

Smooth Reconstruction and Compact Representation of Reflectance Functions for Image-based Relighting

Vincent Masselus Pieter Peers Philip Dutré Yves D. Willems[†]

Department of Computer Science
Katholieke Universiteit Leuven

Abstract

In this paper we present a new method to reconstruct reflectance functions for image-based relighting. A reflectance function describes how a pixel in a photograph is observed depending on the incident illumination on the depicted object. Additionally we present a compact representation of the reconstructed reflectance functions. The reflectance functions are sampled from real objects by illuminating the object from a set of directions while recording photographs. Each pixel in a photograph is a sample of the reflectance function.

Next, a smooth continuous function is reconstructed, using different reconstruction techniques, from the sampled reflectance function. The presented method maintains important high frequency features such as highlights and self-shadowing and ensures visually pleasing relit images, computed with incident illumination containing high and low frequency features.

The reconstructed reflectance functions and incident illumination can be expressed by a common set of basis functions, enabling a significant speed-up of the relighting process. We use a non-linear approximation of higher order wavelets to preserve the smoothness of the reconstructed signal while maintaining good relit image quality. Our method improves on visual quality in comparison with previous image-based relighting methods, especially when animated incident illumination is used.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three dimensional Graphics and Realism I.4.1 [Image processing and Computer Vision]: Digitization and Image Capture

1. Introduction

Image-based relighting is the visualization of real objects with novel incident illumination. This illumination can be captured from the real world or from a virtual environment.

Over the recent years, several image-based relighting methods have been developed, covering a wide range of applications such as special effects in movies and visualizations for archaeology and forensics. Another important application is illumination dependent textures, used in the gaming industry where real-time relighting is required.

Existing image-based relighting techniques produce visually pleasing results for static illumination but suffer often

from poor visual performance on rendering animations with changing incident illumination, or restrict the reflectance properties of the object, or the incident illumination, to contain only low frequencies. The subject of this paper is to overcome these problems.

We conducted a thorough analysis on existing signal reconstruction techniques applied to reflectance functions for image-based relighting. A reflectance function describes how a pixel in a photograph is observed depending on the incident illumination on the depicted object. From our analysis, we conclude that a multilevel B-Spline technique performs best, in terms of visual quality and the ability to represent all features in reflectance functions.

To capture the sampled reflectance functions, the viewpoint is fixed relative to the object and a set of photographs is recorded while the object is illuminated by a light source

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

positioned at a different direction for each photograph. This approach is similar to previous image-based relighting techniques.

The reconstructed reflectance functions, using the multi-level B-Spline technique, can be used to compute relit images of the objects, lit with all-frequency illumination. Furthermore, changing the incident illumination results in a set of relit images, visually coherent over time.

To enable fast relighting and minimize storage requirements, the reflectance functions and the illumination maps can be approximated by a common set of basis functions. The computation for each relit pixel is reduced to multiplying corresponding coefficients of the reflectance function and the incident illumination. We propose to use a non-linear approximation of the reflectance functions using higher order wavelets for this set of common basis functions. This approximation introduces an error of less than 1% during relighting, and allows to compute a relit image with a low average number of coefficients per reflectance function.

2. Previous Work

Relighting has been widely researched over the recent years. We will restrict the following survey of previous work to recent image-based relighting techniques which do not make any assumption involving geometry or reflectance properties and sample the reflectance functions directly by means of point samples.

Debevec et al. [DHT*00] introduced a gantry, called the *Light Stage*, which enables to illuminate a real object from a regular set of directions. The resulting basis images are linearly combined to produce an image of the object lit by an arbitrary light map. Subsequent versions of the Light Stage improve on acquisition time [HCD01] or generalize on a non-regular sampling of directions [MDA02]. Matusik et al. [MPN*02, MPZ*02] extended the Light Stage for viewpoint independent relighting.

Wong et al. [WHON97] extended existing image-based rendering techniques [GGSC96, LH96] with controllable illumination. For a single viewpoint, a set of images of the object, illuminated from different directions is rendered using a global illumination renderer. For each pixel the apparent BRDF at that pixel is stored and compressed using spherical harmonics. To visualize the relit object, the apparent BRDF of each pixel is evaluated for a specific illumination direction. The concept of apparent BRDFs was also used to visualize panoramas with variable illumination [WHF01].

Lin et al. [LWS01] also extended the Lumigraph [GGSC96] and Light Field Rendering [LH96] methods, but to a tri-planar technique. A point light source is positioned on a grid and for each position a light slab is recorded. Afterwards, the object can be rendered from any position illuminated with a point light source positioned on the original grid.

Polynomial Texture Mapping (PTM) [MGW01] is a texturing technique that enables hardware rendering of illumination dependent textures. A PTM is synthesized by recording a series of photographs with the view direction orthogonal to the surface plane and illuminated from different directions. For each texel, the set of the reflected intensities of the different light sources are fitted to a biquadric polynomial.

The following section (3) overviews our method to obtain point samples of the reflectance functions of the pixels. The paper then surveys our two main contributions: *reconstructing smooth reflectance functions* from the captured data (section 4) and a *compact and efficient representation* for these reconstructed reflectance functions (section 5). Finally, we conclude the paper and outline some future work.

3. Sampling Reflectance Functions

3.1. Relighting

Let Ω be the space of all light directions over a hemisphere centered around the object to be illuminated. Angular dependent incident illumination can be expressed as a function $\mathbf{L}_{\text{in}}(\omega)$ of radiance, with $\omega \in \Omega$. The viewpoint to the object is fixed, since we are only interested in the effect of changing the incident illumination on the object, not changing the viewpoint. The radiant illumination through a pixel (x, y) under incident illumination \mathbf{L}_{in} can be expressed as:

$$\mathbf{L}_{\text{out}}(x, y) = \int_{\Omega} \mathbf{R}(\omega, x, y) \mathbf{L}_{\text{in}}(\omega) d\omega, \quad (1)$$

with $\mathbf{R}(\omega, x, y)$ representing the amount of radiance from direction ω that is reflected into pixel (x, y) (Figure 1). This function \mathbf{R} , introduced by Debevec et al. [DHT*00], is called the *reflectance field*.




Figure 1: An object (shown in red) is illuminated by \mathbf{L}_{in} . The reflected radiance towards pixel (x, y) is \mathbf{L}_{out} .

Once the reflectance field $\mathbf{R}(\omega, x, y)$ is acquired, the object can be visualized by evaluating equation 1 for all pixels with any incident illumination \mathbf{L}_{in} .

For a specific pixel, the reflectance field $\mathbf{R}(\omega, x, y)$ is denoted by $\mathbf{R}(\omega)$ and is called the reflectance function of that pixel. The image of the relit object can be calculated by taking the inner product of the incident illumination, and the reflectance function of each pixel.

3.2. Data Acquisition

The reflectance functions $\mathbf{R}(\omega)$ need to be captured, in order to relight an object. A reflectance function is a discontinuous function over Ω , due to self-shadowing boundaries. This makes the exact measurement of $\mathbf{R}(\omega)$ difficult. However, the object can be illuminated from a discrete set of directions ω_i , while a High Dynamic Range (HDR) image is recorded for each direction. Each recorded photograph results in one sample value of the reflectance function for each pixel. We denote a sample value resulting from positioning the light source at direction ω_i as s_{ω_i} . Note that s_{ω_i} can include measurement errors.

Using this set of samples, a discrete-to-continuous reconstruction can be performed to obtain the reflectance function $\mathbf{R}(\omega)$ for each pixel separately.



Figure 2: The setup for our data acquisition. The scene, a set of stones, and the camera are placed on a turntable. A semi-circular brace with 40 light sources is mounted over the turntable. By rotating the turntable 180 degrees in 32 steps and switching on one light source at a time, the object can be illuminated from 64×20 regularly sampled directions.

The data acquisition setup is similar to the one described in [HCD01]. The object and camera are placed on a turntable. A semi-circular brace with 40 light sources is mounted over the turntable. The tilt angle of the light direction is defined by the light source used, the azimuth angle can be changed by rotating the turntable. Using 40 light sources and rotating the turntable 180 degrees in 32 steps, allows to illuminate the object from 64×20 (= 1280) regular sampling directions and record a HDR image for each direction. Note, that in this setup the camera can block a light source. An image of the setup can be seen in figure 2.

4. Reconstruction Methods

Throughout the paper, we use a miniature race car to illustrate the discussed methods. This object was especially chosen because it contains many specular surfaces and fine geometrical details, causing numerous self-occlusion features




Figure 3: A miniature race car featuring diffuse, glossy and specular material properties along with fine detail in geometry. The reflectance functions of this object are sampled at 1280 illumination directions and reconstructed using the zero-order hold technique. Reflectance functions of the selected pixels are shown in the top left and bottom right.

(figure 3). The reflectance functions of this object were sampled from 1280 directions. Four pixels were selected and the accompanying sampled reflectance functions are depicted as well. To visualize a reflectance function defined on a hemisphere, we project it onto a disk.

These four reflectance functions include a wide range of features: a diffuse red pixel on the side of the helmet of the driver, as can be seen in figure 3.1. A yellow pixel on the top of the helmet features a soft highlight (figure 3.2). We also selected two pixels containing high frequency features: a pixel on the side of the car with a sharp highlight and a pixel on the rear of the car complexly occluded by the rear wing (figure 3.3 and 3.4 respectively). At the bottom of the reflectance functions, a dark spot can be seen. This is the camera occluding the light sources.

Several techniques were analyzed to create a continuous function on a hemisphere from a discrete set of samples. In the following sections we review these reconstruction techniques. Previously published techniques comprise *zero-order hold*, *fitting to a biquadric polynomial* and using *spherical harmonics*. Additionally, *interpolation of the sampled data*, using the *wavelet transform* and using *B-Splines* are presented in this paper as techniques to reconstruct reflectance functions. An estimate of the errors of each method are discussed and compared in section 4.1. Some results are discussed in section 4.2.

Zero-order Hold

The set of illumination directions used during data acquisition are plotted on a hemisphere and an angular Voronoi diagram is created using the angles between different directions as a distance measure. A piecewise constant reflectance function for each pixel can be found by assigning the sampled reflectance value of each direction to the corresponding Voronoi cell. Similar methods were used by [DHT*00] for regularly sampled directions and by [MDA02] for irregularly




Figure 4: Reflectance functions of the pixels selected in figure 3, reconstructed using different techniques.

sampled directions. In figure 4.A, the reflectance functions are reconstructed using all 1280 samples and in figure 4.B using a subset of 320 regular samples.

Interpolation

By interpolating the recorded reflectance values, a C^0 continuous reflectance function can be created.

A Delaunay triangulation is constructed on the hemisphere using the sampled directions. Each point in a triangle is interpolated using spherical barycentric coordinates. The resulting reflectance functions are depicted in figure 4.C. They are visually smoother than the results computed by the zero-order hold technique.

Inverse Distance Weighted interpolation (IDW), also called Shepard’s method uses all samples to reconstruct the function in a single direction. Again, the angle between directions was used as a distance measure and the weights of the values of the sampled directions are inverse proportional to the distance of the sampled directions. A variation on this technique calculates interpolated values by only taking into account a pre-defined number of nearest neighbor directions. In figure 4.D a result is shown using only the 32 nearest neighbors.

Fitting to a Biquadric Polynomial

The captured data can also be fitted to a single biquadric polynomial. The resulting reflectance functions, are C^∞ continuous and defined on a global support. This reconstruction method significantly blurs out the high frequency features such as highlights and self-shadowing boundaries. Results of the selected pixels can be seen in figure 4.E. Using a biquadric polynomial can result in negative values for some directions, which are clamped to zero in the visualization.

The method is similar to Polynomial Texture Maps [MGW01]. Although this technique was intended for illumination dependent texture maps, it can be applied to relight objects as well. A major advantage is that only 6 coefficients have to be stored per texel.

Spherical Harmonics

The sampled data can be fitted through a set of spherical harmonic basis functions. This results in a continuous approximation of the reflectance function. However, spherical harmonics are defined on a sphere, posing a problem when using only a hemisphere. We solved this by mirroring the data to the other hemisphere.

Wong et al. [WHON97] applied this technique on sampled reflectance functions of synthetic scenes and used 16 or 25 coefficients, corresponding with four or five bands of the spherical harmonic basis functions respectively. The reconstructed reflectance functions using the first five bands, are very smooth but do not capture all the details, similarly as using a biquadric polynomial (see figure 4.F).

In figure 4.G and 4.H, the reflectance functions are depicted using 225 and 625 spherical harmonic coefficients, for each color channel. The reconstructed reflectance functions are able to capture all features to some extent, but suffer from severe Gibbs ringing or aliasing artefacts. Negative values in the visualizations in figure 4 are clamped to zero.

Wavelet Transform

Wavelets are a flexible tool used in many domains including computer graphics. A good survey can be found in [SDS96, SSC*96].

A sampled signal can be reconstructed to a continuous signal using wavelets, by repeatedly inserting additional samples halfway between each sample point. The magnitude of each new sample point is characterized by the scale function of the wavelet. Intuitively this can be seen as composing a signal using the wavelet transform in which the added wavelet coefficients (high frequency features) are set to zero.

A logical choice would be to use spherical wavelets [SS95]. The highly irregularly spaced samples in the spherical domain are a disadvantage and require a resampling step. The recorded samples, however, are regularly spaced in the latitude-longitude parameterization, due to the setup. We therefore opt for using the wavelet transform in the latitude-longitude parameterization. An additional advantage is that common wavelet implementations can be used without much trouble.

We use two different types of wavelets. The first wavelet is the 5/3 LeGall, or Integer 5/3, which is the shortest symmetrical biorthogonal wavelet with two vanishing moments. Its scaling function is a linear B-Spline. The second wavelet is the well known 9/7 Daubechies wavelet, which is the shortest symmetrical biorthogonal wavelet of order four, and is by construction a cubic B-Spline. Both wavelets are part of the JPEG2000 standard, and are therefore widely implemented. Traditionally these wavelet are mirrored around images boundaries. However, in our implementation we opted for repeating the signal on vertical boundaries, since this fits better to the original spherical domain. We still mirror the wavelets on horizontal boundaries.

Results of using these wavelets can be seen in figure 4.I and figure 4.J. The 5/3 LeGall wavelet gives similar results as linear interpolation. The 9/7 Daubechies wavelet, however, result in a very smooth reflection function, but suffers from aliasing artefacts, which become noticeable when animated incident illumination is used.

B-Splines

Bicubic B-Splines can be used to create a continuous function. These functions are a good trade-off between smoothness (C^2 continuous) and the ability to represent the features in the captured data. The error analysis in section 4.1 will

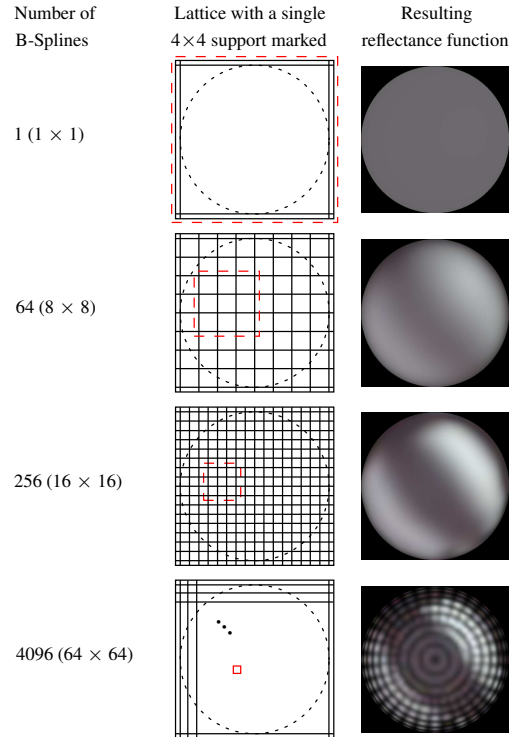


Figure 5: The influence of number of B-Splines on the coverage of a 4×4 support and the resulting reflectance function. The red dashed rectangle shows the size of a 4×4 support for each lattice.

show that reconstruction reflectance functions using multi-level B-Splines performs best with respect to the other reconstruction techniques. Therefore, we will discuss this method in more detail.

B-Splines are not easily defined on a sphere. Therefore we opt for representing the sampled data using the paraboloid map parameterization [HS99]. This representation offers a continuous projection of the hemisphere, in which the boundaries are identical as on a hemisphere and the pole is defined in a single point. Furthermore, The solid angle to projected area ratio is close to constant.

A bicubic B-Spline can be fitted on the data by creating a 4×4 grid of control points on the projected data, as can be seen in the first row of figure 5. The resulting reflectance function is poorly reconstructed because a single bicubic B-Spline cannot represent all features, similarly as using a single biquadratic polynomial (section 4). A possible solution reconstructs the reflectance function using a set of independent B-Splines, each defined on a different 4×4 grid of control points, which are defined on a lattice over the projected data. Figure 5 demonstrates the influence of the resolution of the lattice and the resulting reflectance function. Using few B-

Splines results in a good global fit of the data, but with almost no local detail (figure 5, first three rows) while using more B-Splines produces a reflectance function with a good local fit but lacking global smoothness (figure 5, last row). This problem was also noted by [LWS97].




Figure 6: The construction of the multilevel B-Splines. For each level, the 4×4 grid of control points is subdivided on which a set of B-Splines is defined. These B-Splines are fitted through the difference of the already reconstructed set of B-Splines and the measured values s_{ω_i} . The difference is displayed dark for negative and bright for positive values.

Multilevel B-Splines, presented by [LWS97], allows to fit a smooth approximation through the projected samples without these problems. Multilevel B-Spline interpolation is a hierarchical method that first tries to fit a set of globally smooth B-Splines through the sampled data each defined on a distinct 4×4 grid of control points with large coverage. In each successive step the number of control point in the grid of each B-Spline is doubled in each direction and a new set of B-Splines is created on the four smaller 4×4 grids of control points. These new B-Splines on the smaller grids are fitted through the difference of the sum of the already computed B-Splines and the measured sample values s_{ω_i} , as defined in section 3.1. The method is demonstrated in fig-

ure 6. The hierarchy of B-Spline sets can then be reduced to a single set of B-Splines defined on the 4×4 grids of control points with the smallest coverage used (i.e. highest hierarchy level).

In our implementation we used four starting grids, which produce a smooth global fit of the measured data and we refine to six levels in the hierarchy, which allows for a good local fit. The number of levels in the hierarchy was empirically determined by the magnitude of differences of the measured values and the values of the sampled directions in the already constructed reflectance function. Using six levels allows to fit the data values, while not fitting noise on the data. This method results in a set of 64×64 B-Splines. The reconstructed B-Spline reflectance functions can be seen in figure 4.K.

4.1. Comparison

Suppose we reconstructed the reflectance function $\mathbf{R}_N(\omega)$ using N sampling directions. Reconstructing a reflectance function by a discrete-to-continuous technique introduces an approximation error defined as:

$$\mathbf{E}_{\text{avg}}^2 = \frac{\int_{\Omega} (\mathbf{R}(\omega) - \mathbf{R}_N(\omega))^2 d\omega}{\int_{\Omega} (\mathbf{R}(\omega))^2 d\omega}, \quad (2)$$

with \mathbf{E}_{avg} the average error on the reconstructed reflectance function. \mathbf{E}_{avg} cannot be exactly computed, since $\mathbf{R}(\omega)$ is unknown. However this error can be approximated by sampling an additional large set of M ($\gg N$) directions $\{\omega_1, \omega_2, \dots, \omega_M\}$. Using the M captured samples, denoted as s_{ω_i} , $\mathbf{E}_{\text{avg}}^2$ can be approximated as:

$$\mathbf{E}_{\text{avg}}^2 \approx \frac{\sum_{i=1}^M (s_{\omega_i} - \mathbf{R}_N(\omega_i))^2}{\sum_{i=1}^M s_{\omega_i}^2}. \quad (3)$$

We illuminated the objects from 1280 ($= M$) regularly sampled directions and take a subset of 320 ($= N$) samples to estimate the error on each reconstructed reflectance function.

\mathbf{E}_{avg} will mainly depend on the used reconstruction technique and the kind of reflectance function. Most of the reconstruction techniques will perform well if a pixel represents a diffuse un-occluded surface, thus a low frequency function, and consequently result in a low value for \mathbf{E}_{avg} . However, the error can be high for some reconstruction methods when the reflectance function features a highlight or complex self-shadowing.

Instead of comparing \mathbf{E}_{avg} for a single reflectance function, we compute the average of \mathbf{E}_{avg} over a set of pixels with similar occlusion and material properties. We compare these errors for different reconstruction techniques.

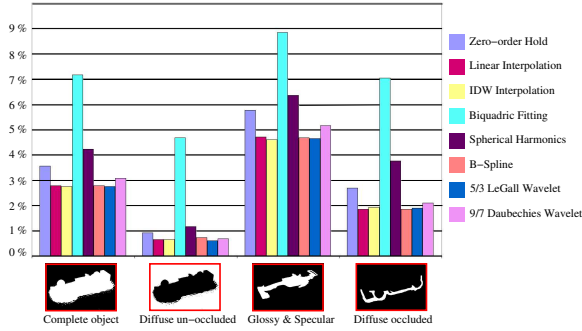


Figure 7: The average relative error of the reconstructed reflectance functions for the object from figure 3 for four different sets of pixels using different reconstruction techniques.

We use four sets of pixels: a set of pixels representing the complete object, pixels representing diffuse un-occluded surfaces, pixels representing diffuse surfaces with complex self-shadowing and pixels located on glossy and specular materials.

The average relative errors E_{avg} of these four sets are plotted in figure 7. The four sets are visualized as a false color images in which the accounted pixels are shown in white.

In general, the error for un-occluded diffuse reflectance functions is low, while reflectance functions featuring high frequency details result in larger errors. Ramamoorthi and Hanrahan [RH01] noted that diffuse unblocked reflectance functions can be represented by the first three spherical harmonics bands, yielding an error of less than 1%. In our experiments, we come to a similar conclusion, having noisy real data (including camera occlusion) and using 5 bands.

Reconstructing the reflectance functions by fitting the data to a biquadric polynomial or using spherical harmonic basis functions results in large errors. This is due to the loss of high frequency features and the use of a set of globally smooth functions. Using the zero-order hold reconstruction technique produces a mediocre result. It preserves most features, but also introduces new high frequency features in the reflectance functions. Our proposed methods have relatively low errors, because the resulting reconstructed reflectance functions are smooth while maintaining most of the features.

The standard deviation on the calculated errors was small for pixels representing diffuse or occluded surfaces. Pixels representing a specular surface produced a large standard deviation which was expected due to the high frequency features in the reflectance functions.

Although these errors give a good indication of the performance of the introduced reconstruction methods, it is still the effect of incident illumination on the object which is most important, not the reflectance function itself. In figure 8, the miniature race car is illuminated by a vertical sliver of light. Using a line of illumination results in images containing

both low frequency illumination features such as long soft shadows (along the direction of the line of illumination) and high frequency illumination features such as very short soft shadows (orthogonal to the line). The reflectance functions were reconstructed using several reconstruction techniques and using 1280 sampled directions.



Figure 8: The miniature race car from figure 3, relit with a vertical sliver of light, using different reconstruction techniques for the reflectance functions.

Using a zero-order hold reconstruction technique results in a pleasing image, however, the shadows and highlights move in a jittered way when the illumination is changed, as can be seen in figure 12 and the accompanying video. Fitting the data with a biquadric polynomial or using 25 spherical harmonics coefficients does not result in a satisfying image. The shadows, especially close to the object, are distorted and the highlights are completely lost. These aliasing effects are due to Gibbs ringing. However, these methods require very few coefficients to represent the data which is of importance for real time relighting. The linear interpolation, B-Spline and wavelet reconstruction techniques deliver a good visual result (IDW interpolation performed almost identical to linear interpolation). The shadows are faithfully recreated, as are the highlights. Multilevel B-Spline reconstruction outperforms linear interpolation and wavelet interpolation when considering dynamic incident illumination (figure 12 and the video).

Since the reconstruction of the reflectance functions is a pre-process to the actual relighting and needs to be done only once, we did not include the pre-processing time of the different reconstruction techniques in the comparison.

Also note that all reconstruction techniques except for the wavelet interpolation can also be applied on irregular sampled measured data.

4.2. Reconstruction Results

The top left of figure 11 displays a set of stones, a scene featuring a lot of self-shadowing effects. This can be noted in the reflectance functions depicted around the figure. On the bottom left of figure 11, a set of coins, together with some selected reflectance functions are shown. This scene was chosen for the highly specular material properties. As can be seen, the highlights have all kinds of shapes which are preserved in the reconstructed reflectance functions. These two scenes are also included on the video.

The right of figure 11 shows a jade statuette. A series of selected reflectance functions, reconstructed using the multilevel B-Spline technique with all 1280 captured samples, are depicted on both sides, containing a whole range of different reflectance features: specular highlights (figure 11.A), diffuse reflection (figure 11.B), subsurface scattering (figure 11.D) and complex occlusion (figure 11.C, G and H).

In these scenes, the multilevel B-Splines technique also performs best in terms of both error and reconstruction smoothness.

5. Fast Relighting

In this section, we assume that the reflectance functions are already reconstructed using one of the techniques presented in the previous section. The reconstructed reflectance functions are provided in an algebraic form. However evaluating the reflectance functions using this algebraic form may require a significant amount of processing time. Additionally, computing an image of a relit object requires to evaluate:

$$\mathbf{L}_{\text{out}} = \int_{\Omega} \mathbf{R}(\omega) \mathbf{L}_{\text{in}}(\omega) d(\omega), \quad (4)$$

for each pixel (equation 1). The time required to compute this inner product for a single pixel is proportional to:

1. The time required to evaluate a reflectance function.
2. The number of samples in \mathbf{L}_{in} .

The computational effort is large for densely sampled incident illumination (such as high resolution environment maps) and can be impractical for real-time or interactive applications.

In the following subsections we speedup this computation by reducing the weight of both factors. *The time required to evaluate a reflectance function* can be reduced by expressing the algebraic form of the reflectance functions and the incident illumination in a common set of basis functions (section 5.1). Further speedup can be attained by a *lossy approximation of \mathbf{L}_{in} and \mathbf{R}* using this same common set of basis functions (section 5.2) and also reducing storage requirement.

5.1. Common Basis Approximation

The functions $\mathbf{L}_{\text{in}}(\omega)$ and $\mathbf{R}(\omega)$ can be expressed using a common set of basis functions $\mathbf{B}_i(\omega)$ and a dual set $\bar{\mathbf{B}}_i(\omega)$ on Ω :

$$\begin{aligned} \mathbf{L}_{\text{in}}(\omega) &= \sum_i \bar{l}_i \mathbf{B}_i(\omega) \\ \mathbf{R}(\omega) &= \sum_i r_i \bar{\mathbf{B}}_i(\omega), \end{aligned}$$

where \bar{l}_i and r_i are the projections of the dual basis functions $\bar{\mathbf{B}}_i(\omega)$ onto $\mathbf{L}_{\text{in}}(\omega)$ and $\mathbf{B}_i(\omega)$ onto $\mathbf{R}(\omega)$ respectively.

Equation 4 can be rewritten using these approximations as:

$$\begin{aligned} \mathbf{L}_{\text{out}} &= \int_{\Omega} \mathbf{R}(\omega) \mathbf{L}_{\text{in}}(\omega) d\omega \\ &= \int_{\Omega} \left(\sum_i r_i \bar{\mathbf{B}}_i(\omega) \right) \left(\sum_j \bar{l}_j \mathbf{B}_j(\omega) \right) d\omega \\ &= \sum_i \sum_j r_i \bar{l}_j \int_{\Omega} \bar{\mathbf{B}}_i(\omega) \mathbf{B}_j(\omega) d\omega. \end{aligned}$$

This can be further simplified using the definition of dual basis functions ($\int_{\Omega} \mathbf{B}_i(\omega) \bar{\mathbf{B}}_j(\omega) d\omega = \delta_{i,j}$) to:

$$\mathbf{L}_{\text{out}} = \sum_i r_i \bar{l}_i. \quad (5)$$

This reduces the inner product computation to multiplying the corresponding coefficients from the reflectance function and the incident illumination. The incident illumination needs to be projected only once on the dual basis functions and the resulting \bar{l}_i can then be re-used for each pixel. Note that both the incident illumination and the reflectance functions can be expressed in the same basis when \mathbf{B}_i is an orthonormal set of basis functions.

Incident illumination is usually represented as an environment map with a finite resolution of N samples. Therefore, \mathbf{L}_{in} can be represented by a finite number of N dual basis functions $\bar{\mathbf{B}}_i$ (i.e. N coefficients \bar{l}_i). From formula 5 follows that, if \mathbf{R} is expressed using the N corresponding basis functions \mathbf{B}_i , then \mathbf{L}_{out} can be exactly computed without introducing an approximation error.

In practice, a reflectance function is first expressed in a discrete representation. The resolution and parameterization of this discrete representation should be equal to the resolution and the parameterization of the incident illumination. This discrete representation can now be more easily expressed in a specific basis (e.g. spherical harmonics, wavelets, ...). The coefficients for each reflectance function can be pre-computed, reducing the *evaluation of the reflectance functions* to a simple lookup operation instead of evaluating the algebraic form.

5.2. Lossy Approximation

Expressing \mathbf{R} and \mathbf{L}_{in} in a common set of basis functions does not reduce the number of multiplications in the inner product computation, which is equal to the number of samples in \mathbf{L}_{in} . A solution is to select a common basis set in which many coefficients are near-zero. A lossy approximation is then applied by leaving out these near-zero coefficients, introducing a small error. Using a lossy approximation method makes it possible to reduce the number of coefficients dramatically without introducing too much error.

In the following sections we discuss spherical harmonics and wavelets as a candidate for such a common set of basis functions to be used with lossy approximation.

Spherical Harmonics

Spherical harmonics can be used to represent reflectance functions of a pixel in a compact way. The main idea is to use a number of low frequency bands to represent the data, leaving out high frequency “details”.

Low frequency reflectance functions can be well approximated using spherical harmonics [RH01], if only the lower bands are used. Using more bands will allow to add more detail in the resulting reflectance functions. A significant drawback of spherical harmonics is Gibbs ringing or aliasing which occurs around high frequency features, such as highlights and self-shadowing boundaries. The resulting reflectance functions are similar as when reconstructing the reflectance functions directly with spherical harmonics (figure 4.G and figure 4.H).

Kautz et al. [KSS02] used the first 15 bands to represent synthetic reflection properties without taking into account self-shadowing effects. Sloan et al. [SKS02] extended this technique to include self-shadowing for synthetic scenes by using two sets of spherical harmonics, one for the reflectance properties and one for the self-occlusion features. Using two sets is only possible if the geometry is known in advance.

Wavelet Compression

Wavelets are well known for their use in image compression and have also been used for representing and compressing incident illumination. Ng et al. [NRH03] used a cube-parameterization and a non-linear Haar wavelet approximation [DeV98] on the incident illumination in an image-based relighting context. Using a non-linear approximation of the incident illumination significantly speeds-up the computation of formula 5. However, it does not reduce the storage requirements for the reflectance field. A more logical choice would be to use non-linear wavelet approximation on the reflectance functions itself, reducing computational and storage requirements. Furthermore, we do not restrict ourselves to only the Haar wavelet. Higher order 5/3 LeGall and 9/7 Daubechies biorthogonal wavelets are also considered. These wavelets have better properties with respect to image compression [UB03] than the Haar wavelet.




Figure 9: Sobolev H^1 -norm versus the number of wavelet coefficients.

We performed a non-linear approximation in the same parameterization as was used for the wavelet reconstruction in section 4 and a resolution of 256×64 . An increase in resolution will not alter the compression ratio much, since the increase in information is limited. For all examples, we started from the reflectance functions reconstructed with the proposed multilevel B-Spline technique.

For error measurements we used the scene from figure 3. To include both spatial as well as temporal errors, we created an image in which each column represents the same column in a relit image of the miniature race car but at different moments during relighting with animated incident illumination (a sliver of light rotating around the object (figure 12)).

In figure 9 the relative Sobolev H^1 -error is plotted in function of the number of wavelet coefficients. The relative L^2 -error resulted in a similar graph and was therefore omitted. In general, the Haar wavelet needs more coefficients to achieve the same relative error. For example, 4096 Haar wavelet coefficients are needed for a $\pm 2\%$ error, as opposed to only 256 5/3 LeGall or 9/7 Daubechies wavelet coefficients. The 9/7 Daubechies wavelet slightly outperforms the 5/3 LeGall wavelet in terms of error: for 256 coefficients the Haar wavelet results in a 7.05% relative error, the 5/3 LeGall wavelet in a 2.12% relative error and the 9/7 Daubechies wavelet in a 1.88% relative error.

Using the same number of coefficients for each reflectance function in the reflectance field is not optimal. Some reflectance functions can be compressed using less coefficients and still result in an acceptable error during relighting. To determine the number of wavelet coefficients to be used for a specific reflectance function, our algorithm selects the largest wavelet coefficients such that the error on the L^2 -norm on the compressed reflectance function itself is bounded by an error-threshold. This L^2 -norm can be easily computed by decomposing the reflectance function into the primal and dual wavelet space and summing the multiplied corresponding coefficients. The L^2 -norm is sufficient for determining the number of wavelet coefficients, although it is only an indication of the error on the inner product of the reflectance function and the (at processing time) *unknown* incident illumination.

It is important that the non-linear approximation, for both fixed and variable number of coefficients, maintains as much as possible the smoothness of the original reflectance functions. Introducing additional discontinuities will result in visually disturbing features when animating the incident illumination. This can be seen in the last three rows in figure 12 and in the accompanying video, both relit from reflectance functions compressed with a variable number of coefficients (0.1% error-threshold). It is clear from this figure and the video that, although the Haar wavelet achieves very good compression ratios, it fails to maintain smoothness. The reason is the low number of vanishing moments in the Haar wavelet. The 5/3 LeGall and 9/7 Daubechies wavelet perform much better and still have very good compression ratios. See figure 10 for a comparison of the Sobolev H^1 -norm versus different error-threshold values and the respective number of wavelet coefficients. The correlation between the number of coefficients for a specific wavelet and the error threshold is hard to predict. When using a variable number of coefficients, more wavelet coefficients are assigned to high-detail reflectance functions, whereas low-detail functions are compressed using less wavelet coefficients. These details contribute little to the error, hence the small difference in error between compressing reflectance functions using a variable number of wavelet coefficients and using a fixed number of wavelet coefficients.

In terms of error and smoothness the 9/7 Daubechies wavelet is preferred (followed by the 5/3 LeGall wavelet), achieving a compression ratio of 1 : 34 for a 0.2% error-threshold and a H^1 -error of less than 1% on a 256×64 discretization of a multilevel B-Spline reconstruction. Non-linear approximation using the Haar wavelet will only be better suited when compressing reflectance functions that are non-smoothly reconstructed (e.g. zero-order hold technique).

6. Conclusion

In this paper, we propose a set of techniques to reconstruct continuous reflectance functions from a set of measured reflectance samples. Several existing methods to reconstruct a continuous reflectance function, such as zero-order hold, fitting the sampled data to a biquadric polynomial or to spherical harmonic basis functions have been compared to these new techniques.

Interpolation of the sampled data ensures that the reflectance function goes through the sampled values at the sampled directions. We analyzed linear interpolation, inverse distance weighted interpolation and wavelet interpolation. *Fitting the data samples to a set of bicubic B-Splines defined on a lattice* produces a C^2 continuous reflectance function which maintains all sampled features such as highlights and self-shadowing and allows to change the illumination while the resulting images form a coherent series of images over time.




Figure 10: Sobolev H^1 -norm versus the error-threshold on the inner product for non-linear approximation using a variable number of coefficients per reflectance function.

From our experiments, we found interpolation, multilevel B-Spline reconstruction and wavelets interpolation to perform better than existing methods. In case the incident illumination is animated, the resulting series of relit images is most coherent when using multilevel B-Spline reconstruction. However, multilevel B-Spline reconstruction may result in small ringing effects in case of very high frequency reflectance functions. As an alternative, linear interpolation techniques (linear interpolation and 5/3 LeGall wavelet interpolation) guarantee no ringing effects at the cost of smoothness.

In order to relight the object faster, a common basis approximation can be created, introducing an approximation error. This error can be reduced by selecting an appropriate set of basis functions and by increasing the number of basis functions. We found 9/7 Daubechies wavelets using an average of 481 coefficients reduces the approximation error to less than 1% percent. The 5/3 LeGall wavelet performs well, but does not reach the same compression ratios as the 9/7 Daubechies wavelet. The Haar wavelet introduced new temporal discontinuities and should therefore not be used to compress smooth reflectance functions.

As a conclusion, we propose to reconstruct a continuous reflectance function from a sampled reflectance function, using the multilevel B-Spline technique. This enables to preserve all features in the sampled reflectance function and allows coherent relighting when the incident illumination varies. Using non-linear approximation with 9/7 Daubechies wavelets, the reflectance functions can be stored with a low average number of coefficients. This enables fast relighting, without losing image quality in the relit images and minimizes storage requirements.

7. Future Work

The computational time needed to generate a relit image of an object can be reduced by selecting an appropriate set of common basis functions. At the same time, the data is also





Figure 11: On the left, a set of stones and an arrangement of coins. These scenes were chosen for the complex self-shadowing and high specular material properties, respectively. On the right, a small jade statuette, selected for the sub-surface scattering effects. All scenes were illuminated from 1280 regular directions. The reflectance functions were reconstructed using the multilevel B-Spline technique.

compressed, enabling us to cache all data simultaneously. In the near future, we will research how relighting can be performed on graphics hardware, enabling real-time high detail relighting.

Our proposed technique allows to reconstruct a continuous reflectance field from a set of samples and represent it in a compact way. Although we presented it for real objects, the method can also be applied to synthetic objects for representing radiance transfer.

In our work, we mathematically compared the different techniques by computing an error on the reconstructed reflectance functions and a visual comparison was conducted on the resulting relit images. However, the spatial and temporal aliasing of the series of relit images should be compared with a mathematical error as well. This is a direction for future research.

The wavelets used for interpolating are not necessarily the best choice. We intend to look at other wavelets specifically designed for interpolation. We also intend to investigate completely other reconstruction techniques such as the push-pull method, as described in [GGSC96].

8. Acknowledgments

We would like to thank Frank Suykens for proofreading the paper. This work was partially supported by FWO Grant #G.0354.00 and K.U.Leuven Grant #OT/01-34.

References

[DeV98] DEVORE R. A.: Nonlinear approximation. *Acta Numerica* 7 (1998), 51–150. 9

[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, Computer Graphics Proceedings* (July 2000), Addison Wesley, pp. 145–156. 2, 3

[GGSC96] GORTLER S. J., GRZESZCZUK R., SZELISKI R., COHEN M. F.: The lumigraph. In *SIGGRAPH 96, Computer Graphics Proceedings* (Aug. 1996), Addison Wesley, pp. 43–54. 2, 11

[HCD01] HAWKINS T., COHEN J., DEBEVEC P.: A photometric approach to digitizing cultural artifacts. In *2nd International Symposium on Virtual Reality, Archaeology, and Cultural Heritage, Glyfada, Greece, November 2001*. (2001). 2, 3

[HS99] HEIDRICH W., SEIDEL H.-P.: Realistic, hardware-accelerated shading and lighting. In *SIGGRAPH 1999 Conference Proceedings* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 171–178. 5

[KSS02] KAUTZ J., SLOAN P.-P., SNYDER J.: Fast, arbitrary brdf shading for low-frequency lighting using spherical harmonics. In *Proceedings of the 13th Eurographics workshop on Rendering* (2002), Eurographics Association, pp. 291–296. 9

[LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *SIGGRAPH 96 Conference Proceedings* (1996), Addison Wesley, pp. 31–42. 2

[LWS97] LEE S., WOLBERG G., SHIN S. Y.: Scattered data interpolation with multilevel B-splines. *IEEE Transactions on Visualization and Computer Graphics* 3, 3 (1997), 228–244. 6

[LWS01] LIN Z., WONG T.-T., SHUM H.-Y.: Relighting with the reflected irradiance field: Representation, sampling and reconstruction. *IEEE Computer Vision and Pattern Recognition 1* (Dec. 2001), 561–567. 2

[MDA02] MASSELUS V., DUTRÉ P., ANRYS F.: The free form light stage. In *Rendering Techniques EG 2002* (2002), Annual Conference Series, Eurographics Association, pp. 247–256. 2, 3

[MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *SIGGRAPH 2001, Computer Graphics Proceedings* (2001), Fiume E., (Ed.), ACM Press / ACM SIGGRAPH, pp. 519–528. 2, 4

[MPN*02] MATUSIK W., PFISTER H., NGAN A., BEARDSLEY P., ZIEGLER R., MCMILLAN L.: Image-based 3D photography using opacity hulls. In *SIGGRAPH 2002 Conference Proceedings* (2002), ACM Press, pp. 427–437. 2

[MPZ*02] MATUSIK W., PFISTER H., ZIEGLER R., NGAN A., MCMILLAN L.: Acquisition and rendering of transparent and refractive objects. In *Rendering Techniques EG 2002* (2002), Annual Conference Series, Eurographics Association, pp. 267–277. 2

[NRH03] NG R., RAMAMOORTHY R., HANRAHAN P.: All-frequency shadows using non-linear wavelet lighting approximation. *ACM Trans. Graph.* 22, 3 (2003), 376–381. 9

[RH01] RAMAMOORTHY R., HANRAHAN P.: An efficient representation for irradiance environment maps. In *SIGGRAPH 2001, Computer Graphics Proceedings* (2001), pp. 497–500. 7, 9

[SDS96] STOLLNITZ E. J., DEROSE T. D., SALESIN D. H.: *Wavelets for computer graphics: theory and applications*. Morgan Kaufmann Publishers, Inc., 1996. 5

[SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics* 21, 3 (July 2002), 527–536. 9

[SS95] SCHRÖDER P., SWELDENS W.: Spherical wavelets: Efficiently representing functions of the sphere. In *SIGGRAPH 95, Computer Graphics Proceedings* (Aug. 1995), Addison Wesley, pp. 161–172. 5

[SSC*96] SCHRÖDER P., SWELDENS W., COHEN M., DEROSE T., SALESIN D.: Wavelets in computer graphics, SIGGRAPH 96 course notes, 1996. 5

[UB03] UNSER M., BLU T.: Mathematical properties of the JPEG2000 wavelet filters. *IEEE Transactions on Image Processing* 12, 9 (Sept. 2003), 1080–1090. 9

[WHF01] WONG T.-T., HENG P.-A., FU C.-W.: Interactive relighting of panoramas. *IEEE Computer Graphics and Applications* 21, 2 (2001), 32–41. 2

[WHON97] WONG T.-T., HENG P.-A., OR S.-H., NG W.-Y.: Image-based rendering with controllable illumination. In *Eurographics Rendering Workshop 1997* (June 1997), Springer Wien, pp. 13–22. 2, 4




Figure 12: A column of pixels is selected from an image of the miniature race car. The incident illumination rotates while the resulting consecutive relit columns of pixels are displayed for different reconstruction and compression techniques.