

# Heterogeneous Subsurface Scattering Using the Finite Element Method

Adam Arbree, Bruce Walter, and Kavita Bala

**Abstract**—Materials with visually important heterogeneous subsurface scattering, including marble, skin, leaves, and minerals, are common in the real world. However, general, accurate and efficient rendering of these materials is an open problem. In this paper, we describe a finite element (FE) solution of the heterogeneous diffusion equation (DE) that solves this problem. Our algorithm is the first to use the FE method to solve the difficult problem of heterogeneous subsurface rendering. To create our algorithm, we make two contributions. First, we correct previous work and derive an accurate and complete heterogeneous diffusion formulation. This formulation has two key elements: an accurate model of the reduced intensity (RI) source, the diffusive source boundary condition (DSBC), and its associated render query function. Second, we solve this formulation accurately and efficiently using the FE method. Using these results, we can render subsurface scattering with a simple four step algorithm. To demonstrate that our algorithm is simultaneously general, accurate and efficient, we test its performance on a series of difficult scenes. For a wide range of materials and geometry, it produces, in minutes, images that nearly match path traced references, that required hours.

**Index Terms**—subsurface scattering, finite element method, diffusion equation, diffusive source boundary condition



## 1 INTRODUCTION

The subsurface scattering of light creates the distinctive appearance of many ubiquitous materials, such as marble, skin, minerals and leaves. However, because of the complexity of the scattering within these materials, efficient and accurate subsurface rendering is challenging and previous approaches have limitations. Monte Carlo (MC) algorithms [20], [26], [38] accurately solve the general subsurface rendering problem but their hours-per-image cost makes them impractical for most applications; algorithms using the dipole diffusion bidirectional surface scattering reflectance distribution function (BSSRDF) [27] are fast, even real-time, but can only model homogeneously scattering materials; and, though capture and re-render systems [12], [16], [18], [37] can quickly generate high-quality images, they can only redisplay captured material models.

Because of the limitations of these methods, current research seeks an improved result: an efficient and accurate rendering algorithm for general heterogeneous materials. In this paper, we achieve this result by using a finite element (FE) algorithm to compute the solution to the heterogeneous diffusion equation (DE). However, creating our solution required two contributions to previous work:

- a complete and correct rendering formulation for the heterogeneous diffusion problem; and
- a general FE algorithm that solves this problem efficiently.

Our contributions strike a balance between performance and accuracy. By solving an approximate scat-

tering model, we ensure efficiency but, by carefully selecting this approximation and by using the FE method, we preserve accuracy. As a result, in a few minutes, our algorithm produces images (see Figure 1, top row) that nearly match path-traced references which required hours of computation (see Figure 1 difference images, top right). Moreover, because our algorithm is a general solution, it can reproduce a wide range of difficult subsurface effects. For example, it can render complex high-frequency aggregates, such as marble (see Figure 1, bottom left); the sharp edges between regions of smooth, noise-free translucency (see Figure 1, bottom center); and, even in a challenging scene lit entirely by subsurface illumination, the complex, structured and high-resolution subsurface detail found in many natural materials (see Figure 1, bottom right).

Our contributions address two gaps found in previous discussions of heterogeneous subsurface rendering. First, we correctly derive our heterogeneous diffusion formulation with an accurate model of the reduced intensity (RI) source, the diffusive source boundary condition (DSBC), and its corresponding render query function. Because we derive the DSBC and the query function from first principles, we are able to correct the errors found in previous discussions [39], [47] and produce a heterogeneous diffusion formulation suitable for high quality rendering applications. Second, though FE algorithms are used for a wide range of mathematical problems, our algorithm improves upon the quality of previous subsurface renderers by introducing the FE method for rendering subsurface scattering.

In particular, our FE algorithm significantly extends the two recent works most closely related to our own. First, we show that our FE solution improves the

• *Authors are with Cornell University, Ithaca, NY 14853.  
E-mail: {arbree,bjw,kb}@graphics.cornell.edu*

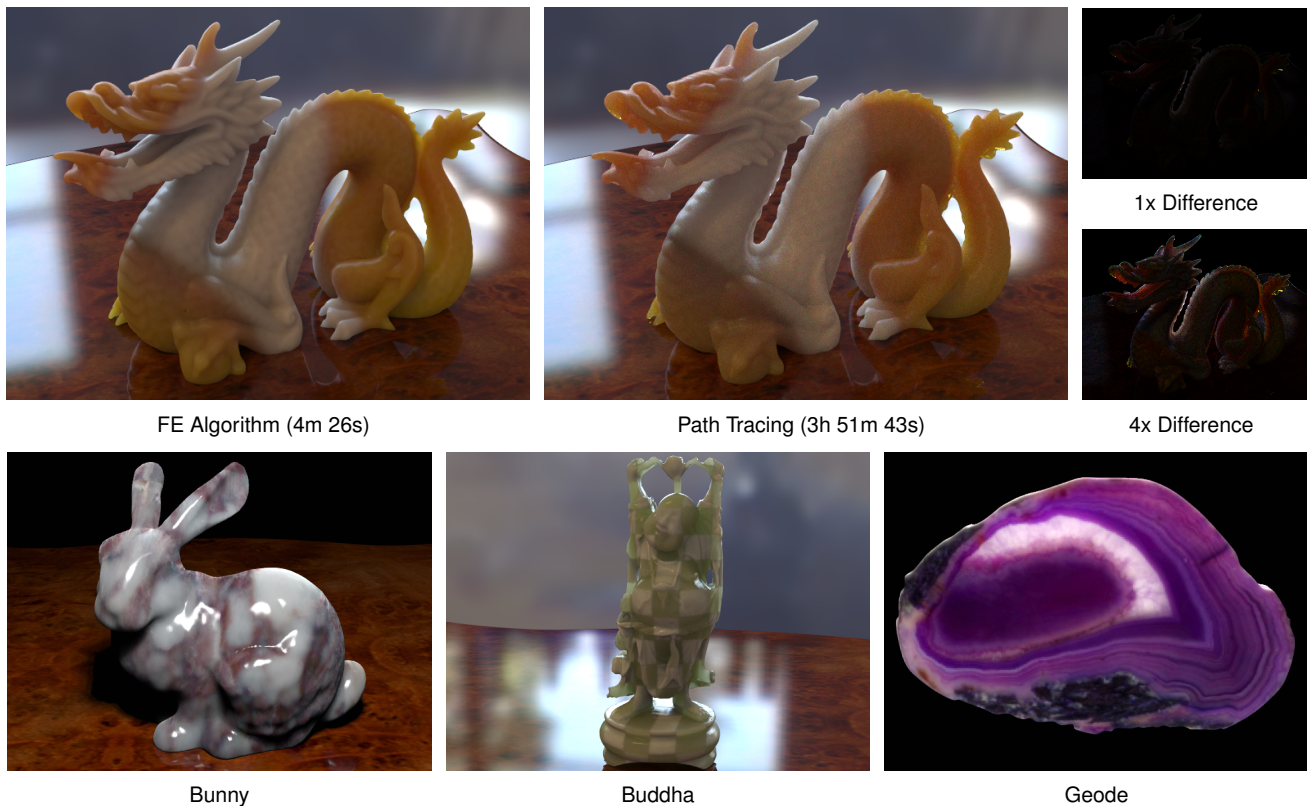


Fig. 1: Our new finite element (FE) algorithm for heterogeneous subsurface scattering can reproduce a wide range of materials in arbitrary geometry. Our results (see Dragon results; top row) are nearly identical to path traced references (see black difference image; top right) but render in a few minutes instead of a few hours. Our new algorithm can (see marble Bunny; left) model complex aggregate materials; (see checkerboard Buddha; middle right) capture sharp detail in heterogeneous material properties while still ensuring noise-free translucency in smooth regions; and (see back-lit Geode; right) scale well to detailed, high-frequency materials even in the difficult case when all illumination passes through the medium.

ambitious and heterogeneous capture, modeling and rendering system recently created by Wang et al. [47]. Though their algorithm performs excellently for their re-render application, we demonstrate that the quality of their results depends on their capture optimization correcting rendering errors. For non-captured materials, their finite difference (FD) rendering algorithm is less accurate and its dependence on PolyGrid [43] meshes places effective limits on the scattering geometry. In contrast, our new algorithm produces high quality images of many materials. Moreover, because our FE solution supports fast, off-the-shelf meshing algorithms and intelligent adaptive mesh refinement, we can achieve these results quickly in a wide range of scattering geometries..

Second, we extend similar, but much more limited, FE results discussed in previous work in medical imaging [39]. Though this work also discusses a FE solution to the DE, it omits the details necessary to create a useful heterogeneous rendering algorithm. The authors only consider a simplified FE solution in a 2D homogeneous disk and, because of this simple test and the different requirements of their imaging application, they conclude that their formulation cannot be used for a wide range of materials and geometry. However, we find the opposite result for the rendering problem.

By using an improved 3D FE solution tailored for rendering, we can achieve high-quality results in a wide range of difficult scenes.

The rest of this paper describes and tests our new heterogeneous subsurface rendering algorithm. First, in the next section, we review previous work on subsurface scattering. Second, in Section 3, we describe our new contributions and rendering algorithm. Afterward, Section 4 reviews diffusion theory to prepare for Sections 5, 6 and 7 that build on this theory to derive our algorithm. Then, in Section 8, we review the details of our implementation and in Section 9 we compare this implementation to previous solutions. Finally, in Section 10, we conclude with limitations and future work.

## 2 PREVIOUS WORK

Due to the breadth of research on subsurface rendering, a complete review lies beyond the scope of this paper. This section quickly reviews most previous work in three broad categories: Monte Carlo (MC) algorithms, dipole diffusion algorithms and capture and re-render systems. Afterward, the two algorithms most closely related to our problem, general heterogeneous subsurface rendering, are reviewed in more detail.

## 2.1 Monte Carlo Algorithms

By directly simulating basic scattering physics, MC rendering algorithms, including path tracing [20], [26], [38], photon mapping [13], [25] and Metropolis sampling [36], can generate individual photons paths in the scattering material. With sufficiently many paths, they can produce images of arbitrary accuracy; however, even for homogeneous materials, this simulation is generally considered impractical and heterogeneous materials are yet more expensive (for example, our path-traced references required many CPU/days of computation, see Section 9.4).

## 2.2 Dipole Diffusion Algorithms

Because of the high cost of MC renderers, practical subsurface algorithms are approximate. Jensen et al. [27] introduced the first such approximation: the dipole diffusion bidirectional surface scattering reflectance distribution function (BSSRDF). The dipole diffusion BSSRDF is derived from an analytic scattering approximation [23], [29] that, though fast, limits both the BSSRDF's quality [32] and fundamentally restricts it to homogeneous materials (though an artificial heterogeneous appearance can be created with textures [31], [44]). However, because it is easy to implement and inexpensive, the dipole diffusion BSSRDF has been widely used. It was first used to accelerate MC algorithms [24], [27] and then several papers extended these algorithms to improve their quality [7], [10], [11], [32]. Recently, the dipole diffusion BSSRDF has been used to create real-time algorithms [5], [6], [8], [21], [22], [31], [33], [34]; included in precomputed radiance transfer systems [41], [48]; and scaled to large scenes using a Light cuts based algorithm [2], [46].

## 2.3 Capture and Re-render Systems

Capture and re-render systems—both for general subsurface scattering materials [18], [37] and specifically for human skin [12], [16]—compute specialized material models from photographs of a physical scattering object. After this capture computation, the internal model can be warped into new geometry and re-rendered. Because these models are computed from real photographs, the re-rendered images have high quality. However, though these systems can capture both homogeneous and heterogeneous materials, these systems can only redisplay their captured models. They cannot render general heterogeneous materials.

## 2.4 Heterogeneous Algorithms

To address the limitations of the categories above, we present our general heterogeneous rendering algorithm. This algorithm computes subsurface scattering by numerically solving the diffusion equation (DE). Though Stam [42] first suggested this approach and Haber et al. [19] later developed numerical solver

for homogeneous materials, only two works have discussed the general heterogeneous problem: the finite difference (FD) solver recently introduced by Wang et al. [47] and the work by Schweiger et al. [39] in the medical imaging field of optical tomography (OT).

### 2.4.1 Finite Difference Algorithm

As part of an impressive, capture and re-rendering system, Wang et al. [47] developed an interactive, FD algorithm for solving the heterogeneous DE. Their goal was the interactive editing of captured material grids and, for this task, their algorithm is excellent. However, as a general, high quality solution, we have found that their FD algorithm has several limitations (see Section 9.5 for a detailed discussion and a comparison to our results). Foremost, errors in their DE formulation can result in incorrect solutions. These errors went unnoticed in their original work because their capture computation corrects for them by altering the captured material parameters<sup>1</sup>.

However, even if we correct their algorithm using our formulation, two additional limitations remain. First, the FD algorithm requires that a PolyGrid [43] be constructed in the scattering geometry. As noted by the authors, this is a difficult process that can require hand optimization to avoid errors. Second, their FD solution—especially with the additional approximations required to achieve interactive performance—has approximation errors. Most importantly, the deformation of the FD algorithm's PolyGrid cannot be exactly modeled in their solver. Therefore, large deformations result in rendering errors or, for some materials, the divergence of the iterative solution.

### 2.4.2 Optical Tomography

Outside of computer graphics, the medical imaging field of OT also develops algorithms for numerically solving the DE. On one hand, OT tackles the same problem addressed by Wang et al. [47]: given a series of photographs of human tissue lit by visible light, determine the scattering properties of the tissue (see [17], [30] for a review of the state-of-the-art). On the other hand, the OT problem is significantly different than the rendering problem. OT applications are less concerned with efficiency, need solutions in much smaller regions than a whole image and consider only materials within the smaller range of scattering parameters found in human tissue.

However, one work [39] discusses, among several others, a DE problem formulation similar to ours and they consider solving this formulation using the finite element (FE) method. But, though their discussion introduces topics similar to ours, it also omits the details

1. Before writing this paper, we discussed this issue with Wang et al. [47] and they kindly acknowledged that these corrections do occur (see Section 9.5 and [3], [47]).

necessary to create a useful rendering algorithm. They do not explore heterogeneous scattering in arbitrary objects and compute only a simplified finite element solution in a 2D homogeneous disk. Because of this simple test, they cannot determine if their formulation is accurate or efficient for a wide range of materials and geometry. Finally, because they address a different problem, their equations lack two elements essential in our rendering solution: a boundary Fresnel term and a rendering query function that converts the final solution into radiance.

### 3 ALGORITHM

Our algorithm improves upon all previous subsurface renderers. Our results are nearly as accurate as Monte Carlo (MC) images, but much less expensive and, unlike both dipole diffusion algorithms and capture and re-render systems, our algorithm can quickly reproduce a broad range of heterogeneous materials. We achieve these results using a finite element (FE) algorithm for solving the heterogeneous diffusion equation (DE). Though the FE method has been used in a wide range of applications spanning many fields, our method is the first to solve subsurface scattering with the FE method. Because of the specific requirements of the rendering problem, our solution must be carefully formulated to ensure accurate results. Unfortunately when developing our solution, we found errors in previous diffusion formulations [39], [47] that we correct in the discussion below (see also [3]). In addition, we further extend the 2D FE solution discussed Schweiger et al. [39] by describing a complete 3D heterogeneous FE solution tailored to the rendering problem. The produce of our careful and tailored derivations is a fast, accurate, and general FE subsurface rendering algorithm.

#### 3.1 Overview of Key Results

This final algorithm is defined by three key results: the diffusive source boundary condition (DSBC), the finite element diffusion equation (FEDE) and the render query function.

##### *Result #1: Diffusive Source Boundary Condition*

Our algorithm approximates subsurface scattering using the DE. This approximation has three components: the DE itself that approximates scattering in the interior of the scattering volume; a boundary condition (BC) that approximates solution's boundary behavior; and a reduced intensity (RI) source model that approximates the radiance entering the scattering media. Most previous work presents these three components in a unified formulation based on the approximations that derive the dipole diffusion algorithm [27]. However, for general heterogeneous problems, we demonstrate that the RI source model used in this formulation, the embedded source model

[11], is both inaccurate and inefficient. To address these limitations, Section 5 derives our first key result: the DSBC (see Equation (14)). The DSBC combines an accurate BC approximation with an improved RI source model to produce a high-quality heterogeneous subsurface scattering approximation.

##### *Result #2: Finite Element Diffusion Equation*

Next, in Section 6, we solve this approximation accurately using the FE method. This solution has two steps. First, the scattering domain is discretized into a mesh that, by construction, implicitly defines a finite basis for functions in the scattering domain. Second, we discretize the constraints of the DE and DSBC using this basis. The result is a large matrix equation whose solution is the coefficient vector of the best approximation, in this basis, of the final subsurface scattering solution. This matrix equation is our second key result, the FEDE (see Equation (21)) .

##### *Result #3: Render Query Function*

Finally, we describe how to compute exitant subsurface radiance, the quantity needed for rendering, from this solution. Because of the approximations of the DE, DSBC and their FE solution, this query function must be carefully constructed to ensure accuracy. Section 7 describes these issues and derives our final key result, the render query function (see Equation (24)).

#### 3.2 Finite Element Rendering Algorithm

Using these results, subsurface rendering is a four-step process:

- Step 1: Mesh the scattering volume
- Step 2: Construct the FEDE for that mesh
- Step 3: Solve the resulting linear system
- Step 4: Compute radiance using the query function

Moreover, because of the generality of our FE solution, this basic algorithm is not only simple, but adaptable. Though in Section 8 we describe the specific details of our full implementation, without change to its basic structure, it can be altered to fit the problem at hand or to incorporate existing tools. Many mesh types can be used—allowing fast meshing algorithms (see Section 8.1) and adaptive refinement (see Section 8.5)—without change to the assembly of the FEDE in Step 2 (see Section 8.2 and Figure 6 for the pseudocode of a general algorithm) and, once the FEDE is assembled, the solution can be computed with almost any matrix algorithm.

### 4 REVIEW OF DIFFUSION THEORY

Before we derive our algorithm, we review some elements of diffusion theory. Our algorithm uses the diffusion equation (DE) to approximate the exact scattering model, the volumetric radiance transfer equation (VRTE). But the basic approximation of the

$\Omega$	= Scattering volume	$\sigma_a(\mathbf{x})$	= Absorption coefficient	$\Gamma_s(\mathbf{x})$	= Incoming external flux
$\partial\Omega$	= Boundary of $\Omega$	$\sigma_s(\mathbf{x})$	= Scattering coefficient	$\Gamma_d^{\text{in}}(\mathbf{x})$	= Inward solution flux
$L(\mathbf{x}, \vec{\omega})$	= Radiance leaving $\mathbf{x}$ in $\vec{\omega}$	$\sigma_t(\mathbf{x})$	= Total extinction coefficient	$\Gamma_d^{\text{ref}}(\mathbf{x})$	= Boundary reflected flux
$L_d(\mathbf{x}, \vec{\omega})$	= Diffusive radiance	$\sigma_{tr}(\mathbf{x})$	= Reduced extinction coefficient	$\mathbb{H}$	= Arbitrary function space
$\phi(\mathbf{x})$	= Radiant fluence	$\kappa_d(\mathbf{x})$	= Diffusion coefficient	$\theta$	= Arbitrary function in $\mathbb{H}$
$\vec{E}(\mathbf{x})$	= Vector irradiance	$\eta$	= Relative index of refraction	$\beta_i$	= $i^{\text{th}}$ basis function
$Q(\mathbf{x}, \vec{\omega})$	= Emittance of the media	$F_r(\eta, \vec{\omega})$	= Fresnel reflectance	$\mathbf{F}$	= Finite element matrix
$Q_{ri}(\mathbf{x})$	= Reduced intensity source	$F_t(\eta, \vec{\omega})$	= Fresnel transmittance	$\mathbf{D}, \mathbf{M}, \mathbf{S}$	= Components of $\mathbf{F}$
$Q_0(\mathbf{x})$	= 0 <sup>th</sup> moment of $Q(\mathbf{x}, \vec{\omega})$	$F_{dr}(\eta)$	= Average Fresnel reflectance	$\vec{r}$	= Finite element right hand side
$p(\vec{\omega}, \vec{\omega}')$	= Scattering phase function	$F_{dt}(\eta)$	= Average Fresnel transmittance	$\vec{q}, \vec{g}$	= Components of $\vec{r}$
$\mu$	= Mean cosine of $p(\vec{\omega}, \vec{\omega}')$	$A(\eta)$	= Fresnel boundary term	$\vec{a}$	= Finite element solution vector

Fig. 2: Summary of notation used throughout the paper.

DE, the diffusion approximation (DA), forces two further approximations: the reduced intensity (RI) source model and the boundary condition (BC) [23]. Because our first key result, the diffusive source boundary condition (DSBC), provides these secondary approximations, we prepare for this discussion by introducing them here.

#### 4.1 Volumetric Radiative Transfer Equation

The light scattered within a randomly scattering medium is defined by three functions that specify the probability of photon interactions<sup>2</sup>. The absorption and scattering coefficients,  $\sigma_a(\mathbf{x})$  and  $\sigma_s(\mathbf{x})$  respectively, give the the number of photons scattered and absorbed per unit distance and the phase function  $p(\vec{\omega}, \vec{\omega}')$  specifies the probability that a scattered photon leaves  $\vec{\omega}$  into  $\vec{\omega}'$ . Given these three functions, the VRTE defines the differential radiance  $L(\mathbf{x}, \vec{\omega})$  in the medium leaving  $\mathbf{x}$  in direction  $\vec{\omega}$  [23]

##### Volumetric Radiative Transfer Equation

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

where  $\sigma_t(\mathbf{x}) = \sigma_a(\mathbf{x}) + \sigma_s(\mathbf{x})$  and  $Q(\mathbf{x}, \vec{\omega})$  is the light source function.

#### 4.2 The Diffusion Equation

To derive the DE, one makes the DA. The DA approximates the VRTE's solution by a simpler function with a near isotropic angular distribution: a linear expansion in the 0<sup>th</sup> and 1<sup>st</sup> angular moments of the radiance: the scalar fluence  $\phi(\mathbf{x})$  and the vector irradiance  $\vec{E}(\mathbf{x})$  respectively.

##### Diffusion Approximation

$$L(\mathbf{x}, \vec{\omega}) = \frac{1}{4\pi}\phi(\mathbf{x}) + \frac{3}{4\pi}\vec{E}(\mathbf{x}) \cdot \vec{\omega} \quad (2)$$

2. Figure 2 summarizes the notation used in this paper.

where

$$\phi(\mathbf{x}) = \int_{4\pi} L(\mathbf{x}, \vec{\omega}) d\vec{\omega} \quad \text{and} \quad \vec{E}(\mathbf{x}) = \int_{4\pi} L(\mathbf{x}, \vec{\omega}) \cdot \vec{\omega} d\vec{\omega}$$

Though, its an involved process (see Ishimaru [23] for the details), essentially the substitution of Equation (2) into Equation (1) yields the:

##### Diffusion Equation

$$-\vec{\nabla} \cdot (\kappa_d(\mathbf{x})\vec{\nabla}\phi(\mathbf{x})) + \sigma_a(\mathbf{x})\phi(\mathbf{x}) = Q_0(\mathbf{x}) + Q_{ri}(\mathbf{x}) \quad (3)$$

This substitution introduces three new quantities. Two,  $\kappa_d(\mathbf{x})$  and  $Q_0(\mathbf{x})$ , are simply redefinitions of the input parameters:

$$\kappa_d = \frac{1}{3[(1-\mu)\sigma_s + \sigma_a]} \quad \text{and} \quad Q_0(\mathbf{x}) = \int_{4\pi} Q(\mathbf{x}, \vec{\omega}) d\vec{\omega}$$

However, the third term, the RI source  $Q_{ri}(\mathbf{x})$  is the first of the two secondary approximations required by the DA.

#### 4.3 Reduced Intensity Source

The RI source is necessary because sufficiently near the boundary incoming radiance violates the DA<sup>3</sup>. Far from the boundary, the DA's required isotropy is naturally ensured because any radiance deep in the material will have randomly scattered many times smoothing away any initial angular properties. However, sufficiently near the boundary, some radiance will not yet have scattered and this radiance can be highly anisotropic. Since this anisotropy can cause approximation errors, this yet-to-be-scattered radiance is removed from the diffusion problem and the DA is applied only to the remaining diffusive radiance  $L_d(\mathbf{x}, \vec{\omega})$ . However, this radiance cannot be neglected entirely and, to account for its removal, the DE's derivation introduces a new artificial source function, the RI source  $Q_{ri}(\mathbf{x})$ , that supplies any diffusive radiance that would have

3. A similar problem also occurs near sources embedded in the medium. However, since these sources rarely occur in physical scenes and, because analytic RI source models exist for many of light source geometries [15], [23], we review only the boundary case here.

been generated before removal [23]. Unfortunately, an exact photon-by-photon RI source computation would erase the efficiency originally gained by using the DE to model bulk scattering and, as a practical matter, the RI source model must be approximated.

#### 4.4 Boundary Condition

Besides the RI source model, the DA requires a second approximation. For finite domains, the DE requires a BC. With the introduction of the RI source, the DE was restricted to only the diffusive radiance. Since by definition this radiance must have scattered at least once, the physical BC requires that  $L_d(\mathbf{x}, \vec{\omega})$  be zero along all incoming directions.

$$\forall(\vec{n} \cdot \vec{\omega}) < 0, L_d(\mathbf{x}, \vec{\omega}) = 0 \quad (4)$$

However, this condition is discontinuous between incoming and outgoing directions and, since this discontinuity violates the DA, a DA-compatible approximate condition must be used instead. Previous work has already described an accurate BC approximation derived from boundary Fresnel effects [15], [23]. We discuss this approximation in Section 5.4 as part of the derivation of the DSBC.

### 5 DIFFUSIVE SOURCE BOUNDARY CONDITION

The boundary condition (BC) and reduced intensity (RI) source model approximations strongly affect the diffusion solution on the boundary of the scattering domain. Because only the boundary is visible in an image, these approximations effectively determine the quality of the final rendering algorithm. In our heterogeneous application, where fine boundary material detail must be accurately modeled, these approximations must be as accurate as possible.

In this section, we demonstrate that the RI source model used by previous dipole algorithms, the embedded source model, is both inefficient and inaccurate. To address these issues, this section discusses our first key result, the diffusive source boundary condition (DSBC), that replaces the embedded approximation.

#### 5.1 Embedded Source Model

The embedded source model (Figure 3(a)) is derived from an analytic RI source model for a collimated beam incident on a homogeneous, semi-infinite slab. For this problem, the RI source can be approximated by a series of omni-directional point sources embedded inside the the medium. For arbitrary problems, the surface is divided into patches. By locally approximating each patch by an instance of the homogeneous slab problem, the collection of patches can be converted into a large collection of point sources that approximate the RI source throughout arbitrary regions of the scattering volume [11].

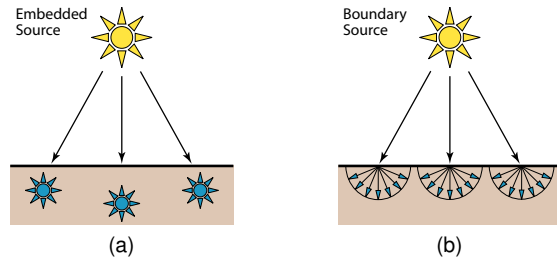


Fig. 3: Two methods of approximating the reduced intensity source  $Q_{ri}(\mathbf{x})$ . The embedded source model (a) approximates  $Q_{ri}(\mathbf{x})$  with point sources while our model, the boundary source model (b), approximates the source as a diffusive flux arriving at the boundary.

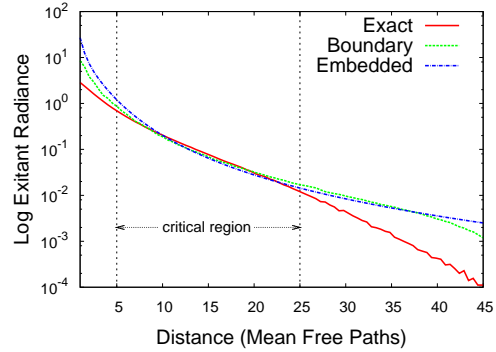


Fig. 4: Comparison of the boundary and embedded source models (top) and their error (bottom) for a collimated beam incident on a semi-infinite slab of homogeneous materials: exact solution (solid red), diffusive source boundary condition (dotted green) and embedded source model (dashed blue).

#### 5.2 Boundary Source Model

Unlike the embedded source model, the DSBC (Figure 3(b)) does not create a volumetric representation of the RI source. Instead, the DSBC is a 2D surface source model. Because of scattering and absorption, the RI source decays exponentially and, due to this rapid falloff, almost all of the power of the RI source power lies close the boundary. Therefore, it is a small approximation to represent the true RI source by its 2D projection onto the boundary. This projection is modelled as a isotropic diffusive flux that arrives at the boundary.

#### 5.3 Limitations of the Embedded Source Model

In heterogeneous problems, the embedded source model has two drawbacks. Foremost, the embedded source approximation is less accurate. To demonstrate this result, we compare both models for the canonical example of a infinitesimal, collimated beam normally incident on a semi-infinite, homogeneously scattering slab. Because it is derived from this example, the embedded source model is maximally accurate for this configuration. In contrast, this example is poorly approximated by the DSBC: a directed, collimated beam cannot be well represented by an isotropic diffuse flux. Since this example compares the embedded model's best case to worst case for the DSBC, we can

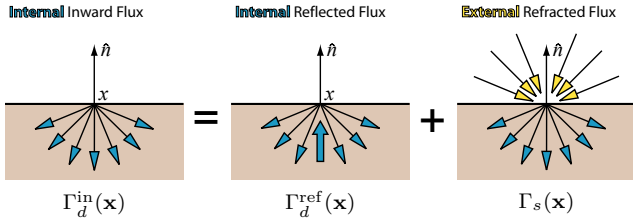


Fig. 5: Diagrams illustrating the three components of the diffusive source boundary condition (DSBC) (Equation (14)). The condition forces  $\Gamma_d^{\text{in}}(\mathbf{x})$ , the internal inward flux at the boundary, to be equal to the sum of  $\Gamma_d^{\text{ref}}(\mathbf{x})$ , the internal flux reflected at the boundary, and  $\Gamma_s(\mathbf{x})$  the exterior light refracted into the material.

use its surprising result—both methods are equally accurate—to predict the relative performance of the two methods across a wide range of scenes.

In Figure 4 we plot, outward from the point of incidence of the beam, a reference embedded solution, our FE solution using the DSBC<sup>4</sup> and an exact Monte Carlo (MC) result. Unfortunately, both near and far from the incident source, both methods have errors. However, these errors result from the breakdown of the diffusion approximation (DA) not the RI source model. In the critical middle region (between 5 and 25 mean free paths) where the DA is accurate, both source methods closely approximate the true result. Since in other more complex scenes the accuracy of the embedded source model will only degrade and the quality of the DSBC will only improve, we can conclude that the DSBC is the more accurate choice overall.

Though improved accuracy is already a compelling argument for the DSBC, the embedded source model has a second drawback: increased cost. Because the embedded source model is volumetric, it depends on both the incoming illumination as well as the heterogeneous material properties near the boundary. This added material dependence increases cost. In complex heterogeneous materials, even if the incident illumination is smooth, the embedded source model must densely sample regions with high-frequency material detail. In contrast, because its computation is independent of the boundary material properties, the DSBC need only sample the smooth incident illumination function.

## 5.4 Derivation of DSBC

Given these advantages, the DSBC is the clear choice for our heterogeneous diffusion problem. To derive<sup>5</sup> the DSBC we augment the Fresnel boundary condition discussed in previous work [23] with a new term to

4. The semi-infinite slab is replaced with a cube 100 mean free paths along each edge and the infinitesimally thin beam is replaced by a spot 0.2 mean free paths in diameter.

5. Though previous work has mentioned equations similar to the DSBC [39], [47], we review the full derivation here because these works omit derivations and contain errors.

model diffusive boundary flux representing the RI source and then apply the DA to the resulting equation.

### 5.4.1 Fresnel Boundary Condition

Regardless of the RI source model, the diffusion BC should model Fresnel interface effects. At every boundary point  $\mathbf{x}$ , the basic condition forces the net inward flux of the solution  $\Gamma_d^{\text{in}}(\mathbf{x})$  to equal the total flux of internal radiance reflected back inward by the Fresnel interface  $\Gamma_d^{\text{ref}}(\mathbf{x})$ .

$$\Gamma_d^{\text{in}}(\mathbf{x}) = \Gamma_d^{\text{ref}}(\mathbf{x}) \quad (5)$$

Using Figures 5(a) and 5(b) as guides, these fluxes can be computed by integrating over the small blue arrows in each figure<sup>6</sup>.

$$\Gamma_d^{\text{in}}(\mathbf{x}) = \int_{(\vec{n} \cdot \vec{\omega}) < 0} L_d(\mathbf{x}, \vec{\omega})(-\vec{n} \cdot \vec{\omega}) d\vec{\omega} \quad (6)$$

$$\Gamma_d^{\text{ref}}(\mathbf{x}) = F_{dr}(\eta) \int_{(\vec{n} \cdot \vec{\omega}) < 0} L_d(\mathbf{x}, -\vec{\omega})(-\vec{n} \cdot \vec{\omega}) d\vec{\omega} \quad (7)$$

In Equation (7), we approximate total reflected internal flux by scaling the total internal flux by  $F_{dr}(\eta)$ , the average Fresnel reflection coefficient.

### 5.4.2 Reduced Intensity Flux

Next, the RI source is added by augmenting the basic condition with a third term.

$$\Gamma_d^{\text{in}}(\mathbf{x}) = \Gamma_d^{\text{ref}}(\mathbf{x}) + \Gamma_s(\mathbf{x}) \quad (8)$$

The new term  $\Gamma_s(\mathbf{x})$  approximates the RI source from the integrated surface radiance refracted into the material at  $\mathbf{x}$ .

$$\Gamma_s(\mathbf{x}) = e^{-\frac{\sigma_a(\mathbf{x})}{\sigma_s(\mathbf{x})}} \int_{(\vec{n} \cdot \vec{\omega}) > 0} F_t(\eta, \vec{\omega}) L(\mathbf{x}, -\vec{\omega})(\vec{n} \cdot \vec{\omega}) d\vec{\omega} \quad (9)$$

In Equation (9), the exponential term approximates the absorption that has occurred as refracted radiance traveled in the medium before scattering and becoming part of the RI source.

### 5.4.3 Substitution of the Diffusion Approximation

Next, using an identity from previous work [23],

$$\int_{(\vec{s} \cdot \vec{\omega}) > 0} L_d(\mathbf{x}, \vec{\omega})(\vec{s} \cdot \vec{\omega}) d\vec{\omega} = \frac{1}{4} \left[ \phi(\mathbf{x}) - 2\kappa_d(\mathbf{x})(\vec{s} \cdot \vec{\nabla})\phi(\mathbf{x}) \right] \quad (10)$$

6. When creating these expressions, one must be careful to use a consistent definition of  $\vec{\omega}$  and  $\vec{n}$  to ensure that resulting boundary condition has the correct signs. To be consistent throughout this section, we have chosen to always: define  $\vec{n}$  as pointing out of the material and define  $\vec{\omega}$  as pointing away from  $\mathbf{x}$ .

we can express Equations (6) and (7) in terms of fluence by substituting the DA (Equation (2)) for  $L_d(\mathbf{x}, \vec{\omega})$  and simplifying.

$$\Gamma_d^{\text{in}}(\mathbf{x}) = \frac{1}{4} \left[ \phi(\mathbf{x}) + 2\kappa_d(\mathbf{x})(\vec{n} \cdot \vec{\nabla})\phi(\mathbf{x}) \right] \quad (11)$$

$$\Gamma_d^{\text{ref}}(\mathbf{x}) = \frac{1}{4} F_{dr}(\eta) \left[ \phi(\mathbf{x}) - 2\kappa_d(\mathbf{x})(\vec{n} \cdot \vec{\nabla})\phi(\mathbf{x}) \right] \quad (12)$$

Finally, substituting Equations (11) and (12) into Equation (8) and then using the average Fresnel transmission ratio  $A(\eta)$

$$A(\eta) = \frac{1 + F_{dr}(\eta)}{1 - F_{dr}(\eta)} \quad (13)$$

to simplify the resulting expression yields

*Result 1: Diffusive Source Boundary Condition*

$$\phi(\mathbf{x}) + 2A(\eta)\kappa_d(\mathbf{x})(\vec{n} \cdot \vec{\nabla})\phi(\mathbf{x}) = \frac{4}{F_{dt}(\eta)}\Gamma_s(\mathbf{x}) \quad (14)$$

## 6 FINITE ELEMENT DIFFUSION EQUATION

In the last section, we derived the diffusive source boundary condition (DSBC) to complete the approximation of our heterogeneous scattering problem. In this section, we solve this problem using the finite element (FE) method. The result is our second key result: the finite element diffusion equation (FEDE).

### 6.1 Overview

The derivation of our FE solution is the same for any 2<sup>nd</sup>-order partial differential equation (PDE), such as the diffusion equation (DE), and has two steps. The first step converts the PDE into a new problem, its weak form. Given a space of functions  $\mathbb{H}$ , the weak form defines an optimization problem over all functions in  $\mathbb{H}$ . This problem's optimal function, the weak solution, is a generalized solution of the original PDE: either it is the PDE's exact solution or it is that solution's closest approximation in  $\mathbb{H}$ . In the second step, one assumes that  $\mathbb{H}$  was constructed from the span of a finite basis (see Section 8.1 for an example). Then, given this basis, the coefficients of the weak solution can be expressed as the solution vector of a linear system. This linear system is the FE solution to the original PDE.

All FE solutions have three important properties [1], [14], [35]. First, if all the PDE's coefficients are bounded and have bounded derivatives<sup>7</sup>, then the FE matrix is positive definite and always has a solution. Second, the components of the FE matrix are non-zero only when the support of the basis functions overlaps. If the basis functions are spatially compact (as most

7. This condition is stronger than necessary. It is sufficient for the coefficients to lie in a 1<sup>st</sup>-order Sobolev space but the introduction of these spaces is beyond the scope of this discussion. See [14] and [1] for a full treatment.

common bases are), then the FE system is sparse and can be computed efficiently. Third, the FE solution is the optimal approximation (as measured by the  $L_2$  norm) in the given basis to the PDE's true solution. Together these three properties ensure that the FE solution is both accurate and efficient.

In the rest of this section, we derive the FE solution to the DE, the FEDE. Our derivation is similar to the general elliptic weak form [14] but we tailor the general form specifically to the diffusion problem. To make the equations in this derivation more compact, we omit, in any term, both the independent spatial variable  $\mathbf{x}$  and the index of refraction variable  $\eta$ .

### 6.2 Weak Form

First we derive the weak form of the DE. In this step, let  $\mathbb{H}$  be an arbitrary space of functions and let  $\theta$  be any function in  $\mathbb{H}$ . Then multiply Equation (3) by  $\theta$  and integrate over the scattering domain  $\Omega$ .

$$-\int_{\Omega} (\vec{\nabla} \cdot (\kappa_d(\mathbf{x})\vec{\nabla}\phi(\mathbf{x})))\theta \, dx + \int_{\Omega} \sigma_a\phi\theta \, dx = \int_{\Omega} Q_0\theta \, dx \quad (15)$$

By definition the left hand side of the weak form must be a bilinear functional over the functions  $\phi$  and  $\theta$ . To achieve this, transform Equation (15) in two steps. First use the

#### Divergence Theorem

$$\int_{\Omega} \vec{\nabla} u \cdot \mathbf{v} \, dx = \int_{\partial\Omega} u(\mathbf{v} \cdot \vec{n}) \, ds - \int_{\Omega} u(\vec{\nabla} \cdot \mathbf{v}) \, dx \quad (16)$$

to integrate the first term by parts.

$$\int_{\Omega} \kappa_d \vec{\nabla}\phi \cdot \vec{\nabla}\theta \, dx - \int_{\partial\Omega} \kappa_d(\vec{n} \cdot \vec{\nabla})\phi\theta \, ds + \int_{\Omega} \sigma_a\phi\theta \, dx = \int_{\Omega} Q_0\theta \, dx \quad (17)$$

Second use the DSBC (Equation (14)) to eliminate the gradient term in Equation (17). This yields the

#### Diffusion Weak Form

Find  $\phi$  such that  $\forall \theta \in \mathbb{H}$ ,

$$\int_{\Omega} \kappa_d \vec{\nabla}\phi \cdot \vec{\nabla}\theta \, dx + \int_{\Omega} \sigma_a\phi\theta \, dx + \frac{1}{2A} \int_{\partial\Omega} \phi\theta \, ds = \int_{\Omega} Q_0\theta \, dx + \frac{2}{AF_{dt}} \int_{\partial\Omega} \Gamma_s\theta \, ds \quad (18)$$



### 6.3 Matrix Equation

To finish the derivation, we convert Equation (18) into a linear system. First, assume that  $\mathbb{H}$  has a finite basis:

$$\mathcal{B}(\mathbf{x}) = \{\beta_0(\mathbf{x}), \beta_1(\mathbf{x}), \dots, \beta_{n-1}(\mathbf{x})\}$$

Given this basis, the weak form is equivalent to ensuring that Equation (18) only hold for each  $\beta_i \in \mathcal{B}$ . If  $n = |\mathcal{B}|$ , then this is a system of  $n$  equations for  $\phi$ .

$$\begin{aligned} \int_{\Omega} \kappa_d \vec{\nabla} \phi \cdot \vec{\nabla} \beta_i \, dx + \int_{\Omega} \sigma_a \phi \beta_i \, dx + \frac{1}{2A} \int_{\partial\Omega} \phi \beta_i \, ds \\ = \int_{\Omega} Q_0 \beta_i \, dx + \frac{2}{AF} \int_{\partial\Omega} \Gamma_s \beta_i \, ds \quad \forall \beta_i \in \mathcal{B} \end{aligned} \quad (19)$$

Next, since  $\phi \in \mathbb{H}$ , there exist some constants  $a_i$  such that

$$\phi = \sum_{i=0}^{n-1} a_i \beta_i. \quad (20)$$

Expanding  $\phi$  in Equation (19) using Equation (20) results in an system of linear equations for the solution's coefficient vector  $\vec{a}$ . This linear system is

*Result 2: Finite Element Diffusion Equation*

$$F\vec{a} = (D + M + S)\vec{a} = (\vec{q} + \vec{g}) = \vec{r} \quad (21)$$

where

$$\begin{aligned} D_{ij} &= \int_{\Omega} \kappa_d \vec{\nabla} \beta_i \cdot \vec{\nabla} \beta_j \, dx & \vec{q}_i &= \int_{\Omega} Q_0 \beta_i \, dx \\ M_{ij} &= \int_{\Omega} \sigma_a \beta_i \beta_j \, dx & \vec{g}_i &= \frac{2}{AF} \int_{\partial\Omega} \Gamma_s \beta_i \, ds \\ S_{ij} &= \frac{1}{2A} \int_{\partial\Omega} \beta_i \beta_j \, ds \end{aligned}$$

## 7 QUERY FUNCTION

Finally, to finish the derivation of our finite element (FE) subsurface rendering algorithm, we derive our third key result: the render query function. This function converts the fluence solution of the finite element diffusion equation (FEDE) into exitant subsurface radiance. Due to the approximations of the diffusion equation (DE) and its FE solution, an accurate query function must address two computational issues.

First, as the isotropy condition of the diffusion approximation (DA) breaks down, the DE solution can begin to contain erroneous angular variation. To avoid these errors, this variation must be averaged. Since, as the DA becomes more accurate, the true solution and the average solution converge [23], this averaging tends to remove artifacts without considerably increasing overall error. Second, as part of our FE solution, the fluence  $\phi(\mathbf{x})$  was projected onto a finite

mesh basis. Since for most bases this projection is a more accurate approximation of  $\phi(\mathbf{x})$  than its gradient is an approximation of  $\vec{\nabla} \phi$ , accurate query functions should not compute fluence gradients.

Together the above issues define our query function. To smooth the erroneous angular variation, the initial query averages the subsurface radiance refracted outward from the scattering material.

$$L(\mathbf{x}, \vec{\omega}) = \frac{F_t(\eta, \vec{\omega})}{\pi} \int_{(\vec{\omega} \cdot \vec{n}) > 0} L_d(\mathbf{x}, \vec{\omega})(\vec{\omega} \cdot \vec{n}) \, d\vec{\omega} \quad (22)$$

Next, this initial radiance query is converted to a fluence query by substituting DA (Equation (2)) and simplifying with Equation (10).

$$L(\mathbf{x}, \vec{\omega}) = \frac{F_t(\eta, \vec{\omega})}{4\pi} \left[ \phi(\mathbf{x}) - 2\kappa_d(\mathbf{x})(\vec{n} \cdot \vec{\omega})\phi(\mathbf{x}) \right] \quad (23)$$

Finally, we use the diffusive source boundary condition (DSBC) (Equation (14)) to remove the less accurate the gradient term yielding:

*Result 3: Query Function*

$$L(\mathbf{x}, \vec{\omega}) = \frac{F_t(\eta, \vec{\omega})}{4\pi} \times \left[ \left( 1 + \frac{1}{A(\eta)} \right) \phi(\mathbf{x}) - \frac{4}{F_{dt}(\eta)A(\eta)} \Gamma_s(\mathbf{x}) \right] \quad (24)$$

## 8 IMPLEMENTATION

In the last three sections, we derived the three key results that underlay the basic four-step finite element (FE) rendering algorithm described in Section 3. This section reviews several implementation details that ensure the algorithm's performance and quality.

### 8.1 Tetrahedral Mesh Basis

For our FE basis, we chose a piecewise-linear basis on a tetrahedral mesh. To construct this basis, we associate each mesh vertex with a basis function. Each basis function equals one at its associated vertex and linearly decreases to zero across the one ring of tetrahedra that include this vertex. Elsewhere, the basis function equals zero. The 2D analog of these functions are commonly called "tent" basis functions. We chose to use tetrahedral cells, rather than hexahedral cells like Wang et al. [47], because there are many, well-studied algorithms for quickly generating high-quality, tetrahedral meshes within a triangular surface mesh. For other cell types, this is a more difficult problem. For our examples, we used Tetgen [40] to generate our meshes.

```

SparseMatrix f_mat;
Vector r_vec;

matrix.zero();
rhs.zero();
foreach Tet t in mesh {
  foreach QuadPt pt in Tet {
    foreach Basis i in Tet {
      foreach Basis j in Tet {
        f_mat[i,j] +=
          pt.wt * Kd(pt) * dot(grad(i,pt), (grad(j,pt)));
        f_mat[i,j] +=
          pt.wt * sigA(pt) * value(i,pt) * value(j,pt);
      }
      r_vec[i] += pt.wt * src(pt) * value(i,pt);
    }
  }

  foreach Face f of Tet {
    if(f on boundary) {
      foreach QuadPt pt on f {
        foreach Basis i in Tet {
          foreach Basis j in Tet {
            f_mat[i,j] +=
              (0.5/A) * pt.wt * value(i,pt) * value(j,pt);
          }
          r_vec[i] +=
            (2/A) * pt.wt * gamma(pt) * value(i,pt);
        }
      }
    }
  }
}

```

Fig. 6: Pseudo-code for the assembly algorithm. Here  $Kd(pt)$ ,  $sigA(pt)$ ,  $src(pt)$  and  $gamma(pt)$  are functions that return the values of  $\kappa_d(\mathbf{x})$ ,  $\sigma_a(\mathbf{x})$ ,  $Q_0(\mathbf{x})$  and  $\Gamma_s(\mathbf{x})$  respectively. The functions  $grad(i,pt)$  and  $value(i,pt)$  return the gradient and value respectively of the  $i^{\text{th}}$  basis function. Finally,  $dot$  computes a dot product and  $pt.wt$  is the weight of the quadrature point.

## 8.2 Assembly of Finite Element Matrix

Given a tetrahedral mesh of the domain, the finite element diffusion equation (FEDE) (Equation (21)) can be written as a sum of integrals over the tetrahedra volumes and faces. For example, if  $\mathcal{T}$  is the set of all tetrahedra and  $\Omega_t$  is the volume of tetrahedron  $t$ , an entry in  $D_{ij}$  (see Equation (21)) can be written as

$$D_{ij} = \sum_{t \in \mathcal{T}} \int_{\Omega_t} \kappa_d \vec{\nabla} \beta_i \cdot \vec{\nabla} \beta_j \, dx \quad (25)$$

However, the terms in this sum are only non-zero if the supports of  $\beta_i$  and  $\beta_j$  overlap. For our piecewise-linear basis, this happens only if vertices  $i$  and  $j$  are part of the same tetrahedron. Given this restricted domain of integration, the assembly of Equation (21) can be expressed as a loop over all tetrahedra. Further, we compute integrals like Equation (25) using quadrature. In this case, computing each integral term reduces to computing the sum of the values of the integrand at a series of quadrature points. Our results in Section 9 use 2<sup>nd</sup> order 3D Gaussian quadrature. When using quadrature, the general assembly algorithm reduces by a small set of nested loops over tetrahedra, quadrature points and basis functions (see Figure 6 for the complete pseudo-code).

## 8.3 Material and Source Projection

During the assembly process, the scattering parameters,  $\kappa_d(\mathbf{x})$  and  $\sigma_a(\mathbf{x})$ , and the source functions,  $\Gamma_s(\mathbf{x})$  and  $Q_0(\mathbf{x})$ , are sampled onto a regular spacing of quadrature points. If these terms have sufficiently high frequencies this sampling will cause aliasing that is visible in the solution. To avoid these errors, the source and material terms are first projected onto the FE basis and then sampled. For our piecewise linear basis, this projection corresponds to sampling the functions at the vertices of the tetrahedral mesh and interpolating within each cell.

## 8.4 Solving the Linear System

To solve the FEDE, any sparse matrix algorithm could be used but our implementation uses the conjugate gradient algorithm with the symmetric successive over-relaxation (SSOR) preconditioner. As mentioned in Section 6.1, if the material coefficients and source functions are bounded and have bounded derivatives, this solution exists. However, if these conditions are violated, the derivatives required to construct the FEDE may not exist or the FEDE matrix may be singular. Additionally, using quadrature to construct the FEDE could result in a singular equation even if the exactly computed system would have been invertible. Conveniently, using the pre-projection discussed above avoids both of these issues. However, even without projection, in all our tests of the prototype, including tests with discontinuous material and source functions, our conjugate gradient solver converged.

## 8.5 Adaptive Mesh Refinement

To ensure accuracy across a range of scales, our implementation uses adaptive mesh refinement to capture subsurface detail. However, to support this refinement we must address two important issues: non-conforming meshes and refinement heuristics.

### 8.5.1 Non-conforming Meshes

Adaptively refining a mesh creates T-junction vertices when neighboring tetrahedra have differing levels of refinement. These vertices are problematic because, when present, the span of the resulting basis contains discontinuous functions. These discontinuous functions violate implicit assumptions required to use the divergence theorem (see Equation (17)) to derive the weak form. When T-junctions are present, continuity must be imposed by constraining the coefficients of the t-junction basis functions. Creating and enforcing these constraints is a complex, but solved, problem in FE analysis beyond the scope of this paper (for the details of a complete implementation see [4]). However FE tools exist that create these constraints automatically. Our implementation uses LibMesh [28].

### 8.5.2 Refinement Heuristics

To refine our meshes, we implemented two simple refinement heuristics. Before the FEDE is assembled, our implementation initially refines the mesh using the following conditions, once each, in order:

- 1) all tetrahedra visible from the camera; and then
- 2) all tetrahedra with large changes in the material coefficients.

To choose the tetrahedra in the second condition,  $\sigma_a(\mathbf{x})$  and  $\sigma_s(\mathbf{x})$  are computed at all the vertices of the mesh. Then, for each tetrahedron, the largest relative difference among all pairs of vertices is determined. All tetrahedra with a local difference greater than one deviation above the mean are refined.

We considered a third condition that, similar to the material refinement condition, would refine cells along shadow edges. However, even for scenes with point sources, this condition did not increase accuracy. Subsurface scattering naturally blurs all shadow edges slightly. After both of the above conditions were applied, the surface tetrahedra were already sufficiently small to capture these blurred edges accurately. However, our algorithm supports arbitrary refinement and, if necessary, new refinement heuristics could be trivially implemented.

### 8.6 Single Scattering

Finally, single scattering—light that scatters exactly once in the medium—is usually poorly modeled by the diffusion equation (DE). As noted in previous work [27], accuracy can be improved if the single scattering is simulated more accurately [20] and the diffusion solver is used only for the remaining multiple scattering. To account for the removal of the single scattering in our subsurface computation, we subtract the total power of the single scattering component from the total power of the reduced intensity (RI) source.

## 9 RESULTS

This section demonstrates that our finite element algorithm is an efficient and accurate rendering method for general heterogeneous subsurface scattering problems. The analysis in this section assesses the results of the new renderer on a set of four distinct test scenes. This section has five parts. The first two describe the test scenes and the details of the rendering computation. The last three sections use these results to highlight the advantages of the new renderer. Section 9.3 itemizes the costs of the finite element (FE) algorithm and notes that the new algorithm adds at most 3 minutes to rendering cost; Section 9.4 demonstrates that the test images are nearly identical to exact images produced with a Monte Carlo (MC) path tracer; and Section 9.5 discusses the improvements the new algorithm makes over the most advanced previous method, Wang et al. [47].

Model	Initial Tets	Time <sup>†</sup>	Refined Tets	Time <sup>†</sup>
Bunny	427,918	40.8s	427,918	0s
Dragon	375,919	59.5s	1,084,599	73s
Buddha	429,158	79.0s	1,369,965	103s
Geode	849,763	60.0s	1,317,804	103s

(a) Mesh sizes, generation cost and refinement cost. Costs range from 40-123 seconds.

Model	Base Costs		FE Costs		Total
	Source	Surface	Assembly	Solve <sup>†</sup>	
Bunny	4s	32s	9s	5s	50s
Dragon	28s	59s	35s	71s	193s
Buddha	29s	40s	44s	137s	250s
Geode	152s	108s	38s	88s	386s

(b) Rendering costs by model and category. Rendering times vary from 1 to 6 minutes.

TABLE 1: Summary of the rendering costs and model parameters for our four test scenes: Bunny, Dragon, Buddha and Geode. Operations marked with <sup>†</sup> are run on a single processor.

### 9.1 Scenes

The prototype implementation was tested using the scenes shown in Figure 1. Each has a different geometry and material. The test scenes were chosen to demonstrate the range of effects a general, high-quality heterogeneous solver can reproduce.

**Bunny** uses a marble texture to simulate scattering in a complex, aggregate material. The bunny scene uses a relatively smaller mesh and simple lighting. This scene emphasizes that in small scenes the FE algorithm only adds a few seconds to the image cost when compared to Wang et al. [47].

**Dragon** is modeled using an optically thinner material similar to translucent plastic. This model shows that our solution can capture smooth changes in color and opacity.

**Buddha** contains a checkerboard of homogeneous marble and jade-like materials. The material is difficult because the solver must simultaneously capture the sharp edges in the material properties but also correctly simulate the smooth translucency in thinner geometry.

**Geode** pushes the limits of our FE algorithm. It demonstrates that the FE approach can easily scale to capture complex, high-frequency scattering even in a very difficult lighting environment where all light passes through the material.

### 9.2 Details of the Rendering Computation

All results were generated on a 8 × 2.66Ghz Xeon workstation with 8GB of RAM and all images are 640×480 pixels. Tables 1(a) and 1(b) summarize the model size, mesh creation costs and rendering time for each scene. All of these costs are fully parallelized except the mesh creation, the mesh refinement and the FE matrix solution costs. Since parallel implementations of these operations were not available, these

times are performed on a single core. Each image is 16x anti-aliased but, for our FE solver, anti-aliasing is essentially free. Once our FE solution is computed, the computation of each sample reduces to a single evaluation of the render query function.

The surface and single scattering components of our images are computed using a combination of Multidimensional Lightcuts (MDLC) [45] and an analytical single scattering approximation [20]. Our FE implementation is split between Java, which provides an implementation of MDLC, and C++, which provides an interface to the LibMesh [28] library. The LibMesh library provides our conjugate gradient matrix solver and mesh iteration and refinement functions.

The Geode model is lit by a small area source, Buddha and Dragon are lit by the Kitchen environment map [9] and the Bunny is lit by a small spherical source. All images, unless otherwise noted, include global illumination approximated by 100,000 virtual indirect sources.

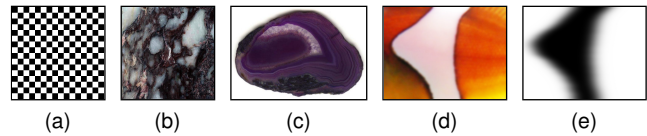
### 9.2.1 Reference Path Tracer

Our reference path tracer is highly optimized. It is fully parallelized and importance samples all subsurface paths. Further, we dramatically accelerate the path tracer by modeling the heterogeneous material as a large, piecewise constant grid of homogeneous cells. To create the path tracing grid, we assign each grid cell the scattering parameters of its center point. Since the interior of each grid cell is homogeneous, the path tracer can generate path segments within each cell without ray marching or repeated material queries. We chose the grid size so that this change had no effect on the final image quality. For all scenes, a grid with 256 cells in the longest dimension was sufficient. We chose the size of the remaining dimensions so the cells were cubical<sup>8</sup>. Without this optimization, the cost of the reference path tracer would have been prohibitive.

### 9.2.2 Material Parameters

The material models used in the test scenes were synthesized by orthographically projecting an image through the scattering geometry and using the pixel values to determine values for the scattering coefficients  $\sigma_a$  and  $\sigma_s$ . The parameters used in this generation, as well as the source images, are provided in Table 2. Each material coefficient is specified by three colors and an image. The first two colors, range min and max, are used to rescale the image to the dynamic range of the scattering parameter and the last color sets the parameter’s overall scale. To compute a scattering parameter at a particular point, one finds the corresponding pixel in the rescaled image, inverts it and multiplies by the base scale. The inversion is

8. To ensure the sharp edges in the Buddha scene, we aligned this rendering grid with the larger checker grid in the Buddha’s material.



Param.	Range Min	Range Max	Base Scale	Tex.
Dragon $\sigma_a$	(.05, .05, .05)	(1.0, 1.0, 1.0)	(0.75, 1.25, 1.75)	(d)
Dragon $\sigma_s$	(.25, .25, .25)	(1.0, 1.0, 1.0)	(16.7, 16.7, 16.7)	(e)
Geode $\sigma_a$	(.01, .01, .01)	(1.0, 1.0, 1.0)	(5.0, 5.0, 5.0)	(c)
Geode $\sigma_s$	constant	constant	(5.0, 5.0, 5.0)	–
Buddha $\sigma_a$	constant	constant	(1.63, 1.18, 4.5)	–
Buddha $\sigma_s$	(.05, .05, .05)	(1.0, 1.0, 1.0)	(16.7, 16.7, 16.7)	(a)
Bunny $\sigma_a$	(.05, .05, .05)	(3.5, 3.5, 3.5)	(7.8, 7.8, 7.8)	(b)
Bunny $\sigma_s$	(.60, .60, .60)	(1.0, 1.0, 1.0)	(13.1, 15.7, 18.0)	(b)

TABLE 2: Parameters used to procedurally generate the volume scattering textures. Section 9.2.2 describes our procedural material model.

necessary because the material parameters specify how much light is removed during scattering whereas, in the image, the colors specify how much light should be added to each pixel.

## 9.3 Performance

Table 1 breaks down the cost of the new FE algorithm by component. From new scene to final image, the system requires between 2 and 10 minutes. This cost can be roughly divided into 2 parts: meshing costs and rendering costs. First, Table 1(a) gives mesh sizes and the costs of mesh generation and refinement. Because the FE algorithm is agnostic to how the scattering domain is discretized, it can use the most efficient meshing algorithms available. For unstructured triangular surface meshes, tetrahedralization is fast (1-2 minutes for all examples) and produces a high quality discretization [40].

Second, Table 1(b) gives the cost to render each image after the mesh has been computed. For the four test scenes, these costs total between 50s and 6 minutes. These costs are further split into two categories. The base costs—computing the boundary source and rendering the surface component—are performed by a separate surface rendering method and are independent of the subsurface rendering algorithm. Only the costs of assembling the finite element diffusion equation (FEDE) and solving it are inherent to our new FE algorithm. For all test scenes, these costs total less than 3 minutes and, except for the Buddha image, are less than half of the rendering cost. Thus, in most scenes, the cost of adding subsurface scattering roughly equals the cost of rendering high-quality images without subsurface scattering. For a simple scenes where this base cost is smaller, like the Bunny, the our FE algorithm adds only 15s to the total render time.

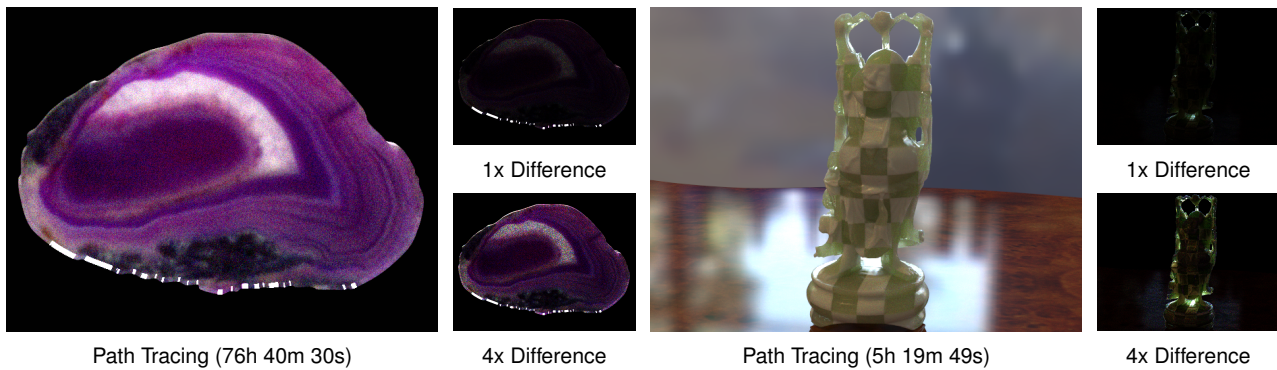


Fig. 7: Path traced reference images for Geode and Buddha from Figure 1. Their absolute difference and 4x magnified difference are presented on the right of each image.

#### 9.4 Comparison with Monte Carlo

Demonstrating the quality of our algorithm, Figures 1 and 7 compare our Dragon, Buddha and Geode images with the same images produced by an exact, MC path tracer. Compared to previous work in subsurface rendering, there is impressive agreement between these two sets of images. However, neither set of images is perfect. Despite thorough optimization (Section 9.2.1 above), the path tracing algorithm is still expensive. To avoid excessive computation, the path tracing was performed progressively and was stopped as soon as the noise fell below a level permitting a reasonable comparison. Computing noise-free images would at least double the multiple-hour cost of the path traced references. As expected, the new algorithm has a considerable advantage in performance. It generates noise-free images in a few minutes.

To facilitate the quality comparison, absolute error images are provided. These error images are nearly black, so 4x magnified versions reveal differences. There are two principle differences. First, particularly prominent in the Buddha and Geode images, the path tracer is able to capture highlights from caustic paths that do not scatter within the material. However, since these caustic paths are not part of the subsurface scattering, this is an error of our MDLC surface render, not our FE algorithm. Second, because the diffusion equation (DE) treats all diffusive radiance as nearly isotropic, it tends to overestimate the frequency of scattering in thin geometry and near the surface and, as a consequence, it underestimates contributions from low order scattering events. This slightly darkens regions when they are lit from behind as in the optically thinner parts of the Dragon and Geode and slightly lightens highly absorptive regions viewed directly, as in the darker checkers on the Buddha. However, overall, these differences are only visible in the magnified errors images and result from fundamental limitations in the diffusion approximation (DA) Given the orders-of-magnitude difference in performance, these results demonstrate that our FE algorithm is suitable even for high quality rendering applications.

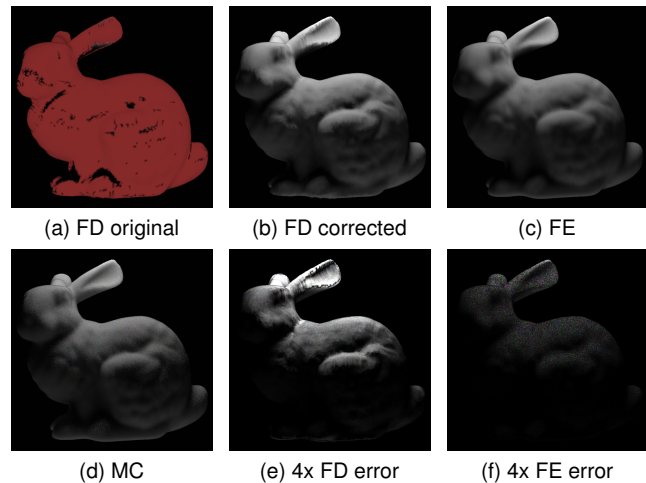


Fig. 8: Images of a constant scattering white bunny lit by two area lights, fill below and key above: Figure 8(a) as described in Wang et al. [47] with negative radiance areas highlighted in red; Figure 8(b) corrected using derivation in Section 5; Figure 8(c) our FE algorithm; Figure 8(d) MC reference; Figure 8(e) 4x absolute error of Figure 8(b); and Figure 8(f) 4x absolute error of Figure 8(c).

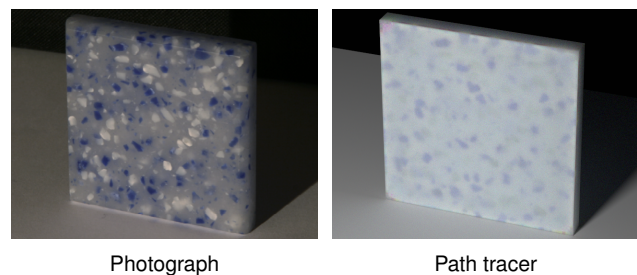


Fig. 9: Comparison of the photograph artificial stone slab captured by Wang et al. [47] to a path traced rendering of the resulting material parameters. The authors have confirmed that this comparison is accurate.

#### 9.5 Comparison to Wang et al. [2008]

As a final test, we compare our new renderer to the best existing approach, Wang et al. [47]. Since the focus of our new algorithm is quality, rather than performance, the comparison is made to a software version of their iterative finite difference (FD) algorithm. This software renderer omits the multi-resolution approximation originally required to achieve interactive performance.

For our comparisons, we allow the FD algorithm to update fully each step and iterate until the solution converges. However, despite these improvements, we demonstrate that the FD algorithm has lower quality and requires an expensive PolyGrid [43] mesh that introduces rendering errors and diverges for some materials.

To facilitate our comparison, Wang et al. [47] kindly provided their measured material data and a PolyGrid [43] bunny model. As an initial test, several images of the bunny with a white, homogeneously scattering material (see Figure 8) were created. Unfortunately, as noted in Section 2.4.1 and [3], their original diffusion formulation was incorrect and, as a result, their original algorithm sometimes computes negative radiance solutions. As shown in red in Figure 8(a), this happens almost everywhere on the homogeneous bunny. In the original work, this error was corrected by altering the material parameters during the capture optimization and, as a result, this error did not manifest itself in any of the authors’ original results. Figure 9 demonstrates this correction. It compares a path traced rendering of the captured artificial stone material with a photograph of the original object. The lack of saturation in the MC result suggests that the acquired parameters have significantly less absorption, making the image brighter. Wang et al. [47] have confirmed the mismatch in Figure 9. In order to make the remaining comparisons in this section fair, we use our correct formulation in further comparisons with their approach.

The rest of the images in Figure 8 directly compare, for meshes of equal size, Wang et al.’s FD algorithm, the new FE algorithm and a path traced reference (Figures 8(b), 8(c) and 8(d) respectively). Figures 8(e) and 8(f) display the error of the FD and FE methods. Since both algorithms depend on the DE, neither solution can produce an exact answer. However, because the FD algorithm relies on a special PolyGrid mesh [43], it has at least three additional sources of error.

- 1) The diffusive source boundary condition (DSBC) is enforced only approximately by using special, smaller PolyGrid cells on the boundary.
- 2) The overall distortion of the PolyGrid can be only approximately modeled during the FD solution.
- 3) Creating the uniformly connected PolyGrid required of the FD solver requires deleting some nodes to along all boundary edges of the grid. This introduces error in the solution near these edges and especially at the grid corners.

Additionally, for certain materials, the iterative FD algorithm can also be unstable. In Figure 10, the Bunny model is rendered with three similar materials using both the FD and FE algorithms. The first material (left) is homogeneous and the two methods are mostly in agreement. However, in the center and right columns, a checker board is introduced by scaling the mean free

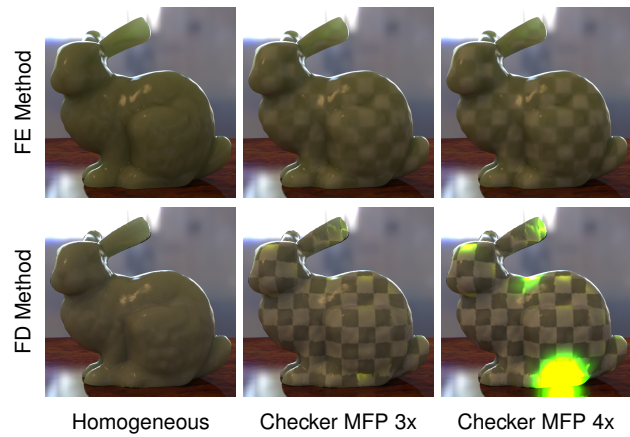


Fig. 10: An example of divergence in the iterative FD solver used by Wang et al. [47]. The bunny’s material is varied from a homogeneous green material (left) toward the material used in the Buddha (see Figure 1) by scaling the mean free path (MFP) of the material in the lighter squares. The FD algorithm diverges beyond a 3x scale (middle; see spots in ear, head and foot) and prominently at a 4x scale (right).

path of the material in alternating sections. As this happens, the FD method diverges. In the middle and right columns, the mean free path has been reduced by factors of 3x and 4x respectively. For these cases, the FD algorithm was stopped after 500 iterations to prevent the fluence from overflowing. In both images, the solution is beginning to diverge in the ears, head and foot. Our algorithm does not have this instability.

Clearly, unlike Wang et al. [47], our algorithm is not interactive but, to compare performance, we must note that in Wang et al. [47] only part of the complete rendering calculation is interactive. The costs of mesh construction were not presented in [47] but the author’s noted they required hand optimization which could take hours. Our automated method requires a few minutes (see Table 1(a)). When considering rendering costs, the surface and source computations are independent of either algorithm and could, as demonstrated in [47], be done in real-time on the GPU. We forgo this optimization for our results because it introduces considerable error and would prevent our overall solution from matching the path traced reference. In the end, the only directly comparable costs are our matrix assembly and its solution and, for a simple scene (see Figure 1 and Table 1, Bunny) like those in Wang et al. [47], these steps add only 15s to the image cost for our mainly single threaded, CPU implementation.

## 10 CONCLUSION

In this paper, we presented an efficient, general and high-quality rendering algorithm for complex heterogeneous materials. To create this algorithm, we corrected errors in previous work and derived an accurate diffusive scattering problem. Then, by accurately solving that problem with the finite element (FE) method, we reduced subsurface rendering to a

simple and adaptable four-step algorithm. To validate this algorithm, we created a general implementation and tested its accuracy on a series of four difficult scenes. These results demonstrate that our algorithm can render images in a few minutes that are nearly identical to accurate path traced images produced in hours and that this algorithm significantly improves upon both the quality and generality of the best previous methods.

### 10.1 Limitations and Future Work

Of course, despite its advantages, our algorithm has its own limitations and future work remains. First, our solution is directly limited in accuracy by the diffusion approximation (DA) as a scattering model. Similar to Li et al. [32], interesting future work could consider detecting where the diffusive simulation is inaccurate and compute these inaccurate regions with a more accurate algorithm. Second and related, there is the more general question of rendering diffusive and non-diffusive aggregates of material, such as a crystal, which might have both translucent *and* transparent regions randomly dispersed throughout. Third, by formulating subsurface scattering as a general FE problem, we open the door to the application of the full breadth of FE theory, including robust adaptive refinement, higher order bases, and multi-resolution solvers. New work could explore how to better leverage these tools to improve the quality and performance of algorithms like ours. Finally, there is immediate future work in considering how to parallelize and enhance the performance of our algorithm both on and off the GPU.

### REFERENCES

- [1] R. T. Ackroyd, *Finite Element Methods for Particle Transport: Applications to Reactor and Radiation Physics*. Research Studies, 1997.
- [2] A. Arbree, B. Walter, and K. Bala, "Single-pass scalable subsurface rendering with lightcuts," *Computer Graphics Forum*, vol. 27, no. 2, pp. 507–516, April 2008. [Online]. Available: <http://dx.doi.org/10.1111/j.1467-8659.2008.01148.x>
- [3] —, "Diffusion formulation for heterogeneous subsurface scattering," Cornell University Computing and Information Science, Tech. Rep. <http://hdl.handle.net/1813/14199>, 2009.
- [4] W. Bangerth and O. Kayser-Herold, "Data structures and requirements for *hp* finite element software," Institute for Scientific Computation, Texas A&M University, Tech. Rep. ISC-07-04-MATH, 2007.
- [5] N. A. Carr, J. D. Hall, and J. C. Hart, "Gpu algorithms for radiosity and subsurface scattering," in *HWWS '03: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 2003, pp. 51–59.
- [6] C.-W. Chang, W.-C. Lin, T.-C. Ho, T.-S. Huang, and J.-H. Chuang, "Real-time translucent rendering using gpu-based texture space importance sampling," *Computer Graphics Forum*, vol. 27, no. 2, pp. 517–526, apr 2008.
- [7] Y. Chen, X. Tong, J. Wang, S. Lin, B. Guo, and H.-Y. Shum, "Shell texture functions," in *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*. New York, NY, USA: ACM, 2004, pp. 343–353.
- [8] C. Dachsbacher and M. Stamminger, "Translucent shadow maps," in *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 2003, pp. 197–201.
- [9] P. Debevec, "Image-based lighting," *IEEE Comput. Graph. Appl.*, vol. 22, no. 2, pp. 26–34, 2002.
- [10] C. Donner and H. W. Jensen, "Light diffusion in multi-layered translucent materials," in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*. New York, NY, USA: ACM, 2005, pp. 1032–1039.
- [11] —, "Rendering translucent materials using photon diffusion," in *Eurographics Symposium on Rendering 2007*, 2007, pp. 243–251.
- [12] C. Donner, T. Weyrich, E. d'Eon, R. Ramamoorthi, and S. Rusinkiewicz, "A layered, heterogeneous reflectance model for acquiring and rendering human skin," *ACM Trans. Graph.*, vol. 27, no. 5, pp. 1–12, 2008.
- [13] J. Dorsey, A. Edelman, H. W. Jensen, J. Legakis, and H. K. Pedersen, "Modeling and rendering of weathered stone," in *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1999, pp. 225–234.
- [14] L. C. Evans, *Partial Differential Equations*. American Mathematical Society, 1998.
- [15] T. J. Farrell, M. S. Patterson, and B. Wilson, "A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo," *Medical Physics*, vol. 19, no. 4, pp. 879–888, 1992.
- [16] A. Ghosh, T. Hawkins, P. Peers, S. Frederiksen, and P. Debevec, "Practical modeling and acquisition of layered facial reflectance," *ACM Trans. Graph.*, vol. 27, no. 5, pp. 1–10, 2008.
- [17] A. P. Gibson, J. C. Hebden, and S. R. Arridge, "Recent advances in diffuse optical imaging," *Physics in Medicine and Biology*, vol. 50, no. 4, pp. R1–R43, 2005.
- [18] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H.-P. Seidel, "Disco: acquisition of translucent objects," in *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*. New York, NY, USA: ACM, 2004, pp. 835–844.
- [19] T. Haber, T. Mertens, P. Bekaert, and F. V. Reeth, "A computational approach to simulate subsurface light diffusion in arbitrarily shaped objects," in *GI '05: Proceedings of Graphics Interface 2005*. School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada: Canadian Human-Computer Communications Society, 2005, pp. 79–86.
- [20] P. Hanrahan and W. Krueger, "Reflection from layered surfaces due to subsurface scattering," in *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 1993, pp. 165–174.
- [21] X. Hao, T. Baby, and A. Varshney, "Interactive subsurface scattering for translucent meshes," in *I3D '03: Proceedings of the 2003 symposium on Interactive 3D graphics*. New York, NY, USA: ACM, 2003, pp. 75–82.
- [22] X. Hao and A. Varshney, "Real-time rendering of translucent meshes," *ACM Trans. Graph.*, vol. 23, no. 2, pp. 120–142, 2004.
- [23] A. Ishimaru, *Wave Propagation and Scattering in Random Media*. Academic Press, 1978.
- [24] H. W. Jensen and J. Buhler, "A rapid hierarchical rendering technique for translucent materials," in *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 2002, pp. 576–581.
- [25] H. W. Jensen and P. H. Christensen, "Efficient simulation of light transport in scenes with participating media using photon maps," in *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 1998, pp. 311–320.
- [26] H. W. Jensen, J. Legakis, and J. Dorsey, "Rendering of wet materials," in *Rendering Techniques '99*, 1999, pp. 273–282.
- [27] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan, "A practical model for subsurface light transport," in *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 2001, pp. 511–518.
- [28] B. Kirk, J. W. Peterson, R. H. Stogner, and G. F. Carey, "libMesh: A C++ Library for Parallel Adaptive Mesh Refinement/Coarsening Simulations," *Engineering with Computers*, vol. 22, no. 3–4, pp. 237–254, 2006.
- [29] J. J. Koenderink and A. J. van Doorn, "Shading in the case of translucent objects," *Proceedings of the SPIE: Human Vision and Electronic Imaging VI*, vol. 4299, pp. 312–320, 2001. [Online]. Available: <http://link.aip.org/link/?PSI/4299/312/1>

- [30] V. Kolehmainen, "Novel approaches to image reconstruction indiffusion tomography," Ph.D. dissertation, Kuopio University, 2001.
- [31] H. P. A. Lensch, M. Goesele, P. Bekaert, J. Kautz, M. A. Magnor, J. Lang, and H.-P. Seidel, "Interactive rendering of translucent objects," in *PG '02: Proceedings of the 10th Pacific Conference on Computer Graphics and Applications*. Washington, DC, USA: IEEE Computer Society, 2002, p. 214.
- [32] H. Li, F. Pellacini, and K. E. Torrance, "A hybrid monte carlo method for accurate and efficient subsurface scattering," in *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, jun 2005, pp. 283–290.
- [33] T. Mertens, J. Kautz, P. Bekaert, F. V. Reeth, and H.-P. Seidel, "Efficient rendering of local subsurface scattering," in *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*. Washington, DC, USA: IEEE Computer Society, 2003, p. 51.
- [34] T. Mertens, J. Kautz, P. Bekaert, H.-P. Seidel, and F. V. Reeth, "Interactive rendering of translucent deformable objects," in *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*. Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 2003, pp. 130–140.
- [35] K. W. Morton, *Numerical Solution of Convection-Diffusion Problems*. Chapman and Hall, 1996.
- [36] M. Pauly, T. Kollig, and A. Keller, "Metropolis light transport for participating media," in *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*. London, UK: Springer-Verlag, 2000, pp. 11–22.
- [37] P. Peers, K. vom Berge, W. Matusik, R. Ramamoorthi, J. Lawrence, S. Rusinkiewicz, and P. Dutré, "A compact factored representation of heterogeneous subsurface scattering," in *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*. New York, NY, USA: ACM, 2006, pp. 746–753.
- [38] M. Pharr and P. Hanrahan, "Monte carlo evaluation of non-linear scattering equations for subsurface reflection," in *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 75–84.
- [39] M. Schweiger, S. R. Arridge, M. Hiraoka, and D. T. Delpy, "The finite element method for the propagation of light in scattering media: Boundary and source conditions," *Medical Physics*, vol. 22, no. 11, pp. 1779–1792, 1995.
- [40] H. Si and K. Gaertner, "Meshing piecewise linear complexes by constrained delaunay tetrahedralizations," in *Proceedings of the 14th International Meshing Roundtable*, sep 2005, pp. 147–163.
- [41] P.-P. Sloan, B. Luna, and J. Snyder, "Local, deformable precomputed radiance transfer," in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*. New York, NY, USA: ACM, 2005, pp. 1216–1224.
- [42] J. Stam, "Multiple Scattering as a Diffusion Process," in *Eurographics Rendering Workshop 1995*. Eurographics, 1995, pp. 41–50.
- [43] M. Tarini, K. Hormann, P. Cignoni, and C. Montani, "Polycube-maps," in *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*. New York, NY, USA: ACM, 2004, pp. 853–860.
- [44] X. Tong, J. Wang, S. Lin, B. Guo, and H.-Y. Shum, "Modeling and rendering of quasi-homogeneous materials," in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*. New York, NY, USA: ACM, 2005, pp. 1054–1061.
- [45] B. Walter, A. Arbree, K. Bala, and D. P. Greenberg, "Multidimensional lightcuts," in *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*. New York, NY, USA: ACM, 2006, pp. 1081–1088.
- [46] B. Walter, S. Fernandez, A. Arbree, K. Bala, M. Donikian, and D. P. Greenberg, "Lightcuts: a scalable approach to illumination," in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*. New York, NY, USA: ACM, 2005, pp. 1098–1107.
- [47] J. Wang, S. Zhao, X. Tong, S. Lin, Z. Lin, Y. Dong, B. Guo, and H.-Y. Shum, "Modeling and rendering of heterogeneous translucent materials using the diffusion equation," *ACM Trans. Graph.*, vol. 27, no. 1, pp. 1–18, 2008.
- [48] R. Wang, J. Tran, and D. Luebke, "All-frequency interactive relighting of translucent objects with single and multiple scattering," in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*. New York, NY, USA: ACM, 2005, pp. 1202–1207.

**Adam Arbree** Adam's Bio.

**Bruce Walter** Bruce's Bio.

**Kavita Bala** KB's Bio