

Tomographic Reconstruction of Transparent Objects

Borislav Trifonov, Derek Bradley, Wolfgang Heidrich[†]

The University of British Columbia

Abstract

The scanning of 3D geometry has become a popular way of capturing the shape of real-world objects. Transparent objects, however, pose problems for traditional scanning methods. We present a visible light tomographic reconstruction method for recovering the shape of transparent objects, such as glass. Our setup is relatively simple to implement, and accounts for refraction, which can be a significant problem in visible light tomography.

Keywords: Hardware – Modeling, Acquisition and Scanning; Methods and Applications – Optics; Modeling – Object Scanning/Acquisition.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation; Digitizing and scanning; I.4.1 [IMAGE PROCESSING AND COMPUTER VISION]: Digitization and Image Capture; Imaging geometry.

1. Introduction

3D scanning has become a valuable tool for digitizing real-world artifacts such as sculptures and other works of art. The most common and affordable scanning methods work with visible light. Both passive lighting methods, such as stereo, and active lighting methods, such as laser scanning, work only for mostly opaque, primarily diffuse objects. Even a relatively small amount of translucency can reduce the precision of the scan [LPC*00], as can interreflections arising from somewhat specular surfaces. These properties prevent the use of current visible light methods for scanning the shape of transparent objects, such as glass or plastics. However, transparent objects are ubiquitous in modern environments, so it is desirable to have a mechanism for scanning their shape.

One possible solution is to spray-paint or otherwise cover the object with an opaque layer, so that traditional visible light scanning methods can be used. However, this approach is problematic, since spray-painting is considered too destructive for many objects, such as artwork. In addition, this approach introduces the typical artifacts that arise in scanning

opaque objects, such as holes due to occlusion near complex geometric details [LPC*00].

Tomographic methods avoid the occlusion problem and the need to paint the object. They reconstruct the 3D geometry of semi-transparent objects from a number of shadow images corresponding to different positions of a source of electromagnetic radiation. Tomography can be used if the medium is semi-transparent to the wavelength of electromagnetic radiation used for acquisition, and if refraction is negligible. This latter requirement typically mandates the use of x-ray radiation, as is the case in medical or engineering computed tomography (CT). Unfortunately, x-ray tomography relies on expensive and bulky equipment, and cannot be used in many environments due to safety considerations.

Recently, there has been work in the computer graphics community on visible light tomography to recover gaseous objects, such as flames [IM04, IM05]. In this case, refraction is negligible.

In this paper, we present a tomographic method for recovering the shape of objects made of glass or other transparent media of various refractive indices. We suspend the object in a fluid within a glass cylinder that has been centered on a turntable. The refractive index of the fluid is matched to be similar to that of the object, thereby minimizing refraction

[†] email: {trifonov,bradleyd,heidrich}@cs.ubc.ca



Figure 1: Left: glass object. Center: one of 360 input images. Right: reconstructed geometry, including some left-over support structure that the object was placed on during acquisition. The reconstruction is the raw output of our method, without further processing.

tion. Refractions still occur at the interfaces between air, the fluid container, and the fluid itself, but we devise a calibration mechanism to account for these.

The primary contributions of this paper are the setup for recovering the 3D shape of transparent objects through tomography with visible light, and a calibration method for determining light paths within the fluid and object. Our tomographic reconstruction method is a variant of algebraic reconstruction techniques (ART). We first review related work, before we present the details of the setup and acquisition in Section 3, and of the reconstruction in Section 4. Section 5 details the results, and we conclude in Section 6.

2. Related Work

Capturing the appearance of transparent objects has so far predominantly been achieved by environment matting techniques (e.g. [ZWCS99, WFZ02, PD03, MPZ*02]). However, these are image-based methods, and we desire a full 3D reconstruction of the geometry.

The body of literature on 3D geometry acquisition is vast, and originates from research communities as diverse as computer graphics, computer vision, medical imaging, and physics. In the following, we review some of the work most directly related to ours, either in terms of methodology, or in terms of intended application.

3D Scanning with Visible Light. The most commonly used 3D scanning techniques today are based on the interaction of visible light with the object to be scanned. These techniques can be roughly grouped into two categories. *Active* lighting methods such as laser stripe scanning [Cyb] or encoded light patterns [BER76, VO90, RHHL01, HHR01] require fine control over the illumination. In contrast, *passive* methods such as shape from stereo [SS02] or shape from shading [ZTCS99, Woo80] do not. Active methods typically produce higher quality results, but require a laboratory setting and are more invasive.

All of these methods assume opaque objects, and work best with mostly diffuse surfaces. Even relatively minor translucency can reduce the precision of the triangulation used for finding surface points. Some recent work in computer vision specifically focuses on the reconstruction of transparent objects [BEN03, KS05, MI05, BV99]. However, while these methods are of great theoretical value, they cannot at this point produce reconstructions of the quality that we achieve. In particular, reconstruction with these methods has not been demonstrated for shapes that differ from their visual hull, and the resulting triangle meshes have many holes.

Recent work by Ihrke et al [IGM05] focuses on reconstruction of transparent fluids. Their method cannot be applied to solids, since the object to be scanned needs to be colored with fluorescent dye.

X-ray Tomography and Tomographic Reconstruction. X-ray computed tomography (see Kak and Slaney [KS01] for an overview) is a common technique in medical imaging and engineering, for example in quality control of mechanical parts. Tomography reconstructs 3D volume densities from a sequence of images representing the shadow of the object as illuminated by an x-ray source from various directions. Tomography works on the assumption that the object is semi-transparent to the radiation, or in other words, the radiation along any ray is never *completely* absorbed by the material. A second assumption is that there is no refraction or scattering, and the radiation passes straight through the material. It is difficult to meet these assumptions with visible light, and this is the primary reason for accepting the high cost and overhead of using x-ray radiation in medical and engineering tomography. One of the major contributions of this paper is a physical setup that fulfills these assumptions for transparent objects and visible light.

Two principal approaches exist for reconstructing volume densities from the line integrals captured in the shadow images. A well known result is the Fourier Slice Theorem [Bra56, KS01], which states that the volume densities

can be recovered from a number of 1D inverse Fourier transforms of lines in the shadow images. The Fourier Slice Theorem allows for very efficient reconstruction of the volume densities. However, it is limited to the case of parallel illumination or, with modifications, point sources.

Algebraic Reconstruction Techniques (ART) do not require parallel or point source illumination [GBH70, MYW99, KS01]. These methods iteratively solve for the volume densities by comparing the absorption of a current estimate to the measured images, and updating the volume accordingly. In our work, we use a variant of ART, which we describe in detail in Section 4.

Optical Tomography. The tomographic assumption of no refraction has prevented the use of visible light in most scenarios. However, visible light tomography has been used successfully to observe gaseous phenomena such as plasmas [IPD94], or, more recently, flames [IM04, IM05]. For these effects, refraction is small enough that it can be neglected for ray lengths corresponding to the size of the object.

In microscopy, the problem of tomography in the presence of refraction has been resolved by placing samples in a gel of similar refractive index [SAP*02]. However, it is not obvious how to extend that setup to a macroscopic domain. In this paper we show how refractive index compensation can be done at a macroscopic scale, and how to calibrate such a system.

In recent years, there has also been a strong interest in optical tomography in the presence of strong scattering, for example in biological tissue [Arr99]. However, at present these methods remain numerically unstable, and very inefficient. Moreover, they are unnecessarily complex for transparent objects, where light only refracts at the interface of two materials, but does not scatter.

Visual Hull and Voxel Coloring. Tomography methods share a number of characteristics with visual hull methods [Lau94, Pot87, MBR*00] in terms of setup and processing. However, visual hull methods are inferior to tomography methods, as they can only reproduce the visual hull, while tomography can, up to a certain numerical precision, reproduce the true shape. Voxel coloring [SD97] produces tighter bounding surfaces than the visual hull. However, because it only works on opaque objects, occlusion would prevent a high quality reproduction of the hole in Figure 1.

3. Setup and Acquisition

Our acquisition system consists of a video camera and a glass cylinder centered on a turntable. The glass cylinder acts as a container for a fluid. To scan a transparent object, we place it inside this fluid, and roughly center it within the cylinder. A diffuse surface acts as a background against

which we measure absorption in the fluid and the target (Section 3.1). In order to reduce exposure times, and hence speed up acquisition, the surface should be as brightly illuminated as possible. A very precise uniformity of the illumination is, however, not required.

The refractive index of the fluid is adjusted to be similar to that of the target (Section 3.2), and refraction within the cylinder is negligible as a result. This property is the key component of our setup that allows us to use tomographic acquisition. However, light still refracts while entering and leaving the cylinder. We calibrate for this effect by measuring the distribution of rays corresponding to camera pixels inside the cylinder (Section 3.3).

With this setup, colored transparent objects can be scanned by taking images under different orientations of the turntable (Section 3.4). Depending on the desired resolution of the reconstruction (described in Section 4), the number of images taken varies between several dozen and a few hundred. Clear transparent objects can also be scanned, but this requires adding a contrast agent to the fluid (Section 3.5).

3.1. Geometry of the Setup

Figure 2 shows the geometry of the setup both as a drawing and as a photograph. We use a precision glass cylinder with a 15 cm inner diameter as a container. It is centered on a turntable using a precision stand made from ABS plastic, produced with a rapid prototyping machine (Stratasys Vantage-i). The camera we use is a Prosilica EC 1350C firewire machine vision camera with 1.5 MPixels, and a 25 mm c-mount lens. The camera can be programmed to capture 12 bit raw, linearly quantized, Bayer pattern images. As a backdrop, we use a uniformly illuminated, diffuse white surface. All components are mounted on an optical table.

3.2. Refractive Index Matching

Our method requires the target object to be contained in a fluid of roughly the refractive index of the target. Minor differences in the refractive index do not pose a problem, as long as they do not significantly change the path length of the rays within the fluid and the refractive object. We used simulation data sets to determine that a discrepancy of up to 5-7% between the fluid and the target is still acceptable. For larger discrepancies the absorption observed from different view points is inconsistent with the assumed ray geometry, so that the method does not converge. This is easy to detect numerically in the form of large changes in some voxels in every iteration.

There are several ways to produce fluids with the appropriate refractive index, for example through mixing different fluids or preparing solutions of water with various salts. In our




Figure 2: Geometry of the setup.

experiments, we use a solution of potassium thiocyanate[†] (potassium salt) in water, which produces a refractive index of up to 1.5 with a solution of around 80% [Bud94], and up to about 1.55 with a super-saturated solution. Potassium thiocyanate does not react with glass or plastic, and is easily washed off without leaving any residue or causing any damage to the object. To match the refractive index of the solution to the target, we manually change the concentration of the solution while the target is suspended in it, and visually observe the refraction against a patterned background. As such, knowledge of the exact numerical value of the refractive index is not required.

Both the solution and the object can exhibit dispersion, in which case the refractive index can only be matched for a small range of wavelengths. We address this situation during the reconstruction (Section 4) by recording a raw Bayer pattern image with the Prosilica camera, and using only the green pixels. The other primaries could also be used, but a Bayer pattern contains twice as many green pixels as red or blue ones, and thus the green channel of an image provides more data. During the index matching process, we found it useful to observe the refraction through a green filter (Lee filter #58).

[†] Potassium thiocyanate is a skin irritant, and toxic if ingested or inhaled, so careful handling is required.

The potassium thiocyanate solution can produce refractive indices similar to those of everyday glass or plastic objects. High quality lead or crystal glass, however, has a higher optical density. For these objects, different solutions or mixtures can be used [Bud94, MM98].

3.3. Geometric Calibration

In order to run a tomographic reconstruction on the image data we acquire, we need to know the light path corresponding to each pixel. Since the refractive index matching eliminates refractions at the interface of the target and the fluid, a light path can be described by a single ray segment within the cylinder.

We recover this ray distribution inside the cylinder with a two-plane parameterization step [LH96, GGSC96]. For every ray corresponding to a camera pixel, we measure its intersection point with two parallel planes within the cylinder. To this end, we used the rapid prototyping machine to build a precision plastic frame that fits exactly inside the cylinder. Into this frame we can slide two parallel, vertical panels with calibration patterns (see Figure 3).




Figure 3: Images for geometric calibration. Left: front plane. Right: back plane

We create a calibration pattern using the ARTag marker system [Fia05], which consists of black squares with a binary encoded pattern on the inside. This system proved to be very robust under the kind of image distortions caused by refraction in the cylinder. The software that is provided for the ARTag system [ART] recovers the position of the square corners, together with an identification of each corner.

By interpolating this corner information, we can determine the 2D coordinates on the panel that correspond to each camera pixel. The information from the two panels can be combined with knowledge about the geometry and position of the panels to determine the 3D ray segment that corresponds to each pixel (see Figure 4).

Due to the geometry of the panels and light rays, the ray distribution is recovered only for a subset of the complete




Figure 4: Top: a visualization of the ray distribution recovered from the calibration step. Bottom: the reconstruction region is a subset of the full cylinder.

cylinder; the outer regions along the side are not covered (Figure 4). As a result, the reconstruction region, that is, the region in which it is possible to recover volume densities, is restricted to a smaller cylindrical region inside the glass container, marked blue in Figure 4. However, it is actually not desirable to use the outer regions of the cylinder, since refraction there is too sensitive to minor variations in the geometry, which can result in distortions.

3.4. Acquisition of Colored Objects

Once the ray distribution in the cylinder has been calibrated, we can acquire colored, transparent objects by placing them inside the cylinder and taking images from different sides, using the turntable to rotate the cylinder and the target inside. One of the images from such a sequence is shown in the center of Figure 1. The color of the object provides the contrast between the fluid and the target, which can be recovered as different volume densities (Section 4). We also take an image of the cylinder without the object, so that we can derive per-pixel illumination values. This allows us to compute the ray absorption for each pixel in each image of the object.

Due to the different densities of the rays on the near and the far side of the cylinder (see Figure 4), we usually perform a full 360° acquisition, rather than only 180° . Depending on

the desired resolution of the reconstruction, we take images every $1^\circ - 5^\circ$, resulting in 72-360 images. In each image, up to approximately 350,000 green pixels of the Bayer pattern contribute to the reconstruction of larger objects.

Since we use a high-quality glass cylinder, and since it is precisely centered on the turntable with a custom mount, the ray distribution we have recovered in the previous step is valid for any rotation of the turntable.

3.5. Acquisition of Clear Objects

Because the potassium thiocyanate solution is a clear fluid, a separate water-soluble dye is required as a contrast agent for scanning an object made of clear glass or similar materials. We use food coloring (Club House red or blue) for this purpose, which we chose because it colors the fluid without introducing noticeable scattering. Adding the contrast agent reverses the situation from Section 3.4: now the fluid absorbs significantly, while the target does not.

This new situation introduces the need for an additional photometric calibration step. Since our reconstruction region is smaller than the full cylinder, as shown in Figure 4, absorption occurs in the ring outside the reconstruction region (see grey ring in Figures 4 and 5). We call this ring the *environment*.




Figure 5: Acquisition geometry using a contrast agent.

We need to compensate for the effects of this absorption. To this end, we take a photo of the colored fluid without the target object. This way, we determine the accumulated absorption

$$A_{cyl} = e^{-\int_a^b \alpha(t) dt} = e^{-(d-a)\alpha}$$

along the ray for all pixels in the image. See Figure 5 for an explanation of the quantities.

From this equation we can determine the differential absorption α , since the length $d - a$ of the ray segment is known from the cylinder geometry and the geometric calibration step. The value of α ought to be the same for all rays due to the homogeneity of the medium. However, the measurement and quantization errors introduce small variations, so we use the average value over all rays.

Once α is known, the absorption along a ray due to the environment is described by the relationship

$$A_{env} = e^{-(b-a)\alpha} \cdot e^{-(d-c)\alpha} = e^{-(b-a+d-c)\alpha},$$

in which the distances $b - a$ and $d - c$ are again known from the geometric calibration and the choice of reconstruction region (Figure 5).

The rest of the image acquisition can now proceed as in Section 3.4, but in the resulting images, every pixel has to be scaled by $1/A_{env}$ for that pixel, in order to compensate for the absorption of the environment.

The amount of color added determines the contrast one can achieve, and hence the resolution of features that can be reconstructed. More color increases the signal-to-noise ratio, but it also darkens the measurements due to larger overall absorption. As a result, one either has to use a brighter light or a longer exposure. For a given light intensity and a given longest exposure time, one should use as much color as possible without creating black pixels in the longest exposure image.

4. Tomographic Reconstruction

As in x-ray tomography, the pixel values in our images are line integrals of the absorption along the corresponding ray. This is called an *integral projection* [KS01]. Different material properties along a ray manifest themselves as different absorptions along ray segments and thus lead to different intensities in the integral projection images.

As mentioned in Section 2, there are several ways of reconstructing volume densities. The Fourier Slice Theorem [KS01, Bra56] states that the inverse Fourier transform of a scanline in an orthographic integral projection image corresponds to a 1D slice of densities in the volume. This observation allows for very efficient reconstruction of the volume densities, but it is limited to the case of parallel illumination or, with modifications, point sources.

In our application, the rays in a single integral projection image do not usually form an orthographic or projective transformation. We therefore have to use more general, albeit slower, methods such as Algebraic Reconstruction Techniques (ART) [GBH70].

4.1. Algebraic Reconstruction

Our reconstruction method is modeled after the Simultaneous Algebraic Reconstruction Technique (SART) of Mueller et al. [MYW99], and works on a voxel grid of densities. To simplify computations, all operations are performed in log space, where the absorption along a ray is described as

$$\begin{aligned} \log A &= - \int_b^c \alpha(t) dt \\ &\approx - \sum_i \alpha_i. \end{aligned}$$

In this equation, α is the density, or local opacity, of a point in the volume as before, b and c are the entry and exit points of the ray in the reconstruction region from Figure 5, and the α_i correspond to the densities at uniformly spaced points t_i along the ray. More specifically, the t_i are intersection points of the ray with a stack of axis-aligned slices through the volume.

Starting with an empty (i.e. zero-valued) voxel grid, the algorithm proceeds iteratively by randomly selecting one of the source images. For the selected image, the current volume of densities is forward projected in a step very similar to conventional volume rendering. The resulting log ray absorptions for all pixels are compared to the target values from the captured image. The difference for every pixel is then back-projected through the volume, and all the values of all voxels along each ray are updated, before the next image is selected. At the end of this process, the volume grid contains the recovered volume densities, from which the shape of the object can be extracted as an iso-surface.

Forward Projection. The forward projection is a volume rendering step, used to compute the ray absorptions assuming the current estimate of the volume densities for a specific camera position. The implementation of this step is somewhat similar to the shear warp algorithm [LL94], but modified to take into account the non-uniform distribution of the rays and using a higher quality reconstruction filter.

The volume is sliced perpendicular to the major coordinate direction that is most parallel to the viewing direction, and the slices are processed back-to-front. For each slice and each ray, we compute the distance of the ray to all voxels j in the proximity of the ray, and use it to look up a weight w_j in a 1D lookup table. The table contains precomputed line integrals of rays passing through a 3D, rotationally-symmetric Kaiser-Bessel filter kernel at various distances. Given the weight w_j , the absorption caused by the ray passing through this slice is

$$\frac{\sum_j w_j \alpha_j^{(k)}}{\sum_j w_j},$$

where $\alpha_j^{(k)}$ is the estimate for the density of voxel j before considering the current image. The division can also be deferred until all slices have been processed, as pointed out by Mueller et al [MYW99].

The Kaiser-Bessel function is a high-quality reconstruction kernel frequently used for tomographic purposes [Lew90]. We use a support radius of two voxel spacings. The pre-integration approach allows us to efficiently use this high-

quality filter, and accounts for the different angles at which the rays intersect the slice.

We process all rays in this fashion before moving on to the next slice. After processing all slices we have obtained the value of $\log A^{(k)}$, the ray absorption produced by the current estimate of the volume densities.

Back Projection and Volume Update. In the second step we update the volume by uniformly distributing the error $\Delta A = \log A - \log A^{(k)}$ along each ray. We again do this by processing the volume slice-by-slice. For each ray and each slice, we re-use the contribution w_j of each voxel to the ray, as computed during the forward projection. We then update the density of each voxel as follows:

$$\alpha_j^{(k+1)} = \alpha_j^{(k)} + \lambda \frac{\sum_j w_j \Delta A}{\sum_j w_j}.$$

The parameter λ is a relaxation factor. If it is too small, then convergence will be slow. If it is too large, then the method becomes unstable and the last image processed will have a disproportional impact on the output geometry. In practice, we achieve good results with values between 0.01 and 0.08.

Optimizations. We have implemented two optimizations to this basic algorithm. First, since the volume is processed in slice order, and the slices can either be parallel to the x-y or the y-z plane depending on the view, we can accelerate the processing by sorting the volume according to the traversal order. For every iteration, we first optimize the memory layout of the volume data such that it is best for traversal along the z-axis, and process all the views that require this traversal direction in random order. Then, we rearrange the memory layout to optimize traversals along the x direction, and process the remaining half of the images in random order.

A second optimization is to parallelize the processing of slices for dual processor or dual core machines. Since the impact of each slice is additive on a logarithmic scale, the various slices can be processed independently of each other in both the forward and the back projection.

Dealing with Inaccurate Refractive Index Matching. A possible source of reconstruction error is inaccurate matching of optical densities between the fluid and the object. The result of such a mismatch is a change in the path the light actually takes, which in turn changes the length of the path and the amount of absorption along it.

One observation is that this effect is more dramatic if the ray intersects the object at an acute angle, than if it intersects at a more oblique angle. We use this observation to locally adapt the parameter λ for each ray. Specifically, after each iteration through the set of images, we compute the gradient field for the current density estimates. During forward projection in the subsequent iteration, we determine the maximum angle between the ray direction and the local gradient along a ray. During back projection, we multiply λ with the cosine of

this angle. As a consequence, rays at acute angles have a significantly lower impact than rays at oblique angles.

5. Results

Figure 1 shows the reconstruction of a colored glass object from 360 images, using a voxel grid of size $475 \times 276 \times 475$, corresponding to a 0.12 mm voxel resolution. This object is particularly challenging for conventional scanning methods, since the small hole in the center is practically impossible to capture with any method but tomography. Figure 6 shows a cut-away view of the reconstruction that we achieved with 5 iterations through the set of images.




Figure 6: Cut-away view of the reconstructed hole on the object from Figure 1.

The iso-surface extraction relies on finding a good iso-value to separate the object from the fluid. Since we are dealing with two materials of more or less constant density, finding such an iso-value is simple. Figure 7 shows a histogram of the densities for the object in Figure 1. The left peak corresponds to the fluid, while the right peak corresponds to the object. The clearly defined valley between these peaks represents the location of the iso-value.

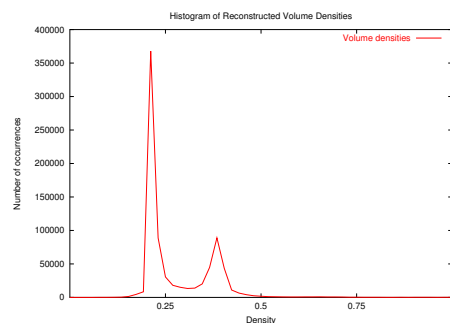


Figure 7: Histogram of volume densities.

The acquisition time for this dataset was about 45 minutes, and the reconstruction time about 90 minutes. The object does contain high-frequency noise as can be witnessed in Figure 1. Some of the base that the object rested on during acquisition is contained in the model. This is a common problem with 3D scans, and is easily handled in a manual cleanup phase.

Figure 12 shows pieces of a glass chess set that we acquired with our method. The pieces are made of clear glass, so we use food coloring to provide contrast, as described in Section 3.5. The concentration of the food coloring can be chosen very easily by visual inspection. The more coloring is used, the better intricate details can be reconstructed, since the contrast is higher. However, adding the food coloring reduces the overall brightness of the acquisition images, such that longer exposure times up to 8 seconds are required for this dataset. This effect increases the acquisition time to about 100 minutes for 360 images. This could be countered by using a higher-intensity light source for illuminating the diffuse background. Figure 8 shows an image from an acquisition sequence of the queen and the bishop of the chess set.




Figure 8: One of 360 images from the acquisition sequence for the queen and the bishop.

To evaluate the precision of the reconstruction, one would ideally have ground truth geometry for comparison. We attempted to obtain such geometry by spray-painting the chess pieces, and scanning them with a Cyberware scanner. Figure 9 shows some of the results. These images are already the result of merging several scans, and hours of post-processing. It is obvious that the scans lack detail compared to our results, and that they suffer from holes, wrong topology due to interreflections, and other artifacts.

Since the laser scans are not fit for ground-truth comparisons, we used digital calipers to obtain quantitative results on precision. We measured the distance of various feature points on the real object, and compared those values to the distances in the scan. We found that the precision of the reconstruction was at least in the order of 0.5 mm, where the size of a grid cell was 0.12 mm.

The volume resolution used for these images was $243 \times$




Figure 9: 3D scans obtained from a Cyberware laser scanner after spray-painting the chess pieces. Despite the obvious problems, these results are a combination of multiple scans that were aligned in hours of manual post-processing.

248×243 . Some problems in the reconstruction are visible near the bottom of the objects, where we separated two objects scanned at the same time. Air bubbles in the glass cause dents in some of the models higher up, especially when the bubble is close to the surface. These artifacts are similar to problems with metal objects such as inlays in medical CT scans. Some tomographic reconstruction methods can deal with this issue, but our implementation of SART currently ignores it.

Another artifact can be seen at the top of the king, where the cross is disconnected from the body. This separation is because the cross had previously broken off, and had to be glued back on. The seam is visible in the source images (Figure 10), and affects the reconstruction. This kind of artifact can be reduced by thresholding rays that are too dark, and not using them for reconstruction. In this case, the reconstruction will have missing geometry where the problem occurs. In the case of colored glass acquisition, ignoring dark rays results in a reconstruction of the visual hull of the problem region.

Another example is shown in Figure 11. With a reconstruction volume of $243 \times 344 \times 243$ and a final surface of 1.06 million polygons, this jar is our largest dataset. The high quality of the reconstruction becomes obvious when zooming in on the thread at the top of the jar.

Discussion. We believe the results presented in this paper are very promising. The physical setup and calibration are quite simple, and hence can be reproduced easily. The reconstructed geometry shows some artifacts, as described above. This is the case for practically every scanning method. However, we believe that these artifacts can easily be cleaned up with standard procedures. A big advantage of tomographic methods is that the resulting triangle mesh is an iso-surface of a density field, and hence closed, watertight, and manifold. In addition, we believe that the artifacts can be reduced or eliminated algorithmically in the reconstruction step. We consider this future work.




Figure 10: In the reconstruction (left), the cross of the king is separated from the body due to a fracture in the glass (right). The problem is reduced by not using rays that are too dark (bottom).




Figure 11: A glass jar (right), its reconstruction (left), and a detail image of the thread at the top (center).

Due to the dimensions of our prototype setup, the reconstruction region is currently relatively small: about 8 cm in diameter. These dimensions were chosen because glass cylinders are not readily available at larger diameters than about 15 cm, and we were initially concerned about the optical quality of the cylinder. However, based on our experience with the calibration method, we believe that a clear plastic container could easily be used instead, in which case the reconstruction region could be increased by a factor of 10 or more at a comparable cost.

In the future we would also like to further improve the robustness under mismatched refractive indices. Ideally, one would be able to dispense with the fluid completely. More realistically, we hope to replace the potassium thiocyanate solution with plain water, and add color as a contrast agent.

6. Conclusions

In this paper, we introduced a method for tomographic reconstruction of transparent objects by placing them into a fluid of a roughly equal index of refraction, and calibrating the ray-distribution in this medium, as seen from an observing camera.

The accuracy of the reconstructed geometry is comparable to commercial methods for acquiring opaque objects. Similar to those methods, some high-frequency noise is visible in the results, but several very effective methods are available for eliminating this noise as a postprocess.

At 45-100 minutes and an additional 60 minutes of software processing, the acquisition times are quite reasonable, especially considering that the result is already a closed, manifold surface, so that any additional post-processing is simplified compared to other object acquisition methods. For these reasons, we believe that the proposed method is a viable approach for acquiring transparent objects.

References

- [Arr99] ARRIGE S.: Optical tomography in medical imaging. *Inverse Problems* 15 (1999), R41–R93.
- [ART] ARTag web site.
<http://www.cv.iit.nrc.ca/research/ar/artag/>.
- [BEN03] BEN-EZRA M., NAYAR S.: What does motion reveal about transparency? In *Proc. of ICCV'03* (2003), vol. 2, pp. 1025–1032.
- [BER76] BITTNER J., EHRLICH G., RHEINGOLD E.: Efficient generation of the binary reflected gray code and its applications. *Communications of the ACM* 19, 9 (1976), 517–521.
- [Bra56] BRACEWELL R.: Strip integration in radio astronomy. *Australian Journal of Physics* 9 (1956), 198–217.
- [Bud94] BUDWIG R.: Refractive index matching for liquid flow investigations. *Experiments in Fluids* 17, 5 (1994), 350–355.
- [BV99] BONET J. D., VIOLA P.: Roxels: Responsibility weighted 3d volume reconstruction. In *Proc. of ICCV'99* (1999), pp. 418–425.
- [Cyb] CYBERWARE.: <http://www.cyberware.com>.
- [Fia05] FIALA M.: ARTag, a fiducial marker system using digital techniques. In *Proc. of CVPR* (2005), vol. 2, pp. 590–596.
- [GBH70] GORDON R., BENDER R., HERMAN G.: Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology* 29 (1970), 471–481.
- [GGSC96] GORTLER S., GRZESZCZUK R., SZELISKI R., COHEN M.: The Lumigraph. In *Proc. of ACM SIGGRAPH* (1996), pp. 43–54.

- [HHR01] HALL-HOLT O., RUSINKIEWICZ S.: Stripe boundary codes for real-time structured-light range scanning of moving objects. In *Proc. of ICCV* (2001).
- [IGM05] IHRKE I., GOLDLUECKE B., MAGNOR M.: Reconstructing the geometry of flowing water. In *Proc. of ICCV'05* (2005), pp. 1055–1060.
- [IM04] IHRKE I., MAGNOR M.: Image-based tomographic reconstruction of flames. In *Proc. ACM/EG Symposium on Animation (SCA'04)* (Aug. 2004), pp. 367–375.
- [IM05] IHRKE I., MAGNOR M.: Adaptive Grid Optical Tomography. In *IMA Vision, Video, and Graphics (VVG'05)* (July 2005), pp. 141–148.
- [IPD94] INGESSON L., PICKALOV V., DONNÉ A.: First tomographic reconstructions and a study of interference filters for visible-light tomography on rtp. *Review of Scientific Instruments* 66, 1 (Jan. 1994), 622–624.
- [KS01] KAK A., SLANEY M.: *Principles of Computerized Tomographic Imaging*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 2001. Reprint of 1988 book published by IEEE Press.
- [KS05] KUTULAKOS K., STEGER E.: A theory of refractive and specular 3d shape by light-path triangulation. In *Proc. of ICCV 2005* (2005), pp. 1448–1455.
- [Lau94] LAURENTINI A.: The visual hull concept for silhouette based image understanding. *IEEE PAMI* 16, 2 (1994), 150–162.
- [Lew90] LEWITT R.: Multidimensional digital image representations using generalized kaiser-bessel window functions. *Journal of the Optical Society of America* 7, 10 (1990), 1834–186.
- [LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proc. of ACM SIGGRAPH* (1996), pp. 31–42.
- [LL94] LACROUTE P., LEVOY M.: Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proc. SIGGRAPH '94* (1994), pp. 451–458.
- [LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D.: The Digital Michelangelo Project: 3D scanning of large statues. In *Proc. of ACM SIGGRAPH* (2000), pp. 131–144.
- [MBR*00] MATUSIK W., BUEHLER C., RASKAR R., GORTLER S., MCMILLAN L.: Image-based visual hulls. In *Proc. of ACM SIGGRAPH* (2000), pp. 369–374.
- [MI05] MIYAZAKI D., IKEUCHI K.: Inverse polarization raytracing: Estimating surface shapes of transparent objects. In *Proc. of CVPR '05* (2005), vol. 2, pp. 910–917.
- [MM98] METCALFE G., MANASSEH R.: Polydisperse sedimentation visualised by a refractive-index matching technique. In *Proc. 26th Australian & New Zealand Chemical Engineering Conference* (1998). Available at <http://resources.hightt.cmit.csiro.au/RManasseh/a983/a983.html>.
- [MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and rendering of transparent and refractive objects. In *Proc. of the Eurographics Symposium on Rendering* (2002), pp. 267–278.
- [MYW99] MUELLER K., YAGEL R., WHELLER J.: Anti-aliased 3D cone-beam reconstruction of low-contrast objects with algebraic methods. *IEEE Transactions on Medical Imaging* 18, 6 (1999), 519–537.
- [PD03] PEERS P., DUTRÉ P.: Wavelet environment matting. In *Proc. of the Eurographics Symposium on Rendering* (2003), pp. 157–166.
- [Pot87] POTMESIL M.: Generating octree models of 3D objects from their silhouettes in a sequence of images. In *Proc. CVGIP* (1987), pp. 1–29.
- [RHHL01] RUSINKIEWICZ S., HALL-HOLT O., LEVOY M.: Real-time 3D model acquisition. In *Proc. of ACM SIGGRAPH* (2001), pp. 438–446.
- [SAP*02] SHARPE J., AHLGREN U., PERRY P., HILL B., ROSS A., HECKSHER-SORENSEN J., BALDOCK R., DAVIDSON D.: Optical projection tomography as a tool for 3d microscopy and gene expression studies. *Science* 296 (Apr. 2002).
- [SD97] SEITZ S., DYER C.: Photorealistic scene reconstruction by voxel coloring. In *Proc. of CVPR* (1997), pp. 1067–1073.
- [SS02] SCHARSTEIN D., SZELISKI R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* 47, 1 (May 2002), 7–42.
- [VO90] VUYLSTEKE P., OOSTERLINCK A.: Range image acquisition with a single binary-encoded light pattern. *IEEE PAMI* 12, 3 (Feb. 1990), 148–164.
- [WFZ02] WEXLER Y., FITZGIBBON A., ZISSERMAN A.: Image-based environment matting. In *Proc. of the Eurographics Symposium on Rendering* (2002), pp. 279–290.
- [Woo80] WOODHAM R.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1 (1980), 139–144.
- [ZTCS99] ZHANG R., TSAI P.-S., CRYER J., SHAH M.: Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 21, 8 (Aug. 1999), 690–706.
- [ZWCS99] ZONGKER D., WERNER D., CURLESS B., SALESIN D.: Environment matting and compositing. In *Proc. of ACM SIGGRAPH* (1999), pp. 205–214.