

A Theory of Refractive Photo-Light-Path Triangulation

Visesh Chari

Projet WILLOW, INRIA Paris
23 avenue d'italie, Paris, France

visesh.chari@inria.fr

Peter Sturm

Projet STEEP, INRIA Rhône Alpes
655, avenue de l'Europe, Montbonnot, France

peter.sturm@inria.fr

Abstract

3D reconstruction of transparent refractive objects like a plastic bottle is challenging: they lack appearance related visual cues and merely reflect and refract light from the surrounding environment. Amongst several approaches to reconstruct such objects, the seminal work of Light-Path triangulation [17] is highly popular because of its general applicability and analysis of minimal scenarios. A light-path is defined as the piece-wise linear path taken by a ray of light as it passes from source, through the object and into the camera. Transparent refractive objects not only affect the geometric configuration of light-paths but also their radiometric properties. In this paper, we describe a method that combines both geometric and radiometric information to do reconstruction. We show two major consequences of the addition of radiometric cues to the light-path setup. Firstly, we extend the case of scenarios in which reconstruction is plausible while reducing the minimal requirements for a unique reconstruction. This happens as a consequence of the fact that radiometric cues add an additional known variable to the already existing system of equations. Secondly, we present a simple algorithm for reconstruction, owing to the nature of the radiometric cue. We present several synthetic experiments to validate our theories, and show high quality reconstructions in challenging scenarios.

1. Introduction

Reconstruction of transparent objects has gathered interest in the last few years [5, 8, 11, 12, 13, 14, 17]. Several cues have been proposed to reconstruct such objects. Methods could be broadly classified into approaches that rely on physical (material) properties of transparent objects, and approaches that try to extend traditional shape acquisition approaches to the case of transparent objects. Among the approaches relying on material properties, geometric and radiometric cues are the most prominent inputs to reconstruction algorithms.

Transparent objects referred to in the recent computer vi-

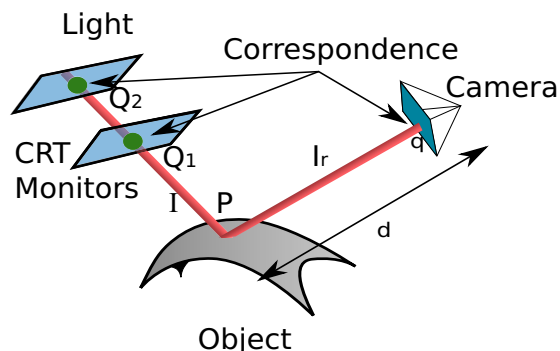


Figure 1: Given a transparent refractive object, a light source and a camera, the pixel q receives light from sources at positions Q_1 and Q_2 . The method of light-path triangulation [17] reconstructs point P by triangulating the line passing through Q_1 and Q_2 with the line of sight of q . Thus two measurements Q_1 and Q_2 are necessary. In contrast, our method uses only one of the points Q_1 or Q_2 while using the intensity ratio I_r/I as the additional radiometric cue. Thus, in contrast to two observations required by [17] to compute depth d , we need only one observation. Both methods estimate the normal at P as well. In both works Q_1/Q_2 is estimated by using a calibrated computer monitor as light source (CRT monitor in our case).

sion literature could be categorized into two kinds: specular and translucent. In this paper, we focus on specular transparent objects, that is objects for which incoming light is partly reflected off the surface and partly entering the object, after undergoing a refraction at the surface. For such objects, geometric cues are an important source of information. Following the popular work of Kutulakos *et al.* [17], several approaches have sought to analyze geometric cues to obtain an accurate reconstruction from multiple images [10, 20]. While these works consider specular surfaces (both mirrors and transparent surfaces), they do not utilize the fact that transparent objects leave a shape dependent radiometric signature in images.

Since transparent surfaces also reflect light falling on

them, photometric stereo and related algorithms have also been proposed to reconstruct the exterior of transparent objects [20, 13]. A recent focus on specularities for reconstruction is one such example [20]. While reconstruction using such methods gives good results, we argue that modeling transparent surfaces as pure mirrors discards important information about their physical properties (refractive index, internal structure, etc.).

In this paper, we present an approach that combines the merits of utilizing both geometric and radiometric cues. Our approach is along the same lines as [17], with one important difference. For every “light-path” that we capture, we record *both* geometric information (position and direction of light rays captured by and originating from light source, depending on requirement) *and* radiometric information (radiance of light at the beginning and end of each light-path). A first main result is that this allows to reconstruct a transparent object from a single configuration of the acquisition setup, i.e. without moving the object, the camera, or the CRT monitor (Figure 1). More generally, we extend the minimal case scenarios of reconstruction in several ways beyond those proposed in Kutulakos *et al* [17], reducing the number of measurements required for reconstruction.

The paper is structured as follows. Section 2 places previous works in perspective. Sections 3 and 4 present basic theory and the basic reconstruction approach proposed in this paper. Possible extensions, e.g. towards using images obtained after two reflections or refractions, are explained in section 5. Sections 6 and 7 provide practical details as well as experimental results. Finally, section 8 gives an overall summary of our approach and points to future directions of research.

Notation: In the rest of the paper, we use the following convention. Bold symbols like \mathbf{E} denote vectors with direction and magnitude, while $\hat{\mathbf{E}}$ respectively E denote the corresponding (unit) direction and magnitude (or scalar).

2. Related Work

In the past, several approaches have used either geometric or radiometric cues to reconstruct transparent surfaces. Geometric approaches typically measure the deviation from perspective imaging produced by a refractive transparent object, and recover the shape as a solution that explains the observation [21]. In [17], the authors present a minimal case study of the conditions in which refractive surfaces can be reconstructed. They **re-cast** transparent object reconstruction as the task of reconstructing the path of **each individual light-path** that is captured by a camera after refraction through a transparent surface. Earlier approaches have focused on parameterizing the object to be reconstructed and then obtaining its parameters by explaining the distortion it produces in several cameras [4]. Other examples of shape recovery from distortion analysis include the more recent

work by [13], which analyzes the specific case of a single dynamic transparent surface that distorts a known background and is observed by multiple cameras. Finally, some recent approaches have also looked at learning based solutions that minimize a certain criterion of optimality [6, 19].

Apart from geometry, radiometric information also turns out to be very important in the case of transparent objects since they simultaneously reflect and refract light. Since they reflect light like a specular surface, many recent photometric approaches have tried to reconstruct transparent surfaces by studying their specularities. While [20] provides a low cost approach to reconstruction by studying specular highlights, [13] shows how to reconstruct transparent surfaces with inhomogeneous refractive interiors, by measuring highlights multiple times to remove extraneous effects like scattering, interreflections etc. One important aspect of radiometry is polarization. When unpolarized light is reflected or transmitted across a dielectric refractive surface, it gets partially polarized. This degree of polarization is connected to the shape of the object, and several authors have attempted to explore this connection. In [9] the authors measure the polarization angle using multiple images from a single view taken with different orientations of a polarizer. They then use Fresnel theory to derive a relationship between measured polarization state and the angle of reflection. In [11] authors use a predict-observe-minimize loop to estimate the shape of the object using a technique called polarization ray tracing.

Since refractive objects present a challenging reconstruction problem, many authors have resorted to using active approaches for reconstructions. Methods like fluorescent immersion range scanning [8] and tomographic reconstructions present alternate approaches that are of practical value when objects are of manageable size.

3. Reflection and Refraction Caused by Transparent Surfaces

In this section, we describe an image formation model for transparent surfaces, that forms the basis of our reconstruction approach. The model is based on well-known optical phenomena.

Let X be a calibrated point light source emitting unpolarized light. By calibrated, we mean here that we know the “amount” of light emitted in every direction (see section 6 for practical approaches/considerations).

Consider light emitted by X in a particular direction. If the light ray hits a transparent surface, part of its energy gets reflected in a mirror direction and part enters the transparent object, after undergoing a refraction at the surface. Both the reflection and the refraction happen within a **plane of refraction** $\hat{\pi}$ that is spanned by the point of intersection of the light ray and the surface, and the surface normal at that point. The **geometric** aspects of reflection and refraction

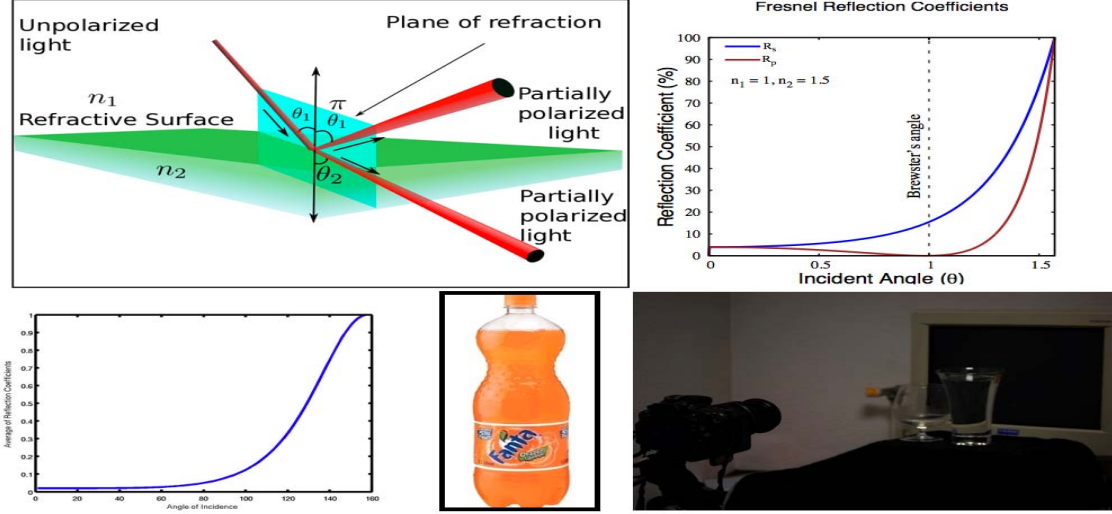


Figure 2: (left top) Description of the general theory behind our approach. While the acquisition is similar to that of Kutulakos *et al* [17], we also include radiance in our measurements (depicted by the changing color of rays while they travel from the illuminant to the camera pixel). (right top) Reflection coefficients computed using Fresnel equations. Courtesy [1] (left bottom) Average of reflection coefficients w.r.t incident angle. Note that the curve suggests unique solution for single bounce cases. (centre bottom) Sample objects used for reconstruction: A Fanta bottle and wine glasses. (right bottom) Simple acquisition setup. A camera facing a monitor with transparent objects in between.

are given by Snell’s laws. Let θ_1 be the angle between the incoming light ray and the surface normal (cf. Figure 2). The reflected light ray forms the opposite angle θ_1 with the normal whereas the angle θ_2 , formed by the refracted ray and the normal, is given by:

$$\sin \theta_2 = \sin \theta_1 \frac{n_1}{n_2} \quad (1)$$

where n_1 and n_2 are the refractive indices of the media on both sides of the surface (n_2 for the inside, n_1 for the outside, where the light source is located).

The **radiometric** aspects are as follows. Let I be the irradiance of the light source in the considered direction. The irradiance of the reflected ray (respectively refracted ray), is given by [1]:

$$I_r = \frac{\sin^2(\theta_1 - \theta_2)}{2 \sin^2(\theta_1 + \theta_2)} \left(1 + \frac{\cos^2(\theta_1 + \theta_2)}{\cos^2(\theta_1 - \theta_2)} \right) I \quad (2)$$

$$I_t = I - I_r \quad (3)$$

where subscript r is used for “reflected” and t for “transmitted” (refracted).

Note that even for unpolarized light sources, the reflected and refracted light will in general be polarized. This is important when studying what happens in case of a second “bounce”, i.e. if light gets reflected or refracted a second time (see section 5.2).

4. Surface Reconstruction Approach

In this section, we combine the geometric and radiometric properties of a light-path, given by equations (1), (2) and (3), in order to reconstruct the surface point P (Figure 1), and thus the object. We use a fully calibrated setup, consisting of a camera and a CRT monitor. The main result of this section is to show that images acquired *without* any displacement of the object, camera, or monitor, allow to reconstruct the object’s shape, and to provide a method for doing so. This was not possible in [17].

4.1. Surface Depth Reconstruction: Reflection

Surface reconstruction is done for individual camera pixels, by estimating the depth of the surface along the lines of sight of pixels. Regularization, such as surface smoothness, may then be added in a subsequent optimization stage. In the following, we first deal with the case of reflections, i.e. the camera acquires images of the monitor, reflected in the object to be reconstructed. Below, we handle the analogous case of refraction.

Let us now consider a single pixel and the acquired intensity I_r . The pixel’s line of sight is known by calibration. Let d be the depth of the object along this line of sight, P be the intersection point of the object surface and the line of sight, and \hat{n} the surface normal at that point. Further, given the matching (cf. Section 4.3), we know the point Q on the monitor whose reflection is seen in the pixel.

Our goal is to compute the depth d . We do so by first computing the incident angle θ_1 between the surface normal

and the incident light ray, after which d is trivial.

Since our setup is radiometrically calibrated, we equate the *radiance ratio* $r = \frac{I_r}{I_t}$ to the reflected and refracted angles using equation 2

$$r = \frac{\sin^2(\theta_1 - \theta_2)}{2 \sin^2(\theta_1 + \theta_2)} \left(1 + \frac{\cos^2(\theta_1 + \theta_2)}{\cos^2(\theta_1 - \theta_2)} \right)$$

Note that although we are considering the case of reflection here, the refracted angle θ_2 nevertheless appears in the equation, due to the object's refractive property.

From the above equation, using (1) and elementary trigonometric identities, one obtains the following constraint in θ_1 (details omitted due to lack of space, ref [1]):

$$\begin{aligned} & s_1^8 \{ 4n_{12}^4 (2r(n_{12}^2 + 1)^2 - (n_{12}^2 - 1)^2) \} + \\ & s_1^6 \{ 4n_{12}^2 ((n_{12}^6 - 3n_{12}^4 - n_{12}^2 + 1) - 2r(n_{12}^2 + 1)(n_{12}^4 + 4n_{12}^2 + 1)) \} \\ & + s_1^4 (2r(n_{12}^8 + 12n_{12}^6 + 30n_{12}^4 + 12n_{12}^2 + 1) \\ & - r^2(n_{12}^4 - 1)^2 - (n_{12}^8 - 4n_{12}^6 + 18n_{12}^4 - 4n_{12}^2 - 1)) \\ & + s_1^2 (2r^2(n_{12}^2 - 1)^2(n_{12}^2 + 1) - 2r(n_{12}^6 + 14n_{12}^4 + 13n_{12}^2 + 2) \\ & - 2(5n_{12}^4 + 2n_{12}^2 - 1)) \\ & + \{ 2r(5n_{12}^2 + 1) - r^2(n_{12}^2 - 1)^2 + 4n_{12}^2 - 1 \} = 0 \quad (4) \end{aligned}$$

Here, $n_{12} = \frac{n_1}{n_2}$ and $s_1 = \sin \theta_1$.

We observe that this is a quartic polynomial in $\sin^2 \theta_1$, i.e. the computation of θ_1 can be considered as being (close to) a closed-form solution.

As mentioned above, given an estimate for θ_1 , the depth d can be computed easily. Among the up to 8 possible real solutions for θ_1 , at most 4 will correspond to a positive depth, i.e. a surface point in front of the camera. Finding a unique θ_1 and thus depth, is made possible in most practical circumstances, as follows. First, the absolute value of θ_1 must be below 90° . Second, θ_1 is typically (much) larger than 30° , due to the practical setup which requires that the camera have both a reflected and a direct view of the monitor. Consider the graph of r as a function of θ_1 for the refractive index of water ($n_2 = 1.33$), in Figure 2 (here, both the camera and the monitor, are in air, i.e. $n_1 = 1$). One observes that for the values of r associated with $\theta_1 \in (30^\circ, 90^\circ)$, there is a unique θ_1 producing these values. To be precise, θ_1 is unique up to sign, but only the positive solution corresponds to a point in front of camera.

Let us summarize the above findings. From images acquired with a completely static setup, we are able to compute the depth of the transparent object, for each pixel in which a reflection is visible. To do so, we need to know the refractive index of the object's material. A unique solution for depth is possible in a large range of practical conditions. In case these are not fulfilled, one may still use bounds on the object depth to get a unique solution.

4.2. Surface Depth Reconstruction: Refraction

The case of refraction can be solved completely analogously to that of reflection. We also obtain a quartic poly-

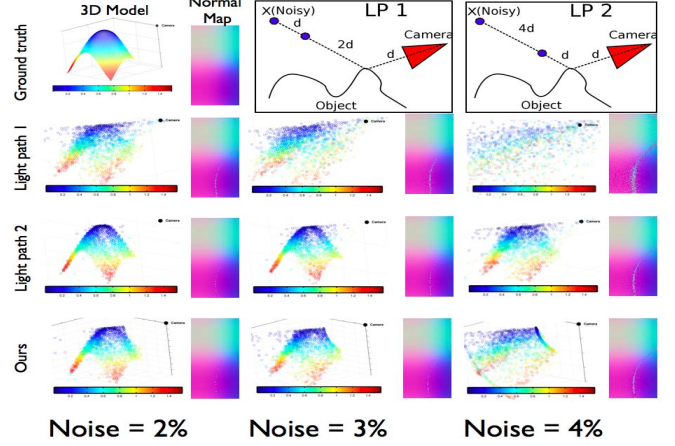


Figure 3: Comparison of simulations between our approach and [17]. In LP-1, corresponding 3D points are normally close to the object, which results in increased error in triangulation and normal estimation. Note that in the same scenario, we have much better normal information because of radiometric information. While LP-2 is robust because correspondences are far away, its highly impractical since use of LCD's for correspondence is problematic (because of light fall-off, scattering etc.). Details in text.

nomial in $\sin^2 \theta_1$ and the rest of the discussion is as above. The main difference is however that in order to observe the refraction, the camera must be inside the same medium as the object or, be located in a medium with the same refractive index as that of the object. For example, a camera looking at a monitor immersed in water (or, more realistically, put next to an aquarium's bottom plate), might allow to reconstruct the water's surface.

4.3. Acquisition

In principle, surface reconstruction is thus possible from a single image. In practice however, we require dense matches between the camera image and the CRT monitor, which is why we acquire multiple images of a sequence of Gray-coded patterns displayed on the monitor [2]. From this, we determine dense matches, i.e. for each individual camera pixel we determine the point on the monitor whose reflection or refraction in the object, is seen by that pixel.

The camera-monitor setup is supposed to be fully calibrated, both geometrically and radiometrically: the camera's intrinsic parameters are known, as well as its position and orientation relative to the monitor. Further, the monitor screen is supposed to be a planar surface. Radiometric calibration of both, camera and monitor, is known too.

5. Possible Extensions

5.1. Unknown Refractive Index

The above approach requires knowledge of the object's refractive index. If this is unknown, one needs additional input to estimate it, for example an additional acquisition, for a different position of the monitor. Then, for each pixel,

we get two constraints of the form (4). The two constraints correspond to two different angles θ_1 , for the two different positions of the monitor. Knowing one of them, one can compute the other. We thus have a total of three unknowns, the refractive index (or, the ratio n_{12} of refractive indices of object and air) and the two angles θ_1 , and three constraints: the two constraints of form (4) and one that links the two angles θ_1 (ref Figure 2 [1]).

Two such acquisitions were used in the geometric approach of [17] for example. However, even when light is only *reflected* off a transparent object surface, equations (4) can be used to solve for relative refractive index $n_{12} = \frac{n_1}{n_2}$, even from a single pixel. Thus, we have *extended* the scenarios where reconstruction is possible to the case of *reflection* off the object surface, when refractive index is unknown. Accurate estimation may be possible due to the refractive index being “shared” by all pixels.

5.2. Double Bounce

Like in pure geometric approaches such as [17], we also study the case of double bounces, e.g. if the camera acquires an image of the monitor through a transparent object, with refractions happening both at its front and back sides. Unlike the “single bounce” case in section 4, the analysis is considerably harder, because of the involvement of polarization. We derive a simple case in [1] to illustrate this fact. While practical considerations make it hard to produce decent results, we nevertheless argue about the theoretical viability of a solution in this scenario.

Let us reconsider equations (2) and (3): they indicate how much energy of the incoming light is reflected respectively refracted. In addition however, for dielectric materials, the reflected and refracted light is polarized, even if the incoming light is not. It can be shown that polarization adds a factor to the radiance ratio that is dependent on the *angle* between the two planes of reflection/refraction encountered along a single light-path. If this is properly taken into account, one may reconstruct the surface even after double bounces and the following conditions arise.

Consider two cameras looking at a transparent object, which refracts light from a known illuminant twice. Further assume that for each camera, two 3D measurements were made per light ray. This corresponds to the case $\langle 2, 2, 2 \rangle$ as per the convention of [17]. In this case, radiance ratio observations are made by each camera for each light-path. It is possible to show that for every light-path, there exists a 1D curve of incident angle pairs such that one angle occurs at each “bounce” of a light-path and the radiance ratio is satisfied [1]. Two cameras give us two such curves. In addition to the geometric constraints expressed in [17] we have one extra constraint per light ray, so it is now possible to solve for the 3D structure of the transparent object given known refractive index.

One reference point ($M = 1$)			
	K = 1	K = 2	K \geq 3
N = 1	✓		
N = 2	✓ ×*		

Two or more reference points ($M \geq 2$)			
	K = 1	K = 2	K \geq 3
N = 1	✓ ×*		
N = 2	✓ ×*	✓ ⁺	
N = 3	✓ ×*	✓ ×	
N \geq 4	✓ ×*	✓ ×*	

Table 1: Tractable triangulation problems. Updated from [17]. M is the number of known 3D points, N is the number of cameras (views) and K is the number of bounces. An * symbol represents the fact that even in the case of only reflection, the *relative refractive index* can be computed. ⁺ indicates theoretical possibility.

5.3. Summary of Minimal Scenarios

Minimal scenarios for the reconstruction of transparent objects that can be solved with the help of radiance measurements, are summarized in Table 1. It is interesting to note that transparent objects have *lesser* minimal requirements for reconstruction than mirror like objects.

In the next section, we present some practical aspects of data acquisition.

6. Practical considerations

The above theory shows that the radiance of a final light ray in a light-path contains information that could be used to reconstruct the entire light-path. In this section we describe important elements of our experiments to collect radiance measurements for reconstruction. Our experiments consist of three parts. 1) We use an illuminant with known geometry to emit unpolarized light in a desired set of directions. 2) Light from the illuminant interacts with the transparent object, and reflects / refracts off its surface towards the camera after one bounce. 3) The camera then captures both the direction and radiance of some of reflected / refracted light rays, which is used for reconstruction.

We use Canon and Nikon DSLRs for our experiments. Since we need to capture the position and radiance of an individual light ray, we adopt the pin-hole model for the camera (smallest aperture and large focal length). This allows each ray to be modeled as a thin cylinder. We arrived on an acceptable range of focal lengths that have desired depth of field and capture both the monitor and object in focus, by trial and error. Finally, for each captured ray, we compute the corresponding pixel on the monitor from which the ray originated using standard methods [3].

Unpolarized illuminant In our experiments, we use a flat CRT monitor (LCD monitors emit polarized light), whose pose is computed with respect to an internally calibrated camera [18]. This is motivated by our need to measure by the illumination and the geometry of our illuminant. We capture the light emitted by each pixel of the monitor in




Figure 4: (Left top) Normal map of “Fanta bottle” sequence. (insets) Note the fine details captured as a result of radiance ratios. (Centre top) Depth map. Blacker colors are closer to camera. (Right top) Two views of the 3D reconstruction, with lighting to highlight shape variations. Phenomenon like scratches on the bottle, inhomogenous refractive index, violation of single bounce through occlusion are some bad effects, but still reliable reconstructions are achieved. Note that since camera is placed far from the object and monitor, large changes in depth cause small changes in angle. This explains some of the “rough”-ness of the reconstruction. Note also that no smoothing or optimization is applied for this reconstruction. (Left bottom) Normal map of “Wine glass” sequence. (Left-centre bottom) Depth map. (Right-centre bottom) Depth map produced if no interreflection removal is performed. Notice the lack of depth variation in one of the glasses (compared using insets with blue borders). Some frequency artifacts can be seen (red inset) due to interreflection removal. (Right bottom) Depth difference between the two cases. Best viewed in color.

several directions in 3D, and fit a smooth model to this data in order to accurately measure the variation in illumination associated with a monitor pixel.

Interreflections A common problem with measuring illumination reflected / refracted off specular transparent objects is interreflections. They not only corrupt the radiance measurement, but also pose a problem to correspondence estimation. In order to remove the effect of interreflections, we use the algorithm of Nayar *et al* [15]. Instead of using a projector to light the scene, we use the CRT monitor instead. We project low frequency checkerboard patterns that are shifted cyclically, and use it to compute the direct and global components of the scene. Figure 5 shows a result for one typical scenario.

Calibration We internally calibrate the camera in order

to compute the direction of the captured rays. We also externally calibrate the monitor w.r.t the camera. We use the monitor as the reference frame of the coordinate system, which is beneficial in the case of multiple cameras. For radiance measurement, we first extract an unprocessed image from the RAW files using `dcraw`¹. We then remove the global component of this image, and then extract direct measurements. Note that only the direct and global components shown in Figure 5 actually come from RAW files.

Correspondence Acquiring correspondence between pixels on the monitor and pixels on the camera that correspond to the same light-path becomes slightly cumbersome when transparent objects are involved [2]. In order to over-

¹Command : `dcraw -r 1 1 1 1 -D 0 -H 1 -q 3 -4 -v` for Canon cameras. We drop the `-D` option for Nikon cameras.

come this, we first remove global components from images used for correspondence by applying the theory of [15] to these binary images, and then use a graph based approach to enforce spatial smoothness in the image while extracting correspondences.

7. Experiments and Results

In the previous sections, we showed that radiance ratios could be used to reconstruct transparent surfaces, which can help in reducing the number of measurements required for reconstruction. We also listed some practical aspects that are relevant to acquiring image data. In this section, we show results of four experiments. The first experiment demonstrates the robustness that radiometry brings to our reconstruction approach. The second experiment demonstrates the accuracy of 3D reconstruction using our method. The third experiment shows our result on an extremely complex scenario, and the fourth experiment shows a very detailed reconstruction.

Experiment 1: Synthetic dataset Figure 3 shows various results for single bounce reconstruction of a sample sinusoidal object when the 3D correspondence is noisy. We compare with two light-path triangulation approaches, one in which the 3D correspondences along a light-path are close to each other and the object (LP-1), and one in which they are far (LP-2). In both cases, noise is added to the farthest 3D correspondence. While one case (LP-1) is sensitive to noise, the second case (LP-2) is robust but impractical. Our approach however, gives a reliable normal map even if the depth is slightly perturbed (compared to LP-2). Note that noise percentage is calculated as ratio of distance between noisy and ground truth data and distance of ground truth and object. Other results w.r.t camera noise and refractive index mismatch are present in the supplementary materials.

Experiment 2: “Water Sequence” Figure 5 (Left column) shows some images acquired in order to reconstruct the surface of water in a plastic bowl. The bowl is around 10 cms in diameter, and is placed about 2.5 meters from the camera. This is a scene with a very simple 3D structure (a plane) and given the smooth surface of the bowl, it also has minimal (but not negligible) interreflections and caustics. Because of the planar nature of the scene, we compute correspondence by simply computing a homography between the reflected image and the direct image of a photograph displayed on the monitor. This homography, adjusted for the internal calibration of the camera and the aspect ratio of the pixels on the monitor, can now be used to compute the normal of the plane [18].

After computing the direct image, we use the homography to compute the ratio of directly observed and reflected radiances. Finally, we hypothesize and test individual pixels for various values of depth, and record the value that

best fits the radiance ratio. Figure 5 shows our result (the ripples in the reconstruction are just quantized depth values). An alternate result was one obtained by using the algorithm of [18]. Comparison of the results gives us a mean squared error of around 0.1 cms (we omit correspondence errors while computing this measure), which shows the accuracy of our reconstruction.

Experiment 3: “Wine Glass Sequence” This sequence is very challenging for approaches that use projected texture for reconstruction because of the large interreflections and caustics present in the scene. Note how global components of the image are present even in places where there is no direct light (Figure 5, red square). For approaches that typically only use geometric methods [17, 13], the subsurface scattering might throw off correspondence measurements. This can heavily influence reconstruction accuracy. On the other hand, robust measurement of the position and direction of light incident on the glasses from the monitor requires a large set of images to be captured while moving the monitor over, say, an optical bench. In our approach, however, both these errors are avoided because we use a *single* measurement per pixel for reconstruction, and use on direct components of the images we capture. Again, like in the case of the “Water Sequence” we hypothesize various depth values along each back-projected pixel, and test its validity using computed radiance ratios. Figure 4 shows depth and normal maps computed using our approach. Notice again how we obtain smooth maps even though no smoothness constraints are imposed on the results.

Experiment 4: “Fanta Bottle Sequence” This sequence highlights the ability of our approach to capture details of a surface. The scene captured consists of a Fanta bottle filled with water. The absence of any large interreflections in this scene results in a very detailed reconstruction as shown in Figure 4. Note that optimizing depth and normal simultaneously would serve to remove the artefacts seen in the figure, especially enforcing the depth-normal consistency (differentiation of depth gives normal).

8. Discussion and Conclusion

Reconstruction of transparent objects remains a challenging problem because of the lack of cues that are normally available for other objects. In this paper, we showed the existence of an approach that combines two of the more widely available cues, namely geometric and radiometric cues. We showed how this leads to simplified acquisition, decently robust algorithm, new minimal solutions, and presented challenging scenarios where our approach yielded accurate reconstructions. Other applications of our approach are in verifying validity of light transport matrices, radiometric calibration etc. We hope to investigate other potential applications that might be useful to the community.




Figure 5: (Left column) Two of 25 images used to compute the direct and global images [15] to remove the effect of interreflections and caustics on radiance measurement. (Left-centre column) Difference between the “direct” component and an image taken (top image) with a white pattern shown on the monitor. Bottom image is the difference. (Right-centre column) Direct (top) and global (bottom) components (RAW images) of Water sequence. Interreflections are simple. (Right column) Correspondence (Y-coordinates) using Gray-codes (top), depthmap (middle) and reconstruction with shading (bottom).

9. Acknowledgements

The first author would like to thank Dikpal Reddy, Sri Kumar Ramalingam, Amit Agrawal, Yuichi Taguchi, Avinash Sharma, Amaël Delaunoy, Antoine Letouzey, Simone Gasparini and Avinash Kumar for several discussions and comments.

References

- [1] Photo-Light-Path Triangulation: Supplementary Material. <http://www.di.ens.fr/~chari/webpage/Research/phototechreport>. 3, 4, 5
- [2] D.G. Aliaga and Y. Xu. An adaptive correspondence algorithm for modeling scenes with strong interreflections. *IEEE Transactions on Visualization and Computer Graphics*, 15:465–480, 2009. 4, 6
- [3] D.G. Aliaga and Y. Xu. A self-calibrating method for photogeometric acquisition of 3D objects. *IEEE-PAMI*, 32:747–754, 2009. 5
- [4] M. Ben-Ezra and S. Nayar. What Does Motion Reveal about Transparency? *ICCV*, 2:1025–1032, 2003. 2
- [5] T. Chen, M. Goesele, and H.-P. Seidel. Mesostructure from specularity. *CVPR*, 2:1825–1832, 2006. 1
- [6] A. Efros, V. Isler, J. Shi, and M. Visontai. Seeing through water. *NIPS*, 2004. 2
- [7] E. Hecht. Optics - 4th edition. *Addison Wesley*, 2002.
- [8] M.B. Hullin, M. Fuchs, I. Ihrke, H.-P. Seidel, and H.P.A. Lensch. Fluorescent immersion range scanning. *SIGGRAPH*, 2008. 1, 2
- [9] C.P. Huynh, A. Robles-Kelly, and E.R. Hancock. Shape and refractive index recovery from single-view polarisation images. *CVPR*, 1229–1236, 2010. 2
- [10] I. Ihrke, K. N. Kutulakos, Hendrik P. A. Lensch, Marcus A. Magnor, Wolfgang Heidrich. Transparent and Specular Object Reconstruction. *Computer Graphics Forum*, 2400–2426, 2010. 1
- [11] D. Miyazaki and K. Ikeuchi. Shape estimation of transparent objects by using inverse polarization ray tracing. *IEEE-PAMI*, 29(11):2018–2030, 2007. 1, 2
- [12] D. Miyazaki, M. Kagesawa, and K. Ikeuchi. Transparent surface modeling by using a pair of polarization images. *IEEE-PAMI*, 26(1):73–82, 2004. 1
- [13] N. Morris and K. Kutulakos. Reconstructing the surface of inhomogeneous transparent scenes by scatter-trace photography. *ICCV*, 2007. 1, 2, 7
- [14] N. Morris and K. Kutulakos. Dynamic refraction stereo. *IEEE-PAMI*, 33(8):1518–1531, 2011. 1
- [15] S.K. Nayar, G. Krishnan, M.D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *SIGGRAPH*, 2006. 6, 7, 8
- [16] M. Oren and S.K. Nayar. A theory of specular surface geometry. *IJCV*, 24(2):105–124, 1997.
- [17] E. Steger and K. Kutulakos. A theory of refractive and specular 3D shape by light-path triangulation. *IJCV*, 76(1):13–29, 2008. 1, 2, 3, 4, 5, 7
- [18] P. Sturm. Algorithms for plane based pose estimation. *CVPR*, 706–711, 2000. 5, 7
- [19] Y. Tian and S. Narasimhan. Seeing through water: Image restoration using model-based tracking. *ICCV*, 2303–2310, 2009. 2
- [20] S.-K. Yeung, T.-P. Wu, C.-K. Tang, T.F. Chan, and S. Osher. Adequate reconstruction of transparent objects on a shoestring budget. *CVPR*, 2011. 1, 2
- [21] G. Wetzstein, G. Roodnick, W. Heidrich and R. Raskar. Refractive Shape from Light Field Distortion. *ICCV*, 2011. 2