

Adaptive Sampling of Reflectance Fields

MARTIN FUCHS

MPI Informatik

VOLKER BLANZ

Universität Siegen

and

HENDRIK P. A. LENSCH and HANS-PETER SEIDEL

MPI Informatik

Image-based relighting achieves high quality in rendering, but it requires a large number of measurements of the reflectance field. This article discusses sampling techniques that improve on the trade-offs between measurement effort and reconstruction quality.

Specifically, we (i) demonstrate that sampling with point lights and from a sparse set of incoming light directions creates artifacts which can be reduced significantly by employing extended light sources for sampling, (ii) propose a sampling algorithm which incrementally chooses light directions adapted to the properties of the reflectance field being measured, thus capturing significant features faster than fixed-pattern sampling, and (iii) combine reflectance fields from two different light domain resolutions.

We present an automated measurement setup for well-defined angular distributions of the incident, indirect illumination. It is based on programmable spotlights with controlled aperture that illuminate the walls around the scene.

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

General Terms: Measurement

Additional Key Words and Phrases: Image-based relighting, reflectance fields

ACM Reference Format:

Fuchs, M., Blanz, V., Lensch, H. P. A., Seidel, H.-P. 2007. Adaptive sampling of reflectance fields. *ACM Trans. Graph.* 26, 2, Article 10 (June 2007), 18 pages. DOI = 10.1145/1243980.1243984 <http://doi.acm.org/10.1145/1243980.1243984>

1. INTRODUCTION

In measurement-based rendering, reflectance field techniques have provided relightable object representations of a remarkable quality. Especially if the viewing position remains fixed and the target illumination is restricted to a distant environment map, these methods have been shown to be very powerful. Due to the principle of linear superposition of light, linear combinations of the input images can simulate linear combinations of lighting environments realistically even with simple interpolation techniques.

For faithful reproduction, however, objects with complex appearance require a huge number of input images, ranging up to the tens of thousands. Existing approaches have been able to achieve this by means of dense fixed-pattern sampling, either in light source positions as in traditional lightstage approaches, or in pixel positions as in the Dual Lightstage method recently presented by Hawkins et al. [2005].

In this article, we address techniques that reduce the measurement effort. With traditional point light setups, reducing the number of illumination measurements would produce aliasing artifacts due to the gaps in the sampled illumination. Those artifacts occur at shadow boundaries (Figure 4) or on shiny objects where the original light sources can be seen as mirror reflections (Figure 3). We show how this can be avoided by using extended light sources with smooth fall-offs as input illuminants (Section 7.1) at the cost of fidelity in high frequency illumination.

Another way to reduce the effort of measurements while retaining high frequency reflections is to distribute the sample directions in an irregular pattern that is adapted to the sensitivity of the reflectance function to changes in incident light direction. We present an adaptive algorithm for this in Section 8 and perform a quantitative and qualitative analysis of its benefits in Section 8.3. Our algorithm starts with a coarse set S_0 of regularly distributed illumination directions (θ_i, ϕ_i) on the upper hemisphere. For each direction, an image

This work has been partially funded by the Max Planck Center for Visual Computing and Communication (BMBF-FKZ011MC01).

Author's address: M. Fuchs, MPI Informatik; email: mfuchs@mpi-inf.mpg.de

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2007 ACM 0730-0301/2007/06-ART10 \$5.00 DOI 10.1145/1243980.1243984 <http://doi.acm.org/10.1145/1243980.1243984>

I_i is recorded. Then, in each iteration j , we analyze the observed slices of the reflectance field and augment S_j by several new sample directions in those directions where the reflectance field is not smooth. The iteration terminates at a fixed number of samples that is specified by the user.

The adaptive sampling technique may still produce the artifacts that we mentioned if the observed object has small, but shiny surfaces that contribute little to the overall error. Therefore, we propose to combine adaptive sampling at narrow-aperture lighting with sparsely-sampled extended light source conditions (Section 9). This technique creates a reflectance field with two levels of physical prefiltering. Sampling theory motivates a discussion of the point distributions generated by the adaptive sampling algorithm that is provided in Section 10.

The implementation of these techniques requires a measurement setup which is able to create light distributions of variable extent at a high directional resolution and in variable sampling densities. Our solution for this is described in Section 5. It consists of a tent of dark cloth in the form of a hemicube which is illuminated from the inside by two projection systems with programmable direction and aperture. Their light indirectly illuminates the scene that we are recording. In contrast to approaches based on point lights, this setup creates extended light sources that can be controlled in size. Fully automated, it can acquire reference data at thousands of illumination directions. We use these data for quantitative evaluations.

The rest of the article is organized as follows. After some definitions of terminology in Section 2, we discuss related work in 4D reflectance fields in Section 3 and address aspects of sampling theory in Section 4. We will then turn to our measurement setup in Section 5 and proceed with rendering aspects in Section 6. Section 7 shows implementation details and results for fixed pattern sampling, Section 8 presents and evaluates our adaptive sampling algorithm, and Section 9 describes a simple method for combining reflectance fields in different sampling resolutions. Section 10 contains a discussion of the achieved sampling distributions in terms of sampling theory and lists limitations of this work. Concluding remarks are found in Section 11.

2. REFLECTANCE FIELDS

Image-based relighting methods capture the optical properties of scenes by acquiring a set of input images in varying light conditions. Because they are data-driven, they can achieve high fidelity without strong model assumptions. Due to the superposition principle of light, the appearance of a scene is completely determined by the light transport between any pair of incoming and reflected light rays.

The light exiting any static scene which is nonfluorescent, non-phosphorescent, does not emit light, and will only be observed from the outside, can be described by an 8D function per-color channel

$$R_{4 \rightarrow 4} : \mathbb{R}^8 \rightarrow \mathbb{R}_0^+, \quad (1)$$

which maps each pair of light rays ($4D \times 4D$) to the corresponding apparent reflectance. This function is called the *8D reflectance field* of the scene [Debevec et al. 2000]. It expresses the relationship between incident and reflected light fields [Levoy and Hanrahan 1996].

While a first approach for efficiently measuring 8D reflectance fields has recently been proposed [Garg et al. 2006], an exhaustive coverage of incident light situations has still not been achieved due to the immense amount of capture effort and data storage required. Often, the dimensionality of the problem is therefore reduced by simplifying assumptions.

If the complexity of the incident light field is reduced to two dimensions, for example, by assuming distant light sources (environment map), a *6D reflectance field* $R_{2 \rightarrow 4}$ is obtained [Matusik et al. 2002]. Such 6D reflectance fields can be parameterized along explicit surfaces [Dana et al. 1999; Weyrich et al. 2005]

Alternatively, a fixed observer can be assumed, reducing the outgoing light field to two dimensions. For a pin-hole camera model, a 2D subset of rays is selected. For more general models, a convolution takes place, which can be subsumed into the reflectance field. If the incident light field still has 4 dimensions, the reflectance field $R_{4 \rightarrow 2}$ again is six-dimensional [Masselus et al. 2003; Sen et al. 2005], if a 2D incident light field is assumed, the light transport is described by the classical *4D reflectance field* $R_{2 \rightarrow 2}$ as originally used by Debevec et al. [2000].

In the following, we discuss 4D reflectance fields. We consider a polar parameterization of the incoming light domain along a rotation angle $\phi \in [0; 2\pi)$ and an elevation angle $\theta \in [0; \pi)$ which equals 0 at the upper pole. Then, images $I(x, y)$ observed by this camera can for any incident light distribution L be described as

$$I(x, y) = \int_0^\pi \int_0^{2\pi} L(\theta, \phi) \cdot R(x, y, \theta, \phi) d\phi \sin \theta d\theta. \quad (2)$$

Then, the 4D image-based relighting problem has two steps.

- (1) *Acquisition*. Given a set of n input images $I_i(x, y)_{0 \leq i < n}$ in light situations L_i , construct a model for $R(x, y, \theta, \phi)$.
- (2) *Rendering*. Given a novel light situation L_{target} , render an image I_{predict} of the scene using the model from the acquisition step. If L_{target} is a linear combination $L_{\text{target}} = \sum_i k_i L_i$ of the light situations used during acquisition, I_{predict} can be computed as $I_{\text{predict}} = \sum_i k_i I_i$ (Linearity of light transport).

3. RELATED WORK

4D reflectance field methods can be categorized according to their model of the reflectance field R , and the restrictions on the acquisition step.

3.1 Work on 4D Reflectance Fields

The most common approach for capturing such 4D reflectance fields for a fixed view point and distant light sources are light stages [Debevec et al. 2000] where a point light source is rotated around the object. Mounting switchable point light sources to fixed positions enables high speed capture [Wenger et al. 2005; Einarsson et al. 2006]. Wenger et al. [2005] also discuss illumination multiplexing [Schechner et al. 2003] to improve the signal-to-noise ratio (SNR). The use of handheld light sources [Masselus et al. 2002] reduces setup complexity at the cost that the measurements cannot be repeated. Unfortunately, if only a small number of point light sources are used to populate the sphere of incident directions, gaps will remain. The resulting reflectance field is prone to aliasing (see Section 7) since point light sources provide poor prefiltering. For this category of reflectance fields models, Masselus et al. [2004] explored interpolation algorithms and compression techniques.

Hawkins et al. [2005] applied Helmholtz reciprocity to the problem. They virtually exchanged the roles of camera and incident light domain, observing a dense sampling of the (θ, ϕ) domain of incident light at the cost of requiring one input image for each pixel position (x, y) in the image domain which limits the resolution in image space.

Exploiting the linearity of light transport, some approaches illuminate the scene with explicit basis functions with extended support.

Generalizing environment mattes [Chuang et al. 2000; Wexler et al. 2002], Peers and Dutré [2003] used adaptively refined wavelets as basis functions and achieve dense angular sampling. The currently employed CRT or TFT monitors do not scale well for larger objects. Distant illumination from the entire sphere of incident light directions is also hard to provide with sufficient SNR even if a cave was used.

The methods discussed have in common that R is represented in each pixel by the same set of basis functions. Matusik et al. [2004] express R independently in each pixel using block-wise constant kernels. Their images are from uncontrolled, but known, illumination conditions. Peers and Dutré [2005] have further improved a per-pixel representation of R : they emit controlled wavelet noise patterns from a monitor, and then fit a wavelet basis to their observed images. This is done on each pixel separately. They obtain a high angular resolution from a moderate number of input images at the cost of a low SNR in dark regions.

Fuchs et al. [2005] have modeled R implicitly, creating rendered images directly as linear combinations of the input images. L_{target} is captured implicitly, too, using the image of a reference object in novel lighting. While the calibration requirements are weak, the limited precision due to displacement of the calibration object makes this approach more useful for low-frequency target illuminations.

3.2 Adaptive Sampling in Other Contexts

In this work, we observe the reflectance field using light sources with varying extent and adaptively refine the set of sampling positions. Related adaptive techniques have been employed in other image-based applications.

Adaptive sampling of the incident light direction in the context of measuring individual homogeneous BRDF has been proposed by Matusik et al. [2003] based on a principal component analysis of a large collection of captured BRDFs. Lensch et al. [2003] computed the next best viewing and lighting directions for capturing spatially varying BRDFs of 3D objects by analyzing the covariance matrices of the parameters obtained by fitting an analytic BRDF model to the current set of measurements of each surface point.

Peers and Dutré [2003] have used adaptively refined wavelets. Instead of capturing a single basis image, Sen et al. [2005] have emitted several illumination patterns simultaneously, parallelizing the acquisition.

Our adaptive scheme is inspired by the work by Schirmacher et al. [1999] who proposed an adaptive light field acquisition method that adaptively subdivides a mesh of light view positions according to an image-space error metric.

4. REFLECTANCE FIELDS AND SAMPLING THEORY

While sampling-based reflectance models discussed in Section 3.1 are robust and achieve a high quality, they require large numbers of input images even for low-frequency target illuminations. This is a result of two problems: gap-free coverage of the incident light domain and aliasing artifacts due to undersampling.

In this article, we will address both by the way our input light sources are constructed. The coverage problem can be mitigated by using extended light sources the aliasing problem can be reduced by using light sources with a smooth fall-off. This is actually a consequence of the results of sampling theory, centered around the sampling theorem by Shannon [1949] which states that a function, which does not contain frequencies over a frequency limit, can be completely reconstructed if sampled at twice the frequency of the

limit. If fewer samples are available, the reconstruction of the function induces aliasing artifacts.

Chai et al. [2000] have discussed optimal reconstruction kernels for light fields of diffuse scenes based on the depth complexity. Reflectance fields, however, can contain arbitrary high frequencies. In order to suppress aliasing artifacts, prefiltering at acquisition time is required as we demonstrate in Figure 3.

One way to achieve such a prefiltering can be deduced from Equation (2). It resembles a convolution, which corresponds to a multiplication in Fourier space. Accordingly, high-frequency content of the reflectance field can physically be suppressed using extended light conditions L_i with smooth fall-offs for its measurement. Thus, while we cannot influence the frequency content of the reflectance field itself, prefiltering can be established through the choice of the illuminations L_i . By this, we limit the maximum frequency in each illumination pattern, and thus observe an approximation of R which is low-pass filtered in the (θ, ϕ) domain (compare to Ramamoorthi and Hanrahan [2001]).

Given the maximum frequency of the input illumination, sampling theory provides a minimum required sampling density for which aliasing-free reconstruction of R can be guaranteed. Conversely, given a fixed sampling density, the corresponding filter kernel can be deduced. Thus, there exists a trade-off between measurement effort (sampling density) and rendering fidelity (reconstructible frequencies).

We exploit this trade-off concentrating the measurement effort on the light directions to which the reflectance field is most sensitive. Our algorithm adaptively refines the distribution of samples (Section 8) according to the local properties of the observed reflectance field. We then combine narrow and extended light sources depending on the local sampling density (Section 9). We approximate an optimal extended light source with indirect lighting (Section 5) and discuss sampling densities in relation to its properties in Section 10.

5. MEASUREMENT SETUP

5.1 Construction

The previous discussion motivates the creation of a measurement setup that supports sampling at variable densities and with overlapping light distributions which creates distant lighting without gaps. We achieve this by shining light on a diffuse enclosure, using the bounce-light for the illumination of the scene we want to measure, as done by Fuchs et al. [2005]. In contrast to their technique, however, we propose a reproducible and automated setup.

We therefore employ so-called moving-head controllable spotlights, which is inspired by Mohan et al. [2005] who, instead of capturing a reflectance field, created synthetic illuminations so that the appearance of a real-world object approximates a user-specified image.

By using two spotlight emitters, our setup covers a large part of the hemisphere. We put the object that we want to measure in the center of the room and illuminate it indirectly by light which bounces off the enclosing tent, creating distant light conditions. As light sources, we use two JB Lighting Varyscan P3 moving-head spotlights containing 250 Watts HTI light sources at a distance of 65cm from the object.

Our setup, as displayed in Figure 1, is built inside a tent of dark, uniform cloth, which covers an area of $4\text{m} \times 3.8\text{m}$ and has a height of 1.9m. The floor is covered by dark carpet. We used dark surfaces, as opposed to the white walls used in previous methods, in order to reduce the power of higher-order bounce light relative to the power of the first-order bounce light used for measurement.



Fig. 1. Photograph of our measurement setup. Spotlight projectors on the floor shine on a tent made from black cloth, and the bounce light illuminates the scene which is arranged on top of the boxes. A tripod-mounted camera records the scene and is protected against lens flares by occluders.

The spotlights can be controlled using the DMX protocol. They can be rotated over 430 degrees pan and 270 degrees tilt, have a programmable iris, shutter, focus and frost effect, which allows us to control the light position in the room, its extent, and the smoothness of the fall-off of the spotlight. The diameter of the emitted light cone can be varied from 4.3° at maximally closed aperture up to 18° at fully open aperture. Due to the fact that the spotlight is located next to the object, the size of the incoming light at the object varies between 3.1° and 4.3° for closed aperture, and between 13° and 18° for open aperture. Figure 2 shows some of the light conditions that can be achieved. In our experiments, we used two settings for the light sources: the wide setting with the iris fully open and the frost effect filter applied, and the setting where the iris was 30% opened and the frost effect switched off.

For image capture, we run HDR sequences [Robertson et al. 2003] on Jenoptik CFcool or C14plus cameras. We choose the length of the HDR sequences depending on the aperture of the spotlights. Typically, we use four exposures with durations of 20, 100, 500, and 1000ms, respectively, if the aperture is fully open and the frost filter is enabled. Choosing the exposure times as multiples of 20ms cancels out 50Hz flickering in the light sources of the spotlights.

5.2 Calibration

In order to apply the setup to reflectance field measurements, we need to determine the relationship between projector pan/tilt settings and light directions, as seen from the object. Also, as our tent is nonspherical, the total brightness of the bounce light arriving at the object depends on the light direction (it is darker in the corners of the tent) and needs to be taken into account. The influence of pan/tilt on the exact size and shape of the light distribution could also be measured, but, in our setup, we may assume that it has no significant effect.

For the calibration, we place a camera with a fish eye lens at the center of the measurement tent, facing the upper hemisphere.

Methods for its precise geometric [Kannala and Brandt 2004] and vignetting [Goldman and Chen 2005] calibration are available; for our experiments, we only performed a manual calibration of the most important parameters for the mapping of environment map directions to camera pixels. As long as the same lens is used for the acquisition the target environment maps, most geometric and vignetting effects will cancel out.

The mapping between pan and tilt positions of the projectors to (θ, ϕ) is obtained by rotating the projector in steps of 10 degrees and recording images of the bounce light from the black cloth. In the captured images, we can compute the average direction of all pixels which have a brightness above a noise threshold, and thus obtain an estimate for the incoming light direction which corresponds to the selected pan/tilt settings. Linear interpolation between these values makes it possible to estimate for any given (θ, ϕ) the corresponding pan/tilt configuration for each of the spotlights.

The nonspherical structure of our measurement environment makes it necessary to compensate for different brightnesses of the bounce light. Photometric calibration is performed by establishing the camera's response curve and measuring the white point. The observed brightness in the calibration images (sum of pixel intensities that are above the noise threshold) gives a relative brightness scaling factor for each direction.

During calibration, we also infer which of the two spotlights is preferred for each (θ, ϕ) , based on the spacing of (θ, ϕ) samples in the calibration images. Due to the field of view of the fish-eye lens, we can only calibrate the upper hemisphere, so we only measure there. However, for most scenes, this is the most significant region.

6. INTERPOLATION SCHEME

Rendering the reflectance field R for an arbitrary environment map may require evaluating $R(\theta, \phi)$ for light angles (θ, ϕ) which have never been observed. For this interpolation problem, Masselus et al. [2004] have investigated several methods. We will now discuss those properties that are required in our setup.

Let $(I_i)_{1 \leq i \leq n}$ be the images of the scene observed for distant light at angles (θ_i, ϕ_i) . Let (θ, ϕ) be the angle for which we want to interpolate. While we cannot guarantee to predict correct values of the reflectance field, we can aim for physically plausible interpolation for which we know that a light configuration exists which creates it. This is the case for linear combinations of the input images with nonnegative coefficients $\lambda_i(\theta, \phi)$ (Section 2):

$$R(x, y, \theta, \phi) = \sum_i \lambda_i(\theta, \phi) I_i(x, y). \quad (3)$$

Equation (2) can now be rewritten as

$$\begin{aligned} I(x, y) &= \int_0^\pi \int_0^{2\pi} L(\theta, \phi) \sum_i \lambda_i(\theta, \phi) I_i(x, y) d\phi \sin\theta d\theta \\ &= \sum_i \underbrace{\int_0^\pi \int_0^{2\pi} L(\theta, \phi) \cdot \lambda_i(\theta, \phi) d\phi \sin\theta d\theta}_{=: k_{i,L}} I_i(x, y). \end{aligned} \quad (4)$$

As the $\lambda_i(\theta, \phi)$ only depend on the structure of R and not on the pixel position (x, y) , we can precompute them. At rendering time, we can then compute the $k_{i,L}$ for the given environment map L and efficiently compute all pixels $I(x, y)$ in the target image as linear combination of the input images (with pixel coordinate-independent linear coefficients k_i).

In addition, we require the following properties of the interpolation weights λ_i :

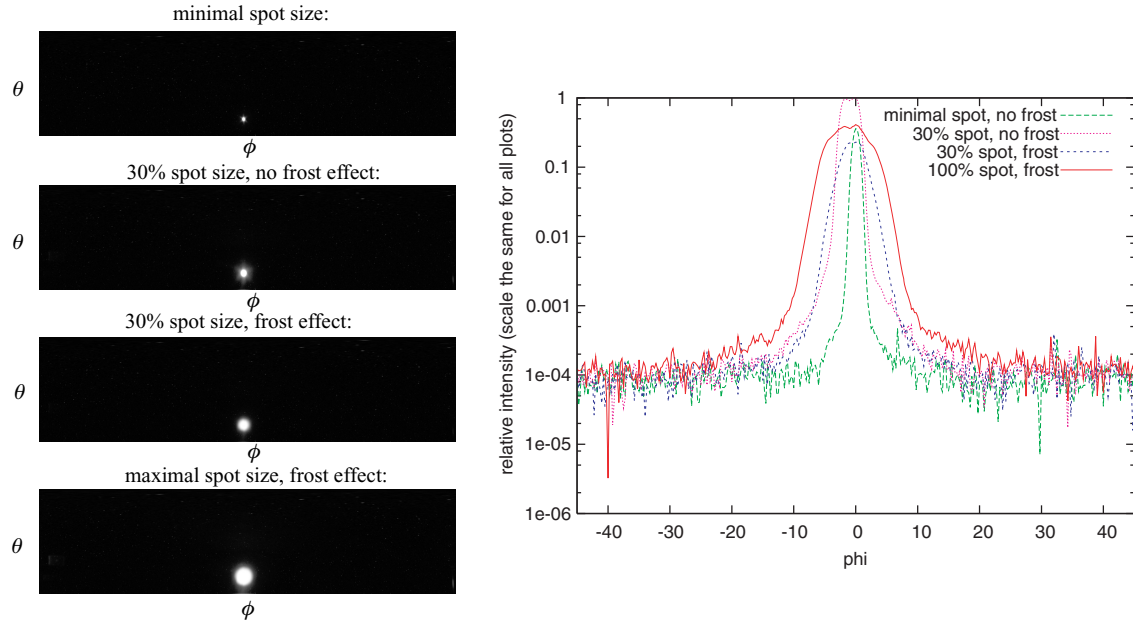


Fig. 2. Illustration of achievable light distributions. To the left, environment maps in polar parameterization of the lighting situations are shown, and, to the right, the intensity distribution along the centrally lit θ angle is plotted for varying ϕ values. Note the high contrast and the smooth fall-off.

—*Local Support*. $\lambda_i(\theta, \phi)$ should be nonzero only if (θ, ϕ) is close to the measurement light direction (θ_i, ϕ_i) of the corresponding image.

—*Partition of Unity*. For each direction (θ, ϕ) , the weights should add up to one:

$$\sum_i \lambda_i(\theta, \phi) = 1 \quad (5)$$

in order to conserve energy.

—*Interpolation*. If the position of the novel direction (θ, ϕ) coincides with a observed direction in (θ_i, ϕ_i) , the corresponding weight $\lambda_i(\theta, \phi)$ should be one.

—*Continuity*. The weights should change C^0 -smooth over changes in directions d .

—*Generality*. The approximation scheme should work with arbitrary, possibly nonregular, distributions of sample directions.

We have chosen a linear interpolation on barycentric coordinates [Masselus et al. 2004] defined on solid angle relations on the sphere throughout the following discussion. While smoother reconstructions are possible, this type of reconstruction can quickly and efficiently be evaluated and is sufficient to achieve the quality we show in our results. A detailed description of the computation of the interpolation weights is given in Appendix A.

7. FIXED-PATTERN SAMPLING

In our setup, we can measure reflectance fields at densities and illumination distributions from arbitrary angles which we will refer to as sampling patterns in the following. Unlike previous approaches with light sources mounted on fixed rigs, we can now freely choose the density both in elevation angle θ and in azimuth ϕ . Unlike handheld approaches, we can precisely distribute and repeat the sampling.

To create patterns at different densities of sampling directions θ , ϕ , and at an approximately uniform spacing, our algorithm places

virtual charges on pseudorandom positions on a sphere and iteratively moves them until an equilibrium is reached. The directions on the upper hemisphere are used as sampling positions. This randomized approach generates an approximately regular lattice with a hexagonal pattern (see Figure 5; the black points from iteration 0 have been generated with this algorithm).

7.1 Results of Fixed-Pattern Sampling

Figure 3 shows some results of fixed sampling with different numbers of input images and spotlight apertures. For the narrow aperture—which corresponds to a light cone of about 8 degrees—we need about 1025 images until the highlights on the ceramic figure become continuous for the vertical stripes case. While few isolated highlights are still visible in the ambient rendering, the general appearance is smooth. If we reduce the number of captured images to 256, strong artifacts occur on the glossy surface. The highlights are not continuous anymore, and, in some places, yellow highlights show up (as the available resolution is not sufficient to separate the contributions from a red and green colored input image), while some other highlights are missing altogether.

If the aperture is fully opened and the frost filter activated, the highlights are blurred out and high-frequency detail is lost. However, for only $n = 256$ input images, the aforementioned artifacts vanish.

For animations where the environment map changes continuously, the large aperture and smooth fall-off induced by the frost filter are highly advantageous as they reduce temporal aliasing. We demonstrate this by rendering an image sequence with a rotating point light source (Figure 4). In the top image, for $n = 1025$ light samples, we perceive a slight jaggedness in the shadow contour. The reason for this becomes visible when we decrease the number of samples to $n = 256$ as in the center image: temporal aliasing artifacts can be seen, the shadow jumps, or its contours fade in and out. This is due to the low number of input light conditions combined

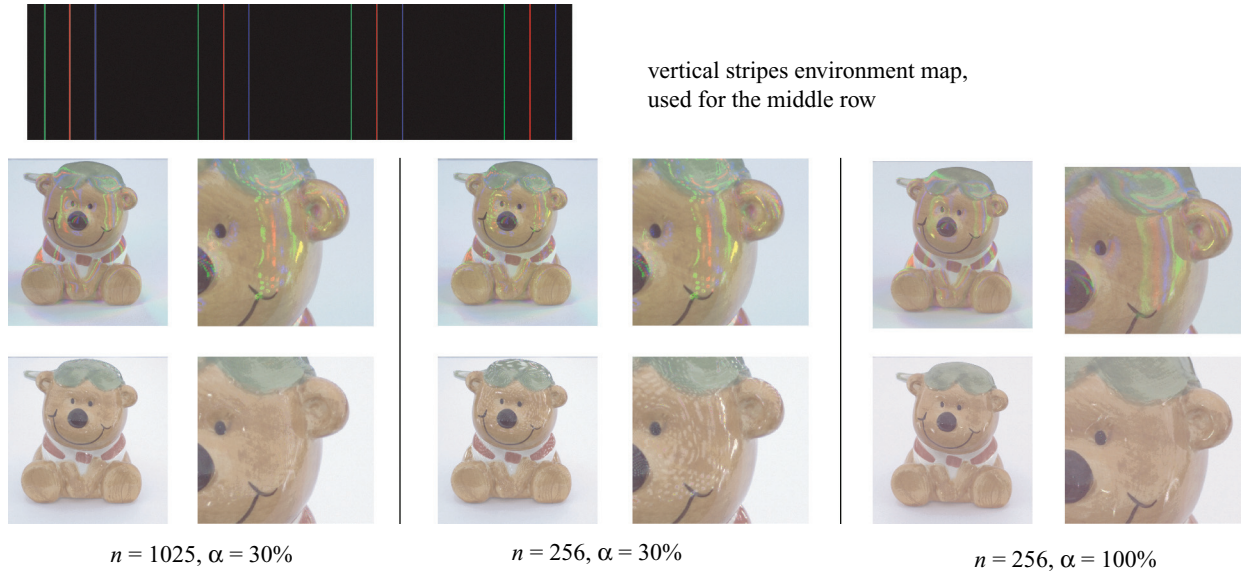


Fig. 3. Regular sampling results for a ceramic teddy for varying numbers of light sampling positions n and varying relative aperture size α . For full aperture, the frost filter was switched on, off otherwise. The middle row of images has been rendered with the stripe illumination on display in the first row, the bottom row has been rendered in ambient white light. The brightness of the resulting images has been rescaled to match each other.

with the narrow angular distribution, causing aliasing problems in the incident light domain.

The bottom image shows again a rendering for only $n = 256$ input images, with wide spotlight aperture and activated frost filter. Although the sampling density is the same as in the center image, the renderings are, though blurry, almost free from aliasing artifacts. The accompanying video file “Temporal Aliasing of Moving Shadows” shows the same scene as Figure 4 but shows the point light in actual motion.

On a side note, Figure 4 also demonstrates that we can resolve the high contrast between lit areas and the cast shadow, which is achieved as a result of using dark cloth as reflective material in our tent which reduces the intensity of scattered, ambient light.

8. ADAPTIVE SAMPLING

8.1 Introduction

The previous results show that, for environments with a low-frequency variation in the angular domain, it is sufficient to work with few input images and to use extended light sources. If the incident illumination contains high frequencies, however, these will be blurred, and detail cannot be preserved. This problem can be alleviated by using a narrower light distribution, but, if the hemisphere of input illumination is still to be covered fully, the sample density on the hemisphere needs to be increased by the square of the factor by which the spot radius is reduced. Using a regular sampling pattern, a significant number of the samples might be wasted, though since they might excessively sample smooth (diffuse) regions of the reflectance field.

We therefore propose an adaptive sampling scheme, which creates a subset chain of sets of incident light directions $S_0 \subset \dots \subset S_j \subset \dots \subset S_n$ with elements from $[0, \frac{\pi}{2}] \times (0, 2\pi)$. After the corresponding images for one set S_j are taken, it is analyzed and new directions are inserted to create S_{j+1} . For now, we work with a

vertical stripes environment map,
used for the middle row

fixed, narrow spot size. In Section 9, we will consider combinations of different spot sizes.

8.2 Algorithm

In this section, we explain in more detail how the sets $S_0 \subset \dots \subset S_j \subset \dots \subset S_n$ of sampling directions are constructed. The first set S_0 contains samples which are regularly distributed on the upper hemisphere as before.

Based on the images I_i for light directions $(\theta_i, \phi_i) \in S_j$, the set S_{j+1} is constructed by including additional sampling directions for which we will acquire the respective images in the following iteration. The detailed algorithm is as follows.

- (1) For each direction (θ_i, ϕ_i) , compute the one-ring N_i

$$N_i := \{k \in \mathbf{N} : (\theta_i, \phi_i) \text{ and } (\theta_k, \phi_k) \text{ share an edge} \\ \text{in the triangulation} \} \quad (6)$$

as the neighborhood of directions around (θ_i, ϕ_i) in the triangulation used for the interpolation (see Section 6 and Appendix A).

- (2) For each direction (θ_i, ϕ_i) , evaluate an error metric describing the smoothness of the reflectance field in the neighborhood of (θ_i, ϕ_i) .

A comparison of I_i to the average of the images $I_k : k \in N_i$ would be a straightforward choice. However, while developing our algorithm, we have observed such a metric causing endless subdivisions at one-sided discontinuities, for example, where a shadow or caustic suddenly appears or disappears due to occlusions in the scene, while neglecting high-frequency effects which are changing rapidly along all directions, for instance moving highlights. This is in contrast to our perceptual preferences as human observers can easily spot jumping highlights in animations or missing highlights in mirroring surfaces but are not as sensitive to the exact location of one-sided discontinuities.

We have therefore chosen a metric which yields low error values if I_i is similar to some of the $I_k : k \in N_i$, but results in a high error if I_i is different to each of the I_k .

In order to achieve this, we interpret the images as vectors and compute the projection error of I_i into $\langle I_k \rangle_{k \in N_i}$, that is, we choose



Fig. 4. Results of relighting in point light conditions. Between the two yellow lines, the same vertical strip of 10 pixels width is rendered repeatedly, while the point light is rotated around the up axis over 60 degrees with 0.5 degrees separation. Note the differences between the images at the shadow boundary. For the narrow aperture, $n = 1025$ images deliver an almost smooth boundary, for $n = 256$ images, the boundary is strongly jagged, and several overlapping shadow regions can be seen to the right of the image. For the same number of input images and for the same incoming light directions but for full aperture and activated frost filter, the jaggedness is almost not perceivable, and, although the shadow to the right of the image is noticeably blurred, it shows a continuous penumbra zone without ghosting artifacts.

linear coefficients μ_k such that

$$\left\| I_i - \sum_{k \in \mathcal{N}_i} \mu_k I_k \right\|^2 \quad (7)$$

is minimized and interpret the value of the entire term as local error. While this heuristic is not perfectly adapted for discontinuities, it produces good results in our experiments, and its value can be very quickly computed using a singular value decomposition on the matrix composed of the I_k .

- (3) For each triangle in the triangulation, assign the sum of the corners' error values as error value.
- (4) For each of the m triangles T_l with the largest error, insert a new direction (θ_l, ϕ_l) at T_l 's circumcenter unless it is close to a direction

already inserted in the current iteration in which case the next most important direction is inserted.

This results in a locally-even distribution [Lloyd 1982]: the newly inserted direction is equally distant to points already present in the local neighborhood as the circumcenter of T_l is on the intersection of the perpendicular bisectors of T_l .

In areas where the local point density suddenly changes, the (θ_l, ϕ_l) may even be outside T_l ; however, inserting a sampling direction there improves the distribution of the sampling directions, and further refinements inside T_l will take place in future iterations if required.

S_{j+1} is then defined as the union of the new sampling directions and S_j . This process is iterated until S_n has been constructed, and the corresponding images have been taken. In each iteration, we add $m = 16$ directions, which is more efficient than adding one direction $m = 1$ at a time. In order to impose an upper limit on sampling

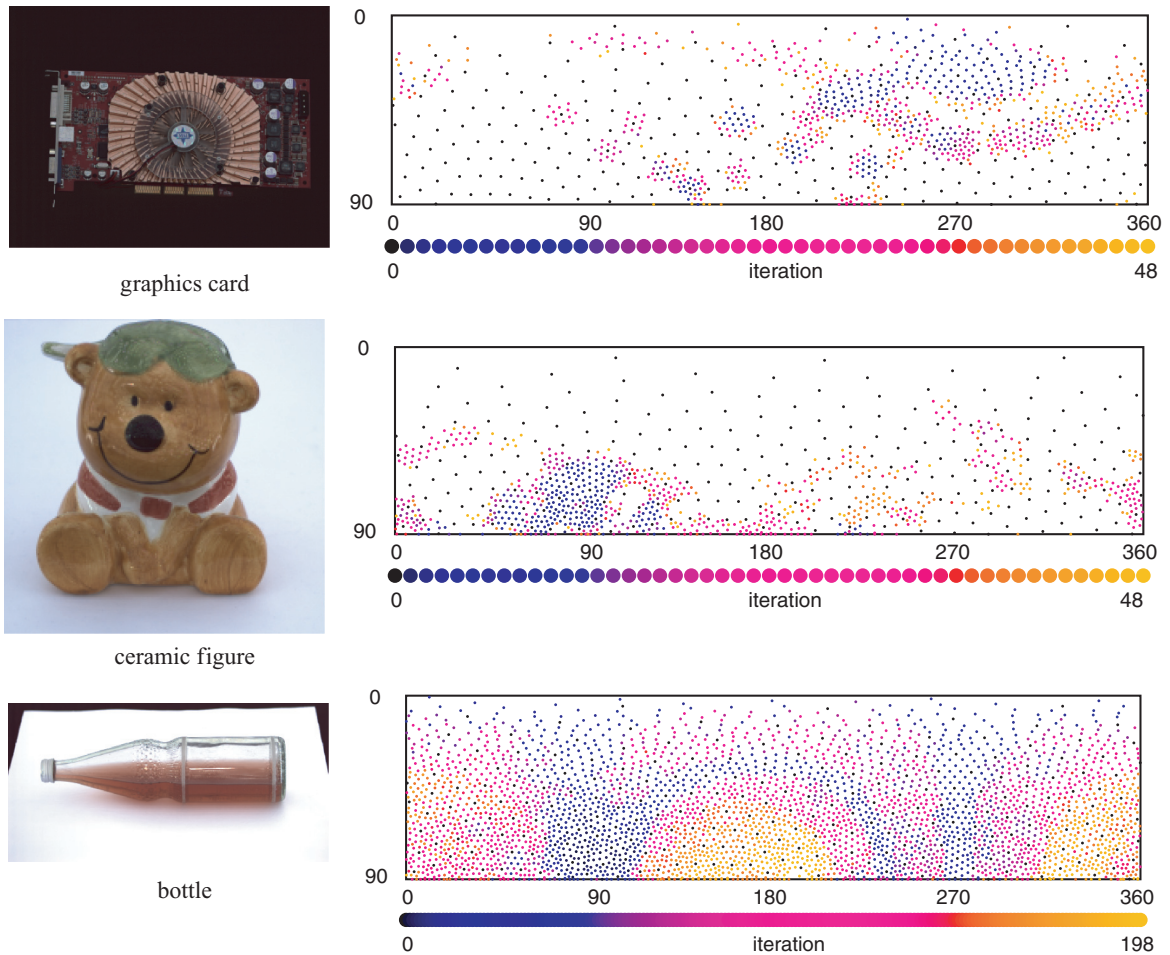


Fig. 5. The plots illustrate the sequence in which light directions have been generated. For initialization, iteration 0 consisted of 256 uniformly distributed samples (drawn in black), later samples grow in regions of interest. The renderings in the left column are for $n = 1024$ total images (after 48 iterations) rendered in ambient light.

density, we skip all triangles where the radius of the circumcircle would be below a threshold of 2 degrees.

8.3 Results of the Adaptive Sampling Algorithm

In order to evaluate our sampling scheme, we have conducted a set of experiments. We have measured reflectance fields of three objects, which we will discuss further:

- (1) a bottle containing a colored liquid which, due to its cylindrical shape and glossy surface material, is expected to have an anisotropic reflectance field, that is, it should demonstrate the benefits of adaptive sampling,
- (2) a graphics card, which has complex self-shadowing and is covered by a diverse set of different materials, and
- (3) a ceramic figure, which serves as a difficult example: on its surface, we find practically every surface normal direction as opposed to the cylindrical bottle. Moreover, its surface is shiny, so each direction in the illuminating environment is almost equally important.

For all these objects, we measured reflectance fields once using fixed sampling configurations and once applying our adaptive sampling scheme. In the adaptive case, we initialized the algorithm with 256 regularly distributed samples and added 16 light situations in each of 48 iterations, yielding 1024 total images. For the bottle data set, we added sample points until the dataset was densely sampled and only less than 16 images could have been inserted, which happened after iteration 198. Figure 5 illustrates the insertion sequence for the different datasets, and shows the objects rendered in ambient light after 48 iterations. As can be seen, clusters of sampling directions occur, indicating regions of directions where the appearance changes drastically.

8.3.1 Quantitative Evaluation. For each of the objects mentioned previously, we have recorded a ground truth measurement, consisting of images in 10,000 light directions densely distributed over the upper hemisphere on a regular grid, which is visualized as the positions of dots in Figure 6.

In order to obtain a numerical measure for the quality of a given reflectance field, we render an image for each of the captured ground truth directions and sum up the squared differences between the

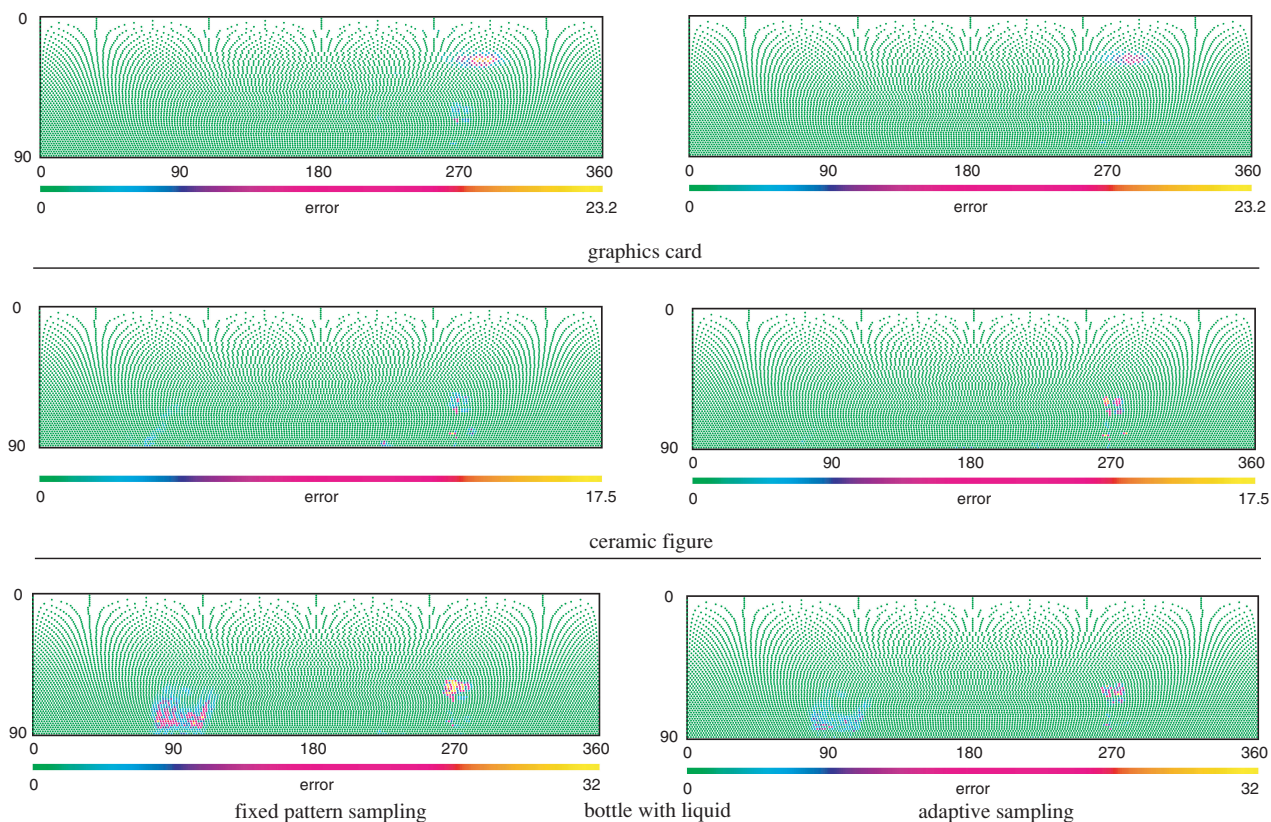


Fig. 6. Spatially plotted errors. Each dot represents the light direction for which a ground truth image was taken, its color gives the sum of squared pixel differences between ground truth and the reconstructed image for this direction. The adaptive sampling column displays values for a reflectance field with 256 initial, regularly distributed samples, into which groups of 16 samples each were inserted over 48 iterations, yielding a total of 1024 images; the fixed-pattern sampling column shows a regularly sampled reflectance field of 1025 input images.

ground truth pixel values and the rendered pixel values. The color-coded differences are plotted in Figure 6 at the corresponding (θ, ϕ) locations. As the aspect ratio of the objects does not match the aspect ratio of the camera, we have cropped the input images to the same region of interest which is used for the renderings in Figure 5. In order to produce the plots efficiently, we have downsampled the images using either a 3×3 or a 4×4 box filter to smaller resolutions: 453×224 pixels for the bottle dataset, 453×341 Pixels for the graphics card, and 200×196 for the ceramic figure. This restricts the precision of the least-squares error metric, but it also increases the robustness against pixel noise generated by the camera. As a result, the absolute numbers of SSD error are only comparable for the same dataset.

In the error plots of Figure 6, we can observe that, for the ceramic figure, the error distribution and the amount of error is comparable both for the adaptive sampling scheme and a fixed sampling scheme with similar amounts of input images. For the other datasets, however, there is considerably less residual error in the adaptively sampled reflectance field than in the fixed pattern case. Compared to Figure 5, we can see that the remaining error zones are to be found in the (θ, ϕ) regions where samples have been densely placed up to the density limit and that these samples have also been placed in early iterations.

In order to evaluate the performance of the adaptive sampling scheme as the iterations proceed, we have summed up the errors depicted in Figure 6 for each iteration, which gives a total numerical error value. As the sampling distribution of the ground truth plots is arranged so as to ensure an equal density over the hemisphere, this is a meaningful error value. Figure 7 shows these global error values as the iterative, adaptive algorithm proceeds, and as reference values for the total error values for control measurements with a fixed sampling pattern.

Even for the ceramic figure as the worst-case scenario, the adaptive sampling scheme stays competitive to fixed-pattern sampling, and, for low numbers of input images, it is consistently better. For the other two reflectance fields, which are more anisotropic, the advantages of adaptive sampling are even more clear. Over a wide range, the same amount of error can be achieved with about half or even less than half of the input images as compared to fixed-pattern sampling.

Especially in the bottle dataset, a plateau is reached at some point close to 700 images, or about 28 iterations. As can be seen in Figure 5, the light directions which have been sampled up to this point (marked in black and dark blue) coincide with the regions which expose the highest remaining error in Figure 6. Also, the final sampling density has already been reached there. Because of

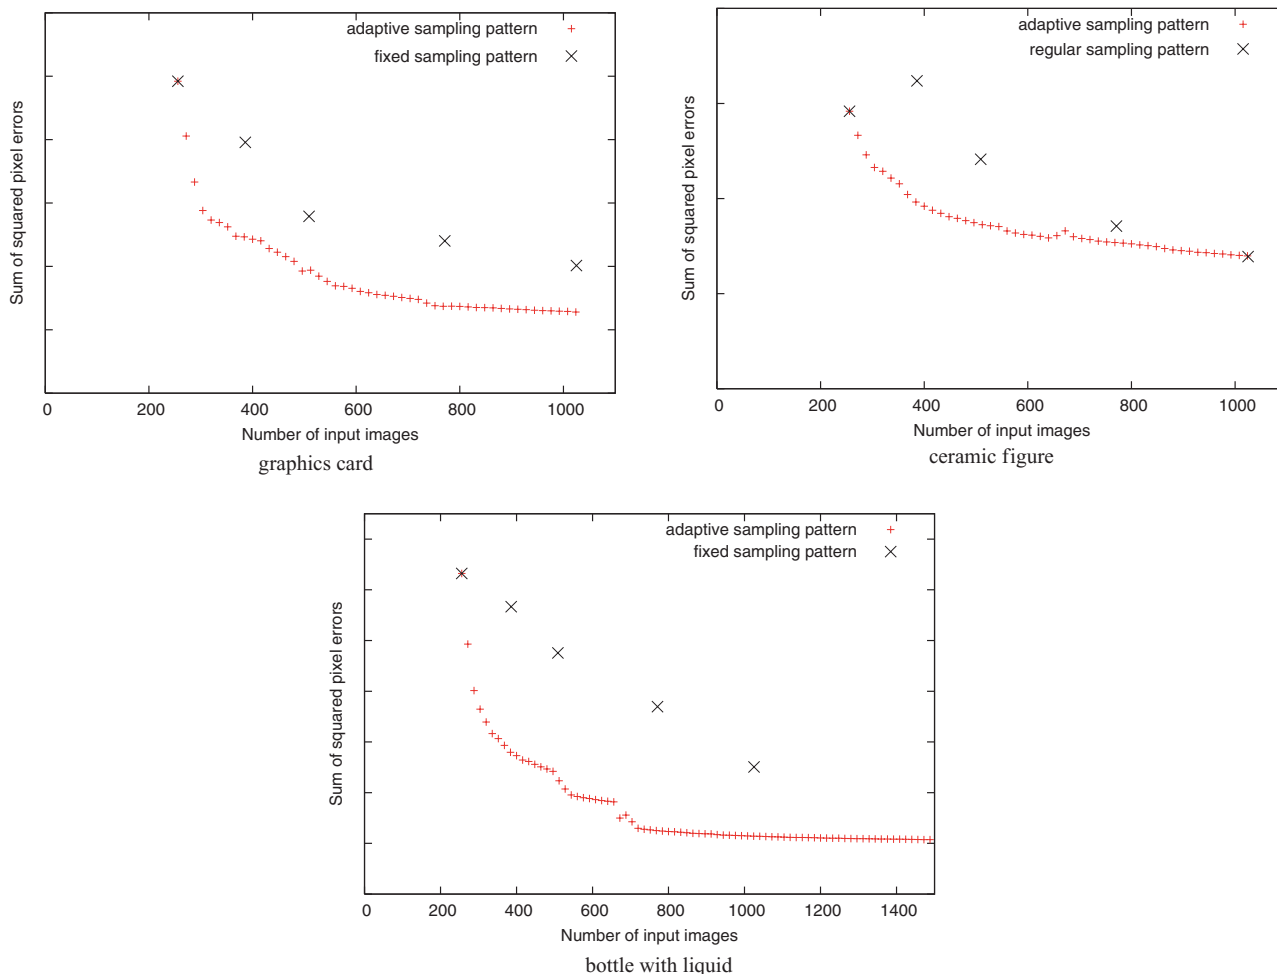


Fig. 7. Plots of the global error value as sum of squared pixel differences between ground truth and adaptive/fixed pattern sampling.

this, the error cannot be reduced much more as samples can only be inserted in areas where the error already was low.

8.3.2 Qualitative Evaluation. So far, we have only shown results on a quantitative evaluation of the different sampling schemes. Equally important, however, are qualitative evaluations, which demonstrate improvements in visible image quality. Therefore, we have created a synthetic high-frequency environment map, which is displayed in the top right of Figure 8. It consists of cells that are filled by one solid color each, using the colors red, green, and blue. Both the frequency of color changes and the purity of the resulting colors can be used as quality indicators in highlight regions of rendered images. Outside highlight regions, the integration over the product of the object BRDF and the environment map will create a diffuse object color since the illumination from a region of the environment will average to white.

Figure 8 shows renderings of the bottle scene in this environment map for different sample numbers both for fixed pattern and adaptive sampling. In the electronic version of this article, you can zoom in until the full resolution is visible; in order to make the specific differences more visible, we provide Figure 9 which enlarges the highlight region on top of the bottle.

The reflectance field composed of $n = 256$ regularly distributed samples, which also serves as input data for the adaptive sampling algorithm, clearly lacks quality: the highlights are visibly separated, and where highlights are present, they show up as composed colors (yellow, magenta, cyan), which means that the same input image was used for two different input light conditions by the interpolation algorithm. This shows that the sampling resolution is much too low.

As the number of samples increases, the fixed pattern scheme closes the gaps between the highlights, but it takes up to $n = 1025$ input images until a clear separation into the colors red, green, and blue takes place. The adaptive sampling scheme, however, reaches full color fidelity much earlier, even after only 128 samples are inserted ($n = 384$).

As can be seen in Figure 10, which shows a detail left to the bottle's center, there are also regions where adaptive sampling reaches full quality slightly later than fixed-pattern sampling. This occurs if the intensity and the image size of highlights is small so these regions are refined later in the algorithm. While the fixed-pattern sampling gradually reduces the distance between light source positions as the sample number increases, the adaptive algorithm may focus on other regions first and refine small highlights only relatively late.

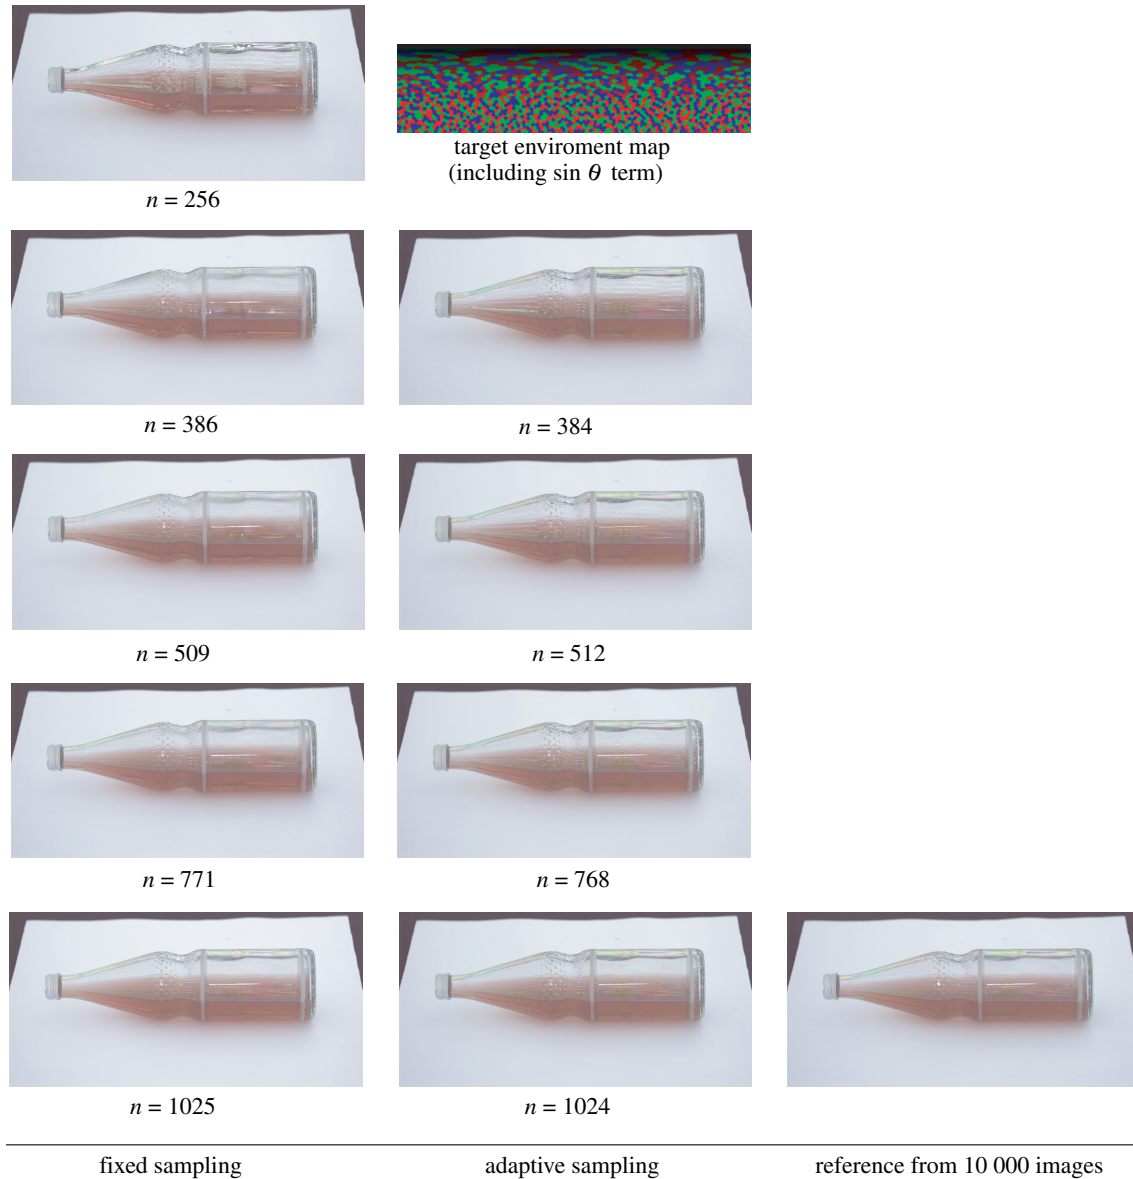


Fig. 8. Renderings of the bottle scene in a high-frequency environment map (top right), for different regular sampling distribution (left column) and different adaptive sampling distributions. The adaptive sampling was initialized by acquiring a regularly sampled reflectance field with $n = 256$ images. To the bottom right, a rendering created from the measurement with 10,000 input images is shown as reference (see Figures 9 and 10 for details).

In Figure 11, renderings for the ceramic figure dataset are displayed. While due to its isotropic reflectance field, the adaptive scheme does not achieve a sufficient density to obtain continuous highlights for $n = 1024$ total images, the highlights on the cap of the figure are reconstructed more precisely than for $n = 1025$ uniformly distributed samples.

8.3.3 Timings/Performance. The acquisition time required to record a dataset in our setup not only varies with the total number of recorded images, but also with the sequence type. We will now list timings at the example of the bottle dataset. For the HDR reconstruction, two standard PCs were used, one of which also ran

the refinement computations and control loop; the projectors and camera were controlled by a laptop.

Acquiring $n = 1025$ sample images in the fixed pattern distribution took 131 minutes, plus 25 minutes of cropping and scaling for the area of interest. Adaptive sampling of $n = 1024$ images, of which 256 are distributed regularly in the initial set, and 16 are added in each following iteration, can be done in 245 minutes including all processing. The acquisition of the 10,000 image ground truth data took 1,023 minutes plus 248 minutes of processing.

The additional time for adaptive sampling is dominantly spent on the following items: the reinitialization of the measurement devices in each iteration, the unordered rotation of the moving heads to the

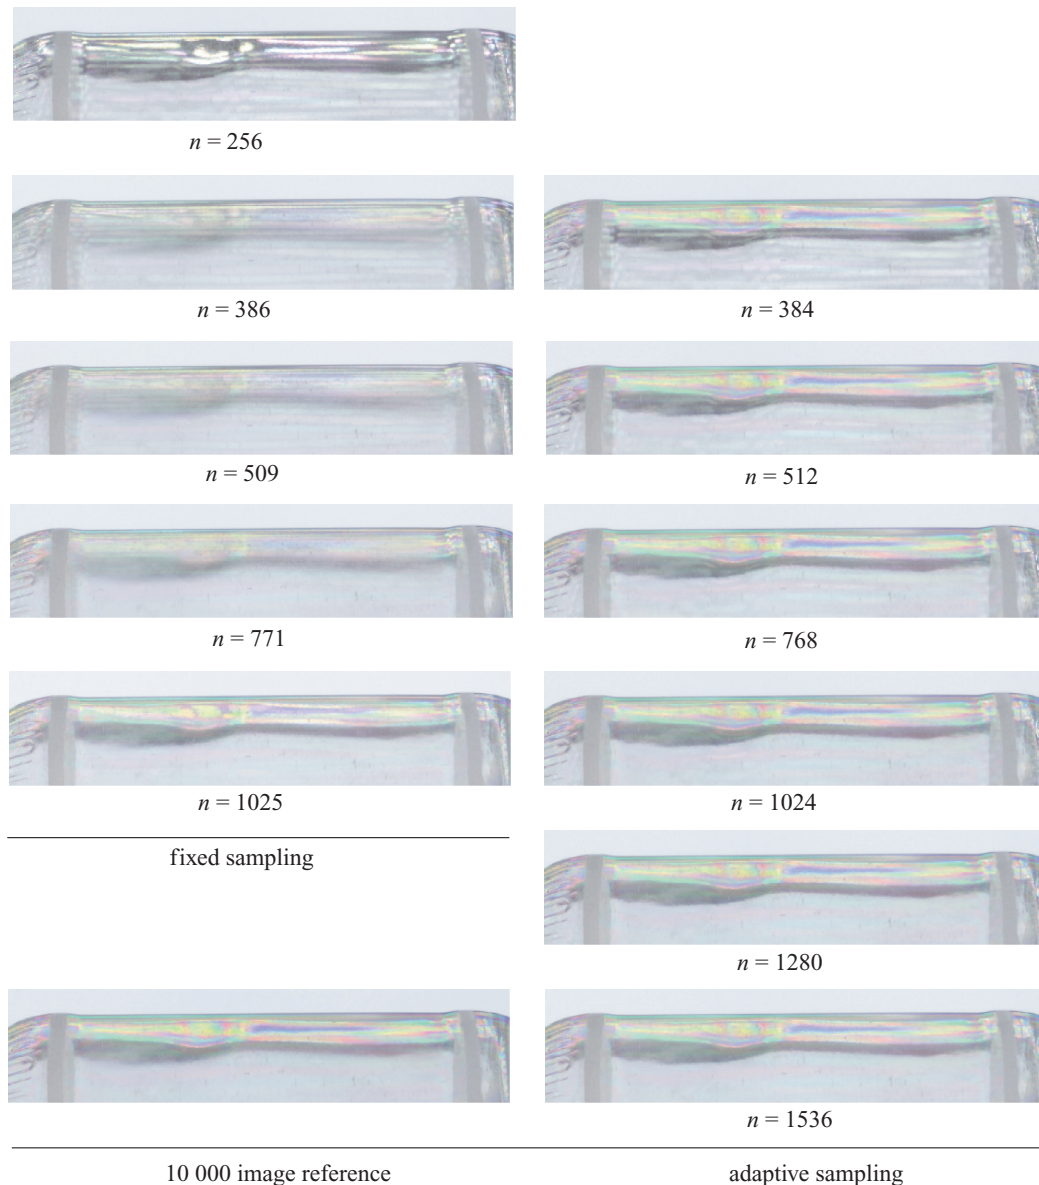


Fig. 9. Detail pictures for Figure 8. Note how the detailed reconstruction of the highlights on the bottle surface is achieved much earlier by adaptive sampling.

unevenly distributed target positions, and the time for evaluating the distance measure in Equation (7). A more efficient implementation with more advanced pipelining and reduced I/O overhead should cover the first two points; the analysis of the reflectance field, point insertions, and triangulation operations in total take about 1.9s/image.

9. TWO-LEVEL REFLECTANCE FIELDS

The previous results already show an improvement over a fixed-pattern sampling as sampling concentrates at the most important light directions. However, as we use a global measurement process, undersampling deficiencies can still occur in areas where small image regions did not contribute sufficiently to the global error measure so as to warrant further subdivision but are sensitive to only sparsely

sampled light directions. Such artifacts may, for example, show up as discontinuous highlights in high curvature regions.

Following the discussions in Sections 4 and 7, we can mask these effects by replacing those samples by samples with a stronger pre-filtering. We therefore propose a two-level representation for the reflectance field which contains samples from two different filter kernels, which we will discuss in more detail in the next paragraphs.

We partition the final set S_n of sampling light directions S_i into two subsets,

$$S_n = S_{\text{narrow}} \cup S_{\text{wide}}, \quad (8)$$

so that S_{narrow} contains densely sampled regions of important directions, and S_{wide} contains sparse samples of directions in smooth regions, filling up the hemisphere. Then, we can create a combined reflectance field of narrow and wide spotlight apertures. For the first

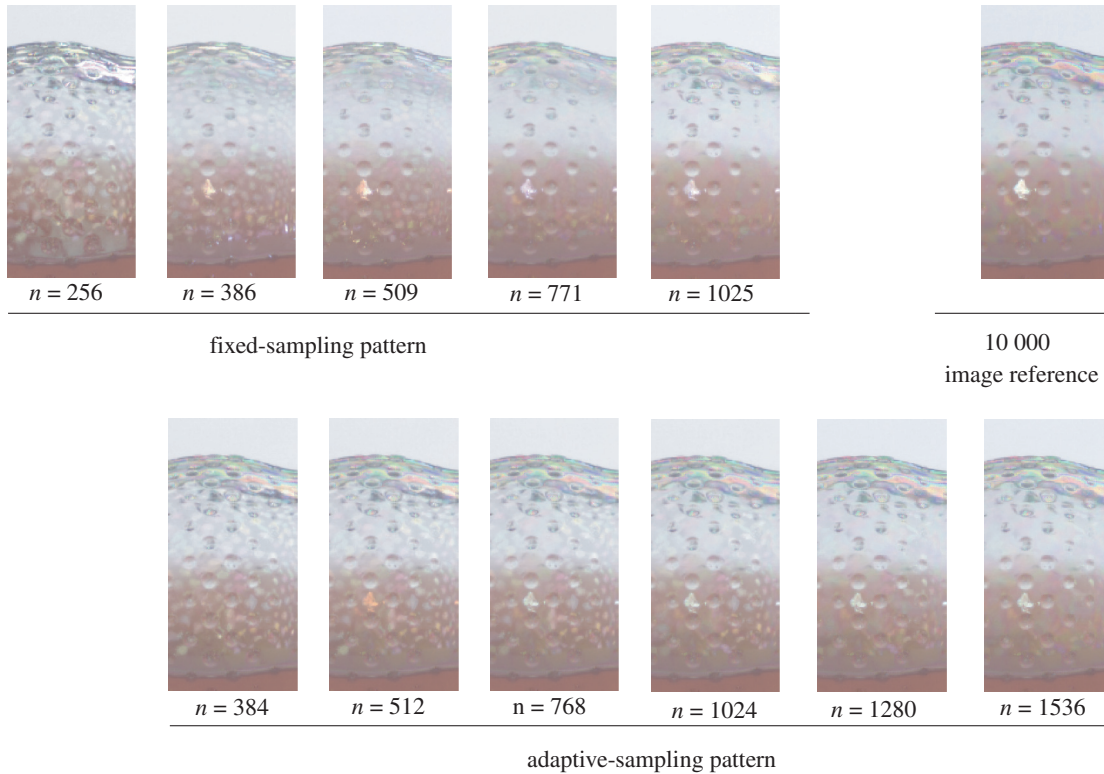


Fig. 10. Detail pictures from Figure 8. Unlike the reflections on the top part of the images which are captured nicely by adaptive sampling, other artefacts are more persistent: These are small colored spots distributed over a large area in the center of each image, best seen in front of the dark liquid (lower half of each image). In adaptive sampling, they disappear later than in fixed-pattern sampling because the algorithm tends to focus on more intense highlights first.

field, we retain the images already taken during the adaptive sampling step, and for the second field, we acquire new images with fully open spotlight aperture and active frost filter. For each of the directions in S_{wide} , we scale the brightness of the full-aperture image so that it matches the average of the respective previous images in narrow lighting. Then, we use it as a replacement in the reflectance field. For rendering, we interpolate the field with the same algorithm as before (see Section 6).

The detailed process is illustrated by an example: we start with one of the adaptively sampled reflectance fields of the bottle scene, and again consider neighborhoods N_i for the directions $(\theta_i, \phi_i) \in S_n$ as in Section 8.2. For each of the directions which were in the initial sample set S_0 , and for which the difference angle to any of the directions in N_i exceeds a threshold (9 degrees in our experiment), we record an additional image lit from the same light direction, but with full spot aperture and active frost filter.

By restricting ourselves to $S_{\text{wide}} \subseteq S_0$, we can sample S_0 once in the wide aperture setting and quickly create renderings for different values of n . This restriction, however, would not be used in an application case.

Figure 12 shows renderings for such fields in the same high-frequency environment map as before and lists the number of total images needed (both from the initial measurement and from the additional measurement in wider light conditions). For $n = 256$, all light directions fulfill the criterion for remeasurement so they are all replaced, and the rendering shows dull highlights as expected.

The next row already shows a large improvement. For $n = 384$, the rendering of the resulting reflectance field shows both the high

quality highlights at the top of the bottle and a continuous illumination in the right-column detail enlargement. While 244 images need to be rerecorded, driving the measurement costs up to 628 images, the result is visually almost of the same quality level as uniformly lit, adaptively or regularly sampled reflectance fields with higher image counts. Since some of the directions have been captured both with narrow and wide spotlights, only 384 images need to be finally stored and accumulated in the rendering step.

The video file *two-level results* shows advantages of the two-level reflectance field over fixed and adaptive sampling as a thin strip environment map rotates. Figure 13 shows a result for a real-world environment map.

10. DISCUSSION

10.1 Evaluation Using Sampling Theory

In this section, we examine the relationship between sampling density and the shapes of the light sources. More specifically, we have computed the power spectra of the light intensity plots on display in Figure 2 as shown in Figure 14 as log-scale plots.

Here, the advantage of using extended light sources with smooth fall-offs becomes apparent. While, for each frequency, some power is present in the spectrum, there is a quick decay as the frequency increases. Thus, we can use the sampling theorem to motivate a maximum sampling frequency.

If, for instance, we choose to ignore frequencies whose contribution is less than 10%, and we use the minimal spot size, we need to



Fig. 11. Renderings of the ceramic figure in the high-frequency environment map shown in Figure 8. As can be seen, the renderings of the colored reflections on the cap is much closer to the reference than in the regular setting for the adaptive sampling case. Still, the narrow light source causes spotty artifacts in the remaining parts as highlights are disconnected. This can be avoided by using a wide aperture for these samples (see Figure 3) which motivates two-level reflectance fields.

deal with frequencies of up to 160 cycles per 360° ¹, requiring that we place the samples at a distance of at most $d_{10\%} = \frac{360^\circ}{160 \cdot 2} = 1.125^\circ$. The table in Figure 14 lists the respective values for the four listed light source configurations.

We can now evaluate the direction distributions that we used by computing their Delaunay triangulations and looking at the angles the edges span. They represent distances between sampling directions which are adjacent in the Voronoi diagram. Figure 15 shows percentile plots of these distributions. As can be seen, the reference data is sufficiently dense so that, when recorded with the narrow light source, the reflectance field data will expose aliasing only in frequencies which contribute less than 10% to the transmitted power. The distances are densely concentrated, corresponding to an approximately even distribution.

For the fixed pattern, the samples are also evenly distributed; for the narrow light source, they are merely dense enough to measure a signal which may still expose aliasing in frequencies which are present with 50% or more of relative power.

The adaptively sampled two-level dataset has an uneven distribution. In the densely sampled area, it is close to the $d_{10\%}$ limit for narrow light sources. At 9° degrees and above, we switch to the wide light-source recordings which require less dense sampling. This threshold separates between the adaptively inserted directions and the start set of directions S_0 (see knee in Figure 15) and is appropriate if we assume that the adaptive sampling algorithm has reliably sampled the most important regions. A more conservative

¹In order to observe the light source fully, we have positioned it a few degrees above the horizon of the recording fish-eye lens. As a result, the measured intensity distribution covers less than 360° arc length. Still, the reported minimal required densities are misestimated by less than 4%.

threshold would be for example, $d_{50\%} = 4.286^\circ$ for the narrow light source as it would impose a limit on the aliasing that could possibly occur.

10.2 Limitations of This Work

The previous discussion of sampling densities in the context of sampling theory includes several simplifications. Especially, it analyzes sampling densities as if the reflectance field was a 1D signal; a 2D analysis is more appropriate, but this significantly more complicated analysis is beyond the scope of this article. Also, as the tent that we used for illumination is not spherical, and the projectors are not precisely in its center, an analysis of sampling effects might also include variations in the light distributions with θ and ϕ .

Unlike authors such as Peers et al. [2005], we do not estimate independent reconstruction functions for each image pixel. This avoids fitting errors, but it restricts us to a global data representation and a global metric for the refinement of the reflectance field sampling. Accordingly, small image regions may expose interesting reflectance features which are not precisely sampled by the adaptive sampling algorithm. Still, we hide aliasing artifacts by sampling sparsely observed regions with large kernels in the two-level approach. As the large kernels overlap, we can be sure that we do not miss features altogether.

The adaptive sampling algorithm is inherently online, that is, it can only take decisions on where to refine the reflectance field based on data it has already observed; thus, important features like isolated highlights may be missed. Still, they will be captured by the wide kernel measurement in the two-level algorithm. Also, we have not exhaustively investigated the refinement criterion. Yet our evaluations show a good performance of the adaptive sampling.

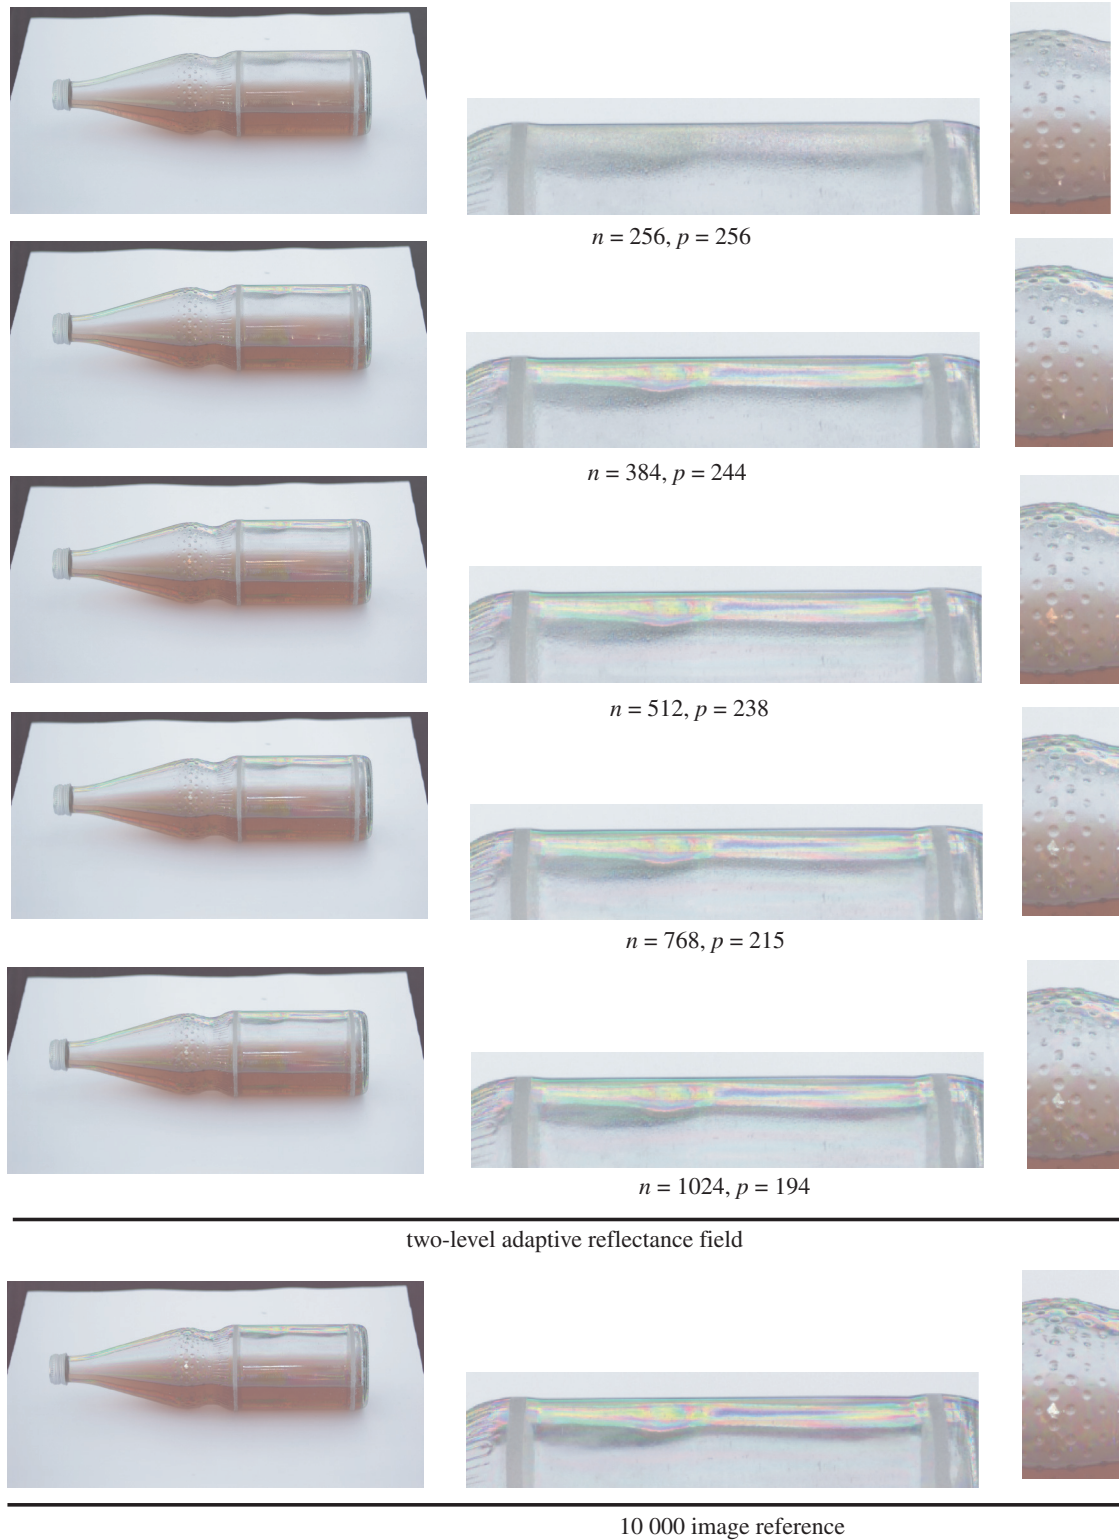


Fig. 12. Reflectance fields of different illuminant spot sizes can be combined into one in the adaptive scheme. Here, n is again the number of images used in the dataset, p is the number of pictures which were rerecorded for wide aperture. For instance, in the top row, which is the initial dataset, all pictures were replaced.

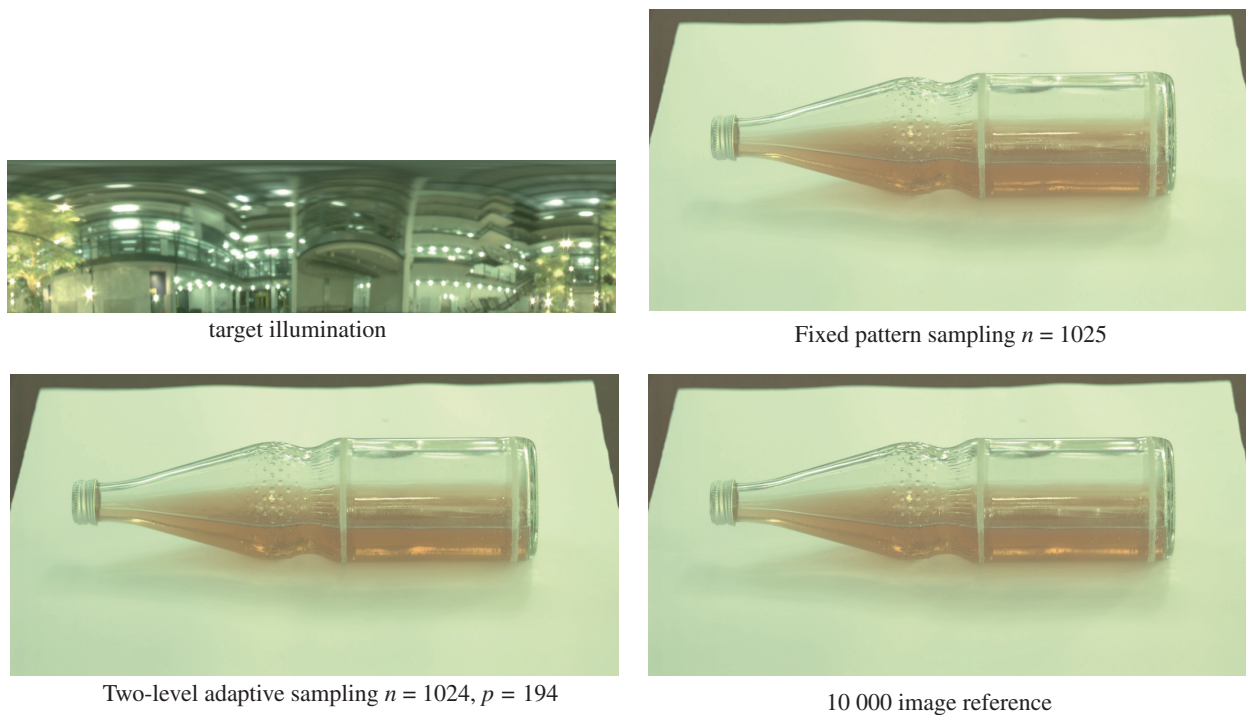
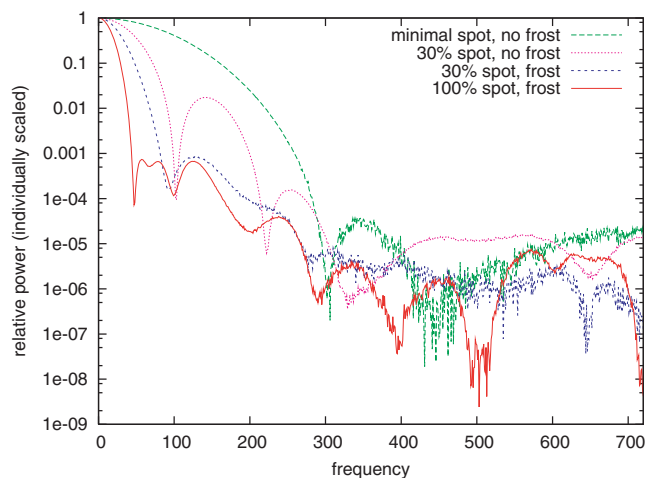


Fig. 13. Rendering of the bottle scene in a real-world environment map. While the differences are subtle, the fused sampling field is closer to the 10,000 reference than a fixed-pattern sampling, especially in the bright red reflection at the lower border of the bottle.



type	$\nu_{10\%}$	$d_{10\%}$
minimal spot	160	1.125°
30% spot ("narrow")	72	2.5°
30% spot, frost	48	3.75°
100% spot, frost ("wide")	29	6.207°

type	$\nu_{50\%}$	$d_{50\%}$
minimal spot	88	2.045°
30% spot ("narrow")	42	4.286°
30% spot, frost	26	6.923°
100% spot, frost ("wide")	17	10.588°

Fig. 14. Power spectrum for a selection of light source configurations computed with a Discrete Fourier Transform on the data sampled at 1,440 locations along ϕ for a fixed θ . Each spectrum is rescaled so that maximally transportable power equals 1. The table to the right lists, for each configuration, the first frequency where the contained power drops below 10%/50%, respectively, and the maximum allowable distance between two sampling directions so that this frequency can be represented according to the sampling theorem.

The evaluation is limited in the sense that we do not take possible target illuminations into account but use an evenly distributed, dense reference reflectance field. The least squares metric for the comparison of reference and rendering is not perceptually motivated.

We have shown that reflectance fields can consist of samples in different resolutions/taken with different filter kernels. In order to

evaluate the improvements with increasing sampling counts (see Figure 12) effectively, we have only combined samples of two different kernels. A more principled way is to use all available light configurations and analyze the corresponding Nyquist Frequencies. After the adaptive algorithm has determined the sampling positions, each sample can then be recorded with an optimal light source.




Fig. 15. Distribution of angles between neighboring sample directions for the bottle dataset in percentile plots. The dashed vertical lines illustrate required sampling densities for the narrow and the wide light configurations if frequencies contributing less than 10% or 50%, respectively, may be ignored. The blue line shows the distribution of the 10,000 reference points set, the grey line shows the distribution of 1,025 sampling directions in a fixed pattern, the red/green line shows the distribution of the adaptive bottle dataset directions after 48 iterations with 1,024 recorded images. The red part is recorded with narrow light sources, the green part is recorded with wide light sources in the two-level field.

Finally, the measurement setup construction involves a trade-off between optimal (close to spherical) geometry and easy construction, which led to the tent design. This nonuniformity is partially mitigated by the light source intensity calibration.

11. CONCLUSION

In image-based relighting, the quality of the renderings depends critically on the characteristics of the illumination basis used for measuring the reflectance field. This article presents a way of numerically evaluating the quality of sampling patterns of different light directions. We have introduced an adaptive sampling scheme and shown, both quantitatively and qualitatively, how it improves on the trade-off between measurement effort, that is, the required number of input images, and measurement quality when compared to fixed-pattern sampling.

The automated measurement setup described in Section 5 makes the required data efficiently available, scaling up to thousands of images. The acquisition system allows for measuring with an extended light source covering any possible direction of the entire hemisphere by indirect illumination reflected off dark cloth. The dark cloth reduces the power of ambient light effectively. We have shown that the use of extended light sources can significantly improve the rendering quality especially at low sampling rates.

In Section 9, we have discussed the combination of narrow and wide light sources in a single reflectance field. In this context, there is potential for future work. For instance, it would be desirable to support more than two different extents of light sources. Further, we would like to avoid sampling the reflectance field from the same light direction twice as this would improve the measurement performance.

A new quality metric is required to correctly compare the quality of a two or multilevel reflectance field with varying light source diameter to one acquired with a spotlight of fixed width. In this

context, the conceptually simple least squares metric used throughout the article could also be replaced with something perceptually more meaningful. Still, the simple scheme presented here already produces useful results for very low image counts.

Our findings demonstrate that adaptive measurement schemes, which take the specific properties of objects into account, have a significant potential to improve the efficiency of reflectance field methods.

APPENDIX

A. COMPUTATION OF INTERPOLATION WEIGHTS

The light directions (θ_i, ϕ_i) used to measure the reflectance field can be interpreted as normalized vectors $d_i \in \mathbb{R}^3$ on the unit sphere, forming a set D . Let $d \in \mathbb{R}^3$ be the unit vector in the direction (θ, ϕ) we want to interpolate the reflectance field R from Equation (2) for.

In the following, we will explain how to select indices j, k, l and interpolation weights $\lambda_j(\theta, \phi)$, $\lambda_k(\theta, \phi)$, $\lambda_l(\theta, \phi)$, such that we can approximate

$$R(x, y, \theta, \phi) \approx \lambda_j(\theta, \phi) \cdot I_j(x, y) + \lambda_k(\theta, \phi) \cdot I_k(x, y) + \lambda_l(\theta, \phi) \cdot I_l(x, y). \quad (9)$$

In a preprocessing step, we compute a triangulation T of the hemispherical hull of D . We do that by computing a Delaunay tetrahedralization of $D \cup \{g\}$, where g is a guardian element below the hemisphere, at the $(\theta, \phi) = (\pi, 0)$ pole. Each tetrahedron which contains a point at infinity is on the outside of D , and contains a triangle t which consists only of points on $D \cup \{g\}$. If and only if t does not contain g , t lies entirely on the upper hemisphere, and thus defined to be part of T . All vertices d which are connected to g by an edge in the tetrahedralization are marked as boundary vertices for further use.

Now, we can intersect the ray defined by d and the origin with all triangles in T . Two cases are possible:

- (1) the ray intersects with the triangle created by (d_j, d_k, d_l) . Let

$$s : \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}_0^+ \quad (10)$$

now be the function which maps three points on the unit sphere to the nonnegative solid angle covered by the inside spherical triangle defined by the points.

We can then define

$$\begin{aligned} \lambda_j(\theta, \phi) &:= \frac{s(d, d_k, d_l)}{s(d_j, d_k, d_l)}, & \lambda_k(\theta, \phi) &:= \frac{s(d_j, d, d_l)}{s(d_j, d_k, d_l)}, \\ \lambda_l(\theta, \phi) &:= \frac{s(d_j, d_k, d)}{s(d_j, d_k, d_l)} \end{aligned} \quad (11)$$

yielding an interpolation with barycentric coordinates on spherical triangle solid angles.

- (2) The ray does not intersect with any triangle. This is possible if d is close to the border of the hemisphere. In that case, we select j and k so that d_j and d_k are boundary vertices so that

$$\phi_j \leq \phi < \phi_k \wedge \exists m : \phi_j < \phi_m < \phi_k \quad (12)$$

By duplicating the vertex d_i , where $\phi_i \leq \phi_m \forall m$, and assigning a value of $\phi := \phi_i + 2\pi$ to it, we can handle the wrap-around condition correctly.

Now, we can define

$$\lambda_j(\theta, \phi) := \frac{\phi_k - \phi}{\phi_k - \phi_j}, \quad \lambda_k(\theta, \phi) := \frac{\phi - \phi_j}{\phi_k - \phi_j}, \quad \lambda_l(\theta, \phi) := 0 \quad (13)$$

as arc interpolation weights.

ACKNOWLEDGMENTS

We would like to thank Ken Clarkson, who wrote the program `hu11` which we used for the triangulation computation; Carsten Stoll for fruitful discussions during the development of the adaptive sampling algorithm; Michael Heinz, Christian Fuchs, Axel Koeppel and Michael Laise for help with the measurement setup; and Rafal Maniuk and Grzegorz Krawczyk for the `pfstools` package we used for the high dynamic range image processing.

REFERENCES

- CHAI, J.-X., CHAN, S.-C., SHUM, H.-Y., AND TONG, X. 2000. Plenoptic sampling. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. ACM Press, New York, NY, 307–318.
- CHUANG, Y.-Y., ZONGKER, D. E., HINDORFF, J., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2000. Environment matting extensions: towards higher accuracy and real-time capture. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. ACM Press, New York, NY, 121–130.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1, 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 *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. ACM Press, New York, NY, 145–156.
- EINARSSON, P., CHABERT, C.-F., JONES, A., MA, W.-C., LAMOND, B., HAWKINS, T., BOLAS, M., SYLWAN, S., AND DEBEVEC, P. 2006. Relighting human locomotion with flowed reflectance fields. In *Proceedings of Eurographics Symposium on Rendering*. 183–194.
- FUCHS, M., BLANZ, V., AND SEIDEL, H.-P. 2005. Bayesian relighting. In *Proceedings of Eurographics Symposium on Rendering*. 157–164.
- GARG, G., TALVALA, E.-V., LEVOY, M., AND LENSCH, H. P. A. 2006. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Proceedings of Eurographics Symposium on Rendering*. 251–262.
- GOLDMAN, D. B. AND CHEN, J.-H. 2005. Vignette and exposure calibration and compensation. In *The 10th IEEE International Conference on Computer Vision*. 899–906.
- HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. A dual light stage. In *Proceedings of Eurographics Symposium on Rendering*. 91–98.
- KANNALA, J. AND BRANDT, S. 2004. A generic camera calibration method for fish-eye lenses. In *Proceedings of the 17th International Conference on Pattern Recognition*. IEEE, 10–13.
- LENSCH, H. P. A., LANG, J., SÁ, A. M., AND SEIDEL, H.-P. 2003. Planned sampling of spatially varying brdfs. *Comput. Graph. For.* 22, 3 (Sept.), 473–482.
- LEVOY, M. AND HANRAHAN, P. 1996. Light field rendering. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'96)*. ACM Press, New York, NY, 31–42.
- LLOYD, S. P. 1982. Least squares quantization in PCM. *IEEE Trans. Inferr. Theory* 28, 128–137.
- MASSELUS, V., DUTRÉ, P., AND ANRYS, F. 2002. The free-form light stage. In *Proceedings of Eurographics Workshop on Rendering*. 247–256.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2003. Relighting with 4d incident light fields. *ACM Trans. Graph.* 22, 3, 613–620.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2004. Smooth reconstruction and compact representation of reflectance functions for image-based relighting. In *Proceedings of Eurographics Symposium on Rendering*. 287–298.
- MATUSIK, W., LOPER, M., AND PFISTER, H. 2004. Progressively-refined reflectance functions from natural illumination. In *Proceedings of Eurographics Symposium on Rendering*. 299–308.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. Efficient isotropic brdf measurement. In *Proceedings of Eurographics Symposium on Rendering*. 241–248.
- MATUSIK, W., PFISTER, H., NGAN, A., BEARDSLEY, P., ZIEGLER, R., AND MCMILLAN, L. 2002. Image-based 3d photography using opacity hulls. *ACM Trans. Graph.* 21, 3, 427–437.
- MOHAN, A., TUMBLIN, J., BODENHEIMER, B., GRIMM, C., AND BAILEY, R. 2005. Table-top computed lighting for practical digital photography. In *Proceedings of Eurographics Symposium on Rendering*. 165–172.
- PEERS, P. AND DUTRÉ, P. 2003. Wavelet environment matting. In *Proceedings of Eurographics Symposium on Rendering*. 157–166.
- PEERS, P. AND DUTRÉ, P. 2005. Inferring reflectance functions from wavelet noise. In *Proceedings of Eurographics Symposium on Rendering*. 173–182.
- RAMAMOORTHI, R. AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press, New York, NY, 117–128.
- ROBERTSON, M. A., BORMAN, S., AND STEVENSON, R. L. 2003. Estimation-theoretic approach to dynamic range improvement using multiple exposures. *J. Electron. Imag.* 12, 2 (April), 219–228.
- SCHECHNER, Y. Y., NAYAR, S. K., AND BELHUMEUR, P. N. 2003. A theory of multiplexed illumination. In *Proceedings of the 9th IEEE International Conference on Computer Vision (ICCV'03)*. IEEE Computer Society, 808.
- SCHIRMACHER, H., HEIDRICH, W., AND SEIDEL, H.-P. 1999. Adaptive acquisition of lumigraphs from synthetic scenes. *Comput. Graph. For.* 18, 3 (Sept.), 151–160.
- SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., AND LENSCH, H. P. A. 2005. Dual photography. *ACM Trans. Graph.* 24, 3 (Aug.), 745–755.
- SHANNON, C. E. 1949. Communication in the presence of noise. In *Proceedings of IRE*. Vol. 37, 10–21.
- WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph.* 24, 3 (Aug.), 756–764.
- WEXLER, Y., FITZGIBBON, A. W., AND ZISSERMAN, A. 2002. Image-based environment matting. In *Proceedings of Eurographics Workshop on Rendering*. 279–290.
- WEYRICH, T., PFISTER, H., AND GROSS, M. 2005. Rendering deformable surface reflectance fields. *IEEE Trans. Visualiz. Comput. Graph.* 11, 1, 48–58.

Received June 2006; revised December 2006; accepted March 2007