



# State of the Art in Efficient Translucent Material Rendering with BSSRDF

Shiyu Liang,<sup>1</sup> Yang Gao,<sup>1</sup> Chonghao Hu,<sup>2</sup> Peng Zhou,<sup>3</sup> Aimin Hao,<sup>1</sup> Lili Wang<sup>1,4</sup> and Hong Qin<sup>5</sup>

<sup>1</sup>State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China  
lsy\_boringfish@163.com, {gaoyangvr, ham, wanglily}@buaa.edu.cn

<sup>2</sup>Beijing University of Posts and Telecommunications, Beijing, China  
hyiker@bupt.edu.cn

<sup>3</sup>Huawei Technologies Co., Ltd., Beijing, China  
zhoupeng88@huawei.com

<sup>4</sup>Peng Cheng Laboratory, Shenzhen, China

<sup>5</sup>Department of Computer Science, Stony Brook University (SUNY at Stony Brook), Stony Brook, NY, USA  
qin@cs.stonybrook.edu

## Abstract

*Sub-surface scattering is always an important feature in translucent material rendering. When light travels through optically thick media, its transport within the medium can be approximated using diffusion theory, and is appropriately described by the bidirectional scattering-surface reflectance distribution function (BSSRDF). BSSRDF methods rely on assumptions about object geometry and light distribution in the medium, which limits their applicability to general participating media problems. However, despite the high computational cost of path tracing, BSSRDF methods are often favoured due to their suitability for real-time applications. We review these methods and discuss the most recent breakthroughs in this field. We begin by summarizing various BSSRDF models and then implement most of them in a 2D searchlight problem to demonstrate their differences. We focus on acceleration methods using BSSRDF, which we categorize into two primary groups: pre-computation and texture methods. Then we go through some related topics, including applications and advanced areas where BSSRDF is used, as well as problems that are sometimes important yet are ignored in sub-surface scattering estimation. In the end of this survey, we point out remaining constraints and challenges, which may motivate future work to facilitate sub-surface scattering.*

**Keywords:** rendering, reflectance & shading models, real-time rendering, subsurface scattering, translucent materials

**CCS Concepts:** Computing methodologies → Rendering; Reflectance modelling; Rasterization

## 1. Introduction

How to create photorealistic images from digital data (models or volumetric parameters) has long been a significant challenge in computer graphics. Researchers have been exploring physically based solutions to simulate light transport, with the aim of minimizing the difference between the appearance of real objects and their rendered counterparts. To achieve this, computation models are required to accurately characterize the distribution of light and the optical properties of materials. Although the bi-directional reflectance distribution function (BRDF) has been widely used to describe light transport on the surface of objects, it is no longer a suitable option when participating media are taken into account.

When considering translucent materials such as human skin, candles and marble, the interaction of light in participating media is crucial for achieving realistic rendering. The light transport simulation within the medium is typically based on the radiative transfer equation (RTE) [Cha60]. This allows for the division of scattering events into scattering, absorption and emission. Path tracing is a popular method used in research to simulate this process, which is still evolving fast in recent years [KGV\*20, DWWH20]. Since several scattering events occur, the estimation of outgoing radiance becomes increasingly complex, which leads to a significant time cost. To accelerate this process without introducing noise, researchers have adopted the diffusion theory [Sta95] framework, which employs approximate solutions of the diffusion equation (DE) to

simulate the light distribution from the medium and geometry assumptions. Since only radiant flux  $\Phi$  incident on the surface is required, it is well suitable for representation with the bidirectional scattering-surface reflectance distribution function (BSSRDF) [NRH\*77], resulting in different diffusion-based methods (based on BSSRDF models). Unlike path tracing, BSSRDF models overlook full light transport within the medium, instead estimating light attenuation between the incident and outgoing points using surface radiant flux. Although these techniques can deliver results comparable to path tracing in simple scenes (such as flat geometry with homogeneous materials) at efficient rates, their quality is constrained by diffusion approximations, which do not strictly adhere to physical laws. However, because of their exceptional acceleration, these methods are widely adopted in real-time applications.

While path tracing is capable of addressing virtually all types of participating media problems, BSSRDF methods are restricted to a subset of them. This is primarily due to their reliance on diffusion-based models, which make several assumptions about translucent objects. These assumptions can lead to shortcuts in materials, such as liquids with low extinction coefficients, resulting in a clear difference between the BSSRDF approaches and other participating media approaches. To distinguish between the methods we discussed in this survey and the general participating media approach, we summarize three major limitations, including:

- **Optically thick:** BSSRDF approaches assume that a large number of scattering events (multiple scattering) occur in the medium. Otherwise, diffusion approximations may result in errors. We exclude some materials in which single scattering can influence appearances significantly, such as low-concentration liquid, low-density cloud and fog [NGD\*06, KF12, SZLG10].
- **Observer outside:** Some participating media problems consider a global effect, meaning the medium involves the entire scene, including the viewpoint or camera [SDS\*16, BSA12]. In contrast, BSSRDF approaches aim only to estimate the radiance leaving the translucent object and are limited to objects that the observer cannot enter.
- **Backlit estimation:** For thick objects, the majority of light will return to the boundary where it refracts. Many BSSRDF approaches only account for this reflectance radiance. In backlit scenarios where light transmits through objects, these methods can perform poorly due to invalid approximations or insufficient information (more details in Section 5.3).

Over the years, numerous techniques for dealing with participating media problems have been developed and well summarized. After Cerezo *et al.* [CPP\*05] conducted a survey about light propagation in the medium more than a decade ago, Nov'ak *et al.* [NGHJ18, NGH\*18] summarized its developments over the next 10 years. Recently, Wu *et al.* [WWY22] introduced the newest groups for homogeneous participating media. However, all of these surveys have only discussed methods based on path tracing, which uses Monte Carlo (MC) to solve the problem, and none of them have focused on sub-surface scattering problems based on diffusion theory. Liang *et al.* [LQTF21] reviewed some important BSSRDF methods but fell short of covering all aspects of them. Thus, it is crucial to categorize various efficient methods of rendering translucent material. In this survey, we provide a comprehensive overview of the BSSRDF

models and explain how they can be applied to solve or accelerate different sub-surface scattering problems. We will investigate not only BSSRDF methods but also other aspects related to them, including discrete methods (finite difference and finite element [FE]), parameter measurement, machine learning and several common problems in BSSRDF. These techniques showcase a range of applications utilizing BSSRDF and provide a more comprehensive solution to sub-surface scattering. We will leave out various path-tracing methods since they have already been sorted out well [NGHJ18].

In this study, we categorize and analyse years of subsurface scattering research with BSSRDF. We investigate the relationships among these methods and look for potential future advances. This survey is structured as follows. Section 2 introduces the background of sub-surface scattering, including its properties, definitions and key diffusion theory, which serves as a firm framework for the subsequent discussion. In Section 3, we describe a major alternative to path tracing, BSSRDF models, which are derived from diffusion theory. Since acceleration solutions are the main focus of this survey, we summarize and classify these methods in Section 4. Section 5 discusses a wider range of efficient rendering and sub-surface scattering applications, including measuring BSSRDF parameters, perfecting the scattering effect and developing advanced machine learning methods. Finally, we discuss some of the remaining limitations and potential obstacles in Section 6, which could be major future development directions.

## 2. Background

To enhance understanding of sub-surface scattering theories, we compile relevant definitions and theories in this section. First, we present an overview of fundamental volumetric light transport, including its properties and formulas, which aims at solving the RTE. Building upon this foundation, diffusion theory proposes an alternative framework to simulate radiance distribution within the medium using certain approximations. Since the solution approximates volumetric light transport without explicit volume construction and can only require incident radiance on the surface, BSSRDF definitions can be employed to effectively articulate the outgoing radiance. Physically based BSSRDF models have benefited from research towards enhancing diffusion theory. We include all relevant definitions in Table 1 to facilitate reference.

### 2.1. Radiative transfer equation

Most research assumes that translucent materials consist of a large number of scattering and absorbing particles. When light travels through the medium, it interacts with these particles through absorption and scattering processes, which can be described by the RTE [Cha60], forming the basis of volumetric light transport. These interactions are estimated using the scattering coefficient  $\sigma_s$  and the absorption coefficient  $\sigma_a$ , which represent the density of scattering and absorbing particles per unit distance, respectively. The density of a scattering event occurring can be obtained by summing up the scattering and absorption coefficients, yielding the extinction coefficient  $\sigma_t = \sigma_s + \sigma_a$ . If a scattering event does occur, the proportion of light that is scattered can be estimated using the single-scattering albedo,  $\alpha = \sigma_s / \sigma_t$ . With these optical properties, the RTE can be

**Table 1:** Nomenclature (top: about optical properties, bottom: about BSSRDF models). We use  $x$  to represent position and  $\omega$  to represent direction, where subscript  $i$  represents incident and  $o$  represents outgoing.

Notation	Description
$p(\omega, \omega')$	phase function
$g$	average cosine of the scattering angle
$\sigma_s$	scattering coefficient
$\sigma_a$	absorption coefficient
$\sigma_t = \sigma_s + \sigma_a$	extinction coefficient
$\sigma'_s = \sigma_s(1 - g)$	reduced scattering coefficient
$\sigma'_t = \sigma'_s + \sigma_a$	reduced extinction coefficient
$\sigma_{tr} = \sqrt{\sigma_a/D}$	effective transport coefficient
$l = 1/\sigma_t$	mean free path (mfp)
$l' = 1/\sigma'_t$	transport mean free path
$\alpha = \sigma_s/\sigma_t$	single-scattering albedo
$\alpha' = \sigma'_s/\sigma'_t$	reduced single-scattering albedo
$L(x, \omega)$	volume radiance
$L_e(x, \omega)$	medium emission
$L_s(x, \omega)$	in-scattered/incoming radiance
$\phi(x)$	fluence/irradiance
$\Phi$	power of the light source
$\vec{E}(x)$	vector irradiance
$Q$	source function
$D$	diffusion coefficient
$A$	reflection parameter
$z_b = 2AD$	linear-extrapolation distance
$\eta = \eta_2/\eta_1$	relative index of refraction
$n$	surface normal
$r =  x_i - x_o $	scattering distance
$S^{(0)}$	reduced intensity transmission
$S^{(1)}$	single scattering term
$S_d$	multiple scattering term
$S = S^{(0)} + S^{(1)} + S_d$	BSSRDF model
$L_o(x_o, \omega_o)$	outgoing radiance
$L_i(x_i, \omega_i)$	incident radiance
$F_r$	Fresnel reflectance
$F_t = 1 - F_r$	Fresnel transmittance
$F_{dr} = \int_{2\pi} F_r \cdot (n \cdot \omega) d\omega$	average diffuse Fresnel reflectance
$R(x_i, x_o) \approx R(r)$	diffuse BSSRDF/reflectance profile
$B(x_o)$	outgoing radiosity
$E(x_i)$	transmitted/incident irradiance

expressed as

$$(\omega \cdot \nabla)L(x, \omega) = -\sigma_t L(x, \omega) + \sigma_s \int_{4\pi} p(\omega, \omega') L(x, \omega') d\omega' + Q(x, \omega), \quad (1)$$

where  $Q$  denotes the volume source distribution, which is estimated as  $Q = \sigma_a L_e(x, \omega)$ ,  $L_e(x, \omega)$  is the medium emission. The phase function  $p(\omega, \omega')$  is the directional density of light scattered at a point in a medium. When integrating Equation (1) and accounting for the hard boundary of objects (which limits the transport distance), the volume rendering equation (VRE) can be obtained as

$$L(x, \omega) = \int_0^z T(x, x') [\sigma_a(x') L_e(x', \omega) + \sigma_s(x') L_s(x', \omega)] dt + T(x, x_s) L(x_s, \omega), \quad (2)$$

where  $L(x_s, \omega)$  is the incident radiance at surface point  $x_s$ ,  $z$  is the distance from  $x$  to the nearest surface along  $\omega$ ,  $T(x, x')$  represents the transmittance along the path from  $x'$  to  $x$ , and  $t$  denotes the distance with  $x' = x - t\omega$ .  $L_s(x, \omega) = \int_{4\pi} p(\omega, \omega') L(x, \omega') d\omega'$  is called in-scattered or incoming radiance, representing the collection of radiance scattered to the direction  $\omega$  from all directions. According to the Bouguer–Lambert law [Lam60],  $T(x, x')$  can be written as

$$T(x, x') = e^{-\tau(t)}, \quad (3)$$

where  $\tau(t) = \int_0^t \sigma_t(x - s\omega) ds$  is the optical thickness. Note that this exponential transmittance model is undergoing a nascent improvement by extending the theory of non-exponential medium to a heterogeneous transmittance model [d18, VJK21]. The VRE forms the basis for path tracing, which can be solved using the integral formulation of MC. For a more detailed discussion, we recommend the report by Novák et al. [NGHJ18, NGH\*18]. However, while MC path tracing can achieve high accuracy, it is slow to converge for complex random-walking procedures. This has motivated research in using analytical formulas to approximate volumetric light transport using a diffusion process [Sta95], leading to the development of BSSRDF models based on diffusion theory.

## 2.2. Diffusion theory

With the complicated integral form, MC path tracing needs a large number of random-walk paths to simulate the transport process in the medium. In contrast, diffusion theory attempts to use *diffusion approximations* to solve RTE. It is grounded on a crucial observation: even if the initial light distribution is extremely anisotropic, it tends to become isotropic in highly scattering media soon because of the multitude of scattering events. Stemming from diffusion approximations, diffusion theory often requires isotropic scattering in the medium, which is determined by the phase function and succinctly described by the average cosine of the scattering angle,  $g$ :


$$g = \int_{4\pi} (\omega \cdot \omega') p(\omega, \omega') d\omega'. \quad (4)$$

If  $g > 0$ , the scattering is predominantly forward, while if  $g < 0$ , it is predominantly backward. When  $g = 0$ , the scattering is considered isotropic ( $p(\omega, \omega') = 1/(4\pi)$ ) for optically thick materials. As per similarity relations [WPW89, ZRB14], different optical properties can lead to the same rendering result after multiple scattering events. This implies that for most materials with  $g > 0$ , a simple scaling of the scattering coefficient by  $1 - g$  can convert the scattering from anisotropic to isotropic. Specifically, in the context of diffusion theory, reduced optical properties are used more frequently than original ones, such as the reduced scattering coefficient  $\sigma'_s = \sigma_s(1 - g)$  and the reduced extinction coefficient  $\sigma'_t = \sigma'_s + \sigma_a$ .

In the case of isotropic scattering, the radiance can be approximated with two terms by considering only the first-order spherical harmonic expansion:

$$L(x, \omega) = \frac{1}{4\pi} \phi(x) + \frac{3}{4\pi} \vec{E}(x) \cdot \omega, \quad (5)$$

where  $\phi(x) = \int_{4\pi} L(x, \omega) d\omega$  is the fluence or irradiance and  $\vec{E}(x) = \int_{4\pi} L(x, \omega) \omega d\omega$  is the vector irradiance. When in the homogeneous



**Figure 1:** The searchlight problem.

medium, assuming that the source function  $Q$  is isotropic, the *diffusion equation* can be obtained by substituting this two-term expansion of the radiance into RTE [JMLH01]:

$$-D\nabla^2\phi(x) + \sigma_a\phi(x) = Q(x), \quad (6)$$

where  $D = 1/(3\sigma'_t)$  represents the diffusion coefficient. The solution  $\phi$  of the DE should be related to a given transport problem with specific boundary conditions. In the scenario of an isotropic point source (a monopole) in an infinite homogeneous medium, the classical diffusion Greens function gives an approximate solution:

$$\phi^m(x) = \frac{\Phi}{4\pi D} \frac{e^{-\sigma_{tr}r(x)}}{r(x)}, \quad (7)$$

where  $\Phi$  is the power of the point light source in the medium and  $r$  is the distance from  $x$  to the light source.  $\sigma_{tr} = \sqrt{\sigma_a/D}$  is the effective transport coefficient. This approximate solution used in early research [JMLH01] is not the only available solution. Different approximations result in solutions either by discarding the transient terms and retaining only the discrete terms [CZ67], or by decoupling diffusion from uncollided flux [Gro56]. The former is used as an important approximation in zero-variance theory [Kd14, KGV\*20] and the latter performs better for materials with high absorption levels [d11]. Recent research has indicated that the spatial correlation between the scattering centres in the medium can also affect diffusion approximations. While classical linear transport theory presumes exponentially distributed free paths, the generalized radiative transfer (GRT) theory [d19] considers spatially correlated scattering centres in arbitrary dimensions, leading to non-exponential light transport. In this instance, moment-preserving diffusion approximations involve the memory of the distance between collisions in free paths. More details and discussions can be found in recent works [d13, BRM\*18, JAG18, d19, KGV\*20], which are crucial for enhancing diffusion approximations in spatially correlated materials.

The solution  $\phi$  in Equation (7) only describes the fluence distribution from a single point source in an infinite medium. To apply it to sub-surface scattering, two major issues must be addressed: the source function and boundary conditions. A representative problem in this context is the searchlight problem (Figure 1). It consists of a flat, semi-infinite and homogeneous medium, and the incident light is only at a surface point from a single direction (the normal of this surface). Since this incident light gives rise to a refracted ray in the medium with exponentially decreasing intensity, the source function aims to replace this refracted beam with point sources inside the medium, making diffusion approximations suitable. On the

other hand, the semi-infinite medium adds a surface constraint to diffusion approximations, which requires the net inward diffuse flux to be zero at each surface point  $x_s$ ,  $\int_{2\pi} L(x_s, \omega)(\omega \cdot n(x_s))d\omega = 0$ . Adding it to the DE, the boundary condition [JMLH01] yields

$$\phi(x_s) - 2AD(n \cdot \nabla)\phi(x_s) = 0, \quad (8)$$

where  $A$  is the reflection parameter. Following the boundary condition, the fluence should be zero at a distance  $z_b = 2AD$  above the surface, commonly referred to as the linear-extrapolation distance. In most methods, a common solution to satisfy the boundary condition is to place a negative mirrored source outside the medium for every positive source inside the medium (see Figure 4). In summary, the diffusion approximation, the boundary condition and the source function together form diffusion theory in sub-surface scattering. Different BSSRDF models have been developed from it to solve the searchlight problem by continuously improving these three aspects. However, the searchlight configuration simplifies the complex geometry of objects and light conditions, which are the primary reasons that the rendering results from the BSSRDF deviate from reality.

### 2.3. BSSRDF definition

The adoption of diffusion theory allows the evaluation of the attenuation of radiance between incident and outgoing points using analytical formulas. To calculate the outgoing radiance  $L_o$  at the outgoing point  $x_o$ , it is necessary to accumulate the incident flux  $\Phi$  from all incident points  $x_i$  that encompass the entire surface of the object. This interaction between different surface locations is well-described using BSSRDF [NRH\*77]. Given the incident flux  $\Phi(x_i, \omega_i)$  at the point  $x_i$  originating from direction  $\omega_i$ , the BSSRDF, denoted by  $S$ , is defined as follows:

$$S(x_o, \omega_o; x_i, \omega_i) = \frac{dL_o(x_o, \omega_o)}{d\Phi(x_i, \omega_i)}. \quad (9)$$

While the BRDF operates under the assumption that light enters and exits at the same point ( $x_i = x_o$ ), the BSSRDF provides a more comprehensive and precise depiction of light transport on the surface of the object. Considering the incident radiance  $L_i$  from all surface area and all incoming directions, the outgoing radiance  $L_o$  is a double integral,

$$L_o(x_o, \omega_o) = \int_A \int_{2\pi} S(x_o, \omega_o; x_i, \omega_i)L_i(x_i, \omega_i)|n \cdot \omega_i|d\omega_i dA(x_i). \quad (10)$$

Estimating the BSSRDF  $S$  in subsurface scattering is challenging due to the complex behaviour of light transport and the various parameters involved. A common approach is to separate  $S$  into three terms:  $S = S^{(0)} + S^{(1)} + S_d$ , where  $S^{(0)}$  is the reduced intensity transmission (no scattering),  $S^{(1)}$  is the single scattering term and  $S_d$  is the multiple scattering term. This division is mainly due to their different dependencies on incident directions [HK93], and each term can be handled by specialized algorithms. However, other separations can also be seen, such as moving  $S^{(1)}$  to  $S_d$  [Chr15] or to  $S^{(0)}$  [FHK14] as a whole. For multiple scattering  $S_d$ , a product of three terms is often used to separate the effect of

directions [JMLH01, dI11]:

$$S_d(x_o, \omega_o; x_i, \omega_i) = \frac{1}{\pi} F_t(\eta, \omega_i) R(x_i, x_o) \frac{F_t(\eta, \omega_o)}{4C_\phi(1/\eta)}, \quad (11)$$

where  $F_t$  is the Fresnel transmittance term,  $4C_\phi$  is a constant needed for normalization and  $\eta$  is the relative index of refraction.  $R(x_i, x_o)$  is called the diffuse BSSRDF or the reflectance profile, representing the ratio of radiance exiting at  $x_o$  from  $x_i$ . For convenience,  $R$  is often considered related to the distance between two points rather than their positions, so  $R(x_i, x_o)$  can be re-written as  $R(\|x_i - x_o\|)$ . In conclusion, if the reflectance profile  $R$  is derived from the diffusion approximation  $\phi$ , the outgoing radiance can be estimated at all surface points. Note that Equation (11) makes assumptions about uniform distributions for incident and outgoing light in  $R$ . Since the modulation of the light directions in the outgoing radiance is only determined by  $F_t$ , it is not valid when taking into account the incident and outgoing directions [FHK14, d'E14].

On the other hand, substituting Equation (11) into Equation (10) and moving terms outside the integration yields:

$$L_o(x_i, \omega_o) = \frac{1}{\pi} F_t(\eta, \omega_o) B(x_o), \quad (12)$$

where  $B(x_o)$  denotes the outgoing radiosity. Furthermore, it can be written as

$$B(x_o) = \int_A E(x_i) R(x_i, x_o) dx_i, \quad (13)$$

where

$$E(x_i) = \int_{2\pi} L_i(x_i, \omega_i) F_t(\eta, \omega_i) (n \cdot \omega_i) d\omega_i \quad (14)$$


is called the transmitted or incident irradiance. The main advantage of this form is that the incoming and outgoing directions are independent, dividing the whole estimation into two passes. All view-independent results (such as incident irradiance  $E$ ) can be pre-computed (Section 4.1) or stored (Section 4.2) in the first pass, significantly improving rendering efficiency. Section 4 will discuss these acceleration methods in greater detail.

### 3. BSSRDF Model

The BSSRDF models are derived from diffusion theory and describe the relationship between the outgoing radiance and incident flux on the surface. Since it ignores the complex light transport in the medium, the introduction of BSSRDF models reduces the rendering time from hours to minutes. The diffusion profile  $R$  in the BSSRDF  $S$  (Section 2.3) is most often derived from diffusion theory or empirical approximations, which is the final required result in many BSSRDF models. After Jensen *et al.* [JMLH01] first derived a diffusion profile from the classical diffusion theory, various theoretical and empirical profiles have been proposed (Figure 2), which are discussed below.

#### 3.1. Theoretical models

Jensen *et al.* [JMLH01] first used two point sources, one positive and one negative, to solve the source function problem in sub-surface scattering with diffusion theory. The two-point source distribution



**Figure 2:** A simple classification of BSSRDF models we discussed. The classification is due to computation complexity and directional dependence (The black labels are theoretical and red labels are empirical). ① [JMLH01], ② [DJ05], ③ [dLE07], ④ [DLR\*09], ⑤ [dI11], ⑥ [dI2], ⑦ [YZXW12], ⑧ [HCJ13], ⑨ [FHK14], ⑩ [d'E14], ⑪ [JZJ\*15], ⑫ [Chr15], ⑬ [FD17].

can approximate the true light distribution well with great accuracy, which is known as *dipole*. To satisfy the boundary conditions, this single dipole is placed near the surface. The positive real light source is located beneath the surface at depth  $z_r = 1/\sigma'_t$ . The negative virtual light source should be mirrored about  $z_b$ , which is located at a distance  $z_v = z_r + 4AD$  above the surface. For the reflection parameter  $A$ , it is estimated as

$$A = \frac{1 + F_{dr}}{1 - F_{dr}}, \quad (15)$$

where  $F_{dr}$  is the average diffuse Fresnel reflectance. Using this configuration of the dipole, the fluence is written as

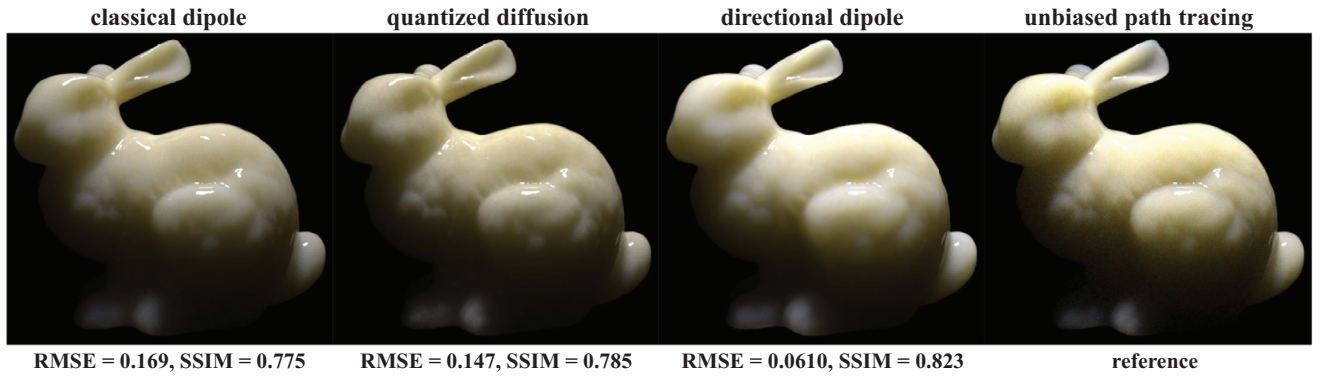
$$\phi(x) = \frac{\Phi}{4\pi D} \left( \frac{e^{-\sigma_{tr}d_r}}{d_r} - \frac{e^{-\sigma_{tr}d_v}}{d_v} \right), \quad (16)$$

where  $d_r = \|x_s - x_r\|$  is the distance between the surface point  $x_s$  and the real source, and  $d_v = \|x_s - x_v\|$  denotes the distance from  $x_s$  to the virtual source. Finally,  $R$  can be estimated as the radiant exitance divided by the incident flux:

$$R(r) = -D \frac{(n \cdot \nabla \phi(x_s))}{d\Phi_i} = \frac{\alpha'}{4\pi} \left[ z_r(\sigma_{tr}d_r + 1) \frac{e^{-\sigma_{tr}d_r}}{d_r^3} + z_v(\sigma_{tr}d_v + 1) \frac{e^{-\sigma_{tr}d_v}}{d_v^3} \right], \quad (17)$$

where  $\alpha'$  is the reduced scattering albedo. The dipole model produces high-quality rendering results in minutes, which is significantly faster than path tracing. However, since the dipole makes many assumptions about object geometry and material, it only fits smooth, semi-infinite homogeneous materials with a single layer, whereas translucent objects in the real world have a more complex appearance. Moreover, this single-depth model is inaccurate under high-frequency illumination, leading to a waxy appearance.

To overcome the drawbacks of the dipole in thin slabs and multi-layered materials, multiple dipoles (*multi-pole*) [DJ05] are used to simulate true light distribution. This infinite array of dipoles is positioned to satisfy boundary conditions at both the top and bottom of the slab. Although the multi-pole can improve the quality of appearance with the combination of multiple layers, it still places the



**Figure 3:** Stanford bunny (made from white grapefruit juice) rendered using different BSSRDF models. While the classical dipole [JMLH01] and QD model [d111] cannot fully capture translucency effects, the directional dipole [FHK14] performs greatly using a ray source, which better corresponds to refracted light from the incident direction on the translucent material surface. Image reproduced from Frisvad et al. [FHK14].

first positive source at a depth of  $z_r$  (1 transport mean free path) in accordance with classical diffusion theory, which limits the slab thickness. To further improve the diffusion theory, d'Eon and Irving [d111] proposed the *Quantized-Diffusion (QD)* model. The QD model utilized Grosjeans approximation [Gro56] to illustrate the RTE with an isotropic point source in the infinite medium, which is more accurate than the diffusion Greens function in the classical diffusion approximation. They presented the modified diffusion coefficient, the reflection parameter and the extended-source function [FPW92] to replace the classical model. This QD model maintains high-frequency details that the dipole fails to capture, but Gaussian weights in it are highly complex and challenging to analyse (in practice, they should be pre-computed). Nevertheless, individual modifications in the fully improved diffusion theory (such as modifying the diffusion coefficient) can enhance previous models. The resulting *better dipole* [d12] maintains the point source rather than the extended source, making it considerably easier to implement. The better dipole provides a higher level of accuracy than the dipole [JMLH01] without extra evaluation cost, but there remains an accuracy gap compared to the QD model.

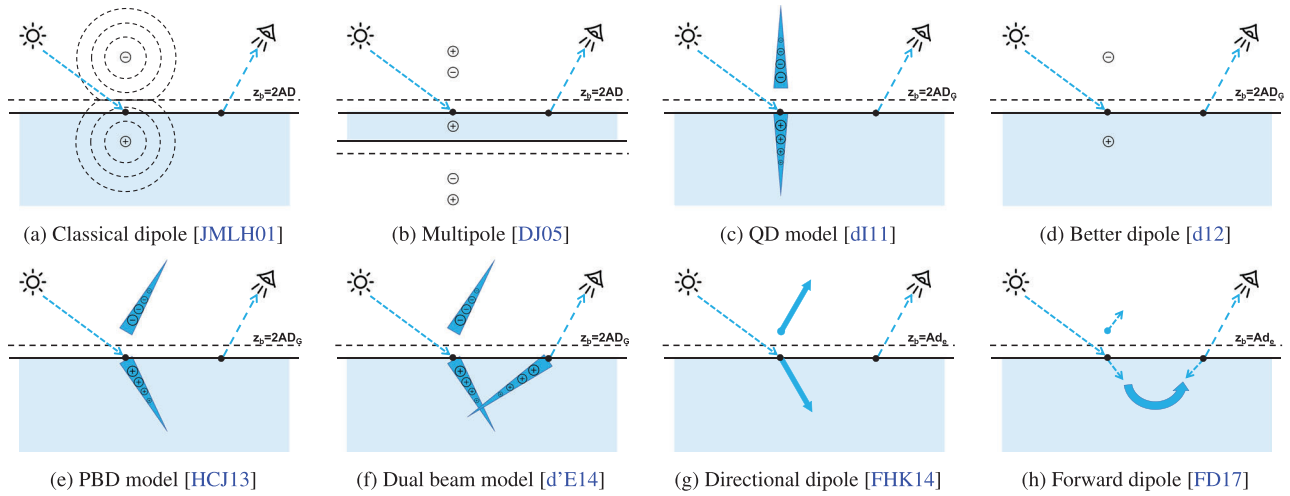
While decomposing  $S_d$  into a direction-independent component  $R$  and using only  $F_r$  to modulate directional influences effectively reduces the complexity of the model, this approach may miss crucial translucency effects. Some BSSRDF models favour a more comprehensive representation of  $S_d$  rather than relying on this simple separation. The direction dipole [FHK14] takes into account the incident direction, which is an important factor overlooked in previous studies. Based on an analytical solution for ray sources [MSG05], the incident direction can be considered in  $S_d$  instead of being modulated by  $F_r$ . While previous models assume that  $R$  is only related to distance, the directional dipole relaxes the assumption of isotropic distribution in incident directions, making the results closer to path-tracing references (Figure 3). Recent *forward scattering dipole* [FD17] proposed a new analytical solution to diffusion theory for highly forward-scattering materials. Rather than light sources, it uses the propagator to simulate light transport in the medium, which takes both incident and outgoing

directions into consideration. The forward scattering dipole has greater accuracy in the  $g \rightarrow 1$  medium than other models. However, since the phase function and propagator used assume strongly forward scattering materials, it cannot be applied to a wide range of materials.

**Hybrid methods.** All the limitations of theoretical models stem from material assumptions, which MC path tracing can eliminate. Some techniques address these restrictions by combining BSSRDF models with MC. These hybrid methods can produce results in less time than path-tracing and with greater accuracy than most BSSRDF models. BSSRDF models cannot work well near the surface, where the light is highly influenced by the incidence direction and material geometry. A simple way to address this issue is to divide the object into two halves, where an isotropic core region is defined within the medium [LPT05]. The photons contribute with path tracing if they do not arrive in this core region; otherwise, the contribution is estimated using the dipole.

Due to the high computational cost of photon random walking, a more common solution used in hybrid approaches within path tracing is photon mapping [JC98]. In these methods, photons are recorded and stored as real light sources when they first scatter in the medium along the refracted direction. The total exitant radiance is estimated by summing all contributions from stored sources using different configurations (the dipole, multi-pole and quadpole) of BSSRDF models, known as *photon diffusion* [DJ07]. Habel et al. [HCJ13] expanded this method to photon beam diffusion, which stores photon beams rather than photon points to improve its accuracy with the improved diffusion theory [d111]. Considering a refracted source beam in the medium, the  $R$  can be estimated using MC integration from photon samples on the beam:

$$\begin{aligned}
 R(x, \omega) &= \int_0^\infty r(x, x_r(t))Q(t)dt \\
 &\approx \frac{1}{N} \sum_{i=1}^N \frac{r(x, x_r(t_i))Q(t_i)}{pdf(t_i|x, \omega)}
 \end{aligned} \tag{18}$$



**Figure 4:** Different setups of theoretical BSSRDF models.  $\oplus$  and  $\ominus$  are positive real and negative virtual light sources, respectively. From the single isotropic point source (a), light sources progressively evolve into multiple sources (b), the extended source with exponentially decreasing intensity (c, e, f), ray source (g) and propagator (h) for improved representations, accompanied by the enhancement of  $z_b$ .

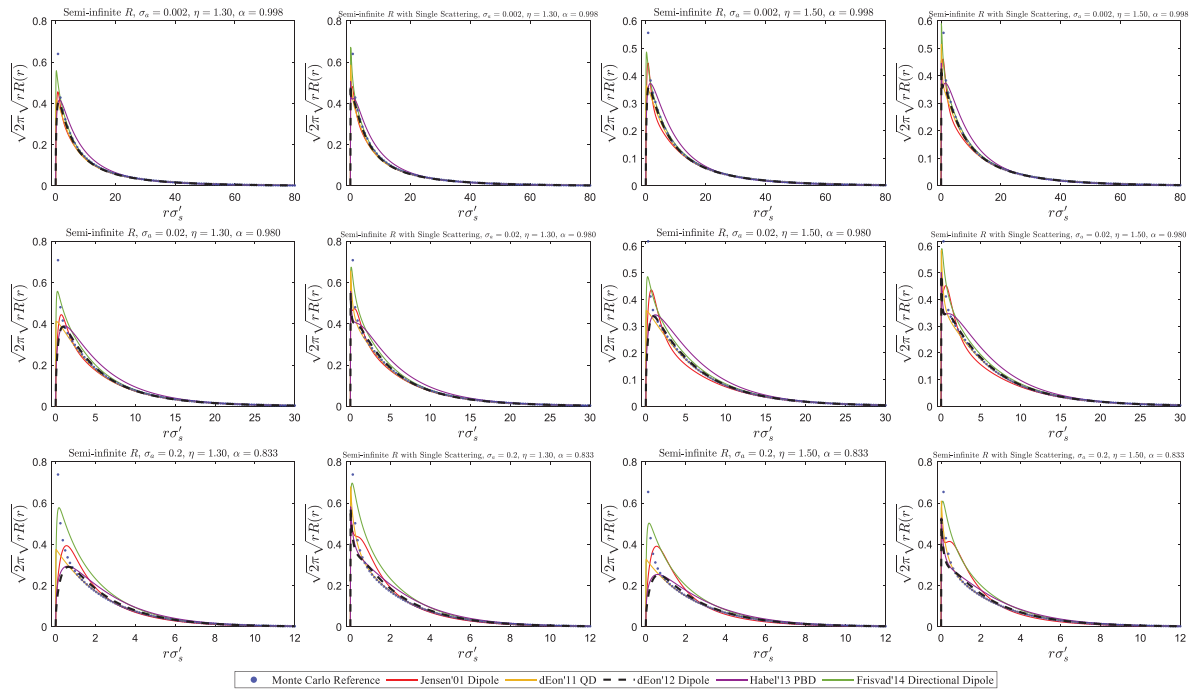
where  $Q(t)$  is the extended source function and  $r(x, x_r(t))$  is the radiant intensity in the improved diffusion theory [dI11].  $pdf(t_i|x, \omega)$  is the PDF of choosing  $t_i$ . Unlike using complex Gaussians to fit the QD model, the  $R$  is estimated by MC and is calculated along the incident direction. d'Eon [d'E14] further extended the photon beam diffusion to account for both the incident and outgoing directions, leading to a complete (both the spatial and angular domains are 4D) 8D BSSRDF.

**Discussion.** The development of the theoretical BSSRDF models, including hybrid methods, has led to significant advances in diffusion theory and the resolution of limitations in the searchlight problem. Figure 4 emphasizes the fundamental differences between these models. Modifications to the diffusion theory are primarily reflected in two areas: the source function and boundary conditions. The source function aims to approximate the incident radiance of the scene light source with real light sources in the medium, enabling the calculation of the fluence  $\phi(x)$  on the material boundaries based on the distance between them. Although the early isotropic single point source (Figure 4a) provides an efficient solution, it cannot handle complex appearances and high-frequency illumination. With multiple sources, both the multi-pole model (Figure 4b) and QD model (Figure 4c) can produce more realistic multi-layer rendering output, albeit at a higher evaluation cost. However, since the extended source in the QD model considers beams with exponentially decreasing intensity, it is more accurate than uniform sources in multi-pole. Recent models have broken the assumption that the reflectance profile  $R$  is independent of the incident and outgoing directions. While the PBD model (Figure 4e) considers incident directions by placing real light sources along refracted incident beams, the dual beam model (Figure 4f) further accounts for outgoing directions by collecting contributions along refracted outgoing beams. In addition to point sources, the ray source (Figure 4g) and even the propagator  $G$  (Figure 4h) can also describe light transport in the medium from oblique beams, leveraging RTE and allowing for modelling more complicated

geometries (such as unflat surfaces). All these considerations make BSSRDF models more accurate and applicable for common materials and scenes by overcoming limitations, enabling the medium to be layered and account for the incident and outgoing directions.

With various real light sources in the medium, the boundary condition is employed to determine the positions of virtual light sources, ensuring that the net inward diffuse flux is zero on the surface. This requirement implies that virtual and real light sources must be symmetric about the linear-extrapolation distance  $z_b$ . In the refining of diffusion theory, both  $A$  and  $D$  in  $z_b$  are approximated more accurately, particularly  $D$ . The QD model (Figure 4c) utilizes Grosjeans modified diffusion coefficient  $D_G$  to replace the original  $D$ . Although employing a single point source (Figure 4d),  $D_G$  can still yield an improved model accuracy. In the directional dipole (Figure 4g),  $d_e$  is considered to be closer to the reference than  $2D$ , providing a superior approximation for the extrapolation distance.

We have implemented several theoretical models that feature reflectance profile  $R(r)$  (where applicable) in the 2D searchlight problem and compared them with the ground truths of MC (as illustrated in Figure 5). For various absorption levels and indices of refraction, we present results that exclude and include single scattering [JMLH01, HCJ13]. It is worth mentioning that the BSSRDF models described by diffusion theory do not attempt to include single scattering [d12]. This fact mainly explains why some results exhibit good and similar performance at larger distances but falter in accuracy near the illumination point ( $r = 0$ ). Since classical diffusion theory tries to incorporate single scattering, the dipole [JMLH01] performs poorly near sources, and even the inclusion of single scattering cannot improve its precision. In contrast, various improved diffusion theories (QD, PBD, directional dipole) with modifications (such as better boundary conditions and source function) have succeeded in matching MC ground truths when adding single scattering. On the other hand, as absorption



**Figure 5:** Comparison between different theoretical BSSRDF models. In each pair of reflectance profiles  $R(r)$ , results are shown without single scattering and with it (on the right). We use a fixed  $\sigma'_s = 1$  and increase  $\sigma_a$ , which alters the single-scatter albedo. Under various absorption levels and refraction indices, the models match the ground truths of the MC to varying degrees.

levels ( $\sigma_a$ ) increase, it is observed that both the dipole and the directional dipole tend to overestimate  $R(r)$  near the point of illumination. In such scenarios, the use of the Grosjean modification (in QD, PBD and better dipole) can still maintain high accuracy, making it an excellent choice for high-absorption materials. In particular, although  $R(r)$  is a useful metric for measuring precision, it may not always be fair or sufficient. Estimating the accuracy of BSSRDF models should consider complex scene conditions. For instance, the directional dipole can perform better than the QD model, as its ray source corresponds more closely to oblique incident beams. The better dipole has also been shown to be more accurate than other models when using spatially varying optical properties [NK18]. A more suitable comparison may involve visual images of rendered objects under specific conditions (object geometry, optical properties and environmental illumination).

### 3.2. Empirical models

Although theoretical BSSRDF models based on diffusion theory significantly reduce computation time, the evaluation of  $S_d$  and the configuration of real sources (points and beams) remain too complex, limiting their ability to be interactive. To address this issue, empirical models fit these theoretical models without considering the physics, resulting in faster rendering time. To simplify the calculation, most of the profiles  $R$  of empirical models are solely connected to the distance  $r$  and are combined with the texture methods (explained in Section 4.2). d'Eon *et al.* [dLE07] employed the weighted sums of Gaussian distribution to fit the diffusion profile  $R$


(dipole or multi-pole).

$$R(r) = \sum_{i=1}^k w_i G(v_i, r), \quad (19)$$

where  $w_i$  is the weight and  $v_i$  is the variances for  $k$  Gaussian distributions. Since  $k$  is proportional to material complexity, multi-layered materials require more Gaussians (around 6) to fit. The main advantage of using Gaussians is their separability, which allows the entire convolution to be translated into two 1D convolutions. However, the number of Gaussians has a considerable impact on rendering cost, hence in some studies  $R$  is approximated using only a single Gaussian [KKCF13, Mik10].

Using Gaussians to fit the dipole model brings great simplification, but it remains costly for real-time rendering for weighted Gaussians, needing about 12 passes in rendering pipelines. To minimize cost, separable sub-surface scattering (SSSS) [JZJ\*15] employs singular value decomposition (SVD) [EY36] to reconstruct the weighted Gaussian kernel  $R$ , which requires only two passes to generate the images. Considering that models should be artist-friendly in production, an easier model simulates close- and far-range scattering using only two Gaussian distributions [JZJ\*15]. Christensen [Chr15] introduced a more user-friendly normalized reflectance profile that only requires surface albedo  $A$  and mean free path length  $l = 1/\sigma'_t$  as parameters. This model aligns directly with brute-force MC references, achieving even greater precision than most theoretical models [dI11, HCF13]. A comparison between them is shown in Figure 6. Since the normalized reflectance profile





**Figure 6:** Comparison of two empirical models using screen-space methods. (a) SSSS [JZJ\*15] +  $15 \times 15$  filtering. (b) Normalized reflectance profile [Chr15] + distance importance sampling [XOKN20] (15 samples). In complex regions such as the beard, SSSS results in too much blur, while the normalized reflectance profile works better.

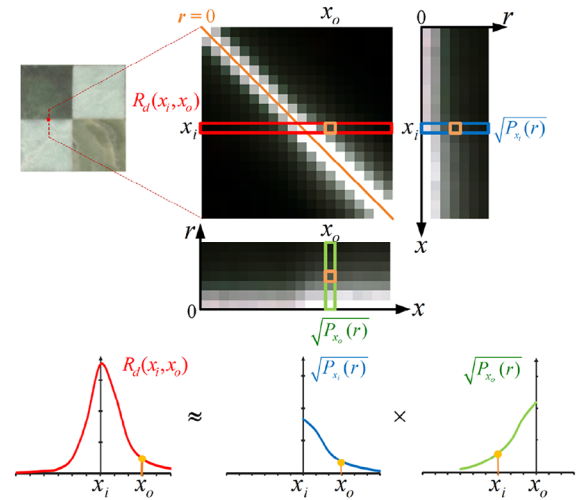
considers surface albedo, it handles complex appearances better but also requires more time.

Not all empirical BSSRDF models are designed to save rendering time. Some are used to solve specific problems. Since the dipole model [JMLH01] cannot handle materials determined by low-order scattering events, Donner *et al.* [DLR\*09] proposed a data-driven BSSRDF model derived from large-scale MC simulations, with the aim of improving the model accuracy in materials with low extinction coefficient. However, it should be noted that the incorrect implementation in their next-event estimator may lead to unusable results in their approach. Based on the QD model [dl11], Yan *et al.* [YZXW12] presented a solution to render translucent materials under complex spherical Gaussian (SG) lights. Given the properties of SG, different parameters are pre-computed and stored in four 2D tables with complex calculations, which are then used to render different materials.

Empirical models most often fit the scattering results directly. Despite the lack of theoretical reasoning, these data-driven profiles show remarkable accuracy (equivalent to or even higher than theoretical models) [Chr15]. While some models aim to achieve higher precision [DLR\*09], the primary objective of most empirical models is to simplify the calculation. Several issues should be considered when applying these models to production. The first is user-friendly parameters. The parameters used in the models should be understood by the artists and can directly affect the scattering results [Chr15]. The other is the sample distribution. A suitable choice of samples (Section 4.2) can improve accuracy and efficiency. One common problem with most empirical BSSRDF models is that the input parameter related to it is only  $r$ . Ignoring incident and outgoing directions is the main factor that limits the accuracy of current empirical models. Therefore, the development of an affordable and user-friendly model that takes into consideration both light and view directions is an important future endeavour. Currently, this area still has a significant gap (Figure 2).

### 3.3. Re-rendering models

A special way to acquire model parameters is through material measurements. These methods first capture specific materials in the real world and obtain their parameters, then re-render them with



**Figure 7:** SubEdit representation illustrated on a measured chessboard material. The full  $R(x_i, x_o)$  matrix can be expressed as the product of two scattering profiles at  $x_i$  and  $x_o$ . Figure reproduced from Song *et al.* [STPP09].

high accuracy. Although most methods aim at searching for optical properties (such as  $\sigma_s$  and  $\sigma_a$ ), some researches have introduced BSSRDF models to fit particular representations of measurable data [STPP09, DJ06]. In this case, special BSSRDF models are defined and used to re-render captured materials.

To make editing heterogeneous materials captured from real-world samples easier, Song *et al.* [STPP09] proposed a new representation of the diffusion profile, *SubEdit*. The diffusion profile  $R(x_i, x_o)$  is separated into two local scattering profiles at incident point  $x_i$  and outgoing point  $x_o$  to decouple the non-local behaviour (Figure 7),

$$R(x_i, x_o) = \sqrt{P_{x_i}(-d)P_{x_o}(d)}, \quad (20)$$

where  $d = x_i - x_o$  denotes the vector between two points.  $P_x(d)$  is the local scattering profile defined at each surface point  $x$ . To reduce storage, it is further approximated with the 1D parameter  $r$  with  $P_x(d) \approx P_x(r)$ , where  $r = ||d||$ . Materials can be modified using this re-rendering approach by altering the local profile at each place, and the heterogeneity is shown with different profiles at  $x_i$  and  $x_o$ . Although it is not the primary target of SubEdit, the required storage of the captured material is much less than the original 4D  $R(x_i, x_o)$ , indicating significant data compression. These scattering profiles can even be decomposed [YTYM20], to further compress the measured data and reduce the memory cost.

The local profile  $P_x(r)$  in SubEdit is radially isotropic for lower storage. Therefore, the strong anisotropic scattering behaviour from narrow discontinuities is difficult to capture if the distance between the measurement points is not small enough. Since the main reason for isotropy is the 1D radial function  $P_x(r)$ , a direct solution is to use the 2D vector  $d = x_i - x_o$  instead of  $r$  (Equation 20) [SW13]. Recently, another approximation has been proposed [YYM20] to represent this anisotropic scattering profile. Following two assumptions that materials are optically thick, and the local profile changes



**Figure 8:** Human faces rendered with the biophysically based model. The illustration depicts 30- and 80-year-olds with a Mediterranean complexion for the male and Caucasian skin for the female. The small inset shows the different scattering profiles for each: A and B correspond to the elderly and young males, while C and D correspond to the elderly and young females. Image reproduced from Iglesias-Guitian et al. [IGAJG15].

smoothly, Equation (20) can be transformed and further approximated as a weighted sum of Gaussians [dLE07]:

$$\begin{aligned} R(x_i, x_o) &= \frac{1}{2} (P_{x_i}(r) + P_{x_o}(r)) \\ &= \frac{1}{2} \sum_h w_{x_i,h} G_{\sigma_h}(r) + \frac{1}{2} \sum_h w_{x_o,h} G_{\sigma_h}(r). \end{aligned} \quad (21)$$

To reflect the radial asymmetry, the weights  $w_{x_o,h}$  are kept on a map centred on  $x_i$ , with  $w_{x_o,h} = W_{x_i,h}(d)$ , where  $d = x_o - x_i$ . And  $w_{x_i,h}$  can be interpreted as  $W_{x_i,h}(0)$ . Compared to previous work [SW13], this re-rendering model can be calculated easily and applied in real time. However, the two assumptions used in this model may pose a difficulty that can increase relative inaccuracy. Although the accuracy can be improved by adding additional Gaussian terms, this comes with a longer rendering time.

**Human skin models.** Since human skin is used in a wide range of applications, it is the most popular field that requires special BSSRDF models. In the real world, skin has complex structures with many biophysically based characteristics, resulting in a multi-layer structure with complex appearances. The simplest structure considers the skin as a two-layer model (epidermis and dermis) [GHP\*08]. This basic skin model can be improved by adding an extra absorbing layer (between two scattering layers) [DWD\*08] for a more complete light interaction. The full light transport in this model requires summing a series of convolutions, which simulates the repeated propagation of light between two layers.

Light absorption in the skin is largely based on various types of chromophores (such as melanin, haemoglobin etc.), which produce different effects for different wavelengths of light. Compared to conventional RGB channels, spectral BSSRDF models prove to be better choices. After a spectral representation of skin properties [DJ06] is used for biophysical considerations, Iglesias-Guitian et al. [IGAJG15] considered skin aging with more skin layers, making this biophysically based model more complex but more accurate. With the wavelength  $\lambda$  in each layer,  $\sigma_a$  is represented as

the mixture of all the absorbing chromophores, and  $\sigma'_s$  is seen as the sum of the Rayleigh and Mie scattering. Due to highly sophisticated computation, this five-layer skin model is able to represent a wide variety of skin behaviour and has the maximum accuracy for skin rendering (Figure 8). Since the stratum corneum and hypodermis have always been ignored or simplified in the basic structure of the skin [DJ06], it can only represent a fraction of the changes mentioned above.

In biophysically based models, the absorption in each skin layer with the wavelength of light  $\lambda$  depends on different volume fractions of various chromophores. Unfortunately, except for complex biophysical parameters, it is also difficult to transform the representation from spectral space to RGB, which may lead to significant errors. However, considering the structure and physical composition of skin in particular, these models with spectral representations can re-render human skin with high accuracy. With measured data, it can reproduce a wide range of heterogeneous and complex skin appearances.

#### 4. Acceleration Methods

The use of BSSRDF models simplifies the sub-surface scattering problem from volume to surface. The radiance estimate at one shading point  $x_o$  must combine the contributions from all  $n$  surface points  $x_i$ . When all the  $n$  shading points on the surface are considered, rendering the entire object results in an  $O(n^2)$  algorithm (Equations 12–14). Jensen and Buhler [JB02] proposed that incident irradiance  $E$  at surface points could be pre-computed as irradiance samples and re-used. The usage of irradiance samples offers significant potential to speed up the estimation even further. Estimation at  $x_o$  from all irradiance samples can be accelerated by reducing the sample count, pre-computing light attenuation from the sample positions, or even storing them in 2D textures. Table 2 shows the main characteristics and the most classic work in each field. In particular, the computational complexity of each method only demonstrates how they select and evaluate irradiance samples. Although  $O(n^2)$  in texture methods means that at the shading point, each sample is

**Table 2:** Summary of the main characteristics of different acceleration methods (top: pre-computation, bottom: texture methods). The table shows their implementation passes (steps), computational complexity, rendering rates (from one dot to five dots, representing the improvement of efficiency) and main limitations.

Methods	Years	Steps (* denotes an extra pre-processing is in need)	Complexity	Rendering rates	Main limitation
<b>Point-patch</b> [JB02]	2002–2021	Sample construction → contribution gathering	$O(n \log n)$	● ○ ○ ○ ○	Large sample count
<b>PRT</b> [WTL05]	2002–2008	* coefficient projection → irradiance reconstruction	$O(n)$	● ● ● ○ ○	Static scene condition
<b>Pre-integration</b> [PB11]	2011–2015	* simple pixel shader	$O(n)$	● ● ● ● ●	Low accuracy
<b>Texture space</b> [dLE07]	2002–2007	Irradiance map → image filtering	$O(n^2)$	● ● ● ○ ○	Many-object problem
<b>Camera-light space</b> [DS03]	2003–2017	Translucent shadow map → view space connection	$O(n^2)$	● ● ● ○ ○	Many-light problem
<b>Screen space</b> [JSG09]	2009–2021	Main pass rendering → post processing	$O(n^2)$	● ● ● ○ ○	Information lost

considered, they can reduce samples by importance sampling or a suitable filtering kernel size.

#### 4.1. Pre-computation

Pre-computation is a kind of solution that pre-computes information and evaluation throughout the scene as much as possible, avoiding complex computation in the run time. Derived from Equation (13), the radiosity  $B_o$  can be evaluated using the discretized form  $B_o = \sum_k E_k F$  for simplicity, where  $F$  is the throughput factor. Since the repeated evaluation of incident irradiance  $E_k$  at surface points is a primary cause of the complexity,  $E_k$  is often pre-computed as irradiance samples and clustered, resulting in a  $O(n \log n)$  algorithm. In addition, the throughput factor  $F$  can also be pre-computed, decreasing the complexity from  $O(n^2)$  to even  $O(n)$ .

##### 4.1.1. Point-Patch methods

A simple option to accelerate sub-surface scattering calculations is to divide the scattering process into two passes [JB02]. In the first pass, several surface points are chosen as samples, and their irradiance is evaluated. The number of these surface samples should be large enough, and the sample distribution should be uniform, making samples sufficiently representative of the entire surface. All the locations and computed irradiance of these samples are pre-computed and saved. In the second pass, the contributions from all samples to the current shading point are summed using the BSSRDF model. Since all of the information is pre-computed, the sample irradiance can be directly obtained from stored data, significantly reducing computation time.

The unsophisticated approach described above has two major issues that need to be considered. The first is the sample representation. A simple point sample is always insufficient to represent an area, even if the distance between samples is small. Although illumination in a small area can be considered identical,  $R(r)$  carries with distance. As a result, instead of using the contribution from single points (Point), the contribution from the adjacent area (Patch) is applied [MKB\*03a, CHH03, SSWN13] as the throughput factor  $F$ . The second issue is the large sample count. In practice, since accumulating contributions from all samples for each shading point is expensive, the number of contributing samples at the current shading point should be reduced. A simple way is to consider only samples in the vicinity of outgoing points [HBV03, HV04, KLC07] owing

to the locality of subsurface scattering. However, sub-surface scattering is more often a global phenomenon. Using only surrounding samples may result in the loss of contributions from other areas. A better solution is clustering, which uses hierarchical structures to organize all samples [JB02, AWB08, WLLC14, MBG\*19]. In the following, we discuss these Patch and clustering techniques in detail:


**Patch techniques:** Different Point-Patch methods can be classified as Point-to-Point, Point-to-Patch and Patch-to-Patch, depending on the throughput factor  $F$ . The simplest Point-to-Point model directly considers the contributions between two points, where  $F$  equals  $R(x_i, x_o)$ . Since the throughput factor of a single point is inaccurate in representing an area, a better throughput factor is to integrate  $R$  from regions rather than points. Assuming the same illumination in a small region  $A$ , the resulting Point-to-Patch throughput factor represents the contribution from  $A_k$  to  $x_o$  [MKB\*03a]:

$$F(A_k, x_o) = \int_{A_k} R(x_i, x_o) dx_i. \quad (22)$$

A further consideration, the Patch-to-Patch model, takes both incident and outgoing regions into account for more complete representations. It is associated with the concept of radiosity technique [GTGB84], which is used for the illumination from a closed reflectance environment. The whole scene is represented as a set of surface patches. The radiosity of each patch is summed from all other patches, which is represented as a linear system, and can be solved by iterative processes (each iteration resembles a single bounce of radiosity transport). The diffuse scattering throughput factor  $F_{ij}$  between two patches  $P_i$  and  $P_j$  follows:

$$F_{ij} = \frac{1}{A_i} \int_{x_i \in P_i} \int_{x_j \in P_j} R(x_j, x_i) dx_j dx_i, \quad (23)$$

where  $A_i$  denotes the area of  $P_i$ . With pre-computation of  $F_{ij}$ , the irradiance, outgoing radiosity and ultimate radiance can be estimated sequentially during three passes [CHH03]. A better option is to keep the initial radiosity of the scene in the matrix  $L$  [SSWN13], and then update this matrix with the form factor matrices to get the final radiosity of the entire scene. Compared to the three-pass algorithm [CHH03], the matrix considered the entire scene with  $n$  light bounces (which included both reflection and scattering) rather than the object with only one bounce, leading to higher accuracy.



**Figure 9:** Multi-clustering in a multi-light problem. Only building one single cluster for irradiance samples  $b_j$  suffers from the increasing number of light sources (left). The same clustering technique can be applied to light sources to reduce the number of paths ( $l_i, b_j, e$ ) (right). Figure reproduced from Arbree et al. [AWB08].

**Clustering techniques:** Directly accumulating contributions from all irradiance samples can be computationally demanding due to their sheer number. However, given that samples from large distance areas contribute with minor variations, they can be grouped into a single sample, thereby reducing the overall sample count. Clustering is a technique that uses hierarchical structures to organize irradiance samples, with each node representing illumination information in all its child nodes, containing the total irradiance, the total area represented and the average location. The outgoing radiance is computed by traversing the hierarchical tree from the root. This traverse stops only when it reaches the leaf node or when the current node is accurate enough, which is checked with an error criterion. Since samples at high-level nodes are used to replace all their leaf nodes, clustering is always faster than adding the contributions from all irradiance samples directly.

At first, only an octree is used to organize samples [JB02]. Since the number of light sources has a linear relationship with the algorithmic complexity, this single cluster of irradiance samples is incapable of handling the scenario with several lights. Some multi-cluster structures are proposed to overcome this problem, where another cluster is often built to organize multiple light sources [WABG06]. The pixel intensity can be estimated using all the connections between two trees. To accommodate sub-surface scattering, a connection triple ( $l_i, b_j, e$ ) [AWB08] is constructed, including light samples  $l_i$ , irradiance samples  $b_j$  and a single eye sample  $e$ . The triple represents a light transport path from the light sources to the camera and is constructed from the selected cuts in two binary clusters (Figure 9). The final radiance can be estimated by summing the contributions of all the triples by about  $e$ .

The hierarchical tree is not the only structure employed in clustering approaches. The matrix format is another useful choice for storing contributions from light sources and surface points, as it can directly estimate the final pixel intensity using matrix multiplication [WLLC14]. This matrix representation is convenient for both clustering and estimation, but it needs more memory to represent extra connections among samples. Furthermore, the multi-light problem is not the only target of multi-clustering. In addition to light sources, the second tree structure can be employed for shading points in the screen space to accelerate rendering [MBG\*19]. The main target of the second tree is to use frequency analysis to


group similar pixels, which is significantly faster than single-tree searching [JB02]. However, it remains a challenge for small objects and high-frequency details.

Combining clustering techniques and Point-Patch methods is an effective way to accelerate the estimation process with fewer accuracy loss. Different multi-clustering structures can handle various scene conditions well, but some challenges remain. First, the choice of samples is a typical problem in clustering techniques. Using pre-defined and fixed irradiance samples is not a good idea, since an inappropriate sample distribution wastes samples in unimportant regions and leaves out details in high-frequency regions. The recent adaptive sampling scheme [NI21] generates irradiance samples dynamically rather than pre-defined. This mesh-based irradiance sampling assigns more samples to surface points with high contributions, resulting in a higher accuracy than other methods. Second, although Point-Patch methods are accurate enough, they are still costly for real-time applications due to the large sample count and 3D hierarchical structures. In the future, a suitable solution is still needed to reduce samples and package structures into simpler forms.

#### 4.1.2. Pre-computed radiance transfer

After pre-computed radiance transfer (PRT) is proposed [SKS02] to compress and store the complex data of environmental illumination, similar ideas were applied in sub-surface scattering. The introduction of PRT can reduce the costly computation of the scattering radiance from complex integral to simple dot products by projecting the irradiance [LGB\*02] or the light transport function [WTL05, XGL\*07, WCPL\*08] with different basis functions. These PRT-based schemes in sub-surface scattering are also described as two-pass algorithms similar to Point-Patch methods. First, the irradiance or light transport is projected onto the chosen basis and then in the second pass, Equation (13) can be re-constructed with the corresponding coefficients. The early methods [LGB\*02] projected the irradiance using spatial basis functions  $\psi_i(x)$ . Since projection coefficients and throughput factors are calculated during the preprocessing stage, the outgoing radiosity can be estimated directly with a matrix vector multiplication.

Instead of irradiance, the light transport function directly represents the integral contribution of surface areas, combining both the irradiance and the throughput factors. Because the outgoing radiosity can be obtained with only a weighted sum, it is a more convenient and suitable factor than irradiance in the PRT projection. In this case, the diffusion profile  $R$  is approximated as a linear sum of the basis functions. The challenge is that finding a fitting basis function is difficult. The wavelet basis [WTL05] cannot provide enough feedback for small changes. Although the non-uniform piecewise polynomial basis [XGL\*07] can solve it, further steps are required to secure a sequence of knots partitioning the spatial domain. A more flexible approach is to use principal component analysis (PCA) to determine the basis functions rather than pre-defined [WCPL\*08]. For a given  $R$ , it can be discretized as a matrix, where SVD can be applied. The PCA analysis allows for greater flexibility in selecting the basis functions, leading to smaller approximation errors. Although the storage of  $R$  in the pre-computation is substantial, it is not required during the rendering time.



**Figure 10:** Pre-integration of diffusion profile  $R$ . The effect of scattering  $D$  is pre-integrated on a spherical, then it is stored into a BRDF lookup indexed by curvature and  $N \cdot L$ . Figure reproduced from Penner and Borshukov [PB11].

Although PRT has limitations in terms of scene conditions, PRT-based methods are suitable for interactive material editing without complete recalculation due to the compact data representation. Since all projection coefficients are stored in textures, scattering results can be changed by updating textures promptly when users modify the material parameters [XGL\*07, WCPL\*08]. Since scene adjustments support both BSSRDF parameters and light conditions (relighting), PRT-based methods are more versatile than Point-Patch methods, as the latter requires constant geometry, light environment and material properties (static scene) after precomputations.

#### 4.1.3. Pre-integration

Since integrating irradiance over the whole surface is costly with the reach of graphics hardware in the early years, it motivates the approximation that the scattering effect can only be considered at the local shading point. Simon [Sim04] tried to wrap diffuse lighting to reduce the contrast of diffuse lighting, with the aim of simply simulating light scattering. An improved method is pre-estimating the integral of  $R$  in Equation (13), which is a completely different idea of pre-computation. The key assumption is that, with the flat boundary under uniform directional light, there are no visible scattering effects [PB11]. Therefore, visible scattering only comes from the mesh curvature, bumps of the normal and shadows. Assuming a spherical surface, the integral of  $R(r)$  reads:

$$D(\theta, r) = \frac{\int_{-\pi}^{\pi} \cos(\theta + x) \cdot R(2r \cdot \sin(x/2)) dx}{\int_{-\pi}^{\pi} R(2r \cdot \sin(x/2)) dx}, \quad (24)$$

where  $\theta$  is the angle formed by the normal  $N$  and light direction  $L$ , and  $r$  is the surface curvature. In the implementation, Equation (24) is pre-computed and stored in a 2D table (Figure 10). As a result, the scattering radiosity in Equation (13) can be estimated as  $D \cdot E(x_o)$  using a simple pixel shader like BRDF, without considering the entire surface.

Other similar curvature-dependent methods can be also seen in pre-integration [KDM11, KTM15]. Except for the difference in the evaluation of the integral of  $R$ , these methods are all independent of the scene conditions, therefore can be pre-computed and saved. The final result can only be estimated using a single pixel. Note

that integrating  $R$  with the mesh curvature is not the only way to approximate subsurface scattering, the consideration of hybrid normals [MHP\*07] can also produce high-quality results. These pre-integrated methods have the lowest cost among the subsurface scattering solutions, making them widely used in production, particularly in mobile applications. However, the computation with only one pixel severely restricts its accuracy, resulting in a significant divergence from the truth.


## 4.2. Texture methods

The storage of samples is at the core of most pre-computation techniques. The samples are typically organized into 3D hierarchical structures that are unsuitable for GPU architecture. To accelerate rendering using the GPU, several methods have been proposed that map information from 3D space to 2D textures, allowing all estimations to be calculated on textures. Although some features may be missing from the mapping, these methods have the potential to significantly improve efficiency, allowing for real-time rendering. Additionally, texture-based methods can reduce the storage requirements for massive samples of information.

### 4.2.1. Texture space

Constructing a one-to-one mapping is a direct way to store 3D surface points in 2D textures. It can be seen as a 2D parameterization of the object's surface. Early mapping techniques involved splitting the mesh into a large number of triangle chunks and projecting them onto the plane [LGB\*02], or using a multi-resolution meshed atlas (MMA) to store surface samples [CHH03]. These methods require a preprocessing stage to map the surface, and the seam artefact of discrete triangle packing should also be considered. A more straightforward way to represent the surface irradiance in a texture is the *irradiance map*, which can avoid the preprocessing of splitting the mesh. As Borshukov and Lewis [BL03] presented, the irradiance of surface points can be rendered to a texture map, and a GPU-based method [GSM04] is used to create this map by transforming all vertex positions into their 2D texture coordinates. Additional textures can be used to store additional information (such as normals and positions) for greater accuracy [GCZ\*08].

The main advantage of texture methods is their estimation of 2D textures, which can significantly accelerate rendering. Using the irradiance map, the integral on the surface (Equation 13) can be replicated in the texture and estimated in two ways: image filtering and importance sampling. Image filtering is best suited to the sum of Gaussian model [dLE07]. For complicated materials (such as three-layer skin) with a total of six Gaussians, each Gaussian convolution can be predicated on the previous result, resulting in various blur levels (Figure 11). The choice of sampling pattern is an important factor for higher efficiency in image filtering, where a straightforward idea is to reduce the sampling density due to the distance from the origin [DS03]. A two-ring sampling scheme [HBH09] is proposed to reduce samples. The inner ring is for mid-distance scattering, and the other is for wide-distance scattering. Only 13 samples with weights (including the centre point) can accurately estimate the scattering, rather than 70 samples in Gaussians.




**Figure 11:** Irradiance map in texture space with two blurred versions (left). The irradiance map is subjected to various Gaussian convolutions, resulting in blurred levels, and the final scattering effect is the weighted sum of these maps. Image reproduced from Jimenez et al. [JSG09].

MC importance sampling is another type that is distinct from image filtering. It chooses samples from the inverse CDF solution for fewer samples. Since most diffusion profiles  $R(r)$  lack analytical solutions for the inverse CDF, a good approximation of PDF  $p(x)$  [CLH\*08] or the inverse CDF [XOKN20] is often sought. Importance sampling allocates more samples to areas with greater contributions than constant sampling in image filtering. As a result, given the same number of samples, importance sampling is generally more accurate. However, a poor approximation may result in unsatisfactory results. Texture-space approaches encounter two common issues. The first problem is texture distortion. While the surface points are parameterized in the texture, the real distance between two points or the surface area of the point may be distorted, leading to incorrect diffusion computation. This problem should be solved by computing texture stretch [dLE07], or using a fit parameterization [YBS04] to modulate the filter. The second limitation is that the object model needs a parameterization to obtain texture coordinates of vertices, which may be lacking in some applications.

#### 4.2.2. Camera-light view

The irradiance map stores information from the entire surface of an object. However, not all surface irradiance is required for scattering estimation. Since all of the outgoing light originates from surface points with incident illumination, the outgoing radiance can be estimated if the positions and irradiance of only all incident points are known. Therefore, instead of storing the entire surface irradiance, only the surface points lit by light sources are necessary for accurate scattering estimation. This irradiance information from incident points is most commonly obtained from the light view. Subsequently, the scattering is estimated using the camera view, resulting in camera-light view approaches.

One early method of storing information from the light view is to use shadow maps [Wil78]. As it is rendered from the direction of the light source, it fits the storage of subsurface scattering nicely with only a little extension, resulting in the *translucent shadow maps* (TSM) [DS03]. Each TSM pixel stores the depth, 3D position, irradiance and surface normal of a single incident point. A two-pass algorithm is used, where the first pass generates the TSM from light directions and the second pass estimates the scattering radiosity from the camera direction. However, a significant issue with the TSM is the simple selection of incident points that contribute to




**Figure 12:** Diagrams of TSM (left) and splats (right). While TSM projects shading points in light view to obtain irradiance samples, splats are a better option to determine which irradiance samples can contribute to shading points.

the shading point  $x_o$ . Selecting surrounding points around projected points results in an unavoidable loss of contributions from other incident points. To address this problem, a novel concept called *splats* [SKP08] is proposed to establish a deeper connection between the light view and the camera view. The splats are a series of quads parallel to the screen (Figure 12) that are generated at each  $x_i$  from the light direction or based on triangular meshes [NRS14], representing the area of contributions from  $x_i$  to the camera. In the camera view, all the contributions of associated splats (where the projected point of  $x_o$  lies) are calculated and accumulated for each pixel. Leaving for quick selection of the corresponding incident points that contribute to  $x_o$ , the splat method is faster and more precise than the TSM.

When discussing camera-light view methods, two major issues should be addressed. First, since the irradiance of incident points is obtained from the light view, it suffers from the many-light problem, as seen in shadow maps. When the camera direction changes, all computations are re-estimated. A solution to this problem is the scattered radiosity map [DCFMB17], which stores rendering results from defined directions. By sampling and averaging these maps, results from any direction can be rendered to avoid repetitive estimations. The distribution of the incident point sample is the other concern. When it comes to splatting, both the number and size of splats should be considered carefully, or an unsuitable configuration may lead to patchy artefacts. The clustering technique is a great solution, which not only speeds up the rendering but also avoids patchy artefacts with adaptive sampling (Figure 13). In hierarchical splatting, all splats are organized in clusters due to the scene conditions, so the splats can be adaptively adjusted [CPZT12]. Unfortunately, problems still remain for sample selection, which is highly dependent on lighting conditions. Therefore, a sudden change in illumination may result in temporal flickering.

#### 4.2.3. Screen space

A key observation is that the invisible surface, such as the backface of objects or the surface outside the viewport, does not appear on the final screen, therefore it is unnecessary to compute the scattering on them [JSG09]. It means that moving the whole process of estimating from texture space to screen space is accurate enough to figure out how scattering works, which has better rendering efficiency than the texture space and light view. Solely using the information on the screen space, the estimation of subsurface



**Figure 13:** Comparison of splatting with and without adaptive sampling. (a) Ground truth. (b) Chen et al. [CPZT12] (40 fps). (c) Shah et al. [SKP08] with similar quality (3 fps). (d) Shah et al. [SKP08] with a comparable rendering speed. At this rate, the results have visual artefacts. Images reproduced from Chen et al. [CPZT12].

scattering is a post-processing step in the screen-space algorithm (Figure 14). The diffusion profile  $R$  is applied to the diffuse part of an image that has been rendered without scattering. Note that the pixels with high depths or gradients have a greater distance, even if they are close to the screen space. Therefore, extra stretch factors should be multiplied by the kernel width.

The advantage of the screen-space method is the independence of both the object and light number, resulting in noticeable acceleration. However, it is accompanied by a lack of information, leading to the following limitations: • The lack of geometry leads to a misjudgment of the real distance between two adjacent points on the screen. • The lack of backface irradiance makes evaluating translucency impossible. • For those points near the screen border, the lack of irradiance outside the viewport causes a mis-calculation.

Some methods attempt to solve these limitations by improving accuracy at the cost of additional storage or time. Backface irradiance can be saved in an extra buffer [Mik10], or stored at different depths

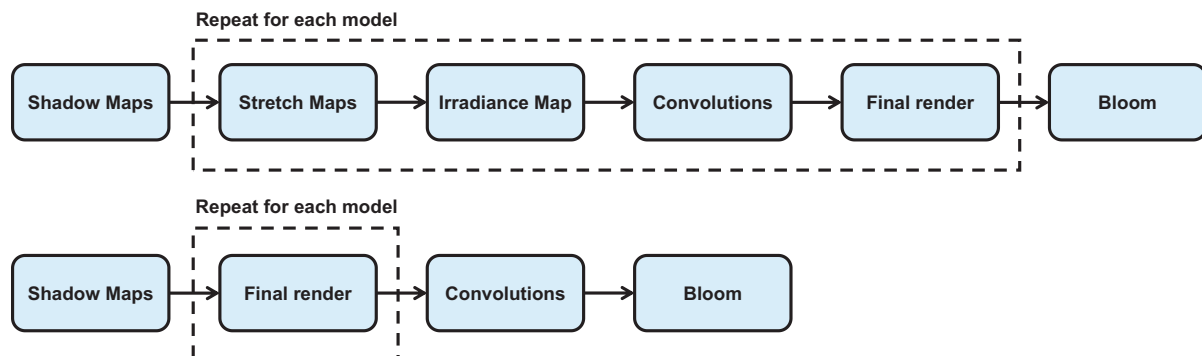
**Table 3:** Characteristics of three texture methods. Per-object, per-light: whether the texture is calculated for each object or each light. Translucency: whether the methods can show translucency effects. Distance: whether the distance between two points is distorted.

Methods	Per-object	Per-light	Translucency	Distance
Texture space	✓	✗	✗	Distortion
Camera-light view	✗	✓	✓	Real
Screen space	✗	✗	✗	Distortion

with numerous layers [MESG11]. These layers can substantially improve rendering accuracy, but the rendering time also increases as the layers increase. Since the screen-space method often fails in the silhouette of the surface, a cross-bilateral filter [Mik10] is proposed to handle it. Due to the small ratio between the projected area and the actual surface in screen space, insufficient sampling density can result in a lack of details. In this case, an additional kernel is a good solution to detect the silhouette.

In recent years, the improvement of screen-space methods has been supplemented with additional methodologies. While the screen-space method lacks translucency effects, the introduction of the splat into the screen space [MR20] can consider them, improving the evaluation of large-distance scattering. Xie *et al.* [XOKN20] combined adaptive sampling to improve the efficiency of the screen-space method. This temporal anti-aliasing (TAA)-based sampling adjusts the sample count with the variance of the rendered images, producing a significant speedup. However, when the history frame is not available, the estimation requires raising both the sample count and the pass time. A better solution is to use control variates (CV) [XO21], which tackle the sample domain instead of the shading domain to avoid this problem. The combination of the screen-space method and other methodologies still has great potential. Since the post-processing step is efficient enough, using additional calculations that can break the limitations is an important direction to further improve its accuracy. Meanwhile, the balance between accuracy and efficiency is also an important consideration.

**Discussion.** Each of the three texture-based approaches has its own pros and cons (Table 3). The irradiance map in texture-space methods considers all of the irradiance samples on the surface and



**Figure 14:** Schemes of the pipeline in texture-space (top) and screen-space (bottom) strategy. Figure reproduced from Jimenez et al. [JSG09].

should be generated for each object in the scene. Although it allows for a thorough search for the irradiance anywhere, in practice only a tiny kernel size is evaluated. These methods suffer from the many-object scene for increasing computation and storage. In addition, since not all surface points can contribute to visible points, they lead to a waste of time and memory. Camera-light view methods transform the storage of irradiance samples from the object surface to the light view, making them independent of the object count. Although they address the many-object problem, the new many-light problem is introduced because each light source requires its own buffers. To build the connection between the light view and the camera view, additional storage (such as position) is also required. However, this information allows them to estimate large-distance scattering well, whereas the texture-space and screen-space methods perform poorly for sampling range and distance distortion. Screen-space methods only estimate the scattering in the rendered images. These methods fit both the many-object and many-light scenes well, which have the highest efficiency of these three kinds. Despite the limitations caused by a lack of information in the screen space, screen-space approaches remain the dominant trend in real-time sub-surface scattering research due to their low cost.

## 5. Other Related Topics

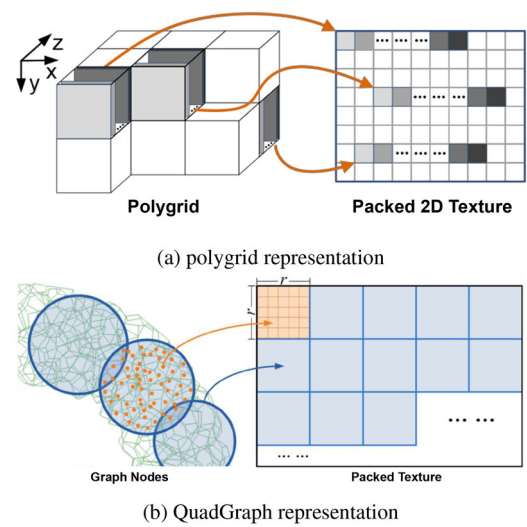
In this section, we explore some additional consequences of effective sub-surface scattering beyond multi-scattering. Discrete methods utilize discretization to extend diffusion from the surface to the volume. BSSRDF measurements capture parameters from real-world materials. On the other hand, translucency, single scattering and heterogeneous materials present significant challenges in the estimation of BSSRDF, with the former two often being neglected despite their critical role. Furthermore, heterogeneous materials pose a general issue in surface BSSRDF models. Additionally, we introduce some advanced machine learning methods, which open up a new avenue in sub-surface scattering.

### 5.1. Discrete methods

A significant limitation of BSSRDF models is that they evaluate scattering from surfaces only, which fails to represent volumetrically heterogeneous materials. This limitation can be overcome by the discrete method, a specialized method for efficient subsurface scattering. The discrete volumes are used for the first time to represent the objects [BLSS93]. As the DE [Sta95] is proposed, Schweiger *et al.* [SAHD95] explored a FE solution to it in medical imaging. The discrete method can be viewed as a compromise between the BSSRDF methods and path tracing, which is derived from the same DE as the BSSRDF models. However, in contrast to BSSRDF methods, it aims at volume instead of surface, which partitions the translucent object into discrete cells (grids or volumes). While representing volumetric parameters such as path tracing, some discrete methods can still maintain the interactive rates.

#### 5.1.1. Finite difference (FD)

In discrete methods, the critical step is to divide the object volume with a continuous domain into discrete cells. To achieve this, the FD method constructs volumetric grids to represent volumetric optical



**Figure 15:** Packing grid into textures. The texture representation makes the linear system be solved on GPU. Then a hierarchical structure is built, which can further accelerate the convergence of the system. Figure reproduced from Refs. [WZT\*08, WWH\*10].

properties. The radiance is stored in each grid node and propagated along the connections between nodes. One effective choice for volumetric grids is the polygrid [WZT\*08]. In a polygrid, the nodes are divided into boundary nodes and interior nodes. Each interior node has six connections and is uniformly distributed. The boundary nodes are connected to an interior node and aligned to the object surface. With the polygrid representation, the DE with boundary conditions can be rewritten as

$$\sum_{j=1}^6 w_{ji} \kappa(v_j) \phi(v_j) - \left( \sum_{j=1}^6 w_{ji} \right) \kappa(v_i) \phi(v_i) - u(v_i) \phi(v_i) = 0, \quad (25)$$

$$\phi(v'_i) + 2C\kappa(v'_i) \frac{\phi(v'_i) - \phi(v'_j)}{d_{ji}} = q(v'_i), \quad (26)$$

where  $v_i$  is the interior node and  $v_j$  is one of six nodes with weight  $w_{ji}$  that are directly connected to it.  $v'_i$  is the boundary node, and  $v'_j$  is the only interior node connected to it with a distance of  $d_{ji}$ .  $q(v'_i)$  represents the incident radiance. The representations of all the nodes can be seen as a system of linear equations that can be solved iteratively with the relaxation scheme [Sta95] until convergence.

Unfortunately, the building of polygrid is complex and only suitable for simple geometry without sharp features, which is a significant limitation. As a solution, QuadGraph [WWH\*10] offers a more effective volumetric grid structure. The QuadGraph is based on tetrahedrons instead of cubes in the polygrid. As a result, it has no limitations on object geometry, making it suitable for materials with arbitrary shapes. On the GPU, the process of solving linear problems can be accelerated (Figure 15). The scattering properties of the entire grid are initially packed into a set of 2D textures, and then a hierarchical structure is created based on the textures. Due to its texture storage, users can directly edit the scattering properties in the texture, making it suitable for material editing.



### 5.1.2. Finite element (FE)

Unlike the FD solution, light propagation in the FE method is based on discrete volumes rather than on grid nodes. Since it eliminates the need for grid construction, the FE method can simplify the pre-processing step and automatically avoid the restrictions of complex geometry. A complete FE scheme with full physical considerations comprises three major calculation results: diffusive source boundary condition (DSBC), finite element diffusion equation and render query function [AWB10]. These results are calculated using a tetrahedral mesh and correspond to three main steps of light transport in discrete volumes (incoming, propagating and outgoing), which are more accurate than the incorrect diffusion formulation in FD solutions [WZT\*08]. However, since these sophisticated procedures involve physical considerations that cannot be executed on the GPU, they require a high computational cost and take minutes to produce noise-free images.

Although the highly accurate FE scheme requires a significant time cost, it can be accelerated using some approximations. With a tetrahedral mesh, the discrete divergence at each vertex  $x_i$  can be approximated using the Voronoi region surrounding it [LSR\*12]. This accelerated method with heterogeneous material representation shortens the calculation with little loss of accuracy, making it suitable for artistic interactions, particularly for material cutting. However, since approximations are used for discrete divergence, appropriate aspect ratios of tetrahedrons are necessary, or the results may deviate from the accurate FE solution [AWB10].

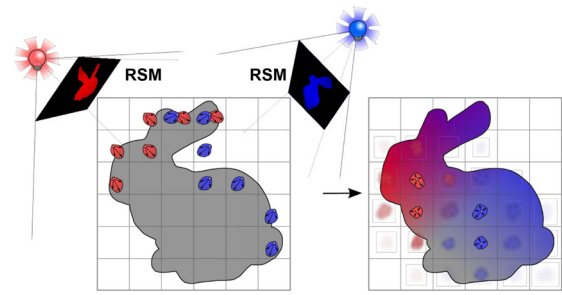
### 5.1.3. Scene-based discretization

All the FD and FE schemes described above require physical considerations to solve the discrete DE with high accuracy. Since the calculation is complex, a simpler method of discretization is to use cubes to represent the entire scene as voxelization. These approaches allow for the estimation of light transport in the medium using lattice-Boltzmann lighting (LBL), which confines light propagation to only six directions along the diagonals  $X$ ,  $Y$  and  $Z$  [BPB\*11].

Since the light propagation volume (LPV) [KD10] naturally fits the discrete representation of materials, it is extended from indirect illumination on the surface to light propagation in the medium, known as SSLPV [BCK\*11]. The SSLPV algorithm mainly consists of two steps (Figure 16). First, in the injection step, the reflecting shadow map (RSM) for each light source is generated to estimate the incident radiance. The intensity of each texel is injected into the LPV grid as an initialization. After that, the flux in each cell iteratively propagates to neighbouring cells. The results of each iteration are accumulated to estimate the final output. This naive SSLPV scheme can be further enhanced by incorporating extra single scattering [KJ14], resulting in a more complete and efficient framework.

## 5.2. BSSRDF measurement

Estimating the BSSRDF generally requires the use of optical properties. This is a significant concern for heterogeneous materials because the characterizing of optical properties at all surface points is complex. A common solution to this problem is to use a camera to capture several photographs of real-world materials and then



**Figure 16:** Two steps of SSLPV. In the inject step (left), virtual point light sources are created on the boundary of the object. Then light is propagated among neighbouring cells (right). Figure reproduced from Børlum et al. [BCK\*11].

establish the appropriate parameters. With the captured data, the parameters can be fitted or sought using different inverse computation models. In this section, we summarize these techniques with photographic estimation on a macroscopic scale. For more detailed and microscope models to acquire optical properties, we recommend the recent survey [FJM\*20].

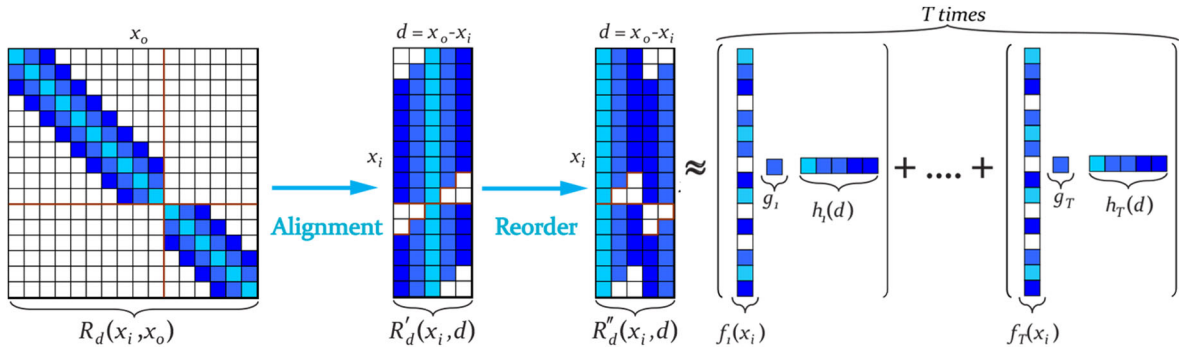
### 5.2.1. Parameters search

To measure the parameters of BSSRDF, a finely focused white beam and a three-CCD video camera provide an early and straightforward configuration [JMLH01]. This setup can be improved with a more complex configuration [GLL\*04, TWL\*05], or using various illumination patterns [TGL\*06, GHP\*08] to simplify and shorten the measuring procedure. For each  $x_i$ , the impulse responses at all other points  $x_o$  are recorded, eventually producing the entire 4D diffuse profile  $R(x_i, x_o)$ . Using the 4D diffuse profile  $R(x_i, x_o)$  directly is not a good choice. This is not only due to high frequencies and large storage requirements, but also because the measurement parameters need to be used for arbitrary geometries. Therefore, further optical properties should be explored, containing  $\sigma_s$ ,  $\sigma_t$ ,  $p(\theta)$  etc. Early methods frequently use the inverse equation of multi-scattering [WMP\*06, MSY09] or single scattering [NGD\*06] to predict properties in different translucent materials. All these methods strictly limit scene conditions, such as illumination, and attempt to do the search using solely inverse scattering models (such as the dipole).

The parameters in the BSSRDF measurement are only approximated using surface information. For a more physical and accurate solution to obtain optical properties, inverse MC volume rendering is used [DMZP14, SKCJ18]. The optimization problem related to inverse scattering is typically solved via gradient descent. It requires a large number of expensive rendering operations, resulting in high computational costs. Therefore, various methods, such as material dictionaries [GZB\*13] and neural networks [CLZ\*20] have been proposed to accelerate this process. Recent research [DLW\*22] hypothesized that these volumetric parameters can be substituted with BSSRDF profiles, leading to fewer samples per iteration. While inverse scattering is gradually replacing traditional BSSRDF measurement, BSSRDF may still be useful in inverse volume rendering.



**Figure 17:** The rendering objects with estimated BSSRDFs. The captured sub-surface scattering parameters are from a single image without any other prior information. From left to right: grape, orange soap and wax. Images reproduced from Munoz et al. [MES\*11].



**Figure 18:** The pipeline of compact  $R(x_i, x_o)$  representation. Figure reproduced from Kurt et al. [KÖP13].

**Image-based transfer.** In the parameter search, some special techniques aim at changing material properties in images [KRFB06, HR13], or transferring them to another object [MES\*11, LDZ\*11, GDH\*18]. Different from measurements, these techniques capture features from only single images. The key idea [KRFB06] is that both the shapes and light environment of a target object can be approximated from input images. Based on that, Munoz *et al.* [MES\*11] used single images to estimate BSSRDF. A linear system is defined and solved due to the one-to-one correspondence between the rendering image and a single input image. Even a single image with no further information can yield reasonable results (Figure 17). However, since the technique only uses information from the source image, either the lack of translucency information or the mis-estimation of 3D shapes in input images may result in an erroneous assessment. Although the recent work [TYS\*19] used the correlation between sub-bands and colour distribution to reproduce more accurate results, it ignored the potential light transport in the medium. Conclusively, transferring translucent material appearances is still challenging.

### 5.2.2. Data compression

The most significant disadvantage of using a 4D diffuse profile  $R(x_i, x_o)$  is that a big dataset is required for a complete heterogeneous representation of the entire object. These spatially varying parameters necessitate impractical memory, which is typically measured in the tens of gigabytes. However, with the

locality and symmetry of the diffuse profile, there are often a large number of repeated or empty elements in  $R(x_i, x_o)$ , which motivates research to compress it. Some approaches attempt to use a compact representation that compresses the data while sacrificing minimal accuracy. Figure 18 depicts a basic process to compress  $R(x_i, x_o)$ .

The measured  $R(x_i, x_o)$  is first represented as a matrix with rows for  $x_i$  and columns for  $x_o$  [PvBM\*06]. To compress the rows, a re-parameterized  $d = x_o - x_i$  is applied first. Then a transformation is used to solve the high frequencies in  $R'(x_i, d)$ . A simple transformation is to re-order the matrix element. Alternatively, a more appropriate and adaptable transformation can be developed using a genetic algorithm (GA) [Kur21]. After that, different matrix factorizations [PvBM\*06, KÖP13] can be applied, splitting the measured data into low-dimensional matrices. Since single factorization reduces the precision of matrix representation, repeated factorization is a better solution. After the factorization has been applied, the model error can be factorized again. This iterative procedure yields a multi-term factorization [KÖP13]:

$$R''(x_i, d) = \sum_{j=1}^T g_j f_j(x_i) h_j(d), \quad (27)$$

where  $T$  denotes the number of terms,  $g_j$  is the scalar core tensor and  $f_j(x_i)$  and  $h_j(d)$  are the univariate tensor functions used in each term. As  $T$  increases, the re-construction using univariate tensor functions becomes increasingly accurate.

### 5.2.3. Fabrication

On the other hand, besides rendering, BSSRDF measurements are also used in physical reproduction, especially in 3D printing. Compared to rendering, fabrication with 3D printers uses a set of basic materials to reproduce translucent materials, which involves both volumetric distribution and material properties. From the measurement of the target material, early works tried to combine a set of base materials with different optical properties to recover real material BSSRDF profiles. Two concurrent techniques used layer pruning [HFM\*10] and clustering [DWP\*10] to search for the optimal stacked material combinations. Papas *et al.* [PRJ\*13] adjusted pigment concentrations to best reproduce the appearance of subsurface scattering with only one clear base material, which makes it fit homogeneous sub-surface scattering better. However, a common limitation among these methods is that, with the small number of base materials used, they may fail to reproduce complex appearances with rich chromatic variations that are outside the colour gamut.


With the improvement of full-colour printing, recent research has focused on high-frequency features at high spatial resolutions. A key challenge lies in the inherent translucency of materials. Due to volumetric light transport between surface points on the coloured materials, simply reproducing colour textures leads to significant blurring of details. Although certain techniques [BAU15] use error diffusion halftoning algorithms to ignore the full BSSRDF of an object, a more accurate technique [ESZ\*17] employs a general MC simulation of volumetric light transport to optimize the volumetric material arrangement. A heuristic is that higher-absorbing material assignments should be steered toward the object surface to counteract heterogeneous scattering. This method is further extended to be capable of reproducing complex geometry [SRB\*19]. While MC simulation is expensive, the most recent technique [RSB\*21] introduces a deep neural network to predict the scattering within the highly heterogeneous medium, performing around two orders of magnitude faster than MC.

### 5.3. Translucency

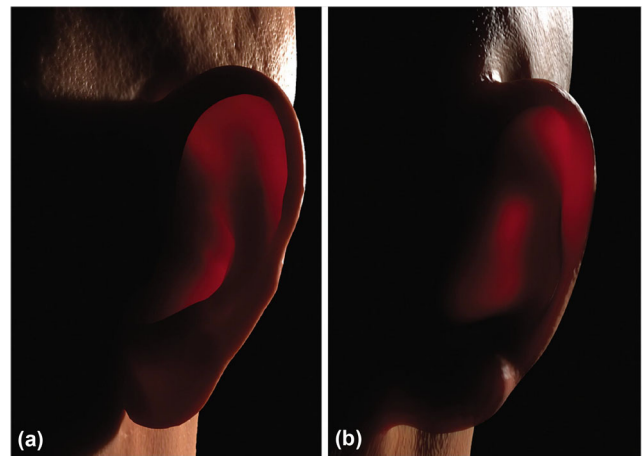
Translucency from the backface is an important effect for thin objects that allows light to transport through with minimal energy loss. The methods of pre-computation (Section 4.1), camera-light view (Section 4.2.2) and discretization (Section 5.1) can consider translucency naturally because they accumulate contributions from all incident points. However, texture methods in texture space and screen space cannot estimate translucency due to limited kernel size or the lack of information from the backface. Therefore, translucency should be added to scattering results in other ways.

One texture-based approach is to extend the TSM [DS03] to incorporate translucency. This involves storing the  $(u, v)$  coordinates and depths of the surface towards the light source and estimating translucency using image filtering on the TSM. Figure 19 illustrates this approach, where shading point  $C$  is projected onto the TSM to obtain its corresponding backface point  $A$ , and translucency is estimated using the irradiance convolution around  $A$ . The convolution kernel is re-written using the sums of Gaussians [dLE07]:

$$R(\sqrt{r^2 + d^2}) = \sum_{i=1}^k w_i e^{-d^2/v_i} G(v_i, r), \quad (28)$$




**Figure 19:** Texture-based estimation of translucency. The translucency at shading point  $C$  is approximated at point  $B$ . The point is projected into TSM, using image filtering to calculate outgoing radiance. Figure reproduced from d'Eon et al. [dLE07].



**Figure 20:** The comparison of thin ear rendering between (a) the multi-pole [DJ05] and (b) fast skin translucency [JWSG10]. Though the looks do not match exactly, they both have a similar red-to-black gradient. Image reproduced from Jimenez et al. [JWSG10].

where  $d$  is the thickness of the object. To simplify the calculation, the convolution at  $C$  is substituted by point  $B$  along the reversed normal of  $A$ . The thickness  $d$  can be directly estimated by  $d = m \cos(\theta)$ , where  $m$  is the depth at  $A$  stored in TSM, and  $\theta$  is the angle between the light direction and the normal at  $A$ . Based on observations that viewers pay more attention to visible surfaces, this method can be further simplified. The information on the backface can be approximated by surface properties on the front face [JWSG10]. Using approximations regarding front-face normal and irradiance, the translucency at shading point  $C$  can be estimated without saving the irradiance of the backface, reducing both the required memory and rendering time. This straightforward method can create realistic results for thin areas of human skin, such as hands and ears (Figure 20). However, the use of shadow mapping can result in some high-frequency characteristics that violate the low frequency of translucency.



**Figure 21:** Single scattering configurations. (a) Smooth Fresnel boundary, (b) diffuse Fresnel boundary.

While texture-based translucency is limited to point light sources with shadow maps, another way to estimate translucency is based on the mesh to solve these problems [MR17]. To separate translucency, the mesh surface  $A$  should be split into a back side  $A_-$  and a front side  $A_+$ , dividing all vertices  $V$  into two subsets,  $V_+ \subset A_+$  and  $V_- \subset A_-$ . Translucency at  $v_o \in A_+$  is calculated by summing all contributions from  $V_-$ . The mesh-based method is independent of the light source, which avoids the constraints of shadow maps (like many-light problems and environmental light). However, it introduces new limitations. When the mesh becomes complicated, the pre-computation of texture storage takes a long time and produces image quality issues that require additional memory to fix.

#### 5.4. Single scattering

Single scattering refers to light transport with only one scattering event in the medium. The division is proposed due to observations that the first-order solution  $L^{(1)}$  of outgoing radiance strongly depends on incident direction, whereas the difference (multi-scattering  $L^{(m)}$ ) between the full solution and  $L^{(1)}$  is approximately independent of it [HK93]. Although single scattering is often regarded as inconsequential in most techniques for optically thick materials, it nonetheless plays an important role in subsurface scattering, particularly near the incident point or with low extinction coefficients.

The single scattering term is typically evaluated as an integral of the incident radiance along the refracted outgoing ray (Figure 21a), as specified by Jensen *et al.* [JMLH01]:

$$\begin{aligned} L_o^{(1)}(x_o, \omega_o) &= \sigma_s(x_o) \int_{2\pi} Fp(\omega'_i \cdot \omega'_o) \int_0^\infty e^{-\sigma_{te}s} L_i(x_i, \omega_i) ds d\omega_i \\ &= \int_A \int_{2\pi} S^{(1)}(x_o, \omega_o; x_i, \omega_i) L_i(x_i, \omega_i) (n \cdot \omega_i) d\omega_i dA(x_i), \end{aligned} \quad (29)$$

where  $s$  is the path length of the refracted outgoing ray,  $F = F_t(\eta, \omega_o)F_t(\eta, \omega_i)$  is the product of the two Fresnel transmission terms, and  $\omega'_i$  and  $\omega'_o$  are the refracted incoming and outgoing directions.  $\sigma_{te} = \sigma_t(x_o) + G\sigma_t(x_i)$  is the combined extinction coefficient, where  $G$  is a geometry factor. The single-scattering BSSRDF  $S^{(1)}$  defined in the second line is a general form of single scattering. Note that with the refracted outgoing direction, not all the refracted incoming directions of  $L_i(x_i, \omega_i)$  can intersect it, implying that a delta function is contained in  $S^{(1)}$  implicitly.


This single scattering function is assumed to have smooth Fresnel boundary conditions and approximate the length of the beam inside the medium with a shadow ray, ignoring refraction. Both of these approximations result in low contributions, which is the main reason why single scattering is omitted in most approaches. Since many objects have diffuse extant boundaries rather than smooth, another function is presented in the same way as the multi-scattering (Equation 18), where the  $r(x, x_o(t))$  term is replaced by a single-scattering response [HCJ13]. The main difference is that the scattering points considered on the refracted incoming ray are independent of the outgoing direction (Figure 21b). Since diffuse single scattering is more significant, it is regarded as a non-negligible part of the diffusion profile. While smooth single scattering works better for objects with smooth surfaces, such as marble and milk, diffuse single scattering is better suited to materials with rough boundaries.

In practice, single scattering can be evaluated with MC integration. Several samples of distance are selected along the refracted ray. The radiance at each sampling point is computed and accumulated. Some early approaches attempted to speed up this procedure by using pre-computation [WTL05, XGL\*07]. However, unlike multi-scattering, it has the difficulty that  $L_o^{(1)}$  is highly related to  $\omega_o$ , which is unknown in the pre-computation step. The texture-based method [FPBB08] can analyse single scattering using ray marching, but it involves some planar approximations in object geometry. A more accurate method is to use mesh-based models [WZHB09]. With several samples  $V$  along the camera ray, all surface points  $P$  satisfying the condition are first found. Single scattering is then obtained with light point  $L$  by creating  $LPV$  routes and summing their contributions. This method can be improved with an inverse idea [Hol15]. In contrast to searching surface points for each camera sample, only the limited segment of the camera ray contributed by each individual triangle is found. This large improvement can deliver higher-quality results in less time (Figure 22), although it still falls short of interactive rates.


#### 5.5. Heterogeneous model

Rendering heterogeneous materials is always challenging, which most approaches that use surface BSSRDF strive to avoid. Despite the fact that spatially varying optical properties are defined on the surface, the BSSRDF models employed assume a homogeneous material with constant parameters when calculating at  $x_o$ . It means that these models can only consider optical properties at a single point and lead to homogeneous profiles, which is a general problem shared by all BSSRDF models. In the early years, a complicated sum of exponentials and a modulation pattern [FGCS05] are used to mimic heterogeneous scattering profiles. The simpler SubE-dit model [STPP09] accounts for the properties both at  $x_i$  and  $x_o$ , but it cannot consider the optical properties in areas between them.

To solve this problem, a heuristic idea is to approximate the region with variable optical characteristics between  $x_i$  and  $x_o$  by a homogeneous section with an average coefficient [SHK17]. First, a scattering region  $E$  is defined, which contains all light paths from  $x_i$  to  $x_o$  (Figure 23). The region  $E$  is modelled as an ellipse, where the length of the main axis is  $d = \|x_i - x_o\|$  and the minor axis can be seen as the spatial spread [PAT\*04]. The estimated averaged coefficients  $\bar{\sigma}_s$



**Figure 22:** Comparison between two mesh-based algorithms for single-scattering effects in approximately equal time. (a) Holzschuch et al. [Hol15] (5.8 min). (b) Walter et al. [WZHB09] (64 samples, 6.2 min). (c) Photon mapping (1.8 M photons, 5.6 min). (d) From left to right are zoom-ins of (a), (b) and (c). Images reproduced from Holzschuch [Hol15].



**Figure 23:** Parameter aggregation in heterogeneous media. The scattering region  $E$  predicts the distribution of light in heterogeneous media (left). The average coefficients are estimated from several samples in the ellipse (right). Figure reproduced from Sone et al. [SHK17].

with the region  $E$  are as follows:

$$\bar{\sigma}_s = \frac{1}{A_E} \int_E \sigma_s(x) dx, \quad (30)$$

where  $A_E$  denotes the region volume. By projecting points in  $E$  onto the surface, the coefficient samples  $\sigma_s(x)$  can be produced through the properties specified on the surface. With this method, Nakamoto and Koike [NK18] experimented with various BSSRDF models to render heterogeneous materials and proved that the better dipole [d12] is the most suitable.

Although the elliptical region is a good approximation, it may omit some light paths. Simply enlarging the region can address the problem, but it introduces a new issue: the loss of illumination features. Elek and Křivánek [EK18] proposed a better data-driven approach to represent the average coefficients between  $x_i$  and  $x_o$ . The parameter aggregation kernel  $k$  (equivalent to  $\bar{\sigma}_s$ ) is represented as a mixture of Gaussians. Since  $k_G$  can be stored in textures, it avoids the complex integration of MC in the aggregation of parameters. On the other hand, while Sone et al. [SHK17] only consider the global component without local effects, this full sub-surface transport is expressed as the product of two local components and one global component, which is more accurate with more local details.

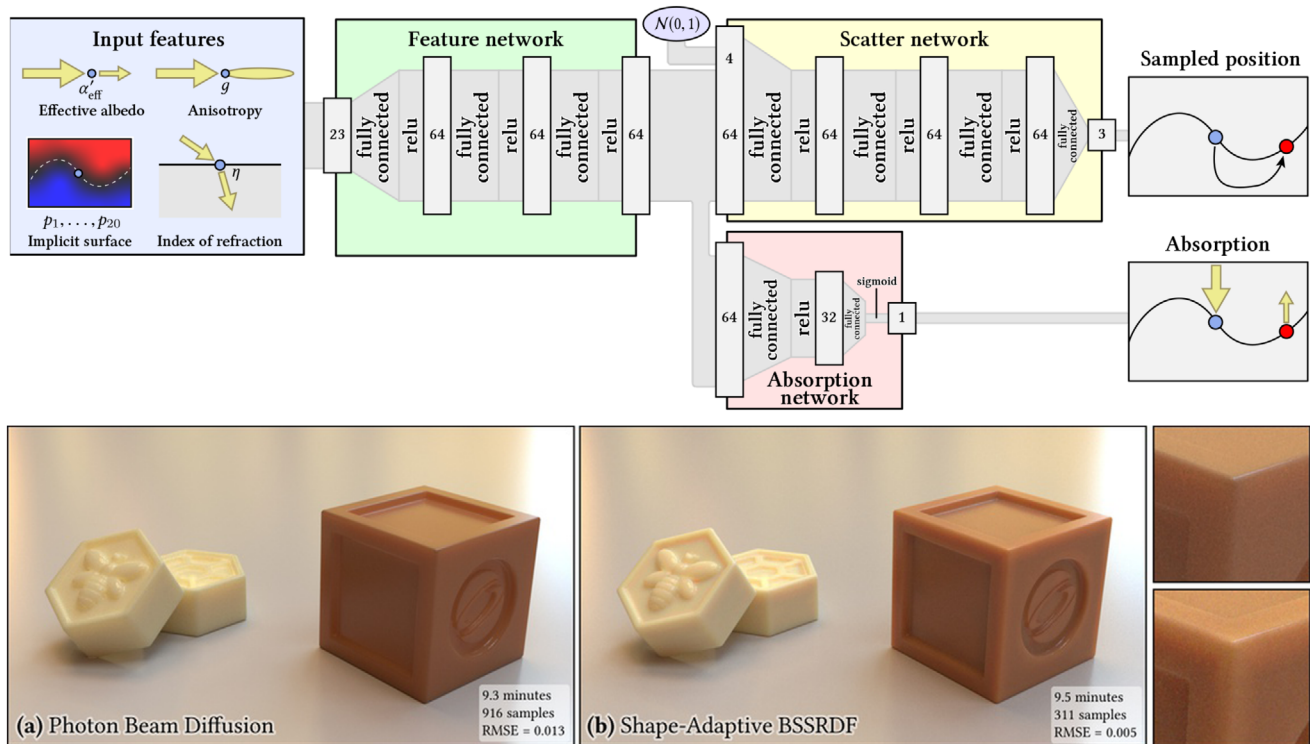
## 5.6. Machine learning

Using machine learning to solve subsurface scattering problems has emerged as an advanced idea in recent years. It does not rely on un-

derstanding the complex physical behaviours and rendering steps, but instead learns the potential relationships from the results. Although training takes some time, the calculation results can be obtained quickly with a trained network. Machine learning has been applied in many aspects of sub-surface scattering, and it still has great potential.

Many traditional BSSRDF models have limitations from diffusion theory, especially regarding the planarity of the surface and the isotropy of light transport. To solve these limitations, Vicini et al. [VKJ19] proposed a shape-adaptive learning model for sub-surface scattering. A total of three neural networks (feature network, scatter network and absorption network) are combined to learn from a reference dataset using MC simulation (Figure 24). The sample distribution and sample absorption at the outgoing point can be generated from the model using only a few features (optical properties and surface geometry) as input. This network-based model can produce results almost as accurate as path-tracing while using less time. However, several restrictions persist, such as the diffuse distribution of outgoing radiance, which may lead to errors in thin anisotropic materials. Leonard et al. [LHW21] presented a sequence of neural network-based conditional variational autoencoders (CVAEs) to solve this problem. The CVAEs can determine the direction of outgoing light from a reference distribution. This means that the local surface geometry does not have to be encoded explicitly, which is a significant improvement.

In addition to traditional usage of rendering materials with known properties, machine learning can also facilitate other aspects of sub-surface scattering. Che et al. [CLZ\*20] presented an inverse transport network (ITN) to estimate the scattering parameters from single images of an unknown translucent object, which is connected to a differentiable MC renderer. Compared to *Analysis by synthesis* in inverse scattering, ITN not only avoids costly optimization problems and the requirements of complete scene conditions, but also solves the difficulties of estimation accuracy without scene conditions and generalization. On the other hand, since MC light transport simulation is expensive in 3D printing fabrication, Rittig et al. [RSB\*21] leveraged a deep neural network to predict the surface appearance of scattering within a highly heterogeneous medium. With similar quality levels, it can significantly optimize 3D print preparation for full heterogeneous materials.



**Figure 24:** The architecture of the learned BSSRDF (top) and its rendering results of the translucent soap block (bottom). The input features are first transformed by a deterministic feature network (green), then shared by a scatter network (outputs 3D sampled position) and an absorption network (outputs radiance absorption). Compared to Photon Beam Diffusion [HCJ13], this learned BSSRDF produces a more realistic appearance and lower error. Figures reproduced from Vicini et al. [VKJ19].

As the combination of machine learning and computer graphics becomes popular in these years, we are pleased to see that some typical problems are gradually being solved. Although the use of machine learning in sub-surface scattering is still in its preliminary stage, more related work is expected to be seen in the future. With the proposal of differentiable MC rendering, applying it to inverse scattering may be a significant improvement. On the other hand, some remaining problems in sub-surface scattering (such as the sample distribution in Point-Patch methods) may also be well-addressed with neural networks.

## 6. Remaining Challenges

The approaches discussed in this study cover most of the subsurface scattering problems. Even though scattering effects can be produced with high accuracy or efficiency, challenges remain that need to be explored further, with some open problems. In this section, we discuss some limitations of current BSSRDF models and suggest future research directions that we believe are worth investigating.

**Limitations of BSSRDF.** Surface BSSRDF models typically rely on simplified assumptions to facilitate the difficult physical procedure of light transport between two points. However, in practice, these assumptions are frequently violated, leading to errors in the resulting images. That is the main reason for some researchers to

re-focus on MC path tracking [Kd14, KGV\*20]. We summarize the main two remaining limitations as follows.

- **Planar geometry.** The smooth surface assumption means the path of light transport can be approximately estimated using the direct distance between  $x_i$  and  $x_o$ . While this approximation is effective in many cases, significant artefacts arise on occasion in high-curvature areas. The surface complexity has a significant impact on two aspects: the path length of light transport and sample distribution. The real light transport in the medium is often longer than  $r = \|x_i - x_o\|$ . For sample selection, projecting samples onto the real surface [MKB\*03b, KKCF13] does not directly satisfy the importance of sampling with PDF. Even machine learning [VKJ19] struggles to obtain a fit PDF. Correctly adjusting the distance and finding a superior PDF on a complicated surface both remain challenges.
- **Isotropy.** The isotropy of light scattering is derived from a series of scattering events and shared by most BSSRDF models (defining  $g = 0$ ). Though scattering anisotropy can be approximated as isotropic with the similarity relation, it is insufficiently accurate, particularly for high anisotropy [FD17]. A broader perspective is that, except for  $g$ , both  $\sigma_s$  and  $\sigma_a$  are also related to radiance direction. The isotropy characteristic assumes that collision particles in the medium are spheres, such that the scattering behaviour is the same regardless of the direction from which the radiance

originates. However, most media in reality are anisotropic systems with non-spherical particles. Jakob *et al.* [JAM\*10] have explored this topic in detail with complex calculations. We hope that a more comprehensive and practical solution can be presented in the future.

**Material perception.** BSSRDF models typically use optical properties to describe potential light transport. However, the human visual system favours a more intuitive way of recognizing translucent materials, which involves research in material perception [GTHP21]. Although the perception of translucent objects depends on optical properties such as the phase function [GXZ\*13], these are complex and not intuitive to users. A more suitable visual perception of translucency comprises several low-level image cues [FB05], including highlights [Mot10], colour [LSX22], object size, edge [GWA\*15] and lighting direction [XWG\*14]. With these intuitive parameters, rather than physical ones, there is significant potential to develop a more user-friendly BSSRDF model. Since the Disney ‘principled’ BRDF [Bur12] serves as a good example, it may be a solution to use more observable and intuitive parameters, such as glossiness (related to highlight) and blurriness (related to edge), which may not necessarily be physically correct.

**Screen space methods.** As the most efficient type of approach for sub-surface scattering, these methods use only information from the screen space and are commonly used in production. However, their accuracy is limited by a lack of geometrical information, specifically the varied depths of selected samples in the screen space. Golubev *et al.* [Gol18] proposed a function extended from the normalized reflectance profile [Chr15] to consider the depth of the sample in the screen space. Unfortunately, it only meets energy conservation, which is not a strict representation of MC integration. We believe that a more detailed consideration of depth will be a great improvement in screen space approaches. On the other hand, the incident direction is also an important factor that influences the scattering effects that are often ignored in screen space methods. Although light directions can be easily determined in the render pipeline, accounting for them in subsurface scattering remains a challenge.

**Translucency model.** Most efforts on the translucency part of sub-surface scattering that we reviewed concentrate on acquiring information from the backface. Few of them pay attention to obtaining an accurate translucency profile  $T(x_i, x_o)$ , as opposed to the reflectance profile  $R(x_i, x_o)$ . Backface irradiance can be obtained using the methods described in Section 4.2. However, a persistent issue is that these methods use reflectance models [JMLH01] to estimate both reflectance and translucency, which is fundamentally unsuitable to address translucency. Although combining the two can be convenient, it is not a good enough method to be accurate. The multi-pole models [DJ05] and QD [dl11] models present transmittance profiles with sophisticated calculations that are hard to implement in practice. Although the weighted sums of Gaussians [dLE07] can fit them well, their parameters are not intuitive. Therefore, there is a need for a user-friendly, accurate and simple translucency model.

**Heterogeneous material measurement.** The acquisition of variable parameters is a complex process in capturing and re-rendering systems. Aside from varied settings, the scene conditions should also be strictly calibrated. Although some capture methods with

only single images have been proposed [MES\*11, CLZ\*20], none of them can be applied to heterogeneous materials. With the resurgence of MC path tracing in recent years, we believe that the representation of volumetric parameters will become an important issue in capture and re-rendering systems. Gkioulekas *et al.* [GLZ16] have taken the first steps toward it by providing a framework to evaluate heterogeneous inverse scattering from simulated measurements. Furthermore, using low-cost methods in the capture is also an aspect that should be considered. Such simplifications can even offer the possibility of extending the measurement from RGB to the full spectrum [IRN\*22].

**Extension to more materials.** Researchers are continually looking for a universal method to render most translucent materials. However, it is challenging for complicated and diverse structures made of special materials. Various modelling and rendering methods are proposed for weathered stone [DEJ\*99], quasi-homogeneous materials [TWL\*05], leaves [HKW07] and human skin [DWd\*08]. These methods have been adjusted to different degrees to better adapt to materials, such as modifications of parameters or special physical behaviours. We believe that specializing in scattering methods is a necessary step toward improving the rendering of specific materials. Yan *et al.* [YSJR17] used the dipole [JMLH01] to simulate multiple scattering between hair strands, demonstrating that BSSRDF models can be used for more than just translucent materials. We expect that the BSSRDF models will also perform well on other small-scale objects, such as granular materials [MPH\*15].

## 7. Conclusion

In this survey, we discuss various issues and solutions related to sub-surface scattering. Although some recent research has re-focused on path tracing, the main direction of radiance estimation still involves using surface information rather than volume in real-time applications. We have summarized different theoretical, empirical and re-rendering BSSRDF models. These models aim for precision, efficiency or special representation. Meanwhile, we have mainly reviewed acceleration techniques in sub-surface scattering. Further related topics can provide a more comprehensive description of efficient translucent material rendering and make the BSSRDF more practical in production. This study aims to be a valuable resource for researchers interested in sub-surface scattering. It can help readers quickly understand the latest developments in this field and the remaining challenges worth investigating. We hope that this report can provide some inspiration for researchers and lead to an improvement in the rendering of translucent materials.

## Acknowledgements

This work was supported by the National Key R &D Program of China (2023YFC3604500), National Natural Science Foundation of China (No. 62002010, 61932003, 62372026), and the Beijing Science and Technology Plan Project (No. Z221100007722001, Z221100007722004, Z231100005923039). This work was also partially supported by USA NSF IIS-1715985 and USA NSF IIS-1812606 (awarded to Hong QIN).

## References

- [AWB08] ARBREE A., WALTER B., BALA K.: Single-pass scalable subsurface rendering with lightcuts. *Computer Graphics Forum* 27, 2 (2008), 507–516.
- [AWB10] ARBREE A., WALTER B., BALA K.: Heterogeneous subsurface scattering using the finite element method. *IEEE Transactions on Visualization and Computer Graphics* 17, 7 (2010), 956–969.
- [BAU15] BRUNTON A., ARIKAN C. A., URBAN P.: Pushing the limits of 3D color printing: Error diffusion with translucent materials. *ACM Transactions on Graphics* 35, 1 (2015), 1–13.
- [BCK\*11] BØRLUM J., CHRISTENSEN B. B., KJELSDEN T. K., MIKKELSEN P. T., NOE K. Ø., RIMESTAD J., MOSEGAARD J.: SSLPV: Subsurface light propagation volumes. In *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics* (Vancouver, British Columbia, Canada, 2011), ACM New York, NY, USA, pp. 7–14.
- [BL03] BORSHUKOV G., LEWIS J.: Realistic human face rendering for “the matrix reloaded”. In *ACM SIGGRAPH 2003 Sketches & Applications* (San Diego, California, USA, 2003), ACM New York, NY, USA, pp. 1.
- [BLSS93] BLASI P., LE SAEC B., SCHLICK C.: A rendering algorithm for discrete volume density objects. *Computer Graphics Forum* 12, 3 (1993), 201–210.
- [BPB\*11] BERNABEI D., PATIL A. H., BANTERLE F., DI BENEDETTO M., GANOVELLI F., PATTANAİK S., SCOPIGNO R.: A parallel architecture for interactively rendering scattering and refraction effects. *IEEE Computer Graphics and Applications* 32, 2 (2011), 34–43.
- [BRM\*18] BITTERLI B., RAVICHANDRAN S., MÜLLER T., WRENNINGE M., NOVÁK J., MARSCHNER S., JAROSZ W.: A radiative transfer framework for non-exponential media. *ACM Transactions on Graphics* 37, 6 (2018), 1–17.
- [BSA12] BILLETER M., SINTORN E., ASSARSSON U.: Real-time multiple scattering using light propagation volumes. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Costa Mesa, California, USA, 2012), pp. 119–126.
- [Bur12] BURLEY B.: Physically-based shading at Disney. In *ACM SIGGRAPH 2012 Courses* (Los Angeles, California, USA, 2012), ACM, New York, NY, USA, pp. 1–27.
- [Cha60] CHANDRASEKHAR S.: *Radiative Transfer*. Dover Publications, New York, USA, 1960.
- [CHH03] CARR N. A., HALL J. D., HART J. C.: GPU algorithms for radiosity and subsurface scattering. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware* (San Diego, California, USA, 2003), vol. 3, Eurographics Association, pp. 51–59.
- [Chr15] CHRISTENSEN P. H.: An approximate reflectance profile for efficient subsurface scattering. In *ACM SIGGRAPH 2015 Talks* (Los Angeles, California, 2015), ACM, New York, NY, USA, pp. 1–1.
- [CLH\*08] CHANG C.-W., LIN W.-C., HO T.-C., HUANG T.-S., CHUANG J.-H.: Real-time translucent rendering using GPU-based texture space importance sampling. *Computer Graphics Forum* 27, 2 (2008), 517–526.
- [CLZ\*20] CHE C., LUAN F., ZHAO S., BALA K., GKIOULEKAS I.: Towards learning-based inverse subsurface scattering. In *2020 IEEE International Conference on Computational Photography* (Saint Louis, Missouri, USA, 2020), IEEE, pp. 1–12.
- [CPP\*05] CEREZO E., PÉREZ F., PUEYO X., SERON F. J., SILLION F. X.: A survey on participating media rendering techniques. *The Visual Computer* 21, 5 (2005), 303–328.
- [CPZT12] CHEN G., PEERS P., ZHANG J., TONG X.: Real-time rendering of deformable heterogeneous translucent objects using multiresolution splatting. *The Visual Computer* 28, 6 (2012), 701–711.
- [CZ67] CASE K. M., ZWEIFEL P. F.: *Linear Transport Theory*. Addison-Wesley Publishing Company, Reading, 1967.
- [DCFMB17] DAL CORSO A., FRISVAD J. R., MOSEGAARD J., BAERENTZEN J. A.: Interactive directional subsurface scattering and transport of emergent light. *The Visual Computer* 33, 3 (2017), 371–383.
- [d'E14] D'EON E.: A dual-beam 3D searchlight BSSRDF. In *ACM SIGGRAPH 2014 Talks* (Vancouver, British Columbia, Canada, 2014), ACM, New York, NY, USA, pp. 1.
- [DEJ\*99] DORSEY J., EDELMAN A., JENSEN H. W., LEGAKIS J., PEDERSEN H. K.: Modeling and rendering of weathered stone. In *ACM SIGGRAPH 1999 Conference Proceedings* (Los Angeles, California, USA, 1999), ACM, New York, NY, USA, pp. 225–234.
- [dI11] D'EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics* 30, 4 (2011), 1–14.
- [DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM Transactions on Graphics* 24, 3 (2005), 1032–1039.
- [DJ06] DONNER C., JENSEN H. W.: A spectral BSSRDF for shading human skin. In *Proceedings of the 17th Eurographics Conference on Rendering Techniques* (Nicosia, Cyprus, 2006), Eurographics Association, pp. 409–418.
- [DJ07] DONNER C., JENSEN H. W.: Rendering translucent materials using photon diffusion. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Grenoble, France, 2007), Eurographics Association, pp. 243–251.



- [dLE07] D'EON E., LUEBKE D., ENDERTON E.: Efficient rendering of human skin. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Grenoble, France, 2007), Eurographics Association, pp. 147–157.
- [DLR\*09] DONNER C., LAWRENCE J., RAMAMOORTHY R., HACHISUKA T., JENSEN H. W., NAYAR S.: An empirical BSSRDF model. *ACM Transactions on Graphics* 28, 3 (2009), 1–10.
- [DLW\*22] DENG X., LUAN F., WALTER B., BALA K., MARSCHNER S.: Reconstructing translucent objects using differentiable rendering. In *ACM SIGGRAPH 2022 Conference Proceedings* (Vancouver, British Columbia, Canada, 2022), ACM, New York, NY, USA, pp. 1–10.
- [DMZP14] DONG B., MOORE K. D., ZHANG W., PEERS P.: Scattering parameters and surface normals from homogeneous translucent materials using photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Columbus, Ohio, USA, 2014), IEEE, pp. 2291–2298.
- [DS03] DACHSBACHER C., STAMMINGER M.: Translucent shadow maps. In *Proceedings of the 14th Eurographics Workshop on Rendering* (Leuven, Belgium, 2003), Eurographics Association, pp. 197–201.
- [DWD\*08] DONNER C., WEYRICH T., D'EON E., RAMAMOORTHY R., RUSINKIEWICZ S.: A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Transactions on Graphics* 27, 5 (2008), 1–12.
- [DWP\*10] DONG Y., WANG J., PELLACINI F., TONG X., GUO B.: Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics* 29, 4 (2010), 1–10.
- [DWWH20] DENG H., WANG B., WANG R., HOLZSCHUCH N.: A practical path guiding method for participating media. *Computational Visual Media* 6, 1 (2020), 37–51.
- [d12] D'EON E.: *A better dipole*. Tech. Rep. 2012. <https://www.eugenedeon.com>.
- [d13] D'EON E.: Rigorous asymptotic and moment-preserving diffusion approximations for generalized linear Boltzmann transport in arbitrary dimension. *Transport Theory and Statistical Physics* 42, 6-7 (2013), 237–297.
- [d18] D'EON E.: A reciprocal formulation of nonexponential radiative transfer. 1: Sketch and motivation. *Journal of Computational and Theoretical Transport* 47, 1-3 (2018), 84–115.
- [d19] D'EON E.: A reciprocal formulation of nonexponential radiative transfer. 2: Monte Carlo estimation and diffusion approximation. *Journal of Computational and Theoretical Transport* 48, 6 (2019), 201–262.
- [EK18] ELEK O., KŘIVÁNEK J.: Towards a principled kernel prediction for spatially varying BSSRDFs. In *Proceedings of the Eurographics 2018 Workshop on Material Appearance Modeling* (2018), Eurographics Association, pp. 45–49.
- [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* 36, 6 (2017), 1–15.
- [EY36] ECKART C., YOUNG G.: The approximation of one matrix by another of lower rank. *Psychometrika* 1, 3 (1936), 211–218.
- [FB05] FLEMING R. W., BÜLTHOFF H. H.: Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception* 2, 3 (2005), 346–382.
- [FD17] FREDERICKX R., DUTRÉ P.: A forward scattering dipole model from a functional integral approximation. *ACM Transactions on Graphics* 36, 4 (2017), 1–13.
- [FGCS05] FUCHS C., GOESELE M., CHEN T., SEIDEL H.-P.: An empirical model for heterogeneous translucent objects. In *ACM SIGGRAPH 2005 Sketches* (Los Angeles, California, 2005), ACM, New York, NY, USA, pp. 24.
- [FHK14] FRISVAD J. R., HACHISUKA T., KJELDSSEN T. K.: Directional dipole model for subsurface scattering. *ACM Transactions on Graphics* 34, 1 (2014), 1–12.
- [FJM\*20] FRISVAD J. R., JENSEN S. A., MADSEN J. S., CORREIA A., YANG L., GREGERSEN S. K. S., MEURET Y., HANSEN P.-E.: Survey of models for acquiring the optical properties of translucent materials. *Computer Graphics Forum* 39, 2 (2020), 729–755.
- [FPBB08] FRANÇOIS G., PATTANAIK S., BOUATOUCH K., BRETON G.: Subsurface texture mapping. *IEEE Computer Graphics and Applications* 28, 1 (2008), 34–42.
- [FPW92] FARRELL T. J., PATTERSON M. S., WILSON B.: A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Medical Physics* 19, 4 (1992), 879–888.
- [GCZ\*08] GONG Y., CHEN W., ZHANG L., ZENG Y., PENG Q.: GPU-based rendering for deformable translucent objects. *The Visual Computer* 24, 2 (2008), 95–103.
- [GDH\*18] GITLINA Y., DHILLON D. S. J., HANSEN J., PAI D. K., GHOSH A.: Practical measurement-based spectral rendering of human skin. In *ACM SIGGRAPH 2018 Posters* (Vancouver, British Columbia, Canada, 2018), ACM, New York, NY, USA, pp. 1–2.
- [GHP\*08] GHOSH A., HAWKINS T., PEERS P., FREDERIKSEN S., DEBEVEC P.: Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics* 27, 5 (2008), 1–10.
- [GLL\*04] GOESELE M., LENSCH H., LANG J., FUCHS C., SEIDEL H.-P., MARKS J.: Disco-acquisition of translucent objects. *ACM Transactions on Graphics* 23 (2004), 835–844.
- [GLZ16] GKIIOULEKAS I., LEVIN A., ZICKLER T.: An evaluation of computational imaging techniques for heterogeneous inverse scattering. In *European Conference on Computer Vision* (Amsterdam, The Netherlands, 2016), Springer, pp. 685–701.

- [Gol18] GOLUBEV E.: Efficient screen-space subsurface scattering using Burley's normalized diffusion in real-time. In *ACM SIGGRAPH 2018 Courses: Advances in Real-Time Rendering in Games Course* (Vancouver, British Columbia, Canada, 2018), ACM, New York, NY, USA, pp. 1–62.
- [Gro56] GROSJEAN C.: A high accuracy approximation for solving multiple scattering problems in infinite homogeneous media. *Il Nuovo Cimento (1955-1965)* 3, 6 (1956), 1262–1275.
- [GSM04] GOSSELIN D., SANDER P. V., MITCHELL J. L.: Real-time texture-space skin rendering. In *ShaderX3: Advanced Rendering Techniques in DirectX and OpenGL*, vol. 3. Charles River Media, Hingham (2004), pp. 171–184.
- [GTGB84] GORAL C. M., TORRANCE K. E., GREENBERG D. P., BATAILLE B.: Modeling the interaction of light between diffuse surfaces. *ACM SIGGRAPH Computer Graphics* 18, 3 (1984), 213–222.
- [GTHP21] GIGILASHVILI D., THOMAS J.-B., HARDEBERG J. Y., PEDERSEN M.: Translucency perception: A review. *Journal of Vision* 21, 8 (2021), 4.
- [GWA\*15] GKIOULEKAS I., WALTER B., ADELSON E. H., BALA K., ZICKLER T.: On the appearance of translucent edges. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Boston, MA, USA, 2015), IEEE, pp. 5528–5536.
- [GXZ\*13] GKIOULEKAS I., XIAO B., ZHAO S., ADELSON E. H., ZICKLER T., BALA K.: Understanding the role of phase function in translucent appearance. *ACM Transactions on Graphics* 32, 5 (2013), 1–19.
- [GZB\*13] GKIOULEKAS I., ZHAO S., BALA K., ZICKLER T., LEVIN A.: Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics* 32, 6 (2013), 1–13.
- [HBH09] HABLE J., BORSHUKOV G., HEJL J.: Fast skin shading. In *ShaderX7: Advanced Rendering Techniques*, vol. 7. Charles River Media, Hingham (2009), pp. 161–173.
- [HBV03] HAO X., BABY T., VARSHNEY A.: Interactive subsurface scattering for translucent meshes. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Monterey, California, USA, 2003), ACM, New York, NY, USA, pp. 75–82.
- [HCJ13] HABEL R., CHRISTENSEN P. H., JAROSZ W.: Photon beam diffusion: A hybrid Monte Carlo method for subsurface scattering. *Computer Graphics Forum* 32, 4 (2013), 27–37.
- [HFM\*10] HAŠAN M., FUCHS M., MATUSIK W., PFISTER H., RUSINKIEWICZ S.: Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics* 29, 4 (2010), 1–10.
- [HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *ACM SIGGRAPH 1993 Conference Proceedings* (Anaheim, California, USA, 1993), ACM, New York, NY, USA, pp. 165–174.
- [HKW07] HABEL R., KUSTERNIG A., WIMMER M.: Physically based real-time translucency for leaves. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Grenoble, France, 2007), Eurographics Association, pp. 253–263.
- [Hol15] HOLZSCHUCH N.: Accurate computation of single scattering in participating media with refractive boundaries. *Computer Graphics Forum* 36 (2015), 48–59.
- [HR13] HAŠAN M., RAMAMOORTHI R.: Interactive albedo editing in path-traced volumetric materials. *ACM Transactions on Graphics* 32, 2 (2013), 1–11.
- [HV04] HAO X., VARSHNEY A.: Real-time rendering of translucent meshes. *ACM Transactions on Graphics* 23, 2 (2004), 120–142.
- [IGAJG15] IGLESIAS-GUITIAN J. A., ALIAGA C., JARABO A., GUTIERREZ D.: A biophysically-based model of the optical properties of skin aging. *Computer Graphics Forum* 34, 2 (2015), 45–55.
- [IRN\*22] ISER T., RITTIG T., NOGUÉ E., NINDEL T. K., WILKIE A.: Affordable spectral measurements of translucent materials. *ACM Transactions on Graphics* 41, 6 (2022), 1–13.
- [JAG18] JARABO A., ALIAGA C., GUTIERREZ D.: A radiative transfer framework for spatially-correlated materials. *ACM Transactions on Graphics* 37, 4 (2018), 1–13.
- [JAM\*10] JAKOB W., ARBREE A., MOON J. T., BALA K., MARSCHNER S.: A radiative transfer framework for rendering materials with anisotropic structure. *ACM Transactions on Graphics* 29, 4 (2010), 1–13.
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics* 21, 3 (2002), 576–581.
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *ACM SIGGRAPH 1998 Conference Proceedings* (Orlando, Florida, USA, 1998), ACM, New York, NY, USA, pp. 311–320.
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *ACM SIGGRAPH 2001 Conference Proceedings* (Los Angeles, California, USA, 2001), ACM, New York, NY, USA, pp. 511–518.
- [JSG09] JIMENEZ J., SUNDSTEDT V., GUTIERREZ D.: Screen-space perceptual rendering of human skin. *ACM Transactions on Applied Perception* 6, 4 (2009), 1–15.
- [JWSG10] JIMENEZ J., WHELAN D., SUNDSTEDT V., GUTIERREZ D.: Real-time realistic skin translucency. *IEEE Computer Graphics and Applications* 30, 4 (2010), 32–41.
- [JZJ\*15] JIMENEZ J., ZSOLNAI K., JARABO A., FREUDE C., AUZINGER T., WU X.-C., VON DER PAHLEN J., WIMMER M., GUTIERREZ D.: Separable subsurface scattering. *Computer Graphics Forum* 34, 6 (2015), 188–197.

- [KD10] KAPLANYAN A., DACHSBACHER C.: Cascaded light propagation volumes for real-time indirect illumination. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Washington, District of Columbia, USA, 2010), ACM, New York, NY, USA, pp. 99–107.
- [Kd14] KŘIVÁNEK J., D'EON E.: A zero-variance-based sampling scheme for Monte Carlo subsurface scattering. In *ACM SIGGRAPH 2014 Talks* (Vancouver, Canada, 2014), ACM, New York, NY, USA, pp. 1.
- [KDM11] KUBO H., DOBASHI Y., MORISHIMA S.: Curvature-dependent reflectance function for interactive rendering of subsurface scattering. *International Journal of Virtual Reality* 10, 1 (2011), 45–51.
- [KF12] KULLA C., FAJARDO M.: Importance sampling techniques for path tracing in participating media. *Computer Graphics Forum* 31, 4 (2012), 1519–1528.
- [KGV\*20] KELLER A., GRITTMANN P., VORBA J., GEORGIEV I., ŠIK M., D'EON E., GAUTRON P., VÉVODA P., KONDAPANENI I.: Advances in Monte Carlo rendering: the legacy of Jaroslav Křivánek. In *ACM SIGGRAPH 2020 Courses* (Virtual Event, USA, 2020), ACM, New York, NY, USA, pp. 1–366.
- [KJ14] KOA M. D., JOHAN H.: ESLPV: Enhanced subsurface light propagation volumes. *The Visual Computer* 30, 6 (2014), 821–831.
- [KKCF13] KING A., KULLA C., CONTY A., FAJARDO M.: BSSRDF importance sampling. In *ACM SIGGRAPH 2013 Talks* (Anaheim, California, USA, 2013), ACM, New York, NY, USA, pp. 1.
- [KLC07] KENG S.-L., LEE W.-Y., CHUANG J.-H.: An efficient caching-based rendering of translucent materials. *The Visual Computer* 23, 1 (2007), 59–69.
- [KÖP13] KURT M., ÖZTÜRK A., PEERS P.: A compact tucker-based factorization model for heterogeneous subsurface scattering. In *Theory and Practice of Computer Graphics* (Bath, United Kingdom, 2013), The Eurographics Association, pp. 85–92.
- [KRFB06] KHAN E. A., REINHARD E., FLEMING R. W., BÜLTHOFF H. H.: Image-based material editing. *ACM Transactions on Graphics* 25, 3 (2006), 654–663.
- [KTM15] KUBO H., TOKOI K., MUKAIGAWA Y.: Real-time rendering of subsurface scattering according to translucency magnitude. In *ACM SIGGRAPH 2015 Posters* (Los Angeles, California, USA, 2015), ACM, New York, NY, USA, pp. 1.
- [Kur21] KURT M.: GenSSS: A genetic algorithm for measured subsurface scattering representation. *The Visual Computer* 37, 2 (2021), 307–323.
- [Lam60] LAMBERT J. H.: Photometria sive de mensura et gradibus luminis, colorum et umbrae. *Sumptibus Vidvae. Typis CP Delfenssen*. E. Klett, Augustae Vindelicorum, 1760.
- [LDZ\*11] LI C., DONG W., ZHOU N., ZHANG X., PAUL J.-C.: Translucent material transfer based on single images. In *SIGGRAPH Asia 2011 Sketches* (Hong Kong, China, 2011), ACM, New York, NY, USA, pp. 1–2.
- [LGB\*02] LENSCH H. P., GOESELE M., BEKAERT P., KAUTZ J., MAGNOR M. A., LANG J., SEIDEL H.-P.: Interactive rendering of translucent objects. In *10th Pacific Conference on Computer Graphics and Applications* (Beijing, China, 2002), IEEE, pp. 214–224.
- [LHW21] LEONARD L., HOEHLIN K., WESTERMANN R.: Learning multiple-scattering solutions for sphere-tracing of volumetric subsurface effects. *Computer Graphics Forum* 40, 2 (2021), 165–178.
- [LPT05] LI H., PELLACINI F., TORRANCE K. E.: A hybrid Monte Carlo method for accurate and efficient subsurface scattering. In *Proceedings of the 16th Eurographics Conference on Rendering Techniques* (Konstanz, Germany, 2005), Eurographics Association, pp. 283–290.
- [LQTF21] LIANG S., QIN Z., TAI Y., FANG D.: Rendering of 3D models based on BSSRDF: A survey. In *3D Imaging Technologies Multi-dimensional Signal Processing and Deep Learning*. Springer, Singapore (Kunming, China, 2021), pp. 283–289.
- [LSR\*12] LI D., SUN X., REN Z., LIN S., TONG Y., GUO B., ZHOU K.: Transcut: Interactive rendering of translucent cutouts. *IEEE Transactions on Visualization and Computer Graphics* 19, 3 (2012), 484–494.
- [LSX22] LIAO C., SAWAYAMA M., XIAO B.: Crystal or jelly? Effect of color on the perception of translucent materials with photographs of real-world objects. *Journal of Vision* 22, 2 (2022), 6.
- [MBG\*19] MILAENEN D., BELCOUR L., GUERTIN J.-P., HACHISUKA T., NOWROUZEZAHRAI D.: A frequency analysis and dual hierarchy for efficient rendering of subsurface scattering. In *Proceedings of the 45th Graphics Interface Conference on Proceedings of Graphics Interface 2019* (Kingston, Canada, 2019), Canadian Human-Computer Communications Society, pp. 1–7.
- [MES\*11] MUNOZ A., ECHEVARRIA J. I., SERON F. J., LOPEZ-MORENO J., GLENCROSS M., GUTIERREZ D.: BSSRDF estimation from single images. *Computer Graphics Forum* 30, 2 (2011), 455–464.
- [MESG11] MUNOZ A., ECHEVARRIA J. I., SERON F. J., GUTIERREZ D.: Convolution-based simulation of homogeneous subsurface scattering. *Computer Graphics Forum* 30, 8 (2011), 2279–2287.
- [MHP\*07] MA W.-C., HAWKINS T., PEERS P., CHABERT C.-F., WEISS M., DEBEVEC P. E.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of the 18th Eurographics Conference on*

- Rendering Techniques* (Grenoble, France, 2007), Eurographics Association, pp. 183–194.
- [Mik10] MIKKELSEN M. S.: *Skin rendering by pseudo-separable cross bilateral filtering*. Tech. Rep. Naughty Dog Inc, 2010.
- [MKB\*03a] MERTENS T., KAUTZ J., BEKAERT P., SEIDEL H.-P., VAN REETH F.: Interactive rendering of translucent deformable objects. In *ACM SIGGRAPH 2003 Sketches & Applications* (San Diego, California, USA, 2003), ACM, New York, NY, USA, pp. 1.
- [MKB\*03b] MERTENS T., KAUTZ J., BEKAERT P., VAN REETH F., SEIDEL H.-P.: Efficient rendering of local subsurface scattering. In *11th Pacific Conference on Computer Graphics and Applications* (Canmore, Alberta, Canada, 2003), IEEE, pp. 51–58.
- [Mot10] MOTOYOSHI I.: Highlight–shading relationship as a cue for the perception of translucent and transparent materials. *Journal of Vision* 10, 9 (2010), 6.
- [MPH\*15] MENG J., PAPAS M., HABEL R., DACHSBACHER C., MARSCHNER S., GROSS M., JAROSZ W.: Multi-scale modeling and rendering of granular materials. *ACM Transactions on Graphics* 34, 4 (2015), 1–13.
- [MR17] MAISCH S., ROPINSKI T.: Spatial adjacency maps for translucency simulation under general illumination. *Computer Graphics Forum* 36, 2 (2017), 443–453.
- [MR20] MAISCH S., ROPINSKI T.: Interactive subsurface scattering for materials with high scattering distances. *Computer Graphics Forum* 39, 6 (2020), 465–479.
- [MSG05] MENON S., SU Q., GROBE R.: Generalized diffusion solution for light scattering from anisotropic sources. *Optics Letters* 30, 12 (2005), 1542–1544.
- [MSY09] MUKAIGAWA Y., SUZUKI K., YAGI Y.: Analysis of subsurface scattering based on dipole approximation. *IPSJ Transactions on Computer Vision and Applications* 1 (2009), 128–138.
- [NGD\*06] NARASIMHAN S. G., GUPTA M., DONNER C., RAMAMOORTHY R., NAYAR S. K., JENSEN H. W.: Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics* 25, 3 (2006), 1003–1012.
- [NGH\*18] NOVÁK J., GEORGIEV I., HANIKA J., KRIVÁNEK J., JAROSZ W.: Monte Carlo methods for physically based volume rendering. In *ACM SIGGRAPH 2018 Courses* (Vancouver, British Columbia, Canada, 2018), ACM, New York, NY, USA, pp. 1–2.
- [NGHJ18] NOVÁK J., GEORGIEV I., HANIKA J., JAROSZ W.: Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum* 37, 2 (2018), 551–576.
- [NI21] NABATA K., IWASAKI K.: Adaptive irradiance sampling for many-light rendering of subsurface scattering. *IEEE Transactions on Visualization and Computer Graphics* 28, 10 (2021), 3324–3335.
- [NK18] NAKAMOTO K., KOIKE T.: Which bssrdf model is better for heterogeneous materials? In *ACM SIGGRAPH 2018 Posters* (Vancouver, British Columbia, Canada, 2018), ACM, New York, NY, USA, pp. 1–2.
- [NRH\*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: *Geometrical considerations and nomenclature for reflectance*. University of Michigan Library, 1977.
- [NRS14] NALBACH O., RITSCHEL T., SEIDEL H.-P.: Deep screen space. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (San Francisco, California, USA, 2014), ACM, New York, NY, USA, pp. 79–86.
- [PAT\*04] PREMOŽE S., ASHIKHMIN M., TESSENDORF J., RAMAMOORTHY R., NAYAR S.: Practical rendering of multiple scattering effects in participating media. In *Proceedings of the 15th Eurographics Conference on Rendering Techniques* (Norrköping, Sweden, 2004), Eurographics Association, pp. 363–374.
- [PB11] PENNER E., BORSHUKOV G.: Pre-integrated skin shading. In *GPU Pro*, vol. 2. A K Peters/CRC Press, Natick (2011), pp. 41–56.
- [PRJ\*13] PAPAS M., REGG C., JAROSZ W., BICKEL B., JACKSON P., MATUSIK W., MARSCHNER S., GROSS M.: Fabricating translucent materials using continuous pigment mixtures. *ACM Transactions on Graphics* 32, 4 (2013), 1–12.
- [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 Transactions on Graphics* 25, 3 (2006), 746–753.
- [RSB\*21] RITTIG T., SUMIN D., BABAEI V., DIDYK P., VOLOBOY A., WILKIE A., BICKEL B., MYSZKOWSKI K., WEYRICH T., KRIVÁNEK J.: Neural acceleration of scattering-aware color 3D printing. *Computer Graphics Forum* 40, 2 (2021), 205–219.
- [SAHD95] SCHWEIGER M., ARRIDGE S., HIRAOKA M., DELPY D.: The finite element method for the propagation of light in scattering media: Boundary and source conditions. *Medical Physics* 22, 11 (1995), 1779–1792.
- [SDS\*16] SHINYA M., DOBASHI Y., SHIRAISHI M., KAWASHIMA M., NISHITA T.: Multiple scattering approximation in heterogeneous media by narrow beam distributions. *Computer Graphics Forum* 35, 7 (2016), 373–382.
- [SHK17] SONE H., HACHISUKA T., KOIKE T.: Parameter estimation of bssrdf for heterogeneous materials. In *Eurographics (Short Papers)* (Lyon, France, 2017), Eurographics Association, pp. 73–76.
- [Sim04] SIMON G.: Real-time approximations to subsurface scattering. In *GPU Gems*, vol. 1. Addison-Wesley Professional, Boston (2004), pp. 263–278.

- [SKCJ18] SENGUPTA S., KANAZAWA A., CASTILLO C. D., JACOBS D. W.: SfSNet: Learning shape, reflectance and illuminance of faces ‘in the wild’. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Salt Lake City, Utah, USA, 2018), IEEE, pp. 6296–6305.
- [SKP08] SHAH M. A., KONTTINEN J., PATTANAIK S.: Image-space subsurface scattering for interactive rendering of deformable translucent objects. *IEEE Computer Graphics and Applications* 29, 1 (2008), 66–78.
- [SKS02] SLOAN P.-P., KAUTZ J., SNYDER J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics* 21, 3 (2002), 527–536.
- [SRB\*19] SUMIN D., RITTIG T., BABAEI V., NINDEL T., WILKIE A., DIDYK P., BICKEL B., KRIVÁNEK J., MYŠKOWSKI K., WEYRICH T.: Geometry-aware scattering compensation for 3D printing. *ACM Transactions on Graphics* 38, 4 (2019), 1–14.
- [SSWN13] SHENG Y., SHI Y., WANG L., NARASIMHAN S. G.: A practical analytic model for the radiosity of translucent scenes. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Orlando, Florida, USA, 2013), ACM, New York, NY, USA, pp. 63–70.
- [Sta95] STAM J.: Multiple scattering as a diffusion process. In *Eurographics Workshop on Rendering Techniques* (Dublin, Ireland, 1995), Springer, pp. 41–50.
- [STPP09] SONG Y., TONG X., PELLACINI F., PEERS P.: SubEdit: A representation for editing measured heterogeneous subsurface scattering. *ACM Transactions on Graphics* 28, 3 (2009), 1–10.
- [SW13] SONG Y., WANG W.: A data-driven model for anisotropic heterogeneous subsurface scattering. In *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference* (Kaohsiung, Taiwan, 2013), IEEE, pp. 1–7.
- [SZLG10] SUN X., ZHOU K., LIN S., GUO B.: Line space gathering for single scattering in large scenes. *ACM Transactions on Graphics* 29, 4 (2010), 1–8.
- [TGL\*06] TARIQ S., GARDNER A., LLAMAS I., JONES A., DEBEVEC P., TURK G.: Efficient estimation of spatially varying subsurface scattering parameters. In *11th International Fall Workshop on Vision, Modeling, and Visualization* (Aachen, Germany, 2006), Akademische Verlagsgesellschaft, pp. 129–136.
- [TWL\*05] TONG X., WANG J., LIN S., GUO B., SHUM H.-Y.: Modeling and rendering of quasi-homogeneous materials. *ACM Transactions on Graphics* 24, 3 (2005), 1054–1061.
- [TYS\*19] TODO H., YATAGAWA T., SAWAYAMA M., DOBASHI Y., KAKIMOTO M.: Image-based translucency transfer through correlation analysis over multi-scale spatial color distribution. *The Visual Computer* 35 (2019), 811–822.
- [VJK21] VICINI D., JAKOB W., KAPLANYAN A.: A non-exponential transmittance model for volumetric scene representations. *ACM Transactions on Graphics* 40, 4 (2021), 1–16.
- [VKJ19] VICINI D., KOLTUN V., JAKOB W.: A learned shape-adaptive subsurface scattering model. *ACM Transactions on Graphics* 38, 4 (2019), 1–15.
- [WABG06] WALTER B., ARBREE A., BALA K., GREENBERG D. P.: Multidimensional lightcuts. *ACM Transactions on Graphics* 25, 3 (2006), 1081–1088.
- [WCPL\*08] WANG R., CHESLACK-POSTAVA E., LUEBKE D., CHEN Q., HUA W., PENG Q., BAO H.: Real-time editing and relighting of homogeneous translucent materials. *The Visual Computer* 24, 7 (2008), 565–575.
- [Wil78] WILLIAMS L.: Casting curved shadows on curved surfaces. *ACM SIGGRAPH Computer Graphics* 12, 3 (1978), 270–274.
- [WLLC14] WU Y.-T., LI T.-M., LIN Y.-H., CHUANG Y.-Y.: Dual-matrix sampling for scalable translucent material rendering. *IEEE Transactions on Visualization and Computer Graphics* 21, 3 (2014), 363–374.
- [WMP\*06] WEYRICH T., MATUSIK W., PFISTER H., BICKEL B., DONNER C., TU C., MCANDLESS J., LEE J., NGAN A., JENSEN H. W., GROSS M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics* 25, 3 (2006), 1013–1024.
- [WPW89] WYMAN D. R., PATTERSON M. S., WILSON B. C.: Similarity relations for the interaction parameters in radiation transport. *Applied Optics* 28, 24 (1989), 5243–5249.
- [WTL05] WANG R., TRAN J., LUEBKE D.: All-frequency interactive relighting of translucent objects with single and multiple scattering. *ACM Transactions on Graphics* 24, 3 (2005), 1202–1207.
- [WWH\*10] WANG Y., WANG J., HOLZSCHUCH N., SUBR K., YONG J.-H., GUO B.: Real-time rendering of heterogeneous translucent objects with arbitrary shapes. *Computer Graphics Forum* 29, 2 (2010), 497–506.
- [WWY22] WU W., WANG B., YAN L.-Q.: A survey on rendering homogeneous participating media. *Computational Visual Media* 8, 2 (2022), 177–198.
- [WZHB09] WALTER B., ZHAO S., HOLZSCHUCH N., BALA K.: Single scattering in refractive media with triangle mesh boundaries. *ACM Transactions on Graphics* 28, 3 (2009), 92.
- [WZT\*08] WANG J., ZHAO S., TONG X., LIN S., LIN Z., DONG Y., GUO B., SHUM H.-Y.: Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics* 27, 1 (2008), 1–18.
- [XGL\*07] XU K., GAO Y., LI Y., JU T., HU S.-M.: Real-time homogenous translucent material editing. *Computer Graphics Forum* 26, 3 (2007), 545–552.

- [XO21] XIE T., OLANO M.: Real-time subsurface control variates: Temporally stable adaptive sampling. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 4, 1 (2021), 1–18.
- [XOKN20] XIE T., OLANO M., KARIS B., NARKOWICZ K.: Real-time subsurface scattering with single pass variance-guided adaptive importance sampling. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 1 (2020), 1–21.
- [XWG\*14] XIAO B., WALTER B., GKIIOULEKAS I., ZICKLER T., ADELSON E., BALA K.: Looking against the light: How perception of translucency depends on lighting direction. *Journal of Vision* 14, 3 (2014), 17.
- [YBS04] YOSHIZAWA S., BELYAEV A., SEIDEL H.-P.: A fast and simple stretch-minimizing mesh parameterization. In *Proceedings Shape Modeling Applications 2004* (2004), IEEE, pp. 200–208.
- [YSJR17] YAN L.-Q., SUN W., JENSEN H. W., RAMAMOORTHY R.: A bssrdf model for efficient rendering of fur with global illumination. *ACM Transactions on Graphics* 36, 6 (2017), 1–13.
- [YTYM20] YATAGAWA T., TODO H., YAMAGUCHI Y., MORISHIMA S.: Data compression for measured heterogeneous subsurface scattering via scattering profile blending. *The Visual Computer* 36, 3 (2020), 541–558.
- [YYM20] YATAGAWA T., YAMAGUCHI Y., MORISHIMA S.: LinSSS: Linear decomposition of heterogeneous subsurface scattering for real-time screen-space rendering. *The Visual Computer* 36, 10 (2020), 1979–1992.
- [YZXW12] YAN L.-Q., ZHOU Y., XU K., WANG R.: Accurate translucent material rendering under spherical Gaussian lights. *Computer Graphics Forum* 31, 7 (2012), 2267–2276.
- [ZRB14] ZHAO S., RAMAMOORTHY R., BALA K.: High-order similarity relations in radiative transfer. *ACM Transactions on Graphics* 33, 4 (2014), 1–12.