

Fluorescent Immersion Range Scanning

Matthias B. Hullin¹Martin Fuchs¹Ivo Ihrke^{2,1}Hans-Peter Seidel¹Hendrik P. A. Lensch¹¹MPI Informatik* ²University of British Columbia[†]

Abstract

The quality of a 3D range scan should not depend on the surface properties of the object. Most active range scanning techniques, however, assume a diffuse reflector to allow for a robust detection of incident light patterns. In our approach we embed the object into a fluorescent liquid. By analyzing the light rays that become visible due to fluorescence rather than analyzing their reflections off the surface, we can detect the intersection points between the projected laser sheet and the object surface for a wide range of different materials. For transparent objects we can even directly depict a slice through the object in just one image by matching its refractive index to the one of the embedding liquid. This enables a direct sampling of the object geometry without the need for computational reconstruction. This way, a high-resolution 3D volume can be assembled simply by sweeping a laser plane through the object. We demonstrate the effectiveness of our light sheet range scanning approach on a set of objects manufactured from a variety of materials and material mixes, including dark, translucent and transparent objects.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning;

Keywords: 3D scanning, transparent surfaces, fluorescent dye

1 Motivation

The acquisition of 3-dimensional surfaces is a problem which occurs frequently in mechanical and optical engineering, computer vision and many other application fields. Today’s most precise methods involve active illumination by means of laser beams, lines or more sophisticated projection patterns. Usually, the light reflected or scattered by the surface is captured by an off-axis camera, so that the depth information can be recovered by triangulation.

The usability of most of these methods is limited by the object material, which defines the reflectance properties. An ideal surface for this class of 3D scanners scatters an incoming ray of light diffusely into all directions so that each impinging light ray results in a well-defined hit point visible from any viewing direction. However, many objects and materials exhibit a reflectance that is highly uncooperative with regard to range scanning. Complex effects such as transparency, subsurface light transport, specular reflectivity, and interreflections between surface points pose a challenge for the design of a general and robust system for shape acquisition.

Some of these effects have been addressed by modified scanning systems, more robust analysis of the scanning pattern [Curless and



Figure 1: *Fluorescent Immersion Range Scanning allows us to capture the 3D surface of challenging materials such as this crystal goblet. The object is placed into a tank filled with a fluorescent liquid which renders an incident light sheet visible. Inside the object, however, no fluorescence takes place. The surface contour is detected at the line where the intensity drops.*

Levoy 1995; Nayar et al. 2006], or e.g. by exploiting other physical properties of light such as polarization [Chen et al. 2007].

In this paper, we take a different approach: Instead of analyzing the intensity of the reflection we observe light rays as they propagate through space *before* they hit the surface. The intersection point with the surface is located where the ray is interrupted, see Figure 2. This holds for all kinds of surfaces except perfect mirrors.

In order to directly observe light rays, we have to embed the object into a participating medium. In fact, sparse sets of light rays have already been used as probes for optical density measurements, e.g. in the context of smoke scanning [Hawkins et al. 2005; Fuchs et al. 2007]. However, media such as smoke or dilute milk exhibit multiple scattering, resulting in a visible halo around the actual ray and a significant decrease in perceived contrast. This is due to the fact that scattering events in these media are *elastic*, i.e. upon scattering of a photon its wavelength is conserved. Therefore, photons can be scattered any number of times before they leave the volume or are absorbed.

Fluorescent media also scatter incoming light, but they interact *inelastically* with photons, absorbing them in a particular wavelength range and emitting photons of larger wavelength. Compared to elastic scattering, fluorescence has a number of advantages:

- The phase function is almost perfectly isotropic.
- Multiple scattering is negligible due to a small overlap of the absorption and emission spectra.
- The spectral separation of fluorescence and directly reflected light is simple using an optical long-pass filter.

Based on this principle of observing the attenuation along a ray in a fluorescent fluid rather than detecting the direct reflection of light from the surface, we present a 3D range scanning system that can detect the surface of objects composed of a wide range of materials. Our system enables the acquisition of dark, glossy, subsurface-scattering or even fully transparent objects (Figures 1 and 15).

Exploring the same principle, we propose a second setup that can perform direct volume slicing for transparent objects with a constant index of refraction (Section 6). After matching the refractive index of the fluid and the object, the light rays in the volume form straight lines. This matching has previously been used in optical tomography approaches [Sharpe et al. 2002; Trifonov et al. 2006].

*e-mail: {hullin,mfuchs,hpseidel,lensch}@mpi-inf.mpg.de

[†]e-mail: ivoihrke@cs.ubc.ca




Figure 2: (a) *Traditional laser range scanning:* A sheet of light is projected onto the object surface and its reflection is observed by a camera (not shown), (b) *Fluorescent immersion range scanning:* A sheet of light is projected onto an object immersed in a participating medium. The light sheet becomes visible and its cut-off is observed by a camera (not shown), enabling triangulation of surface points.

In contrast to tomographic reconstruction, we directly capture individual silhouette slices through the object as the laser sheet is only visible outside the object, in regions where it interacts with the fluorescent fluid. In Figure 18, we show the captured volumes of two intricate glass objects.

2 Related Work

In this article we investigate 3D range scanning of objects with non-cooperative materials, such as refractive surfaces, objects made of materials exhibiting a significant specular component, light absorbing materials such as diffusely reflecting, dark surfaces and mixtures thereof.

3D range scanning has been investigated for several decades, the majority of approaches assuming a diffuse reflectance of the object's surface. The literature on the subject is most commonly divided into active and passive techniques. Active light techniques include laser range scanning, coded structured light systems and time-of-flight scanners. An overview of the state of the art in active light range scanners can be found in [Blais 2004].

The further a surface deviates from the Lambertian reflectance assumption, the less accurate standard 3D range scanning techniques become [Curless and Levoy 1995; Beraldin 2004]. While coating objects with paint or removable powder is an option, clearly there are situations where this is not desirable. Several researchers have addressed the shortcomings of traditional range scanning techniques, extending the class of objects that can be successfully acquired by active light techniques. Curless and Levoy [1995] show that varying reflectance of the scanned surface, among other factors, results in systematic depth errors. They propose to analyze the raw laser range data in tilted space-time images of single scanlines to stabilize the detection process. Trucco and Fisher [1994] and Chen et al. [2007] employ polarization-based analysis to reliably remove specular reflections and subsurface scattering effects, respectively. Clark et al. [1997] investigate the use of a two-camera setup to disambiguate range measurements in the case of specular surfaces. Most laser range scanning systems assume a single peak of laser illumination per scanline of the imaging sensor. Park et al. [2004; 2008] show that multiple peaks frequently occur in the case of specular objects and propose methods to eliminate false measurements by applying local smoothness and global consistency constraints.

The other broad class of object acquisition techniques uses passive sensing. A recent survey of image based 3D acquisition, covering techniques such as structure from motion, stereo, multi-view stereo and shape from shading, is given in [Remondino and El-Hakim 2006]. Passive range sensing techniques have also been extended to account for non-lambertian materials, see e.g. [Jin et al. 2005; Davis et al. 2005; Zickler et al. 2002] for recent work. However, the accuracy of the acquired object geometry varies widely and is in general lower than that achievable with active light techniques.

As discussed before, specialized techniques have been developed to cover different classes of uncooperative materials. Usually these materials either exhibit a significantly non-lambertian reflectance resulting in secondary lighting effects such as interreflections or refraction. In the following, we review work related to refractive object acquisition. For a more detailed review we refer the reader to a recent survey article [Ihrke et al. 2008].

Interest was spawned by the photogrammetry community, with the goal of accurately measuring under-water objects [Höhle 1971]. The problem is posed as bundle adjustment in the presence of a single refracting surface bounding a medium, surrounding the object of interest. The refractive index is assumed to be known. The requirement of a known refractive index can be relaxed when the interface is planar [Maas 1995]. Reconstruction of the surface geometry is described in [Murase 1992]. The author analyzes a distorted pattern at the bottom of a water tank to estimate the shape of a water surface. Morris and Kutulakos [2005] use the same principle to deal with dynamic water surfaces exhibiting an unknown index of refraction. The reconstruction of single refracting surfaces is a special case of multi-interface reconstruction which has been theoretically analyzed by Kutulakos and Steger [2008]. The method is based on ray measurements using an LCD screen mounted onto a precision translation stage.

A different technique for reconstructing refractive objects or fluids is based on inverse ray-tracing. Assuming an object's surface at a certain position, an image is generated using a forward model of light transport. The surface is then optimized such that the synthetically generated image matches the photograph as closely as possible. One example of this technique is introduced by Ihrke et al. [2005] who reconstruct free-flowing water surfaces using measurements of the optical path length inside the water column. Another related approach uses polarized light in conjunction with polarization ray-tracing [Miyazaki and Ikeuchi 2005]. A different polarization-based technique is described by Saito et al. [1999]. The authors measure the polarization state of surface highlights and compute surface shape from this information.

Ben-Ezra and Nayar [2003] use a model-based approach to reconstruct simple shapes such as spheres, cylinders and lens-shapes based on the apparent movement of refracted feature points. Morris and Kutulakos [2007] describe an active light technique called scatter trace photography that can be used to acquire the geometry of objects with difficult material properties. The method is based on analysis of the materials' reflectance properties using a moving light source.

The acquisition of object geometry in the presence of participating media, as presented by Narasimhan et al. [2005], is another field of research that shares certain aspects with our work. Here, the intention is to counteract the degradation in image quality induced by the scattering environment. However, rather than causing objectionable effects, scattering media can also be used as a tool for vision and measurement purposes. For instance, methods for acquiring 3D density distributions within heterogeneous participating or fluorescent media have been presented before [Hawkins et al. 2005; Fuchs et al. 2007; Deusch and Dracos 2001]. In our case, we employ a homogeneous medium in order to ease the detection of the scanned object's surface.

Finally, 3D shapes can be acquired by means of computerized tomography, e.g. using X-rays that do not underlie refraction. Trifonov et al. [2006] transfer the CT methodology to optical wavelengths by immersing transparent objects in a liquid of matching refractive index. Our volume scanning approach is inspired by this work. However, we do not need to perform the numerical inversion of the projection operator.




Figure 3: Light sheet ($\lambda = 532 \text{ nm}$) projected into a glass cylinder filled with dilute milk (left) and fluorescent Eosin Y solution (right). Notice the color shift, and the contrast improvement on the lower edge of the light sheet due to the absence of multiple scattering.

3 Making Light Rays Visible

Our light sheet range scanning approach is based on the visualization of light rays propagating through space in front of the object surface. We achieve this visualization by employing a participating medium that, at every point in the volume, diverts a small portion of the incident illumination towards the camera. Possible candidates are scattering media such as smoke, dilute milk or fluorescent liquids, but fluorescent liquids are particularly well suited for our task of 3D scanning. In this section we compare the characteristics of dilute milk and fluorescent liquids using a test scene composed of different materials (Figures 5 and 6).

3.1 Elastic Scattering in Dilute Milk

Scattering in milk mainly occurs due to the presence of particles such as fat globules. They absorb, reflect and refract incoming photons without altering the wavelength. Light transport in such media is usually described by two parameters, an average phase function per scattering event and the mean free path of a photon in the medium. The single-scattering phase function is intrinsically given by the kind of milk being used (size and shape of particles). The mean free path depends mainly on the scattering cross-section and on the concentration of scatterers in the water, a parameter we can immediately control.

Thus, the concentration of the participating medium inside water determines the scattering properties of the solution, as shown in Figure 6 (top row). For low concentrations of, e.g., dilute milk, we observe light sheets that are mostly due to single scattering. However, the intensity of the scattered light is very weak—much weaker than the direct surface reflections—which brings us back to traditional range scanning. Higher concentrations of the scattering agent increase the sheet’s intensity, but also raise the probability of photons being scattered multiple times. As the concentration is increased further, multiple scattering starts dominating the image formation, making it difficult to locate the light sheet.

3.2 Inelastic Fluorescence Scattering

In the ideal case for our application, multiple scattering and contributions from the surface would be non-existing and single scatter-




Figure 4: Normalized absorption and emission spectra of Eosin Y and transmission curve of the longpass filter. Data obtained from [TU Graz; Thorlabs]




Figure 5: In order to compare elastic and inelastic scattering, we present a test scene composed of an opaque checkerboard pattern, a translucent dolphin and a mirror in a 25 l water tank filled with scattering media of varying concentration. The intersection lines between light sheet and surfaces are marked in red.




Figure 6: (top) Test scene in dilute milk. For very low concentrations, multiple scattering is rather unlikely. In this case, however, the light sheet is very dim and hardly visible in the presence of the much brighter surface features. (bottom) Scene in Eosin Y solution, captured through longpass filter and normalized for print.

ing would determine the image formation. As it turns out, we can get close to this case by using a fluorescent dye either inside water or inside the refractive index matching fluid.

The term fluorescence denotes a process where a photon is absorbed by an electronic system (atom, molecule, ion), which then emits a photon of different wavelength. The change of wavelength always occurs towards long waves (*Stokes shift*), as required by the conservation of energy. This results in a set of very desirable properties: Multiple scattering can almost completely be suppressed by using an excitation wavelength that is close to the long wavelength tail of the absorption spectrum (see Figure 4). Due to the loss of energy, a photon that takes part in one scattering event is unlikely to be scattered again, rendering the fluorescent fluid close to transparent for the emitted wavelength. Using an optical longpass filter, we can conveniently separate the fluorescence from surface reflections. Finally, as a side effect, the phase function is almost perfectly isotropic, since the excited state of the dye molecule typically has a lifetime in the order of a few nanoseconds, during which the molecule can rotate numerous times before it emits a photon into an uncorrelated direction.

3.3 Comparison

- *Choice of concentration.* The scattering behavior of the additive is crucial for the image quality. While it is difficult to find a suitable concentration of milk, we can vary the concentration of Eosin Y over a wide range without affecting the image contrast (Figure 6, bottom series).
- *Intensity.* For milk, the intensity that reaches the camera decreases with the distance to the scattering location. This is not the case for the fluorescent liquid, which is totally transparent for the fluorescence emission.
- *Contrast.* The same holds for the contrast (see also Figure 3).




Figure 7: Our measurement setup consists of a computer controlled laser projector, a CCD camera and a glass tank containing a fluorescent liquid and the object.

- *Surface reflections.* In the case of milk, reflections from the surface to the camera (including subsurface light transport) can become much brighter than the light sheet, see e.g. the dolphin in Figure 6 (top). Using a fluorescent dye, a longpass filter removes most of these contributions because of their shorter wavelength.
- *Directionality.* Dilute milk scatters light anisotropically with strong preference on forward and backward scattering. The phase function of fluorescence events is close to isotropic.

4 Measurement Setup

In the following we discuss our physical measurement setup, the choice of chemicals involved as well as the calibration of our range scanning system.

4.1 General Setup

The setup consists of a computer-controlled laser projector, a CCD camera and a glass tank, see Figure 7. The glass tank is filled with a liquid mixed with fluorescent dye in which we immerse the object to be measured. The fluorescent solution renders sheets of light, projected by the laser system, visible. For our experiments, we used different liquids mixed with the chemical Eosin Y, which exhibits fluorescent characteristics. The type of liquid used for surface scanning, Section 5, differs from the ones employed in volume scanning, Section 6. However, there are some common considerations.

The intensity of emitted light depends on the concentration of the fluorescent agent. A higher concentration increases the scattering coefficient σ_s of the solution, resulting in a higher fluorescent emission intensity. On the other hand, the absorption of light along the laser ray also increases. In order to minimize CCD noise, we find a trade-off that minimizes intensity fall-off along the illuminating rays while keeping exposure times at reasonable levels.

As a light source, we use a frequency doubled Nd:YAG laser ($P = 60\text{mW}$, $\lambda = 532\text{nm}$) and a computer controlled XY pair of galvanometer scanners. Our imaging sensor is a Canon EOS 5D equipped with a 100 mm f/2.8 macro lens. In order to capture only the fluorescence and filter out any reflections of laser light, we equip the lens with an optical longpass filter (Thorlabs FEL0550) with a cut-off wavelength of $\lambda_c = 550\text{nm}$.

The scanned object is fixed onto a precisely manufactured, manual turn-table positioned in the glass tank. By rotating and scanning the object, we acquire surface scans from different orientations, similar to traditional range scanning approaches. The single scans have to be merged in order to produce a full mesh.




Figure 8: Top view on three glass cylinders intersecting a laser plane within a fluorescent liquid (sketch). Three different cases occur: (1) $n_{\text{glass}} > n_{\text{liquid}}$, (2) $n_{\text{glass}} < n_{\text{liquid}}$, (3) good match ($n_{\text{glass}} \approx n_{\text{liquid}}$). The arrow indicates the incoming light direction.

Our two setups mainly differ in the arrangement of the laser projector with respect to the imaging sensor and in the choice of liquid that fills the glass tank.

4.1.1 Surface Scan Specifics

For fluorescent immersion surface scanning, Section 5, we employ plain water as a basis for the fluorescent solution. The concentration of Eosin Y is approximately $0.5\text{ mg} \cdot \text{l}^{-1}$. We position the camera such that its optical axis is approximately aligned with the normal of the bounding plane of the glass tank. This way, we minimize image aberrations and achieve a relatively homogenous sampling rate in the image plane. We position the laser at an angle of approximately $30^\circ - 45^\circ$ to the camera’s optical axis. As in traditional laser scanning there is a trade-off between triangulation accuracy and occlusion problems. Additionally, in our case, we want to avoid grazing angles of the laser plane w.r.t. the front face of the glass tank. Our setup for surface acquisition is shown in Figure 7.

4.1.2 Volume Scan Specifics

In the case of fluorescent immersion volume scanning for refractive objects with constant index of refraction, Section 6, we match the refractive index of the immersing liquid to that of the object before scanning it. This measure straightens the ray paths and enables us to directly observe volumetric slices of the object.

Refractive index matching liquids are available for a wide range of indices. However, they tend to get more and more poisonous as the refractive index increases. Trifonov et al. [2006] suggest the use of potassium thiocyanate to obtain a refractive index of around 1.5, which is sufficient for most light glasses. From a saturated solution (around 80 %) the refractive indices of borosilicate glass ($n = 1.48$) and quartz glass ($n = 1.46$) can easily be matched by gradually adding water until the refractive index of liquid and glass is approximately equal (see Figure 8).

Note, however, that the refractive index can only be matched perfectly for the single wavelength of the laser. Because of the broadband spectrum of the fluorescence light, dispersion may cause a slight disturbance of light rays on their way from the illuminated plane to the camera through glass structures (see Figure 17).

A drawback of potassium thiocyanate is that it weakens the fluorescence of Eosin significantly. We also experimented with glycerol, which, unlike potassium thiocyanate, is non-ionic and therefore induces less chemical stress on the dye molecules, resulting in a higher efficiency (more emission per absorption). However, due to the high viscosity of glycerol, heavy stirring is necessary to blend it with water, which in turn produces numerous small bubbles.

In addition to matching the refractive index of the immersing liquid to the object, we change the laser - camera setup to an approximate 90° setup. The laser projects sheets of light through one side of the glass tank, while the camera observes the object slices through the




Figure 9: An example of a laser sheet being projected onto the calibration target. The sheet of light intersects the planar target in a curve. In the background, part of the calibration pattern is visible. The dots have a distance of 10 mm center to center.

other. This is possible because occlusion is not going to occur for transparent objects immersed in a refractive index matched liquid. The objects are close to invisible to the naked eye due to negligible refraction.

4.2 Calibration

We calibrate our range scanning system using a two-step procedure. First, the viewing rays are calibrated using an image-based approach similar to [Trifonov et al. 2006]. A precisely manufactured, planar calibration target with known world coordinates is affixed to the turntable and recorded from different orientations. This step results in a vector of 3D world coordinates (x, y, z) per pixel (u, v) of our imaging sensor, defining mappings

$$T_\alpha : \mathbb{R}^2 \rightarrow \mathbb{R}^3$$

from image space to world space. α denotes the orientation of the calibration target and T_0 and T_{180} define the front and the back of our reconstruction volume, respectively. Transforming the pixel coordinates (u, v) to world coordinates using the transformations T_α , we obtain a set of 3D points (x, y, z) for each camera pixel. We compute the camera rays within the glass tank by fitting a three-dimensional line, Eq. 3, to these points.

A similar image-based technique can be used in the second step to calibrate the laser sheets inside the recording volume [Trucco and Fisher 1994]. However, since we project sheets of light into a refractive medium, single rays within the sheet form different incident angles with the medium. This results in a deformation of the sheets, which no longer form planes in space, see Figure 9.

In order to intersect camera rays with the curved light sheets and thus triangulate 3D points, we need an algebraic description of the curved laser illumination. To compute this description, we record a sparsely sampled set of laser sheets projected onto the planar calibration target, again in different orientations. The intersection points of the laser sheet with the calibration target in the image plane are computed as described in Section 5. We denote the set of sampled intersection points by $(u_\alpha^i, v_\alpha^i, t_\alpha)$ $i = 1 \dots N$, where N is the number of detected 2D points for laser sheet t_α with the calibration target rotated into position α . We then compute the positions of these sample points in world coordinates by applying the appropriate mapping T_α :

$$(x, y, z, t) = [T_\alpha(u_\alpha, v_\alpha); t_\alpha]. \quad (1)$$

For better readability we drop the sample number i from the equation. Applying Eq. (1) we obtain a number of world points (x, y, z, t) that were hit by a particular laser sheet t . We use these sampled points to interpolate this information in the whole reconstruction volume. We model the behavior of the light sheets inside the glass tank as a quadratic polynomial in y and t , while keeping the dependency on x linear:

$$\begin{aligned} z &= f(x, y, t) = a_0y^2x + a_1y^2 + a_2yx + a_3y + a_4x + a_5, \quad (2) \\ a_i &= b_{i0}t^2 + b_{i1}t + b_{i2}. \end{aligned}$$

Eq. (2) is linear in its 18 coefficients b_{00} to b_{52} . Thus, we obtain one linear equation per sample point (x, y, z, t) and can solve for the polynomial's coefficients in a least squares sense. The choice to keep f linear in its argument x is motivated by the fact that in our setup the laser sheet forms a constant angle with the boundary of the glass tank for this variable. Therefore, the curvature of the laser sheet will not be affected by changes in the x coordinate.

4.3 Triangulation

The task of triangulation is to determine a 3D point from pixel positions (u, v) on the image sensor and a laser sheet t . The first step of the calibration procedure yields a ray for each image pixel:

$$\mathbf{r}(u, v) = \mathbf{p}(u, v) + s\mathbf{d}(u, v). \quad (3)$$

Rearranging Eq. (2) into an implicit form

$$f(x, y, t) - z = 0, \quad (4)$$

and combining it with Eq. (3) we obtain a cubic polynomial in s . Solving it results in three solutions, which we insert into the ray equation, Eq. (3). We check the resulting points against our reconstruction volume and discard the ones that fall outside of it.

4.4 Data Acquisition

The scanning process involves the sweeping of a vertical laser sheet through the acquisition volume. We start the scanning process with the laser sheet $t = 0$ closest to the camera and proceed towards the back of the glass tank, i.e. with increasing t the distance between the laser sheet and the camera increases as well. For each projected laser sheet we take an image with the imaging sensor. The output of this system is a stack of images $I_t(u, v)$, which we analyze to determine the intersection points between the object and the laser sheet.

In the following, we describe a robust surface detection algorithm, Section 5, and a volume slicing method for refractive objects, Section 6.

5 3D Surface Scanning

Using the setup described in Section 4.1.1 we first present a system for performing a surface range scan which operates on surfaces that are otherwise challenging for traditional 3D scanners, namely transparent, translucent or dark materials.

Due to the use of fluorescent Eosin Y the laser sheet is visible to the camera as long as it travels through un-occluded space. As soon as the incident light interacts with the surface, most of its energy is absorbed by or transmitted through the object. Since the object itself is supposed not to be fluorescent, we observe a significant fall-off in the measured intensity along the laser sheet exactly at the intersection point with the object's surface as seen in Figure 10. Our surface detection will therefore concentrate on detecting a negative change in intensity rather than detecting a peak as in traditional range scanning [Curless and Levoy 1995].

5.1 Influence of the Material

This drop in intensity can be observed regardless of the surface properties, as long as the material is not perfectly mirroring. It is, however, superposed by a number of additional effects caused by global light transport. Examples of such effects are specular or diffuse reflections of the light sheet on the object surface, as well as volume caustics of refracted or reflected light rays (see e.g. Figure 1 (middle)). In Section 5.2, we propose a method to deal with these effects.




Figure 10: Laser sheet hitting various surfaces. The intensity along the laser line drops significantly when interacting with the surface. The intersection line is clearly visible for transparent (Figure 1), dark (a), translucent (b), and partially diffuse (c) surfaces. Note that for the first three cases the intensity does not increase right at the intersection line. The die in (c) is transparent but slightly fluorescent itself, resulting in a bright profile.

Furthermore, diffuse surfaces of high albedo exhibit a bright line caused by volumetric scattering in front of the surface. This causes a small systematic error, as discussed in Section 7.2.3.




Figure 11: (a) Space-time slice for one scan line through the crystal goblet. (b) Shows the detected surface in red which is nicely aligned with the local minimum of the gradient component g_t (c). At the surface, the gradient component g_u (d) also has a minimum, but at other places it is more corrupted by stationary features.

5.2 Detection Algorithm

In order to detect the surface, we analyze the captured images $I_t(u, v)$ in a space-time stack, and process them initially per camera scanline. For each scanline v , we obtain a slice $S_v(u, t)$ through the space-time volume as depicted in Figure 11.

It is instructive to look at the gradient g with components (g_u, g_v, g_t) of the recorded intensity of one pixel in the space time volume. $g_t(u, t)$ and $g_u(u, t)$ are shown in Figure 11(c) and (d). We compute the gradients on a median filtered space-time volume (kernel size $3 \times 3 \times 3$). For any surface we can state the following properties:

1. The component g_u has a negative peak.
(We are looking for an intensity drop.)
2. The component g_t has a negative peak.
(It has to move if the laser sheet is moved.)
3. The 2D gradient (g_u, g_v) must not be perpendicular to the laser direction \vec{l} in the image.
(Mistrust features that are aligned with laser rays.)
4. (optional) The detected surface line in the slice must be strictly monotonic.

As can be seen in Figure 11(d), the component g_u can be corrupted by caustics and reflections that intersect with the actual edge in image space, resulting in the vertical lines in the gradient image. This is not the case in the g_t component, where the wanted surface ap-



Figure 12: Scan maps generated from space-time stack. From left to right: depth map $t(u, v)$ before and after segment size filtering, world coordinate map (x, y, z) . Colors are normalized for print.

pears much cleaner. We therefore determine all local minima of g_t that are below a threshold θ_t with sub-pixel precision, reporting a fractional t value for each camera pixel. Blurring the gradient images before the thresholding allows to trade resolution against smoother results.

In the next step, we discard all candidate points who do not meet the first and third criteria with the corresponding thresholds $g_u \leq \theta_u$ and $\angle(\vec{l}, (g_u, g_v)) \leq \theta_\alpha$. In order to close some occasional gaps, we soften this selection process by also keeping points which have an 8-neighbor that meets the criteria.

Finally, we traverse the slice from bottom to top, and collect all first surface points encountered. We require the surface to be monotonic, which is necessary for transparent objects. Otherwise it might happen that distorted features behind the first surface are detected as well. For nontransparent materials, however, the monotonicity constraint can be dropped, resulting in a higher yield of valid surface points.

After performing this routine for each space-time slice, we obtain a map containing a t value per camera pixel (Figure 12). It might still contain a number of patches that correspond to volume caustics due to reflection or refraction. Since those unwanted surfaces are typically small in size, we can eliminate most of them by segmenting the depth map estimate into continuous patches and discarding or filling up segments that contain less than n_{\min} pixels.

The absolute threshold, denoising and segment size filter parameters have to be found to meet demands of the respective measurement setup (camera resolution, measurement geometry, light sheet brightness, etc.). After downscaling all camera images by 50% (corresponding to approx. 3 megapixels), the following parameters performed well for our setup: $\theta_\alpha = 85^\circ$, $n_{\min} = 1000$. The optimal choice of the thresholds θ_t and θ_u depends on the observed contrast in the space-time images. We determined them experimentally for our setup and used the same set of parameters for all surface scans in this paper.

The filtered depth map $t(u, v)$ is converted into a 3D surface mesh using the camera calibration as described in Section 4.3.

6 Single-scan Direct Volume Capturing

In this section we describe a method for capturing *transparent* objects volumetrically, without the need for tomographic reconstruction. We present a technique that directly records one slice through the volume for every laser sheet.




Figure 13: Capturing the mouse model (cf. Figure 18): Two slices of the volume stack after matching the refractive index. The contour of the mouse is clearly visible in the surrounding fluorescent liquid.

6.1 Matching the Refractive Index

So far, our scanning was limited to the first surface, where the light is being refracted and no longer propagates in a straight line (see Figure 8). We change our setup to the one described in Section 4.1.2. By matching the refractive index of the liquid to that of the object the ray direction is preserved even at interfaces between liquid and solid, so that all intersections between a light ray and the object occur along a straight line (cf. Trifonov et al. [2006]).

Using a fluorescent liquid, we can now observe entire slices of clear objects at once when illuminating the tank with a laser sheet. Since the index of refraction is matched, the viewing rays will not bend at the surface. The object overall, and in particular in front or behind the laser plane, therefore is not visible at all. Since the object itself is not fluorescent, only the surrounding liquid lights up while the object remains black (see Figure 13). The object’s surface is located at the boundary between light and dark image regions. In order to capture the entire volume we sweep the laser plane through the entire volume once.

6.2 Volume Processing

Our volume acquisition pipeline is organized as follows:

1. Recording of a space-time stack by sweeping a light plane through the volume
2. Binary segmentation
3. Filtering in spatial and space-time domain (see Figure 14) to reduce noise and artifacts
4. Resampling of the filtered stack into world coordinates (x, y, z)
5. Extraction of the isosurface using Marching Cubes [Lorenson and Cline 1987]

Given an input stack of images, (for exemplary slices see Figure 13), we perform a binary segmentation to separate the object from the background. A stable separation is obtained by using a median filtered reference image in which the laser sheet does not hit the object.

The segmented slices might still contain a set of misclassified pixels which are mostly filtered out in the pipeline illustrated in Figure 14. The following effects have to be accounted for:

Noise. Particles in the liquid, bubbles, and other effects lead to the presence of a "salt and pepper"-type noise which forms connected regions in the order of a few pixels. They can be easily removed by inverting all regions that are smaller than 200 pixels (Figure 14 (b)).

Needle-shaped shadows. Inclusions in the glass, such as bubbles of air or particles, or regions with different refraction index will cast shadows. Similarly, a slight mismatch in the refractive index can cause such shadows, as depicted in Figure 8. We propose two different approaches to detect and to remove these kinds of artifacts.

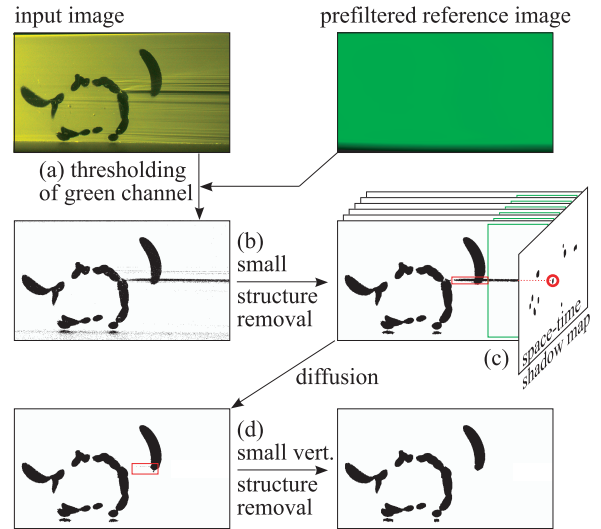


Figure 14: The volume filtering pipeline: A binary segmentation is performed using a simple thresholding at 50% of a median filtered reference light sheet (a). As a first denoising step, all four-connected image regions smaller than 200 pixels are discarded (b). Shadows cast by inclusions are detected as they penetrate the space to the right of the object (green box) which should be un-occluded (c). After tracing them back towards the light source, the traversed space-time volume segments are filled by 3D diffusion. Remaining horizontal features of insignificant height are removed in the last filtering step (d).

Most of these artifacts extend to the far right side of each volume slice. Any feature that is detected in the region outside the object’s bounding box is treated as a shadow region. We project all of these regions onto a shadow map. Then we trace each shadow region back towards the laser, until its entire cross-section is lit. All pixels on this trace are marked as invalid.

After marking all possible shadow pixels they are filled using a simple volumetric diffusion in the space-time stack of images, i.e. even across different laser sheets. The process is illustrated in Figure 14 (c) and (d).

Some small horizontal lines that do not extend to the outside are missed by this approach. We simply remove all horizontal features that are less than five pixels in height which is very small compared to any part of the real object (see Figure 14 (d)).

The filtering so far was performed on a space-time stack of images, which, as the final step, is resampled into world coordinates using the calibration data.

7 Results

To demonstrate the versatility of our approach we have acquired the surface shape of several challenging objects. Figure 15 illustrates the performance of our prototype fluorescent immersion range scanner compared to traditional scans acquired with a commercial laser range scanner (Minolta Vi910). The figure shows a scan from a single direction for both acquisition systems. Please observe, that the data produced by our system is essentially unfiltered, except for the simple image processing operations explained in Sections 5 and 6.

The horse model (top row) consists of two different materials. The first material is a black, rough plastic, while the second one has

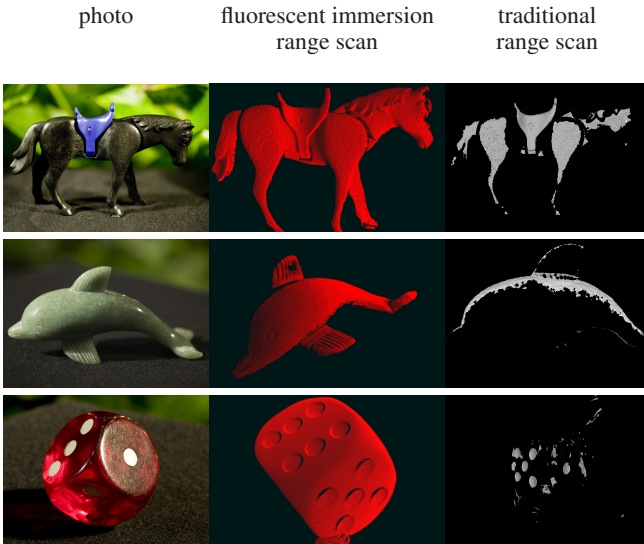


Figure 15: By immersing objects into a fluorescent liquid we have acquired high quality range scans for a set of surfaces that are difficult to acquire with traditional methods.

a higher albedo and is smooth. Both materials are glossy, showing significant surface highlights. The Minolta scanner captures the higher albedo material quite well but has problems with the dark material. In comparison, our approach captures the details of the surface very well and even captures the surface roughness of the dark plastic. The dolphin (middle row) consists of a translucent stone and is an example for a sub-surface scattering material with some crystal structure beneath the surface. Again, the commercial scanner fails to provide a dense depth map, whereas our result reveals significant detail. In this example, however, we observe small holes in the fins of the dolphin which cannot be attributed to occlusion alone. The reason is an overly aggressive thresholding during the construction of the depth map where pixels with a low gradient component g_u are removed. The lines on the fins are thus classified as stationary features. Our third example (bottom row) is a transparent die with diffuse, white dots. Actually, the object’s transparent part is fluorescent as well, but the material’s extinction coefficient happens to be high enough that all fluorescence emission is limited to the surface. The Minolta scanner can only capture the diffuse, white part of the surface, while our method produces a dense depth map for both material types. However, due to the fluorescence of the material, a conventional laser scanner equipped with a green laser would probably be able to capture the full surface of this object. The crystal goblet, Figure 1, is completely transparent and exhibits significant refraction due to its faceted surface. Nevertheless, the depth map recovered by our method (Figure 12) has only a few minor holes and is virtually free of noise. Using only 6 scans from different directions, almost the entire outer surface of the goblet can be reconstructed with high accuracy. By simple superposition of the individual calibrated scans and without the need for advanced mesh alignment methods, we achieve a low noise level in the order of 0.05% of the overall object size (Figure 16).

The volume slicing technique presented in Section 6 is geared towards capturing transparent objects. We demonstrate its effectiveness on two objects. The first is a bowl of quartz glass ($n = 1.46$) with a set of quartz glass cylinders of 10 mm diameter. Our other glass object is a mouse figurine, made from borosilicate glass ($n = 1.48$). Figure 18 shows the results of the two volume scans. Both objects were scanned with a potassium thiocyanate solution as refractive index matching liquid.




Figure 16: For a quantitative analysis of our first-surface scanning method, we took a cross-section of the goblet model (red line in (a)) and fitted an ellipse against it. Given 3000 sample points (b), we obtained an RMS error of 0.046 mm at an object diameter of 43.2 mm.




Figure 17: Slices of the glass cylinders dataset. Notice the difference in edge quality depending on the distance of the scanning laser sheet to the camera: near (left) and far (right).

The bowl is a challenging object because of the numerous overlapping, lens-like cylinders that focus the light, requiring very precise refractive index matching. The size of the acquired volume model is $1100 \times 720 \times 1400$ voxels, the real world size of the object being 92 mm. This results in a voxel resolution of about 0.08 mm in each dimension. However, the overall resolution achieved in practice is a little lower. As can be seen in Figure 18 (c), in the back of the volume (from the camera’s viewpoint) some artifacts appear. They result from the differing wavelength of the laser illumination and the light emitted by Eosin Y. As discussed in Section 4.1.2, the refractive index matching is only accurate for one particular wavelength. We have to match the refractive indices for the wavelength of the illuminating laser light to ensure planar light sheets while scanning the volume, and to suppress edge shadows (see Section 6). This, however, results in a slight mismatch of the refractive index for the light emitted via fluorescence, giving rise to minor artifacts in the back of the scanned volume. A comparison of slice images from different positions in the scanning volume is shown in Figure 17, illustrating the effect. The gap that is visible in the bowl in Figure 18(d) is caused by the fact that this part of the bowl was outside the calibrated volume. We missed it because the bowl is not visible in the refractive index matched liquid without the illumination of the laser sheet.

The second object, a hand-made mouse, has intricate detail and a complicated surface topology. Its length is approximately 40 mm and the acquired volume has a size of $500 \times 300 \times 500$ voxels, resulting in a similar effective resolution as in the previous example.

The filtered volumes can be visualized with traditional volume rendering techniques producing the x-ray images in Figure 18(b). After downsampling the data by a factor of 0.5 in all dimensions, we extract the isosurface to generate surface models (c) which we render using a glass shader in a synthetic environment in (d). For quantitative analysis, we extracted two of the glass cylinders, one close to the camera and the other further away, and fitted analytical models against them. Given 11500 samples on each cylinder, we obtained a RMS error of 0.037 mm for the near and 0.285 mm for the far cylinder.

7.1 Comparison to Existing Methods

While there have been many previous methods specializing on individual types of materials, e.g. for glossy [Chen et al. 2006; Park and Kak 2008], (inhomogeneous) transparent [Morris and Kutulakos 2007] or translucent scenes [Chen et al. 2007], only Morris and Kutulakos [2007] attempted the acquisition of objects with mixtures of different materials. The resolution obtained by our surface scanning method is on par with existing standard techniques for diffuse objects.

Comparing our volume scanning procedure with existing methods such as CT or MRI is problematic, since our algorithm is tailored to recovering a binary density volume. However, given the comparatively small acquisition and reconstruction effort of our method and the high resolution it provides, we believe that for objects made of clear, homogeneous materials our method is an attractive alternative to existing techniques.

7.2 Applicability and Limitations

7.2.1 Immersing Objects in Liquid

For most objects composed of plastic, glass, ceramics, stone, or metals the immersion in a liquid is uncritical. Care has to be taken with porous materials and substances that are soluble in water or prone to corrosion. On the other hand, the same holds for covering the surface with diffuse spray paint which is often alcohol-based and, as a side effect, modifies the surface geometry. Depending on the requirements of the object, the solvent for the Eosin molecules can be chosen from a wide range of fluids with different chemical properties, e.g. (deaerated) water, ethanol, glycerol, other alcohols or nonpolar oils, some of which are used by archaeologists to preserve organic objects.

7.2.2 Mirroring Materials

Strongly mirroring materials are still difficult to acquire. In the case of curved mirroring surfaces, strong caustics may appear in front of the surface, masking out any intersection line behind them. Given a perfect planar mirror, our method might fail because the change of intensity at the surface vanishes. These problems can probably be resolved using more sophisticated light patterns and/or multiple cameras.

7.2.3 High-Albedo Materials

For diffuse materials with high albedo (such as the calibration target in Figure 9), two effects become apparent. First, the surface is brightened up by the laser sheet in front of it (compare e.g. the brightness of the white surface to the one of the holes). Since we are dealing with diffuse reflection, this is a stationary, low-frequency effect which does not affect the detection routine. Even more prominent is the bright line at the intersection between laser sheet and surface, caused by volumetric scattering of the diffusely reflected light sheet. Under the assumption of isotropic scattering it can be shown that this effect can never become brighter than the light sheet itself. Nevertheless, its superposition with the incident light sheet leads to a small shift of the detected edge towards larger distances (by max. half the ray diameter). So far, we have not found a way to circumvent this small systematic error while still keeping the detection method as general as it currently is. For the calibration, however, using a diffuse red calibration target instead of the white one would be a perfect workaround.

8 Conclusion

Fluorescent immersion 3D scanning offers novel means for capturing the geometry of objects that due to their material are difficult to

capture with traditional 3D range scanning techniques. Examples of such surfaces are dark, translucent or even transparent surfaces. The key idea of both the proposed surface and volume scan techniques is to place the object into a tank filled with a fluorescent fluid. Rather than detecting peak intensities as in structured light scanning, the object's surface will appear as a drop of intensity. This detection method is quite robust to the surface reflection properties of the object and produces high quality surfaces and volumes with rather simple acquisition effort.

Using the fluorescent fluid, light rays can be made visible without the disturbing effects of multiple scattering. Visualizing light rays this way might inspire other novel acquisition methods, for example for visualizing and capturing reflection properties of materials.

Acknowledgements

We would like to thank our reviewers for their suggestions and comments. Also, our thanks go to Boris Ajdin, Tongbo Chen, Miguel Granados, Conny Liegl and Thomas Schultz for their assistance with data processing and the video, as well as Dr. Christel Weins, Axel Köppel and Michael Laise for the technical support. This work has been partially funded by the DFG Emmy Noether fellowship (Le 1341/1-1) and the Max Planck Center for Visual Computing and Communication (BMBF-FKZ01IMC01). Ivo Ihrke was supported by a Feodor Lynen fellowship of the Humboldt Foundation.

References

- BEN-EZRA, M., AND NAYAR, S. 2003. What Does Motion Reveal About Transparency? In *Proc. of ICCV'03*, vol. 2, 1025–1032.
- BERALDIN, J.-A. 2004. Integration of Laser Scanning and Close-Range Photogrammetry - the Last Decade and Beyond. In *Proceedings of the XXth ISPRS Congress*, 972–983.
- BLAIS, F. 2004. Review of 20 Years of Range Sensor Development. *Journal of Electronic Imaging* 13, 1, 231–243.
- CHEN, T., GOESELE, M., AND SEIDEL, H.-P. 2006. Mesostructure from Specularity. In *Proc. of CVPR '06*, 17–22.
- CHEN, T., LENSCH, H. P. A., FUCHS, C., AND SEIDEL, H.-P. 2007. Polarization and Phase-Shifting for 3D Scanning of Translucent Objects. In *Proc. of CVPR '07*, 1–8.
- CLARK, J., TRUCCO, E., AND WOLFF, L. B. 1997. Using Light Polarization in Laser Scanning. *Image and Vision Computing* 15, 1, 107–117.
- CURLESS, B., AND LEVOY, M. 1995. Better Optical Triangulation Through Spacetime Analysis. In *Proc. of ICCV'95*, 987–994.
- DAVIS, J., YANG, R., AND WANG, L. 2005. BRDF Invariant Stereo using Light Transport Constancy. In *Proc. of ICCV'05*, 436–443.
- DEUSCH, S., AND DRACOS, T. 2001. Time resolved 3d passive scalar concentration-field imaging by laser induced fluorescence (LIF) in moving liquids. *Meas. Sci. Technol.*, 12, 188–200.
- FUCHS, C., CHEN, T., GOESELE, M., THEISEL, H., AND SEIDEL, H.-P. 2007. Density Estimation for Dynamic Volumes. *Computers & Graphics* 31, 2 (Apr.), 205–211.
- HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. Acquisition of Time-Varying Participating Media. In *Proc. of ACM SIGGRAPH 2005*, ACM, 812–815.

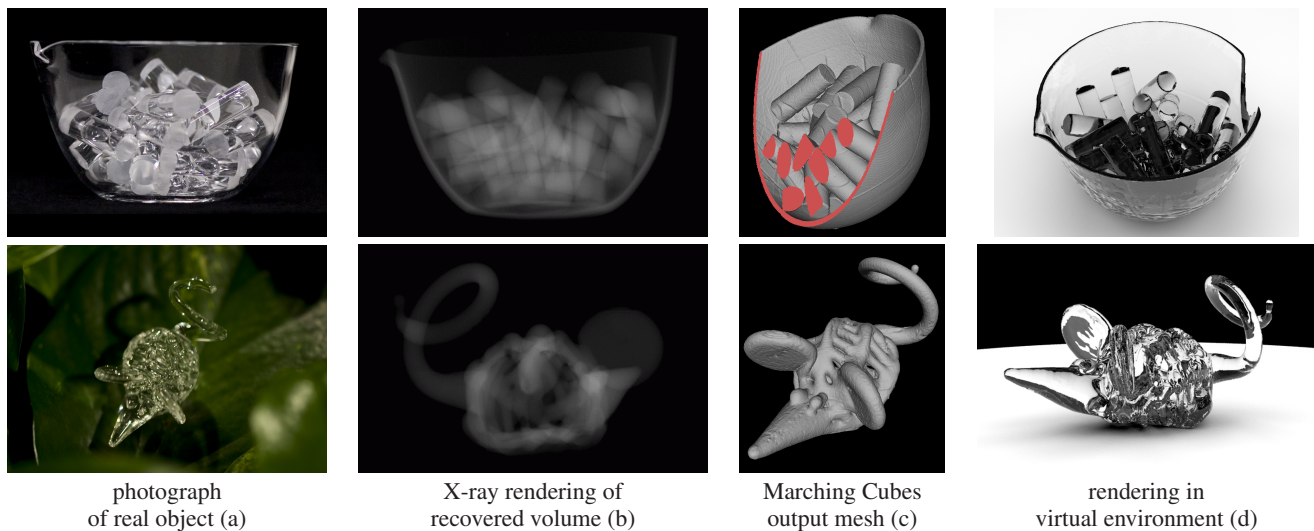


Figure 18: The volume acquisition method enables the acquisition of an entire volume for transparent objects with a constant refractive index. The scan can be performed in one sweep with negligible post-processing. We present volume renderings of the acquired data sets (b), extract the isosurface (c) and render the surface model in a synthetic environment (d).

- HÖHLE, J. 1971. Reconstruction of the Underwater Object. *Photogrammetric Engineering* 37, 948–954.
- IHRKE, I., GOLDLUECKE, B., AND MAGNOR, M. 2005. Reconstructing the Geometry of Flowing Water. In *Proc. of ICCV'05*, 1055–1060.
- IHRKE, I., KUTULAKOS, K. N., LENSCH, H. P. A., MAGNOR, M., AND HEIDRICH, W. 2008. State of the Art in Transparent and Specular Object Reconstruction. In *STAR Proceedings of Eurographics*, 87–108.
- JIN, H., SOATTO, S., AND YEZZI, A. J. 2005. Multi-View Stereo Reconstruction of Dense Shape and Complex Appearance. *International Journal of Computer Vision* 63, 3 (Jul), 175–189.
- KUTULAKOS, K. N., AND STEGER, E. 2008. A Theory of Refractive and Specular 3D Shape by Light-Path Triangulation. *International Journal of Computer Vision (IJCV)* 76, 1, 13–29.
- LORENSEN, W. E., AND CLINE, H. E. 1987. Marching cubes: A high resolution 3D surface construction algorithm. In *Proc. of ACM SIGGRAPH 87*, 163–169.
- MAAS, H.-G. 1995. New Developments in Multimedia Photogrammetry. In *Optical 3D Measurement Techniques III*, A. Grün and H. Kahmen, Eds. Wichmann Verlag.
- MIYAZAKI, D., AND IKEUCHI, K. 2005. Inverse Polarization Raytracing: Estimating Surface Shapes of Transparent Objects. In *Proc. of CVPR '05*, vol. 2, 910–917.
- MORRIS, N. J. W., AND KUTULAKOS, K. N. 2005. Dynamic Refraction Stereo. In *Proc. of ICCV'05*, 1573–1580.
- MORRIS, N. J. W., AND KUTULAKOS, K. N. 2007. Reconstructing the Surface of Inhomogeneous Transparent Scenes by Scatter-Trace Photography. In *Proc. of ICCV'07*, 1–8.
- MURASE, H. 1992. Surface Shape Reconstruction of a Non-rigid Transparent Object Using Refraction and Motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 10 (October), 1045–1052.
- NARASIMHAN, S. G., NAYAR, S. K., SUN, B., AND KOPPAL, S. J. 2005. Structured Light in Scattering Media. *Proc. of ICCV'05 I*, 420–427.
- NAYAR, S. K., KRISHNAN, G., GROSSBERG, M. D., AND RASKAR, R. 2006. Fast Separation of Direct and Global Components of a Scene Using High Frequency Illumination. In *Proc. of ACM SIGGRAPH 2006*, 935–944.
- PARK, J., AND KAK, A. C. 2004. Specularity Elimination in Range Sensing for Accurate 3D Modeling of Specular Objects. In *Proceedings of 3DPVT'04*, 707–714.
- PARK, J., AND KAK, C. 2008. 3D Modeling of Optically Challenging Objects. *IEEE Trans. on Visualization and Computer Graphics* 14, 2 (March/April), 246–262.
- REMONDINO, F., AND EL-HAKIM, S. 2006. Image Based 3D Modeling: A Review. *The Photogrammetric Record* 21, 115, 269–291.
- SAITO, M., SATO, Y., IKEUCHI, K., AND KASHIWAGI, H. 1999. Measurement of Surface Orientations of Transparent Objects using Polarization in Highlight. In *Proc. of CVPR '99*, vol. 1, 381–386.
- SHARPE, J., AHLGREN, U., PERRY, P., HILL, B., ROSS, A., HECKSHER-SORENSEN, J., BALDOCK, R., AND DAVIDSON, D. 2002. Optical Projection Tomography as a Tool for 3D Microscopy and Gene Expression Studies. *Science* 296, 19, 541–545.
- THORLABS, INC. Transmission curve of FEL0550 longpass filter. <http://www.thorlabs.com/Thorcat/7600/7672-S01.pdf>.
- TRIFONOV, B., BRADLEY, D., AND HEIDRICH, W. 2006. Tomographic Reconstruction of Transparent Objects. In *Proc. of EGSR'06*, 51–60.
- TRUCCO, E., AND FISHER, R. B. 1994. Acquisition of Consistent Range Data Using Local Calibration. In *IEEE International Conference on Robotics and Automation*, 3410–3415.
- TU GRAZ, INSTITUTE OF ANALYTICAL CHEMISTRY. Database of fluorescent dyes. <http://www.fluorophores.org>.
- ZICKLER, T., BELHUMEUR, P. N., AND KRIEGMAN, D. J. 2002. Helmholtz Stereopsis: Exploiting Reciprocity for Surface Reconstruction. *International Journal of Computer Vision (IJCV)* 49, 2-3, 215–227.