

DISCO – Acquisition of Translucent Objects

Michael Goesele

Hendrik P. A. Lensch

Jochen Lang

Christian Fuchs

Hans-Peter Seidel

MPI Informatik*

Abstract

Translucent objects are characterized by diffuse light scattering beneath the object’s surface. Light enters and leaves an object at possibly distinct surface locations. This paper presents the first method to acquire this transport behavior for arbitrary inhomogeneous objects. Individual surface points are illuminated in our DISCO measurement facility and the object’s impulse response is recorded with a high-dynamic range video camera. The acquired data is resampled into a hierarchical model of the object’s light scattering properties. Missing values are consistently interpolated resulting in measurement-based, complete and accurate representations of real translucent objects which can be rendered with various algorithms.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and Scanning, Viewing Algorithms, I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically Based Modeling

Keywords: Subsurface Scattering, BSSRDF, Translucency, Acquisition, Reflection Model

1 Introduction

Light interacts with ideal opaque objects purely locally – it is scattered at the point of incidence according to the bidirectional reflectance distribution function (BRDF) [Nicodemus et al. 1977]. In contrast, light passes straight through transparent objects (e.g., glass) and is only reflected or refracted at material boundaries. Many daily life objects (e.g., milk, skin or marble) are translucent and belong to neither of these categories. They are characterized by multiple light scattering inside the object. This *subsurface scattering* behavior leads to a very distinct appearance which is visually important: light shines through objects, they appear smooth and surface details are hidden (see Figure 2 for an example of the same object with and without subsurface scattering).

Translucent objects can be rendered using a variety of physical simulation techniques. Most recent rendering systems are based on the dipole approximation to a diffusion model [Jensen et al. 2001]. This approach enables interactive evaluation while providing good visual quality. Furthermore, physically correct parameters can be determined for homogeneous materials by a point-based measurement setup [Jensen et al. 2001]. To our knowledge, there exists however no technique that can derive the necessary input data for a real object with spatially varying properties. Several image-based acquisition systems are able to capture objects that exhibit subsurface scattering. But, they do not take the specific properties of translucent objects into account and are unable to provide a comprehensive model of such an object. They are especially not able



Figure 1: A model of an alabaster horse sculpture acquired by our DISCO method. Differences in material are clearly visible when the model is lit from behind. Figure 2 shows a photograph of the sculpture and Section 9 discusses the result in detail.

to render translucent objects from arbitrary viewpoints and under arbitrary lighting conditions where the incident illumination varies per surface point (e.g., illumination with textured point lights, or projected shadow borders).

We therefore propose a method to acquire a comprehensive model of a translucent object with diffuse light propagation inside the object and local illumination variation due to shadows or directionally dependent light sources. To this end we sequentially illuminate a dense set of locations on an object’s surface and observe its impulse response function with a high-dynamic range video camera. The acquired data is resampled into a hierarchical data structure taking advantage of the specific properties of translucent material such as the exponential fall-off near the point of incidence and the smooth global response. We introduce methods to consistently interpolate holes (mainly caused by occlusion) in the local and global data and to reduce the variance due to noise in the acquisition process. We implemented the rendering method of Lensch et al. [2003b] that allows us to display the acquired objects interactively. Finally, we discuss how the acquired data could be incorporated into many current rendering algorithms as a model for real objects exhibiting subsurface scattering.

Our main contribution is our DISCO method (Digital Imaging of Subsurface sCattering Objects) which acquires a general model of real translucent objects. The model can be rendered under arbitrary viewing and illumination conditions including local illumination. We describe our complete pipeline including acquisition, post-processing and display and validate our approach using several objects with diverse properties.

The remainder of this paper is structured as follows: After discussing previous work (Section 2) we summarize the theory of subsurface scattering, briefly describe our measurement approach, and state our modeling assumptions in Section 3. We describe details of our acquisition in Section 4. The resampling and postprocessing steps are presented next, followed by a discussion of applicable rendering methods in Section 8. Results of our measurement-based modeling are presented in Section 9 before we conclude and discuss future work.

*email: {goesele,lensch,lang,cfuchs,hpseidel@mpi-sb.mpg.de}



Figure 2: Photographs of an alabaster horse model with and without subsurface scattering under identical illumination conditions. The object was covered with fine dust to “turn off” subsurface scattering. Much surface detail is hidden by the translucency.

2 Previous Work

The theory of radiation in scattering media is a well studied problem [Ishimaru 1978] and is of interest in many applications (e.g., wave propagation for communication, medicine, remote sensing). The goal in medicine and remote sensing is to draw conclusions about the interior of a region by analyzing the scattered response to incoming radiation which requires solving a difficult inverse problem. The predominant goal in computer graphics is to acquire enough data to achieve a realistic visual model of an object’s response to incident illumination. While knowledge of the interior structure of an object is required to perform exact simulations of the physical behavior, it is neither mandatory nor necessarily efficient. Our approach is entirely based on external measurements and depends not on the inversion.

2.1 Models for Translucent Objects

Translucent objects can be modeled by defining their basic physical properties (e.g., the absorption and scattering cross sections σ_a and σ_s [Ishimaru 1978]) for each point inside their volume. Rendering such an object could then make use of physical simulations or appropriate approximations.

Alternatively, it is sufficient to record the visible effects of subsurface scattering. Translucent objects can be modeled using a bidirectional scattering-surface reflectance distribution function $S(x_i, \omega_i; x_o, \omega_o)$ (BSSRDF) [Nicomemus et al. 1977] that relates irradiance at a surface position x_i to the reflected radiance at x_o for arbitrary incoming and outgoing directions ω_i and ω_o . The complexity of this 8D function makes handling it quite cumbersome. However, for the case of optically dense material, the directional dependency is negligible since the response is dominated by multiple scattered particles. The BSSRDF model of Jensen et al. [2001] is therefore split into a directionally dependent single scattering term $S^{(1)}$ and a directionally independent multiple scattering term S_d . Omitting the single scattering term, the BSSRDF of a translucent object can be collapsed to the 4D diffuse subsurface reflectance function $R_d(x_i, x_o)$ that depends only on the incoming and outgoing surface positions.

2.2 Acquisition Techniques

BRDF acquisition techniques [Marschner et al. 1999; Lensch et al. 2003a] can only recover the local reflectance properties but are unable to model the distant light transport characteristic for subsurface scattering [Jensen et al. 2001]. Many image-based acquisition and rendering techniques are able to record and reproduce some aspects of translucent objects: Techniques based on light fields [Levoy and

Hanrahan 1996] or lumigraphs [Gortler et al. 1996] achieve this for a given set of viewing and lighting conditions. Surface light fields [Miller et al. 1998; Wood et al. 2000] record objects under a single fixed illumination for arbitrary viewpoints. In contrast to surface light fields, reflectance fields [Debevec et al. 2000] as well as polynomial texture maps [Malzbender et al. 2001] capture an object illuminated by a set of distant point light sources seen from a fixed viewpoint.

Strongly translucent and transparent objects with reflection and refraction effects can be recorded for a single viewpoint using environment matting [Zongker et al. 1999; Chuang et al. 2000]. Matusik et al. [2002] combine environment matting and reflectance fields to acquire an image-based representation of transparent and refractive objects including their 3D shape. Even though some of the above techniques record objects under varying illumination, they record only cases where the whole object is illuminated by distant light sources. Hence, as noted by Debevec et al. [2000], they are unable to faithfully reproduce the effects of local illumination variation such as shadow boundaries projected onto an object.

Masselus et al. [2003] capture the reflectance field of a scene for a fixed viewpoint and arbitrary illumination parameterized as a 4D incident light field. This allows to light the scene with arbitrary light sources including local illumination within the relatively coarse resolution limit of the incident light field.

Jensen et al. [2001] measure the scattering parameters for common materials such as marble or skin. They illuminate a single point on the surface, capture the reflected radiance with a digital camera, and compute the absorption cross section σ_a and the reduced scattering cross section σ'_s . Their results allow modeling of homogeneous objects made from these materials. Arbitrary inhomogeneous objects are beyond the scope of their measurement approach.

In contrast, we take advantage of the specific properties of translucent materials and acquire the diffuse subsurface reflectance function $R_d(x_i, x_o)$ that has no angular variance at incident and exiting surface locations. This allows us to densely sample the incoming and outgoing surface locations. We furthermore go beyond a pure image-based representation and transform the acquired data into a hierarchical data structure which reduces the storage cost considerably and allows us to interpolate missing data points consistently.

2.3 Rendering

A variety of rendering techniques such as finite element methods [Rushmeier and Torrance 1990; Sillion 1995; Blasi et al. 1993], bidirectional path tracing [Hanrahan and Krueger 1993; Lafortune and Willems 1996], photon mapping [Jensen and Christensen 1998; Dorsey et al. 1999], Monte Carlo simulations [Pharr and Hanrahan 2000; Jensen et al. 1999], or diffusion [Stam 1995; Stam 2001] are able to simulate the effects of subsurface scattering based on physical principles. Subsurface scattering can also be integrated in the framework of precomputed radiance transfer [Sloan et al. 2003].

Jensen et al. [2001] introduced a practical BSSRDF model consisting of a dipole approximation of the diffuse, multiple scattering term and an extension of the Hanrahan and Krueger BRDF model [Hanrahan and Krueger 1993] for the local, directional dependent effects caused by single scattering. The simplicity of the model implies drastically reduced rendering times compared to a full Monte Carlo Simulation. Although the dipole approximation is only correct for the surface of a homogeneous, infinite half-space, it yields visually convincing results. The degree of realism can be further improved by adding a modulation texture to simulate spatially varying materials.

Rendering with BSSRDFs is expensive since contributions from the whole surface must be taken into account. Jensen and Buh-

ler [2002] sample therefore the irradiance first and hierarchically evaluate the diffusion approximation in a second pass. Hao et al. [2003] use a local illumination model and precompute the transfer factors for all vertex neighborhoods. Mertens et al. [2003a] use a clustering hierarchy which is re-built in real time to render deformable translucent objects interactively. They also propose a local subsurface scattering algorithm that integrates the samples in image space at interactive rates using hardware acceleration [Mertens et al. 2003b]. The translucent shadow map [Dachsbacher and Stamminger 2003] augments a shadow map by the irradiance entering the object and the surface normal. This information is used to evaluate the dipole approximation by hierarchically filtering the translucent shadow map in real time on graphics hardware. Lensch et al. [2003b] compute the local and global response due to subsurface scattering in a preprocessing phase. The local response is stored as filter kernels for an illumination map, the global response is modeled by vertex-to-vertex throughput factors. Both terms are evaluated in parallel at rendering time. Carr et al. [2003] built an adaptive link structure in a preprocessing phase encoding the overall response to incoming light and evaluate it at rendering time using 3 passes on the graphics board.

Apart from Jensen et al. [2001], none of the above rendering methods are based on extensive sets of measured input data. The main focus of DISCO is therefore to capture the exact behavior of real translucent objects including effects that cannot be captured by simplified models such as the dipole approximation (e.g., cracks in an object, volumetrically varying properties, hollow objects). The acquired model is comprehensive and versatile. We will also discuss how it can be integrated into a variety of the above rendering algorithms and present the acquired models rendered by the method of Lensch et al. [2003b].

3 Background

The BSSRDF S provides a general model for basic light transport. Omitting spectral effects, polarization, and fluorescence, the outgoing radiance $L^\rightarrow(x_o, \omega_o)$ at a surface location x_o and a direction ω_o is related to the incident radiance $L^\leftarrow(x_i, \omega_i)$ at all surface locations x_i and for all directions ω_i via

$$L^\rightarrow(x_o, \omega_o) = \int_A \int_{\Omega_+(x_i)} L^\leftarrow(x_i, \omega_i) S(x_i, \omega_i; x_o, \omega_o) d\omega_i dx_i.$$

A denotes the surface of the object and $\Omega_+(x_i)$ is the hemisphere of directions on the outside of the surface at x_i .

Under the assumption of multiple light scattering inside the object which removes the angular dependencies at incoming and exiting surface locations we can replace S with a 4D diffuse subsurface reflectance function $R_d(x_i, x_o)$:

$$L^\rightarrow(x_o, \omega_o) = \frac{1}{\pi} F_{t,o}(\eta, \omega_o) \int_A R_d(x_i, x_o) \cdot \int_{\Omega_+(x_i)} L^\leftarrow(x_i, \omega_i) F_{t,i}(\eta, \omega_i) \langle N_i \cdot \omega_i \rangle d\omega_i dx_i. \quad (1)$$

The Fresnel transmittance factors $F_t(\eta, \omega)$ model what fraction of the flux or radiosity is transmitted through the surface boundary for a direction ω and a given relative index of refraction η . The factor $1/\pi$ converts radiosity into exitant radiance. The area foreshortening is represented by $\langle N_i \cdot \omega_i \rangle$.

The goal of our work is to measure $R_d(x_i, x_o)$ per color channel for all incoming and outgoing surface locations x_i and x_o . If we illuminate a target object at a single surface point x'_i with known incident radiance $L^\leftarrow(x'_i, \omega_i)$, we can observe $L^\rightarrow(x_o, \omega_o)$ on the object's surface. Inverting Equation 1 becomes trivial and we can record the impulse response $R_d(x'_i, x_o)$ for all x_o .

3.1 Hierarchical Model

Storing and rendering the full 4D function R_d with high resolution is impractical. The nature of optically dense translucent objects makes R_d however well suitable for a hierarchical modeling approach: R_d will typically vary strongly in the proximity of x_i due to the exponential fall-off of radiosity inside the scattering volume. In distant areas, R_d varies quite smoothly. Sharp features are usually caused by inhomogeneous material properties within a small volume on and right below the object's surface at x_i and x_o . All photons pass through these volumes and their influence is significant (see Feng et al. [1993] for a study of photon path distributions within scattering material).

Lensch et al. [2003b] introduced therefore a three-part model in their rendering system: The irradiance for all surface areas is collected in a texture atlas in the form of a diffuse light map. Spatially varying filter kernels $K_{(u,v)}$ that are convolved with the light map model an object's impulse response in the immediate vicinity of the incoming surface location x_i . The global term is modeled by a vertex-to-vertex throughput factor matrix \mathbf{F} . Energy transport is performed by multiplication with the per-vertex irradiance. The global response is linearly interpolated between the vertices of the underlying triangle mesh (equivalent to Gouraud interpolation). An optional modulation texture T_ρ adds surface appearance detail by modulating the global response.

The general strategy is to use a model with high sampling density in the vicinity of the incoming impulse and a much coarser sampling in distant, smooth areas (still modulated by a detail texture T_ρ). The sample location could for example also be determined by adaptively subdividing the input mesh. Subdivision allows one to choose the size of the densely sampled area adaptively and to guarantee that global and corresponding local samples coincide whenever possible.

We implemented our renderer based on the technique of Lensch et al. [2003b] and will describe all further processing steps with respect to this rendering technique. The problems and the proposed solutions are however general and transferring them to an alternative hierarchical data structure should be straightforward.

3.2 Measurement Overview

In our measurements, we illuminate individual surface points of an object with a narrow laser beam as shown in Figure 3. A digital camera observes the reflected radiance. Given the geometry of both setup and object as well as the Fresnel transmittance factors $F_t(\eta, \omega)$, we can invert Equation 1 and record samples of the diffuse subsurface reflectance function $R_d(x_i, x_o)$. $R_d(x_i, x_o)$ can be sampled densely by changing the laser's point of incidence and the camera pose.

In practice, the target object is illuminated with a narrow beam of light with finite width. Regarding the influence of the incident angle ω_i on the outgoing radiance $L^\rightarrow(x_o, \omega_o)$ we can therefore distinguish between two cases. If x_o is located outside the beam incident at x_i we can assume that the full energy (modulated by the Fresnel factor $F_t(\eta, \omega_i)$) enters the object at x_i and is scattered diffusely inside. The angle of incidence ω_i influences then only $F_t(\eta, \omega_i)$, the area foreshortening has to be omitted.

Within the area of incidence of the beam, the area foreshortening is definitely important since it scales the irradiance at each point. The shape of the peak around the incident beam will also change to some degree depending on the beam width and the angle of incidence ω_i .

An unknown factor in our measurements are the Fresnel factors $F_t(\eta, \omega)$. They represent only an approximation to the real behavior of a surface [Schlick 1994] and require at least knowledge of the relative index of refraction η . We follow here Jensen et al. [2001] and set $\eta = 1.3$ but acknowledge that this is not exact.



Figure 3: Acquisition setup and acquisition order. *Left and middle:* The laser projector illuminates the target object (resting on the turntable). The scene is observed by the HDR video camera from several positions. The two spotlights visible on both sides of the camera in the right image are used to illuminate the object during silhouette detection. *Right:* Acquisition order in pseudo code.

Godin et al. [2001] have shown that 3D laser range scanning of translucent objects suffers from a systematic bias since the peak location is moved for varying viewing and lighting directions. Since we also rely on peak detection, we expect that our measurements are influenced by this behavior to some degree.

Overall, the ideal measurement setup would therefore illuminate an object with an infinitesimal beam of light along the surface normal in order to avoid all these problems. But as this is not practical we have to compromise. One good criterion to limit the amount of error introduced is to consider only samples where the angle of incidence is below a threshold. The introduced noise must be dealt with in later stages of the pipeline.

4 Acquisition Setup

Figure 3 shows an overview over the DISCO setup: The target object is illuminated by a laser projection system that sweeps a laser beam over the object’s surface. The object is placed on a turntable in order to illuminate as many surface regions as possible. The object’s response to the illumination is recorded by a high-dynamic range video camera. The camera is manually placed at different locations relative to the projection system in order to minimize occlusions and to record the object from all sides. We rely on the repeatability of the laser projector and the turntable to ensure that identical surface positions are lit for all camera positions. We record one image per view (a combination of turntable and camera position) where the object is fully illuminated by additional spotlights for registration of the object with respect to the camera. The whole acquisition runs at 15 fps (half the camera speed) as we discard all frames during which the laser position was changed. Additional breaks occur when the laser color is changed (due to the warmup curve of the lasers) or when the turntable is rotated.

The acquisition order in pseudo code is shown in Figure 3 (right). We will again make use of this order in our postprocessing steps where many computations need only to be done once per view.

4.1 Laser System

The custom-built laser system consists of three individual lasers – a red diode laser and green and blue solid-state lasers with wavelengths of 635 nm, 532 nm, and 476 nm and 10 mW optical power per laser. All laser beams are fed into a single optical fiber to ensure exact alignment of the resulting beams. A collimator at the other end of the fiber focuses the beam to a size of about 2 mm within the working range. A 2D galvanometer scanner deflects the beam with high precision along a regular angular grid to achieve a sample spacing of about 1 mm. Care is taken in the whole laser system to reduce laser speckle to a minimum in order to avoid measurement artifacts.

4.2 High-Dynamic Range Video Capture

High-dynamic range (HDR) video capture can be achieved with specialized cameras and/or software (see Kang et al. [2003] for an overview). A specific requirement of our measurement setup is the quite extreme dynamic range for which we need linear radiance values without interfering quantization artifacts or blooming. In addition, care must be taken that the interesting features in the vicinity of the laser spot are not masked by lens flare or other artifacts caused by the high scene contrast. The use of a high quality lens is therefore mandatory.

Cameras with standard logarithmic CMOS chips are in principle well suited for the task at hand due to the exponential fall-off of subsurface scattering. Linear response can be achieved by a simple HDR calibration step [Debevec and Malik 1997; Robertson et al. 1999]. These cameras suffer however often from strong quantization artifacts as the images are quantized to 8–12 bits. Newer technologies such as the Photonfocus LINLOG technology should improve this situation.

In DISCO, we use a Silicon Vision Lars III high-dynamic range video camera equipped with a Jenoptik Lametar 2.8/25 mm lens. The camera records grayscale images with 768×496 pixel resolution at up to 30 fps and returns linear radiance values over a range of approximately 120 dB. The basic principle of this camera is that each pixel decides at fixed time steps (powers of two) whether it already received sufficient irradiance. This corresponds to an estimate whether the pixel will be overexposed at the next time step given constant irradiance. Each pixel records the exposure time and the amount of charge collected within this time from which linear radiance values with high precision throughout the dynamic range are computed. The acquired image streams are then compressed and stored on a dedicated RAID array.

4.3 Geometry Acquisition

Optical 3D scanning of translucent objects is challenging due to the non-local light reflection [Levoy et al. 2000; Godin et al. 2001; Matusik et al. 2002] – even a human observer can have difficulties to visually detect fine shape details (Figure 2). We cover therefore the test objects with a thin layer of white dust to achieve an almost Lambertian reflection (see Figure 2). The objects are then scanned with a Minolta VI-910 laser scanner and final triangle meshes are generated using commercial geometry-processing tools.

4.4 Geometric and Photometric Calibration

We performed a geometric camera calibration [Bouguet 2003] to recover the intrinsic parameters of the HDR camera. Due to the lack of features on translucent objects and the ease of silhouette detection caused by the global light transport, we use a silhouette based registration algorithm [Lensch et al. 2001] to recover the pose of the target object relative to the camera. Given a set of laser hit points




Figure 4: The global term for the starfruit model. *Left*: A single column of the global throughput factor matrix. The green vertex in the center marks the point of incidence, red areas denote missing values. *Right*: The same vertex after interpolation. Note that values on the ridge are interpolated consistently.

on the object’s surface and the corresponding deflection settings of the laser projector, we are able to recover the position of the laser projector relative to the setup.

For the photometric calibration, we rely on the overall linearity of the camera output and assume that the laser power is constant over time. We then need to perform a white balancing step for the individual lasers taking the spectral response of the camera into account. To this end, we sequentially illuminate a white, scattering calibration target with the three lasers and sum up the contribution of all image pixels.

5 Efficient Data Access

A typical acquisition yields an uncompressed data volume of several hundreds of gigabytes. (Data compression can reduce the size of the raw data to a few tens of gigabytes while our final models are typically only a few hundreds of megabytes.) It is therefore mandatory to use efficient algorithms to access the input data when estimating our hierarchical model. In this section, we describe the measures we took to speed up this access before we describe the post-processing of the global and local term in Sections 6 and 7, respectively.

Each complete acquisition consists of a small number of views (combinations of camera position and turntable rotation, typically 20–30 views). Most of the essential information for the further processing steps is constant per view and needs to be computed only once as the camera observes the object always from the same perspective. This information includes the position of the object and the laser projector relative to the camera. We precompute the Fresnel term and assign all pixels in the input image a confidence value based on the viewing and lighting directions that is used as weighting factor for the input data. We also reject at this stage all pixels that are close to a specular highlight or seen under grazing angles and generate the mapping from our texture atlas into the input images. The mapping from vertices to image coordinates and the visibility are precomputed.

In estimating our hierarchical model, we evaluate all images for a given turntable and laser projector position. These images show the object under identical illumination conditions and contain all information that is available for a specific illumination condition. We first decide whether the image tuple is valid, i.e., whether the laser spot is visible, and can then efficiently resample the data using the precomputed information. Our current implementation processes the input data streams with up to 50 fps on a PC with 3 GHz Xeon CPU which is more than three times the speed of the acquisition.

6 Global Term Post-Processing

An object’s diffuse subsurface reflectance function $R_d(x_i, x_o)$ (see Section 3) away from the point of incidence is represented with a



Figure 5: The global term for the starfruit model. *Left*: Irradiance. *Middle*: The global term before interpolation. Missing data leads to artifacts. *Right*: The global term after interpolation. The artifacts are reduced and missing data is filled in.

global term in our hierarchical model. The discrete representation of this global term is the matrix of throughput factors \mathbf{F} . In our method the shape of the object is represented as a triangular surface mesh. The variation of the diffuse subsurface reflectance function is approximated linearly between vertices in response to incident light at a vertex. This discretization approach corresponds to a *Galerkin* method with linear triangular elements, i.e., linear shape functions over triangular elements with hat functions at incident light locations. The throughput factor matrix \mathbf{F} is then filled with subsurface reflectance functions – one function per vertex which is stored in a column of \mathbf{F} . Light transport is of course symmetric and so is \mathbf{F} . The task in post-processing is to estimate the throughput factor matrix \mathbf{F} based on the acquired measurements. Post-processing combines individual observations for a specific point of incidence x_i , it interpolates within a single subsurface reflectance function (a column of \mathbf{F}) and between subsurface reflectance functions of neighboring vertices. During interpolation the distribution of energy within the material needs to be taken into account. In homogeneous material the energy falls off exponentially away from the point of incident light according to the diffusion approximation. However, our method is aimed at inhomogeneous objects with behavior which deviates from the smooth diffusion approximation. Figure 4 shows an individual subsurface reflectance function for the star fruit model before and after interpolation of missing values. Figure 5 shows renderings with the full matrix \mathbf{F} before and after interpolation.

6.1 Data Resampling

The acquired data is in a suitable format for a subsurface reflectance function representation because of our choice of measurement method. We record the response of the object to incident light at a point on the surface. A high-dynamic range image of the object is already a scaled subsurface reflectance function for light entering at a point x_i . However, the data consists of samples at discrete locations of the image plane. It also only covers the part of the surface which is in a view. Knowing the 3D geometry of the measurement setup allows us to resample the data on the object’s surface. The surface is represented as a triangular surface mesh where each triangle represents a similar surface area. Resampling then is the look-up of the bi-linearly interpolated image intensity at each vertex location. The position of the incident light needs also to be resampled and is assumed to contribute to the three vertices of the enclosing triangle. We weight the light according to the barycentric coordinates of the point inside the triangle. We combine measurements of the same subsurface reflectance function observed from different viewpoints and with varying laser colors. The result of the resampling are columns of the throughput factor matrix \mathbf{F} in *RGB* color space with some missing entries due to unobserved surface area and completely missing columns for vertices on the surface which were never lit. The interpolation of the matrix of throughput factors which addresses these cases is described next.

6.2 Interpolation

Interpolation of the throughput factors within a column of the matrix is the task of function interpolation on an irregularly meshed surface. In the diffusion approximation for homogeneous media the function is a sum of exponential decays which non-trivially depend on distance. Jensen et al. [2001] report difficulties in fitting the diffusion approximation with dipole lighting to the measured responses of homogeneous material samples. The function is more complicated for inhomogeneous objects and, in our case the interpolation has to potentially fill in large unseen areas. We conclude that function fitting seems inappropriate. Instead, we use filtering of throughput factors on the mesh (similar to mesh fairing [Taubin 1995]) with different averaging filter kernels. The edge length in the mesh is taken into account as by Desbrun et al. [1999]. The filter operates on the logarithm of the transfer coefficients because of the exponential decay of the subsurface reflectance function. The filtering operates on color information which originates from separate images with sequential laser illumination. Therefore, the noise in the color information will affect luminance and chrominance. We choose the $CIE YU^*V^*$ model [CIE 1986] since it yields linear luminance values and allows interpolation of chrominance values in a near-linear space with respect to the human visual system.

On the symmetric matrix \mathbf{F} we fill all the missing entries $\mathbf{F}_{r,k}$ in columns k where we observe the incident light point $\mathbf{F}_{k,k}$. Filling is performed by weighted averaging fixing the observation values. In particular, we solve the following iteration (similar to Perona and Malik [1990] but on a mesh domain)

$$\mathbf{F}_{r,k}^{t+1} = \mathbf{F}_{r,k}^t + \alpha * (1 - c_r) * \sum_{n \in \mathcal{N}} \psi \left(\frac{\mathbf{F}_{n,k} - \mathbf{F}_{r,k}}{e_{n,r}}, \sigma \right)$$

The neighbors n of a vertex r are its one-ring neighborhood \mathcal{N} connected with edges of length $e_{n,r}$. The weighting function ψ can be understood as the derivative of a error norm ρ with the scale parameter σ [Sapiro 2001]. We employ least squares with $\psi(x, \sigma) = 2 * x / \sigma^2$ and median filtering with $\psi(x) = \text{sign}(x)$. The choice of *confidence* c_r in the current transfer coefficient $\mathbf{F}_{r,k}$ controls the update rate and $c_r = 1$ keeps the existing measurement fixed. We observed satisfactory results with this approach. For highly non-smooth transfer matrices image inpainting techniques (e.g., [Bertalmio et al. 2000]) may be able to interpolate structure as well.

Subsurface reflectance functions for vertices for which the laser did not reach the triangle surface fan are interpolated iteratively from neighboring vertices. This is justified since a throughput factor $\mathbf{F}_{k,c}$ far from the point of illumination is typically similar to throughput factor $\mathbf{F}_{k,n}$ connecting the same vertex k with incident light positions n close to the position of illumination c . The neighborhood \mathcal{N} is the one-ring of the vertex c . This approach will break down for the diagonal entry $\mathbf{F}_{k,k}$ and close-by points. The neighborhood \mathcal{N} of diagonal elements $\mathbf{F}_{k,k}$ are the neighboring diagonal $\mathbf{F}_{n,n}$ of the one-ring. We define the vertices of the one-ring neighborhood as close points and blend diagonal interpolation and far interpolation. We define diagonal *neighbors* of a close point as the distance-weighted average response of the vertices of the one-ring neighborhood of diagonal neighbors. (This is similar to the interpolation of Green's functions of an elastic solid described by Pai et al. [2001]). Far and diagonal interpolation are illustrated in Figure 6.

We ensure symmetry of the throughput factor matrix by setting $\mathbf{F}^{t+1} = \frac{1}{2}(\mathbf{F}^t + (\mathbf{F}^t)^T)$ between each kind of interpolation, as well as, at the beginning and the end of the post-processing. The interpolation for missing entries in a column of \mathbf{F} and for missing complete columns is achieved with the same filtering framework. The differences between the two tasks are limited to different neighborhood definitions. Our filtering framework could also be easily extended




Figure 6: Far and diagonal interpolation of throughput factor matrix \mathbf{F} . The throughput factor shown in red is interpolated based on the neighboring factors shown in black.




Figure 7: *Left*: Texture atlas with borders for the horse model. *Right*: Recovered detail texture (color coded version of the green channel: blue areas become darker, red areas become brighter, green areas remain unchanged).

to diffusion solved by forward Euler iterations and diffusion solved with implicit schemes [Desbrun et al. 1999].

6.3 Detail Texture

The detail represented in the throughput factor matrix \mathbf{F} is limited by the resolution of the underlying mesh. We follow Lensch et al. [2003b] in increasing the resolution of the subsurface reflectance functions with two kind of textures. We estimate illumination dependent local high-resolution filter kernels (see Section 7) and a global illumination independent modulation texture T_ρ . The texture values of T_ρ are derived during resampling from object regions distant from the point of incidence of the illumination. For a texel in T_ρ that is visible and smoothly lit in an input image we record its RGB value t_{uv} . We select the appropriate mesh triangle and interpolate t_{int} based on the throughput factors at the corresponding vertex locations. The weighted average of the ratio t_{uv}/t_{int} over all appropriate input images is stored in T_ρ . At rendering time, we multiply the global response with T_ρ . The construction ensures that the radiosity remains unchanged at mesh vertices and is consistently modulated at all other locations. Figure 7 shows as an example the recovered detail texture for the horse model.

7 Local Term Post-Processing

We follow Lensch et al. [2003b] and model local light propagation by fixed sized spatially varying filter kernels $K_{(u,v)}$ (7×7 pixels in all our examples) on a texture atlas. The filter kernels convolve the irradiance at texel locations for local energy transfer. We use a texture atlas with a border and with a low guaranteed upper limit for texture stretch (on the order of 2). Figure 7 shows an example texture atlas of size 512×512 pixel with borders. The filter kernels are estimated in the texture atlas domain based on images where the laser spot is visible. In general, the peak of the laser spot will fall between the discrete pixel locations in the texture atlas. In the




Figure 8: Building the filter kernels. *Left*: Location of the laser peak mapped to its’ nearest neighbor texel. The plot shows monochrome luminance values from the 8×8 texel neighborhood around the highlight. The horizontal axis is the distance in mm to a neighboring texel in the texture atlas. The values show a fall-off with distance but no clear structure. *Middle*: If we plot the same data over the distance to the sub-pixel peak location, a fall-off resembling the sum of two exponential functions appears. *Right*: Plot of the dipole approximation for the material properties of marble from Jensen et al. [2001] for shape comparison.

following, we describe the resampling necessary to ensure accurate filter kernels.

7.1 Data Resampling

Figure 8 (left) shows an example of the peak of a filter kernel mapped to the nearest neighbor texel. All texels in an 8×8 pixel neighborhood around the peak location are plotted according to their 3D distance to the peak texel. A clearly structured fall-off is not visible. If we however plot them according to their distance to the subpixel location of the peak (see center of Figure 8), the general shape of the fall-off of subsurface scattering material is recognizable (see Figure 8 (right) for a comparison to the dipole approximation with arbitrary parameters). The measurement data also reveals strong variations due to surface detail at a given texel location. Surface detail is associated with a texel location and should not be resampled to preserve the sharpness of features.

In order to recover plausible filter kernels, we therefore shift the peak location to all four neighboring pixel locations while keeping the surface details spatially fixed. To separate the illumination from the surface detail, we fit the approximation

$$m(d) = c_1 \cdot e^{\alpha_1 d} + c_2 \cdot e^{\alpha_2 d}$$

to the data. d is the 3D distance from the peak location, c_1 and c_2 are fit per color channel, α_1 , and α_2 are identical for all color channels. In order to achieve a stable fit we first set c_2 to 0 and fit α_1 and c_1 in log space to all data points with $d > \xi_1$. We then fix α_1 and c_1 and fit the remaining parameters to all data points with $d < \xi_2$ ($\xi_1 = 3$ mm and $\xi_2 = 1$ mm for the example in Figure 8). This typically results in a stable fit of the exponents. In the next step, we fit only the color values c_1 and c_2 to all data points. The difference between the measured data points and $m(d)$ is then encoded in a 8×8 pixel multiplicative correction texture C .

In order to compute the four filter kernels for the neighboring vertices, we evaluate $m(d)$ for their respective center location and divide the result values by C . If some pixels in the neighborhood were not visible in the input images, we can interpolate them using $m(d)$.

7.2 Interpolation

Interpolation of filter kernels is performed by vector filtering over the surface of the object. Each 7×7 filter kernel is represented as a 49-vector. We use the same filtering framework as for the transfer coefficients \mathbf{F} but on the texture atlas domain instead of the triangular mesh. The neighborhood \mathcal{N} in the texture atlas contains a texel’s four-neighbors except for texture atlas boundaries. The boundaries

need to be treated specially in order to ensure the correct neighborhood information. Each boundary texel links therefore to a corresponding non-boundary texel in a different map of the atlas that is used instead of the boundary texel during filtering. Filter kernels also differ in size since different areas in the atlas are not isometric. The functional approximation allows for an easy interpolation independent of size at texture map boundaries.

Beside the vector of function parameters, we also propagate the multiplicative correction textures C . Typically, most of the multiplicative textures overlap in our examples due to their size and the laser spot sample spacing so that little interpolation is necessary.

8 Rendering

The acquired and processed dataset can be directly rendered using the approach of Lensch et al. [2003b]. While they computed the throughput factors, the filter kernels and the modulation texture in a pre-processing step based on the dipole approximation, we measure this data. The actual rendering step is then straightforward: We first render an illumination map of the object. The global and local energy transfer are evaluated in parallel using the CPU and both parts are rendered together. A more detailed description can be found in [Lensch et al. 2003b].

The model is however not limited to this particular rendering technique. Although it does not contain the basic parameters needed for a physical simulation, it can still be treated as a “black box” by a Monte Carlo raytracer or a photon mapping system. Both require only a probabilistic evaluation of the local and global throughput factors which can be sped up using an inverted cumulative density function [Pitman 1992]. The model can also be directly integrated in the preprocessing phase of precomputed radiance transfer [Sloan et al. 2003], the local illumination framework of Hao et al. [2003] or the adaptive link technique by Carr et al. [2003]. The model can substitute an evaluation of the dipole approximation [Jensen et al. 2001] although the size of the data structure makes an evaluation on current graphics hardware difficult.

9 Results

We acquired the following objects to validate our approach: An alabaster horse sculpture, a rubber duck and a starfruit (carambola). The alabaster horse has strong variation in its subsurface scattering properties and complex geometry. It consists of regions with varied translucency and contains cracks inside the material. The base material of the rubber duck is uniform, strongly scattering rubber. The beak and the eyes are painted on the outside with glossy, opaque paint. Incoming light is so strongly diffused at the rubber surface



Figure 9: The test objects under indoor illumination (top row) and illuminated by all three lasers (bottom row).

	Horse	Duck	Starfruit
# input views	24	25	20
# input images	1.065.744	541.125	401.220
input size (compressed)	31G	14G	12G
acquisition time	20.5h	11.25h	8h
# vertices	8924	5002	5001
# filter kernels	82.390	115.151	112.538
processing time (resampling)	7.8h	3.6h	3.4h

Table 1: Some statistics about the acquired models.

that the assumption of diffuse, multiple scattering remains valid although the duck is empty inside. The starfruit is an example of a translucent biological object with concave geometry and a relatively specular surface. Figure 9 shows all objects under indoor illumination, as well as, illuminated by the three lasers (i.e., the object’s impulse response). Synthesized images of the measured models are presented in Figures 1, 4, 5, 10, 11, 12, and 13. Table 1 summarizes important properties of these models.

DISCO captures models with local variation and significant detail. The model of the alabaster horse sculpture in Figure 1 shows nicely areas of differing translucency. There is a more opaque region at the head around the eyes while the muzzle is quite translucent. The left side of the head (when viewed from the muzzle) has various opaque areas towards the support and in the mane. There is a crack running top to bottom in the center at the neck as well. The model captures these volume effects. Figure 10 demonstrates how the local light transport adds surface detail and gives an impression of the varying optical densities at the object’s surface. The structure is also visible in the global throughput factor matrix albeit smoother. The side-by-side comparison of a rendering of the model and a photograph shows that our method recovers fine structure detail (see lower images in Figure 10). The highly translucent veins as well as the strongly diffuse patches which are visible in the photograph are present in our model. The slightly brownish region in the center is also captured well by our model, both, in color and shape.

The duck model in Figures 11 and 12 shows how our method can deal with completely heterogenous objects without representing the material distribution within the object explicitly. The rubber duck is made of regular material overall but head and body appear nearly separate when lit from behind. The wings block more light probably because of the extra thickness. The beak and eyes are




Figure 10: The horse model. *Top*: Irradiance, local light transport by filter kernels, and global light transport due to the throughput factor matrix. *Bottom left*: Combined rendering of local and global terms. The local light transport adds surface detail and gives an impression of the material properties at the object’s surface whose basic structure is also visible in the global term. *Bottom right*: Photograph of the real object under similar conditions. A slide projector was used to produce the sharp shadow boundary.




Figure 11: The duck model illuminated with a spot light. *Left*: Local filter kernels. *Center*: Global throughput factors. *Right*: Combined rendering.

marked with an opaque paint layer. The model captures these major deviation from homogeneous behavior while being smooth overall. The head of the rendered duck lacks features since we are reaching the limits of the dynamic range of our video camera.

The starfruit is visually quite interesting because of its shape. The model in Figure 13 is of good fidelity despite missing capture data due to the geometrically complex shape. Additional input images would fill gaps better than our data interpolation technique.

9.1 Discussion

The quality of the input data, both, in terms of surface coverage and in terms of noise, determines the amount of post-processing necessary and ultimately the quality of the final model. Surface coverage is limited by the classical stereo vision occlusion problem. There is also a trade-off between additional imagery and increased acquisition time. The intensity of the low-cost lasers of our projection system varies due to noise requiring additional smoothing during the interpolation. This variation as well as the observed intensity drift due to thermal effects lead currently to color artifacts in the final renderings. These issues can be solved by either locking the lasers’ intensity or by monitoring and calibrating for the variations.

Our DISCO method is limited to a specific class of objects with




Figure 12: Details of the duck model. *Left:* Head of the duck illuminated with a small spotlight from the back. The beak and the eyes are clearly visible. The spot near the tail is caused by diffuse reflection from the back of the head during the acquisition. *Middle:* Body of the duck model illuminated with a small spotlight from the back. *Right:* Photograph of the duck taken under similar conditions. The images show the specific property of the real duck that light illuminating the head is scattered mainly inside the head and light illuminating the body is scattered mainly inside the body.



Figure 13: Starfruit model with global transfer function and local term (7×7 filter kernels). *Left:* The hierarchical model of the starfruit is illuminated from the front right. *Right:* The same model illuminated from the back. The global transfer function alone is rendered in Figures 4 and 5.

strong subsurface scattering and diffuse surface reflection. The diffuse subsurface reflectance function is unable to represent the angular dependency of the direct surface reflection. This is noticeable in the rendering of the horse as well as the starfruit. In weakly scattering media such as murky water or honey, the diffuse light transport assumption does not hold. But as this assumption is widely used in current rendering systems, our approach can be used for both qualitative and perceptual validation of these models beyond the results shown in Figure 8.

10 Conclusion and Future Work

We presented DISCO – the first method to acquire the subsurface scattering behavior of optically dense translucent objects. The comprehensive, hierarchical models can be used by a variety of rendering approaches. Missing information is consistently interpolated and noise artifacts of the acquisition are reduced. We validated the approach by acquiring three translucent objects with strongly differing behavior.

For the future, we would like to improve the quality of the acquired models by improving the acquisition setup as discussed in Section 9.1. An acquisition planning step which determines the lit surface positions should be able to improve the quality of the models, reduce the post-processing effort and speed up the acquisition. Furthermore, we would like to validate the method by acquiring well-defined test targets and comparing the results both to other measurements (e.g., point-based measurements for homogeneous materials), as well as, to simulation results. We expect that the DISCO approach is not only useful to digitize translucent objects but can also help to analyze strengths and weaknesses of current rendering approaches.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable comments. Thanks also to Heiko Wanning for 3D scanning and Kristina Scherbaum for some illustrations. This work was funded in part by the DFG Schwerpunktprogramm V3D2 and by the European Union within the scope of the ViHAP3D Project (IST-2001-32641).

References

- BERTALMIO, M., SAPIRO, G., CASELLES, V., AND BALLESTER, C. 2000. Image Inpainting. In *SIGGRAPH 2000*, 417–424.
- BLASI, P., SAËC, B. L., AND SCHLICK, C. 1993. A Rendering Algorithm for Discrete Volume Density Objects. *Computer Graphics Forum 13*, 3, 201–210.
- BOUGUET, J.-Y., 2003. Camera Calibration Toolbox for Matlab. See <http://www.vision.caltech.edu/bouguetj>.
- CARR, N. A., HALL, J. D., AND HART, J. C. 2003. GPU Algorithms for Radiosity and Subsurface Scattering. In *Graphics Hardware 2003*, 51–59.
- CHUANG, Y.-Y., ZONGKER, D. E., HINDORFF, J., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2000. Environment Matching Extensions: Towards Higher Accuracy and Real-time Capture. In *SIGGRAPH 2000*, 121–130.
- CIE, 1986. Colorimetry. Publication CIE No. 15.2.
- DACHSBACHER, C., AND STAMMINGER, M. 2003. Translucent Shadow Maps. In *Rendering Workshop 2003*, 197–201.
- DEBEVEC, P., AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *SIGGRAPH 1997*, 369–378.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the Reflectance Field of a Human Face. In *SIGGRAPH 2000*, 145–156.
- DESBRUN, M., MEYER, M., SCHRÖDER, P., AND BARR, A. H. 1999. Implicit Fairing of Irregular Meshes using Diffusion and Curvature Flow. In *SIGGRAPH 1999*, 317–324.
- DORSEY, J., EDELMAN, A., JENSEN, H. W., LEGAKIS, J., AND PEDERSEN, H. K. 1999. Modeling and Rendering of Weathered Stone. In *SIGGRAPH 1999*, 225–234.
- FENG, S., ZENG, F., AND CHANCE, B. 1993. Monte Carlo Simulations of Photon Migration Path Distributions in Multiple Scattering Media. In *Photon Migration and Imaging in Random Media and Tissues, Proc. of SPIE Vol. 1888*, 78–89.
- GODIN, G., RIOUX, M., BERARDIN, J.-A., LEVOY, M., COURNOYER, L., AND BLAIS, F. 2001. An Assessment of Laser Range Measurement on Marble Surfaces. In *5th Conf. on Optical 3D Measurement Techniques*.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The Lumigraph. In *Computer Graphics Proceedings (ACM SIGGRAPH 96)*, 43–54.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from Layered Surfaces due to Subsurface Scattering. In *SIGGRAPH 1993*, 165–174.

- HAO, X., BABY, T., AND VARSHNEY, A. 2003. Interactive Sub-surface Scattering for Translucent Meshes. In *Proc. of the 2003 Symposium on Interactive 3D Graphics*, 75–82.
- ISHIMARU, A. 1978. *Wave Propagation and Scattering in Random Media*. Academic Press.
- JENSEN, H. W., AND BUHLER, J. 2002. A Rapid Hierarchical Rendering Technique for Translucent Materials. In *SIGGRAPH 2002*, 576–581.
- JENSEN, H. W., AND CHRISTENSEN, P. H. 1998. Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps. In *SIGGRAPH 1998*, 311–320.
- JENSEN, H. W., LEGAKIS, J., AND DORSEY, J. 1999. Rendering of Wet Materials. In *Rendering Workshop 1999*, 281–290.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A Practical Model for Subsurface Light Transport. In *SIGGRAPH 2001*, 511–518.
- KANG, S. B., UYTENDAELE, M., WINDER, S., AND SZELISKI, R. 2003. High Dynamic Range Video. *ACM Trans. Graph.* 22, 3, 319–325.
- LAFORTUNE, E. P. F., AND WILLEMS, Y. D. 1996. Rendering Participating Media with Bidirectional Path Tracing. In *Rendering Workshop 1996*, 91–100.
- LENSCH, H., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Silhouette-based Algorithm for Texture Registration and Stitching. *Graphical Models* 63, 4, 245–262.
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Image-based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Trans. Graph.* 22, 2, 234–257.
- LENSCH, H. P., GOESELE, M., BEKAERT, P., KAUTZ, J., MAGNOR, M. A., LANG, J., AND SEIDEL, H.-P. 2003. Interactive Rendering of Translucent Objects. *Computer Graphics Forum* 22, 2, 195–206.
- LEVOY, M., AND HANRAHAN, P. 1996. Light Field Rendering. In *Computer Graphics Proceedings (ACM SIGGRAPH 96)*, 31–42.
- LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. 2000. The Digital Michelangelo Project. In *SIGGRAPH 2000*, 131–144.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial Texture Maps. In *SIGGRAPH 2001*, 519–528.
- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF Measurement Including Human Skin. In *Rendering Workshop 1999*, 139–152.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2003. Relighting with 4D Incident Light Fields. *ACM Trans. Graph.* 22, 3, 613–620.
- MATUSIK, W., PFISTER, H., ZIEGLER, R., NGAN, A., AND MCMILLAN, L. 2002. Acquisition and Rendering of Transparent and Refractive Objects. In *Rendering Workshop 2002*, 267–278.
- MERTENS, T., KAUTZ, J., BEKAERT, P., SEIDEL, H.-P., AND VAN REETH, F. 2003. Interactive Rendering of Translucent Deformable Objects. In *Rendering Workshop 2003*, 130–140.
- MERTENS, T., KAUTZ, J., BEKAERT, P., VAN REETH, F., AND SEIDEL, H.-P. 2003. Efficient Rendering of Local Subsurface Scattering. In *Pacific Graphics 2003*, 51–58.
- MILLER, G. S. P., RUBIN, S., AND PONCELEON, D. 1998. Lazy Decompression of Surface Light Fields for Precomputed Global Illumination. In *Rendering Workshop 1998*, 281–292.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometrical Considerations and Nomenclature for Reflectance. National Bureau of Standards.
- PAI, D. K., VAN DEN DOEL, K., JAMES, D. L., LANG, J., LLOYD, J. E., RICHMOND, J. L., AND YAU, S. H. 2001. Scanning Physical Interaction Behavior of 3D Objects. In *SIGGRAPH 2001*, 87–96.
- PERONA, P., AND MALIK, J. 1990. Scale-Space and Edge Detection using Anisotropic Diffusion. *IEEE PAMI* 12, 629–639.
- PHARR, M., AND HANRAHAN, P. 2000. Monte Carlo Evaluation of Non-linear Scattering Equations for Subsurface Reflection. In *SIGGRAPH 2000*, 75–84.
- PITMAN, J. 1992. *Probability*. Springer.
- ROBERTSON, M. A., BORMAN, S., AND STEVENSON, R. L. 1999. Dynamic Range Improvement Through Multiple Exposures. In *Proc. of the Int. Conf. on Image Processing (ICIP'99)*, 159–163.
- RUSHMEIER, H. E., AND TORRANCE, K. E. 1990. Extending the Radiosity Method to Include Specularly Reflecting and Translucent Materials. *ACM Trans. Graph.* 9, 1, 1–27.
- SAPIRO, G. 2001. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press.
- SCHLICK, C. 1994. An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum* 13, 3, 233–246.
- SILLION, F. X. 1995. A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters. *IEEE Trans. Visualization and Computer Graphics* 1, 3, 240–254.
- SLOAN, P.-P., HALL, J., HART, J., AND SNYDER, J. 2003. Clustered Principal Components for Precomputed Radiance Transfer. *ACM Trans. Graph.* 22, 3, 382–391.
- STAM, J. 1995. Multiple Scattering as a Diffusion Process. In *Rendering Workshop 1995*, 51–58.
- STAM, J. 2001. An Illumination Model for a Skin Layer Bounded by Rough Surfaces. In *Rendering Workshop 2001*, 39–52.
- TAUBIN, G. 1995. A Signal Processing Approach to Fair Surface Design. In *SIGGRAPH 1995*, 351–358.
- WOOD, D. N., AZUMA, D. I., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D. H., AND STUETZLE, W. 2000. Surface Light Fields for 3D Photography. In *SIGGRAPH 2000*, 287–296.
- ZONGKER, D. E., WERNER, D. M., CURLESS, B., AND SALESIN, D. H. 1999. Environment Matting and Compositing. In *SIGGRAPH 1999*, 205–214.