

High-Order Similarity Relations in Radiative Transfer

Shuang Zhao
Cornell University

Ravi Ramamoorthi
University of California, Berkeley

Kavita Bala
Cornell University




Figure 1: We introduce a new approach utilizing high-order similarity relations, which can be used to accelerate Monte Carlo rendering of translucent materials. (a) Reference path-traced rendering of a Corinthian capital made of a translucent material with a complicated phase function. (b) Image rendered with the same algorithm but using a reduced scattering coefficient and an isotropic phase function: although a 3.6X speedup is obtained, the resulting accuracy is poor (see the included relative error visualization with the color mapping shown above the renderings). (c) Image rendered using the same reduced scattering coefficient as (b) and a phase function provided by our method: with a slightly higher speedup, significantly better accuracy is obtained. (d) Plots of the phase functions used in (a, b, c). Our theory permits finding a tabulated function (the orange curve) accurately reproducing the reference appearance.

Abstract

Radiative transfer equations (RTEs) with different scattering parameters can lead to identical solution radiance fields. Similarity theory studies this effect by introducing a hierarchy of equivalence relations called “similarity relations”. Unfortunately, given a set of scattering parameters, it remains unclear how to find altered ones satisfying these relations, significantly limiting the theory’s practical value. This paper presents a complete exposition of similarity theory, which provides fundamental insights into the structure of the RTE’s parameter space. To utilize the theory in its general high-order form, we introduce a new approach to solve for the altered parameters including the absorption and scattering coefficients as well as a fully tabulated phase function. We demonstrate the practical utility of our work using two applications: forward and inverse rendering of translucent media. Forward rendering is our main application, and we develop an algorithm exploiting similarity relations to offer “free” speedups for Monte Carlo rendering of optically dense and forward-scattering materials. For inverse rendering, we propose a proof-of-concept approach which warps the parameter space and greatly improves the efficiency of gradient descent algorithms. We believe similarity theory is important for simulating and acquiring volume-based appearance, and our approach has the potential to benefit a wide range of future applications in this area.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—Three-Dimensional Graphics and Realism

Keywords: radiative transfer, subsurface scattering

Links: [DL](#) [PDF](#) [WEB](#)

1 Introduction

Many real-world materials including marble, jade, and human skin exhibit distinctive appearances arising from subsurface scattering of light. Understanding, simulating, and measuring this phenomenon has been an active research area in computer graphics for decades.

The physics of subsurface scattering is normally modeled with the radiative transfer framework [Chandrasekhar 1960]. The core of this framework is the radiative transfer equation (RTE) which governs how frequently light scatters and how it gets redirected or absorbed (when scattering occurs) via a set of scattering parameters.

The parameter space of the RTE contains infinitely many equivalence classes such that different parameters in each class lead to identical solution radiance fields, under the assumption that these radiance fields have bounded angular frequencies. Similarity theory, introduced to applied physics by Wyman et al. [1989a], studies this property by deriving a hierarchy of equivalence relations called “similarity relations”. Higher-order similarity relations offer finer partitions of the space, and parameters in the resulting equivalence classes can produce identical radiance with higher frequencies.

Previously, only a special case of the simplest order-1 similarity relation has been used in computer graphics. Furthermore, given a set of scattering parameters, computing an altered set adhering to (higher-order) similarity relations remains a challenge, especially

for the altered phase function which is not uniquely determined by the relations.

In this paper, we present a complete exposition of similarity theory and introduce practical algorithms to utilize this theory in its general higher-order forms. Our theoretical contributions include:

- Introducing to graphics the full derivation of similarity relations (Section 4.1) and discussing their connection to diffusion theory (Section 4.2).
- Developing novel algorithms to determine the existence of and to solve for the parameters (including absorption/scattering coefficients and a tabulated phase function) satisfying similarity relation of any given order (Section 5).

Our theory can lead to practical applications in forward and inverse rendering of translucent media. The presence of equivalence classes is a significant challenge in inverse rendering. This is because different parameters can provide very similar appearances, causing the optimization problem to be ill-conditioned. Section 6 introduces a proof-of-concept method reparameterizing the search space, so that gradient descent algorithms become much more effective in the new space.

Forward rendering is our main application (Section 7). We develop a simple procedure that takes a set of scattering parameters and outputs an altered set in a few seconds. Replacing the original parameters with the altered ones can accelerate Monte Carlo rendering of optically dense and forward-scattering media (over 3X speedups can be achieved for volume path tracing). A key benefit offered by our method is that no changes need to be made to core rendering algorithms: only material scattering parameters, which are inputs to the renderer, are modified.

2 Related Work

Because of the breadth of research on modeling, rendering, and measuring subsurface scattering, a complete survey is beyond the scope of this paper. In this section, we quickly review prior work in five categories: radiative transfer, Monte Carlo methods, diffusion methods, similarity theory, and inverse rendering.

Radiative Transfer. Radiative transfer is used in many areas including astrophysics, neutron transport, and computer graphics [Chandrasekhar 1960; Ishimaru 1978]. Recently, Jakob et al. [2010] introduced a generalized framework to better handle scattering media with oriented structures, such as volumetric fabrics. In this paper, we focus on the classical form of the RTE.

Monte Carlo Methods. Monte Carlo methods solve the full RTE directly. Volumetric path tracing and its variations [Kajiya and Von Herzen 1984; Lafortune and Willems 1996; Pauly et al. 2000] provide unbiased estimators for the solution radiance by randomly constructing light paths and evaluating their contributions. In addition, various techniques such as volumetric photon mapping [Hachisuka et al. 2012] and many-lights methods [Dachsbacher et al. 2014] have been developed, which offer faster convergence than path tracing methods but often at the cost of introducing bias in the results.

Diffusion Methods. Diffusion methods replace the RTE with the diffusion equation (DE) by applying a first-order approximation to directional radiance [Ishimaru 1978]. Many approaches have been proposed to solve the DE including approximated analytical solutions [Jensen et al. 2001; D'Eon and Irving 2011] and finite element (or finite difference) based methods [Stam 1995; Wang et al. 2008; Arbre et al. 2011]. The diffusion approximation requires the resulting radiance field to be smooth, which is usually

violated near material boundaries and in optically thin regions. Consequently, hybrid methods [Li et al. 2005; Donner and Jensen 2007; Habel et al. 2013] combine Monte Carlo methods and diffusion for better accuracy.

Similarity Theory. Similarity theory was introduced by Wyman et al. [1989a; 1989b] in applied physics. The authors derived a set of relations between two sets of scattering parameters so that the resulting RTEs have identical solution radiance fields when their directional frequencies are bounded. A highly simplified order-1 form of this theory is used extensively in diffusion methods and has been applied for accelerating Monte Carlo simulation of light transport [Chatigny et al. 1999; Frisvad et al. 2007]. However, very limited work has been done, in both computer graphics and applied physics, to utilize such relations at higher orders.

Inverse Rendering. Inverse rendering methods solve for the material properties in a scene given the desired appearance and have many applications in appearance acquisition. Multiple methods have been developed to recover subsurface scattering properties [Wang et al. 2008; Dobashi et al. 2012; Papas et al. 2013; Gkioulekas et al. 2013b]. We show in this paper that similarity theory can be helpful for solving the inverse volume rendering problem.

3 Overview

In this section, we first briefly revisit the basic concept of radiative transfer and present a mathematical description of similarity theory. Then, we describe the computational challenges for forward and inverse rendering of translucent media and our plan to tackle these challenges.

Radiative Transfer. The radiative transfer equation (RTE) [Chandrasekhar 1960], in the form usually used in graphics, is

$$(\boldsymbol{\omega} \cdot \nabla)L(\boldsymbol{\omega}) = -\sigma_t L(\boldsymbol{\omega}) + \sigma_s \int_{\mathbb{S}^2} f(\boldsymbol{\omega}' \cdot \boldsymbol{\omega})L(\boldsymbol{\omega}') d\boldsymbol{\omega}' + Q(\boldsymbol{\omega})$$

where the spatial dependencies are dropped for notational convenience. This equation describes that the directional derivative of the radiance field L is determined by its value via the *attenuation coefficient* σ_t , an integral of L at the same location via the *scattering coefficient* σ_s and the *phase function* $f(\cdot)$, and the *source term* Q . In addition, $\sigma_t = \sigma_a + \sigma_s$ where σ_a is called the *absorption coefficient*. For participating media with no internal source, Q vanishes, yielding

$$(\boldsymbol{\omega} \cdot \nabla)L(\boldsymbol{\omega}) = -\sigma_t L(\boldsymbol{\omega}) + \sigma_s \int_{\mathbb{S}^2} f(\boldsymbol{\omega}' \cdot \boldsymbol{\omega})L(\boldsymbol{\omega}') d\boldsymbol{\omega}'. \quad (1)$$

Note that the phase function $f(\cdot)$ is assumed to be a 1D function of $(\boldsymbol{\omega} \cdot \boldsymbol{\omega}')$. Generalizing $f(\cdot)$ to a full 4D function (of $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$) is possible [Jakob et al. 2010] but beyond the scope of this paper.

Similarity Theory. Theoretically, assuming the solution radiance L in the RTE (1) is band-limited (in the spherical harmonics, or SH, domain), there exist altered parameters σ_t^* , σ_s^* , and $f^*(\cdot)$, such that the corresponding altered RTE

$$(\boldsymbol{\omega} \cdot \nabla)L(\boldsymbol{\omega}) = -\sigma_t^* L(\boldsymbol{\omega}) + \sigma_s^* \int_{\mathbb{S}^2} f^*(\boldsymbol{\omega}' \cdot \boldsymbol{\omega})L(\boldsymbol{\omega}') d\boldsymbol{\omega}' \quad (2)$$

has a solution which equals that of (1) exactly. Similarity theory [Wyman et al. 1989a] describes the relations between the altered parameters and the original ones, which are presented in Section 4. Based on these relations, the parameter space can be partitioned into multiple equivalence classes. In practice, when the assumption does not hold perfectly, parameters in one equivalence class produce approximately identical appearances.

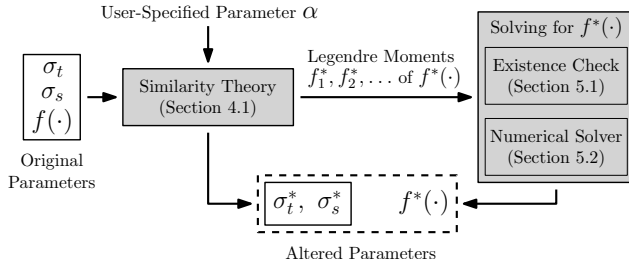


Figure 2: Our pipeline to speedup forward rendering of translucent media. It takes the original scattering parameters as well as a user-specified $\alpha \in (0, 1)$ and outputs the altered parameters.

Inverse Rendering: Challenges. Given the complicated and highly non-linear relation between the scattering parameters and the resulting appearance, inverse rendering is usually modeled as an optimization problem where the parameter space needs to be explored locally [Wang et al. 2008; Hašan et al. 2010; Dobashi et al. 2012; Gkioulekas et al. 2013b]. This process, however, can be extremely expensive as the search space is often high-dimensional, and exploring it requires iteratively solving the forward problem, which is challenging by itself. The presence of the equivalence classes makes the inverse problem even more difficult as it creates ambiguities between different sets of parameters which can cause the optimization to be ill-conditioned.

Inverse Rendering: Our Approach. Similarity theory suggests the locations in the parameter space where the ambiguities occur. Section 4.3 illustrates such an example using a simple 2D space. Based on this understanding, we introduce a proof-of-concept method (Section 6) which warps the space in a non-linear manner so that it becomes much easier for gradient based methods to find good solutions.

Forward Rendering: Challenges. Forward rendering of translucent materials is our main application. It requires solving (1) which in general has no closed-form solution, and accurate numerical solutions often involve Monte Carlo simulations. To render a normal sized object made of optically dense materials, such as milk and marble, hundreds or thousands of subsurface scatterings need to be simulated on each light path, yielding slow performance. Among such materials, the highly forward-scattering ones are particularly difficult to handle. They include phase functions that send light into very concentrated regions, causing path tracing based algorithms to produce high noise, and photon mapping or many-lights methods to require a massive number of photons or virtual lights to avoid intense artifacts or energy loss.

Forward Rendering: Our Approach. We tackle the challenge of rendering optically dense and forward-scattering materials using similarity theory. Particularly, we look for equivalent parameters $\sigma_t^*, \sigma_s^*, f^*(\cdot)$ with $\sigma_t^* < \sigma_t$ because a smaller attenuation coefficient means fewer scattering events to simulate and less computation required. The most basic version of this idea, which reduces σ_s and sets $f^*(\cdot)$ to isotropic, has been used in graphics but can result in poor accuracy (Section 4.2).

Figure 2 previews the pipeline of our method. The user provides the original scattering parameters $\sigma_t, \sigma_s, f(\cdot)$ as well as an extra parameter $\alpha \in (0, 1)$ controlling the tradeoff between performance and accuracy. The first component of our pipeline computes the altered absorption and scattering coefficients based on similarity theory (Section 4.1). The altered phase function $f^*(\cdot)$, however, is not directly given. Instead, similarity theory specifies the desired Legendre moments f_1^*, f_2^*, \dots of $f^*(\cdot)$. The second component of the pipeline then numerically solves for $f^*(\cdot)$ as a tabulated function given those moments (Section 5).

Our approach is easy to implement (pseudocode is in Section 7.1 and our MATLAB implementation is available as supplementary material) and straight-forward to use: the user can simply replace the original scattering parameters with the altered ones (which are the outputs of our pipeline). The base rendering method does not need to be changed. Thorough experimental evaluations of our method are in Section 8.

4 Similarity Theory

Similarity theory was originally introduced to applied physics by Wyman et al. [1989a; 1989b]. It studies the equivalence classes of the RTE’s parameter space by introducing a hierarchy of equivalence relations called *similarity relations*.

Note that, given the original scattering parameters, the similarity relations do not directly provide the values of all altered parameters, and computing these values (the altered phase function in particular) is non-trivial. We introduce a novel approach to solve for these parameters in Section 5.

In this section, we first present the full derivation of similarity relations (Section 4.1) following the original version proposed by Wyman et al. [1989a]. Then, Section 4.2 discusses the connection between similarity theory and approaches based on first-order approximations of the RTE (such as diffusion methods). Finally, Section 4.3 shows an example of capturing the structure of a simple RTE’s parameter space using the derived relations.

4.1 Derivation of Similarity Relations

We present the derivation (following [Wyman et al. 1989a]) of a set of relations between the original parameters $\sigma_a, \sigma_s, f(\cdot)$ and the altered ones $\sigma_a^*, \sigma_s^*, f^*(\cdot)$ such that the original RTE (1) and its altered version (2) have identical solution radiance L , based on the assumption that L has bounded directional frequency. The resulting relations are in (18).

Rearranging the terms in the original RTE (1) yields

$$(\boldsymbol{\omega} \cdot \nabla)L(\boldsymbol{\omega}) + I(\boldsymbol{\omega}) = 0 \quad (3)$$

where

$$I(\boldsymbol{\omega}) := \sigma_t L(\boldsymbol{\omega}) - \sigma_s \int_{\mathbb{S}^2} f(\boldsymbol{\omega}' \cdot \boldsymbol{\omega}) L(\boldsymbol{\omega}') d\boldsymbol{\omega}'.$$

Similarly, the altered RTE (2) can be rewritten as

$$(\boldsymbol{\omega} \cdot \nabla)L(\boldsymbol{\omega}) + I^*(\boldsymbol{\omega}) = 0. \quad (4)$$

Then having one solution L satisfying both (3) and (4) implies that for all $\boldsymbol{\omega} \in \mathbb{S}^2$,

$$I(\boldsymbol{\omega}) = I^*(\boldsymbol{\omega}). \quad (5)$$

To derive the similarity relations from (5), one can represent $I(\boldsymbol{\omega})$ and $I^*(\boldsymbol{\omega})$ in SH and equate the corresponding SH coefficients. To write $I(\boldsymbol{\omega})$ in SH, the radiance field L and the phase function $f(\cdot)$ need to be expanded, yielding

$$L(\boldsymbol{\omega}) = \sum_{n=0}^{\infty} \sum_{m=-n}^n a_{mn} Y_n^m(\boldsymbol{\omega}), \quad (6)$$

$$\begin{aligned} f(\boldsymbol{\omega}' \cdot \boldsymbol{\omega}) &= \sum_{n=0}^{\infty} \frac{2n+1}{4\pi} f_n P_n(\boldsymbol{\omega}' \cdot \boldsymbol{\omega}) \\ &= \sum_{n=0}^{\infty} \sum_{m=-n}^n f_n Y_n^m(\boldsymbol{\omega}) \bar{Y}_n^m(\boldsymbol{\omega}') \end{aligned} \quad (7)$$

where Y_n^m is the SH basis function, P_n is the Legendre polynomial of degree n ,

$$f_n = 2\pi \int_{-1}^1 f(t)P_n(t)dt \quad (8)$$

is the n -th Legendre moment of $f(\cdot)$, and the bar superscript denotes complex conjugation. The second equality in (7) follows the Addition Theorem [Arfken et al. 1985]. For heterogeneous materials, both a_{mn} and f_n have spatial dependencies.

Given (6) and (7), it has been shown (see Appendix A) that

$$I(\boldsymbol{\omega}) = \sum_{n=0}^{\infty} \sum_{m=-n}^n a_{mn} \sigma_{tr,n} Y_n^m(\boldsymbol{\omega}), \quad (9)$$

$$I^*(\boldsymbol{\omega}) = \sum_{n=0}^{\infty} \sum_{m=-n}^n a_{mn} \sigma_{tr,n}^* Y_n^m(\boldsymbol{\omega}). \quad (10)$$

where

$$\begin{aligned} \sigma_{tr,n} &:= \sigma_t - \sigma_s f_n = \sigma_a + \sigma_s(1 - f_n), \\ \sigma_{tr,n}^* &:= \sigma_t^* - \sigma_s^* f_n^* = \sigma_a^* + \sigma_s^*(1 - f_n^*) \end{aligned} \quad (11)$$

are called the *reduced attenuation coefficients* of order n . Note that a_{mn} appears in both (9) and (10) since $I(\boldsymbol{\omega})$ and $I^*(\boldsymbol{\omega})$ are assumed to share the same radiance field L .

From (9) and (10), the SH coefficients of $I(\boldsymbol{\omega})$ and $I^*(\boldsymbol{\omega})$ can now be equated, which leads to $a_{mn} \sigma_{tr,n} = a_{mn} \sigma_{tr,n}^*$ for all $n \geq 0$ and $-n \leq m \leq n$. Namely,

$$\begin{aligned} a_{mn}(\sigma_{tr,n} - \sigma_{tr,n}^*) &= \\ a_{mn}[(\sigma_a - \sigma_a^*) + (\sigma_s(1 - f_n) - \sigma_s^*(1 - f_n^*))] &= 0. \end{aligned} \quad (12)$$

Since (12) needs to hold for all n and m , consider a special case where $n = m = 0$. It holds that

$$a_{00} = \int_{\mathbb{S}^2} Y_0^0(\boldsymbol{\omega})L(\boldsymbol{\omega})d\boldsymbol{\omega} = \frac{\phi}{2\sqrt{\pi}} \quad (13)$$

where $\phi := \int_{\mathbb{S}^2} L(\boldsymbol{\omega})d\boldsymbol{\omega}$ is the *fluence*. Because $f(\boldsymbol{\omega}' \cdot \boldsymbol{\omega})$ and $f^*(\boldsymbol{\omega}' \cdot \boldsymbol{\omega})$, as functions of $\boldsymbol{\omega}'$, are probability densities over \mathbb{S}^2 , it holds that $f_0 = f_0^* = 1$. Then, (12) becomes $\frac{\phi}{2\sqrt{\pi}}(\sigma_a - \sigma_a^*) = 0$. Because ϕ is generally non-zero, this implies

$$\sigma_a = \sigma_a^*. \quad (14)$$

In general, for $n \geq 1$, given (14), (12) becomes

$$a_{mn}[\sigma_s(1 - f_n) - \sigma_s^*(1 - f_n^*)] = 0. \quad (15)$$

To ensure that (15) holds for any L (namely for arbitrary a_{mn}), we need to have

$$\sigma_s(1 - f_n) = \sigma_s^*(1 - f_n^*). \quad (16)$$

Similarity Relations. Wyman et al. [1989a] showed that the only solution adhering to (14) and (16) for all n, m is the trivial one: $\sigma_a^* = \sigma_a$, $\sigma_s^* = \sigma_s$, and $f^*(\cdot) \equiv f(\cdot)$. Thus, in general, there is no ‘‘perfect’’ similarity relation. However, when L is band-limited in SH domain, we have

$$a_{mn} = 0 \quad \text{for } n > N, -n \leq m \leq n \quad (17)$$

where N is a constant capturing the maximal angular frequency. If $N = 1$, for example, L is called *linearly anisotropic*. When (17)

holds, (16) only needs to be enforced for $1 \leq n \leq N$, yielding the *similarity relation of order N* :

$$\begin{aligned} \sigma_a &= \sigma_a^*, \\ \sigma_s(1 - f_n) &= \sigma_s^*(1 - f_n^*) \quad \text{for } 1 \leq n \leq N. \end{aligned} \quad (18)$$

In practice, (17) may not hold everywhere inside the medium. In this case, the altered RTE (2) will have a solution radiance approximating that of the original (1).

The Legendre moment constraints in lower-order similarity relations are subsets of those in higher-order ones. Therefore, (18) essentially provides a *hierarchy of equivalence relations* that partition the parameter space of a RTE with different granularity.

Spatial Dependency. Similarity relation (18) needs to be satisfied at every location \mathbf{x} (which has been dropped for notational convenience) within the material volume. Namely, for heterogeneous materials where σ_s , σ_a , and $f(\cdot)$ are spatially varying, the altered parameters σ_s^* , σ_a^* , and $f^*(\cdot)$ should also have spatial dependencies so that (18) is satisfied independently for each \mathbf{x} .

Generalized Forms. The similarity relations (18) force the absorption coefficient σ_a to remain unchanged. When $\sigma_a^* \neq \sigma_a$, the solution radiance L to the original (1) and the altered RTE (2) are normally not identical. Wyman et al. [1989b] showed that if one weakens the requirement of identical L and only asks for equal resulting fluence $\phi = \int_{\mathbb{S}^2} L(\boldsymbol{\omega})d\boldsymbol{\omega}$, generalizations of (18) can be obtained. In particular, the *generalized order-1 and order-2 similarity relations* are respectively

$$\sigma_a \sigma_{tr,1} = \sigma_a^* \sigma_{tr,1}^*, \quad \frac{\sigma_a \sigma_{tr,1} \sigma_{tr,2}}{\sigma_s(1 - f_2)} = \frac{\sigma_a^* \sigma_{tr,1}^* \sigma_{tr,2}^*}{\sigma_s^*(1 - f_2^*)}. \quad (19)$$

When $\sigma_a^* = \sigma_a$, (19) reduces to (18) with $N = 1$ and $N = 2$.

Since most graphics applications care about radiance L (which determines an object’s appearance) instead of fluence ϕ , we focus on the standard similarity relations (18) in the rest of this paper. Please refer to Section 3 of the supplementary document for more information on the generalized versions.

4.2 Discussion: Relation to First-Order Methods

In many prior works [Chatigny et al. 1999; Frisvad et al. 2007], it was common to set

$$\sigma_a^* = \sigma_a, \quad \sigma_s^* = \sigma_s(1 - f_1), \quad f^*(\boldsymbol{\omega}' \cdot \boldsymbol{\omega}) = \frac{1}{4\pi}. \quad (20)$$

It is easy to verify that the altered phase function $f^*(\cdot)$ has $f_1^* = 0$. Thus, $\sigma_s^*(1 - f_1^*) = \sigma_s^* = \sigma_s(1 - f_1)$, and (20) satisfies the order-1 similarity relation.

The altered parameters in (20) are also used by diffusion methods [Jensen et al. 2001; Arbree et al. 2011; D’Eon and Irving 2011] where σ_s^* is called the *reduced scattering coefficient*. In fact, the order-1 similarity relation and the diffusion approximation share the same assumption that L is linearly anisotropic (namely (17) holds with $N = 1$), so they offer similar levels of accuracy. We present an alternative derivation of the diffusion equation (DE) using this assumption in Section 2 of the supplementary document.

On the other hand, methods based on first-order approximations, including (20), have limited accuracy. These methods assume that the radiance field is linearly anisotropic and can perform poorly in optically thin or close-to-boundary regions where the assumption is often violated. Furthermore, phase functions with similar first moments can lead to dramatically different appearances (examples are shown in Figures 1 and 9 as well as by




Figure 3: Equivalence classes of a 2D parameter space. White dots indicate the reference parameter points: (0.9, 50) for the left plot and (0.5, 25) for the right. Dashed lines contain all points belonging to the same equivalence class (defined by the order-1 similarity relation) as the references. Low-error regions on the error surfaces (in false color) match the predicted equivalence classes, confirming the theory.

Gkioulekas et al. [2013a]). Unfortunately, first-order methods care only about the first moment of a phase function and cannot capture all these varying appearances, as demonstrated in Sections 8.2 and 8.3.

To our knowledge, there is no prior work in computer graphics which considers higher-order similarity relations. In Section 5, we show how to solve for the altered parameters satisfying the similarity relation of any given order.

4.3 Example: Equivalence Classes

We now illustrate the structure of the parameter space of a simple RTE and discuss how such structure can be exploited to benefit forward and inverse rendering of volumetric media.

Consider a RTE with a fixed absorption coefficient and a Henyey-Greenstein (HG) phase function [1941]. Then its parameter space is 2D: one for the scattering coefficient σ_s and the other for the HG parameter g . Because the first Legendre moment of an HG function with parameter g is simply g itself, two parameter points (g, σ_s) and (g', σ'_s) belong to the same equivalence class defined by the order-1 similarity relation¹ if $\sigma'_s(1 - g) = \sigma_s(1 - g')$. Figure 3 shows the equivalence classes of two reference parameters (0.9, 50) and (0.5, 25) plotted as dashed lines.

To validate the equivalence classes defined by similarity theory, we created multiple renderings of a homogeneous, unit-sized cube (under side lighting) using different scattering parameters. Denote the image rendered with parameters g and σ_s as $\mathbf{I}(g, \sigma_s)$. The insets in Figure 3 show images rendered with the reference parameters. For each point (g, σ_s) in the space, define an error function

$$d(g, \sigma_s) := \|\mathbf{I}(g, \sigma_s) - \mathbf{I}_0\|_2 \quad (21)$$

where \mathbf{I}_0 is the image rendered with the reference parameters. The error surfaces defined by d are visualized as plot backgrounds. We can see that the shapes of low-error regions match the curves predicted by similarity theory well. Note that the error values over the dashed lines are not exactly zero, since the equivalence classes are defined based on the assumption that the radiance is linearly anisotropic, which is normally not the case near boundaries.

Inverse Rendering. The presence of these low-error regions causes inverse rendering to be very challenging. One reason is that the shapes of these regions are not convex, so that many optimization algorithms are not guaranteed to find a global optimum. Furthermore, within the low-error regions, the gradient is

¹ Similarity relations beyond order-1 are not useful for this parameter space, as each equivalence class would contain only a single point.

fairly small and can be easily dominated by Monte Carlo or measurement noise. Based on this understanding, Section 6 presents a simple method that reparameterizes the search space, causing gradient based methods to be much more effective.

Forward Rendering. We can exploit the structure of the parameter space to benefit forward rendering applications (Section 7). To render an object with scattering parameters coming from the upper-right region of Figure 3, for instance, we can instead use a set of parameters in the same equivalence class but located at the bottom-left of the space, as both sets of parameters lead to approximately the same appearance. By picking a smaller scattering coefficient, the material’s optical density is reduced, causing light to scatter less frequently. Consequently, fewer scattering events need to be simulated, and speedups can be obtained.

5 Solving for Altered Parameters

Deriving the similarity relations (18) is only half the story. Applications such as forward rendering require full sets of parameters: the values of σ_a^* , σ_s^* , and a complete phase function $f^*(\cdot)$. Unfortunately, only $\sigma_a^* = \sigma_a$ is given directly by (18). The other relations, $\sigma_s(1 - f_n) = \sigma_s^*(1 - f_n^*)$, are constraints.

To determine σ_s^* , we consider the ratio between σ_s^* and σ_s :

$$\alpha := \sigma_s^* / \sigma_s. \quad (22)$$

In our forward rendering pipeline (Figure 2), this ratio is selected by the user. Given α , we have $\sigma_s^* = \alpha\sigma_s$, and the only parameter that remains unknown is the altered phase function $f^*(\cdot)$. For fixed order N , the Legendre moments of $f^*(\cdot)$ need to satisfy

$$f_0^* = 1, \quad f_i^* = 1 - \frac{1 - f_i}{\alpha} \quad \text{for } 1 \leq i \leq N. \quad (23)$$

Unfortunately, computing $f^*(\cdot)$ given f_0^*, \dots, f_N^* is non-trivial. As a phase function, $f^*(\cdot)$ needs to be nonnegative. Discarding all moments higher than order- N by setting $f^*(t) = \sum_{n=0}^N \frac{2n+1}{4\pi} f_n^* P_n(t)$, however, generally does not offer nonnegativity. In addition, given the Legendre moment constraints, a nonnegative $f^*(\cdot)$ may not exist at all.

Wyman et al. [1989a] proposed a simple approach to provide $f^*(\cdot)$. This method does not allow N and α to be specified simultaneously. Instead, it takes N as the user input and constructs a phase function $f^*(\cdot)$ with $f_N^* = 0$. This method, therefore, offers insufficient flexibility: in many applications including forward rendering, we need to control both N and α to achieve good performance and accuracy.

In this section, we introduce a general technique to find $f^*(\cdot)$ for any given N and α . Section 5.1 presents existence conditions of $f^*(\cdot)$ given f_0^*, \dots, f_N^* . Section 5.2 introduces an algorithm to solve for $f^*(\cdot)$ numerically as a tabulated (piecewise-constant) function.

5.1 Existence of the Altered Phase Function

We now show the sufficient and necessary conditions for the existence of a nonnegative function $f^*(\cdot)$ with its Legendre moments f_0^*, \dots, f_N^* given.

For $f^*(\cdot)$, its n -th monomial moment is $\gamma_n^* := \int_{-1}^1 f(t) t^n dt$. Since $P_n(\cdot)$ is a polynomial of degree n , the Legendre moments f_0^*, \dots, f_N^* and monomial moments $\gamma_0^*, \dots, \gamma_N^*$ of $f^*(\cdot)$ uniquely determine each other. Given these Legendre moments, the corresponding monomial moments can be computed by solving a linear system (see Section 4 of the supplementary document for

details). It follows that determining whether $f^*(\cdot)$ exists given its Legendre moments f_0^*, \dots, f_N^* is equivalent to checking its existence given the monomial moments $\gamma_0^*, \dots, \gamma_N^*$. The latter is called the *truncated Hausdorff moment problem* and has been studied in probability theory [Curto and Fialkow 1991]. In fact, $f^*(\cdot)$ exists if and only if certain Hankel matrices formed using the monomial moments are positive semi-definite. The following theorem provides a formal description of this result (see Theorems 4.1 and 4.3 in Curto and Fialkow's work [1991] for the proof).

Theorem 1. Given $\gamma_0^*, \gamma_1^*, \dots, \gamma_N^*$ with $\gamma_0^* > 0$. For each n , let \mathbf{U}_n , \mathbf{V}_n , and \mathbf{W}_n be $n \times n$ Hankel matrices such that

$$U_n(i, j) = \gamma_{i+j-2}^*, \quad V_n(i, j) = \gamma_{i+j-1}^*, \quad W_n(i, j) = \gamma_{i+j}^* \quad (24)$$

for $1 \leq i, j \leq n$. Then, a nonnegative function $f^*(\cdot)$ with monomial moments $\gamma_0^*, \gamma_1^*, \dots, \gamma_N^*$ exists if and only if:

- (odd case) when $N = 2k + 1$,

$$\mathbf{U}_{k+1} - \mathbf{V}_{k+1} \succcurlyeq 0, \quad \mathbf{U}_{k+1} + \mathbf{V}_{k+1} \succcurlyeq 0 \quad (25)$$

where " $\succcurlyeq 0$ " denotes *positive semi-definiteness* of a matrix;

- (even case) when $N = 2k$,

$$\mathbf{U}_{k+1} \succcurlyeq 0, \quad \mathbf{U}_k - \mathbf{W}_k \succcurlyeq 0. \quad (26)$$

Based on Theorem 1, a function `IFEXISTS()` can be easily implemented which takes the desired Legendre moments f_0^*, \dots, f_N^* and returns a Boolean indicating if $f^*(\cdot)$ exists.

5.2 Computing the Altered Phase Function

Although Theorem 1 allows us to efficiently check the existence of the altered phase function $f^*(\cdot)$, it does not provide a practical way to find one (if it exists). Next, we introduce an algorithm to solve for $f^*(\cdot)$ numerically. The resulting $f^*(\cdot)$ can then be used for physically-based rendering applications (Section 7).

We represent $f^*(\cdot)$ as the linear combination of k basis functions $g_1(\cdot), \dots, g_k(\cdot)$:

$$f^*(t) = \sum_{i=1}^k c_i g_i(t). \quad (27)$$

Note that this can lose generality as the bases may not be able to represent all nonnegative functions on $[-1, 1]$. The selection of k is discussed in the end of this subsection.

Given (27) and (8), the n -th Legendre moment of $f^*(\cdot)$ equals

$$f_n^* = 2\pi \int_{-1}^1 \left(\sum_{i=1}^k c_i g_i(t) \right) P_n(t) dt = \sum_{i=1}^k c_i g_{i,n} \quad (28)$$

where

$$g_{i,n} := 2\pi \int_{-1}^1 g_i(t) P_n(t) dt$$

is the n -th Legendre moment of $g_i(\cdot)$.

Let $\mathbf{f}^* := (f_0^* f_1^* \dots f_N^*)^T$, $\mathbf{c} := (c_1 c_2 \dots c_k)^T$, and

$$\mathbf{G} := \begin{pmatrix} g_{1,0} & g_{2,0} & \dots & g_{k,0} \\ g_{1,1} & g_{2,1} & \dots & g_{k,1} \\ \vdots & \vdots & \ddots & \vdots \\ g_{1,N} & g_{2,N} & \dots & g_{k,N} \end{pmatrix},$$

then (28) with $n = 0, 1, \dots, N$ can be rewritten as

$$\mathbf{f}^* = \mathbf{G} \mathbf{c}. \quad (29)$$

To summarize, if $g_1(\cdot), g_2(\cdot), \dots, g_k(\cdot)$, which determine \mathbf{G} , are given, we need to solve for \mathbf{c} such that $f^*(\cdot)$ is nonnegative and the Legendre moment constraints (29) are satisfied.

In our implementation, we pick the boxcar basis functions:

$$g_i(t) := \begin{cases} 1 & -1 + \frac{2i-2}{k} \leq t < -1 + \frac{2i}{k} \\ 0 & \text{otherwise} \end{cases}, \quad (30)$$

for $i = 1, 2, \dots, k$. We choose this basis for its simplicity and flexibility. Under (30), $f^*(\cdot)$ becomes piecewise-constant, and its value in the i -th piece simply equals c_i . It follows that $f^*(\cdot)$ is nonnegative if and only if $\mathbf{c} \geq \mathbf{0}$ (defined component-wise). We then would like to find a nonnegative k -dimensional vector \mathbf{c} satisfying (29). This system, however, is normally under-constrained since k , the number of pieces in $f^*(\cdot)$, can be much greater than $(N+1)$, the amount of Legendre moment constraints. Therefore, we regularize the system by introducing a smoothness term $\mathbf{S} \mathbf{c}$ with $\mathbf{S} \in \mathbb{R}^{(k-2) \times k}$ being a 1D Poisson matrix

$$\mathbf{S} = \begin{pmatrix} -1 & 2 & -1 & & & & \\ & -1 & 2 & -1 & & & \\ & & & \dots & & & \\ & & & & -1 & 2 & -1 \end{pmatrix}.$$

This smoothness term captures the second derivative of $f^*(\cdot)$, and we would like $\|\mathbf{S} \mathbf{c}\|_2$ to be minimized. We choose the 2-norm since the solution \mathbf{c} is robust to the choice of k , and the resulting phase function $f^*(\cdot)$ tends to send light into a wide range of directions, which is a desirable feature.

Let $\mathbf{Q} := \mathbf{S}^T \mathbf{S}$, then $\|\mathbf{S} \mathbf{c}\|_2^2 = \mathbf{c}^T \mathbf{Q} \mathbf{c}$, and we need to find the minimizer to the following quadratic programming problem:

$$\min_{\mathbf{c}} (\mathbf{c}^T \mathbf{Q} \mathbf{c}) \quad \text{subject to} \quad \mathbf{c} \geq \mathbf{0}, \quad \mathbf{G} \mathbf{c} = \mathbf{f}^*. \quad (31)$$

Since \mathbf{Q} is positive semi-definite, the global optimum can be found in polynomial time [Kozlov et al. 1980]. We solve (31) using the Gurobi Optimization Libraries [2013].

Selecting k . Given \mathbf{f}^* , if the existence condition (25) or (26) is violated, the quadratic programming problem in (31) is guaranteed to be infeasible. However, the converse is not necessarily true: when (31) is infeasible, there could still exist some $f^*(\cdot)$ which cannot be represented using the k basis functions $g_1(\cdot), g_2(\cdot), \dots, g_k(\cdot)$. In this case, we need a larger set of bases. One possibility to determine the value of k is to start with some relatively small k_0 and double it whenever (31) is infeasible, but there should be a solution (according to Theorem 1). In practice, however, we found that simply setting $k = 360$ is sufficient to find the solutions in all our experiments.

6 Application: Inverse Rendering

In this section, we first describe an inverse rendering problem which we believe is a good example to demonstrate the practical usefulness of similarity theory. Then, we introduce a reparameterized search space in which gradient descent algorithms converge to a good solution much faster.

The Problem. Consider the problem of acquiring the material parameters of a cube made of a scattering medium. Given a photograph of the cube lit by an area light from the side (identical to the setting used in Section 4.3), we assume that the cube has a HG phase function and its absorption coefficient σ_a is given. The goal is to find the scattering coefficient σ_s and HG parameter g .




Figure 4: Search spaces for an inverse rendering problem: (a) the original space; (b) the reparameterized space. The plotted region in (a) maps to the area enclosed by the dashed lines in (b). Using the original space, the stochastic gradient descent (SGD) algorithm starting from point S is trapped at point P, which is far from the real solution T. Using the reparameterized space, the algorithm is able to find point R that is much closer to the real solution.

Our Solution. Given the highly complicated and non-linear relation between the parameters and the resulting rendered image, exact analytical solutions do not exist for this problem. Instead, we use the error function d defined in (21) with I_0 set to the input image, and solve for g, σ_s such that the error is minimized. This optimization problem can be solved using the stochastic gradient descent (SGD) algorithm. The nondeterminism is caused by the fact that the gradients need to be obtained through Monte Carlo simulations and can be noisy.

Unfortunately, the use of a relatively low-frequency (soft) lighting results in large regions where $d(g, \sigma_s)$ is close to zero (as shown in Figure 3), and the gradient values within these regions are very small and can be dominated by Monte Carlo or measurement noise. Consequently, after hitting this region, it becomes very difficult for SGD to make further progress. Figure 4-a shows such an example in which I_0 is generated using parameters at point T. If we choose point S as the initial guess and execute SGD, the solution point moves to the low-error region (indicated in purple) very quickly, but then gets “trapped” there. Figure 5-b demonstrates that the solution (point P) found by this process generalizes poorly to high-frequency (hard) lighting conditions.

To address this problem, we reparameterize the search space (Figure 4-b) by replacing the horizontal axis by the reduced scattering coefficient $\sigma'_s := (1 - g)\sigma_s$. Under this reparameterization, both axes have the same units, and the error surface becomes significantly more regular (Figure 4-b). To search for the solution in the new space, we perform two one-dimensional SGDs. First, we fix σ_s and look for σ'_s that minimizes the error. Since each σ'_s corresponds to an equivalence class given by the order-1 similarity relation, this step allows us to select a class with minimal error. Then, we search for the best parameter point within this class by keeping σ'_s fixed and performing another 1D search to find the best σ_s . Figure 4-b shows an example starting from point S (the same initial guess as in Figure 4-a) where the first 1D search finds Q and the second returns R. Figure 5-c shows that this new solution matches the ground truth better under both the original (low-frequency) and the novel (high-frequency) lighting.

7 Application: Forward Rendering

Optically dense and forward-scattering materials are very common in the real-world [Frisvad et al. 2007; Gkioulekas et al. 2013b], but they are challenging to render. Our main application is to offer speedups to Monte Carlo rendering of these materials without modifying the core rendering algorithms. The basic idea is to find altered parameters with $\sigma_s^* < \sigma_s$ (namely, with $\alpha < 1$), so that fewer scattering events need to be handled.




Figure 5: Images rendered using the real solution (a) as well as solutions found by executing SGD on the original search space (b) and the reparameterized one (c). Visualizations of per-pixel relative error (using the color mapping in Figure 1) are included in (b, c). The images in the top row are used during the optimization process, and those in the bottom with a novel lighting are for validation. The solution found using the reparameterized space shown in (c) leads to better results in both configurations.

We introduce a practical algorithm (Algorithm 1) based on the theory introduced in Sections 4 and 5. To use this algorithm, the user can simply input the original scattering parameters and an extra number α which controls the balance between performance and accuracy, and the algorithm outputs a set of altered parameters corresponding to an optically thinner and less forward-scattering material. Rendering images using these altered parameters (without modifying the renderer) costs only a fraction of the computation required to render with the original ones. If the original parameters are spatially varying, the algorithm needs to be performed at each spatial location (or for each homogeneous region).

Next, we provide a detailed description of our technique: Section 7.1 presents an overview of the method (whose pipeline has been previewed in Figure 2); Section 7.2 describes the practical aspects of the user-specified parameter α ; Section 7.3 introduces an “overfitting” problem and discusses how to avoid it. Detailed experimental evaluations of our technique are in Section 8.

7.1 Overview

Given the original scattering parameters $\sigma_a, \sigma_s, f(\cdot)$ and the user-specified parameter α , our method sets $\sigma_a^* = \sigma_a$ and $\sigma_s^* = \alpha\sigma_s$ (line 2 of Algorithm 1). Then, we determine the order N of the similarity relation to satisfy. Greater N normally provides better accuracy but may lead to unsatisfiable constraints. Thus, we use the `IFEXISTS()` routine implementing Theorem 1 (line 6) to find the highest order N of which the similarity relation is satisfiable (lines 4 to 10). In practice, we bound N with $N_0 = 5$ (line 5) because our experiments indicate that going beyond this order does not provide observable improvement for resulting quality.

Occasionally, higher-order relations (greater N) yield worse accuracy, and we call this problem “overfitting”. This is another reason that we pick $N_0 = 5$ as relations beyond this order tend to overfit. Section 7.3 presents a simple method to reject the overfitting results (line 14 of Algorithm 1).

Algorithm 1 contains the pseudocode of the entire pipeline. Despite the sophistication of the underlying theory, the algorithm itself is very easy to implement. Our MATLAB implementation is available as supplementary material.

Algorithm 1 Computing the altered scattering parameters.

```

1: function COMPUTEPARAMETERS( $\sigma_a, \sigma_s, f(\cdot), \alpha$ )
2:    $\sigma_a^* \leftarrow \sigma_a, \sigma_s^* \leftarrow \alpha \sigma_s$ 
3:   Compute  $f_1^*, \dots, f_{N_0}^*$  using (23)
4:    $N \leftarrow 1$  ▷ Compute the order  $N$ 
5:   while  $N < N_0$  do
6:     if not IFEXISTS( $1, f_1^*, \dots, f_{N+1}^*$ ) then ▷ Theorem 1
7:       break
8:     end if
9:      $N \leftarrow N + 1$ 
10:  end while
11:  repeat ▷ Estimate  $f^*(\cdot)$  numerically
12:    Solve the order- $N$  version of (31) to obtain  $f^*(\cdot)$ 
13:     $N \leftarrow N - 1$ 
14:  until  $f^*(\cdot)$  is NOT overfitting ▷ Section 7.3
15:  return  $\sigma_a^*, \sigma_s^*, f^*(\cdot)$ 
16: end function

```

7.2 User-Specified Parameter

We now describe the way in which the user-specified parameter α in Algorithm 1 affects performance and accuracy, and discuss how to pick the value of this parameter properly.

Balancing Performance and Accuracy. In practice, the parameter α controls the tradeoff between performance and accuracy: small α offers good performance but potentially poor accuracy; large α provides good accuracy but at the cost of slower performance.

When $\sigma_a^* \ll \sigma_s^*$ (namely the single-scattering albedo is high), it holds that $\alpha = \sigma_s^*/\sigma_s \approx (\sigma_s^* + \sigma_a^*)/(\sigma_s + \sigma_a) = \sigma_t^*/\sigma_t$. Therefore, along a unit distance within the altered material, the expected number of scattering events is roughly a factor α of that within the original material. For Monte Carlo methods where each scattering event is explicitly simulated, this usually means that α is linearly related to the rendering time.² Experimental results demonstrating this effect are in Section 8.1.

Picking α . Although α can theoretically take any value in $(0, 1)$, not every value in this range leads to high-quality results. Assuming that the original material is forward-scattering ($f_1 > 0$), we found that the resulting accuracy decreases rapidly when $f_1^* < 0$. So we require $f_1^* \geq 0$ which implies

$$\alpha \geq 1 - f_1. \quad (32)$$

For objects with optically thin regions, relatively large α values are required to produce high-quality results. We observed that setting $\alpha = \max(0.3, 1 - f_1)$ worked quite well for all our experiments, even under conditions that are highly unforgiving to errors, such as back lighting. To fine-tune this parameter for greater speedups, a small number of test renderings can be performed as in Figure 7 but using a low resolution and a small number of random samples, so that it introduces little overhead and the resulting α can be reused for high-quality renderings or generating animated sequences. In this case, we suggest starting with a smaller α value such as $\max(0.1, 1 - f_1)$ and increasing α iteratively until the test renderings converge visually.

²With the presence of Russian roulette, the expected number of scattering events can be bounded, which may reduce the rendering time for large α and makes the relation between α and the rendering time more complicated. Because performing Russian roulette can introduce a great amount of noise when the single-scattering albedo σ_s/σ_t is close to 1, how to do it properly is non-trivial and beyond the scope of this paper. Thus, we did not use this technique when creating our renderings.

7.3 Overfitting

Since overfitting does not happen very frequently and normally only causes subtle visual differences, we use a simple method to determine whether a solution $f^*(\cdot)$ is likely to overfit (based on its “support”) and reject the overfitting solutions. Examples of overfitting are in Section 8.2 (Figure 8) and Section 6 of the supplementary document.

Given a phase function $f(\cdot)$, we define its (*normalized*) support $\text{nz}(f)$ to be a fraction between 0 and 1 that equals the portion of the domain where $f(\cdot)$ is greater than zero. Namely,

$$\text{nz}(f) := \frac{|\{t : f(t) > 0\}|}{2}.$$

For a tabulated phase function $f^*(t)$ with k bins, $\text{nz}(f^*) = k'/k$ where k' is the number of bins in which $f^*(t)$ is positive.

Intuitively, to approximate a translucent material with an altered one where scattering occurs less frequently, one needs to allow light to scatter into a wider range of directions. Our experiments indicate that given the moment constraints in (29), solution phase functions with larger supports provide better results. Therefore, we regularize $f^*(\cdot)$ by minimizing the 2-norm of the smoothness term, which favors support instead of sparsity, in Section 5.2.

On the other hand, for fixed α , the altered phase function needs to be more concentrated (namely, to have smaller support) for satisfying higher-order similarity relations. Occasionally, these relations become barely satisfiable, resulting in phase functions with very limited supports that tend to overfit.

To remedy this problem, we threshold the support of altered phase functions. Given the original phase function $f(\cdot)$, if the altered phase function $f^*(\cdot)$ has $\text{nz}(f^*) < \beta \text{nz}(f)$ for some constant $\beta \in (0, 1]$, we reject it (line 14 of Algorithm 1). In our experiments, we used $\beta = 0.65$.

8 Experimental Results

In this section, we first show how the choice of α in Algorithm 1 balances performance and accuracy (Section 8.1). Next, Section 8.2 demonstrates that considering similarity relations beyond order-1 can provide much better accuracy for complex phase functions. Then, we show that higher-order analysis is required to capture perceptually significant cues proposed by Gkioulekas et al. [2013a] (Section 8.3). Finally, Section 8.4 exhibits rendered results for a variety of translucent media. All our renderings are created using the Mitsuba physically-based renderer [2010].

8.1 Performance versus Accuracy

We now show how the choice of α affects the performance of Monte Carlo path tracing and the resulting accuracy through experiments.

We created two scenes each of which contains a translucent object lit under high-frequency environment lighting [Debevec 1998]. Both objects have HG phase functions with $g = 0.95$, and the sizes of these objects are several hundreds times the mean free path. Then, we generated 7 sets of altered parameters using Algorithm 1 with α ranging from 0.05 to 0.75.

Figure 6 shows the rendering quality (evaluated using the HDR-VDP-2 perceptual metric [Mantiuk et al. 2011] where a higher score means better quality) and the execution time (using standard volume path tracing) as functions of α . Note that, in practice, the speedup does not exactly equal $1/\alpha$ as the rendering time is affected by many factors varying among different α values, such




Figure 6: The rendering quality scores (evaluated using the HDR-VDP-2 metric) and the execution times when changing the value of α . Data points on the purple curves marked with ‘b’, ‘c’, and ‘d’ respectively correspond to renderings in Figure 7-bcd.




Figure 7: Renderings of a heterogeneous dragon: (a) ground truth; (b, c, d) renderings using the altered parameters generated using Algorithm 1 with different α values. As in Figure 1, the relative error visualizations are included.

as cache performance (which is higher for greater α since the rendering algorithm tends to access data with better locality).

Figure 7 shows some of the rendered images corresponding to the purple curves in Figure 6. More renderings are presented in Section 5 of the supplementary document. We can see that when $\alpha = 0.05$, while a 8.0X speedup can be achieved, the resulting accuracy is unsatisfactory (Figure 7-b). On the other hand, with $\alpha = 0.1$ or 0.2 (which respectively offer speedups of 5.5X and 3.4X), significantly better accuracy can be obtained (Figure 7-cd).

8.2 Higher-Order Similarity Relations

Although first-order approximations work adequately for simple phase functions (such as single-lobe HG), they do not have sufficient representative power to capture higher moments of $f(\cdot)$.

We took a phase function proposed by Gkioulekas et al. [2013a] which is a linear combination of two distributions:

$$f(\cos \theta) = 0.9 \text{HG}(0.95, \cos \theta) + 0.1 \text{vMF}(-75, \cos \theta) \quad (33)$$

where $\text{HG}(g, \cdot)$ and $\text{vMF}(\kappa, \cdot)$ denote the HG function with parameter g and the von Mises-Fisher distribution with parameter κ , respectively. Then, we solved the quadratic programming problem in (31) to construct three altered versions of (33) adhering to the order-1, order-4, and order-5 similarity relations, respectively. The first Legendre moments of these altered phase functions are all zero. Figure 8-a plots (33) and its altered versions.

Figure 8-bcde contains a homogeneous dragon rendered using these phase functions. We can see that the accuracy offered by the order-1 version (Figure 8-b) is not ideal as it is not able to





Figure 8: A complicated phase function and its three altered versions respectively satisfying the order-1, order-4, and order-5 similarity relations are plotted in (a). Renderings of a homogeneous dragon (using the plotted phase functions) under side lighting (left) and front lighting (right) are in (b, c, d, e). The order-1 version yields poor accuracy; the order-5 version works adequately but not as well as the order-4 one under both lighting conditions.

capture higher-order features of the original phase function. The renderings produced using the order-4 version (Figure 8-d), on the other hand, match the ground truth very well. The order-5 version (the orange curve in Figure 8-a) has a fairly low support and does not generalize to different lighting conditions as well as the order-4 one (Figure 8-e), although the visual difference is subtle. Thus, this solution is overfitting and will be rejected by Algorithm 1 (which instead returns the order-4 version). Please refer to Section 6 of the supplementary document for more examples on overfitting.

8.3 Spanning the 2D Perception Space

Next, we evaluate our method on a family of phase functions [Gkioulekas et al. 2013a] spanning a 2D perception space. The two axes of this space capture the optical density (vertical) and the level of “glass-like” appearance (horizontal).

We picked 40 representatives from this family and rendered an image for each of them (with the absorption and scattering coefficients fixed). Figure 9-c shows two of these renderings. Then, we

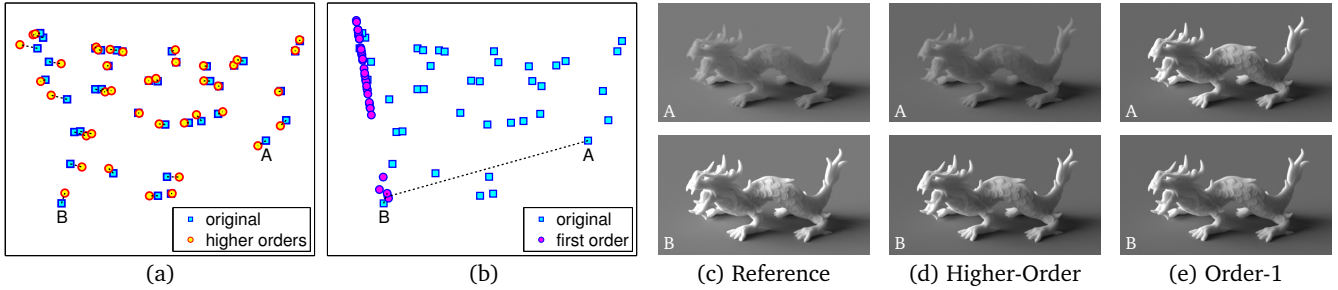


Figure 9: 2D embeddings: (a) altered parameters satisfying up to order-5 similarity relations can well maintain the structure of the original embedding; (b) satisfying only the order-1 relation causes the projections to collapse to a 1D line. The dashed lines in (a, b) connect the projections of images rendered with the original and the altered parameters. The remaining columns show renderings of two phase functions (marked with A and B) which have similar first moments: (c) reference renderings, (d, e) images rendered using altered parameters adhering to higher-order relations and the order-1 relation, respectively. As demonstrated in (e), first-order approximations do not have sufficient representative power to distinguish these phase functions (such as A and B), causing them to be mapped to similar locations in (b).

applied classical multidimensional scaling (MDS) to create a 2D embedding of the renderings where a linear transform T assigns each image a 2D coordinate. The blue squares in Figure 9-ab illustrate this 2D embedding.

Next, for each phase function, we computed an altered set of parameters satisfying similarity relations up to order-5 (using Algorithm 1) and rendered an image accordingly. Two of these rendered images are in Figure 9-d. Then, we projected those images into the previously created 2D space (using the same operator T). Figure 9-a shows that the projections can well maintain the structure of the original embedding. The small offsets between corresponding points in the two embeddings cause little visible difference, and some of them are caused by the Monte Carlo noise.

On the other hand, if we use a configuration which satisfies only the order-1 relation, the resulting renderings are missing important visual cues (as illustrated in Figure 9-e), causing their projections to collapse to a 1D line (Figure 9-b).

Figure 9 demonstrates that higher-order analysis is crucial to accurately capture perceptually significant visual cues. Section 7 of the supplementary document contains all the images used to create the three embeddings in Figure 9.

8.4 Rendered Results

Next, we demonstrate that the altered scattering parameters generated by Algorithm 1 produce appearances that accurately match the ground truth under a variety of scene configurations.

Figures 1 and 10 exhibit rendered images created using volume path tracing. Each reference rendering contains an object with a spatially invariant phase function. Consequently, only one altered version needs to be computed for each result, which takes less than a second using our MATLAB implementation of Algorithm 1.

Figure 1 contains a Corinthian capital made of a homogeneous material with a complicated phase function (the one marked as 'A' in Figure 9). It also has a refractive interface modeled using the microfacet model [Walter et al. 2007]. Altered parameters generated by our method offer a 3.7X speedup, and the resulting images match the ground truth very well.

In the first row of Figure 10, we show a sculpture made of a highly heterogeneous material where the white regions are about 5 times as dense as the green ones, which is easily visible under back lighting as shown in (a). The sculpture has the complex

phase function described in (33) and a rough dielectric interface. Our method offers a 2.7X speedup, and the results match the ground truth very well in both the thick and the thin regions.

The second row of Figure 10 contains renderings of a highly scattering smoke volume which has a HG phase function with $g = 0.95$. Parameters provided by our method, which satisfy the order-5 similarity relation, lead to a 3.5X speedup while maintaining good accuracy even at highly thin regions.

In the third row, we show rendered images of a homogeneous bust made of a material with a complicated phase function and a rough dielectric interface. Parameters provided by our method speeds up the rendering process by 3.4X, and the resulting images match the reference quite well.

9 Conclusion and Future Work

In this paper, we present a complete exposition of similarity theory, providing fundamental insights into the structure of the RTE's parameter space. Furthermore, we develop a novel approach to solve for the altered parameters satisfying the similarity relation of any given order. Since the altered phase function is not fully specified by the relations, we present the sufficient and necessary conditions for its existence and introduce a numerical algorithm to find one (when it exists) that can produce high-quality renderings.

We use two applications, forward and inverse rendering of translucent media, to demonstrate the practical utility of our theory. For inverse rendering, we present a reparameterized search space to overcome the challenges caused by the presence of equivalence classes. For forward rendering, our main application, we develop an approach (Algorithm 1) to offer speedups to Monte Carlo rendering of optically dense and forward-scattering media without having to modify the core rendering algorithms. Despite the sophistication of the theory, our method is very easy to implement and introduces negligible overhead to the full rendering pipeline.

There are many areas of future work. For inverse rendering, we would like to study how similarity theory can improve solving material appearance acquisition problems with high-dimensional search spaces on real data. For forward rendering, we plan to explore heuristics for choosing the value of α adaptively, which is a generalization of the hybrid framework, so that greater speedups can be obtained. In addition, the phase functions used in our results have no spatial variation. For spatially varying phase functions, computing many altered versions can be costly, and inter-

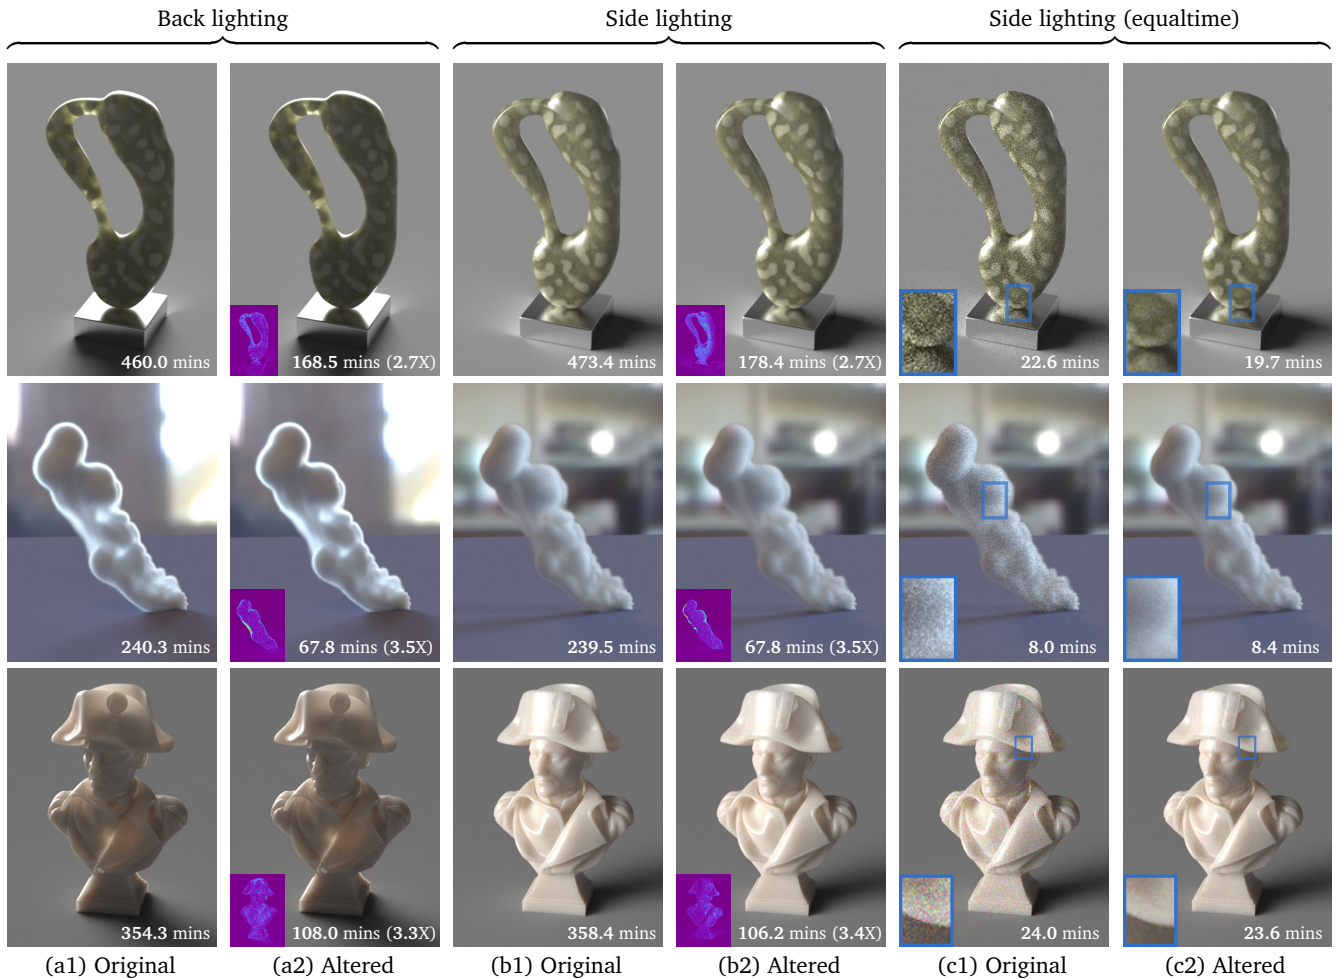


Figure 10: Path-traced renderings with various scene configurations. Columns (a, b) contain images rendered using the original and the altered parameters with the same number of sample paths per pixel. Column (c) uses the same scene configuration as (b) and shows images rendered in similar time with both parameters (see the insets to assess noise). The relative error maps (using the color scheme in Figure 1) are included in (a2, b2).

polating them without violating the similarity relation constraints is a non-trivial problem. Thus, we intend to develop approaches to efficiently compute and properly interpolate spatially varying altered phase functions. Finally, for the theory, we would like to investigate how the generalized similarity relations, such as (19), could benefit computer graphics applications.

Acknowledgments

We are grateful to Ioannis Gkioulekas, Steve Marschner, and Bruce Walter for their insightful suggestions. We also thank the anonymous reviewers for their helpful comments. Funding for this work was provided by NSF IIS grants 1011832, 1011919, 1161645, and the Intel Science and Technology Center for Visual Computing.

References

- ARBREE, A., WALTER, B., AND BALA, K. 2011. Heterogeneous subsurface scattering using the finite element method. *IEEE Trans. on Visualization and Computer Graphics* 17, 7, 956–969.
- ARFKEN, G. B., WEBER, H.-J., AND RUBY, L. 1985. *Mathematical methods for physicists*. Academic press New York.
- CHANDRASEKHAR, S. 1960. *Radiative transfer*. Courier Dover Publications.
- CHATIGNY, S., MORIN, M., ASSELIN, D., PAINCHAUD, Y., AND BEAUDRY, P. 1999. Hybrid Monte Carlo for photon transport through optically thick scattering media. *Applied optics* 38, 28, 6075–6086.
- CURTO, R. E., AND FIALKOW, L. A. 1991. Recursiveness, positivity, and truncated moment problems. *Houston J. Math* 17, 4, 603–635.
- DACHSBACHER, C., KÁŽIVÁNEK, J., HAÁAN, M., ARBREE, A., WALTER, B., AND NOVÁK, J. 2014. Scalable realistic rendering with many-light methods. *Computer Graphics Forum* 33, 1, 88–104.
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of SIGGRAPH 1998*, 189–198.

- D'EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph.* 30, 4, 56:1–56:14.
- DOBASHI, Y., IWASAKI, W., ONO, A., YAMAMOTO, T., YUE, Y., AND NISHITA, T. 2012. An inverse problem approach for automatically adjusting the parameters for rendering clouds using photographs. *ACM Trans. Graph.* 31, 6, 145:1–145:10.
- DONNER, C., AND JENSEN, H. W. 2007. Rendering translucent materials using photon diffusion. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques, EGSR'07*, 243–251.
- FRISVAD, J. R., CHRISTENSEN, N. J., AND JENSEN, H. W. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM Trans. Graph.* 26, 3, 60:1–60:10.
- GKIOULEKAS, I., XIAO, B., ZHAO, S., ADELSON, E. H., ZICKLER, T., AND BALA, K. 2013. Understanding the role of phase function in translucent appearance. *ACM Trans. Graph.* 32, 5, 147:1–147:19.
- GKIOULEKAS, I., ZHAO, S., BALA, K., ZICKLER, T., AND LEVIN, A. 2013. Inverse volume rendering with material dictionaries. *ACM Trans. Graph.* 32, 6, 162:1–162:13.
- GUROBI, 2013. Gurobi optimization libraries. www.gurobi.com.
- HABEL, R., CHRISTENSEN, P. H., AND JAROSZ, W. 2013. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. In *Computer Graphics Forum*, vol. 32, 27–37.
- HACHISUKA, T., JAROSZ, W., BOUCHARD, G., CHRISTENSEN, P., FRISVAD, J. R., JAKOB, W., JENSEN, H. W., KASCHALK, M., KNAUS, C., SELLE, A., AND SPENCER, B. 2012. State of the art in photon density estimation. In *ACM SIGGRAPH 2012 Courses, SIGGRAPH '12*, 6:1–6:469.
- HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph.* 29, 4, 61:1–61:10.
- HENYEV, L. G., AND GREENSTEIN, J. L. 1941. Diffuse radiation in the galaxy. *The Astrophysical Journal* 93, 70–83.
- ISHIMARU, A. 1978. *Wave propagation and scattering in random media*, vol. 2. Academic press New York.
- JAKOB, W., ARBREE, A., MOON, J. T., BALA, K., AND MARSCHNER, S. 2010. A radiative transfer framework for rendering materials with anisotropic structure. *ACM Trans. Graph.* 29, 4, 53:1–53:13.
- JAKOB, W., 2010. Mitsuba renderer. www.mitsuba-renderer.org.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001*, 511–518.
- KAJIYA, J. T., AND VON HERZEN, B. P. 1984. Ray tracing volume densities. *SIGGRAPH Comput. Graph.* 18, 3, 165–174.
- KOZLOV, M. K., TARASOV, S. P., AND KHACHIYAN, L. G. 1980. The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics* 20, 5, 223–228.
- LAFORTUNE, E. P., AND WILLEMS, Y. D. 1996. Rendering participating media with bidirectional path tracing. In *Rendering Techniques* 96. Springer, 91–100.
- LI, H., PELLACINI, F., AND TORRANCE, K. E. 2005. A hybrid Monte Carlo method for accurate and efficient subsurface scattering. In *Proceedings of EGSR 2005*, 283–290.
- MANTIUK, R., KIM, K. J., REMPEL, A. G., AND HEIDRICH, W. 2011. HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans. Graph.* 30, 4, 40:1–40:14.
- PAPAS, M., REGG, C., JAROSZ, W., BICKEL, B., JACKSON, P., MATUSIK, W., MARSCHNER, S., AND GROSS, M. 2013. Fabricating translucent materials using continuous pigment mixtures. *ACM Trans. Graph.* 32, 4, 146:1–146:12.
- PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Proceedings of EGWR 2000*, 11–22.
- STAM, J. 1995. Multiple scattering as a diffusion process. In *Rendering Techniques* 95. 41–50.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of EGSR 2007*, 195–206.
- WANG, J., ZHAO, S., TONG, X., LIN, S., LIN, Z., DONG, Y., GUO, B., AND SHUM, H.-Y. 2008. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Trans. Graph.* 27, 1, 9:1–9:18.
- WYMAN, D. R., PATTERSON, M. S., AND WILSON, B. C. 1989. Similarity relations for anisotropic scattering in Monte Carlo simulations of deeply penetrating neutral particles. *Journal of Computational Physics* 81, 1, 137–150.
- WYMAN, D. R., PATTERSON, M. S., AND WILSON, B. C. 1989. Similarity relations for the interaction parameters in radiation transport. *Applied optics* 28, 24, 5243–5249.

A Expansion of $I(\omega)$ in Spherical Harmonics

In this section, we prove (9) based on the orthonormality of SH.

$$\begin{aligned}
 I(\omega) &= \sum_{n,m} \sigma_t a_{mn} Y_n^m(\omega) - \\
 &\quad \int_{\mathbb{S}^2} \left[\sum_{n,m} a_{mn} Y_n^m(\omega') \right] \left[\sum_{i,j} \sigma_s f_i Y_i^j(\omega) \bar{Y}_i^j(\omega') \right] d\omega' \\
 &= \sum_{n,m} \sigma_t a_{mn} Y_n^m(\omega) - \\
 &\quad \sum_{n,m} \sum_{i,j} \left(a_{mn} Y_i^j(\omega) \sigma_s f_i \underbrace{\int_{\mathbb{S}^2} Y_n^m(\omega') \bar{Y}_i^j(\omega') d\omega'}_{= \delta_{ni} \delta_{mj}} \right) \\
 &= \sum_{n,m} \sigma_t a_{mn} Y_n^m(\omega) - \sum_{n,m} a_{mn} Y_n^m(\omega) \sigma_s f_n \\
 &= \sum_{n,m} a_{mn} \sigma_{tr,n} Y_n^m(\omega).
 \end{aligned}$$

where $\sigma_{tr,n}$ follows the definition in (11).