

A Compact Tucker-Based Factorization Model for Heterogeneous Subsurface Scattering

Murat Kurt¹, Aydın Öztürk² and Pieter Peers³

¹International Computer Institute, Ege University

²Department of Computer Engineering, Izmir University

³The College of William & Mary




Figure 1: For visual comparison on a statue under spot lighting, (a) a heterogeneous chessboard (8×8) was rendered with a full Monte Carlo path tracing algorithm (reference image); (b) and (c) were rendered using Peers et al. [PvBM*06] and our factored subsurface scattering model, respectively. (d), (e) and (f) show closeups of highlighted regions. Below each image we also report the RMSE value (lower is better) and PSNR value (higher is better). For better comparison, false-color differences were scaled by a factor of 5.

Abstract

This paper presents a novel compact factored subsurface scattering representation for optically thick, heterogeneous translucent materials. Our subsurface scattering representation is a combination of Tucker-based factorization and a linear regression method. We first apply Tucker factorization on the intensity profiles of the heterogeneous subsurface scattering responses. Next, we fit a polynomial model for characterizing the differences between the different color channels with a linear regression procedure. We show that our method achieves good compression while maintaining visual fidelity. We validate our heterogeneous subsurface scattering representation on various real-world heterogeneous translucent materials, geometries and lighting conditions.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

Translucent materials, such as wax and marble, have a unique soft appearance. Although, fast rendering solutions [JB02] and compact analytical representations [JMLH01, dI11] for homogenous translucent materials exist, no such compact analytical representation for heterogeneous translucent materials has been presented. Moreover, representing subsurface scattering effects of heterogeneous translucent materials is a challenging task, since the structural deficiencies, impurities and composite structures found inside the object volume (e.g., the veins in marble) make the acquisition and representation of heterogeneous translucent materials a non-trivial task. Acquiring heterogeneous data is a convenient way that yields realistic results [PvBM*06, STPP09]. However, due to the enormous storage requirements (gigabytes), efficient compression algorithms are necessary.

In this paper, we propose a novel factorization based subsurface scattering model for representing heterogeneous translucent materials. First, we factorize the intensity profiles of the heterogeneous subsurface scattering responses. After that, we fit a polynomial model for characterizing the differences between the different color channels with a linear regression procedure. Consequently, our subsurface scattering representation is a hybrid model that consists of a Tucker-based representation and a polynomial model.

Factorization is a popular tool for representing material appearance, such as BRDFs [LRR04, BÖK11], BTFs [VT04, WWS*05] and subsurface scattering effects of heterogeneous translucent materials [PvBM*06]. We show that our novel Tucker-based factored subsurface scattering model can represent heterogeneous translucent materials more compactly and visually more plausible. As can be seen in Figure 1, our novel material model representation can be used with any geometry, while providing a compact and visually plausible representation of heterogeneous subsurface scattering.

2. Related Work

This paper is closely related to the following two subdomains in computer graphics: representation of translucent materials and factorization of multidimensional functions.

Representation of Translucent Materials: *The diffusion dipole approximation* for homogeneous subsurface scattering was introduced to computer graphics in the seminal work by Jensen et al. [JMLH01]. Although Jensen et al.'s BSSRDF model is an analytical model and it can only represent homogenous translucent materials, many researchers have extended this model to represent other types of translucent materials and/or render translucent materials at real-time frame rates. For example, Mertens et al. [MKB*05], Donner and Jensen [DJ05], d'Eon et al. [dLE07], and Jimenez et al. [JSG09, JWSG10] have extended this model to represent

human skin. Jakob et al. [JAM*10] derived a new anisotropic dipole approximation model for representing anisotropic homogenous medium. d'Eon and Irving's [dI11] subsurface scattering representation is more accurate than Jensen et al.'s BSSRDF model, and that is valid for a larger range of scattering and absorption coefficients. However, these representations have not been designed to represent heterogeneous translucent materials accurately.

Goesele et al. [GLL*04] presented DISCO, a laser-based acquisition system and a compact model for representing heterogeneous translucent materials. However, Goesele et al.'s representation is an object model representation and it closely depends on the underlying geometry. Tong et al. [TWL*05] proposed a representation for quasi-homogenous translucent materials, i.e., translucent materials with uniformly distributed heterogeneous elements. Fuchs et al. [FGCS05] represented heterogeneous translucent materials with a linear combination of exponential fall-off functions. Song et al.'s [STPP09] *SubEdit* representation allows interactive editing and rendering of heterogeneous translucent materials, at the cost of sacrificing efficiency in compactness. However, none of these representations have been designed to represent real-world heterogeneous translucent materials accurately and efficiently.

The work of Peers et al. [PvBM*06] is most closely related to ours. In this work, Peers et al. proposed a compact representation that factorizes the remainder (i.e., heterogeneities) after division by a homogeneous approximation of the translucent material using a Non-negative Matrix Factorization (NMF)-based algorithm. Our method differs from this work in two ways. First, we employ a Tucker-based factorization to compactly represent the mean heterogeneous subsurface scattering response over the color channels per surface point. While not as efficient as the specialized NMF-based factorization of Peers et al., it provides a faster and more convenient algorithm. Second, we exploit the similarities between the different color channels, and characterize the differences between channels by fitting a polynomial model with a linear regression procedure. The resulting combined model yields a more compact and more accurate model.

Factorization: In computer graphics, factorization has been a useful tool for the compact representation of BRDFs [KM99, MAA01, LRR04, SZC*07, BÖK11], spatially varying BRDFs [LBAD*06], Bidirectional Texture Functions (BTFs) [VT04, WWS*05] and heterogeneous subsurface scattering [PvBM*06].

An example of popular factorization methods is NMF, which has been used in importance sampling of BRDFs [LRR04] and representation of heterogeneous subsurface scattering [PvBM*06]. Another popular factorization method is based on tensor decomposition [KB09]. Tensor approximations [KB09] have been used in volume simulations [WWS*05], BTF representations [VT04, WWS*05],

BRDF representations [SZC*07, BÖK11] and importance sampling of BRDFs [BÖK11]. Following the work of Bilgili et al. [BÖK11], we will use Tucker factorization based on tensor decomposition for representing measured heterogeneous subsurface scattering data compactly and accurately.

3. Background

The behavior of translucent materials is described by the Bidirectional Scattering Surface Reflectance Distribution Function (BSSRDF) [NRH*77] $S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o)$ which relates outgoing radiance $L_o(x_o, \vec{\omega}_o)$ at a point x_o in a direction $\vec{\omega}_o$ to the incoming radiance $L_i(x_i, \vec{\omega}_i)$ at a location x_i and an incoming direction $\vec{\omega}_i$:

$$L_o(x_o, \vec{\omega}_o) = \int_A \int_{\Omega^+} L_i(x_i, \vec{\omega}_i) S(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) (\vec{\omega}_i \cdot \vec{n}) d\vec{\omega}_i dx_i, \quad (1)$$

where A is the area around the point x_o , Ω^+ is the hemisphere around x_i and \vec{n} is the surface normal at point x_i . Eq. (1) can be separated into a local component and a global component. While the local component represents the light immediately reflected from a surface, the global component represents the light scattering in the material volume. As in [PvBM*06], we model the global component by the diffuse BSSRDF S_d which can be further decomposed as [DJ05]:

$$S_d(x_i, \vec{\omega}_i; x_o, \vec{\omega}_o) = \frac{1}{\pi} F_i(x_i, \vec{\omega}_i) R_d(x_i, x_o) F_o(x_o, \vec{\omega}_o), \quad (2)$$

where F_o and F_i are directionally dependent components, and $R_d(x_i, x_o)$ is a four dimensional (4D) spatial subsurface scattering component. As in Goesele et al. [GLL*04], Peers et al. [PvBM*06] and Song et al. [STPP09], we focus on accurately representing the 4D spatial component R_d of heterogeneous translucent materials, and we ignore the local component and the directional dependencies (F_o and F_i).

4. Subsurface Scattering Representation

In this section we will first introduce our new factorization and regression framework before briefly discussing the impact of the various parameters.

4.1. Pre-Processing

We represent the measured BSSRDF as a matrix $R_d(x_i, x_o)$ where x_i and x_o are incoming and outgoing surface locations. The measured 4D input data is linearized to a two-dimensional (2D) matrix. Figure 2(a) shows an example of a subsurface scattering matrix R_d of a heterogeneous material containing a light blocking vein. The effects of the light blocking vein are expressed by horizontal and vertical discontinuities in R_d .

First we reorganize the subsurface scattering matrix $R_d(x_i, x_o)$ by a change of variables $d = x_o - x_i$ (Figure 2(b)). The reparameterized subsurface scattering matrix $R'_d(x_i, d)$

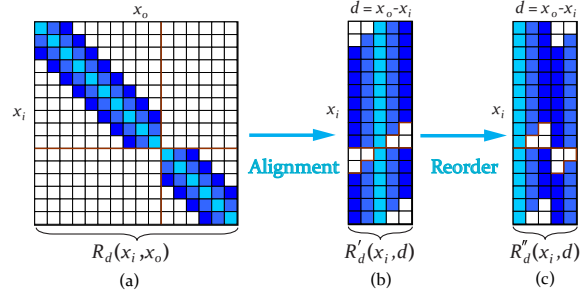


Figure 2: (a) - (b) We reformat the BSSRDF matrix (a) by first aligning the diagonal by a change of variables to $R'_d(x_i, d)$. (b) - (c) Next, we reorder the elements in each row by shifting and wrapping around each row such that the element with the maximum magnitude moves to the first position. After that, we divide each row with its maximum value, finally yielding the subsurface scattering matrix $R''_d(x_i, d)$. To reconstruct $R'_d(x_i, d)$ from $R''_d(x_i, d)$, we only need to store the coordinate of the maximum element and its corresponding value in each row.

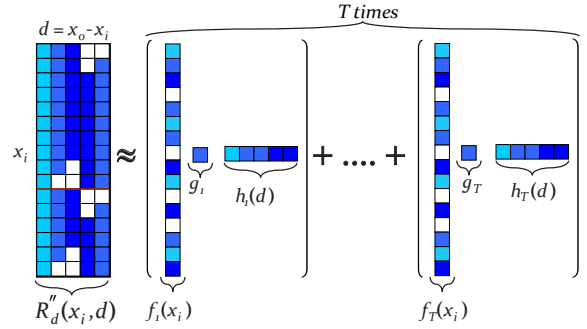


Figure 3: Similar to Bilgili et al. [BÖK11], we use an error modeling approach using Tucker factorization for representing measured subsurface scattering matrix $R''_d(x_i, d)$.

can be factorized instead of $R_d(x_i, x_o)$. To increase the effectiveness of the factorization, we apply a similar approach as described in [XYL*09]. We shift each row independently such that the maximum element in the each row will become the first element and divide each row by its maximum value (Figure 2(c)). Note that $R''_d(x_i, d)$ is better suited for efficient factorization than $R'_d(x_i, d)$, because these operations align similar values along the same columns. Another advantage of the shifting of the rows is that this also allows us to efficiently compensate for any shift in the peak that can occur due to measurement or calibration issues.

4.2. Factorization

For an efficient and compact factorization, we apply the error modeling approach using the Tucker factorization [Tuc66] on $R''_d(x_i, d)$. According to the error modeling approach, we first factorize $R''_d(x_i, d)$ matrix with Tucker factorization. After the first Tucker factorization, we get our Tucker-based model with two vectors ($f_1(x_i)$ and $h_1(d)$), and a scalar value (g_1). Then, we factorize the model errors again with Tucker

Sample Material	Physical Size (cm ²)	Resolution (pixels)	Kernel Size (pixel)	Original Size (Gb)	No of T	No of P	Factored Size (Mb)	CR	RMSE
Chessboard (4 × 4)	12.6 × 12.6	277 × 277	29 × 29	2.74	15	7	24.11	1/116	0.0102
Chessboard (8 × 8)	25.1 × 25.1	222 × 222	39 × 39	2.75	10	4	10.2	1/276	0.0242
Marble (close up)	2.6 × 2.6	128 × 128	39 × 39	6.0	15	4	4.19	1/1466	0.0072
Densely Veined Marble	13.0 × 13.0	213 × 211	29 × 29	4.92	15	4	11.0	1/458	0.0152

Table 1: Statistics of the factored heterogeneous subsurface scattering materials. The table also summarizes some statistics of our subsurface scattering model with typically selected values for model parameters T and P .

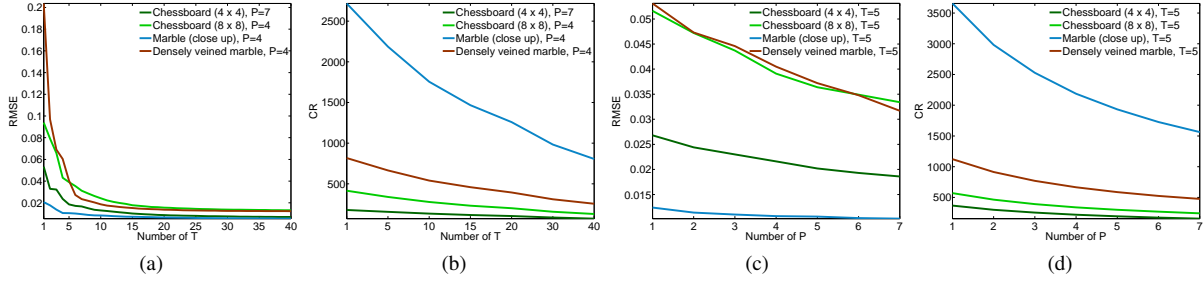


Figure 4: (a) RMSE of our Tucker-based model for various T parameters; (b) the compression ratio (CR) of our model for various T parameters; (c) RMSE of our Tucker-based model for various P parameters; (d) CR of our model for various P parameters.

factorization. This process is repeated for a predetermined number of times. As can be seen in Figure 3, the final subsurface scattering model will be the sum of estimation of model errors and the first factorization of $R'_d(x_i, d)$. Please refer to [BÖK11] for an in depth discussion on the error modeling approach. The resulting subsurface scattering model can be formalized as:

$$R'_d(x_i, d) \approx \sum_{j=1}^T g_j f_j(x_i) h_j(d), \quad (3)$$

where T is the total number of terms, g_j is the scalar core tensor, $f_j(x_i)$ and $h_j(d)$ are the univariate tensor functions, x_i incoming surface location and $d = x_o - x_i$. Since we apply Tucker factorization to 2D subsurface scattering data, the applied factorization method is similar to Singular Value Decomposition (SVD) method. However, our Tucker-based factorization algorithm can also be applied to higher dimensional data efficiently, which can be seen as an advantage over the classical SVD approach.

While, our Tucker-based factorization by itself (without the linear regression presented in Section 4.3) is not as compact as Peers et al.'s [PvBM*06] subsurface scattering representation, it is less complex and computationally more efficient as shown in Table 2.

4.3. Linear Regression

Prior work in compressing heterogeneous subsurface scattering, treated each color channel independently. However, while not identical, the scattering profiles for the different color channels are very similar. We exploit this correlation, by only applying the Tucker-based factorization to the mean values of the measured subsurface scattering values of the

three color channels, and approximate the deviations from the mean by fitting a polynomial regression model of degree P . In the linear regression procedure, we estimate the linear coefficients for each row of measured subsurface scattering matrix. Then, the corresponding models for each color channel can be written as:

$$R_{dr}(x_i, x_o) \approx \sum_{p=0}^P \beta_{rpx_i} R'_d(x_i, d)^p, \quad (4)$$

$$R_{dg}(x_i, x_o) \approx \sum_{p=0}^P \beta_{gpx_i} R'_d(x_i, d)^p, \quad (5)$$

$$R_{db}(x_i, x_o) \approx \sum_{p=0}^P \beta_{bpx_i} R'_d(x_i, d)^p, \quad (6)$$

where P is the degree of the polynomial, and β_{rpx_i} , β_{gpx_i} and β_{bpx_i} are the parameters of the model. Linear least square optimization techniques were used to fit the model to subsurface scattering data and the subsurface scattering values for the underlying color channel were estimated from the fitted model.

Our linear regression-based method exploits coherency between the color channels and provides a more compact representation without significant loss of accuracy. This approach can potentially also be applied other factorization-based compression methods.

4.4. Parameter Analysis

Our subsurface scattering representation has two parameters: T and P . T is the number of terms in the Tucker factorization, P is the degree of the polynomial approximation used in the subsurface scattering model.

The Tucker-based factorization plays an important role in

the accuracy of the factored subsurface scattering, since it is applied to the mean values. As can be seen in Figure 4(a), when the first 10 terms are used, our Tucker-based subsurface scattering model can represent the general shape of the subsurface scattering accurately. Since we apply Tucker-based factorization to the mean values, the number of Tucker terms T directly impacts the compression rate. On the other hand, our linear regression based method exploits coherency between the color channels and provides a more compact representation without significant loss of accuracy. As can be seen in Figure 4(c), errors decrease linearly as the degree of the polynomials increase.

As the Tucker factorization is only applied to the mean intensity of the color channels (i.e., size $\sim T$), while a polynomial fit is applied per color channel (i.e., size $\sim 3P$), it is more efficient to prioritize minimizing P then it is to minimize T .

5. Results

To visualize our results, we implemented a rendering scheme similar to Peers et al. [PvBM*06] in the Mitsuba rendering system [Jak13]. We verified our Tucker-based subsurface scattering model on several real-world subsurface scattering materials, ranging from fairly homogeneous to highly translucent heterogeneous materials. Table 1 gives an overview of the factored heterogeneous translucent materials, and lists a number of statistics of our subsurface scattering model based on typical values for P and T . Figure 5 shows the results of our model with $T = 15$ and $P = 4$ for marble (close up) and densely veined marble materials. These results show that our method provides high compression ratio, while maintaining visual fidelity.

We also compared our model with Peers et al.'s [PvBM*06] subsurface scattering model on a selection of heterogeneous translucent materials in Figure 1 and Figure 6. Both are rendered under a spot light with a full Monte Carlo path tracing algorithm to better illustrate the effects of the subsurface scattering. We set the parameters of our representation to $T = 10$ and $P = 4$ for chessboard (8×8) material in Figure 1, and $T = 40$ and $P = 7$ for chessboard (4×4) material in Figure 6. We also report the root-mean-square error (RMSE), the peak signal-to-noise ratio (PSNR) [Ric02] and false-color difference images of zoom-ins to better visualize the differences between the representations. All of these results show that our Tucker factorization based subsurface scattering model represents heterogeneous translucent materials more accurately at comparable data sizes.

Furthermore, we compared measured and factored subsurface responses of selected surface points in Figure 7. In Figure 7, the dashed square is approximately equals to *kernel size* and it illustrates the relative size of the responses. Furthermore, Figure 7 also compares Peers et al.'s subsurface scattering model and our subsurface scattering model at

comparable data sizes. Therefore, we set parameters of our representation to $T = 40$ and $P = 7$ for representing chessboard (4×4) material, and $T = 10$ and $P = 7$ for representing chessboard (8×8) material. In this setting, our model needs 38.9 MB and 13.6 MB storage, and it gives 0.007 and 0.0212 RMSE for representing chessboard (4×4) and chessboard (8×8) materials, respectively. On other hand, Peers et al.'s representation needs 38.7 MB and 10.8 MB storage, and it gives 0.0134 and 0.0384 RMSE for representing chessboard (4×4) and chessboard (8×8) materials, respectively. As can be seen in Figure 1, Figure 6 and Figure 7, another advantage of our representation is that the modeling errors of our representation are uniformly distributed—a preferred property for a subsurface scattering representation. Our representation also compensates for any shift in the peak that can occur due to measurement or calibration issues. Finally, our heterogeneous subsurface scattering representation is more accurate than Song et al.'s [STPP09] heterogeneous subsurface scattering representation (*SubEdit*) at comparable data sizes, since Song et al. have reported in their paper that their representation has a slightly larger relative error (2 – 7%) than Peers et al.'s [PvBM*06] representation at comparable data sizes.

Our subsurface scattering model has two parameters that can be tweaked, namely T and P . As a result, compared to Peers et al.'s [PvBM*06] subsurface scattering representation, our subsurface scattering representation is more flexible. Higher T and P values mean more fitting accuracy and more storage needs. On the other hand, our Tucker-based subsurface scattering model still reproduces visually plausible measured heterogeneous subsurface scattering materials even with lower T and P values. This illustrates the flexibility in exchanging accuracy versus size of our factored representation.

Finally, Table 2 lists a comparison of computation times for different subsurface scattering models. We selected parameters of our Tucker-based model for the chessboard (4×4) dataset ($T = 10$; $P = 7$) and the chessboard (8×8) dataset ($T = 5$; $P = 4$), so that RMSE values of the compared models are the same. All computations are performed on a dual Intel Xeon X5650 CPU @ 2.67 GHz with 48 GB RAM. These results show that our subsurface scattering representation is faster ($36 \times -88 \times$) and more compact ($1.30 \times -1.83 \times$) than Peers et al.'s [PvBM*06] representation at the same RMSE values.

6. Conclusions and Future Work

In this paper we have presented a compact and efficient factorization-based representation for the spatial component of heterogeneous subsurface scattering. Our subsurface scattering representation is composed of Tucker factorization and a linear regression procedure. We have illustrated efficiency and compactness of our factorization based subsurface scattering model on a number of real-world heteroge-




Figure 5: In all of these images, our Tucker-based subsurface scattering model was used to render different geometries under spot light illumination. We set $T = 15$, and $P = 4$ for our factored subsurface scattering model in these materials. (a) and (c) are marble (close up) material, (b) and (d) are densely veined marble material.

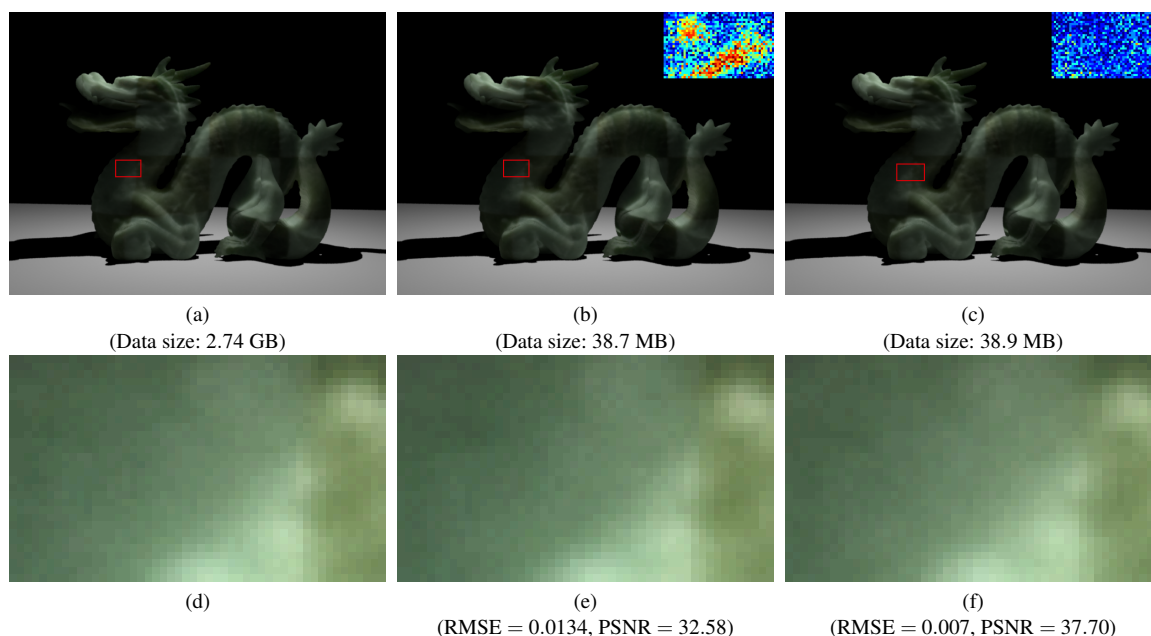


Figure 6: For visual comparison on a dragon under spot lighting, (a) a heterogeneous chessboard (4×4) was rendered with a full Monte Carlo path tracing algorithm (reference image); (b) and (c) were rendered using Peers et al. [PvBM*06] and our factored subsurface scattering model, respectively. (d), (e) and (f) show closeups of highlighted regions. Below each image we also report the RMSE value (lower is better) and PSNR value (higher is better). For better comparison, false-color differences were scaled by a factor of 20.

neous translucent materials. We have demonstrated that our compact factored representation can be applied to any geometry and it can be easily integrated into a standard global illumination rendering system, resulting in convincing images. Furthermore, we compared our subsurface scattering model with Peers et al.’s [PvBM*06] factored model and we showed that our compact subsurface scattering model can represent heterogeneous subsurface scattering effects accurately and efficiently.

In the future we are interested in investigating the use of perceptual metrics and reparametrizations for even better representations of the subsurface scattering profiles. Finally, we are interested in exploring rendering algorithms to employ our subsurface scattering representation directly in real-time applications.

Representation	Chessboard (4×4)	Chessboard (8×8)
Peers et al.	3.756	1.135
Our model	0.0427	0.0314
Representation	Chessboard (4×4)	Chessboard (8×8)
Peers et al.	38.7	10.8
Our model	21.1	8.33

Table 2: (Top) average albedo computation times (in seconds) and (bottom) data sizes (in MB) of various subsurface scattering representations at the same RMSE values.

Acknowledgements

This work was supported by the Scientific and Technical Research Council of Turkey (Project No: 111E208). Pieter Peers was partially supported by a grant from the National Science Foundation (IIS-1217765). The authors would like to thank Diego Gutierrez for his generous support to this

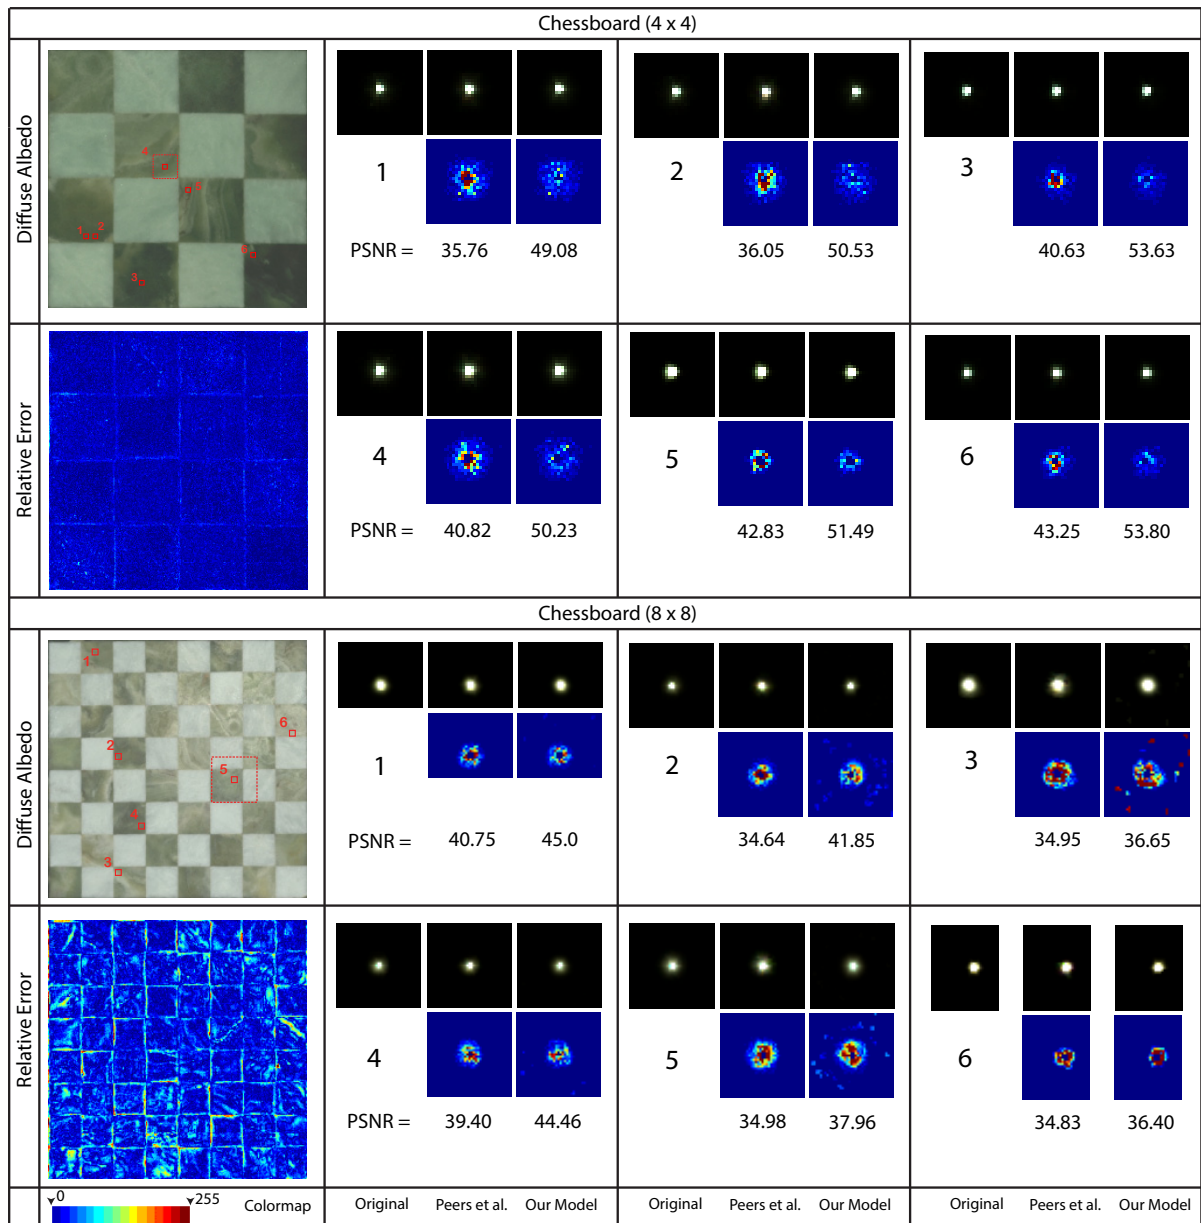


Figure 7: A comparison between Peers et al.'s factored model and our model at comparable data sizes. For each material, the diffuse albedo map, a relative error distribution plot of our representation and a selection of measured responses with the corresponding models approximations are shown. The locations of the responses are marked on the diffuse albedo map. We also computed false-color difference images between measured responses and corresponding approximations of models. For better comparison, false-color differences were scaled by a factor of 20. We also reported PSNR values (higher is better).

project. Finally, the authors would also like to thank the anonymous reviewers for their helpful suggestions and comments.

References

[BÖK11] BILGILI A., ÖZTÜRK A., KURT M.: A general brdf representation based on tensor decomposition. *Computer Graph-*

ics Forum 30, 8 (December 2011), 2427–2439. 2, 3, 4

[dI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM TOG* 30, 4 (July 2011), 56:1–56:14. (Proc. SIGGRAPH '11). 2

[DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM TOG* 24, 3 (July 2005), 1032–1039. (Proc. SIGGRAPH '05). 2, 3

[dLE07] D'EON E., LUEBKE D. P., ENDERTON E.: Efficient

- rendering of human skin. In *Proc. of Eurographics Symposium on Rendering* (2007), pp. 147–157. 2
- [FGCS05] FUCHS C., GOESELE M., CHEN T., SEIDEL H.-P.: An empirical model for heterogeneous translucent objects. In *ACM SIGGRAPH 2005 Sketches* (2005), SIGGRAPH '05. 2
- [GLL*04] GOESELE M., LENSCH H. P. A., LANG J., FUCHS C., SEIDEL H.-P.: DISCO: acquisition of translucent objects. *ACM TOG* 23, 3 (Aug. 2004), 835–844. (Proc. SIGGRAPH '04). 2, 3
- [Jak13] JAKOB W.: Mitsuba renderer, 2013. <http://www.mitsuba-renderer.org>. 5
- [JAM*10] JAKOB W., ARBREE A., MOON J. T., BALA K., MARSCHNER S.: A radiative transfer framework for rendering materials with anisotropic structure. *ACM TOG* 29, 4 (July 2010), 53:1–53:13. (Proc. SIGGRAPH '10). 2
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM TOG* 21, 3 (July 2002), 576–581. (Proc. SIGGRAPH '02). 2
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proc. SIGGRAPH '01* (2001), pp. 511–518. 2
- [JSG09] JIMENEZ J., SUNDSTEDT V., GUTIERREZ D.: Screen-space perceptual rendering of human skin. *ACM Transactions on Applied Perception* 6, 4 (Oct. 2009), 23:1–23:15. (Proc. APGV '09). 2
- [JWSG10] JIMENEZ J., WHELAN D., SUNDSTEDT V., GUTIERREZ D.: Real-time realistic skin translucency. *IEEE Computer Graphics and Applications* 30, 4 (2010), 32–41. 2
- [KB09] KOLDA T. G., BADER B. W.: Tensor decompositions and applications. *SIAM Review* 51, 3 (Sept. 2009), 455–500. 2
- [KM99] KAUTZ J., MCCOOL M. D.: Interactive rendering with arbitrary brdfs using separable approximations. In *Proc. of Eurographics Workshop on Rendering* (Granada, Spain, 1999), pp. 247–260. 2
- [LBAD*06] LAWRENCE J., BEN-ARTZI A., DECORO C., MATUSIK W., PFISTER H., RAMAMOORTHY R., RUSINKIEWICZ S.: Inverse shade trees for non-parametric material representation and editing. *ACM TOG* 25, 3 (July 2006), 735–745. (Proc. SIGGRAPH '06). 2
- [LRR04] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHY R.: Efficient BRDF importance sampling using a factored representation. *ACM TOG* 23, 3 (2004), 496–505. (Proc. SIGGRAPH '04). 2
- [MAA01] MCCOOL M. D., ANG J., AHMAD A.: Homomorphic factorization of brdfs for high-performance rendering. In *Proc. SIGGRAPH '01* (2001), ACM, pp. 171–178. 2
- [MKB*05] MERTENS T., KAUTZ J., BEKAERT P., REETH F. V., SEIDEL H.-P.: Efficient rendering of local subsurface scattering. *Computer Graphics Forum* 24, 1 (2005), 41–49. 2
- [NRH*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERS T.: *Geometrical Considerations and Nomenclature for Reflectance*. Monograph, National Bureau of Standards (US), Oct. 1977. 3
- [PvBM*06] PEERS P., VOM BERGE K., MATUSIK W., RAMAMOORTHY R., LAWRENCE J., RUSINKIEWICZ S., DUTRÉ P.: A compact factored representation of heterogeneous subsurface scattering. *ACM TOG* 25, 3 (July 2006), 746–753. (Proc. SIGGRAPH '06). 1, 2, 3, 4, 5, 6
- [Ric02] RICHARDSON I. E.: *Video Codec Design: Developing Image and Video Compression Systems*. John Wiley & Sons, Inc., New York, NY, USA, 2002. 5
- [STPP09] SONG Y., TONG X., PELLACINI F., PEERS P.: Subedit: a representation for editing measured heterogeneous subsurface scattering. *ACM TOG* 28, 3 (July 2009), 31:1–31:10. (Proc. SIGGRAPH '09). 2, 3, 5
- [SZC*07] SUN X., ZHOU K., CHEN Y., LIN S., SHI J., GUO B.: Interactive relighting with dynamic brdfs. *ACM TOG* 26, 3 (2007), 27:1–27:10. (Proc. SIGGRAPH '07). 2, 3
- [Tuc66] TUCKER L. R.: Some mathematical notes on three-mode factor analysis. *Psychometrika* 31, 3 (Sept. 1966), 279–311. 3
- [TWL*05] TONG X., WANG J., LIN S., GUO B., SHUM H.-Y.: Modeling and rendering of quasi-homogeneous materials. *ACM TOG* 24, 3 (July 2005), 1054–1061. (Proc. SIGGRAPH '05). 2
- [VT04] VASILESCU M. A. O., TERZOPOULOS D.: Tensor-textures: multilinear image-based rendering. *ACM TOG* 23, 3 (2004), 336–342. (Proc. SIGGRAPH '04). 2
- [WWS*05] WANG H., WU Q., SHI L., YU Y., AHUJA N.: Out-of-core tensor approximation of multi-dimensional matrices of visual data. *ACM TOG* 24, 3 (2005), 527–535. (Proc. SIGGRAPH '05). 2
- [XYL*09] XU D., YAN S., LIN S., HUANG T. S., CHANG S.-F.: Enhancing bilinear subspace learning by element rearrangement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31, 10 (2009), 1913–1920. 3