Relighting with 4D Incident Light Fields

Vincent Masselus* Pieter Peers* Philip Dutré* Yves D. Willems *

Department of Computer Science
Katholieke Universiteit Leuven

Abstract

We present an image-based technique to relight real objects illuminated by a 4D incident light field, representing the illumination of an environment. By exploiting the richness in angular and spatial variation of the light field, objects can be relit with a high degree of realism.

We record photographs of an object, illuminated from various positions and directions, using a projector mounted on a gantry as a moving light source. The resulting basis images are used to create a subset of the full reflectance field of the object. Using this reflectance field, we can create an image of the object, relit with any incident light field and observed from a fixed camera position.

To maintain acceptable recording times and reduce the amount of data, we propose an efficient data acquisition method.

Since the object can be relit with a 4D incident light field, illumination effects encoded in the light field, such as shafts of shadow or spot light effects, can be realized.

CR Categories: I.2.10 [Artificial Intelligence]: Vision and Scene Understanding—intensity, color, photometry and thresholding I.3.7 [Computer Graphics]: Three Dimensional Graphics and Realism—color, shading, shadowing and texture I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—radiometry, reflectance, scanning

Keywords: Image-Based Techniques, Relighting, Reflectance Field, Light Field.

1 Introduction

Image-based relighting has become a well studied and popular topic in computer graphics. Several fast techniques have been developed to generate relit images of real objects. All these image-based methods use incident light maps, also known as environment maps, that represent the incoming light at a single point.

These techniques deliver visually pleasing results. However, there is no spatial variation in the illumination used, since a light map encodes incident illumination at one point. Relighting an object using such a light map is only correct for the exact point at which the light map was recorded. The resulting pictures showing the re-illuminated object are visually satisfying; however, some effects are impossible to generate.

*e-mail:{vincentm,pieterp,phil,ydw}@cs.kuleuven.ac.be

Figure 1: An arrangement of chess pieces, illuminated using a traditional light map (left), and illuminated by an incident light field containing two spot lights, computed using our algorithm (right).

Suppose that we want to capture the illumination effect of a spot light in such an incident light map. Only when the beam of the spot light hits the point at which the light map is recorded, will the object be illuminated. An example of this is given on the left of figure 1. Otherwise, the illumination map does not capture the illumination of the spot light and the relit object would be completely dark.

In reality, the object may be partly illuminated by the spot light or even multiple spot lights could be used, as shown on the right of figure 1. To accomplish these effects with a relighting algorithm, we need to accomodate for angular as well as spatial variation in the incident illumination. In other words, we need to represent the illumination incident on a volume, not in a single point.

A light field represents the illumination leaving a volume for any position and in any direction, and is described by a 4D function. Light fields have been used mainly in image-based rendering techniques for displaying (real) objects from any viewpoint without knowledge of geometry or material properties. Such light fields are called radiant light fields or exitant light fields. Similarly, incident light fields can be constructed, in which a viewer can walk and look around in a limited space and see the surrounding environment represented by the light field.

To relight an object with a 4D incident light field, we need to capture the reflectance field of the object for any incident lighting configuration. The reflectance field of an object determines the illumination transfer through a volume containing the object. In this paper, we will address the problem of capturing the reflectance field of an object, and relighting the object with a given incident light field. To illustrate our technique, we use artificial incident light fields, although real environments could be used as well.

To acquire the reflectance field of an object, we photograph the object illuminated with a number of basis incident light fields. These basis fields constitute a linear basis for expressing incident light fields that we want to use during the relighting stage. Each illumination condition is created by a standard LCD projector, which is mounted on a movable gantry. As a result, we have a set of basis images, which altogether capture the total reflectance field of the object.

To relight an object we express the given light field as a linear combination of the basis light fields. A final image of the relit object can be created using the acquired reflectance field.

In the next section we give an overview of some related work and in section 3, we specify the parameterizations used for incident light fields and reflectance fields. In section 4, we outline our technique. We propose a mathematical framework in 4.1 and propose a method to capture a 6D slice of the full reflectance field. We discuss the data acquisition (4.2) of the basis images, and the relighting (4.3) of the real objects. Because of the sheer amount of data and acquisition time, we developed an effective speedup in the data capturing process. This is outlined in section 5. In section 6 we discuss the results obtained with our technique. Section 7 concludes the paper and indicates directions for future research.

2 Related work

Image-based rendering is a widely researched area within the field of computer graphics. Basically, these methods try to capture and visualize the 5D plenoptic function [Adelson and Bergen 1991]. The Lumigraph [Gortler et al. 1996] and Light Field Rendering [Levoy and Hanrahan 1996] showed that this 5D function could be represented as a 4D function, the light field, when the viewer is constrained to the outside of a bounding volume of the object. Wood et al. [2000] introduced the idea of surface light fields, a 4D representation of appearance of objects using some rough approximation of the geometry but generating photo-realistic results.

Although these image-based rendering methods produce high quality images at dazzling speeds, they often require geometry information and are restricted to fixed scenes and fixed illumination. Image-based relighting investigates the generation of images of objects or scenes with variable illumination without knowledge of the geometry of the object or scene.

Since illumination behaves linearly with relation to intensity, Nimeroff et al. [1994] noticed that a weighted sum of basis images can result in an image of the relit object. In their technique, basis images are generated using a set of illumination conditions defined by the theory of steerable functions. Each basis image is a rendering of the object from a fixed viewpoint using one of these illumination conditions, making this technique only practical for virtual objects or scenes.

Wong et al. [1997] introduced the concept of an apparent BRDF. For each pixel on the image plane, an apparent BRDF is created from a set of images of the object lit by a directional light source. The object can be rendered with any directional illumination by evaluating the apparent BRDF for each pixel using this direct illumination. In Wong et al. [2001] spatial variation in the illumination is possible and effects of point light sources and even spot lights can be achieved in panoramic scenes using the apparent BRDF method. However, geometrical information is required, rendering this technique less practical for real objects.

Lin et al. [2001] developed a method to relight real objects, inspired by the Lumigraph. In their method, the viewpoint is fixed and the illumination is variable. A set of images of the object is taken from a fixed viewpoint, lit by a point light source which is mechanically placed at known locations on a plane. By resampling the data, the real object can be correctly relit with any point light source positioned on that plane and even directional light sources are possible. Spatial variation in the illumination is, however, not possible.

Koudelka et al. [2001] proposed an image-based technique to relight objects with arbitrary positioned point light sources. The surface geometry of the object can be obtained by estimating the depth for each visible point, using two sets of basis images. With this approximation of the geometry and one of the basis sets of images, a relit image can be rendered.

Debevec et al. [2000] introduced the reflectance field of an object as an 8D function. It determines how an incident light field on the object is transformed into a radiant light field, leaving the object. By restricting the incident light field to incident directional light and reducing the radiant light field to one viewpoint, this 8D function can be reduced to a 4D slice of the 8D reflectance field. For real objects, an approximation of this 4D slice was constructed by sampling incident directional light on the object and taking photographs of it. This sampling was done using a movable light source mounted on a gantry, called a Light Stage, emitting directional light to the object. Several other versions of this Light Stage have been developed [Hawkins et al. 2001; Masselus et al. 2002] improving data acquisition time or portability and scalability. In the same paper it was pointed out that using more viewpoints leads to an approximation of a 6D slice, restricting the reflectance field to only incident directional light. This was fully implemented by Matusik et al. [2002a; 2002b]. Using a set of cameras, light sources and plasma screens, a huge amount of data was collected and processed, enabling the user to view a directionally relit object from any viewpoint.

In this paper, we also record a 6D slice of the reflectance field, but we keep the camera fixed and use 4D incident light fields.

In [Debevec et al. 2002], a technique was introduced to relight objects in real life. The illumination of an environment is recreated by the use of a Light Stage 3.0, which is a geodesic dome containing color changing light sources aimed at the center of the dome. The objects inside the dome can then be illuminated as if they were in the environment. This paper hinted at the possibility of using video projectors instead of directional uniform emitting light sources to lit objects with spatial varying illumination. We pursued this idea of using projectors to capture a 6D slice of the reflectance field.

3 Light fields and reflectance fields

As with the Lumigraph [Gortler et al. 1996] and Light Field Rendering [Levoy and Hanrahan 1996], we enclose a part of the environment by a surrounding surface and represent all light incident on this volume as a 4D function, which is called the incident light field. We choose a hemisphere as the bounding surface since it is natural when considering incident light fields and is closely related to the way we acquire our data set (see section 4.2). In figure 2 an incident light field on an irregular volume and the same incident light field on a hemisphere is displayed.

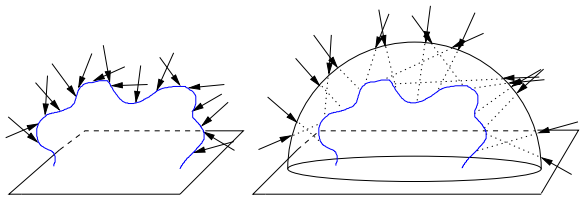

Figure 2: On the left, an incident light field is shown on an irregular volume, on the right, the same light field represented on a hemisphere.

On a hemisphere, we parameterize a position with the azimuth angle ϕ_p and tilt angle θ_p. For a chosen point on the hemisphere, a direction can be parameterized in a local frame, with ϕ_a and θ_a as the azimuth and tilt angle (Figure 3). This parameterization of the plenoptic function was introduced by [Ihm et al. 1997] and was used to render light fields in real time using graphics hardware.

A light map, representing light incident on a single point, is a special case of an incident light field. In a light map, the value of the incident light field is only dependent on ϕ_p and θ_p, and independent of ϕ_a and θ_a.

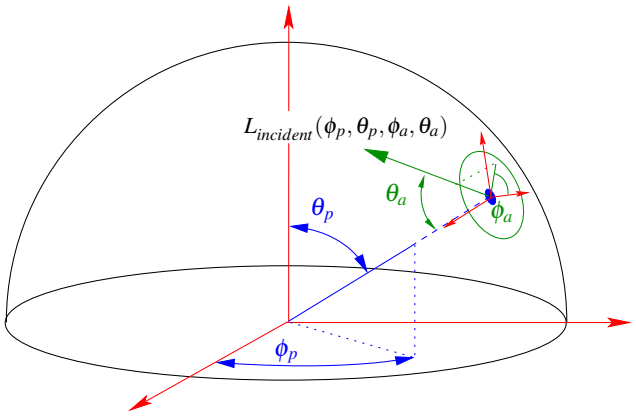

Figure 3: Parameterization of an incident light field on a hemisphere. ϕ_p and ϕ_a range from 0 to 2π. θ_p and θ_a range from 0 to $\pi/2$ with 0 being the direction perpendicular to the plane containing ϕ.

Let Ω be the 4D space of all possible incident directions on all points of a hemisphere:

$$\Omega = [0,2\pi] \times [0,\pi/2] \times [0,2\pi] \times [0,\pi/2].$$

Within this space Ω, we can define a variable Θ as:

$$\Theta = (\phi_p, \theta_p, \phi_a, \theta_a).$$

Using this parameterization, a 4D incident light field $\mathbf{L_{incident}}$ can be expressed as:

$$\mathbf{L_{incident}} = \mathbf{L_{incident}}(\phi_p, \theta_p, \phi_a, \theta_a) = \mathbf{L_{incident}}(\Theta).$$

As introduced by Debevec et al. [2000], the reflectance field $\mathbf{R}$ is an 8D function. It determines the light transfer between light entering a bounding volume at a direction and position $\Theta_{incident}$, and leaving at $\Theta_{exitant}$:

$$\mathbf{R} = \mathbf{R}(\Theta_{incident}, \Theta_{exitant}).$$

We restrict the reflectance field to a position on the view plane, so the radiant light field is the 2D function $\mathbf{L_{exitant}}(x,y)$, with x and y parameterizing the image plane (Figure 4).

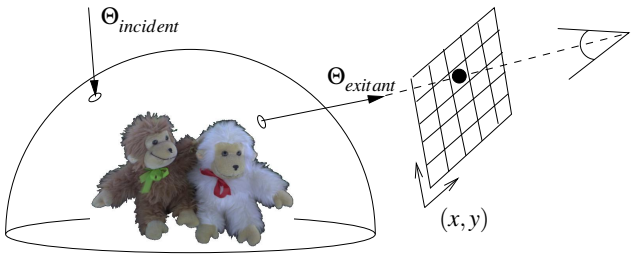

Figure 4: A 6D slice of the reflectance field determining the light transfer from $\Theta_{incident}$ towards (x,y) on the image plane.

Since we only consider a fixed camera position, our reflectance field $\mathbf{R}$ can be simplified to a 6D function:

$$\mathbf{R} = \mathbf{R}(\phi_p, \theta_p, \phi_a, \theta_a, x, y) = \mathbf{R}(\Theta_{incident}, x, y).$$

The reflectance field $\mathbf{R}(\Theta_{incident}, x, y)$ determines how an incident light field $\mathbf{L_{incident}}(\Theta)$ is transfered through the volume. We can express the total value of $\mathbf{L_{exitant}}(x,y)$ by the relighting equation as follows:

$$\mathbf{L_{exitant}}(x,y) = \int_\Omega \mathbf{R}(\Theta,x,y)\mathbf{L_{incident}}(\Theta)d\mu(\Theta), \quad (1)$$

with $\mu(\Theta)$ a measure in space Ω.

4 Acquiring and using the 6D reflectance field

In this section we will develop a mathematical framework to express a light field as a linear combination of basis functions. The relighting equation can be approximated using this basis (section 4.1). In section 4.2, we develop a setup to illuminate and capture the object with the light field basis. Using this data we are able to relight the object with an incident light field in section 4.3.

4.1 Mathematical framework

We create a set of N disjunct partitions Ω_i of Ω. The union of all these partitions Ω_i cover Ω completely:

$$\bigcup_i \Omega_i = \Omega \quad \text{and} \quad \Omega_i \bigcap \Omega_j = \emptyset \quad \text{if} \quad i \neq j, \forall\, i,j \in \{1,...,N\}. \quad (2)$$

We define a set of N basis functions using these partitions Ω_i. Each basis function $\mathbf{B_i}$ is set to be the constant value 1 on the interval of Ω_i and 0 outside Ω_i:

$$\mathbf{B_i}(\Theta) = \begin{cases} 1 & \text{if} \quad \Theta \in \Omega_i \\ 0 & \text{if} \quad \Theta \notin \Omega_i. \end{cases} \quad (3)$$

With these N basis functions, we can approximate the light field $\mathbf{L_{incident}}(\Theta)$ as a linear combination:

$$\mathbf{L_{incident}}(\Theta) \approx \sum_{i=1}^{N} l_i \mathbf{B_i}(\Theta).$$

Substituting in equation 1 delivers:

$$\begin{aligned} \mathbf{L_{exitant}}(x,y) &= \int_\Omega \mathbf{R}(\Theta,x,y)\mathbf{L_{incident}}(\Theta)d\mu(\Theta) \\ &\approx \int_\Omega \mathbf{R}(\Theta,x,y)\left(\sum_{i=1}^{N} l_i \mathbf{B_i}(\Theta)\right)d\mu(\Theta) \\ &\approx \sum_{i=1}^{N} l_i \int_\Omega \mathbf{R}(\Theta,x,y)\mathbf{B_i}(\Theta)d\mu(\Theta) \\ &\approx \sum_{i=1}^{N} l_i \int_{\Omega_i} \mathbf{R}(\Theta,x,y)d\mu(\Theta), \end{aligned}$$

or

$$\mathbf{L_{exitant}}(x,y) \approx \sum_{i=1}^{N} l_i \mathbf{R_i}(x,y), \quad (4)$$

where

$$\mathbf{R_i}(x,y) = \int_{\Omega_i} \mathbf{R}(\Theta,x,y)d\mu(\Theta). \quad (5)$$

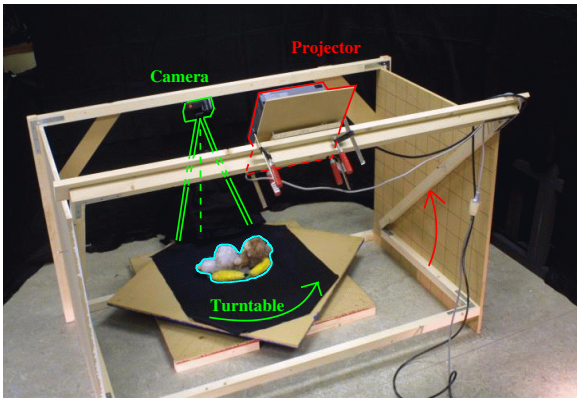

Figure 5: This picture shows our setup for measuring the reflectance field of an object (outlined in blue). The projector (outlined in red), moves on a circular path over the object. The object itself is positioned on a turntable, to which the camera is attached (outlined in green).

4.2 Data acquisition

In order to relight an object with an arbitrary incident light field expressed in terms of our basis defined in equation 3, we need to record all the terms of $\mathbf{R_i}$ (Equation 5). Each $\mathbf{R_i}$ is the transfer of an incident light field with radiance 1 W/m^2sr in the partition Ω_i. We can obtain the value for $\mathbf{R_i}(x,y)$ by taking a photograph of the object with the correct illumination $\mathbf{B_i}$.

We use an LCD projector as a light source, mounted on a moveable gantry, and at the same time, the object is placed on a turntable. The gantry is able to move the projector from 0 to 90 degrees, and the turntable is able to rotate the full 360 degrees (Figure 5). Since the position of the camera is fixed relative to the object, it is attached to the turntable as well. The combination of both movements allows us to place the projector in any position on the hemisphere relative to the object, thereby acquiring a sampling of (ϕ_p, θ_p) positions.

In each of the positions, the projector emits a series of white square patterns of constant intensity over some part of its projecting angle. Each bundle thus projected accounts for a single illumination of the object due to $\mathbf{B_i}$, and effectively covers the entire (ϕ_a, θ_a) range, relevant for our object. Figure 6 gives a schematic overview of the acquisition.

The gantry moves the projector to each position on the hemisphere, aimed at its center, and the angular range of the projector is used to emit a pyramid of light. Not all directions for a projector position can be sampled, since the angular range of a standard video projector is not a complete hemisphere. However, only the reflectance of the object itself is of interest, and that requires only a subset of all directions for each position. We can record all required information as long as the object can be illuminated completely. This constrains the size of the recordable objects by the angular range of the projector.

The digital camera as well as the projector are calibrated, using a GretagMcBeth ColorChecker DC, to account for color shifts. The turntable and gantry are currently operated manually. This results in the need for human intervention each time the position of the gantry or turntable needs to be changed. However, the data acquisition could be automated by using a motorized turntable and gantry, rendering human intervention obsolete.

We transform the pictures into high dynamic range (HDR) images, using the inverse response curve of the digital camera [Debevec and Malik 1997].

In our implementation, the images have a resolution of 1440 × 960 pixels, recorded with a Canon EOS D30 digital camera. The

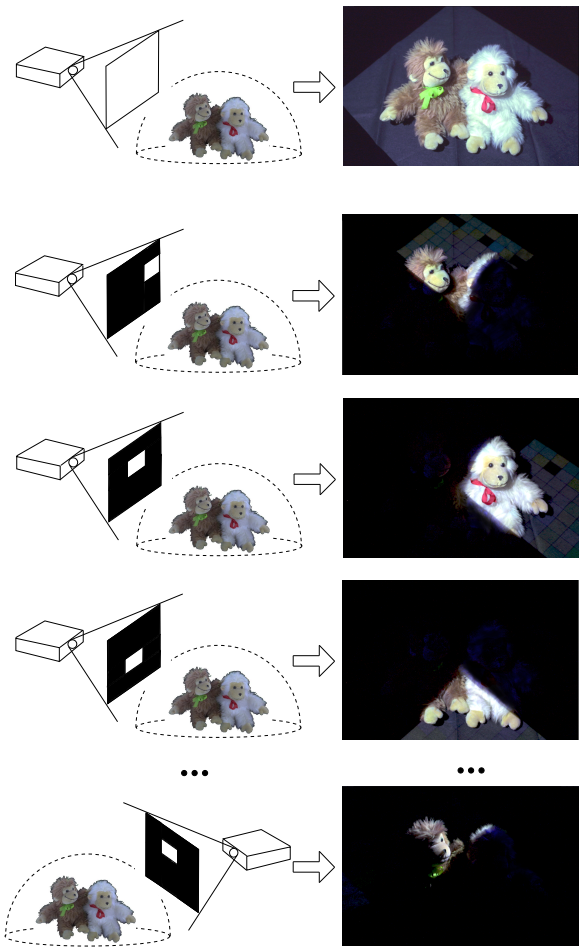

Figure 6: The topmost picture shows the object illuminated by the projector emitting white light over its entire angular range. This would be similar to data acquisition by the Light Stage. Our technique has the projector emit various basis light fields, resulting in a partially illuminated object.

number of basis light fields $\mathbf{B_i}$ will directly relate to the level of detail in the relit image (i.e. the sampling density of the incident light field). Using more projector positions and more patterns requires more basis images and, as such, a longer acquisition time.

In our setup, we used 32×7 projector positions and for each position 16×16 patterns. This enables fairly detailed relit objects while the amount of data and time to capture it, is still within limits. Using images with a resolution of 1440×960, RLE compression and a noise filter on the images, the total amount of data is about 10Gb. With Huffman compression, this can further be reduced to 6Gb. Using an effective speedup, described in section 5, the acquisition time of the basis images can be reduced to 11 minutes for each projector position, or 41 hours for the complete data capture. We expect this acquisition time could be significally reduced if the gantry and turntable would be automated and if a digital video camera were to be used.

Areas outside the bundle of light projected on the object should in theory receive no direct light at all. However, emitting 'black' pixels using a projector still results in some amount of emitted light. This has to be corrected in the recorded images. Therefore, for every projector position, an additional image is recorded while emitting a completely black pattern. From the difference of each basis image and the additional ambient image, the effect of only the white square pattern can be computed.

4.3 Relighting

The object can be illuminated with any incident light field once the basis images are recorded. Following equation 4, we need to find l_i for each interval Ω_i, given a target incident light field $\mathbf{L_{incident}}$. By projecting $\mathbf{L_{incident}}$ on basis $\mathbf{B_i}(\Theta)$ we find:

$$\mathbf{L_{incident}}(\Theta) \approx \sum_{i=1}^{N} l_i \mathbf{B_i}(\Theta),$$

thus

$$\int_{\Omega} \mathbf{L_{incident}}(\Theta)\mathbf{B_j}(\Theta)d\mu(\Theta) \approx \sum_{i=1}^{N} \int_{\Omega} l_i \mathbf{B_i}(\Theta)\mathbf{B_j}(\Theta)d\mu(\Theta)$$

$$\int_{\Omega_j} \mathbf{L_{incident}}(\Theta)d\mu(\Theta) \approx \sum_{i=1}^{N} \int_{\Omega_j} l_i \mathbf{B_i}(\Theta)d\mu(\Theta).$$

Using equation 2:

$$l_j = \frac{\int_{\Omega_j} \mathbf{L_{incident}}(\Theta)d\mu(\Theta)}{\int_{\Omega_j} d\mu(\Theta)},$$

where l_j is the mean of the incident light field $\mathbf{L_{incident}}$ over interval Ω_j. Since we use artificial light fields, we can compute these values with a global illumination renderer. The relit object is then given by:

$$\mathbf{L_{exitant}}(x,y) = \sum_{i=1}^{N} l_i \mathbf{R_i}(x,y).$$

A weighted sum of the pixels (x,y) in the recorded basis images is created since $\mathbf{R_i}(x,y)$ represents the pixel intensity. To render a complete image, we do this for all pixels (x,y). The weights in the weighted sum remain constant for all pixels, thus we can create a weighted sum for the basis images, which will result in the relit image.

5 Data acquisition speedup

Capturing the reflectance field of an object is an enormous task. A large amount of photographs has to be recorded. We propose a practical technique to reduce this number of photographs. Suppose we need to project $M \times M$ square patterns, tiled in a grid, for a single projector position.

The idea is based on the assumption that the bundles of light have a limited, local influence and that we can emit more than one square pattern for each photograph. An image of each individual square pattern can be reconstructed afterwards using the photographs of the multiple emitted square patterns.

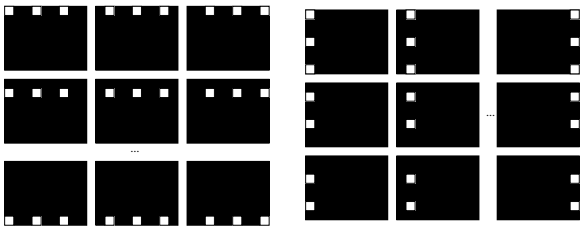

Figure 7: On the left, the set of horizontal patterns is shown, on the right, the set of vertical patterns. Each set of patterns is split up into 3 sets, in such a way that the patterns are spaced out evenly.

We create a horizontal set of patterns H as shown on the left of figure 7. Each H is constructed in such a way that each individual square pattern does not influence another square pattern in the same horizontal pattern H, once projected. A similar set, the vertical set of patterns V is created as well, as depicted on the right of figure 7. An original square pattern $\mathbf{B_i}$ occurs twice: once in a horizontal set H_x and once in a vertical set V_y. By construction we know that all other squares in the pattern do not influence $\mathbf{B_i}$ *and* that the intersection of H_x and V_y equals $\mathbf{B_i}$.

The image of projecting an individual square $\mathbf{B_i}$ can be reconstructed by taking for each pixel the minimum of the corresponding pixel values in the images resulting from projecting H_x and V_y. This reconstruced image is an approximation of the photograph resulting from emitting $\mathbf{B_i}$. In figure 8, an illustration of this technique is given.

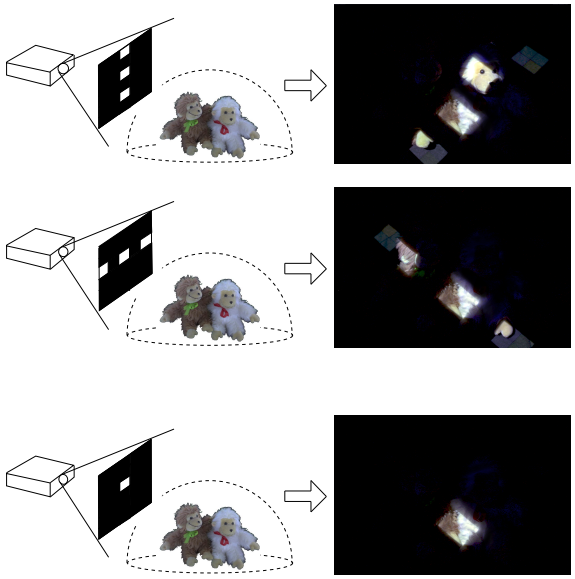

Figure 8: An illustration of the use of patterns on the two toy monkeys. On top, one vertical and one horizontal pattern are projected and the resulting images are shown. At the bottom, we project a pattern with one white square and the resulting image is displayed. This resulting image at the bottom is approximated by taking the minimum per pixel of the two top images which are the result of the projection of the horizontal and vertical patterns above.

We assume that a square pattern has local influence and that the square patterns in each horizontal and vertical pattern do not interfere with each other once projected (except for the square pattern they share).

If these assumptions are not met, this technique will fail. For example, if the object were a diffuse concave bowl pointed toward the camera, there would be a lot of interreflection in the bowl. Using the minimum on pixel values would not result in the effect of a single square pattern due to this indirect light. In this case, this technique would fail and one would be forced to project each $\mathbf{B_i}$ independently. We however, had no problems using it, since the objects we photographed had only limited interreflections. By introducing a distance between the squares in the sets of patterns, the amount of allowed interreflection can be chosen. In our implementation, we split up the sets of patterns for each row or column into 3 separate sets. This reduced the amount of photographs from 16×16 to $3 \times 2 \times 16$. In general we can reduce the complexity of our data acquisition from $\mathcal{O}(M^2)$ to $\mathcal{O}(M)$, with M^2 being the number of squares to be projected.

6 Results

To illustrate the 4D light fields we use for relighting, we use a visualisation as shown in figure 9. For the full grid of 32×7 (ϕ_p, θ_p) points on the hemisphere, we have 16×16 (ϕ_a, θ_a) values. Thus, each little square shown in the mosaic shows the incoming illumination at a specific point on the hemisphere and each pixel in such a square corresponds to a direction.

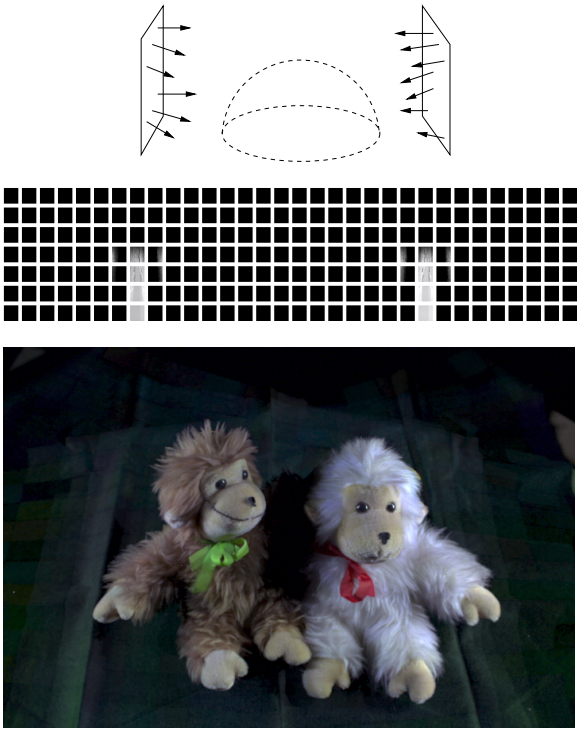

Figure 9: The environment consists of two white light sources facing each other. The light field is captured in the middle of the two light sources. In the middle, our representation of the incident light field is displayed. At the bottom the result of relighting the two monkeys with this light field is shown.

If the monkeys are shifted to the left in the environment, the left light source would illuminate the monkeys more than the right light source. This effect can be seen in figure 11a. In figure 11b the object is moved forward in the environment, resulting in illumination coming mainly from the back. In both cases the light field is shown.

More results on another object can be seen in figure 12. Here, a toy dog is placed in the same environment as the monkeys. Different concentrations of light can be seen on the sleigh, resulting from moving the object near one of the light sources, especially note the shadow on the sleigh (right image).

In the previous examples, the effects of spatial variation in the incident illumination are very subtle. In figure 13, more imaginative examples are given. The incident illumination consists of spot lights aimed at the chess pieces in the scene, and thus there is a high-frequency spatial variation. Therefore, spatial variation in the data aquisition process was raised to 32×32. The spot lights result in clearly visible spatial variation in the illumination. Due to this high-frequency spatial variation in the illumination, aliasing artefacts are visible. The emitted patterns can be recognized on the ground plate, as can be clearly seen in the zoom-ins on the right of figure 13.

Using our data set, we can also relight the object with a light map. A comparison is made in figure 10 between our technique and a relighting technique using a light map like in [Debevec et al. 2000]. The environment consists of three shafts of light. With our technique, the shafts are clearly noticeable. Using only a light map of this environment produces the lower relit image in figure 10. These results have no spatial variation in the illumination.

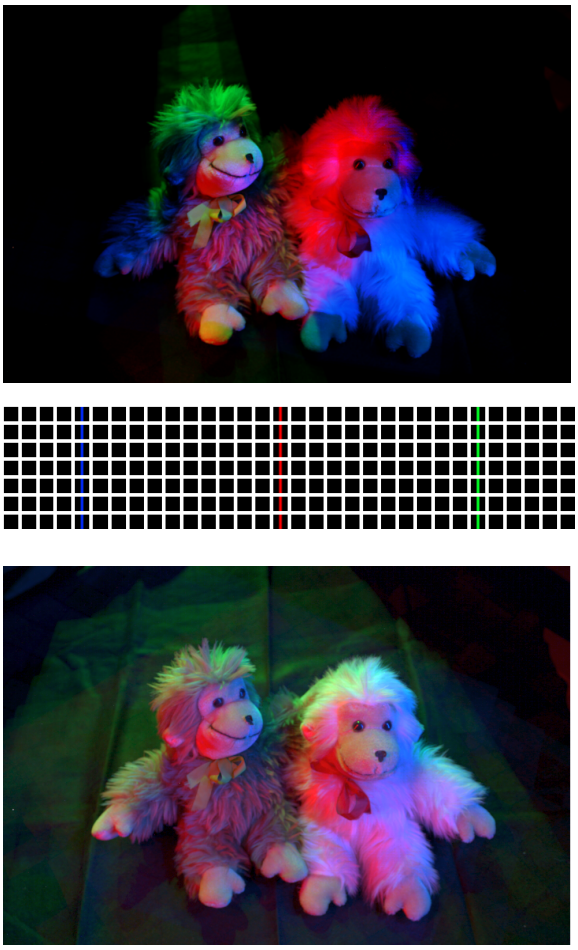

Figure 10: An artificial environment in which three shafts of light are cast onto the object. On top, we see the monkeys relit with our technique. The shafts are clearly noticeable and the colors are merged correctly. Notice also the complex scattering of the light in the fur of the monkeys. The light field is displayed in the middle. At the bottom, the result of relighting the object with a corresponding light map of this environment. Using the light map, no spatial variation in the illumination can be visualized in the relit image.

7 Conclusion and future work

We presented a technique to capture a 6D slice of the full reflectance field, enabling us to relight objects viewed from a fixed point with illumination varying spatially as well as angularly.

Real objects can be merged into a virtual scene using a differential technique similar to the method proposed by [Debevec 1998]. Using our relighting technique, the real object could be relit with an incident light map of that virtual environment and objects near the position where the real object is to be merged can result in local reflection on the real object.

Due to the low resolution of the captured reflection field and the use of box patterns, some aliasing effects are visible in the resulting relit images, as can be clearly seen in figure 13. This could be easily accommodated by capturing the reflectance field at a much finer

resolution. However, this would further increase the already huge amount of data and acquisition time. To remove the aliasing effects, Gaussian patterns could be used instead of box patterns, resulting in smoother images but at the cost of blurring out the incident light field. In future research, the use of other patterns, such as Gaussian or gray code inspired patterns will be investigated to reduce these aliasing effects and the amount of data to be captured.

We restricted our examples to artificial environments to relight our objects since real environments are difficult to capture. We expect that 4D incident light fields captured in a real environment could be used as well with our technique. We can acquire a 4D incident light field of a real environment by taking photographs with a fish-eye lense of the upper hemisphere in an environment at known positions. Thus taking 2D samples of the 4D light field. This technique would be very similar to the one proposed by [Aliaga and Carlbom 2001]. After resampling this acquired light field, we can create an image of our object illuminated using this incident light field.

Our technique captures a 6D slice of the 8D reflectance field. If multiple camera views are used, the full reflectance field could be captured. Because of the huge amount of data, sparser data sets to represent the reflectance field will need to be investigated, before the 8D reflectance field can be captured at decent resolutions.

8 Acknowlegdements

We would like to thank the reviewers for the useful comments and hints to improve the paper. Many thanks as well to all the people in our research group, especially Frank Suykens for his countless clever suggestions, Bart Adams and Karl vom Berge for the proofreading and Ares Lagae for helping with the movie-clips. The first author was supported by FWO Grant #G.0354.00, the second author was supported by K.U.Leuven Grant OT/01/34.

References

ADELSON, E. H., AND BERGEN, J. R. 1991. The plenoptic function and the elements of early vision. In *Computational models of visual processing*, Landy and Movshon, Eds., MIT press.

ALIAGA, D. G., AND CARLBOM, I. 2001. Plenoptic stitching: A scalable method for reconstructing 3D interactive walkthroughs. In *SIGGRAPH 2001, Computer Graphics Proceedings*, ACM Press / ACM SIGGRAPH, E. Fiume, Ed., Annual Conference Series, SIGGRAPH, 443–450.

CHEN, W.-C., BOUGUET, J.-Y., CHU, M. H., AND GRZESZCZUK, R. 2002. Light field mapping: Efficient representation and hardware rendering of surface light fields. In *SIGGRAPH 2002 Conference Proceedings*, ACM Press / ACM SIGGRAPH, J. Hughes, Ed., Annual Conference Series, SIGGRAPH, 447–456.

DEBEVEC, P., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97 Conference Graphics Proceedings*, Addison Wesley, T. Whitted, Ed., Annual Conference Series, ACM SIGGRAPH, 369–378.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *SIGGRAPH 2000, Computer Graphics Proceedings*, Addison Wesley, K. Akeley, Ed., Annual Conference Series, ACM SIGGRAPH, 145–156.

DEBEVEC, P., WENGER, A., TCHOU, C., GARDNER, A., WAESE, J., AND HAWKINS, T. 2002. A lighting reproduction approach to live-action compositing. In *SIGGRAPH 2002 Conference Proceedings*, ACM Press/ACM SIGGRAPH, J. Hughes, Ed., Annual Conference Series, SIGGRAPH, 547–556.

DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH 98 Conference Proceedings*, Addison Wesley, M. Cohen, Ed., Annual Conference Series, ACM SIGGRAPH, 189–198. ISBN 0-89791-999-8.

GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *SIGGRAPH 96, Computer Graphics Proceedings*, Addison Wesley, H. Rushmeier, Ed., Annual Conference Series, ACM SIGGRAPH, 43–54.

HAWKINS, T., COHEN, J., AND DEBEVEC, P. 2001. A photometric approach to digitizing cultural artifacts. In *In 2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage, Glyfada, Greece, November 2001*.

IHM, I., PARK, S., AND LEE, R. K. 1997. Rendering of spherical light fields. In *5th Pacific Conference on Computer Graphics and Applications (Pacific Graphics '97)*, Y. G. Shin and J. K. Hahn, Eds., Annual Conference Series, IEEE Computer Society, 59–69.

J NIMEROFF, E. S., AND DORSEY, J. 1994. Efficient re-rendering of naturally illuminated environments. In *Eurographics Rendering Workshop 1994*, Springer-Verlag, Darmstadt, Germany, EG.

LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *SIGGRAPH 96 Conference Proceedings*, Addison Wesley, H. Rushmeier, Ed., Annual Conference Series, ACM SIGGRAPH, 31–42.

LIN, Z., WONG, T.-T., AND SHUM, H.-Y. 2001. Relighting with the reflected irradiance field: Representation, sampling and reconstruction. *IEEE Computer Vision and Pattern Recognition 1* (Dec.), 561–567.

MASSELUS, V., DUTRÉ, P., AND ANRYS, F. 2002. The free-form light stage. In *Rendering Techniques EG 2002*, Annual Conference Series, EG, 247–255.

MATUSIK, W., PFISTER, H., NGAN, A., BEARDSLEY, P., ZIEGLER, R., AND MCMILLAN, L. 2002. Image-based 3D photography using opacity hulls. In *SIGGRAPH 2002 Conference Proceedings*, ACM Press/ACM SIGGRAPH, J. Hughes, Ed., Annual Conference Series, SIGGRAPH, 427–437.

MATUSIK, W., PFISTER, H., ZIEGLER, R., NGAN, A., AND MCMILLAN, L. 2002. Acquisition and rendering of transparant and refractive objects. In *Rendering Techniques EG 2002*, Annual Conference Series, EG, 267–277.

MELISSA L KOUDELKA, PETER N BELHUMEUR, S. M., AND KRIEGMAN, D. J. 2001. Image-based modeling and rendering of surfaces with arbitrary brdfs. In *Computer Vision and Pattern Recognition (CVPR'01), Computer Graphics Proceedings*, Annual Conference Series.

WONG, T.-T., HENG, P.-A., OR, S.-H., AND NG, W.-Y. 1997. Image-based rendering with controllable illumination. In *Eurographics Rendering Workshop 1997*, Springer Wien, J. Dorsey and P. Slusallek, Eds., Eurographics, 13–22.

WONG, T.-T., HENG, P.-A., AND FU, C.-W. 2001. Interactive relighting of panoramas. *IEEE Computer Graphics and Applications 21*, 2 (Mar./Apr.), 32–41.

WOOD, D. N., AZUMA, D. I., ALDINGER, K., CURLESS, B., DUCHAMP, T., SALESIN, D. H., AND STUETZLE, W. 2000. Surface light fields for 3D photography. In *SIGGRAPH 2000, Computer Graphics Proceedings*, ACM Press / ACM SIGGRAPH, K. Akeley, Ed., Annual Conference Series, SIGGRAPH, 287–296.

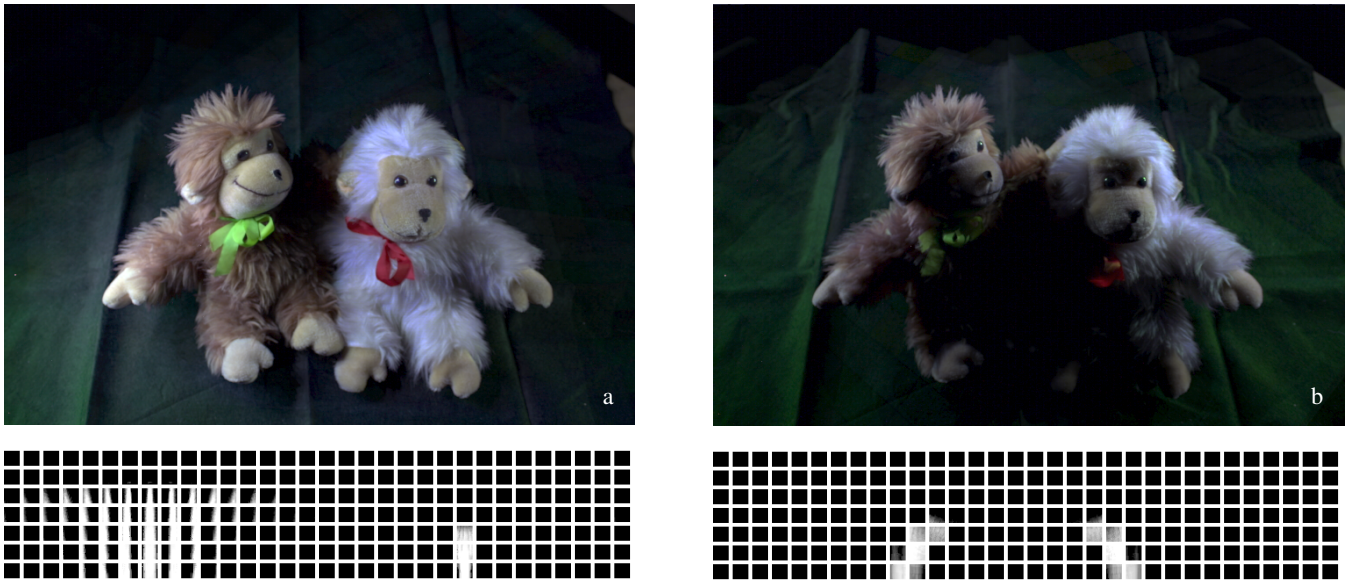

Figure 11: Two toy monkeys relit in an environment with two light sources from figure 9. In image (a), the monkeys are shifted leftwards relative to the incident light field, so the brown monkey is brighter, since it is closer to a light source. In image (b), the monkeys are shifted forward relative to the incident light field, so the illumination comes mainly from the back.

Figure 12: A toy dog on a wooden sleigh relit with the same environment used to relight the objects in figure 9. In image (a), the sleigh is shifted leftwards relative to the incident light field, in (b) the object is in the middle and in (c), it is shifted to the right.

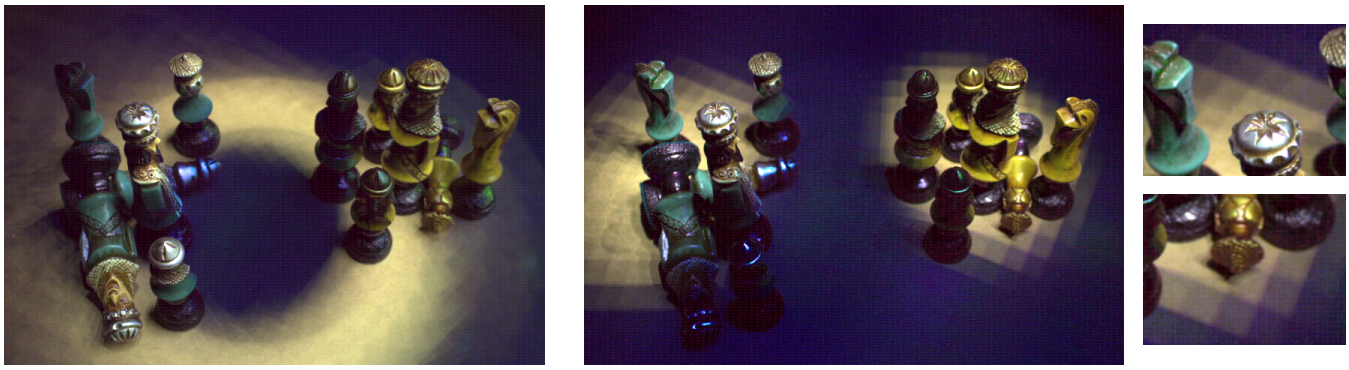

Figure 13: A configuration of chess pieces relit with our algorithm. For this example 32×32 patterns were used to capture the angular variance in the reflectance field. On the left we see the result of a set of spot lights configured in a circle. In the middle, only two spot lights are used, aimed at the two groups of chess pieces. The two zoom-ins on the right show aliasing artefacts.