

Image-Based Reconstruction of Spatial Appearance and Geometric Detail

HENDRIK P. A. LENSCH, JAN KAUTZ and MICHAEL GOESELE

Max-Planck-Institut für Informatik

and

WOLFGANG HEIDRICH

The University of British Columbia

and

HANS-PETER SEIDEL

Max-Planck-Institut für Informatik

Real-world objects are usually composed of a number of different materials that often show subtle changes even within a single material. Photorealistic rendering of such objects requires accurate measurements of the reflection properties of each material, as well as the spatially varying effects. We present an image-based measuring method that robustly detects the different materials of real objects and fits an average bidirectional reflectance distribution function (BRDF) to each of them. In order to model local changes as well, we project the measured data for each surface point into a basis formed by the recovered BRDFs leading to a truly spatially varying BRDF representation. Real-world objects often also have fine geometric detail that is not represented in an acquired mesh. To increase the detail, we derive normal maps even for non-Lambertian surfaces using our measured BRDFs. A high quality model of a real object can be generated with relatively little input data. The generated model allows for rendering under arbitrary viewing and lighting conditions and realistically reproduces the appearance of the original object.

Categories and Subject Descriptors: I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Color, Shading, Shadowing and Texture*; I.4.1 [**Image Processing and Computer Vision**]: Digitization and Image Capture

General Terms: Measurement

Additional Key Words and Phrases: BRDF measurement, spatially varying BRDFs, normal map acquisition, photometric stereo, shape from shading

1. INTRODUCTION

The use of realistic models for all components of image synthesis is a fundamental prerequisite for photorealistic rendering. This includes models for the geometry, light sources, and cameras, as well as materials and micro structure. As more and more visual complexity is demanded, it is less and less feasible to generate these models manually. Automatic

Author's address: Hendrik P. A. Lensch, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, lensch@mpi-sb.mpg.de.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0730-0301/20YY/0100-0001 \$5.00

and semi-automatic methods for model acquisition are therefore becoming increasingly important.

In this paper we concentrate on the acquisition of realistic materials. In particular, we describe an acquisition process for spatially varying BRDFs that is efficient, reliable, and requires little manual intervention. Other methods described in the literature (see Section 3 for an overview) either focus on homogeneous materials, or make assumptions on the type of material to be measured (e.g. human faces). In our work, we measure spatially varying BRDFs without making any additional assumptions. Further, we use the derived reflection properties to compute normal maps even for non-Lambertian surfaces. In particular, our contributions are

- efficient, reliable, and mostly automatic calibration schemes for the light source position relative to the geometry,
- a robust and efficient BRDF fitting process that clusters the acquired samples into groups of similar materials and fits a Lafortune model [Lafortune et al. 1997] to each group,
- a method that projects every sample texel into a basis of BRDFs obtained from the clustering procedure. This projection accurately represents the material at that point and results in a compact representation of a truly spatially varying BRDF,
- an algorithm that uses the reconstructed BRDF at every point together with the measured reflectance samples to optimize the orientation of the surface normal yielding a normal map.

We require only a relatively small number of high dynamic range photographs (about 15-25 images for one object), thereby speeding up the acquisition phase.

As a result of the fitting, clustering, and projection process, we obtain a compact representation of spatially varying materials that is well suited for rendering purposes (see Figure 11 for an example). The method works both for objects consisting of a mixture of distinct materials (e.g. paint and silver, see Figure 13), or for smooth transitions between material properties.

2. OVERVIEW AND MOTIVATION

Geometry and reflection properties have to be acquired in order to realistically reproduce a real world object. The basic ideas have been presented in [Lensch et al. 2001] which are further refined and extended in this paper. In particular, this includes extended sections on the setup and data resampling. Furthermore, we added new techniques for calibrating the light source position and for the acquisition of normal maps.

The reconstruction of spatially varying BRDFs requires to sample a 6D function in some way. A very naive approach would be to measure the BRDF of every single point on the surface separately, e.g. using a gonireflectometer. It is clear that this method would be tedious if not impossible. A more suitable approach has been taken by Debevec et al. [2000] who constructed the *light stage* where a point light source spins around the object while a video camera takes several hundreds of images for a fixed view. Hereby, the BRDF of each visible surface point is captured for exactly one viewing direction as a function of incident lighting.

In this paper we present a technique for measuring spatially varying BRDFs for any viewing direction using far fewer images. The idea of our algorithm is based on the fact that most man-made (and even many natural) objects consist of mixtures of only a very

small set of materials. We heavily exploit the coherence in the reflection properties of surface points belonging to the same material in order to drastically reduce the number of BRDF samples (i.e., the number of images) that have to be collected. The BRDFs of the basic materials are measured similar to the image-based measurement technique proposed by Marschner et al. [1998; 1999].

What remains is to represent the subtle variation and details which are present even within the same material. Real objects exhibit a lot of local variations in the material properties (such as imperfections within the material, dust, or smooth transitions to other materials). Therefore, a single, constant BRDF per cluster will not result in realistic appearance. In order to model these subtle details it is necessary to assign a separate BRDF to each point on the surface. Given the sparse input data, a direct measurement of these per-texel BRDFs is unfortunately impossible. However, one could think of the BRDF of a single point as a mixture of the BRDFs reconstructed for the basis materials. In this case, one only has to determine how much the basic materials contribute to the BRDF of each surface point. This task can be performed even with sparse input data.

Interreflections within the object are not yet considered in our algorithm. Our approach tries however to find the optimal approximation to the measured data, and in many cases comes up with a reasonable apparent BRDF even for concave regions.

The presentation of the overall procedure is organized as follows. Previous work concerning BRDF measurements and appearance acquisition is reviewed in the following section. Section 4 describes in detail the acquisition setup and the individual acquisition steps: capturing of HDR images, image to model registration, calibration of the point light source, and so on. The captured data is then resampled and reorganized to allow efficient access to all data that belongs to a single point on the surface (see Section 5).

From the resampled data the spatially varying BRDFs are reconstructed in two subsequent steps, a clustering and a projection step. The reflection properties of each material can be represented well by a single BRDF (Section 6). Thus, these basic materials of the object have to be found and separated. In Section 7 we explain how to group the surface points into clusters consisting of the same basic materials.

Local variation resulting in a per-texel BRDF is described as a weighted sum of the BRDFs of the basic materials. A set of basic materials is determined for every cluster, but the weights of these materials can vary from point to point. These weights are determined by projecting the measured data of each point into the basis of clustered materials (see Section 8).

The acquired reflectance samples can further be used to also obtain a more precise estimate of the surface normal at every point compared to the normal provided by the triangle mesh. In Section 9 the measured BRDFs are used to derive normal maps based on the shading information even for non-Lambertian materials.

Section 10 briefly describes our rendering method. In Section 11 we present our results and then we conclude in Section 12.

3. RELATED WORK

The representation of real-world materials has recently received much attention in the computer graphics community. The approaches can be grouped into three categories: light field and image database methods with static illumination, dense sampling of the light and viewing directions to generate a tabular representation of the BRDF, and finally the fitting of

reflection models, often based on a sparse set of samples. This last approach is the one we take and extend to spatially varying BRDFs.

3.1 Light Field Techniques

In the first category, there has been a number of approaches ranging from a relatively sparse set of images with a geometric model [Debevec et al. 1996] over the Lumigraph [Gortler et al. 1996] with more images and a coarser model to the light field [Levoy and Hanrahan 1996] with no geometry and a dense image database. Recently surface light fields [Wood et al. 2000; Miller et al. 1998] have become popular, which feature both a dense sampling of the directional information and a detailed geometry. In our work we use an algorithm similar to the function quantization approach proposed by Wood et al. [2000] to resample the image data into a compact representation which we extended to model effects due to the incident light.

3.2 Sampling the Reflection Field

In contrast to light field approaches, bidirectional texture functions [Dana et al. 1999] also allow for changes in the lighting conditions, although at very high storage costs. Another lighting dependent texture called polynomial texture maps has been proposed by Malzbender et al. [2001]. In order to allow sparser sampling Liu et al. [2001] synthesized new samples by reconstructing the surface structure of the observed material.

Debevec et al. [2000] describe a method for acquiring the reflectance field of human faces. In one part of their work they fit a specialized reflection model for human skin to the measured data (consisting of about 200 images). Both specular and diffuse parameters of the reflection model can vary rapidly across the surface, but other parameters like the de-saturation of the diffuse component at grazing angles are constant and only apply to human skin. In our work we try to avoid making assumptions on the kind of material we are measuring by using general BRDF models.

Matusik et al. [2002] acquire the reflectance field of different slightly transparent objects, but store the data in image space. The acquired data is directly used for rendering. This technique was extended to transparent and refractive objects [Matusik et al. 2002]. Furukawa et al. [2002] also acquired the reflectance field of different objects and proposed various compression techniques, again without fitting a BRDF model to the data.

3.3 BRDF Measurement

The traditional approach to measure reflectance properties is to use specialized devices (gonioreflectometers), that position both a light source and a sensor relative to the material. These devices can only obtain one sample for each pair of light and sensor position and are therefore relatively slow.

More recently, image-based approaches have been proposed. These methods are able to acquire a large number of samples at once. For example, Ward Larson [1992] uses a hemispherical mirror to sample the exitant hemisphere of light with a single image. Instead of using curved mirrors, it is also possible to use curved geometry to obtain a large number of samples with a single image. This approach is taken by Lu et al. [1998], who assume a cylindrical surface, and Marschner et al. [1998; 1999] who obtain the geometry using a range scanner. Our method is similar in spirit to the method of Marschner et al., but we are also dealing with spatially varying BRDFs and we are fitting a reflection model rather than using a tabular form in order to achieve a compact representation.

A number of researchers have described methods for fitting reflection models to measured sample data [Debevec et al. 2000; Koenderink et al. 1996; Lafortune et al. 1997; McAllister 2002; Sato et al. 1997; Ward Larson 1992; Yu et al. 1999]. Of these methods, the ones by Ward Larson [1992] and Lafortune et al. [1997] do not consider spatial variations. Sato et al. [1997] fit a Torrance-Sparrow model [Torrance and Sparrow 1967] to the data, and consider high-frequency variations for the diffuse part but only per-triangle variations for the specular part. This is also the case for the work by Yu et al. [1998; 1999], which also takes indirect illumination into account. Boivin and Gagalowicz [2001] reconstruct arbitrary reflection properties for whole patches in a scene using just one image. McAllister [2002] fits a spatially varying Lafortune model to very densely sampled planar materials. The achieved results are impressive but the technique requires flat surfaces and an automated setup to get a dense sampling of the reflection properties. In [Ramamoorthi and Hanrahan 2001] and [Gibson et al. 2001] inverse rendering algorithms are proposed that reconstruct the reflection properties and the incident light field at the same time. Ramamoorthi and Hanrahan [2001] as well as Westin et al. [1992] project BRDF data into a spherical harmonics basis instead of fitting an explicit reflection model.

We use in our work the Lafortune model because it is compact, and capable of representing interesting BRDF properties such as off-specular peaks and retro-reflection.

3.4 Normal Maps

Reflection properties together with measured photometric data can also be used to derive geometric information of the original object. Photometric stereo approaches (see [Zhang et al. 1999] for an overview) have been developed to extract geometric information from a set of pictures with different lighting conditions. The shading information is used to compute surface normals and depth values.

Rushmeier et al. calculate normal directions at every visible surface point from a set of images showing the same view of an object illuminated by a point light source placed at different but known positions for each image [Rushmeier et al. 1997; Rushmeier et al. 1998; Bernardini et al. 2001]. Besides the normal direction the diffuse color (albedo map) is reconstructed by solving a linear system of equation for each point. The technique assumes however the surface to be purely diffuse (Lambertian) and simply discards the brightest radiance samples at each point to circumvent specular highlights. Hereby valuable data is ignored. Additionally, as the surface gets more and more glossy it is hard to guarantee that all remaining samples show a purely diffuse reflection which is required to compute the correct normal direction. In our work we use measured non-Lambertian reflection properties to compute normal maps for arbitrary materials.

4. ACQUISITION

We obtain the 3D models with a structured light 3D scanner and a computer tomography scanner both generating dense triangle meshes. The triangle meshes are smoothed [Garland and Heckbert 1997; Kobbelt 1996], manually cleaned, and decimated.

All images are acquired in a measurement lab (see Figure 1) using a Kodak DCS 560 professional digital camera. An HMI metal halide bulb serves as point light source for the BRDF measurements. The interior of the photo studio is covered with dark and diffusely reflecting felt to minimize the influence of the environment on the measurements. The details of the lab environment are described elsewhere [Goesele et al. 2000].




Fig. 1. The acquisition setup in a photo studio covered with dark felt (from left to right): point light source, metal spheres for light source tracking, object to be measured, digital still camera.




Fig. 2. Image series captured for one camera/lighting position. a) Silhouette image and reconstructed silhouette. b) Two images for recovering the light source position (see Section 4.1). c) Photograph samples with varying exposure time for HDR image reconstruction.

Several views of each object are captured with different camera and light source positions. Light source and camera are positioned manually, which is however easily possible since only a few different views are required. Several constraints should be met when selecting the views to obtain the best quality:

- all surface points should be visible in more than one image,
- the position of the camera and the light source should be varied in order to provide different pairs of viewing and lighting directions, and
- at least one highlight should be observed in each material.

These considerations may slightly increase the number of images that are required to reproduce objects with a larger number of different basis materials. In our experience their number is typically small.

For each view we acquire three sets of images: one image of the object’s silhouette to register the 3D model with the image (Figure 2a), and two images to recover the light source position as explained in Section 4.1 (Figure 2b). We then acquire a high dynamic




Fig. 3. *Left figure: the ring flash mounted on the camera yields a highlight in the center of the spheres. Right figure: rays from the camera to the light source highlights will be reflected to the point light.*

range image of the object lit by the point light source by taking a series of photographs with varying exposure time [Debevec and Malik 1997; Robertson et al. 1999] (Figure 2c). We use an entropy-based dark frame subtraction algorithm [Goesele et al. 2001] to remove dark current CCD noise.

Once per session a high dynamic range image of a gray card with known camera and light position is taken in order to allow for an absolute calibration of the reflectance. In addition, a series of calibration images of a checkerboard pattern is taken whenever the lens settings are changed. The calibration method proposed by Zhang [1999] is used to recover the intrinsic camera parameters.

To register the images with the 3D model we use a silhouette-based method [Lensch et al. 2000; 2001] that yields the camera position relative to the object for a single view. Given the 3D model of the object the registration is performed by aligning the captured silhouette with the silhouette from a virtual view of the 3D model. The final view is found by minimizing the difference between the two silhouettes.

4.1 Recovering the Light Source Position

In order to recover the position of the point light source a geometric approach was used which requires no user interaction. Six steel spheres (manufactured for a ball bearing with $6\mu\text{m}$ precision) of known, equal diameter are used (see Figure 1). A metal fixture manufactured using a CNC milling machine ensures that the spheres are lying exactly on a straight line at an equal and known distance.

For each view two images of the spheres are acquired. One view shows only the reflection of the point light source in the spheres (Figure 2b top). For the second view a ring flash mounted on the camera lens produces a highlight on the center of each sphere (Figure 2b bottom). The exact centers of these reflections in the images are determined by automatically fitting ellipses to the highlights.

Given the intrinsic parameters of the camera, the pixel coordinates of the reflections of the ring flash define rays in space on which the centers $\vec{p}_{\text{sphere},i}$ of the spheres are located (see Figure 3). If the camera is in the origin this can be described by

$$\vec{p}_{\text{sphere},i} = \lambda_i \cdot \vec{f}_i, \quad (1)$$

where \vec{f}_i is the normalized ray direction towards the highlight in sphere i , λ_i is the ray parameter. A second constraint is that all spheres are offset by the same distance by the

vector \vec{d} , i.e.,

$$\vec{p}_{\text{sphere},i} = \vec{p}_{\text{sphere},0} + i \cdot \vec{d}. \quad (2)$$

Combining Equation 1 and 2 yields

$$\lambda_i \cdot \vec{f}_i = \vec{p}_{\text{sphere},0} + i \cdot \vec{d} \quad (3)$$

$$\Leftrightarrow -i \cdot \vec{d} = \vec{p}_{\text{sphere},0} - \lambda_i \cdot \vec{f}_i. \quad (4)$$

The fact that $-i \cdot \vec{d}$ is located on a line through $\vec{p}_{\text{sphere},0}$ with direction \vec{f}_i can also be expressed differently as

$$\vec{f}_i \times \vec{p}_{\text{sphere},0} + \vec{f}_i \times (i \cdot \vec{d}) = \vec{0}. \quad (5)$$

We obtain some scaled solution for $\vec{p}_{\text{sphere},0}$ and \vec{d} as the least squares solution of the system of linear equation formed by all highlights when constraining one component of \vec{d} to be one¹. Since the distance between the spheres $\|\vec{d}\|$ is actually known we can rescale the solution to obtain the final positions of the spheres.

Now we send rays from the camera to the positions of the light source highlights and reflect the rays off the spheres. The light source position is located at the intersection of the rays (see Figure 3). To increase the stability of this method more than the three necessary spheres are used and a least squares approximation is computed. The light source position could be determined with up to 2cm precision where both the camera and the light source were approximately 1m apart from the object.

After measuring the light source position the metal spheres are covered with black cloth in order to avoid disturbing the acquisition of the HDR images.

5. RESAMPLING OF REFLECTANCE VALUES

After acquisition of the geometric model, high dynamic range image recovery, and registration, it is necessary to merge the acquired data for further processing. For each point on the model's surface we collect all available information into a data structure we call a *lumitexel*.

One lumitexel, denoted by \mathcal{L} , is generated for every visible surface point. Each lumitexel stores the geometric and photometric data of one point, i.e., its position \vec{x} and the normal \vec{n} in world coordinates. Linked to the lumitexel is a list of reflectance samples \mathcal{R}_i , each representing the measured reflectance r of the surface point captured by one image plus the direction of the light \vec{U} and the viewing direction \vec{V} which are further transformed into the local coordinate frame of the surface point spanned by \vec{n} and a deterministically constructed tangent and bi-normal yielding \vec{u} and \vec{v} .

A lumitexel can be seen as a very sparsely sampled BRDF. We define the error between a given BRDF f and a lumitexel \mathcal{L} as:

$$E_f(\mathcal{L}) = \frac{1}{|\mathcal{L}|} \sum_{\mathcal{R}_i \in \mathcal{L}} s \cdot I(f(\vec{u}_i, \vec{v}_i)u_{i,z}, r_i) + D(f(\vec{u}_i, \vec{v}_i)u_{i,z}, r_i), \quad (6)$$

¹For stability reasons one may try two different components of \vec{d} set to one, e.g. $\vec{d}_z = 1$ will not work if the spheres are oriented parallel to the image plane.

where $|\mathcal{L}|$ stands for the number of reflectance samples linked to the lumitexel, $I(r_1, r_2)$ is a function measuring the intensity difference, and $D(r_1, r_2)$ measures the color-difference. We introduce the weight s , to be able to compensate for noisy data (e.g. a slightly wrong normal resulting in a wrong highlight). In order to emphasize the color-difference we always set $s \leq 1$. Note that the cosine between the normal and the local light direction u_z is already included in our reflectance samples r such that the BRDF f has to be multiplied by it.

5.1 Assembling Lumitexels

Collecting all reflectance samples for a lumitexel requires a resampling of the input images for the particular point on the surface. First, one has to determine the set of surface points for which a lumitexel should be generated. In order to obtain the highest quality with respect to the input images, the sampling density of the surface points must match that of the images.




Fig. 4. The correspondence between pixel position and point position \vec{x} on the object is computed by tracing a ray through the image onto the object. At every \vec{x} a local normal \vec{n} can be computed from the triangle's vertex normals.

Every triangle of the 3D model is projected into each image using the previously determined camera parameters. The area of the projected triangle is measured in pixels and the triangle is assigned to the image I_{best} in which its projected area is largest. For every pixel within the triangle in I_{best} a lumitexel is generated.

The position \vec{x} of the surface point for the lumitexel is given by the intersection of the ray from the camera through the pixel with the mesh (see Figure 4). The normal \vec{n} is interpolated using the triangle's vertex normals.

A reflectance sample \mathcal{R}_j is now constructed for each image I_j in which \vec{x} is visible from the camera position and the surface point is lit by the point light source. The vectors \vec{u}_j and \vec{v}_j can be directly calculated. The associated reflectance is found by projecting \vec{x} onto the image plane and retrieving the color c_j at that point using bilinear interpolation. Note, that for I_{best} no bilinear interpolation is necessary and c_{best} can be obtained without resampling since \vec{x} exactly maps to the original pixel by construction. The reflectance r_j of the reflectance sample \mathcal{R}_j is obtained by scaling c_j according to the brightness of the light source and the squared distance from the light source to \vec{x} .

5.2 Discarding Data at Depth Discontinuities

In order to increase the quality of the resampled data it is sometimes necessary to discard some of the input data. Especially near depth discontinuities the resampled data is prone to registration errors. If the 3D model is not perfectly aligned with the 2D image the part of the surface that is visible in one pixel may not correspond to the surface part predicted by the 3D model. In the case of depth discontinuities the visible part and the predicted part will not even be adjacent. Reflectance samples would then be assigned to completely wrong surface points or lumitexels.

Furthermore, since a sensor element of the camera always integrates over a finite area, the reflectance values reported at depth discontinuities are never reliable even if the registration with the 3D model were perfect. Thus, it is necessary to discard the image data at depth discontinuities.

The depth discontinuities are detected using the following approach: A depth map of the registered 3D model is rendered and subsequently blurred using an averaging filter. This changes the depth values of pixels near depth discontinuities while pixels showing a flat surfaces will not be affected. Regions where the filtered depth map deviates more than a small threshold from the original one will not be considered for further processing. The threshold can be computed given the filter size and the difference of two adjacent depth values that should be detected as a discontinuity.



Fig. 5. Left: Dark stripes on the dress and on the hand are due to depth discontinuities. Right: By discarding samples at depth discontinuities and shadow boundaries these effects have been removed.

The same approach also applies to the shadowing problem. Here, depth discontinuities result in shadow boundaries whose position can only be determined up to some uncertainty. Hence, also pixels near shadow boundaries have to be discarded. They can be determined by a filtered shadow map.

The results of removing samples at depth discontinuities are displayed in Figure 5. Note that the dark stripes on the dress and across the hand have been removed by this step.

6. BRDF FITTING

In this section we first detail the Lafortune BRDF model [Lafortune et al. 1997] that to fit to our given lumitexels. Then we explain how this fit is performed using Levenberg-Marquardt optimization.

6.1 Lafortune Model

BRDFs are four-dimensional functions that depend on the local viewing and light direction. The dependence on wavelength is often neglected or simply three different BRDFs are used for the red, green, and blue channel. We use the latter approach.

Instead of representing a measured BRDF as a 4D table the measured samples are in our case approximated with a parameterized BRDF model. This has two advantages: Firstly, the BRDF requires much less storage since only the parameters are stored and secondly, we only require a sparse set of samples that would not be sufficient to faithfully represent a complete tabular BRDF.

Many different BRDF models have been proposed (e.g. [Torrance and Sparrow 1967; Ward Larson 1992]) with different strengths and weaknesses. Our method may be used together with any parameterized BRDF model. We have chosen the computationally simple but general and physically plausible Lafortune model [1997] in its isotropic form, i.e., the orientation of the tangent and bi-normal within the tangential plane has no influence on the resulting reflectance:

$$f(\vec{u}, \vec{v}) = \rho_d + \sum_i [C_{x,i}(u_x v_x + u_y v_y) + C_{z,i} u_z v_z]^{N_i}. \quad (7)$$

This model uses only a handful of parameters: \vec{u} and \vec{v} are the local light and viewing directions, ρ_d is the diffuse component, N_i is the specular exponent, the ratio between $C_{x,i}$ and $C_{z,i}$ indicates the off-specularity of lobe i of the BRDF. The sign of $C_{x,i}$ makes the lobe i either retro-reflective (positive $C_{x,i}$) or forward-reflective (negative $C_{x,i}$). The albedo of the lobe i is given by the magnitude of the parameters $C_{x,i}$ and $C_{z,i}$. From now on we will denote the BRDF with $f(\vec{a}; \vec{u}, \vec{v})$, where \vec{a} subsumes all the parameters of the model, i.e., ρ_d , $C_{x,i}$, $C_{z,i}$, and N_i . In the case of only one lobe \vec{a} is 12-dimensional (4 parameters for each color channel).

6.2 Non-Linear Fitting

The Lafortune BRDF is non-linear in its parameters, which means that we have to use a non-linear optimization method to fit the parameters to the given data. As in the original work by Lafortune et al. [1997], we use the Levenberg-Marquardt optimization [Press et al. 1992] to determine the parameters of the Lafortune model from our measured data.

We ensure that the fitting process works well and does not get stuck in undesired local minima by initializing the fitting routine with parameters that correspond to an average BRDF.

The Levenberg-Marquardt optimization outputs not only the best-fit parameter vector \vec{a} , but also a covariance matrix of the parameters, which provides a rough idea of the parameters that could not be fit well. This information is used in our splitting and clustering algorithm, as explained in the next section.

7. CLUSTERING

In this section we will explain how we cluster the given lumitexels so that each cluster C_i corresponds to one material of the object. Given a set of BRDFs $\{f_i\}$, each cluster C_i consists of a list of all the lumitexels \mathcal{L}_i for which f_i provides the best approximation. Determining these clusters is a problem closely related to vector quantization [Gersho and Gray 1992] and k -means clustering [Lloyd 1982; MacQueen 1967], both of which work in affine spaces. Unfortunately, we do not have an affine space when clustering BRDF samples since there is no meaningful distance measure for BRDF samples with arbitrary viewing and lighting directions. Therefore we are employing a modified Lloyd [1982] iteration method.

The general idea is to first fit a BRDF f to an initial cluster containing all the data. Then we generate two new BRDF models f_1 and f_2 using the covariance matrix from the fit (explained in more detail below) representing two new clusters. The lumitexels \mathcal{L}_i from the original cluster are then distributed according to the errors $E_{f_1}(\mathcal{L}_i)$ and $E_{f_2}(\mathcal{L}_i)$ into the new clusters. We then recursively choose another cluster, split it, and redistribute the lumitexels and so on. This is repeated until the desired number of materials is reached, as detailed in Section 7.4.

7.1 Lumitexel Selection

The fitting procedure described in Section 6 performs a relatively large number of operations per reflectance sample. Thus, it is expensive to fit a BRDF using all lumitexels (and all reflectance samples contained in the lumitexels) generated by the assembling procedure. Instead, it is sufficient to consider only a few thousand lumitexels at the beginning. Later on, we increase the number for an accurate fit.

A first, naive approach to choosing this subset for fitting selects every n -th lumitexel regardless of its reliability or possible contribution. However, as stated in [Yu et al. 1999] and [Schirmacher et al. 1999], for a robust estimation of the specular part of a BRDF it is very important to include reflectance samples within the specular lobe of the material. Unfortunately, these brightest pixels statistically also carry the largest error.

Following these ideas we select more lumitexels in areas where a highlight is likely to occur. These areas are determined by the surface normal, the light source position and a synthetic BRDF with a broad highlight.

7.2 Splitting

Fitting just a single BRDF to the initial cluster of course is not sufficient if the concerned object consists of more than one material. Rather, we have to recursively split the clusters to account for the different materials comprising the object. We decide which cluster to split, by computing the following error for all clusters C_j :

$$E(C_j) = \sum_{\mathcal{L}_i \in C_j} E_f(\mathcal{L}_i). \quad (8)$$

The cluster C_j with the largest error will be split into two new clusters each with a different BRDF. Further materials can be extracted by further splitting the clusters.

But how do we split a cluster? The BRDF fit to a cluster represents the average material of the lumitexels in that cluster. Fitting the BRDF using the Levenberg-Marquardt algorithm (see Section 6) will also provide us with the covariance matrix of the parameters.




Fig. 6. Split-recluster-fit process (SRF). The initial BRDF is split into two new BRDFs using the covariance matrix. The lumitexels from the initial cluster are distributed according to their distance to the BRDFs. Then we fit the BRDF again to each new cluster. We now iterate the reclustering and fitting until the resulting BRDFs and clusters have converged.

The eigenvector belonging to the largest eigenvalue of this matrix represents the direction in which the variance of the samples is highest, and is therefore a good choice for the direction in which the parameter space is to be split.

Let \vec{a} be the fit parameter vector of the BRDF $f(\vec{a}; \vec{u}, \vec{v})$ for cluster C . Vector \vec{e} denotes the eigenvector belonging to the largest eigenvalue λ of the corresponding covariance matrix. We then construct two new BRDFs:

$$f_1(\vec{a} + \tau\lambda\vec{e}; \vec{u}, \vec{v}) \quad \text{and} \quad f_2(\vec{a} - \tau\lambda\vec{e}; \vec{u}, \vec{v}), \quad (9)$$

where τ is a scaling factor to adapt λ to a moderate value. Two new clusters C_1 and C_2 are generated by distributing every lumitxel \mathcal{L}_i of cluster C either to C_1 if $E_{f_1}(\mathcal{L}_i) < E_{f_2}(\mathcal{L}_i)$, or to C_2 otherwise. In the next step, f_1 and f_2 are fit to best approximate the lumitexels in the new clusters.

7.3 Reclustering

Because the parameters of a BRDF fit to a multi-material cluster are not necessarily the center of the parameters of the contained materials and due to improper scaling of λ and other reasons like noise, the performed split will not be optimal and the two new clusters may not be clearly separated, e.g. in the case of two distinct materials some lumitexels belonging to one material may still be assigned to the cluster of the other material.

A better separation can be achieved by iterating the procedure of distributing the lumitexels \mathcal{L}_i based on $E_{f_1}(\mathcal{L}_i)$ and $E_{f_2}(\mathcal{L}_i)$, and then fitting the BRDFs again. The iteration stops when the number of lumitexels in the generated cluster does not change any more. In our experiments this reclustering operation leads to a clear separation of materials and is done after each split. The split-recluster-fit (SRF) process is visualized in Figure 6.

When more than two clusters have been generated by successive binary splits and a new material is clearly distinguished, it is helpful to clean the other clusters, which were not involved in the last split, from all lumitexels belonging to the newly discovered material. This can be done in a global reclustering step by redistributing all initial lumitexels \mathcal{L}_i to the cluster C_j with

$$j = \underset{k}{\operatorname{argmin}} E_{f_k}(\mathcal{L}_i). \quad (10)$$

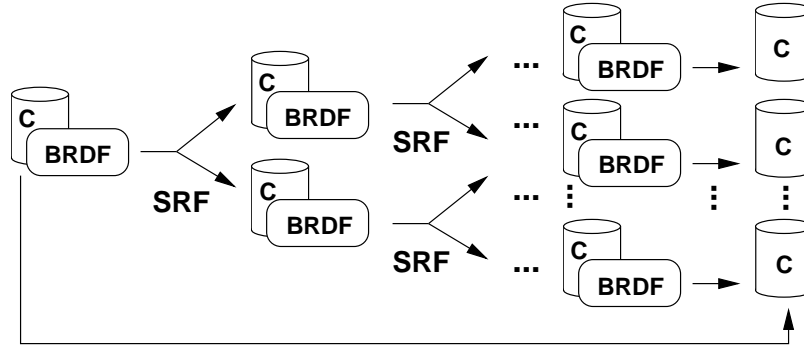


Fig. 7. The complete splitting and reclustering algorithm including the global reclustering, which is similar to the recluster-fit iteration, only that all lumitexels are distributed among all clusters.




Fig. 8. The clustering process at work. In every image a new cluster was created. The object was reshaped using only the single BRDFs fit to each cluster before the projection into a basis of multiple BRDFs.

And again, the BRDFs of all involved clusters have to be refit. This global reclustering is repeated several times to clearly separate the materials. We stop this iteration when the percentage of change is smaller than some ϵ , or a maximum number of iterations is reached. The complete splitting and reclustering algorithm is depicted in Figure 7 and the processing on a real model is shown in Figure 8.

7.4 Termination of the Splitting Process

We still have to decide when to stop the splitting process. To do this we require the user to input the estimated number of materials $|M|$. We stop the splitting and clustering process after at least $|M|$ clusters have been created. More clusters can be generated to compensate for the often noisy and not absolutely accurate reflectance samples (e.g. slightly wrong normals, noise in the images, misregistration, etc.).

This means that we do not necessarily have a one to one mapping between actual materials and clusters. This is not crucial since the projection, which we will present in the next section, uses a weighted sum of several BRDFs to accurately represent every lumitexel.

8. PROJECTION

The representation of an object by a collection of only a few clusters and BRDFs makes the virtual object look artificial as can be seen in Figure 9. The main reason for this is that real surfaces exhibit changes in the reflective properties even within a single material. These




Fig. 9. Left: The result of the clustering process still does not look realistic since there is no variation of the material within one cluster. Right: Spatial variation derived by projecting the reflectance samples of each lumitexel in a basis formed by the clustered materials.

changes cannot be represented by a single BRDF per cluster since all lumitexels within the cluster are assigned the same BRDF parameters.

To obtain truly spatially varying BRDFs we must find a specific BRDF for each lumitexel. But the sparse input data does not allow to fit a reliable or even meaningful BRDF to a single lumitexel because each lumitexel consists of only a few reflectance samples. In addition, we would need to acquire a highlight in every lumitexel to reliably determine the specular part, as already explained in Section 7.1.

The solution is to project each lumitexel into a basis of BRDFs (see Section 8.1). The BRDF f_{π_i} of a lumitexel \mathcal{L}_i is then represented by a linear combination of m BRDFs f_1, f_2, \dots, f_m :

$$f_{\pi_i} = t_1 f_1 + t_2 f_2 + \dots + t_m f_m, \quad (11)$$

with t_1, t_2, \dots, t_m being positive scalar weights. This forces the space of solutions (i.e., the possible BRDFs for a pixel) to be plausible since the basis BRDFs are already fit reliably to a large number of reflectance samples.

Given the BRDFs, the weights have to be determined for each lumitexel. Let $r_{j=1\dots|\mathcal{L}_i|}$ be the reflectance values of the lumitexel \mathcal{L}_i . The weights are found by a least square optimization of the following system of equations using singular-value decomposition:

$$\begin{pmatrix} r_1 \\ r_2 \\ \vdots \\ r_{|\mathcal{L}_i|} \end{pmatrix} = \begin{pmatrix} \tilde{f}_1(\vec{u}_1, \vec{v}_1) & \tilde{f}_2(\vec{u}_1, \vec{v}_1) & \dots & \tilde{f}_m(\vec{u}_1, \vec{v}_1) \\ \tilde{f}_1(\vec{u}_2, \vec{v}_2) & \tilde{f}_2(\vec{u}_2, \vec{v}_2) & \dots & \tilde{f}_m(\vec{u}_2, \vec{v}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{f}_1(\vec{u}_{|\mathcal{L}_i|}, \vec{v}_{|\mathcal{L}_i|}) & \tilde{f}_2(\vec{u}_{|\mathcal{L}_i|}, \vec{v}_{|\mathcal{L}_i|}) & \dots & \tilde{f}_m(\vec{u}_{|\mathcal{L}_i|}, \vec{v}_{|\mathcal{L}_i|}) \end{pmatrix} \begin{pmatrix} t_1 \\ t_2 \\ \vdots \\ t_m \end{pmatrix}, \quad (12)$$

with $\tilde{f}(\vec{u}, \vec{v}) := f(\vec{u}, \vec{v})u_z$. Compared to the non-linear fitting of BRDF model parameters (see Section 6.2), we now have a linear problem to solve with a smaller degree of freedom and even more constraints. Equation 12 shows only the system for one color channel, whereas the weights t_i have to be the same for all channels. In contrast to this, BRDF parameters would require a distinct set of parameters per channel.

We find the weights t_i as the non-negative least square solution of the system of equations. Negative values are avoided because they may result in an oscillating BRDF that represents only the given reflectance samples accurately but will produce unpredictable values for other viewing and light directions.

8.1 Basis BRDFs

The next question is how to determine the set of basis BRDFs. Since the changes of the surface properties within one material tend to be small, a distinct set of basis BRDFs is assigned to each cluster. It is therefore sufficient to store just the scalar weights per lumitexel instead of the full set of BRDF parameters.

Finding the optimal set of BRDFs f_1, f_2, \dots, f_m , that minimizes the error

$$E_\pi(C) = \frac{1}{|C|} \sum_{\mathcal{L}_i \in C} E_{f_{\pi_i}}(\mathcal{L}_i) \quad (13)$$

for a cluster C , where f_{π_i} denotes the least square projection of the lumitexel \mathcal{L}_i as defined in Equation 11, is a problem of principal function analysis (PFA) (see [Wood et al. 2000]). Principal function analysis is closely related to principal component analysis (PCA) with the important difference that functions f_m are optimized instead of vectors. Unfortunately, the PFA does not reduce to a simple eigenvalue problem as PCA does. To minimize $E_\pi(C)$, we again perform a least square optimization using the Levenberg-Marquardt method, this time fitting m BRDFs simultaneously. Within each iteration we recompute the projection f_{π_i} of lumitexel \mathcal{L}_i into the currently estimated basis.

As for every optimization problem the initial parameters (BRDFs) are quite important. For a given cluster C , we use the following BRDFs as a basis:

- f_C , the BRDF fit to the cluster C ,
- the BRDFs of spatially neighboring clusters to match lumitexels at cluster boundaries,
- the BRDFs of similar clusters with respect to the material, and
- two BRDFs based on f_C , one with slightly increased and one with decreased diffuse component ρ_d and exponent N .

In our experiments it turned out that this initial basis together with the projection already produces very good results with small errors. In most cases the PFA computed almost negligible changes to the initial BRDFs. This is to be expected because the initially chosen basis constructed through splitting and clustering already approximates the material properties quite well.

9. ACQUIRING NORMAL MAPS

The reconstructed spatially varying BRDFs can further be used to add geometric detail in the form of normal maps. The resolution of the acquired geometry of an object is typically limited by the applied 3D scanning device. Additional processing of the 3D data like

combining multiple scans, smoothing the surface to remove noise and mesh simplification to reduce the complexity of the model further erases fine scale geometric detail.

Although our projection method compensates somewhat for imprecise normals, one can observe in Figures 10 and 13 that much geometric detail was lost during the geometry acquisition, especially in the hair region.

Our method of measuring reflection properties as described in Section 8 can easily be extended to measure normal maps even for surfaces that are not perfectly diffuse.

In Section 5 the normal \vec{n} and a deterministically constructed tangent and bi-normal are used to transform the lighting and viewing direction, \vec{U} and \vec{V} from world coordinates into the local coordinate frame, yielding \vec{u} and \vec{v} at each surface point. In the case of isotropic materials the direction of the normal is sufficient to define this transformation since the BRDF is independent of the orientation of the tangent within the tangential plane by definition. Thus, the transformation into the local coordinate frame can be carried out by just two rotations about the y and z axis:

$$\vec{u} = R_y(-\alpha)R_z(-\beta)\vec{U}, \quad \text{and} \quad (14)$$

$$\vec{v} = R_y(-\alpha)R_z(-\beta)\vec{V}, \quad (15)$$

where α and β are azimuth and zenith of the normal.

As already mentioned in Section 5 an initial estimate of the normal at every surface point/lumitexel is provided by the triangle mesh. Based on these inexact normals the basic materials of the object are separated by the clustering process (Section 7). Subsequently, for each lumitexel the weighting coefficients for the basis materials are determined by projection to obtain starting values for these coefficients.

Given that enough reflectance samples are provided at every point (more than two) it is possible to extract the direction of the normal (α, β) for every lumitexel. This is done by minimizing the error between the measured reflectance samples r_i and the evaluated reflectance values

$$f_\pi(R_y(-\alpha)R_z(-\beta)\vec{U}_i, R_y(-\alpha)R_z(-\beta)\vec{V}_i), \quad (16)$$

for which once again the Levenberg-Marquardt algorithm is applied. Both the direction of the normal and the optimal weights for the basis BRDFs can be found in the same step. However, since the overall distribution of the normals is altered by computing normal maps one has to recompute the BRDF for the single clusters afterwards and has to perform the projection once again. If desired, the process of normal fitting, recomputation of the BRDF, and projection can be iterated. In our experience this is normally not necessary.

The quality of the results is presented in Figure 10 where normal maps have been recovered for all materials. Unfortunately, the method produces some artifacts in concave regions where interreflections become important (lower right of Figure 10). Since interreflections are not yet considered in our algorithm the normal directions are noisy and the BRDFs are not very accurate. A feature sensitive smoothing of the normals followed by a BRDF projection may slightly improve the results.

Since a non-linear optimization is performed for every lumitexel the recovery of the normal map is a time-consuming step. For the bird model it took around three hours but was completely automatic.




Fig. 10. *Left: Normals of the original mesh. Right: Normals optimized using spatially varying BRDFs*

10. RENDERING

As explained in Section 5.1, we know the position of every lumitexel, as well as the triangle it belongs to and the 2D coordinates within that triangle.

This information can then be used to generate an index texture for the full object. For every texel, that texture contains an index to the cluster it belongs to. Then we generate a weight texture map for every cluster that stores the weights resulting from the projection into the basis BRDFs. The parameters for the basis BRDFs of every cluster are stored in a small table. Additionally, we may have a normal map that was reconstructed using the technique from Section 9.

Raytracing such an object is very simple, since for every point on the object that is raytraced we can simply look up the cluster the texel belongs to. Then we evaluate the basis BRDFs for the local light and viewing direction and compute the weighted sum using the weight texture map. So rendering simply is a matter of evaluating a few Lafortune BRDFs per pixel, and weighting the results.

If no normal map was reconstructed, mip-mapping can easily be used. Since the weighted sum is just a linear operation, the weights of neighboring texels can simply be averaged to generate the next coarser mip-map level.

Another, more accurate method is to explicitly construct lumitexels at coarser mip-map levels that combine the reflectance samples of the finer ones. The linear weights are computed to best fit these larger collections of radiance samples. This also allows to reconstruct normal maps for every mip-map level.

If the original images are of high resolution and hence the object is sampled very densely, point-based rendering using forward projection is a viable alternative. It completely avoids the generation of texture maps and the resulting data can be used with almost no further processing. This method is used to display our results.

11. RESULTS

We applied our algorithm to four different objects consisting of different materials with varying reflection properties in both the diffuse and the specular part. The model of the angels was generated by extracting an isosurface of a computer tomography scan. The geometry of all other models was captured using a structured light 3D scanner. Some statistics about the meshes and the number of acquired views are listed in Table I. Acquisition of 20 views (each needing about 15 photographs) takes approx. 2.5h. This is the only step that requires user input. The high dynamic range conversion and the registration with the 3D model takes about 5h but is a completely automated task. The clustering and the final projection takes about 1.5h for all models, and is again automatic. An additional 3h (angels), 3h (bird), 4h (bust), 2.5h (elk) were needed for the normal fitting. While timings for the conversion is dependent on the number of input images, the clustering time mainly depends on the number of texels in the selected subset. The remaining steps of projection and normal fitting depend on the number of overall reflectance samples.

In Figure 8 you can see how five successive split operations partition the lumitexels of the bird into its five materials. The splits were performed as described in Section 7. Only the per-cluster BRDFs determined by the clustering process are used for shading, making the object look rather synthetic. After performing the projection step every lumitexel is represented in a basis of four BRDFs, now resulting in a much more detailed and realistic appearance, see Figure 9.

The bust in Figure 11 shows another reconstructed object with very different reflection properties. The bronze look is very well captured.

Another model is shown in Figure 12 where the reconstruction of a wooden elk using 25 views is compared to an actual photograph. The registration worked well enough to reproduce the fine detail wooden structure. Unfortunately, the bad quality of the geometric model caused some noticeable differences at the rim of the antler. It also causes the darker stripe right behind the eye, where applying a 10x10 filter to the shadow map could not remove enough points around the shadow edge. These artifacts will disappear if a better quality geometry models is available. Due to the slightly more complex geometry 25




Fig. 11. A bronze bust rendered with a spatially varying BRDF acquired with our reconstruction method (without normal fitting). The reconstructed BRDFs allow for rendering the model with arbitrary lighting.



Fig. 12. A comparison of a photograph (left) of a wooden elk with the reconstructed model (right). Four clusters have been generated which faithfully reproduce the appearance. Note how the wooden structure is preserved.

model	T	V	L	R	C	B
angels	47000	27	1606223	7.6	9	6
bird	14000	25	1917043	6.3	5	4
bust	50000	16	3627404	4.2	3	4
elk	50000	25	1659945	5.4	4	4

Table I. This table lists the number of triangles (T) of each model, the number of views (V) we used to reconstruct the spatially varying BRDFs, the number of acquired lumitexels (L) and the average number of reflectance samples (R) per lumitexel, the number of partitioned material clusters (C), and the number of basis BRDFs (B) per cluster.

model	1-RMS	C-RMS	P-RMS	F-RMS	NF-RMS
angels	.2953	.1163	.1113	.1111	0.0703
bird	.1513	.0627	.0387	.0387	0.0269
bust	.1025	.0839	.0583	.0581	0.0113
elk	.0768	.0556	.0275	.0274	0.0202

Table II. This table lists for each model the RMS error for a single average BRDF (1-RMS), the RMS error when using per-cluster BRDFs, the RMS error after projecting every lumitexel into the basis of BRDFs, the RMS error after doing a PFA on the basis BRDFs and projecting every lumitexel into the new basis, and finally the RMS error after fitting the normals.

different views are actually not enough to cover the entire surface. This is why some holes appear close to the wheels. Clearly, more images would remove these problems.

In Figure 13 you can see another comparison between an object rendered with an acquired BRDF (using the projection method) and a photograph of the object. They are very similar, but differences can be seen in highlights and in places where not enough reflectance samples were captured. Capturing more samples will increase the quality. The difference in the hair region in the right picture is due to missing detail in the triangle mesh. This detail can be recovered by computing normal maps as explained in Section 9, the result is shown in Figure 13 bottom.

Another difference is due to the fact that the diffuse color of one lumitexel may not be represented in any of the constructed clusters because the number of lumitexels belonging to the same material can be so small that they nearly vanish in the mass of lumitexels of the cluster they are currently assigned to. This effect can for example be observed at the mouth of the larger angel which in reality exhibits a much more saturated red, see Figure 13.

In Table II we list RMS errors computed between all the reflectance samples of a model and the reconstructed BRDFs. You can see that the error considerably decreases when going from one average BRDF to per-cluster BRDFs and then to per-pixel BRDFs (using projection). As already mentioned the PFA only slightly changes the RMS error.

Generally, it can be said that for all the models only a few clusters were needed to accurately represent all the materials since the projection takes care of material changes. In our experiments even Lafortune BRDFs consisting of a single lobe were sufficient to form good bases for the clustering and projection.

The fitting of normal maps works very well using the reconstructed BRDFs. Fine geometric detail could be recovered and the quality of the overall model was increased even for non-Lambertian surfaces. This is also indicated by the RMS error in Table II which has been noticeably reduced by the normal fitting step.




Fig. 13. Left side: Photograph of model. Right side: Model with acquired BRDF rendered from the same view with similar lighting direction. The difference in the hair region is due to missing detail in the triangle mesh. Bottom: After computing a normal map the missing detail has been recovered. Note how the highlights around the Christmas tree and on the left wing matches the original.

One can however observe noise in regions where too few reflectance samples or views have been acquired. Although three reflectance samples should be enough to determine the normal direction in the ideal case, more samples have to be provided to get reliable results.

Furthermore, in concave regions of materials with high reflectance interreflections become very important. Since interreflections are not yet considered in our algorithm the presented method unfortunately produces noisy normals and inaccurate BRDFs in those regions. This problem may be solved by applying techniques capable of dealing with in-

terreflections such as [Nayar et al. 1990; 1991] or [Forsyth and Zisserman 1989].

Due to the lack of a test object that had a single base color but varying specularity, we experimented with artificially generated data. We generated samples for five different Lafortune BRDFs with the same

$$k_d = (0.15, 0.3, 0.1), \text{ and different exponents}$$

$$N = 36.2, 42.2, 48.2, 54.2, 60.2,$$

where $-C_x = -C_y = C_z = \sqrt[N]{(N+2)/2\pi}$ corresponding to the modified Phong model [Lewis 1993]. To each reflectance sample we added up to 5% noise. Overall there were about 46000 reflectance samples (on average 11 per texel) with random lighting directions and five different viewing directions. Our clustering algorithm was able to clearly distinguish these five materials although they had the same color but different specular lobes. The resulting k_d 's and exponents had less than 0.2% error.

Reliable classification of these materials requires however that part of a highlight is visible in at least one of the reflectance sample at every surface point. Depending on the narrowness of the specular lobe and the geometry of the object, a lot of images may be required to achieve this for real objects. For example, on the angels model the specular part of the blue skirt was overestimated at some texels due to undersampled highlights.

12. CONCLUSIONS AND FUTURE WORK

We have presented an algorithm and demonstrated a system for reconstructing a high-quality spatially varying BRDF from complex solid objects using only a small number of images. The same input images are also used to improve geometric detail. Combining the resulting optical and geometric data allows for accurately shaded, photorealistic rendering of these objects from new viewpoints and under arbitrary lighting conditions.

Several objects consisting of different materials have been acquired to demonstrate the quality and accuracy of our approach. The resulting spatially varying BRDFs accurately represent the original materials. The normal fitting algorithm recovered geometric detail that was not represented by the original mesh obtained from the 3D scanner.

The reconstruction of a new model is a relatively simple task and requires only a moderate amount of human work, mainly during the acquisition process. Except for this, all the data processing and fitting algorithms are automatic. Fortunately, the number of input views required by our algorithm is rather small.

Compared to previous approaches for representing real-world objects, like surface light fields or reflection fields which needed up to 600 images [Wood et al. 2000] our method requires less input data and even the size of the output data ($\sim 25\text{MB}$) is considerably smaller. All we have to store per texel are the linear blending weights for the basic materials.

Representing the spatially varying BRDFs as texture maps allows to modify the object's geometry after the acquisition. Since the BRDF is not changed with the geometry the object can be altered or animated while preserving the material properties and thus the realistic appearance.

To speed up the rendering time of objects with spatially varying BRDFs we want to investigate the possibility to do hardware acceleration. Here, it is again useful to represent the recovered data as texture maps. Using techniques from [Kautz and Seidel 2000] or from [McAllister et al. 2002] the rendering should be easily implemented on modern graphics hardware since the Lafortune model is fairly simple to evaluate.

13. ACKNOWLEDGMENTS

We would like to thank Kolja Kähler, Christian Rössl, Mario Botsch and the IMP Erlangen for acquiring the 3D meshes used in this paper and Nico Schäfer for acquiring some of the BRDF data. Thanks also to Hartmut Schirmacher for proofreading and valuable discussion. This work was in part supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the BC Advanced Systems Institute, and the Institute for Robotics and Intelligent Systems (IRIS). Part of this work has been funded by the DFG Schwerpunktprogramm V3D2, an NVidia fellowship, and by the European Community within the scope of the ViHAP3D Project IST-2001-32641 “Virtual Heritage: High-Quality 3D Acquisition and Presentation.”

REFERENCES

- BERNARDINI, F., MARTIN, I. M., AND RUSHMEIER, H. 2001. High-quality texture reconstruction from multiple scans. *IEEE Transactions on Visualization and Computer Graphics* 7, 4 (October - November), 318–332. ISSN 1077-2626.
- BOIVIN, S. AND GAGALOWICZ, A. 2001. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Proceedings of SIGGRAPH 2001*, E. Fiume, Ed. Computer Graphics Proceedings, Annual Conference Series. ACM Press / ACM SIGGRAPH, 107–116. ISBN 1-58113-292-1.
- DANA, K., VAN GINNEKEN, B., NAYAR, S., AND KOENDERINK, J. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics* 18, 1 (January), 1–34.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the Reflectance Field of a Human Face. In *Proc. SIGGRAPH*. 145–156. ISBN 1-58113-208-5.
- DEBEVEC, P. AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. SIGGRAPH*. 369–378.
- DEBEVEC, P., TAYLOR, C., AND MALIK, J. 1996. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *Proc. SIGGRAPH*. 11–20.
- FORSYTH, D. A. AND ZISSERMAN, A. 1989. Mutual Illumination. In *Proceedings of Computer Vision and Pattern Recognition (CVPR '89)*. IEEE Computer Society Press, 466–473.
- FURUKAWA, R., KAWASAKI, H., IKEUCHI, K., AND SAKAUCHI, M. 2002. Appearance based object modeling using texture database: Acquisition compression and rendering. In *Thirteenth Eurographics Workshop on Rendering*. 267–276.
- GARLAND, M. AND HECKBERT, P. August 1997. Surface Simplification Using Quadric Error Metrics. In *Proc. SIGGRAPH*. 209–216.
- GERSHO, A. AND GRAY, R. 1992. *Vector Quantization and Signal Compression*. Kluwer Acad. Publishers.
- GIBSON, S., HOWARD, T., AND HUBBOLD, R. 2001. Flexible image-based photometric reconstruction using virtual light sources. *Computer Graphics Forum* 20, 3. ISSN 1067-7055.
- GOESELE, M., HEIDRICH, W., LENSCH, H. P., AND SEIDEL, H.-P. 2000. Building a photo studio for measurement purposes. In *Proceedings of Vision, Modeling, and Visualization*. 231–238.
- GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Entropy based dark frame subtraction. In *Image Processing, Image Quality, Image Capture Systems Conference (PICS)*. 293–298.
- GORTLER, S., GRZESZCZUK, R., SZELINSKI, R., AND COHEN, M. 1996. The Lumigraph. In *Proc. SIGGRAPH*. 43–54.
- KAUTZ, J. AND SEIDEL, H.-P. August 2000. Towards Interactive Bump Mapping with Anisotropic Shift-Variant BRDFs. In *Eurographics/SIGGRAPH Hardware Workshop*. 51–58.
- KOBELT, L. 1996. Discrete fairing. In *Proc. of the 7th IMA Conf. on the Mathematics of Surfaces*. 101–131.
- KOENDERINK, J., VAN DOORN, A., AND STAVRIDIS, M. 1996. Bidirectional Reflection Distribution Function expressed in terms of surface scattering modes. In *Proc. 4th Europ. Conf. on Computer Vision*. 28–39.
- LAFORTUNE, E., FOO, S., TORRANCE, K., AND GREENBERG, D. 1997. Non-Linear Approximation of Reflectance Functions. In *Proc. SIGGRAPH*. 117–126.
- LENSCH, H., HEIDRICH, W., AND SEIDEL, H.-P. 2000. Automated Texture Registration and Stitching for Real World Models. In *Pacific Graphics '00*. 317–326.
- LENSCH, H., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Silhouette-based algorithm for texture registration and stitching. *Graphical Models* 63, 4 (July), 245–262.
- LENSCH, H. P. A., KAUTZ, J., GOESELE, M., HEIDRICH, W., AND SEIDEL, H.-P. 2001. Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*. Eurographics, 103–114. ISBN 3-211-83709-4.
- LEVOY, M. AND HANRAHAN, P. 1996. Light Field Rendering. In *Proc. SIGGRAPH*. 31–42.
- LEWIS, R. 1993. Making shaders more physically plausible. In *4th Eurographics Workshop on Rendering*. 47–62.
- LIU, X., YU, Y., AND SHUM, H.-Y. 2001. Synthesizing bidirectional texture functions for real-world surfaces. In *Proceedings of ACM SIGGRAPH 2001*. Computer Graphics Proceedings, Annual Conference Series. ACM Press / ACM SIGGRAPH, 97–106. ISBN 1-58113-292-1.
- LLOYD, S. 1982. Least squares quantization in PCM. *IEEE Trans. on Information Theory* IT-28, 129–137.

- LU, R., KOENDERINK, J., AND KAPPERS, A. 1998. Optical Properties (bidirectional reflectance distribution functions) of velvet. *Applied Optics* 37, 25 (Sept.), 5974–5984.
- MACQUEEN, J. 1967. Some methods for classification and analysis of multivariate observations. In *Proc. of the 5th Berkeley Symp. on Mathematical Statistics and Probability*. Vol. 1.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. In *Proceedings of ACM SIGGRAPH 2001*. Computer Graphics Proceedings, Annual Conference Series. ACM Press / ACM SIGGRAPH, 519–528. ISBN 1-58113-292-1.
- MARSCHNER, S. 1998. Inverse rendering for computer graphics. Ph.D. thesis, Cornell University.
- MARSCHNER, S., WESTIN, S., LAFORTUNE, E., TORRANCE, K., AND GREENBERG, D. 1999. Image-based BRDF Measurement Including Human Skin. In *10th Eurographics Workshop on Rendering*. 131–144.
- MATUSIK, W., PFISTER, H., NGAN, A., BEARDSLEY, P., AND MCMILLAN, L. 2002. Image-Based 3D Photography Using Opacity Hulls. In *Proceedings SIGGRAPH*. 427–437.
- MATUSIK, W., PFISTER, H., ZIEGLER, R., NGAN, A., AND MCMILLAN, L. 2002. Acquisition and Rendering of Transparent and Refractive Objects. In *Thirteenth Eurographics Workshop on Rendering*. 277–288.
- MCALLISTER, D. 2002. A Generalized Representation of Surface Appearance. Ph.D. thesis, University of North Carolina.
- MCALLISTER, D., LASTRA, A., AND HEIDRICH, W. 2002. Efficient Rendering of Spatial Bi-directional Reflectance Distribution Functions. In *Proceedings Graphics Hardware*.
- MILLER, G., RUBIN, S., AND PONCELEON, D. 1998. Lazy decompression of surface light fields for precomputed global illumination. In *9th Eurographics Workshop on Rendering*. 281–292.
- NAYAR, S., IKEUCHI, K., AND KANADE, T. 1991. Surface reflection: physical and geometrical perspectives. *IEEE Trans. PAMI* 13, 7, 611–634.
- NAYAR, S. K., IKEUCHI, K., AND KANADE, T. 1990. Shape from Interreflections. In *International Conference on Computer Vision (ICCV '90)*. Osaka, Japan, 2–11.
- PRESS, W., TEUKOLSKY, S., VETTERLING, W., AND FLANNERY, B. 1992. *Numerical Recipes in C: The Art of Scientific Computing (2nd ed.)*. Cambridge University Press. ISBN 0-521-43108-5.
- RAMAMOORTHI, R. AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proceedings of SIGGRAPH 2001*, E. Fiume, Ed. Computer Graphics Proceedings, Annual Conference Series. ACM Press / ACM SIGGRAPH, 117–128. ISBN 1-58113-292-1.
- ROBERTSON, M. A., S., AND STEVENSON, R. L. 1999. Dynamic Range Improvement Through Multiple Exposures. In *Proc. of the Int. Conf. on Image Processing (ICIP'99)*. IEEE, 159–163.
- RUSHMEIER, H., BERNARDINI, F., MITTLEMAN, J., AND TAUBIN, G. 1998. Acquiring input for rendering at appropriate levels of detail: Digitizing a pietà. *Eurographics Rendering Workshop 1998*, 81–92. ISBN 3-211-83213-0. Held in Vienna, Austria.
- RUSHMEIER, H., TAUBIN, G., AND GUÉZIEC, A. 1997. Applying Shape from Lighting Variation to Bump Map Capture. In *8th Eurographics Workshop on Rendering Workshop*. 35–44.
- SATO, Y., WHEELER, M., AND IKEUCHI, K. August 1997. Object Shape and Reflectance Modeling from Observation. In *Proc. SIGGRAPH*. 379–388.
- SCHIRMACHER, H., HEIDRICH, W., RUBICK, M., SCHIRON, D., AND SEIDEL, H.-P. 1999. Image-Based BRDF Reconstruction. In *Proc. of the 4th VMV Conference*. 285–292.
- TORRANCE, K. AND SPARROW, E. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America* 57, 9.
- WARD LARSON, G. 1992. Measuring and Modeling Anisotropic Reflection. In *Proc. SIGGRAPH*. 265–272.
- WESTIN, S., ARVO, J., AND TORRANCE, K. 1992. Predicting Reflectance Functions From Complex Surfaces. In *Proc. SIGGRAPH*. 255–264.
- WOOD, D., AZUMA, D., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D., AND STUETZLE, W. 2000. Surface Light Fields for 3D Photography. In *Proc. SIGGRAPH*. 287–296.
- YU, Y., DEBEVEC, P., MALIK, J., AND HAWKINS, T. 1999. Inverse Global Illumination: Recovering Reflectance Models of Real Scenes From Photographs. In *Proc. SIGGRAPH*. 215–224.
- YU, Y. AND MALIK, J. July 1998. Recovering Photometric Properties of Architectural Scenes from Photographs. In *Proc. SIGGRAPH*. 207–218.
- ZHANG, R., TSAI, P.-S., CRYER, J., AND SHAH, M. 1999. Shape from Shading: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21, 8 (August), 690–706.

ZHANG, Z. 1999. Flexible Camera Calibration By Viewing a Plane From Unknown Orientations. In *Int. Conf. on Computer Vision*. 666–673.

Received Month Year; revised Month Year; accepted Month Year