

Tutorial Notes
EUROGRAPHICS 2007
Tutorial 6

Capturing Reflectance - From Theory to Practice

September 3, 2007

Hendrik P. A. Lensch
MPI Informatik
URL: <http://www.mpi-inf.mpg.de/~lensch>

Michael Goesele
University of Washington
URL: <http://www.cs.washington.edu/homes/goesele>

Gero Müller
Bonn University
URL: <http://cg.cs.uni-bonn.de/staff/person.asp?id=51>

Abstract

One important problem in photorealistic or predictive rendering nowadays is to realistically model the light interaction with objects. Measurements can capture the reflection properties of real world surface, i.e., they are one way of obtaining realistic reflection properties.

For arbitrary (non-fluorescent, non-phosphorescent) materials, the reflection properties can be described by the 8D reflectance field of the surface, also called BSSRDF. Since densely sampling an 8D function is currently not practical various acquisition methods have been proposed which reduce the number of dimensions by restricting the viewing or relighting capabilities of the captured data sets. In this tutorial we will mainly focus on three different approaches, the first allowing to reconstruct opaque surfaces from a very small set of input images, the second allows for arbitrary surfaces but under the assumption of distant light sources and the last which allows for relighting an arbitrary scene with arbitrary spatially varying light patterns.

After a short introduction explaining some fundamental concepts regarding measuring and representing reflection properties, the basics of data acquisition with photographs will be addressed. The tutorial present the set of current state-of-the art algorithms for acquiring and modeling 3D objects. The tutorial investigates the strengths and limitations of each technique and sorts them by their complexity with regard to acquisition costs. Besides describing the theoretical contributions we will furthermore point out the practical issues when acquiring reflectance fields in order to help interested users to build and implement their own acquisition setup.

Syllabus

- 8:30 **Introduction** (Lensch)
material properties
classification of techniques
- 8:45 **Acquisition Basics** (Goesele)
light sources
cameras
HDR
- 9:15 **Reflectance Sharing** (Goesele)
image-based BRDF measurement
spatially varying BRDFs
- 9:45 **BREAK**
- 10:00 **Reflectance Fields for Distant Lights** (Müller)
BTFs
light stage
acquisition, compression, synthesis and rendering
- 10:40 **Near-field Reflectance Fields** (Lensch)
relighting with 4D reflectance fields
dual photography
- 11:15 **Conclusion, Q/A** (all)

Resume of the Presenters

Michael Goesele is a postdoctoral research associate in the computer graphics and vision group at the University of Washington. In 1999, he joined the computer graphics group at the MPI Informatik and received his PhD from Saarland University in 2004. His research is focused on a broad range of acquisition techniques for computer graphics. Among others, he recently published two papers at ACM SIGGRAPH about the acquisition of light sources (Accurate Light Source Acquisition and Presentation) and translucent objects (DISCO – Acquisition of Translucent Objects). He has given several lectures and tutorials (e.g. at Eurographics 2002 and SIGGRAPH 2005) about the topics covered in the tutorial.

Gero Mueller currently works as a research assistant and Ph.D. student in the computer graphics group of Prof. Reinhard Klein at the University of Bonn, Germany. He received his diploma in computer science from the University of Bonn in 2002. His main research interests are realistic material representations, in particular BTFs. He has authored and co-authored several papers about this topic. At Eurographics 2004 he presented a state-of-the-art report covering the acquisition, compression, synthesis and rendering of BTFs and gave tutorials about the topic at various events (e.g. at Siggraph 2005).

Hendrik P. A. Lensch is the head of an independent research group "General Appearance Acquisition and Computational Photography" at the MPI Informatik in Saarbrücken, Germany. The group is part of the Max Planck Center for Visual Computing and Communication. He received his diploma in computers science from the University of Erlangen in 1999 and after joining the computer graphics group at MPI received his PhD from Saarland University in 2003. Dr. Lensch spent two years (2005-2006) as a visiting assistant professor at Stanford University, USA. His research interests include 3D appearance acquisition, image-based rendering and computational photography. For his work on reflectance measurement he received the Eurographics Young Researcher Award 2005. He was awarded an Emmy Noether Fellowship by the German Research Foundation in 2007. He has given several lectures and tutorials at various conferences including SIGGRAPH courses on realistic materials in 2002 and 2005.

Annotated Bibliography

Introduction

The goal of this annotated bibliography is to provide an overview over the most important publications in the areas covered by the course. Our goal was especially to help newcomers to the field to quickly become familiar with the main papers and serve as a starting point for further literature study. This is naturally always a subjective choice and we claim therefore by no means that the list of selected papers is complete and apologize for any important papers we missed.

General References

- [1] Richard S. Hunter and Richard W. Harold. *The Measurement of Appearance*. Wiley, 2. ed., 5. print. edition, 1987.

In this book, the various effects of reflections off surfaces are carefully described and analyzed. The authors provide valuable and intuitive insights on how to distinguish the appearance of two different materials. The book furthermore illustrates how the appearance of real world surfaces can be measured giving examples of techniques commonly applied in print industry. The main focus is on measuring the appearance of planar surfaces.

- [2] Fred E. Nicodemus, Joseph C. Richmond, Jack J. Hsia, I. W. Ginsberg, and T. Limperis. *Geometrical Considerations and Nomenclature for Reflectance*. National Bureau of Standards, 1977.

This report introduces the basic concepts of BSSRDFs, BRDFs, and related functions to describe reflectance. It also defines the nomenclature for all of them and describes their relationships such as the derivation of the BRDF from the BSSRDF.

BRDFs

- [1] James F. Blinn. Models of Light Reflection for Computer Synthesized Pictures. In *SIGGRAPH '77: Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198. ACM Press, 1977.

This paper introduces the empirical Blinn-Phong model (based on the earlier Phong model [18]). It can model more realistic reflections using three parameters (diffuse and specular coefficient, specular exponent). The specular lobe is computed based on the halfway vector.

- [2] Samuel Boivin and André Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In Eugene Fiume, editor, *Proceedings of SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 107–116. ACM Press / ACM SIGGRAPH, August 2001. ISBN 1-58113-292-1.

This paper tries to solve the difficult problem of measuring BRDF in indoor scenes from a single observation. The hope is that the global illumination and grouping of measurements of multiple surface points provide sufficient constraints to estimate a per-patch BRDF. At first a simple diffuse BRDF model is assumed. If the observed error is still insufficient a specular lobe is added. In case of failure, further tests involve anisotropic or mirroring BRDFs.

- [3] R. Cook and K. Torrance. A reflection model for computer graphics. *ACM Transactions On Graphics*, 1(1):7–24, 1982.

The Cook-Torrance model is a modification of earlier reflectance models. The main assumption is that the surface is composed of tiny, perfectly reflective, smooth microfacets oriented at different directions. The facets are assumed to be V-shaped and their distribution is isotropic. The model takes into account the fact that the light might be blocked by other microfacets (shadowing). Similarly, it also considers the fact that the viewer does not see some of the microfacets since they are blocked by the other microfacets (masking effect). The model takes into account an average Fresnel term (polarization is not considered) when modelling the reflectance of individual microfacets. However, it does not allow for multiple light bounces between the microfacets. The orientation of the facets is assumed to have some distribution - Cook and Torrance use the Beckman distribution function.

- [4] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. In *Proc. SIGGRAPH*, pages 145–156, July 2000. ISBN 1-58113-208-5.

While this paper actually introduced the concept of reflectance fields it also contains a section where a BRDF model is fit to the measured data of each texel. The spatially varying BRDF yields some compression compared to the full reflectance field data set.

- [5] Paul Debevec, Chris Tchou, Andrew Gardner, Tim Hawkins, Charis Poullis, Jessi Stumpfel, Andrew Jones, Nathaniel Yun, Per Einarsson, Therese Lundgren, Marcos Fajardo, and Philippe Martinez. Estimating Surface Reflectance Properties of a Complex Scene under Captured Natural Illumination. Technical Report ICT-TR-06.2004, USC ICT, December 2004.

This reports combines the idea of clustered BRDFs with global inverse illumination. For a number of representative spots/materials the BRDF is captured using standard image-based BRDF measurement techniques under controlled illumination conditions. In order to capture the spatially varying BRDF of a building the incident light onto the building is captured by an environment map which serves as a illumination source in a global illumination framework. Based on the differences between the synthesized images and the captured HDR images the weight for combining the cluster BRDFs are updated for each texel individually.

- [6] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec. Linear light source reflectometry. *ACM Trans. Graphics.*, 22(3):749–758, 2003.

In this paper the fully spatially varying BRDF and a transmission term is estimated for rather flat documents. The illumination is provided by a linear light source which has to be considered during the BRDF estimation. The same data is also used to scan the 3D geometry of the surface.

- [7] Athinodoros S. Georghiades. Recovering 3-d shape and reflectance from a small number of photographs. In *Eurographics Symposium on Rendering: 14th Eurographics Workshop on Rendering*, pages 230–240, June 2003.

Georghiades addresses the problem of estimating shape and reflection properties at the same time. Given a set of images of the scene illuminated by a point light source of unknown position the approach sets up an optimization problem that solves for the diffuse component of the BRDF and the actual surface normal per pixel as well as a global specular component and the light source positions in the individual images. As in other shape-from-shading approaches assuming a continuous surface introduces a regularization term that allows for solving the large optimization problem.

- [8] X. He, K. Torrance, F. Sillion, and D. Greenberg. A comprehensive physical model for light reflection. *Computer Graphics*, 25(Annual Conference Series):175–186, 1991.

This paper presents a reflectance model that accounts for the phenomena that can be explained using both geometrical optics and wave optics (diffraction, interference). The

model supports arbitrary polarization of incident light, but the simplifications for unpolarized light are also presented. In general, the reflectance is modelled as a sum of three components: specular, directional diffuse, and uniform diffuse. The specular component accounts for mirror-like reflection. It depends on the Fresnel reflectivity, roughness, and shadowing factors. The directional diffuse contribution of the reflectance function is the most complex term. It accounts for diffraction and interference effects. It depends on surface statistics (the effective roughness and the autocorrelation length). The uniform-diffuse contribution is a result of multiple microfacet reflections and sub-surface reflections. It is expressed as a simple function of wavelength. The resulting isotropic reflectance model for unpolarized light is a function of four parameters. Each of the parameters has some physical meaning and (at least theoretically) can be measured separately.

- [9] Eric P. F. Lafortune, Sing-Choong Foo, Kenneth E. Torrance, and Donald P. Greenberg. Non-linear Approximation of Reflectance Functions. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126. ACM Press/Addison-Wesley Publishing Co., 1997.

The Lafortune model presented in this paper is an extension of the Phong model [18] with a diffuse term and multiple lobes. Each lobe consists of a weighted dot product between viewing and lighting direction raised to some power. This empirical model can handle off-specular peaks, backscattering and anisotropy and is frequently used to model the reflection properties of real, measured materials.

- [10] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, pages 103–114. Eurographics, June 2001. ISBN 3-211-83709-4.

This paper introduces the concept of capturing cluster BRDFs and expressing the spatial variation by per-textel weighted sums of cluster BRDFs. Making use of the idea of image-based BRDF measurements samples from multiple surface points are combined when determining the cluster BRDFs. This results in more reliable, that is, more plausible BRDF parameters and at the same time reduces the number of required input images. Drastically different materials distributed in the same patch can be reproduced faithfully.

- [11] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-Based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Transactions on Graphics*, 22(2):234–257, April 2003.

This paper extends the previous work towards estimating per-textel normals. Starting from a scanned and smoothed 3D geometry model the per-textel BRDF is estimated. In a photometric stereo approach the current estimate of the BRDF is used to update the surface normal.

- [12] S. Marschner, S. Westin, E. Lafortune, and K. Torrance. Image-based measurement of the Bidirectional Reflection Distribution Function. *Applied Optics*, 39(16):2592–2600, 2000.

Marschner et al. describe an image-based BRDF measurement system. They use a material sample with different surface normals. Each point with a different surface normal gives a different BRDF measurement. Their system uses a spherical sample of homogeneous material. A fixed camera takes images of the sample under illumination from an orbiting light source. The system, although limited to only isotropic BRDF measurements, is both fast and robust. Furthermore, they extend their method to surface geometry acquired with a laser range scanner to acquire reflectance of a human face.

- [13] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg. Image-based BRDF Measurement Including Human Skin. In *10th Eurographics Workshop on Rendering*, pages 131–144, June 1999.

This paper applied the idea of image-based BRDF measurement to objects of arbitrary geometry. A 3D scan of the object provides the geometric information. Multiple images illuminated by a flash light are combined in order to estimate a single BRDF for the object.

- [14] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Trans. Graph.*, 22(3):759–769, 2003.

The authors built an automatic measurement setup to densely capture isotropic BRDFs using spherical material samples. They analyze the data and construct a low-dimensional data-driven BRDF model using non-linear dimensionality reduction techniques.

- [15] D. McAllister, A. Lastra, and W. Heidrich. Efficient rendering of spatial bi-directional reflectance distribution functions. *Graphics Hardware 2002*, 2002.

The authors present the first real-time rendering framework for BTFs. They used the Lafortune model to approximate the spatially varying BRDFs which leads to an extreme compact representation amenable to hardware implementation. Since the Lafortune model does not approximate meso-scale shadowing and masking effects well, it is only suitable for materials with minor depth variation (SVBRDFs).

- [16] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–226. Eurographics Association, 2005.

This paper extends the measurement setup of [14] to anisotropic BRDFs. It furthermore fits the parameters of several BRDF models to the measured materials and analyzes the fitting quality.

- [17] K. Nishino, Z. Zhang, and K. Ikeuchi. "determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis". In *in Proc. of Eighth IEEE International Conference on Computer Vision ICCV '01*, pages 599–606, July 2001.

Nishino et al. address the complicated problem of reconstructing BRDF and incident illumination at the same time. Specular highlights observed in the individual images are projected into a global environment map to estimate incident illumination. In the next step the BRDF is estimated. Spatial variation is restricted to the diffuse component.

- [18] Bui Tuong Phong. Illumination for Computer Generated Pictures. *Commun. ACM*, 18(6):311–317, 1975.

This paper introduces the Phong model – one of the earliest empirical lighting models for computer graphics. The model consists of a diffuse term and one specular lobe. It is neither energy conserving nor reciprocal and is only well-suited to approximate plastic-like materials. Improvements and extensions of the model include [1, 9].

- [19] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In Eugene Fiume, editor, *Proceedings of SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 117–128. ACM Press / ACM SIGGRAPH, August 2001. ISBN 1-58113-292-1.

This paper as well presents a solution to the problem of estimating BRDF and illumination from the same set of images for which an iterative algorithm has been developed. The paper is mostly well-known for the use of spherical harmonics to represent environment maps as well as BRDFs. This representation allows for computing the convolution of the BRDF with the environment map by a simple dot product.

- [20] Y. Sato, M. Wheeler, and K. Ikeuchi. Object Shape and Reflectance Modeling from Observation. In *Proc. SIGGRAPH*, pages 379–388, August 1997.

In this paper shape and reflectance properties are captured using the same sensor but different illumination. There is no explicit registration step necessary to match 3D geometry and 2D images. The diffuse component of the BRDF is estimated per pixel while the specular component is constant per patch.

- [21] G. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(Annual Conference Series):265–273, 1992.

This paper presents one of the first methods to speed up the BRDF measurement process. Ward's measurement device (imaging gonio-reflectometer) consists of a hemispherical mirror and a CCD camera with a fisheye lens. The main advantage of his system is that the CCD camera can take multiple, simultaneous BRDF measurements. Each photosite of the imaging sensor contains a separate BRDF value. Moving the light source and material over all incident angles enables the measurement of arbitrary BRDFs. Ward also presents a BRDF model that is based on the elliptical Gaussian distribution. The model is carefully designed to be physically plausible - it supports energy conservation and reciprocity. It is also relatively simple and can be evaluated efficiently. The parameters of the model have physical meaning and theoretically can be measured independently.

- [22] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse Global Illumination: Recovering Reflectance Models of Real Scenes From Photographs. In *Proc. SIGGRAPH*, pages 215–224, August 1999.

This paper considers indoor scenes. A few input images are aligned with a geometry model of the room and the furniture. Given the positions of the light sources, the light transport in the room can be simulated incorporating global illumination effects. The estimated BRDF minimizes the error between the measured values and a global illumination forward solution.

- [23] Y. Yu and J. Malik. Recovering Photometric Properties of Architectural Scenes from Photographs. In *Proc. SIGGRAPH*, pages 207–218, July 1998.

The BRDFs of buildings in outdoor scenes are estimated considering the incident illumination from the sun and the sky. Only the diffuse component is allowed to vary freely across the surface.

BTFs

- [1] Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. Reflectance and texture of real-world surfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157, 1997.

This paper introduced Bidirectional Texture Functions to the computer graphics community. The authors present the CURET reflectance and texture database which made BTF and BRDF measurements publicly available for the first time. The sampling density of the BTFs (205 images per material) was not yet sufficient for high-quality rendering.

- [2] Jefferson Y. Han and Ken Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.*, 22(3):741–748, 2003.

A promising approach for capturing several BTF samples at once using a kaleidoscope is presented. A problem with the approach is that it is quite sensitive to imperfections in the mirrors and their configuration because the light is reflected several times within the kaleidoscope.

- [3] Xinguo Liu, Yaohua Hu, Jingdan Zhang, Xin Tong, Baining Guo, and Heung-Yeung Shum. Synthesis and Rendering of Bidirectional Texture Functions on Arbitrary Surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 10(3):278–289, 2004.

This paper can be regarded as a follow up paper to the BTF synthesis paper of Tong et al. from Siggraph 2002. It uses SVD to compress the BTF data and shows how the BTF can be synthesized and rendered from this compressed representation while achieving a significant speed up compared to the original method. It is also shown how the BTF can be rendered with graphics hardware.

- [4] D. McAllister, A. Lastra, and W. Heidrich. Efficient rendering of spatial bi-directional reflectance distribution functions. *Graphics Hardware 2002*, 2002.

The authors present the first real-time rendering framework for BTFs. They used the Lafortune model to approximate the spatially varying BRDFs which leads to an extreme compact representation amenable to hardware implementation. Since the Lafortune model does not approximate meso-scale shadowing and masking effects well, it is only suitable for materials with minor depth variation (SVBRDFs).

- [5] Jan Meseth, Gero Müller, and Reinhard Klein. Reflectance field based real-time, high-quality rendering of bidirectional texture functions. *Computers and Graphics*, 28(1):103–112, February 2004.

This paper addresses the problem of using parametric functions for representing BTFs with significant meso-structure. They propose to fit parametric functions not to the whole per-texel apparent BRDF but to the per-view per-texel reflectance functions which use to be relatively smooth functions.

- [6] G. Müller, G. H. Bendels, and R. Klein. Rapid synchronous acquisition of geometry and btf for cultural heritage artefacts. In *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST)*, pages 13–20. Eurographics Association, Eurographics Association, November 2005.

Based on a camera array of 151 of-the-shelf digital cameras a method for rapidly acquiring the geometry and reflectance of objects with significant meso-scale geometry is presented. It combines image-based 3D reconstruction and BTF compression and rendering techniques.

- [7] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein. Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum*, 24(1):83–109, March 2005.

This comprehensive overview discusses from acquisition, over synthesis to rendering of BTFs most of the topics covered in the BTF-part of this tutorial. The relevant publications in the field of BTFs up to the year 2005 are introduced.

- [8] G. Müller, R. Sarlette, and R. Klein. Data-driven local coordinate systems for image-based rendering. *Computer Graphics Forum*, 25(3), September 2006.

In this paper a data-driven technique for computing local coordinate systems from image-based reflectance measurements is presented. These coordinate systems allow to align the per-texel reflectance measurements which results in increased compression performance with negligible run-time overhead.

- [9] M. Sattler, R. Sarlette, and R. Klein. Efficient and realistic visualization of cloth. *Proceedings of the Eurographics Symposium on Rendering 2003*, 2003.

In this paper the first BTF real-time rendering framework based on statistical data analysis is presented. It describes the whole pipeline from measurement using a fully automatic setup over compression to rendering including image-based illumination and large scale shadows. It also introduces the BTF Database Bonn which still offers the most detailed publicly available BTF data.

- [10] Peter-Pike Sloan, Xinguo Liu, Heung-Yeung Shum, and John Snyder. Bi-Scale Radiance Transfer. *ACM Transactions on Graphics*, 22(3):370–375, 2003.

The authors combine Precomputed Radiance Transfer with BTFs to achieve striking real-time renderings of BTF-covered objects realistically lit by environment maps. They represent the BTF by projecting the data per sampled view direction into the Spherical Harmonics basis.

- [11] Frank Suykens, Karl vom Berge, Ares Lagae, and Philip Dutré. Interactive Rendering of Bidirectional Texture Functions. In *Eurographics 2003*, pages 463–472, September 2003.

This BTF compression and rendering method approximates the BTF data per texel using a sophisticated factorization scheme called Chained Matrix Factorization. The idea is to factorize the data with different parameterizations which are suitable for the different significant features of the per-texel apparent BRDFs. Thereby the data can be reliably represented with a much smaller number of factors compared to standard matrix factorization based on SVD.

- [12] M. A. O. Vasilescu and Demetri Terzopoulos. Tensortextures: Multilinear image-based rendering. In *Proceedings of SIGGRAPH*, August 2004.

This work introduces tensor representations for image-based datasets. In contrast to the classic matrix representation multi-linear tensors allow a so-called strategic dimensionality reduction. This means that e.g. more components can be spent for encoding the view variation which results in perceptually more satisfying reconstructions.

- [13] Hongcheng Wang, Qing Wu, Lin Shi, Yizhou Yu, and Narendra Ahuja. Out-of-core tensor approximation of multi-dimensional matrices of visual data. *ACM Trans. Graph.*, 24(3):527–535, 2005.

This paper improves the 3D tensor representation of Vasilescu et al. by arranging the data in higher-dimensional tensors (e.g. 5D). This allows to exploit the coherence along other dimensions like between the rows and columns of the measured images. The method achieves very high compression rates, generalizes to higher-dimensional datasets like time-varying BTFs and can be implemented as an out-of-core technique. The reconstruction costs are a disadvantage of the method.

Near-field Reflectance Fields

- [1] Billy Chen and Hendrik P. A. Lensch. Light source interpolation for sparsely sampled reflectance fields. In Günther Greiner, Joachim Hornegger, Heinrich Niemann, and Marc Stamminger, editors, *Vision, Modeling, and Visualization 2005 (VMV'05)*, pages 461–469, Erlangen, Germany, November 2005. Aka.

Captured reflectance fields are typically sparsely sampled in the light direction domain. In this paper, a method is presented that allows for smoothly moving light sources in near-field reflectance fields. The system treats high frequency illumination effects such as highlights and shadows separately from slowly moving effects such as the cosine fall-off and interreflections, for which linear blending is sufficient to reproduce the appearance of intermediate light source positions. Highlights and shadows are detected using intrinsic images and then moved according to the detected optical flow. The technique further exploits the properties of near-field reflectance fields to perform virtual 3D scanning.

- [2] Yanyun Chen, Xin Tong, Jiaping Wang, Stephen Lin, Baining Guo, and Heung-Yeung Shum. Shell texture functions. *ACM Transactions on Graphics*, 23(3):343–353, August 2004.

This paper presents an appearance representation approach that is particularly suited for heterogeneous translucent objects. The translucent object is divided into a homogeneous diffusely scattering core surrounded by volume of heterogeneous translucent material. The shell texture function (STF) provides an intermediate data structure representing the light transport and the mesostructure of the outer shell. For each voxel in the shell volume the irradiance due to light impinging from arbitrary directions is precomputed and stored in a 5D data structure.

- [3] Frédo Durand, Nicolas Holzschuch, Cyril Soler, Eric Chan, and François X. Sillion. A frequency analysis of light transport. *ACM Transactions on Graphics*, 24(3):1115–1126, August 2005.

This paper analyzes the different effects of occluders, reflectors, or the propagation of light in free space on the spatial and angular frequency content of the transformed light field. The authors propose a signal-processing framework and show a large set of instructive examples. They further show how the analysis of the frequency content of the light field can be used to control sampling rates or the choice of reconstruction kernels in rendering, pre-computed radiance transfer, and inverse rendering.

- [4] Gaurav Garg, Eino-Ville Talvala, Marc Levoy, and Hendrik P. A. Lensch. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pages 251–262, June 2006.

Capturing dense light transport matrices so far required a sequential sensing of individual incident light rays. In this paper, two techniques are combined in order to allow for fast acquisition of arbitrarily complex reflectance fields. The first is the use of H-matrices which subdivide a matrix hierarchically until each sub-block can be represented sufficiently well using a low-rank approximation of the block. The second ingredient is an symmetric acquisition system where cameras and projectors are coupled by a beam splitter allowing for emitting light and sensing light along exactly the same rays. This turns the captured reflectance field into a symmetric tensor whose sub-blocks can be determined in parallel given that they are of low rank. The paper features one of the first full 8D reflectance fields, at a rather low resolution, though.

- [5] Michael Goesele, Hendrik P. A. Lensch, Jochen Lang, Christian Fuchs, and Hans-Peter Seidel. DISCO – Acquisition of Translucent Objects. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)*, 23(3), 2004.

This is the first paper that captured the diffuse reflectance R_d of a real translucent object with inhomogeneous material properties. The object is pointwise illuminated and its impulse response function is captured with a HDR camera. A hierarchical model of transfer functions is computed from a large number of input images. Rendering can be performed in real time using an earlier approach.

- [6] Akira Ishimaru. *Wave Propagation and Scattering in Random Media*. Academic Press, 1978.

This book describes the physical principles of single and multiple scattering in various types of media and derives the mathematical formulations.

- [7] Henrik Wann Jensen and Juan Buhler. A Rapid Hierarchical Rendering Technique for Translucent Materials. In *SIGGRAPH 2002*, pages 576–581, 2002.

The authors propose a hierarchical evaluation technique to speed up the rendering of translucent objects using the dipole model [8]. This is the first of a whole series of papers proposing various rendering techniques to speed up evaluation of the dipole model – see e.g. [5] for a list of such publications.

- [8] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A Practical Model for Subsurface Light Transport. In *SIGGRAPH 2001*, pages 511–518, 2001.

This paper introduces the dipole model as an approximation for translucent objects in computer graphics. It describes the derivation of the model, its use for rendering, and compares the results to Monte-Carlo simulations. The authors describe also a measurement setup to determine the required parameters for real materials and provide a table of measured values. The dipole model is used in many publications including [7] as a fast method to evaluate the effects of subsurface scattering.

- [9] Shree K. Nayar, Gurunandan Krishnan, Michael D. Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, pages 935–944, New York, NY, USA, 2006. ACM Press.

In this paper a very efficient method is presented for separating the direct and the global component of the light reflected by a scene due to illumination by a projector. The key observation is that global light transport is due to multiple scattering and therefore dampens high frequency in spatially varying illumination patterns. The technique makes use of multiple shifted high frequency patterns and provides a very simple formula to perform the separation from the minimum and maximum intensity observed for each pixel in the sequence of shifted patterns. The separation results to some extent depend on the frequency of the illumination pattern.

- [10] Pieter Peers, Karl vom Berge, Wojciech Matusik, Ravi Ramamoorthi, Jason Lawrence, Szymon Rusinkiewicz, and Philip Dutré. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics*, 25(3):746–753, July 2006.

This paper presents a method for transferring the reflection properties of one heterogeneous translucent object onto novel geometry. In an initial acquisition the diffuse subsurface reflectance is measured on a planar slab of material by illuminating individual points. The effect of subsurface scattering is assumed to be localized having a well controlled support. In order to compress the captured reflectance function the illumination peaks are aligned to one column and a set of homogeneous BSSRDFs is determined to describe the general shape. Dividing the measured samples by the homogeneous approximation results in a representation of the heterogeneous effects which can be factored in a compact way. When transferring the reflectance function to novel geometry only the light transport in a local neighborhood is considered.

- [11] Steven M. Seitz, Yasuyuki Matsushita, and Kiriakos N. Kutulakos. A theory of inverse light transport. In *ICCV '05: Proceedings of the Tenth IEEE International Conference on Computer Vision*, pages 1440–1447, Washington, DC, USA, 2005. IEEE Computer Society.

Given a captured near-field reflectance field between a projector and a camera, this paper analyzes how the reflectance field can be inverted in order to render the scene after the first, the second, or after multiple light indirections. The results indicate that it is sometimes possible to remove multiple scattering effects from captured reflectance fields. Note that the inversion of the reflectance field is possible only for a couple of special cases.

- [12] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. Dual photography. *ACM Transactions on Graphics*, 24(3):745–755, August 2005.

This paper presents an acquisition system for capturing near-field reflectance fields, i.e. measuring the light transport on a ray-to-ray basis. Using an adaptive algorithm, the reflectance field between a camera and a projector is measured such that the influence of

every projector pixel to every camera pixel is determined, yielding a 4D light transport matrix. Exploiting Helmholtz reciprocity, the light transport direction can be inverted. Instead of sending out light from the projector it is turned virtually into a sensing camera capturing the scene as if illuminated by a virtual projector at the location of the original camera. The adaptive and parallel capturing scheme accelerates the acquisition time for sparse light transport matrices by three orders of magnitude.

- [13] Xin Tong, Jiaping Wang, Stephen Lin, Baining Guo, and Heung-Yeung Shum. Modeling and rendering of quasi-homogeneous materials. *ACM Transactions on Graphics*, 24(3):1054–1061, August 2005.

This paper features an acquisition system and a model for capturing and rendering quasi-homogeneous materials. The model consists of a homogeneous subsurface reflectance function augmented by two functions modeling the mesostructure effects locally, i.e. independently for the incident and the exitant point of the light transport. The subsurface scattering effect is captured by sweeping a line stripe laser over the surface from various directions. In addition, a full BTF is acquired.

Presenters' Slides