Eurographics Symposium on Rendering (2005)
Kavita Bala, Philip Dutré (Editors)

Bayesian Relighting

Martin Fuchs, Volker Blanz, and Hans-Peter Seidel †

MPI Informatik

Abstract

We present a simple method for relighting real objects viewed from a fixed camera position. Instead of setting up a calibrated measurement device, such as a light stage, we manually sweep a spotlight over the walls of a white room, illuminating the object indirectly. In contrast to previous methods, we use arbitrary and unknown angular distributions of incoming light. Neither the incident light nor the reflectance function need to be represented explicitly in our approach.

The new method relies on images of a probe object, for instance a black snooker ball, placed near the target object. Pictures of the probe in a novel illumination are decomposed into a linear combination of measured images of the probe. Then, a linear combination of images of the target object with the same coefficients produces a synthetic image with the new illumination. We use a simple Bayesian approach to find the most plausible output image, given the picture of the probe and the statistics observed in the dataset of samples.

Our results for a variety of novel illuminations, including synthetic lighting by relatively narrow light sources as well as natural illuminations, demonstrate that the new technique is a useful, low cost alternative to existing techniques for a broad range of objects and materials.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and RealismI.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture

1. Introduction

Example-based relighting of real objects or scenes for novel illuminations that were captured in natural environments has proven to be a powerful approach in computer graphics, producing a broad range of impressive results [ZWCS99, DHT*00,HWT*04,TSE*04,MDA02,MLP04,NN04]. Complex effects such as subsurface scattering, interreflection, shadowing and refraction are captured automatically by these techniques.

Most of these methods explicitly estimate the reflectance function at each visible point of the object or scene. This may be achieved by calibrated point light sources, such as a light stage [DHT*00, KBMK01, HWT*04], where the setup of the lights provides full control of the direction of incident illumination, or by methods that recover the light direction of a point light that is located in various positions in the room [MDA02]. A light stage can also be used to display light patterns, such as point-sampled spherical harmonics [GTW*04], for rapid capturing of moving objects.

An alternative setup is to display controlled light patterns behind the scene [PD03, MLP04], as it is known from environment matting [ZWCS99], or multiple binary point lights that are demultiplexed into single-light responses [SNB03].

For uncontrolled, natural lighting during measurements, Matusik et al. [MLP04] proposed a method that iteratively fits a reflectance function to the measured data, optimizing each point of the scene independently by quadratic programming and a decomposition of the reflectance into rectangular kernels, given the angular distribution of incident distant light in each sample image. This incident light is known explicitly either from the illumination setup (a monitor), or from images of a metallic sphere. This method has been used for relighting a city view, based on images recorded over three days.

For relighting, most authors capture novel, natural illuminations with mirror spheres used as light probes. The large variation of radiances observed on the spheres is often captured by high dynamic range (HDR) imaging. From the image of the sphere, the angular distribution of the incident distant light can be recovered. Rendering the new scene exploits the principle of superposition of light in an elegant way [DHT*00]: if the process of capturing images is approximately linear, an image of an object illuminated by an environment can be decomposed into a weighted sum of images from pre-recorded illuminations. Specifically, if the sample images are recorded with directional light, the weights for the final image are obtained from the overall novel radiance in the neighborhood of each sampled direction.

Linear superposition of light has also been used in light-

† e-mail:{mfuchs|blanz|hpseidel}@mpi-sb.mpg.de

Figure 1: *Four relighting examples (top row) as linear combination of 272 images, the coefficients being defined by novel images of a probe object (bottom, left image of each pair) which are reconstructed with the sampled probe images (right).*

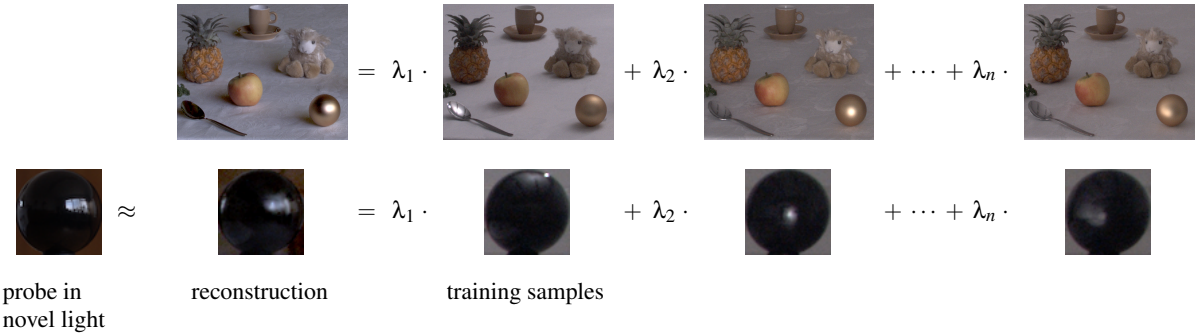

Figure 2: *Relit scene in novel illumination as linear combination of $n = 272$ previously recorded images (top row). The coefficients λ_j are found by reproducing an image of a probe object (bottom row), using a maximum-a-posteriori estimate.*

ing design algorithms that rerender virtual scenes [NJS*94, DKN*95] or relight natural objects [ADW04a, ADW04b] by combining a set of basis images. In image-based lighting design, the user paints portions of the target image in the desired color, and quadratic programming or simulated annealing algorithms compute the lighting setup that reproduces the desired appearance, which can be used to determine the settings for lighting the real object artistically [ADW04a, ADW04b].

Our novel approach exploits the linear nature of light even more by treating the reflectance function and the angular distribution of incident (distant) light only implicitly, and by transferring the linear coefficients between images of a probe object, which defines a novel illumination to be used for relighting, and images of the target object that will be relighted (Figure 2). The system uses a database of images both of the probe and target, taken in a variety of sample illuminations. Whatever the incident light maps in the sample images were, the principle of superposition implies that a linear combination of these sample light maps produces an image that is a linear combination of sample images with the same coefficients. As a consequence, images of different objects placed in the same environment can be decomposed into linear combinations with the same coefficients. Our algorithm decomposes a given image of the probe in a novel illumination into a linear combination of images from arbitrary, unknown sample illuminations. Then, we transfer the coefficients to the target object, and obtain a linear combination predicting the appearance of the target in the novel illumination.

In a Bayesian approach, our algorithm takes into account the statistical properties of the sample data in order to find the most plausible image of the target object (maximum-a-posteriori estimate), given the image of the probe object in a novel illumination. This statistical criterion avoids artifacts that would occur in a direct reconstruction, for example due to noise. Specifically, we form a tradeoff between reproducing the probe image as faithfully as possible, and maximizing the prior probability of the illumination that might have given rise to this image, based on the distribution learned from examples. This distribution can be estimated even without explicitly knowing the incident light maps $L(\omega)$ by restricting ourselves to the linear span of sample light maps, and performing a principal component analysis on the probe images rather than the unknown light maps.

2. Implicit Relighting

For a fixed viewpoint and a non-local incident light distribution $L(\omega)$, with ω denoting an incident light direction $\omega = (\theta, \phi)$, the radiance observed in a point $\mathbf{x}$ of an image is

$$I(\mathbf{x}) = \int_{\Omega} L(\omega) \cdot R(\mathbf{x}, \omega) d\omega \qquad (1)$$

where $R(\mathbf{x}, \omega)$ is the reflectance field [DHT*00]. R subsumes effects such as shadowing, foreshortening of incident light due to the unknown surface normal, interreflections among

surface elements, and subsurface scattering. The red, green and blue color channels are treated separately throughout this paper, and the indices for color channels are suppressed. In the classic light stage approach, R is sampled by directional light $L_i(\theta,\phi) = \delta(\theta-\theta_i)\cdot\delta(\phi-\phi_i)$ from discrete directions θ_i, ϕ_i, expressed in terms of the Dirac delta function $\delta()$. According to Equation (1), the images are then direct measurements of R: $I_i(\mathbf{x}) = R(\mathbf{x},\theta_i,\phi_i)$.

In contrast, our approach uses arbitrary, unknown incident light distribution $L_i(\omega)$, $i = 1,\ldots,n$ for sampling. The goal of our measurement is to sample a basis of the most relevant subspace of the vector space of functions L for subsequent relighting of natural objects. Due to the principle of superposition, linear combinations

$$L(\omega) = \sum_{i=1}^{n} \lambda_i \cdot L_i(\omega) \tag{2}$$

are physically valid incident light distributions for all $\lambda_i \geq 0$. Negative coefficients may still result in positive $L(\omega)$ if some L_i are overlapping distributions, as we discuss in Section 4. These linear combinations produce new images

$$\begin{aligned} I(\mathbf{x}) &= \int_\Omega \left(\sum_{i=1}^{n} \lambda_i \cdot L_i(\omega) \right) \cdot R(\mathbf{x},\omega) d\omega \\ &= \sum_{i=1}^{n} \lambda_i \cdot \int_\Omega L_i(\omega)\cdot R(\mathbf{x},\omega) d\omega = \sum_{i=1}^{n} \lambda_i \cdot I_i(\mathbf{x}). \end{aligned} \tag{3}$$

For relighting, the light stage technique uses weights λ_i that are computed by integrating the incident radiance L_{novel} over neighborhoods of the sampled directions [DHT*00, MDA02].

We propose a novel technique that uses Equation (3) first for estimating λ_i from an image of the probe object in the novel illumination, and then applies (3) again to relight the target object for the new image I (Figure 2). More specifically, let the image S of the probe object be

$$\begin{aligned} S(\mathbf{x}) = \sum_{i=1}^{n} \lambda_i \cdot S_i(\mathbf{x}) &= \sum_{i=1}^{n} \lambda_i \cdot \int_\Omega L_{i,\text{probe}}(\omega)\cdot R_{\text{probe}}(\mathbf{x},\omega) d\omega \\ &= \int_\Omega L_{\text{probe}}(\omega)\cdot R_{\text{probe}}(\mathbf{x},\omega) d\omega \end{aligned} \tag{4}$$

with sample images S_i of the probe. L_{probe} and R_{probe} denote the incident light distribution and reflectance field of the probe object, while the variables I, L and R refer to the images, light distribution and reflectance field of the target object. We now assume that the incident light map L and L_{probe} for the target and probe are equal throughout the process, or at least that both are formed by the same linear combinations of samples:

$$L = \sum_{i=1}^{n} \lambda_i \cdot L_i\,, \quad L_{\text{probe}} = \sum_{i=1}^{n} \lambda_i \cdot L_{i,\text{probe}}\,. \tag{5}$$

Given an image S of the probe in a novel illumination, we can then find the expansion (Equation 4) that is optimal in the least squares sense, minimizing the total error over all pixels (x,y)

$$E(\lambda_i, S) = \| \sum_{i=1}^{n} \lambda_i \cdot S_i - S \|^2 = \sum_{x,y} \left(\sum_{i=1}^{n} \lambda_i \cdot S_i(x,y) - S(x,y) \right)^2 \tag{6}$$

and obtain the set of coefficients λ_i required for relighting the target (Equation 3). The minimum of $E(\lambda_i, S)$ can be found by solving a simple linear system using standard methods such as the pseudo-inverse matrix. In the following section, we describe a more appropriate technique that employs regularization.

3. Bayesian Relighting

Solving directly for the parameters λ_i by a pseudo-inverse would produce overfitting artifacts, as shown in Figure 3 (where $\eta = 0$): first, the images of the probe object are noisy, so the system would attempt to reproduce this noise. Second, the samples do not span the full space of possible illuminations, so a least-squares reconstruction of the novel illumination would involve extreme coefficients λ_i far from the convex hull of examples. This implies that noise in the sample images of the probe would be scaled with large factors λ_i of opposite sign, causing amplified noise in the result.

Therefore, we take a maximum-a-posteriori approach (MAP, see [DHS01]) to relighting: given an image S_{novel} of the probe image in a novel illumination, we find the image I of the target object that maximizes the conditional probability $p(I|S_{\text{novel}})$ (posterior probability), based on an estimate of the prior probability p of lighting conditions from the sample set. We do not need to know the incident light map $L(\omega)$ explicitly, but only in terms of a linear combination of sample illuminations. Based on the prior probability, a regularization parameter controls how conservative our estimate will be.

To estimate the prior, we perform a principal component analysis (PCA) on the set of probe samples S_i, $i = 1\ldots n$: let $\overline{S} = \frac{1}{n}\sum_i S_i$, and $\mathbf{A}$ be the matrix formed by the columns $(S_i - \overline{S})$. PCA is based on a diagonalization of the covariance matrix: $\mathbf{C} = \frac{1}{n}\mathbf{A}\mathbf{A}^{\mathrm{T}} = \mathbf{U}\,\mathrm{diag}(\sigma_i^2)\,\mathbf{U}^{\mathrm{T}}$, where σ_i are the standard deviations of the data along the orthogonal principal component vectors $\mathbf{u}_i$ given by the columns of $\mathbf{U}$. This diagonalization is achieved by a Singular Value Decomposition [PTVF92]

$$\mathbf{A} = \mathbf{U}\mathbf{W}\mathbf{V}^T \tag{7}$$

with a diagonal matrix $\mathbf{W} = \mathrm{diag}\,(w_i)$, $\sigma_i = \frac{1}{\sqrt{n}} w_i$, and an orthogonal matrix $\mathbf{V}$. A probe image S can be written as a linear combination of the principal components

$$S = \sum_i c_i \mathbf{u}_i + \overline{S} = \mathbf{U}\mathbf{c} + \overline{S}, \tag{8}$$

where c_i are the linear coefficients. The estimated normal distribution of samples is, with a normalization factor ν_p,

$$p(\mathbf{c}) = \nu_p e^{-\frac{1}{2}\sum_i \frac{c_i^2}{\sigma_i^2}} \tag{9}$$

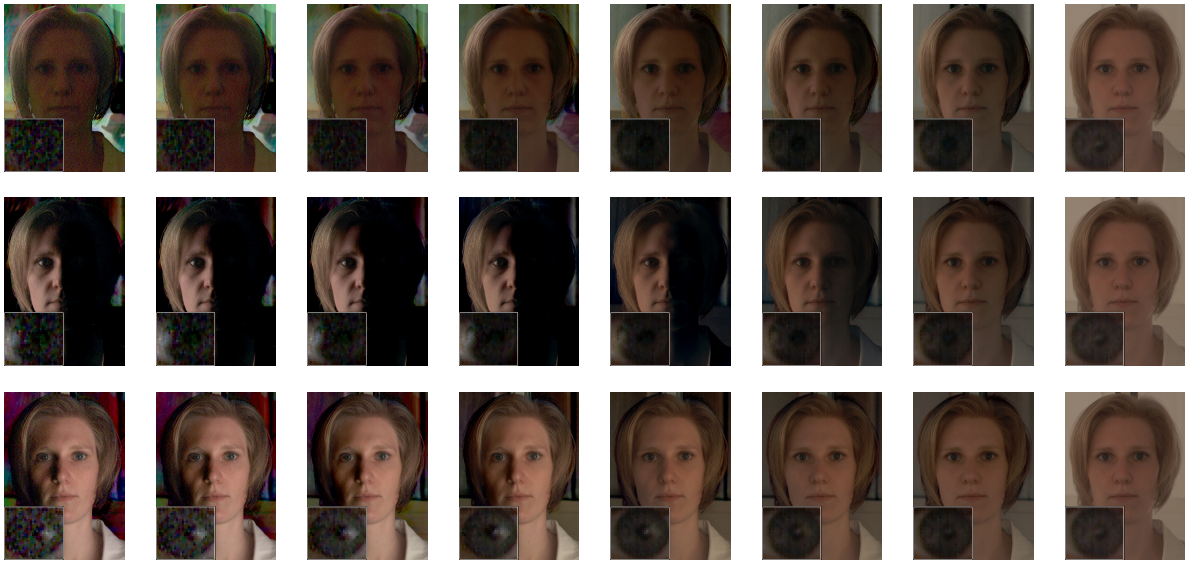

$\eta = 0$ $\eta = 0.0001$ $\eta = 0.001$ $\eta = 0.01$ $\eta = 0.1$ $\eta = 0.3$ $\eta = 1$ $\eta = 100$

Figure 3: *Each row shows a predicted image with close-up on the test subject's right eye. The parameter η is used for regularization in the renderings. The highlight in the eye and the shadow distribution demonstrate the tradeoff between a detailed lighting (low values of η) and low noise (high values η) in the rendering. For these renderings, $n = 625$ input images were used.*

within the linear span of examples. Since the coefficients of a linear combination of probe images S_i also describe the combination of light maps $L_i(\omega)$, p also captures the estimated probability density of light distributions within the span of L_i. With additive Gaussian pixel noise in the probe images S, the likelihood of an incident light map $L(\omega)$ producing S is

$$p(S|\mathbf{c}) = \nu_l \cdot \prod_{x,y} e^{-\frac{1}{2\sigma_N^2}\left(\sum c_i u_{i,x,y} + \overline{S}_{x,y} - S_{x,y}\right)^2} = \nu_l \cdot e^{-\frac{1}{2\sigma_N^2}\left\|\mathbf{U}\mathbf{c}+\overline{S}-S\right\|^2}, \tag{10}$$

with a standard deviation σ_N and a normalization factor ν_l. The norm $\|.\|^2$ denotes the sum of squared pixel differences.

According to Bayes' theorem, the posterior probability is

$$p(\mathbf{c}|S) \sim p(S|\mathbf{c}) \cdot p(\mathbf{c}), \tag{11}$$

which is maximized if a cost function given by the negative, rescaled logarithm is minimized:

$$E(\mathbf{c},S) = \left\|\mathbf{U}\mathbf{c}+\overline{S}-S\right\|^2 + \eta \sum_i \frac{c_i^2}{\sigma_i^2}, \tag{12}$$

where $\eta = \sigma_N^2$ is a regularization parameter that can be used to control how conservative the estimate is supposed to be, which depends on the anticipated measurement noise and the properties of the sampled illuminations, such as their angular distribution and angular overlap. The more complete and smooth the basis of samples, the smaller an η we may choose without producing artifacts. Figure 3 illustrates the effect of different values of η in our system.

E is minimal if $\frac{\partial E}{\partial c_i} = 0$ for all i:

$$\frac{\partial E}{\partial c_i} = 2\langle \mathbf{u}_i, \sum_k c_k \mathbf{u}_k + \overline{S} - S\rangle + 2\eta \frac{c_i}{\sigma_i^2} = 0, \tag{13}$$

which is achieved for

$$\mathbf{c} = \operatorname{diag}\left(\frac{\sigma_i^2}{\sigma_i^2+\eta}\right)\mathbf{U}^T(S-\overline{S}). \tag{14}$$

The conservative best fit can be rewritten in terms of the original basis, using $\mathbf{V}^T\mathbf{V} = \mathrm{id}$:

$$S_{\mathrm{MAP}} = \mathbf{U}\mathbf{c}+\overline{S} = \mathbf{A}\mathbf{V}\mathbf{W}^{-1}\mathbf{c}+\overline{S} = \mathbf{A}\tilde{\mathbf{c}}+\overline{S} \tag{15}$$

$$\text{where} \quad \tilde{\mathbf{c}} = \mathbf{V}\operatorname{diag}\left(\frac{1}{\sqrt{n}}\cdot\frac{\sigma_i}{\sigma_i^2+\eta}\right)\mathbf{U}^T(S-\overline{S}). \tag{16}$$

Using the definition of $\overline{S}$, we obtain

$$S_{\mathrm{MAP}} = \sum_i \lambda_i S_i\,, \quad \lambda_i = \tilde{c}_i + \frac{1}{n}(1 - \sum_k \tilde{c}_k). \tag{17}$$

These coefficients λ_i also provide the maximum-a-posteriori prediction for the target object image at a novel illumination:

$$I_{\mathrm{MAP}} = \sum_i \lambda_i I_i\,. \tag{18}$$

4. Sampling and Relighting of Objects

For collecting images of the probe and target objects at different illuminations, we use an inexpensive setup with widely available equipment. Probe and target objects should be relatively close together to make sure that they are illuminated in the same way in the sense of Equation (5). Probe

Figure 4: *Arbitrary probe object (left, 128×128 pixels), sparse sampling: an arbitrary part of the scene can be used as a probe object in the case of a fixed camera setup. The top row shows the reconstruction for sparse sampling ($n = 250$), the bottom row ground truth. While the sparseness causes blurred highlights and shadows, the reconstruction does not produce multiple blended shadow boundaries or comparable artifacts as point-based lighting sometimes does.*

and target can be captured either in the same picture, as we did, or in separate pictures taken with two cameras. We used an Olympus C5050Z digital camera for still images at a resolution of 2576×1925 pixels, an Imperx MDC 1004 video camera at 1004×1004 pixels for the data set shown in Figure 3, and an HDRC VGAx high dynamic range 640×480 video camera, courtesy of IMS-CHIPS[†], for the face data set shown in the video. For all cameras, images were captured in raw format, and a linearization and Bayer reconstruction were performed. The renderings in this paper are subject to an sRGB non-linear transform, approximating a gamma value of 2.2.

The illumination in our measurement was indirect light from the white walls and ceiling of a seminar room in our lab (Figure 5). Walking around the room, we illuminate different parts of the room with a hand-held HMI light source (Joker-Bug 800 by K5600). The method should work with any bright light source, and as the illumination may change during exposure, long exposures do not deteriorate the measurements for static objects. We avoid to hit the objects or the camera directly by using a reflector and pointing the light away from the measurement setup. In the seminar room, ceiling and walls were far enough to approximately satisfy the assumption of distant light.

While our approach does not require calibrated illumination with a known distribution $L(\omega)$, and neither ambient light nor smaller objects or darker regions in the room affect the measurements, there are two issues to take care of: first, the incident light should cover as much of the sphere around the objects as possible across different measurements. Regions that were left out cannot be incident light directions

Figure 5: *Measurement setup for the still camera: an Olympus C5050Z digital camera records objects on a table which are indirectly lit by a hand-held spotlight pointed at the white walls, ceiling and floor. The probe object, a black snooker ball, is mounted on the small tripod next to the table.*

in relighting. Second, the illumination patches, which essentially define the basis $L_i(\omega)$ of light distributions, should be overlapping and smooth: if the scene is illuminated by point lights or by small patches of indirect light from the walls, novel probe images with specular reflections between those that were measured cannot be reconstructed, and the new light directions will be missed altogether. Therefore, we started off by illuminating large portions of the room from a larger distance in overlapping patches, and then lit overlapping sequences of smaller and smaller patches.

The probe object can be any object that is sensitive to illumination changes, as Figure 4 demonstrates. For most measurements, we chose to use a sphere, since the point-by-point mapping between sampled images and images of the object at novel illumination can be established easily due to symmetry, without fixing the object to the camera. For this mapping, which is needed to find the linear combination of sam-

[†] http://www.ims-chips.de/

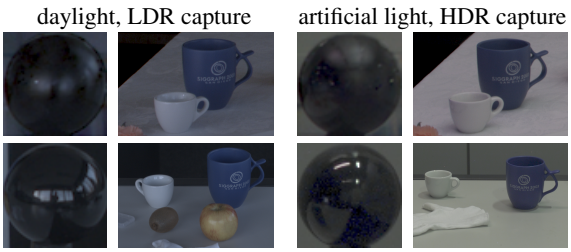

Figure 6: *Reconstruction of two light situations (top row) and ground truth images (bottom row) of a scene in similar arrangement. The left situation is captured in low dynamic range in daylight, the right situation is taken as multi-exposure image in artificial illumination. While the predicted images approximate the lighting condition only with a sparse set ($n = 272$) of input images, they match the overall brightness and distribution of highlights and shadows.*

Figure 7: *Results of relighting a test subject from $n = 75$ still images. Some artifacts arise because of movement (see edges), but lighting remains realistic.*

ples that reproduces the novel probe image best, we select the sphere by a bounding box in the images, and apply a scale and translation operation, assuming orthographic projection of the sphere. We achieve good results scaling the sphere to 64×64 pixels, and masking the non-sphere parts of the images.

In capturing incident light distributions, the dynamic range of the camera is an important issue. Most authors record images of a metallic sphere with high-dynamic range imaging to avoid saturated – and therefore underestimated – highlights on the sphere. In order to reduce the radiance at highlights, we prefer to use a black snooker ball, which reflects only a small portion of the incident light to the camera [TSE*04]: according to the Fresnel formulas, the specular reflectance of the non-metallic snooker ball is 1.0 at tangent directions, and falls off rapidly to a value of 0.04 in the center (at an index of refraction of $n = 1.5$). As a result, our probe object produces relatively dim specular reflections that are likely to be within the dynamic range of a digital camera in images that, at the same time, capture the target scene appropriately.

Our entire relighting process takes the following steps:

Training Step: Record a set of n images at different illumination with fixed cameras and static objects, define a bounding box around the probe sphere in the first image, crop and scale the probe in all images, perform a PCA on the probe images and store the result.

Prediction Step: Given a photograph of the probe in a novel illumination, we crop and scale the probe again from the image, compute $\mathbf{c}$, $\tilde{\mathbf{c}}$ and λ_i (Equations 14,15,17) and form the weighted sum of sampled images I_i (Equation 17).

The variation in overall brightness in our sample sets turned out to be sufficient to cover the variations in novel illuminations without rescaling.

Unlike most previous methods, our linear combinations may involve negative coefficients λ_i: in the classical light stage approach [DHT*00] and most subsequent methods, linear coefficients are positive, since they are weights proportional to the incident radiance in the neighborhoods of discrete light directions. Matusik et al. [MLP04] enforce constraints on the model coefficients.

In our setting, negative coefficients arise from the overlapping base functions. Still, they do not imply physically invalid results: consider an image with two lights A and B, and one with A only. The difference image reproduces the situation with B only, and all resulting pixel values represent valid positive radiances. However, negative color values may still occur in our least-squares framework within the range of approximation errors. In a second iteration, we alleviate this problem by fitting against an image consisting of the inverse of negative result pixels on the probe (with a smaller value for η), and adding the resulting coefficients to the previous results.

5. Results

Figure 1, 2 and 8 show images of objects that were relighted with our system, demonstrating the high spatial detail that can be achieved with still camera measurements.

The training step for these images contained 272 sample images, taken in about 30 minutes. The computation time for the PCA on the probe pictures took about 20 seconds per color channel on a PC with 3 GHz Intel Xeon Processor. Determining the coefficients λ_i for some target image S_{target} takes less than 2 seconds, and reconstructing an image takes between 1 (for 64×64 images) and 16 seconds (for 708×560 pixels). These numbers are for the data set from Figure 2, but are comparable to the others. The timings are performed after transferring the pictures from the camera's Compact Flash Card, and reconstructing the Bayer pattern in each of the input images.

Figure 8 shows the wide range of material appearances which are captured by our approach: cloth (napkin), polished metal (cutlery), glossy objects (orange, plate), transparent objects (wine inside glass) and even near-field caustics (jelly). All of these are plausibly relighted. For a ground truth comparison, we reproduced lighting of two scenes, as shown in Figure 6.

For live objects, such as human faces, the acquisition time is an issue. Figure 7 demonstrates that from $n = 75$ sam-

Figure 8: *Various materials relighted, $n = 272$. The rows a show reproduced natural lighting (as in Figure 1, second column), the rows b show synthetic lighting of a directional dominant light source in an otherwise totally dark room. The objects are from four different data acquisitions: (i) dish, napkin and wine, (ii) jelly, (iii) spoon, sheep and juice, (iv) oranges.*

ple illuminations, which were captured in about 8 minutes, interesting effects of human skin structure can already be captured. However, especially for the perceptually important visual properties of human eyes, a setup where the eyes are open is more appropriate; therefore, we also performed experiments with two faster cameras. In one experiment, the test subject (Figure 3) was recorded for 25 seconds with a video camera, yielding 625 input images which allow us to recreate even specular highlights in the eyes. In another experiment (see supplemental video of a smiling person holding a toy animal), we used an HDR video camera yielding 1000 frames in a comparable time.

Even though the cameras gave us abundant image data in a short period of time, we had to record for about 25 seconds to sweep the light source's cone over the wall manually, covering a sufficient set of light conditions. Residual movements of the test subject who was recorded with the still camera (Figure 7), which become visible as relief-like artifacts in the supplementary video, are less prominent in the video acquisition setup (Figure 3) due to shorter recording intervals.

In the background of the synthetic images, behind the target objects, our technique tends to produce ghost images that show the experimenter and the spotlight, as the experimenter becomes part of the distant incoming light environment. We masked the background in the video, and cropped the images to the object region in the Figures.

Although it is designed for relighting with natural illuminations, our approach can also be used for synthetic relighting, based on renderings of a snooker ball. For Figures 8 (row *b*), 3, and the rotating light source in the video, we rendered a sphere with Phong BRDF and an additional Fresnel term for a refractive index $n_{\text{refract}} = 1.5$. The Phong exponent gives us an easy control of the distribution of incoming light; by choosing a low exponent, an extended light source is simulated. Synthetic images of the ball created with a ray-tracer or global illumination techniques could be used as well.

For illumination design, as shown in Figure 1 (right), the user draws patterns of incoming light with standard imaging software into an image of the probe object, which is then reconstructed by our Bayesian method for transferring the lighting on the target. This is unlike previous methods [ADW04a, ADW04b], where the lighting was designed in the target image directly. For practical applications, both approaches are useful, but they address different design purposes.

6. Conclusion

The contributions of our method are a new theoretical approach for relighting, and a low-cost system that requires no light stage or other sophisticated setup or equipment. From a maximum-a-posteriori approach, we have derived a simple mathematical formula which makes the relighting algo-

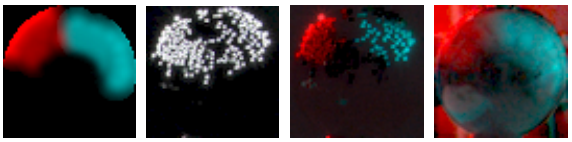

(i) (ii) (iii) (iv)

Figure 9: *Comparison to impulse-response sampling: in order to reconstruct a target (i), sampling $n = 271$ non-uniformly distributed points (ii) inside the target image's active area creates holes and non-smooth artifacts (iii), even though we sampled light directions favorably. In contrast, our method with extended incoming light sources (iv) gives a continuous reconstruction.*

rithm easy to implement. We hope that our technique helps to make relighting more available to a broad range of users.

In contrast to previous work, our method does not apply point-light illuminations for sampling [DHT*00, HWT*04, MDA02], but low spatial frequency illuminations. Both concepts work well for diffuse objects in arbitrary lighting, but involve different tradeoffs for specular objects and cast shadows, given a limited set of sample illuminations. Our method tends to blur highlights and shadow edges, as can be seen in Figure 4, but reconstructs extended light sources well. Impulse-response methods reproduce extended sources by individual points, as illustrated in Figure 9. Also, in animations with moving directional light sources, sharp specular highlights fade in and out, while our method produces smoothly moving, but slightly broader highlights (as seen in the eyes in the supplemental video).

As a conceptual advantage, our implicit approach learns the mapping between probe and output images directly, rather than investing in the estimate of intermediate information, such as incident light [Deb98] or reflectance [MLP04]. We have proposed a new, more general notion of a light probe object, which makes the method interesting for new applications in fixed camera setups. We presented a result employing toy figures for that purpose.

Our method fits seamlessly into existing acquisition pipelines that measure incident light distributions explicitly, as the mapping of the light distribution to the snooker ball is straight-forward. However, this is not our primary goal, since we propose a different measurement process for determining illumination that is equally simple as the conventional method of capturing a mirror sphere. It is easy to improve the speed of the illumination sampling by technical means, as the measurement setup is uncalibrated.

From a given set of sampled illumination conditions, our statistical approach enables us to predict a relighted image in an optimal sense without explicit knowledge of the object or lighting properties, making it a consequent implementation of learning-based computer graphics.

References

[ADW04a] ANRYS F., DUTRÉ P., WILLEMS Y. D.: Image based lighting design. In *The 4th IASTED Int. Conf. on Visualization, Imaging, and Image Proc.* (2004). 2, 7

[ADW04b] ANRYS F., DUTRÉ P., WILLEMS Y. D.: *Lighting Design by Simulated Annealing*. Technical report CW 393, Departement Computerwetenschappen, KU.Leuven, 2004. 2, 7

[Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH '98:* (1998), ACM Press, pp. 189–198. 8

[DHS01] DUDA R., HART P., STORK D.: *Pattern Classification*, 2nd ed. John Wiley & Sons, New York, 2001. 3

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *SIGGRAPH 2000* (2000), ACM Press/Addison-Wesley Publishing Co., pp. 145–156. 1, 2, 3, 6, 8

[DKN*95] DOBASHI Y., KANEDA K., NAKASHIMA T., YAMASHITA H., NISHITA T.: A quick rendering method using basis functions for interactive lighting design. In *Computer Graphics Forum, Vol. 14, No. 3 EG'95* (1995), pp. 229–240. 2

[GTW*04] GARDNER A., TCHOU C., WENGER A., HAWKINS T., DEBEVEC P.: Postproduction re-illumination of live action using interleaved lighting. In *SIGGRAPH 2004 (Poster)* (2004). 1

[HWT*04] HAWKINS T., WENGER A., TCHOU C., GARDNER A., GÖRANSSON F., DEBEVEC P.: Animatable facial reflectance fields. In *Rendering Techniques 2004* (2004), Eurographics Assoc., pp. 309–319. 1, 8

[KBMK01] KOUDELKA M., BELHUMEUR P. N., MAGDA S., KRIEGMAN D. J.: Image-based modeling and rendering of surfaces with arbitrary brdfs. In *Proceedings of IEEE CVPR 2001* (2001), pp. 568–575. 1

[MDA02] MASSELUS V., DUTRÉ P., ANRYS F.: The free-form light stage. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering* (2002), Eurographics Association, pp. 247–256. 1, 3, 8

[MLP04] MATUSIK W., LOPER M., PFISTER H.: Progressively-refined reflectance functions from natural illumination. In *Rendering Techniques 2004* (2004), Eurographics Association, pp. 299–308. 1, 6, 8

[NJS*94] NIMEROFF, J., SIMONCELLI, E., DORSEY J.: Efficient re-rendering of naturally illuminated environments. In *5th Eurographics Workshop on Rendering* (1994). 2

[NN04] NISHINO K., NAYAR S. K.: Eyes for relighting. *ACM Trans. Graph. 23*, 3 (2004), 704–711. 1

[PD03] PEERS P., DUTRÉ P.: Wavelet environment matting. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering* (2003), Eurographics Association, pp. 157–166. 1

[PTVF92] PRESS W. H., TEUKOLSKY S. A., VETTERLING W. T., FLANNERY B. P.: *Numerical Recipes in C*. Cambridge University Press, Cambridge, 1992. 3

[SNB03] SCHECHNER Y. Y., NAYAR S. K., BELHUMEUR P.: A theory of multiplexed illumination. In *Proc. Int. Conf. on Computer Vision (ICCV)* (2003), pp. 808–815. 1

[TSE*04] TCHOU C., STUMPFEL J., EINARSSON P., FAJARDO M., , DEBEVEC P.: Unlighting the parthenon. In *SIGGRAPH 2003, Sketches and Applications* (2004). 1, 6

[ZWCS99] ZONGKER D. E., WERNER D. M., CURLESS B., SALESIN D. H.: Environment matting and compositing. In *SIGGRAPH '99* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 205–214. 1