Workshop on Material Appearance Modeling (2018)
H. Rushmeier and R. Klein (Editors)

Towards a Principled Kernel Prediction for Spatially Varying BSSRDFs

Oskar Elek and Jaroslav Křivánek

Charles University, Prague

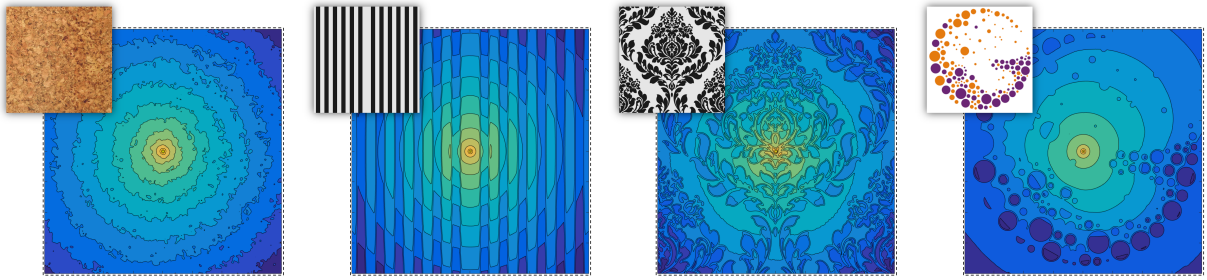

Figure 1: *Predictions of the spatially varying sub-surface scattering kernels (false-color logarithmic maps) for heterogeneous materials (defined by 2.5D scattering albedo textures shown in the respective thumbnails). The proposed framework consists of two key steps: 1) a principled statistical aggregation of the material properties, and 2) a suitable decomposition of the sub-surface scattering kernel in order to preserve both the local and global illumination features.*

Abstract

While the modeling of sub-surface translucency using homogeneous BSSRDFs is an established industry standard, applying the same approach to heterogeneous materials is predominantly heuristical. We propose a more principled methodology for obtaining and evaluating a spatially varying BSSRDF, on the basis of the volumetric sub-surface structure of the simulated material. The key ideas enabling this are a simulation-data driven kernel for aggregating the spatially varying material parameters, and a structure-preserving decomposition of the sub-surface transport into a local and a global component. Our current results show significantly improved accuracy for planar materials with spatially varying scattering albedo, with added discussion about extending the approach for general geometries and full heterogeneity of the material parameters.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—; I.3.3 [Computer Graphics]: Picture/Image Generation—;

1. Introduction

Numerous volumetric materials are conveniently modeled using the notion of diffusive surface-to-surface transport, i.e., via a *bi-directional scattering-surface reflectance distribution function* (BSSRDF) [NRH*77]. In spite of the fact that the visual industry is steadily moving towards using path-based prediction methods [Jar13,FWKH17], BSSRDFs remain a relatively low-cost and controllable representation for a broad class of translucent materials with well defined boundary, especially solids and some fluids.

There is however a subset of such materials that defy this convenient description: those with heterogeneous sub-surface structure, such as stone [PvBM*06], layered and organic materials [DJ05, DWd*08], or those used in 3D fabrication [ESZ*17]. In these cases, a spatially varying model is needed, an SV-BSSRDF – the measurement and efficient representation of which is challenging in its own right [PvBM*06, SW13]. Although related to this line of work (especially Peers et al. [PvBM*06] and their handling of high-frequency material features), this paper addresses a complementary problem: deriving, or predicting, the sub-surface scattering kernel (i.e., the SV-BSSRDF) from the optical properties of the material *directly*.

Designing an accurate BSSRDF for homogeneous materials from the first principles is already a complex task in itself, as documented by recent works [JMLH01, DLR*09, dI11, FHK14, FD17]. Formulating and applying an SV-BSSRDF carries at least two additional challenges on top of that:

- the application of a heterogeneous sub-surface scattering kernel *cannot* be mathematically described as a convolution with the incident illumination (implying that image-based methods will not generally work well); and
- the SV-BSSRDF kernel is generally defined in terms of the *global parameter distribution* under the material surface, rather than just the *local* material properties at the interaction sites.

The existing solutions to these mostly rely on empirical heuristics: the assumption of smoothly varying transport [JMLH01, DJ05], composition of factorized 'local' kernels [STPP09, SW13], or stochastic averaging of the material parameters [SHK17]. Although these approaches can be applied in specific conditions, they share the following two limitations: lack of a principled way to aggregate the material parameters, and over-blurring of local high-frequency features. Our aim is to address exactly these two points.

2. Methodology

The approach to deriving and applying the SV-BSSRDF proposed here is conceptually most related to the recent work of Sone et al. [SHK17], and proceeds in six steps (the first two of which constitute pre-processing):

1. For each pair of incident and exitant points on the surface, estimate an *unweighted, medium-independent* distribution of scattered sub-surface paths.
2. Fit a parametric model to the path distributions, which can then be efficiently evaluated on-the-fly.
3. Aggregate the material properties using the fitted model from Step 2.
4. Separate the sub-surface transport into a local and a global component.
5. Use point-estimated material properties to evaluate the local component, and the aggregate from Step 3 for the global component.
6. Use standard Monte-Carlo integration to solve the transport for every surface point of interest.

While the above scheme is general, we adopt a set of assumptions to demonstrate the validity of this approach. We rely on the standard medium parametrization (optical density σ_t, single-scattering albedo α, scattering anisotropy factor g and refractive index η), and assume that only the scattering albedo varies spatially. We also define the albedo by a 2.5D texture (i.e., assume its invariance w.r.t. depth), and assume a flat infinite half-space geometry illuminated by directionally uniform lighting. These assumptions are common [SHK17] and we discuss the possibility of relaxing most of them in Sec. 4. What follows is a detailed description of the currently implemented key steps of our approach.

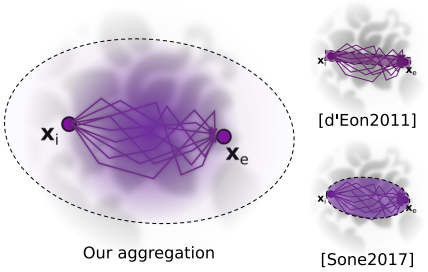

Figure 2: *Our framework regards the material aggregation kernel as a smooth 'footprint' of the sub-surface path distribution. This contrasts with the existing heuristics which perform the aggregation within an empirically bounded region, such as a line [dI11] or ellipse [SHK17].*

Basis BSSRDF Model. Instead of building on a principled BSSRDF (e.g., one of the models reviewed in Sec. 1), we opt for a data-driven approach akin to Christensen and Burley [Chr15]. This is to eliminate any inaccuracies that typically arise in deriving such solutions from first principles. Similarly to Christensen and Burley we perform a brute-force analog MC simulation to obtain ground truth sub-surface scattering profiles, fixing the scattering anisotropy g and refractive index η to constant values (specifically $g = 0.3$ and $\eta = 1.5$, corresponding to the photo-polymer materials we mainly aim to model, see Elek, Sumin [ESZ*17]). To the resulting tabulated profiles (parametrized only by the scattering albedo α and normalized distance d at this point) we fit an exponential mixture model, using non-linear fitting with the Simulated Annealing method.

We found that the double-exponential model proposed by Christensen and Burley [Chr15] is only accurate for *index-matched* materials (i.e., with refractive index identical with the ambient one). Consequently, we use a mixture of five exponential functions that better captures the more complex shapes of the sub-surface scattering kernels that are characteristic in the presence of total internal reflection off the material interface. In addition, the long-tailed BSSRDF kernels that arise in our materials are typically rather poorly fitted using the standard L_p norms, with actual impact on the resulting prediction accuracy. We therefore opted for a relative (ratio) objective function, which yields accurate fits visually indistinguishable from the reference MC simulations.

Aggregation Kernel Estimate. We follow a similar rationale as Sone et al. [SHK17]: the *sub-surface distribution* of light paths between the incident / exitant points should serve as a template for the statistical aggregation of the material parameters. Sone et al. propose to aggregate (average) the material parameters in the vicinity of these points. However, rather than using their empirical averaging within a sharply bounded region estimated by the path integral framework [PARN04], we again opted for a data-driven approach.

We chose to represent the parameter aggregation kernel

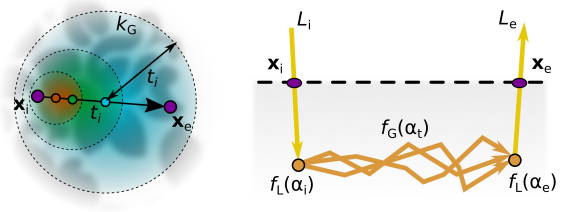

Figure 3: ***Left:*** *Aggregation of the material parameters, using three color-coded Gaussians (see Eq. 1).* ***Right:*** *Decomposition of the transport kernel into local and global components f_L and f_G (see Eq. 2).*

k (see Sone et al. [SHK17], Eq. 3) for the incident / exitant point pair $\mathbf{x}_i$ and $\mathbf{x}_e$ by a Gaussian mixture model:

$$k(\mathbf{x}_i, \mathbf{x}_e) = \sum_{i=1}^{N} w_i \cdot k_G(\mu_i, \sigma_i^2), \tag{1}$$

where k_G is the isotropic 2D Gaussian, w_i normalized non-negative weights, $\mu_i \in \mathbb{R}^2$ and $\sigma_i^2 \in \mathbb{R}^+$ the respective means and variances. The particular choice of basis is less important here than its ability to represent smooth, continuous distributions of paths under the material surface (see Fig. 2). For simplicity, we constrained the Gaussian components' means μ_i to lie along the line connecting $\mathbf{x}_i$ and $\mathbf{x}_e$, and to have the same width (standard deviation) as the mean's distance from $\mathbf{x}_i$. This is sketched in Fig. 3 left. In all our experiments we use $N = 3$ and $w = t = \{4, 2, 1\}/7$, determined by manually fitting to the path distributions yielded by the analog MC simulation previously used for the basis BSSRDF fitting.

The parameter aggregation using the kernel from Eq. 1 is efficiently implemented as N weighted lookups into a Gaussian pyramid built upfront over the 2D albedo texture. This makes the aggregation a constant-time operation (as N is fixed), avoiding costly integration over the elliptical aggregation region as proposed by Sone et al. [SHK17].

Transport Separation. The previous step yields an aggregate estimate of the material parameters within a region of size proportional to the assumed scale of sub-surface transport (see Fig. 3 left). Intuitively, this corresponds to finding an expected local material composition. The problematic part of using this aggregate to directly estimate point-to-point transport (as proposed and evaluated by Sone et al. [SHK17]) is that such an approach neglects local interactions near the entry and exit points $\mathbf{x}_i$ / $\mathbf{x}_e$.

As shown in Fig. 3 right, we instead propose to separate the sub-surface transport into a local and a global interaction components. This corresponds to a separate handling of the first and last vertices of the sub-surface paths, a proposition that seems justified since these specific interactions will be spatially well correlated in translucent materials. The transport between the interaction sites $\mathbf{x}_i$ and $\mathbf{x}_e$ (i.e., the SV-BSSRDF kernel S_V) can then be defined as a product of two local components f_L and one global component f_G:

$$\begin{aligned} S_V(\mathbf{x}_i, \mathbf{x}_e) &= f_L(\mathbf{x}_i) \cdot f_G(\mathbf{x}_i, \mathbf{x}_e) \cdot f_L(\mathbf{x}_e) \\ &= \frac{\alpha_i}{\alpha_t} \cdot S(\alpha_t, |\mathbf{x}_e - \mathbf{x}_i|) \cdot \frac{\alpha_e}{\alpha_t}, \end{aligned} \tag{2}$$

where S is the basis (homogeneous) BSSRDF. Since we currently consider only the albedo variation, the components in Eq. 2 are parametrized only by the local albedos α_i / α_e and the global albedo α_t aggregated using Eq. 1.

The definition of the local components f_L introduces the notion of a *relative albedo*, that is, w.r.t. the aggregate. This is required for proper normalization of the SV-BSSRDF: indeed, for a locally homogeneous medium ($\alpha_i, \alpha_e \approx \alpha_t$) Eq. 2 simplifies to the basis BSSRDF S.

Evaluation and Sampling. Computing the sub-surface transport using the SV-BSSRDF defined in Eq. 2 amounts to a standard Monte-Carlo integration: for every surface point of interest $\mathbf{x}_e$ we sample a set of points $\mathbf{x}_i$ and the incident radiance therein, and for each compute the aggregate parameters (α_t) and evaluate the transport kernel (Eq. 2).

3. Results

We evaluate our composite SV-BSSRDF for a variety of configurations in Fig. 4: two didactic scenes (STRIPES1 and STRIPES2), two mono-chromatic scenes with multi-scale details (FLORAL and ESCHER) and two scenes to demonstrate color bleeding (CORK and LOGO). Each scene is a flat medium defined by a 400^2 albedo texture viewed from the top, 10 mean free paths (MFP) wide, and illuminated in the center by a circular light with the radius of 1 MFP. With the exception of the path-traced references, all images were generated using 1024 SV-BSSRDF samples per pixel.

We compare against two different approaches: the factorized model of Song et al. [STPP09] and the aggregate model of Sone et al. [SHK17]. The former method performs geometric averaging two homogeneous BSSRDF kernels, one for each of $\mathbf{x}_i$ / $\mathbf{x}_e$, based on purely local estimates of the medium parameters. The latter approach corresponds to evaluating only the global term f_G of Eq. 2, with the parameters aggregated within a bounded elliptical region. It should be noted that while the factorized model has been designed for compression rather than direct derivation of the SV-BSSRDF kernel, in our case it can be used as a conceptual stand-in for other local methods [JMLH01, dI11].

The results generally follow intuition to a good degree. The factorized model tends to overestimate the transport (e.g., in STRIPES2 or ESCHER), since it lacks the information about the material distribution between the entry and exit points; for the same reason it cannot reproduce the characteristic translucent color bleeding (LOGO). The aggregate model overcomes these issues, but suffers from strong over-blurring of local details across the board, since it lacks a specific mechanism to preserve them. As shown in Fig. 5

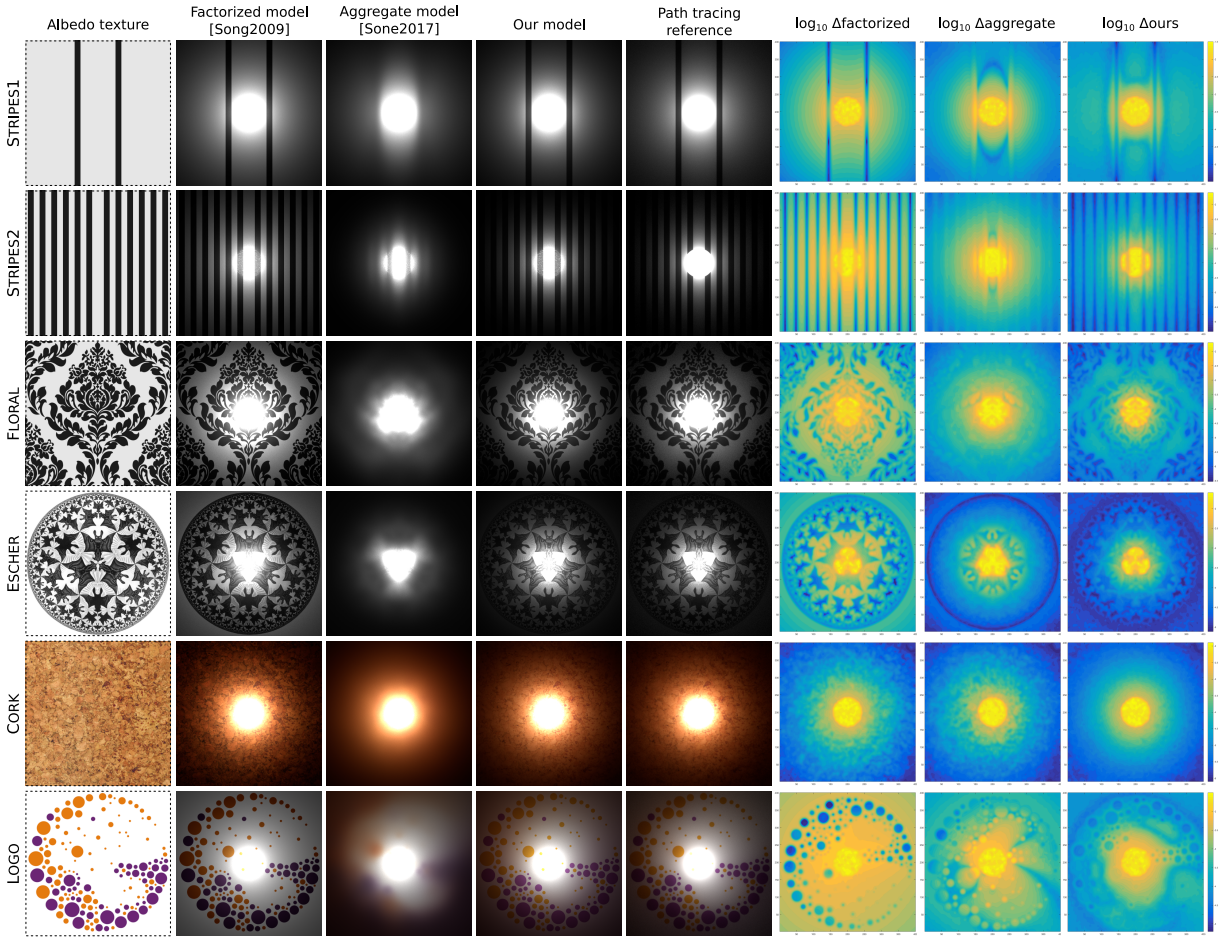

Figure 4: *Main evaluation of the proposed method, discussed in Sec. 3. The renders are tone-mapped with the standard Gamma operator and manually exposed to emphasize the indirect sub-surface transport. The false-color plots show the absolute difference from the reference, and were blurred with a small Gaussian kernel to eliminate spike noise in the visualization.*

this effect also varies with the material's optical thickness, whereas our model can preserve details for a wide range of medium densities σ_t.

Evaluation Cost. Our method is the cheapest option out of the compared three, since it requires only a single evaluation of the expensive baseline BSSRDF, as opposed to two evaluations needed by the factorized method. This is true for the aggregate model as well, which however needs a costly Monte-Carlo integration for the aggregation step, whereas in our case this is a constant-time operation since we utilize Gaussian pyramids for this purpose (cf. Sec. 2).

4. Conclusion

We present a novel methodology for deriving and applying a spatially varying BSSRDF model on the basis of the simulated material's optical properties. Compared to the existing works, the proposed method provides a more principled way of aggregating the material parameters, and well preserves both the local and global transport features thanks to a tailored decomposition of the SV-BSSRDF model. The results obtained for the investigated configurations are very promising, although there are numerous opportunities for future improvement; amongst others the following.

Principled Aggregation Kernel. At the moment, the parameter aggregation kernel is fit manually, as a proof of concept. A better course of action would be numerically fitting a meta-parametric model to the distributions of (contribution-weighted) sub-surface paths for every pair of surface points and required medium parametrizations. Such a meta-model would then output the aggregation kernel on-the-fly, for instance again in the form of a Gaussian mixture model.

Spatial Variation of Material Parameters. While the aggregation of other medium properties beyond the scattering albedo is desirable, arguably the most relevant is handling

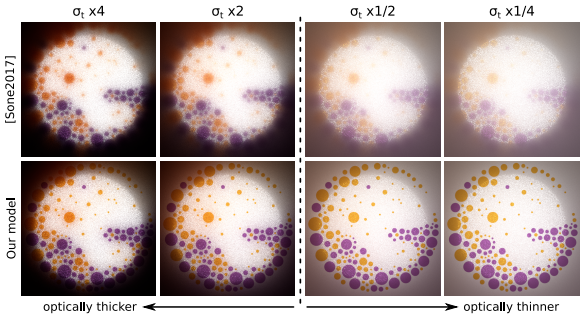

Figure 5: *Color bleeding in materials with varying optical thicknesses σ_t (relative to baseline), demonstrating the more consistent detail preservation of our model compared to Sone et al. [SHK17]. The light source radius is 4 mean free paths, otherwise the configuration is identical with Fig. 4.*

the optical density σ_t. We conjecture that rather than using a simple average [SHK17], it should be used as a scaling factor during the albedo aggregation. This of course needs to be verified by future experiments.

Importance Sampling. Our implementation currently uses a uniform distribution of sample points during the evaluation, resulting in significant noise even after 1k samples. The main challenge in deriving an efficient importance-sampling scheme for the proposed model is that rather than given explicitly, the SV-BSSRDF is *implicitly* defined by the pair of sampling points. One possible approach to tackle this circular dependence could be a multi-level sampling scheme, where the sample point were found iteratively, based on increasingly more accurate estimates of the transport kernel.

Other Improvements. Since the basis BSSRDF determines the accuracy of the global transport, research in this direction is highly relevant. Also, to make the aggregation work for more general geometries and full 3D material parameters distributions, additional investigation of higher-dimensional aggregation kernels and surface-adaptive sampling techniques [FHK14, SHK17] is, too, warranted.

Acknowledgements. We are grateful to Sebastian Herholz and Per Christensen for their feedback, Eugene D'Eon for sharing his simulation code, and Wenzel Jakob for developing the open-source renderer Mitsuba used to generate the reference images. This project has received funding from the European Union's Horizon 2020 research and innovation programme, under the Marie Sklodowska-Curie grant agreement No 642841 (DISTRO), and was further supported by the Czech Science Foundation grant 16-18964S and the Charles University grant SVV-2017-260452.

References

[Chr15] CHRISTENSEN P. H.: An approximate reflectance profile for efficient subsurface scattering. In *ACM SIGGRAPH Talks* (2015), ACM. 2

[dI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH) 30*, 4 (July 2011), 56:1–56:14. 2, 3

[DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics (Proc. SIGGRAPH) 24*, 3 (2005), 1032–39. 1, 2

[DLR*09] DONNER C., LAWRENCE J., RAMAMOORTHI R., HACHISUKA T., JENSEN H. W., NAYAR S.: An empirical BSSRDF model. *ACM Transactions on Graphics (Proc. SIGGRAPH) 28*, 3 (2009), 30:1–30:10. 2

[DWd*08] DONNER C., WEYRICH T., D'EON E., RAMAMOORTHI R., RUSINKIEWICZ S.: A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia) 27* (2008), 140:1–140:12. 1

[ESZ*17] ELEK O., SUMIN D., ZHANG R., WEYRICH T., MYSZKOWSKI K., BICKEL B., WILKIE A., KŘIVÁNEK J.: Scattering-aware texture reproduction for 3D printing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 36*, 6 (2017), 241:1–241:15. 1, 2

[FD17] FREDERICKX R., , DUTRÉ P.: A forward scattering dipole model from a functional integral approximation. *ACM Trans. Graph. 36*, 4 (July 2017). 2

[FHK14] FRISVAD J. R., HACHISUKA T., KJELDSEN T. K.: Directional dipole model for subsurface scattering. *ACM Transactions on Graphics 34*, 1 (2014), 5:1–5:12. 2, 5

[FWKH17] FONG J., WRENNINGE M., KULLA C., HABEL R.: Production volume rendering: Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses* (New York, NY, USA, 2017), SIGGRAPH '17, ACM, pp. 2:1–2:79. 1

[Jar13] JAROSZ W.: The perils of evolutionary rendering research: Beyond the point sample. Eurographics Symposium on Rendering, Invited Talk, 2013. 1

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proc. SIGGRAPH* (2001). 2, 3

[NRH*77] NICODEMUS F., RICHMOND J., HSIA J., GINSBERG I., LIMPERIS T.: *Geometrical considerations and nomenclature for reflectance*. Tech. rep., Final Report National Bureau of Standards, Washington, DC. Inst. for Basic Standards., 1977. 1

[PARN04] PREMOŽE S., ASHIKHMIN M., RAMAMOORTHI R., NAYAR S. K.: Practical rendering of multiple scattering effects in participating media. In *Proc. EGSR* (2004). 2

[PvBM*06] PEERS P., VOM BERGE K., MATUSIK W., RAMAMOORTHI R., LAWRENCE J., RUSINKIEWICZ S., DUTRÉ P.: A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 25*, 3 (2006), 746–53. 1

[SHK17] SONE H., HACHISUKA T., KOIKE T.: Parameter estimation of BSSRDF for heterogeneous materials. In *Eurographics Short Papers* (2017). 2, 3, 5

[STPP09] SONG Y., TONG X., PELLACINI F., PEERS P.: Subedit: A representation for editing measured heterogeneous subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 28*, 3 (July 2009), 31:1–31:10. 2, 3

[SW13] SONG Y., WANG W.: A data-driven model for anisotropic heterogeneous subsurface scattering. In *Proc. of Signal and Information Processing Association Annual Summit and Conference* (2013). 1, 2