

# Reflectance and Shape Estimation with a Light Field Camera under Natural Illumination

Thanh Trung Ngo<sup>1</sup>  
trungbeo@gmail.com

Hajime Nagahara<sup>2</sup>  
nagahara@ids.osaka-u.ac.jp

Ko Nishino<sup>3</sup>  
kon@drexel.edu

Rin-ichiro Taniguchi<sup>4</sup>  
rin@kyudai.jp

Yasushi Yagi<sup>1</sup>  
yagi@am.sanken.osaka-u.ac.jp

<sup>1</sup> The Institute of Scientific and Industrial Research  
Osaka University, Japan

<sup>2</sup> Institute for Dataability Science  
Osaka University, Japan

<sup>3</sup> Department of Computer Science  
Drexel University, Philadelphia, USA

<sup>4</sup> Graduate School of Information Science and Electrical Engineering  
Kyushu University, Fukuoka, Japan

---

## Abstract

Reflectance and shape are two important components in visually perceiving the real world. Inferring the reflectance and shape of an object through cameras is a fundamental research topic in the field of computer vision. While three-dimensional shape recovery is pervasive with varieties of approaches and practical applications, reflectance recovery has only emerged recently. Reflectance recovery is a challenging task that is usually conducted in controlled environments, such as a laboratory environment with a special apparatus. However, it is desirable that the reflectance be recovered in the field with a handy camera so that reflectance can be jointly recovered with the shape. To that end, we present a solution that simultaneously recovers the reflectance and shape (i.e., dense depth and normal maps) of an object under natural illumination with commercially available handy cameras. We employ a light field camera to capture one light field image of the object, and a 360-degree camera to capture the illumination. The proposed method provides promising results in simulation and real-world experiments.

## 1 Introduction

In the field of computer vision, we need to understand the geometry and material of an object to obtain information about the object. The visual perception of the object depends on the illuminating environment, which poses a challenging and interesting task for computer vision to understand three components: the geometry, material, and illumination. This is in fact an inverse rendering problem, the complexity of which is extremely high [24]. To relax this complexity, computer-vision researchers usually assume to know one or two components and they then can recover the remaining one(s). For example, researchers can assume the reflectance is as simple as Lambertian reflectance, and the shape can then be recovered

knowing the illumination [23, 30] or even without knowing the illumination [9]. Meanwhile, other researchers assume to know the shape and the reflectance and can then recover the illumination [9]. In our research, we found that the illumination is not a serious problem and can be easily captured with a handy 360-degree camera [22]. We can thus relax the illumination and focus on recovering the reflectance and three-dimensional (3-D) shape of the object.

3-D shape recovery has been well studied in the field of computer vision with varieties of approaches and practical applications. The approaches can work under both controlled [9, 23] and uncontrolled [19, 26, 34] environments. A large number of methods assume Lambertian reflectance so that they can recover the scene with [6, 31] or without [10, 9] consideration of the illumination. This also allows the methods to work with a mobile or handy camera [27]. Other specific reflectance models, such as those of dielectric reflectance [24, 25], mirrored reflectance [9], or a combination of Lambertian reflectance and specular reflectance [28], are also assumed to relax the shape recovery.

However, reflectance recovery is challenging because real-world material reflectance is difficult to represent mathematically with a model. Researchers therefore try to approximate real-world material reflectance with varieties of models for specific types of materials, such as a Lambertian model for diffuse materials, dielectric model [33] for a ceramic or plastic, Phong reflectance model [21] for specular material, Torrance-Sparrow reflectance model [16, 29] for a rough surface, data-driven reflectance model [12], directional statistics bidirectional reflectance distribution function (DSBRDF) model [9, 15] for more general reflectance. Measuring material reflectance is also difficult. The measurement is usually made in a laboratory under well-controlled conditions. The widely used method relies on image-based bidirectional reflectance distribution function (BRDF) measurement [2, 10, 12, 13] by capturing many images of the material sample with different known light directions.

Moreover, joint shape and reflectance recovery using images is even more challenging, particularly under uncontrolled illuminations. To relax the problem, most methods assume a Lambertian material and try to recover the spatially varying BRDF and shape of the object [11]. Other methods use a specific reflectance model such as a mirror model [8]. A dichromatic model (a combination of diffuse and specular reflectance models) is also used [32] with a known point light source and it is thus only used in a laboratory. Oxholm and Nishino [19] simultaneously recovered the general reflectance and shape of an object under natural illumination. Their method employs a general reflectance model (DSBRDF [15]) and thus works with a wide range of real-world materials. However, it is only used for a well-setup environment of multiple calibrated cameras or a single camera that moves to several positions about the object. The method may be hard to employ in practice, which limits its real-world application. In their work, shape is represented by the surface normal which is constrained only by photo-consistency among sparse and wide-baseline stereo corresponding points. Inspired by this pioneering work, we conduct our research with a light field camera to reduce the preparation effort and encourage practical use. In particular, we further constrain the shape by introducing depth and reformulate the probabilistic framework to update depth. The depth and normal are separately updated owing the practical fact that depth and normal maps may not be equivalent and they are directly updated with different cues, multi-view stereo correspondence and radiometry, respectively. However, the depth and normal are strongly correlated, and a new constraint is presented to tighten them. The constraint is referred to as *depth-normal consistency constraint*, a geometry constraint, in our paper. As a result, we can estimate not only reflectance and normal but also depth simultaneously.

The contributions of our paper are summarized as follows. Our work is the first work that recovers a general isotropic reflectance and depth map under natural illumination with a

light field camera. The production of a dense depth map needs a huge computation cost. To reduce the computational cost, we present a multi-stage algorithm, where the earlier stage solves a coarser problem with a simpler and faster solution.

## 2 Image Generative Model

Our system includes a light field camera that captures the light field image of an object and a 360-degree camera that captures the environmental illumination. The light field and 360-degree cameras are calibrated and registered so that their geometrical relationship is known in advance. The world coordinate system coincides with the coordinate system of the light field camera, where  $Oz$  is the optical axis. We use the *plane and tangent direction presentation* [85] for the light field image, where a scene ray is described by a four-dimensional function,  $\mathbf{I}(x, y, s, t)$ .  $(x, y)$  denotes parameter of a sub-aperture that a scene ray passes through, while  $(x, y, 0)$  gives 3-D location of the sub-aperture on the  $Oxy$  plane, and  $(s, t)$  represents the tangent direction of the scene ray. The center sub-aperture location coincides with the origin  $O$ . The environmental illumination  $\mathbf{L}$  is captured by the 360-degree camera and is represented by a panorama image, of which each pixel is assumed as a directional light source from infinite distance. In the world coordinate system,  $\mathbf{L}$  is parameterized by the incoming light direction  $\boldsymbol{\omega}_i$ ,  $\mathbf{L}(\boldsymbol{\omega}_i)$ , where  $\boldsymbol{\omega}_i = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$ , and  $\theta$  and  $\phi$  are respectively the zenith and azimuth angles of the light direction.

The target object is assumed smooth and the surface geometry is represented by a surface function  $F(x, y, z) = 0$ . The surface reflectance is assumed homogeneous and isotropic and is represented by a single model with parameters  $\mathcal{R}$ . In our algorithm, we only consider a set of object points visible in the center sub-aperture image,  $\boldsymbol{\Omega} = \{\mathbf{P}\}$ . The object shape is represented by a depth map  $\mathcal{Z} = \{Z_{\mathbf{P}} | \mathbf{P} \in \boldsymbol{\Omega}\}$  and a normal map  $\mathcal{N} = \{\mathbf{N}_{\mathbf{P}} | \mathbf{P} \in \boldsymbol{\Omega}\}$ .

The surface normal  $\mathbf{N}_{\mathbf{P}}$  of an object point  $\mathbf{P} = (X_{\mathbf{P}}, Y_{\mathbf{P}}, Z_{\mathbf{P}})$  can be derived from the first partial derivative:

$$\mathbf{N}_{\mathbf{P}} = \frac{-(F_x(X_{\mathbf{P}}, Y_{\mathbf{P}}), F_y(X_{\mathbf{P}}, Y_{\mathbf{P}}), F_z(X_{\mathbf{P}}, Y_{\mathbf{P}}))^T}{\sqrt{F_x(X_{\mathbf{P}}, Y_{\mathbf{P}})^2 + F_y(X_{\mathbf{P}}, Y_{\mathbf{P}})^2 + F_z(X_{\mathbf{P}}, Y_{\mathbf{P}})^2}}. \quad (1)$$

This object point can be observed on the different sub-aperture images  $(x, y)$  and directions  $(s, t)$  depending on its depth such that

$$(s, t, 1)^T = (\mathbf{P} - (x, y, 0)^T) / Z_{\mathbf{P}}. \quad (2)$$

We see that  $s, t$  depends on the sub-aperture location  $(x, y, 0)^T$ ; hence,  $s = s(x, y)$  and  $t = t(x, y)$ . From (2), we obtain a relationship between image points on the center sub-aperture and on a surrounding sub-aperture image:

$$s(x, y) = s(0, 0) - x / Z_{\mathbf{P}}, \quad t(x, y) = t(0, 0) - y / Z_{\mathbf{P}}. \quad (3)$$

We define an observing direction  $\boldsymbol{\omega}_{\mathbf{P}}(x, y)$  for  $\mathbf{P}$  from the sub-aperture  $(x, y)$ ; hence,  $\boldsymbol{\omega}_{\mathbf{P}}(x, y)$  and  $(s, t, 1)^T$  are collinear:

$$\boldsymbol{\omega}_{\mathbf{P}}(x, y) = -(s(x, y), t(x, y), 1)^T / \sqrt{s(x, y)^2 + t(x, y)^2 + 1}. \quad (4)$$

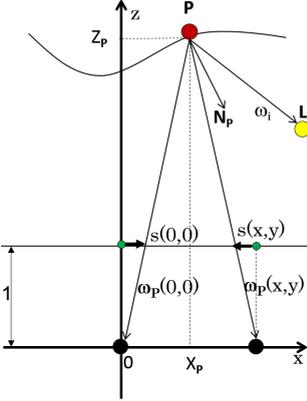


Figure 1: A 2-D illustration of the world coordinate system and light field representation.

Because the direction  $(s(0,0), t(0,0))$  of the point  $\mathbf{P}$  is constant from the center sub-aperture, the observing direction from a surrounding sub-aperture depends only on its location  $(x,y,0)$  and the point depth  $\mathcal{Z}_{\mathbf{P}}$ .

Given the geometry and reflectance function  $\rho$ , the actual reflectance  $\mathbf{I}(x,y,s,t)$  on each sub-aperture image captured by the camera can be modeled mathematically by integration:

$$\mathbf{E}(x,y,s,t) = \int \rho(\tau(\boldsymbol{\omega}_i, \boldsymbol{\omega}_{\mathbf{P}}(x,y), \mathbf{N}_{\mathbf{P}}); \mathcal{R}) \mathbf{L}(\boldsymbol{\omega}_i) \max(0, \mathbf{N}_{\mathbf{P}} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \quad (5)$$

where  $\tau$  is a function that transforms  $\boldsymbol{\omega}_i$  and  $\boldsymbol{\omega}_{\mathbf{P}}$  in relation with the surface normal  $\mathbf{N}_{\mathbf{P}}$  into halfway parameterization variables  $\theta_d$  and  $\theta_h$  [24].

### 3 Probabilistic Estimation Framework

Similar to the work of Oxholm and Nishino [19], we use a similar probabilistic framework to formulate our estimation problem. The geometry and reflectance estimation can be formulated as a maximum a posteriori for the whole light field:

$$p(\mathcal{N}, \mathcal{Z}, \mathcal{R} | \mathbf{I}) \propto p(\mathbf{I} | \mathcal{N}, \mathcal{Z}, \mathcal{R}) p(\mathcal{N}, \mathcal{Z}) p(\mathcal{R}), \quad (6)$$

where  $p(\mathbf{I} | \mathcal{N}, \mathcal{Z}, \mathcal{R})$  is a likelihood that quantifies how the geometry and reflectance match the light field; it can also be referred as a measure of photo-consistency among all sub-aperture images.  $p(\mathcal{N}, \mathcal{Z})$  is a geometrical constraint on the shape (normal and depth), and  $p(\mathcal{R})$  is a practical reflectance prior. To cope with the high dynamic range of the reflectance, we also process reflectance in the log-intensity domain [9, 19]. We assume the captured and modeled image intensities differ by some scalar  $\mu$  and Gaussian noise with variance  $\sigma^2$ :

$$\log(\mathbf{I}(x,y,s,t)) = \log(\mu \mathbf{E}(x,y,s,t)) + N(\mathbf{0}, \sigma^2) \quad (7)$$

$$= \log(\mathbf{E}(x,y,s,t)) + N(\mu \mathbf{1}, \sigma^2). \quad (8)$$

#### 3.1 Image likelihood

We first describe the first term of (6), the image likelihood for the light field of the object. It is a joint likelihood of individual likelihoods for light rays  $(u,v,s,t)$  from an object point  $\mathbf{X}$ . The individual likelihood is evaluated by a dissimilarity  $d(\cdot)$  between the captured and modeled intensities for the light ray:

$$p(\mathbf{I}(x,y,s,t) | \mathbf{N}_{\mathbf{X}}, \mathcal{Z}_{\mathbf{X}}, \mathcal{R}) = N(d(\mathbf{I}(x,y,s,t), \mathbf{E}(x,y,s,t)); \mu, \sigma^2).$$

The joint likelihood for the whole object is then computed for all the object points using all the sub-aperture images:

$$p(\mathbf{I} | \mathcal{N}, \mathcal{Z}, \mathcal{R}) = \prod_{\mathbf{X} \in \Omega} \prod_{(x,y)} p(\mathbf{I}(x,y,s,t) | \mathbf{N}_{\mathbf{X}}, \mathcal{Z}_{\mathbf{X}}, \mathcal{R}). \quad (9)$$

Although the baselines between camera sub-apertures are narrow, there are still mismatches between the sub-aperture images even when the depth map is perfect. We therefore employ a robust score function for  $d(\cdot)$  to compute the matching score between images. The Welsh function is a robust function that is suited to our situation:

$$d(\mathbf{I}(x,y,s,t), \mathbf{E}(x,y,s,t); \sigma, \mu) = \sqrt{\frac{n_{\lambda} \sigma^2}{2} \left[ 1 - \exp\left(-\frac{\|\log \mathbf{I}(\cdot) - \log \mathbf{E}(\cdot) - \mu \mathbf{1}\|^2}{n_{\lambda} \sigma^2}\right) \right]}, \quad (10)$$

where  $n_{\lambda}$  is the number of color channels, which is three for  $R, G, B$  in our experiment.

### 3.2 Geometrical constraints

We use the alternating optimization scheme to estimate all unknowns. The geometrical constraint is formulated depending on whether the normal or depth is updated:

$$p(\mathcal{N}, \mathcal{Z}) \propto \begin{cases} p(\mathcal{N}|\mathcal{Z}) & \text{when updating depth,} \\ p(\mathcal{Z}|\mathcal{N})p(\mathcal{N}) & \text{when updating normal.} \end{cases} \quad (11)$$

The detailed constraints are described below, and include (a) the depth-normal consistency constraint, (b) the smoothness constraint on the surface normal  $p_s$ , (c) the surface gradient constraint  $p_g$ , and (d) the occluding boundary constraint  $p_b$ . The prior on the surface normal is then  $p(\mathcal{N}) = p_s p_g p_b$ . However, we do not apply any practical prior to the depth, as described in (11).

(a) *Depth-normal consistency constraint*: This constraint results from the object surface being smooth and the surface normal being perpendicular to the surface gradients. The constraint that quantifies how much the normal matches the depth is

$$p(\mathcal{Z}|\mathcal{N}) \propto \prod_{\mathbf{X} \in \Omega} \prod_{\mathbf{Y} \in ne(\mathbf{X})} \exp\left(-\beta_{zn}(\mathbf{N}_{\mathbf{X}} \cdot (\mathbf{X} - \mathbf{Y}))^2\right), \quad (12)$$

where  $\beta_{zn}$  controls the constraint strength and  $ne(\mathbf{X})$  is a set of neighboring points of  $\mathbf{X}$  in  $\Omega$ . Similarly, the constraint that quantifies how much the depth matches the normal is

$$p(\mathcal{N}|\mathcal{Z}) \propto \prod_{\mathbf{X} \in \Omega} \prod_{\mathbf{Y} \in ne(\mathbf{X})} \exp\left(-\beta_{nz}(\mathbf{N}_{\mathbf{X}} \cdot (\mathbf{X} - \mathbf{Y}))^2\right), \quad (13)$$

where  $\beta_{nz}$  controls the constraint strength.

(b) *Normal smoothness constraint [14]*:

$$p_s(\mathcal{N}) \propto \prod_{\mathbf{X}} \prod_{\mathbf{Y} \in ne(\mathbf{X})} \exp\{-\beta_s \arccos^2(\mathbf{N}_{\mathbf{X}} \cdot \mathbf{N}_{\mathbf{Y}})\}, \quad (14)$$

where  $\beta_s$  controls the strength of the constraint.

(c) *Surface gradient constraint [14]*: This constraint ensures the resulting gradient is the same as in the observed image, and it is built on the center sub-aperture image only:

$$p_g(\mathcal{N}) \propto \prod_{\mathbf{X}} \prod_{\mathbf{Y} \in ne(\mathbf{X})} \exp\left(-\beta_g \|(\log \mathbf{E}_{\mathbf{Y}}(\cdot) - \log \mathbf{E}_{\mathbf{X}}(\cdot)) - (\log \mathbf{I}_{\mathbf{Y}}(\cdot) - \log \mathbf{I}_{\mathbf{X}}(\cdot))\|^2\right), \quad (15)$$

where  $\beta_g$  controls the strength of this constraint.

(d) *Occluding boundary constraint [14]*: At the occluding boundary, the surface normal should be oriented orthogonally to the observing direction:

$$p_b(\mathcal{N}) \propto \prod_{\mathbf{X} \in \mathbf{B}} \exp\left(-\beta_b \arccos^2(\mathbf{N}_{\mathbf{X}} \cdot \boldsymbol{\omega}_{\mathbf{X}}(0,0))\right), \quad (16)$$

where  $\mathbf{B}$  is the set of boundary pixels in the center sub-aperture image, and  $\beta_b$  controls the strength of this constraint.

### 3.3 Reflectance constraints

We use the DSBRDF [9] to model an isotropic reflectance of the target material in our algorithm. This reflectance is modeled as a sum of lobes for each color channel  $\lambda$ :

$$\rho_\lambda(\theta_d, \theta_h; \kappa, \gamma, \mathbf{c}_\lambda) = \sum_r c_{r,\lambda} \left( \exp \left[ \kappa_r(\theta_d) \cos^{\gamma_r(\theta_d)} \theta_h \right] - 1 \right), \quad (17)$$

where  $\kappa$ ,  $\gamma$  respectively control the overall brightness and specularity of the lobes and chromaticity  $\mathbf{c}_\lambda = \{c_{r,\lambda}\}$  modulates the color channel  $\lambda$  with a constraint that  $\sum_\lambda c_{r,\lambda} = 1$ . There are three lobes and three channels in our case. We thus need two chromaticity values per lobe per channel, and six values for  $\kappa$  or  $\gamma$  per lobe. Totally, the DSBRDF is represented by 42 parameters including six chromaticity parameters. In practice, Lombardi and Nishino represented a DSBRDF with fewer parameters employing functional principal component analysis for  $\kappa$  and  $\gamma$  [9].  $\kappa$  and  $\gamma$  of each DSBRDF can be represented by a point in a high-dimension subspace  $\Psi = \{\psi_i | i \in \{1, \dots, 36\}\}$ . The subspace is constructed with a large number of measured reflectance materials (e.g., using the MERL BRDF database [10]). The earlier parameter  $\psi_i$  with smaller  $i$  statistically has more power to represent a practical BRDF. Accuracy and compactness can be traded off using a subset of the first  $n_\Psi$  parameters. In experiments representing the MERL BRDF database with this DSBRDF model, they showed that with just  $n_\Psi=13$  parameters, the MERL BRDF database could be fitted well [9]. In our experiments, we use  $n_\Psi=14$ . In this compact representation, each reflectance is represented by  $\mathcal{R} = \{\mathbf{c}, \Psi\} | \mathbf{c} = \{\mathbf{c}_\lambda\}, \Psi = \{\psi_i\} | i \in \{1, \dots, n_\Psi\}$ .

Further, we use two constraints on the DSBRDF reflectance, one for the chromaticities and one for the coefficients so that  $p(\mathcal{R}) = p(\Psi)^{\beta_\Psi} p(\mathbf{c})^{\beta_c}$ .  $p(\Psi) = N(\Psi, \Sigma)$  is the same as in [10], where the covariance matrix  $\Sigma$  is learned from the MERL database, and  $p(\mathbf{c})$  is the same as in [9].  $\beta_\Psi, \beta_c$  control the priors' strength.

## 4 Multi-stage Algorithm

There are many unknowns for the depth  $\mathcal{Z}$ , normal  $\mathcal{N}$ , and reflectance  $\mathcal{R}$  to be estimated simultaneously. It is usually difficult for the estimation to converge stably. To efficiently manage the stability and reduce the computational cost, we make the estimation in stages. The idea for this algorithm originates from the generative model of the light field image in (5) with different simplification levels.

### 4.1 Stage 1: Depth Estimation from Plane Sweeping

When the shape is not accurate, we assume the reflectance and light functions are constant for all observing directions, then we only estimate the depth. The probabilistic framework in (6) is simplified by relaxing reflectance and light:

$$p(\mathcal{Z} | \mathbf{I}) \propto p(\mathbf{I} | \mathcal{Z}). \quad (18)$$

In this case, the object point has similar intensity crossing the sub-aperture images:  $I(x, y, s, t) \approx I(0, 0, s(0, 0), t(0, 0))$ . We use plane sweeping to find an optimal depth by matching intensity between the center and all other sub-aperture images.

A graph cut [8] is employed to obtain a smoother depth map.

## 4.2 Stage 2: Depth and Normal Estimation with Lambertian

We only assume the reflectance function is constant for all observing directions, and we then simultaneously update the depth  $\mathcal{Z}$  and surface normal  $\mathcal{N}$ . The probabilistic framework in (6) is simplified by relaxing the reflectance:

$$p(\mathcal{N}, \mathcal{Z} | \mathbf{I}) \propto p(\mathbf{I} | \mathcal{N}, \mathcal{Z}) p(\mathcal{N}, \mathcal{Z}). \quad (19)$$

This can be done by setting a special case of DSRDF reflectance in (17) that there is one lobe is 1 and  $\gamma$  and  $\kappa$  are set to 0 and 1, respectively. The estimation method is similar to that presented in the next subsection with Lambertian reflectance.

## 4.3 Stage 3: Depth, Normal, and Reflectance Estimation

In this final stage, we use the DSRDF to constrain the real-world isotropic material reflectance. All unknowns, depth, normal, and reflectance parameters, are updated without approximation. It is costly to handle, but the computational cost is reduced appreciably by the first two stages. An iterative optimization scheme using a probabilistic framework is employed that alternates between (a) updating the Gaussian noise ( $\sigma, \mu$ ), (b) updating the reflectance, (c) updating the surface normal, and (d) updating the depth is employed.

(a) *Update Gaussian noise:* To update the Gaussian noise defined in (8) for current reflectance and shape, we simply compute the standard deviation  $\sigma$  and mean  $\mu$  of all the errors for object points in the light field.

(b) *Update reflectance:* To update the reflectance  $\mathcal{R}$  assuming that the geometry is constant, the maximum a posteriori estimate in (6) becomes

$$p(\mathcal{R} | \mathbf{I}) \propto p(\mathbf{I} | \mathcal{N}, \mathcal{Z}, \mathcal{R}) p(\Psi)^{\beta_\Psi} p(\mathbf{c})^{\beta_c}. \quad (20)$$

However, observations between sub-aperture images do not differ much, and we only use the light field from the center sub-aperture image to reduce the computational cost.

(c) *Update surface normal:* In this step, the surface normals are updated relying on the photo-consistency among sub-aperture images, normal-depth consistency, and surface normal prior assuming that the reflectance and surface depth are known. The objective function in (6) is formulated as

$$p(\mathcal{N} | \mathbf{I}) \propto p(\mathbf{I} | \mathcal{N}, \mathcal{Z}, \mathcal{R}) p(\mathcal{Z} | \mathcal{N}) p_s(\mathcal{N}) p_g(\mathcal{N}) p_b(\mathcal{N}). \quad (21)$$

(d) *Update depth:*

We continue to update the surface depth relying on the photo-consistency between sub-aperture images and normal-depth consistency, assuming that the reflectance and surface normal are known. The objective function in (6) is formulated as

$$p(\mathcal{Z} | \mathbf{I}) \propto p(\mathbf{I} | \mathcal{N}, \mathcal{Z}, \mathcal{R}) p(\mathcal{N} | \mathcal{Z}). \quad (22)$$

The three stages are illustrated in Figure 2.

## 5 Experiments

We evaluated the proposed method by conducting several synthesized and real-world experiments. Because there is no related work for the light field camera, we did not include a

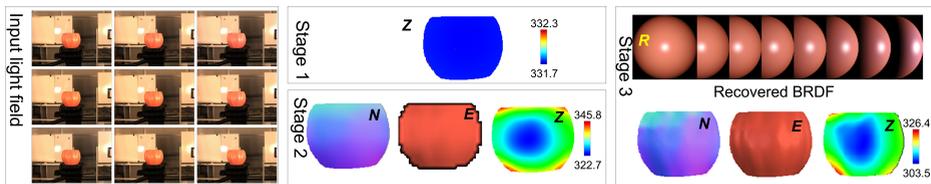


Figure 2: An illustration of the three stage results with an input light field. In stage 1, we estimate only depth  $\mathcal{Z}$ . In stage 2, we estimate the normal  $\mathcal{N}$  and depth  $\mathcal{Z}$ . In stage 3, we estimate the normal  $\mathcal{N}$ , depth  $\mathcal{Z}$ , and reflectance  $\mathcal{R}$ . Rendered images  $\mathbf{E}$  from (5) using the recovered shape and reflectance for stages 2 and 3 are also presented.

comparison in our experiment. The synthesis experiments were quantified by the surface normal error and the depth error relative to the ground-truths. Meanwhile, the real experiments were evaluated qualitatively.

Parameters in our experiments were set as  $\beta_s = 10$ ,  $\beta_g = 0.05$ ,  $\beta_{zn} = 10$ ,  $\beta_{zn} = 0.01$ ,  $\beta_b = 5$ ,  $\beta_\Psi = 0.001$ , and  $\beta_c = 0.001$ . The resolution of natural illumination is  $256 \times 128$ .

## 5.1 Simulation Experiments

Because the diffusion level of a material and distance from an object to the camera might affect reconstruction performance, we evaluated the proposed method with synthetic light field images against variations of (1) the diffusion level of material and (2) the depth from the object to the camera. The object geometries, reflectance, and natural illumination were taken from [10], the MERL BRDF database [11], and the light probe gallery created by the Institute for Creative Technologies of the University of Southern California [12], respectively. The camera model and geometry were generated by Pbrt [13].  $5 \times 5$  sub-aperture images were generated with the synthesized light field camera, with the minimum baseline between two viewpoints being 10 mm. The field of view was  $30^\circ$  and the sub-aperture image resolution was  $640 \times 480$ . The distance from the object to the camera ranged from 240 to 440 mm.

In the first simulation experiment, we synthesized different BRDFs in terms of the diffusion level by linearly combining BRDFs of *steel* and *white fabric* from the MERL database [11], while the distance from the object to the camera was 400 mm. We used *Grace cathedral illumination*, which is shown in Figure 3. The performances of the proposed method are summarized in Figure 4.



Figure 3: Natural illumination.

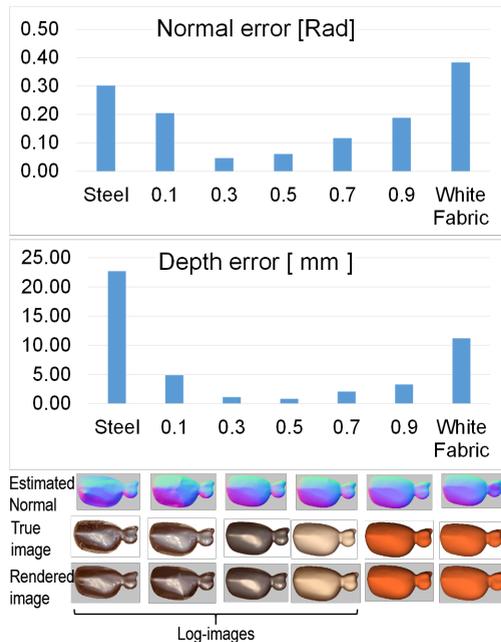


Figure 4: Simulation experiment with different synthesized materials.

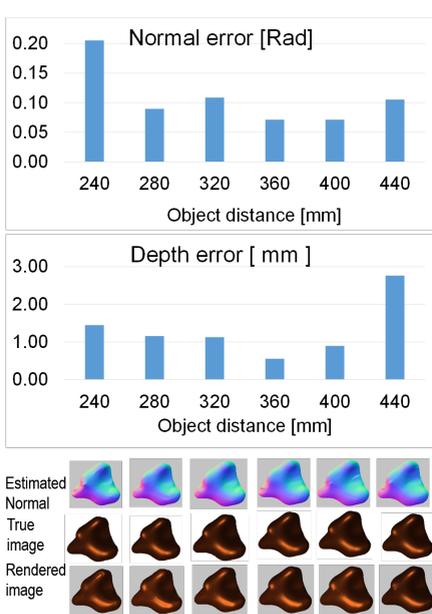


Figure 5: Simulation experiment with different distances of the object to the camera.

The results reveal that the proposed method can work with a certain range of diffuse and specular materials. For an extremely diffuse material, such as fabric, the narrow baseline between viewpoints is not wide enough to see parallax between sub-aperture images and the surface normal and depth are thus not recovered so well. The performance improves when the material is more specular. However, the proposed method does not work well for an extremely specular material, such as steel. The problem is that we employ a precomputed reflectance look-up table to accelerate the computation. The limited resolution of the look-up table cannot deal with extremely specular material with a sharp specular lobe. In the future work, we plan to employ different techniques to avoid such limitation.

In the second simulation experiment, we carried out a simulation experiment with different distances from the object to the camera. The synthesized images were generated with *blob7* [10], *gold metallic paint* [11], and Grace cathedral illumination. The performances of the proposed method are shown in Figure 5.

The results reveal that when the object is far from the camera, the performance is poor because the parallax is relatively small. The performance improves when the object moves closer to the camera. However, there is appreciable occlusion between viewpoints when the object is too close the camera. The situation is similar to that of a wide-baseline multi-view stereo reconstruction [12]. Fortunately, the problem is not practical for a light field camera and we do not focus on solving it.

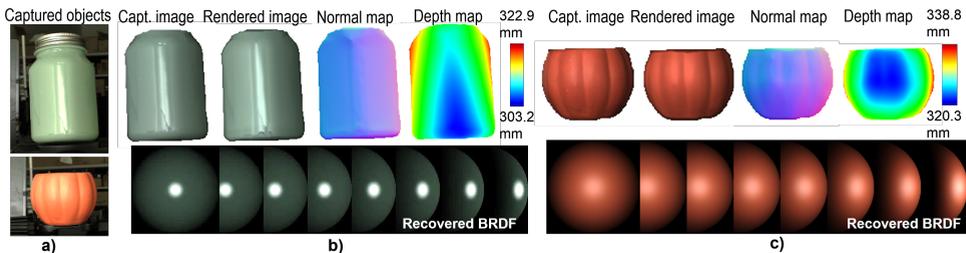


Figure 6: Real-world experiments: a) examples of captured images, b) reconstruction for a highly specular plastic bottle, and c) reconstruction for a pottery jar. Input object and light images were white-balanced. Reflectance function is visualized with a white point-light source and a sphere of the recovered material.

## 5.2 Real Experiments

We evaluated the proposed method in real-world experiments with different objects. We employed a Lytro ILLUM camera to capture and make high-dynamic-range images of objects, and a Theta S camera [24] to capture natural illumination. The environment of our experiment is shown in Figure 3.b. The camera focal length was set at 50 mm. The sub-aperture image size was  $625 \times 434$ . After calibration of the Lytro camera, the baseline from a viewpoint to its nearest neighbor was about 0.5 mm. Similar to the case in simulation experiments, we used  $5 \times 5$  sub-aperture images. An example of  $3 \times 3$  sub-aperture images is shown at the left of Figure 2.

We performed real experiments with two objects, a highly specular plastic bottle and a pottery jar. The results are shown in Figure 6. The results reveal that the surface normals and reflectance functions are well recovered so that the rendered object images and captured images are similar. We also see that the material of the bottle is much more specular than that of the pottery jar. Overall, we find that the proposed method works well with real-world objects against a narrow baseline of the light field camera.

## 6 Conclusion and Future Works

We presented the recovery of the shape and reflectance of an object with a light field camera under natural illumination. The advantage of the proposed method is that it is practical to deploy in reality with minimal effort to acquire input images and more information on the object is recovered. We employ a multi-stage algorithm to handle the high complexity. In experiments, we currently get good results with several real-world materials. In future work, we plan to improve the quality of reconstruction for more extreme specular materials, such as a mirror and metal to demonstrate that the proposed method can work well with various types of material.

## 7 Acknowledgment

We thank Glenn Pennycook, MSc, from Edanz Group ([www.edanzediting.com/ac](http://www.edanzediting.com/ac)) for editing a draft of this manuscript.

## References

- [1] Jonathan Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. 1(c):1–19, 2013. ISSN 0162-8828. doi: 10.1109/TPAMI.2014.2377712. URL <http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA586648>.
- [2] Moshe Ben-ezra, Jiaping Wang, and Bennett Wilburn. An LED-only BRDF Measurement Device. (i):1–8, 2015. URL <papers3://publication/uuid/630BCBE8-038A-4752-A1C5-DD1304D3EA13>.
- [3] Amael Delaunoy and Marc Pollefeys. Photometric Bundle Adjustment for Dense Multi-View 3D Modeling. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1486–1492, 2014. ISSN 10636919. doi: 10.1109/CVPR.2014.193.
- [4] Sam Van der Jeught and Joris J J Dirckx. Real-time structured light profilometry: a review. *Optics and Lasers in Engineering*, 87:18–31, 2016. ISSN 0143-8166. doi: <http://dx.doi.org/10.1016/j.optlaseng.2016.01.011>. URL <http://www.sciencedirect.com/science/article/pii/S0143816616000166>.
- [5] Clement Godard, Peter Hedman, Wenbin Li, and Gabriel J. Brostow. Multi-view Reconstruction of Highly Specular Surfaces in Uncontrolled Environments. *Proceedings - 2015 International Conference on 3D Vision, 3DV 2015*, pages 19–27, 2015. doi: 10.1109/3DV.2015.10.
- [6] Tomoaki Higo, Yasuyuki Matsushita, and Katsushi Ikeuchi. Consensus photometric stereo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1157–1164. Ieee, jun 2010. ISBN 978-1-4244-6984-0. doi: 10.1109/CVPR.2010.5540084. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5540084>.
- [7] Micah K. Johnson and Edward H. Adelson. Shape estimation in natural illumination. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2553–2560, 2011. ISSN 10636919. doi: 10.1109/CVPR.2011.5995510.
- [8] Vladimir Kolmogorov, Pascal Monasse, and Pauline Tan. Kolmogorov and Zabih’s Graph Cuts Stereo Matching Algorithm. *IPOl Journal*, 4:220–251, 2014. doi: 10.5201/ipol.2014.97. URL <http://dx.doi.org/10.5201/ipol.2014.97>.
- [9] Stephen Lombardi and Ko Nishino. Reflectance and Illumination Recovery in the Wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):129–141, 2016. ISSN 0162-8828. doi: 10.1109/TPAMI.2015.2430318. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7102772>.
- [10] S R Marschner, S H Westin, E P Lafortune, and K E Torrance. Image-based bidirectional reflectance distribution function measurement. *Applied optics*, 39(16):2592–2600, 2000. ISSN 0003-6935. doi: 10.1364/AO.39.002592.
- [11] Matt Pharr, Wenzel Jakob and Greg Humphreys. Physically Based Rendering. URL <http://www.pbrt.org/>.
- [12] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A Data-driven Reflectance Model. *ACM Trans. Graph.*, 22(3):759–769, jul 2003. ISSN 0730-0301. doi: 10.1145/882262.882343. URL <http://doi.acm.org/10.1145/882262.882343>.

- [13] Yasuhiro Mukaigawa, Seiichi Tagawa, Jaewon Kim, Ramesh Raskar, Yasuyuki Matsushita, and Yasushi Yagi. Hemispherical confocal imaging using turtleback reflector. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6492 LNCS(PART 1):336–349, 2011. ISSN 03029743. doi: 10.1007/978-3-642-19315-6\_26.
- [14] Trung Ngo Thanh, Hajime Nagahara, and Rin-ichiro Taniguchi. Shape and light directions from shading and polarization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2310–2318, 2015.
- [15] Ko Nishino and Stephen Lombardi. Directional statistics-based reflectance model for isotropic bidirectional reflectance distribution functions. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 28(1):8–18, jan 2011. ISSN 1520-8532. URL <http://www.ncbi.nlm.nih.gov/pubmed/21200406>.
- [16] Michael Oren and Shree K. Nayar. Generalization of Lambert’s reflectance model. *Proceedings of the 21st annual conference on Computer graphics and interactive techniques - SIGGRAPH '94*, pages 239–246, 1994. doi: 10.1145/192161.192213. URL <http://portal.acm.org/citation.cfm?doid=192161.192213>.
- [17] Geoffrey Oxholm and Ko Nishino. Shape and Reflectance from Natural Illumination. *Computer Vision – ECCV 2012*, 7572:528–541, 2012. doi: 10.1109/CVPR.2014.277.
- [18] Geoffrey Oxholm and Ko Nishino. Multiview Shape and Reflectance from Natural Illumination. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2163–2170, 2014.
- [19] Geoffrey Oxholm and Ko Nishino. Shape and Reflectance Estimation in the Wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):376–389, 2016.
- [20] Gustavo Patow and Xavier Pueyo. A Survey of Inverse Rendering Problems. *Computer Graphics Forum*, 22(4):663–687, 2003. ISSN 01677055. doi: 10.1111/j.1467-8659.2003.00716.x.
- [21] Bui Tuong Phong. Illumination for Computer Generated Pictures. *Communications of the ACM*, 18(6):311–317, 1975. ISSN 00010782. doi: 10.1145/360825.360839.
- [22] Ricoh Company. Theta S camera. URL <https://theta360.com/en/about/theta/s.html>.
- [23] Robert J. Woodham. Photometric Method For Determining Surface Orientation From Multiple Images. *Optical Engineering*, 19(1), 1980. URL <http://opticalengineering.spiedigitallibrary.org/article.aspx?articleid=1221923>.
- [24] Szymon M. Rusinkiewicz. A New Change of Variables for Efficient BRDF Representation. *Rendering Techniques, Proceedings of Eurographics Rendering Workshop*, pages 11–22, 1998.
- [25] William A.P. Smith, Ravi Ramamoorthi, and Silvia Tozza. Linear depth estimation from an uncalibrated, monocular polarisation image. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9912 LNCS: 109–125, 2016. ISSN 16113349. doi: 10.1007/978-3-319-46484-8\_7.
- [26] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- [27] Michael W Tao and Pratul P Srinivasan. Depth from Shading, Defocus, and Correspondence Using Light-Field Angular Coherence. *CVPR*, pages 1940–1948, 2015.

- [28] Michael W Tao, Ting-chun Wang, Jitendra Malik, and Ravi Ramamoorthi. Depth Estimation for Glossy Surfaces with Light-Field Cameras. In *ECCV Workshop Light Fields Computer Vision*, pages 1–14, 2014.
- [29] K E Torrance and E M Sparrow. Theory for Off-Specular Reflection From Roughened Surfaces\*. *J. Opt. Soc. Am.*, 57(9):1105–1114, sep 1967. doi: 10.1364/JOSA.57.001105. URL <http://www.osapublishing.org/abstract.cfm?URI=josa-57-9-1105>.
- [30] USC Institute for Creative Technologies. High-Resolution Light Probe Image Gallery. URL <http://gl.ict.usc.edu/Data/HighResProbes/>.
- [31] George Vogiatzis and Roberto Cipolla. Multiview Photometric Stereo. *PAMI*, 30(3):548–554, 2008.
- [32] Ting-chun Wang, Manmohan Chandraker, Alexei A. Efros, and Ravi Ramamoorthi. SVBRDF-Invariant Shape and Reflectance Estimation from Light-Field Cameras Differential Stereo. *Cvpr*, pages 5451–5459, 2016. ISSN 10636919. doi: 10.1109/CVPR.2016.588.
- [33] Lawrence B Wolff. Generalizing Lambert’s Law For Smooth Surfaces. In *ECCV*, pages 40–53, 1996.
- [34] Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Transactions on Graphics*, 35(6):1–12, 2016. ISSN 07300301. doi: 10.1145/2980179.2980248. URL <http://dl.acm.org/citation.cfm?doid=2980179.2980248>.
- [35] Changyin Zhou and Shree K. Nayar. Computational cameras: Convergence of optics and processing. *IEEE Transactions on Image Processing*, 20(12):3322–3340, 2011. ISSN 10577149. doi: 10.1109/TIP.2011.2171700.