

A radiative transfer framework for non-exponential media

BENEDIKT BITTERLI, Dartmouth College, USA
SRINATH RAVICHANDRAN, Dartmouth College, USA
THOMAS MÜLLER, Disney Research, ETH Zürich, Switzerland
MAGNUS WRENNINGE, Pixar Animation Studios, USA
JAN NOVÁK, Disney Research, Switzerland
STEVE MARSCHNER, Cornell University, USA
WOJCIECH JAROSZ, Dartmouth College, USA




Fig. 1. We show a cloud rendered with a traditional exponential transmittance (left) and with a non-exponential, long-tailed transmittance curve (right). The non-exponential transmittance leads to both deeper light penetration as well as a softer appearance near the surface, allowing for a richer appearance.

We develop a new theory of volumetric light transport for media with non-exponential free-flight distributions. Recent insights from atmospheric sciences and neutron transport demonstrate that such distributions arise in the presence of correlated scatterers, which are naturally produced by processes such as cloud condensation and fractal-pattern formation. Our theory formulates a non-exponential path integral as the result of averaging stochastic classical media, and we introduce practical models to solve the resulting averaging problem efficiently. Our theory results in a generalized path integral which allows us to handle non-exponential media using the full range of Monte Carlo rendering algorithms while enriching the range of achievable appearance. We propose parametric models for controlling the statistical correlations by leveraging work on stochastic processes, and we develop a method to combine such unresolved correlations (and the resulting

non-exponential free-flight behavior) with explicitly modeled macroscopic heterogeneity. This provides a powerful authoring approach where artists can freely design the shape of the attenuation profile separately from the macroscopic heterogeneous density, while our theory provides a physically consistent interpretation in terms of a path space integral. We address important considerations for graphics including reciprocity and bidirectional rendering algorithms, all in the presence of surfaces and correlated media.

CCS Concepts: • **Computing methodologies** → **Ray tracing**;

Additional Key Words and Phrases: non-exponential transport, volume rendering, radiative transfer, participating media, global illumination

ACM Reference Format:

Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for non-exponential media. *ACM Trans. Graph.* 37, 6, Article 225 (November 2018), 17 pages. <https://doi.org/10.1145/3272127.3275103>

1 INTRODUCTION

Standard techniques in graphics for rendering scenes containing participating media rely on the classical radiative transfer equation (RTE) [Chandrasekhar 1960]. A central assumption of the classical RTE is that the medium is composed of tiny, *statistically independent* scatterers. This independence leads to a “memoryless” Poisson process and the familiar exponential falloff of light (see Figure 2a). While this model has proved useful in a wide range of applications, it does not accurately describe media where there are any kind of correlations between individual scatterers.

Authors’ addresses: Benedikt Bitterli, Dartmouth College, Department of Computer Science, 9 Maynard St. Hanover, NH, 03755, USA; Srinath Ravichandran, Dartmouth College, Department of Computer Science, 9 Maynard St. Hanover, NH, 03755, USA; Thomas Müller, Disney Research, ETH Zürich, Stampfenbachstrasse 48, Zürich, 8006, Switzerland; Magnus Wrenninge, Pixar Animation Studios, USA; Jan Novák, Disney Research, Stampfenbachstrasse 48, Zürich, 8006, Switzerland; Steve Marschner, Cornell University, Computer Science Department, 313 Gates Hall, Ithaca, NY, 14853, USA; Wojciech Jarosz, Dartmouth College, Department of Computer Science, 9 Maynard St. Hanover, NH, 03755, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2018/11-ART225 \$15.00
<https://doi.org/10.1145/3272127.3275103>

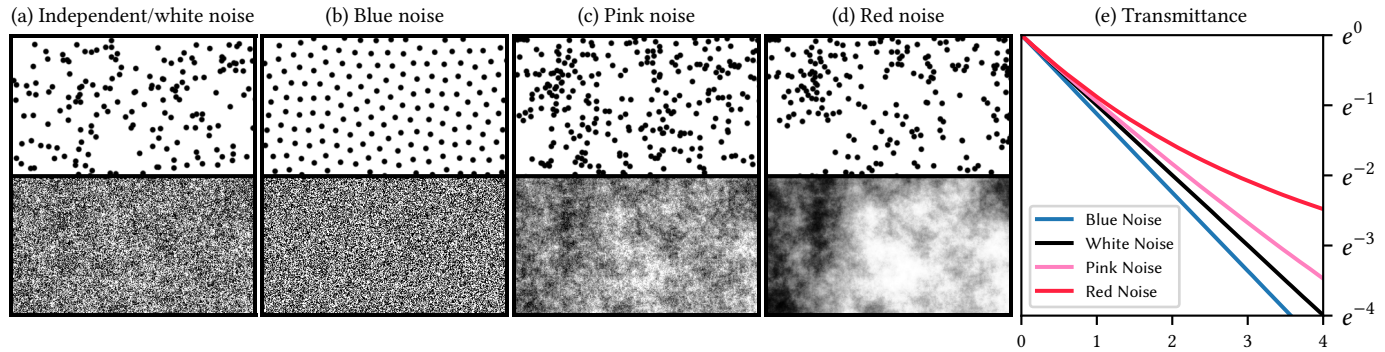


Fig. 2. We show media with discrete scatterers (top) and continuous densities (bottom) of different distributions (a)–(d), and the average transmittance measured in these media as a function of mean free paths (e). Independently placed scatterers ((a), white noise) lead to the classical exponential transmittance. Negatively correlated scatterers ((b), blue noise) lead to faster-than-exponential extinction. Positively correlated scatterers ((c)–(d), pink/red noise) lead to slower-than-exponential extinction.

1.1 The case for correlated/non-exponential transport

A wealth of media with correlated scatterers can be found in the class of colloids, i.e. mixtures of insoluble microscopic particles dispersed in another substance. This covers a wide variety of materials such as foods (e.g. milk, honey, soups, juices), agents (gels, foams, oils, ink, paint, toothpaste), gases (mist, dust, fog, steam, smoke, clouds), outer space (stellar clouds, planetary rings, comets) and even some glasses and plastics. Interactions between these particles is dictated by a variety of forces, such as electrostatic forces, van der Waals forces or excluded volume repulsion. The inter-particle correlations caused by these forces has been widely reported in literature [Beresford-Smith et al. 1985; Goicochea 2013; Grier and Behrens 2001; Jamali 2015; Rahmani 2013; Smith et al. 2004; Xu et al. 1998] and is the subject of research in a diverse set of fields. If significant correlations exist on the scale of the mean-free path of the medium [Davis and Marshak 2004], non-exponential free-flight distributions arise, where positive [Borovoi 2002; Kostinski 2001] and negative [Shaw et al. 2002] correlations lead to slower-than-exponential and faster-than-exponential free-flights, respectively. For the specific case of clouds, growing evidence from atmospheric sciences shows that correlations in the positions of water droplets exist at scales ranging at least from centimeters to kilometers [Davis et al. 1999; Kostinski and Jameson 2000].

In Figure 2 we perform a simple Monte Carlo experiment where we trace photons with random origins and directions through 2D media consisting of explicitly modeled absorbing particles. We gather statistics about the averaged transmittance (e) by tabulating along the horizontal axis the fraction of photons that survive for a given distance. In each case the average number of particles is identical, and it is the statistical correlations that give rise to different light attenuation behavior. Intuitively, particles in positively correlated media (c,d) “clump” together and leave larger gaps than expected, so photons that traverse these gaps skew the free-flight distribution towards the tail, allowing light to penetrate further on average. The opposite happens in negatively correlated media. Ignoring these correlations and instead assuming statistical independence leads to an inaccurate estimation of light transport with notably different visual appearance; c.f. Figure 1.

We wish to account for such violations of independence in a physically plausible manner, and enable practical rendering algorithms that enrich the level of control over the appearance of participating media. We first formulate the general problem as an ensemble average in stochastic media (Section 2), and solve it approximately with a new non-exponential path integral that supports surfaces and heterogeneity (Section 3), which we then make reciprocal (Section 3.1).

This path integral is general, but impractical to solve in the presence of heterogeneity, and we introduce a new decomposition for heterogeneous non-exponential media (Section 3.2) that superimposes non-exponential free-flights onto a classical heterogeneous medium. This results in a more flexible and directable system where artists can author classical media properties but additionally choose between different free-flight distributions, all while providing a consistent interpretation of the light transport that is independent of the chosen Monte Carlo rendering algorithm. Our new path integral shows how to easily change any Monte Carlo path-sampling algorithm to accommodate arbitrary free-flight distributions. To leverage this new flexibility, we use the mathematical formalisms of stochastic Gaussian processes and fractal noise [Barnsley et al. 1988] to obtain artist-directable parametric models for non-exponential free-flight distributions (Section 4).

1.2 Related work

Atmospheric sciences & neutron transport. Motivated by empirical observations [Davis et al. 1999; Kostinski and Jameson 2000] of multi-scale correlations of liquid water content in clouds, the atmospheric science community has developed many statistical models [Borovoi 2002; Davis and Marshak 2004; Davis and Mineev-Weinstein 2011; Davis and Xu 2014; Kostinski 2001; Shaw et al. 2002] to explain how such correlations lead to non-exponential aggregate transport behavior. These methods typically reduce the cloud to a fully homogenized slab while introducing statistically equivalent non-exponential free-flights. We adopt the fractal variability model proposed by Davis and colleagues [2011; 2014], but additionally provide a way to superimpose the non-exponentiality from such unresolved fluctuations onto resolved macro-scale heterogeneity. Likewise, motivated by correlations and non-uniformity of pebble-bed reactors, Larsen and

Vasquez [2011; 2014a; 2014b] recently derived generalized RTE-like models for neutron transport that allows non-exponential, angle-dependent free flight distributions by introducing an additional “memory” parameter tracking the distance since the last interaction. We discuss these models in more detail in Section 2 and show that they are a special case of our more general model.

Non-exponential behavior also arises when there is “cross-talk” between neutrons (photons) of different energy-levels (wavelengths). This is often ignored in graphics – unless fluorescence/in-elastic scattering [Gutierrez et al. 2008; Jarabo and Arellano 2018] needs to be considered – but is quite common in neutron transport where multi-energy simulations are standard practice. D’Eon [2016] provides an excellent overview of previous work on non-exponential free-flights outside of graphics. Unfortunately, it is difficult to directly leverage these prior formulations in graphics, since they do not construct a theory considering both volumes and “solid” surfaces that can readily be solved using a chosen Monte Carlo rendering algorithm.

Graphics. There has also been some work in graphics that explored non-classical transport, often in the context of approximating complex geometry as a continuous participating medium. Moon et al. [2007] introduced the concept of non-exponential transport to graphics, and several approaches [Meng et al. 2015; Moon et al. 2008; Müller et al. 2016] have since considered the problem of accelerating multiple scattering within discrete random (granular) media. These methods all addressed isolated rendering problems (via data-driven tabulation or by simplifying to the classical RTE) and do not provide a theory for how to handle non-classical transport in general. Jakob et al. [2010] and follow-up work [Dupuy et al. 2016; Heitz et al. 2015] took a step in this direction by enhancing the classical RTE to account for angular structure/correlations using the microflake model. These theories, however, still ignore the *spatial* correlations that give rise to non-exponential free-flight behavior in media. Our theory shows how to achieve this by modeling unresolved spatial correlations and density fluctuations (volumetric “roughness”) statistically, akin to how microfacet models [Blinn 1977; Cook and Torrance 1981] represent surface roughness statistically.

Artistic control. The film industry has recently started exploring the use of non-exponential behavior for artistic control and accelerated multiple scattering [Bouthors et al. 2008; Nowrouzezahrai et al. 2011; Wrenninge et al. 2013, 2011]. Common strategies include artificially lengthening the mean-free path for shadow rays of higher-order scattering [Bouthors et al. 2008; Wrenninge et al. 2011] or replacing exponential transmittance with a sum of exponentials [Wrenninge et al. 2013]. While originally developed in an ad-hoc way, we discuss in Section 4.4 how such sums of exponentials have a physically grounded interpretation. Pixar’s RenderMan likewise allows specifying separate “shallow” and “deep” scattering mean-free paths to preserve fine-scale surface details while controlling long-range subsurface scattering. D’Eon [2013] derived diffusion equations for non-exponential free-flights which could conceivably be used in this context. Most recently, Wrenninge et al. [2017] applied non-exponential free-flights from Davis and Xu [2014] to homogeneous path-traced subsurface scattering. While such artistic techniques have proven useful in production, the industry has gradually shifted

to physically based light transport using path tracing-based approaches [Christensen and Jarosz 2016; Fong et al. 2017; Keller et al. 2015], reinforcing the need for a more flexible theory of light transport in scenes with surfaces and correlated heterogeneous media which can be solved in a consistent way using any desired Monte Carlo rendering algorithm.

Concurrent Work. Concurrent to our work, d’Eon [2018] extend the work of Larsen and Vasques [2011] to include surfaces while ensuring reciprocity of the resulting RTE. He achieves this by deriving an uncorrelated free-flight PDF that satisfies reciprocity over two-segment paths. Our f_p function (Section 3) is equivalent to d’Eon’s uncorrelated free-flight PDF; however, we are able to show that the correlated and uncorrelated free-flight PDF are linked because of the underlying physical process of ensemble averaging. Reciprocity in our model is then simply a result of modelling the underlying physical process. In contrast, the uncorrelated free-flight PDF of d’Eon is a mathematical construct and is not a result of any particular property of the medium. In addition, the RTE derived by d’Eon is only *weakly* reciprocal, i.e. the path throughput is not reciprocal for individual path segments, which complicates rendering with bidirectional methods. Our path integral on the other hand is fully reciprocal; it also supports heterogeneity.

Jarabo et al. [2018] also extend the Larsen and Vasques model to include surfaces and statistical heterogeneity. Their RTE is general and allows for different correlated and uncorrelated source terms and free-flights. Although this RTE could be made reciprocal, Jarabo et al. do not explore the mathematical relationship between correlated and uncorrelated free-flight PDFs, and do not discuss reciprocity. In addition, although their RTE supports heterogeneity in theory, they do not develop a practical model and are only able to render homogeneous results. Finally, they do not provide a path integral formulation of their theory, which makes integration with bidirectional methods difficult.

In the field of neutron transport, Camminady et al. [2017] extend the work of Larsen and Vasques [2011] to support piecewise homogeneous media. They do so by combining a 1D free-flight PDF with the accumulated extinction across interfaces of the medium. Their model can be viewed as a special case of our heterogeneity model, which supports continuously varying heterogeneity. Unlike Camminady et al., we also provide a physical intuition of this heterogeneity model in the context of stochastic media and develop efficient and unbiased heterogeneous distance sampling methods for certain classes of media.

2 BACKGROUND

The classical RTE can be viewed as the expected value of a stochastic process, where in each realization light paths interact with scattering particles distributed independently at random. This gives rise to the classical exponential model.

Similar to the classical RTE, we formulate the non-exponential rendering problem as the average behavior of light paths interacting with a stochastic medium (Figure 3, top). To formulate the problem, we need to first pick a model of how each realization of the medium is expressed: We can either model realizations as discrete point distributions of scatterers (Figure 2, top), or describe each realization

as a heterogeneous classical medium (Figure 2, bottom). The latter model is more convenient to work with mathematically, and we will use this model for the rest of this paper. However, we include an alternative derivation of our path integral using the discrete model in the supplemental material to demonstrate that both models are equivalent for the purposes of our theory.

In the following, we denote a realization of the medium by the symbol μ , and use it as a subscript for quantities that depend on the realization. For example, we associate a heterogeneous extinction coefficient $\sigma_\mu(\mathbf{x})$ with each realization. We use the notation $\langle \cdot \rangle$ to express ensemble averages; for example $\sigma(\mathbf{x}) = \langle \sigma_\mu(\mathbf{x}) \rangle$ is the macroscopic extinction coefficient, averaged over all realizations of the medium.

Within each realization, light transport is exponential, and the transmittance, free-flight PDF and optical depth are defined as usual:

$$\text{Tr}_\mu(\mathbf{x}, \mathbf{x}_t) = e^{-\tau_\mu(\mathbf{x}, \mathbf{x}_t)}, \quad \text{and} \quad p_\mu(\mathbf{x}, \mathbf{x}_t) = e^{-\tau_\mu(\mathbf{x}, \mathbf{x}_t)} \sigma_\mu(\mathbf{x}_t), \quad (1)$$

$$\text{with} \quad \tau_\mu(\mathbf{x}, \mathbf{x}_t) = \int_0^t \sigma_\mu(\mathbf{x}_s) ds, \quad (2)$$

and Tr_μ and p_μ further satisfy the well-known integral relationship

$$\text{Tr}_\mu(\mathbf{x}, \mathbf{x}_t) = \int_t^\infty p_\mu(\mathbf{x}, \mathbf{x}_s) ds. \quad (3)$$

We use the notation $\mathbf{x}_t = \mathbf{x} + t \cdot \omega$ to express points along rays. Within each realization, light transport is described by the standard path integral [Veach 1997]. The measurement I is the product of emission $L_e(\mathbf{x}_0, \mathbf{x}_1)$, sensor response $W(\mathbf{x}_{k-1}, \mathbf{x}_k)$ and path throughput $g(\bar{\mathbf{x}})$ for length- k path $\bar{\mathbf{x}} = \mathbf{x}_0 \cdots \mathbf{x}_k$, integrated over the set of paths \mathcal{P} :

$$I_\mu = \int_{\mathcal{P}} L_e(\mathbf{x}_0, \mathbf{x}_1) g_\mu(\bar{\mathbf{x}}) W(\mathbf{x}_{k-1}, \mathbf{x}_k) d\bar{\mathbf{x}}, \quad (4)$$

with

$$g_\mu(\bar{\mathbf{x}}) = \left[\prod_{i=1}^{k-1} f_\mu(\mathbf{x}_i) \Sigma_\mu(\mathbf{x}_i) \right] \left[\prod_{i=0}^{k-1} \text{Tr}_\mu(\mathbf{x}_i, \mathbf{x}_{i+1}) G(\mathbf{x}_i, \mathbf{x}_{i+1}) \right] \quad (5)$$


and

$$\Sigma_\mu(\mathbf{x}) = \begin{cases} \sigma_\mu(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{V} \\ 1 & \text{if } \mathbf{x} \in \partial\mathcal{V} \end{cases} \quad (6)$$

$$f_\mu(\mathbf{x}_i) = \begin{cases} f_{m,\mu}(\mathbf{x}_{i-1}, \mathbf{x}_i, \mathbf{x}_{i+1}) \cdot \alpha_\mu(\mathbf{x}) & \text{if } \mathbf{x}_i \in \mathcal{V} \\ f_s(\mathbf{x}_{i-1}, \mathbf{x}_i, \mathbf{x}_{i+1}) & \text{if } \mathbf{x}_i \in \partial\mathcal{V}, \end{cases} \quad (7)$$

$$G(\mathbf{x}_i, \mathbf{x}_{i+1}) = \frac{D(\mathbf{x}_i) V(\mathbf{x}_i, \mathbf{x}_{i+1}) D(\mathbf{x}_{i+1})}{\|\mathbf{x}_i - \mathbf{x}_{i+1}\|^2}, \quad D(\mathbf{x}_i) = \begin{cases} 1 & \text{if } \mathbf{x}_i \in \mathcal{V} \\ \cos \theta_i & \text{if } \mathbf{x}_i \in \partial\mathcal{V} \end{cases} \quad (8)$$

where $f_{m,\mu}$ and f_s are the phase function and BRDF, respectively, α_μ is the albedo, \mathcal{V} represents the set of all medium points, $\partial\mathcal{V}$ the set of all surface points and $V(\mathbf{x}_i, \mathbf{x}_{i+1})$ is the visibility function. We illustrate these terms in the figure below:



Looking at Equation (5) and the illustration above, it may seem trivial to simply insert a numerically averaged transmittance into the classical RTE and expect non-exponential behavior. In the supplemental material we prove, however, that only an exponential transmittance satisfies energy conservation in the classical RTE, requiring us to develop a more flexible theory.

Instead, we treat the non-exponential rendering problem as computing the expected value of this measurement, i.e. the ensemble average $\langle I_\mu \rangle$. In the following, we assume that emission and measurement happen only on surfaces and do not depend on the realization; this is for notational convenience, and we present a general derivation in the supplemental material. This assumption reduces the problem to computing the average path throughput:

$$\langle g_\mu(\bar{\mathbf{x}}) \rangle = \left\langle \left[\prod_{i=1}^{k-1} f_\mu(\mathbf{x}_i) \Sigma_\mu(\mathbf{x}_i) \right] \left[\prod_{i=0}^{k-1} \text{Tr}_\mu(\mathbf{x}_i, \mathbf{x}_{i+1}) G(\mathbf{x}_i, \mathbf{x}_{i+1}) \right] \right\rangle. \quad (9)$$

This ensemble average captures all correlation effects in the medium, but is impractical to compute in general: All terms of the path throughput above depend on the realization μ and are correlated with each other, which makes simplification challenging without making further assumptions.

In the field of neutron transport, Larsen and Vasques [2011] considered a simplified version of this problem, in which

- the medium is infinite and contains no surfaces,
- the medium is *statistically homogeneous* or stationary, i.e. the statistics of the medium do not change spatially,
- the phase function and albedo are independent of the realization, and
- photon free-flights only depend on the distance to the previous event (scattering or emission).

Using these assumptions, Larsen and Vasques successfully derive a generalized linear Boltzmann equation (GBE) that describes non-exponential transport. Although they do not derive a path integral and their model has significant limitations, we take inspiration from their approach and apply some of their assumptions to our more general problem (Equation (9)).

Our goal is to derive a new path integral that supports non-exponentiality and lifts some of the restrictions of the Larsen and Vasques model that make it unsuitable for graphics. In particular, our path integral should include bounded media and surface transport within the same framework, in addition to supporting macroscopic heterogeneity. This path integral should also be practical to compute and fully reciprocal to allow the use of all rendering algorithms used in graphics.

The derivation of our path integral proceeds in three steps: We first import two critical assumptions from the Larsen and Vasques model and apply them to Equation (9) to obtain a new non-exponential path integral. This path integral is not reciprocal in the traditional sense, and we address this issue in Section 3.1 to derive a fully reciprocal version. Up until this point, we support statistical heterogeneity in its most general form, which makes practical rendering difficult; to address this, we introduce a practical heterogeneity model in Section 3.2 that preserves non-exponentiality while enabling efficient rendering algorithms for the heterogeneous case.




Fig. 3. We formulate the problem of non-exponential rendering as the ensemble average of the path throughput over all realizations of a stochastic medium (top). We approximate the solution to this problem by computing the ensemble average separately over each segment (bottom).

3 A NON-EXPONENTIAL PATH INTEGRAL

Before we can proceed with the derivation of our proposed path integral, we first need to formally define the notion of an ensemble average. We associate with each realization μ of the medium a probability measure $dP\{\mu\}$ that describes its likelihood. An ensemble average of an arbitrary function $f_\mu(\mathbf{x})$ is then a measurement over \mathcal{R} , the space of all realizations (the *ensemble*):

$$\langle f_\mu(\mathbf{x}) \rangle = \int_{\mathcal{R}} f_\mu(\mathbf{x}) dP\{\mu\}. \quad (10)$$

The ensemble average preserves all properties of the integral, such as linearity.

Returning to the problem of computing $\langle g_\mu(\bar{\mathbf{x}}) \rangle$, we now introduce our first assumption: Phase function and albedo are uncorrelated with the realization μ . The linearity property of $\langle \cdot \rangle$ then allows us to move these terms outside of the ensemble average:

$$\langle g_\mu(\bar{\mathbf{x}}) \rangle \approx \prod_{i=1}^{k-1} f(\mathbf{x}_i) \prod_{i=0}^{k-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) \left\langle \prod_{i=0}^{k-1} \text{Tr}_\mu(\mathbf{x}_i, \mathbf{x}_{i+1}) \Sigma_\mu(\mathbf{x}_{i+1}) \right\rangle. \quad (11)$$

This simplifies the problem to computing the average transmittance and extinction coefficient across the entire path. For notational convenience, we have also moved the geometry term (which is independent of the realization) and added a $\Sigma(\mathbf{x}_k)$ term (which is 1, since we have assumed surface sensors).

The second assumption we import from Larsen and Vasques is that photon free-flights only depend on the distance to the previous event. In other words, the “memory” of each photon does not extend beyond a single segment of the path; intuitively, the photon “jumps” into a different realization every time it interacts with the medium.

This assumption effectively decorrelates segments of the path from each other, and this allows us to compute the ensemble average independently for each segment instead of over the path as a whole (Figure 3, bottom):

$$\left\langle \prod_{i=0}^{k-1} \text{Tr}_\mu(\mathbf{x}_i, \mathbf{x}_{i+1}) \Sigma_\mu(\mathbf{x}_{i+1}) \right\rangle \approx \prod_{i=0}^{k-1} \langle T_\mu(\mathbf{x}_i, \mathbf{x}_{i+1}) \rangle. \quad (12)$$

We have moved the ensemble average into the product and substituted a *transport kernel* T_μ for the segment term. What should this term be? We know that it should resemble $\text{Tr} \times \Sigma$, the segment term of the classical path integral. We observe that the classical segment term actually represents two different functions, depending on the end point of the segment: If \mathbf{x}_{i+1} lies in the medium, the segment

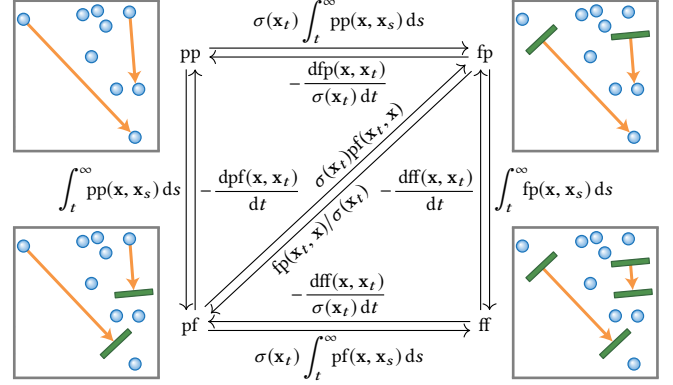


Fig. 4. The four transport functions pp, fp, pf, and ff correspond to ensemble averaging photons starting or ending in free-space or on particles, and they are related to one another mathematically via integration, differentiation, and normalization.

term becomes the free-flight PDF p_μ ; if it does not, Σ_μ becomes 1 and the segment term reduces to the transmittance Tr_μ instead.

However, simply taking the ensemble average of the classical segment term is not sufficient, since this does not take into consideration the start point of the segment: If \mathbf{x}_i lies in the medium, it represents a scattering event—in other words, \mathbf{x}_i must coincide with a scatterer. The ensemble may include many realizations of the medium in which there is no scatterer at \mathbf{x}_i , and averaging over all realizations is then not correct. Instead, we need to compute an average *conditioned* on \mathbf{x}_i coinciding with a scatterer.

This results in *four* different functions (see Figure 4) that represent the four possible combinations of \mathbf{x}_i and \mathbf{x}_{i+1} coinciding with a particle or lying in free-space, or equivalently, all possible combinations of transmittance or free-flight PDF averaged conditionally or unconditionally. These four functions form the ensemble transport kernel:

$$T(\mathbf{x}, \mathbf{y}) = \langle T_\mu(\mathbf{x}, \mathbf{y}) \rangle = \begin{cases} \text{ff}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{x} \in \partial\mathcal{V} \wedge \mathbf{y} \in \partial\mathcal{V} \\ \text{fp}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{x} \in \partial\mathcal{V} \wedge \mathbf{y} \in \mathcal{V} \\ \text{pf}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{x} \in \mathcal{V} \wedge \mathbf{y} \in \partial\mathcal{V} \\ \text{pp}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{x} \in \mathcal{V} \wedge \mathbf{y} \in \mathcal{V}, \end{cases} \quad (13)$$

where we have labeled the functions according to the type of transport they represent, i.e. “free space-to-free space”, “free space-to-particle”, “particle-to-free space” and “particle-to-particle”. To represent transport starting on a particle (i.e. pp and pf), we need a way to express ensemble averages over scatterers only. The probability density of encountering a scatterer at \mathbf{x}_i is proportional to $\sigma_\mu(\mathbf{x}_i)$ (assuming constant particle radii¹), and we can express the expectation of some function $f(\mathbf{x}_i)$ conditioned on \mathbf{x}_i coinciding with a scatterer as the *weighted* ensemble average

$$\langle f(\mathbf{x}_i) \rho_\mu(\mathbf{x}_i) \rangle \quad \text{with} \quad \rho_\mu(\mathbf{x}) = \frac{\sigma_\mu(\mathbf{x})}{\sigma(\mathbf{x})}. \quad (14)$$

¹More precisely, it is proportional to the number density. $\sigma_\mu(\mathbf{x}_i)$ is proportional to the number density times particle cross-section, the latter of which is determined by the particle radius.

With this tool, we can now define the four transport functions:

$$\text{ff}(\mathbf{x}, \mathbf{y}) = \langle \text{Tr}_\mu(\mathbf{x}, \mathbf{y}) \rangle \quad (15a)$$

$$\text{fp}(\mathbf{x}, \mathbf{y}) = \langle \text{Tr}_\mu(\mathbf{x}, \mathbf{y}) \sigma_\mu(\mathbf{y}) \rangle \quad (15b)$$

$$\text{pf}(\mathbf{x}, \mathbf{y}) = \langle \text{Tr}_\mu(\mathbf{x}, \mathbf{y}) \rho_\mu(\mathbf{x}) \rangle \quad (15c)$$

$$\text{pp}(\mathbf{x}, \mathbf{y}) = \langle \text{Tr}_\mu(\mathbf{x}, \mathbf{y}) \sigma_\mu(\mathbf{y}) \rho_\mu(\mathbf{x}) \rangle, \quad (15d)$$

and our path integral becomes

$$\langle g_\mu(\bar{\mathbf{x}}) \rangle \approx \prod_{i=1}^{k-1} f(\mathbf{x}_i) \prod_{i=0}^{k-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) \prod_{i=0}^{k-1} T(\mathbf{x}_i, \mathbf{x}_{i+1}). \quad (16)$$

It is not at all clear that this throughput is reciprocal. Indeed, we will show that the transport kernel T can give different results if the segment is reversed, and we address this problem in Section 3.1. However, before then it is worth analyzing the four different functions of T in more detail and develop some intuition for their meaning.

$\text{pp}(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $\text{fp}(\mathbf{x}_i, \mathbf{x}_{i+1})$ are averages of free-flight PDFs and are themselves free-flight PDFs: They correspond exactly to the density of particle-to-particle and free-space-to-particle flights. On the other hand, $\text{pf}(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $\text{ff}(\mathbf{x}_i, \mathbf{x}_{i+1})$ are averages of transmittances and are themselves a transmittance.

In the classical RTE, the transmittance and free-flight PDF are tightly related (Equation (3)). Because the functions we introduced are themselves averages of the classical transmittance and free-flight PDF, and the averaging operator is linear, we can easily prove (see supplemental) that the same holds for the averaged functions:

$$\text{pf}(\mathbf{x}, \mathbf{x}_t) = \int_t^\infty \text{pp}(\mathbf{x}, \mathbf{x}_s) ds, \quad \text{ff}(\mathbf{x}, \mathbf{x}_t) = \int_t^\infty \text{fp}(\mathbf{x}, \mathbf{x}_s) ds. \quad (17)$$

Perhaps more surprisingly, we can also prove that pf and fp are directly related. Intuitively, $\text{fp}(\mathbf{x}, \mathbf{y})$ represents the PDF of hitting a scatterer at \mathbf{y} starting from free-space point \mathbf{x} , whereas $\text{pf}(\mathbf{x}, \mathbf{y})$ represents the probability of reaching free-space point \mathbf{y} starting from a scatterer at \mathbf{x} . These functions describe reverse events in a sense, and indeed we can show

$$\text{fp}(\mathbf{x}, \mathbf{y}) = \langle \text{Tr}_\mu(\mathbf{x}, \mathbf{y}) \sigma_\mu(\mathbf{y}) \rangle \quad (18)$$

$$= \sigma(\mathbf{y}) \cdot \left\langle \text{Tr}_\mu(\mathbf{y}, \mathbf{x}) \frac{\sigma_\mu(\mathbf{y})}{\sigma(\mathbf{y})} \right\rangle = \sigma(\mathbf{y}) \cdot \langle \text{Tr}_\mu(\mathbf{y}, \mathbf{x}) \rho_\mu(\mathbf{y}) \rangle \quad (19)$$

$$= \sigma(\mathbf{y}) \cdot \text{pf}(\mathbf{y}, \mathbf{x}) \quad (20)$$

These four functions are not arbitrary; they represent probabilities (densities) of events between two points, and are tightly linked to each other—given one function, all three others can be derived. The relationships between them are a direct result of averaging transmittances and free-flight PDFs (17) and performing unconditioned averages or averages correlated with scatterer locations (20). We illustrate these relationships and their inverses in Figure 4.

Comparison to the GBE. It is worth pointing out that in the case of an infinite, statistically homogeneous medium without surface boundaries, our proposed path integral is equivalent to the General Boltzmann Equation of Larsen and Vasques [2011]. This is because then all path vertices lie in the medium and our transport kernel only evaluates pp , the particle-to-particle free-flight PDF; this is the same function used in the GBE.

3.1 Reciprocity

So far, our transport kernel is not reciprocal, i.e. $T(\mathbf{x}, \mathbf{y}) \neq T(\mathbf{y}, \mathbf{x})$ in the general case. There are two reasons for this: Free-flight PDFs (such as pp) are not generally reciprocal, and swapping the end-points of a surface-to-medium segment causes a different case statement to be evaluated (fp in one case, pf in the other).

However, we will show that this problem can be addressed through a simple rearranging of terms. We proceed by deriving a reciprocal transport kernel through mathematical manipulation and then offer an intuitive interpretation of this new kernel.

We begin by defining

$$\Sigma(\mathbf{x}) = \begin{cases} \sigma(\mathbf{x}) & \text{if } \mathbf{x} \in V \\ 1 & \text{else.} \end{cases} \quad (21)$$

Now consider the modified transport kernel \tilde{T}

$$\tilde{T}(\mathbf{x}, \mathbf{y}) = \frac{T(\mathbf{x}, \mathbf{y})}{\Sigma(\mathbf{y})} = \begin{cases} \text{ff}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{x} \in \partial\mathcal{V} \wedge \mathbf{y} \in \partial\mathcal{V} \\ \frac{\text{fp}(\mathbf{x}, \mathbf{y})}{\sigma(\mathbf{y})} = \text{pf}(\mathbf{y}, \mathbf{x}) & \text{if } \mathbf{x} \in \partial\mathcal{V} \wedge \mathbf{y} \in \mathcal{V} \\ \text{pf}(\mathbf{x}, \mathbf{y}) & \text{if } \mathbf{x} \in \mathcal{V} \wedge \mathbf{y} \in \partial\mathcal{V} \\ \frac{\text{pp}(\mathbf{x}, \mathbf{y})}{\sigma(\mathbf{y})} & \text{if } \mathbf{x} \in \mathcal{V} \wedge \mathbf{y} \in \mathcal{V}. \end{cases} \quad (22)$$

Different to before, \tilde{T} will evaluate pf for a segment that has one end-point on a surface and another in the medium, regardless of direction. Additionally, \tilde{T} evaluates $\text{pp}(\mathbf{x}, \mathbf{y})/\sigma(\mathbf{y})$ for a segment with both end points in the medium. We can easily prove (see supplemental) that this term is equivalent to averaging the transmittance along a segment, conditioned on both end points lying on a particle. This means that \tilde{T} evaluates only averages of transmittance for all four possible configurations of a segment. The classical transmittance is reciprocal, and so are averages of transmittances; hence, \tilde{T} is reciprocal.

To complete the derivation, we use the relation $T(\mathbf{x}, \mathbf{y}) = \Sigma(\mathbf{y})\tilde{T}(\mathbf{x}, \mathbf{y})$ in Equation (16) to obtain our proposed reciprocal, non-exponential path throughput:

$$\langle g_\mu(\bar{\mathbf{x}}) \rangle \approx \left[\prod_{i=1}^{k-1} f(\mathbf{x}_i) \Sigma(\mathbf{x}_i) \right] \left[\prod_{i=0}^{k-1} \tilde{T}(\mathbf{x}_i, \mathbf{x}_{i+1}) G(\mathbf{x}_i, \mathbf{x}_{i+1}) \right]. \quad (23)$$

We have rearranged the Σ terms into the vertex terms so that this equation looks very similar to the classical path integral. Indeed, the only difference to the exponential case is the replacement of transmittance Tr with our modified transport kernel \tilde{T} . This makes it straightforward to simulate our proposed non-exponential transport model with existing (bidirectional) rendering methods. We describe the necessary changes to existing renderers in detail in Section 5.

3.2 Heterogeneity

In Section 2, we formulated the original averaging problem in a statistically heterogeneous medium. As a result of this, all equations so far explicitly support heterogeneity in its most general form—for example, the macroscopic density could vary, or the correlation between scattering particles could change from one part of the medium to the next.

Although this generality can be useful, it also means that the transport kernel $\tilde{T}(\mathbf{x}, \mathbf{y})$ is a 6D function. In the general case, this

makes it impractical to derive analytical forms of \tilde{T} or to tabulate it numerically, and rendering with this form of the path integral is difficult. To address this, we introduce an explicit decomposition of the medium into a 1D transport kernel and a 3D density field that enables practical rendering methods while preserving non-exponentiality.

Our current model and the model of Larsen and Vasques can be viewed as lying on two opposite sides of a spectrum: The Larsen and Vasques model assumes a fully homogenized model with identical statistics everywhere, whereas our model makes no assumptions about how correlations are distributed spatially.

Assuming statistical homogeneity is convenient, because ensemble-averaged transmittances and free-flight PDFs become functions of only distance. Parametric models of transmittance can be readily derived for certain classes of stochastic media (Section 4) if they are statistically homogeneous, and for all other media, the required 1D functions can be easily tabulated numerically. On the other hand, heterogeneous media find wide-spread use in graphics, and any path integral not supporting such media is heavily limited.

This motivates us to seek a middle ground between the two models that combines the benefits of both: We decompose the description of the medium into a heterogeneous, macroscopic density $\sigma(\mathbf{x})$ and a 1D transport kernel $T(t)$.

In other words, we assume that the medium is statistically heterogeneous, but that the underlying correlations are the same throughout the medium. Our model imposes that for a medium with unit macroscopic density ($\sigma(\mathbf{x}) = 1$), averages of transmittance and free-flight PDF become functions only of distance:

$$\text{ff}(\mathbf{x}, \mathbf{x}_t) = \text{ff}(t) = \langle \text{Tr}(t) \rangle \quad (24a)$$

$$\text{fp}(\mathbf{x}, \mathbf{x}_t) = \text{fp}(t) = \langle \text{Tr}(t) \sigma_\mu(\mathbf{x}_t) \rangle \quad (24b)$$

$$\text{pf}(\mathbf{x}, \mathbf{x}_t) = \text{pf}(t) = \langle \text{Tr}(t) \rho_\mu(\mathbf{x}_t) \rangle \quad (24c)$$

$$\text{pp}(\mathbf{x}, \mathbf{x}_t) = \text{pp}(t) = \langle \text{Tr}(t) \sigma_\mu(\mathbf{x}_t) \rho_\mu(\mathbf{x}_t) \rangle \quad (24d)$$

We can view these functions as forming a “base” transport kernel $\tilde{T}_{\text{base}}(t)$ that captures only correlations, but not heterogeneity.

For a medium with non-unit macroscopic density but the *same correlations* as the medium above, we express its transport kernel in terms of the base kernel and the macroscopic optical depth:

$$\tilde{T}(\mathbf{x}_i, \mathbf{x}_{i+1}) = \tilde{T}_{\text{base}}(\tau_i) = \begin{cases} \text{ff}(\tau_i) & \text{if } \mathbf{x}_i \in \partial\mathcal{V} \wedge \mathbf{x}_{i+1} \in \partial\mathcal{V} \\ \text{pp}(\tau_i) & \text{if } \mathbf{x}_i \in \mathcal{V} \wedge \mathbf{x}_{i+1} \in \mathcal{V} \\ \text{pf}(\tau_i) & \text{else,} \end{cases} \quad (25)$$

with $\tau_i = \tau(\mathbf{x}_i, \mathbf{x}_{i+1})$. We obtained this equation by inserting the equivalent of (22) for \tilde{T}_{base} . Because pp is a PDF, it incurs a Jacobian factor of $\sigma(\mathbf{x}_{i+1})$ from the transformation. However, this factor cancels the division by σ from (22), leaving us with a remarkably simple transport kernel.

This decomposition can be interpreted as separating the medium into explicitly modeled density variations and statistically modeled density fluctuations that exist beyond what is modeled explicitly. This model has several advantages: Reducing the transport kernel to 1D makes it practical to estimate it numerically or derive it analytically. Additionally, the decomposition of the medium into a

heterogeneous density and a “plug-in” function describing correlations allows us to easily superimpose non-exponentiality onto existing heterogeneous media without additional effort. Finally, relating the microscopic and macroscopic model via the optical depth enables practical sampling methods that are even unbiased in some special cases (Section 5).

Physical interpretation. This heterogeneity model has a simple physical interpretation, which we illustrate in Figure 5. The left column in this figure represents an instance of the “base” medium, which is statistically homogeneous and has a mean extinction of 1. The middle column of this figure represents an instance of a medium with the same correlations, but a mean extinction of 2. Through Equation (25), we can relate transmittance in this medium to transmittance in the base medium via the optical depth; this can be interpreted as “squishing” the base medium horizontally and “stretching” it vertically, so that it has twice the amplitude and twice the frequency. In the general case of a macroscopically heterogeneous medium (right column), the amount of this stretching is determined by the local macro-scale extinction coefficient $\sigma(\mathbf{x})$.

4 MODELING NON-EXPONENTIAL ATTENUATION

Now that we have a theory than can admit non-exponential behavior, the next remaining question is how we should obtain, represent, and design the four transport functions (24), and how we can physically interpret the corresponding light transport.

It is common to model surface appearance at multiple scales, e.g. by representing large-scale variation using explicit geometry and displacements, while modeling fine-scale roughness statistically using a BRDF. We propose to leverage our theory to provide a similar decomposition for volumes, where we model large-scale heterogeneous variation explicitly with spatially-varying macroscopic media properties ($\alpha(\mathbf{x})$, $\sigma(\mathbf{x})$, $f_m(\mathbf{x})$), but we can additionally account for scatterer correlations or unresolved fine-scale heterogeneity *statistically* by modifying the transport functions (24). Inspired by phenomenological [Ashikhmin and Shirley 2000; Phong 1975], data-driven [Ashikhmin et al. 2000; Bagher et al. 2016; Matusik et al. 2003], and statistical microfacet [Blinn 1977; Cook and Torrance 1981] models for surface roughness, we can likewise obtain transmittance functions in different ways:




Fig. 5. We express volumetric micro-roughness σ_μ (left) statistically and model macro-scale extinction σ explicitly (black, homogeneous: center, heterogeneous: right) and combine them (blue). When we pass the base transport functions through macroscopic optical depth $\tau(s)$, the statistical fluctuations/correlations that σ_μ (left) represents are scaled vertically and squeezed horizontally (middle, right) by the local extinction coefficient $\sigma(s)$.

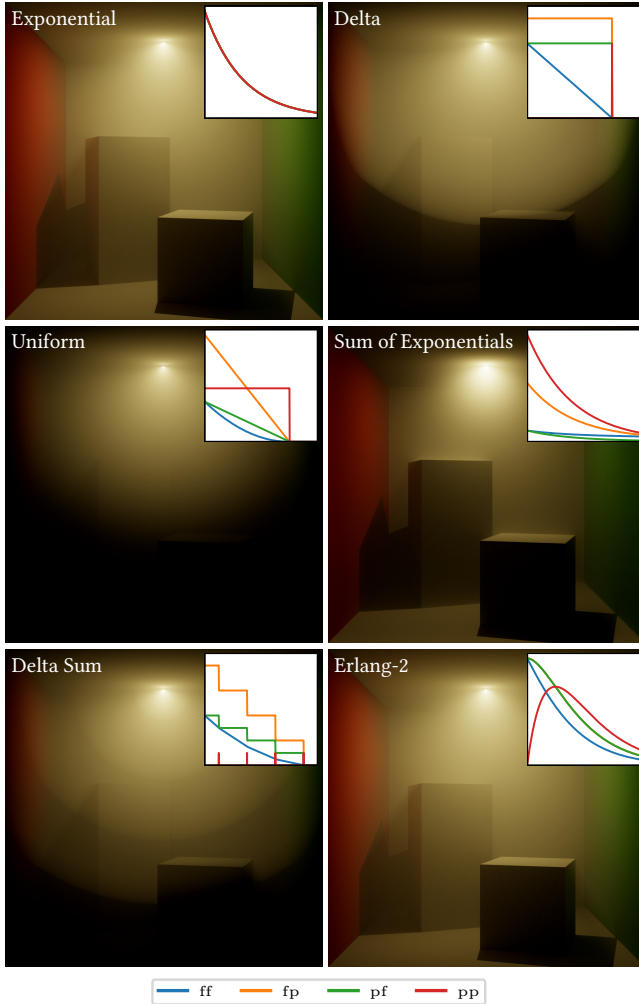


Fig. 6. We show a homogeneous medium in a Cornell Box rendered with six different free-flight PDFs. We set pp to be the classical exponential (top left), a delta function $\delta_{3/4}$ (top right), uniform $\mathcal{U}(0, 3/4)$ (middle left), a sum of two exponentials (middle right), a delta sum $\text{III}(4)$ (bottom left), and the Erlang-2 distribution $\mathcal{E}(2)$ (bottom right). Compared to the exponential baseline (top left), we can achieve a wide range of appearances through simple parametric transmittance curves.

- (1) The artist-driven “phenomenological” way (Section 4.1), where we directly design a free-flight PDF/transmittance free-hand or with simple parametric models;
- (2) The “data-driven” physically based way (Section 4.2), where we instantiate a distribution of physical scatterers, and obtain the transmittance induced by these scatterers numerically (through sampling); and
- (3) The “statistical” physically based way (Section 4.3) that seeks analytic parametric models for these functions driven by some statistical description of the distribution of the physical medium scatterers (Section 4.4).

4.1 Transmittance via directly designed free-flight PDFs

A simple phenomenological approach is to directly prescribe the free-flight PDF (either Equation (24a) or (24c)) to something other than an exponential. Table 1 lists the PDF $\text{pdf}(x)$ and CDF $\text{cdf}(x)$ of common statistical distributions. We can take any such distribution defined on the positive real line, and directly set e.g. $\text{fp}(\tau) = \text{pdf}(\tau)$, we then have $\text{ff}(\tau) = 1 - \text{cdf}(\tau)$, and the other functions follow from the relationships in Figure 4.

We experimented with a variety of such distributions, including a step, linear ramp, sum of impulses, sum of exponentials, and Erlang-2. Figure 6 shows renderings and the corresponding parametric transport curves. It would also be possible to allow artists to design custom curves using a familiar graph editor interface.

The sum of two exponentials is particularly useful as the two exponents allow separately controlling the falloff of light near the start and tail of the distribution. Wrenninge et al. [2013] used this falloff for shadow rays to allow light to penetrate deeper, approximating multiple scattering. Our theory allows us to incorporate this in a consistent way for arbitrary bounces and light transport algorithms.

Inspired by this idea, we also took a classical monopole diffusion profile [Habel et al. 2013; Jensen et al. 2001] and interpreted it as a free-flight distribution. The radial falloff of a classical monopole is proportional to an Erlang-2 distribution with rate parameter $\lambda = \sqrt{3(1 - \alpha)(1 - g\alpha)}$ dependent on albedo α and the average cosine g . Simulating single scattering from a point light in a homogeneous medium with this transmittance profile would produce results similar to an all-bounce monopole diffusion approximation. Applied as a transmittance profile in our theory, this lets light penetrate further into the medium at each bounce, and allows us to approximate multiple-scattered transport using fewer explicitly simulated bounces.

4.2 Data-driven transmittance via ensemble averaging

While directly designing transmittance functions in this top-down approach provides high-level artistic controls, it may not correspond to any physically realizable distribution of scattering particles. To obtain physically justifiable profiles, we can instead numerically approximate the ensemble-averaging process by generating stochastic realizations of heterogeneous media or discrete point distributions.

Without loss of generality, we will consider the function $\text{ff}(t)$ knowing we can obtain the other three functions via integration, normalization or differentiation using the relationships in Figure 4.

Ensemble-averaging discrete media/particles. One physically based approach would be to explicitly construct a discrete collection of scattering *particles* (Figure 2a-d, top) and approximate the transmittance using a Monte Carlo ensemble average:

$$\text{ff}(t) = \langle V_\mu(\mathbf{x}, \mathbf{x}_t) \rangle \approx \frac{1}{N} \sum_{i=1}^N V_{\mu_i}(\mathbf{x}, \mathbf{x}_t), \quad (26)$$

where V_μ denotes the visibility within each random realization, and the ensemble average depends only on distance t if the stochastic point process is statistically isotropic and homogeneous.

Table 1. The probability density functions pdf(x), cumulative distribution function cdf(x), and characteristic functions $\varphi(r)$ for a variety of statistical distributions. $\Gamma(\alpha)$ and $\gamma(s, x)$ are the complete and lower incomplete gamma functions, and $\delta(x)$ and $H(x)$ are the Dirac delta and Heaviside step functions. When μ is used as a parameter it specifies the mean of the corresponding PDF. We omit the implicit values of 0 or 1 in PDFs and CDFs with bounded support. When hand-designing attenuation via free-flight distributions, pdf(τ) describes the free-flight distribution pp(τ), and transmittance is obtained via pf(τ) = 1 - cdf(τ); fp and ff are obtained via normalization and integration. When designing attenuation via $1/f$ noise, pdf($\bar{\sigma}_\mu$) describes the distribution of $\bar{\sigma}_\mu$ due to fBm micro-fluctuations, transmittance is obtained from $\varphi(i\tau)$, and the free-flight distribution becomes $-d\varphi(i\tau)/d\tau$. The last two columns list sampling routines for these two approaches.

Distribution	Parameters	pdf(x)	cdf(x)	$\varphi(r)$	Sampling	
					$x_i \propto \text{pdf}(x)$	$\propto -d\varphi(ir)/dr$
Delta δ_μ	mean: $\mu > 0$	$\delta(x - \mu)$	$H(x - \mu)$	$e^{ir\mu}$	$x_\delta = \mu$	$-\frac{\ln \xi}{x_\delta}$
DeltaSum III(n)	impulses: $\{\mu_1, \dots, \mu_n\}$	$\frac{1}{n} \sum_{i=1}^n \delta(x - \mu_i)$	$\frac{1}{n} \sum_{i=1}^n H(x - \mu_i)$	$\frac{1}{n} \sum_{i=1}^n e^{ir\mu_i}$	$x_{III} = \text{uniform from } \{\mu_1, \dots, \mu_n\}$	$-\frac{\ln \xi}{x_{III}}$
Uniform $\mathcal{U}(a, b)$	range: $0 < a < x < b < \infty$	$\frac{1}{b-a}$	$\frac{x-a}{b-a}$	$\frac{e^{irb} - e^{ira}}{ir(b-a)}$	$x_u = (b-a)\xi + a$	$-\frac{\ln \xi_2}{x_u}$
Linear $\mathcal{L}(b)$	range: $0 < x < b < \infty$	$\frac{2}{b} \left(1 - \frac{x}{b}\right)$	$1 - 2\frac{x}{b} + \frac{x^2}{b^2}$	—	$x_l = b(1 - \sqrt{\xi})$	$-\frac{\ln \xi_2}{x_l}$
Normal $\mathcal{N}(\mu, v)$	mean: $0 \leq \mu < \infty$ variance: $v > 0$	$(2\pi v)^{-\frac{1}{2}} e^{-\frac{(x-\mu)^2}{2v}}$	$\frac{1}{2} \left[1 + \text{erf}\left(\frac{x-\mu}{\sqrt{2v}}\right)\right]$	$e^{ir\mu - r^2 v/2}$	$x_n = \sqrt{-2 \ln \xi_1} \cos(2\pi \xi_2)$	$-\frac{\ln \xi_2}{x_n}$
Gamma $\Gamma(\mu, \alpha)$	mean: $0 \leq \mu < \infty$ shape: $\alpha > 0$	$\frac{1}{\Gamma(\alpha)} \left(\frac{\alpha}{\mu}\right)^\alpha x^{\alpha-1} e^{-x\alpha/\mu}$	$\frac{1}{\Gamma(\alpha)} \gamma(\alpha, x\alpha/\mu)$	$(1 - ir\mu/\alpha)^{-\alpha}$	Marsaglia's method	$(\xi^{-1/\alpha} - 1)\alpha/\mu$
Erlang-2 $\mathcal{E}(\lambda)$	rate: $\lambda > 0$ shape: $k \in \mathbb{N}$	$\lambda^2 x e^{-\lambda x}$	$e^{-\lambda x} + \lambda x e^{-\lambda x}$	$(1 - ir/\lambda)^{-2}$	$x_{\mathcal{E}} = -\frac{1}{\lambda} \ln(\xi_1 \xi_2)$	$(\xi^{-1/2} - 1)\lambda$

This ensemble average amounts to computing the probability of having no scattering particles over a distance t along a random ray. If the positions of the particles are statistically independent, then this tabulation procedure will converge to the exponential distribution (see Figure 2a,e), since it is computing the probability of no events occurring over a distance t within a Poisson process. However, it is also possible to instantiate points with specially crafted spatial correlations that e.g. induce or inhibit clumping. Figure 2e visualizes the transmittance curves obtained from 2D discrete points sets with spectral power falloffs corresponding to “blue” (b), “pink” (c) and “red” (d) noise respectively. Data-driven curves like these are a “gold standard” in the sense that they can handle any distribution of scatterers we can explicitly construct. But they require expensive sampling and tabulation for every set of parameters, which quickly becomes intractable for large collections of particles in higher dimensions.

Ensemble-averaging continuous densities. We can forego instantiating discrete particles and instead model the spatial correlations via a heterogeneous micro-scale density field $\sigma_\mu(\mathbf{x})$ (Figure 2a-d, bottom). Given a fixed realization of $\sigma_\mu(\mathbf{x})$, the transmittance defined in Equation (1) can be written equivalently as:

$$\text{Tr}_\mu(\mathbf{x}, \mathbf{x}_t) = e^{-t \bar{\sigma}_\mu(\mathbf{x}, \mathbf{x}_t)}, \quad \text{where} \quad \bar{\sigma}_\mu(\mathbf{x}, \mathbf{x}_t) = \frac{1}{t} \int_0^t \sigma_\mu(\mathbf{x}_s) ds \quad (27)$$

is the micro-scale density field $\sigma_\mu(\mathbf{x})$ averaged, or “blurred”, along a line segment of length t between \mathbf{x} and \mathbf{x}_t .

For a statistically homogeneous medium, the ensemble-average transmittance becomes:

$$\text{ff}(t) = \langle \text{Tr}_\mu(t) \rangle = \left\langle e^{-t \bar{\sigma}_\mu(\mathbf{x}, \mathbf{x}_t)} \right\rangle \approx \frac{1}{N} \sum_{i=1}^N e^{-t \bar{\sigma}_{\mu_i}(\mathbf{x}, \mathbf{x}_t)}. \quad (28)$$

The Monte Carlo estimate is averaging transmittance through $\bar{\sigma}_\mu(t) = \bar{\sigma}_\mu(\mathbf{x}, \mathbf{x}_t)$ for a fixed t , but over random realizations μ_i of the medium. Figure 7 shows random 1D transects $\sigma_\mu(t)$ (left) in a medium where

$\sigma_\mu(\mathbf{x})$ is modeled using a single octave of Perlin noise [Perlin 1985], and transects (middle) of the corresponding “blurred” micro-scale density $\bar{\sigma}_\mu(t)$.

While computing ensemble-averaged transmittance this way avoids the complexities of instantiating billions of discrete particles, it still makes parametric control cumbersome due to the need for tabulation.

4.3 Probabilistic ensemble-averaging

When performing ensemble averaging, $\bar{\sigma}_\mu$ (or equivalently τ_μ) acts as a random variable with some probability density pdf($\bar{\sigma}_\mu | t$) describing its variability for each distance t . We visualize this as histograms on the right-hand side of Figure 7.

Longer-than-exponential tails. This probabilistic view allows us to explain why in Figure 2 the ensemble-averaged transmittance (red, pink) results in a longer tail than the exponential (black). Jensen’s inequality states that for a random variable X and a concave function f : $\langle f(X) \rangle \geq f(\langle X \rangle)$. Substituting the exponential for f , and $\bar{\sigma}_\mu(t)$ for X , we should therefore expect:

$$\left\langle e^{-t \bar{\sigma}_\mu(t)} \right\rangle \geq e^{-t \langle \bar{\sigma}_\mu(t) \rangle}. \quad (29)$$

This will be an equality iff $\bar{\sigma}_\mu(t) = \langle \bar{\sigma}_\mu(t) \rangle$ for all t , which would mean the medium had no density fluctuations (was homogeneous) to begin with.

Link to characteristic functions. We can now write the ensemble average as the integral:

$$\text{ff}(t) = \left\langle e^{-t \bar{\sigma}_\mu(t)} \right\rangle = \int_0^\infty e^{-t \bar{\sigma}_\mu} \text{pdf}(\bar{\sigma}_\mu | t) d\bar{\sigma}_\mu. \quad (30)$$

This takes a form remarkably similar to the *characteristic function* (CF), which is the Fourier transform of a random variable’s PDF:

$$\varphi_X(r) = \left\langle e^{irX} \right\rangle = \int_{\mathbb{R}} e^{irx} \text{pdf}_X(x) dx, \quad (31)$$




Fig. 7. Left: 1D transects of a Perlin noise micro-scale density function $\sigma_\mu(t)$ for random rays and realizations, a single sample of which is highlighted in black. σ_μ is a random variable, the distribution of which, $\text{pdf}(\sigma_\mu)$ is illustrated by the adjacent back histogram. Middle: The line-averaged micro-scale density $\bar{\sigma}_\mu(t)$ is likewise a random variable, but the distribution $\text{pdf}(\bar{\sigma}_\mu(t))$ now potentially depends on the length of the line segment t as shown by the histograms for $t = 2.5$ (red) and $t = 22.5$ (green). Right: When the micro-scale density $\sigma_\mu(t)$ is modeled as $1/f^\beta$ fractal noise, the histograms of line-averaged micro-scale density $\bar{\sigma}_\mu(t)$ (thin transparent lines) take a closed form where the spectral exponent β directly controls the variance $v_{\bar{\sigma}_\mu(t)}$ (thick opaque lines) as a function of distance t , where lower values of β lead to more rapid decay.

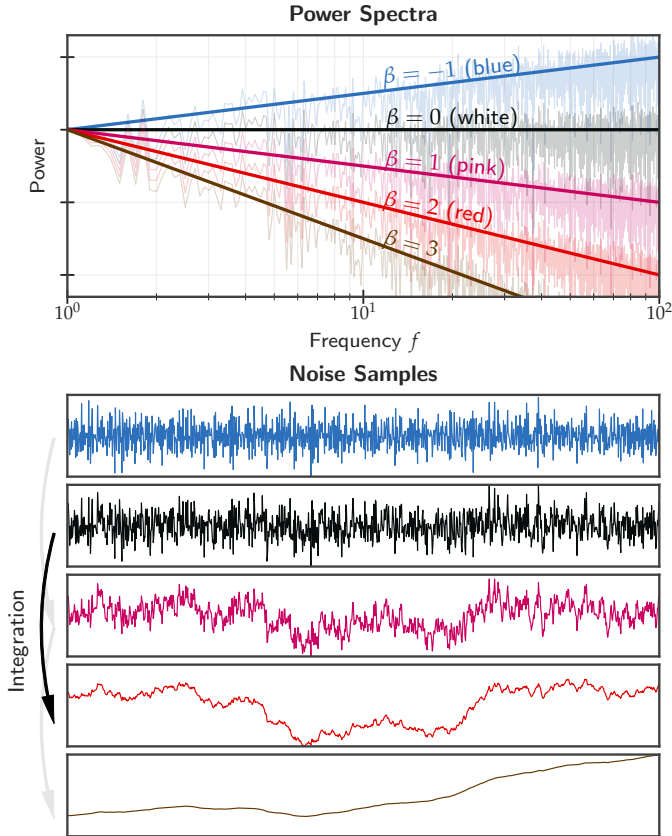


Fig. 8. $1/f^\beta$ noise (bottom) is characterized by a spectral exponent β which determines the slope of its power spectrum on a log-log plot (top). Qualitatively, different values of β lead to noises of different “roughnesses” (bottom), and integrating a noise with spectral exponent β (e.g. white noise) produces another noise with spectral exponent $\beta + 2$ (e.g. red noise).

where $i = \sqrt{-1}$ is the imaginary constant. Comparing Equation (31) to (30), we see that for a fixed t the ensemble-averaged transmittance

is simply the CF of the random variable $\bar{\sigma}_\mu$ (by passing in ti for r):²

$$\text{ff}(t) = \varphi_{\bar{\sigma}_\mu(t)}(ti). \quad (32)$$

This is a very powerful tool, because, as long as we can express the variation of $\bar{\sigma}_\mu$ (or τ_μ) using a distribution with a well-known CF (see Table 1), we can obtain closed-form expressions for the ensemble-averaged transmittance! Note that the parameters of the distribution $\text{pdf}(\bar{\sigma}_\mu | t)$ may depend on the distance t , as we saw in Figure 7.

Discussion. So what should this distribution, and its dependence on t be? We could simply choose a convenient PDF and allow an artist to set its (potentially t -dependent) parameters by hand. While this would be fully parametric and analytic, it is unclear how (or whether) such a model would correspond to any micro-scale density fluctuations or correlations. Alternatively, we could ensure correspondence by fitting the distribution and its t -dependent parameters to tabulated data from realizations of $\bar{\sigma}_\mu(t)$, like in Figure 7. Next, we will instead leverage recent work by Davis and colleagues [2011; 2014] to obtain a parametric, closed-form transmittance function whose parameters provide a physical interpretation of the underlying micro-scale extinction field as a fractal medium.

4.4 Closed-form average transmittance in fractal media

We will follow Davis and Mineev-Weinstein [2011] and model the variability in $\sigma_\mu(\mathbf{x})$ as a 3D fractal with 1D transects $\sigma_\mu(x)$ characterized by $1/f^\beta$ fractal noise. Figure 8 provides a visual explanation: Qualitatively, β determines how “rough” the noise will be (bottom), while quantitatively it dictates that the falloff of the function’s power spectrum will be $1/f^\beta$ (top). This directly controls the spatial correlations, where $\beta = 0$ means uncorrelated “white” noise, while $\beta > 0$ and $\beta < 0$ produce positive and negative correlations respectively.

We model $\sigma_\mu(x)$ statistically as a noise defined by: a mean μ_σ , a constant C controlling its overall amplitude, and its spectral exponent $-1 \leq \beta \leq 1$ ranging from “blue” ($\beta = -1$) to “pink” ($\beta = 1$). For convenience, we will combine these into a parameter vector

²We could also express this using *moment generating functions* (MGFs), which are defined similarly to the CF, but without the imaginary constant.

$\psi = \{\mu_\sigma, C, \beta\}$. This allows us to write the ensemble-averaged transmittance (30) as:

$$\text{ff}(t | \psi) = \int_0^\infty e^{-t\bar{\sigma}_\mu} \text{pdf}(\bar{\sigma}_\mu | t, \psi) d\bar{\sigma}_\mu, \quad (33)$$

where the PDF is now determined by the distance t and parameters $\psi = \{\mu_\sigma, C, \beta\}$ defining the medium.

Given this fractal noise model, what can we say about $\text{pdf}(\bar{\sigma}_\mu | t, \psi)$?

Gaussian $1/f^\beta$ noise. A key property of Gaussian $1/f^\beta$ noise is that it produces fractional Brownian motion (fBm) [Mandelbrot and Ness 1968] with $\beta + 2$ via integration (arrows in Figure 8).³ We therefore know that $\tau_\mu(t)$ will be a noise with $1 \leq \beta_\tau \leq 3$ and its distribution (as well as that of $\bar{\sigma}_\mu$) will be a Gaussian. Davis and Mineev-Weinstein [2011] formally derived the dependence of this Gaussian on the distance t and medium parameters ψ :

$$\text{pdf}(\bar{\sigma}_\mu | t, \psi) = \mathcal{N}(\mu_\sigma, v_{\bar{\sigma}_\mu}(t)), \quad \text{with } v_{\bar{\sigma}_\mu}(t) = (C\mu_\sigma)^{\beta+1} t^{\beta-1}. \quad (34)$$

Figure 7 visualizes random transects of $\bar{\sigma}_\mu(t)$: the medium σ_μ averaged over a distance t . This figure shows the same information as the second row of Figure 7, but this time for media modeled as fractal noise with $\beta = -1, 0, 1$. The superimposed analytic curves plot the standard deviation $\sqrt{v_{\bar{\sigma}_\mu}(t)}$ from Equation (34) as a function of t , matching the behavior of the random transects well.

Using the characteristic function of the normal distribution (Table 1) with these parameters, we can obtain a closed-form expression for the ensemble-averaged transmittance via Equation (32):

$$\text{ff}_\mathcal{N}(t | \psi) \approx e^{-\mu_\sigma t + v_{\bar{\sigma}_\mu}(t)t^2/2} = e^{-\mu_\sigma t + (C\mu_\sigma t)^{\beta+1}/2} \quad (35)$$

Unfortunately, this is only an approximation because the Gaussian is supported on the entire real line, but values of $\sigma_\mu < 0$ are non-physical. This model for noise is therefore only reasonable when μ_σ is set sufficiently high and v set sufficiently low so that negative extinction coefficients are unlikely to occur.

Gamma-distributed $1/f^\beta$ noise. To counteract the artifacts arising due to these negative intrusions, Davis and Mineev-Weinstein [2011] proposed modeling the distribution of extinction values σ_μ with the strictly non-negative Gamma distribution $\text{pdf}(\bar{\sigma}_\mu | t, \psi) = \Gamma(\mu_\sigma, \alpha(t))$, with parameters set to match the mean and variance of the Gaussian model. Combining Equation (32) with Table 1 gives:

$$\text{ff}_\Gamma(t | \psi) = \varphi_{\Gamma, \psi}(ti) = \left(1 + \frac{\mu_\sigma t}{\alpha(t)}\right)^{-\alpha(t)}, \quad (36)$$

$$\text{where } \alpha(t) = \frac{\mu_\sigma^2}{v_{\bar{\sigma}_\mu}(t)} = \frac{(\mu_\sigma t)^{1-\beta}}{C^{1+\beta}} \quad (37)$$

is the Gamma model's shape parameter that enforces a variance dictated by Equation (34).

³ β is directly related to the Hurst parameter $H = (\beta - 1)/2$ more commonly use in the fractal literature [Barnsley et al. 1988] when describing fBm and the "persistence" parameter of fractal Perlin noise [Ebert et al. 2002] used in graphics.

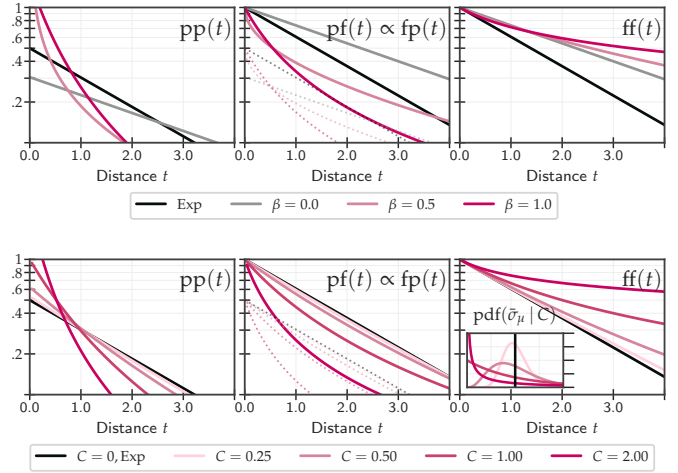


Fig. 9. Ensemble-averaged transmittances (ff, pf) and free-flight PDFs (pp, fp) arising from Gamma-distributed fractal $1/f^\beta$ noise (36) with $\bar{\sigma} = 1/2$. Top: for a fixed $C = 3/2$ and different values of β ; Bottom: for a fixed $\beta = 1$ and different values of C . Since $\text{pf} \propto \text{fp}$, we only show fp in faded dashed lines in the middle plots. The inset on the bottom right shows the Gamma distribution controlling $\text{pdf}(\bar{\sigma}_\mu)$ for pink noise.

Discussion. Equation (36) provides a simple 3-parameter model to produce non-exponential transmittance functions by specifying the mean extinction μ_σ , the overall amplitude of variation C , and the color/spectral falloff/correlation of the noise β . Figure 9 visualizes the four resulting transport functions for fixed C and varying β (top), and for fixed $\beta = 1$ pink noise and varying C (bottom).

As we would hope, both the Gaussian (35) and Gamma (36) models reduce to a simple exponential when variance $v \rightarrow 0$ (and hence $\alpha \rightarrow \infty$). This will happen if the medium is actually homogeneous ($C = 0$), or if we have white noise ($\beta = 0$), both of which correspond to independent scatterers. In the latter case of white noise, we obtain exponential falloff, but with a modified extinction coefficient.

"Pink" $1/f$ noise (when $\beta = 1$) is another interesting case because it has a "scale-invariant" property where the distribution $\text{pdf}(\bar{\sigma}_\mu)$ in Equation (35) and (36) no longer depends on t (see Figure 7). This reduces the transmittance completely to the CF of $\text{pdf}(\bar{\sigma}_\mu)$:

$$\text{ff}_{\text{pink}}(t) = \int_0^\infty e^{-t\bar{\sigma}_\mu} \text{pdf}(\bar{\sigma}_\mu) d\bar{\sigma}_\mu = \varphi_{\bar{\sigma}_\mu}(ti). \quad (38)$$

With the Gamma model, for instance, the variance that plugs into Equation (36) reduces to $v_{\bar{\sigma}_\mu} = C^2 \mu_\sigma^2$ when $\beta = 1$, and so $\text{pdf}(\bar{\sigma}_\mu) = \Gamma(\mu_\sigma, \alpha)$ where $\alpha = C^{-\beta-1}$ no longer depends on t .

Equation (38) shows that ensemble averaging transmittance over a heterogeneous medium can be equivalently interpreted as averaging transmittance across realizations that are each *homogeneous*, but with a random extinction coefficient drawn from the distribution $\text{pdf}(\bar{\sigma}_\mu)$. This also gives a physical interpretation to other choices of non-exponential functions which are themselves characteristic functions of some distribution (e.g. the sum of exponentials [Wrenninge et al. 2013] is the CF of a sum of two deltas $\text{III}(2)$): these all




Fig. 10. In contrast to the deterministic case (Figure 5), the statistical behavior of rescaled fractal noise depends on its power spectrum via Equation (40). For instance, the standard deviation of blue noise ($\beta = -1$) does not depend on σ , while for pink noise ($\beta = 1$) it is proportional to σ . The colored intervals visualize how the standard deviation is affected by $\sigma(s)$ for $\beta = -1, 0, 1$.

correspond to scale-invariant pink noise, but where $\text{pdf}(\bar{\sigma}_\mu)$ dictates the distribution of blurred extinction values the fractal takes on.

4.5 Combining macro- and micro- properties

We previously introduced our heterogeneity model (Section 3.2) and gave a physical interpretation of it in terms of “squeezing” and “stretching” the micro-density function by the macroscopic optical depth. However, with ensemble averaged transmittance from fractal media, we can only reason about the micro-density function in a statistical sense, and this physical interpretation changes slightly. Firstly, due to the self-similar nature of $1/f^\beta$ fractal noise, horizontal squeezing and stretching has no effect on the wavelength content and can be ignored. Instead, we can see from Equations (35,36) that scaling the distance t by some constant c is equivalent to scaling the mean μ_σ by the same constant:

$$\text{ff}(ct \mid \{\mu_\sigma, C, \beta\}) = \text{ff}(t \mid \{c\mu_\sigma, C, \beta\}). \quad (39)$$

This means that inserting the macroscopic optical depth adjusts the mean μ_σ of the micro-scale medium to locally match the density σ at the macro-scale. Changing just the mean μ_σ of the noise, however, changes its standard deviation in a β -dependent way since:

$$v_{\bar{\sigma}_\mu}(t \mid \{c\mu_\sigma, C, \beta\}) = c^{\beta+1} \cdot v_{\bar{\sigma}_\mu}(t \mid \{\mu_\sigma, C, \beta\}). \quad (40)$$

This suggests that the way the density match is achieved has a different interpretation based on the color of the noise (Figure 10). As in the deterministic case (Figure 5), scaling the distance by a factor of c in “pink” noise ($\beta = 1$) corresponds to *scaling* the micro-fluctuations vertically by the constant c , but as β decreases the standard deviation is scaled less, until at “blue” noise ($\beta = -1$) it is not scaled at all ($c^0 = 1$), suggesting a vertical *shift/translation* of the noise instead.

5 IMPLEMENTATION

In this section, we give a brief outline of the modifications that need to be made to an existing rendering algorithm to support our path integral. We outline the procedure for evaluation and sampling of \tilde{T} in Algorithm 1. We assume that we are given functions ff , pp and pf computed in a base medium (i.e. with macroscopic extinction of 1), as well as procedures $\text{sampleFP}(\xi)$ and $\text{samplePP}(\xi)$ to produce samples distributed proportional to fp and pp , respectively. We also require functions to evaluate the PDF and CDF for a given sample.



Fig. 11. A heterogeneous medium rendered with regular tracking (left) and our unbiased delta tracking (right) for a Gamma distributed pink noise medium. Both methods are unbiased and have comparable noise levels, but delta tracking renders 23× faster.

Our sampling algorithm proceeds by generating a sample from fp or pp , depending on whether \mathbf{x}_i lies on a surface or not. Due to our heterogeneity model, we can interpret the generated sample as an optical depth τ^* in the macroscopic medium and solve for the distance s along the ray such that $\tau(\mathbf{x}_i, \mathbf{x}_i + s \cdot \boldsymbol{\omega}) = \tau^*$. We do this analytically in homogeneous media, or numerically using regular tracking/raymarching in heterogeneous media.

The probability (density) of the generated vertex is obtained from the sample CDF (PDF) if s lies beyond (before) the distance to the nearest surface. The sample weight is then simply \tilde{T}/pdf . If fp and pp are sampled perfectly (e.g. using formulas from Table 1), then this routine is an optimal distance sampling strategy.

5.1 Unbiased Delta Tracking

The sampling method in Algorithm 1 relies on regular tracking (slow) or raymarching (biased) to sample distances. Methods based on delta tracking [Coleman 1968; Novák et al. 2018; Raab et al. 2008; Woodcock et al. 1965] offer potentially better performance while remaining unbiased, but unfortunately rely on transmittance being an exponential (namely $\text{Tr}(a+b) = \text{Tr}(a)\text{Tr}(b)$). However, our heterogeneity model allows us to use delta tracking for a subclass of non-exponential media.

For stochastic media characterized by $1/f$ (“pink”) fractal noise, Equation (38) shows that the macroscopic transmittance is equivalent to ensemble averaging random *homogeneous* realizations. This means that an unbiased estimation of transmittance is possible by simply averaging scaled exponentials:

$$\text{ff}_{\text{pink}}(t) = \int_0^\infty e^{-t\bar{\sigma}_\mu} \text{pdf}(\bar{\sigma}_\mu) d\bar{\sigma}_\mu \approx \frac{1}{N} \sum_i^N e^{-t\bar{\sigma}_{\mu_i}}, \quad (41)$$

where the scaling factors $\bar{\sigma}_{\mu_i}$ are drawn $\propto \text{pdf}(\bar{\sigma}_\mu)$.

Likewise, since our heterogeneity model inserts the macroscopic optical depth into the base transmittance, for a pink noise medium, we can globally multiply the heterogeneous medium by a constant factor and compute the classical transmittance, averaged over the distribution of scaling factors $\text{pdf}(\bar{\sigma}_\mu)$. We can trivially turn this into a probabilistic process, in which we sample a scale factor $\propto \text{pdf}(\bar{\sigma}_\mu)$, multiply the macroscopic density by this factor and compute the classical transmittance (or free-flight distance) through the resulting medium. This allows us to use any existing unbiased distance sampling algorithm, while obtaining a non-exponential result.

In Figure 11, we demonstrate a cloud with Gamma distributed pink noise ($\text{pdf}(\bar{\sigma}_\mu) \propto \Gamma(\alpha, \alpha)$, $\text{ff}(t) = (1 + t/\alpha)^{-\alpha}$ [Davis and Xu

Algorithm 1: Transport kernel evaluation and sampling for our method

```

1 function  $\tilde{T}_{\text{base}}(\tau, \text{surfaceStart}, \text{surfaceEnd})$ 
2   if  $\text{surfaceStart} \wedge \text{surfaceEnd}$  then return  $\text{ff}(\tau)$ 
3   else if  $\neg \text{surfaceStart} \wedge \neg \text{surfaceEnd}$  then return  $\text{pp}(\tau)$ 
4   else return  $\text{pf}(\tau)$ 
5 function  $\text{sample}(\mathbf{x}_i, \omega, \xi)$ 
6    $\text{surfaceStart} \leftarrow \mathbf{x}_i \in \partial \mathcal{V}$ 
7    $\tau^* \leftarrow \text{surfaceStart} ? \text{sampleFP}(\xi) : \text{samplePP}(\xi)$ 
8    $s \leftarrow \text{invertTau}(\tau^*)$ 
9    $z \leftarrow \text{raytrace}(\mathbf{x}_i, \omega)$ 
10   $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i + \min(s, z) \cdot \omega$ 
11   $\text{surfaceEnd} \leftarrow s < z$ 
12  if  $\text{surfaceEnd}$  then
13     $\text{pdf} = 1 - (\text{surfaceStart} ? \text{sampleFP}_{\text{CDF}}(\tau^*) : \text{samplePP}_{\text{CDF}}(\tau^*))$ 
14  else
15     $\text{pdf} = \sigma(\mathbf{x}_{i+1}) \cdot (\text{surfaceStart} ? \text{sampleFP}_{\text{PDF}}(\tau^*) : \text{samplePP}_{\text{PDF}}(\tau^*))$ 
16   $\text{weight} = \tilde{T}_{\text{base}}(\tau^*, \text{surfaceStart}, \text{surfaceEnd}) / \text{pdf}$ 
17  return  $\{\text{weight}, \mathbf{x}_{i+1}\}$ 

```

2014]) rendered with regular tracking as in Algorithm 1 as well as classical delta tracking, for which the medium was scaled by a scaling factor sampled from a Gamma distribution at every bounce. Delta tracking converges to the same result as regular tracking, but 23× faster.

6 RESULTS

We implemented our theory in two existing rendering systems, PBRTv3 [Pharr et al. 2016] and Tungsten [Bitterli 2018], following the outline in Section 5. We will release the code of both of our implementations.

In Figure 6, we show a homogeneous medium rendered with five different designed transmittance curves and compare the results to an exponential medium. The non-exponential transmittances are simple parametric functions, such as linear or quadratic curves, but lead to a wide range of appearances that cannot be reproduced by the exponential. These curves were designed artificially and do not have a physical interpretation in terms of correlated scatterers, but their simplicity and ease of control makes them powerful tools for artistic control of the medium.

Figure 12 shows a similar homogeneous medium with a linearly decreasing, non-exponential transmittance. The same scene is rendered with three different rendering algorithms—path tracing, light



Fig. 12. We render the same scene with different rendering algorithms and obtain identical results regardless of transport direction. This supports the fact that our method is reciprocal.




Fig. 13. We verify our model and compare it to ground truth in a Monte Carlo experiment in a semi-infinite medium. Please see Section 6 for details.

tracing and bidirectional path tracing—that were modified to support our theory. All three algorithms produce identical results, which supports the fact that our theory is reciprocal.

Figure 1 shows a heterogeneous cloud, rendered with a traditional exponential transmittance, and a long-tailed non-exponential transmittance derived from the Davis and Mineev-Weinstein model. The long tail of the non-exponential transmittance allows light to penetrate deeper into the cloud, which leads to a brighter appearance near the bottom and a softer look near the surface, giving an overall richer appearance.

In Figure 14, we compare different non-exponential transmittances to an exponential reference in a scene with subsurface scattering. The mean free path of the non-exponential free-flight PDF was matched to that of the exponential; however, despite having the same MFP, the non-exponential transmittance is able to achieve a wide range of appearances.

Finally, in Figure 15 we compare a heterogeneous cloud rendered using a transmittance derived from fractal noise (bottom two rows) using the Davis and Mineev-Weinstein model with varying parameters (β , C). The fractal noise model allows a range of different appearances and mean free paths. For comparison, we also provide the same cloud rendered using an exponential with varying mean free paths (top row).

We include a comprehensive supplemental material with interactive image viewer, full-resolution HDR images of all our renderings, additional derivations as well as additional renderings of our scenes with the exponential model with different mean free paths for reference. We encourage the reader to refer to our supplemental material for a more careful comparison.

Monte Carlo Simulation. We additionally implemented our model in a simple Monte Carlo simulation (Figure 13) for verification. We created a semi-infinite stochastic “red” noise medium and measured the mean penetration depth (after a fixed number of bounces) of photons emitted from a pencil beam at normal incidence. We compare the ground truth ensemble average against results of our path integral (“Ours”), the classical RTE and the model of Larsen and Vasques [2011] (“GBE”) computed on a homogenized version of this medium. We additionally compare the ensemble average of a modified RTE in which photons “jump” to a different realization of the medium at every scattering event, conditioned on coinciding with a scatterer in the new realization (“Ours (model)”). The curves of our path integral match that of this modified RTE exactly, which




Fig. 14. We compare a dense homogeneous medium rendered with different non-exponential transmittances to an exponential reference. The mean free path of the non-exponential free-flight PDF was matched to that of the exponential. Despite matching the MFP, we are able to achieve a wide range of appearances.

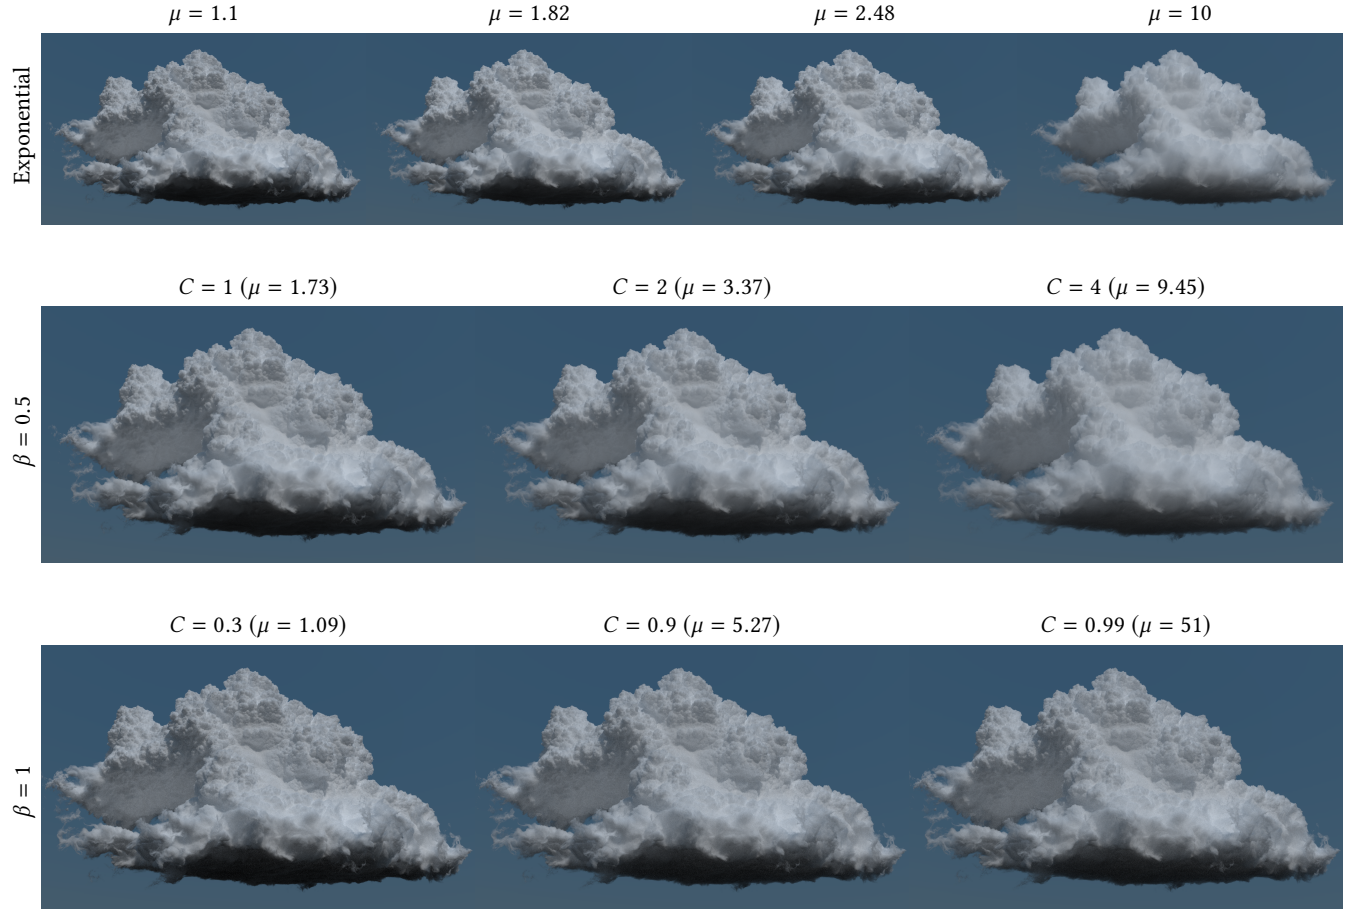


Fig. 15. We compare a heterogeneous cloud rendered using a transmittance derived from fractal noise (bottom two rows) using the Davis and Mineev-Weinstein model with varying parameters (β , C). The fractal noise model allows a range of different appearances and mean free paths. For comparison, we also provide the same cloud rendered using an exponential with varying mean free paths (top row).

verifies that our path integral computes precisely the ensemble average of this physical model. Additionally, the curves of our method are closer to ground truth than the classical exponential model or the model of Larsen and Vasques, which demonstrates that more sophisticated models of correlated scatterers and surfaces lead to more accurate results.

Although our method leads to more accurate results than previous work, we still observe some discrepancies between our model and ground truth. To better understand these differences, we implemented a modified version of our method that uses a free-flight PDF conditioned on the free-flight distance of the previous segment on the path (“Ours (two-bounce)”). Unlike the other methods, this model does not reset correlations at every interaction, but correlates free-flights of adjacent path segments. The disadvantage of this approach is that currently this conditioned free-flight PDF cannot be obtained analytically, and it must be tabulated in an expensive pre-process instead. Although not a practical approach today, this two-segment correlated method is closer to ground truth than our base model, which suggests that modelling correlations between more than two vertices can further improve accuracy.

7 CONCLUSION

We introduced a new theory of volumetric light transport that allows for media with non-exponential free-flight distributions. Such distributions are the consequence of correlations between scatterers, which can arise from physical processes in e.g. the atmospheric sciences and neutron transport. We derived a new non-exponential path integral that approximates the general problem of ensemble averaged transport in stochastic media and supports surfaces and heterogeneity in a combined framework, while being reciprocal and practical to compute. Both our surface and heterogeneity models have simple physical interpretations and are easy to retrofit into existing volumetric rendering workflows.

We presented a wide range of tools to leverage this newfound flexibility, ranging from simple parametric transmittance curves to powerful mathematical formalisms for Gaussian processes and fractal noise. The latter allowed us to obtain closed-form, ensemble averaged transmittances for scatterers distributed with different colors of noise.

Our theory can be implemented with only minor changes to existing rendering algorithms, and our results demonstrate that this allows for a rich range of volumetric appearances while still allowing for fully bidirectional transport. For the special case of pink noise media, we even adapted algorithms for unbiased distance sampling, which opens the door to efficient and unbiased rendering of non-exponential heterogeneous media.

7.1 Limitations and Future Work

There are several limitations to our work that leave ample grounds for future research.

Our theory currently assumes that there is a single, globally fixed free-flight PDF. This means that we cannot handle multiple overlapping media with different transmittance functions, even though this could be a useful tool in practice. Although we can take advantage of unbiased distance sampling methods for certain media,




Fig. 16. While Davis and Mineev-Weinstein’s [2011] 2-parameter Gamma-distributed $1/f^\beta$ noise model allows changing the spectral exponent from $\beta = 1$ (pink) down to $\beta = -1$ (blue) noise, this leads to non-physical free-flight PDFs (pp and fp) which take on negative values and which cannot be properly normalized. We therefore only use their model for $0 \leq \beta \leq 1$.

it is unclear how to do this in the general case. More research is required to adapt efficient, unbiased distance sampling methods for all non-exponential heterogeneous media to our theory.

While our theory leads to a richer volumetric appearance space, navigating this space is still a challenge. Simply maintaining the same single-scattering mean-free path does not necessarily provide the best perceptual appearance match between a chosen non-exponential model and classical exponential transport. In fact, there are two distinct mean-free paths in our framework (one for pp and one for fp), so this match is not even unique. Our experience suggests that the classical parameters (e.g. mean free-path) and the parameters of the correlations are not orthogonal, so adjusting them individually to reach a desired appearance is cumbersome. Reparametrizing the appearance space into perceptually uniform and orthogonal appearance dimensions, as has previously been investigated for color [Smith and Guild 1931], glossy materials [Wills et al. 2009], and translucency [Gkioulekas et al. 2013; Zhao et al. 2014], would greatly facilitate future adoption of our model.

Another interesting avenue for future research is the synergy between non-exponential transmittance and diffusion theory. In Section 4.1, we took the classical diffusion monopole and interpreted it as a free-flight distribution. Extending this concept to more sophisticated diffusion models [d’Eon 2013] could allow for efficient, approximate multi-bounce rendering. Matching the resulting diffusion mean-free path may also provide a better starting point for a perceptually motivated appearance space.

While Davis and Mineev-Weinstein’s [2011] 2-parameter fractal $1/f^\beta$ noise model theoretically allows changing the spectral exponent from $\beta = 1$ (pink) down to $\beta = -1$ (blue) noise, we found that, in our framework considering both media and surfaces, it produces non-physical free-flight PDFs and transmittances which go negative when $\beta \leq 0$ (see Figure 16). We suspect this arises because Davis and Mineev-Weinstein operate exclusively on ff and make simplifying assumptions without enforcing the necessary constraints on ff’s first and second derivatives to make our other three transport functions physically plausible. A more detailed investigation is warranted, but since $\beta \leq 0$ produces approximately exponential behavior, we have not found this reduction in the parameter range to be a concern.

Somewhat surprisingly, the observed transport behavior in a participating medium can be non-exponential even for classical

media where the scatterers are independently distributed. Since the weighted average of exponentials is not an exponential, the spectral averaging at the sensor will lead to non-exponential appearance if the medium properties vary across the sensor's spectral response. The ability of our theory to model this non-exponentiality could benefit inverse problems in computer vision and appearance capture, which rely on accurate forward models.

ACKNOWLEDGMENTS

We thank Adrian Jarabo, Carlos Aliaga and Diego Gutierrez for helpful discussions, and the reviewers for their insightful feedback. This work was partially supported by a grant from Microsoft Azure, as well as the National Science Foundation (Grant CNS-1205521).

REFERENCES

- Michael Ashikhmin and Peter Shirley. 2000. An Anisotropic Phong BRDF Model. *J. Graph. Tools* 5, 2 (Feb. 2000), 25–32. <https://doi.org/10.1080/10867651.2000.10487522>
- Michael Ashikhmin, Simon Premoze, and Peter Shirley. 2000. A Microfacet-based BRDF Generator. In *Proc. SIGGRAPH*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 65–74. <https://doi.org/10.1145/344779.344814>
- Mahdi M. Bagher, John Snyder, and Derek Nowrouzezahrai. 2016. A Non-Parametric Factor Microfacet Model for Isotropic BRDFs. *ACM Trans. Graph. (Proc. SIGGRAPH)* 35, 5 (July 2016), 159:1–159:16. <https://doi.org/10.1145/2907941>
- Michael F. Barnsley, Robert L. Devaney, Benoit B. Mandelbrot, Heinz-Otto Peitgen, Dietmar Saupe, Richard F. Voss, Yuval Fisher, and Michael McGuire. 1988. *The Science of Fractal Images* (1st ed.). Springer Publishing Company, Incorporated. <https://doi.org/10.1007/978-1-4612-3784-6>
- Bryan Beresford-Smith, Derek YC Chan, and D John Mitchell. 1985. The electrostatic interaction in colloidal systems with low added electrolyte. *Journal of colloid and interface science* 105, 1 (1985), 216–234.
- Benedikt Bitterli. 2018. Tungsten Renderer. (2018). <https://github.com/tunabrain/tungsten/>.
- James F. Blinn. 1977. Models of Light Reflection for Computer Synthesized Pictures. *Proc. SIGGRAPH* 11, 2 (July 1977), 192–198. <https://doi.org/10.1145/965141.563893>
- Anatoli Borovoi. 2002. On the extinction of radiation by a homogeneous but spatially correlated random medium: comment. *J. Opt. Soc. Am. A* 19, 12 (Dec 2002), 2517–2520. <https://doi.org/10.1364/JOSAA.19.002517>
- Antoine Bouthors, Fabrice Neyret, Nelson Max, Eric Bruneton, and Cyril Crassin. 2008. Interactive multiple anisotropic scattering in clouds. In *Proceedings of the Symposium on Interactive 3D Graphics and Games*. ACM, ACM, 173–182. <https://doi.org/10.1145/1342250.1342277>
- Thomas Camminady, Martin Frank, and Edward W. Larsen. 2017. Nonclassical Particle Transport in Heterogeneous Materials. In *International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering*.
- S. Chandrasekhar. 1960. *Radiative Transfer*. Dover Publications.
- Per H. Christensen and Wojciech Jarosz. 2016. The Path to Path-Traced Movies. *Foundations and Trends in Computer Graphics and Vision* 10, 2 (October 2016), 103–175. <https://doi.org/10.1561/06000000073>
- W. A. Coleman. 1968. Mathematical verification of a certain Monte Carlo sampling technique and applications of the technique to radiation transport problems. *Nuclear Science and Engineering* 32, 1 (April 1968), 76–81. <https://doi.org/10.13182/NSE68-1>
- Robert L. Cook and Kenneth E. Torrance. 1981. A Reflectance Model for Computer Graphics. *Proc. SIGGRAPH* 15, 3 (Aug. 1981), 307–316. <https://doi.org/10.1145/965161.806819>
- Anthony B. Davis and Alexander Marshak. 2004. Photon propagation in heterogeneous optical media with spatial correlations: enhanced mean-free-paths and wider-than-exponential free-path distributions. *Journal of Quantitative Spectroscopy and Radiative Transfer* 84, 1 (2004), 3–34. [https://doi.org/10.1016/S0022-4073\(03\)00114-6](https://doi.org/10.1016/S0022-4073(03)00114-6)
- Anthony B. Davis, Alexander Marshak, H. Gerber, and Warren J. Wiscombe. 1999. Horizontal structure of marine boundary layer clouds from centimeter to kilometer scales. *Journal of Geophysical Research: Atmospheres* 104, D6 (1999), 6123–6144. <https://doi.org/10.1029/1998JD200078>
- Anthony B. Davis and Mark B. Mineev-Weinstein. 2011. Radiation propagation in random media: From positive to negative correlations in high-frequency fluctuations. *Journal of Quantitative Spectroscopy and Radiative Transfer* 112, 4 (March 2011), 632–645. <https://doi.org/10.1016/j.jqsrt.2010.10.001>
- Anthony B. Davis and Feng Xu. 2014. A Generalized Linear Transport Model for Spatially Correlated Stochastic Media. *Journal of Computational and Theoretical Transport* 43, 1–7 (2014), 474–514. <https://doi.org/10.1080/23324309.2014.978083>
- Eugene d'Eon. 2013. Rigorous asymptotic and moment-preserving diffusion approximations for generalized linear Boltzmann transport in d dimensions. *CoRR* abs/1312.1412 (2013). <http://arxiv.org/abs/1312.1412>
- Eugene d'Eon. 2016. *A Hitchhiker's Guide to Multiple Scattering*. Self-published. <http://www.eugenedeon.com/project/a-hitchhikers-guide-to-multiple-scattering/>
- Eugene d'Eon. 2018. A reciprocal formulation of non-exponential radiative transfer. 1: Sketch and motivation. (2018). [arXiv:1803.03259](https://arxiv.org/abs/1803.03259) <https://arxiv.org/abs/1803.03259>
- Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress Towards the Unification of Microfacet and Microflake Theories. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations (EGSR '16)*. Eurographics Association, Goslar Germany, Germany, 55–63. <https://doi.org/10.2312/sre.20161210>
- David S. Ebert, F. Kenton Musgrave, Darwyn Peachey, Ken Perlin, and Steven Worley. 2002. *Texturing and Modeling: A Procedural Approach* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Julian Fong, Magnus Wrenninge, Christopher Kulla, and Ralf Habel. 2017. Production Volume Rendering: SIGGRAPH 2017 Course. In *ACM SIGGRAPH 2017 Courses (SIGGRAPH '17)*. ACM, New York, NY, USA, Article 2, 79 pages. <https://doi.org/10.1145/3084873.3084907>
- Ioannis Gkioulekas, Bei Xiao, Shuang Zhao, Edward Adelson, Todd Zickler, and Kavita Bala. 2013. Understanding the Role of Phase Function in Translucent Appearance. *ACM Trans. Graph.* 32, 5, Article 131 (Oct. 2013), 12 pages.
- Armando Gama Goicochea. 2013. A Model for the Stability of a TiO₂ Dispersion. *ISRN Materials Science* 2013 (2013).
- David G Grier and Sven H Behrens. 2001. Interactions in colloidal suspensions. In *Electrostatic effects in soft matter and biophysics*. Springer, 87–116.
- Diego Gutierrez, Francisco Seron, Adolfo Muñoz, and Oscar Anson. 2008. Visualizing Underwater Ocean Optics. *Computer Graphics Forum (Proc. Eurographics)* 27, 2 (2008), 547–556.
- Ralf Habel, Per H. Christensen, and Wojciech Jarosz. 2013. *Classical and Improved Diffusion Theory for Subsurface Scattering*. Technical Report. Disney Research Zürich.
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX Microflake Distribution. *ACM Trans. Graph.* 34, 4 (July 2015), 48:1–48:11.
- Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A Radiative Transfer Framework for Rendering Materials with Anisotropic Structure. *ACM Trans. Graph.* 29, 4, Article 53 (July 2010), 13 pages. <https://doi.org/10.1145/1778765.1778790>
- Safa Jamali. 2015. *Rheology of Colloidal Suspensions: A Computational Study*. Ph.D. Dissertation.
- Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A Radiative Transfer Framework for Spatially-Correlated Materials. *ACM Trans. Graph. (Proc. SIGGRAPH)* 37, 4 (July 2018).
- Adrian Jarabo and Victor Arellano. 2018. Bidirectional Rendering of Vector Light Transport. *Computer Graphics Forum* (2018), n/a–n/a. <https://doi.org/10.1111/cgf.13314> awaiting publication.
- Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A Practical Model for Subsurface Light Transport. In *Proc. SIGGRAPH*. ACM, New York, NY, USA, 511–518. <https://doi.org/10.1145/383259.383319>
- A. Keller, L. Fascione, M. Fajardo, I. Georgiev, P. Christensen, J. Hanika, C. Eisenacher, and G. Nichols. 2015. The Path Tracing Revolution in the Movie Industry. In *ACM SIGGRAPH 2015 Courses (SIGGRAPH '15)*. ACM, New York, NY, USA, Article 24, 7 pages. <https://doi.org/10.1145/2776880.2792699>
- Alexander B. Kostinski. 2001. On the extinction of radiation by a homogeneous but spatially correlated random medium. *J. Opt. Soc. Am. A* 18, 8 (Aug 2001), 1929–1933. <https://doi.org/10.1364/JOSAA.18.001929>
- A. B. Kostinski and A. R. Jameson. 2000. On the Spatial Distribution of Cloud Particles. *Journal of the Atmospheric Sciences* 57, 7 (2000), 901–915.
- Edward W. Larsen and Richard Vasques. 2011. A generalized linear Boltzmann equation for non-classical particle transport. *Journal of Quantitative Spectroscopy and Radiative Transfer* 112, 4 (2011), 619–631.
- Benoit B. Mandelbrot and John W. Van Ness. 1968. Fractional Brownian Motions, Fractional Noises and Applications. *SIAM Rev.* 10, 4 (1968), 422–437. <https://doi.org/10.1137/1010093>
- Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003. A Data-Driven Reflectance Model. *ACM Trans. Graph. (Proc. SIGGRAPH)* 22, 3 (July 2003), 759–769. <https://doi.org/10.1145/882262.882343>
- Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. 2015. Multi-scale Modeling and Rendering of Granular Materials. *ACM Trans. Graph.* 34, 4, Article 49 (July 2015), 13 pages. <https://doi.org/10.1145/2766949>
- Jonathan T. Moon, Bruce Walter, and Steve Marschner. 2008. Efficient Multiple Scattering in Hair Using Spherical Harmonics. *ACM Trans. Graph. (Proc. SIGGRAPH)* 27, 3, Article 31 (Aug. 2008), 7 pages. <https://doi.org/10.1145/1360612.1360630>
- Jonathan T. Moon, Bruce Walter, and Stephen R. Marschner. 2007. Rendering Discrete Random Media Using Precomputed Scattering Solutions. In *Rendering Techniques (Proc. EGSR)*. 231–242.

- Thomas Müller, Marios Papas, Markus Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient Rendering of Heterogeneous Polydisperse Granular Media. *ACM Trans. Graph.* 35, 6, Article 168 (Nov. 2016), 14 pages. <https://doi.org/10.1145/2980179.2982429>
- Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 37, 2 (May 2018).
- Derek Nowrouzezahrai, Jared Johnson, Andrew Selle, Dylan Lacewell, Michael Kaschalk, and Wojciech Jarosz. 2011. A Programmable System for Artistic Volumetric Lighting. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4, Article 29 (Aug. 2011). <https://doi.org/10.1145/2010324.1964924>
- Ken Perlin. 1985. An Image Synthesizer. *Proc. SIGGRAPH* 19, 3 (July 1985), 287–296. <https://doi.org/10.1145/325165.325247>
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Bui Tuong Phong. 1975. Illumination for Computer Generated Pictures. *Commun. ACM* 18, 6 (June 1975), 311–317. <https://doi.org/10.1145/360825.360839>
- Matthias Raab, Daniel Seibert, and Alexander Keller. 2008. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 591–606. https://doi.org/10.1007/978-3-540-74496-2_35
- Y. Rahmani. 2013. *Micromechanics and rheology of hard and soft-sphere colloidal glasses*. Ph.D. Dissertation.
- Raymond A Shaw, Alexander B Kostinski, and Daniel D Lanterman. 2002. Super-exponential extinction of radiation in a negatively correlated random medium. *Journal of Quantitative Spectroscopy and Radiative Transfer* 75, 1 (2002), 13 – 20. [https://doi.org/10.1016/S0022-4073\(01\)00287-4](https://doi.org/10.1016/S0022-4073(01)00287-4)
- B. Smith, J. Vasut, T. Hyde, L. Matthews, J. Reay, M. Cook, and J. Schmoke. 2004. Dusty plasma correlation function experiment. *Advances in Space Research* 34, 11 (2004), 2379 – 2383. <https://doi.org/10.1016/j.asr.2003.07.073>
- T Smith and J Guild. 1931. The C.I.E. colorimetric standards and their use. *Transactions of the Optical Society* 33, 3 (1931), 73. <https://doi.org/10.1088/1475-4878/33/3/301>
- Richard Vasques and Edward W. Larsen. 2014a. Non-classical particle transport with angular-dependent path-length distributions. I: Theory. *Annals of Nuclear Energy* 70, Supplement C (2014), 292 – 300. <https://doi.org/10.1016/j.anucene.2013.12.021>
- Richard Vasques and Edward W. Larsen. 2014b. Non-classical particle transport with angular-dependent path-length distributions. II: Application to pebble bed reactor cores. *Annals of Nuclear Energy* 70, Supplement C (2014), 301 – 311. <https://doi.org/10.1016/j.anucene.2013.12.020>
- Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph.D. Dissertation. Stanford, CA, USA.
- Josh Wills, Sameer Agarwal, David Kriegman, and Serge Belongie. 2009. Toward a Perceptual Space for Gloss. *ACM Trans. Graph. (Proc. SIGGRAPH)* 28, 4, Article 103 (Sept. 2009), 15 pages. <https://doi.org/10.1145/1559755.1559760>
- E. R. Woodcock, T. Murphy, P. J. Hemmings, and T. C. Longworth. 1965. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In *Applications of Computing Methods to Reactor Problems*. Argonne National Laboratory.
- Magnus Wrenninge, Chris Kulla, and Viktor Lundqvist. 2013. Oz: The Great and Volumetric. In *ACM SIGGRAPH 2013 Talks (SIGGRAPH '13)*. ACM, New York, NY, USA, Article 46, 1 pages. <https://doi.org/10.1145/2504459.2504518>
- Magnus Wrenninge, Ryusuke Villemin, and Christophe Hery. 2017. *Path Traced Sub-surface Scattering using Anisotropic Phase Functions and Non-Exponential Free Flights*. Technical Report.
- Magnus Wrenninge, Nafees Bin Zafar, Ollie Harding, Gavin Graham, Jerry Tessendorf, Victor Grant, Andrew Clinton, and Antoine Bouthors. 2011. Production Volume Rendering 2: Systems. In *ACM SIGGRAPH 2011 Courses*. ACM, New York, NY, USA.
- Wen Xu, Alex Nikolov, Darsh T Wasan, Alex Gonsalves, and Rajendra P Borwankar. 1998. Fat particle structure and stability of food emulsions. *Journal of food science* 63, 2 (1998), 183–188.
- Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order Similarity Relations in Radiative Transfer. *ACM Trans. Graph. (Proc. SIGGRAPH)* 33, 4, Article 104 (July 2014), 12 pages. <https://doi.org/10.1145/2601097.2601104>