

Solving partial differential equations in participating media

BAILEY MILLER, Carnegie Mellon University, USA

ROHAN SAWHNEY, NVIDIA, USA

KEENAN CRANE, Carnegie Mellon University, USA

IOANNIS GKIOULEKAS, Carnegie Mellon University, USA

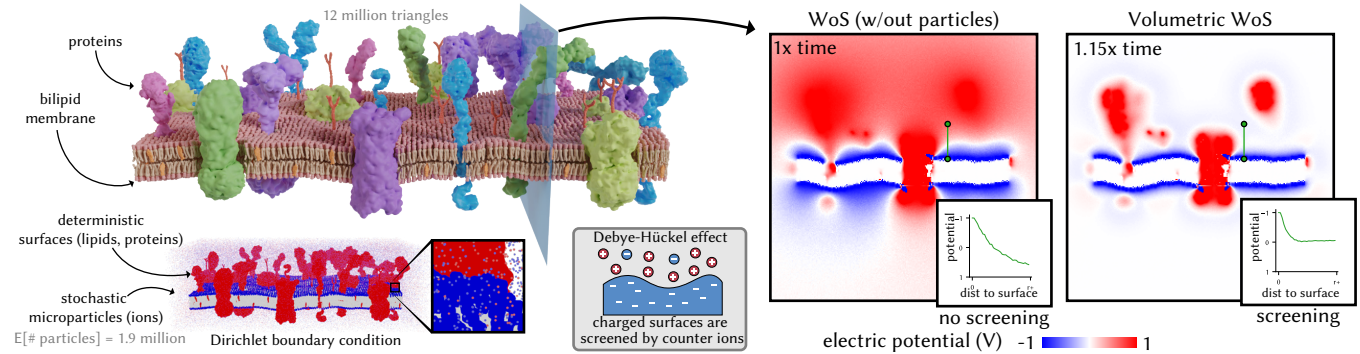


Figure 1. Real physical systems, such as biomembranes, have extraordinarily complex geometry, often managed by *homogenizing* particulate substances that mediate physical interactions. For example, in the setup we show, a homogenized PDE is commonly used to model *electrostatic screening* due to ions in solution. In reality, however, homogenization can be highly inaccurate, as particles often exhibit a scale similar to geometric features. We develop Monte Carlo methods that directly account for both particle and boundary geometry, maintaining efficiency without any limiting assumptions or geometric simplification.

We consider the problem of solving partial differential equations (PDEs) in domains with complex microparticle geometry that is impractical, or intractable, to model explicitly. Drawing inspiration from volume rendering, we propose tackling this problem by treating the domain as a participating medium that models microparticle geometry *stochastically*, through aggregate statistical properties (e.g., particle density). We first introduce the problem setting of PDE simulation in participating media. We then specialize to *exponential media* and describe the properties that make them an attractive model of microparticle geometry for PDE simulation problems. We use these properties to develop two new algorithms, *volumetric walk on spheres* and *volumetric walk on stars*, that generalize previous Monte Carlo algorithms to enable efficient and discretization-free simulation of linear elliptic PDEs (e.g., Laplace) in participating media. We demonstrate experimentally that our algorithms can solve Laplace boundary value problems with complex microparticle geometry more accurately and more efficiently than previous approaches, such as ensemble averaging and homogenization.

CCS Concepts: • **Computing methodologies** → **Physical simulation; Rendering.**

Additional Key Words and Phrases: Partial differential equations, Monte Carlo methods, walk on spheres, participating media, volume rendering

Authors' Contact Information: Bailey Miller, bmmiller@andrew.cmu.edu, Carnegie Mellon University, Pittsburgh, PA, USA; Rohan Sawhney, rsawhney@nvidia.com, NVIDIA, Santa Clara, CA, USA; Keenan Crane, kmcrane@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, PA, USA; Ioannis Gkioulekas, igkioule@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, PA, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2025 Copyright held by the owner/author(s).
ACM 1557-7368/2025/8-ART
<https://doi.org/10.1145/3731152>

ACM Reference Format:

Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2025. Solving partial differential equations in participating media. *ACM Trans. Graph.* 44, 4 (August 2025), 20 pages. <https://doi.org/10.1145/3731152>

1 Introduction

How do we simulate the electrostatic potential due to ions around a bilipid membrane (Figure 1)? Or the concentration of ozone forming in clouds due to the photochemical effect (Figure 13)? These are natural phenomena where we know how to model the underlying physical processes mathematically, using partial differential equations (PDEs) with appropriate boundary conditions. Yet despite continued advances—in computer graphics and other sciences—on numerical PDE solvers, both phenomena remain challenging to simulate *accurately* and *efficiently*. The root of the challenge is the same in both phenomena: Even though the problem geometry at the macroscopic level is relatively simple—we can accurately model the outer shell of the cloud in Figure 13, or the lipids and proteins in Figure 1, with good efficiency—at the microscopic level the geometry becomes extremely complex—we need to specify the position of every water droplet forming the cloud, or every ion around the membrane, a daunting task. These are but two examples of a much broader class of natural phenomena that involve similarly complex *microparticle geometry*, and thus are likewise challenging to simulate. This class includes liquid flow in porous or colloidal media, diffusive processes across biology, or heat transfer in granular media.

Our goal in this paper is to develop numerical methods that enable accurate and efficient simulation of PDEs in the presence of complex microparticle geometry, as in the above examples. Obviously, we are not the first to consider this problem. Perhaps most

prominent among previous approaches are *homogenization methods* [Marchenko and Khruslov 2008], which deal with complex microparticle geometry by modeling only the *asymptotic* behavior of the PDE solution, for example, as particles become infinitesimally small and infinitely dense. Such limiting cases effectively eliminate (“homogenize”) the microparticle geometry, and are thus tractable to simulate. Unfortunately, the resulting PDE solutions are not always good approximations of the true behavior: Real problem settings never exactly match the homogenization assumptions, and the degree of deviation from these assumptions can vary drastically across the domain of each individual problem, making it difficult to provide accuracy guarantees or means to control the approximation error.

The need to develop numerical methods that can deal with complex microparticle geometry also arises in other areas of computer graphics. For example, in rendering, such methods are needed to simulate photorealistic images of scenes that include materials such as clouds, smoke, fog, and biological tissue. Accurate simulation of light transport in such scenes requires accounting for the multiple interactions of light with the microparticle geometry of these materials, thus posing the same computational challenge as in PDE simulation. To overcome this challenge, the predominant approach in rendering has been to represent materials with complex microparticle geometry as *participating media*: Rather than enumerate explicit configurations of microscopic particles, such media model the microparticle geometry *stochastically*, through bulk statistical properties such as the average particle density at different parts of the material. Combining this representation with Monte Carlo rendering methods—such as path tracing—has led to the development of *volume rendering* algorithms [Novák et al. 2018]—such as volumetric path tracing—as a highly successful methodology for the simulation of light transport in materials with complex microparticle geometry. In particular, volume rendering algorithms compute the average (in the sense of statistical expectation) light transport behavior inside such materials, without requiring the limiting assumptions of homogenization methods. This computational capability has facilitated applications in areas well beyond computer graphics, including remote sensing, chemistry, material science, and medicine.

In this paper, we set out to bring the same computational capability to PDE simulation in problem settings with complex microparticle geometry. To achieve this goal, we focus on *Monte Carlo algorithms* for simulation of *linear elliptic PDEs* [Sawhney and Crane 2020]—for example *walk on spheres* and *walk on stars*—to leverage their striking similarity to Monte Carlo rendering algorithms, as well as other advantages they provide over alternative grid-based simulation methods. We first formally introduce the problem of PDE simulation in participating media (Section 3), then specialize to so-called *exponential media*—the most common type of participating media, assuming independent particles described statistically by the *Poisson Boolean model* for stochastic geometry (Section 4).

We then develop two new algorithms, *volumetric walk on spheres* (Section 5) and *volumetric walk on stars* (Section 6), which generalize their non-volumetric namesakes to support simulation of linear elliptic PDEs in participating media. Our development mimics that of volume rendering algorithms for exponential media, taking advantage of the close similarities between Monte Carlo algorithms

for simulation and rendering (Appendix B). We demonstrate the accuracy and efficiency of our algorithms (Section 7) through comparisons with homogenization and other baseline approaches (ensemble averaging). Lastly, we show example applications of our algorithms through simulations of natural phenomena with complex microparticle geometry (Section 8)—in particular the two phenomena at the start of this section, the electrostatic potential of bilipid membranes (Figure 1), and the photochemical effect in clouds (Figure 13). We provide an open-source implementation on the project website.¹

2 Related work

Our work bridges ideas from PDE simulation, rendering, and stochastic geometry. We review related literature across these areas.

Monte Carlo PDE simulation. Our work continues the development of *Monte Carlo algorithms* for PDE simulation, generalizing algorithms such as *walk on spheres* [Muller 1956] and *walk on stars* [Sawhney et al. 2023; Miller et al. 2024b] to problems involving participating media. Despite their long history in applied mathematics and other areas [Sabelfeld and Simonov 2016], Monte Carlo algorithms for PDE simulation were only recently introduced to computer graphics [Sawhney and Crane 2020]. Since their introduction, they have been gaining popularity as an alternative to traditional grid-based methods (finite elements and boundary elements [Hunter and Pullan 2001; Costabel 1987]), thanks to the critical advantages they provide—output sensitivity, parallelism, robustness to imperfect geometry, and compatibility with varied geometric representations [Sawhney and Crane 2020, Section 1]. Graphics research in the past five years has seen the rapid development of these algorithms to support much broader types of PDEs [Rioux-Lavoie et al. 2022; Bati et al. 2023; De Lambilly et al. 2023; Sugimoto et al. 2024; Sawhney et al. 2022] and boundary conditions [Nabizadeh et al. 2021; Miller et al. 2024b; Sawhney et al. 2023; Sugimoto et al. 2023], as well as improve efficiency [Qi et al. 2022; Li et al. 2024, 2023; Bakboug and Peers 2023; Miller et al. 2023].

We aim to further extend the capabilities of Monte Carlo PDE simulation algorithms, to enable simulation of problems with extremely complex microparticle geometry. Such problems arise in the modeling of natural phenomena such as flow effects in porous or colloidal media [Kadivar et al. 2021], diffusive effects in biology [Rothschild 1992; Brydges and Federbush 1980] (Figure 1), and radiative-diffusive photochemical effects in clouds [Faust 1994] (Figure 13). Simulating these phenomena has been a long-standing challenge also for grid-based simulation methods, typically necessitating the use of *homogenization methods*, as we review below. We introduce a fundamentally different approach to addressing this challenge, based on representations of microparticle geometry as *participating media*. Our approach is inspired by the similarity between Monte Carlo algorithms for PDE simulation and rendering, as we elaborate next.

Participating media and volume rendering. Participating media have a long history in computer graphics, and especially rendering [Drebin et al. 1988], as a methodology for tractably modeling microparticle geometry. Rather than pin down an exact particle configuration, they model *stochastic* configurations where particle

¹https://imaging.cs.cmu.edu/volumetric_walk_on_spheres

properties (location, shape) are random variables determined by the medium properties. Different statistical models give rise to different types of media, including *exponential media* with independent spherical or anisotropic particles [Jakob et al. 2010; Heitz et al. 2015], or non-exponential media with correlated (e.g., repulsive or attractive) particles [Bitterli et al. 2018; Jarabo et al. 2018; d'Eon 2018, 2019].

Participating media have enabled light transport simulation in microparticle geometry of extreme complexity [Meng et al. 2015; Müller et al. 2016; Moon et al. 2007], through tailored *volume rendering* algorithms (such as volumetric path tracing) [Novák et al. 2018]. These algorithms generalize Monte Carlo rendering algorithms for deterministic geometry (such as path tracing) [Veach 1998], to account for geometry stochasticity inside participating media. This generalization requires only replacing geometric queries used by the deterministic algorithms (ray casting) with routines that instead sample random query outcomes from distributions determined by the medium properties (free-flight distance sampling). The development of efficient and accurate such routines has been a fruitful research area [Novák et al. 2014; Raab et al. 2006; Kutz et al. 2017; Kettunen et al. 2021; Miller et al. 2019; Georgiev et al. 2019].

Volume rendering algorithms simulate *expected* light transport in participating media, and are used well beyond computer graphics, for example in remote sensing [Levis et al. 2017, 2015; Salesin et al. 2024a,b], chemistry [Berne and Pecora 2000; Weitz and Pine 1993], and medical imaging [Alterman et al. 2021; Bar et al. 2019]. Recent work has extended volume rendering algorithms to stochastic *macroscopic* geometry [Miller et al. 2024a; Vicini et al. 2021; Seyb et al. 2024], for example arising due to acquisition noise or incomplete surface information [Sellán and Jacobson 2022, 2023]. We aim to enable similar capabilities in PDE simulation, by developing volumetric Monte Carlo simulation algorithms. This development is aided by the structural similarity between Monte Carlo algorithms for simulation and rendering—as in rendering, all we need to do is replace geometric queries in simulation algorithms (closest point queries) with appropriate sampling routines (closest point sampling).

Homogenization of PDEs. Homogenization methods provide an alternative methodology for analyzing and estimating solutions to linear elliptic PDEs in domains perforated by microparticle geometry. These methods consider the limit case as particles become infinitesimally small and at the same time infinitely dense. At that limit, the PDE solution asymptotically approaches the solution to a PDE with additional *screening* in a domain without the microparticle geometry—Marchenko and Khruslov [2008] provide rigorous statements. Starting with the work of Papanicolaou and Varadhan [1980], this asymptotic behavior has been shown to hold for different types of microparticle geometry, including periodic Cioranescu and Murat [1982] (Cioranescu and Murat [1997] provides an English translation) and stochastic with independent or correlated particles [Caffarelli and Mellet 2009; Calvo-Jurado et al. 2015; Giunti et al. 2018]. Analogous asymptotic results also hold for more general PDEs outside the scope of our paper, for example homogenizing the Stokes flow equation into the Darcy-Brinkman equation [Giunti and Höfer 2019; Whitaker 1986; Brinkman 1949].

Unfortunately, the solution to the homogenized PDE can be a poor approximation to the true solution for problems that are far




Figure 2. We consider domains $\Omega := V \setminus O$ equal to the difference between a deterministic volume V and a random configuration of particles O contained in the volume. The particle configuration follows the Poisson Boolean model (PBM) whose density $\lambda(x)$ is defined over the volume.

from the homogenization limit, for example with particles of modest size or small density. Additionally, the approximation error can vary considerably at different parts of the PDE domain, for example near detailed geometry versus far from the boundary. Our volumetric method overcomes these issues by considering the expected, rather than asymptotic, PDE solution, and remains accurate across particle and density scales, and throughout the entire domain. We revisit homogenization and show experimental comparisons in Section 7.2. Lastly, homogenization methods have proven successful in graphics for simulation problems involving microscopic and multi-scale geometry beyond particle perforations [Yuan et al. 2024; Desbrun et al. 2013; Sperl et al. 2020; Kharevych et al. 2009].

3 Background and problem statement

We first review boundary value problems (BVPs) with linear elliptic PDEs and Dirichlet boundary conditions, and the walk on spheres algorithm for solving such BVPs. We then introduce the problem of solving BVPs in participating media—the focus of our paper.

3.1 Notation

We work in three dimensions and use the Euclidean norm for distances. Given any set $Q \subset \mathbb{R}^3$ and point $x \in \mathbb{R}^3$, we define the functions $\text{closest}(x, Q) := \arg\min_{y \in Q} \|x - y\|$ and $\text{dist}(x, Q) := \|x - \text{closest}(Q, x)\|$ returning the closest point and shortest distance (resp.) to x among points in Q . We denote by $B(x, r)$ the ball with center x and radius r . For any two points $x, y \in \mathbb{R}^3$, we define $\text{dir}(x, y) := (y - x) / \|y - x\|$ as the unit vector pointing from x to y .

Throughout the paper, we use three sets: the *domain* Ω , the *volume* V , and the *microparticle geometry* O , with boundaries $\partial\Omega$, ∂V , and ∂O (resp.). We detail their roles and relationships in Sections 3.2 and 3.4. For any point $x \in \mathbb{R}^3$, we denote its closest points on the boundary of these sets as $y^{\partial\Omega}(x) := \text{closest}(x, \partial\Omega)$, $y^{\partial V}(x) := \text{closest}(x, \partial V)$, and $y^{\partial O}(x) := \text{closest}(x, \partial O)$; we also denote the corresponding shortest distances $r^{\partial\Omega}(x) := \|x - y^{\partial\Omega}(x)\|$, $r^{\partial V}(x) := \|x - y^{\partial V}(x)\|$, and $r^{\partial O}(x) := \|x - y^{\partial O}(x)\|$. Given a sequence of points x_0, x_1, \dots (e.g., a random walk), we use the shorthand notation $y_k^{\partial\Omega} := y^{\partial\Omega}(x_k)$ and $r_k^{\partial\Omega} := r^{\partial\Omega}(x_k)$ whenever the meaning is clear from context; we do likewise for closest points and shortest distances relative to ∂V and ∂O , and other quantities that depend on such sequences later in the paper.

We denote by Δ the Laplace operator on \mathbb{R}^3 , and by $P : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ the rotationally symmetric Poisson kernel of the zero-Dirichlet

Laplace equation on a ball [Sawhney et al. 2023, Appendix A], parameterized with a radius r .

3.2 The Laplace equation

To simplify exposition, throughout Sections 3–5 we consider a prototypical BVP involving the *Laplace equation with Dirichlet boundary conditions*. Our methods can extend to other linear elliptic PDEs that can be simulated with walk on spheres (e.g., screened Poisson equation). It is also possible to extend to other boundary conditions (e.g., Neumann, Robin) that can be simulated using the *walk on stars* algorithm [Sawhney et al. 2023; Miller et al. 2024b]. We delay discussion of this case till Section 6, and for now focus on the BVP:

$$\begin{aligned} \Delta u(x) &= 0 & \text{in } \Omega, \\ u(x) &= g(x) & \text{on } \partial\Omega. \end{aligned} \quad (1)$$

Here, $\Omega \subset \mathbb{R}^3$ is the domain of the BVP, $g : \partial\Omega \rightarrow \mathbb{R}$ is the *Dirichlet boundary data*, and $u : \Omega \rightarrow \mathbb{R}$ is the solution we want to estimate.

3.3 Walk on spheres

When the domain Ω is deterministic, the *walk on spheres* (WoS) algorithm [Sawhney and Crane 2020; Muller 1956] computes stochastic estimates to the solution u of the BVP (1), through recursive single-sample Monte Carlo integration. The starting point of the derivation of WoS is to represent the solution $u(x_0)$ at a point $x_0 \in \Omega$ as an integral of the solution over a sphere centered at x_0 . A single-sample Monte Carlo estimate of this integral requires estimating the solution $u(x_1)$ at a point x_1 sampled on the sphere. WoS estimates $u(x_1)$ by iterating the same single-sample Monte Carlo estimation, resulting in a recursive procedure that performs a random walk x_0, x_1, \dots . The walk terminates when it reaches a point x_k within a small distance $\varepsilon > 0$ from the domain boundary $\partial\Omega$, where the solution is approximated by the known Dirichlet boundary data g at the boundary point closest to x_k —a so-called ε -shell approximation.

Concretely, we can express the solution u to the BVP (1) at a point $x \in \Omega$ using the *boundary integral equation* (BIE) [Costabel 1987]:

$$u(x) = \int_{\partial B(x, r^{\partial\Omega}(x))} P(r^{\partial\Omega}(x)) u(y) dA(y), \quad (2)$$

where dA is the surface area measure. Starting at a point x_0 , recursive single-sample Monte Carlo estimation of this equation with the ε -shell approximation results in the WoS estimator:

$$\langle u(x_k) \rangle := \begin{cases} g(y_k^{\partial\Omega}), & r_k^{\partial\Omega} < \varepsilon, \\ \frac{P(r_k^{\partial\Omega})}{P(r_k^{\partial\Omega})} \langle u(x_{k+1}) \rangle, & \text{otherwise.} \end{cases} \quad (3)$$

At each step, WoS performs a *closest point query* to determine the boundary point $y_k^{\partial\Omega} \in \partial\Omega$ closest to the current walk point x_k , then sets $r_k^{\partial\Omega} := \|x_k - y_k^{\partial\Omega}\|$. The next walk point x_{k+1} is sampled on the sphere $\partial B(x_k, r_k^{\partial\Omega})$ with uniform probability $p(r_k^{\partial\Omega}) := 1/4\pi(r_k^{\partial\Omega})^2$.

3.4 Boundary value problems in participating media

Our focus is instances of the BVP (1) where the domain Ω is perforated by *stochastic microparticle geometry*, known only up to a probabilistic model specified through the properties of a *participating medium*. The stochasticity acts as a convenient abstraction to alleviate the complexity of exact modeling of microscopic geometry

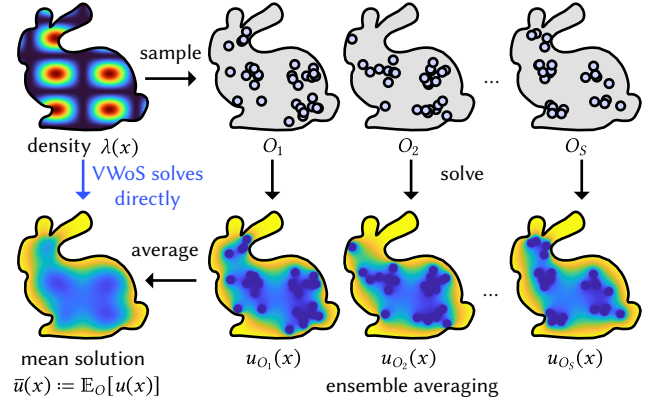


Figure 3. *Ensemble averaging* is a simple but expensive method to estimate the mean solution of a PDE in a participating medium, by first sampling many random particle configurations (*top row*), then solving the PDE on each sampled domain (*bottom row*), and finally averaging the computed solutions. Our *volumetric walk on spheres* algorithm directly estimates the mean solution without expensive ensemble averaging.

(for example, in domains that comprise an intractably large number of particles, such as tissue, clouds, colloidal suspensions, or porous rock formations). Before we specialize the type of participating medium in Section 4, we use this section to set up such BVPs and explain the notion of solution we are interested in.

As we show in Figure 2, we consider the microparticle geometry $O \subset \mathbb{R}^3$ to be the union of a configuration of particles. We assume that all particles are inside a *deterministic* domain $V \subset \mathbb{R}^3$ with boundary ∂V ; following rendering, we term V the *participating medium* or *volume*. Then, the domain of the BVP (1) is equal to the part of the volume not occupied by particles:

$$\Omega := V \setminus O. \quad (4)$$

Additionally, as now any point in V may be on $\partial\Omega$, we extend the domain of the Dirichlet boundary data, $g : V \rightarrow \mathbb{R}$.

The configuration of particles is *random*, with a distribution determined by the medium properties (e.g., particle density, Section 4). Thus, the BVP solution u is a random variable, and we want to compute its expected value, which we term the *mean solution* \bar{u} :

$$\bar{u}(x) := \mathbb{E}_O[u(x)] = \int_{\mathcal{P}(V)} p(O) u(x) dO. \quad (5)$$

Here, the probability $p(O)$ of a particle configuration O , the integration domain, and the integration measure depend on the stochastic microparticle geometry model, as we detail in Section 4.

Ensemble averaging. We can estimate \bar{u} using *ensemble averaging* (Figure 3), which involves: 1. sampling many particle configurations O_s from p ; 2. computing the solution u_{O_s} for each O_s using, e.g., WoS; 3. averaging the computed solutions. The resulting estimate:

$$\langle \bar{u}(x) \rangle_{EA} := \frac{1}{S} \sum_{s=1}^S u_{O_s}(x), \quad (6)$$

is consistent as $S \rightarrow \infty$ and unbiased. Though simple to implement and invaluable for validating correctness of alternative methods, ensemble averaging is impractically expensive: It requires repeatedly sampling large particle configurations O_s , and performing expensive PDE solves in the resulting complex domains. The high computational cost of ensemble averaging is well documented in rendering [Bar et al. 2019; Bitterli et al. 2018], where it has motivated the development of *volume rendering* algorithms [Novák et al. 2018] that simulate light transport in participating media without ensemble averaging. These algorithms typically specialize to specific models of stochastic microparticle geometry, most commonly the *Poisson Boolean model*. Motivated by the success of volume rendering algorithms, our work uses this model to develop simulation algorithms for unbiased estimation of the mean solution \bar{u} that elide ensemble averaging. We detail the Poisson Boolean model in Section 4, then develop our algorithms in Section 5. We elaborate on the relative merits of our algorithms compared to ensemble averaging in Section 7, where we also show experimental comparisons.

4 Poisson Boolean model and exponential media

To model the stochastic microparticle geometry in participating media, we use the *Poisson Boolean model* (PBM), which is commonplace in scientific and engineering applications [Kadivar et al. 2021; Speidel et al. 2018]. In computer graphics, this model underlies volume rendering algorithms for *exponential media* [Novák et al. 2018]. The widespread use of the PBM is for reasons of both modeling accuracy—it is appropriate for phenomena involving *independent* particles—and computational convenience—it is endowed with a wealth of mathematical properties that facilitate simulation.

We first provide the definition of the PBM, then explain how to use it to perform *closest point sampling*, the key sampling procedure we will need in Section 5 to generalize WoS to participating media. The results we present have close analogues in volume rendering, as we detail in Appendix B. We focus on the simplest form of the PBM, which assumes that the microparticle geometry comprises *spherical particles of a fixed radius*. Our presentation follows Last and Penrose [2017, Chapters 16–17] and Chiu et al. [2013, Chapters 2–3], and we refer to these textbooks for more detailed treatments, including generalizations to other types of particles.

DEFINITION 1: POISSON BOOLEAN MODEL

We consider a function $\lambda : V \rightarrow \mathbb{R}_{\geq 0}$ such that $\int_V \lambda(x) dx < \infty$, and a scalar $R \in \mathbb{R}_{\geq 0}$. A stochastic microparticle geometry $O \subset \mathbb{R}^3$ follows the *Poisson Boolean model* with *density* λ and *size* R if it equals a union of balls $O := \bigcup_{n=1}^N B(c_n, R)$ such that the set of centers $C_O := \{c_n \in V\}_{n=1}^N$ is a *Poisson point process* on V with rate function λ . Equivalently:

- the number of balls is a Poisson-distributed integer random variable, $N \sim \text{Poisson}(\int_V \lambda(x) dx)$;
- conditionally on N , the centers are independent and distributed proportionally to the density² λ , $p(c_n | N) \propto \lambda(c_n)$.

²The density λ is also known as the *intensity function* of the Poisson point process.

We often consider the special case of the *homogeneous* PBM with constant density $\lambda(x) \stackrel{h}{=} \lambda$, and refer to the general case as the *heterogeneous* model. We write $O \sim \Phi(\lambda, R)$ for a particle configuration that follows the PBM with rate λ and size R .

Closest point distribution and sampling.

As WoS interacts with the domain through closest point queries, using it in participating media requires reasoning about the *random* closest point $y^{\partial O}(x) \in \partial O$ between a point $x \in V$ and the boundary ∂O of the stochastic microparticle geometry O . This random point follows the so-called *closest point distribution*, whose probability density function (PDF) p_x^{cp} is known analytically when O follows the PBM [Last and Penrose 2017, Section 16.3]. We discuss this distribution in more detail in Appendix A; here we focus instead on *sampling* from it, which we need to perform WoS in participating media (Section 5).

As O is a union of spherical particles, we perform closest point sampling by first sampling the *closest particle center* $c(x) \in C_O$ to x . We then determine the closest point $y^{\partial O}(x)$ to x , and associated shortest distance $r^{\partial O}(x)$, from $c(x)$ as (Figure 4):

$$r^{\partial O}(x) = \|x - c(x)\| - R, \quad y^{\partial O}(x) = x + r^{\partial O}(x) \text{dir}(x, c(x)). \quad (7)$$

To sample $c(x)$, we use the *polar representation* of the Poisson point process describing C_O [Last and Penrose 2017, Section 7.4].




Figure 4. Computing the closest point $y^{\partial O}(x)$ from the closest center $c(x)$.

PROPOSITION 1: POLAR REPRESENTATION OF POISSON POINT PROCESSES

We assume that the set of centers $C_O := \{c_n \in V\}_{n=1}^N$ is a Poisson point process on V with rate function λ . For any point $x \in V$, we let $c(x) \in C_O$ be its random closest center, and $r^c(x) := \|x - c(x)\|$ the random shortest distance-to-center. Then:

1. The shortest distance-to-center $r^c(x)$ has PDF:

$$p_x^{\text{dc}}(r) := \exp(-\Lambda(x, r)) \int_{\partial B(x, r)} \lambda(y) dA(y), \quad (8)$$

where we define:

$$\Lambda(x, r) := \int_{B(x, r)} \lambda(y) dy. \quad (9)$$

2. The closest center $c(x)$ has *conditional* PDF given $r^c(x) = r$:

$$p_x^{\text{cc}}(y | r) := \frac{\lambda(y)}{\int_{\partial B(x, r)} \lambda(y) dA(y)}, \quad (10)$$

for $y \in \partial B(x, r)$, and zero everywhere else.³

In the homogeneous case, Equations (8) and (10) simplify to:

$$p_x^{\text{dc}}(r) \stackrel{h}{=} \exp(-4/3\pi(r)^3\lambda) 4\pi(r)^2\lambda, \quad (11)$$

$$p_x^{\text{cc}}(y | r) \stackrel{h}{=} \frac{1}{4\pi r^2}. \quad (12)$$

³More precisely, the PDF is a Dirac delta on $\partial B(x, r)$ with respect to the area measure.




Figure 5. The Poisson Boolean model enables efficient closest point sampling—or equivalently closest particle center sampling (Figure 4). In a homogeneous medium (left), we first sample an exponential random variable to determine the (cubed) distance to the closest center, then uniformly sample the center itself on a sphere of appropriate radius. In a heterogeneous medium (right), we use thinning to first sample multiple centers—in order of increasing distance—from a medium homogenized through the addition of fictitious density, randomly accept or reject the sampled centers, then use the first accepted one as the sampled closest center.

As we detail in Appendix B, the PDF p_x^{dc} is analogous to the *free-flight distribution* in volume rendering of exponential media [Novák et al. 2018; Bitterli et al. 2018].

Equation (8) (or (11) in the homogeneous case) implies that the *cubed distance* $(r^c(x))^3$ from x to $c(x)$ is an *exponential random variable* with rate $\Lambda(B(x, r))$. This property and Equation (10) (or (12) in the homogeneous case) allow sampling $c(x)$ by first sampling its cubed distance from x , then sampling a point on the corresponding sphere around x [Chiu et al. 2013, Section 2.5] (Figure 5):

- In the homogeneous case, we first sample an exponential random variate $\xi \sim \text{Expo}[4/3\pi\lambda]$, equal to the cubed distance from x to $c(x)$. We then sample $c(x)$ uniformly on $\partial B(x, \sqrt[3]{\xi})$.
- In the heterogeneous case, we use the *thinning* method for sampling heterogeneous Poisson processes [Lewis and Shedler 1979]. Given a *majorant density* $\bar{\lambda} \geq \lambda(x)$, $\forall x \in \mathbb{R}^3$, we sequentially sample exponential random variates $\xi_1, \xi_2, \dots \sim \text{Expo}[4/3\pi\bar{\lambda}]$. For each ξ_s , we sample a point c_s uniformly on $\partial B(x, \sqrt[3]{\sum_{t=1}^s \xi_t})$, and randomly accept or reject it with acceptance probability $\lambda(c_s)/\bar{\lambda}$. Then, we set $c(x)$ as the first accepted point c_s .

After sampling $c(x)$, we set $r^{\partial O}(x)$ and $y^{\partial O}(x)$ using Equation (7). This procedure can return $r^{\partial O}(x) < 0$, corresponding to a case where x is inside a particle, $x \in O$ —it will be convenient for the algorithm we develop in Section 5.1 to distinguish this case by setting $y^{\partial O}(x) := x$. We summarize the sampling procedure for the heterogeneous case in Algorithm 1, which reduces to that for the homogeneous case when λ is constant and we use $\bar{\lambda} := \lambda$.

Intuitively, we can interpret Algorithm 1 as follows: We fill the volume V with *fictitious density* until the sum of real and fictitious density equals $\bar{\lambda}$ everywhere. We then sample particles from the resulting homogeneous medium using exponential sampling, in order of increasing distance from x . We reject particles due to the fictitious density until we find the first particle due to the real density.

Algorithm 1 Closest point sampling in the Poisson Boolean model.

Input: A query point x , a majorant density $\bar{\lambda}$, a struct implementing the PBM density λ , the PBM particle size R .

Output: Closest point $y^{\partial O}$.

```

1: function SAMPLECLOSESTPOINT( $x, \bar{\lambda}, \lambda, R$ )
2:    $\triangleright$  Initialize cubed radius
3:    $b \leftarrow 0$ 
4:   while TRUE do
5:      $\triangleright$  Sample exponential variate
6:      $\xi \leftarrow \text{SAMPLEEXPONENTIAL}(4/3\pi\bar{\lambda})$ 
7:      $\triangleright$  Increment cubed radius
8:      $b += \xi$ 
9:      $\triangleright$  Uniformly sample a point on the unit sphere
10:     $v \leftarrow \text{SAMPLEUNITSPHERE}()$ 
11:     $\triangleright$  Compute particle center
12:     $c \leftarrow x + \sqrt[3]{b} v$ 
13:     $\triangleright$  Compute acceptance probability
14:     $\alpha \leftarrow \lambda.\text{EVALUATE}(c) / \bar{\lambda}$ 
15:     $\triangleright$  Accept or reject the particle center
16:    if  $\alpha > \text{SAMPLEUNIFORM}(0, 1)$  then break
17:     $\triangleright$  Compute distance to closest point
18:     $r^{\partial O} \leftarrow \|x - c\| - R$ 
19:     $\triangleright$  Check if  $x$  is inside the sampled particle
20:    if  $r^{\partial O} < 0$  then return  $x$ 
21:     $\triangleright$  Compute sampled closest point
22:     $y^{\partial O} \leftarrow x + r^{\partial O} \text{dir}(x, c)$ 
23:    return  $y^{\partial O}$ 
24: end function

```

As we detail in Appendix B, Algorithm 1 is analogous to the *delta tracking* algorithm in volume rendering for sampling free-flight distances in heterogeneous exponential media [Coleman 1968; Raab et al. 2006], which has a similar intuitive interpretation in terms of fictitious density [Miller et al. 2019; Novák et al. 2014].

Conditional closest point sampling. Our derivation in Section 5 additionally requires sampling the closest point $y^{\partial O}(x) \in \partial O$ *conditionally* on events of the form $Q \cap O = \emptyset$ for various sets $Q \subset V$ —i.e., knowing that no point of the microparticle geometry is in \emptyset . We prove in Appendix A that we can do so using Algorithm 1, after replacing λ with the *conditional density* $\lambda(\cdot | Q) : V \rightarrow \mathbb{R}_{\geq 0}$:

$$\lambda(x | Q) := \begin{cases} 0, & x \in Q^{\oplus R}, \\ \lambda(x), & \text{otherwise,} \end{cases} \quad (13)$$

where $Q^{\oplus R} := \{x \in \mathbb{R}^3 : \text{dist}(x, Q) \leq R\}$ is the dilation of Q by a ball of radius R . Intuitively, zeroing out the density inside $Q^{\oplus R}$ ensures that centers sampled closer to Q than R will be rejected, thus guaranteeing $Q \cap O = \emptyset$.

5 Volumetric walk on spheres

We now derive a recursive Monte Carlo estimator for the mean solution \bar{u} of the BVP (1) in exponential participating media. Our estimator generalizes the standard WoS estimator (3), by leveraging




Figure 6. (Left) Walk on spheres jumps along spheres whose radius it determines through closest point queries with respect to the deterministic domain boundary. It terminates when it reaches the ε -shell of the boundary. (Right) Volumetric walk on spheres also jumps along spheres, but determines their radius through closest point *sampling* with respect to both stochastic microparticle geometry and the deterministic domain. It terminates when it reaches the ε -shell of either the microparticle geometry or the domain.

one of its critical properties [Sawhney and Crane 2020, Section 1]: WoS interacts with the BVP domain Ω only through closest point queries to determine, at each point x , its closest point $y^{\partial\Omega}$ on $\partial\Omega$. Thus, when Ω includes stochastic microparticle geometry, it suffices to replace these closest point queries with closest point sampling from appropriate distributions (Figure 6). The resulting *volumetric walk on spheres* (VWoS) estimator (Equation (19) and Algorithm 2) enables direct simulation of the mean solution without ensemble averaging, while remaining structurally very close to WoS.

5.1 Derivation

At a high level, our derivation of VWoS mimics that of WoS: We first derive an integral equation (Equation (18)) for the mean solution \bar{u} —an analogue of the BIE Equation (2)—then apply recursive single-sample Monte Carlo estimation. There is, however, an important difference with WoS: As we recurse, the integral equation changes, to *condition* on the history accumulated during previous steps. To build intuition about this conditioning, it is convenient to separately derive the initial and subsequent steps of the VWoS recursion.

Initial step. To evaluate the mean solution \bar{u} at an initial point $x_0 \in V$, we use expectation on both sides of the BIE (2):

$$\bar{u}(x_0) = \mathbb{E}_O \left[\int_{\partial B(x_0, r_0^{\partial\Omega})} P(r_0^{\partial\Omega}) u(x_1) dA(x_1) \right]. \quad (14)$$

Besides u , the only random quantity in the right-hand-side integral is the distance $r_0^{\partial\Omega} := \|x_0 - y_0^{\partial\Omega}\|$ between x_0 and $y_0^{\partial\Omega}$, the random point on $\partial\Omega$ closest to x_0 . We use the shorthand $p_0^{\text{cp}} := p_{x_0}^{\text{cp}}$ for the PDF of $y_0^{\partial\Omega}$ —the closest point distribution in Section 4 modified to account for the deterministic medium boundary ∂V . Using the law of total expectation, we rewrite Equation (14) as:

$$\bar{u}(x_0) = \int_V p_0^{\text{cp}}(y_0) \int_{\partial B(x_0, r_0)} P(r_0) \bar{u}(x_1 | y_0^{\partial\Omega} = y_0) dA(x_1) dy_0, \quad (15)$$

where $r_0 := \|x_0 - y_0\|$ and $\cdot | \cdot$ is probabilistic conditioning.

To derive an estimator for $\bar{u}(x_0)$, we first use single-sample Monte Carlo for both integrals in Equation (15).

1. We sample a point $y_0 \sim p_0^{\text{cp}}$ by: using Algorithm 1 to sample a closest point $y_0^{\partial O}$ on the stochastic ∂O ; using a closest point query to determine the closest point $y_0^{\partial V}$ on the deterministic ∂V ; then setting $y_0 := \text{closest}(x_0, \{y_0^{\partial O}, y_0^{\partial V}\})$.
2. We sample a point x_1 uniformly on the sphere $\partial B(x_0, r_0)$.

How we proceed depends on the sampled distance r_0 . If r_0 is smaller than a threshold ε , we use an ε -shell approximation to set $\bar{u}(x_0)$ equal to the Dirichlet boundary data $g(x_0)$, and terminate.⁴ Importantly, this approximation also covers the case where $r_0 = 0$, which means that Algorithm 1 determined that x_0 is inside a particle (Section 4). Otherwise, we need to recursively estimate $\bar{u}(x_1 | y_0^{\partial\Omega} = y_0)$, as we explain below. Thus, we arrive at the estimator:

$$\langle \bar{u}(x_0) \rangle := \begin{cases} g(x_0), & r_0 < \varepsilon, \\ \frac{P(r_0)}{p(r_0)} \langle \bar{u}(x_1 | y_0^{\partial\Omega} = y_0) \rangle, & \text{otherwise.} \end{cases} \quad (16)$$

Subsequent steps. We now consider the k -th step of the VWoS recursion. After consecutive applications of the law of total expectation and single-sample Monte Carlo—analogously to Equations (14) and (16) (resp.)—we must *condition* on closest points on the domain boundary sampled in all previous steps. Intuitively, each sampled closest point pins down part of the stochastic microparticle geometry that subsequent closest point sampling procedures must remember and continue to respect. To simplify notation, we define the walk *memory* M_k accumulated at step k as the logical conjunction of closest point sampling outcomes in previous steps:

$$M_k := y_0^{\partial\Omega} = y_0 \wedge \cdots \wedge y_{k-1}^{\partial\Omega} = y_{k-1}, \quad M_0 := \emptyset. \quad (17)$$

We elaborate on the interpretation of M_k , and how to use it for conditional closest point sampling, in Section 5.2.

Then, at the k -th step of the VWoS recursion, we must estimate the conditional mean solution $\bar{u}(x_k | M_k)$. Exactly analogously to the case $k = 0$ and Equation (15), we use the BIE (2) and the law of total expectation to derive an integral equation for $\bar{u}(x_k | M_k)$, which we term the *boundary integral equation in participating media*.

BOUNDARY INTEGRAL EQUATION IN PARTICIPATING MEDIA

At the k -step of a random walk with memory M_k as in Equation (17), the conditional mean solution \bar{u} to the Laplace equation of (1) in a participating medium V satisfies:

$$\bar{u}(x_k | M_k) = \int_V p_k^{\text{cp}}(y_k | M_k) \cdot \int_{\partial B(x_k, r_k)} P(r_k) \bar{u}(x_{k+1} | M_{k+1}) dA(x_{k+1}) dy_k, \quad (18)$$

where $r_k := \|x_k - y_k\|$.

Exactly analogously to the derivation of the estimator (16) for Equation (15), we use single-sample Monte Carlo and the ε -shell approximation to derive an estimator for Equation (18), which we term the *volumetric walk on spheres* (VWoS) estimator.

⁴This ε -shell approximation differs from that in WoS (Equation (3)): It evaluates g directly at x_0 , which is possible due to the extension of g to all of V . In practice, for small particle sizes R , the two approximations behave very similarly.

VOLUMETRIC WALK ON SPHERES ESTIMATOR

A recursive single-sample Monte Carlo estimator for Equation (18) is given by:

$$\langle \bar{u}(x_k | M_k) \rangle := \begin{cases} g(x_k), & r_k < \varepsilon, \\ \frac{p(r_k)}{p(r_k)} \langle \bar{u}(x_{k+1} | M_{k+1}) \rangle, & \text{otherwise,} \end{cases} \quad (19)$$

where: the memories M_k and M_{k+1} are defined in Equation (17); $r_k := \|x_k - y_k\|$ with y_k sampled from $p_k^{\text{cp}}(\cdot | M_k)$; and the next point x_{k+1} is sampled uniformly on $\partial B(x_k, r_k)$.

Using the definition of M_0 in Equation (17), Equations (18) and (19) subsume Equations (15) and (16) (resp.) by setting $k := 0$.

Algorithm 2 summarizes an implementation of VWoS, highlighting differences from WoS (Equation (3)). The two estimators are structurally near-identical, with only two differences:

- Whereas WoS determines sphere radii using deterministic closest point queries, VWoS does so using conditional closest point sampling (Algorithm 2, line 8).
- Whereas WoS is *memoryless*—each step is independent of previous steps—VWoS has *full memory*—each step depends on all previous steps (Algorithm 2, lines 14 and 20).

Thanks to its close similarity to WoS, VWoS maintains the advantages of WoS [Sawhney and Crane 2020, Section 1], and is easy to implement within existing WoS libraries [Sawhney and Miller 2024], requiring only an implementation of memory (including associated sampling and updating procedures), which we discuss next.

5.2 Closest point sampling with memory

Intuitively, closest point sampling with memory M_k requires respecting empty space (inside walk spheres) and sampled particles accumulated during previous walk steps. We first formalize this intuition about conditioning on M_k , then present procedures for memory updating and closest point sampling with memory in VWoS.

Understanding conditioning on memory. As M_k is a conjunction of multiple *closest point events* (Equation (17)), we first consider each individual such event. Conditioning on $y_l^{\partial\Omega} = y_l$, for any $l = 0, \dots, k-1$, has two implications (Figure 7(b, c)):

- C1. No point of the microparticle geometry O is closer to walk point x_l than $r_l = \|x_l - y_l\|$, i.e., the ball $B(x_l, r_l)$ is empty.
- C2. If y_l is not on the medium boundary ∂V , then it is on the stochastic boundary ∂O , i.e., the microparticle geometry O includes a particle $B(c_l, R)$ centered at $c_l := y_l + R \text{dir}(x_l, y_l)$.

Formally, from C1–C2, the closest point event $y_l^{\partial\Omega} = y_l$ is equivalent to the event $B(x_l, r_l) \cap O = \emptyset \wedge c_l \in C_O$.

Next we consider the full memory M_k . It is convenient to associate with M_k two sets summarizing the information from all closest point events it includes. We define the *empty-ball memory* $E(M_k)$ and *sampled-particle memory* $C(M_k)$ as the unions of the empty balls (C1) and sampled particles (C2) (resp.) for all eligible events in M_k :

$$E(M_k) := \bigcup_{l=0}^{k-1} B(x_l, r_l), \quad C(M_k) := \bigcup_{l: y_l \notin \partial V} B(c_l, R). \quad (20)$$

$C(M_k)$ will typically include fewer than k (and maybe even zero) particles, as only steps where the sampled closest point is not on

Algorithm 2 The volumetric walk on spheres estimator.

Note: Comments in orange highlight changes to walk on spheres.

Input: A query point x , a parameter ε , a majorant density $\bar{\lambda}$, a struct implementing the PBM density λ , the PBM particle size R .

Output: A single-sample estimate of the mean solution $\bar{u}(x)$.

```

1: function INITIALIZEESTIMATION( $x, \varepsilon, \bar{\lambda}, \lambda, R$ )
2:    $\triangleright$  Initialize empty memory
3:    $M \leftarrow \text{MEMORY.INITIALIZE}()$ 
4:   return VOLUMETRICWALKONSPHERES( $x, \varepsilon, \bar{\lambda}, \lambda, R, M$ )
5: end function
6: function VOLUMETRICWALKONSPHERES( $x, \varepsilon, \bar{\lambda}, \lambda, R, M$ )
7:    $\triangleright$  Sample closest point conditionally on current memory
8:    $y \leftarrow \text{SAMPLECLOSESTPOINTWITHMEMORY}(x, \bar{\lambda}, \lambda, R, M)$ 
9:    $\triangleright$  Compute radius of next walk sphere
10:   $r \leftarrow \|x - y\|$ 
11:   $\triangleright$  Check for  $\varepsilon$ -shell approximation
12:  if  $r < \varepsilon$  then return  $g(x)$ 
13:   $\triangleright$  Update memory:
14:   $M.\text{UPDATE}(x, y)$ 
15:   $\triangleright$  Uniformly sample a point on the unit sphere
16:   $v \leftarrow \text{SAMPLEUNITSPHERE}()$ 
17:   $\triangleright$  Compute next walk point
18:   $x \leftarrow x + rv$ 
19:   $\triangleright$  Continue from next walk point with updated memory
20:  return VOLUMETRICWALKONSPHERES( $x, \varepsilon, \bar{\lambda}, \lambda, R, M$ )
21: end function
```

∂V contribute a particle (C2). Additionally, $C(M_k)$ may include the same particle multiple times: at each step $l \in \{1, k\}$, the closest point y_l may be on the boundary of a particle sampled at a prior step $l' < l$, and fixed as deterministic geometry for subsequent steps.

Using $E(M_k)$ and $C(M_k)$, we can express the implications of conditioning on M_k as direct generalizations of C1–C2 (Figure 7(a)):

- M1. The space $E(M_k)$ is empty.
- M2. The microparticle geometry O includes $C(M_k)$.

Formally, from M1–M2, the memory M_k is equivalent to the event $E(M_k) \cap O = \emptyset \wedge \{c_l : y_l \notin \partial V\} \subset C_O$.

In practice, we implement memory as in Algorithm 3:⁵ We use two list data structures that contain the center-radius pairs (x_l, r_l) and particle centers c_l in Equation (20), along with procedures for memory updates at each walk step (Algorithm 3, line 17), containment queries on $E(M_k)$ (Algorithm 3, line 25), and closest point queries on $\partial C(M_k)$ (Algorithm 3, line 31)—we use the two queries for closest point sampling with memory, as we discuss next.

Sampling procedure. From the above discussion, sampling a closest point y_k conditionally on M_k requires that we do not sample in the empty space $E(M_k)$ (M1), and that we consider the boundary of previously sampled particles $C(M_k)$ as *deterministic* (M2). Realizing both requirements algorithmically is straightforward. First, we use thinning on the density of the medium to remove the empty space,

⁵For efficiency, this implementation does not allow duplicate particles in $C(M_k)$.




Figure 7. (a) To sample the closest point $y^{\partial\Omega}$ at x conditionally on the memory M accumulated during a walk, we determine two points: First, we sample the random closest point $y^{\partial\Omega}$ on the stochastic microparticle geometry, but with the PBM density zeroed out inside the spheres formed during the walk. Second, we query the closest point $y^{\partial V}$ on the deterministic boundary of the medium and previously sampled particles. Then, we select the closest of these two points to x , $y^{\partial\Omega} := \text{closest}(x, \{y^{\partial\Omega}, y^{\partial V}\})$. After sampling, we add to M a new empty sphere (b), and a new particle if $y^{\partial\Omega}$ was not on the boundary of the medium or previously sampled particles (c).

and augment its geometry to include the sampled particles:

$$\lambda_k(x) := \lambda(x | E(M_k)), \quad \partial V_k := \partial(V \setminus C(M_k)). \quad (21)$$

Second, we sample $y_k \sim p_k^{\text{cp}}(\cdot | M_k)$ by: using Algorithm 1 with λ_k to sample a closest point $y_k^{\partial\Omega}$ on the stochastic boundary $\partial\Omega$; using a closest point query to determine the closest point $y_k^{\partial V}$ on the deterministic boundary ∂V_k ; then setting $y_k := \text{closest}(x_k, \{y_k^{\partial\Omega}, y_k^{\partial V}\})$. Algorithm 4 summarizes this procedure. Used with empty memory ($M := \emptyset$), this procedure is equivalent to how we sample $y_0 \sim p_0^{\text{cp}}(\cdot)$ at the initial step of VWoS (Equation (16)).

6 Volumetric walk on stars

We now turn our attention to a generalization of the BVP (1) that prescribes mixed Dirichlet and Neumann boundary conditions. Our discussion is brief, as most of the concepts we introduced in Sections 3 and 5 for the Dirichlet-only case extend to this case.

6.1 Boundary value problem

We focus on estimating solutions to the BVP:

$$\begin{aligned} \Delta u(x) &= 0 & \text{in } \Omega, \\ u(x) &= g(x) & \text{on } \partial\Omega^{\text{Dir}}, \\ \frac{\partial u}{\partial n}(x) &= 0 & \text{on } \partial\Omega^{\text{Neu}}. \end{aligned} \quad (22)$$

Compared to Equation (1), here we partition $\partial\Omega$ into a subset $\partial\Omega^{\text{Dir}}$ where the solution has prescribed values u (Dirichlet boundary conditions), and a subset $\partial\Omega^{\text{Neu}}$ where the solution has prescribed normal derivatives $\partial u / \partial n$ (Neumann boundary conditions).

6.2 Walk on stars

When the domain Ω is deterministic, the *walk on stars* (WoSt) algorithm [Sawhney et al. 2023; Miller et al. 2024b] estimates the BVP solution u in a manner analogous to WoS—recursive single-sample Monte Carlo estimation of an appropriate BIE. In particular, instead

Algorithm 3 Implementation of memory.

```

1: struct MEMORY
2:   attribute  $\mathcal{E}$        $\triangleright$ List of center-radius pairs for empty balls
3:   attribute  $C$        $\triangleright$ List of centers of sampled particles
4:    $\triangleright$  Initialize empty memory
5:   function INITIALIZE()
6:      $\mathcal{E}, C \leftarrow \emptyset$ 
7:   end function
8:    $\triangleright$  Add an empty ball of center  $x$  and radius  $r$ 
9:   function ADDEMTYSPPHERE( $x, r$ )
10:     $\mathcal{E} \leftarrow \mathcal{E} \cup \{(x, r)\}$ 
11:  end function
12:   $\triangleright$  Add a sampled particle of center  $c$ 
13:  function ADDSAMPLEDPARTICLE( $c$ )
14:    if  $c'$  in  $C$  then return       $\triangleright$ Skip if  $c$  already in list
15:     $C \leftarrow C \cup \{c\}$ 
16:  end function
17:   $\triangleright$  Update memory given walk point  $x$  and its closest point  $y$ 
18:  function UPDATE( $x, y$ )
19:     $\triangleright$  Add walk step to list of empty balls
20:    ADDEMTYSPPHERE( $x, \|x - y\|$ )
21:     $\triangleright$  Add sampled particle if  $y$  is not on medium boundary
22:    if  $y \notin \partial V$  then ADDSAMPLEDPARTICLE( $y + R \text{ dir}(x, y)$ )
23:  end function
24:   $\triangleright$  Perform a containment query on dilated empty space
25:  function ISINSIDEDILATEDEMPTYSPHERES( $x, R$ )
26:    for ( $x', r$ ) in  $\mathcal{E}$  do
27:      if  $\|x - x'\| < r + R$  then return TRUE
28:    return FALSE
29:  end function
30:   $\triangleright$  Perform a closest point query on sampled particle boundary
31:  function GETCLOSESTPOINT( $x$ )
32:     $\triangleright$  Query closest particle center
33:     $c_{\text{nn}} \leftarrow \text{NEARESTNEIGHBOR}(x, C)$ 
34:     $\triangleright$  Return corresponding closest point on particle boundary
35:    return  $c_{\text{nn}} - R \text{ dir}(x, c_{\text{nn}})$ 
36:  end function

```

of Equation (2) for the Dirichlet-only problem, WoSt starts from the following BIE for the mixed Dirichlet-Neumann problem (22):

$$u(x) = \int_{\partial\text{St}(x, r^{\text{St}}(x))} P(r^{\text{St}}(x)) u(y) dA(y). \quad (23)$$

The integration domain is the boundary of a *star-shaped region* $\text{St}(x, r^{\text{St}}(x))$ defined as follows [Sawhney et al. 2023, Section 4]: For any point $x \in \Omega$, we denote by $y^{\text{Dir}}(x) := \text{closest}(x, \partial\Omega^{\text{Dir}})$ its closest point on $\partial\Omega^{\text{Dir}}$, and by $y^{\text{sil}}(x)$ its closest point on the *visibility silhouette* of $\partial\Omega^{\text{Neu}}$. We denote by $r^{\text{Dir}}(x) := \|x - y^{\text{Dir}}(x)\|$, $r^{\text{sil}}(x) := \|x - y^{\text{sil}}(x)\|$ the distances to these points, and by $r^{\text{St}}(x) := \min\{r^{\text{Dir}}(x), r^{\text{sil}}(x)\}$ their minimum. Then, the star-shaped region equals $\text{St}(x, r^{\text{St}}(x)) := B(x, r^{\text{St}}(x)) \cap \Omega$. Importantly, $\partial\text{St}(x, r^{\text{St}}(x))$

Algorithm 4 Closest point sampling with memory.

```

1:  $\triangleright$  Implementation of conditional density in Equation (21)
2: struct CONDITIONALDENSITY
3:   attributes  $\lambda, R, M$   $\triangleright$  PBM parameters and memory
4:   function INITIALIZE( $\lambda', R', M'$ )  $\triangleright$  Initialize attributes
5:      $\lambda \leftarrow \lambda', R \leftarrow R', M \leftarrow M'$ 
6:   end function
7:  $\triangleright$  Evaluate conditional density at query point  $x$ 
8:   function EVALUATE( $x$ )
9:     if  $M$ .ISINSIDE DILATED EMPTY SPHERES( $x, R$ ) then
10:       return 0  $\triangleright x$  is inside empty space
11:     else return  $\lambda$ .EVALUATE( $x$ )
12:   end function
13:  $\triangleright$  Implementation of updated medium geometry in Equation (21)
14: struct UPDATEDMEDIUMGEOMETRY
15:   attributes  $\partial V, M$   $\triangleright$  Medium geometry and memory
16:   function INITIALIZE( $\partial V', M'$ )  $\triangleright$  Initialize attributes
17:      $\partial V \leftarrow \partial V', M \leftarrow M'$ 
18:   end function
19:  $\triangleright$  Perform a closest point query on updated medium boundary
20:   function GETCLOSESTPOINT( $x$ )
21:      $\triangleright$  Query closest points on medium and particle boundaries
22:      $y \leftarrow \partial V$ .GETCLOSESTPOINT( $x$ )
23:      $y' \leftarrow M$ .GETCLOSESTPOINT( $x$ )
24:     return  $\text{closest}(x, \{y, y'\})$   $\triangleright$  Return closest of two points
25:   end function
Input: A query point  $x$ , a majorant density  $\tilde{\lambda}$ , a struct implementing the PBM density  $\lambda$ , the PBM particle size  $R$ , a struct implementing memory  $M$ .
Output: Closest point  $y$ .
26: function SAMPLECLOSESTPOINTWITHMEMORY( $x, \tilde{\lambda}, R, M$ )
27:    $\triangleright$  Create conditional density
28:    $\lambda_{\text{cond}} \leftarrow \text{CONDITIONALDENSITY.INITIALIZE}(\tilde{\lambda}, R, M)$ 
29:    $\triangleright$  Create updated medium geometry
30:    $\partial V_{\text{cond}} \leftarrow \text{UPDATEDMEDIUMGEOMETRY.INITIALIZE}(\partial V, M)$ 
31:    $\triangleright$  Sample a closest point conditioning on empty balls
32:    $y^{\partial O} \leftarrow \text{SAMPLECLOSESTPOINT}(x, \tilde{\lambda}, \lambda_{\text{cond}}, R)$ 
33:    $\triangleright$  Query closest point on updated deterministic boundary
34:    $y^{\partial V} \leftarrow \partial V_{\text{cond}}$ .GETCLOSESTPOINT( $x$ )
35:    $\triangleright$  Return the closest of sampled and deterministic closest points
36:    $y \leftarrow \text{closest}(x, \{y^{\partial O}, y^{\partial V}\})$ 
37:   return  $y$ 

```

can include parts of $\partial\Omega^{\text{Neu}}$ but not $\partial\Omega^{\text{Dir}}$ —except, potentially, for the closest point $y^{\text{Dir}}(x)$ when $r^{\text{St}}(x) = r^{\text{Dir}}(x)$.

Using recursive single-sample Monte Carlo estimation on Equation (23), WoSt performs a random walk x_0, x_1, \dots that terminates using the same ε -shell approximation as in WoS. At each walk point x_k , the expression for the WoSt estimator is the same as in Equation (3), except replacing $r_k^{\partial\Omega}$ with $r_k^{\text{St}} := r^{\text{St}}(x_k)$. At each walk step, WoSt performs a closest point query to determine $y_k^{\text{Dir}} := y^{\text{Dir}}(x_k)$,

a *closest silhouette point query* to determine $y_k^{\text{sil}} := y^{\text{sil}}(x_k)$, and directional sampling to determine the next walk point x_{k+1} .

6.3 Participating media

We next consider the mixed Dirichlet-Neumann BVP (22) when the domain Ω includes stochastic microparticle geometry. We follow the setup of Section 3.4, defining Ω as in Equation (4), and assume that $O \sim \Phi(\lambda, R)$. We further assume that the Dirichlet boundary coincides with the deterministic medium boundary, and the Neumann boundary comprises the boundary of the particles; that is:

$$\partial\Omega^{\text{Dir}} := \partial V \text{ and } \partial\Omega^{\text{Neu}} := \partial O. \quad (24)$$

We choose this problem specification to simplify exposition, but we can extend our method to the case where $\partial\Omega^{\text{Dir}}$ and $\partial\Omega^{\text{Neu}}$ each include both deterministic and stochastic boundaries.

Given the close similarity between Equations (2) and (23), and between WoS and WoSt, we can adapt the derivation in Section 5.1 exactly analogously to the mixed Dirichlet-Neumann BVP. The result is a *volumetric walk on stars* (VWoSt) estimator that has the same form as the VWoS estimator (19), but with r_k being the radius used to form the star-shaped region $\text{St}(x_k, r_k)$. As explained in Section 6.2, r_k equals the minimum of $\|x_k - y_k^{\text{Dir}}\|$ and $\|x_k - y_k^{\text{sil}}\|$. From Equation (24), y_k^{Dir} is the *deterministic* closest point $y_k^{\partial V}$ on ∂V , and y_k^{sil} is the *random* closest silhouette point on ∂O . VWoSt determines y_k^{sil} through *closest silhouette point sampling*, conditional on the memory M_k of all such points sampled in previous steps.

6.4 Closest silhouette point sampling with memory

As the microparticle geometry O comprises spherical particles of the same radius R , the closest silhouette point on ∂O will lie on the same particle as the closest point on ∂O (Figure 8). Thus, we perform closest point sampling with memory as in Section 5.2 to determine the closest point $y_k^{\partial O}$, then compute from it y_k^{sil} and r_k^{sil} analytically. From these values, we compute the radius r_k for the star-shaped region $\text{St}(x_k, r_k) = B(x_k, r_k) \cap O$.

Because r_k can be greater than the shortest distance $r_k^{\partial O} = \|x_k - y_k^{\partial O}\|$ to ∂O , forming the star-shaped region $\text{St}(x_k, r_k)$ requires also determining any additional particles that are closer to x_k than r_k . We do so by using Algorithm 1 with conditional density $\lambda(\cdot | E(M_k))$ (Equation (21)) to continue sampling particles beyond the closest one, in order of increasing distance from x_k , until we exceed the distance r_k . As these particles become fixed geometry for subsequent steps, we include them in the memory M_k —and in particular, in the sampled-particle memory $C(M_k)$ (Equation (20)).




Figure 8. Sampling to form a star-shaped region.

7 Experimental evaluation

Our experimental evaluation includes: 1. comparisons of VWoS and VWoSt with ensemble averaging to both validate unbiasedness and assess performance (Section 7.2); 2. comparisons with homogenization methods (Section 7.2); and 3. analysis of the impact of memory (Section 7.3). Our experiments demonstrate that VWoS and VWoSt

provide performance and accuracy benefits over both ensemble averaging and homogenization, though the magnitude of these benefits depends on the experimental setup.

Implementation details. We implement VWoS and VWoSt in ZOMBIE [Sawhney and Miller 2024], making only minor modifications to its WoS and WoSt routines. Our implementation supports mixed boundary conditions on the medium boundary ∂V , and both Dirichlet and Neumann conditions on the microparticle geometry boundary ∂O . We represent density as either an analytic function or a dense grid. When using analytic functions, we compute a majorant for the largest deterministic empty ball at each point; when using a grid, we compute a global majorant. We set the ε -shell parameter to be at least one order of magnitude smaller than the particle radius. Tables 1 and 2 provide details (e.g., medium properties, boundary conditions, algorithmic parameters) for each experiment.

We implement ensemble averaging also in ZOMBIE, with the following optimizations: 1. We share sampled particle configurations across evaluation points, to amortize sampling costs. 2. We construct an, also shared, bounding value hierarchy (BVH) for each configuration, to accelerate closest point queries. For the BVH, we use FCPW [Sawhney 2021] and extend it to support spherical primitives.

7.1 Comparison to ensemble averaging

We compare our VWoS and VWoSt algorithms against ensemble averaging, in BVPs involving various volume geometries and medium parameters. These comparisons aim to both validate the consistency and unbiasedness of our algorithms—by ensuring that their mean solution estimates match those from ensemble averaging—and quantify the performance improvements they provide.

Volumetric walk on spheres. Figure 9 shows comparisons between VWoS and ensemble averaging, in BVPs with Dirichlet-only boundary conditions. In the mushroom domain of Figure 9(a–c), we vary particle size and density across nearly two orders of magnitude. In the ginseng root domain of Figure 9(d–f), we evaluate the solution near geometric features whose scale varies by a factor of 4.

In all experiments, VWoS estimates match the mean solution estimates from ensemble averaging. Moreover, VWoS is more than 3× faster than ensemble averaging in terms of the time it takes to perform the same number of walks—and thus reach the same variance—when accounting for both particle configuration sampling and solve time in ensemble averaging. This performance difference becomes more stark in experiments requiring sparse evaluation points or dense media; we revisit this point in Section 8.1 where, depending on the density of evaluation points, VWoS is one-to-several orders of magnitude faster than ensemble averaging. The slow performance of ensemble averaging is because of three reasons:

1. Ensemble averaging has poor output sensitivity, as it must sample entire particle configurations even to compute the solution at only one point. Thus, ensemble averaging wastes considerable compute sampling particles in large parts of the volume that have little to no impact to its output.
2. Ensemble averaging has poor sample amortization, as it must sample a new particle configuration for each walk, or at least

batch of walks. Sharing sampled particle configurations across evaluation points mitigates this issue only partially.

3. Ensemble averaging has poor geometric query performance, as it must perform closest point queries against entire particle configurations. Even with logarithmic-complexity query implementations (e.g., with a BVH), large particle numbers in dense media introduce considerable computational overhead.

By contrast, in VWoS, each walk samples few particles, on demand, and only in its locality; the result is good output sensitivity, no sample waste, and no overhead in geometric queries. Thus, VWoS greatly improves performance relative to ensemble averaging, analogous to the performance improvements volume rendering algorithms provide over ensemble averaging for light transport simulation.

Volumetric walk on stars. Figure 10 shows comparisons between VWoSt and ensemble averaging, in BVPs with mixed Dirichlet-Neumann boundary conditions. We perform experiments on the same scene using two densities, which are Gaussian-shaped along one dimension and constant along the other two.

In both cases, VWoSt matches the estimates of ensemble averaging and improves performance (measured as in the VWoS experiments above). In Figure 10(a), the concentrated density impacts both methods: it reduces configuration sampling efficiency in ensemble averaging, and requires large majorants for closest silhouette point sampling in VWoSt. In Figure 10(b), where density is less concentrated, sampling configurations in ensemble averaging is much cheaper. However, walks are more expensive in ensemble averaging than in VWoSt, because they perform closest silhouette point queries against entire configurations.

7.2 Comparison to homogenization

Homogenization methods [Giunti et al. 2018] transform a BVP with stochastic microparticle geometry into a *homogenized BVP* that involves a modified PDE in deterministic geometry. The solution u_h of the homogenized BVP converges to the mean solution \bar{u} of the original BVP only *asymptotically* at the limit of infinitesimally small and infinitely dense particles ($R \rightarrow 0$, $\lambda \rightarrow \infty$, while λR remains constant). For the BVP (1), the homogenized BVP is:⁶

$$\begin{aligned} \Delta u_h(x) - 4\pi\lambda R u_h(x) &= 0 & \text{in } V, \\ u_h(x) &= g(x) & \text{on } \partial V. \end{aligned} \quad (25)$$

The homogenized BVP can then be solved efficiently using standard WoS. However, for any finite values of R and λ , homogenization provides *biased* estimates of the mean solution, with bias increasing as R or λ deviate more from the asymptotic case. Bias can also vary at different domain points, e.g., near geometrically thin versus thick parts. By contrast, VWoS is *unbiased* and provides accurate estimates *robustly* across particle properties and domain points.

We demonstrate these advantages experimentally in Figure 9. In the mushroom domain Figure 9(a–c), increasing the particle size leads to noticeable bias in the homogenization solution. In the ginseng root domain Figure 9(d–f), for a fixed particle radius, bias is higher at thin parts of the root (where the particle radius is comparable to the domain size) than at thick parts (where the particle radius

⁶This homogenization procedure is distinct from using a homogeneous participating medium to perform closest point sampling (Algorithm 1).

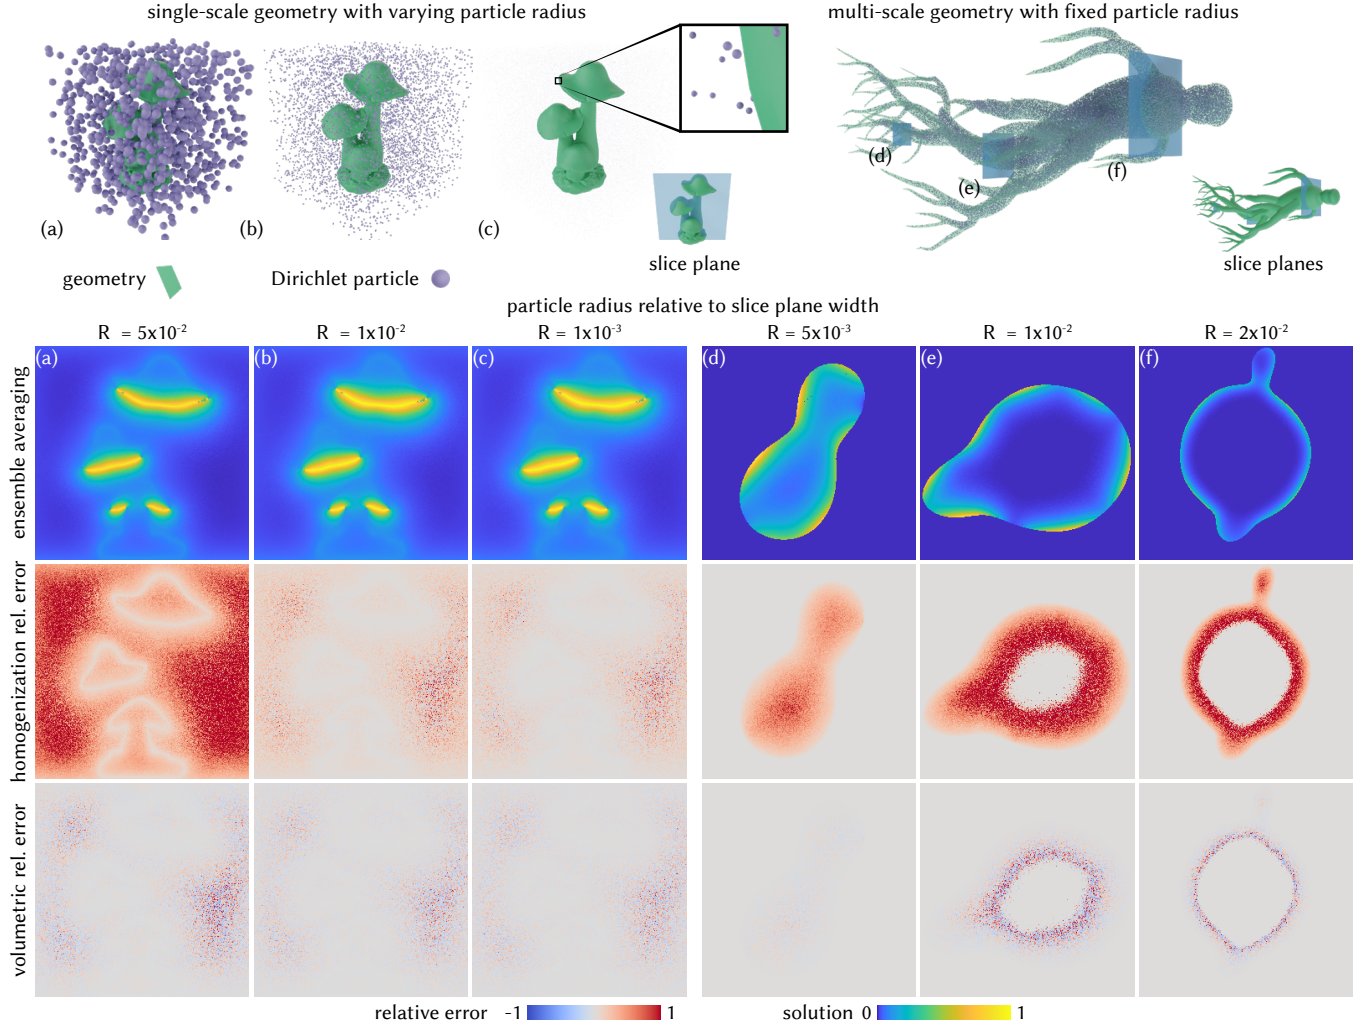


Figure 9. We compare the outputs of VVoS (*fourth row*) and homogenization (*third row*) to the reference mean solution (*second row*) computed with ensemble averaging, in BVPs with Dirichlet-only boundary conditions. Across a range of medium parameters and boundaries, VVoS reliably produces unbiased mean solution estimates, whereas homogenization introduces noticeable bias as the particle size increases (a, b), or at geometrically thin parts of the volume (d, e). For each experiment, we visualize (*first row*) a representative sampled configuration of the microparticle geometry.

is much smaller than the domain size). Additionally, bias is higher near the medium boundary than away from it. Thus, even though homogenization is $2 - 20\times$ faster than VVoS in these experiments, the solution estimates it produces have considerable bias that is hard to control across problem settings and complex domains.

7.3 Analysis of impact from memory

Walk memory M is a consequence of the law of total expectation in the derivation of the volumetric BIE Equation (18). Memory arises likewise in the derivation of the volume rendering equation (Appendix B); however, a memoryless approximation is commonplace in volume rendering algorithms [Bitterli et al. 2018], especially for exponential media. Motivated by this precedent, we experimented

with a memoryless variant of VVoS (i.e., always performing unconditional closest point sampling). We found that this approximation leads to highly inaccurate solution estimates *and* much worse runtime performance: Walks rarely terminate inside the volume while still sampling step sizes that decrease super-exponentially with medium density; hence average walk length increases dramatically. Thus, memory appears to be more important for our algorithms than for rendering ones—we expand on this point in Appendix B.

To help assess the performance impact of memory, we first visualize in Figure 11 statistics relating walk length to the sizes (numbers of balls) N_k^E of the empty-ball memory $E(M_k)$, and N_k^C of the sampled-particle memory $C(M_k)$, for the experiments in Section 7.1. N_k^E always equals the walk length k —each walk step adds a new

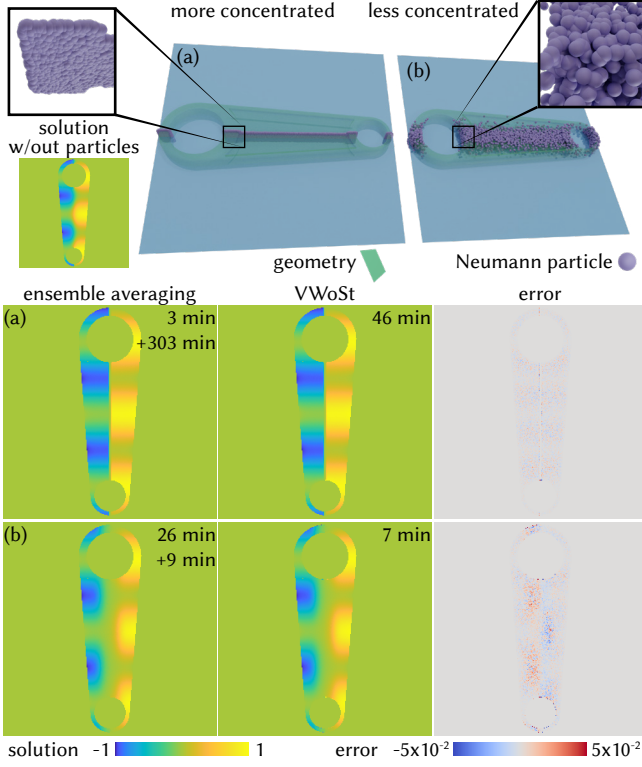


Figure 10. We compare the output of VWoSt (second column) to the reference mean solution (first column) computed with ensemble averaging, in BVPs with Neumann boundary conditions on the microparticle geometry. We also report the runtimes of VWoSt and ensemble averaging (the + numbers are the time to sample particle configurations for ensemble averaging). We experiment with more concentrated (second row) and less concentrated (third row) particle densities. The error images (third column) show that VWoSt correctly estimates the mean solution nearly $5\times$ faster than ensemble averaging. For each experiment, we visualize (first row) a representative sampled configuration of the microparticle geometry.

empty ball. By contrast, the growth of N_k^C with walk length depends on the density and type of boundary conditions on the particles: For Dirichlet boundary conditions, each walk step adds at most one—and often zero—new particle during a closest point query (Section 5.2). For Neumann boundary conditions, each walk step may add several particles during a closest silhouette point query (Section 6.4). In both cases, higher particle density leads to faster growth of N_k^C .

Inspired by Seyb et al. [2024], we also experiment with a variant of VWoSt that uses *finite* memory—only storing the most recent empty ball and sampled particle—to mitigate the overhead of conditional sampling as memory grows. Figure 12 compares this variant with standard VWoSt that uses full memory. Using finite memory introduces significant bias and provides marginal to no performance improvements, suggesting that finite memory of any size does not offer a favorable bias-performance trade-off—larger sizes of finite memory will provide only smaller performance improvements. Performance

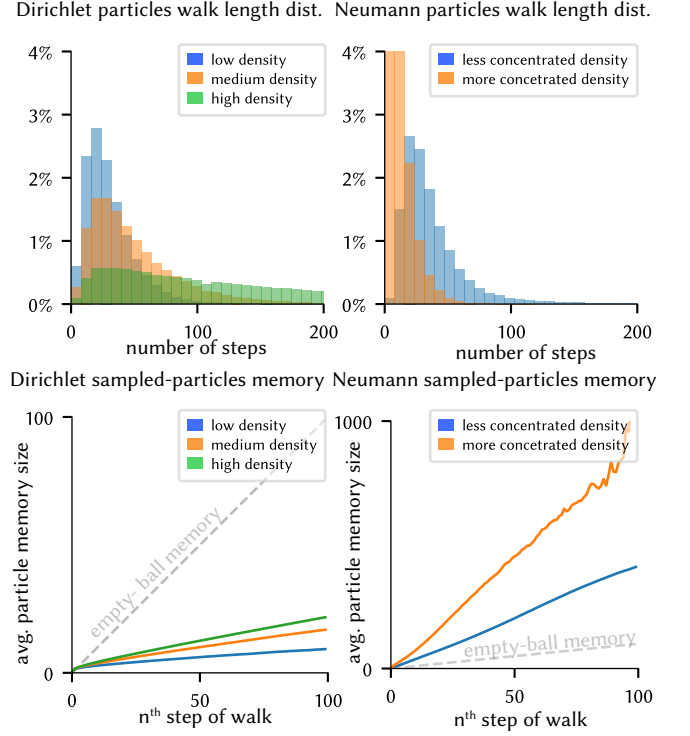


Figure 11. Statistics for walk length (first row) and memory size (second row) for the mushroom (Figure 9(a–c), Dirichlet boundary conditions) and connector (Figure 10(a–b), Neumann boundary conditions) domains. Though the size of the empty-ball memory always equals walk length, the size of the sampled-particle memory can grow slower (in the Dirichlet case) or faster (in the Neumann case) than walk length. In both cases, increased density leads to longer walks and faster growth of sampled-particle memory.

does not improve more significantly because computational cost is dominated by geometric queries against the medium boundary, which are repeated at each walk step even if step size is very small. The overhead from these repeated queries becomes worse as particle density, and therefore walk length, increases (Figure 11, top row). However, if this overhead is reduced—for example, by leveraging the spatial coherence of steps in volumetric walks—then it may become useful to reconsider the use of finite memory. Additionally, finite memory could help improve parallelism on memory-constrained GPU devices, due to improved memory management. We defer further exploration of finite memory to future work.

8 Simulation of natural phenomena

To showcase the ability of our algorithms to deal with a variety of simulation settings with complex microparticle geometry, we use them for simulation of two model natural phenomena.

8.1 Electrostatics near biological membranes

In Figure 1, we use VWoSt to model electrostatic potentials in systems with complex cellular geometries. Simulating such potentials and their effect on biomolecules is common in biochemistry [Davis

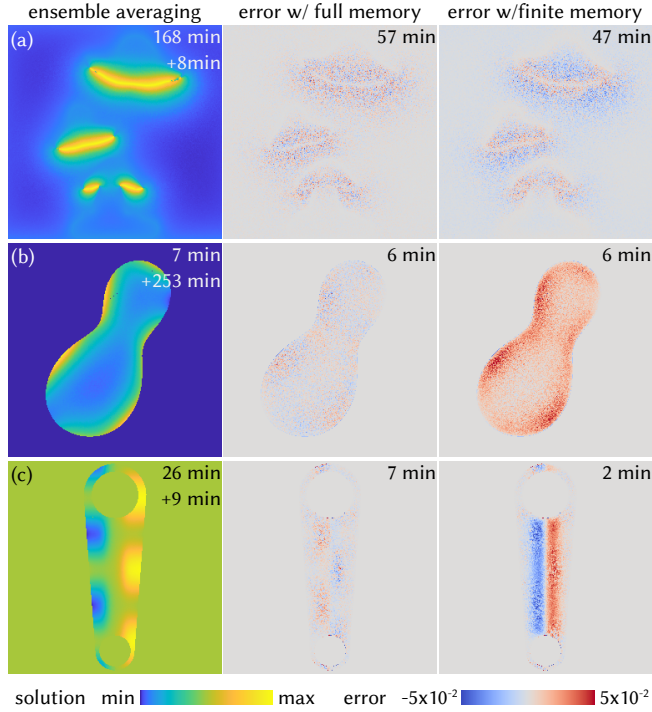


Figure 12. We compare the bias-performance trade-off of finite memory of size one (third column) versus full memory (second column) using ensemble averaging (first column) for reference. We show error images and report runtimes for each method (the + numbers are the time to sample particle configurations for ensemble averaging). Using finite memory improves runtime only marginally (a, c) or not at all (b), yet always introduces significant bias in solution estimates. These results suggest that memory is crucial for estimation accuracy, and finite memory does not offer a favorable bias-performance trade-off for either VWoS (a, b) or VWoSt (c).

and McCammon 1990; Gilson et al. 1988]. However, accurately accounting for the complex molecular geometry involved in these simulations poses significant computational challenges. To manage this complexity, homogenization approaches abstract away the fine details of molecular geometry using simplified models based on asymptotic cases. One such model, the Debye-Hückel framework [Andelman 1995], is commonly used to describe the screening of electrostatic potentials by ions. Though effective in some scenarios, these models often lead to inaccurate predictions, especially near membrane surfaces where the size of ions relative to the geometry becomes critical [Davis and McCammon 1990].

The setup in Figure 1 assumes constant particle density (expected 1.9 million particles per configuration). The concentrations of sodium (Na^+) and chloride (Cl^-) ions depend on the distance to the nearest charged surface: Cl^- ions are more concentrated near positively charged surfaces, and Na^+ ions are more concentrated near negatively charged surfaces. The boundary conditions ensure neutrality (zero charge) away from the surface, and the charge reaches a maximum value of ± 0.25 V on the surface. Compared to running WoS in this setup with the participating medium removed, VWoS

increases runtime by only 15%. In return for this slight overhead, VWoS produces solutions that greatly improve modeling accuracy of electrostatic potentials due to particle-membrane interactions.

This experiment highlights the advantages of VWoS over ensemble averaging, and in particular its ability to simulate complex microparticle geometry without the need to repeatedly sample entire particle configurations. For example, in Figure 1, computing the electrostatic potential on a 256×256 slice plane takes 11 sec with ensemble averaging, and only 1 sec with VWoS—a speedup of more than an order of magnitude. The performance improvement is even greater when solution estimates are needed only along a line (e.g., to compute the screening plots in the insets), in which case VWoS is more than 10,000 \times faster than ensemble averaging.

8.2 Photochemical effect in clouds

As we discuss in Section 4 and Appendix B, the PBM we use to derive VWoS and VWoSt also underlies volume rendering algorithms such as volumetric path tracing (VPT) [Novák et al. 2018]. These algorithms are commonplace in scientific applications that model light transport inside participating media such as clouds or tissue. Combining VWoS and VWoSt with volume rendering creates opportunities for *coupled simulation* of light transport and other physical phenomena (e.g., diffusive effects) in these media. Such a combination continues recent work on coupled Monte Carlo simulation in deterministic geometry—e.g., combining WoS and path tracing to simulate multimodal heat transfer [Bati et al. 2023].

As an example of such an opportunity, in Figure 13, we couple VWoS and VPT to simulate a simplified atmospheric photochemical system. In such systems, sunlight drives the production of pollutants such as ozone, which then diffuse in the atmosphere [Madronich and Flocke 1999; Seinfeld and Pandis 2016]. Clouds in particular play a prominent role in atmospheric photochemistry [Hall et al. 2018; Bianco et al. 2020], but simulation of the photochemical effect in them is challenging, due to their incredibly complex microparticle geometry (billions of water droplets [Gryspeerdt et al. 2022]).

Coupled volumetric rendering and PDE simulation can help overcome this challenge. As a proof-of-concept demonstration, we use a cloud model from FBRT-v3 [Pharr et al. 2018] to set up a photochemical-effect simulation problem. We use VPT to model sunlight (light transport), and VWoS to model ozone concentration as a Dirichlet Laplace BVP (1)—an approximation to more accurate models of ozone diffusion [Hanna et al. 1982]. At each point in the cloud, we set the Dirichlet boundary data equal to the mean incident *fluence* \bar{H}_i (W m^{-2}) due to a directional light source modeling the Sun.

Altogether, we can estimate ozone concentration inside the cloud by running VWoS walks that, upon termination, switch to VPT light paths to estimate incident fluence. Importantly, when tracing these paths, we must account for memory about the medium accumulated during the walk: as light transport and diffusion occur in the same medium, the fluence estimation must be conditioned on the same information as the ozone concentration estimation, i.e., $\bar{u}(x_k | M_k) = \bar{H}_i(x_k | M_k)$. We handle this conditional evaluation by thinning the medium density with respect to previously sampled empty balls, and occluding rays hitting previously sampled particles. As Figure 13 shows, accounting for memory when coupling




Figure 13. We couple VVoS and VPT to model diffusion and light transport (resp.) in a proof-of-concept atmospheric photochemical system: a cloud [Pharr et al. 2018] inside which fluence due to multiply scattered sunlight generates ozone, which then diffuses throughout the cloud (*first and second column*). We estimate ozone concentration (*fourth and fifth column*) by performing VVoS walks that, upon termination, use VPT paths to estimate a Dirichlet boundary condition equal to fluence (*third column*). As both the VVoS walk and VPT path interact with the same microparticle geometry, the memory accumulated from the walk must carry over to the path. Correctly coupling with memory has a non-trivial impact on the estimated ozone concentration (*fifth column*).

VVoS and VPT captures shadowing effects that drastically impact the mean solution. Additionally, if we scale the heterogeneous cloud density, the ozone concentration changes in a manner resembling light transport variation in optically dense versus thin media.

9 Limitations and future work

We introduced the problem of *PDE simulation in participating media* as a framework for computationally modeling natural phenomena involving complex microparticle geometry. We described how to rigorously model such geometry stochastically using the *Poisson Boolean model*, which enabled us to develop sampling procedures that serve as stochastic counterparts of common geometric queries. We additionally developed the *volumetric walk on spheres* (VVoS) and *volumetric walk on stars* (VVoSt) algorithms that use these sampling procedures for Monte Carlo estimation of the solution to this simulation problem. Our algorithms directly generalize the standard walk on spheres and walk on stars algorithms, maintaining their structure and sharing their attractive features.

Clearly, these contributions are just first steps, aimed to set the groundwork for exploring the interaction between PDEs and participating media. By bridging ideas from rendering, PDE simulation, and stochastic geometry, the problem setting we introduced presents opportunities for future research in several exciting directions.

Sampling algorithms. A core component of our VVoS algorithm is closest point sampling with the thinning procedure of Algorithm 1. This procedure often becomes the performance bottleneck of our algorithm, for example in volumes with highly concentrated density. A similar bottleneck behavior occurs in volume rendering, where an analogous algorithm, *delta tracking*, is used for free-flight distance sampling [Coleman 1968; Raab et al. 2006]—we elaborate on this analogy in Appendix B. These performance issues have motivated

extensive rendering research on improved sampling algorithms (as well as algorithms for the closely related problem of transmittance estimation [Georgiev et al. 2019; Kettunen et al. 2021; Kutz et al. 2017]), for example by making better use of intermediate sampled points [Novák et al. 2014] or using multiple-importance sampling [Miller et al. 2019]. Other approaches address performance issues by using adaptive or progressive majorants [Misso et al. 2023] instead of a global one. All these approaches can potentially be adapted to VVoS, and help bring about similar performance improvements.

Microparticle geometry models. We focused on the simplest form of a Boolean model that uses independent, spherical, and fixed-size particles. However, many phenomena involve participating media with particles that violate these assumptions: they can have non-spherical particles (biological materials such as seeds and pollen, soil and other sediment); particles of wideband size distributions (powders, sand grains); or non-independent particles that form regular structures (crystals), or can repel and attract each other (ions, colloids). Extending our methods to these more varied participating media requires using more general Boolean models for microparticle geometry. Poisson Boolean models with non-spherical particles or particles of random sizes share many of the same properties we used in Section 4 [Last and Penrose 2017, Chapter 17]. Thus, we expect that simulating these models should require only modest modifications to VVoS. Other Boolean models sample correlated particle centers from non-Poisson point processes Chiu et al. [2013, Chapter 5] (for example, Matérn [Stoyan and Stoyan 1985] or Gibbs [Sabatini and Villa 2024]). Simulation of such models is more challenging, due to the lack of efficient closest point sampling procedures. Efforts to generalize VVoS to more general Boolean models can benefit from insights from volume rendering, where there already exist variants of VPT for participating media with multi-sized [Frisvad et al. 2007],

anisotropic [Jakob et al. 2010; Heitz et al. 2015], or correlated [Jarabo et al. 2018; Bitterli et al. 2018; d'Eon 2018, 2019] particles.

Other stochastic geometry models. Recent work in computer graphics has shown it can be beneficial to model macroscopic object-level geometry stochastically, for example to account for uncertainty in 3D acquisition or facilitate shape optimization [Vicini et al. 2021; Miller et al. 2024a]. Particularly appealing in this context are stochastic implicit surface representations based on Gaussian processes, thanks to their rich mathematical properties [Sheffield 2007] and relationship to mature geometry processing algorithms [Sellán and Jacobson 2022, 2023]. Algorithms for rendering such representations have emerged recently [Seyb et al. 2024], and can help inform research on analogous algorithms for PDE simulation.

More general PDEs. Our volumetric WoS and WoSt algorithms can solve the same types of PDEs as their counterparts for deterministic geometry—we focused on the Laplace equation, but our algorithms are straightforward to apply to Poisson and screened Poisson equations. However, many phenomena suitable for modeling using participating media involve PDEs out of scope for our algorithms, such as linear elasticity [James and Pai 1999], heat conduction and convection [Hahn and Özisik 2012], Stokes flow [Du et al. 2020], and Navier-Stokes flow [Stam 1999]. Recent work has introduced Monte Carlo simulation algorithms for some of these PDEs [Rioux-Lavoie et al. 2022; Sugimoto et al. 2024; Bati et al. 2023], often directly extending WoS or WoSt. Using our theory to develop volumetric variants of these algorithms would greatly expand the type of phenomena we can simulate using participating media.

Acknowledgments

This work was supported by National Science Foundation (NSF) award 2212290; National Institute of Food and Agriculture award 2023-67021-39073; a gift from nTopology; NSF Graduate Research Fellowship DGE2140739 and an NVIDIA graduate fellowship for Bailey Miller; a Packard Foundation Fellowship for Keenan Crane; and Alfred P. Sloan Research Fellowship FG202013153 for Ioannis Gkioulekas. Rohan Sawhney thanks Ken Museth for supporting this work. The protein model (Figure 1) is from user *QuadroFlow* on TURBOSQUID. The ginseng root model (Figures 9(d–f), 12(b)) is from user *3dror* on TURBOSQUID. The mushroom model (Figures 9(a–c), 12(a)) and the connector model (Figures 10, 12(c)) are from the THING10K dataset [Zhou and Jacobson 2016]. The cloud model (Figure 13) is from the PBRT-v3 scene repository [Pharr et al. 2018].

References

Marina Alterman, Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2021. Imaging with Local Speckle Intensity Correlations: Theory and Practice. *ACM Trans. Graph.* 40, 3, Article 30 (July 2021), 22 pages. <https://doi.org/10.1145/3447392>

David Andelman. 1995. Electrostatic properties of membranes: the Poisson-Boltzmann theory. In *Handbook of biological physics*. Vol. 1. Elsevier, 603–642.

Ghada Bakboui and Pieter Peers. 2023. Mean Value Caching for Walk on Spheres. In *Eurographics Symposium on Rendering*. The Eurographics Association.

Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Trans. Graph.* 38, 4, Article 39 (July 2019), 22 pages. <https://doi.org/10.1145/3306346.3322950>

Mégane Bati, Stéphane Blanco, Christophe Coustet, Vincent Eymet, Vincent Forest, Richard Fournier, Jacques Gautrais, Nicolas Mellado, Mathias Paulin, and Benjamin Piau. 2023. Coupling Conduction, Convection and Radiative Transfer in a Single

Path-Space: Application to Infrared Rendering. *ACM Trans. Graph.* 42, 4, Article 79 (jul 2023), 20 pages. <https://doi.org/10.1145/3592121>

Bruce J Berne and Robert Pecora. 2000. *Dynamic light scattering: with applications to chemistry, biology, and physics*. Courier Corporation.

Angela Bianco, Monica Passananti, Marcello Brigante, and Gilles Mailhot. 2020. Photochemistry of the cloud aqueous phase: A review. *Molecules* 25, 2 (2020), 423.

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 (Dec. 2018), 17 pages. <https://doi.org/10.1145/3272127.3275103>

Hendrik C Brinkman. 1949. A calculation of the viscous force exerted by a flowing fluid on a dense swarm of particles. *Flow, Turbulence and Combustion* 1, 1 (1949), 27–34.

David C Brydges and Paul Federbush. 1980. Debye screening. *Communications in Mathematical Physics* 73, 3 (1980), 197–246.

Luis A Caffarelli and Antoine Mellet. 2009. Random homogenization of an obstacle problem. In *Annales de l'IHAP Analyse non linéaire*, Vol. 26. 375–395.

Carmen Calvo-Jurado, Juan Casado-Díaz, and Manuel Luna-Laynez. 2015. Homogenization of the Poisson equation with Dirichlet conditions in random perforated domains. *J. Comput. Appl. Math.* 275 (2015), 375–381.

Sung Nok Chiu, Dietrich Stoyan, Wilfrid S Kendall, and Joseph Mecke. 2013. *Stochastic geometry and its applications*. John Wiley & Sons.

Doina Cioranescu and Fracis Murat. 1982. Un terme étrange venu d'ailleurs I & II [A strange term coming from nowhere I & II]. *Nonlinear partial differential equations and their applications, Collège de France Séminar Vols. II and III, Vols 60, 70* (1982), 98–138.

Doina Cioranescu and François Murat. 1997. A strange term coming from nowhere. *Topics in the mathematical modelling of composite materials* (1997), 45–93.

Wayne 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 (1968), 76–81.

Martin Costabel. 1987. Principles of boundary element methods. *Computer Physics Reports* 6, 1-6 (1987), 243–274.

Malcolm E Davis and J Andrew McCammon. 1990. Electrostatics in biomolecular structure and dynamics. *Chemical Reviews* 90, 3 (1990), 509–521.

Auguste De Lambilly, Gabriel Benedetti, Nour Rizk, Chen Hanqi, Siyuan Huang, Junnan Qiu, David Louapre, Raphael Granier De Cassagnac, and Damien Rohmer. 2023. Heat Simulation on Meshless Crafted-Made Shapes. In *Proceedings of the 16th ACM SIGGRAPH Conference on Motion, Interaction and Games* (Rennes, France) (MIG '23). Association for Computing Machinery, New York, NY, USA, Article 9, 7 pages. <https://doi.org/10.1145/3623264.3624457>

Eugene d'Eon. 2018. A reciprocal formulation of nonexponential radiative transfer. 1: Sketch and motivation. *Journal of Computational and Theoretical Transport* 47, 1-3 (2018), 84–115.

Eugene d'Eon. 2019. A reciprocal formulation of nonexponential radiative transfer. 2: Monte Carlo estimation and diffusion approximation. *Journal of Computational and Theoretical Transport* 48, 6 (2019), 201–262.

Mathieu Desbrun, Roger D Donaldson, and Houman Owahdi. 2013. Modeling across scales: Discrete geometric structures in homogenization and inverse homogenization. *Multiscale analysis and nonlinear dynamics* (2013), 19–64.

Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. 1988. Volume rendering. *SIGGRAPH Comput. Graph.* 22, 4 (June 1988), 65–74. <https://doi.org/10.1145/378456.378484>

Tao Du, Kui Wu, Andrew Spielberg, Wojciech Matusik, Bo Zhu, and Eftychios Sifakis. 2020. Functional optimization of fluidic devices with differentiable stokes flow. *ACM Trans. Graph.* 39, 6, Article 197 (Nov. 2020), 15 pages. <https://doi.org/10.1145/3414685.3417795>

Bruce C Faust. 1994. Photochemistry of clouds, fogs, and aerosols. *Environmental science & technology* 28, 5 (1994), 216A–222A.

Jeppe Revall Frisvad, Niels Jørgen Christensen, and Henrik Wann Jensen. 2007. Computing the scattering properties of participating media using Lorenz-Mie theory. *ACM Trans. Graph.* 26, 3 (July 2007), 60–es. <https://doi.org/10.1145/1276377.1276452>

Iliyan Georgiev, Zackary Misso, Toshiya Hachisuka, Derek Nowrouzezahrai, Jaroslav Krivánek, and Wojciech Jarosz. 2019. Integral formulations of volumetric transmittance. *ACM Trans. Graph.* 38, 6, Article 154 (Nov. 2019), 17 pages. <https://doi.org/10.1145/3355089.3356559>

Michael K Gilson, Kim A Sharp, and Barry H Honig. 1988. Calculating the electrostatic potential of molecules in solution: method and error assessment. *Journal of computational chemistry* 9, 4 (1988), 327–335.

Arianna Giunti and Richard M Höfer. 2019. Homogenisation for the Stokes equations in randomly perforated domains under almost minimal assumptions on the size of the holes. In *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, Vol. 36. Elsevier, 1829–1868.

Arianna Giunti, Richard Höfer, and Juan J. L. Velázquez. 2018. Homogenization for the Poisson equation in randomly perforated domains under minimal assumptions on the size of the holes. *Communications in Partial Differential Equations* 43, 9 (2018), 1377–1412. <https://doi.org/10.1080/03605302.2018.1531425>

- Edward Gryspeerd, Daniel T McCoy, Ewan Crosbie, Richard H Moore, Graeme J Nott, David Painemal, Jennifer Small-Griswold, Armin Sorooshian, and Luke Ziemba. 2022. The impact of sampling strategy on the cloud droplet number concentration estimated from satellite data. *Atmospheric Measurement Techniques* 15, 12 (2022), 3875–3892.
- David W Hahn and M Necati Özisik. 2012. *Heat conduction*. John Wiley & Sons.
- Samuel R Hall, Kirk Ullmann, Michael J Prather, Clare M Flynn, Lee T Murray, Arlene M Fiore, Gustavo Correa, Sarah A Strode, Stephen D Steenrod, Jean-Francois Lamarque, et al. 2018. Cloud impacts on photochemistry: building a climatology of photolysis rates from the Atmospheric Tomography mission. *Atmospheric Chemistry and Physics* 18, 22 (2018), 16809–16828.
- Steven R Hanna, Gary A Briggs, and Rayford P Hosker. 1982. Handbook on Atmospheric Diffusion. *Department of Energy* (1982).
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGX microflake distribution. *ACM Trans. Graph.* 34, 4, Article 48 (July 2015), 11 pages. <https://doi.org/10.1145/2766988>
- Peter Hunter and Andrew Pullan. 2001. Fem/bem notes. *Department of Engineering Science, The University of Auckland, New Zealand* (2001).
- Wenzel Jakob, Adam Arbree, Jonathan T. Moon, Kavita Bala, and Steve Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. In *ACM SIGGRAPH 2010 Papers* (Los Angeles, California) (SIGGRAPH '10). Association for Computing Machinery, New York, NY, USA, Article 53, 13 pages. <https://doi.org/10.1145/1833349.1778790>
- Doug L James and Dinesh K Pai. 1999. Artdefo: accurate real time deformable objects. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 65–72.
- Adrian Jarabo, Carlos Aliaga, and Diego Gutierrez. 2018. A radiative transfer framework for spatially-correlated materials. *ACM Trans. Graph.* 37, 4, Article 83 (July 2018), 13 pages. <https://doi.org/10.1145/3197517.3201282>
- Mohammadreza Kadivar, David Tormey, and Gerard McGranaghan. 2021. A review on turbulent flow over rough surfaces: Fundamentals and theories. *International Journal of Thermofluids* 10 (2021), 100077.
- Markus Kettunen, Eugene D'Eon, Jacopo Pantaleoni, and Jan Novák. 2021. An unbiased ray-marching transmittance estimator. *ACM Trans. Graph.* 40, 4, Article 137 (July 2021), 20 pages. <https://doi.org/10.1145/3450626.3459937>
- Lily Kharevych, Patrick Mullen, Houman Owahdi, and Mathieu Desbrun. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph.* 28, 3, Article 51 (July 2009), 8 pages. <https://doi.org/10.1145/1531326.1531357>
- Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Trans. Graph.* 36, 4, Article 111 (July 2017), 16 pages. <https://doi.org/10.1145/3072959.3073665>
- Günter Last and Mathew Penrose. 2017. *Lectures on the Poisson process*. Vol. 7. Cambridge University Press.
- Aviad Levis, Yoav Y Schechner, Amit Aides, and Anthony B Davis. 2015. Airborne three-dimensional cloud tomography. In *Proceedings of the IEEE International Conference on Computer Vision*. 3379–3387.
- Aviad Levis, Yoav Y Schechner, and Anthony B Davis. 2017. Multiple-scattering microphysics tomography. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6740–6749.
- Peter A W Lewis and Gerald S Shedler. 1979. Simulation of nonhomogeneous Poisson processes by thinning. *Naval research logistics quarterly* 26, 3 (1979), 403–413.
- Zilu Li, Guandao Yang, Xi Deng, Christopher De Sa, Bharath Hariharan, and Steve Marschner. 2023. Neural Caches for Monte Carlo Partial Differential Equation Solvers. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 34, 10 pages. <https://doi.org/10.1145/3610548.3618141>
- Zilu Li, Guandao Yang, Qingqing Zhao, Xi Deng, Leonidas Guibas, Bharath Hariharan, and Gordon Wetzstein. 2024. Neural Control Variates with Automatic Integration. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 10, 9 pages. <https://doi.org/10.1145/3641519.3657395>
- Sasha Madronich and Siri Flocke. 1999. The role of solar radiation in atmospheric chemistry. *Environmental photochemistry* (1999), 1–26.
- Vladimir A Marchenko and Evgueni Ya Khruslov. 2008. *Homogenization of partial differential equations*. Vol. 46. Springer Science & Business Media.
- 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>
- Bailey Miller, Hanyu Chen, Alice Lai, and Ioannis Gkioulekas. 2024a. Objects as Volumes: A Stochastic Geometry View of Opaque Solids. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 87–97.
- Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A null-scattering path integral formulation of light transport. *ACM Trans. Graph.* 38, 4, Article 44 (July 2019), 13 pages. <https://doi.org/10.1145/3306346.3323025>
- Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2023. Boundary Value Caching for Walk on Spheres. *ACM Trans. Graph.* 42, 4, Article 82 (July 2023), 11 pages. <https://doi.org/10.1145/3592400>
- Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2024b. Walkin' Robin: Walk on Stars with Robin Boundary Conditions. *ACM Trans. Graph.* 43, 4, Article 41 (July 2024), 18 pages. <https://doi.org/10.1145/3658153>
- Zackary Misso, Yining Karl Li, Brent Burley, Daniel Teece, and Wojciech Jarosz. 2023. Progressive null-tracking for volumetric rendering. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 31, 10 pages. <https://doi.org/10.1145/3588432.3591557>
- Jonathan T Moon, Bruce Walter, and Stephen R Marschner. 2007. Rendering discrete random media using precomputed scattering solutions. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. 231–242.
- Mervin E Muller. 1956. Some Continuous Monte Carlo Methods for the Dirichlet Problem. *Annals of Mathematical Statistics* 27, 3 (Sept. 1956), 569–589.
- 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 (Dec. 2016), 14 pages. <https://doi.org/10.1145/2980179.2982429>
- Mohammad Sina Nabizadeh, Ravi Ramamoorthi, and Albert Chern. 2021. Kelvin transformations for simulations on infinite domains. *ACM Trans. Graph.* 40, 4, Article 97 (July 2021), 15 pages. <https://doi.org/10.1145/3450626.3459809>
- Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo methods for volumetric light transport simulation. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 551–576.
- Jan Novák, Andrew Selle, and Wojciech Jarosz. 2014. Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph.* 33, 6, Article 179 (Nov. 2014), 11 pages. <https://doi.org/10.1145/2661229.2661292>
- George C Papanicolaou and SR Srinivasa Varadhan. 1980. Diffusion in regions with many small holes. In *Stochastic Differential Systems Filtering and Control: Proceedings*. Springer, 190–206.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2018. Scenes for pbtr-v3. <http://pbtr.org/scenes-v3.html> Accessed: 23/01/2025.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.
- Yang Qi, Dario Seyb, Benedikt Bitterli, and Wojciech Jarosz. 2022. A bidirectional formulation for Walk on Spheres. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 51–62.
- Matthias Raab, Daniel Seibert, and Alexander Keller. 2006. Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Illumination 2006*. Springer, 591–605.
- Damien Rioux-Lavoie, Ryusuke Sugimoto, Tümay Özdemir, Naoharu H. Shimada, Christopher Batty, Derek Nowrouzezahrai, and Toshiya Hachisuka. 2022. A Monte Carlo Method for Fluid Simulation. *ACM Trans. Graph.* 41, 6, Article 240 (Nov. 2022), 16 pages. <https://doi.org/10.1145/3550454.3555450>
- Brian J Rothschild. 1992. Application of stochastic geometry to problems in plankton ecology. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 336, 1277 (1992), 225–237.
- Silvia Sabatini and Elena Villa. 2024. On a special class of gibbs hard-core point processes modeling random patterns of non-overlapping grains. *Stochastic Analysis and Applications* 42, 2 (2024), 386–413.
- Karl K. Sabelfeld and Nikolai A. Simonov. 2016. *Stochastic Methods for Boundary Value Problems*. De Gruyter, Berlin, Boston. <https://doi.org/doi:10.1515/9783110479454>
- Katherine Salesin, Kirk D Knobelspiesse, Jacek Chowdhary, Peng-Wang Zhai, and Wojciech Jarosz. 2024a. Unifying radiative transfer models in computer graphics and remote sensing, Part I: A survey. *Journal of Quantitative Spectroscopy and Radiative Transfer* 314 (2024), 108847.
- Katherine Salesin, Kirk D Knobelspiesse, Jacek Chowdhary, Peng-Wang Zhai, and Wojciech Jarosz. 2024b. Unifying radiative transfer models in computer graphics and remote sensing, Part II: A differentiable, polarimetric forward model and validation. *Journal of Quantitative Spectroscopy and Radiative Transfer* 315 (2024), 108849.
- Rohan Sawhney. 2021. FCPW: Fastest Closest Points in the West.
- Rohan Sawhney and Keenan Crane. 2020. Monte Carlo geometry processing: a grid-free approach to PDE-based methods on volumetric domains. *ACM Trans. Graph.* 39, 4, Article 123 (Aug. 2020), 18 pages. <https://doi.org/10.1145/3386569.3392374>
- Rohan Sawhney and Bailey Miller. 2024. *Zombie: A Grid-Free Monte Carlo Solver for PDEs*. <https://github.com/rohan-sawhney/zombie>
- Rohan Sawhney, Bailey Miller, Ioannis Gkioulekas, and Keenan Crane. 2023. Walk on Stars: A Grid-Free Monte Carlo Method for PDEs with Neumann Boundary Conditions. *ACM Trans. Graph.* 42, 4, Article 80 (July 2023), 20 pages. <https://doi.org/10.1145/3592398>
- Rohan Sawhney, Dario Seyb, Wojciech Jarosz, and Keenan Crane. 2022. Grid-free Monte Carlo for PDEs with spatially varying coefficients. *ACM Trans. Graph.* 41, 4, Article 53 (July 2022), 17 pages. <https://doi.org/10.1145/3528223.3530134>

- John H Seinfeld and Spyros N Pandis. 2016. *Atmospheric chemistry and physics: from air pollution to climate change*. John Wiley & Sons.
- Silvia Sellán and Alec Jacobson. 2022. Stochastic Poisson Surface Reconstruction. *ACM Trans. Graph.* 41, 6, Article 227 (Nov. 2022), 12 pages. <https://doi.org/10.1145/3550454.3555441>
- Silvia Sellán and Alec Jacobson. 2023. Neural Stochastic Poisson Surface Reconstruction. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 74, 9 pages. <https://doi.org/10.1145/3610548.3618162>
- Dario Seyb, Eugene d'Eon, Benedikt Bitterli, and Wojciech Jarosz. 2024. From microfacets to participating media: A unified theory of light transport with stochastic geometry. *ACM Trans. Graph.* 43, 4, Article 112 (July 2024), 17 pages. <https://doi.org/10.1145/3658121>
- Scott Sheffield. 2007. Gaussian free fields for mathematicians. *Probability theory and related fields* 139, 3 (2007), 521–541.
- Leo Speidel, Heather A Harrington, S Jonathan Chapman, and Mason A Porter. 2018. Topological data analysis of continuum percolation with disks. *Physical Review E* 98, 1 (2018), 012318.
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized yarn-level cloth. *ACM Trans. Graph.* 39, 4, Article 48 (Aug. 2020), 16 pages. <https://doi.org/10.1145/3386569.3392412>
- Jos Stam. 1999. Stable fluids. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., USA, 121–128. <https://doi.org/10.1145/311535.311548>
- Dietrich Stoyan and Helga Stoyan. 1985. On one of Matérn's hard-core point process models. *Mathematische Nachrichten* 122, 1 (1985), 205–214.
- Ryusuke Sugimoto, Christopher Batty, and Toshiya Hachisuka. 2024. Velocity-Based Monte Carlo Fluids. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 8, 11 pages. <https://doi.org/10.1145/3641519.3657405>
- Ryusuke Sugimoto, Terry Chen, Yiti Jiang, Christopher Batty, and Toshiya Hachisuka. 2023. A Practical Walk-on-Boundary Method for Boundary Value Problems. *ACM Trans. Graph.* 42, 4, Article 81 (July 2023), 16 pages. <https://doi.org/10.1145/3592109>
- Eric Veach. 1998. *Robust Monte Carlo methods for light transport simulation*. Stanford University.
- Delio Vicini, Wenzel Jakob, and Anton Kaplanyan. 2021. A non-exponential transmittance model for volumetric scene representations. *ACM Trans. Graph.* 40, 4, Article 136 (July 2021), 16 pages. <https://doi.org/10.1145/3450626.3459815>
- David A Weitz and David J Pine. 1993. Diffusing-wave spectroscopy. *Dynamic light scattering* (1993), 652–720.
- Stephen Whitaker. 1986. Flow in porous media I: A theoretical derivation of Darcy's law. *Transport in porous media* 1 (1986), 3–25.
- Chun Yuan, Haoyang Shi, Lei Lan, Yuxing Qiu, Cem Yuksel, Huamin Wang, Chenfanfu Jiang, Kui Wu, and Yin Yang. 2024. Volumetric Homogenization for Knitwear Simulation. *ACM Trans. Graph.* 43, 6, Article 207 (Nov. 2024), 19 pages. <https://doi.org/10.1145/3687911>
- Qingnan Zhou and Alec Jacobson. 2016. Thingi10K: A Dataset of 10,000 3D-Printing Models. *arXiv preprint arXiv:1605.04797* (2016).

A Details on the Poisson Boolean model

We elaborate on the closest point distribution and the related spherical contact distribution in the PBM. We also prove the correctness of conditional closest point sampling using Algorithm 1 with the conditional density of Equation (13). We use the following property of the Poisson point process [Last and Penrose 2017, Definition 3.1].

PROPOSITION 2: PROPERTIES OF THE POISSON POINT PROCESS

We assume that the point set $C_O := \{c_n \in V\}_{n=1}^N$ is a sample of a *Poisson point process* with rate function λ . Then:

P1. For any set $Q \subset V$,

$$\Pr\{Q \cap C_O = \emptyset\} = \exp\left(-\int_Q \lambda(y) dy\right). \quad (26)$$

P2. For any two sets $Q, S \subset V$, if $Q \cap S = \emptyset$, then the events $Q \cap C_O = \emptyset$ and $S \cap C_O = \emptyset$ are independent.

Properties P1 and P2 are analogues of the properties of exponential and independent (resp.) increments of the Poisson process on

the real line. P1 also explains the term *exponential media* for volumes with stochastic microparticle geometry following the PBM.

Spherical contact and closest point distributions. We consider the random closest point $y^{\partial O}(x) \in \partial O$ and random shortest distance $r^{\partial O}(x) := \|x - y^{\partial O}(x)\|$ between a point $x \in V$ and the boundary ∂O . The random variables $r^{\partial O}(x)$ and $y^{\partial O}(x)$ follow the *spherical contact distribution* and *closest point distribution* [Last and Penrose 2017, Section 16.3], whose PDFs we denote p_x^{sc} and p_x^{cp} (resp.).

We can derive p_x^{sc} from p_x^{dc} (Equation (8)), using the relationship $r^{\partial O}(x) = r^c(x) - R$ (Figure 4). Doing so requires conditioning on $r^c(x) \geq R$ to ensure a positive value for $r^{\partial O}(x)$. Requiring $r^c(x) \geq R$ is equivalent to $x \notin O$, or $B(x, R) \cap C_O = \emptyset$. Then, from Equations (8) and (26) and the definition of conditional probabilities:

$$p_x^{\text{sc}}(r) := p_x^{\text{dc}}(r + R \mid B(x, R) \cap C_O = \emptyset) \quad (27)$$

$$= \exp(-\Lambda(x, r + R) + \Lambda(x, R)) \int_{\partial B(x, r+R)} \lambda(y) dA(y). \quad (28)$$

We can likewise derive p_x^{cp} from p_x^{cc} (Equation (10)), this time using the relationship $y^{\partial O}(x) = c(x) - R \text{dir}(x, c(x))$ (Figure 4). From Equations (8) and (26) and conditioning on $r^{\partial O}(x) = r$,

$$p_x^{\text{cp}}(y \mid r) := \frac{\lambda(y + R \text{dir}(x, y))}{\int_{\partial B(x, r+R)} \lambda(y) dA(y)} \frac{(r + R)^2}{r^2}. \quad (29)$$

In the homogeneous case, Equations (28) and (29) simplify to:

$$p_x^{\text{sc}}(r) \stackrel{h}{=} \exp\left(-4/3\pi\left((r + R)^3 - R^3\right)\lambda\right) 4\pi(r + R)^2\lambda, \quad (30)$$

$$p_x^{\text{cp}}(y \mid r) \stackrel{h}{=} \frac{1}{4\pi r^2}. \quad (31)$$

Comparing Equations (28) and (29) with Equations (8) and (10) (resp.), we note that the distributions for $r^{\partial O}(x)$ and $y^{\partial O}(x)$ are closely related to those of $r^c(x)$ and $c(x)$. However, the simpler expressions for the center-based quantities, and in particular the exponential property of $(r^c(x))^3$, greatly simplify sampling a closest center relative to directly sampling a closest point (Section 4).

Proof of conditional sampling. The statement $Q \cap O = \emptyset$ is equivalent to $Q^{\oplus R} \cap C_O = \emptyset$. Using Properties P1 and P2 and the definition of conditional probabilities, we can update Equation (8) as:

$$p_x^{\text{dc}}(r \mid Q^{\oplus R} \cap C_O = \emptyset) = \exp\left(-\Lambda(x, r) + \int_{B(x, r) \cap Q^{\oplus R}} \lambda(y) dy\right) \cdot \left(\int_{\partial B(x, r)} \lambda(y) dA(y) - \int_{\partial B(x, r) \cap Q^{\oplus R}} \lambda(y) dA(y)\right), \quad (32)$$

which equals the unconditional p_x^{dc} of Equation (8) computed using the conditional density $\lambda(\cdot \mid Q)$ in Equation (13).

B Volume rendering

It is instructive to compare and contrast VWoS with volume rendering algorithms for exponential media, and in particular its closest analogue, *volumetric path tracing* (VPT). We first overview VPT, then discuss analogies with VWoS. Our discussion of VPT and volume rendering is brief, and we refer to Novák et al. [2018] and Pharr et al. [2023, Chapter 15] for more detailed treatments.

VPT overview. In rendering, the domain Ω plays the role of the scene. When the scene geometry is deterministic, path tracing computes the incident radiance L_i at a point x and direction ω recursively, using the conservation law for radiance along a ray:

$$L_i(x, \omega) = L_o(y^{\partial\Omega}(x, \omega), -\omega), \quad (33)$$

where L_o is outgoing radiance. The point $y^{\partial\Omega}(x, \omega)$ is the first intersection with the scene boundary $\partial\Omega$ of a ray with origin x and direction ω . We term the distance $t^{\partial\Omega}(x, \omega) := \|x - y^{\partial\Omega}(x, \omega)\|$ the *ray distance* along the ray with origin x and direction ω . Noting that $y^{\partial\Omega}(x, \omega) = x + t^{\partial\Omega}(x, \omega)\omega$, Equation (33) becomes:

$$L_i(x, \omega) = L_o(x + t^{\partial\Omega}(x, \omega)\omega, -\omega). \quad (34)$$

When the scene geometry is deterministic, we can compute $y^{\partial\Omega}(x, \omega)$ and $t^{\partial\Omega}(x, \omega)$ by performing a *ray casting query*.

In a participating medium, the scene includes both deterministic and stochastic geometry as in Equation (4). Volumetric path tracing then computes the incident *mean radiance* \bar{L}_i , defined analogously to the mean solution \bar{u} in Equation (5).⁷ Taking the expectation of Equation (34) and applying the law of total expectation gives us:

$$\bar{L}_i(x, \omega) = \int_0^\infty p^t(t) \bar{L}_o(x + t\omega, -\omega | t^{\partial\Omega}(x, \omega) = t) dt. \quad (35)$$

Here, we used the fact that in a participating medium the ray distance $t^{\partial\Omega}(x, \omega)$ is a continuous random variable, indicating its PDF as p^t . From Equation (4), it follows that:

$$t^{\partial\Omega}(x, \omega) = \min\{t^{\partial O}(x, \omega), t^{\partial V}(x, \omega)\}, \quad (36)$$

where: 1. $t^{\partial O}(x, \omega)$ is the random ray distance to the boundary ∂O of the microparticle geometry; 2. $t^{\partial O}(x, \omega)$ is the deterministic ray distance to the boundary ∂V of the volume.

The random ray distance $t^{\partial O}(x, \omega)$ follows the *linear contact distribution* [Last and Penrose 2017, Section 16.3] with associated PDF $p_{x,\omega}^{\text{ff}}$ and tail distribution function $T_{x,\omega}^{\text{ff}}$. The rendering literature uses the terms *free-flight distance* for $t^{\partial O}(x, \omega)$, *free-flight distribution* for $p_{x,\omega}^{\text{ff}}$, and *transmittance* for $T_{x,\omega}^{\text{ff}}$ [Bitterli et al. 2018]. From Equation (36), it follows that we can rewrite Equation (35) as:

$$\begin{aligned} \bar{L}_i(x, \omega) = & \int_0^{t^{\partial V}} p_{x,\omega}^{\text{ff}}(t) \bar{L}_o(x + t\omega, -\omega | t^{\partial\Omega} = t) dt \\ & + T_{x,\omega}^{\text{ff}}(t^{\partial V}) \bar{L}_o(x + t^{\partial V}\omega, -\omega | t^{\partial\Omega} = t^{\partial V}), \end{aligned} \quad (37)$$

where we used the shorthands $t^{\partial\Omega}$, $t^{\partial O}$, and $t^{\partial V}$ to simplify notation. Equation (37) is the *volume rendering equation* (VRE). To estimate \bar{L}_i , VPT first samples a free-flight distance $t \sim p_{x,\omega}^{\text{ff}}$ and then:

1. estimates $\bar{L}_o(x + t\omega, -\omega | t^{\partial\Omega} = t)$ if $t^{\partial\Omega} < t^{\partial V}$, effectively applying single-sample Monte Carlo to the integral term;
2. estimates $\bar{L}_o(x + t^{\partial V}\omega, -\omega | t^{\partial\Omega} = t)$ if $t^{\partial\Omega} \geq t^{\partial V}$.

Estimation then proceeds recursively, albeit *without memory*: VPT approximates $\bar{L}_o(x + t\omega, -\omega | t^{\partial\Omega} = t) \approx \bar{L}_o(x + t\omega, -\omega)$ —a so-called *renewal assumption* [Seyb et al. 2024]—and continues to iterate Equation (37) (after first using the in-scattering equation to convert from outgoing to incident radiance).

⁷Notably, the rendering literature typically does not distinguish between mean radiance and radiance, treating the two quantities as interchangeable.

To make free-flight distance sampling tractable, classical VPT algorithms assume *exponential media*, or equivalently, that the microparticle geometry follows the PBM $\Phi(\lambda, R)$. Using the PBM properties in Proposition 2 (analogously to the derivation of p_x^{sc} in Equation (28)), as well as an assumption that particle size R is appropriately small, the transmittance and free-flight distribution become [Last and Penrose 2017, Section 16.3]:

$$T_{x,\omega}^{\text{ff}}(t) = \exp\left(-\int_0^t 4\pi R^2 \lambda(x + s\omega) ds\right), \quad (38)$$

$$p_{x,\omega}^{\text{ff}}(t) = \exp\left(-\int_0^t 4\pi R^2 \lambda(x + s\omega) ds\right) 4\pi R^2 \lambda(x + t\omega), \quad (39)$$

or in the homogeneous case, analogously to Equation (30):

$$T_{x,\omega}^{\text{ff}}(t) \stackrel{\text{h}}{=} \exp\left(-4\pi R^2 \lambda t\right), \quad (40)$$

$$p_{x,\omega}^{\text{ff}}(t) = \exp\left(-4\pi R^2 \lambda t\right) 4\pi R^2 \lambda. \quad (41)$$

We note that the closest analogue of the free-flight distribution $p_{x,\omega}^{\text{ff}}$ in VWoS is the spherical contact distribution p_x^{sc} (Equation (28)). However, under the small R assumption, $p_{x,\omega}^{\text{ff}}$ is more similar to the shortest distance-to-center distribution p_x^{dc} (Equation (8)).

It follows that, in the homogeneous case, sampling the free-flight distance requires simply sampling an exponential random variate with rate $4\pi R^2 \lambda t$. In the heterogeneous case, sampling can be done by thinning [Lewis and Shedler 1979], also known as *delta tracking* in the rendering literature: We progressively increase the free-flight distance by increments sampled assuming homogeneous density $\bar{\lambda} \geq \lambda$, until we accept a value with probability $\lambda/\bar{\lambda}$.

Comparison with VWoS. The above discussion helps highlight the many analogies between our VWoS algorithm for PDE simulation in participating media, and VPT for rendering in participating media.

- Both algorithms are derived by first transforming a recursive equation through the law of total expectation, then applying single-sample Monte Carlo: the BIE (2) becomes the BIE in participating media (18) for VWoS; and the radiance conservation law (34) becomes the VRE (37) for VPT.
- Both algorithms replace deterministic geometric queries with sampling operations: VWoS replaces closest point queries with closest point sampling, and VPT replaces ray-casting queries with free-flight distance sampling.
- These sampling operations require characterizing associated PDFs: the shortest distance-to-center p_x^{dc} in VWoS, and the free-flight distribution (or linear contact distribution) $p_{x,\omega}^{\text{ff}}$ in VPT.
- Both algorithms use the PBM to make these distributions easy to sample: the cubed shortest distance in VWoS and the free-flight distance in VPT become exponential random variables.
- Both algorithms use thinning for sampling under heterogeneous densities: WoS uses Algorithm 1, and VPT uses delta tracking.

At the same time, there are important differences.

- Perhaps the most salient difference relates to *memory*: VWoS has full memory by conditioning, during recursion, on the outcomes of sampling operations at previous steps. VPT has no memory, “forgetting” those previous outcomes. Ignoring memory is a reasonable approximation in VPT because linear segments along a

Table 1. Scene parameters for experiments in Sections 7 and 8. We report the maximum density and corresponding mean free ball radius (average shortest distance to particle centers) of the participating media. We also report the maximum extent (maximum length across all dimensions) of the scenes.

scene	ϵ -shell width	particle size R	maximum density λ	mean free ball radius	maximum scene extent
Figure 1, membrane	1×10^{-4}	1×10^{-3}	1×10^5	1.2×10^{-2}	4.0×10^0
Figure 9(a), mushroom	1×10^{-4}	5×10^{-2}	1×10^2	1.2×10^{-1}	2.3×10^0
Figure 9(b), mushroom	1×10^{-4}	1×10^{-2}	5×10^2	7.0×10^{-2}	2.3×10^0
Figure 9(c), mushroom	1×10^{-4}	1×10^{-3}	5×10^3	3.2×10^{-2}	2.3×10^0
Figure 9(d), ginseng	1×10^{-4}	2×10^{-3}	1×10^6	5.5×10^{-3}	4.0×10^{-1}
Figure 9(e), ginseng	1×10^{-4}	2×10^{-3}	1×10^6	5.5×10^{-3}	2.0×10^{-1}
Figure 9(f), ginseng	1×10^{-4}	2×10^{-3}	1×10^6	5.5×10^{-3}	1.0×10^{-1}
Figure 10(a), connector	1×10^{-3}	2×10^{-2}	5.5×10^5	6.8×10^{-3}	1.8×10^0
Figure 10(b), connector	1×10^{-3}	2×10^{-2}	2.5×10^7	1.9×10^{-3}	1.8×10^0
Figure 13, cloud	1×10^{-4}	5×10^{-3}	5×10^4	1.5×10^{-2}	2.0×10^0

Table 2. Dirichlet boundary data g on the medium boundary ∂V and particle boundary ∂O , for experiments in Sections 7 and 8. x_i is the i -th coordinate of the point $x \in \mathbb{R}^3$. $y^{\partial V}(x)$ and $r^{\partial V}(x)$ are the closest point and shortest distance (resp.) between x and ∂V .

scene	medium boundary	particle boundary
Figure 1, membrane	$\text{texture}(x)$	$-0.25\text{texture}(y^{\partial V}(x)) \exp(-200r^{\partial V}(x)^2)$
Figure 9, mushroom	$\text{texture}(x)$	0
Figure 9, ginseng	$0.5(\cos(2 \exp(-2(x_2 - 1.75))x_0) \cos(2 \exp(-2(x_2 - 1.75))x_1) - 1.75)$	0
Figure 10, connector	if $x_0 < 0$ then $0.5 \cos(10x_2) + 0.5$ else $0.5 \cos(10x_2) - 0.5$	0
Figure 13, cloud	0	fluence $\bar{H}_1(x)$

light path are unlikely to overlap, making it unnecessary for a segment to condition on the empty space carved by a previous one. By contrast in VWoS, balls along a random walk are very likely to overlap, requiring memory for accurate estimation (Section 7.3).

- A more subtle difference becomes evident when we compare the BIE in participating media (18) with the VRE (37): The former conditions on closest *points*, whereas the latter conditions on shortest ray *distances*. This difference stems from the fact that whereas conditioning on the shortest ray distance fixes a unique closest intersection point, conditioning on the shortest distance only fixes the closest point to lie on a sphere, requiring additional area sampling on that sphere (Proposition 1 and Algorithm 1).

Reconciling these differences creates future research opportunities, for example towards volume rendering algorithms that use memory for improved estimation accuracy, or towards VWoS variants that use only shortest distance sampling for improved efficiency.

Lastly, we note that volume rendering applications typically specify participating media through their *extinction coefficient* $\sigma_t(x) := 4\pi R^2 \lambda(x)$ that appears in Equations (38)–(41), rather than their density $\lambda(x)$. The closed-form relationship between the two allows us to reuse abundant publicly available volume models for rendering also in PDE simulation, as we do in Section 8.2.

C Scene parameters

We report scene parameters (Table 1) and boundary conditions (Table 2) for all experiments in Sections 7 and 8.