

A Compact Factored Representation of Heterogeneous Subsurface Scattering

Pieter Peers¹ Karl vom Berge¹ Wojciech Matusik² Ravi Ramamoorthi³
Jason Lawrence⁴ Szymon Rusinkiewicz⁴ Philip Dutré¹
Katholieke Universiteit Leuven¹ MERL² Columbia University³ Princeton University⁴



Figure 1: A factored composite wax material model applied to the Stanford dragon. The material is composed of two kinds of wax with different scattering properties. Left: illuminated by an area light source from above. Middle: the material's diffuse albedo (no subsurface scattering). Right: illuminated from above by a texture projection light.

Abstract

Many translucent materials exhibit heterogeneous subsurface scattering, which arises from complex internal structures. The acquisition and representation of these scattering functions is a complex problem that has been only partially addressed in previous techniques. Unlike homogeneous materials, the spatial component of heterogeneous subsurface scattering can vary arbitrarily over surface locations. Storing the spatial component without compression leads to impractically large datasets. In this paper, we address the problem of acquiring and compactly representing the spatial component of heterogeneous subsurface scattering functions. We propose a material model based on matrix factorization that can be mapped onto arbitrary geometry, and, due to its compact form, can be incorporated into most visualization systems with little overhead. We present results of several real-world datasets that are acquired using a projector and a digital camera.

Keywords: Subsurface scattering, Non-negative matrix factorization, Image-based acquisition

1 Introduction

In recent years, subsurface scattering has received much attention in computer graphics. Initial research focused on the visualization and simulation of subsurface scattering materials [Hanrahan and Krueger 1993; Dorsey et al. 1999; Jensen et al. 2001; Lensch et al. 2003; Mertens et al. 2003]. With the rapid advancement of visualization algorithms, however, the need for measuring subsurface scattering properties of physical materials has increased.

Homogeneous subsurface scattering, such as that observed in milk, can be easily measured and fitted to an analytical model like the dipole approximation [Jensen et al. 2001]. But it is still an open

problem to acquire and represent general non-homogeneous subsurface scattering, observed in materials such as veined marble. Tong et al. [2005] identified two approaches to represent subsurface scattering: an *object model representation* and a *material model representation*. The object model approach captures the subsurface properties coupled to a specific geometry [Goessele et al. 2004]. The material model approach captures the subsurface scattering properties independent of the underlying geometrical shape. Unlike object model representations, a material model representation can be applied to any geometry at the cost of additional computations or approximations [Tong et al. 2005]. This paper focuses on the acquisition and compact representation of a material model for the spatial component of highly heterogeneous subsurface scattering materials.

We present a novel compact representation based on non-negative matrix factorization. A non-negative factored representation offers several advantages: the result of any light transport calculation will remain positive, the non-negative terms enable importance sampling, and could potentially allow user-guided editing. In the context of large multi-dimensional datasets, factorization has been a popular tool [Fournier 1995; Heidrich and Seidel 1999; Kautz and Seidel 2000; Kautz and McCool 1999; Latta and Kolb 2002; Suykens et al. 2003; Lawrence et al. 2004]. However, these techniques are designed for representing BRDFs and are not suited for representing the spatial component of heterogeneous subsurface scattering. A key observation is that most heterogeneous materials are a mix of a limited number of homogeneous basis materials. The spatial responses of these homogeneous basis materials are deformed by spatial discontinuities in the material. We will show that by eliminating the effect of the homogeneous subsurface scattering, a more suitable form for factorization can be found.

Synthetic datasets are free from calibration errors or measurement noise and are often too clean. It is therefore essential to validate the developed factorization method on real-world data. Since we are interested in obtaining a *material model*, we restrict the acquisition to planar samples in order to avoid geometrical spatial dependencies. Rather than the commonly used camera-laser systems, we employ a camera-projector system to acquire the spatial subsurface scattering component of a planar material sample. Our approach

Copyright © 2006 by the Association for Computing Machinery, Inc.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.

© 2006 ACM 0730-0301/06/0700-0746 \$5.00

speeds up acquisition by sampling multiple points in parallel, with minimal tradeoffs.

2 Related Work

The work presented in this paper is closely related to the following sub-domains in computer graphics: factorization of multi-dimensional functions, modeling and simulation/visualization of subsurface scattering, and image-based acquisition and representations of translucent objects and materials.

Factorization. Factorization has been a successful tool in decomposing BRDFs into sets of lower dimensional factors and terms, enabling interactive rendering [Fournier 1995; Heidrich and Seidel 1999; Kautz and Seidel 2000; Kautz and McCool 1999; Latta and Kolb 2002], and improved importance sampling [Lawrence et al. 2004]. Suykens et al. [2003] extend the idea of homomorphic factorization to interactively render BTFs. Recently, Lawrence et al. [2006] used constrained matrix factorization to represent and edit spatially varying BRDFs. These techniques rely on the assumption that the BRDF is a separable function, which we will show is not the case for subsurface scattering effects.

Rendering, Modeling and Simulation. A number of approaches simulate subsurface scattering by explicitly modeling the physical properties of each point inside a volume. In computer graphics, subsurface scattering was pioneered by Hanrahan and Krueger [1993]. Dorsey et al. [1999] and Pharr and Hanrahan [2000] successfully simulate heterogeneous materials using Monte Carlo methods by explicitly modeling the interactions inside the volume. Jensen et al. [2001] presented a practical model for homogeneous materials based on a dipole approximation. With the advent of powerful graphics hardware, interactive rendering techniques for translucent materials have been developed [Mertens et al. 2003; Lensch et al. 2003]. More recently Chen et al. [2004] presented a technique to render synthetic heterogeneous materials by partially pre-computing light transport near the surface and using a dipole approximation for the material's inner core. None of these approaches are suited to model or visualize real-world heterogeneous translucent materials.

Acquisition and Representation. A number of image-based techniques are able to capture and render translucent objects. Light-fields [Levoy and Hanrahan 1996] or lumigraphs [Gortler et al. 1996] capture the view-dependent appearance of objects under fixed lighting conditions. Surface light-fields [Miller et al. 1998; Wood et al. 2000] further improve on this approach. Reflectance fields [Debevec et al. 2000; Masselus et al. 2003] are capable of representing an object from a fixed viewpoint but under variable lighting conditions. However, none of these systems are geared towards acquiring and compactly representing subsurface scattering.

The recently proposed DISCO acquisition system [Goesele et al. 2004], an example of an object model representation, measures a 4D subsurface scattering function over an object with respect to incoming and outgoing surface points. The system assumes a smooth global subsurface scattering component, which is interpolated over the object's mesh. Tong et al. [2005] presented a technique to capture quasi-homogeneous materials, i.e., translucent materials with evenly-distributed heterogeneous elements. The authors base their system on the key observation that quasi-homogeneous materials exhibit a difference in subsurface scattering properties at local and global scales: locally, the non-uniformity in physical material properties leads to heterogeneous subsurface scattering, while at a global scale, the even distribution of heterogeneous elements leads to homogeneous scattering. Neither approach, however, is suited to accurately represent general heterogeneous translucent materials.

Fuchs et al. [2005] developed an empirical model for representing heterogeneous subsurface scattering based on the DISCO system [Goesele et al. 2004]. The response at each surface location is a linear combination of exponentials, modulated by a single spatially-varying texture map. To further compensate for anisotropic scattering behavior, the domain around each response is radially subdivided and the approximations over the different domains are blended together. Our method differs from this work in two ways. First, by using a data-driven approach based on matrix factorization, it can handle more general scattering behaviors than those representable by exponential fall-offs. Second, by using several spatial textures (i.e., factorization terms), it can model a greater range of spatial scattering variations, such as heterogeneities due to veins in marble.

3 Background

The behavior of subsurface scattering materials is described by the general *bidirectional surface scattering reflectance distribution function* (BSSRDF) $S(x_i, \omega_i; x_o, \omega_o)$ [Nicolodemus et al. 1977], which relates outgoing radiance $L(x_o, \omega_o)$ at a point x_o in a direction ω_o to the incident illumination $L(x_i, \omega_i)$ at a location x_i and incoming direction ω_i :

$$L(x_o, \omega_o) = \int_A \int_{\Omega} S(x_i, \omega_i; x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i,$$

where A is an area around the point x_o , and Ω denotes the upper hemisphere around x_i . We separate this integral into a local component L_l , which accounts for light immediately reflected from a surface, and a global component L_g :

$$L(x_o, \omega_o) = L_l(x_o, \omega_o) + L_g(x_o, \omega_o).$$

The local and global component L_l and L_g are defined as:

$$\begin{aligned} L_l(x_o, \omega_o) &= \int_{\Omega} f_s(x_o, \omega_o, \omega_i) L(x_o, \omega_i) (N_o \cdot \omega_i) d\omega_i, \\ L_g(x_o, \omega_o) &= \int_A \int_{\Omega} S_d(x_i, \omega_i; x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i, \end{aligned}$$

where f_s can be interpreted as a spatially varying BRDF, S_d is called the diffuse BSSRDF, and N_o is the surface normal at x_o . The diffuse BSSRDF is often decomposed into a product of lower-dimensional functions. A commonly used decomposition for S_d [Donner and Jensen 2005] is:

$$S_d(x_i, \omega_i; x_o, \omega_o) = \frac{1}{\pi} f_i(x_i, \omega_i) R_d(x_i, x_o) f_o(x_o, \omega_o),$$

where f_o and f_i are directionally dependent components, and $R_d(x_i, x_o)$ is the spatial subsurface scattering component. As in [Goesele et al. 2004] we will focus exclusively on the acquisition of the 4D spatial component R_d of heterogeneous subsurface scattering materials, ignoring any directional dependency (f_o and f_i). The acquisition of these angular components can be accomplished in a manner similar to that described by Tong et al. [2005].

4 Acquisition

In this section we discuss the acquisition of the spatial subsurface scattering component of a real-world sample. First, we describe the setup (a projector-camera pair) and compare this system to a camera-laser setup. Then we discuss the calibration required and finally, the acquisition process itself.

4.1 Acquisition Setup

The acquisition setup consists of a digital video camera and a projector. We position the sample approximately half a meter away from the projector, taking special care to ensure that the sample surface is perpendicular to the projector. The camera is positioned above the projector, such that no specular highlights from the sample are visible to the camera. During acquisition, we assume that there is no local component L_l . We acknowledge that this is inaccurate; but it does not affect the developed factorization method significantly.

In our setup we use an Allied Vision Technologies (AVT) Dolphin F-145C CCD FireWire camera equipped with a 12.0 – 36.0mm C-mount lens. We capture all photographs at the maximum resolution of 1392×1040 using 10 bits per pixel. For the projector we use an Optoma EP739H DLP projector with a 1 : 2000 contrast ratio and a digital DVI connection to avoid AD/DA conversion artifacts.

Previous approaches [Goesele et al. 2004; Tong et al. 2005; Fuchs et al. 2005] use a laser to sample the spatial subsurface scattering component. Replacing the laser with a projector offers the following advantages:

- **Multi-chromatic:** A laser system samples the color spectrum very sparsely (e.g., 3 wavelengths). Depending on the spectral response of the sample, a significant portion of the reflectance properties could be lost. Furthermore, laser speckle can occur due to the coherent nature of laser light. These concerns are absent when using a projector.
- **Complex patterns:** A projector setup can emit any illumination pattern without resorting to time multiplexing or using a complex additional lens system.
- **Cost:** A projector costs considerably less than a laser system. In addition, a projector can be bought off the shelf, and does not require special controller hardware.

A projector, however, also has disadvantages with respect to a laser system:

- **Refresh rate:** By construction, a projector refreshes the image at regular intervals. To avoid artifacts, only the exposure times that are a multiple of this refresh rate should be used. This places a lower bound on the exposure times when capturing high dynamic range (HDR) images. In principle, a laser system is not constrained by a minimum exposure time. However, unless a high speed camera (e.g., [Wenger et al. 2005]) is used, the projector's refresh rate is faster than the maximum frame rate of most high-resolution digital cameras.
- **Limited contrast ratio:** A noticeable amount of illumination is visible when emitting a black image from a projector. This background illumination places an upper limit on the exposure times before a photograph is over-saturated. In our setup, this was found to be in excess of 4 seconds, a point at which thermal noise in the camera already introduces significant artifacts.

4.2 Calibration

The camera and projector require careful calibration before the subsurface scattering of a sample can be accurately acquired.

Radiometric response. We directly read out the raw Bayer data from the camera, which has a linear radiometric response. A linear regression is applied to determine the slope for each pixel separately. After combining the raw Bayer photographs into a high dynamic range image, high-quality demosaicing [Malvar et al. 2004] is performed.

Intrinsic and extrinsic camera and projector parameters.

Both the intrinsic and extrinsic camera and projector parameters must be calibrated. Rather than calibrating each device independently (e.g., [Zhang 2000]), a simultaneous calibration process is used. First, the sample is covered by a white diffuse layer. Next, a series of calibration patterns are projected on the diffuse sample. Each calibration pattern consists of a set of regularly spaced dots; these dots are slightly shifted in consecutive calibration patterns. For each dot, the center position in the camera image is computed. Using the known positions of the dots in both camera and projector space, a homography can be recovered between both devices.

Camera and projector lens vignetting. In order to compensate for camera and projector lens vignetting, each recorded image is divided by an HDR photograph of a projected white image on the diffuse sample surface.

Color calibration is limited to white balancing that relates the observed white of the projector to a reference white.

Refresh rate of the projector. DLP projectors use a color wheel to induce color sensations. To avoid interference, the camera exposure times must be a multiple of the rotational frequency of the color wheel. The rotational frequency was empirically determined to be 30 Hertz.

Black level of the projector. To measure the black level, a large number (~ 100) of HDR Bayer images of the sample are recorded while emitting only a black pattern from the projector. For each pixel, the average and variance of the pixel intensity are computed. Each pixel in a newly recorded HDR photograph that falls below a confidence limit of 99% is set to zero. The average is subtracted from the remaining pixels to eliminate as much as possible the effect of background illumination.

4.3 Data Acquisition

As with a laser setup, individual surface points can be illuminated by a projector. By recording HDR photographs, 2D slices of the BSSRDF are obtained. This approach, however, results in impractically long acquisition times when a dense spatial resolution is required.

Theoretically, illuminating a single surface point yields an infinitely large response. In practice, the extent of this observed footprint is limited by the dynamic range of the recorded HDR photographs. This observation can be exploited to accelerate the acquisition process by emitting multiple beams in parallel and extracting the individual responses afterwards. Rather than illuminating a single surface point, a regular grid is projected onto the sample. The spacing in the grid is manually set to avoid influence between neighboring footprints.

The acquisition pipeline is as follows: For each emitted pattern, an HDR Bayer image is recorded. Next, the influence of the projector background illumination, lens vignetting, and thermal camera noise is removed and the resulting Bayer image is demosaiced. Using the camera-projector homography, we compensate for the off-center camera position so that there is a one-to-one mapping between camera and projector pixels. Finally, individual responses are extracted and white balanced.

Instead of shifting the consecutive grid patterns sequentially, we use a random order to minimize the effect of temporal fluctuations in projected intensity. For all the examples in this paper, exposures ranging from 1/30 to 2 seconds were used to capture HDR images. In conjunction with the 10 bit raw camera data, this yields a theoretical dynamic range of 100,000 : 1. However, after removing camera




Figure 2: The subsurface scattering matrix R_d . (a) An illustration of a homogeneous material containing a light blocking vein. (b) The corresponding subsurface scattering matrix R_d . The effect of the light blocking vein is expressed by horizontal and vertical discontinuities. (c) A sheared reparameterization of R_d .

noise and subtracting the black level of the projector, an average dynamic range of 10,000 : 1 is maintained. Acquisition timings range from 3 to 6 hours, depending on the spatial resolution and the grid spacing. For comparison, illuminating each surface point sequentially requires 50 hours or more.

5 Factorization

In this section we introduce our compact factored form, suitable for representing and rendering the spatial component of heterogeneous subsurface scattering materials acquired using the setup described in the previous section.

5.1 Motivation

In order to find a suitable factored form for the spatial component of the BSSRDF, we first store $R_d(x_i, x_o)$ in a matrix. Studying the general form of this matrix helps to determine a suitable factorization form. The global subsurface scattering matrix R_d is a linearized matrix of 2D surface points, where rows are over x_i and columns over x_o . Each element contains the ratio of light transferred from x_i to x_o , and vice versa. For simplicity we will assume that both x_i and x_o are parameterized over the surface with the same resolution.

The construction and form of the subsurface scattering matrix R_d is illustrated in Figure 2. For illustration purposes, only 2D subsurface scattering functions are considered; that is, both x_i and x_o are on a single line (1D). The conclusions drawn from this example, however, are analogous when using linearized 2D coordinates for x_i and x_o . Figure 2.a shows a material exhibiting homogeneous subsurface scattering, except in the middle, where a vein obstructs light propagation. Two responses are depicted; the left is unaffected by the vein, while the right response is cut off. In Figure 2.b the resulting subsurface scattering matrix $R_d(x_i, x_o)$ is shown. This is a banded matrix, in which the homogeneous response is shifted along the diagonal. In the middle of this matrix a discontinuity can be observed, caused by the light blocking vein. In general, heterogeneities are expressed as horizontal and vertical discontinuities in the diagonal structure, which is the result of the subsurface light transport.

Using classical factorization methods, such as singular value decomposition (SVD) or non-negative matrix factorization (NMF), directly on the matrix $R_d(x_i, x_o)$ does not yield satisfactory results. The matrix $R_d(x_i, x_o)$ is dominated by a diagonal structure. Matrices with these kinds of structures usually have a high rank, and consequently do not decompose very well into lower order terms. A standard solution to such a problem is to reparameterize the matrix so that it becomes more suitable for factorization. An obvious reparameterization would be to shear the matrix, such that the diagonal becomes aligned to one of the axes. In Figure 2.c, a reparameterized $R'_d(x_i, d)$ is shown, with $d = x_o - x_i$. The diagonal structure that hindered factorization is now rectified. However, the heterogeneities are now expressed in features along the horizontal and




Figure 3: A marble material sample illustrating the structure of the subsurface scattering matrix R_d . The responses are measured on a single line, indicated in green. The discontinuities are not as pronounced in the subsurface matrix $R_d(x_i, x_o)$, but are clearly visible in the reparameterized matrix $R'_d(x_i, d)$ (marked in red in the reparameterized matrix). Regardless of the geometrical form of the heterogeneity (e.g., the blue marked vein), the effects on R_d are always expressed as horizontal and vertical discontinuities.

skew diagonal direction, reducing the effectiveness of a classical factorization.

We verified the structure of the subsurface scattering matrix $R_d(x_i, x_o)$ on a slab of marble, shown in Figure 3. The marble responses are measured for different points on a single line, indicated in green. The discontinuities are not as pronounced as expected in the subsurface matrix $R_d(x_i, x_o)$, because the veins through the marble do not block the light completely. However, the horizontal and skew diagonal discontinuities, highlighted in red, are clearly visible in the reparameterized matrix $R'_d(x_i, d)$. This empirically confirms our conclusions of the experiment in Figure 2, regarding the structure and form of the subsurface scattering matrix R_d .

5.2 Factorization Formula

In this paper, we will not use a reparameterization to make the subsurface matrix $R_d(x_i, x_o)$ more suitable for factorization, but will instead divide out as much as possible the diagonal structure that hinders efficient factorization.

From the reparameterized matrix $R'_d(x_i, d)$, an average response function $g(d)$ can be computed by simply taking the average response or by taking the maximum value over each column (Figure 4.a). This average response function $g(d)$ is a good approximation for the homogeneous subsurface scattering kernel, and can be used to create a homogeneous approximation $G(x_i, x_o)$ of $R_d(x_i, x_o)$ (Figure 4.b):

$$G(x_i, x_o) = g(d) = g(x_o - x_i).$$

An interesting result can be obtained by dividing the subsurface scattering matrix $R_d(x_i, x_o)$ *component-wise* by $G(x_i, x_o)$, effectively dividing out the diagonal subsurface scattering features. This is depicted in Figure 5 (left). The dashed elements of $R_d(x_i, x_o)/G(x_i, x_o)$ in Figure 5 mark elements that are the result of a division by zero. As a consequence, a matrix suitable for factorization can be obtained by careful selection of the values of these elements in R_d/G . In Figure 5, the dashed elements are filled in such that the resulting matrix can be easily separated into a 2-term factorization.

Moving the G term to the other side yields the final factored form:

$$R_d(x_i, x_o) \approx (F(x_i, t)H(t, x_o)) \star G(x_i, x_o),$$

where $\star$ denotes a *component-wise* matrix multiplication defined by: $(A \star B)_{ij} = a_{ij}b_{ij}$. The width of F and height of H are user




Figure 4: By computing the average response $g(d)$, a homogeneous approximation $G(x_i, x_o)$ of the subsurface scattering matrix $R_d(x_i, x_o)$ can be obtained.

selected (i.e., number of terms), and determine the compression ratio and the degree of approximation. Although $G(x_i, x_o)$ has the same dimensions as $R_d(x_i, x_o)$, the storage requirements are minimal, since only $g(d)$ needs to be stored.

As mentioned previously, the conclusions from this 2D example also hold for general 4D subsurface scattering functions. It is important to realize that, regardless of the specific geometrical form of a heterogeneity (e.g., the blue marked vein on the marble slab in Figure 3), the effects on the subsurface scattering matrix $R_d(x_i, x_o)$ are always expressed as horizontal and vertical discontinuities.

5.3 Update Rules

As detailed in the previous section, the homogeneous subsurface scattering matrix G can be easily determined if R_d is known. The F and H terms, however, cannot be determined directly and need to be factored out. In this paper, we employ a variant of iterative non-negative matrix factorization [Lee and Seung 2000] to decompose the matrix $R_d(x_i, x_o)$. F and H are initialized to random values, and subsequently iteratively updated, using the multiplicative update rules presented below, until convergence is reached.

More specifically, the factorization should minimize the following error function, keeping in mind the non-negativity constraint:

$$\mathcal{E}^2 = \frac{1}{2} \|W \star (R_d - ((FH) \star G))\|_F^2,$$

where $\|\cdot\|_F$ is the Frobenius norm and W a weighting matrix. By following a similar strategy as Lee and Seung [2000], multiplicative update rules can be derived:

$$\begin{aligned} H &\leftarrow H \star \frac{F^T (W^2 \star G \star R_d)}{F^T (W^2 \star G \star ((FH) \star G))}, \\ F &\leftarrow F \star \frac{(W^2 \star G \star R_d) H^T}{(W^2 \star G \star ((FH) \star G)) H^T}, \end{aligned}$$

where the division is *component-wise* and $W^2 = W \star W$. A formal derivation and proof of convergence can be found in [Peers and Dutré 2006].

Due to sub-pixel errors in the camera-projector calibration, it is possible that the peak of the average exponential fall-off function $g(d)$ is slightly off-center. This will have a significant impact on the error $\mathcal{E}^2$, and the factorization process will spend a large percentage of all iteration steps trying to keep the peak as close as possible to the center. A suitable weighting matrix W can minimize the effect of a misaligned peak, and convergence can be attained in fewer iterations. Setting $w_{ij} = 1/\sqrt{g_{ij}}$ has this desired effect. Moreover, it simplifies the update rules:

$$H \leftarrow H \star \frac{F^T R_d}{F^T ((FH) \star G)},$$




Figure 5: The homogeneous approximation $G(x_i, x_o)$ can be used to eliminate the diagonal subsurface scattering features in $R_d(x_i, x_o)$ by a component-wise division. Dashed elements are the result of a division by zero. The values of these elements can be set such that the resulting matrix is better suited for factorization.

$$F \leftarrow F \star \frac{R_d H^T}{((FH) \star G) H^T}.$$

To regularize the update rules above, we follow the method of Cichocki et al. [2006]. Furthermore, to avoid dividing by zero, a small ϵ -value is added to each element in $g(d)$.

5.4 Clustering

We have thus far assumed that the subsurface scattering sample consists of a homogeneous material, where the heterogeneities are the result of geometrical deficiencies. This works well for materials such as marble but fails when the sample contains a combination of basis materials, such as a chess board made of two interlocking materials (see Figure 7.b).

We solve this problem by first clustering the responses (rows of R_d) and computing a subsurface scattering kernel $g(d)$ for each cluster separately. The resulting matrix G will be a mix of the different kernels $g(d)$. Ideally we would like to cluster on subsurface scattering properties (i.e., exponential fall-off) and not on spatial heterogeneities. To minimize the effect of these spatial heterogeneities we cluster according to the feature vectors $f_{x_i}(r)$:

$$f_{x_i}(r) = \max_{\alpha} R_d(x_i, p_{x_i}(\alpha, r)),$$

where $p_{x_i}(\alpha, r)$ is a polar mapping of x_o around x_i . This function $f_{x_i}(r)$ will be an approximation of the exponential fall-off function of the response at position x_i , and is further normalized to eliminate the effect of different response albedos. In our implementation, we use a k -means clustering algorithm. The number of clusters k is selected by the user, and should correspond to the number of distinct materials with different scattering characteristics that are present in the dataset. Since the clustering is performed beforehand and the resulting homogeneous approximation G can be computed without much additional overhead, factorization will not incur a penalty.

6 Visualization

Our compact material model for subsurface scattering can be easily incorporated into an existing renderer. The results in this paper are generated by an adapted photon map [Jensen 2001] implementation of PBRT [Pharr and Humphreys 2004]. Photons entering a translucent material are stored in a special subsurface scattering map. Photons not absorbed by a subsurface scattering material are traced as usual. During rendering, when a ray hits a translucent material, all photons from the subsurface scattering map, within the range of the subsurface scattering response, are gathered and weighted by the response at that point. The weighting can be efficiently computed using the factored representation. In our current implementation we use standard texture mapping to apply a material model to a mesh.

But low-distortion texture mapping would be preferable to minimize approximation errors. A texture synthesis technique, such as Liu et al. [2004], can be adapted to better texture arbitrary geometries.

7 Results and Discussions

We verified our factorization method on several real-world subsurface scattering examples, ranging from fairly homogeneous to highly translucent heterogeneous materials. Table 1 gives an overview of the acquired and factored materials, the parameters used, the compression ratio, and relative errors. Although the kernel size is a slight exaggeration of the subsurface scattering footprint, it can still be used to show the extent of the subsurface scattering for each example (i.e., $\text{Kernel Size} \times \text{Physical Size/Resolution}$). For example, the wax candle material used in Figure 7.a has a kernel radius of approximately 4.5cm, while the kernel radius for marble is approximately 0.75cm. On average, a compression ratio of 1/50 is obtained. Note, that this ratio compares the size of factored results with the total storage requirements of only the kernels of each response. Finally, the relative errors are computed by comparing, for each surface point, the measured response with the factored response, normalized by the energy of the response.

The approximation error is determined by the number of terms, the number of clusters, and the number of iterations. Similarly to standard non-negative matrix factorization, the error decays exponentially, and drops significantly during the first 50 iterations. The optimal number of terms and clusters depends on the specific material. However, due to the rapid error decay, a quick verification of the selected parameters is possible after only a few (i.e., 50) iterations. In general, we observe that the number of clusters is related to the degree of heterogeneity in the material. It is possible that different materials are present in a single pixel, resulting in a larger number of required clusters than initially anticipated. The number of terms, on the other hand, is related to the degree of discontinuity due to the heterogeneities. A subsurface scattering material exhibiting sharp discontinuities requires fewer terms than a material with smooth features.

A selection of acquired materials and factored results is shown in Figure 7. For each material, we include a photograph of the original sample. We also show the diffuse albedo and a relative error distribution plot. In addition, for each material a selection of measured responses and their factored counterparts are shown. The locations of these responses are marked on the diffuse albedo map. The dashed square illustrates the relative size of the responses. The range of errors in the relative error distribution plots can be inferred from Table 1. The relative error distribution plot and the reconstructed responses show that our factorization method is able to faithfully reproduce the heterogeneous discontinuities. An additional advantage is that the acquired data is denoised by our factorization method.

Figures 1 and 6 show visualizations of compactly factored material models mapped onto a novel geometry. The material model used in the Stanford dragon example of Figure 1 is the candle material (first column in Figure 7). This material consists of two components: red wax, exhibiting little subsurface scattering, and very translucent yellow wax. Heterogeneities are caused by the internal structure of the candle, and at the interface between the two materials. For comparison we also show the dragon textured by the diffuse albedo map of the candle material (Figure 1, middle). The right example of Figure 1 is illuminated from above by a texture projection light. The emitted pattern is still distinguishable on the red wax and on the ground plate, while on the highly translucent yellow wax the pattern is almost completely washed out due to the subsurface scattering. The effect of the heterogeneous subsurface

scattering is most visible at the interface of the red and yellow wax. In Figure 6, the well known Buddha model is visualized using the material models for layered white onyx and cracked crystal onyx (respectively, the third and fourth columns in Figure 7). For both examples we also included a visualization under a texture projection light to better illustrate the effects of the subsurface scattering.

8 Conclusions

In this paper we have presented a compact factored representation for the spatial component of heterogeneous subsurface scattering materials that is based on non-negative matrix factorization. Our representation consists of a number of terms that modulate incoming and outgoing radiance and a piece-wise linear homogeneous approximation of the subsurface scattering core. We observe that the homogeneous subsurface scattering kernel can be divided out, leaving just the discontinuities caused by the heterogeneities present in the material. These residual discontinuities can be efficiently factorized using a non-negative factorization approach.

We have illustrated the efficiency of our factorization technique on a number of real-world heterogeneous subsurface scattering material datasets. These datasets are obtained using a projector-camera pair. To speed up acquisition, multiple responses are recorded in parallel. Finally, we have demonstrated that the compact factored representation of our material models can be easily integrated into a standard global illumination rendering system, resulting in convincing images.

For future work, we would like to look into out-of-core factorization methods. Currently our implementation requires the whole subsurface scattering matrix to fit in the main memory, limiting the resolution at which the samples can be acquired. Finally, a piece-wise linear homogeneous approximation is not always well suited to divide out the subsurface scattering kernel. For example, a sample might fall right on an edge between two different materials. In this case the resulting sample response would be a linear combination of the responses of the two materials. Using a linear combination of multiple subsurface scattering kernels would yield a better approximation. However, it is not clear how this case might be factorized in a computationally stable manner.

Acknowledgments

This research was funded by K.U.Leuven grant GOA/2004/05, NSF grants #0446916 and #0347427, and the Sloan Foundation. Special thanks to Saskia Mordijck and Matt Brand.

References

- CHEN, Y., TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2004. Shell texture functions. *ACM Transactions on Graphics* 23, 3, 343–353.
- CICHOCKI, A., AMARI, S., AND ZDUNEK, R. 2006. Extended SMART algorithms for non-negative matrix factorization. In *Proceeding of the Eighth International Conference on Artificial Intelligence and Soft Computing*.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 145–156.
- DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics* 24, 3, 1032–1039.
- DORSEY, J., EDELMAN, A., JENSEN, H. W., LEGAKIS, J., AND PEDERSEN, H. K. 1999. Modeling and rendering of weathered stone. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 225–234.
- FOURNIER, A. 1995. Separating reflection functions for linear radiosity. In *Eurographics Workshop on Rendering*, 296–305.




Figure 6: The Stanford Buddha model rendered using the material models for layered white onyx and cracked crystal onyx (respectively, the third and fourth column in Figure 7). For both materials a visualization under textured and uniform illumination is shown.

- FUCHS, C., GOESELE, M., CHEN, T., AND SEIDEL, H.-P. 2005. An empirical model for heterogeneous translucent objects. Research Report MPI-I-2005-4-006, Max-Planck-Institut für Informatik, May.
- GOESELE, M., LENSCH, H. P. A., LANG, J., FUCHS, C., AND SEIDEL, H.-P. 2004. DISCO: acquisition of translucent objects. *ACM Transactions on Graphics* 23, 3, 835–844.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 43–54.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, 165–174.
- HEIDRICH, W., AND SEIDEL, H.-P. 1999. Realistic, hardware-accelerated shading and lighting. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 171–178.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 511–518.
- JENSEN, H. W. 2001. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters.
- KAUTZ, J., AND MCCOOL, M. D. 1999. Interactive rendering with arbitrary BRDFs using separable approximations. In *SIGGRAPH '99: ACM SIGGRAPH 99 Conference abstracts and applications*, 253.
- KAUTZ, J., AND SEIDEL, H.-P. 2000. Towards interactive bump mapping with anisotropic shift-variant BRDFs. In *HWWS '00: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware*, 51–58.
- LATTA, L., AND KOLB, A. 2002. Homomorphic factorization of BRDF-based lighting computation. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, 509–516.
- LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHY, R. 2004. Efficient BRDF importance sampling using a factored representation. *ACM Transactions on Graphics* 23, 3, 496–505.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics*.
- LEE, D. D., AND SEUNG, H. S. 2000. Algorithms for non-negative matrix factorization. In *NIPS*, 556–562.
- LENSCH, H. P. A., GOESELE, M., BEKAERT, P., KAUTZ, J., MAGNOR, M. A., LANG, J., AND SEIDEL, H.-P. 2003. Interactive rendering of translucent objects. *Computer Graphics Forum*, 195–205.
- LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 31–42.
- LIU, X., HU, Y., ZHANG, J., TONG, X., GUO, B., AND SHUM, H.-Y. 2004. Synthesis and rendering of bidirectional texture functions on arbitrary surfaces. *IEEE Transactions on Visualization and Computer Graphics* 10(3), 278–289.
- MALVAR, H. S., WEI HE, L., AND CUTLER, R. 2004. High-quality linear interpolation for demosaicing of Bayer-patterned color images. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2003. Relighting with 4D incident light fields. *ACM Transactions on Graphics* 22, 3, 613–620.
- MERTENS, T., KAUTZ, J., BEKAERT, P., SEIDEL, H.-P., AND REETH, F. V. 2003. Interactive rendering of translucent deformable objects. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, 130–140.
- MILLER, G., RUBIN, S., AND PONCELEN, D. 1998. Lazy decomposition of surface light fields for precomputed global illumination. In *EGRW '98: Proceedings of the 9th Eurographics workshop on Rendering*, 281–292.
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. Geometric considerations and nomenclature for reflectance. Tech. Rep. 161, National Bureau of Standards (US).
- PEERS, P., AND DUTRÉ, P. 2006. Update rules for a weighted non-negative FH*G factorization. Technical Report CW440, Katholieke Universiteit Leuven, April.
- PHARR, M., AND HANRAHAN, P. 2000. Monte Carlo evaluation of non-linear scattering equations for subsurface reflection. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 75–84.
- PHARR, M., AND HUMPHREYS, G. 2004. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc.
- SUYKENS, F., VOM BERGE, K., LAGAE, A., AND DUTRÉ, P. 2003. Interactive rendering with bidirectional texture functions. *Computer Graphics Forum* 22, 3, 463–472.
- TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Modeling and rendering of quasi-homogeneous materials. *ACM Transactions on Graphics* 24, 3, 1054–1061.
- WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics* 24, 3, 756–764.
- 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 '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 287–296.
- ZHANG, Z. 2000. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 11, 1330–1334.

Sample Material		Physical Size (cm ²)	Resolution (pixels)	Kernel Size (pixels)	No of Terms	No of Clusters	Original Data Size	Factored Data Size	Ratio	Min/Avg/Max Relative Error
Candle (Red and Yellow Wax)	(a)	19.8 × 14.3	212 × 154	49 × 49	12	100	898Mb	13Mb	1/69	0.002/0.031/0.050
White and Green Marble (4 × 4)	(b)	12.6 × 12.6	277 × 277	39 × 39	20	100	1.4Gb	39Mb	1/37	0.003/0.010/0.060
White and Green Marble (8 × 8)		25.1 × 25.1	222 × 222	39 × 39	8	10	859Mb	12Mb	1/72	0.003/0.019/0.094
Vertical Lines (White Onyx)	(c)	15.2 × 15.2	229 × 229	39 × 39	20	10	914Mb	26Mb	1/35	0.004/0.017/0.100
Cracked Material (Crystal Onyx)	(d)	18.5 × 17.6	270 × 260	45 × 45	12	5	1.6Gb	22Mb	1/74	0.011/0.040/0.560
Marble (close up)		2.6 × 2.6	128 × 128	39 × 39	16	1	286Mb	6.6Mb	1/43	0.006/0.013/0.056
Densely Veined Marble		13.0 × 13.0	213 × 211	29 × 29	8	10	433Mb	9.9Mb	1/44	0.003/0.096/0.163
Slightly Veined Marble		17.9 × 17.9	207 × 207	29 × 29	20	1	413Mb	21.0Mb	1/19	0.001/0.005/0.024

Table 1: Statistics and details regarding acquired and factored subsurface scattering materials.






























	Candle (a)		Checker board (4 × 4) (b)		White Onyx (c)		Crystal Onyx (d)	
Original Sample								
Diffuse Albedo								
Relative Error								
1.								
2.								
3.								
4.								
5.								
6.								
	Original	Factored	Original	Factored	Original	Factored	Original	Factored

Figure 7: A selection of acquired and factored materials. For each material, a photograph of the original sample, the diffuse albedo map, a relative error distribution plot, and a selection of measured responses with the corresponding factored approximations are shown. The locations of the responses are marked on the diffuse albedo map. The dashed square illustrates the relative size of the responses. Table 1 gives additional information regarding the range of the relative errors.