

Tomographic Reconstruction of Transparent Objects

by

Borislav Danielov Trifonov

B.Sc., The University of South Florida, 2002

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

The Faculty of Graduate Studies

(Computer Science)

The University of British Columbia

December, 2006

© Borislav Danielov Trifonov 2006

Abstract

This thesis presents an optical acquisition setup and application of tomographic reconstruction to recover the shape of transparent objects. Although various optical scanning methods have been used to recover the shape of objects, they are normally intended for opaque objects, and there are difficulties in applying them to transparent ones. An alternative is to use X-ray computed tomography, but this requires a specialized setup, and computer graphics laboratories are not expected to have such equipment. Additionally, our setup avoids other problems of optical scanning, such as caused by occlusions, and is able to recover the internal geometry of the objects.

Table of Contents

Abstract	ii
Table of Contents	iii
List of Figures	v
Acknowledgements	vi
1 Introduction	1
1.1 Objectives	3
1.2 Basic assumptions	3
1.3 Overview	5
2 Related Work	6
2.1 Visible light scanning	6
2.2 X-ray computed tomography	7
2.3 Visual hull and voxel coloring	8
2.4 Optical tomography	8
3 Acquisition	10
3.1 Physical setup	10
3.2 Minimizing refraction	11
3.3 Calibration	13
	iii

Table of Contents

3.4 Acquisition	14
4 Reconstruction	17
4.1 SART	17
4.2 Projection	18
4.3 Backprojection	19
4.4 Implementation	20
5 Results	22
6 Conclusions and Future Work	31
Bibliography	33

List of Figures

3.1	The acquisition setup.	10
3.2	Front and rear calibration images.	13
3.3	Ray distribution and reconstruction region from calibration. .	14
3.4	Geometry for acquisition of clear objects.	16
4.1	Cross section and splatted views of the Kaiser-Bessel filter. .	19
5.1	Example of colored object.	22
5.2	Histogram of reconstructed volume densities.	23
5.3	Reconstructed colored object.	24
5.4	Clear object and colored fluid projection.	24
5.5	Queen and bishop.	25
5.6	Queen and bishop reconstructions.	26
5.7	King projection and reconstruction with defect	26
5.8	Laser scans of the painted bishop and king.	27
5.9	Chess pieces and reconstructions.	28
5.10	Jar and reconstruction.	29

Acknowledgements

I would like to acknowledge the support of my supervisor, Wolfgang Heidrich, and thank him for the initial idea for this project, as well as his help and patience during this research and implementation. I also owe gratitude to the other faculty members, including my advisor David Lowe; George Tsiknis for whom I worked as a teaching assistant; and thesis reader Michiel van de Panne. Derek Bradley was responsible for building the calibration software around the marker system, and Matthew Trentacoste created the camera control software. Last but not least, I want to thank Abhijeet Ghosh and the other graphics lab members for making my time at UBC more interesting.

Chapter 1

Introduction

Obtaining the 3D shape of real-world objects is a major area of graphics research. Such scanning can serve a variety of purposes, such as allowing existing models for movies and games to be simply scanned instead of relying on a human artist to create them in modeling software, but, most importantly, it can be used to digitize the full geometry of works of art for archiving, virtual museums, and so on.

Most commonly, laser scanners or stereoscopic imaging can be used to obtain 3D scans. Unfortunately, these methods rely on the assumption that the surface is opaque and diffuse. Near areas of significant detail, occlusion can prevent the whole outer surface from being reconstructed, and any internal geometry is inaccessible to the scanners. Glass and transparent (including colored) plastics cannot be digitized using such techniques directly. The objects can be painted, but that involves extra work, and may be too destructive for works of art. Moreover, the disadvantages mentioned above for surface scanning methods will now apply to the painted objects, losing the potential information transparency provides.

The full geometry of any solid object can be recovered by the use of transmission-based scanning. X-ray computed tomography is the most well known version of this. A narrow-spectral band X-ray source is used, with photon energy optimized to produce maximum contrast given the material

and size of the object to be scanned. A series of projections are taken in a planar or helical orbit around the object to produce views from different angles, where the value of the projection image at each point is mainly dependent on the absorption along the corresponding ray (scattering and refraction are usually assumed to be minimal). Various efficient and numerically stable reconstruction algorithms exist that can be used to produce a 3D volume of densities from the projection data. X-ray computed tomography has the disadvantage of requiring expensive equipment usually lacking in computer graphics laboratories, as well as operators trained in its use and safety procedures, including shielding and use of dosimeters.

Tomography can be modified to use with visible light for scanning non-refractive gaseous transparent objects. In the case of opaque objects, related algorithms can be used to recover the visual hull of the object. A different type of tomography has been done with infrared light in highly scattering mediums, and it is conceivable to immerse a transparent object in a scattering liquid or smoke for such scanning, but these algorithms are inefficient and numerically unstable.

Performing transmission tomography with visible light for transparent objects is problematic due to strong refraction at the solid-air interface. We solve this problem by immersing the object in a transparent cylinder filled with a fluid of a similar refractive index to the object, so that the refraction occurs at the cylinder-air and cylinder-fluid interface, and can be accounted for. Our calibration method determines the ray paths inside the fluid and through the object.

1.1 Objectives

The goal of this research was to develop a practical, non-destructive, and easily reproducible setup for digitizing the 3D shape of real-world transparent objects through a procedure that consists of:

- Visible light imaging.
- A calibration procedure for determining the path of light rays through the region to be reconstructed.
- An optimized version of a tomographic reconstruction method (simultaneous algebraic reconstruction technique, or SART).

The resulting project was published in [31].

1.2 Basic assumptions

There are several assumptions about the nature of the objects that are to be reconstructed. First, it is important that refractive index mismatch between the object and the fluid it is immersed in is small, on the order of 5%. If larger, the resulting reconstruction loses accuracy and the algorithm may even fail to converge. A second assumption is that all camera rays intersecting the object go through; that is, there are no opaque regions. The presence of such regions creates holes in the reconstructed geometry, and corrupts volume data in the vertical range which each of the regions spanned. However, as discussed later, such effects can be minimized. Third, both objects and fluid must have low scattering, as algebraic reconstruction does not take this effect into account; different algorithms are used for tomography with strong scattering. Fourth, we assume the refractive index

of the object can be matched closely by a relatively safe and easy to obtain fluid. Our use of potassium thiocyanate allows us to match a number of glasses, but not things such as high-lead content crystal glass. Various fluids, such as oils (possibly in solution with alcohols), can be used to match a wider range of transparent materials, including plastics.

There are a few practical issues to be considered. We assume the objects will not contain any parts that may be affected by the refractive index matching fluid. In our case, the fluid is potassium thiocyanate solution, and it is corrosive to metals; this would exclude, for example, art objects that contain metal or metalized areas. Another assumption is that there are no refractive heterogeneities within the object; from this follows that internal hollow regions fully disconnected from the outside are not acceptable, since the matching fluid cannot fill them when the objects are immersed. Additionally, objects must fit in the cylinder used to contain the fluid, well clear of the sides, where distortion due to lensing is significant from the camera's perspective. We also assume that colored transparent objects will not be so dark as to limit the contrast significantly with practical lighting and exposure times. The setup needs to be such that rays are minimally divergent through the reconstruction region, as that can produce sampling artifacts. Although it is possible to deal with such sampling issues in the reconstruction algorithm, we assumed that would not be the case in our setup as it is not difficult to set up the camera in such a way that ray divergence would be limited, and sampling density within a given slice perpendicular to the beam would be only somewhat non-uniform.

1.3 Overview

In Chapter 2, we discuss related work in capturing the shape of objects in the real world. Then, in Chapter 3, we describe our physical setup and acquisition process, followed by the algebraic reconstruction method from tomography we use in Chapter 4. A presentation of results in Chapter 5 and a wrap-up in Chapter 6 conclude the thesis.

Chapter 2

Related Work

Although 3D geometry acquisition is usually associated with computer graphics, it has a longer history in other fields such as medical imaging and engineering. A number of techniques are related to our work.

2.1 Visible light scanning

3D scanning with visible light can be grouped into passive and active methods. Passive ones often originate from computer vision research, and the most well known are ones that use stereoscopy, where stereo disparity between images can be used to determine depth information [28]. Shape from stereo relies on either matching image areas by correlation, or finding corresponding features in the images. Due to limited precision, problems with occlusion, and matching ambiguities, stereo is more suited to vision applications such as robot navigation than obtaining accurate shapes of objects.

Another passive method is shape from shading [35, 36], which usually use multiple light sources instead of multiple light views. In general, such approaches are limited by the need to know the surface reflectance, and usually it is assumed to be diffuse. Neither shape from shading nor stereoscopic methods can deal with specularities and translucencies.

As opposed to these passive approaches, active lighting methods require specialized illumination. They use either encoded patterns of light [26, 34,

37]) or lasers [4]. In these methods, a calibrated structured light source projects encoded (so the software can distinguish them) vertical planes into the scene, and a camera images the contour lines. From the horizontal shift of a point on each stripe, the software can compute a 3D position. As in the case of passive methods, specularity and translucency can cause gross errors. A different approach is needed.

Environment matting techniques [26, 34, 37] are able to capture appearance of transparent objects, but are not able to get the actual 3D shape.

2.2 X-ray computed tomography

There are several medical imaging techniques that produce 3D volumes from objects, including MRI and PET scanning. Computed tomography [13], however, is the most commonly used, and is also used in engineering applications, such as to image defects in materials. The most frequently encountered type of computed tomography, transmission CT, assumes that refraction and scattering are negligible, and that density variations influence the transmitted brightness along a ray. A narrow band of X-ray wavelengths are chosen so that all rays at least partially pass through the object to be imaged, yet sufficient contrast remains.

The most common approach for reconstructing volumetric data from a set of projections is based on the Fourier Slice Theorem [2, 13]. One-dimensional Fourier transforms of lines parallel to the plane of rotation are filtered and backprojected to recover 2D slices of the objects. Each line must be illuminated by either parallel X-ray beams, or fan beams from a point source. Usually, the CT scanning involves multiple orbits to produce a volume from a collection of the reconstructed slices.

Another approach that is becoming more common is Algebraic Reconstruction Techniques (ART) [7, 13], which essentially solves a set of linear equations by iteratively updating current estimated voxel densities with those of the projection image. ART methods normally proceed in a ray by ray basis, but an alternative, Simultaneous ART (SART), proceeds one projection at a time. A variant of the latter is what we use in our application. A significant advantage is that SART can handle cone beams, and is not restricted to a carefully aligned, equally spaced set of projections.

Statistical methods have also been applied to computed tomography, but have generally been very inefficient [14].

2.3 Visual hull and voxel coloring

Visual hull reconstruction techniques share some similarities with tomography [16, 21, 27], but they can only produce the visual hull. A more promising approach for opaque objects may be [20], where the reflected light is taken into account. Voxel coloring [29] solves a correspondence problem as visible light scanning methods, but has the tomography-like constraint for the camera positioning and reconstruction region. Although it is an improvement over the visual hull, it is still unable to capture occluded details. None of these methods work with transparent objects.

2.4 Optical tomography

Optical transmission tomography has been used to acquire the shape of non-refracting, partially-transparent objects such as plasmas [12] and flames [9, 10]. However, all solid objects have significant refraction, making such meth-

ods unsuitable. Emission-based reconstruction of fluids containing fluorescent dyes was studied by [11]. Tomography not based on straight transmission has also been studied. Optical tomography in biological tissues with high scattering generally relies on statistical methods or non-linear optimization [1], and is very inefficient and overly complex in the case of no scattering. Related methods include the use in microscopy of phase information in the interferometry-like optical coherence tomography [30]. In microscopy applications refractive index matching has recently been applied to transmission tomography on a small scale [5], with filtered backprojection used for reconstruction; however this lacks scalability as it is limited to the known ray paths within the microscope's imaging field. Our calibration procedure and the use of SART allows us to have a much more flexible, macroscopic setup.

Chapter 3

Acquisition

The simplicity of our setup is apparent from Figure 3.1, consisting of an optical table on which we mounted a camera on the left, a transparent cylinder holding the object and refractive index matching fluid, and a brightly lit diffuse background surface on the right. After a simple calibration step to obtain the ray paths in the cylinder, a number of projections are taken for different rotations of the turntable (several exposures for each angle). To scan non-colored transparent objects, dye is added to the fluid.

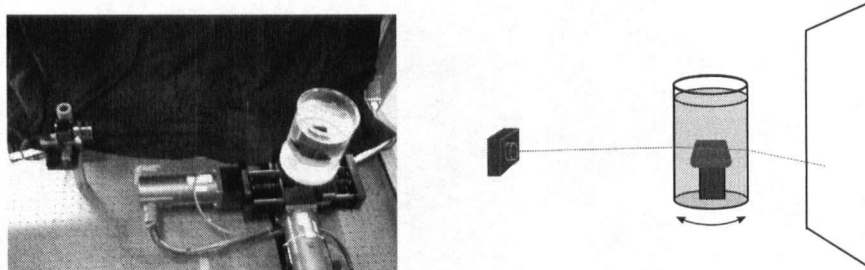


Figure 3.1: The acquisition setup.

3.1 Physical setup

Since we were concerned with optical quality, we used a precision glass cylinder from a scientific supplier, with a diameter of 15 cm. Due to lensing

distortion near the sides, that gave a cylindrically shaped usable reconstruction region of about 9 cm diameter, and this is the limit of objects we can scan, as they must fit fully within that region.

A base centering the cylinder on the turntable, and the object support stand and calibration panel holders were made of plastic using a rapid prototyping machine [4]. We used a 1.5 megapixel machine vision camera that could capture 12-bit linearly quantized, Bayer mosaicked images. In order to get a dynamic range beyond the 12 bits of the camera, multiple exposures were used and combined into high dynamic range images using HDRGen [33].

The background was a diffuse white surface which we illuminated with a strong light at an angle from the side, so as to avoid any reflections on the cylinder surface. Higher background brightness has the advantage of speeding up acquisition by reducing needed exposure times. Since the surface was not completely uniformly lit, we used a calibration image with the cylinder without an object in it so that we could factor out the unevenness.

3.2 Minimizing refraction

Since the refractive index of the object to be scanned needs to be approximately matched by the fluid it is immersed in, it was necessary to find a practical fluid with a refractive index that would allow some adjustment in the target range. As we were most interested in glass rather than plastic objects, we examined possibilities with refractive index of 1.5 to 1.6. Borosilicate glasses are commonly around 1.5, with more common glasses somewhat higher. Some types of glasses, such as lead crystal, have a very high refractive index and we did not attempt scanning such materials.

A number possible matching fluids exist [22], including benzene, and

various mixtures of alcohols and other hydrocarbons, as well as different oils. Most of these are either prohibitively toxic for use in a typical graphics laboratory, or difficult or expensive to obtain. One simple solution was to use common mineral oil, but this has the disadvantage of having a fixed index of refraction, without the possibility of adjustment that a solution has. We did test the cheaper alternative of vegetable oil, but were unsuccessful in finding a dye that would not cause significant scattering when dissolved in the oil. A very concentrated sugar solution can reach the refractive index of glass, but due to the high viscosity of the syrup, it is difficult to work with. With these considerations in mind, we chose to use a solution of potassium thiocyanate in water [3], which, while corrosive and an irritant, was deemed sufficiently safe with careful handling.

By varying the concentration, a range of refractive indices can be matched. At 80%, it has a refractive index of 1.5, which is suitable for borosilicate glasses and some plastics. To obtain a higher index of about 1.55, we created a super-saturated solution by heating it to dissolve more of the salt, after which the solution was allowed to cool (the refractive index varies somewhat with temperature). The onset of crystal formation and resulting fall of refractive index was slow enough to allow time for complete acquisitions. The exact index of refraction achieved was not measured due to lack of instrumentation.

It was found that the potassium thiocyanate solution had significant dispersion, which is in addition to the dispersion caused by the objects. It was thus necessary to limit the wavelength of light used to a small portion of the spectrum, by only using the green pixels from the Bayer mosaicked image from the camera. Additional narrowing was accomplished by the use of a green filter mounted in front of the camera lens.

3.3 Calibration

For tomographic reconstruction, it is necessary to know the path of each ray through the reconstruction region. Similar to lumigraph/lightfield rendering [8, 17], we parameterized the rays by two planes. By placing the planes inside the cylinder, we do not need to be concerned with any effects outside this region, such as refraction at the cylinder-air and cylinder-fluid interfaces, as long as these remain symmetric under turntable rotation. As we only perform calibration once, it was critical to center the cylinder precisely.

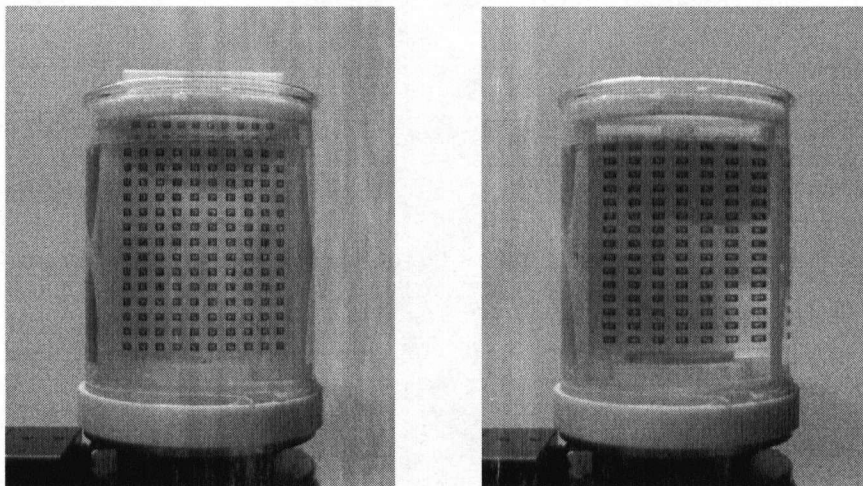


Figure 3.2: Front and rear calibration images.

The planes and their positioning structure were made with the rapid prototyping machine, and we attached calibration grids to their front. The calibration pattern and recognition was done using the ARTag system [6], which was able to detect almost all markers in the image despite the strong lensing distortion, given relatively even illumination. The system identifies

the corners of the square markers, and ray coordinates on each plane are interpolated between these points. No low pass filtering was necessary as the refractive distortion varies slowly with respect to the marker density.

For each camera pixel, using the coordinates of the ray's intersection on the two planes, and knowing the plane geometry, it is possible to determine the path of the ray segment within the cylinder. The region formed by the intersection of ray beams from all turntable orientations forms the reconstruction region within which the objects must fit. A decimated representation of the ray segments is shown in Figure 3.3 (vertical decimation is increased for clarity), along with the reconstruction region formed by the intersections of beams from all views.

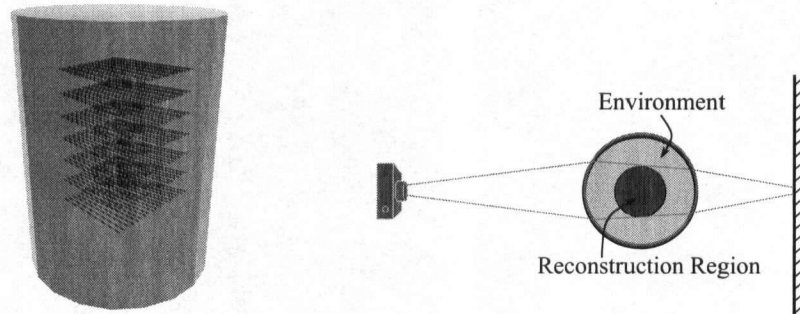


Figure 3.3: Ray distribution and reconstruction region from calibration.

3.4 Acquisition

Scanning consists of imaging a projection from a number of different rotations. The number of projections needed depends on the resolution of the object, and following [24], we use on the order of 0.67 times the horizontal

volume resolution. In order to improve results given the higher ray density at the rear of the reconstruction region, images are taken around a full rotation rather than just 180° .

The set of exposures from which each projection is created is adjusted to get the full contrast range of the region of the image occupied by the object, so that at the shortest exposure time the darkest pixels are black, and at the longest, the lightest ones are saturated. Since HDRGen failed to correctly derive the camera curves for some views, the same camera parameters were used for all projections in a set, even though in theory they should not vary.

In order to scan clear objects, it is necessary to add a contrast agent (food coloring) to the refractive index matching fluid. This creates a problem, since now there is light absorption along each ray outside the reconstruction region. Referring to Figure 3.4, the absorption along a ray is given by

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

where a and b are the intersection points of the ray with the cylinder and α is the absorption coefficient. Since we have an image of the empty cylinder, from which we have factored out the background image, we can determine α for each ray. In practice, there is some variation over the cylinder due to measurement errors, so we average the value obtained from all rays.

If b and c are the intersection points of the ray and the reconstruction region (which can be computed from the calibration data), the absorption due to the ring of fluid outside this region is

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

for each ray. The pixel values associated with the rays can be simply divided by their corresponding A_{env} to extract an image of the reconstruction region, so that tomography can be applied as in the case of colored objects.

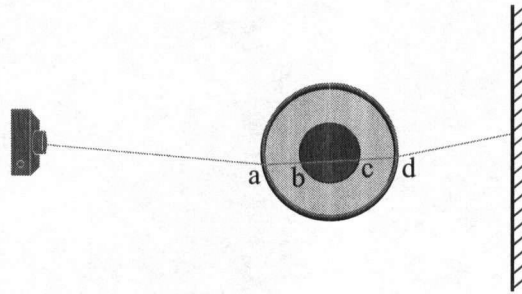


Figure 3.4: Geometry for acquisition of clear objects.

After acquisition, the Bayer mosaicked images (and ray data) are re-sampled to a resolution matching that of the volume to be reconstructed, and cropped to the smallest reconstruction region that fits the object for efficiency reasons.

Chapter 4

Reconstruction

Although the Fourier Slice Theorem based reconstruction is very efficient, it is for either parallel rays or those from point sources. In our case, the ray distribution does not match either case, necessitating the use of a more general method such as Algebraic Reconstruction Techniques (ART). Specifically, Simultaneous Algebraic Reconstruction (SART) was chosen; although it is somewhat slower than ART, it produces less sampling artifacts when ray density varies. No visible sampling artifacts manifested themselves in our testing on simulated data, so that SART was sufficient without taking explicit account of sampling nonuniformities, thereby simplifying the algorithm.

4.1 SART

A ray is attenuated exponentially by absorption along its path. In the discrete case, we assume that within each discrete region, attenuation is constant, and the integration becomes summation. If we convert the images and operations to log space, we get absorption

$$\log A = - \int_b^c \alpha(t) dt \approx - \sum_i \alpha_i.$$

where the region under consideration is between a and b along the ray, and α_i are the densities at each discrete region along the ray. Through

several iterations through the set of projections (randomized each time), we perform a forward projection through the volume, compute an error image, and update voxels during a backprojection step.

Our approach derives from the SART version described in [24, 25]. The volume is sliced along the axis that is most perpendicular to the direction from which the current projection has been taken (minimizing the angle between the slices and the image plane). This allows us to walk the ray front slice by slice, accumulating filter-weighted values and later backprojecting the correction image. After several iterations through the volume, density values at each voxel converge to the reconstruction region from the acquisition step, and an isosurface can be then extracted.

4.2 Projection

The first step is object order volume rendering similar to [15]; however, sampling is done differently. Walking the ray front through each volume slice, at each ray-slice intersection, the voxel within a filter window are weighted by a Kaiser-Bessel filter [18] (Figure 4.1). This radially symmetric filter was precomputed and pre-integrated (splatted) in MATLAB; thus, determining the filter weight is a fast table look-up. A filter radius of two was sufficient for anti-aliasing, while a larger radius only increased blurring. The resulting (log) absorption for ray i through the volume is

$$\log A_i^{(k)} = \frac{\sum_n w_{in} \alpha_n^{(k)}}{\sum_n w_{in}},$$

where the $\alpha_n^{(k)}$ denotes the current density estimate of voxel v_n , and w_{in} is the filter weight of the ray for that voxel. For efficiency, the sum in the denominator is accumulated in parallel, and the division performed after

all slices have been processed. The resulting log absorption rendering is subtracted from the log projection, giving the pixels of the correction image for the next step, $\Delta A_i = \log A_i - \log A_i^{(k)}$.

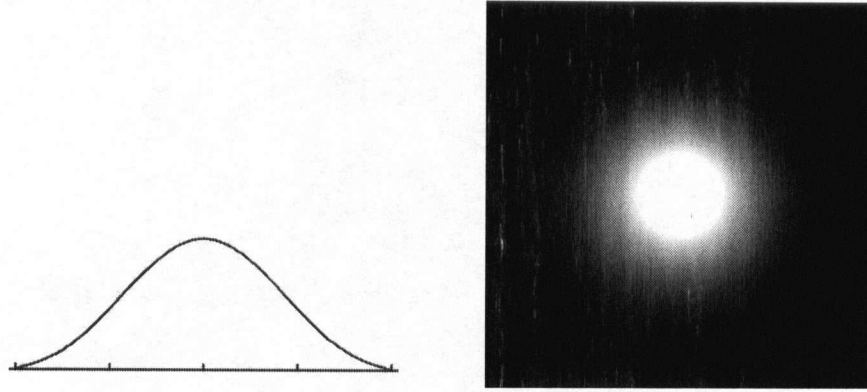


Figure 4.1: Cross section and splatted views of the Kaiser-Bessel filter.

4.3 Backprojection

To apply the corrections to the volume, the correction value for each ray is multiplied by a relaxation parameter λ , and distributed to each voxel along the ray:

$$\alpha_j^{(k+1)} = \alpha_j^{(k)} + \lambda \sum_i \frac{\sum_j w_{ij} \Delta A_i}{\sum_j w_{ij}}.$$

As in the forward projection step, this proceeds slice by slice, and within each slice, for a given ray, all voxels within the filter window are found, and the weighted corrections are applied to them. As in equation (4) in [24], normalization can be deferred after all slices have been processed, and the

rearranged equation becomes

$$\alpha_j^{(k+1)} = \alpha_j^{(k)} + \lambda \frac{\sum_i \frac{A_i - \sum_n w_{in} \alpha_n^{(k)}}{\sum_n w_{in}} w_{ij}}{\sum_i w_{ij}}.$$

Choosing λ affects the number of iterations required for reconstruction. Larger values speed up convergence, but if too large, the algorithm will give too much weight to the last projection that was processed, and will not converge. An initial value of 0.04 to 0.1 was useful, and best results from the fewest iterations were achieved by reducing λ after each iteration through the set of projections. In most cases, after three iterations there was no further improvement (measured by the total correction applied each time), and even two iterations produced good results.

4.4 Implementation

An optimization for SART proposed in [25] is to cache filter weights determined during the projection step so they can be used for the backprojection one. Due to the need for a large amount of memory, they propose going through the volume slices in slabs shifting vertically, where the slab thickness depends on the vertical travel of the ray through the volume, as well as the filter weight. Due to the lensing of rays through our volume, we decided it may be possible that slabs would be too thick for this to be a significant improvement.

Instead, we traverse whole slices at a time and the filters lookup operation is performed separately during backprojection. In order to have good cache coherency, the volume is laid out in memory in the order in which it will be accessed. Since slice orientation (parallel to x-y or y-z plane) is determined by which one is most perpendicular to the direction from which

the corresponding projection was taken, a layout that is optimal for half the projections is sub-optimal for the other half. Our image order is randomized for an iteration, but afterwards the sequence is sorted in several bins so that the volume may be reorganized in memory just a few times.

An additional optimization was parallelization using OpenMP. Running it on a dual-CPU system with two hardware threads per CPU, it was possible to parallelize most processing-intensive parts of the algorithm, since each slice can be processed independently, only needing accumulation to compute the correction image.

Using SIMD instructions on the CPU is unfortunately not possible for this algorithm's core, given this type of filter kernel sampling (other than for converting floating point values to filter table indices), since current CPUs do not provide the scatter and gather operations that are necessary for vectorization in this case (GPUs do, however).

In order to improve results with mismatches between the refractive index of the fluid and object, it is possible to weigh down rays that are most likely to be incorrect. If a ray intersects an object at an acute angle, it will be more affected by refraction. To find these rays efficiently, after each iteration through the set of images, gradients are computed for the density estimate at each voxel. The cosine of the maximum angle a ray encounters during projection between its direction and the gradients of the voxels it intersects is multiplied to λ , lowering the ray's likely erroneous contribution. In order to have a noticeable effect, this necessitates the use of more than the two or three iterations that are otherwise sufficient.

Chapter 5

Results

Initially, we tested our implementation on simulated data from a ray tracer, which allowed us to determine that refractive index mismatch up to around 5% still produced acceptable results (small features not swamped by artifacts). Raising the mismatch resulted in both global deformation and increased surface roughness. The volume resolutions for the synthetic and acquired data were set to correspond to the projection image resolutions.

To test reconstruction of colored objects, without a dye in the fluid, we used a red glass object that is shown, along with one of the projections, in Figure 5.1. A total of 360 projections were taken.

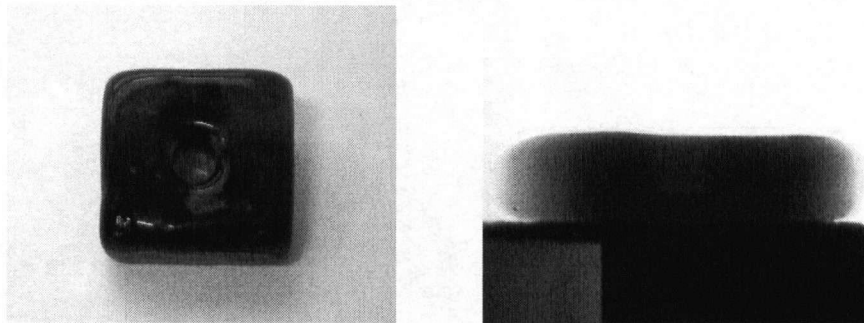


Figure 5.1: Example colored transparent object and one of the projections.

The object was reconstructed on a volume of $475 \times 276 \times 475$ voxels (a voxel corresponding to about 0.12 mm), taking about an hour and a half

on our dual 3.6 GHz system for five iterations through all projections. The marching cubes algorithm [19] was used to extract an isosurface from the volumetric data. In an ideal case of a uniformly absorbing object material, the voxels would have two possible values, corresponding to the fluid and to the material. In practice, we get a histogram with two peaks. In order to extract an isosurface from the volume, an iso-value is chosen in the valley between the peaks (Figure 5.2).

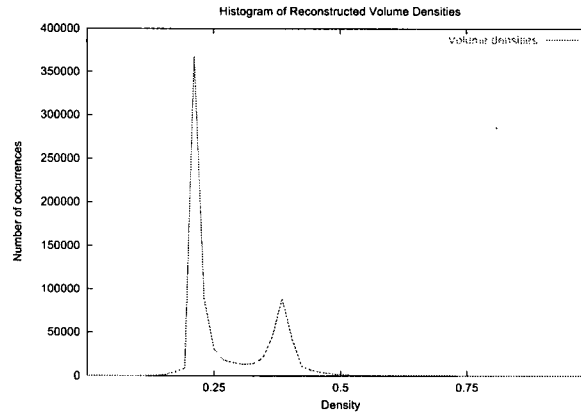


Figure 5.2: Histogram of reconstructed volume densities.

The reconstruction of the colored object can be seen in Figure 5.3. Note the reconstructed internal geometry of the hole shown in the cut-away view, a feature unique to tomographic approaches. The bottom of the object necessitates a simple cleanup; it is an artifact due to the opaque base on which the object rested during acquisition.

To reconstruct clear objects, food coloring dye was used to make the fluid absorptive. The amount of dye to use is a compromise between increasing contrast with more dye, while retaining enough brightness, given our light

source, to avoid the need for very long exposure times. Figure 5.4 shows one projection of a queen and bishop pieces from a colorless glass chess set. All pieces were reconstructed by using three iterations.

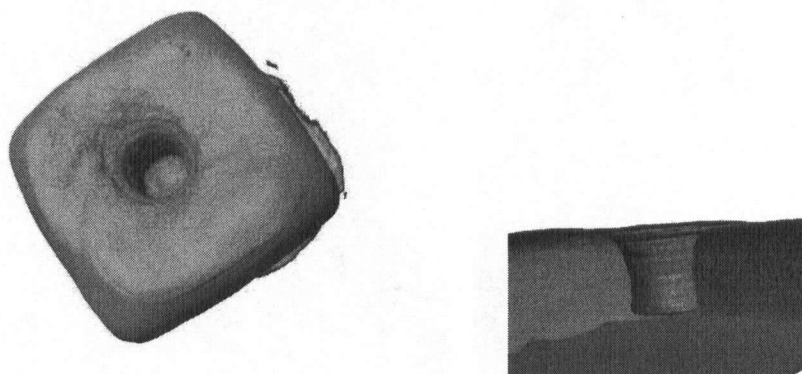


Figure 5.3: Reconstructed object from Figure 5.1 and cut-away view showing internal geometry.



Figure 5.4: Clear object and colored fluid projection.

The queen and bishop and their reconstructions, manually separated from the $243 \times 248 \times 243$ volume, can be seen in Figure 5.5 and Figure 5.6,

respectively. A number of artifacts are visible in this example. Since the objects were positioned beside each other, they had to be manually separated, thus the reconstructed bases are not clean. Additionally, imperfections in the glass, especially air bubbles, resulted in holes/dents in the recovered surfaces, because they appear dark in the projections and all rays that cross them incorrectly lower the density of voxels along their paths.

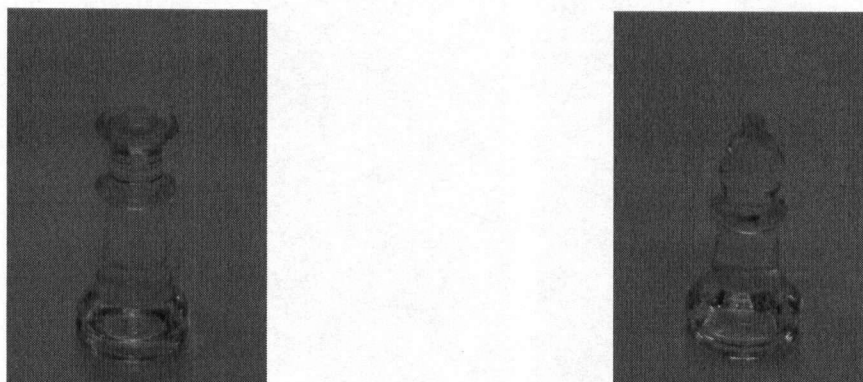


Figure 5.5: Queen and bishop.

This problem is most obvious in the king piece. The cross was broken from the body and subsequently glued. Figure 5.7 clearly shows how the dark seam corrupts reconstruction in the region around it. We were able to get some improvement by clamping the darkness of the pixel value to that which would result from the highest absorption possible along the given ray (which is possible since we compute the absorption coefficient of the fluid during reconstruction). Moreover, we determined that the limited divergence of rays through the volume (which was also the case with our physical setup), did not cause the visible sampling artifacts discussed by [24], so we simplified our implementation by not dealing with it explicitly.

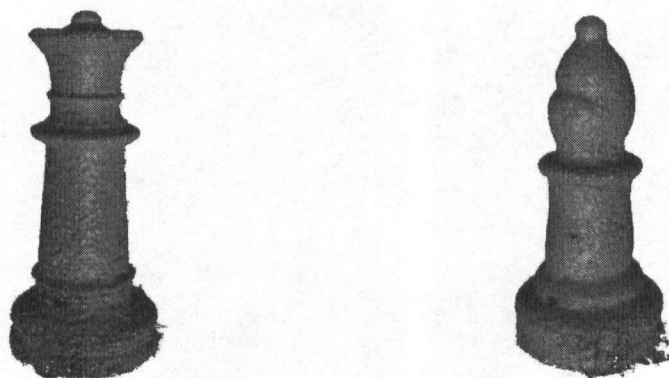


Figure 5.6: Queen and bishop reconstructions.

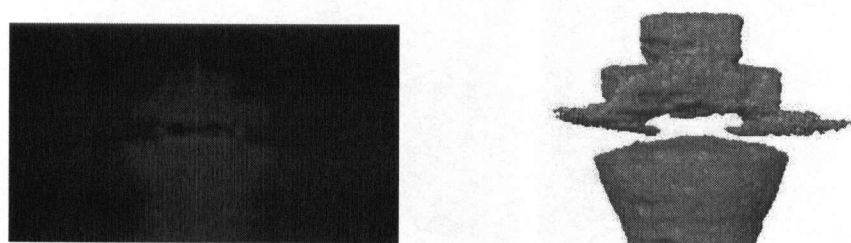


Figure 5.7: A projection showing the defect in the king, and the resulting artifact in reconstruction.

In order to test the accuracy, we tried spray-painting the figures and using the Cyberware laser scanner (Figure 5.8), a long process of merging various scans for each figure, but the results have artifacts and a significant lack of detail. Using digital calipers instead, we were able to determine an accuracy of reconstruction around 0.5 mm on the voxel grid pitch of 0.12 mm for the rectangular object.

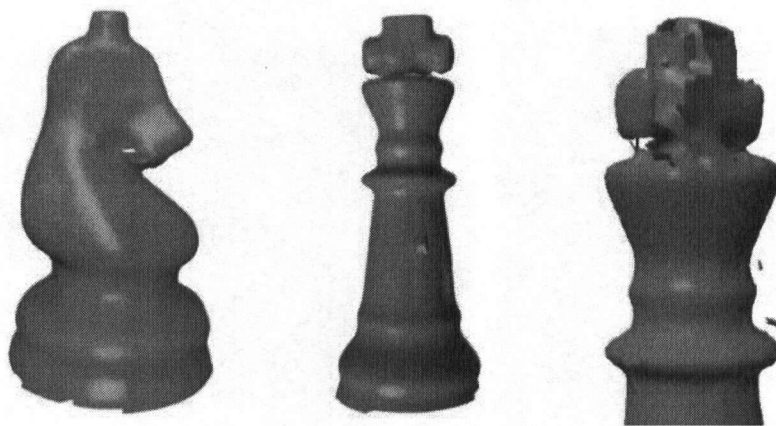


Figure 5.8: Laser scans of the painted bishop and king.

The photographs and reconstructions of the other chess pieces are shown in Figure 5.9.

To test a larger dataset, we performed reconstruction of the jar in Figure 5.10 in a $243 \times 344 \times 243$ volume. The threads on the jar's throat are about one millimetre in the thinnest parts, and the noise floor is visibly below that scale in the reconstruction.

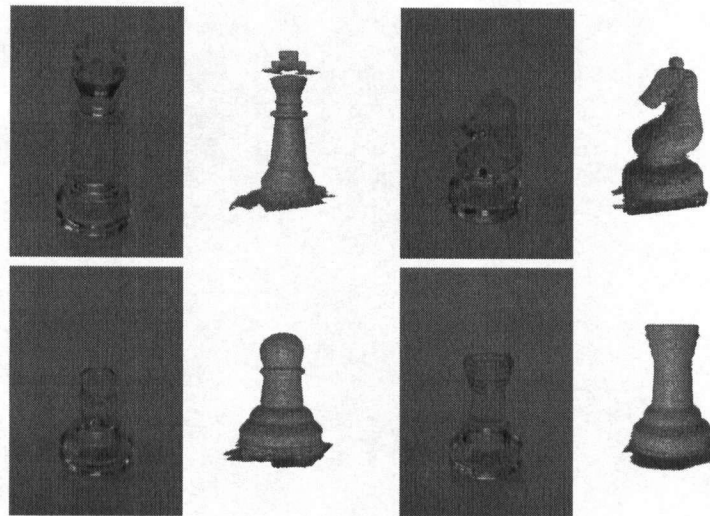


Figure 5.9: Chess pieces and reconstructions.

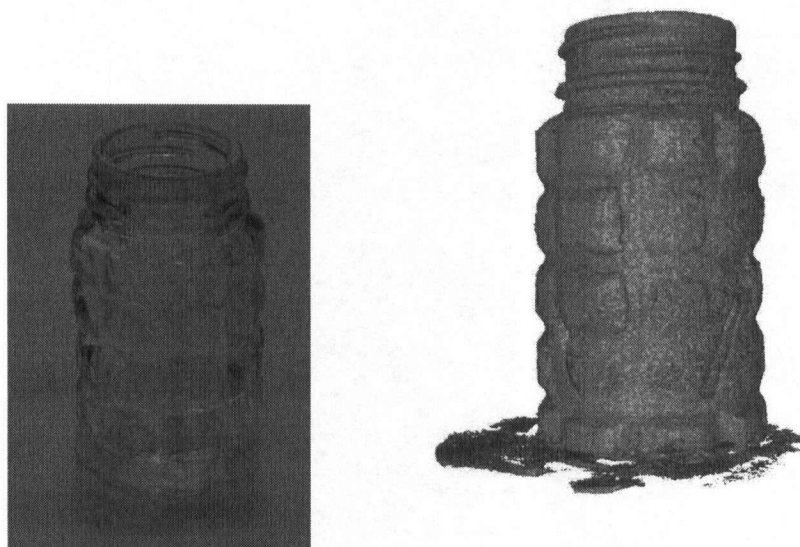


Figure 5.10: Jar and reconstruction.

The weighting of ray correction values by the angle between the ray and local gradients to deal with small refractive index mismatches had limited success: it caused some improvement in the reconstruction quality, but also had a penalty in terms of speed. Reconstruction time for smaller volume resolutions is significantly faster – on the order of 20 minutes for a cubic volume 128 voxels per side on our system. In these cases time is dominated by the acquisition step, due to the need for fairly long exposure times, and multiple exposures, used for each projection. The former could be sped up by the use of a brighter light source, and the latter by a high dynamic range camera. For resolutions of 512 voxels per dimension, reconstruction time is

several hours for a single iteration, so for such and higher resolutions, the use of graphics hardware to accelerate reconstruction similar to [23] may seem attractive, though implementation would not be as straightforward due to some of the differences in our adaptation of SART.

Chapter 6

Conclusions and Future Work

We created a practical, non-destructive system for acquiring the 3D shape of real-world objects through the use of tomographic reconstruction with visible light acquisition, using a simple and inexpensive setup for refractive index matching, and a calibration method of finding the ray paths through the reconstruction region. The system is capable of a fraction of a percent accuracy, although some post-processing may be needed to remove some high frequency noise. Acquisition and reconstruction together take from one to two hours depending on number of projections and reconstructed volume resolution. Thus, we have achieved the objectives outlined in the introduction.

For improved resolution of reconstructions, simply increasing projection resolution is not sufficient. A number of possible changes may be necessary: more precise refractive index matching, higher dynamic range images, finer calibration, and more projections.

Possibilities for future work include the exploration of the use of other fluids for matching the refractive index, allowing a wider range of transparent materials to be used; using a transparent plastic instead of glass cylinder to allow large objects to be scanned without making price a prohibitive

factor; explicitly taking local ray sampling density into account to allow more freedom in the geometry of the camera and object setup; and the reduction of the number of needed projections and increase in efficiency by the use of priors, similar to their application in X-ray tomography by [32].

Bibliography

- [1] S. Arrige. Optical tomography in medical imaging. *Inverse Problems*, 15:R41–R93, 1999.
- [2] R. Bracewell. Strip integration in radio astronomy. *Australian Journal of Physics*, 9:198–217, 1956.
- [3] R. Budwig. Refractive index matching for liquid flow investigations. *Experiments in Fluids*, 17(5):350–355, 1994.
- [4] Cyberware. <http://www.cyberware.com>.
- [5] M. Fauver and E. J. Seibel. Three-dimensional imaging of single isolated cell nuclei using optical projection tomography. *Optics Express*, 13(11):4210–4223, 2005.
- [6] M. Fiala. ARTag, a fiducial marker system using digital techniques. In *Proc. of CVPR*, volume 2, pages 590–596, 2005.
- [7] R. Gordon, R. Bender, and G. Herman. Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology*, 29:471–481, 1970.
- [8] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The Lumigraph. In *Proc. of ACM SIGGRAPH*, pages 43–54, 1996.

- [9] I. Ihrke and M. Magnor. Image-based tomographic reconstruction of flames. In *Proc. ACM/EG Symposium on Animation (SCA'04)*, pages 367–375, August 2004.
- [10] I. Ihrke and M. Magnor. Adaptive Grid Optical Tomography. In *IMA Vision, Video, and Graphics (VVG'05)*, pages 141–148, July 2005.
- [11] Ivo Ihrke, Bastian Goldluecke, and Marcus Magnor. Reconstructing the geometry of flowing water. In *International Conference on Computer Vision 2005*, pages 1055–1060, 2005.
- [12] L. Ingesson, V. Pickalov, and A. Donné. First tomographic reconstructions and a study of interference filters for visible-light tomography on rtp. *Review of Scientific Instruments*, 66(1):622–624, January 1994.
- [13] A. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 2001. Reprint of 1988 book published by IEEE Press.
- [14] J.S. Kole. Statistical image reconstruction for transmission tomography using relaxed ordered subset algorithms. *Physics in Medicine and Biology*, 50:1533–1545, March 2005.
- [15] P. Lacroute and M. Levoy. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Proc. SIGGRAPH '94*, pages 451–458, 1994.
- [16] A. Laurentini. The visual hull concept for silhouette based image understanding. *IEEE PAMI*, 16(2):150–162, 1994.

- [17] M. Levoy and P. Hanrahan. Light field rendering. In *Proc. of ACM SIGGRAPH*, pages 31–42, 1996.
- [18] R. Lewitt. Multidimensional digital image representations using generalized kaiser-bessel window functions. *Journal of the Optical Society of America*, 7(10):1834–186, 1990.
- [19] W. Lorensen and H. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proc. of ACM SIGGRAPH*, pages 163–169, 1987.
- [20] D. L. Marks, R. A. Stack, D. J. Brady, and D.C. Munson Jr. Cone-beam tomography with a digital camera. *Applied Optics*, 40(11):1795–1805, 2001.
- [21] W. Matusik, C. Buehler, R. Raskar, S. Gortler, and L. McMillan. Image-based visual hulls. In *Proc. of ACM SIGGRAPH*, pages 369–374, 2000.
- [22] G. Metcalfe and R. Manasseh. Polydisperse sedimentation visualised by a refractive-index matching technique. In *Proc. 26th Australian & New Zealand Chemical Engineering Conference*, 1998. Available at <http://resources.hihett.cmit.csiro.au/RManasseh/a983/a983.html>.
- [23] K. Mueller and R. Yagel. Rapid 3D cone-beam reconstruction with the simultaneous algebraic reconstruction technique (sart) using 2d texture mapping hardware. *IEEE Transactions on Medical Imaging*, 19(12):1227–1237, 2000.

- [24] K. Mueller, R. Yagel, and J. Wheller. Anti-aliased 3D cone-beam reconstruction of low-contrast objects with algebraic methods. *IEEE Transactions on Medical Imaging*, 18(6):519–537, 1999.
- [25] K. Mueller, R. Yagel, and J. J. Wheller. Fast implementation of algebraic methods for 3d reconstruction from cone-beam data. *IEEE Transactions on Medical Imaging*, 18(6):538–548, 1999.
- [26] P. Peers and P. Dutré. Wavelet environment matting. In *Proc. of the Eurographics Symposium on Rendering*, pages 157–166, 2003.
- [27] M. Potmesil. Generating octree models of 3D objects from their silhouettes in a sequence of images. In *Proc. CVGIP*, pages 1–29, 1987.
- [28] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47(1):7–42, May 2002.
- [29] S. Seitz and C. Dyer. Photorealistic scene reconstruction by voxel coloring. In *Proc. of CVPR*, pages 1067–1073, 1997.
- [30] P. H. Tomlins and R. K. Wang. Theory, developments and applications of optical coherence tomography. *Journal of Physics D: Applied Physics*, 38(15):2519–2535, 2005.
- [31] B. Trifonov, D. Bradley, and W. Heidrich. Tomographic reconstruction of transparent objects. In *Proc. of the Eurographics Symposium on Rendering*, pages 51–60, 2006.
- [32] V. L. Vengrinovich, Yu B Denkevich, and G-R Tillack. Reconstruction of three-dimensional binary structures from an extremely limited num-

ber of cone-beam x-ray projections. choice of prior. *Journal of Physics D: Applied Physics*, 32:2505–2514, 1999.

- [33] G. Ward. <http://www.anywhere.com>.
- [34] Y. Wexler, A. Fitzgibbon, and A. Zisserman. Image-based environment matting. In *Proc. of the Eurographics Symposium on Rendering*, pages 279–290, 2002.
- [35] R. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [36] R. Zhang, P.-S. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 21(8):690–706, August 1999.
- [37] D. Zongker, D. Werner, B. Curless, and D. Salesin. Environment matting and compositing. In *Proc. of ACM SIGGRAPH*, pages 205–214, 1999.