

Modeling and Rendering of Quasi-Homogeneous Materials

Xin Tong

Jiaping Wang*

Stephen Lin

Baining Guo

Heung-Yeung Shum

Microsoft Research Asia

*CAS Institute of Computing Technology

Abstract

Many translucent materials consist of evenly-distributed heterogeneous elements which produce a complex appearance under different lighting and viewing directions. For these *quasi-homogeneous* materials, existing techniques do not address how to acquire their material representations from physical samples in a way that allows arbitrary geometry models to be rendered with these materials. We propose a model for such materials that can be readily acquired from physical samples. This material model can be applied to geometric models of arbitrary shapes, and the resulting objects can be efficiently rendered without expensive subsurface light transport simulation. In developing a material model with these attributes, we capitalize on a key observation about the subsurface scattering characteristics of quasi-homogeneous materials at different scales. Locally, the non-uniformity of these materials leads to inhomogeneous subsurface scattering. For subsurface scattering on a global scale, we show that a lengthy photon path through an even distribution of heterogeneous elements statistically resembles scattering in a homogeneous medium. This observation allows us to represent and measure the global light transport within quasi-homogeneous materials as well as the transfer of light into and out of a material volume through surface mesostructures. We demonstrate our technique with results for several challenging materials that exhibit sophisticated appearance features such as transmission of back illumination through surface mesostructures.

Keywords: subsurface scattering, reflectance and shading models, rendering

1 Introduction

Many non-homogeneous translucent materials are composed of various substances that are evenly distributed throughout its volume. These *quasi-homogeneous* materials present a formidable challenge in realistic rendering because of their complex spatially variant subsurface scattering properties [Hanrahan and Krueger 1993]. Furthermore, their appearance is often complicated by surface mesostructures [Koenderink and Doorn 1996] that not only produce surface reflections, but also affect how light enters and exits a material volume. That computer graphics has still not been able to convincingly render a slice of bread, as notably mentioned by the late Alain Fournier [Fiume 2001], can largely be attributed to such intricacies in material structure and optical properties.

Quasi-homogeneous materials, or any other non-homogeneous material, can be adequately described by the bidirectional scattering-surface reflectance-distribution function (BSSRDF) [Nicodemus et al. 1977]. Unfortunately, acquiring a BSSRDF is extremely difficult. While a technique is available [Jensen et al.



Figure 1: A slice of bread rendered using a quasi-homogeneous material model acquired from a physical sample. The geometric model of the rendered slice was designed by an artist. Composed of many evenly distributed cavities and ingredients, bread is a common example of a quasi-homogeneous material.

2001] for measuring the BSSRDF of homogeneous materials, it is not clear how to extend this technique to non-homogeneous materials. Another approach to modeling quasi-homogeneous materials is to create synthetic volumetric models with subsurface scattering properties. This, however, has proven to be a highly non-trivial task [Chen et al. 2004]. Synthetic volumetric models furthermore require expensive light transport simulation for rendering [Dorsey et al. 1999; Pharr and Hanrahan 2000].

In this paper, we present a technique for modeling and rendering quasi-homogeneous materials that is based on a material representation which can be readily acquired from physical samples. This material representation can be applied to geometric models of arbitrary shape, and the resulting objects can be efficiently rendered without expensive subsurface light transport simulation. To emphasize its broad applicability to arbitrary geometric models, we call this representation a *material model* to distinguish it from *object models* which are typically used in image-based modeling and rendering of real-world objects (e.g., [Goessele et al. 2004]). A material model describes how light is scattered by the material, while an object model captures the actual appearance of a specific physical object. Unlike a material model, an object model measured from a given translucent object cannot be applied to objects of other shapes, because the appearance of each surface point is not a local quantity but the result of light propagation within the whole object.

The key observation behind our material model is that the subsurface scattering characteristics of quasi-homogeneous materials can be effectively analyzed through examination at two different scales. At a local level, the heterogeneity of a volume leads to non-homogeneous subsurface scattering. At a larger scale, because of the even distribution of materials within a quasi-homogeneous volume, small neighborhoods centered at different points in the volume

*This work was done while Jiaping Wang was a visiting student at Microsoft Research Asia.

have a statistically similar material composition and distribution. For global subsurface scattering at this larger scale, photon propagation is similar to that in a homogeneous material composed of particles whose scattering properties approximate the overall scattering of these volume neighborhoods.

Based on this observation of local heterogeneity and global homogeneity, we derive a material representation that consists of four components: a homogeneous subsurface scattering component for global light transport, a local reflectance function, a mesostructure entrance function, and a mesostructure exiting function. The last three components account for local scattering effects that include inhomogeneous local subsurface scattering, surface reflections from mesostructure, and locally inhomogeneous mesostructure effects on light entering and exiting the material. The four components of our model can be acquired by capturing both laser stripe and halogen lamp images of a physical sample under multiple illumination and viewing directions, and then breaking down these images into appropriate components.

We demonstrate results for several challenging materials that exhibit sophisticated appearance features, including bread as shown in Fig. 1. Although some of these materials have been measured as bidirectional texture functions (BTFs) in the CURET database [Dana et al. 1999], a BTF exhibits global light transport specific to the shape and size of the imaged sample, and consequently does not represent light scattering within a volume of arbitrary geometry. Moreover, since a BTF does not measure the transmission of back-light through materials, a translucent object synthesized with BTFs will exhibit incorrect rendering results, especially in places where a significant amount of light leaves the material volume through mesostructures. This problem is addressed in our model by the mesostructure entrance and exiting functions.

In the remainder of the paper, we present our modeling and rendering technique for quasi-homogeneous materials as follows. Section 2 reviews related work. Our material model is described in Section 3, and the system for its acquisition is outlined in Section 4. The rendering method for the material model is presented in Section 5. Results are reported in Section 6, followed by conclusions and discussion of future work in Section 7.

2 Related Work

Material Models: The most widely used material model in computer graphics is the bidirectional reflectance distribution function (BRDF), which was first derived in [Nicomemus et al. 1977] as a simplification of the BSSRDF. Since this simplification excludes light transport in a material volume, the BRDF is intended for homogeneous opaque materials. To model from physical material samples, numerous techniques have been presented for measuring BRDFs (e.g., [Ward 1992; Matusik et al. 2003]). For non-homogeneous opaque materials, spatially-variant BRDFs are generally used, such as in [Gardner et al. 2003]. The BTF introduced in [Dana et al. 1999] furthermore handles surface mesostructures, and can be measured in a variety of ways (e.g., [Dana 2001; Han and Perlin 2003]). The polynomial texture map (PTM) [Malzbender et al. 2001] also deals with spatially-variant reflectance and mesostructures but only for a fixed viewpoint. All BRDFs, BTFs, and PTMs are not suitable for translucent materials because these material models do not provide sufficient information on subsurface scattering. For example, the transmission of backlight through materials is not captured.

Subsurface scattering was introduced to computer graphics by [Hanrahan and Krueger 1993]. For homogeneous translucent materials, a practical subsurface light transport model was proposed in [Jensen et al. 2001]. Concurrent work [Koenderink and van Doorn 2001] also studied translucent materials using a diffusion approximation. For homogeneous material models, some efficient rendering techniques have been presented [Jensen and Buhler 2002;




Figure 2: Light scattering geometry for quasi-homogeneous materials.

Mertens et al. 2003; Hao et al. 2003; Lensch et al. 2003].

For non-homogeneous materials, early techniques [Dorsey et al. 1999; Pharr and Hanrahan 2000] create synthetic volumetric models and employ expensive light transport simulations for rendering. The recently proposed shell texture function (STF) [Chen et al. 2004] gains efficiency by partially precomputing light transport near the material surface and approximating the material deep within objects as a homogeneous core. Capturing STFs from real-world materials, however, remains very difficult. In [Chen et al. 2004], a material sample is segmented into homogeneous sub-volumes according to its CT data, and the optical properties of each sub-volume are obtained from tables of measured homogeneous materials [Jensen et al. 2001]. Segmenting CT data is a well-known challenge in medical imaging that is tedious to perform. Furthermore, few measurement tables are available, and some materials such as marble exhibit variations in albedo that cannot be determined from tables.

Object Models: Image-based models of objects are easier to acquire than material models, and naturally include subsurface scattering, mesostructures and other appearance features. Because of the high dimensionality of viewing and lighting conditions, these methods record only a subset of possible object appearances. Surface light fields [Wood et al. 2000; Chen et al. 2002] capture the appearance of an object from different viewing angles under a fixed lighting environment, while reflectance fields [Debevec et al. 2000; Masselus et al. 2003] represent an object from a fixed viewpoint with different lighting directions. Matusik et al. [2002a] present a system for capturing reflectance fields of distant lighting from arbitrary viewpoints. For specifically acquiring transparent or highly translucent objects, environment matting techniques have been presented to capture light refraction [Matusik et al. 2002b; Zongker et al. 1999; Chuang et al. 2000].

For inhomogeneous translucent objects, the DISCO acquisition system [Goesele et al. 2004] recovers a 4D subsurface scattering function with respect to incoming and outgoing surface points, while disregarding the directional dependence on lighting and viewpoint. Surface geometries are assumed to be smooth for accurate interpolation of global light transport. Because of its smooth surface restriction, DISCO does not address the issue of correctly capturing how light enters and leaves an object through mesostructures.

To avoid object dependence, we recover from images a material model that can be employed with objects of arbitrary geometry. Additionally, the transmission of light through mesostructures is measured.

3 Material Representation

The bidirectional scattering-surface reflectance-distribution function (BSSRDF) provides a general model for light scattering from surfaces, and from it we present our representation for quasi-homogeneous materials. The outgoing radiance $L(x_o, \omega_o)$ at a surface point x_o in direction ω_o can be computed by integrating the contribution of incoming radiance $L(x_i, \omega_i)$ for all incident directions ω_i over the surface A ,

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

Here, $dx_i = (n \cdot \omega_i) dA(x_i)$, where $(n \cdot \omega_i)$ is the cosine factor and $dA(x_i)$ is a differential area at x_i . We separate the integral into local and global contributions

$$L(x_o, \omega_o) = \int_{A_o} \int_{\Omega} S(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i + \int_{B_o} \int_{\Omega} S(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i$$

where the local contribution is integrated over a small disk A_o of radius r_s around x_o and the global contribution is aggregated from the rest of the surface $B_o = A - A_o$, illustrated in Fig. 2. Since the incident radiance $L(x_i, \omega_i)$ may be regarded as locally uniform over A_o , the local contribution $L_l(x_o, \omega_o)$ can be written as

$$L_l(x_o, \omega_o) = \int_{A_o} \int_{\Omega} S(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i \approx \int_{\Omega} R(x_o, \omega_i, \omega_o) L(x_i, \omega_i) d\omega_i$$

where

$$R(x_o, \omega_i, \omega_o) = \int_{A_o} S(x_i, \omega_i, x_o, \omega_o) dx_i$$

is the local reflectance function.

For the global contribution

$$L_g(x_o, \omega_o) = \int_{B_o} \int_{\Omega} S(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i, \quad (1)$$

the incoming radiance $L(x_i, \omega_i)$ undergoes significant scattering and we wish to approximate its impact at x_o by an averaging effect as follows:

$$\int_{B_o} \int_{\Omega} \frac{1}{\pi r_s^2} \left[\int_{A_i} S(x_i + r_i, \omega_i, x_o, \omega_o) dr_i \right] L(x_i, \omega_i) d\omega_i dx_i$$

where A_i is a disk of radius r_s around x_i and $L(x_i, \omega_i)$ is locally uniform over A_i . We can further average over A_o and obtain

$$\int_{B_o} \int_{\Omega} \frac{1}{\pi^2 r_s^4} \left[\int_{A_i} \int_{A_o} S(x_i + r_i, \omega_i, x_o + r_o, \omega_o) dr_i dr_o \right] L(x_i, \omega_i) d\omega_i dx_i.$$

Notice that

$$\int_{A_i} \int_{A_o} S(x_i + r_i, \omega_i, x_o + r_o, \omega_o) dr_i dr_o L(x_i, \omega_i)$$

is the radiance contribution from area A_i to area A_o . Since a quasi-homogeneous material is homogeneous at a large scale and is optically thick, we can apply the dipole diffusion approximation to this area-to-area contribution:

$$\int_{A_i} \int_{A_o} S(x_i + r_i, \omega_i, x_o + r_o, \omega_o) dr_i dr_o L(x_i, \omega_i) \approx \pi^2 r_s^4 F_o(A_o, \omega_o) R_d(x_i, x_o) L(x_i, \omega_i) f(\omega_i)$$

where $R_d(x_i, x_o)$ is the diffuse reflectance given by the dipole approximation, $F_o(A_o, \omega_o)$ is the average outgoing Fresnel term over A_o in direction ω_o , and $f(\omega_i)$ is the mesostructure entrance function of the material, which essentially represents an average product of the incoming Fresnel term and cosine foreshortening factor over the surface area. The mesostructure entrance function is expressed independently of location since its average effect within a local region A_i is taken to be uniform over the material surface.

Based on the above analysis for a local region A_o , we know that $L_g(x_o, \omega_o)$ is proportional to the dipole diffusion term

$$L_g(x_o, \omega_o) \propto \int_{B_o} \int_{\Omega} R_d(x_i, x_o) L(x_i, \omega_i) f(\omega_i) d\omega_i dx_i.$$

The above relation can be rewritten as

$$L_g(x_o, \omega_o) = \int_{B_o} \int_{\Omega} f_v(x_o, \omega_o) R_d(x_i, x_o) L(x_i, \omega_i) f(\omega_i) d\omega_i dx_i \quad (2)$$




Figure 3: Capture device of the material measurement system

by introducing a mesostructure exiting function $f_v(x_o, \omega_o)$ that modulates the average diffuse contribution of the dipole diffusion term according to the view-dependent effects at x_o , the local variation within A_o , and the outgoing Fresnel factor at x_o . From Eq. (1) and Eq. (2), the mesostructure exiting function can be expressed as

$$f_v(x_o, \omega_o) = \frac{\int_{B_o} \int_{\Omega} S(x_i, \omega_i, x_o, \omega_o) L(x_i, \omega_i) d\omega_i dx_i}{\int_{B_o} \int_{\Omega} R_d(x_i, x_o) L(x_i, \omega_i) f(\omega_i) d\omega_i dx_i}. \quad (3)$$

In summary, the outgoing radiance from a quasi-homogeneous material can be formulated as

$$L(x_o, \omega_o) = \int_{\Omega} R(x_o, \omega_i, \omega_o) L(x_i, \omega_i) d\omega_i + \int_{B_o} \int_{\Omega} f_v(x_o, \omega_o) R_d(x_i, x_o) L(x_i, \omega_i) f(\omega_i) d\omega_i dx_i. \quad (4)$$

Based on this equation, we propose a model for quasi-homogeneous materials that consists of four components: the local reflectance function $R(x_o, \omega_i, \omega_o)$, the mesostructure exiting function $f_v(x_o, \omega_o)$, the mesostructure entrance function $f(\omega_i)$, and the global dipole term $R_d(x_i, x_o)$. As we shall see, all of these components can be measured from real materials.

4 Measurement System

In measuring a material model, a laser beam could be used to illuminate a material sample, but its highly concentrated illumination results in large intensity differences between specular radiance and global scattering radiance. An extensive dynamic range presents an impediment to accurate radiance measurement, as noted in [Goe-sele et al. 2004]. To diminish this problem, we separately capture global and local scattering using different types of lighting. Since global scattering does not contain specular reflections and exhibits an exponential falloff with increasing scattering distance, we employ a laser to heighten its visibility. Specifically, we use a laser stripe instead of a laser beam to significantly accelerate the capture process and substantially reduce the amount of image data. Although a laser stripe does not illuminate a local circular area as shown for A_i in Fig. 2, it can nevertheless be used for global scattering measurements. In acquiring local scattering, we take a BTF measurement approach and use lamps instead of lasers because their uniform illumination produces a relatively narrow dynamic range and reduces obscuration of texture details by specular reflections.

4.1 Capture Device

Device setup: Shown in Fig. 3, our capture device consists of four components— a turntable, an imaging array, a lighting array and a laser scan unit— that are each controlled by a computer. The cameras, lights and laser unit are mounted on two concentric arcs

centered around the turntable, which can be rotated by a stepping motor. On the flat surface of the turntable is placed the material sample to be measured.

The imaging array consists of eight Dragonfly cameras evenly mounted on a stationary arc. The bottom camera, as well as the bottom lamp, are not used in the capture process, because the top surface of a thick sample is occluded from them. From each of the 24-bit color cameras, images of resolution 1024×768 are acquired at 3.75 Hz using different shutter speeds and the maximum aperture setting to create high dynamic range (HDR) images.

The lighting array is mounted on the right arc and consists of eight halogen lamps, which provides nearly directional light for samples on the turntable. The light arc is driven by another stepping motor that rotates it around the turntable center.

The laser scan unit contains one of three interchangeable 10mw lasers (red:650nm, green:532nm, blue:473nm) with a line lens attached to produce a laser stripe. In our current implementation, only one laser at a time can be attached to the laser control unit, so image measurement is iterated three times for the red, green and blue lasers. The laser scan unit can be manually fixed at any position on the light arc, and a controller adjusts the orientation of the laser to sweep a stripe over the surface of the sample.

Device calibration: For geometric calibration, the origin of our coordinate system is set at the center of the turntable, with the XY plane aligned with the flat surface. Before capture, we calibrate the position of the light sources manually as in [Bouquet and Perona 1998]. The intrinsic and extrinsic parameters of the cameras are calibrated using the method described in [Zhang 1999].

For photometric calibration, we first determine the response curves of the cameras using the method in [Mitsunaga and Nayar 1999]. Then for relative color calibration of the cameras, we place a standard color pattern on the turntable and capture images from each camera with light from a given lamp. To calibrate lamp illumination colors, we conversely capture an image from a given camera for lamp.

To calibrate the laser stripes, we compute the angle between the laser plane, defined as the plane containing the laser stripe and laser emission point, and the XY plane. This is done by projecting a laser stripe onto two planes of different heights and calculating the offset of the laser stripe in the two planes from two recorded laser scan images. From this offset, we can derive the angle between the laser plane and the XY plane. Since the distance of the laser emitter from the turntable is much larger than the material sample diameter, we regard all the laser planes in one scan as being parallel.

4.2 Material Model Acquisition

Image data is measured using the laser stripe and lamp illumination in the process outlined in Fig. 4. From the acquired data set, appearance components that correspond to the elements of our material model are separated to enable inverse rendering of dipole diffusion parameters and measurement of the texture functions. In the laser stripe images, global subsurface scattering is clearly revealed, but its appearance is influenced by mesostructure entrance and exiting functions. To decouple these reflectance components, our technique takes advantage of their distinct directional dependencies: the mesostructure entrance function depending on light direction, the mesostructure exiting function on view direction, and dipole diffusion being directionally independent. Subtracting the global components from the halogen lamp images leaves the local reflectance function, which consists of local non-homogeneous subsurface scattering and mesostructure surface reflections.

In the image capture procedure, all laser images are rectified onto the reference plane of the material model, defined as the top surface of the sample, which is taken to be globally flat but can contain local mesostructures. In the rectification process, pixels on the laser

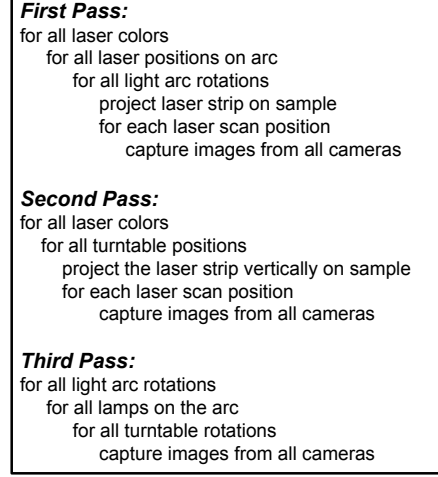


Figure 4: Pseudocode for image acquisition in the capture system.

stripe are identified according to a prescribed intensity threshold, and to these pixels a line is fit. We do not use HDR images in the rectification process. The obtained set of rectified images $I(p, v, l)$ record for incoming flux from laser stripe l the measured radiance in direction v from 3D location $p(x, y)$ on the reference plane, where p lies within a user-specified distance interval from l such that radiance from p results from global scattering and has an intensity that is sufficiently above the camera noise level.

The measured material sample is assumed to be optically thick such that a negligible amount of light passes through the bottom surface. As generally done in BTF capture systems, only a central portion of the sample is measured to avoid edge effects such as light leakage through the sides.

Mesostructure Entrance Function: In the first step of the measurement process, we scan the sample with laser stripes emitted from different light directions. As described in the first pass of Fig. 4, we mount the laser at several positions along the light arc. For each position, we rotate the light arc to sampled angles and capture the laser scan images at sampled viewing directions.

For a laser stripe projected from the vertical direction, we set its mesostructure entrance function value to one. From other incident light directions, the mesostructure entrance function is computed in relation to the vertical direction as

$$f(\omega_i) = \frac{\sum_v \sum_p \sum_{l \in \omega_i} I(p, v, l)}{\sum_v \sum_p \sum_{l \in \omega_0} I(p, v, l)}$$

where the numerator aggregates images with laser lines from direction ω_i , and the denominator sums images with laser lines from the vertical direction ω_0 . Although images from a single view are sufficient to solve for this function, we utilize several sparse viewing samples (4×4 in our implementation) for greater robustness.

Dipole Diffusion Parameters: As described in the second pass of Fig. 4, laser scan images from different viewing directions and under vertically projected illumination are captured to compute the dipole diffusion parameters. After we rotate the turntable to its desired sampling position, the light arc is rotated accordingly so that the sample is always scanned along the same direction. The appearance of the sample is captured by all cameras simultaneously.

To fit the global dipole model to this data set, we utilize the approximation given in [Jensen et al. 2001] which sets the refraction index to 1.3 in dipole fitting. Although this value may not be exact, it also works well for the materials we captured. The global dipole term can then be determined from two scattering properties: the

reduced albedo α' and reduced extinction coefficient σ'_t . Directly fitting these two parameters to measured data has proven to be ill-conditioned [Jensen and Buhler 2002]. Instead, we first measure the total diffuse reflectance R of the material sample and derive the reduced albedo α' . Then we fit the reduced extinction coefficient σ'_t from the captured images.

In our case, the total diffuse reflectance is the ratio between the outgoing flux of the sample surface and incoming flux of a vertical laser stripe, expressed for stripe l from direction ω_0 as

$$\sum_p \sum_v I(p, v, l) = R \frac{k_1}{k_0} \phi$$

where k_1 is the laser stripe length on the sample, ϕ is the laser flux measured from a reference surface of known reflectance with laser stripe length k_0 . Also we assume the intensity of the laser strip to be constant along the line. The effect of the mesostructure exiting function is effectively removed by summing corresponding image pixels from the different viewpoints. For a more robust estimate, we capture multiple laser stripe projections at different positions and then average the resulting diffuse reflectances. With the total diffuse reflectance, we solve for the reduced albedo α' of the material as in [Jensen and Buhler 2002].

With the derived reduced albedo α' , the material's reduced extinction coefficient σ'_t can be solved by fitting the dipole model in [Jensen et al. 2001] to the measured data. In fitting the dipole model, we sum the same image set over the viewing directions to eliminate mesostructure exiting function effects from the measured data:

$$I(p, l) = \sum_v I(p, v, l) = \sum_{p' \in l} R_d(p, p') \frac{\phi}{k_0}. \quad (5)$$

We additionally integrate the outgoing radiance of all points p_d that are a distance d from the laser line to obtain the function:

$$I(d) = \sum_{p_d} \sum_{p' \in l} R_d(p_d, p') \frac{\phi}{k_0}. \quad (6)$$

Since the double summation in Eq. (6) has no closed-form solution, we take the approach of [Gardner et al. 2003] and create a virtual version of the laser stripe with the same incoming flux $\frac{k_1 \phi}{k_0}$ and render $I_D(d)$ from it using Eq. (6) and σ'_t values sampled from 0.0 to 10.0, which covers most of the materials in which we are interested. Then we take the σ'_t value whose solution $I_D(d)$ to Eq. (6) most closely matches the measured $I(d)$ according to $\|I_D(d) - I(d)\|^2$. To efficiently determine the σ'_t in our current implementation, a coarse-to-fine search is employed. The effectiveness of this approach is demonstrated in Fig. 5, where the blue points indicate $I(d)$ values measured from a real sponge, and the red curve represents $I_D(d)$ values computed with the dipole approximation from the recovered values of α' and σ'_t .

Mesostructure Exiting Function: From the image set captured for estimating the dipole parameters, the mesostructure exiting function can be solved using the dipole model in Eq. (5) as

$$f_v(p, \omega_0) = \sum_l I(p, v, l) / \sum_l I(p, l),$$

which essentially represents the image residual after factoring out global dipole diffusion. We note that since multiple scattering in global light transport is directionally independent, Eq. (3) can in practice be determined independently of incident light direction.

Local Reflectance Function: In the last step, we capture halogen lamp images of the material sample with different lighting directions and viewpoints, as outlined in the third pass of Fig. 4. The




Figure 5: Comparison of dipole diffusion intensities measured from sponge data (blue), and values computed from recovered scattering values (red). The green bar represents the user-specified range used to estimate the dipole diffusion parameters.

acquired images contain both local and global lighting effects. To obtain an LRF that represents only local scattering, we compute the global component from the determined dipole parameters and mesostructure entrance/exiting functions, then subtract this global component from the captured data:

$$R(x_o, \omega_i, \omega_o) = I(x_o, \omega_i, \omega_o) - \sum_{x_i \neq x_o} f_v(x_o, \omega_o) R_d(x_i, x_o) L(x_i, \omega_i) f(\omega_i).$$

Although the dipole diffusion model is not intended to model light transport through inhomogeneous local volumes, we allow some portion of the local scattering to be represented by it, instead of being fully modelled in the LRF. This avoids the need to specify the extent of area B_o in Eq. (4). Subtracting the dipole component in the local area may lead to negative values, but these are not errors and are used to obtain the correct rendering result. For any negative values in the final rendering result due to noise or measurement errors, we simply truncate them to zero before tone mapping. Since the sampled laser and lamp positions are not identical due to the structure of our device, mesostructure entrance function values at the lamp directions are bilinearly interpolated from the four nearest laser directions.

5 Rendering

As a material model, the quasi-homogeneous material acquired from real samples can be easily applied to arbitrary objects. For this purpose, we reorganize the captured data as a 2D texture $T(x, y)$ defined on the reference plane. Each texel $T(x_i, y_i)$ consists of the reflectance function values and mesostructure exiting function at (x_i, y_i) . Similar to the method described in [Liu et al. 2004], we synthesize the texture $T(x, y)$ onto the target surface mesh while maintaining consistency in surface appearance under different lighting and viewing directions. The scattering properties are directly assigned to the object for simulating global light transport within its volume.

Quasi-homogeneous objects formed from this material model are straightforward to render. To evaluate the radiance of a surface point x_o under a given lighting condition, we first attenuate the incoming radiance on object surface according to the mesostructure entrance function. The diffuse light transport beneath the surface is then evaluated using the dipole approximation and modified by the mesostructure exiting function texture mapped onto the surface. For the local scattering effects, we directly integrate the local reflectance function on x_o over the incoming radiance at x_o . As a result, both appearance variations over the object surface and the scattering effects inside the object volume are evaluated with our material model.




Figure 6: Acquired physical samples. Top row: Lamp images; Middle row: Mesostructure exiting functions; Bottom row: Local reflectance functions.

To render the LRF, existing hardware-accelerated BTF rendering methods (e.g., [Sloan et al. 2003; Suykens et al. 2003; Mueller et al. 2003; Liu et al. 2004]) can be utilized, while an efficient dipole algorithm such as in [Jensen and Buhler 2002; Lensch et al. 2003] can be used to compute global light transport. With these methods, the local and global contributions can be rapidly computed and added together to give the rendering result.

6 Experimental Results

We implemented the measurement system described in Section 4 and captured several material samples. System specifications in capturing the texture function components of a material model are listed in Table 1. Data processing is performed on a PC with a 3.2GHZ Intel Xeon CPU and 4GB memory. Measurement of the dipole term occurs in about one minute, and it takes about four hours for image rectification and HDR image construction.

Fig. 6(a) shows an image of all the acquired samples illuminated by a halogen lamp, without color calibration. Figs. 6(b) and 6(c) show their corresponding mesostructure exiting functions and local reflectance functions under specific viewing and lighting directions. We note that the mesostructure exiting function exhibits spatial variations that correspond to the local reflectance function. The scattering properties of the captured materials are shown in Table 2.

To validate our material model, we created a synthetic volume shown in Fig. 7 composed of two materials with the following scattering properties in the RGB color bands: $\alpha' = (0.4248, 0.4248, 0.99)$ and $\sigma'_t = (1.0, 1.5, 2.13)mm^{-1}$ for one material, and $\alpha' = (0.98, 0.98, 0.98)$ and $\sigma'_t = (1.5, 1.5, 1.5)mm^{-1}$ for the other. From rendered images of this synthetic volume, we simulated the measurement process and recovered the following dipole diffusion parameters: $\alpha' = (0.8224, 0.8125, 0.9853)$ and $\sigma'_t = (1.1701, 1.29605, 1.53456)mm^{-1}$. With the acquired material model, we generated its appearance under a new lighting condition as shown in Fig. 7. Slight blurring exists in our rendered result due to interpolation of the mesostructure exiting function, but our model nevertheless provides a good approximation to direct rendering from the original synthetic volume.

For real materials whose images contain noise and registration errors, the quasi-homogeneous material model can also be acquired and rendered with accurate results, as shown for a sponge in Fig. 8. Fig. 5 demonstrates the measurement performance for this example.

In addition, our system was used to recover the scattering properties of whole milk, for which it obtained measured values that are in close agreement with those reported in [Jensen et al. 2001].

We integrated the quasi-homogeneous material model into a ray tracer, where the global light transport effect is computed by the dipole approximation as in [Jensen and Buhler 2002]. The rendering results of the different models under local and global illuminations are shown in Figs. 9, 10, 11, and 12.

In Fig. 9, we compare our rendering of an acquired sponge block model to renderings with BTF data and with a combination of BTF data and a dipole diffusion model. A high dynamic range environment map is used for illumination. We note that the appearance of the sponge under backlighting is convincingly generated by the quasi-homogeneous material model but cannot be accurately produced with a BTF. The sponge material model is also used in rendering the bunny in Fig. 10 under different viewing and lighting directions.

Fig. 11 displays rendering results of a teapot modelled with rice-cake material under different lighting and viewing conditions. Note that both the global light transport effects, which are exhibited on the boundary of the teapot body and in the shadowed area, and local lighting effects are well captured and rendered.

Fig. 12 shows the rendering result of bread slices under different lighting and viewing directions. The bread is modelled with the acquired material model, while the crust of the bread and the plate are shaded with a conventional BRDF model and color texture maps.

7 Discussion and Conclusion

The basis of our quasi-homogeneous material model is its difference in subsurface scattering properties at local and global scales. The complexity of local scattering due to material heterogeneity leads to a couple of assumed conditions in rendering an object made with the material. One is that the scale of mesh geometry variations exceeds the scale of local material heterogeneity, such that all light transport effects due to object shape are approximately modelled by dipole diffusion. Another rendering condition is illumination uniformity in a local area, which arises from our BTF representation of local reflectance functions. For illumination to be locally uniform, the distance of illuminants from the object must be much larger than the radius of a local heterogeneous area. With an assumption of locally uniform illumination, sharp changes in incident illumination will cause some inaccuracy in our result. This problem could be addressed by prefiltering irradiance on the surface, or measuring point response functions within local areas.

An assumed condition in the capture process is that the material sample is optically thick enough so that a negligible amount of light exits through the bottom of the sample. In our current implementation, this transmitted light is absorbed with a black cloth beneath the sample and is ignored. Since transmitted light is not accounted for in our measurement system, it may lead to some inaccuracies in our acquired material model.

In future work, we plan to investigate techniques for accelerating our rendering method to obtain real-time performance. Several effective compression techniques have previously been presented for texture data [Chen et al. 2002; Mueller et al. 2003; Liu et al. 2004], and we also intend to examine how our material model data can be compressed to facilitate its use.

Acknowledgements

The authors would like to thank Yanyun Chen and Kyros Kutulakos for helpful discussions, and Kun Xu for his help in capturing and processing data. The bread slice was modelled by Mindong Xie. The authors are grateful to the anonymous reviewers for their helpful suggestions and comments.

Table 1: Capture System Specifications

	Mesostructure entrance function	Mesostructure exiting function	Local reflectance function
view resolution	N/A	12×7	12×7
light resolution	12×7	N/A	12×7
image resolution	256×256	256×256	256×256
final data size	1KB	63MB	5.2GB
capturing time	3.0h	1.2h	2.0h
processing time	0.5h	0.5h	0.5h

Table 2: Captured Dipole Diffusion Parameters

	Reduced albedo α'	Reduced extinction σ'_t [mm^{-1}]
Sponge	(0.9985, 0.9971, 0.7577)	(1.6397, 1.5927, 1.3894)
Bread	(0.9805, 0.9569, 0.9179)	(0.9175, 0.9355, 0.8961)
RiceCake	(0.9986, 0.9766, 0.8400)	(0.8554, 0.8394, 0.7943)

References

- BOUGUET, J.-Y., AND PERONA, P. 1998. 3d photography on your desk. In *Proc. Int. Conf. on Computer Vision*, 43–50.
- CHEN, W.-C., BOUGUET, J.-Y., CHU, M. H., AND GRZESZCZUK, R. 2002. Light field mapping: Efficient representation and hardware rendering of surface light fields. *ACM Trans. on Graphics* 21, 3, 447–456.
- CHEN, Y., TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2004. Shell texture functions. *ACM Trans. on Graphics* 23, 3, 343–353.
- CHUANG, Y.-Y., ZONGKER, D. E., HINDORFF, J., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2000. Environment matting extensions: towards higher accuracy and real-time capture. In *Proc. SIGGRAPH 2000*, 121–130.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. on Graphics* 18, 1, 1–34.
- DANA, K. J. 2001. Brdf/btf measurement device. In *Proc. Int. Conf. on Computer Vision*, vol. 2, 460–466.
- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proc. SIGGRAPH 2000*, 145–156.
- DORSEY, J., EDELMAN, A., LEGAKIS, J., JENSEN, H. W., AND PEDERSEN, H. K. 1999. Modeling and rendering of weathered stone. In *Proc. SIGGRAPH 1999*, 225–234.
- FIUME, E. 2001. Dedication to Alain Fournier. In *Proc. SIGGRAPH '01*, 10.
- GARDNER, A., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2003. Linear light source reflectometry. *Proc. SIGGRAPH '03*, 749–758.
- GOESELE, M., LENSCH, H. P. A., LANG, J., FUCHS, C., AND SEIDEL, H.-P. 2004. Disco: acquisition of translucent objects. *ACM Trans. on Graphics* 23, 3, 835–844.
- HAN, J. Y., AND PERLIN, K. 2003. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. on Graphics* 22, 3, 741–748.
- HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proc. SIGGRAPH 1993*, 165–174.
- HAO, X., BABY, T., AND VARSHNEY, A. 2003. Interactive subsurface scattering for translucent meshes. In *ACM Symposium on Interactive 3D Graphics*, 75–82.
- JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. In *Proc. SIGGRAPH 2002*, 576–581.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proc. SIGGRAPH 2001*, 511–518.
- KOENDERINK, J. J., AND DOORN, A. J. V. 1996. Illuminance texture due to surface mesostructure. *Journal of the Optical Society of America* 13, 3, 452–463.
- KOENDERINK, J., AND VAN DOORN, A. 2001. Shading in the case of translucent objects. *Proceedings of SPIE* 4299, 312–320.
- 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* 22, 2, 195–206.




Figure 7: Ground truth comparison with a synthetic material. Left: Rendered from a captured material model; Right: Rendered from ground truth synthetic data.




Figure 8: Ground truth comparison with a real sponge. Left: Rendered from a captured material model; Right: Actual photograph.

- 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 Trans. on Visualization and Computer Graphics* 10, 3, 278–289.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. *Proc. SIGGRAPH 2001*, 519–528.
- MASSELUS, V., PEERS, P., DUTRÉ, P., AND WILLEMS, Y. D. 2003. Relighting with 4d incident light fields. *ACM Trans. on Graphics* 22, 3, 613–620.
- MATUSIK, W., PFISTER, H., NGAN, A., BEARDSLEY, P., ZIEGLER, R., AND MCMILLAN, L. 2002. Image-based 3d photography using opacity hulls. In *Proc. SIGGRAPH '02*, 427–437.
- MATUSIK, W., PFISTER, H., ZIEGLER, R., NGAN, A., AND MCMILLAN, L. 2002. Acquisition and rendering of transparent and refractive objects. In *Proc. Eurographics Workshop on Rendering*, 267–278.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. on Graphics* 22, 3, 759–769.
- MERTENS, T., KAUTZ, J., BEKAERT, P., SEIDEL, H.-P., AND REETH, F. V. 2003. Interactive rendering of translucent deformable objects. In *Proc. Eurographics Workshop on Rendering*, 130–140.
- MITSUNAGA, T., AND NAYAR, S. 1999. Radiometric self calibration. In *Proc. Computer Vision and Pattern Recognition*, 1374–1380.
- MUELLER, G., MESETH, J., AND KLEIN, R. 2003. Compression and real-time rendering of measured BTFs using local PCA. In *Proc. Vision, Modeling and Vis.*
- NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. *Geometrical Considerations and Nomenclature for Reflectance*. National Bureau of Standards (US).
- PHARR, M., AND HANRAHAN, P. M. 2000. Monte Carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proc. SIGGRAPH 2000*, 275–286.
- SLOAN, P.-P., LIU, X., SHUM, H.-Y., AND SNYDER, J. 2003. Bi-scale radiance transfer. *ACM Trans. on Graphics* 22, 3, 370–375.
- SUYKENS, F., VOM BERGE, K., LAGAE, A., AND DUTRÉ, P. 2003. Interactive rendering with bidirectional texture functions. In *Proc. Eurographics 03*.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. In *Proc. SIGGRAPH '92*, 265–272.
- WOOD, D., AZUMA, D., ALDINGER, W., CURLESS, B., DUCHAMP, T., SALESIN, D., AND STUETZLE, W. 2000. Surface light fields for 3D photography. *Proc. SIGGRAPH 2000*, 287–296.
- ZHANG, Z. 1999. Flexible camera calibration by viewing a plane from unknown orientations. In *Proc. Int. Conf. on Computer Vision*, 666–673.
- ZONGKER, D. E., WERNER, D. M., CURLESS, B., AND SALESIN, D. H. 1999. Environment matting and compositing. In *Proc. SIGGRAPH 1999*, 205–214.




Figure 9: A thin sponge with back illumination. Left: Using a quasi-homogeneous material model; Center: Using a BTF; Right: Using a combination of a BTF and a dipole diffusion model for subsurface scattering.




Figure 10: A sponge bunny under different lighting and viewing directions.



Figure 11: A teapot made of rice cake, shown under different lighting and viewing directions. Global light transport is exhibited near the silhouettes of the teapot and in the shadowed areas.



Figure 12: Slices of bread under different lighting and viewing directions. The bread is modelled with an acquired material model, while the crust of the bread and the other objects in the scene are shaded with a conventional BRDF model and color texture maps.