Acquiring axially-symmetric transparent objects using single-view transmission imaging

Jaewon Kim Ilya Reshetouski Abhijeet Ghosh
Imperial College London
{jaewon.kim15,ghosh}@imperial.ac.uk

Abstract

We propose a novel, practical solution for high quality reconstruction of axially-symmetric transparent objects. While a special case, such transparent objects are ubiquitous in the real world. Common examples of these are glasses, tumblers, goblets, carafes, etc., that can have very unique and visually appealing forms making their reconstruction interesting for vision and graphics applications. Our acquisition setup involves imaging such objects from a single viewpoint while illuminating them from directly behind with a few patterns emitted from an LCD panel. Our reconstruction step is then based on optimization of the object's geometry and its refractive index to minimize the difference between observed and simulated transmission/refraction of rays passing through the object. We exploit the object's axial symmetry as a strong shape prior which allows us to achieve robust reconstruction from a single viewpoint using a simple, commodity acquisition setup. We demonstrate high quality reconstruction of several common rotationally symmetric as well as more complex n-fold symmetric transparent objects with our approach.

1. Introduction

Significant advances have been achieved in vision and graphics in the area of 3D scanning of objects including non-lambertian and specular surfaces. However, accurate modeling and reconstruction of transparent objects has continued to be a challenging problem due to the transmissive and refractive properties of such objects. Generally, the problem involves the estimation of optical properties such as index of refraction as well as shape estimation (multiple surfaces) of a target object. Previous work in vision and graphics has tried to tackle the more general problem of 3D reconstruction of transparent objects, achieving results with limited quality in practice with relatively simple optical setups [11], or requiring more complicated optical setups for volumetric reconstruction with multiview acquisition [12] or immersion in various types of liquids to simplify the image formation model [20, 5].

Figure 1: Examples of acquired axially-symmetric transparent objects rendered lit with two frontal area light sources.

In this work, we focus on high quality reconstruction for the special case of axially-symmetric transparent objects. We note that such transparent objects are ubiquitous in the real world with common examples being glasses, goblets, tumblers, carafes, etc. (see Fig. 1). Such everyday objects can have very unique and visually appealing forms which makes their reconstruction interesting for various vision and graphics applications. We propose a practical approach towards such reconstruction using a very simple setup with commodity components involving a camera and an LCD panel and exploit the inherent (vertical) axial-symmetry of the objects to acquire data from just a single viewpoint for efficient acquisition. Our reconstruction method takes an analysis-by-synthesis approach where we optimize the shape and refractive index of the target object by comparing the acquired and simulated (using ray-tracing) transmission/refraction of patterns emitted by the LCD panel in the background. We demonstrate high quality reconstruction for two classes of axially-symmetric transparent objects - those exhibiting complete rotational symmetry, and more complex objects exhibiting n-fold axial-symmetry - using our practical acquisition approach. Unlike works that tar-

get a volumetric reconstruction, we target a surface (mesh based) reconstruction which has the advantage of achieving reconstruction at near camera resolution (∼2K) without being memory limited for such high resolution.

The rest of the paper is organized as follows: we first cover some relevant previous work in Section 2. We then describe our practical acquisition setup and procedure in Section 3, before describing our analysis-by-synthesis approach for reconstructing rotationally symmetric transparent objects in Section 4. We then present the extension of the approach for more complex objects exhibiting n-fold axial symmetry in Section 5. Finally, we present additional reconstruction results and analysis of our method in Section 6.

2. Related Work

While there exists a vast literature on general 3D scanning, we will restrict the discussion here to acquisition and modeling of transparent objects. A recent survey on the topic can be found in [7]. We review some closely related works in the following:

Our acquisition setup is inspired by environment matting [24]. However, compared to an image based representation obtained with environment matting, we estimate a full 3D representation of a transparent object including shape and refractive index. Matusik et al. [14] proposed a data-driven technique for acquisition and rendering of transparent objects. Their approach is based on acquiring an approximate geometry using shape-from-silhouette reconstruction, coupled with view dependent reflectance field acquired using multiple cameras and light sources. While effective in creating a realistic rendering result, the approach is not suitable for simulating accurate light transport with the acquired model.

Model based reconstruction has previously been explored for transparent objects with simple known parametric shapes [2]. Closer to our approach, shapes of transparent objects have been estimated using direct ray measurements in [11, 4, 17] or time-of-flight distortion [18]. These methods however requires for each ray to cross no more than two interfaces, restricting its application to solid transparent objects that satisfy this condition. In comparison, our approach allows reconstruction of hollow axially-symmetric transparent objects with up to four interfaces.

Researchers have also investigated volumetric reconstruction of transparent objects. Here, immersion of the object in a liquid with matching index of refraction has been explored to prevent refractive bending of rays, enabling tomographic reconstruction from transmissive imaging [20, 5]. Additionally, immersion in a fluorescent liquid has been proposed for estimating the outer surface of a transparent object [5]. Fluorescent dye has also been employed for extinction based tomographic reconstruction of liquid volumes [6]. While very general solutions, these approaches may not always be desirable due to requiring immersion in various liquids.

Volumetric reconstruction of refractive transparent objects is a challenging inverse problem, known from the optics literature [21]. In particular, there were earlier attempts to recover axially-symmetric objects [13], but the proposed solutions were limited by complicated optical setup which are not easy to scale for larger objects and the methods proved to work in practice only for low levels of refraction inside the object. Recently, Ma et al. [12] have proposed a more practical setup for volumetric reconstruction of transparent objects using refractive tomography based on transport of intensity formulation. Their measurement approach involves a collimated beam source transmitting through a transparent object while the resulting refractions are imaged on screens placed at two different focal depths to estimate entry and exit ray correspondence for tomography. The object then needs to be rotated to observe such projections from multiple viewpoints for volumetric reconstruction. In comparison, our approach just requires acquisition from a single viewpoint using a very simple optical setup and exploits the object's inherent axial-symmetry as a strong shape prior for reconstruction. We also note that in contrast to tomographic approaches which tend to smooth out sharp features, our method enables high resolution reconstruction while preserving high frequencies in the reconstructed shape.

Also related is background oriented schlieren imaging which has also been applied to the problem of tomorgraphic reconstruction of gas flows and liquids [1]. Wetzstein et al. [23] have proposed a single image acquisition approach for reconstructing thin transparent surfaces using distortion of light field background illumination. The approach was extended to background illumination with light field probes as a way of imaging refractions in a transparent volume [22, 9]. These above approaches are however suitable for volumes with small magnitudes of refraction.

Also related to our method is the work of Miyazaki and Ikeuchi [15] who proposed an inverse raytracing framework (with polarization imaging) for reconstructing the outer surface of transparent objects from surface reflection. Finally, Morris and Kutulakos [16] have proposed scatter trace photography as solution for reconstructing the outer surface of complex transparent objects with inhomogeneous interiors. Their method works by separating the first surface reflection from the complex secondary bounces inside such objects. Our method purely operates on transmission/refraction through a clear transparent object while ignoring any surface reflection effects.

3. Acquisition and Setup

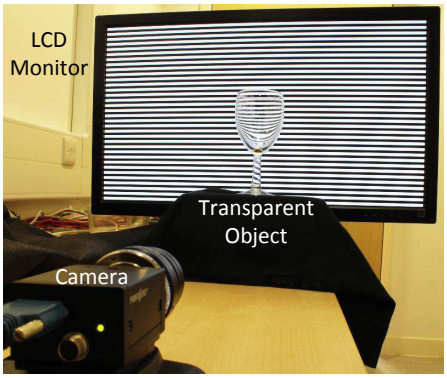

Figure 2: Acquisition setup consisting of a camera and an LCD monitor.

We employ a very simple "environment matting" like acquisition setup which consists of the target transparent object placed between a camera and an LCD panel (see Fig. 2). We employ a Point Grey Grasshopper3 (GS3-U3-41C6C-C 4.1MP CMOS) machine vision camera ($2K$ resolution), and an LG 27" LCD monitor (1920×1080 resolution) for our experiments. We place the object close to the monitor at about 5 cm distance in order to keep both the object and the panel in focus during acquisition. We then emit a few lighting patterns on the LCD panel while observing their distortion due to transmission through the object. This allows us to estimate a ray deflection map which we later employ for estimating the object's shape and refractive index.

In order to obtain accurate pixel-screen correspondence with a small set of measurements, we employ a combination of horizontal and vertical linear gradients (and their inverses) and a few high frequency gray codes (Fig.3, a). We employ the linear gradients to compute approximate[1] screen coordinates for camera rays and then refine the position estimate using the high frequency gray codes. We also employ the horizontal gradients later in our pipeline for inner shape estimation. This results in a capture sequence of 13 patterns from a single viewpoint - one constant white screen illumination, four patterns consisting of X and Y linear gradients and their inverses, and four patterns each of the X and Y high frequency gray codes. Each photograph is taken with 400 ms exposure time at full 2K resolution in 16 bit pixel depth, and the acquisition finishes in under 10 seconds. We also perform camera calibration using Bouguet's calibration toolbox [3].

In the following, we explain how the acquired data is useful for reconstructing axially-symmetric transparent objects exhibiting complete rotational symmetry (Section 4), and more complex n-fold axial-symmetry (Section 5).

[1]Employing only gradients is not precise enough in practice due to any potential optical non-linearities such as screen gamma/falloff and/or global illumination effects.

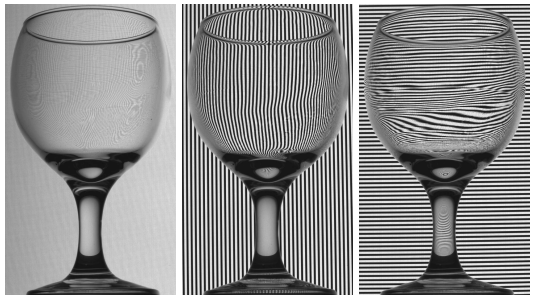

(a) Sample photographs with gradient and gray code patterns

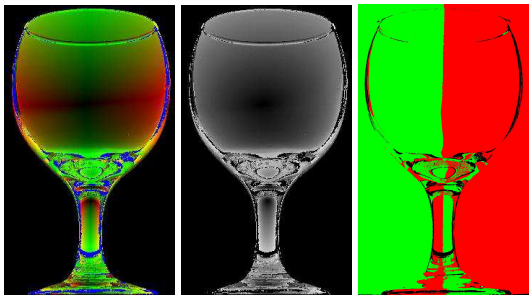

(b) Def. vector (c) Def. amp (d) Flip image

Figure 3: Acquired data used for computation of a deflection map. (a) Sample photographs taken with horizontal X gradient, and high frequency vertical and horizontal patterns respectively. (b) Visualization of deflection vectors with R, G encoding x, y deflection respectively, and B encoding missing data. (c) Deflection amplitude. (d) Flip image computed from deflection map used for identifying solid (2-interface) vs hollow (4-interface) sections.

4. Rotationally symmetric objects

Given the acquired data, our reconstruction pipeline for completely rotationally symmetric objects proceeds as follows: we first detect the outer silhouette of the object against the background and employ the 2D silhouette to reconstruct the outer surface through its rotation about the symmetry axis. We assume that the object consists of solid or hollow sections, with 2- or 4-interfaces respectively for ray traversal, which are detected next. The algorithm then focuses on the estimation of the object's refractive index as well as inner shape estimation of any hollow sections using an inverse rendering procedure. An overview of the reconstruction pipeline can be seen in Fig. 4.

4.1. Outer shape and rotation axis

We first detect the outer silhouette of the object using the two images acquired with a horizontal gradient and its inverse pattern on the LCD screen. The outer silhouette can

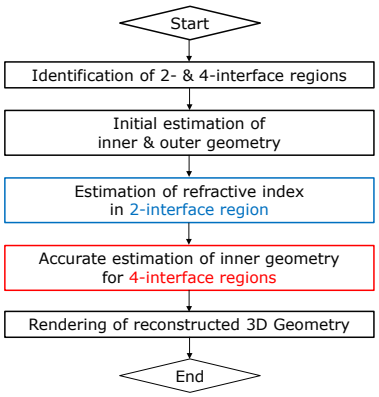

Figure 4: Proposed pipeline for reconstructing an axially-symmetric transparent object with solid and hollow sections.

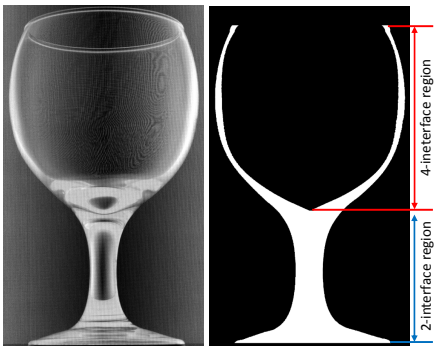

(a) Enhanced boundary (b) Approx. geometry

Figure 5: Initial estimation of inner and outer geometry. An edge enhanced image (a) is generated by subtracting a photograph under horizontal gradient (Fig. 3 top-left) from its inverse. Initial inner and outer geometry (b) is then estimated by processing (a).

be easily computed from the magnitude of the difference of these two gradient images (shown in Fig. 5, a) which enhances the vertical edges. The next step is to estimate the 3D position and orientation of the symmetry axis. In order to do this, we first find the $2D$ axis of symmetry on the camera image plane. Given the 2D outer silhouette, this can be done using line fitting. As the $2D$ symmetry axis is a projection of the object's $3D$ symmetry axis, we can calculate $3D$ position of the axis using known position of the LCD screen and the known distance from the screen to the center of rotation of the object (pre-calibrated). The LCD plane and the object's rotation axis are parallel to each other in our setup allowing us to project the axis from 2D to 3D. Finally, using the object's $3D$ axis of rotation, $2D$ silhouette and the camera's intrinsics, we can estimate the object's circular radius for each cross-section along the axis of rotation. We sample the rotation axis at the resolution of the 2D silhouette and create a rotational cross-section for each scan-line to define the outer shape.

4.2. Deflection map and flip-image

Before proceeding further with the reconstruction, we need to segment the object into separate solid and hollow sections with 2- and 4-interfaces respectively. The solid sections are purely described by the outer shape of the object. However, the hollow 4-interface sections need additional estimation of the inner shape and thickness. In order to do this, we first compute a ray deflection map of how the transparent object distorts the intersection of camera rays with points on the LCD screen due to refraction. Such deflection information is computed from the camera-screen point correspondence obtained using the gradient and gray code patterns. To effectively capture the deflection information, the data includes both 2D direction and magnitude (Fig. 3 b, c).

Using the above estimated deflection map and axis of symmetry, we compute a so called flip image (Fig. 3, d) necessary for our reconstruction pipeline. The flip image marks rays that cross over from one side of the symmetry axis to the other during propagation between screen and camera. Given the vertical axis of symmetry of the objects we acquire, this results in marking of rays that cross over from the left side of the axis (marked green) to the right side (marked red) or vice versa. A flip in the color coding with respect to the background then represents a solid section, while the absence of flip represents a hollow section for subsequent processing. This works well in our case due to the relatively high refractive index of the transparent objects ($1.3 - 1.6$) along with their circular cross-sections that cause the solid sections to flip rays about the symmetry axis.

4.3. Initialization of inner shape

After segmenting the object into solid and hollow sections, the next step is to initialize the inner shape of the object's hollow sections. This can be done using inner silhouette detection in conjunction with the computed flip image. First, we approximately determine the cross-sections along the rotation axis where there is a change in the number of interfaces using the previously obtained flip image. Next, we obtain an initial estimate of the thickness of the inner walls of the determined hollow sections. Initializing the wall thickness simply amounts to determining the inner radius of circular cross-sections orthogonal to the symmetry axis which we determine as follows:

1. Given the outer silhouette, we once again employ the edge enhanced image shown in Fig. 5 (a) to detect the inner silhouette. This is done by detecting the inner left and right edges along a scan line (cross-section) in the hollow 4-interface section of this enhanced image as local maxima of horizontal intensity gradients.

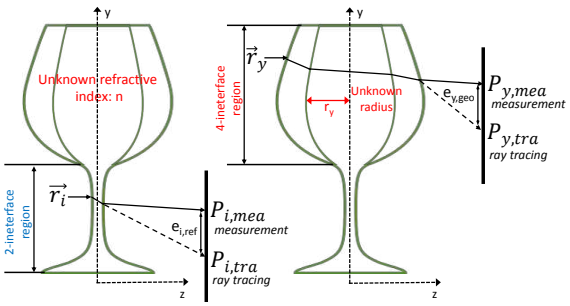

Figure 6: Ray diagrams for the unknown refractive index and inner geometry estimation based on inverse ray tracing.

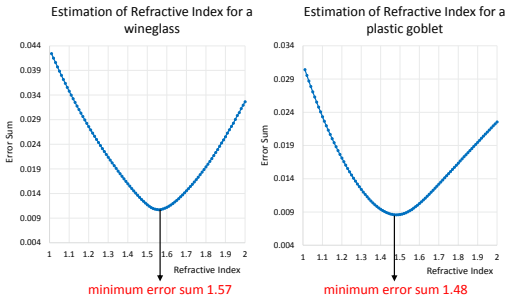

Figure 7: The refractive index is estimated by searching for a value that minimizes the error between ray traced and measured ray deflections. The error plots present the estimated η for a wineglass (left), and a transparent plastic goblet (right).

2. The thickness of a cross-section is then estimated as the mean of the left and right wall thickness. Finally, for smoothness of the estimated wall thickness across cross-sections, we fit a 4-th degree polynomial curve to the estimated cross-section thicknesses. This is done in order to filter out any noise in the inner silhouette detection and to obtain a smooth profile of the inner shape (Fig. 5, b).

These steps provide us with an initial estimate of the 3D shape of the object which we then refine using an inverse rendering procedure that estimates the object's refractive index and the final optimized shape.

4.4. Refractive index estimation

We take a sequential approach of first estimating the unknown refractive index η, followed by refining the final (inner) geometry. We restrict the refractive index estimation to the solid 2-interface section of the transparent object where the 3D shape is completely described by the outer silhouette. For each ray corresponding to pixels in the ROI (segmented using flip image), ray tracing is performed to compute the trajectory of refraction with a chosen η as shown in Fig. 6 (left). Then, the sum of error, $e_{i,ref}$, is calculated between the computed trajectory for the chosen η and a measured deflection map.

$$\underset{n}{arg\,min} \sum_{i\in[x,y]} e_{i,ref} = \underset{n}{arg\,min} \sum_{i\in[x,y]} |P_{i,mea} - P_{i,tra}| \tag{1}$$

This process is iterated over a sufficiently wide range of values for $\eta \in [1.0, 2.0]$ while searching for the value that minimizes error sum in Equation 1. Fig. 7 (left) shows the iterations for the wineglass in Fig. 3 with an estimated $\eta = 1.57$, which is in the range of crown and flint glass[2].

[2]Department of Physics and Astronomy, Georgia State University http://hyperphysics.phy-astr.gsu.edu/hbase/tables/indrf.html

4.5. Refinement of inner shape

Once the refractive index is determined, the last step of the pipeline is to refine the initial estimate of the inner shape of the hollow sections to obtain the final estimated 3D shape of the transparent object. The problem now reduces to estimating the final unknown radius r_y for each circular cross-section given a known refractive index using inverse rendering (Fig. 6, right). While the solid 2-interface section was used for estimating a single value of refractive index for the entire object, in this case the inverse rendering error sum is computed separately for each cross-section for estimating the final cross-section radii. Again, the sum of minimum error $e_{y,geo}$ is chosen as the best estimation of unknown radii for each cross-section of the hollow 4-interface section as given by Equation 2.

$$\underset{r_y}{arg\,min} \sum_{y} e_{y,geo} = \underset{r_y}{arg\,min} \sum_{y} |P_{y,mea} - P_{y,tra}| \tag{2}$$

We provide the initial wall thickness previously estimated in Section 4.3 as an initial guess to the shape optimization procedure for efficient convergence. Figure 8 (a) shows the initial estimate of the wall thickness in blue and the final estimated result for the inner radii in red. Note that the singular curvature between lines 650 and 700 is the result of a more accurate estimation for the actual shape of the bottom of the glass bowl after the optimization. This singular curvature creates the bright oval pattern in the corresponding area of the reference photograph in Fig. 8 (b) with constant backlighting. Additionally, various other noticeable patterns are created around the interface between the solid and hollow sections and near the solid base of the wine glass in the photograph that arise due to the high curvature and geometry of these sections of the glass. The rendering result in Fig. 8 (c) shows a faithful reproduction of much of these patterns in addition to the overall shape and appearance compared to the photograph. We present additional results of such reconstruction in Section 6.

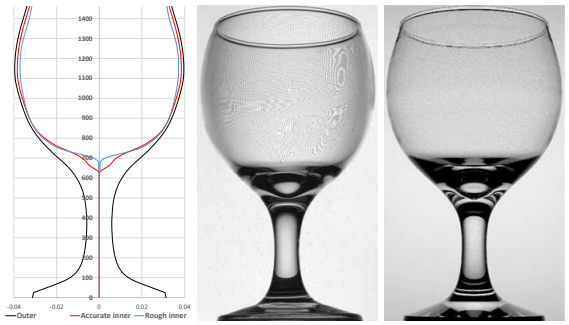

(a) Inner geometry (b) Photograph (c) Rendering

Figure 8: Final estimated inner geometry and rendering result. (a) The estimated inner radii for the hollow section is marked in red compared to the initial estimate shown in blue. The final object rendering (c) is a good match with reference photograph (b).

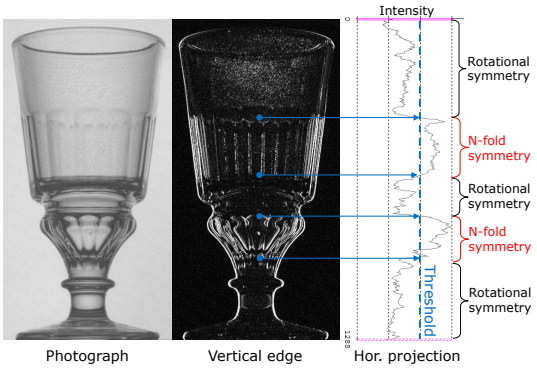

Figure 9: Separation of rotational and n-fold symmetry sections using amplified vertical edges.

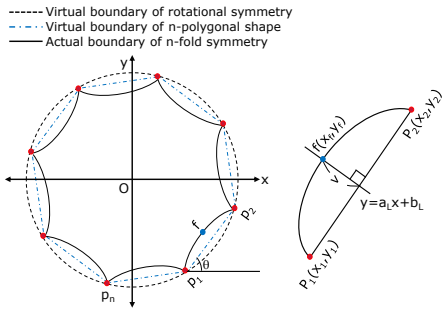

Figure 10: Notation for modeling curved faces in an n-fold symmetric region.

5. N-fold symmetric objects

The reconstruction method for a rotationally symmetric object as described in the previous section can be extended for the reconstruction of an n-fold symmetric transparent object with similar solid and hollow sections. We assume that the outer cross-section is n-fold symmetric while the inner cross-section is still circular which is true for many everyday n-fold symmetric objects such as goblets, tumblers, etc. We observe that in most cases such objects can be categorized as a combination of sections with complete rotational symmetry and sections with n-fold symmetry as shown in Fig. 9. We propose a reconstruction method for the n-fold symmetric sections using a polygonal shape for the outer cross-sections. The method requires a single user input which is the vertex number of a polygon. The first step of the method is separating any rotationally symmetric sections from the n-fold symmetric sections. We observe that for vertical axially symmetric objects, an n-fold symmetric section has stronger vertical edges than a rotationally symmetric section as shown in Fig. 9 (center). A horizontal projection over the vertical edges intensifies the edge signals for the n-fold symmetric sections and allows us to separate them using a simple threshold (Fig. 9, right).

Unlike a rotationally symmetric section for which the outer cross-section is circular, a specific modeling method is required for generating the outer geometry of an n-fold symmetric section. We model the n-fold geometry as curved lines between vertices of an n polygonal shape which is inscribed by the virtual circular shape of rotational symmetry (Fig. 10). The curved n-fold shape also includes the planar n-polygonal shape as a special case when the curvature, v, is zero. We model the curved lines as a quadratic function with vertex f which is parameterized by the local curvature, v. The fixed position of each n-polygonal vertex, p_i is specified by the user input of the vertex number of the polygon. The curved line between points p_1, p_2, and f is described by the standard quadratic equation $ax^2+bx+c=y$. To generate the various faces of the n-fold curve, the quadratic is rotated by θ which is the slope angle of line p_1p_2. Equation 3 then expresses the constant terms of the rotated quadratic equation $ax'^2+bx'+c=y'$, where x_i and y_i are x and y coordinates of p_1, p_2, and f. Note that the relative orientation of the n-polygon with respect to the camera is obtained by rotating the model by 1° intervals about the symmetry axis (within a search range $[0-90]$) and comparing the 2D projection of each rotation with the reference photograph.

$$
\begin{aligned}
a &= \frac{(B_1-B_3)(A_2-A_3)-(B_2-B_3)(A_1-A_3)}{(A_1^2-A_3^2)(A_2-A_3)-(A_2^2-A_3^2)(A_1-A_3)} \\
b &= \frac{(B_1-B_3)-a(A_1^2-A_3^2)}{(A_1-A_3)} \\
c &= B_1-aA_1^2-bA_1 \\
&where\ A=x_i'cos\theta+y_i'sin\theta, B=y_i'cos\theta-x_i'sin\theta
\end{aligned}
\tag{3}
$$

Examples of curved shape modeling for the n-fold symmetric sections are shown in Fig. 11 including faces with concave, flat, and convex shapes respectively. Note that the

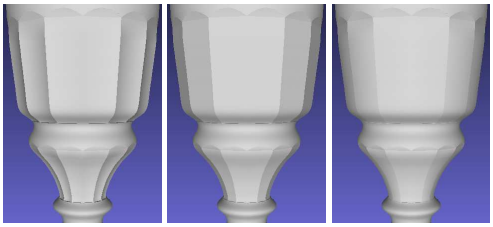

Figure 11: Examples of curved face modeling for the n-fold symmetric cross-sections of the object in Figure 9 including concave, planar, and convex modeling respectively.

local curvature of the n-fold outer faces is not known initially. Hence, we propose to iteratively estimate the inner shape and the outer n-fold face curvature in a sequential manner. We first initialize the outer geometry with planar faces for the n-fold sections and employ this outer shape approximation to estimate the inner shape (with circular cross-sections) using inverse rendering as described in Section 4. This results in a semi-accurate reconstruction of n-fold symmetric sections (Fig. 12, first column) with larger residual errors in sections that have considerable curvature in the outer faces (e.g., the lower n-fold region shown in the bottom of the first column), and low residual errors in sections with negligible curvature in the outer faces (e.g., the upper n-fold section). In a second phase, we then optimize for the face curvature, v, of the n-fold faces by minimizing the residual error between ray tracing and measured deflections (Fig. 12, second column). For a candidate curvature value, the unknown vertex point f of the rotated quadratic equation is specified by Equation 4.

$$
\begin{aligned}
&(1+a_L^2)x_f^2 + (2B'a_L - 2A')x_f + A'^2 + B'^2 - v^2 = 0 \\
&y_f = a_L x_f + b_L \\
&where\ A' = \frac{x_1 + x_2}{2}, B' = b_L - \frac{y_1 + y_2}{2}
\end{aligned}
\tag{4}
$$

Since the equation is quadratic its two solutions suggest two possible positions for the vertex, one creating a concave shape and the other convex. We try a range of curvature candidates spanning from concave to convex outer geometry to search for a configuration with minimal residual error. We then iterate the sequential estimation of inner geometry and outer face curvature until the change in residual errors is lower than a threshold, giving us the final reconstruction result (Fig. 12, last column). The final rendering result is comparable to the reference photograph in Figure 9. We present an additional example in the next section.

6. Results

We have tested our acquisition and reconstruction method for a wide range of axially symmetric transparent objects.

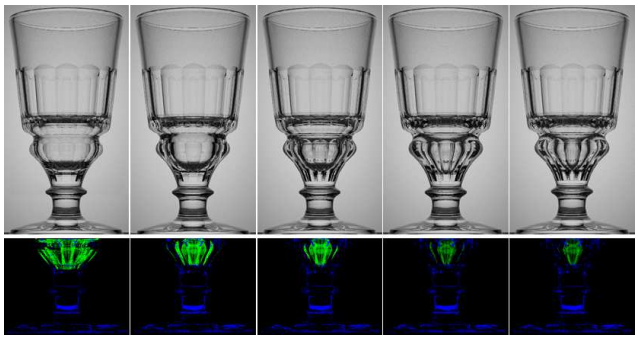

Figure 12: Rendering results (top row) and corresponding residual errors (bottom row) with iterative estimation of inner geometry and curvature of n-fold faces. The first and the second columns are initial inner geometry estimation with flat surface and initial curvature estimation for n-fold faces, respectively. The fourth and the fifth columns are final estimation for inner geometry and face curvature, respectively.

Fig. 1 presents a realistic rendering of various acquired glass objects lit with two frontal area light sources using Mitsuba [8]. We also present rendering comparisons to photographs of the acquired objects under constant screen background illumination as seen in Figs. 8 and 12. A few additional representative results can be seen in Fig. 13. Note that the cocktail glass (a) has two distinct hollow sections including a small spherical section at the bottom that is well reconstructed using our approach. (b) presents a rotationally symmetric champagne glass, and (c) a pint glass. Here, the pint glass has a few labels on the surface that cause occlusions for transmission measurements in our setup. Despite this, our reconstruction is fairly accurate demonstrating the robustness of the approach for shape estimation. (d) presents a more complex example of n-fold axial symmetry: a glass tumbler with lozenge-pattern. Here, we provided the number of lozenges in a cross-section as an input to the reconstruction algorithm. Finally, we present the reconstruction result of a shot glass containing some vodka in (e). Here, we first acquired an empty shot glass in order to accurately recover the inner geometry of the glass. We then recaptured the glass filled with vodka in order to estimate the liquid's refractive index ($\eta = 1.36$) for the rendering. As can be seen, unlike some previous approaches which require immersion of a transparent object in various liquids, our approach scales to also acquiring axially symmetric transparent objects containing transparent liquids.

The refractive index estimation requires around 10 minutes on a machine with Intel iCore7 2.5 GHz quad-core processor and 16 GB RAM. The inner geometry optimization requires around 30 minutes for a rotationally symmetric object, while an n-fold symmetric object takes around 90 minutes due to the iterative estimation of both inner and outer shape. In order to quantitatively verify the accuracy

(a) Cocktail glass (b) Champg. gl. (c) Pint glass (d) Tumbler with lozenge pattern (e) Shot glass containing vodka

Figure 13: Additional examples of reconstructed transparent objects. Left: photographs. Right: renderings.

(unit: mm)

	Rotational symmetry	N-fold symmetry
Avg. Error	0.351	0.859
St. dev.	0.171	0.321

Table 1: Reconstruction accuracy compared to physical measurements with vernier calipers.

of our reconstruction, we made physical measurements of the cross-section thicknesses using a high precision vernier caliper (Mitutoyo 500-196-30, 0.01mm res.). We did this because of difficulties in comparison against alternate acquisition approaches. For example, scanning with powder coating would only enable acquisition of the outer surface as the inner surface would be occluded for many of the acquired objects, while dyeing would still require somehow eliminating refractions for reconstruction (e.g., immersing in refractive index matching liquid). Instead, we made several cross-sectional caliper measurements (at 10 different positions) of six different objects (three each in the two categories of rotationally and n-fold symmetric) to compute the mean and std. deviation of the reconstruction error (Table 1). The mean error is less than 1mm in both cases. As expected, we achieve higher accuracy for completely rotationally symmetric objects due to their simpler cross-sections. We also obtain consistent estimates of η for various transparent solids and liquids [10, 19]. We include a few additional results in the supplemental material.

Discussion and Limitations: Our approach is specific to axially symmetric objects and we take advantage of this symmetry for robust single view reconstruction. Given the strong symmetry assumption, any minor asymmetries in the actual object are not reconstructed. We further rely on model based reconstruction for objects with n-fold symmetry. However, due to an extensive set of possible variations, it is somewhat necessary to adopt the surface modeling for individual types of such symmetry (e.g., lozenge pattern). It might be possible to extend the approach somewhat to reconstruct more general shapes with multiview acquisition. However, there are limits to how many interfaces can be resolved in the general case [11]. While less general than tomography based volumetric reconstruction, our approach can produce very high quality results which preserve high frequency shape features due to near camera resolution estimation of discrete refractive boundaries. Our renderings under constant background illumination are a good qualitative match to photographs. However, there are some noticeable differences due to us not accurately modeling the angular fall-off of the LCD screen illumination or its polarization characteristics in the renderings. We rely on being able to observe ray deflections through a solid 2-interface section for estimating the refractive index and currently only employ the green channel (central wavelength) data for this purpose. Our approach will require modifications for acquisition of birefringent transparent objects. The accuracy of our reconstruction is also limited by the resolution of the LCD screen which impacts the deflection map resolution.

7. Conclusions

We have presented a very practical approach for high quality reconstruction of axially symmetric transparent objects. Such objects are quite common in the real world and can have very unique, aesthetic and complex shape and appearance. Our approach employs a simple environment matting style setup for efficient single view acquisition and robust reconstruction of such transparent objects including estimation of shape and refractive index. We demonstrate high quality reconstruction results for a wide range of rotationally symmetric and n-fold symmetric everyday objects. For these classes of objects, we achieve significantly better qualitative results compared to prior work targeting more general transparent object reconstruction.

8. Acknowledgements

We thank Santa Ghosh for support and assistance. This work was supported by an EPSRC Early Career Fellowship EP/N006259/1, Royal Society Wolfson Research Merit Award, and 'The Cross-Ministry Giga KOREA Project' granted from the Ministry of Science, ICT and Future Planning Korea.

References

[1] B. Atcheson, I. Ihrke, W. Heidrich, A. Tevs, D. Bradley, M. Magnor, and H.-P. Seidel. Time-resolved 3D capture of non-stationary gas flows. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 27(5), 2008. 2

[2] M. Ben-Ezra and S. K. Nayar. What does motion reveal about transparency? In *9th IEEE International Conference on Computer Vision (ICCV)*, 2003. 2

[3] J.-Y. Bouguet. Camera Calibration Toolbox for Matlab, 2005. 3

[4] K. Han, K.-Y. K. Wong, and M. Liu. A fixed viewpoint approach for dense reconstruction of transparent objects. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2

[5] M. B. Hullin, M. Fuchs, I. Ihrke, H.-P. Seidel, and H. P. A. Lensch. Fluorescent Immersion Range Scanning. *ACM Trans. on Graphics (SIGGRAPH'08)*, 27(3):87:1 – 87:10, 2008. 1, 2

[6] I. Ihrke, B. Goldluecke, and M. Magnor. Reconstructing the Geometry of Flowing Water. In *15th IEEE International Conference on Computer Vision (ICCV)*, pages 1055–1060, 2005. 2

[7] I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich. Transparent and specular object reconstruction. *Computer Graphics Forum*, 29(8):2400–2426, 2010. 2

[8] W. Jakob. Mitsuba Physically Based Renderer, 2010. 7

[9] Y. Ji, J. Ye, and J. Yu. Reconstructing gas flows using light-path approximation. In *Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition*, CVPR '13, pages 2507–2514, 2013. 2

[10] S. N. Kasarovaa, N. G. Sultanovaa, C. D. Ivanovb, and I. D. Nikolovb. Analysis of the dispersion of optical plastic materials. *Optical Materials*, 29(11):1481 – 1490, 2007. 8

[11] K. N. Kutulakos and E. Steger. A theory of refractive and specular 3d shape by light-path triangulation. In *10th IEEE International Conference on Computer Vision (ICCV), 17-20 October, Beijing, China*, pages 1448–1455, 2005. 1, 2, 8

[12] C. Ma, X. Lin, J. Suo, Q. Dai, and G. Wetzstein. Transparent Object Reconstruction via Coded Transport of Intensity. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 1, 2

[13] Y. Maruyama, K. Iwata, and R. Nagata. Determination of axially symmetrical refractive index distribution from directions of emerging rays. *Appl. Opt.*, 16(9):2500–2503, Sep 1977. 2

[14] W. Matusik, H. Pfister, R. Ziegler, A. Ngan, and L. McMillan. Acquisition and Rendering of Transparent and Refractive Objects. In P. Debevec and S. Gibson, editors, *Eurographics Workshop on Rendering*. The Eurographics Association, 2002. 2

[15] D. Miyazaki and K. Ikeuchi. Shape estimation of transparent objects by using inverse polarization raytracing. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(11):2018–2030, 2007. 2

[16] N. J. Morris and K. N. Kutulakos. Reconstructing the Surface of Inhomogeneous Transparent Scenes by Scatter Trace Photography. In *11th IEEE International Conference on Computer Vision (ICCV)*, 2007. 2

[17] Y. Qian, M. Gong, and Y. Hong Yang. 3D reconstruction of transparent objects with position-normal consistency. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2

[18] K. Tanaka, Y. Mukaigawa, H. Kubo, Y. Matsushita, and Y. Yagi. Recovering transparent shape from time-of-flight distortion. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2

[19] The-Foundry. Refractive Index Table. 8

[20] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic Reconstruction of Transparent Objects. In *Proc. Eurographics Symposium on Rendering*, pages 51–60, 2006. 1, 2

[21] C. M. Vest. Tomography for properties of materials that bend rays: a tutorial. *Appl. Opt.*, 24(23):4089–4094, 1985. 2

[22] G. Wetzstein, R. Raskar, and W. Heidrich. Hand-Held Schlieren Photography with Light Field Probes. In *IEEE International Conference on Computational Photography (ICCP)*, 2011. 2

[23] G. Wetzstein, D. Roodnick, R. Raskar, and W. Heidrich. Refractive Shape from Light Field Distortion. In *15th IEEE International Conference on Computer Vision (ICCV)*, 2011. 2

[24] D. E. Zongker, D. M. Werner, B. Curless, and D. H. Salesin. Environment matting and compositing. In *Proceedings of ACM SIGGRAPH 99*, pages 205–214, July 1999. 2