partial range is useful for post-processing images captured with the same camera and flash and that task does not require calibration of the full range of light emission. The MAE for the CAMERA-LED dataset was 8.97. The estimated geometric parameters of the light source is reported in Table 1. For SLR-FLASH, our method recovers a more flat RID curve as expected, since its built-in diffuser causes it to behave like an isotropic light source. The MAE in this case is 15.35. A likely reason for this is that the flash in this case is a near area-light source. Nevertheless, the geometric parameters are accurately estimated (Table 1) – the estimated light axis is almost parallel to the optical axis and the light position appears accurate to within 1-2 cms based on our best guess for the camera center.

³for image intensities that lie in the range [0,255]




Figure 8. Results on CAMERA-LED and SLR-FLASH (top and bottom row resp.): (a) Camera-light rig. (b) Input image. (c) Shading image generated using model. (d) Residuals and MAE between rendered and gradient-filtered image (not shown). (e) The recovered RID curve.

6. Conclusions

We have presented a practical method for calibrating a non-isotropic near point light rigidly attached to a camera. Our method uses images of a weakly textured planar scene from multiple viewpoints and estimates light position and orientation using shading cues and specular highlights. The results on real data are encouraging and it seems that the method can be feasible in situations where special hardware/objects cannot be used. We expect that higher accuracy can be obtained by jointly optimizing all model parameters in a final stage. The investigation of more accurate models for non-point near light source is another interesting direction of future work.

References

- [1] A. Abrams, C. Hawley, and R. Pless. Heliometric stereo: Shape from sun position. In *Proc. of ECCV*, pages 357–370, 2012.
- [2] J. Ackermann, S. Fuhrmann, and M. Goesele. Geometric point light source calibration. In *Vision, Modeling, and Visualization*, pages 161–168, 2013.
- [3] N. G. Alldrin and D. J. Kriegman. A planar light probe. In *Proc. of CVPR*, pages 2324–2330, 2006.
- [4] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *Int'l Journal of Computer Vision (IJCV)*, 72(3):239–257, 2007.
- [5] M. Ben-Ezra, J. Wang, B. Wilburn, X. Li, and L. Ma. An led-only brdf measurement device. In *Proc. of CVPR*, 2008.
- [6] J. F. Blinn. Models of light reflection for computer synthesized pictures. *Proc. of SIGGRAPH*, 11(2), July 1977.
- [7] P. Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. of SIGGRAPH*, pages 189–198, 1998.
- [8] O. Drbohlav and M. Chantler. Can two specular pixels calibrate photometric stereo? In *Proc. of ICCV*, pages 1850–1857, 2005.
- [9] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*, volume 2. Cambridge Univ Press, 2000.
- [10] T. Higo, Y. Matsushita, N. Joshi, and K. Ikeuchi. A hand-held photometric stereo camera for 3-d modeling. In *Proc. of ICCV*, pages 1234–1241, 2009.
- [11] B. K. P. Horn. Obtaining shape from shading information. *The Psychology of Computer Vision*, pages 115–155, 1975.
- [12] K. Kato, F. Sakaue, and J. Sato. Extended multiple view geometry for lights and cameras from photometric and geometric constraints. In *Proc. of ICPR*, pages 2110–2113, 2010.
- [13] M. Liao, L. Wang, R. Yang, and M. Gong. Light fall-off stereo. In *Proc. of CVPR*, pages 1–8, 2007.
- [14] V. Maier, V. F. Polec, S. G. Pavel, and A. Filip. Optimal analytical identification of the photometric curves for the indoor direct lighting luminaires. In *Optimization of Electrical and Electronic Equipment (OPTIM)*, pages 67–72, 2012.
- [15] M. Powell, S. Sarkar, and D. Goldgof. A simple strategy for calibrating the geometry of light sources. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 23(9):1022–1027, 2001.
- [16] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *Proc. of ECCV*, pages 859–872, 2008.
- [17] R. Rykowski and H. Kostal. Novel approach for led luminous intensity measurement. In *Proc. SPIE*, volume 6910, 2008.
- [18] D. Schnieders and K.-Y. K. Wong. Camera and light calibration from reflections on a sphere. *Computer Vision and Image Understanding (CVIU)*, 117(10):1536–1547, 2013.
- [19] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. A biquadratic reflectance model for radiometric image analysis. In *Proc. of CVPR*, pages 230–237, 2012.
- [20] R. Simons and A. Bean. *Lighting engineering: applied calculations*. Routledge, 2012.
- [21] T. Takai, K. Niinuma, A. Maki, and T. Matsuyama. Difference sphere: an approach to near light source estimation. In *Proc. of CVPR*, pages I: 98–105.
- [22] M. Weber and R. Cipolla. A practical method for estimation of point light-sources. In *Proc. of BMVC*, pages 1–10, 2001.
- [23] K.-Y. K. Wong, D. Schnieders, and S. Li. Recovering light directions and camera poses from a single sphere. In *Proc. of ECCV*, pages 631–642, 2008.
- [24] R. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Eng.*, 19(1), 1980.
- [25] T. Wu and C. Tang. Dense photometric stereo using a mirror sphere and graph cut. In *Proc. of CVPR*, pages I: 140–147, 2005.
- [26] L.-F. Yu, S.-K. Yeung, Y.-W. Tai, D. Terzopoulos, and T. Chan. Outdoor photometric stereo. In *Proc. of ICCP*, pages 1–8, 2013.
- [27] W. Zhou and C. Kambhampettu. Estimation of the size and location of multiple area light sources. *Proc. of ICPR*, 3:214–217, 2004.
- [28] W. Zhou and C. Kambhampettu. A unified framework for scene illuminant estimation. *Image and Vision Computing*, 26(3):415–429, 2008.