

A Fixed Viewpoint Approach for Dense Reconstruction of Transparent Objects

Kai Han¹, Kwan-Yee K. Wong¹, Miaomiao Liu²

¹The University of Hong Kong*, Hong Kong,

²NICTA[†] and CECS, ANU, Canberra

¹{khan, kykwong}@cs.hku.hk, ²miaomiao.liu@nicta.com.au

Abstract

This paper addresses the problem of reconstructing the surface shape of transparent objects. The difficulty of this problem originates from the viewpoint dependent appearance of a transparent object, which quickly makes reconstruction methods tailored for diffuse surfaces fail disgracefully. In this paper, we develop a fixed viewpoint approach for dense surface reconstruction of transparent objects based on refraction of light. We introduce a simple setup that allows us alter the incident light paths before light rays enter the object, and develop a method for recovering the object surface based on reconstructing and triangulating such incident light paths. Our proposed approach does not need to model the complex interactions of light as it travels through the object, neither does it assume any parametric form for the shape of the object nor the exact number of refractions and reflections taken place along the light paths. It can therefore handle transparent objects with a complex shape and structure, with unknown and even inhomogeneous refractive index. Experimental results on both synthetic and real data are presented which demonstrate the feasibility and accuracy of our proposed approach.

1. Introduction

Reconstructing a 3D model of an object from its 2D images has always been a hot topic in the field of computer vision. It has many important applications in robotics, augmented reality, video games, movie production, reverse engineering, etc. Despite the problem of 3D model reconstruction has virtually been solved for opaque objects with a diffuse surface, reconstruction of transparent objects is still very challenging and remains an open problem. The viewpoint dependent appearance of a transparent object quickly

*This project is supported by a grant from the Research Grant Council of the Hong Kong (SAR), China, under the project HKU 718113E.

[†]NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the ARC through the ICT Centre of Excellence program.




Figure 1. Real reconstruction setup and examples of transparent objects.

renders reconstruction methods tailored for diffuse surfaces useless. The literature is relatively sparse when it comes to shape recovery of transparent objects, and most of the existing methods are still highly theoretical. In fact, even with restrictive assumptions and special hardware setups, state-of-the-art methods can only handle transparent objects with very simple shape. Meanwhile, it is not difficult to see that there exist many transparent objects in our world (e.g., liquids, glass, plastics, crystals and diamonds). Hence, the study of 3D model reconstruction cannot be considered completed without taking transparent objects into accounts.

As mentioned previously, the difficulty of reconstructing a transparent object originates from its viewpoint dependent appearance. A transparent object may alter a light path by reflection, refraction, absorption and scattering at both its exterior surface as well as its interior structure. A number of existing work attempted to reconstruct a transparent object by exploiting specular highlights produced on the object surface [9, 10]. This approach considers only reflection of light taken place at the object surface, and greatly simplifies the problem by making it not necessary to consider the complex interactions of light as it travels through the object. However, refraction of light is indeed an important and unique characteristic of transparent objects. It provides in-

formation on surface shape and should not be ignored. On the other hand, methods based on reflection of light often work only under very restrictive assumptions and precisely controlled environments, making them not very practical.

In this paper, we focus our study in dense surface reconstruction of transparent objects. We introduce a fixed viewpoint approach for recovering the surface of a transparent object based on refraction of light. Like those methods that are based on specular highlights, our fixed viewpoint approach does not need to model explicitly the complex interactions of light as it travels through the object. We present a simple setup that allows us alter the incident light paths before light rays enter the object, and develop a method for recovering the surface of a transparent object based on reconstructing and triangulating these incident light paths. Compared with existing methods, our proposed method has the following benefits:

- It does not assume any parametric form for the shape of a transparent object.
- It can handle a transparent object with a complex structure, with an unknown and even inhomogeneous refractive index.
- It considers only the incident light paths before light rays enter a transparent object, and makes no assumption on the exact number of refractions and reflections taken place as light travels through the object.
- The proposed setup is simple and inexpensive.

The rest of the paper is organized as follows. Section 2 briefly reviews existing techniques in the literature for shape recovery of transparent objects. Section 3 describes our proposed approach for dense surface reconstruction of transparent objects in detail. Experimental results on both synthetic and real data are presented in Section 4, followed by conclusion in Section 5.

2. Related Work

In [19, 20], Murase reconstructed a rippling water surface from the average observed coordinates of an underwater pattern under orthographic projection. Morris and Kutulakos [17] solved a similar problem with an unknown refractive index of the liquid using two calibrated cameras and a known reference pattern. In [6], Hata *et al.* used structured light and Genetic Algorithm to estimate the shape of a transparent paste drop on a board. Ben-Ezra and Nayar [1] assumed a parametric form for the shape of a transparent object and estimated the shape parameters under the assumptions of known camera motion and distant background. In [25], Wetzstein *et al.* proposed a single image approach for reconstructing thin refractive surfaces using

light field probes. In [12, 13], Kutulakos and Steger characterized the class of reconstructible specular scenes, and developed algorithms for depth map computation in the cases where refraction/reflection of light occurs exactly once and twice respectively. In [4], Eren *et al.* determined the surface shape of glass objects using laser surface heating and thermal imaging. In [18], Morris and Kutulakos introduced *scatter-trace* of a pixel and recovered the exterior surface of a transparent object using the non-negligible specular reflection component. Similarly, Yeung *et al.* [26] exploited specular highlights and proposed a dual-layered graph-cut method to reconstruct the surface of a solid transparent object. In [8], Ihrke *et al.* dyed water with a fluorescent chemical and presented a level set method for reconstructing a free flowing water surface from multi-video input data by minimizing a photo-consistency error computed using ray-tracing. Miyazaki and Ikeuchi [16] proposed an iterative method to estimate the front surface shape of a transparent object by minimizing the difference between observed polarization data and polarization raytracing result under the assumptions of a known refractive index, illumination distribution and back surface shape. In [24], Trifonov *et al.* introduced a visible light tomographic reconstruction method by immersing the transparent object into a fluid with a similar refractive index. The 3D shape was recovered by building the light paths within the fluid and object. In [7], Hullin *et al.* embedded the transparent objects into fluorescence. The object surfaces were reconstructed by detecting the intersections of the visible laser sheets with the visual rays. Similar light sheet range scanning approach was introduced by Narasimhan *et al.* in [21] to acquire object geometry in the presence of scattering media. In [22], O’Toole *et al.* developed the Structured Light Transport (SLT) technique. Based on SLT, they implemented an imaging device that allows one-shot indirect-invariant imaging for reconstructing transparent and mirror surfaces using structured light. In [23], Shan *et al.* introduced a framework for optimizing a refractive height field from a single image under the assumptions of an orthographic camera, known background, single refractive material and differentiable height field. In [3], Chari and Sturm introduced a method that integrates radiometric information into light path triangulation for reconstruction of transparent objects from a single image. In [15], Ma *et al.* reformulated the transport of intensity equation in terms of light fields, and presented a technique for refractive index field reconstruction using coded illumination. In [14], Liu *et al.* proposed a frequency based method for establishing correspondences on transparent and mirror surfaces, reconstruction can then be done using any stereo methods. In [11], Ji *et al.* reconstructed the refractive index field of a gas volume by establishing ray-to-ray correspondences using a light field probe and solving the light paths through the refractive index field with variational method

based on Fermat’s Principle.

Note that existing solutions for surface reconstruction of transparent objects often work only under restrictive assumptions (e.g., known refractive index, single refractive material, exact number of refractions, non-negligible reflection of light, orthographic projection) and special hardware setups (e.g., light field probes, laser surface heating with thermal imaging, dying liquids with fluorescent chemical, immersing objects into liquids with similar refractive indexes), or for a particular class of objects (e.g., with known parametric model/average shape). There exists no general solution to this challenging and open problem. In this paper, we develop a fixed viewpoint approach for dense surface reconstruction of transparent objects based on altering and triangulating the incident light paths before light rays enter the object. We present a simple setup that allows us alter the incident light paths by means of refraction of light. Under this proposed setup, the segment of a light path between the first entry point on the object surface and the optical center of the camera remains fixed. This allows us ignore the details of the complex interactions of light inside the object. Compared with existing methods, our proposed approach (1) assumes neither a known nor homogeneous refractive index of the object; (2) places no restriction on the exact number of refractions and reflections taken place along a light path; (3) assumes no parametric form for the shape of the object. This allows our approach to handle transparent objects with a complex structure. Besides, our proposed setup is also very simple, and does not depend on any special and expensive hardware.

3. Shape Recovery of Transparent Objects

3.1. Notations and Problem Formulation

To solve the surface reconstruction problem, we consider a set of light paths originating from a reference pattern placed behind a transparent object, passing through the object and eventually reaching the image plane. We partition every such light path into two parts, namely (i) the *path before contact* (PBC) which originates from the reference pattern and ends at the *first entry point* (FEP) on the object surface (see the red paths in Figure 2(a)) and (ii) the *path after contact* (PAC) which originates from the FEP, passes through the interior of the object and finally terminates at the optical center of the camera (see the green paths in Figure 2(a)). We can now reformulate the surface reconstruction problem into estimating the FEP. The approach we take to tackle this problem is by altering the PBC while fixing the PAC for each light path. This enables us ignore the details of the complex interactions of light inside the object, and recover the FEP by triangulating the PBCs. In the next section, we introduce a simple setup that allows us alter the PBCs by means of refraction of light.




Figure 2. (a) A light path through an object is partitioned into two parts, namely i) the *path before contact* (PBC) which originates from the reference pattern to the *first entry point* (FEP) on the object surface (i.e., the red paths) and ii) the *path after contact* (PAC) that originates from FEP, passes through the interior of the object and terminates at the optical center of the camera (i.e., the green paths). (b) The first entry point can be recovered by filling the tank with a liquid to alter the PBC and triangulating the two PBCs.

3.2. Setup and Assumptions

In our proposed setup, a camera is used to capture images of a transparent object in front of a reference pattern. The camera and the object are kept fixed with respect to each other. This will ensure the PACs remain unchanged for all the image points of the object. The reference pattern is used here to reconstruct the PBCs, and is placed in two distinct positions. As mentioned before, our approach is based on altering the PBCs. To achieve this, we employ a water tank to immerse part of the object surface into a liquid so as to alter the PBCs by means of refraction of light (see Figure 2(b)). Two images of the transparent object are acquired for each position of the reference pattern, one without liquid in the tank and one with liquid in the tank. By establishing correspondences between image points of the object and points on the reference pattern, we can construct two PBCs for each image point, one in air and one in the liquid respectively. The FEP can then be recovered by triangulating these two PBCs.

Note that our proposed approach does not require the prior knowledge of the refractive indices of the object or that of the liquid. If, however, the refractive index of the liquid is known a priori, it is possible to also recover the surface normal at each FEP. The only assumption made in our approach is that the PACs remain unchanged when the object is partially immersed into the liquid. This assumption implies that a light ray should not reenter into the liquid medium after its FEP. This generally holds true for transparent objects with a convex shape, and for objects with holes completely enclosed inside the object. This assumption allows us to handle object with inhomogeneous refractive index. In practice, our method can also handle objects with shallow concavities as long as the previous assumption is

satisfied.

3.3. Dense Refraction Correspondences

Before we can triangulate PBCs to recover the FEP, we first need to construct the PBCs from the images. Since the two distinct positions of the reference pattern have been calibrated, it is straightforward to construct the PBC for an image point by locating a correspondence point on the reference pattern under each of the two distinct positions in the same medium (i.e., with/without liquid in the tank). It is obvious that the quality of the correspondences will have a direct effect on the quality of the reconstruction. There exist many methods for establishing correspondences, such as Gray Code, Phase Shift, etc. However, these methods often can only provide sparse correspondences with limited precisions (e.g., a small patch of pixels is mapped to a small region on a reference plane due to finite discretization). In this work, we would like to establish close-to point-to-point correspondences between the image and the reference pattern. We employ a portable display screen (e.g., an iPad) to serve as the reference pattern, and show a sequence of a thin stripe sweeping across the screen in vertical direction and then in horizontal direction (see Figure 3). We capture an image for each of the positions of the sweeping stripe. For each image point, its correspondence on the reference pattern can then be solved by examining the sequence of intensity values of the image point for each sweeping direction and locating the peak intensity value. The position of the stripe that produces the peak intensity value in each sweeping direction then gives us the position of the correspondence on the reference pattern. In order to improve the accuracy, we also fit a quadratic curve to the intensity profile in the neighborhood of the peak value to locate the peak.




Figure 3. The upper row shows images of a transparent hemisphere captured in front of a gray background (from left to right: reference pattern at higher position and without water, reference pattern at higher position and with water, reference pattern at lower position and without water, and reference pattern at lower position and with water). The lower row shows images of the hemisphere captured in front of a sweeping stripe (in the same order).




Figure 4. PBC construction and FEP estimation. The correspondences of an image point q on the reference pattern under position 0 and position 1 define a PBC. Given two PBCs in two different media, the FEP for q can be obtained by triangulating the PBCs.

3.4. Light Path Triangulation

Suppose high quality correspondences have been established between the images and the reference pattern under each of the two distinct positions and in each of the two media (i.e., with and without liquid in the tank). We can then construct two PBCs for each image point of the transparent object using the calibrated positions of the reference pattern. The FEP can then be recovered as the point of intersection between the two PBCs. Below we derive a simple solution for the FEP based on the established correspondences of an image point.

Consider an image point q of the transparent object. Suppose \mathbf{M}_0 and \mathbf{M}_1 denote its correspondences on the reference pattern under position 0 and position 1 with liquid in the tank, respectively. Similarly, let \mathbf{N}_0 and \mathbf{N}_1 denote its correspondences on the reference pattern under position 0 and position 1 without liquid in the tank, respectively (see Figure 4). The PBCs for q in liquid and in air, respectively, can be expressed as

$$L_M : \mathbf{M}(s) = \mathbf{M}_0 + s\mathbf{U}, \quad (1)$$

$$L_N : \mathbf{N}(t) = \mathbf{N}_0 + t\mathbf{V}, \quad (2)$$

where $\mathbf{U} = \frac{\mathbf{M}_1 - \mathbf{M}_0}{\|\mathbf{M}_1 - \mathbf{M}_0\|}$ and $\mathbf{V} = \frac{\mathbf{N}_1 - \mathbf{N}_0}{\|\mathbf{N}_1 - \mathbf{N}_0\|}$. Under a perfect situation, the FEP for q is given by the point of intersection between L_M and L_N .

Due to noise, however, L_M and L_N often may not intersect with each other at a point. In this situation, we seek the point $\mathbf{M}_c = \mathbf{M}(s_c)$ on L_M and the point $\mathbf{N}_c = \mathbf{N}(t_c)$ on L_N such that the distance between \mathbf{M}_c and \mathbf{N}_c is a minimum. Note that geometrically, for the distance between \mathbf{M}_c and \mathbf{N}_c to be a minimum, the line segment joining \mathbf{M}_c and \mathbf{N}_c must be orthogonal to both L_M and L_N . This orthogonal

constraint can be expressed as

$$(\mathbf{M}_c - \mathbf{N}_c) \cdot \mathbf{U} = 0, \quad (3)$$

$$(\mathbf{M}_c - \mathbf{N}_c) \cdot \mathbf{V} = 0. \quad (4)$$

Substituting (1) and (2) into the above equations gives

$$(\mathbf{M}_0 - \mathbf{N}_0 + s_c \mathbf{U} - t_c \mathbf{V}) \cdot \mathbf{U} = 0, \quad (5)$$

$$(\mathbf{M}_0 - \mathbf{N}_0 + s_c \mathbf{U} - t_c \mathbf{V}) \cdot \mathbf{V} = 0. \quad (6)$$

Solving the above simultaneous equations gives us s_c and t_c , and hence the points \mathbf{M}_c and \mathbf{N}_c . The distance between \mathbf{M}_c and \mathbf{N}_c can be taken as a quality measure of the reconstruction. If the distance is below a specified threshold, the mid-point between \mathbf{M}_c and \mathbf{N}_c can be taken as the FEP for q . Note that if \mathbf{U} and \mathbf{V} are parallel, there will not be a unique solution for the above equations. This corresponds to the case where the two PBCs coincide with each other. This case, fortunately, will be irrelevant to our proposed approach.

3.5. Surface Normal Reconstruction

Recall that for the purpose of surface reconstruction, neither the refractive index of the object nor that of the liquid is needed. If, however, the refractive index of the liquid is known a priori, it is possible to recover the surface normal at each FEP (see Figure 4). Let θ_1 and θ_2 denote the incident angles of the PBCs in the liquid and air, respectively, at the surface point P , and θ_0 denote the refracted angle. Suppose the refractive index of the object, liquid and air are given by λ_0 , λ_1 and λ_2 , respectively. By Snell's Law, we have

$$\lambda_0 \sin \theta_0 = \lambda_1 \sin \theta_1 = \lambda_2 \sin \theta_2. \quad (7)$$

Let $\Delta\theta = \cos^{-1}(\mathbf{U} \cdot \mathbf{V})$ denote the angle between the two PBCs, substituting this into (7) gives

$$\lambda_1 \sin \theta_1 = \lambda_2 \sin(\theta_1 + \Delta\theta). \quad (8)$$

With known refractive indices λ_1 and λ_2 for the liquid and air, respectively, the incident angle θ_1 can be recovered by

$$\theta_1 = \tan^{-1} \left(\frac{\lambda_2 \sin \Delta\theta}{\lambda_1 - \lambda_2 \cos \Delta\theta} \right). \quad (9)$$

The surface normal \mathbf{n}_p at P is then given by

$$\mathbf{n}_p = \mathbf{R}(\theta_1, \mathbf{V} \times \mathbf{U})\mathbf{U}, \quad (10)$$

where $\mathbf{R}(\theta, \mathbf{a})$ denotes a Rodrigues rotation matrix for rotating about the axis \mathbf{a} by the angle θ .

4. Experimental Evaluation

We now demonstrate the effectiveness of our approach on synthetic and real surfaces. In the remainder of this section, we present both quantitative and qualitative reconstruction results.

4.1. Synthetic Data

For our synthetic experiments, we used *Pov-Ray* to simulate the entire experimental setup. In particular, we modelled the transparent object as a semi-ellipsoid with the following parametric equation

$$\begin{cases} \left(\frac{x}{12.5}\right)^2 + \left(\frac{y}{12.5}\right)^2 + \left(\frac{z}{5}\right)^2 = 1, \\ z > 0. \end{cases} \quad (11)$$

We further assumed the transparent ellipsoid had a refractive index $\lambda = 1.5$.¹ A reference plane displaying a set of thin stripe sweeping patterns was placed at two different positions. A synthetic perspective camera was used to capture the refraction of the reference pattern through the transparent object immersed in air ($\lambda = 1.0$) and liquid ($\lambda = 1.3$) respectively. We adopted the strategies described in Section 3.3 to obtain dense refraction correspondences. More than 7M refraction correspondences were used in our synthetic experiment.

We constructed a pair of PBCs for each FEP based on the retrieved refraction correspondences. The transparent surface was then recovered from the ray triangulation of PBC pairs. Note that surface normals can be computed given PBC pairs and the refractive index of the medium. Figure 5 depicts the reconstructed FEP cloud as well as surface normals. We also provide the depth map of the reconstructed object for accuracy evaluation².

In practice, reconstruction errors originate from the inaccuracy in finding the refraction correspondences on the reference patterns. The errors may be increased by decreasing the relative distance between the pair of reference patterns. Therefore, we carried out the joint analysis by adding 2D zero-mean Gaussian noise to the extracted dense correspondences on the pair of reference patterns together with varying the relative distance between the reference patterns. The noise level ranged from 0.1 to 1.0. The relative distance between the reference patterns varied from 4 to 20. We fixed one reference pattern, namely setting one pattern at $z = 10$ in our experiment, and varied the position of the other one at various locations. Here we placed the other pattern at $z = 6, 15, 20, 25, 30$. The reconstruction accuracy is evaluated based on the root mean square error (RMSE) between the ground truth surface and the reconstruction. We further provide the angular distance between our reconstructed normals and the computed ground truth normals from the analytical equation of the semi-ellipsoid. Figure 6 (a-b) depicts the RMSE and the angular errors as a function of noise level and the relative distances between the pair of reference patterns. It shows that the reconstruction errors decrease with

¹The transparent object can be inhomogeneous, namely the refractive index varied across the interior of the object.

²The depth map is defined as the z component for each 3D point.




Figure 5. Evaluation on the synthetic semi-ellipsoid. (a) Ground truth. (b) Our reconstructed results. The first column shows the ground truth point cloud and the reconstructed FEP cloud; the second column shows the angle between each PBC pairs; the third column shows the depth map; the fourth column shows the normal map.

the increase of the distance between relative reference patterns.

We further conducted the analysis of the reconstruction error with respect to the medium refractive index. Another two media with different refractive indices were tested in the experiment, namely $\lambda = 1.5$ and $\lambda = 1.7$. The reference patterns were placed at $z = 6$ and $z = 10$ respectively. Figure 6 (c-d) shows that the reconstruction results are improved by increasing the refractive indices of the medium.

4.2. Experiments on Real Data

To evaluate the accuracy of our approach on real data, we performed the experiments on a smooth glass *hemisphere*, and a diamond-shape *ornament* with piecewise planar surfaces depicted in Figure 1. We acquired images with a *Canon EOS 40D* camera equipped with a 24 mm lens and used a 9.7-inch iPad with a resolution of 2048×1536 as the reference plane, on which we displayed stripe patterns that let us extract the dense refraction correspondences following the strategy in Section 3.3. In order to construct PBCs, the reference plane was placed at two different positions in a water tank. Under each position, we first took one set of images of the sweeping stripe patterns refracted by the object directly. We then filled the tank with water, having a refractive index $\lambda = 1.33$, to alter the PBCs of the surface and took another set of images. In brief, four sets of images with a resolution of 3888×2592 were captured for each object. This yielded dense correspondences (see Table 1). The pose of the reference plane relative to the camera was calibrated with Matlab Calibration Toolbox [2].

	<i>Hemisphere</i>	<i>Ornament</i>
Captured Images	2800	2200
Correspondences	1180300	546173
Reconstructed FEPs	1115748	519162
Reconstructed normals	1115748	519162

Table 1. Statistical data for real experiment. We show the number of captured encoding pattern images, refraction correspondences, the reconstructed FEPs, and the reconstructed normals for our dense reconstruction of *hemisphere* and *ornament*.

A pair of PBCs were constructed from the extracted refraction correspondences for each FEP, which led to the surface reconstruction by ray triangulation. We treated the reconstructed FEPs with small PBC angle ($< 1^\circ$), or out of the depth range between the camera and reference planes as noise points. The normal for each FEP was then recovered with the knowledge of refractive indices 1.0 and 1.33 for air and water respectively. In Figure 7, we show our reconstructed 3D FEP cloud, angles between each pair of PBCs, depth map, and surface normals for *hemisphere* and *ornament* respectively. Note that large reconstruction error occurs in regions with small PBC angles.

Since no ground truth was available, a sphere was fitted from the FEP cloud to evaluate the reconstruction accuracy for the *hemisphere*. We compare the fitted sphere radius with the measurement, which are 26.95 mm and 27.99 mm respectively, suggesting a high accuracy of the reconstruction. Table 2 shows the reconstruction errors of the *hemisphere*. In order to evaluate the reconstruction of the




Figure 6. Reconstructed FEP position error and normal angular error as a function of noise level and relative distance between each pair of reference patterns. (a)-(b) shows the average FEP position error and normal angular error respectively over 500 rounds with different random noise and relative distance between reference patterns. (c)-(d) shows the average FEP position error and normal angular error respectively over 500 rounds with different random noise and refractive indices of the media.

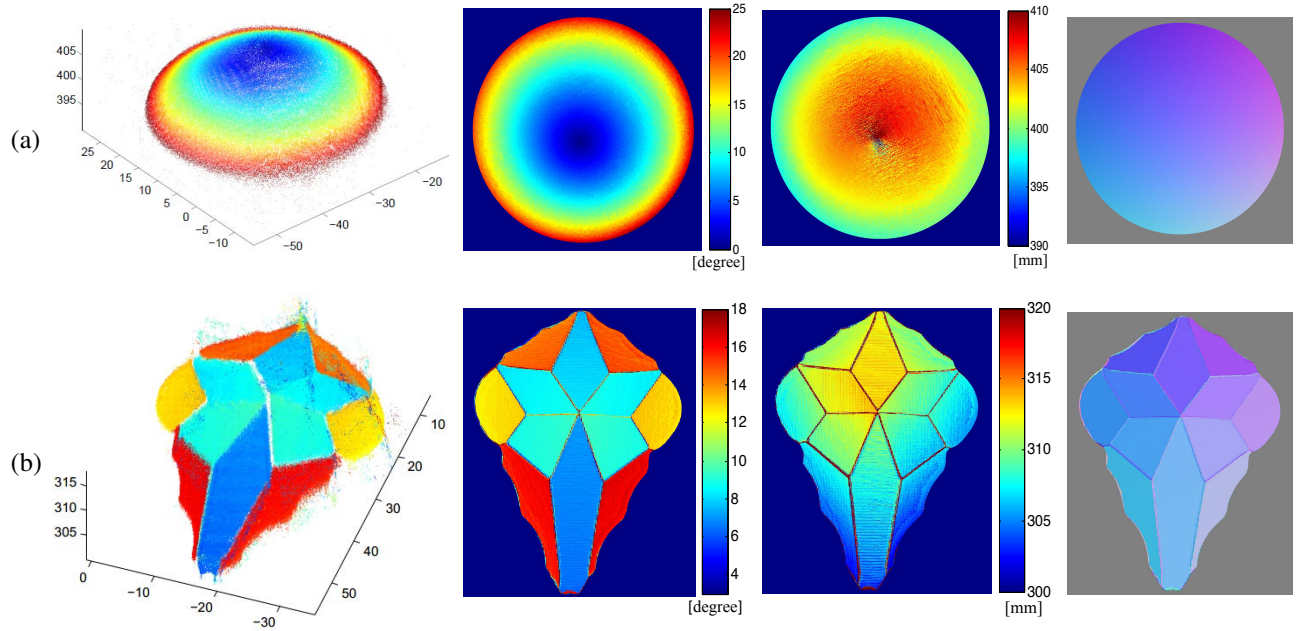


Figure 7. (a) *Hemisphere* reconstruction results. (b) *Ornament* reconstruction results. The first column shows the reconstructed FEPs; the second column shows the angle between each PBC pairs; the third column shows the depth map; the fourth column shows the reconstructed normal map.

	Position (mm)	Normal (degree)
Mean error	0.5903	6.9665
Median error	0.4179	6.9215


Table 2. Reconstruction error measurements for *hemisphere*. The position error is defined as the difference between the distance from the fitted center to each FEP and the length of fitted radius. The normal error is defined as the angle between the ray from the fitted center to each FEP and the reconstructed normal for each FEP.

ornament, we first used RANSAC [5] to fit a plane for each facet. The reconstruction error for each facet was measured by the distances from the reconstructed FEPs to the fitted plane, as well as the angles between the reconstructed nor-

mals of the facet and the normal of the fitted facet. The results shown in Table 3 suggest that our proposed approach can accurately reconstruct the piecewise planar *ornament*.

5. Conclusion

In this paper, we develop a fixed viewpoint approach for dense surface reconstruction of transparent objects. We introduce a simple setup that allows us alter the incident light paths by immersing part of the object surface in a liquid, while keeping the rest of the light paths fixed as light rays travel through the object. This greatly simplifies the problem by making it not necessary to model the complex interactions of light inside the object, and allows the object surface to be recovered by triangulating the incident light paths. Our approach can handle transparent objects with a



Facet label	1	2	3	4	5	6
Mean normal error (degree)	6.2654	9.9585	7.5905	9.6871	3.6591	6.8677
Median normal error (degree)	6.1677	9.5906	7.5109	9.7511	3.4741	6.7908
Mean position error (mm)	0.7250	0.6814	0.6675	0.6767	0.5881	1.0442
Median position error (mm)	0.6108	0.5945	0.5755	0.5721	0.5133	0.6333
RANSAC position inliers (%)	40.33	42.97	44.06	43.38	49.02	40.64

Table 3. Reconstruction error measurements of *ornament*. Left figure: shows the labels for each facet of the ornament. Right table: shows the various error metric used in reconstruction error evaluation for *ornament*. Due to its piecewise property, we fitted each facet using RANSAC with an inlier threshold of $0.5mm$ and then measured the distances from the FEPs to the fitted plane and also the angle difference between the reconstructed normals of each facet region and the fitted facet normal.

complex structure, with an unknown and even inhomogeneous refractive index. The only assumption to the objects is that the light paths should not reenter the liquid medium once they enter the object. If the refractive index of the liquid is known a priori, our method can also recover the surface normal at each reconstructed surface point.

References

- [1] M. Ben-Ezra and S. K. Nayar. What does motion reveal about transparency? In *ICCV*, volume 2, pages 1025–1032, 2003.
- [2] J.-Y. Bouguet. Camera calibration toolbox for matlab. http://www.vision.caltech.edu/bouguetj/calib_doc/.
- [3] V. Chari and P. Sturm. A theory of refractive photo-light-path triangulation. In *CVPR*, pages 1438–1445, 2013.
- [4] G. Eren, O. Aubreton, F. Meriaudeau, L. A. S. Secades, D. Fofi, A. T. Naskali, F. Truchetet, and A. Ercil. Scanning from heating: 3D shape estimation of transparent objects from local surface heating. *Optics Express*, 17(14):11457–11468, 2009.
- [5] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [6] S. Hata, Y. Saitoh, S. Kumamura, and K. Kaida. Shape extraction of transparent object using genetic algorithm. In *ICPR*, volume 4, pages 684–688, 1996.
- [7] M. B. Hullin, M. Fuchs, I. Ihrke, H.-P. Seidel, and H. P. A. Lensch. Fluorescent immersion range scanning. In *SIGGRAPH*, pages 87:1–10, 2008.
- [8] I. Ihrke, B. Goidluecke, and M. Magnor. Reconstructing the geometry of flowing water. In *ICCV*, volume 2, pages 1055–1060, 2005.
- [9] I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich. State of the art in transparent and specular object reconstruction. In *Eurographics STAR*, pages 87–108, 2008.
- [10] I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich. Transparent and specular object reconstruction. *Computer Graphics Forum*, 29:2400–2426, 2010.
- [11] Y. Ji, J. Ye, and J. Yu. Reconstructing gas flows using light-path approximation. In *CVPR*, pages 2507–2514, 2013.
- [12] K. N. Kutulakos and E. Steger. A theory of refractive and specular 3D shape by light-path triangulation. In *ICCV*, volume 2, pages 1448–1455, 2005.
- [13] K. N. Kutulakos and E. Steger. A theory of refractive and specular 3D shape by light-path triangulation. *IJCV*, 76:13–29, 2008.
- [14] D. Liu, X. Chen, and Y.-H. Yang. Frequency-based 3d reconstruction of transparent and specular objects. In *CVPR*, pages 660–667, 2014.
- [15] C. Ma, X. Lin, J. Suo, Q. Dai, and G. Wetzstein. Transparent object reconstruction via coded transport of intensity. In *CVPR*, pages 3238–3245, 2014.
- [16] D. Miyazaki and K. Ikeuchi. Inverse polarization raytracing: estimating surface shapes of transparent objects. In *CVPR*, volume 2, pages 910–917, 2005.
- [17] N. J. W. Morris and K. N. Kutulakos. Dynamic refraction stereo. In *ICCV*, volume 2, pages 1573–1580, 2005.
- [18] N. J. W. Morris and K. N. Kutulakos. Reconstructing the surface of inhomogeneous transparent scenes by scatter-trace photography. In *ICCV*, pages 1–8, 2007.
- [19] H. Murase. Surface shape reconstruction of an undulating transparent object. In *ICCV*, pages 313–317, 1990.
- [20] H. Murase. Surface shape reconstruction of a nonrigid transport object using refraction and motion. *PAMI*, 14(10):1045–1052, 1992.
- [21] S. G. Narasimhan, S. K. Nayar, B. Sun, and S. J. Koppal. Structured light in scattering media. In *ICCV*, pages 420–427, 2005.
- [22] M. O’Toole, J. Mather, and K. N. Kutulakos. 3d shape and indirect appearance by structured light transport. In *CVPR*, pages 3246–3253, 2014.
- [23] Q. Shan, S. Agarwal, and B. Curless. Refractive height fields from single and multiple images. In *CVPR*, pages 286–293, 2012.
- [24] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic reconstruction of transparent objects. In *SIGGRAPH 2006 Sketches*, 2006.
- [25] G. Wetzstein, D. Roodnick, W. Heidrich, and R. Raskar. Refractive shape from light field distortion. In *ICCV*, pages 1180–1186, 2011.
- [26] S.-K. Yeung, T.-P. Wu, C.-K. Tang, T. F. Chan, and S. Osher. Adequate reconstruction of transparent objects on a shoestring budget. In *CVPR*, pages 2513–2520, 2011.