

Graphics-processing-unit-accelerated Monte Carlo simulation of polarized light in complex three-dimensional media

Shijie Yan¹,^a Steven L. Jacques²,^b Jessica C. Ramella-Roman³,^c and Qianqian Fang¹,^{d,*}

^aNortheastern University, Department of Electrical and Computer Engineering, Boston, Massachusetts, United States

^bUniversity of Washington, Department of Bioengineering, Seattle, Washington, United States

^cFlorida International University, Department of Biomedical Engineering, Miami, Florida, United States

^dNortheastern University, Department of Bioengineering, Boston, Massachusetts, United States

Abstract

Significance: Monte Carlo (MC) methods have been applied for studying interactions between polarized light and biological tissues, but most existing MC codes supporting polarization modeling can only simulate homogeneous or multi-layered domains, resulting in approximations when handling realistic tissue structures.

Aim: Over the past decade, the speed of MC simulations has seen dramatic improvement with massively parallel computing techniques. Developing hardware-accelerated MC simulation algorithms that can accurately model polarized light inside three-dimensional (3D) heterogeneous tissues can greatly expand the utility of polarization in biophotonics applications.

Approach: Here, we report a highly efficient polarized MC algorithm capable of modeling arbitrarily complex media defined over a voxelated domain. Each voxel of the domain can be associated with spherical scatters of various radii and densities. The Stokes vector of each simulated photon packet is updated through photon propagation, creating spatially resolved polarization measurements over the detectors or domain surface.

Results: We have implemented this algorithm in our widely disseminated MC simulator, Monte Carlo eXtreme (MCX). It is validated by comparing with a reference central-processing-unit-based simulator in both homogeneous and layered domains, showing excellent agreement and a 931-fold speedup.

Conclusion: The polarization-enabled MCX offers biophotonics community an efficient tool to explore polarized light in bio-tissues, and is freely available at <http://mcx.space/>.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JBO.27.8.083015](https://doi.org/10.1117/1.JBO.27.8.083015)]

Keywords: Monte Carlo method; light transport; polarization; optical imaging.

Paper 220009SSR received Jan. 13, 2022; accepted for publication Apr. 8, 2022; published online May 9, 2022.

1 Introduction

Polarized light has been found to be highly sensitive to medium structures and hence has been widely adopted in optical imaging to probe microstructural features inside biological tissues.¹⁻⁴ For example, the polarization status of the backscattered light can be measured to characterize the superficial layer of skin for cancer diagnostic purposes.^{5,6} The measurements of tissue birefringence permit quantification of abnormalities of the retinal nerve fiber layer⁷ and cornea,⁸

*Address all correspondence to Qianqian Fang, q.fang@neu.edu

as well as three-dimensional (3D) reconstruction of nerve fiber orientations inside human brains⁹ and orientations of collagen within the uterine cervix.¹⁰ By measuring the unequal absorption of left-handed and right-handed circularly polarized light, circular dichroism can rapidly determine the folding properties of proteins.¹¹ Quantification of collagen and birefringent media alignment can improve evaluation of a therapeutic strategy and its outcome in scar management.¹² Polarized light imaging (PLI) uses linearly co-polarized images subtracted by those of cross-polarized light to create a differential image based on the small population of superficially scattered photons that still retain much of the incident polarization state.⁵ PLI subtracts the large randomized population of multiple-scattered photons that produce a blinding background of diffuse light. The resulting difference image enhances the contrast of superficial tissue layers and rejects deeper tissue structures, enabling wide-field screening of epidermal or epithelial layers. Mueller matrix polarimetry and the use of various decomposition methods can also be used to pinpoint different regions and structures within biological tissue.^{2,3} Accurately simulating polarized light transport inside complex tissues allow quantitative investigations of the depth response of polarized light and the perturbations produced by local tissue abnormalities.

The propagation of polarized light inside scattering media can be described by the vector radiative transfer equation (VRTE).¹³ Analytically solving the VRTE is not viable in complex media such as human tissues. Owing to its high flexibility and simplicity in programming, the Monte Carlo (MC) method, among other numerical techniques,^{14–19} has been one of the limited approaches available to quantitatively model interactions of polarized light with scattering media. Depending on the vectorial representations of polarization states, polarized light MC algorithms can be largely categorized into two formalisms—Jones calculus and Mueller calculus.^{2,4,20} The Jones calculus used in the electric-field MC (EMC) algorithm traces the amplitudes and phases of two orthogonal electric-field components (Jones vector) and is therefore well suited for simulating light coherence effects.²¹ On the other hand, the Mueller calculus describes the state of polarization using the Stokes vector.^{22–26} The Stokes vector does not contain the absolute phase of the electrical field but allows to model unpolarized, partially polarized and fully polarized light. It can be obtained by measuring four intensity values. Similarly, Mueller matrices can be obtained through 16 intensity measurements and using Mueller matrix decomposition, quantities, such as tissue retardation, depolarization and de-attenuation can be obtained.⁴

A well-known limitation of MC methods is the long computation time. Due to the rapid emergence of massively parallel computing techniques, benefit largely from the fast advances in many core processors such as graphics processing units (GPUs), MC simulations of polarized light have seen significant speed improvement over the past decade. Several groups have reported parallel EMC implementations.^{27–29} Wang et al.²⁷ presented a compute unified device architecture (CUDA)³⁰-based EMC to model coherent light in a single homogeneous slab and achieved over 370× speedup compared to the central processing unit (CPU) counterpart.²¹ Ding et al.²⁹ extended the GPU-based EMC algorithm to consider multi-layered media at the expense of reduced speedup (~45×). In addition, Li et al.³¹ presented a CUDA-based polarized light MC algorithm to model interstitial media embedded with spherical and cylindrical scatterers. It employed a single-kernel scheme and was hundreds of times faster than its CPU version.²⁶ In 2019, Oulhaj et al.³² reported a GPU-accelerated MC algorithm to efficiently compute the sensitivity profile for polarized light inside homogeneous media. The reported GPU implementation was verified against a widely used CPU-based code developed by Ramella-Roman et al.²⁴ and reported over 150× speedup.

Although these studies have demonstrated significantly improved simulation speed, most of these simulators only support layered domains and can not address the needs in modeling increasingly complex media.⁴ In the simulation of biological tissues with irregular-shaped structures, employing simplifications in domain geometries could introduce significant errors. For example, in the MC modeling of human brains, noticeable differences have been observed between layered-slab models and more anatomically realistic models such as voxel-based and mesh-based brain models.³³

In this work, we present an open-source and GPU-accelerated MC simulator to model polarized light inside 3D heterogeneous media. This MC algorithm utilizes a 3D voxelated grid to

represent spatially varying distributions of spherical scatterers, characterized by their radii and densities. We use Muller calculus to update the Stokes vectors of simulated photon packets, from which we can compute various polarimetry related measurements along the surface of the domain. In the remainder of this paper, we first briefly review the steps of the meridian-plane MC algorithm.²⁴ Then we detail our GPU-implementation of this algorithm, as part of our enhanced open-source MC software—Monte Carlo eXtreme (MCX),³⁴ including the preprocessing steps to encode the distribution of particles into a 3D array data structure and optimization strategies to better use GPU resources. In Sec. 3, we validate the proposed GPU-based polarization-enabled MCX (pMCX) against the widely used CPU MC simulator “meridianMC” written by Ramella-Roman et al.³⁵ and quantify the speed improvement using several benchmarks of homogeneous and heterogeneous domains. Finally, we summarize the key findings and discuss future directions.

2 Methods

2.1 Meridian-Plane Polarized Light MC

The meridian-plane polarized light MC algorithm³⁵ largely follows the standard MC photon transport simulation steps,³⁶ including “launch,” “move,” “absorb,” “scatter,” and “detection.” At the “launch” stage, the initial weight, position, and direction vector are defined for each photon packet depending on the profile of the incident beam. To describe the polarization state, the Stokes vector is defined with respect to the initial meridian plane for every simulated photon packet.²⁴ The meridian plane is defined by the plane spanned by the photon propagation direction and the z axis, as shown in Fig. 2 in Ref. 24. The Stokes vector \vec{S} consists of four quantities $[I, Q, U, V]$, where I ($I \geq 0$) describes the total light intensity, Q ($-1 \leq Q \leq 1$) controls the mixing between horizontally ($Q = 1$) and vertically ($Q = -1$) linearly polarized light, U ($-1 \leq U \leq 1$) controls the mixing between $+45$ deg ($U = 1$) and -45 deg ($U = -1$) linearly polarized light, and V ($-1 \leq V \leq 1$) controls the mixing between right ($V = 1$) and left ($V = -1$) circularly polarized light.¹

After the “launch” step, the photon packet starts propagating inside the simulation domain. In lossy media, the packet weight is monotonically reduced along the photon’s paths and the weight loss is accumulated into the local grid element (such as a voxel or tetrahedral element³⁷). When arriving at an interaction site, the photon packet changes direction due to scattering. To compute the new direction cosines, the scattering zenith angle θ ($0 \leq \theta \leq \pi$) and azimuth angle ϕ ($0 \leq \phi < 2\pi$) are statistically sampled. Compared to the standard MC, the scattering step in a polarized light simulation requires additional computation to properly update \vec{S} . First, the probability density function (also known as the scattering phase function) of polarized light has a bivariate dependence on both θ and ϕ . For incident light with a Stokes vector $\vec{S}_{in} = [I_{in}, Q_{in}, U_{in}, V_{in}]$, the phase function $P(\theta, \phi)$ ²⁴ is

$$P(\theta, \phi) = s_{11}(\theta)I_{in} + s_{12}(\theta)[Q_{in} \cos(2\phi) + U_{in} \sin(2\phi)], \quad (1)$$

where $s_{11}(\theta)$ and $s_{12}(\theta)$ are elements from the scattering matrix $M(\theta)$ from a homogeneous spherical particle, computed via the Mie theory³⁸

$$M(\theta) = \begin{bmatrix} s_{11}(\theta) & s_{12}(\theta) & 0 & 0 \\ s_{12}(\theta) & s_{11}(\theta) & 0 & 0 \\ 0 & 0 & s_{33}(\theta) & s_{34}(\theta) \\ 0 & 0 & -s_{34}(\theta) & s_{33}(\theta) \end{bmatrix}. \quad (2)$$

A rejection method²⁴ is employed to select both angles θ and ϕ . Once θ and ϕ are determined, the Stokes vector must be rotated relative to the new meridian plane using $M(\theta)$ to update the polarization states.

To efficiently perform the rejection method and Stokes vector rotation, the elements of the scattering matrix $M(\theta)$ of all user-specified spherical scatter species are pre-computed over a

discretized set of θ . A photon packet is terminated when it escapes from the simulation domain or, if enabled, fails to survive a Russian roulette.³⁶ It is noteworthy that the original meridian-plane polarized light MC assumes refractive-index matched domain boundaries.²⁴ The Stokes vector of the escaping photon is rotated relative to the meridian plane of the detector placed immediately outside the domain boundaries, before being accumulated to generate desired output quantities. A detailed description of the formulas used in the meridian plane MC algorithm can be found in the literature.²⁴

2.2 Implementing Meridian Plane MC in MCX

The original CPU-based meridian-plane MC program,³⁵ referred to as “mcMeridian” (stok1.c) hereinafter, is dedicated to modeling homogeneous infinite slab geometries. In contrast, the CUDA-based MCX is capable of modeling arbitrarily heterogeneous media represented by a 3D voxelated domain.³⁴ In a non-polarization MCX simulation, the domain is represented by a 3D integer array with each number representing the index or label of the tissue at each voxel. The actual optical properties of the tissue label are stored in a look-up table, with four element per tissue type: absorption coefficient μ_a (1/mm), scattering coefficient μ_s (1/mm), anisotropy g , and refractive index n . To simulate polarized light, the scattering properties of each type of scatterer must be included in addition. To simplify the computation, here we only consider spherical scatterers. The radius r (μm), refractive index n_{sph} , and volumetric number density ρ ($1/\mu\text{m}^3$) of the spherical particle scatters can be specified for each tissue type. When spherical particle properties are specified, the corresponding μ_s , g , and elements of the scattering matrix $M(\theta)$ are pre-computed using the Mie theory³⁹ on the host (i.e., CPU). As shown in Eq. (2), the scattering matrix of homogeneous spherical scatterers consists of four independent floating-point numbers $s_{11}(\theta)$, $s_{12}(\theta)$, $s_{33}(\theta)$, and $s_{34}(\theta)$. In our implementation, the scattering parameters are sampled at 1000 evenly spaced points between 0 and π , as done in mcMeridian.³⁵ The pre-computed optical properties and scattering matrix data are then transferred to the device (i.e., GPU), as shown in Fig. 1.

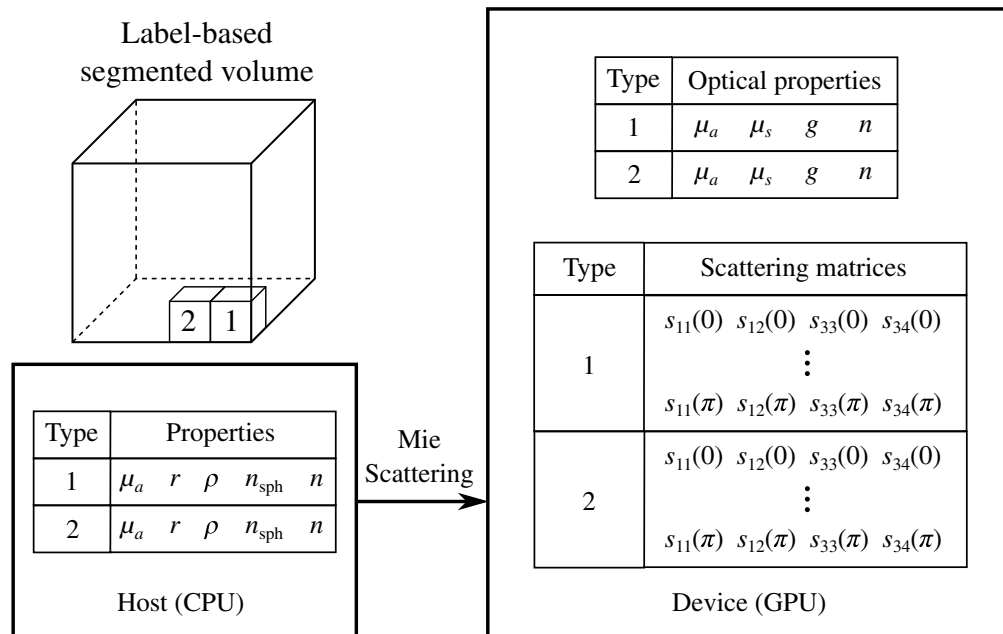


Fig. 1 Media representation and media data preprocessing in a polarization-enabled Monte Carlo simulation.

3 Results and Discussions

In this section, we first validate the aforementioned pMCX in a homogeneous slab using the single-threaded CPU-based implementation (mcMeridian²⁴) as a reference, which has been extensively used by the community and validated by experimental studies.²⁵ The speed improvement is also quantified. It is noteworthy that mcMeridian³⁵ simulates an infinite slab media geometry in the x/y directions, whereas in an MCX simulation, a photon is confined inside a bounding box with user-specified dimensions.³⁴ To ensure that our speed comparison is valid, we modified the source code of mcMeridian and added an implicit bounding box to match the dimension settings in pMCX. In the first benchmark, the simulation domain is a $20 \times 20 \times 10$ mm³ homogeneous slab, the Mie scattering parameters of the embedded spherical scatterer are $\mu_a = 0$ mm⁻¹, $r = 1.015$ μ m, $\rho = 1.152 \times 10^{-4}$ μ m⁻³, and $n_{\text{sph}} = 1.59$. The refractive index of the background medium is $n = 1.33$. A monochromatic pencil beam source is positioned at the bottom center (10,10,0) mm of the domain, pointing toward the $+z$ axis and emitting horizontally polarized light at wavelength $\lambda = 632.8$ nm. The initial Stokes vector of the incident beam is $\vec{S} = [1, 1, 0, 0]$. The backscattered photons are collected by a square-shaped area detector (20×20 mm²) placed on the boundary at $z = 0$ mm. In this benchmark, 10^8 photon packets are simulated on a desktop running Ubuntu 18.04 with an Intel i7-6700K CPU and an NVIDIA RTX 2080 GPU.

In Fig. 2, we compare the distribution of backscattered $[I, Q, U, V]$ components using contour plots in MATLAB (MathWorks, Inc., Natick, Massachusetts, United States) and observed excellent agreement between mcMeridian and pMCX solutions. For further quantitative analysis, the root-mean-square errors of I, Q, U, V (in \log_{10} scale) between mcMeridian and pMCX are computed on the matching detector area (20×20 mm²), reporting 0.0076, 0.0881, 0.1015, and 0.0938, respectively. We measure the total runtimes, including input data preprocessing, photon transport simulation, and output image generation, with mcMeridian reporting 18111.44 s using the Intel CPU and pMCX reporting 19.45 s on the NVIDIA RTX 2080 GPU, suggesting a 931 \times speedup. In addition, we also benchmark simulation speeds when storing the scattering matrix data over different GPU memory locations, including global, shared, and constant memories.³⁰ The global memory implementation reports the fastest speed at 8401 photons/ms, followed by the shared memory (4965 photons/ms) and constant memory (2182 photons/ms) implementations. Although the shared memory is known to be the fastest among the three memory types, it has a very small size, up to 48 KB per block.³⁰ For storing the scattering matrix of a single species of scatterer at 1000 angular steps, a total of 16 KB memory is needed. Allocating a large amount of shared memory can lead to drastically reduced active block number, which explains the lower speed compared to the global memory case. On the other hand, constant memory also has a small size (64 KB).³⁰ It is most efficient when a memory value is being reused many times after a single read. However, the use of the rejection method requires random access to the buffer which fails to be accelerated by the constant memory due to high “cache-miss.”³⁰

In the next benchmark, we further validate our pMCX simulator by comparing with an extended mcMeridian (with added support of layered media) in a two-layer domain. The slab-shaped simulation domain has a size of $100 \times 100 \times 50$ mm³ with the thickness of the superficial layer, d_e , ranging between 0 and 10 mm. The Mie scattering parameters are $\lambda = 632.8$ nm, $\mu_a = 0.001$ mm⁻¹, sphere radius $r = 0.05$ μ m, number density $\rho = 19.11$ μ m⁻³, $n_{\text{sph}} = 1.59$, and $n = 1.33$ for the superficial layer. The bottom layer has $r = 0.3$ μ m, $\rho = 2.198 \times 10^{-2}$ μ m⁻³, and the other parameters are the same as the superficial layer. These choices of r and ρ yield a reduced scattering coefficient $\mu'_s = \mu_s \cdot (1 - g) = 1$ mm⁻¹ for both layers, along with the same absorption μ_a and hence approximately the same reflected intensity I for all d_e . A pencil beam is located on the surface of the slab at (50,50,0) mm, pointing toward the $+z$ axis and emitting horizontally polarized light, with the initial Stokes vector $\vec{S} = [1, 1, 0, 0]$. A total of 10^7 photon packets are simulated for both mcMeridian and pMCX.

In Fig. 3, we plot the total reflected I and Q components as a function of the superficial layer thickness d_e . We do not include total U and V plots because they are nearly zeros across all d_e values, this is expected as U and V distributions sum to zeros due to symmetric positive and




Fig. 2 Contour plots of the absolute backscattered (a) I , (b) Q , (c) U and (d) V (in \log_{10} scale) generated by mcMeridian (black solid lines) and pMCX (white dashed lines).

negative components. The outputs from mcMeridian (black solid lines) and those from pMCX (red circles) once again show excellent agreements. The plot of Q increases with the thickness d_e of the superficial layer, which contains smaller spherical scatters and hence stronger backscattering than the deeper layer. With thickness d_e growing from 0 to 1 cm, the total I value shows a minuscule increase by 0.17% (from 0.9080 to 0.9095), as shown in the inset in Fig. 3(a), as a result of sub-diffusive scattering. The two-phase transition of Q matches our expectations: when the superficial layer is very thin, the reflectance values are close to the value as if the domain is entirely filled with the bottom medium (green dashed line); as we increase d_e , the reflectance values asymptotically approach those determined by the media in the superficial layer (blue dashed line).

Finally, we show simulation of a slab-shaped medium embedded with a spherical inclusion, showcasing pMCX's capability of modeling heterogeneous domains. In this benchmark (Fig. 4), the simulation domain is a $10 \times 10 \times 1.2$ mm³ slab with a spherical inclusion of radius 0.5 mm centered at (6,6,0.6) mm. The inclusion and the slab share identical absorption coefficient $\mu_a = 0.005$ mm⁻¹, reduced scattering coefficient $\mu'_s = 1$ mm⁻¹, and refractive index $n = 1.33$. However, the Mie scatters inside both domains are different. The background medium is filled with scatterers of radius $r = 0.05$ μm and volume density $\rho = 19.11$ μm^{-3} ; the inclusion is filled with scatterers of radius $r = 1$ μm and volume density $\rho = 1.11 \times 10^{-3}$ μm^{-3} . The choices of r and ρ values in either domain was computed based on the Mie theory to ensure

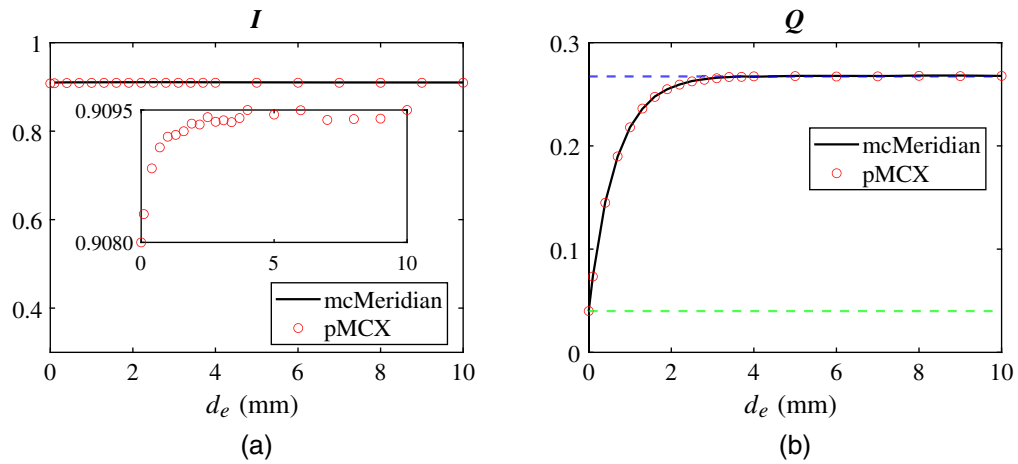


Fig. 3 Validation of pMCX in a two-layer domain. We plot the backscattered (a) I and (b) Q components computed by pMCX and mcMeridian as the superficial layer thickness (d_e) increases from 0 to 10 mm. Two dashed lines in (b) indicate back-scattered Q values computed from a homogeneous slab filled only with the medium of the bottom layer (green) and that of the superficial layer (blue). The inset in (a) shows a zoom-in view of the y axis of the I component obtained by pMCX to demonstrate subtle variations due to sub-diffusive scattering effect.




Fig. 4 Distributions of (a) I , (b) Q , (c) U and (d) V backscattered from a $10 \times 10 \times 1.2$ mm³ slab. A spherical inclusion of radius 0.5 mm is centered at (6,6,0.6) mm.

their reduced scattering coefficients are the same. All spherical scatterings have a refractive index of, $n_{\text{sph}} = 1.59$. A $10 \times 10 \text{ mm}^2$ uniform planar light source is placed at the bottom ($z = 0 \text{ mm}$) surface, pointing toward the $+z$ axis and emitting horizontally polarized light with the wavelength $\lambda = 632.8 \text{ nm}$. A cyclic boundary condition is applied to the four bounding box facets at $\pm x / \pm y$ directions to approximate an infinite slab and infinite-plane source. The incident Stokes vector is $\vec{S} = [1, 1, 0, 0]$. A total of 2×10^8 photons are simulated on an NVIDIA RTX 2080 GPU. We compare the distributions of backscattered $[I, Q, U, V]$ at $z = 0 \text{ mm}$, as shown in Fig. 4. Because the inclusion and background slab share the same μ_a, μ_s' , and n , a regular diffuse optics forward model without polarization capability would generate no contrast to the inclusion. However, our pMCX simulation has revealed distinct image contrasts in I, Q , and V images at the correct inclusion locations, suggesting the potential to detect tissue microstructure differences using polarized light. It is noteworthy that the noise level in each of the images is partially related to the anisotropy g determined based on the background scatterer parameters—a larger spherical radius results in a higher g value and less back-scattered photons. The significantly higher amplitude of inclusion contrast in Q image compared to that in I further demonstrates the advantages of using polarized imaging in detecting scattering differences compared to traditional diffuse optics where only I is typically measured.

4 Conclusion

In summary, we report a massively parallel implementation of polarized MC algorithm in our MCX simulator for modeling the propagation of polarized light inside complex media filled with spherical scatterers. Enabled by its built-in voxel-based geometric representation, the pMCX can handle arbitrarily heterogeneous media. We have described the preprocessing steps to encode the scattering properties of spherical particles with various radii and volume densities into a 3D voxel-based data structure. We provide validation and speed benchmarks ranging from simple homogeneous to complex heterogeneous domains. In all benchmarks, our pMCX solver reports excellent match with the widely used reference solver mcMeridian, while providing a speedup nearly three orders of magnitude. In addition, we observed different GPU memory utilization efficiency among global, constant, and shared memories, with the global memory implementation yielding the highest speed and least restriction. It is noteworthy that a limitation in both mcMeridian and pMCX simulations is that all media boundaries are assumed to have matched refractive indices. We plan to further extend this work to update the Stokes vector across mismatched boundaries.

Disclosures

No conflicts of interest, financial or otherwise, are declared by the authors.

Acknowledgments

This research is supported by the National Institutes of Health (Grant Nos. R01-GM114365, R01-CA204443, and R01-EB026998); by STROBE, a National Science Foundation Science & Technology Center (DMR 1548924); and by the National Science Foundation Engineering Research Center for Precise Advanced Technologies and Health Systems for Underserved Populations (PATHS-UP) (Grant No. 1648451).

Code, Data, and Materials Availability

The pMCX algorithm has been incorporated in our widely disseminated MC simulator and is freely available at <http://mcx.space/>.

References

1. D. A. Boas, C. Pitris, and N. Ramanujam, *Handbook of Biomedical Optics*, CRC Press, Boca Raton, Florida (2011).

2. J. C. Ramella-Roman, I. Saytashev, and M. Piccini, "A review of polarization-based imaging technologies for clinical and preclinical applications," *J. Opt.* **22**, 123001 (2020).
3. V. V. Tuchin, "Polarized light interaction with tissues," *J. Biomed. Opt.* **21**(7), 071114 (2016).
4. C. He et al., "Polarisation optics for biomedical and clinical applications: a review," *Light Sci. Appl.* **10**, 194 (2021).
5. S. L. Jacques, J. C. Ramella-Roman, and K. Lee, "Imaging skin pathology with polarized light," *J. Biomed. Opt.* **7**(3), 329–340 (2002).
6. G. L. Liu, Y. Li, and B. D. Cameron, "Polarization-based optical imaging and processing techniques with application to the cancer diagnostics," *Proc. SPIE* **4617**, 208–220 (2002).
7. S. Zotter et al., "Measuring retinal nerve fiber layer birefringence, retardation, and thickness using wide-field, high-speed polarization sensitive spectral domain OCT," *Invest. Ophthalmol. Visual Sci.* **54**, 72–84 (2013).
8. I. Saytashev et al., "Self validating Mueller matrix micro-mesoscope (SAMMM) for the characterization of biological media," *Opt. Lett.* **45**, 2168–2171 (2020).
9. H. Wiese et al., "Polarized light imaging of the human brain: a new approach to the data analysis of tilted sections," *Proc. SPIE* **9099**, 90990U (2014).
10. J. Chue-Sang et al., "Use of Mueller matrix polarimetry and optical coherence tomography in the characterization of cervical collagen anisotropy," *J. Biomed. Opt.* **8**, 086010 (2017).
11. N. J. Greenfield, "Using circular dichroism spectra to estimate protein secondary structure," *Nat. Protoc.* **1**(6), 2876–2890 (2006).
12. J. C. Ramella-Roman et al., "Preferential alignment of birefringent tissue measured with polarization sensitive techniques," *Proc. SPIE* **9303**, 93030I (2015).
13. M. I. Mishchenko, "Vector radiative transfer equation for arbitrarily shaped and arbitrarily oriented particles: a microphysical derivation from statistical electromagnetics," *Appl. Opt.* **41**, 7114–7134 (2002).
14. S. Chandrasekhar, *Radiative Transfer*, Dover Publications, New York (1960).
15. K. Evans and G. Stephens, "A new polarized atmospheric radiative transfer model," *J. Quant. Spectrosc. Radiat. Transfer* **46**(5), 413–423 (1991).
16. S. A. Prah, M. J. C. van Gemert, and A. J. Welch, "Determining the optical properties of turbid media by using the adding-doubling method," *Appl. Opt.* **32**, 559–568 (1993).
17. A. D. Kim and A. Ishimaru, "A Chebyshev spectral method for radiative transfer equations applied to electromagnetic wave propagation and scattering in a discrete random medium," *J. Comput. Phys.* **152**(1), 264–280 (1999).
18. M. L. Adams and E. W. Larsen, "Fast iterative methods for discrete-ordinates particle transport calculations," *Prog. Nucl. Energy* **40**(1), 3–159 (2002).
19. A. D. Kim and J. B. Keller, "Light propagation in biological tissue," *J. Opt. Soc. Am. A* **20**, 92–98 (2003).
20. H. G. Akarçay et al., "Monte Carlo modeling of polarized light propagation: Stokes vs. Jones. Part I," *Appl. Opt.* **53**, 7576–7585 (2014).
21. M. Xu, "Electric field Monte Carlo simulation of polarized light propagation in turbid media," *Opt. Express* **12**, 6530–6539 (2004).
22. S. Bartel and A. H. Hielscher, "Monte Carlo simulations of the diffuse backscattering Mueller matrix for highly scattering media," *Appl. Opt.* **39**, 1580–1588 (2000).
23. D. Côté and I. A. Vitkin, "Robust concentration determination of optically active molecules in turbid media with validated three-dimensional polarization sensitive Monte Carlo calculations," *Opt. Express* **13**, 148–163 (2005).
24. J. C. Ramella-Roman, S. A. Prah, and S. L. Jacques, "Three Monte Carlo programs of polarized light transport into scattering media: part I," *Opt. Express* **13**, 4420–4438 (2005).
25. J. C. Ramella-Roman, S. A. Prah, and S. L. Jacques, "Three Monte Carlo programs of polarized light transport into scattering media: part II," *Opt. Express* **13**, 10392–10405 (2005).
26. T. Yun et al., "Monte Carlo simulation of polarized photon scattering in anisotropic media," *Opt. Express* **17**, 16590–16602 (2009).
27. Y. Wang et al., "GPU accelerated electric field Monte Carlo simulation of light propagation in turbid media using a finite-size beam model," *Opt. Express* **20**, 16618–16630 (2012).

28. A. Doronin, C. Macdonald, and I. V. Meglinski, "Propagation of coherent polarized light in turbid highly scattering medium," *J. Biomed. Opt.* **19**(2), 025005 (2014).
29. C. Ding et al., "Electric field Monte Carlo simulation of polarized light propagation in multi-layered media," *Proc. SPIE* **10461**, 104610E (2017).
30. NVIDIA Corp., "CUDA C programming guide," Version 10.0 (2018).
31. P. Li et al., "GPU acceleration of Monte Carlo simulations for polarized photon scattering in anisotropic turbid media," *Appl. Opt.* **55**, 7468–7476 (2016).
32. H. Oulhaj et al., "Diffuse optical tomography with polarized light: a GPU-accelerated polarization-sensitive Monte Carlo simulations for efficient sensitivity kernel computation," *Proc. SPIE* **11074**, 110740Q (2019).
33. A. P. Tran, S. Yan, and Q. Fang, "Improving model-based functional near-infrared spectroscopy analysis using mesh-based anatomical and light-transport models," *Neurophotonics* **7**(1), 015008 (2020).
34. Q. Fang and D. A. Boas, "Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units," *Opt. Express* **17**(22), 20178–20190 (2009).
35. J. C. Ramella-Roman, S. A. Pahl, and S. L. Jacques, "Meridian plane Monte Carlo," <https://omlc.org/software/polarization/Meridian.html>
36. L. V. Wang, S. L. Jacques, and L. Zheng, "MCML-Monte Carlo modeling of light transport in multi-layered tissues," *Comput. Methods Progr. Biomed.* **47**(2), 131–146 (1995).
37. Q. Fang, "Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates," *Biomed. Opt. Express* **1**(1), 165–175 (2010).
38. C. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*, Wiley Science Paperback Series (1998).
39. S. Pahl, "ANSI C code for Mie scattering," <https://omlc.org/software/mie/index.html>

Shijie Yan is a PhD student in electrical engineering at Northeastern University. He received his BE degree in information science and engineering from Southeast University, China, in 2013, and his MS degree in electrical and computer engineering from Northeastern University in 2017. His research interests include Monte Carlo photon transport simulation algorithms, parallel computing, and GPU programming and optimization.

Steven L. Jacques is an affiliate professor of bioengineering at the University of Washington in Seattle, Washington, United States. He received his PhD from the University of California, Berkeley, in 1984, then served as a research associate/instructor at Massachusetts General Hospital/Harvard Medical School, as an associate professor at the University of Texas M.D. Anderson Cancer Center, and as a full professor at Oregon Health & Science University (OHSU). He has worked in the field of biomedical optics and laser–tissue interactions for 36 years.

Jessica C. Ramella-Roman is an associate professor in the Bioengineering Department, Florida International University (FIU) in Miami, Florida, United States. She received her PhD from Oregon Health & Science University (OHSU) in 2004. She was a postdoctoral fellow at Johns Hopkins University Applied Physics Laboratory and became an assistant professor at Catholic University of America in 2005, and an associate in 2012 before joining FIU in 2013. Her research focuses on the development of imaging modalities based on spectroscopy and polarization including multimodal applications of nonlinear microscopy.

Qianqian Fang is an associate professor in the Bioengineering Department, Northeastern University, Boston, Massachusetts, United States. He received his PhD from Thayer School of Engineering, Dartmouth College, in 2005. He then joined Massachusetts General Hospital and became an instructor of radiology in 2009 and assistant professor of radiology in 2012, before he joined Northeastern University in 2015 as an assistant professor. His research interests include translational medical imaging devices, multi-modal imaging, image reconstruction algorithms, and high performance computing tools to facilitate the development of next-generation imaging platforms.